Nortel Technical **Journal**

Produced by Nortel's R&D community

FOCUS ON **NETWORK SECURITY**

This issue of the *Nortel Technical Journal* focuses on the architecture, technology innovations, and leading-edge products and services that will help operators and enterprises move beyond traditional approaches to network security, address new multi-dimensional security challenges, and evolve their networks to provide a truly trusted network environment.

Contents

Issue 3, February 2006

The *Nortel Technical Journal* aims to provide a vehicle for Nortel's global R&D community to share with one another and with external audiences the wealth of technology innovation under way across our network of laboratories.

Published by:
Office of the Chief Technology Officer

Guest editor for this issue:
Rod Wallace

Executive editor:
Peter Carbone

Managing editor:
Velma LeBlanc

The articles in this issue do not necessarily reflect a plan of record or commitment to product. The *Nortel Technical Journal* is intended simply as a vehicle through which we can exchange information, share views and perspectives, and expand meaningful dialog within our R&D community and with external audiences.

We welcome your feedback and look forward to your comments, suggestions for improvement, and ideas for future topics. Please send them by email to: norteltechnicaljournal@nortel.com

Printed in Canada on recycled and recyclable paper using vegetable-based inks.

Addressing the new security challenges of real-time, multimedia, and mobile networks

by Rod Wallace

The vision of a truly trusted network environment requires a new approach to network security – one that addresses the challenges for network-wide security in today's real-time, multimedia, and mobile network environments. Nortel – with its strong experience in developing secure carrier-grade networks in both service provider and enterprise environments – is addressing these new security challenges holistically, with a comprehensive security architecture, industry-leading technology innovations, and a suite of leading-edge products and services. In this endeavor, Nortel is one of very few vendors that can leverage a deep understanding of real-time networking performance requirements and knowledge of how all security technologies, devices, and products must interoperate over all types of networks and across all varieties of media and applications.

In 1988, 23-year-old Cornell University student Robert Morris unleashed the Internet's first worm – a 99-line piece of software code that crippled thousands of computers and gave him access to a "federal-interest computer." Since then, thousands of worms, viruses, and other such malicious software (malware) have been created and released by a growing collection of underground hacking communities and cybercriminals intent on exploiting the vulnerabilities of IP-based infrastructures, which have been open to security attacks since the genesis of the Internet.

Indeed, the number of incidents continues to rise exponentially. According to an Internet Security Threat report compiled by Symantec Corp. (a global leader in information security that also maintains one of the world's largest databases of security vulnerabilities), the number of denial of service (DoS) attacks in the first six months of 2005 grew by more than 680 percent, causing widespread and increasingly disruptive overloading of computers and networks. Similarly, in its annual security report, security firm Sophos plc. (an information security leader that also tracks global data on threats) reported that the number of new virus, worm, trojan horse, and spyware threats rose by 48 percent between January and November 2005, as more and more criminals turned to cybercrime.

Not only are the number of security incidents and their associated costs rising every year, but the exploits themselves are becoming more sophisticated. Attacks can be transmitted through simple zip or jpeg files, for instance, and they have the ability to evolve and mutate, making them difficult to discover. Some recent attacks even include auto-update functionality, allowing them to download the final attack code at the last minute. Exploit code is also easier to create: toolkits and libraries that have been created to allow attackers to develop malicious code quickly and make it user-friendly are allowing a larger hacker population to make use of them.

At the same time, exploits now spread faster than ever and can be constructed and launched by less knowledgeable attackers. Today, 80 percent of exploits become available within 60 days of a vulnerability being released, with an increasing number of exploits created in fewer than six days. This trend raises the concern about emerging "zero-day attacks," where networks will have to defend themselves against new exploits that come with no advance warning.

Security attacks are also becoming more versatile in searching out multiple vulnerabilities and ways to infiltrate target systems. New malware can now check for a number of vulnerabilities in systems, as well as mutate and spread via multiple means, such as email attachments, malicious images included in websites, and active attacks against vulnerable protocols and services. Some viruses can even survive a complete wipe-and-install of the computer hard disk by hiding in video card flash RAM or nonvolatile BIOS (basic input/basic output) space.

In this environment, attacks on networks may come through the edge firewall, vulnerable application protocols exposed to the outside, teleworker or road-warrior PCs, email, social engineering, and even mobile phones and PDAs. Moreover, along with the growing frequency of threats and their sophistication, the network environment itself is changing rapidly. Consider the following:

• Wireless has become the predominant mode of accessing the network. In fact, wireless devices now outnumber fixed devices globally. In Asia alone, for example, some 14 million new wireless subscribers are being added every month.

• Service providers are evolving to provide converged services such as quadruple play (voice, video/TV, data, and mobile access over the same infrastructure) or converged wireline/wireless access through cellular networks, WiFi hotspots, and Bluetooth microcells.
• Communication applications themselves are starting to share information with other applications – one common example being voice mail that is sent directly to an email inbox.
• Devices, too, are converging. Laptop PCs are merging with today's personal organizers to create connected mobile devices that are capable of being always on and always connected, and running numerous business applications, including email, document management, data storage, voice, and video.
• New business models are also emerging, giving rise to further security challenges. For example, more and more companies are forming online business partnerships and running their joint businesses over multiple and shared infrastructures, which raises the challenge of needing to exchange important information, such as subscriber lists, without sharing private information about individuals. As well, other companies, such as Skype and Virgin Mobile, are leveraging the ability to unbundle the network from services and applications, allowing them to run a global service provider business without having to own any network infrastructure.

In short, the changing network environment, combined with the ongoing sophistication and rapidly changing nature of security attacks, makes it impossible to effectively combat threats at any single point in the network.

New approach for network security

The traditional wisdom and approach to network security has centered on point solutions – such as firewalls and anti-virus software – and intervention that has largely been manual. Although this approach worked well when the IP network was wired, data-only, and best effort, it is no longer sufficient for today's mobile, multimedia, and real-time network environments. An architectural approach, one that employs defense-in-depth methodology, is necessary to address the new multi-dimensional challenge of network security.

As described in the article about Nortel's Network Security Architecture (page 7), Nortel believes that effectively combating emerging threats requires comprehensive protection at the device, network, and application layers of the network. Specifically, it requires protection of:
• the individual devices, through hardening, patching, access control, use of secure protocols, and enforcement of a security policy;
• the network core, through protecting the perimeter by traditional port-based controls and active detection of malicious traffic, protecting the network interior by detecting and isolating threats that bypass perimeters, and securing all key control and management channels; and
• the applications, by ensuring effective protection of user data and enabling business continuity.

Effective network protection also requires centralized network intelligence that can correlate related events that, while they may appear insignificant on their own, taken together paint a clear picture of an attack under way. Based on this correlation, the network can take immediate action to isolate the sources and targets of such attempts. In the future, tying security to user identity and information rights management will further protect the network and sensitive data from attacks.

Nortel – with its depth of real-time networking experience across carrier and enterprise, and fixed and mobile networks – has a solid understanding of the security issues and vulnerabilities inherent in a converged world of communications, and a clear line of sight into the kinds of security capabilities that networks require today and into the future. In particular, Nortel is developing its security solutions around the following philosophy and principles:
• Security must be endemic to the network, and not just focused on point solutions.
• Security solutions must be application-aware – meaning, the secure network needs to understand which applications are delay-sensitive (voice, video), and must be tuned to perform the necessary security checks without impacting real-time performance requirements and without interfering with quality of service or the user's quality of experience.
• The network should be secured using a layered defense approach – also called defense in depth. This multilayer approach – which ensures end-to-end network security from the endpoints to

Nortel Global Services

Nortel's Global Services organization is a critical element of Nortel's overall security strategy. For companies needing additional assistance in maintaining the layered defense security posture of their networks – whether enterprise, wired, or wireless – Nortel works with customers to:
• develop sound security policies;
• develop a sound, auditable posture – using cost-effective tools – for adhering to industry regulations, such as Sarbanes-Oxley (SOX) in the U.S.; market security guidelines such as ISO 17799 and British Standard (BS) 7799; or the customer's own security policies;
• evaluate the level of risk through regular risk assessments;
• perform incident response and cyber-forensics;
• design and integrate secure VoIP, multimedia, and mobility solutions; and
• manage a myriad of security functions, such as intrusion prevention/detection systems and VPNs.

Investigating the feasibility of remote-timing attacks on modern cryptosystems

by Marcus Leech

To those in the security industry, it comes as no surprise that computer applications and systems require continuous patching. Indeed, security is a field that requires constant vigilance against increasingly knowledgeable cybercriminals, who – lured by the potential for financial gain or notoriety – persistently probe into the deepest corners of system operations to uncover vulnerabilities and create ways to exploit them.

For this reason, designers and researchers at Nortel maintain a close watch on emerging theories and activities surrounding potential vulnerabilities, and work to both elucidate the problems and determine what, if any, impact they may have on the design of communications systems and networks.

An interesting recent example is the ongoing debate in the industry over the feasibility of remote-timing attacks on modern cryptographic systems.

Timing attacks are a type of side-channel attack that involves analyzing the length of time required to execute a function. For instance, it takes longer to compute certain bit strings than others, and this time period provides important information to an attacker trying to uncover the cryptographic key that will unlock access to, say, a financial company's system.

To draw an analogy, imagine a would-be attacker observing, over a period of time, people entering their debit card personal identification numbers (PINs). By observing a large population of users entering their PINs, the attacker could likely distinguish among those who use a 4-digit PIN versus a 5- or 6-digit PIN. While the attacker would not be able to obtain the PINs themselves, he or she could isolate the "easy" targets – those who use 4-digit numbers – by measuring the average PIN entry time at the keypad. This timing information, known as "information leakage," would be considered useful to the attacker, helping him or her to identify vulnerable users.

In a similar way, timing analysis for cryptographic systems involves observing and analyzing the time it takes cryptographic primitives that are used inside the mathematical algorithm to compute. The thought is that timing information obtained through remote-timing attacks would yield actual key bits, or even bit sequences, from a particular segment of the user population.

While some cryptographic mathematicians suggest that remote-timing attacks are possible, others in the communications networking camp argue that such attacks are not.

To better understand the issue and determine the potential impact on communications networks, Nortel researchers – with deep expertise in both cryptography and networking – conducted extensive simulations and quantitative analysis to determine if a potential attacker would be able to measure very subtle nanosecond-scale timing differences over real-world Internet channels.

These simulations showed that in today's Internet such remote-timing attacks are largely infeasible. Unless the extremely fast network links of most of today's communications networks have very tightly constrained jitter – that is, it always takes, say, 37.2 nanoseconds to transmit data in one direction or the other – then a tremendous number of queries would be required to gather enough data to determine the cryptographic key. In addition, the time required to make this number of queries would likely exceed the lifetime of a cryptographic key, which is typically changed frequently.

With this information, system designers have solid guidelines on key-change policies that are considered prudent for cryptographic systems, particularly for systems that use symmetric encryption algorithms.

Marcus Leech *is Senior Advisor, Security Standards, in the CTO Office. He has expertise in cryptography, software design, and systems security, and is a former Security Area Director in the Internet Engineering Task Force (IETF).*

the network core – guards against single points of security failure by using multiple options to provide protection.

• The network should be secured using standards-based solutions (which minimizes integration costs), should enable state-of-the-art security capabilities, and should address future security needs.

• The network must be adaptive, able to proactively respond to attack, and provide real-time situational awareness. Manual intervention will not suffice when a swift response is needed.

• A company's security plans must encompass all users, processes, and technologies, as well as the impact each has on the network's security.

• Appropriate attention should be paid not only to the design of a secure network, but also to maintaining network security through best practices, such as risk assessments. Maintaining security requires a plan to work closely with other network vendors when vulnerabilities in network protocols or specific products are discovered.

In effect, the approach to network security should be aimed at creating a trusted place where people will go with their businesses, their personal information, their lives. More than just making sure that others can't intercept their communications, people need to trust the network in a very holistic sense.

This vision of a secure, "trusted" network builds on Nortel's underlying

belief that security can encompass much more than simply protecting the network from attacks, and certainly more than a "necessary evil" or an impediment to productivity, as many in the industry believe.

On the contrary, a converged network is capable of ultimately providing better security than would ever have been possible in the separate IP or telephony worlds. Indeed, a secure, trusted network can be a powerful enabler for new services, increased productivity, and simplifying and enhancing the user experience.

To this end, Nortel's Network Security Architecture – essentially a roadmap – captures Nortel's vision of how network security, and the technologies and products that enable it, will evolve (page 7).

This architecture stands as the company's framework for understanding, interpreting, and addressing the industry needs in this space. It can serve as a model for our customers' network security planning and, to date, has generated a tremendous amount of interest and feedback from customers, industry security experts, and decision makers in governments.

The premise of Nortel's vision is that once a network is provisioned with a basic level of end-to-end security – which focuses mainly on "keeping the bad guys out" – it can then be evolved to "enable the good." In other words, it can be a network that not only understands what is happening at the time of a threat and proactively responds without waiting for human intervention, but also one that grants access to only the information that individuals are entitled to.

Looking past this horizon, a secure trusted network could enable exciting and powerful new applications, such as networks that are able to automatically and easily manage the identities of individuals, are aware of a person's presence and location, and can send the right – and protected – information to users when, where, and how they want it.

While the architecture represents the segmentation of all the security technologies and the evolution path Nortel believes they will follow, the company also has, with its layered defense approach, a "recipe" for how these technologies should be implemented in a structured, building-block fashion. This layered defense approach is outlined in some of the articles in this issue of the *Nortel Technical Journal*, along with other key initiatives, technology innovations, and solutions under way in Nortel's R&D labs that are helping customers move forward in providing a secure, trusted network environment.

• The article on page 7 details Nortel's Network Security Architecture, highlighting the three-phase evolution – from a basic secure network, to an autonomic network, through to an authenticated network.

• The article on page 13 provides insight into how today's network security challenges are bringing about a convergence in the traditionally separate roles of a company's Chief Information Officer (CIO) responsible for IT security and Chief Security Officer (CSO) responsible for physical security, and underscores the importance of networks in providing real-time situational awareness.

• The article on page 20 describes the pioneering work of Nortel researchers to develop, for implementation across all of its portfolios, a set of baseline security requirements – a uniform set of network-wide security requirements that form the "must-haves" for basic network security. Recognizing the need for the industry to adopt a common set of requirements for the benefit of providers, enterprises, and vendors, the Nortel development team rallied the industry to standardize on these baselines so that all users can benefit.

• The article on page 28 describes how Nortel is helping enterprises choose the right security solutions that will work together, across all layers of their networks, to make their infrastructures not only highly resilient against attacks and compliant with today's many regulations, but also able to evolve easily over time with new capabilities and applications that boost productivity.

• The article on page 37 details how Nortel is addressing the real-time security requirements of Voice over IP (VoIP) and multimedia traffic in both carrier and enterprise networks, and is ensuring that secure solutions for converged networks are tuned to multi-service performance requirements.

• The article on page 46 discusses the challenges of securing wireless networks across all of today's different access technologies, and highlights Nortel's world-leading technology innovations – among them, its best-in-class Mobile IPsec technology – to meet the unique security needs of subscriber nomadicity.

• The article on page 54 addresses the stringent security specifications of governments – another area where Nortel is driving innovation and leadership. This article looks at how Nortel Government Solutions (formerly Nortel PEC Solutions following the 2005 acquisition of U.S.-based PEC Solutions) is working closely with the U.S. Government to integrate best-in-class technologies and products to secure its critical national infrastructures and design the trusted networks that are key to secure information sharing among government entities. While this article highlights the U.S. Government as an example, the network concepts discussed are also relevant to other government infrastructures around the world.

• And, the final article in this issue (page 62) provides some perspectives on where the industry may be headed, more detail on what a trusted communications environment could enable, and how it will profoundly change the way people, businesses, and societies communicate.

History of leadership

Across all these initiatives, Nortel is building on its leadership in developing secure carrier-grade networks for enterprises and service providers, and extending its already considerable edge in the area of security solutions. Nortel's

Nortel: Helping to shape industry direction for network security

by Marcus Leech

Nortel participates in more than 85 different international, regional, and national standards development organizations (SDOs). Nearly all of these standards bodies have as part of their focus one or several components dealing specifically with security – either through defining new security protocols, or defining protocol elements, operational procedures, and systems engineering standards that involve security.

International and national technical standards are often critical in the development of any telecommunications or networking product. Products need to "talk" to other products over media and transmission systems whose characteristics must be standardized to allow maximum interoperability. Standards compliance can reduce or eliminate customer challenges in deploying equipment from multiple vendors. Indeed, in the network security area, when "many eyes" have evaluated and suggested changes to a protocol standard, that can lead to a better standard – one that can be trusted from a security perspective.

For more than a decade, Nortel has actively participated in the security standards development process, by both promoting various security-related standards and accepting key management roles in important standards bodies. Nortel continues to work across all major SDOs, including the International Telecommunication Union (ITU-T), European Telecommunications Standards Institute (ETSI), Third Generation Partnership Project (3GPP), U.S. Alliance for Telecommunications Industry Solutions (ATIS), and the Internet Engineering Task Force (IETF).

For example, among its many security standards-related activities, Nortel:
• took a lead role in defining baseline security requirements and in priming their development into industry standards, through such bodies as the ITU-T, ANSI, and 3GPP (see article on page 20);
• chairs the IETF PKI4IPSEC working group – an effort designed to ease the burden of certificate management for IPsec deployments;
• held management roles in other IETF security work, including Authenticated Firewall Traversal (AFT), Internet Tracing (ITRACE), and Multicast Multimedia Security (MSEC);
• made key contributions to the early work of the Link-Layer Security Task group in IEEE 802.1, which produced the 802.1AE MACSec security standard;
• holds the vice-chairman position of the important SA3 (Security) working group of the 3GPP; and
• contributed to the early work in the IETF on Secure Inter-domain Routing.

In addition to standards activities, Nortel consults and works closely with such key security advisory government agencies as the National Security Telecommunications Advisory Committee (NSTAC) in the U.S. and the National Infrastructure Security Coordination Centre (NISCC) in the U.K.

Nortel is also actively involved in key professional security advisory organizations, including the Carnegie Mellon University's Computer Emergency Response Team (CERT) and the global Forum of Incident Response and Security Teams (FIRST). As well, Nortel is a board member of the Internet Security Alliance, which promotes information security practices, policies, and technologies to enhance the security of the Internet.

Nortel also participates in other bodies, including:
• ICSA, an industry-recognized authority for security certifications for firewalls, VPNs, and other security products;
• the U.S. Network Reliability and Interoperability Council (NRIC);
• the U.S. National Coordinating Center (NCC), which facilitates voluntary collaboration and information sharing among Government and industry ICSA members;
• the System Administration, Networking, and Security (SANS) Institute, a cooperative research and education organization of more than 150,000 system administrators, security professionals, and network administrators who share lessons learned and work to find solutions to emerging security challenges; and
• the Trusted Computing Group (TCG), an industry-led initiative formed to develop and promote standards-based security specifications for protecting and interworking devices and equipment across multivendor network security solutions.

As well, Nortel coordinates and participates with its partners in security standards forums. For instance, Nortel and Symantec standards leaders had a number of sessions together about TCG that led to Nortel joining its Trusted Network Connect work group. Symantec and Nortel also co-hosted the November 2005 IETF-64 meeting in Vancouver, Canada.

Marcus Leech *is Senior Advisor, Security Standards, in the CTO Office. Marcus was also area director for security in the IETF for four years.*

solutions have long been relied upon by more than 80 percent of the top 100 U.S. banks and continue to support billions of transactions at the world's largest and most important stock exchanges. For instance, Nortel:
• has developed a suite of award-winning, best-in-class security products, including VPN Gateways and Routers, as well as Application Switches and Switched Firewalls;
• led, and continues to lead, industry

consortia and standards in defining leading-edge architectures, including the baseline security requirements for the management plane, which have been adopted by key standards organizations around the world (see page 20);
• has held the #1 position in the dedicated SSL VPN appliance market for three years running, and holds the #2 position in the overall SSL VPN market (Infonetics Research);
• received Frost and Sullivan's 2005 Award for Product Line Strategy Leadership for secure VoIP equipment for carrier and enterprise markets; and
• holds a number of patents around innovative security technologies, including mobility enhancements to IPsec. As well, Nortel demonstrated, in September 2005, the world's first successful prototype of a 10-gigabit-per-second optical switch with integrated 256-bit encryption, using the Advanced Encryption Standard (AES), to ensure that the highest levels of security are met for confidential and sensitive communications (see page 67).

In addition, Nortel's solutions incorporate best-of-breed technologies and products developed by leading vendors with whom Nortel has established partnerships and strategic relationships. A key pillar of Nortel's solutions development strategy is to partner with those companies developing the best of the best and then integrate those solutions into an end-to-end network offering.

Nortel is working with several key partners, including Check Point Software Technologies Ltd., IBM, and Symantec, to develop various technologies, such as key intrusion prevention/detection technology, firewall technology, secure XML processing, secure event management, endpoint security solutions, and next-generation security systems, among many others.

Also important to Nortel's overall security strategy is its Global Services organization, which is working to help carrier and enterprise customers better understand their security requirements, pinpoint any specific vulnerabilities, and design individual solutions to eliminate weakness in their networks (page 2).

Continuing to shape industry direction

In addition to developing a broad spectrum of end-to-end security solutions, Nortel also understands that the technology challenges for network security are not only complex, but also never-ending. For example, unlike engineering for network reliability, where reliability issues can be permanently engineered out of a system, Nortel understands that network security threats are constantly evolving and technology must constantly adapt.

For this reason, Nortel:
• maintains watch on emerging theories and activities surrounding potential vulnerabilities, and works to both elucidate problems and determine any impact on the design of systems and networks (see sidebar page 3);
• shares its insights and security best practices with customers, industry consortia, and external security advisory organizations;
• works to influence industry directions through national and international standards bodies. Nortel works across all major standards development organizations, including the International Telecommunication Union (ITU-T), Third Generation Partnership Project (3GPP), U.S. Alliance for Telecommunications Industry Solutions (ATIS), and the Internet Engineering Task Force (IETF), to develop standardized security solutions (see sidebar page 5);
• works with key government critical-infrastructure advisory agencies, such as the National Security Telecommunications Advisory Committee (NSTAC) in the U.S. and the National Infrastructure Security Coordination Centre (NISCC) in the U.K., to ensure the resiliency of government networks upon which national communications structures rely;
• is actively involved in key professional security advisory organizations, including the Carnegie Mellon University's Computer Emergency Response Team (CERT) and the global Forum of Incident Response and Security Teams (FIRST), to understand current vulnerabilities and their potential impact on Nortel products;
• has a strong focus on best practices, which include vulnerability assessment and management, company-wide education on vulnerability scanning, and programs for post-incident analysis and support. New initiatives will focus on feedback mechanisms for the design community and lead to more robust solutions; and
• offers emergency support to service providers and enterprises through its Security Incident Response team, which delivers immediate assistance to deal with security problems or attacks. The team can work either on-site with network administrators or can hook into the attacked network remotely through one of Nortel's three network management centers, which provide managed services and outsourcing support to customers worldwide.

Combining these efforts with its innovative technologies, suite of products and services, and forward-looking architecture, Nortel is helping its customers to not only address today's increasingly complex network security environment, but also cost-effectively evolve their infrastructures toward a truly trusted communications environment. ■

Rod Wallace *is Leader of Security Solutions and Services. Rod is also a key contributor to the National Security Telecommunications Advisory Committee (NSTAC) and the Cybersecurity Workgroup within the FCC's Network Reliability and Interoperability Council. In addition, Rod leads Nortel's participation in NSTAC and the Alliance for Telecommunications Industry Solutions (ATIS), and is a board member of the Internet Security Alliance.*

Nortel's Network Security Architecture: New dimensions in network security

by Matt Broda

Nortel has developed a forward-looking architectural approach that will help operators and enterprises systematically and seamlessly evolve their networks toward the secure, trust-based network environments envisioned for the future. Through its Network Security Architecture, Nortel lays out its vision for how network security and associated enabling technologies and products will evolve through three defined network phases – from a basic secure network, to an autonomic network that provides enhanced situational awareness and reduces the need for human intervention, to an authenticated network that can automatically adjust its security posture to individual users.

Nortel's Network Security Architecture reflects Nortel's holistic vision of how network security will evolve to bring increased levels of trust and enhance user productivity. The objective of the architecture is to create the framework that guides Nortel's security technology evolution and to serve as a model for customers' security planning strategies as they evolve their networks.

This architectural framework evolves through three phases: basic, autonomic, and authenticated. This three-phase evolution, however, is not strictly sequential, since customers may achieve different levels of security at different times in their networks, based on their own unique requirements. Rather, these phases should be viewed as different maturity levels with associated functionality that, in some cases, already exists in solutions. As networks and solutions evolve from today's primary focus on "keeping the bad guys out," we can expect increased focus on taking human involvement out of the attack-response process and on "enabling the good."

In each of the three phases, the architecture maps across the three key network domains: the applications that run on the network; the network itself; and the various connected user devices, including wired and wireless, handheld and desktop.

Application domain: Applications and services – which users access through the network – are what make networks useful, by providing information, enhancing productivity, and enabling effective communication. Applications include network file storage and backup, document repositories, directories, customer lists, and services such as multimedia, email, voicemail, instant messaging, Voice over IP (VoIP), and conferencing. The focus in this domain is to protect data stored in the network as well as data in transit, while at the same time protecting users from unsolicited traffic, such as email SPAM, instant messaging SPAM (called SPIM), or VoIP SPAM (called SPIT).

Network domain: This domain refers both to the "in the cloud" infrastructure – the LAN, MAN, WAN elements – that enables communication between devices, as well as to the traffic flowing through that infrastructure. Securing the network domain involves defending the network from attacks and protecting device-management and network-management traffic, and includes:

- network perimeter protection of key security zones;
- safeguarding the network interior from insider attacks, and providing a secondary layer of defense and detection for perimeter protection;
- authentication and integrity of signaling control and management traffic, such as user provisioning, network routing, and topology control data; and
- encryption of sensitive data, such as user passwords, security configuration data, and lawful intercept configuration data.

Network boundaries, defined by zones of trust, need to be protected from intrusions and attacks. This protection is accomplished by sanitizing and selectively allowing connections to cross zone boundaries, and is supplemented by monitoring network traffic patterns within the zones to detect anomalies and attack patterns.

Securing the control and signaling traffic involves making sure that any network routing protocols as well as application control and signaling protocols, such as Session Initiation Protocol (SIP) and H.323 signaling for VoIP networks, are adequately protected from attacks.

Device domain: Every individual endpoint and network device that provides or uses network services needs specific security functionality to address user and administrator authentication and authorization, ensure the application of secure communication protocols, and protect the integrity of device software and configuration. Devices can range from call servers, media switches, and firewalls to all user devices, such as

network-connected personal digital assistants (PDAs) and user workstations.

Phase One: Basic secure network

The first "basic" phase of the Network Security Architecture focuses mainly on "keeping the bad guys out," which means ensuring that unauthorized users as well as malicious and unsolicited traffic is kept out of sensitive zones in the network. It also means protecting valuable data stored on the network from compromise, whether accidental or deliberate.

In the basic secure network (Figure 1), components across all three domains (application, network, and device) work together using a "layered defense" approach to address different kinds of threats, and to ensure that there is no single point of security failure.

At the application domain, what users care about most is that their valuable data – credit card numbers, personal address books, documents, and financial records as examples – is stored securely in the network, accessible only by authorized users, protected from accidental or malicious loss, and properly backed up, as well as transported in an equally secure manner.

Nortel employs a number of technologies and solutions to address application-level security. For instance, IPsec (Internet Protocol Security) and SSL VPN (Secure Socket Layer virtual private network) functionality are used to protect business-application content between business locations and mobile or remote users. Nortel has also pioneered several innovative solutions to enable VoIP traffic to run effectively over both types of VPNs. Other application domain solutions include: unsolicited traffic protection, with applications such as content filtering solutions and sophisticated filtering within Nortel's Application Switch (formerly the Alteon switch); emerging secure storage solutions; and firewall functionality that provides advanced understanding of the application protocol and protects application-level traffic from both unsolicited traffic and application-level attacks.

The articles in this issue discuss some of these solutions, including the Nortel Secure Network Access (SNA) technology (page 31), and the company's work with Websense Inc. to protect mobile users (page 48). Looking forward, Nortel is working on technologies and partnerships to address other requirements in this space.

In the basic phase, solutions that address the network domain focus on protecting both the perimeter and interior of the network. Solutions to protect the network perimeter include enforcing network segmentation – by providing various types of protection with firewalls, deep packet inspection, intrusion prevention/detection, malware scanning, and content inspection – at the boundaries of network zones. Protecting the network interior involves enabling detection and mitigation of threats that originate inside secure network zones, as well as providing a secondary defense for threats that make it through perimeter defenses. Nortel's Switched Firewalls, Network Application Switches (NASs), and access control lists (ACLs) and VPN functionality, which are implemented across a number of Nortel products, provide strong support for network protection. Nortel is working with its partners on perimeter protection products that will inspect traffic even more closely, as well as on additional capabilities for various multimedia-specific border protection devices.

Nortel's Threat Protection System (TPS) (page 35) also provides strong intrusion prevention and detection functionality across the network domain. TPS sensors situated throughout the network send event activity, such as an intrusion alert, to a centralized defense center, which then correlates these events, alerts operators when suspicious activity or likely attacks are detected, and automatically takes steps to eliminate the threat. Intrusion prevention and detection capabilities are also embedded in several individual products. For instance, Nortel recently demonstrated on its NAS a prototype of Symantec's first-attack protection, which enables the switch to recognize the latest viruses and attack types.

In addition, the basic phase provides endpoint security, giving network operators the ability to set security policies that will be enforced on devices before allowing them to join the network. If a device fails a compliance check, the device may be admitted with restricted access (i.e., quarantined), it may be sent to a remediation portal (to update malware definitions and re-scan, for example), or it may be blocked from the network altogether. The Nortel SNA endpoint security solution provides endpoint compliance enforcement on VPN tunnel connections (leveraging Nortel's Tunnel Guard technology), on endpoints connected directly to LANs, and on those connecting over wireless LANs (WLANs). The Nortel SNA solution ensures that all wired and wireless users, connected locally or remotely, have valid identities and comply with security policies.

As well, network and device management in this first phase of the architecture focuses on the basic management capabilities needed to provide device security configuration and basic network security set-up and monitoring. These capabilities create the basis for Phases Two and Three of the architecture, where the network is enhanced to enable advanced situational awareness, autonomic threat response, and user-focused security policy management.

Baseline security requirements: In addition to the aforementioned protections, Nortel developed over several years a set of baseline security requirements that constitute a uniform set of network-wide security measures and best practices for implementation across product portfolios. This baseline provides protection across both network and device domains. It includes all the security "primitives" that must be implemented in the network devices to enforce user access and protect devices from attack, as well as the necessary

Figure 1. Phase One – Basic secure network

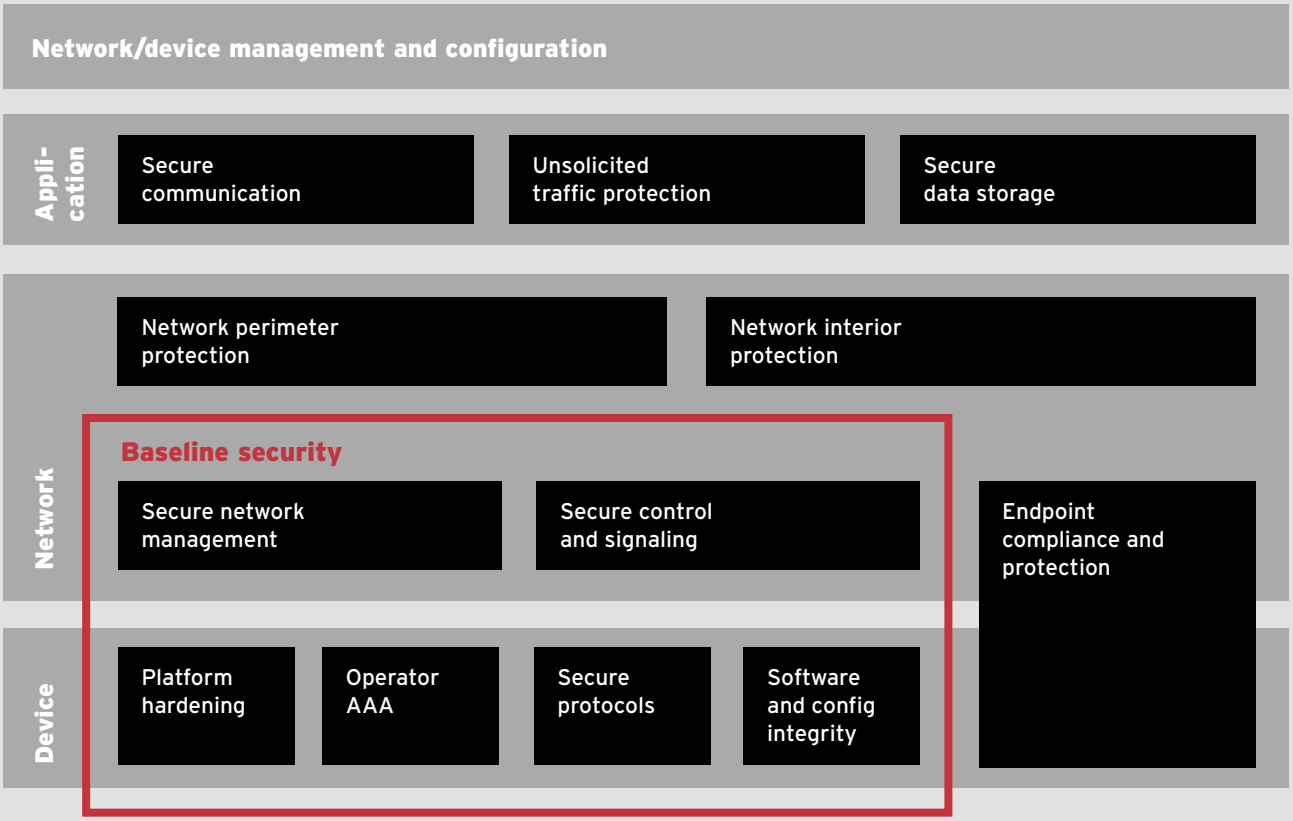

Phase One of Nortel's Network Security Architecture focuses on "keeping the bad guys out," by ensuring that unauthorized users as well as malicious and unsolicited traffic are kept out of the sensitive zones in the network, and that the valuable data stored on the network is protected. In addition to protective measures in the application, network, and device domains, a basic secure network includes key baseline security requirements – a uniform suite of standardized security measures.

measures to secure network management, signaling/control, and user traffic in the network. Nortel was one of the first vendors to develop a suite of baseline security requirements for its own portfolios, and rallied the industry to standardize on these baselines in both national and international standards bodies (page 20).

Phase Two: Autonomic network

The next phase of the architecture focuses on reducing the complexity and need for human involvement in day-to-day security management activities. In a world where security attacks can impact the reliability of the network within seconds, the network needs to have the intelligence and inherent situational awareness to be able to automatically protect itself and recover from threats, on a second-by-second basis. As well, the network needs to be able to automatically adapt as more and more applications actively seek to directly control the network configuration. For instance, an application that sets up a peer-to-peer whiteboard or other collaboration session, or a videoconference session, needs to signal the network to dynamically allocate more bandwidth or other resources to ensure quality of service. In such cases, the network must be able to quickly and automatically recognize those application requests that are legitimate and compliant with the network policy and those that are not.

An autonomic network, then, is able to respond without human intervention more quickly to attacks across all three domains (device, network, and application) to contain any damage immediately and operate seamlessly under many different attack conditions.

An autonomic network can also automatically adjust to the needs of specific applications to create an optimal operating environment for them. Furthermore, such a network will provide more meaningful and actionable security reporting to operators.

The ultimate goal in Phase Two is to implement a common approach through which network elements and applications alike can generate and report relevant events to a centralized event correlation and remediation function that is responsible for collecting, normalizing, and processing network-wide security events (Figure 2). With

Figure 2. Phase Two – Autonomic network

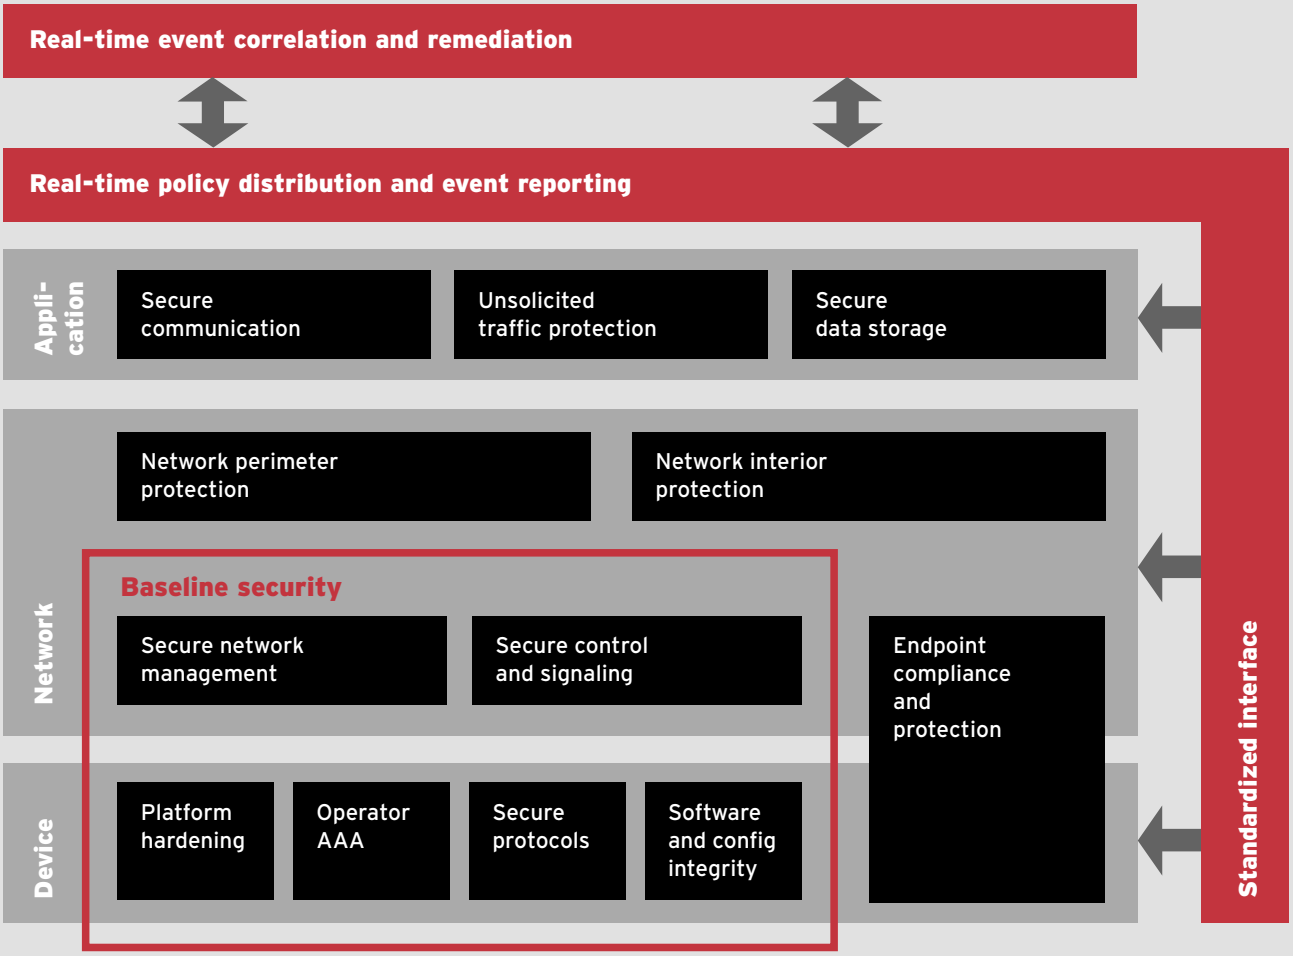

In Phase Two of Nortel's Network Security Architecture, the network evolves to include a real-time event correlation and remediation capability, as well as a real-time policy distribution and event reporting function. These capabilities give the network the intelligence and the standardized interface it needs to correlate event information from all devices in the network, to provide real-time situational awareness, and to automatically protect itself and recover from threats without human intervention.

this information provided on a real-time basis, the network can reliably identify attack patterns and anomalies, effectively detect and mitigate such attacks as distributed denial of service (DDoS) and coordinated hacking attempts, reconfigure network devices in real time, and take corrective actions to protect key resources.

A security information manager (SIM) function is central to realizing autonomic network capabilities. This function enables the intelligent correlation of events across the entire network, followed by autonomic action under attack circumstances. The SIM function would then report both the correlated information and network response, in real time, to network security operators.

This capability would serve to simplify the information with which security operators need to contend, by eliminating false positives, correlating and normalizing related information, and presenting the most meaningful view for the given audience.

In this way, network personnel can obtain a meaningful real-time report of the state of their network security "health," enabling them to monitor the working conditions of their various defenses and alerting them when any part of the security system fails.

Nortel is actively pursuing various solutions in this space. For example, Nortel researchers are working with GuardedNet (recently acquired by Micromuse Inc., which has since been acquired by IBM) on an interface, based on the IPFix protocol, that will enable Nortel's Ethernet Routing Switch to interwork with a third-party security information management platform. [The IPFix protocol, defined by the Internet Engineering Task Force (IETF), enables high-level flow information to be sent to a correlation center.]

The increasing focus on compliance

Figure 3. Phase Three – Authenticated network

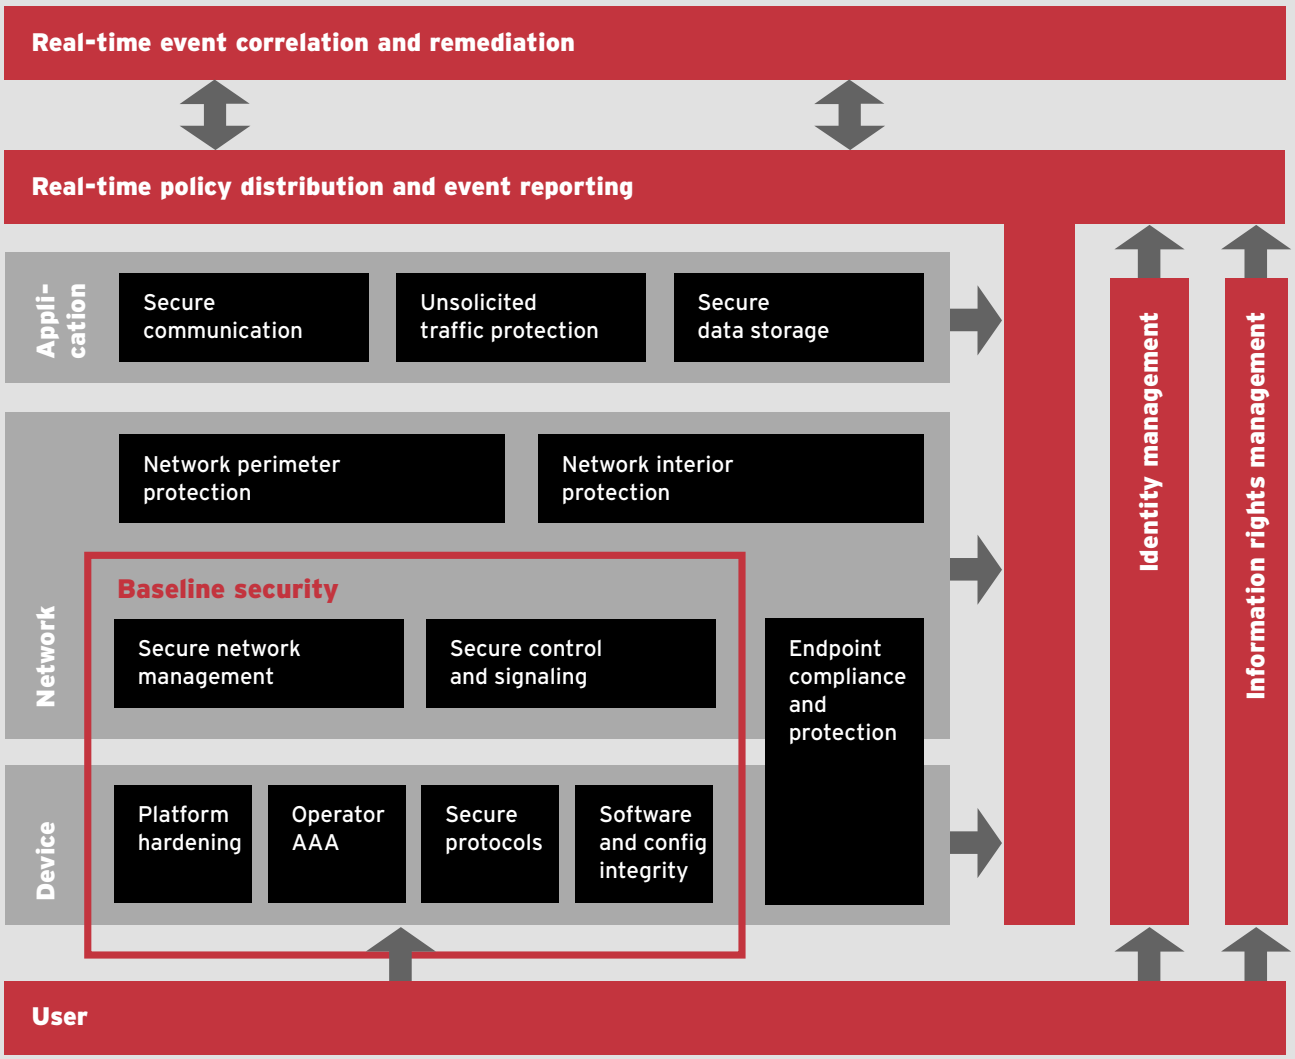

As networks evolve to Phase Three of Nortel's architectural vision, increased emphasis can be placed on "enabling the good," by building on the capabilities introduced in the first two phases. Key additions in this phase center on identity management, which enables the network to recognize who users are, what devices they are using at any given time, and where they are located; and on information rights management, which determines which security policies to enforce given the particular circumstances under which a user is accessing the network, and allows users to access only that information to which they are entitled. Following this phase of evolution, the network framework is in place to support a truly trusted communications environment across carrier and enterprise domains.

with existing and emerging regulations worldwide, as well as the need to create and enforce their own corporate policies, also makes it important for customers to be able to assess and audit the security of their networks. To this end, Nortel's solutions include a third-party auditing tool (called the ComplianceAuthority from SecureInfo Corporation) that assists operators in identifying and addressing non-compliance with security best practices. Through this tool's "dashboard-type" display, personnel can determine whether the network, on any given day, is compliant with such legislation as Canada's Personal Information Protection and Electronic Documents Act (PIPEDA) and the U.S. Sarbanes-Oxley (SOX) legislation and Health Insurance Portability and Accountability Act (HIPAA), as well as with customers' own corporate policies and such industry standards as International Standards Organization (ISO) 17799 and British Standard (BS) 7799.

Phase Three: Authenticated network

As more mobile users use a variety of devices to access the network, the traditional concept of a physical perimeter no longer applies: in many cases, the

perimeter is, effectively, wherever the users are. Thus, the third phase of the architecture focuses on tying together traditional network security approaches with network awareness of both the user's access device and the user's physical environment, to bring an entirely new dimension to network security.

In other words, Phase Three involves enabling the network to recognize who users are, what devices they are using at any given time, and where they are located. Phase Three also involves enabling the network to adjust its security posture to individual users, determine which security policies to enforce given the particular circumstances under which a user is accessing the network, and allow users to access only that information to which they are entitled under the circumstances.

Enabling these capabilities are two fundamental new elements introduced in Phase Three: identity management, and information rights management (Figure 3).

Identity management involves enabling network knowledge of user identities through individual profiles – that is, who the users are and how they are authenticated – across different network domains. This identity profile consists of various credentials, such as a unique user ID and authentication information, which could be a password or an answer to a personal question ("something you know"), a smart card or other physical device ("something you have"), or a biometric sample such as a fingerprint or retinal scan ("something you are"). With some forms of identity management, users must remember multiple passwords and multiple credentials (e.g., a different password or identifier for voice mail, email, and website access). In the Nortel preferred identity management solution, the network would store and manage a single identity profile (although multi-factor authentication could be applied), and would be able to recognize a user regardless of device type or physical location.

As identity management capabilities become more commonplace, multiple companies and network operators could establish agreements to work as a federation, allowing subscribers to connect across the domains of many different networks using a federated identity. (For more on federated identity management, see pages 60 and 65.)

Information rights management uses the identity profile to more effectively enforce the security policy of an enterprise or organization with which a user is associated, and ensure that individuals can access only the information to which they are entitled. Furthermore, the network becomes aware of the access rights and dispositions on the various types of information and can intelligently control what happens to the information and documents, preventing malicious or accidental compromise.

For the network to communicate this identity and policy control information across its various domains and among the different network elements, the control interfaces introduced in Phase Two need to evolve to provide a capability for real-time policy decision-making.

The key to this capability lies in the interface between security functions at the user, device, network, and application levels and those in external or higher-level enterprise security/policy infrastructures. This interface needs to work hand-in-hand with integrated management and real-time control of user access to network services and data, as well as personalization and customization of these services and data according to a user's identity, location, and access device type, allowing federation across different domains based upon a set of predetermined corporate policies.

For example, a worker who accesses his corporate network from the office would be given full access to the network. But when that same employee accesses the corporate network from another location, such as a public Internet kiosk at the airport, the network automatically recognizes that the user is accessing the network from an unsecured location. In such locations, the network might limit access to a subset of the corporate information (such as only non-confidential or public information), and enforce further restrictions, such as limiting the ability to download confidential information onto a local device. A short time later, the employee takes a call from his boss using a mobile device that connects through a WiFi or WLAN access point. The network recognizes that the call is confidential and either secures the data appropriately or notifies both parties that they are using an unsecured connection. Similarly, a contractor would be limited in access rights, relative to regular employees, and be permitted access only to certain sites, such as those housing public non-proprietary information or billing processes. Or, a visiting customer, on-site for a meeting, could plug her laptop into the conference room port and the network would immediately recognize her as a customer, with limited permission to connect through the corporate network – to the public network, for example, but not to corporate sites.

These examples illustrate how, with identity management and information rights management functionality integrated with traditional security mechanisms, the network is able to automatically adjust its security posture to individual users – simply and transparently.

With these three architectural phases – basic, autonomic, and authenticated – Nortel's Network Security Architecture takes a forward-looking view of network security, looking beyond today's industry trends and evolving customer needs to anticipate and prepare the framework for systematically evolving the network from one that provides basic security to an end-to-end communications environment that is secure and trusted. ■

Matt Broda *is Leader, Strategic Security, CTO Office.*

The convergence of IT and physical security systems

by Scott Hurd and Tim Williams

The growing convergence of networks driven by IP packet technology and the digitization of all kinds of business information and communications have created new opportunities for enterprises to implement powerful and efficient security structures that integrate the currently separate domains of IT and physical security. More than just a melding of technology, however, security convergence also requires secure-by-design integrated security policies and processes, as well as a security-aware and vigilant workforce. As a leader across the whole span of convergence technologies and solutions – from optical-enabled packet networks and converged wireless/wireline applications to integrated triple-play (voice, data, and video) multimedia services – Nortel is ideally positioned to unleash the power of the network to help organizations build end-to-end converged security solutions designed to thwart increasingly sophisticated attacks on their business assets and operations.

With enterprises more vulnerable than ever to both internal and external attacks on their networks and facilities, security and risk management have leapt to the forefront of almost every aspect of business activity. The increased integration and interoperability of business operations, the growing use of mobile communications and remote network access, and the reliance on external partners, suppliers, and customers that have enabled the emergence of a global, on-demand economy have at the same time raised the security risk for many enterprises and organizations. A significant event or attack affecting one part of the infrastructure (such as the supply chain or corporate network) can have a devastating effect on local operations around the world. Moreover, companies are now not only affected by the vulnerabilities and threats facing their own organizations, but also by those of others.

Traditional security and risk management practices have examined and confronted potential threats to an organization by defining them as either IT-related or physical, typically the responsibility of the Chief Information Officer (CIO) and Chief Security Officer (CSO), respectively. Under this division of responsibilities, IT security encompasses those aspects required to protect the network, including virus protection, logical access controls to various networks and applications, patching of software to stay ahead of security vulnerabilities flagged by manufacturers, and a myriad of related concerns. Physical security, on the other hand, relates to all aspects of facility security, including security staff, closed-circuit television surveillance monitors, building access controls, and related processes and equipment.

Such classifications often result in standalone solutions that address specific threats and mitigate damage without regard to the impact on the organization as a whole. What's more, general enterprise security efforts that provide protection for people, proprietary information, and property in all forms – as well as crisis management, incident response, investigations of white collar crime and related losses, counterfeiting, product diversion, and many other related matters – end up segmented between the IT and physical security domains. As a result of this isolated perspective, security solutions often fail to account for all of the business vulnerabilities in a comprehensive, coordinated fashion.

Advances in technology and the convergence of networks and applications are opening up new opportunities to bring these domains together. Today, for example, sensors exist that can monitor and detect hundreds of different physical attributes, including temperature and pressure, chemical composition, biological agents and radioactivity, and radio frequency identification (RFID) signals from tagged assets. In one case, Nortel discussions with one advanced healthcare institution identified more than 30,000 network addressable sensing and control points in an existing building's HVAC (heating, ventilation, and air conditioning) control systems. With digital technology, this physical world of sensors can be transformed into bits and bytes that can be communicated, stored, and processed by IT systems and networks.

In truth, areas that have long been the purview of the physical security realm (such as closed circuit monitors and building access controls) are only now beginning to benefit from integration with the longstanding strengths of the IT domain (high-speed networks capable of carrying real-time audio and video, centralized databases of user identities and access entitlements, etc.). A simple example

of integration would be networking a collection of distributed closed-circuit monitors to a central network operations center, where the video feeds can be correlated with other corporate information and managed more efficiently and cost-effectively (see sidebar on page 15).

The progression to such integrated enterprise-wide security systems is still in its infancy, at the discussion and planning stage. Nortel was one of the first companies to realize the potential for unified IT/physical corporate security enabled by emerging convergence technologies. As a company that has leadership across the whole convergence space – from network and network element convergence to applications and services convergence – Nortel is one of the few in the industry with the end-to-end view required to plan and implement a unified enterprise security system.

Regardless of the type of asset, whether digitized information or physical items, the general principles for securing that asset are common: understand the value of the asset, together with any risks or threats associated with its loss, and put appropriate controls in place to ensure its effective protection and continued fitness for use to accomplish the required business objectives. Consequently, aligning the organizations, objectives, and finances of the various groups responsible for the different security functions within a business is key to successful convergence in the security space.

Early in 2005, Nortel took a significant step toward convergence in the security space through the organizational alignment of the traditional IT security function (e.g., information and network security management) and the corporate security function (e.g., protection of people, forensic investigations, and all manner of IT- and criminal-related investigations). These functions are now fully aligned with respect to strategy, investment, and implementation, and Nortel has already realized significant benefits in terms of efficiency and cost avoidance.

What keeps us awake at night

The need for a more powerful integrated enterprise security approach is being driven by increasingly sophisticated attacks on business assets and operations – incursions that are growing in both number and severity. With ubiquitous Internet access and a plethora of chat boards, hacker sites, and reusable technology on which to draw, the task of the hacker – whether seeking simply to create a nuisance or leaning more toward the serious threat of corporate espionage – has never been easier. Given the ready availability of reusable code on which to build ever-more dangerous applications, networks have been bombarded with a growing tide of both new and known malicious software (malware).

In 2005, Nortel's anti-virus team isolated and submitted more than 400 malware samples to Symantec Corp., a worldwide leader in the development of security solutions and the supplier of the anti-virus tools used within Nortel. More than 50 percent of these samples had not been reported through any other means, highlighting not only the world-class expertise of Nortel's security team, but also the company's commitment to working with others in the security industry to improve overall network security.

Typical avenues of attack include infected email and instant messaging, extending now even to cellphone text messaging and other forms of personal communication, where the malware is activated when the receiver clicks on a link or opens an attachment. As well, the interconnectedness of the Internet has made it easier for attackers to scan for vulnerable hosts that have not kept their operating systems up to date with security patches. Indeed, recent experiments have shown that an unprotected Windows PC connected to the network will become infected with all manner of viruses, trojans, adware, and spyware within just a few minutes of being plugged in.

For the most part, the majority of malware has presented itself as a nuisance – in the form of SPAM-bots, denial of service attacks, and ego-building self-replicating worms. In the past year, the industry has even witnessed "virus wars" between competing groups of malware builders, each maligning the other's code, and seeking to gain the greatest public recognition possible for their misdeeds. Unfortunately, with the rapid proliferation of various strains of malicious code in the fertile breeding ground of the Internet underground, malware attacks have begun to pose a far more serious threat.

Trojans, for instance, are typically easily constructed by piecing together various code fragments obtained freely on the Internet, but their intent is far more sinister than their relatively benign constituents. As these readily available building blocks proliferate, we can expect increasingly greater threats to privacy, intellectual property, and general information security. What's more, as defenses against these forms of malware evolve and improve, so too does the sophistication of the attack methods. Over the past few months, for example, the industry has seen malware that can write itself to EPROM chips on a system's video card, meaning the malware will survive even a complete erasure and reinstallation of the operating system.

Making the situation even more worrisome, such attacks on corporate assets are not confined to cyberspace. For instance, hackers posing as janitors recently attempted to steal £220 million from a bank in London, England, by installing devices on computer keyboards and obtaining the log-in information they needed to break into the bank's network. In another instance, hackers posing as street people camped outside a corporation after hours and searched

Harnessing the power of network convergence to improve corporate security: Some scenarios

The emergence of next-generation converged networks – including optical-enabled packet networks, wireless technology, multimedia communications, and digital video – is making it possible, for the first time, to integrate IT and physical security systems in real-time security solutions that work across an organization to safeguard its business assets and operations.

While IT/physical security convergence is still at the discussion and planning stage for the industry as a whole, one can think of concrete applications where such integration would deliver value.

For example, with converged network technology, it would be possible to network a collection of closed-circuit monitors to a central network operations center (located, say, in New York), where security personnel can view real-time video footage sent over the IP packet network from cameras located at the company's remote sites.

In one possible scenario using this technology, a person walks through a building (say, in Dallas), sets a briefcase down, and walks away. Using artificial intelligence capabilities to analyze the real-time video feed, the networked security system quickly locks onto the briefcase, tags it as "abandoned," and alerts the security team. A security guard, on duty at the New York operations center, receives the notification on a laptop computer or PDA and quickly calls security personnel at the Dallas location who can investigate the suspicious object. Unlike video surveillance systems that simply archive the footage, this real-time surveillance system – built to operate over the IP network and existing IT structure – would enable a quicker response to the possible threat, while reducing costs and raising productivity by centralizing management operations.

In another possible application, building access solutions could be tied to a corporate identity directory and smart-card technology to verify a person's identity and track his or her physical location and movement within a building.

Consider, for example, a case where a password is used illegally. Intuitively, one would report the problem to the IT department which, acting alone, may not be able to prevent break-ins to buildings simply by changing the password. However, by working in tandem with a converged security system, the IT and physical security departments could monitor who was in the building when the break-in occurred to resolve the breach much sooner or, if measures had already been in place, could have proactively prevented it from occurring at all.

Similarly, a person's use of IT systems (such as a PC) could be correlated with the building access logs. If someone is shown to be logged in to a system, but not physically in the building, a security exception flag would be raised, and if no reasonable explanation existed, the incident would be investigated.

In another instance, security personnel (or police or fire departments) could be issued cellphones running a soft client that allows them to access network directories and other information and communicate across the network to remote locations or operations centers – unlike walkie-talkie-type systems where the range is limited and communication is restricted to other personnel carrying walkie-talkies. This type of integrated security system, empowered by the network and IT infrastructures, could prove invaluable in coordinating response and recovery efforts during emergency situations.

Moving forward, many other applications will emerge as network convergence technologies continue to mature, and as organizations and vendors such as Nortel explore a whole range of possible opportunities for integrating IT and physical security systems in a more powerful and efficient security network. To this end, an understanding of how networks and technologies work together – a defining strength of Nortel – is critical. With its enormous breadth and depth of networking knowledge across all types of networks – carrier and enterprise, wireless and wireline – Nortel is a leader in delivering solutions that leverage the capabilities of a converged network infrastructure. Moreover, through its Global Services organization, Nortel can aid customers in analyzing their security requirements and designing an end-to-end solution for securing their entire network.

through dumpsters for computer printouts, employee notes, and other material that revealed employee passwords, which were then used to attack that company's network. And high-capacity USB memory sticks have become commonplace, enabling hackers to download gigabytes of information onto a pocket-sized device. As a result, it is quite common for financial institutions to fill USB ports on company-issued laptops and PCs with epoxy to thwart the practice.

Although organizations naturally tend to think of technological solutions first when combating these threats, effective security is not solely dependent on technology. Addressing the growing corporate-wide security risks includes a strong organizational element focused across three critical and equally important areas – process, technology, and people.

Process

A significant component of effective security lies in proper process definition, adoption, and implementation – typically based on best practices, either generated internally or industry-wide.

Figure 1. Sample security process mapping

	Corporate security	Facilities	Information services	Human resources	Legal	Internal audit	Corporate communications	Line management
Compliance								
Policies and procedures	●	■	■	■	■	■	■	■
Standards and guidelines	■	●	●	●	●	■	–	●
Intellectual Property reviews	●	–	■	–	■	■	–	■
Physical security audits	●	■	■	–	–	●	–	■
Due diligence								
Investigations	●	■	■	■	■	■	■	■
Education and awareness	●	●	●	■	■	■	■	●
Consultation and advice	●	■	●	■	■	■	–	■
Security operations								
Key employee protection	●	■	–	■	■	–	–	■
Crisis management (security-related)	●	■	■	■	■	■	■	■
Security guard management	■	●	–	–	–	–	–	■
Location emergency response	■	●	■	■	■	–	■	■
Law enforcement liaison	●	■	–	–	–	–	■	■

● Process owner
■ Cross-functional team member
– Not applicable

An important first step in implementing an integrated corporate-wide security plan is to examine all the security-related elements in an organization and their relationships to one other. The process begins by mapping out all of the security elements within a particular organization (shown in this example on the vertical axis of the chart), then assigning process ownership and identifying cross-functional relationships across each of the corporate departments (the horizontal axis of the chart). This exercise helps identify any gaps in security coverage or confusion over roles. With this mapping, more detailed documents can be drawn up that clearly spell out the specific and cross-functional responsibilities of various departments. This process helps ensure that employees understand their responsibilities in the context of the entire security process, and also that key facilitators in an emergency situation immediately recognize their responsibilities and respond in a coordinated fashion.

Secure processes are secure by design, not by accident. Specious reasoning has no place in the secure enterprise, as illustrated in the exchange between Homer and Lisa Simpson below:

Homer: Not a bear in sight. The Bear Patrol must be working like a charm!
Lisa: That's specious reasoning, Dad.
Homer: Thank you, dear.
Lisa: By your logic I could claim that this rock keeps tigers away.
Homer: Oh, how does it work?
Lisa: It doesn't work.
Homer: Uh-huh.
Lisa: It's just a stupid rock.
Homer: Uh-huh.
Lisa: But I don't see any tigers around, do you?
[Homer thinks of this, then pulls out some money]
Homer: Lisa, I want to buy your rock.

* Dialog TM and © FOX

Just because a security threat has not yet compromised the enterprise environment does not mean that "the security patrol" is working. The real – and sometimes very difficult – challenge for the CIO/CSO is to balance the overall security risk against the investment required to mitigate the many vulnerabilities that exist in the enterprise. This task is complicated by the existence of many separate legacy systems that must be updated and integrated into an overall security plan.

To develop an integrated security plan, an organization needs to review key business processes, assess and define information access and workflow needs, and formulate a detailed plan and roadmap, including budgets, timelines, and performance benchmarks. The processes reviewed should include all security services required by an organization's particular business model (for example, background investigations, due diligence, and executive protection).

A critical first step is to identify the security threats facing various parts of the organization, and then define and assign ownership of each responsibility (Figure 1).

The payoff from initiating and instituting a comprehensive, well-thought-out process to safeguard enterprise assets – whether people, information, intellectual property, physical facilities, or business operations – can be significant, as Nortel itself found out. In August of 2003, when more than 50 million people and enterprises were affected by a major electrical power blackout across the northeastern United States and Canada, Nortel maintained close to 85 percent of its productivity during the nine-day state of emergency, because it had invested some 12 months in constructing an integrated business continuity plan prior to the blackout. During the outage, clearly designated responsibilities and directives allowed employees to reroute incoming telephone calls to centers outside the affected area. This in-place process allowed the company to communicate with employees, customers, and suppliers to ensure minimal effect on its business operations.

Technology

As part of the overall security plan, existing physical and IT infrastructures need to be assessed to determine what technologies and components may be required to meet not only current needs but also future expansion plans. This assessment should include security system design, implementation, and integration, as well as ongoing measurements and evaluation of existing expenditures to ensure effectiveness. The complexity in building a successful defense lies in the requirement to allow legitimate traffic to pass the security defenses unimpeded, while reliably blocking malicious traffic, which in many cases will be masquerading as legitimate (Figure 2).

To be sure, technology plays an extremely important role in keeping the enterprise safe. According to the Cooperative Association for Internet Data Analysis (CAIDA), the spread of the SQL Slammer worm on January 25, 2003 infected 74,855 hosts worldwide in the first 30 minutes of its attack. While this may seem to be a relatively small number of hosts globally, the resulting network traffic caused by the worm's attempts to propagate itself around the Internet almost instantly brought down large portions of the Internet and many corporate networks. This rapid outbreak pattern happens at machine speeds – far too fast for human intervention and response – and therefore requires technological solutions that can either detect and prevent attacks, or respond in near-real-time to mitigate the effects of an outbreak.

As companies increase the number of business partners with whom they work, and simultaneously increase the mobility and geographic distribution of trusted workers entering the virtual enterprise, the traditional "security perimeter" is all but disappearing. In essence, the enterprise simply "becomes" the Internet – a widely interconnected mixture of network elements, hosts, and applications, each of which must be sufficiently hardened to withstand all manner of external attack, while preserving the ability to offer acceptable, secure performance to meet its intended purpose.

As security attacks grow and evolve over time, defenses must adjust to the ever-changing threat landscape. The sophistication of modern attacks against all levels of the business environment is the key reason that a layered defense strategy is essential: it creates a much higher likelihood that both known and new threats will be rendered harmless by one or more layers of the defensive strategy.

Through its Network Security Architecture (page 7), Nortel has adopted this layered defense approach as a basis for delivering leading security products, applications, and solutions for enterprises (see article on page 28), VoIP/multimedia communications for enterprises and carriers (see article on page 37), and mobile broadband communications for wireless operators (see article on page 46).

The convergence of technologies and business processes onto an IP-based infrastructure is also making it possible to change the way physical security

and many of its related processes are handled. For example, the digital security cameras and alarm systems protecting buildings and other assets can be monitored remotely – sometimes from hundreds, and in a few cases, thousands of kilometers away – using LANs and WANs to carry the data to a central operations center, in much the same way as other business processes make use of these networks. In fact, security industry leaders believe many roles currently fulfilled by security personnel can be automated in this fashion, saving organizations considerable time and money. In addition, manufacturers of card-key entry systems are now beginning to integrate building and network access controls (such as passwords, fingerprints, or other biometrics) on one card – enabling physical and IT security information across many locations to be centrally coordinated using LANs and WANs in the converged IP infrastructure.

People

While process and technology are critical for building an effective security structure, it is people who form the first and last line of defense. It makes little sense to invest in sophisticated access controls, passwords, or related technology if employees can be easily duped into allowing others to circumvent or have access to these controls. As more and more sophisticated exploits appear (social engineering, phishing, trojans, etc.), technological defenses and process controls can provide only a certain level of protection.

In fact, the number of attacks by social engineers – "people hackers" who con employees into sharing all manner of competitive and other confidential information – continues to rise, driven by those seeking unfair advantage in today's fiercely competitive markets. Typically, they ask for seemingly unimportant information that may be widely distributed internally, but is not meant to be released externally (such as organization charts, telephone directories or email address books, log-on information, "sample" purchase order forms, and product information, including launch plans). Also on the rise is phishing, where an apparently legitimate email tricks people into giving out identification numbers, passwords, and

Figure 2. Defending against network threats

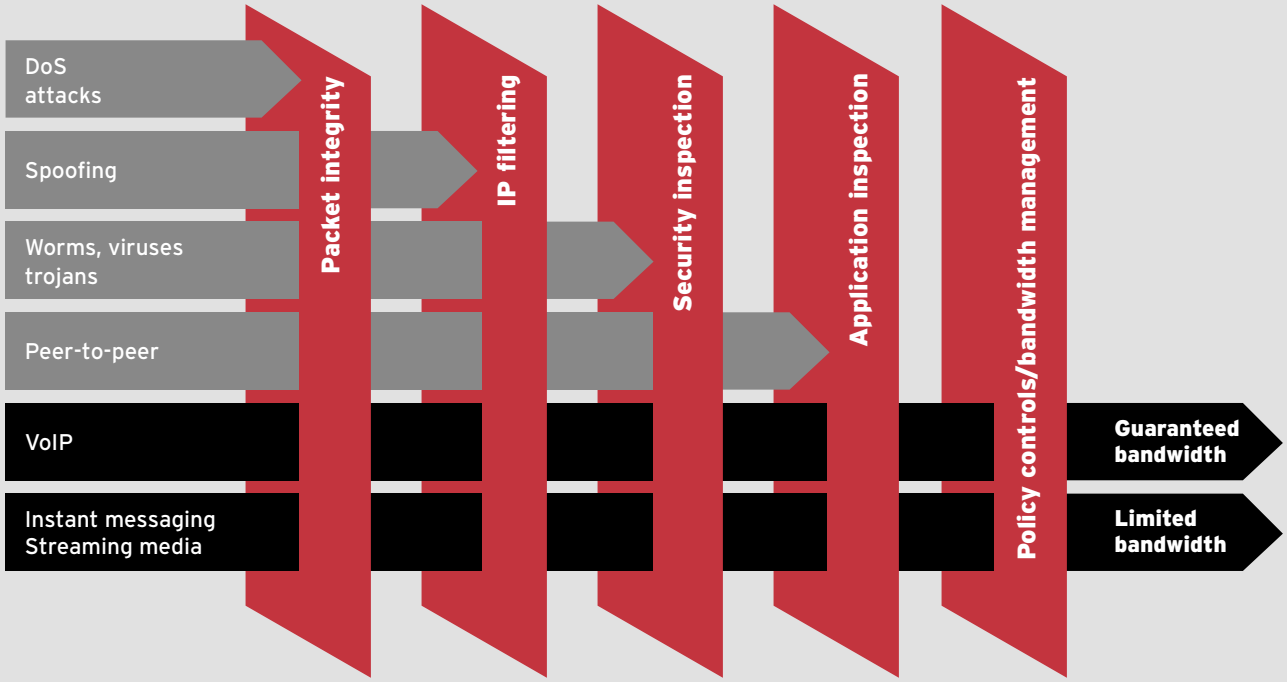

To protect an enterprise's core network, systems, and applications, various defense mechanisms can be employed to allow desirable traffic to propagate as defined by the enterprise's security policy, while blocking undesirable traffic.

In the figure, the vertical slices represent various security mechanisms – testing packet integrity, filtering IP traffic, inspecting packet contents both for security (testing for known and suspected malicious content) and for application "fitness for use," and managing policy controls and bandwidth requirements.

In the above example, the different types of attacks – such as denial of service (DoS) attacks, spoofing, worms/viruses/trojans, and undesired peer-to-peer traffic – are blocked by different layers of the defense, while the "good" traffic – such as VoIP, streaming media, and instant messaging – pass through unimpeded (although security policies are applied to guarantee VoIP bandwidth or limit streaming bandwidth, as required).

other pieces of personal information that can be used to compromise security measures and gain entry into corporate networks.

According to Kevin Mitnick (a reformed hacker turned security consultant who was caught and jailed for phone fraud and computer intrusion), humans are the weakest link in any security system, a statement illustrated with the following anecdote from an article in *CSO Magazine* (CSO Online, October 2002):

The intrepid social engineer calls up the network operations center of a cell phone company during a snowstorm. After befriending the operators, he asks them: "I left my SecureID card on my desk. Will you fetch it for me?" he asks. Of course, the network operators are too busy to do that, so they do the next best thing: They read off the ever-changing code on their own token, allowing the hacker to break in and steal the company's source code. In this example, the caller is able to "prove" his identity by telling the network operators his office number, the department where he worked and the name of his supervisor – all information that the attacker had gleaned from previous phone calls to the company. Mitnick's message is that organizations need to treat phone lists, org charts, technical procedure manuals, and other information as highly confidential in order to protect themselves from social engineering attacks.

This anecdote illustrates the severe vulnerability posed by a well-intentioned but security-challenged staff member – the end result, a serious security breach, despite having in place well-planned and implemented technological defenses, such as two-factor authentication (something you have, such as a secureID token, combined with something you know, such as a unique personal identification number).

The only defense against these forms of attack is a properly trained and security-aware workforce that is on guard for any deviations from policy or process, treats even seemingly benign exceptions as potentially suspicious, views all incoming communications (email, instant messages, or web-based information access) with a skeptical eye, and quickly reports any incidents that do occur.

Indeed, research shows that up to 70 percent of incidents are reported by employees, rather than discovered by any other systematic means, such as intrusion detection or monitoring systems.

As real-life examples have shown, failure to provide adequate security can result in significant business loss, breach of regulatory requirements (e.g., export control and/or privacy legislation), loss of brand reputation and shareholder value, and debilitating business disruption. This reality demands effective processes, secure-by-design technologies, and a vigilant workforce in a converged IT/physical/corporate security system to stay one step ahead in the "spy vs. spy" showdown to secure the enterprise. ■

Scott Hurd *is Information Services Director, Nortel Global Security.*
Tim Williams *is Vice President, Corporate and Systems Security.*

Baseline security: Leading the industry toward a standard foundation for infrastructure security

by Mike Lee

Nortel was one of the first vendors to develop a consistent set of baseline security requirements and best practices. First used in its own product portfolios, these baseline requirements today form the foundation for the security baseline standards being defined by both national and international standards bodies, including the Alliance for Telecommunications Industry Solutions (ATIS) and the International Telecommunication Union (ITU-T). These standards are key to the creation of a more resilient, secure, and trusted network environment.

Several years ago, when service providers and enterprises began converging their networks onto single IP-based infrastructures, Nortel – from its vantage point as a pioneer and leader in building secure, reliable networks – foresaw the need for an industry-wide set of baseline security requirements. Anticipating the coming challenges for network security, Nortel set out to develop a common and consistent set of security features that would form the "must-have" suite of protective measures and best practices needed to provide a solid security foundation for future next-generation, converged networks.

Nortel began by developing a common baseline of security requirements for its own products, and then championed the adoption of these requirements by the industry, through industry forums, customer engagements, and standards development organizations. While the standards focus has been on public networks, the principles and techniques can also be applied to private or enterprise networks.

Nortel took the view that sharing its basic infrastructure security strategy with the industry made good business sense, for several reasons.

For one, a standardized security baseline addresses the network complexity that was developing as operators and enterprises responded to growing network security concerns with differing but related requirements. To illustrate, one request for proposal (RFP) issued not long ago by a major North American wireline service provider listed more than 2,000 security requirements – a significant change from what previously might have been a handful of items. Moreover, different operators were adopting different approaches to addressing various security issues. While similar in intent, their security choices differ vastly in implementation and have introduced a host of challenges for both providers and vendors.

Second, providers face significant cost and deployment challenges in having to integrate inconsistent and often incompatible security feature sets from multiple vendors, which ironically exposes customer networks to even greater security vulnerabilities. For example, if one piece of equipment uses IPsec (Internet Protocol Security) exclusively as its encryption protocol and another piece of equipment relies on the Transport Layer Security (TLS) protocol, the two boxes won't be able to "talk," and therefore would be unable to provide seamless encryption – even though the underlying encryption technologies used by both protocols are equally as strong. Similarly, two different authentication technologies, such as RADIUS and Kerberos, while equally good, are incompatible.

Third, a standardized security baseline would address the challenges that vendors face in having to "cover the waterfront" and support all customer requirements, which is difficult and costly to do because it requires manufacturers to develop a super-set of security technologies across all products, with the added burden of keeping up with the rapid pace of new attacks and security bulletins. Additionally, a security baseline would facilitate the generation of inter-carrier interconnection security agreements.

Nortel knew that a common set of security specifications would enable service providers to more easily procure and build secure infrastructures comprising multiple vendor platforms, while enabling vendors to lower the complexity and costs of development.

Different traffic plane requirements

As a starting point, Nortel defined a three-pronged approach to developing baseline security standards. First, it would tackle the new security needs

of network management (management plane issues), followed by the needs of signaling (signaling plane issues), and then of user traffic (media plane issues).

Indeed, as networks have converged on IP-centric infrastructures, these three planes are no longer separated physically, as they were in traditional telecommunications networks.

In the past, protecting the overall network from intrusion by hackers and other threats was relatively straightforward, because purpose-specific traffic was separated onto different and isolated elements in the network. Operations, administration, and maintenance (OAM) traffic, for instance, traveled over a separate management plane in the network on point-to-point connections that could be accessed only by legitimate operators in the customers' network operations centers (NOCs). Similarly, signaling traffic took a separate communications path through the network via distinct signaling network elements, such as CCS7 (Common Channel Signaling System 7) elements. The general public, then, had access only to user traffic and was unable to penetrate either the management or signaling planes of the network (Figure 1).

On the whole, past telecommunications networks were considered relatively safe from widespread malicious activity. When the comparatively rare intrusion or malicious act did occur, it was often in the form of an error on the part of operators, or it was a typical type of fraudulent activity – such as an attempt to change service profiles or alter billing records – and was easily detected.

By contrast, in next-generation converged networks, all packet types are sent over common network elements. Because different traffic types are no longer separated, such mechanisms as encryption are now needed to virtually separate and protect the traffic, in order to provide even the most basic security, which involves ensuring data confidentiality, integrity, authentication, authorization, and auditability – known in the industry by the mnemonic 'CIAAA.'

These CIAAA requirements were the starting point for the Nortel team, as it defined the "must-have" security needs for all three traffic planes.

In doing so, Nortel has adhered to several key security principles. Specifically, that:

- customers expect the network elements and products to be inherently secure at manufacture;
- a layered defense (also known as defense in depth) security process should be applied to ensure network safety,

Figure 1. Convergence of network traffic types

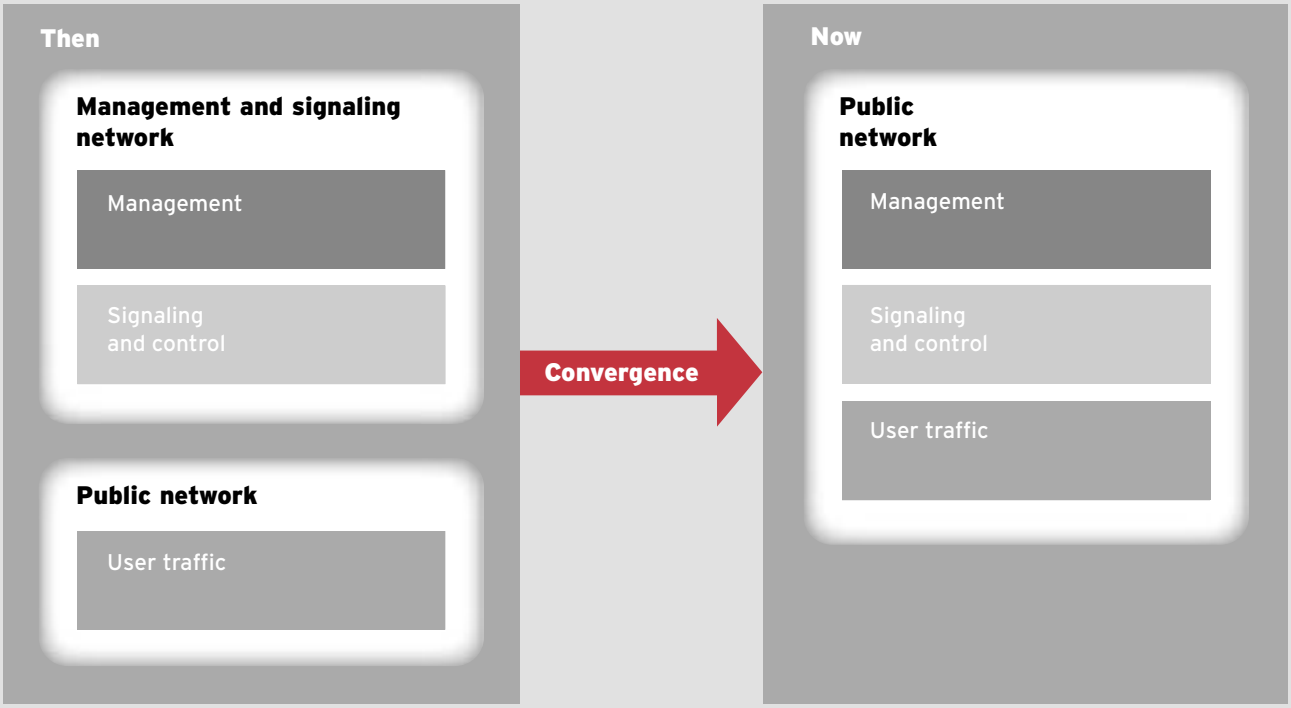

In telecommunications networks of the past, public traffic and network management and signaling traffic were sent on separate networks. In this way, the management and signaling networks were relatively easy to secure and isolate from malicious activity. Today, by contrast, management, signaling, and user traffic are sent on the same "pipe," and threats that were once confined to user traffic are now threats to all traffic planes. The Nortel-defined set of baseline security requirements provides the higher levels of security now required for all planes.

through multiple layered security mechanisms with no single point of failure;
• security capabilities should be integrated into network products at the design stage;
• a common infrastructure security strategy is needed across all product lines to ensure interoperability; and
• standard protocols should be used to help ensure interoperability with third-party products.

Securing the management plane

Guided by these principles, Nortel targeted the management plane as the first priority, since it was here that attackers could most easily cause large-scale network disruption. For instance, an attacker – whether a disgruntled insider or a hacker from the general public – could reconfigure network equipment, cause denial of service (DoS) and network outages, destroy critical data, or perpetrate fraud through stealing or manipulating service or billing records.

Nortel's CTO Office led a cross-corporate security team with representatives across the product groups. This team examined the requirements and potential vulnerabilities of the management plane, sifted through the merits of the many existing technologies, solutions, and best practices, and identified those that would best meet the needs of an end-to-end network.

Indeed, an unsecured management plane has many vulnerabilities that can be exploited. These vulnerabilities include:
• use of unsecured protocols for network management, such as SNMPv1&2, FTP, Telnet, TFTP, and HTTP;
• use of weak, locally managed static passwords for operator authentication (an example of a weak password is "admin" for an administration password, or a three- or four-letter dictionary word, such as "cat" or "dog");
• lack of access control (or authorization) mechanisms to define network access permissions for different levels of administrators;
• security information, such as secret keys and passwords, that are transmitted in the clear or stored in plaintext;
• backdoor programs left in network management systems; and
• unnecessary functions, such as a Telnet administration function that has been left on, running on the operating systems of management devices and presenting targets for attack.

To protect the network against these vulnerabilities, the team identified several key requirements for management plane security, several of which are considered mandatory – including operating system hardening, virus-free software, cryptographic protection, secure remote access, operator authentication and access control, standardized security logs, and vulnerability assessment – as well as two best-practice recommendations (intrusion prevention/detection and firewall protection) that are considered optional, depending on individual customer needs (Figure 2). While many of these requirements need to be built in at the design stage, some – such as operating system hardening and anti-virus protection – also require ongoing monitoring to ensure that the latest patches are incorporated.

Operating system hardening: Computers and network elements are vulnerable to any number of attacks, including backdoor programs, password grabber and cracking tools, exploitation of defects in operating system services, and DoS attacks. Once a system has been compromised, an intruder can modify or destroy information, disclose sensitive information, install malicious code to gather information, or use the compromised server to attack other systems.

A key measure to counter these attacks is to harden commercial operating systems using procedures such as turning off unused services, ensuring removal of default passwords, and ensuring that security patches are up to date. Such services are essentially sound practices followed during installation and configuration of operating systems. The management plane security baseline requires that all operating systems used for management purposes undergo operating system hardening.

Virus-free software: "Virus" is a term used to categorize several types of malicious software programs, or malware, including viruses, worms, and trojan horse programs. Securing the management plane, then, requires that all network software be scanned using anti-virus software to ensure that it is virus-free to the maximum reasonable extent possible before installation. All software used in Nortel products, whether developed in-house or sourced from a third-party, is checked for viruses using a Nortel-developed virus-detection process before being incorporated into product.

Cryptographic protection: Encryption of data provides a high degree of protection against malicious insiders, while allowing legitimate operators access via encryption keys. The management plane security baseline requires encryption and authentication for all management traffic to ensure data confidentiality and integrity. Since different customers prefer different security protocols for cryptographic services, the baseline standard accommodates several, including IPsec, Secure Shell (SSH), Secure Socket Layer/Transport Layer Security (SSL/TLS), and Simple Network Management Protocol Version 3 (SNMPv3), all of which can provide data integrity and confidentiality for network management traffic.

Nortel includes all of these security protocols across its product lines, and the CTO team provides consultation to individual product groups on a network or product basis, to not only help guide implementation of the protocol that will work best for a specific customer but also to ensure end-to-end encryption solutions.

Secure remote access for operators: A particular security concern in the management plane is to ensure that network management data and processes are accessible over the public Internet only by legitimate operators, who often need to administer the network from a remote location. In this case, strong security

is needed to authenticate these remote operators, as well as protect the confidentiality and integrity of all data both to and from them.

To provide this level of security, the management plane baseline requires the use of secure virtual private networks (VPNs) based on IPsec. IPsec VPNs, such as Nortel's VPN Router portfolio (formerly known as Contivity), provide secure encrypted tunnels for data traffic to and from all remote operators.

Centralized network operator authentication/access control: "Authentication" refers to proof of identity of the party accessing the network, and Nortel's customers require strong authentication of all network operators. Moreover, once an operator has been allowed onto the network, access control policies limit the network resources that network operators can administer.

The management plane security baseline requires a centralized authentication/authorization system with enforcement of strong passwords for all Nortel products. To achieve this level of protection, Nortel uses a system based on RADIUS/LDAP (Lightweight Directory Access Protocol) to automate centralized authentication within Nortel solutions. The use of Pluggable Authentication Mechanism (PAM) is also recommended to allow other customer-specified authentication mechanisms, such as Kerberos, to be incorporated more easily.

Security audit logs maintain an audit trail of operator activities and events, and provide a basis for accountability, reconstruction of security incidents, problem analysis, and long-term trend analysis. (The raw data collected is called the "audit log," and the verifiable path of events through the audit logs is referred to as the "audit trail.")

Audit log information helps to identify the root cause of a security problem and prevent future incidents. For instance, audit logs can be used to reconstruct the sequence of events that led up to a problem, such as an intruder gaining unauthorized access to system resources, or a system malfunction caused by an incorrect configuration or faulty implementation.

To be effective, logs must contain enough security information to conduct an after-the-fact investigation or analysis of security incidents. As well, to be useful across end-to-end network solutions, logs need to be in the same format.

Nortel's management plane baseline calls for a common log format to be used across its product portfolio, along with a common comprehensive log content specification detailing the security events that need to be logged (e.g., administrator log-in and configuration changes). Syslog is the recommended logging mechanism for storage and transfer of logs, because it is a common mechanism compatible with all third-party log analyzer systems.

Vulnerability assessment of products is used to discover security weaknesses and areas of risk before a product is deployed. While product verification testing focuses on ensuring that systems pass defined test scenarios, vulnerability testing is designed to try to make the system fail by circumventing security controls, capturing confidential data, obtaining unauthorized access, and performing other attacks.

The management plane baseline requires that vulnerability assessment be

Figure 2. Baseline security at a glance

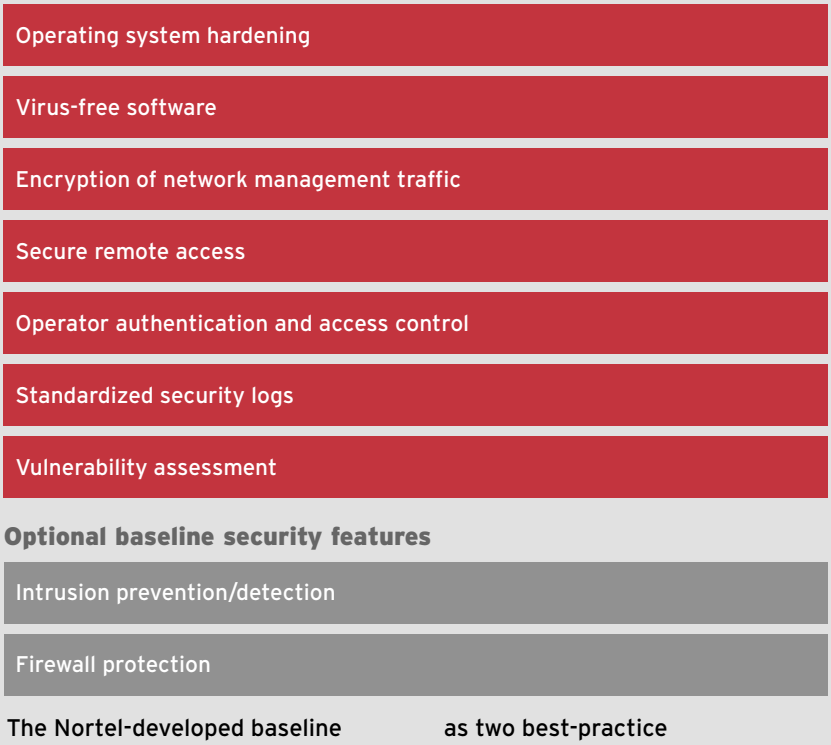

The Nortel-developed baseline security requirements for the management plane have been used to form the foundation for industry-wide security baseline standards – a common and consistent set of security requirements that form the "must-have" suite of protective measures for basic security. These requirements include several that are considered mandatory, as well as two best-practice recommendations (intrusion prevention/detection and firewall protection) that are considered optional, depending on individual customer needs. Nortel has also defined requirements for the signaling and media planes, and is sharing these recommendations and best practices with the industry at large through key standards bodies.

conducted routinely in order to identify and better understand threats and vulnerabilities, to determine an acceptable level of risk, and to mitigate identified issues. In this effort, Nortel created a comprehensive company-wide program that includes training for all product groups on how to perform this assessment and resolve issues.

In addition to these seven measures that Nortel defined as mandatory, the team also identified two best practices – intrusion prevention/detection and firewall protection – that, while considered optional with respect to the baseline requirements defined by Nortel and being formally standardized, are considered important by many customers. **Intrusion prevention/detection systems (IPS/IDS)** can be incorporated in a network solution to provide even stronger defense. For example, Nortel's Threat Protection System (TPS) (page 35) can be used to warn network administrators of the possibility of a security incident, such as a compromised server or DoS attack.

IPS/IDS can be broadly categorized according to the following criteria:

- *Incident prevention/detection timeframe:* real-time or off-line, depending on whether system logs and network traffic are analyzed as events take place or in batch mode during off hours;
- *Type of installation:* network-based or host-based. A network-based IPS/IDS typically involves multiple monitors (often pre-configured appliances) installed at choke points in the network (where all traffic between two points can be monitored). A host-based IPS/IDS requires that software that monitors network connections and user activity on servers that need to be protected be installed directly on those servers; and
- *Type of reaction to incidents:* whether the IPS/IDS actively intervenes to head off attacks (such as by modifying firewall rules or router filters), or simply notifies staff or other network systems of

Table. Summary of baseline requirements to mitigate against common infrastructure vulnerabilities

Network security threat	Nortel-recommended mitigation measures
Masquerade	• Use of RADIUS for centralized password management ensures strong authentication of operators. • Encryption of management traffic prevents snooping of administrator passwords.
Denial of service (DoS)	• Addition of firewalls provides first stage of defense-in-depth strategy to prevent DoS attacks. • Nortel standard security logs provide strong traceability for forensic analysis.
Hacking	• Use of RADIUS for centralized password management ensures strong authentication of operators. • Encryption of management traffic prevents control or modification of network resources by hackers.
Sabotage	• Use of RADIUS for centralized password management ensures only legitimate operators have access to the system. • Encryption of management traffic prevents control or modification of network resources by attackers. • Nortel standard security logs provide strong traceability for forensic analysis.
Intrusion	• Use of RADIUS for centralized password management ensures only legitimate operators have access to the system. • Encrypted management traffic ensures attackers cannot control network elements. • Nortel standard security logs enable intrusion analysis and incident recovery. • Use of intrusion prevention/detection systems.
Backdoor programs	• Operating system hardening ensures backdoor programs are removed or disabled. • Encryption of management traffic and RADIUS authentication limit access to equipment to only legitimate operators. • Nortel standard security logs provide strong traceability for forensic analysis.
Disgruntled employees	• Encryption of management traffic allows only authorized insiders to view network data and/or modify network element operation. • RADIUS authentication ensures that only legitimate operators can access/modify equipment. • Nortel standard security logs provide strong traceability for forensic analysis. • Firewall placement provides segmentation between network zones and limits scope of any attack.

Table continued

Network security threat	Nortel-recommended mitigation measures
Snooping	• Encryption of management traffic prevents snooping. • Nortel standard security logs provide strong traceability. • Use of firewalls limits scope of any attack.
Modification of data	• Encryption of management traffic prevents modification of management data by attackers. • Nortel standard security logs provide strong traceability. • Use of firewalls limits scope of any attack.
Proliferation of unsecured protocols (unsecured protocols include ICMP, Telnet, SNMPv1&2, DHCP, TFTP, NTP, DNS, and HTTP)	• Replacement of unsecured protocols with secure, encrypted protocols. • Telnet, FTP are replaced by SSH or IPsec. • IPsec used to encapsulate other unsecured protocols. • Use of TLS for HTTP traffic. • Use of SNMPv3 for SNMP traffic.
Use of weak, locally managed, static passwords	• Use of centralized RADIUS server enforces strong, centrally managed passwords, as per Nortel standard password guidelines.
Unprotected security information (e.g., unencrypted password files, passwords sent in the clear, firewall rule sets, and cryptographic keys)	• RADIUS protocol transmits passwords across network and stores passwords in a hashed format. • Encryption of management traffic ensures critical data is sent securely across network.
Non-hardened network elements and operating systems	• Operating system hardening.
Management ports and interfaces unnecessarily exposed to the public network	• Encryption of management traffic. • Firewall segmentation of network.
Industrial espionage	• Encryption of management traffic and strong authentication of operators via RADIUS prevents unauthorized access to equipment.

the problem.

For proper intrusion prevention/detection measures on the network management plane, Nortel recommends that both network and host-based IPS/IDS be implemented in the network solution.

Firewall protection: A firewall is a set of safeguards that enforce a security policy between two networks. Firewalls are sometimes called "policy enforcement points" that implement an organization's corporate security policy, which is typically expressed as a rule set in the configuration language of the particular firewall.

Traditionally, firewalls were used to isolate private networks (intranets) from public networks (the Internet). Nortel recommends the use of firewalls in all network solutions in order to segment the management, signaling, and user traffic into different security domains. In this role, the firewall controls the type of traffic that transits the boundary between different security domains. Depending on the type of firewall (application versus packet filtering), this control can also be extended to include filtering of the application content of the data flow. Typically, firewall placement, type, and filtering rules are designed for a particular network implementation.

Management plane baseline standardization

These nine network safeguards constitute Nortel's management plane baseline security requirements, which were formally documented in the company's systems requirements document (SRD) and are being implemented across all product portfolios. (The table on page 24 summarizes the recommended measures to mitigate against key network infrastructure vulnerabilities.)

Nortel then initiated an activity at the U.S. National Security Telecommunications Advisory

Committee (NSTAC) to standardize these requirements and drive them into the industry at large.

Early activity at NSTAC was so successful that Nortel was asked to act as a technical editor on a new U.S. standard for management plane security being implemented through the U.S. Alliance for Telecommunications Industry Solutions (ATIS). Working with a large government and industry team at ATIS – a team that included members from the U.S. Department of Defense, as well as from Verizon, Sprint, MCI, BT, and others – Nortel drove the baseline requirements into the American National Standards Institute (ANSI) T1.276-2003 *A Baseline of Security Requirements for the Telecommunications Industry.*

Following this effort, Nortel then worked to drive the management plane baseline requirements into international standards bodies, including the 3GPP (Third Generation Partnership Project) TS 32.371 Security Concepts and Requirements, and the International Telecommunication Union (ITU-T) M.3016.x series of standards. Nortel was also the editor of the Security Requirements for NGN Release 1, produced within the ITU-T SG13 Focus Group on NGN.

In addition, Nortel has driven the management plane baseline best practices into the U.S. Network Reliability and Interoperability Council (NRIC), where they have been accepted as formal NRIC best-practice recommendations.

Securing the signaling/control plane

With standardization of management plane security well under way, Nortel then focused its attention on determining the fundamental requirements for securing communication between signaling elements in multimedia networks – e.g., those that use H.323 and Session Initiation Protocol (SIP) protocols.

Following an approach similar to the one it used for the management plane, the Nortel team identified several key security requirements and associated technologies to address the requirements for the signaling/control plane. These include:

Data confidentiality and integrity: The signaling/control plane baseline requires the use of IPsec or TLS protocols to provide data confidentiality and integrity – that is, to protect all SIP signaling messages from unauthorized reception and modification.

Authentication: Authentication verifies the identities of those involved in a communications exchange. The signaling/control plane baseline requires the enforcement of bidirectional authentication based on X.509 certificates through IPsec or TLS protocols for all machine-to-machine SIP signaling exchanges. For SIP user agents (SIP soft clients, SIP phones, and SIP integrated access devices), the signaling/control plane baseline requires authentication based on HTTP digest over a secure protocol, and recommends the use of X.509 certificates.

Access control (authorization): Access control is based on lists of known IP addresses with which a network element or server will allow communication. The signaling/control plane baseline recommends the use of access control lists for SIP client and server applications, enforced by packet filtering software.

Audit logs: Security audit logs maintain an audit trail of network element and server events, and are used to identify causes of security problems, prevent future incidents, and provide information for evidence. The signaling/control plane baseline provides a list of SIP signaling and control events to be logged.

Signaling/control plane baseline standardization

As part of its overall baseline security strategy, Nortel is taking these signaling/control plane specifications to the industry for standardization. The first focus for standardization of signaling security standards was within the ATIS Packet Technologies and Systems Committee (PTSC). Within PTSC, Nortel holds a vice-chair position and also chairs the PTSC Security Subcommittee where this standardization is occurring.

The PTSC committee is establishing a family of five signaling and control plane security standards, with Nortel acting as a technical editor for two of these. This activity is under way and the first of the five proposed standards – the Generic Signaling and Control Plane Security Requirements for Multimedia Networks – is currently being validated by ATIS members. The expectation is that these ATIS-produced standards will be fed into the ITU-T, for adoption into ITU-T Recommendations.

Securing the media plane

Nortel as well as others in the industry are also working to define baseline requirements for the media plane, which carries user traffic. When defined, the media plane security baseline will focus on security requirements for real-time user traffic on multimedia networks that use H.323 and SIP protocols.

Before these requirements can be formalized, however, several challenges need to be addressed. Chief among these is the current industry discussion about how much security is actually needed on the media plane for real-time voice and multimedia, which traditionally in public applications was not secured except in special environments, such as military applications. This debate centers on whether user traffic encryption is necessary: because user traffic forms the bulk of all network traffic, the added digital signal processing steps needed to encrypt all user data could potentially impact circuit performance and introduce new overhead in the network and in the endpoints, which could potentially require additional hardware (such as a dedicated encryption chip) in the endpoint devices and lead to increased costs.

Nortel believes that encrypting user traffic is important, not only because it will protect users from such attacks as eavesdropping, but also because it will

be key to creating trusted, secure end-to-end networks in the future, as well as help to enable such enhanced capabilities as network-wide identity management (see page 65).

To secure real-time traffic, Nortel recommends the use of Secure Real-time Transport Protocol (SRTP), a protocol defined by the Internet Engineering Task Force (IETF) that operates on top of IP. When implemented properly, encryption can be supported while keeping delays to a minimum. Nortel is currently applying its significant heritage and leadership in understanding the requirements of real-time networking to several technology innovations in this area.

A second challenge with encrypting user traffic is the need for greater processing power at the network endpoints – a challenge that Nortel is also working to address. Potential solutions, among others, could include encryption chips embedded directly into the end devices, or suitable high-speed signal processors that boost processing power.

At the same time, though, Nortel recognizes that user traffic encryption solutions must be sensitive to law enforcement needs. Indeed, another challenge is the need to comply with the legal intercept requirements of some governments to enable law enforcement to access public user voice traffic under court order. If this traffic is encrypted by the carrier, there is an expectation that the network must also have the ability to decrypt it when required for the customer and authorized law enforcement agencies. To do this, the system must be able to store, track, and secure the encryption keys needed to decipher the code. Here, Nortel is exploring such potential solutions as key-sharing mechanisms.

Media plane standardization

As it did with both the management and signaling plane security baselines, Nortel is taking its media plane security recommendations to the industry at large. The first focus for standardization is within the ATIS Performance, Reliability and QoS Technical Committee (PRQC). The PRQC is establishing a family of media plane security standards, and Nortel is acting as a technical editor. These media plane standards will specify appropriate levels of protection, such as endpoint authentication, the use of SRTP, key exchange mechanisms, and other security measures such as the use of firewalls designed specifically for handling multimedia traffic. Activity at the PRQC is currently under way, with the first standardization expected in the 2006 timeframe.

By establishing and implementing baseline infrastructure security requirements for the network management, signaling/control, and media planes, and driving the industry toward standardized adoption of these baselines, Nortel is demonstrating its significant leadership in shaping next-generation networks and enabling the secure, trusted networks of the future. ■

Mike Lee *is Senior Security Architect in the CTO Office, and is contributor and technical editor on the baseline security standards in the ATIS, 3GPP, and ITU-T standards bodies.*

Technologies to secure the enterprise

by Brad Black, Pat Patterson, and Tony Rybczynski

For enterprise customers, Nortel offers complete security solutions that provide adaptive, end-to-end network protection and information security – enabling enterprises to not only protect against growing threats to their operations and employee productivity, but also meet the heightened consumer privacy and accuracy requirements of regulatory compliance. Building on Nortel's Network Security Architecture, these solutions leverage a layered defense approach to security with endpoint, perimeter, core network, and communications security, as well as security management solutions components. Along with its own technology innovations, Nortel leverages best-in-class technologies from such partners as Symantec, Check Point, RSA, Opsware Inc., and GuardedNet (recently acquired by Micromuse, which has since been acquired by IBM). This article focuses on Nortel's comprehensive portfolio of security products and applications for the enterprise.

Today's enterprises are enjoying the many benefits of a richer set of communications capabilities, with fewer boundaries between them and their business partners, customers, and remote employees. However, these benefits can at times be outweighed by the various security risks of doing business on public networks and open intranets.

According to the CSI/FBI 2005 Computer Crime and Security Survey, total losses reported by the 659 respondents were more than US$130 million – with viruses ($42.8 million), unauthorized access ($31.2 million), and theft of proprietary information ($30.9 million) leading the way. Since most enterprises don't report losses from security breaches, the impact is undoubtedly much greater. What's more, costs that are reported don't reflect the damage that negative publicity can cause to an organization that has been attacked. And while attacks from outside receive most of the publicity, in fact threats to today's enterprise networks are more likely to come from internal sources than external sources. A 2005 Deloitte Touche Global Security Survey found that 35 percent of respondents said they had encountered attacks from inside their organization within the last 12 months (up from 14 percent in 2004), compared to 26 percent who said the attacks came from external sources.

Regulatory compliance is another key driver behind today's focus on security. The need to protect information from theft or tampering is being driven by the heightened requirements for enhanced communications privacy and the accuracy of reporting financial assets. According to Nemertes Research, 75 percent of today's increase in enterprise security spending is driven by compliance demands.

Nortel's approach

With growing threats and the requirements for regulatory compliance, enterprises must make the right business decisions to appropriately protect both their assets and sensitive information (payroll, research and development, etc.), as well as their customers' privacy. Ultimately, a solid approach to network security involves:

- addressing the technical, business, and human aspects of security;
- putting the appropriate processes in place; and
- choosing the right security solutions.

Such an approach not only ensures security of the network, but also enhances overall network reliability, resiliency, business continuity, and business productivity.

A properly designed and implemented security policy is the starting point for building a security solution. A clear security policy is an absolute requirement for all types of enterprises and should be a living document and process – one that is enforced, implemented, and updated to reflect the latest changes in an enterprise's infrastructure and service requirements.

Once a security policy is developed, the next major step is to ensure that that the policy is implemented using a layered defense approach (see sidebar on page 29). Nortel's layered defense approach, based on Nortel's Network Security Architecture (page 7), is designed to ensure that there are no single points of security failure in a network, while enabling a network to adapt to future threats. To ensure complete protection, a layered defense approach relies on applying successive zones of trust at multiple areas in the network. To enable easy integration and simplified operations that reduce overall network security costs, Nortel's layered defense approach is built on open, standards-based solutions that leverage the security capabilities and products from best-of-breed security vendors, such as Symantec, Check Point, and RSA.

This layered defense approach includes solution components for the endpoint, perimeter, core network, and

Layered defense approach provides multiple levels of protection against security threats

Similar to the way that it builds highly reliable networks, Nortel implements security in a layered fashion across all network devices and processes. This approach ensures that there are no single points of security failure in a network and that networks will be able to adapt to future threats. If a primary layer of security is breached, successive layers of protection are in place to thwart an attack.

This security strategy is applied across all levels in the layered defense approach, including core network, perimeter, communications, and endpoint security. (Security applied at each of these layers is described in the main article.)

As shown in the diagram, the layered defense solution components are mapped to a prototypical enterprise across zones of trust, each with its own set of security protections. The circle, square, diamond, and triangle symbols around each zone indicate the types of security solution components (core network, perimeter, communications, endpoint) that are leveraged by that zone. These zones typically include a data center, secure multimedia zone, local area network (LAN), security operations center, demilitarized zone (DMZ), and branch offices and teleworker sites.

• **Data center:** houses all critical applications and files and is protected by a Nortel Switched Firewall, Nortel Threat Protection System (TPS) intrusion sensors, and a Nortel Application Switch that performs load balancing and intelligent traffic management for multiple servers.

• **Secure multimedia zone**: provides security for IP telephony and multimedia applications and is protected by Nortel's Secure Multimedia Controller (which provides the firewall function and signaling encryption), TPS intrusion sensors, Application Switch, and Multimedia Communication Server and/or Communication Server 1000 that process the IP telephony and multimedia applications.

• **LAN**: provides local area networking for phone- and PC-based applications and is protected by the TPS intrusion sensors and the Nortel Secure Network Access Switch, and contains Nortel Ethernet and Ethernet Routing switches that facilitate VLAN separation between voice and data traffic.

• **Security operations center:** manages various security applications within the network and is protected by a Nortel Switched Firewall, and houses Nortel's TPS Defense Center, Configuration Manager, Enterprise Policy Manager, and

Network example of a layered defense approach

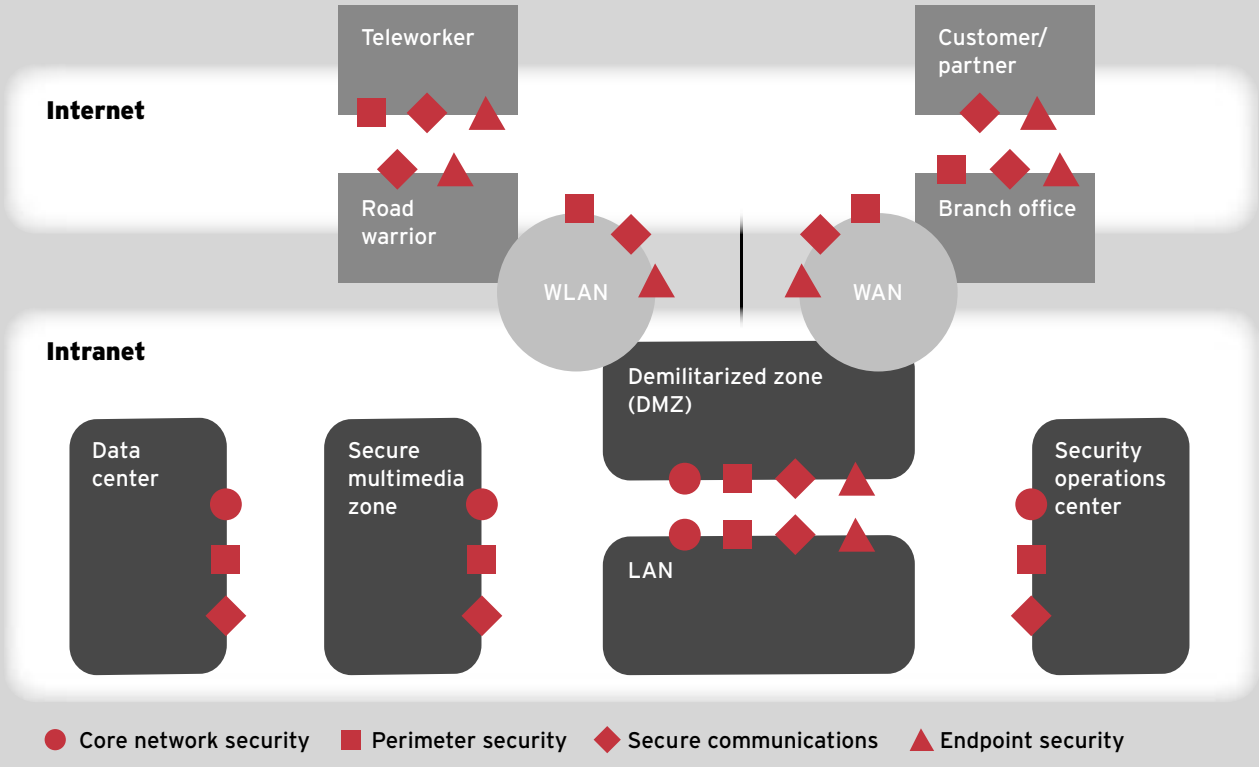

Security Event Manager.

• **Demilitarized zone** (DMZ): is the point of connection between the Internet, wide area network (WAN), wireless LAN (WLAN), and intranet and where services, such as the enterprise's website, are hosted. Typically, non-employees (customers, suppliers, contractors, etc.) entering the intranet are restricted to accessing only certain types of non-sensitive information. The DMZ is protected by Nortel's Switched Firewall, TPS intrusion sensors, and Application Switch. The DMZ also has a VPN Gateway, where valid remote users (teleworkers, road warriors, customers, and partners) can access the enterprise intranet securely. Traffic can come in over the Internet from a branch office through a VPN Router or, in some cases, a private frame-relay-based WAN. Traffic can also come in over a WLAN, which has its own protection in the form of a Wireless Security Switch, WLAN Access Point, and Wireless Mesh Network (when that access method is used).

• **Branch offices and teleworker sites**: are protected by VPN Routers that simultaneously secure devices from attack via an embedded firewall and encrypt all communications crossing the Internet or WAN, protecting them from theft of information or VoIP eavesdropping.

communications security layers, as well as platform security and management security that apply across all these layers. The remainder of this article looks at each of these components individually and highlights Nortel's offerings in each area.

Platform security

Individual devices – such as communication and application servers, user workstations, laptops, and mobile devices – that provide or use network services need to be secured. Each of these devices has specific security requirements that address user and administrator authentication and authorization to protect the device from external attacks, ensure security communication protocols, and protect the integrity of the device software and configuration. Platform security is a fundamental requirement in the basic phase of Nortel's Network Security Architecture (page 7).

Endpoint security

As employees, business partners, and customers make more use of the enterprise network to meet their business objectives, enterprises need more control of the endpoints that are used to access the network. Because so many threats are now coming from internal users on the network, this control must encompass wired and wireless endpoints not only within the network but also at remote locations, where there is less control over a user's device. The goal of endpoint security is to ensure the user has a valid identity and the connected device complies with the organization's security policy.

Endpoint security is a key functional element of Nortel's layered defense approach to security and integral to Nortel's Network Security Architecture. Nortel leads the industry in providing a unified security solution that is tightly integrated with the converged network infrastructure and supports authentication, compliance testing, and policy-driven user controls across the complete range of user devices (remote and local, voice and data, wired and wireless, and client and clientless).

The Nortel Secure Network Access (SNA) solution – which implements consistent, intelligent, user-based policy management across the enterprise network – delivers this capability by facilitating authentication and continually checking to ensure that the latest anti-virus or firewall applications, definitions, and OS software patches are installed and running on any and all devices both before and after users are authorized to access the network. The device-agnostic Nortel SNA solution provides proactive, continuous traffic analysis for any possible security attacks while also helping to meet regulatory compliance requirements (see sidebar page 31).

Perimeter security

Today's perimeter is not relegated to the organization's edge: it can be an internal perimeter around departments, or secure multimedia zones that protect multimedia and IP telephony call servers, or a perimeter at the external edge of the network around a data center or a single user. Today's perimeter security is all about keeping the "good stuff" in and the "bad stuff" out by securing the boundaries between zones that have different levels of trust. Accordingly, Nortel's perimeter security products are designed to ensure effective and efficient secure network zone boundaries, enabling businesses to protect their information assets without restricting their business agility.

The primary products include:

• the Nortel Switched Firewall (NSF), which is the primary Nortel enterprise perimeter firewall for use at the gateway to the corporate network or primary data center;

• the Nortel Secure Multimedia Controller (SMC), which is used to provide protection directly in front of key multimedia and Voice over IP (VoIP) applications inside the enterprise;

• the Nortel VPN Router, which is used

Nortel Secure Network Access solution: Proactive identity and policy compliance enforcement against security attacks

The Nortel Secure Network Access (SNA) solution provides enterprises with the controls required to ensure that all wired and wireless users – whether connected locally or remotely – have valid identities and comply with security policies, and that these identities are decoupled from the device being used. The Nortel SNA solution is tightly integrated with the converged IP network infrastructure and supports authentication, compliance testing, and policy-driven user controls across the complete range of user devices.

The general industry approach is to provide endpoint security with a client-based solution working into an in-line security portal. The Nortel SNA technology is unique in that it supports both a clientless and client-based approach, and utilizes a highly scalable control point: the Nortel SNA Switch.

A clientless approach reduces operational costs associated with having to deploy client software for every type of device in the enterprise environment, and opens the solution to non-employee devices (e.g., guest or contractor laptops). Using this approach, the Nortel solution provides flexibility, reduces administrative overhead, and minimizes performance impacts.

The Nortel SNA Switch provides endpoint authentication and interrogation of all devices that access the enterprise network system, including the device's operating system, patches, anti-virus software, personal firewall status, registry settings, and other system configuration components. This authentication and interrogation is performed out-of-path in the data transfer phase, providing a highly scalable solution that minimizes latency impacts on real-time applications.

The Nortel SNA solution is a comprehensive, device-agnostic endpoint security system that implements consistent, intelligent, user-based policy management across the enterprise network – providing proactive, "always on" traffic enforcement for any possible security attacks, while periodically reassessing the system for any configuration changes that deviate from the security policy. In its high-availability mode, all sessions are synchronized between two controllers working in a cluster configuration, eliminating single points of failure.

User authentication can take place through a variety of options including Media Access Control (MAC) address filtering, IEEE 802.1x/EAP authentication for LAN-attached devices, SSL-protected username password exchanges, or user authentication integral to IPsec for remote users.

In any case, authentication can be based on one, two, or three factors – colloquially referred to as what you know (password), what you have (physical token), and who you are (fingerprint or other biometric information).

Once authentication is complete, endpoint security enforces security policy compliance by ensuring that the required security definitions – such as anti-virus applications, personal firewalls, software patches, and other safeguards – are in effect before any user is granted access to network or application resources. It also provides for seamless quarantine and remediation of the infected resources by providing a mechanism that updates a client machine to meet host-integrity policy requirements.

This updating can be accomplished by dynamically assigning clients to virtual local area networks (VLANs) and applying context-based user filters – for example, to a "green" VLAN for full network access, a "yellow" VLAN when corrective action is required, and a "red" VLAN when network access is denied.

The Nortel SNA solution builds on Nortel's proven Tunnel Guard technology, which has been part of Nortel's integrated IPsec/SSL VPN technology, to deliver a unified approach to endpoint security across remote and local, voice and data, wired and wireless, and client and clientless environments. Nortel's VPN technology has been used to make secure connectivity available to more than 100 million users worldwide.

At the heart of this technology is a Nortel Software Requirement Set (SRS) of entities and rules that define security policy requirements on the network. The management interface enables administrators to configure SRS entities and rules on the intelligent Nortel SNA Switch (for local user access), as well as on VPN Routers and Gateways (for remote and WLAN user access). The SNA agent is either part of the multi-OS VPN client or is downloaded as an applet to a clientless device. It provides the checks that the PC or client device complies with the defined SRS entities, and reports SRS rule failures.

The Nortel technology also supports an open application programming interface (API) that third-party software vendors, such as Symantec, can use to perform more detailed self-checking and automatic software updates on the endpoint device. This capability demonstrates Nortel's commitment to a security ecosystem that delivers added value to customers through best-of-breed vendor partnerships.

Nortel's endpoint security architecture is also consistent with Trusted Network Connect (TNC), an open software architecture currently being defined by the Trusted Computing Group (TCG) – an industry alliance of prominent networking and security technology vendors of which Nortel is an active

member. TNC could be used by network administrators to enforce security policies for endpoint host connections across multivendor networks.

Moving forward, the Nortel SNA technology will evolve toward autonomic operation, extending signature-based crimeware detection to an innovative zero-hour behavioral protection model that uses real-time analysis of traffic patterns to combat new attacks, disabling malicious code before the code has a chance to propagate in the network. Moreover, as the ecosystem around the Nortel SNA solution evolves, automatic context-based remediation will be provided, whereby automatic software updates are downloaded to the device based on the failed checks and rules.

at the branch perimeter where integrated firewall, VPN, and routing functionality is needed; and

• the Nortel Application Switch Intelligent Traffic Management capability, which inspects and classifies application traffic flows to ensure applications receive the appropriate bandwidth priority and security treatment.

The NSF and SMC are next-generation application-aware firewalls that perform deep packet inspection (Figure 1) to detect and block attacks that directly target applications and data using the packet payload or application messaging. Nortel's market-leading Layer 2-7 deep packet inspection technology, originally developed for the Nortel Application Switch, is used in these products to provide enhanced protection, as well as accelerate the NSF's firewall performance by dramatically increasing the speed at which the inspection can take place.

As well as deep packet inspection, Layer 4-7 content filtering and denial of service (DoS) protection are built into both the NSF and the SMC. The NSF also leverages complementary Layer 2-7 security enabled by the Check Point Next Generation (NG) Application Intelligence engine, providing multiple layers of multi-gigabit protection at a perimeter. For additional perimeter protection, the NSF can load-balance multiple groups of intrusion sensors.

The integration of Nortel's switch-accelerated NSF platform with Check Point NG software provides perimeter protection without compromising application performance, incorporating deep packet inspection with a DoS signature database to identify the most popular attacks, including Teardrop, Smurf, Ping of Death, SQL Slammer, LAND, etc. *(For more detail, see the NSF product brief by going to www.nortel.com and entering nn110160 in the search field for the NSF 5100 series, and nn110161 for the NSF 6000 series.)*

Designed specifically to support real-time applications such as VoIP and multimedia, the NSF provides protection while maintaining user quality of experience. The NSF can also leverage the five 9s reliability of Nortel's Ethernet Routing Switch 8600, which can support up to four firewalls on its Service Delivery Module.

The SMC – part of Nortel's award-winning VoIP Security portfolio – provides a comprehensive approach to enabling a secure multimedia zone that protects business-critical converged or multimedia infrastructure resources from internal and external attacks. The SMC provides stateful firewalling, protection from application attacks (DoS, worms, etc.), and signaling encryption on traffic destined for the multimedia infrastructure, as well as integrated configuration and management with simplified policies and QoS capabilities. (For more detail on the SMC, see page 40.)

In addition to the protection provided by the NSF and SMC, the Nortel VPN Router includes a full-featured stateful inspection firewall that provides gateway and branch office firewall protections in a single platform with a state-of-the-art VPN capability that ensures that encrypted traffic is also firewall-inspected. Firewall user authentication in the Nortel VPN Router goes one step further in providing fine-grained access to network resources based on user permissions, whether over a tunneled or non-tunneled connection.

Finally, to protect the perimeter from network killer attacks (such as high-volume DoS, virus, and worm attacks), the Nortel Application Switch Intelligent Traffic Management (ITM) capability leverages real-time attack mitigation signature updates from Symantec, a world leader in this area. In addition, the Nortel Application Switch supports application delivery capabilities, such as load balancing and bandwidth management.

To mitigate risks from instant messaging and peer-to-peer (P2P) applications, such as Skype (the global Internet telephony company), enterprises can use the ITM capability to either minimize (rate limit/shape) the amount of band-

Figure 1. Deep packet inspection

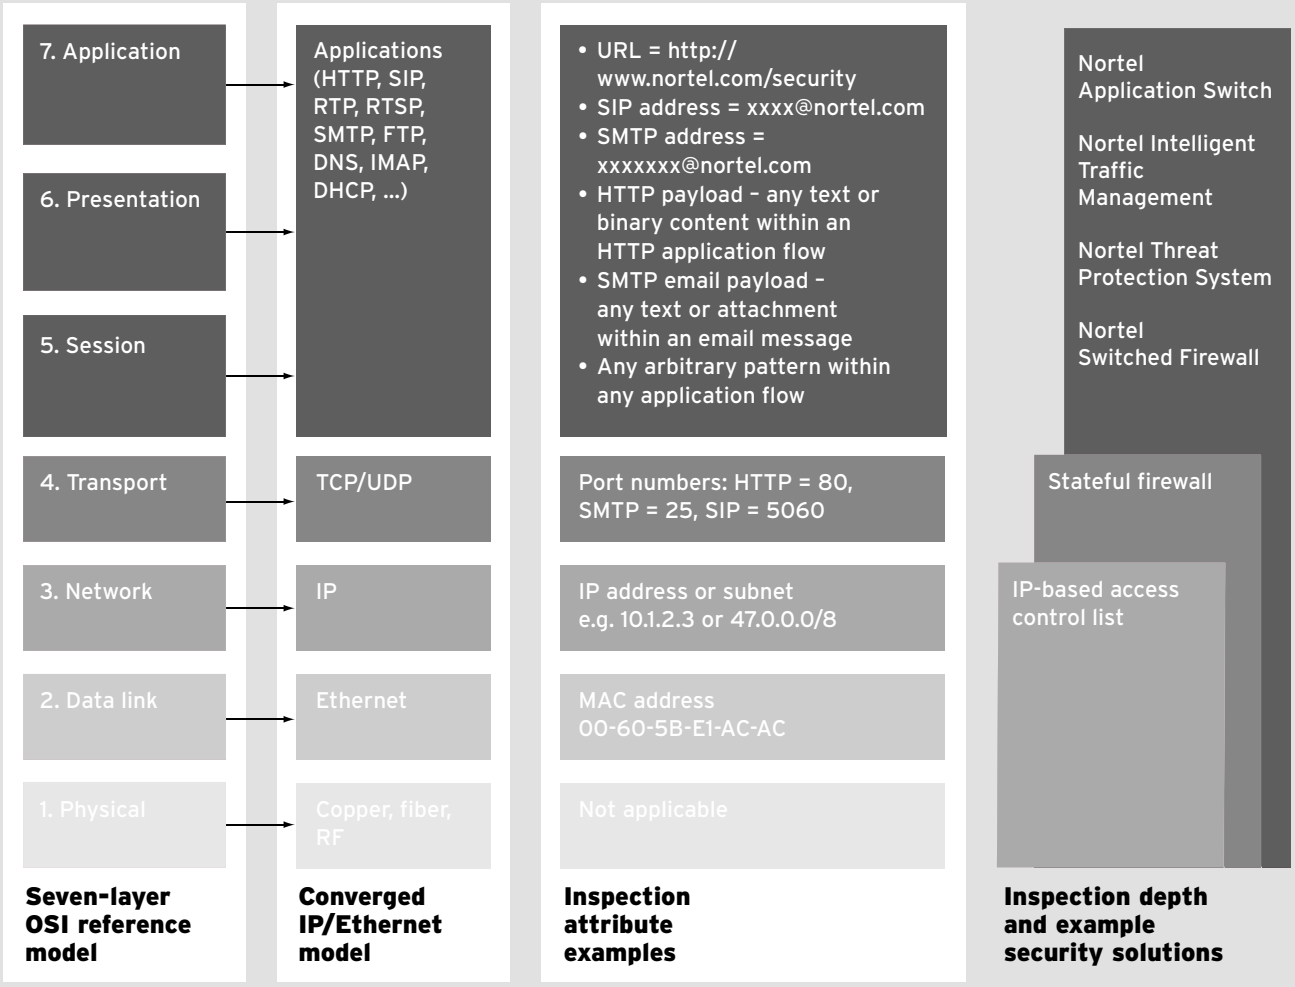

Deep packet inspection refers to the ability of a network device to inspect communication flows for application-level content to provide enhanced levels of security against increasingly sophisticated attack methods. Deep packet inspection techniques are incorporated in several Nortel security solutions, including the Nortel Application Switch, Nortel Intelligent Traffic Management, Nortel Threat Protection System, and Nortel Switched Firewall.

The diagram maps the seven-layer protocol stack of the OSI reference model to the converged IP/Ethernet architecture, and to the attributes that are typically inspected by security systems at each of the levels.

As shown on the right in the diagram, traditional security approaches stopped at the packet header level and could perform only security enforcement, forwarding, and analysis on such attributes as IP addresses. The introduction of stateful firewalls raised inspections a level higher to also examine Transport Layer Protocol port numbers, which enabled access decisions to be made based on the type of traffic – allowing web traffic but blocking email traffic, for example.

As network threats began to masquerade as valid IP hosts or valid applications, however, simple IP access control lists and stateful firewalls were not sufficient to identify malicious traffic – for example, bots, worms, and trojans could masquerade as web traffic to gain access to the network.

By employing deep packet inspection, Nortel security solutions can perform the more complex processing needed to identify this traffic – including protocol conformance checking, security signature pattern matching, and bandwidth management techniques to preserve known valid traffic while limiting unknown traffic – by examining URLs, SIP addresses, SMTP addresses, HTTP payloads, XML tags, etc. Since application payloads may be spread across many IP packets, additional techniques such as stream re-assembly and defragmenting are required to properly identify and match threat signatures within a session.

Deep packet inspection can be performed in-line to drop offending packets as they pass through a perimeter security device, or out-of-path to perform more complex processing to detect intrusions without introducing latency to real-time traffic.

width available for such applications or deny their use altogether – providing perimeter protection around a set of high-value latency-sensitive application servers [such as VoIP, SIP multimedia, enterprise resource planning (ERP), and customer relationship management (CRM) servers] or a data center, as well as protection against internally originated threats. *(For more detail, see the Application Switch product brief by going to www.nortel.com and entering nn104642 in the search field.)*

An additional way to enable perimeter defense is by using a virtual local area network (VLAN) to isolate and separate network traffic. VLANs are supported by many Nortel products, including Ethernet Switches, Ethernet Routing Switches, and Wireless Access Points. For example, enterprises regularly use VLANs to segregate IP telephony traffic from data traffic.

Core network security

Continually monitoring the network for malicious activity is key to ensuring that a network can detect an attack that slips through the endpoint layer of security and can take appropriate action to block the attack and ensure survivability. Monitoring is especially critical in the case of internally generated attacks or infections that may have unwittingly been released into the network by an otherwise innocent user. A prime example is a virus attached to an instant message. With an early warning system, the network can quickly identify the signs of such a virus, then define an effective mitigation tactic and push out a policy to enforcement points in the network to filter out the unwanted traffic. The Nortel Threat Protection System (TPS) provides this warning, as well as event management and mitigation in the network core (Figure 2).

Supporting the TPS in the network core is Nortel's Ethernet Routing Switch 8600, which can be deployed with optional modules including the Service Delivery Module, which supports up to four NSFs, and the Web Switching Module, which provides additional DoS, filtering, and application abuse protections.

A key goal of security in the core is to minimize disruption to network resources and applications from an attack (DoS, worms, virus, etc.). As part of the layered defense approach to security, the enterprise should choose reliable security solutions that ensure the highest level of uptime. In this regard, Nortel has long been designing products with the highest reliability for carrier environments, and brings high availability and reliability to its full suite of enterprise products to ensure the survivability of the network even when under attack.

Also playing a role in core network security is the Nortel Enterprise Policy Manager (EPM), which provides real-time policy-provisioning capabilities across such Nortel devices as Ethernet Routing Switches and VPN Routers on the network to mitigate the swift penetration of any virus or worm identified by the TPS, the Nortel Application Switch, or manually. Policy-based networking enabled by the EPM goes a step beyond authentication and compliance testing into the network, by ensuring that users or a group of users (for example, all employees of a financial department) have access to only authorized applications, and marrying this authorization to individual, departmental, or corporate security policies.

Communications security

Protecting corporate and government information from unauthorized discovery, eavesdropping, or misappropriation while in transit across networks is an important element of the layered defense approach to security. Having made secure connectivity available to more than 100 million users worldwide, Nortel offers enterprises several options within its portfolio to secure traffic from theft or eavesdropping both on and off their networks – allowing Nortel customers to choose the exact solution that meets their organization's security needs while minimizing the security solution's total cost of ownership.

In this context, Nortel VLANs can be used in conjunction with Nortel VPNs to enhance security. Coupled with the endpoint security capabilities described earlier, VPN and VLAN-enabled user devices are also checked before being allowed to join any network.

Nortel VPN technologies include IPsec and Secure Socket Layer (SSL). IPsec provides cryptographic protection at the network layer (OSI Layer 3), while SSL secures web traffic communications at the transport layer (Layer 4) and offers the added benefit of not requiring the support of client software. SSL acceleration, which offloads SSL encryption processing from application servers, is leveraged in all Nortel VPN platforms to enhance the performance and capacity of SSL secure communications.

Nortel is a leader in offering support for both these VPN technologies in a single platform. Nortel VPN platforms ensure that users receive the benefits of unified SSL and IPsec secure access – without having to pay a penalty in terms of scaling, performance, or total cost of ownership. These platforms include the following:

• the Nortel VPN Gateway (also available embedded in the Nortel Application Switch) offers unified SSL and IPsec secure remote access for a wide range of uses, from traditional remote access applications (such as email) to web-based access, extranets, portals, and externalized intranet web applications/resources. Native applications normally supported only in an IPsec environment can be seamlessly supported by the VPN Gateway through an SSL VPN (including toll-quality voice). Added SSL endpoint protections – such as automatic timeout for situations where a user forgets to log out when at a web kiosk, and dynamic access policies to limit application access based on the employee's location – provide enhanced use by enabling mobility without sacrificing security.

Figure 2. Nortel Threat Protection System

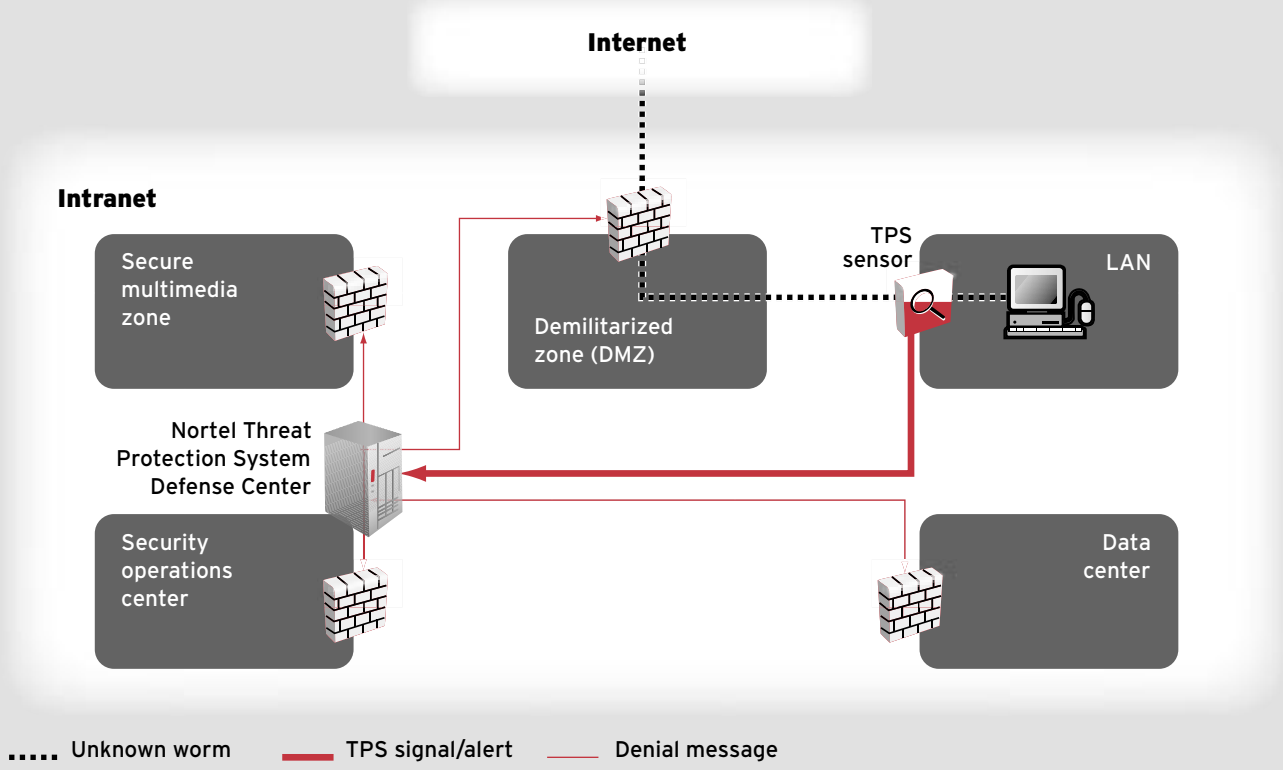

The Nortel Threat Protection System (TPS) works by placing intrusion sensors and real-time threat intelligence (RTI) sensors strategically throughout the network (e.g., wiring closets, server farms, data centers, and secure multimedia zones).

Using the data they gather, the TPS identifies threats and automatically applies a policy to network devices, including Nortel Switched Firewalls, as well as the Nortel Intelligent Traffic Management (ITM) capability, to mitigate various attacks. The TPS can detect known threats via deep packet inspection (see Figure 1), and can detect new unknown threats launched in zero-day attacks by scanning for anomalies in traffic patterns. The Nortel TPS's flexible rule set and instruction language allow the network operator to write rules that target not only the exploit but also the vulnerability – a unique differentiation, since one vulnerability can be exploited in many ways.

In the example shown here, if an employee unknowingly downloads a worm from the Internet onto his or her desktop and any part of the worm gets past the initial line of network defenses, the TPS is designed to detect the attack and send an alert to the TPS Defense Center. The Defense Center then immediately sends out a policy instructing all the Nortel Switched Firewalls in the network to block the worm, stopping the worm from propagating and ensuring the survivability of the network.

The TPS can leverage network asset information (including devices, operating systems, and applications) gathered by RTI sensors to eliminate false alarms, as well as provide a clear view of a possible threat's impact. For example, if an SQL slammer attack is introduced to a network that doesn't have any SQL applications, the TPS will determine that no remedial action is necessary. The real-time asset inventory can also be used to identify vulnerabilities before assets are compromised, as well as provide an audit trail for complying with regulations.

This real-time intelligence in the Nortel TPS not only enhances its active threat protection capability, but also greatly enhances the overall security architecture as it evolves to the autonomic stage.

• Nortel's VPN Router portfolio of award-winning, market-leading unified access VPN platforms can address site-to-site environments, ranging from small office/home office (SOHO) and branch offices to large enterprise and government data centers. In addition, the Nortel VPN Router supports remote access and endpoint protection, and includes an integrated firewall. Using the exclusive Nortel Secure Routing Technology (SRT), VPN Routers enable dynamic routing to be leveraged inside site-to-site tunnels. This SRT capability minimizes the administration costs and headaches of maintaining static routes and enhances the user quality of experience by allowing the best path to be selected for each type of data and destination.

• Nortel Secure Routers support a suite

of advanced security features, including IPsec-based VPNs for site-to-site environments, stateful packet inspection firewalls, authentication, access controls, network address translation (NAT), and VLAN tagging and forwarding protection. The robust routing, high performance, and low-latency capabilities of the Secure Router, combined with built-in security that ensures the privacy and integrity of transmitted data, makes it ideal for supporting the needs of multimedia communications.

• the Nortel WLAN 2300 series Security Switches support wireless security standards, such as WPA/WPA2 for security, or they can forward Point-to-Point Tunneling Protocol (PPTP), SSL, and IPsec tunneling to a VPN Gateway for secure communications across wireless LANs. These mechanisms can even be combined to provide the ultimate level of security. To ensure the mobility benefits of a WLAN are not compromised by the need for secure communications, Nortel WLAN Security Switches support seamless secure roaming across IP subnets, while maintaining the VPN tunnels.

• Nortel's Services Edge Router 5500 platform supports up to 50,000 IPsec tunnels, and meets the requirements of employee or customer secure remote access for large enterprises or VPN service providers.

Security management

Ensuring that a security solution can be effectively managed to maximize benefits and minimize costs is key to all phases of Nortel's Network Security Architecture. Effective security management requires configuration, policy, and event management components, and Nortel provides solutions in all these areas.

Configuration management: Nortel partners with Opsware Inc. to offer complete multivendor network configuration control that tracks, regulates, and automates all configuration and software changes across multivendor network devices – a key capability since devices that are incorrectly configured can be a major weakness in a network's security posture. In addition, Nortel's partnered solution enables IT governance initiatives, automates delivery and enforcement of network change control processes, and provides automated management of security and compliance best-practices.

Policy management: Nortel's Enterprise Policy Manager (EPM) provides centralized, policy-based provisioning to reduce management complexity and operational cost. The EPM also augments network admission control to help protect network resources and control DoS attacks, and enables the network to quickly respond to systemic and swiftly spreading threats before patches are available or virus updates are released. The strength of the EPM is its ability to push filters in real-time to numerous devices on the network. While being able to leverage the user-based policy provisioning of the EPM for an individual user or group of users, enterprises can also construct static policies that would apply to the entire enterprise network to push security policies to various devices. This capability ensures that an enterprise can react quickly when it learns of new threats, without having to rely on individual security policy updates per device.

Event management: Nortel has partnered with leading vendor GuardedNet (recently acquired by Micromuse, which has since been acquired by IBM) to offer a complete, centralized security event management and incident response system called Security Event Management (SEM). SEM greatly enhances the ability to provide clear situational awareness and is key to advancing to the autonomic phase in Nortel's Network Security Architecture. SEM acts to collect, normalize, correlate, and prioritize reports of security policy violations throughout the network and IT infrastructure. Examples of such violations include traffic filter and firewall rule violations, repeated unsuccessful log-in attempts, and anomalous traffic patterns, among others. Gaining a network-wide view of security events provides the basis for security operations, audit compliance, and incident detection, investigation, and response.

A layered defense approach is key to ensuring that an organization removes all single points of security failure and is able to fully leverage the benefits realized from Nortel's state-of-the-art applications and networks. By building multiple approaches to security enforcement into all areas within a network, organizations can deploy a security infrastructure that is highly resilient against attacks, while also providing the privacy/accuracy capabilities needed to remain compliant with so many of today's new regulations.

Furthermore, as the vision of Nortel's Network Security Architecture articulates, network security must evolve over time from today's primarily basic-phase networks to an authenticated phase that enables users to be more productive, as opposed to just protecting them from malicious activity. By following the guidance of Nortel's Network Security Architecture and leveraging Nortel's security products and solutions discussed in this article in a layered defense approach to security, organizations are well-positioned to not only defend their networks against today's threats but to also evolve to protect against tomorrow's threats, while minimizing downtime and maximizing user productivity. ■

Brad Black *is Security Solutions Leader, Enterprise CTO Office.*
Pat Patterson *is Team Leader, Security Solutions.*
Tony Rybczynski *is Director of Strategic Enterprise Technologies, Enterprise CTO Office.*

Protecting VoIP and multimedia communications from growing security threats

by Glen Brownridge, Louis LeVay, and Tony Rybczynski

Voice over IP and converged multimedia solutions can drive down operations costs and increase productivity by providing flexible service packaging and increased end-user mobility options. However, with online security breaches escalating almost daily, enterprises and service providers alike have concerns about their ability to meet the heightened security and reliability requirements of their new IP telephony and multimedia systems. Nortel – with its long history of building secure, reliable, end-to-end networks and its leadership in both VoIP and multimedia technologies – is a leader in securing these systems for both service providers and enterprises. In fact, Nortel's leadership was recognized in September 2005 when it was awarded the 2005 Frost & Sullivan Product Line Strategy Award for IP telephony equipment for the enterprise and carrier markets, based on its comprehensive IP telephony and multimedia security strategy.

Voice over IP (VoIP), or IP telephony as it is also called, has come a long way since the first rudimentary applications provided free, although erratic, phone calls over the unmanaged Internet. Today, voice quality on properly engineered public IP backbones and quality of service (QoS)-enabled private networks can match or even exceed that on public and private switched telephone networks. This capability is made possible in part by newer voice codecs (coder/decoders) that deliver superior PSTN-grade voice quality yet consume just a fraction of the bandwidth required by traditional TDM networks. Indeed, today's technology – hard and soft clients, softswitches and communication servers, media gateways, and application servers – makes it possible to deliver high-quality telephony services on a converged IP packet infrastructure more cost effectively than on the public switched telephone network (PSTN).

With these advantages – and the opportunity to significantly trim operating costs, according to Nortel's business case research – it's no surprise that all major carriers and most enterprises have implemented IP telephony to some degree. For example, Nortel itself achieved a payback in 10 months and a net present value (NPV) of $18 million through its IP telephony and multimedia deployment. What's more, these deployments will continue to grow as more and more networks converge. Infonetics Research projects the global carrier VoIP equipment market will grow at a compound annual rate of 25 percent, from just over $1 billion in 2002 to nearly $4.3 billion in 2006. Similar growth is expected on the enterprise side.

However, while VoIP offers compelling advantages, it also presents a security paradox. The very openness and ubiquity that make IP networks such powerful business tools can also make them security liabilities. The ports and portals that welcome legitimate subscribers and users into the network can also offer opportunities for hackers and others who would misappropriate network resources for personal gain or malicious intent. Using an arsenal of tools and techniques – including IP spoofing, denial of service (DoS) attacks, and backdoor entries, among others (see page 38) – attackers have three primary objectives when seeking to compromise a network: disruption of service, theft of service, and violation of confidentiality.

Securing IP telephony and multimedia

To combat such attacks, Nortel's IP telephony solutions – based on the Communication Server (CS) 1000/2000/2100 and Business Communications Manager (BCM) portfolios – have been designed to meet the stringent operational, reliability, and performance needs of service providers serving residential and business users, as well as those of small businesses, governments, and enterprises serving end users.

These IP telephony solutions can be enhanced to deliver converged, multimedia solutions – including real-time video, secure instant messaging, application sharing, white boarding, and presence – through Nortel's Multimedia Communication Server 5100/5200 (MCS 5100/5200), either as private or hosted solutions. Converged multimedia solutions are emerging as an essential productivity tool for enterprises, while at the same time enriching the communications experience for consumers and business users alike.

In protecting multimedia servers, clients, application servers, and

gateways, Nortel is applying a number of key principles to all of its product developments to protect the integrity of the IP telephony and multimedia communications system and ensure the confidentiality of user information. Specifically:

• Multimedia security solutions must be adaptable to the security policies of the network operator, whether that operator is an enterprise IT group or a service provider.

• While the IP networking infrastructure must be secured from a data perspective, any security mechanisms that are employed must operate in an environment that requires stringent VoIP and multimedia real-time performance and has very demanding requirements for latency/jitter (less than 150 milliseconds end-to-end) and packet loss (approaching 0 percent).

• Business-critical communication servers and associated signaling and control systems must be physically secure and protected from internal and external attacks.

• The simplicity and consistent user experience that Nortel drives to achieve across different devices as well as across wired and wireless connectivity modes must be maintained, and must be transparent to authentication methodologies and any encryption technologies that are used.

• Proactive development and support for standards in all multimedia products will ensure that service providers and enterprises receive the functionality and interoperability they require.

• A holistic approach to security must be taken across the entire multimedia environment to allow inter-service-provider, inter-enterprise, and public-to-private interoperability.

Nortel's strategy for securing IP telephony and more broadly multimedia is to take a layered defense approach to security, a key tenet of Nortel's Network Security Architecture (page 7). A layered defense approach to multimedia, and to networking in general, ensures that there are no single points of security failure in a network. It is achieved by using multiple approaches to security enforcement at multiple areas within a network, and is bolstered by leveraging standards-based solutions that utilize security capabilities and products developed by Nortel and through partnerships with best-of-breed security vendors. This approach is different than the traditional IT approach, which has focused on protecting the perimeter through firewalls.

In addition to the best practices applied to securing the overall IP network infrastructure, Nortel's VoIP and multimedia solutions address specific multimedia security capabilities within the layered architecture, including device-level security; perimeter protection, and endpoint compliance and protection at the network level; and application-level security.

Device-level security

At the device level, Nortel's telephony servers (the CS and BCM portfolios),

VoIP security threats

Network attackers have a broad repertoire of tools and techniques that they can use to launch multi-level attacks against various network resources. These tools and techniques, which are now being used or adapted to attack associated VoIP clients and servers, include:

• **Unauthorized access to network resources:** often the result of attackers taking advantage of weak user authentication and authorization tools in various communications systems, improper allocation of hidden space, shared privileges among applications, or even sloppy employee habits toward security. Indeed, there have been a number of well-publicized cases of hackers taking control of IP telephony clients – due to such system or product weaknesses as the lack of administration passwords (e.g., allowing an attacker to turn on the speaker phone and listen to a conversation), or vulnerabilities associated with unauthenticated configuration server access (e.g., allowing an attacker to configure a phone for unauthorized calling).

• **IP spoofing or session hijacking:** allows an attacker to disable a phone and assume the identity of a VoIP client.

• **Network sniffers** running on a PC can decode data from plaintext packets, enabling hackers to steal user names and passwords and then use that information to launch deeper attacks. Network sniffers can also eavesdrop on voice conversations, which is particularly easy to do over shared media technologies, such as wireless LANs and cable modems.

• **Denial of Service (DoS) attacks:** flood a multimedia server with illegitimate requests, overloading the system and preventing legitimate users from accessing service.

• **Man-in-the-middle assaults:** allow an attacker to intercept messages in a public key exchange between a communication server and a VoIP client and retransmit the messages, substituting his or her own public key, thereby tricking the original entities/users into thinking they are communicating with each other rather than with the attacker (for more on public keys, see page 58).

• **Backdoor entries** to communication servers can be accidentally or intentionally left open, allowing an attacker to launch a DoS attack using these entry points to bring down the entire phone system.

• **Masquerading:** enables a hacker to pose as a legitimate subscriber to illicitly obtain services, or to pose as a valid administrator or engineer to access the network, often to elevate user privileges, which in turn can be used to gain access to sensitive information.

multimedia servers (MCS portfolio), and application servers delivering unified messaging, contact center, and CTI services (Nortel's Application Center portfolio) are protected through the baseline security requirements document, which includes platform hardening (for example, by turning off unused services and closing unused ports), separation of management functions from service-critical functions, and tight access control and user authentication, authorization, and accounting.

Nortel, in fact, has a long history of building secure equipment for a variety of environments, including large public carriers/service providers, military installations, and enterprise businesses ranging from Fortune 50 companies to small businesses. For example, the CS 1000 runs on a real-time operating system that is used in very high-availability applications, such as the Mars Rover, life-support systems, nuclear power plant control systems, and communication satellites.

To the various voice, multimedia, and application servers, Nortel applies baseline security methodologies (see article on page 20) to ensure that backdoor programs that might be exploited by an attack are systematically closed. For example:

- unused ports (e.g., for consoles or remote modem access) are turned off;
- only authorized application software is allowed;
- multiple levels of privileges (e.g., monitor, configure, control) are supported for authenticated operational personnel;
- user passwords are securely stored;
- password formatting and change management are strictly controlled; and
- management traffic (such as billing information) can be optionally encrypted, even for internal transmission, by using Nortel's VPN Routers.

Platform and service hardening on the CS 1000/2000/2100 and MCS 5100/5200 ensures that even if threats manage to pass through other security provisions, their impact will be minimized and mitigated – a capability that has been confirmed by vulnerability testing undertaken by Nortel verification organizations.

Perimeter protection

Network perimeter protection is applied to VoIP and multimedia resources configured in what can be called a secure multimedia zone. This protection ensures that only legitimate multimedia, signaling, and management traffic is allowed into this trusted domain.

A secure multimedia zone provides a "security fence" around the communication and multimedia servers using products such as the Nortel Secure Multimedia Controller (page 40), guarding them from internal as well as external threats. Different approaches are warranted for service providers and enterprises, since service providers generally allow subscriber access to their VoIP systems over the open Internet, while enterprises are primarily concerned with internal deployments over a relatively secure enterprise network and with remote and mobile access over secure tunnel extensions.

Segregating network elements into secure zones enables network operators to balance the protection value of isolating these elements against the necessity for interaction among network segments. A variety of mechanisms – such as packet filters, firewalls, routing restrictions through non-secure virtual local area networks (VLANs) and non-secure VPNs, traffic segregation, and anti-spoofing measures – work together to provide layers of defense around key multimedia components.

Virtual LANs, which are supported on Nortel Ethernet Switches and Ethernet Routing Switches, provide basic network compartmentalization and segmentation capabilities, enabling various functions to be segregated in their own private LANs, with cross-traffic from other VLAN segments strictly controlled or prohibited. The use of VLAN "tags" organize the network architecture into discrete and logically separate areas for access/aggregation, services, call processing, media processing, and network operations. Because each of these network segments naturally has different access needs, traffic types, and user authorizations, segregating their functions protects them without impeding their usefulness.

Virtual separation of VoIP traffic across the WAN can also be achieved using IP-VPNs or MPLS (Multi-Protocol Label Switching) tunnels, leveraging Nortel's VPN Gateways, for example. VPNs and MPLS, which carry network traffic on a logical connection over a shared or public network facility, safeguard access over the wide area to critical VoIP resources.

Access to the IP core can also be controlled by a dedicated firewall and proxies, such as the Nortel Switched Firewall, which ensures that only valid VoIP protocols, Real-time Transport Protocol (RTP) media streams, and corresponding OAM traffic are allowed through. Firewalls and proxies also throttle excessive or inappropriate traffic to thwart DoS attacks. The Nortel BCM, in fact, provides both an integrated VPN gateway and firewall, as part of its office-in-a-box functionality for small and medium businesses.

Endpoint compliance and protection

Endpoint compliance and protection – whether applied locally or remotely to IP wireless or wireless phones, or to soft clients in PCs and PDAs – ensures that only authenticated users and compliant devices, as defined by the operator's security policy, can connect to the network and that these devices are authorized to access certain applications and networking resources.

Endpoint compliance and protection separates user and device authentication, so that services and applications are device-independent and users are not tied to a specific device. User authentication consists of a secure

Nortel's Secure Multimedia Controller: Creating a secure multimedia zone for business-critical communications

Nortel's Secure Multimedia Controller (SMC) is a key element of Nortel's IP telephony and multimedia solutions and a unique Nortel differentiator, providing a comprehensive secure multimedia "zone" that protects business-critical converged and multimedia infrastructure resources from internal and external attacks. The SMC provides filtering and attack protection on traffic destined for the enterprise's multimedia infrastructure, as well as integrated configuration and management features with simplified policies and quality-of-service capabilities.

As shown in the diagram, the SMC is a network element that front ends the secure multimedia zone, and connects it to the enterprise LAN/WAN. The secure multimedia zone can include VoIP servers, such as Nortel's Communication Server 1000 (CS 1000) and Multimedia Communication Server 5100 (MCS 5100) or third-party servers, as well as related application servers and media gateways. IP telephony and multimedia users connect into the enterprise LAN/WAN or remotely over the Internet through a VPN gateway/firewall.

The SMC acts as a stateful firewall that dynamically opens Real-time Transport Protocol (RTP) ports supporting media (user) traffic, based on session requests from local and remote users on the enterprise network, without statically opening a block of User Datagram Protocol (UDP) ports. Opening ports on-demand is a much more secure solution than the alternative of leaving unused ports open to attackers.

As a result, the SMC is able to maximize the availability of critical communication services while protecting against both internal and external threats, including such well-known attacks as TCP SYN Flood, port scan, ICMP Echo, Unknown IP Protocol, reassembly attacks, IP spoofing across the network, and UDP short header attacks. The SMC also enforces flow-based policies, including the maximum allowable connection rate, packet rate, session bandwidth, and maximum number of sessions for Session Initiation Protocol (SIP), H.323 (a legacy voice and video protocol), and UNIStim (the Nortel VoIP thin client signaling protocol).

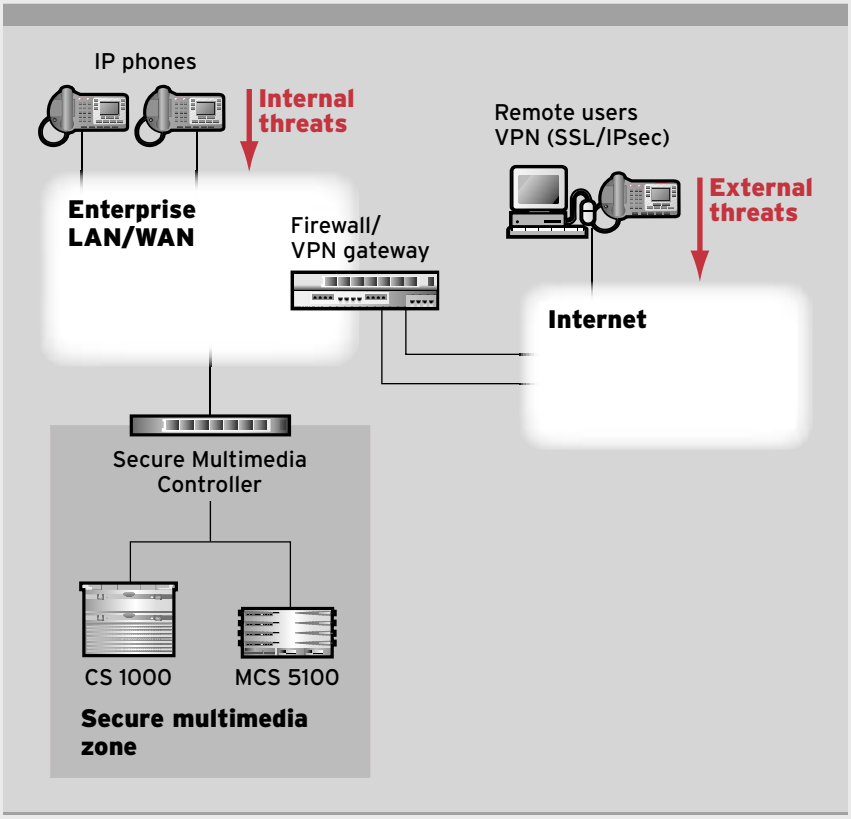

The SMC not only protects multimedia servers against attacks, but also provides a scalable architecture for both low-end and high-end deployments, with the ability to support cryptographic off-load capabilities. Encrypting signaling traffic prevents an eavesdropper from acquiring valuable or personal information of the call signaling, while encrypted keys allow IP phones to authenticate servers to prevent man-in-the-middle attacks from impostor servers that send false signals. For both authentication and encryption, the SMC supports 1024-bit RSA private keys, 16-character fingerprint-based authentication, and master key mechanisms for creating session keys.

The SMC also supports a number of secure session features, including session caching for more efficient reconnect; session/master key renewal with timeout keys for enhanced security; automatic update of private keys on the client; RSA throttling by limiting decryption rates for improved performance; resource conservation via secure signaling pings; and client policies tied to client subnets.

exchange of identifiers (such as user names) between the telephone and the communication server. Authentication determines that users are who they say they are, and links them to authorization rules that determine where they can go in the network and what they can do once they get there.

Device authentication verifies that the devices are legitimate. For example, two communication servers could authenticate each other to ensure that neither one is actually a hacker's PC masquerading as a communication server. Single-factor authentication requires only one proof point, such as the originating address. Two-factor authentication (stronger but slower) requires both something the user has (such as device address or biometric scan) with something the user knows (such as a password or SecurID code). Users will accept a two-stage log-on for one-time registration of their IP phones, but not for every time they want to make a simple phone call. This more stringent level of two-stage authentication should also be used by operators when accessing the network management systems.

Centralized administration of passwords required in the authentication process enables enforcement of password strength and removes the need for local storage of passwords on the network elements.

Currently, multimedia users on MCS are authenticated using industry-standard HTTP digest authentication. The client starts by making an unauthenticated request to the server, and the server responds with a response indicating that it supports digest authentication. The server also sends a nonce, which can be thought of as an opaque token. The client then requests the resource again, sending up the username and a cryptographic hash of the password combined with the nonce value. The server then generates the hash itself, and if it matches the request's hash, the request is allowed. In addition, an identity check can be performed every time the user uses a service; for example, an identity check can be done whenever a call is made or an IM sent, to ensure that the user has not modified the authentication. There also are plans to support transport layer security for multimedia signaling in future releases of the MCS 5200 and the CS 2000.

Additional mechanisms associated with endpoint protection are provided within the enterprise. The Nortel Secure Network Access (SNA) solution (page 31) supports IEEE 802.1x and the Extensible Authentication Protocol (EAP). To reduce the incremental processing burden on backend authentication servers, the Nortel WLAN Security Switch incorporates a unique 802.1x acceleration feature. These mechanisms prevent an unknown device from masquerading as an IP phone, or an IP phone moving to an unauthorized network port. The Nortel Ethernet Switch, Ethernet Routing Switch, and WLAN Security Switch portfolios all support 802.1x/EAP authentication. They also support Media Access Control (MAC) address filtering as an added form of access control.

In the case of IP telephony soft clients, the Nortel SNA solution supports IPsec VPN authentication, as well as submission of usernames and credentials via an SSL (Secure Socket Layer) VPN. IPsec refers to a suite of Internet Engineering Task Force (IETF) security protocols that protect IP communications through encryption, authentication, confidentiality, data integrity, anti-replay protection, and protection against traffic-flow analysis. IPsec is an optional overlay for IPv4 and an integral component of IPv6.

IPsec SSL VPN connections support interrogation of locally and remotely attached devices by leveraging the Nortel SNA solution, to verify compliance with organizational security policies, such as the latest firewall and virus software definitions. In the VPN space, the Nortel SNA leverages Tunnel Guard technology, a unique capability introduced a few years ago on the Nortel VPN portfolio that offers enterprises the flexibility to either enforce endpoint security through an agent installed as part of a Nortel VPN client, or by dynamically downloading the agent to any device attempting to set up an SSL VPN connection.

The SSL protocol is widely used to protect communications to and from the World Wide Web. SSL is built into most browsers and web servers to provide data encryption, server authentication, message integrity, and optional client authentication. Furthermore, SSL interworks with third-party security vendors by offering an open application programming interface (API). If the device does not meet security policy, it can be placed in a remediation VLAN until the device becomes policy-compliant.

Nortel's IP telephony solutions offer two other endpoint protection features of note. First they meet or exceed regulatory requirements associated with emergency calling (911 in North America). Secondly, Nortel's IP sets deliver such applications as corporate directories, visual voicemail, and zone paging through a secure application gateway – unlike other vendors' solutions which create vulnerabilities by building a browser into each set.

Application-level security

At the application level, secure communications provides link-level or end-to-end encryption for signaling and/or multimedia traffic.

Protecting the confidentiality of voice communications is important in making IP telephony a viable option to traditional TDM systems. In Nortel's discussions with end users, enterprises, and service providers, virtually all consider voice calls over enterprise and public circuit-switched voice networks as confidential, implying a high degree of trust in internal IT and service provider staff.

continued on page 44...

Extending VoIP securely across network boundaries

Until now, connections between different VoIP networks (which include multimedia communications) have been made primarily via TDM or analog mechanisms where packet traffic is converted to TDM traffic and then back to packet. While this approach isolates VoIP networks from each other and circumvents many interoperability issues, it also adds unnecessary service limitations, cost, and complexity. It also degrades VoIP quality, since multiple TDM-to-IP transcoding "hops" increase latency and can add distortion. These undesirable effects undermine service quality and the fundamental business case for VoIP – the potential to deliver profitable voice, video, and other real-time communication services over a cost-effective converged infrastructure.

If carriers and enterprises are to realize the full benefits of VoIP, they must be able to directly connect networks together, packet to packet, without converting to TDM. However, while direct packet interconnect improves service quality and lowers complexity and costs, it also raises some new issues:

- How does traffic traverse the stringent border protections, such as firewalls and network address translation (NAT), that protect inherently unsecure VoIP networks?
- How can you block intruders and attacks to safeguard the security of internal network resources while still welcoming legitimate IP telephony traffic without delay?
- How can traffic navigate through IP devices, such as firewalls and NAT, in a carefully controlled manner, while still meeting the end-to-end quality of service (QoS) and

Diagram A. Reference VoIP interconnect architecture

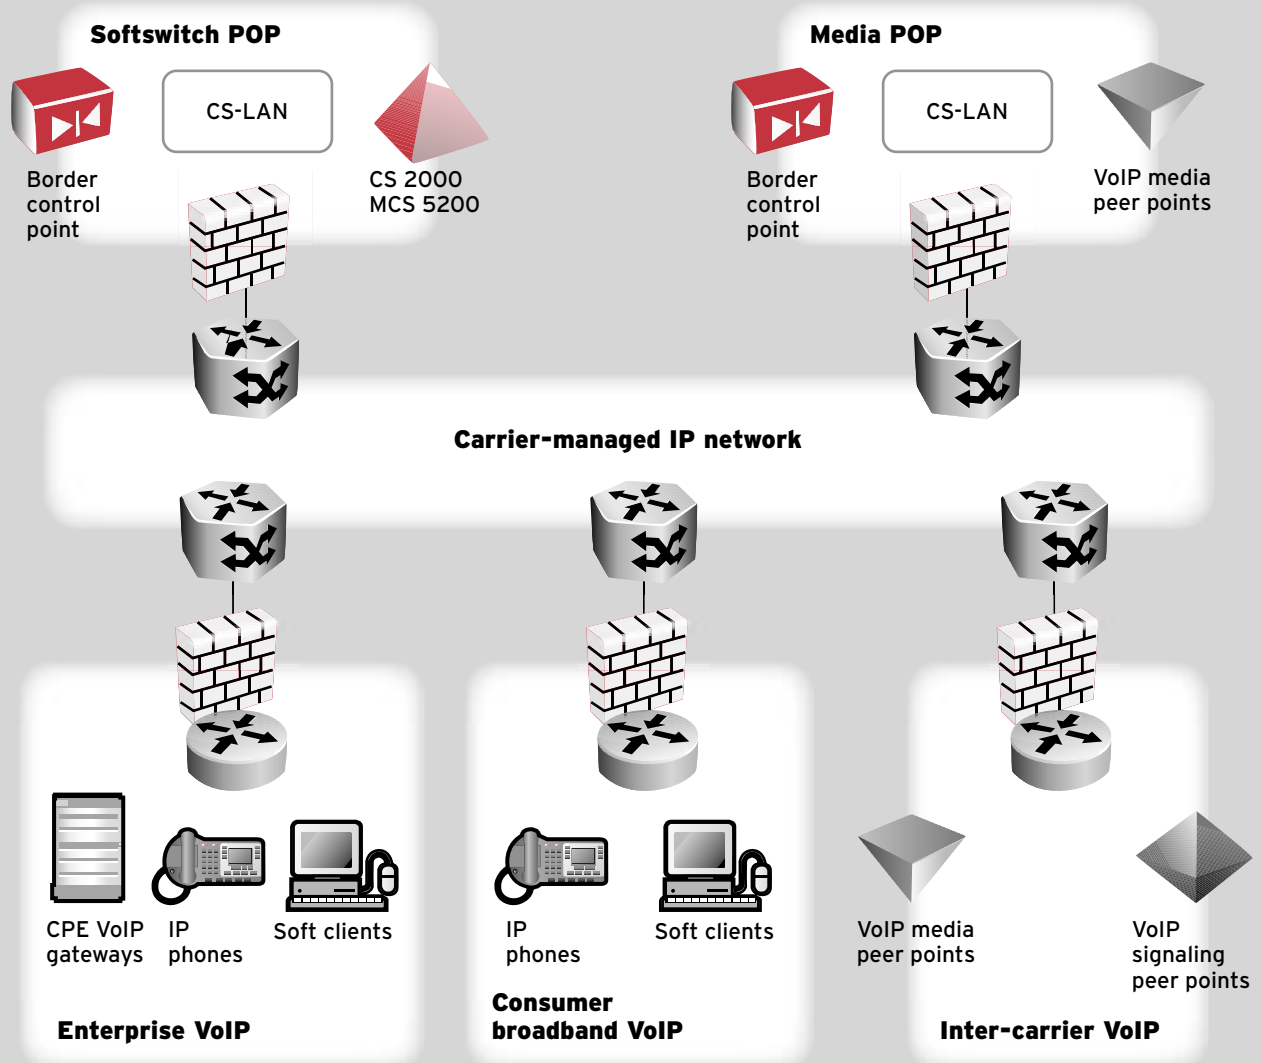

security requirements of voice traffic?

Nortel's distributed VoIP Border Control solution addresses all of these concerns. Nortel's VoIP Border Control is a set of signaling and media control functions that ensure secured, carrier-grade VoIP interconnect communication across any IP-to-IP network border – whether carrier-to-enterprise or carrier-to-carrier. These functions include such "must-have" capabilities as call/session setup, media transport selection, and NAT/firewall traversal and security, as well as such value-added functions as lawful intercept and policy services [such as admission control, QoS, and service level agreement (SLA) assurance].

Diagram A shows a reference network architecture for VoIP interconnect empowered by Nortel's VoIP Border Control solution. In this example, access is provided through a carrier-managed IP network that serves the following VoIP access domains:

• Enterprise VoIP, where operators provide a hosted VoIP service to the enterprise, either directly from IP phones or soft clients or through CPE VoIP gateways;

• Consumer broadband VoIP, where commercial VoIP providers use their own DSL or cable network to provide VoIP to residential subscribers; and

• Inter-carrier VoIP, where VoIP traffic is routed directly between the networks of individual carriers who have access agreements in place with each other.

The signaling traffic from these domains goes to a softswitch in a softswitch point of presence (POP) deployed in regional hub cities, while user traffic goes to a media POP. Each softswitch POP, in turn, serves several media POPs, which are normally deployed close to interconnect customers in remote cities or the suburbs of large cities in order to reduce VoIP back-haul transmission costs. (For simplicity, we show only one softswitch POP and one media POP in the diagram.)

Within the softswitch POP, the softswitches – a Nortel Communication Server 2000 (CS 2000) and/or Nortel Multimedia Communication Server (MCS) 5200 – provide centralized session control and signaling peering; centralized routing, translation, call processing, services and features, database, OAM&P, and billing capabilities; and centralized control of distributed Nortel Border Control Points (BCPs) that are inserted in a call or a session on an as-needed basis. The softswitch controls the BCPs as pooled resources distributed in strategic locations in the network. Depending on the network engineering configuration, BCPs can be located in a softswitch POP or media POP or both. As the organization grows, softswitches and BCPs can be easily added, providing a highly scalable solution.

Nortel's cost-optimized BCPs provide media proxy functions, including softswitch-controlled media firewalling on a per-session basis both for directly addressable endpoints and for obscured endpoints behind the NAT. In this case, the BCP will relay packets received on its reserved IP address and port only if the packets originate from an authorized media endpoint. All other packets will be discarded, mitigating denial of service attacks. Since media is proxied through the BCP and not directly between endpoints, the topology is hidden for media endpoints. This configuration also provides a publicly reachable media peering point between network trust domains and limits the need to route media (or IP packets) directly between endpoints.

Within each softswitch POP and media POP, a physically secure Communication Server LAN (CS-LAN) is used to provide internal communications among network elements.

As shown in Diagram B, redundant Nortel Ethernet Routing Switch 8600s (ERS 8600s – previously Passport 8600) provide high-performance packet filtering to regulate traffic entering or leaving the CS-LAN. The CS-LAN itself is segregated into several subnets. For example, element management systems (EMSs) reside in an OAM&P subnet, all call processing elements (call servers, media gateway controllers, etc.) are contained in a call processing (CallP) subnet, and optional media com-

Diagram B. Network segmentation in a Nortel carrier VoIP architecture

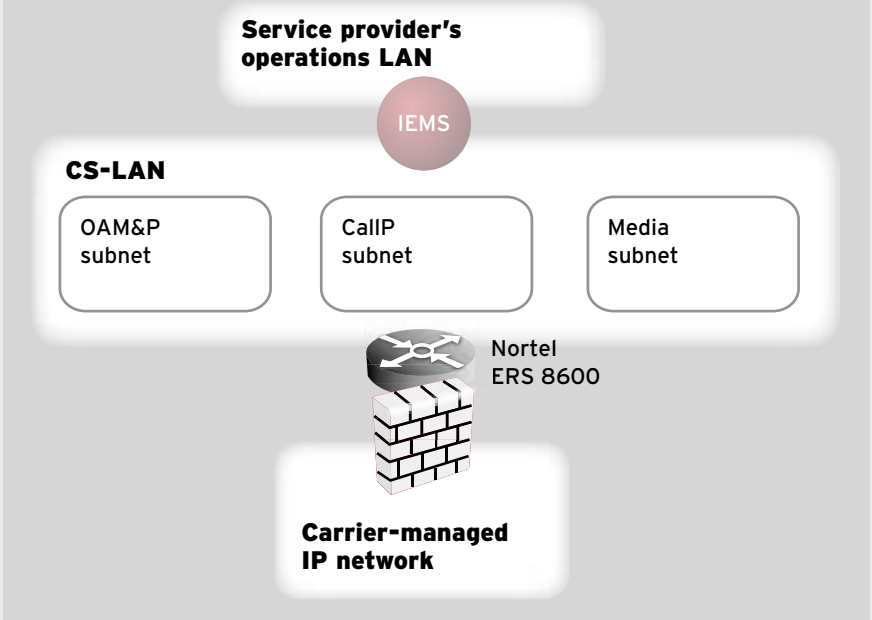

ponents (BCPs, media servers, media gateways) reside in a media subnet.

The subnets are mapped to VLANs on the ERS 8600, providing Layer 2 separation of traffic switched within the CS-LAN. Through the use of VLANs, subnets, and packet filtering, traffic between the subnets is strictly controlled by the ERS 8600 to create a full, logical separation and to guard against unauthorized access. Only valid call signaling, media, and OAM&P traffic can reach the VoIP servers that process that traffic in the CS-LAN. Traffic among the servers is isolated from traffic between external devices (such as remote media gateways or VoIP clients) and individual VoIP servers – a practice that limits the types of traffic to which key VoIP servers are exposed and minimizes the opportunity for attack.

Furthermore, certain types of traffic are regulated by proxies between the CS-LAN subnets. For example, all access to CallP and media subnets from the service provider's operations LAN – also referred to as the NOC (network operations center) – must be proxied by EMSs, such as the Nortel Integrated EMS (IEMS), through secure, authenticated interfaces so that key call processing functions are protected even from internal threats.

With the addition of physical security around CS-LAN interfaces – along with complementary security features, such as authentication, encryption, and stateful firewall capabilities at the demarcation points between remote VoIP access/aggregation networks and the core VoIP network – this network architecture provides a solid level of security both inside and outside the CS-LAN.

Advantages of a distributed solution

Nortel's distributed Border Control solution offers a number of advantages over standalone session border control (SBC) solutions, where both media and signaling traffic are routed through the SBC. Because BCPs, media blades, and ports are network-wide shared resources across multiple softswitches (CS 2000s and/or MCS 5200s), traffic loads can be optimally distributed to provide better resource utilization and network-level resilience. For instance, traffic can be automatically redirected from highly utilized BCPs to less-utilized BCPs to alleviate congestion. With the standalone SBC solution, resource sharing is confined to the per-node level, making it difficult to coordinate multiple SBCs to balance traffic loads and redirect congestion traffic.

Moreover, signaling security (e.g., encryption and key exchange) is done end-to-end between VoIP application endpoints [e.g., CICM (Centrex IP Client Manager), MCS, and IP phones], which is the preferred option for VoIP. Standalone SBC solutions can only provide hop-to-hop security, and they must be in the routing path, leading to scalability challenges.

In addition, the Nortel VoIP Border Control solution is an integral part of Nortel's turnkey VoIP solution. It has been rigorously tested through Nortel's product verification process, integrated with the IEMS, and included within Nortel's well-established product support infrastructure. In contrast, introduction of another new SBC solution would require significant interoperability testing, OSS integration, and additional product support.

For a more detailed description, see *Extending VoIP across network boundaries* (by going to www.nortel.com and entering nn107880-041404 in the search field). This document discusses the functions of VoIP border control, the approaches being taken by major softswitch vendors and others, and which strategy might be right for a customer's VoIP interconnect needs.

The majority of enterprises send data traffic in the clear (without using encryption) across their internal networks, having invested in making this traffic secure through application of a layered defense approach to securing their converged enterprise IT infrastructures.

Indeed, today's switched Ethernet in-building architectures – with voice VLANs, protocol controls such as Address Resolution Protocol (ARP) spoofing prevention, and secure wiring closets – can make in-building IP calls as secure as TDM calls.

To secure calls outside the building into the PSTN, packet traffic is first converted in a media gateway to a TDM call, and security that is equivalent to the TDM environment is maintained. If calls go to another building across a leased line or frame relay/ATM virtual circuit, then again the majority of enterprises traditionally consider this secure enough for data and, more recently, for voice.

At the other end of the spectrum, virtually all customers consider any traffic being handled over the Internet – whether voice or data, and whether accessed over wired or wireless – as being unsecure and subject to potential snooping attacks. This concern leads to the general practice of using SSL to secure critical information, such as user authentication and financial transactions, and IPsec VPN technologies to secure enterprise remote and branch access. Nortel IPsec and SSL VPN solutions totally secure IP telephony calls with teleworkers, road warriors, and branch and remote offices within enterprise networks.

Similarly, the explosion of wireless LANs (WLANs), which are inherently a shared media technology, has brought with it the need to provide over-the-air encryption to avoid any form of eavesdropping. In fact, many enterprises don't differentiate between Internet and WLAN access from a security perspective. WLAN standards such as IEEE 802.11i mitigate voice and data

vulnerabilities associated with running voice over WLANs.

In addition, end-to-end voice media and signaling encryption is a requirement for critical command and control voice communications in military and government security agencies. Secure Real-time Transport Protocol (SRTP) can be used to secure the media in the call, so that even if an attacker intercepts every voice packet, the information is protected.

On the other hand, while end-to-end telephony media encryption based on SRTP is being introduced by Nortel and some other enterprise vendors, end-to-end VoIP encryption is not seen as a requirement by many enterprises when balancing risk versus cost.

Moreover, regulatory wiretapping legislation in the U.S. and elsewhere precludes providing end-to-end voice encryption for consumer VoIP services.

Risks associated with VoIP eavesdropping can be effectively mitigated through various mechanisms, starting with strong authentication to positively identify end users – for example, by using challenge/response-based SIP (Session Initiation Protocol) client authentication and IPsec call signaling authentication between servers, and between servers and gateways.

Encrypting the signaling traffic can foil any attempts at systematic VoIP eavesdropping by protecting any information on the session, including user authentication information and session information, such as who is calling whom.

Signaling encryption can use IPsec between servers and between servers and media gateways, and Secure UNIStim (based on RSA public key encryption and AES-128) and Transport Layer Security (TLS) protocol for SIP-based clients. TLS is an IETF standard that merges SSL and other protocols. TLS enhances SSL with more secure data encryption and is supplanting SSL as a major standard for securing Web/HTTP traffic and VoIP protocols such as SIP. Nortel already broadly supports IP telephony signaling encryption using these mechanisms.

Sitting squarely between low-risk internal IP calls and high-risk Internet and wireless calls is an environment ideally suited to snooping: meet-me conferencing.

Access to conference calls is traditionally controlled by an operator or by passwords that are sent to all participants: anyone who receives the password through social engineering (e.g., impersonating an employee who claims to have misplaced the information and asks for the password, or by picking up a page containing the password at a printer station) can dial in. While some systems have limited chairperson controls managed through the telephone keypad, these controls are hard to remember and seldom used.

Nortel's MCS provides a range of easy-to-use meet-me conferencing security features through a chairperson visual control panel, which allows the chairperson to track who has joined or dropped off the call, and even to require participants to revalidate through a secondary password. All this makes for a significantly more secure environment for both traditional and IP telephony.

For customers wanting assistance to ensure the utmost security in their VoIP and multimedia deployments, Nortel Global Services can provide secure design and integration services that evaluate the existing network architecture, include collaborative working sessions and project planning with staff, provide a detailed security architecture plan, and provide detailed network diagrams for the VoIP and/or multimedia solutions and specific deployment requirements.

Future capabilities

By taking a layered defense approach to multimedia security within the basic phase of Nortel's Network Security Architecture (see page 7), Nortel allows its enterprise and service provider customers to defend their networks against today's threats at the levels appropriate to their applications.

Beyond these capabilities, the autonomic phase in this architecture focuses on network-wide security incident management through a combination of sophisticated correlation of events, logs, and external security sources, flow-based anomaly detection, and vulnerability assessment, as well as the addition of centralized policy control and network intelligence that will enable the network to automatically protect data and multimedia traffic alike, without human intervention.

Building on this control and intelligence, the next architectural phase – the "authenticated network" – focuses on managing user identities and individual entitlements across different network domains. This network-wide capability will enable the network to manage a single identity profile for each user, and recognize that individual regardless of his or her location or device type. Identity management could be extended in a federated model – a system based on agreement between two or more enterprises to allow the user to use the same user name, password, or other personal identification across the member networks in order to conduct transactions. This capability will open the door for truly trusted collaboration among consumers, among enterprises, and between enterprises and their customers. ■

Glen Brownridge *is Senior Product Line Manager, Carrier VoIP.*
Louis LeVay *is Carrier VoIP Architect.*
Tony Rybczynski *is Director of Strategic Enterprise Technologies, Enterprise CTO Office.*

Safeguarding mobile devices and wireless network infrastructure in a broadband mobile world

by Frédéric Bastien, John Garrison, Don Keeler, Emily Nichols, and Paul Tse

For nearly two billion GSM, UMTS, and CDMA wireless subscribers around the world, the advent of mobile broadband applications such as wireless streaming video and real-time collaborative videoconferencing is increasing productivity and bringing a rich world of information whenever and wherever it is needed. Wireless operators, however, are faced with significant challenges in not only protecting mobile devices from hackers, virus-writers, and others with malicious intent, but also securing their own network infrastructures against attacks launched from these mobile devices. Nortel is combining its industry-leading experience across all the major wireless technologies, its broad range of security products, and its standards leadership and understanding of end-to-end security requirements to design and deliver the security solutions that wireless operators will need to meet these challenges.

Mobile broadband is emerging as the new access reality for next-generation communications. Until now, mobile devices were mostly based on proprietary operating systems, had fairly limited processing capabilities, and lacked high-speed access to Internet applications. However, in much the same way that high-speed DSL and cable Internet access have supercharged wireline operators' businesses, Universal Mobile Telecommunications System (UMTS), High Speed Downlink Packet Access (HSDPA), and CDMA 1xEV-DO (evolution data optimized) technologies are giving wireless operators the bandwidth boost needed to do the same. In fact, new mobile devices now enjoy megabyte access to the Internet, can store gigabytes of information on memory cards or embedded memory, have much stronger processing capabilities, and are running industry open-standard operating systems, such as Windows Mobile, PalmOS, and Symbian.

These new capabilities, however, are bringing mobile devices and communications increasingly into the sights of hackers, spammers, and virus writers. Today, wireless broadband service providers have become increasingly aware that they must protect not only their subscribers from such traditional Internet threats as viruses, spam, worms and trojans, but also their own infrastructures from offending mobile device behavior that could compromise subscriber data, corrupt billing records, or even congest network resources – denying not only data service but also voice service to subscribers, with a consequent hit to the wireless operator's revenue line.

Nortel is uniquely positioned to help wireless service providers protect their network infrastructures and mobile devices. Nortel offers network infrastructure products and professional security services for all of the major wireless technologies, providing the opportunity to integrate security detection, mitigation, and preventative measures where they are most effective in the network. Nortel, for example, already provides mobility applications for more than 300 customer networks – with more than 240,000 wireless base stations deployed worldwide. Nortel was also the industry's first supplier with wireless networks operating in all advanced radio technologies (GSM/GPRS, CDMA 2000 1X, UMTS, and WLAN) and is the only end-to-end provider of all next-generation wireless solutions, including HSDPA and 1xEV-DO.

Nortel also pioneered the Wireless Mesh Network solution that incorporates several security features [such as implementation of the Robust Security Network (RSN) as defined in the latest IEEE 802.11i security standard]. This solution extends wireless LAN (WLAN) coverage across a much larger area than such traditional hotspots as airports and Internet cafes, and beyond the conference rooms and kiosks of enterprise campuses. (For more on Nortel's Wireless Mesh Network and its security architecture, see Issue 2 of the *Nortel Technical Journal*, page 20, at www.nortel.com/corporate/news/collateral/ntj2.pdf.)

In addition, Nortel offers a broad range of layered defense security solutions – including firewalls, intrusion detection and intrusion prevention systems, and secure VPN gateways and routers that provide multiple levels of protection across the network (see articles on page 28 and page 37). As well, Nortel's experience in security forensics and its protocols expertise enable effective solutions to be architected and designed to protect a wireless operator's evolving network. Finally, Nortel's Global Services organization helps wireless operators protect themselves from the host of aforementioned threats by

designing and integrating security seamlessly into their wireless infrastructures.

Wireless security challenges

Although many of the underlying security risks are the same as in wireline broadband networks, the security threats associated with mobile data devices and with the wireless network infrastructure elements are more numerous, varied, and often unique, and they must be addressed with tailored solutions.

When a mobile device connects only through a particular service provider's wireless network, it – and by extension, the operator's network infrastructure – can be protected by firewalls, anti-virus, and other in-line security mechanisms deployed in that infrastructure. However, to increase their usefulness, mobile devices often have multiple ways of staying "connected," allowing a single device to be accessible across different environments, including cellular packet data networks (CDMA 1xEV-DO, GPRS, EDGE, UMTS), Bluetooth, WiFi, and mobile-to-mobile connections.

Ease of mobility between network environments and the ability to roam into visited networks means the operator's network has to deal with devices that it neither knows nor manages – across many different types of mobile phones and several different operating systems (unlike the wireline world where the Microsoft OS dominates). For a wireless service provider, the additional threats resulting from connectivity outside of its control make it more difficult to protect its subscribers' mobile devices, which in turn presents

Figure 1. Nortel's wireless gateway hosted security suite

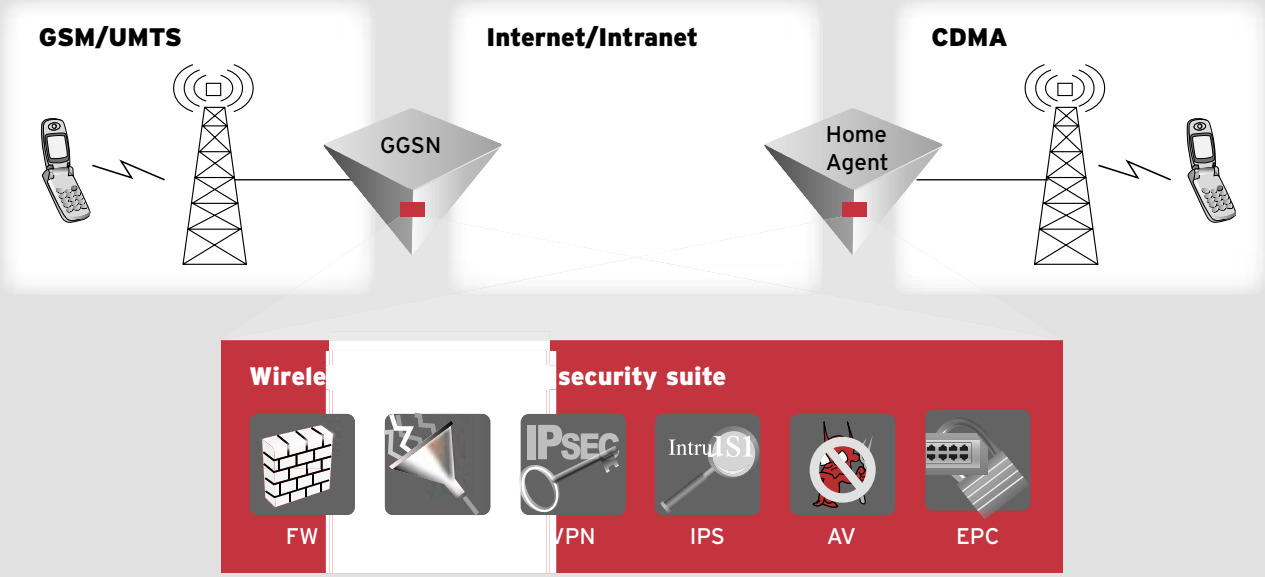

Nortel's hosted security suite is being integrated on Nortel wireless gateway nodes that provide the interconnection between wireless networks and the Internet/intranet – the Gateway GPRS Support Node (GGSN) in GSM/UMTS wireless networks, and the Home Agent (HA) in CDMA wireless networks. Because these wireless gateways see all the data traffic going between the mobile devices, in addition to the traffic between the mobile devices and the Internet, they are an ideal location for hosting the security suite.

The hosted security suite includes a stateful firewall (FW), URL filtering (UF), and virtual private network (VPN). An intrusion protection system (IPS) and anti-virus (AV) protection are being developed. Endpoint compliance (EPC) is a future capability that will allow the network to query the mobile terminal to determine what security software it has in place (such as firewalls, up-to-date signatures, etc.) and then make a decision on whether to grant the mobile device access to the network.

One function of the GGSN and HA is to assign the IP address that will be used by the mobile device. When mobile devices are anchored off the same GGSN node or the same HA node, traffic between mobiles (for music sharing, for example) can pass directly through that node in peer-to-peer fashion without having to go through the Internet/intranet. The GGSN or HA can enable direct mobile-to-mobile connection because it knows, or "owns," the IP addresses of both mobiles, and can therefore bypass having to route the traffic between them through the Internet/intranet.

Unlike security solutions that rely on external nodes to provide such functions as anti-virus protection or intrusion protection to secure this mobile-to-mobile traffic – which requires the traffic to be routed out into the wireless operator's network and back again – Nortel's solution of centralizing an integrated suite of security services on the gateways protects this traffic more effectively and efficiently, while making it easier to upgrade or add new security capabilities.

challenges when introducing secure value-added and revenue-generating applications and services to these subscribers.

What's more, in an environment of "always on" high-speed mobile data, undesirable and unsolicited packet traffic, such as common port scans, that is of no great concern in wireline networks can consume valuable radio resources and create denial of service (DoS) conditions on the access network – affecting not only wireless packet data services but also the wireless provider's voice services that share the same scarce radio resources. Such DoS attacks are especially worrisome for wireless operators since voice services still generate the overwhelming majority of their revenue.

For example, a mobile device that connects over WiFi, Bluetooth, or other non-secured means can become infected with a virus or worm, and that infection may try to propagate when the mobile device reconnects to the wireless service provider's network. Indeed, the world's first mobile phone virus, Cabir, started spreading from handset to handset in early 2005, hidden inside photo or sound messages received through Bluetooth wireless connections. A hacker only needs to detect an unsecured, open Bluetooth interface in any device to push a virus or other form of malware through it.

To a large extent, mobile-to-mobile traffic within wireless networks is not addressable by existing security mechanisms. Traditional firewalls and threat

Nortel joins forces with Websense to protect mobile users from the dangers of the Internet

Nortel has teamed up with Websense, Inc. to deliver an innovative URL filtering and security solution that helps protect GSM/UMTS mobile handsets and devices (including computers using GPRS/UMTS PCMCIA memory cards) from unwanted or even malicious content from spammers, hackers, and overly aggressive marketers who are increasingly targeting wireless subscribers. Nortel's initial deployment is targeted for GSM/UMTS-enabled devices, although Websense's solution will work with any device and is independent of the access technology (CDMA, UMTS, LAN, or WLAN).

Delivering a new level of protection for mobile devices, this solution combines Websense's web security and filtering expertise with Nortel's leadership in end-to-end packet networking technology to put more intelligence into the network and position wireless operators to deliver secure, reliable next-generation services and applications to end users globally.

Specifically, Websense's URL filtering technology is being combined with Nortel's Gateway GPRS Support Node (GGSN) platform, which provides the interconnection between GSM/UMTS wireless data networks and external data packet networks. The GGSN wireless packet core solution integrates wireless and IP value-added services and enables personalized IP packet inspection and filtering for a given user's data session. This capability enables different filters to be applied to different accounts according to the preferences of the user, (i.e., allowing certain content for some users, but blocking it for others).

Websense provides the URL database that is queried by the GGSN. The database will return an allowed/disallowed indication to the GGSN, which will then redirect disallowed requests to a web page indicating that access has been blocked. To ensure that the URL database is kept up to date, it receives automated updates.

A market leader

Websense's best-of-breed web filtering and web security solutions have been deployed with more than 24,000 organizations worldwide. The company was recently named the market share leader in the web filtering software segment of the secure content management (SCM) market for the third consecutive year (IDC, October 2004).

The innovative Nortel/Websense URL filtering and security solution not only protects mobile handsets and devices from content that could interfere with voice and data services, but also blocks the download of malicious code. It allows GSM/UMTS wireless operators to set mobile handset and device Internet access policies for subscribers across three URL filtering categories – security threats, adult material, and undesirable content. Each category contains content in a variety of subcategories to allow for flexible and carefully defined policy settings.

To block threats and unwanted content found on high-risk websites, subscribers can request that the operator automatically restrict or block access, which, for example, could be used for parents to prevent their children from viewing inappropriate or dangerous material, such as adult content.

The Websense web filtering and web security software fully and seamlessly supports the wireless access protocol (WAP) and wireless markup language (WML) for wireless browsers, along with their wireline browser counterparts HTTP and HTML, providing mobile handset users accessing the Internet the same level of protection that traditional browser-based users already receive from Websense software.

This alliance with Websense is yet another step toward building a complete integrated security service that Nortel can deliver to wireless operators. As Nortel continues to enhance the security service's suite of capabilities, Nortel will be able to provide operators with market-leading security solutions their subscribers need and want.

protection systems, for example, can't currently be effectively deployed in the mobile-to-mobile subscriber services path, because the wireless radio links themselves have special characteristics that do not lend themselves to these typical IP security solutions (for example, the traffic is tunneled differently). Based on its broad networking expertise, Nortel is currently exploring with customers a number of innovative solutions that will address this challenge.

In addition, malware can infect a mobile device through the sharing of removable media, such as Secure Digital or CompactFlash memory cards, or through the synchronization of email, calendars, files, and other applications with desktop PCs. With malware able to propagate in either direction between the mobile devices and PCs, cross-contamination can result – with wireless devices infecting PCs in the network, and vice versa.

Malware is also becoming increasingly problematic as mobile devices acquire multimedia web browsing capabilities beyond just text and static images. Increasingly, malware can be attached to short message service (SMS), images, video, instant messages, and other forms of media, making them more susceptible to attacks in much the same way as PC web browsers with multimedia capabilities.

Such threats complicate the ability of the wireless service provider to secure not only its wireless infrastructure but also its ability to offer effective hosted security to its subscribers, because, at least for now, traditional protection mechanisms exist only while connected to the wireless service provider's own network. Without this level of protection, wireless operators will find it difficult to expand on the range of IP-based services that consumers will trust and accept.

Nortel's hosted security suite

Nortel is at the forefront of building a comprehensive hosted security suite that will enable operators to offer their

Figure 2. Carrier endpoint compliance – solution elements

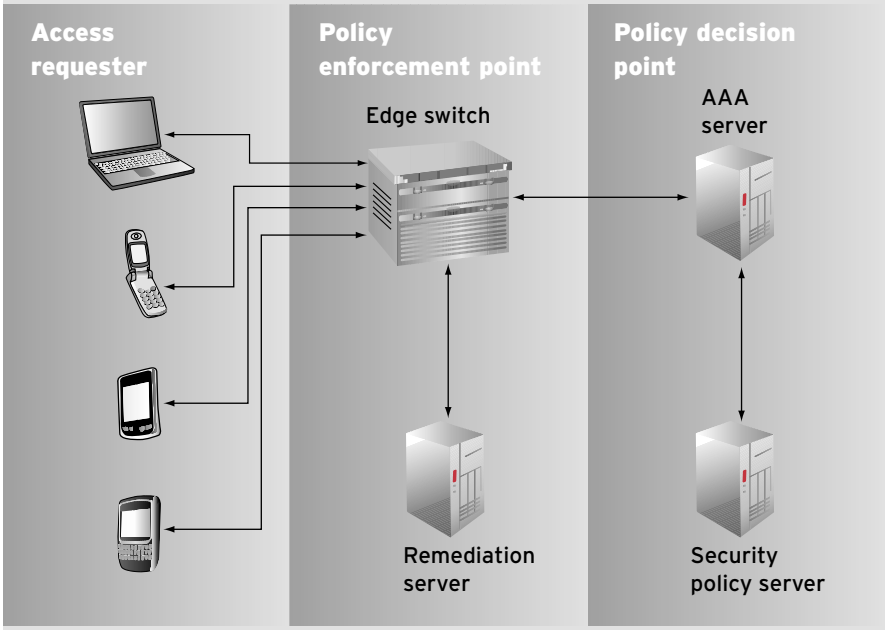

Endpoints are a major risk area both to users (individual subscribers or businesses) and to carriers. Mobile endpoints have diverse operating systems, are packed with the latest multimedia capabilities, and have multiple intrusion paths aside from the carrier's IP access – all of which contribute to increased security risks that aren't addressable from conventional network-based defenses.

Effective endpoint compliance ensures that devices – such as smartphones, PDAs, laptops, and BlackBerry devices – do not pose a security threat to either the network or the user. Typically, the carrier will enforce broad endpoint compliance policies to protect the network, and apply additional policies through security policy controls on individual subscribers and/or subscriber groups based on their needs. For example, devices can be inspected to ensure that they have various security features mandated by a corporate security policy, such as anti-virus programs, firewalls, intrusion detection and intrusion prevention systems, passwords, encryption, or the proper patch level in the case of an attack. Based on this information, the edge switch can either deny or grant access to the network based on the device's compliance to the policy controls. To enable this control, endpoint compliance client software must be installed on the device, either through download or at point of sale.

As shown in the diagram, access to the network is granted or denied by an edge switch, such as Nortel's Packet Data Serving Node (PDSN) or Gateway GPRS Support Node (GGSN), which is also responsible for assigning IP addresses and routing packets. The information on which to make the decision is obtained from an authorization, authentication and accounting (AAA) server, which consults a security policy server to determine whether the subscriber meets the security policy control and subscription requirements; that is, to confirm that the endpoint device has the proper security configuration and the subscriber has the right to access the requested information. If it is determined the subscriber or device does not adhere to network security policies, the edge switch can also interface with a remediation server to help the denied devices achieve compliance.

customers an integrated set of security features (including URL filtering and intrusion protection services) to prevent, detect, and combat attacks on their wireless devices.

Expected to be offered by operators as part of a bundled services package, hosted security services differ from managed security services in that they are centralized in the operator's network, rather than being distributed at customers' sites. Centralizing these services simplifies the operator's task of managing security software updates and other such tasks across the wide variety of mobile devices and wireless operating systems that currently populate the wireless space.

Nortel's hosted security suite (Figure 1) is being designed to incorporate a broad spectrum of security features, including:

• customer-visible security measures, such as anti-virus, anti-spam, URL filtering, and other protection mechanisms;

• background security measures, which look for specific signatures in the headers and run without the knowledge of the user. These measures include DoS and distributed DoS (DDoS) prevention, per-user firewall protection, and intrusion prevention/detection systems;

• "one-time" security services requested by the user – for example, remote automatic backup/restore of information such as directory lists and programs on the mobile device, or handset data reset to restore the device to an earlier backed-up state if it becomes infected with a virus.

This "peace of mind" hosted security suite sits on Nortel's industry-leading wireless gateways – the Gateway GPRS Support Node (GGSN) in GSM/UMTS packet networks and the Home Agent (HA) in CDMA networks. The GGSN and HA are ideal places to apply such advanced security services since, unlike standalone security appliances that would be positioned at peering points, they are aware of the subscriber's identity and policies, are responsible for generating billing information, and are able to act on both Internet-to-mobile and mobile-to-mobile communications.

Locating the security suite on the Nortel GGSN and HA products also takes advantage of their existing security mechanisms, including integrated state-aware subscriber firewalls, DoS/DDoS prevention, and anti-spoofing protection.

A state-aware subscriber firewall limits a subscriber's access based on local policies applied on a per-subscriber basis through IP packet filtering mechanisms. Stateful firewalls use a traffic inspection algorithm to keep track of the status of each connection and regulate traffic flows accordingly (typically, forward traffic, drop traffic, or drop and log traffic). Stateful firewall traffic inspection rules are based on a combination of examining such attributes as source IP address, destination IP address, and TCP/UDP port number.

DoS/DDoS attack prevention uses signature detection (verifying that there is a properly formed header in the message) to recognize an attack. The goal of DoS/DDoS attacks is not to gain unauthorized access to machines or data, but rather to prevent the legitimate users of a service from using it. A DoS/DDoS attack can come in many forms. Attackers may "flood" a network with large volumes of data or deliberately consume a scarce or limited resource, such as process control blocks or pending network connections. Attackers may also disrupt physical components of the network or manipulate data in transit, including encrypted data. DoS attacks are launched from a single node (mobile or otherwise) and continually ping a server, tying it up and turning it into a "zombie" that can't handle other tasks. DDoS attacks are launched simultaneously from many nodes.

Anti-spoofing protection prevents an unauthorized person from impersonating an authorized user to gain access to the network. The anti-spoofing policy ensures that packets are sent by the terminal with the source IP address originally assigned by the gateway. It also ensures that the packet received from the Internet does not have a source address that coincides with the set of addresses known to be located on the access side, which would normally be expected to come from a mobile device and not from the Internet.

URL and web content filtering controls access to websites based on the access privileges of the subscribers, and enables an operator to block website access based on configurable categories selected by the subscriber. The URL database in Nortel's solution is provided by Websense, a leading URL and web content filtering solutions provider (see sidebar on page 48). The major filtering categories are:

• security threats – sites that enable phishing, keylogger, malicious code, spyware, and hacking;

• adult material – sites that are deemed pornographic in nature and suitable for viewing only by adults; and

• undesirable content – sites that promote violence, drug use, gambling, or illegal activity, as examples.

Intrusion prevention systems (IPS) use techniques such as signature detection and, more importantly, anomaly detection to prevent security breaches in the network arising from hacker activity, trojans, protocol exploits, scan attacks, and worms. Sophisticated anomaly detection algorithms help prevent not only the onset of attacks where signatures are known but also new threats without having to wait for signatures to be developed, by correlating such events as network scanning activity that can precede an attack. An IPS based in the network provides more effective protection than intrusion protection only at the user device because there is more information available to detect the anomalous behavior – while an individual user device can see only what's hitting it, a network device can see that same event hitting

many users in sequence. Moreover, network-based security solutions allow the signature and anomaly databases that characterize the latest security threats to be kept current, ensuring that once a new threat is identified, preventative measures can be taken immediately and applied across the subscriber population. Nortel is partnering with third parties to provide these IPS capabilities.

In addition to these security measures, the Nortel GGSN and HA gateways support IPsec VPNs that provide strong triple DES (3DES) or AES (Advanced Encryption Standard) encryption. IPsec VPNs enable operators to offer cryptographically protected delivery of data traffic to the enterprise's network from users' mobile devices, allowing medium and small businesses to benefit from network-based VPN services with management of the VPN and authentication of subscribers outsourced to the network operator. With the GGSN and HA, each VPN has its own routing domain, which ensures that data is not passed between VPNs and allows the VPNs to manage overlapping address spaces.

Advanced security capabilities

In addition to developing the operator-hosted security suite to protect mobile devices, Nortel is also moving forward on a number of innovative technology solutions to protect wireless network infrastructure elements from attacks launched from the mobile devices.

As an overall strategy, Nortel advo-

Mobile IPsec: Secure wireless mobility made simple

by Ron Pon

Combining the proven security features of an IP Security (IPsec) VPN with full mobility, Nortel's ground-breaking Mobile IPsec technology offers a new way to keep IP sessions both secure and uninterrupted as users roam or change networks – a critical component of any online service where private or confidential information is exchanged.

Currently delivered on Nortel's VPN Router portfolio (formerly Contivity), Mobile IPsec is unique in the industry in that it requires no additional deployment of network infrastructure or client software, and almost no configuration on the part of the administrator. Other solutions involve extremely complex deployments of Mobile IP, or overlay solutions that require mobility servers and additional client-side software, resulting in higher operating costs, additional administration, and lower network performance. For its innovative application of IPsec technology in a mobile setting, the Mobile IPsec development team won a Nortel Technology Award of Excellence, which recognizes innovations that deliver clear customer value and contribute to Nortel's overall industry leadership. [IPsec is a set of protocols developed by the Internet Engineering Task Force (IETF) to support the secure exchange of packets at the IP layer. It has become one of the most trusted protocols for secure communications.]

Mobile IPsec is an ideal complement to mobile applications. As users roam from network to network, their VPN connection remains intact, providing uninterrupted access to applications and data. Because the VPN session is persistent, users do not have to log-on again to restart their VPN tunnel. Connectivity for any application is maintained, enabling multimedia, voice, video, email, and file transfer sessions to remain intact.

Mobile IPsec, in effect, mobilizes applications that may not otherwise work for the roaming user. In a wireless LAN (WLAN) environment, for example, users can roam between floors or buildings, docking and undocking their notebooks as they switch between wired and wireless networks, while still keeping their VPN connection alive. The persistent IPsec connection ensures that a user keeps the same virtual IP address even as the physical IP address changes on the notebook.

In a wireless WAN (WWAN) environment, Mobile IPsec enables users to maintain the same VPN session as they roam between GPRS/EDGE, CDMA, GSM/UMTS, and other wireless/wireline networks. What's more, users can maintain their VPN session even when moving between WLAN and WWAN environments.

Should the VPN tunnel fail, Mobile IPsec includes the capability to transfer the session to an alternate switch without requiring the user to log-on and authenticate again. This capability could be used, for example, where an enterprise uses different internal and external software requirement sets (SRSs) – rules that define security policy requirements on the network – to secure its WLAN access and remote access, respectively. It would enable a user session to be transferred from the internal SRS to the external SRS when roaming from the WLAN network to the WWAN network, as well as allow the user to roam between WLANs connected to different internal SRSs across a large campus.

Through its innovative application of technology in Mobile IPsec, Nortel is providing new capabilities that work toward making our increasingly mobile world safe and easy.

Ron Pon *is Senior Network Security Architect in the CTO Office, and the original inventor of this technology. Ron also leads Nortel's participation in the Trusted Computing Group, an organization driving what could well be the next revolution in secure computing.*

cates a layered defense approach that is designed to provide multiple levels of protection as traffic traverses different network layers and domains. (In this context, "network layers" do not refer to the ISO protocol stack layers.) While Internet-sourced attacks on mobile devices, the carrier infrastructure, and on the Internet/carrier data network boundary must also be considered in designing a secure network, they are well covered by familiar security solutions deployed at network perimeters (see article on page 28) and are not discussed in this article, which focuses on the issues specific to mobility networks.

The first layer of defense is at the access edge layer, where the goal is to protect the wireless operator's network assets from offending mobile devices and software by enforcing security policies before network connections are allowed. Threat protection at this layer must effectively contend with high-speed, low-latency, rapid-transaction connections, as well as with the hundreds of different mobile device types ubiquitously accessing the network – each potentially having slightly different application-usage and threat-vector profiles. At the device level, effective threat protection uses Nortel's endpoint compliance architecture paired with various endpoint protection capabilities (such as Nortel services edge routers' firewalls, anti-virus protection, automated OS and application patching, and device passwords).

Since all traffic in the network passes through the access edge, Nortel provides additional protection within the carrier's network by deploying a Nortel firewall at the services edge, and by introducing intrusion detection/prevention along the subscriber services path via Nortel's Threat Protection System (see page 35). Nortel can further enhance access edge security by engineering network address translation (NAT) and endpoint compliance strategies into the overall network design.

Here, the combination of Nortel's security expertise and awareness of the network infrastructure gives operators the ability to overcome unique security and regulatory complications as voice and data services converge. For example, as mobile Voice over IP (VoIP) implementations expand, the need to protect the network infrastructure from data-enabled threats increases. Protection deployments, however, must also take into account the requirements of such voice services as emergency 911 calls, where service cannot be denied under any circumstances, irrespective of associated threat vectors.

The second layer of defense is at the services edge which, because of its position adjacent to the Internet and carrier application layers, is the main threat-blocking point in the network architecture for IP-based traffic.

Effective protection at the services edge requires stateful awareness of all subscriber sessions and traffic in order to control traffic into and out of the Internet domain. However, the benefits that operators receive from the deployment of public network access in terms of lower capital outlays and reduced operating expenses must be balanced against the increased security exposures associated with forgoing the use of trusted networks for interfacing to an Internet point of presence. Fortunately, not all traffic from the access edge crosses the services edge, and also there are fewer mobile device types with connections crossing the services edge. As a result, carriers have the opportunity to tailor protection profiles to a limited number of specific services edge device types, allowing them to focus their security defense architecture.

At this level, deploying a combination of dynamic stateful firewalls and intrusion prevention and intrusion detection systems, such as Nortel's Threat Protection System, along the subscriber services path (for anomaly detection) will protect the carrier's network against port scans, worms, trojan horses, viruses, and DoS attacks. To provide a fuller defense, NAT, Layer 2 traffic management services, anti-virus scanning, content filtering, anti-spam blocking, and IPsec implementations can be added. For example, a firewall policy that allows mobile-initiated services to access the wireless network could be combined with a threat protection system that is allowed to close the associated ports on command – providing the operator with services flexibility, low administrative overhead for services provisioning, and dynamic protection against known (signatured) and zero-day (anomalous) attacks.

The third layer of defense is at the application point of presence. Because this layer processes only application-specific traffic, security measures are tailored to specific threats. Specific authentication and authorization rules, combined with traffic protection at the services edge, provide a comprehensive threat capture and prevention mechanism. As voice and multimedia services converge onto packet-based networks, they are treated like other data applications, such as email and HTTP web browsing. These converged services require strict authentication, authorization, and traffic protection measures (protocol correctness, buffer overflow prevention, and distinction between control and content). These requirements become more challenging in a mobile environment where there is movement between network domains and between IP addresses, and where accessibility through roaming networks must be maintained.

The fourth layer of defense is a restricted defense center zone – a private network maintained by the wireless operator's security operations center – where mission-critical services [such as home location registers, home subscriber services (HSS), domain name system (DNS), and authorization, authentication, and accounting (AAA)] as well as private and sensitive data are protected by maintaining separation between application and bearer traffic. Host-based intrusion protection at the server level, combined with tight coupling to endpoint compliance strategies, provide the most effective security

protection at this level.

Using integrated network designs and implementations, Nortel's security solutions for wireless networks can identify security threats and take appropriate actions to counter the threat. By leveraging Nortel's awareness of the network, offending devices can be selectively and effectively quarantined. The combination of Nortel's wireless services edge routing products, standard authorization products, and Nortel's Threat Protection System enables threats along the subscriber services path to be detected, notifications sent to the authentication servers, and offending data sessions redirected and suspended while remedial action is taken by the operator.

Infrastructure security: management plane

Securing the management plane is also a critical part of an end-to-end security solution. Network management nodes contain management policies and databases that are critical to the operation of the network.

To protect this part of the network, Nortel is working to ensure that its wireless products comply with Nortel's company-wide set of baseline security requirements (see article on page 20). These baseline security requirements – covering everything from platform and OS hardening (turning off unused services, closing unused ports, etc.), to strong authentication and encryption capabilities, and support for a security audit trail – apply to all types of network-connected devices, including mobile devices, network infrastructure elements, and application servers. For example, platform hardening on a smartphone may include staying up to date with security patches and running a mobile Internet security solution that includes a firewall, anti-virus, and intrusion prevention technologies. Implementing these Nortel baseline measures significantly reduces the risk of a security breech for wireless customers, and demonstrates Nortel's commitment to providing secure network solutions.

With the growth of wireless broadband services, mobile network operators are facing similar security challenges as their wireline counterparts. However, the wireless security threat vectors are broader and require specific solutions that cannot be delivered without in-depth knowledge of wireless network architectures and operations.

Nortel is leveraging both its security expertise coming from the enterprise and wireline businesses and its extensive field experience in designing, building, and optimizing mobile networks. The combination of these strengths provides a portfolio of security solutions and services tailored for mobile operators, encompassing a suite of GGSN- and PDSN/HA-hosted security services to protect wireless end users, a wireless-specific layered defense strategy to protect the network infrastructure itself, and baseline security measures to protect the management plane of the wireless network from both internal and external attacks. ■

Frédéric Bastien *is Product Line Manager Leader for GSM/UMTS Systems.*
John Garrison *is Product Line Manager for CDMA Packet Data Solutions.*
Don Keeler *is Product Line Manager for GSM/UMTS Core Network Evolution.*
Emily Nichols *is Product Line Manager for GSM/UMTS System Security.*
Paul Tse *is a Practice Advisor in Security Professional Services.*

Secure information sharing for the U.S. Government

by Thomas Casey, Alan Harbitter, Margaret Leary, and Ian Martin

Effective sharing of information among disparate communities at all levels of government – national, state, regional, and local – has become a top priority of the U.S. Government, whose leaders repeatedly assert that more effective information sharing is a prerequisite to advance homeland security efforts. While a common assumption is that the prime barriers to sharing information stem from policy and cultural issues, a complex technology challenge must also be tackled: security. Nortel's U.S. Federal Government-focused subsidiary, Nortel Government Solutions, offers a wide range of information assurance solutions that enable secure information sharing for its government clients. While this article describes secure technology solutions within the context of U.S. Government initiatives, these approaches have broader applicability for governments worldwide, as well as for other large organizations and enterprises around the world.

More effective information sharing is a high-priority concern of agencies in all sectors of the U.S. Government. In fact, it has become one of the highest priority issues in health care, public safety, and homeland security endeavors.

Without a doubt, strong motivation, commitment, and focus exists throughout the U.S. Government to link disparate networks and data repositories to enable more effective sharing of information – both horizontally between peer departments and agencies, such as public safety agencies, the military, customs and immigration, and intelligence departments; and vertically among local, regional, state, and federal entities.

One key barrier to establishing secure information-sharing networks can be described as cultural – grappling with attitudes among some information owners that holding information protects it and doesn't threaten their control. To overcome this obstacle, a sense of trust must be fostered among sharing parties.

Equally important, however, is the need to implement the technologies, policies, and controls necessary to establish "electronic trust" – that is, to ensure that information-sharing networks can provide a communications environment in which the sharing parties are confident that the information being sent electronically not only is protected from interception and unauthorized disclosure, but also has originated from a known and trusted source.

Addressing this challenge requires system designers to thoroughly understand the nature of, and prerequisites for, security for information-sharing networks. It also involves the integration of virtually every information security tool available today, as well as knowledge of the many new technologies and processes on the horizon.

Nortel Government Solutions (formerly Nortel PEC Solutions) has a long history of working with the U.S. Government market in planning, developing, and supporting customized mission-critical solutions that use a wide range of advanced information security technologies. For instance, Nortel has designed and implemented secure network systems for the U.S. Departments of Defense, Homeland Security, Justice, Treasury, and Veterans Affairs. (Many of the security technologies discussed in this article are core capabilities of Nortel Government Solutions.)

The information-sharing challenge

A secure and trusted information-sharing environment is a prerequisite to enabling users to interact with and share information easily and seamlessly across many different networks and databases nationwide. This capability can significantly improve the effectiveness of many functions, such as intelligence gathering and public safety efforts. The following scenarios illustrate the benefits of being able to tie together information from a variety of sources and enable secure interactions between users and applications, between applications themselves, and among different users. These scenarios also highlight the key capabilities that a secure information-sharing infrastructure must provide, including proper identification and authorization of all parties, and the secure transfer and storage of information.

User-to-application scenario: Consider, for example, a local law enforcement officer at a standard traffic stop. Basic protocol dictates that the officer request and verify the individual's driver's license and vehicle registration. However, the officer could also check a wide range of other computer applications, such as immigration databases, terrorist watch lists, criminal information and intelligence repositories, and counter-drug intelligence databases that may be owned by external organizations, such as the U.S. Federal Bureau of Investigation, the Drug

Enforcement Administration, and the Department of Homeland Security. To do this, these externally owned applications must be able to recognize the officer in order to determine if he or she has the correct credentials to receive the information. Then, the information, which is likely to be sensitive from an intelligence and privacy perspective, must be secured while in transit. Finally, the device on which the officer receives the information must be able to store that information securely.

Application-to-application scenario: Using the same traffic stop scenario, it might be more time-efficient for the officer to interact with a single application that can then reach out to disparate sources to gather all the requisite information. The owners, then, of all the applications that are sending and receiving the information need to establish agreements, which include the process by which users are authorized and cleared, how sensitive information is handled, and how information flow is audited. These agreements would be negotiated and monitored in order to establish electronic trust among the owners of the information-sharing applications. As with the first scenario, the requirements for secure sharing include identification and protection of information both during transit and while in storage.

User-to-user scenario: Text messaging and the sharing of multimedia information, including video and images, are now becoming commonplace in mission-critical government applications. For these transactions to be secure, each party must have the ability to positively identify peers and to access secure communications channels to exchange information.

Recognizing that nationwide electronic trust is critical to enabling these information-sharing capabilities, the U.S. Government has passed significant legislation and has set mandates surrounding security requirements (page 56).

Applying technology to the information-sharing challenge

Information security is best addressed from a holistic, architectural perspective, which traditionally first visualizes the universe of security services through three fundamental areas: Confidentiality, Integrity, and Availability, sometimes represented by the mnemonic "CIA." A complete security architecture provides services that address all three areas:

- Confidentiality – preserves authorized restrictions on access and disclosure, including means for protecting personal privacy and proprietary information.
- Integrity – guards against improper information modification or destruction, and ensures information non-repudiation and authenticity.
- Availability – ensures timely and reliable access to, and use of, information.

Nortel's Network Security Architecture (page 7) addresses all three areas at all domains in the network (application, network, and device domains). Several security technologies are used to address these areas, as Figure 1 highlights. While not all of these technologies are addressed directly in this article, two of the newer technologies are highlighted, specifically public key infrastructures (PKIs) on page 58 and federated identity management on page 60 – capabilities that are part of the Phase Three evolution of Nortel's Network Security Architecture.

To demonstrate more clearly how selected technologies secure information

Figure 1. Elements of a comprehensive information-sharing security framework

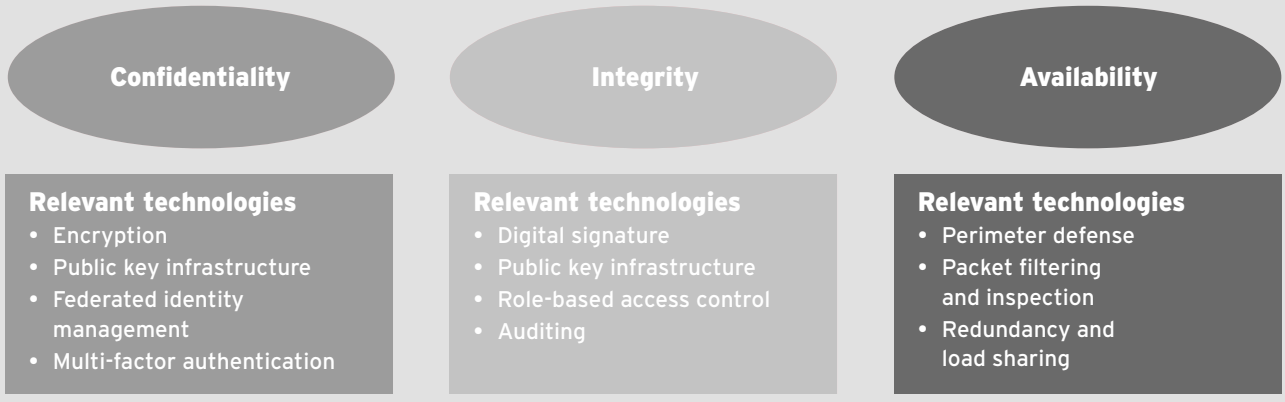

A comprehensive information-sharing security framework must: preserve authorized restrictions on access and disclosure, including the protection of personal privacy and proprietary information (confidentiality); guard against improper information modification or destruction, including ensuring information non-repudiation and authenticity (integrity); and ensure timely and reliable access to, and use of, information (availability). Several security technologies are used to address each area, as shown in the diagram.

sharing, we discuss three use-case scenarios: user-to-application, application-to-application, and user-to-user.

User-to-application: The user-to-application use-case scenario, shown in Figure 2, illustrates how users would interact with identification and authentication services in a large network with multiple information providers. These authentication and credential verification steps follow the general sequence used in a federated identity architecture and support information sharing on a large scale.

Step 1: As Figure 2 shows, the use-case begins with a police officer (subscriber) authenticating to her home network – the city's local police department (acting as the credential service provider, or CSP). The authentication and management of the officer is the police department network's responsibility. As the officer joins, leaves, or changes status, the police network's credential service (CS) tracks her status and adjusts her credentials accordingly. If, for example, the officer's clearance status changes to allow her to view classified intelligence information, her local identity server would add that new credential to her profile.

Step 2: The officer requests access to information that is owned by an organization other than her police department. Police department application software, operating on the officer's behalf, recognizes that the requested information is stored in an agency application (AA), which is managed by

U.S. Government initiatives for secure information sharing

Over the past few years, the U.S. Government has passed significant legislation and set mandates to put in place the nationwide "electronic trust" capability that is critical for secure information sharing. These initiatives include:

FISMA: The Federal Information Security Management Act of 2002 is the primary legislation that governs federal information security in the U.S. FISMA placed the responsibility for information security with the head of each Federal Agency, and called for the National Institute of Standards and Technologies (NIST) to develop the Federal Information Processing Standards (FIPS) and other Special Publications (SP) guidelines. In response, NIST has published standards for security categorization (FIPS 199), minimum security requirements (FIPS 200 draft), recommended security controls (SP 800-53), and security certification and accreditation (SP 800-37).

HSPD-12: On August 27, 2004, President Bush issued the Homeland Security Presidential Directive (HSPD)-12, which directed the creation of a common identification standard for all federal employees, contractors, and affiliates. Citing the risks of wide variations in identity proofing and identity management, the directive mandated that NIST promulgate standards for identity-proofing personnel as well as technical requirements for a common smart-card-based Personal Identity Verification (PIV) system. Within the year, NIST developed the FIPS 201 document, which establishes identification standards for all Federal employees and contractors who require physical access to Federal facilities and logical access to Federal information systems.

The NIST standard consists of two parts: PIV-I addresses identity proofing requirements, and PIV-II specifies the technical requirements for identity credentials. NIST provides further related technical implementation guidance for smart-card architecture and interfaces, in SP 800-73 (Interfaces for PIV), and for biometric interfaces in SP 800-76 (Biometric Data Specification for PIV). Federal agencies are expected to begin implementing PIV-I no later than October 25, 2005, with the smart-card rollout starting as soon as late 2006.

Real ID Act: Having addressed identity management for online government applications and for its employees and contractors, the Real ID Act, passed May 11, 2005, assigns additional responsibilities for identity management to the Department of Homeland Security. The intent of this Act is to allow the DHS to set minimum issuance standards for state-issued driver's licenses and identification documents. Minimum document requirements specified in the bill include name, date of birth, gender, license number, digital photograph, address, and signature. The bill also specifies that the cards be read by a common machine-readable technology. While the Act does not require individual States to adhere to the Department of Transportation's driver's license specifications, the fact that citizens will be unable to use non-compliant licenses for federal purposes (including travel on a U.S. airplane, opening a bank account, or collecting social security payments) provides incentives for the individual States to comply.

HIPAA: Recognizing the need to encourage the widespread use of electronic information sharing in healthcare, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires, among other provisions, that the Department of Health and Human Services establish national standards for electronic healthcare transactions and address the security and privacy of health data. HIPAA established minimum security standards that healthcare organizations must follow when maintaining, storing, or processing healthcare information. Additionally, the HIPAA privacy regulations represent the first national standards that address protecting the privacy of individual health data, directing that organizations entrusted with healthcare information protect it against intentional or inadvertent misuse or disclosure.

the (fictional) Department of Shared Information (DSI).

Step 3: At this point, the officer must make her identity known to the DSI network. Software running on her behalf passes an assertion to the DSI stating that the police department, using means accepted by the DSI as sufficiently secure, has authenticated the officer. The assertion also contains identity information about a subscriber and verified attributes (such as role and job function). Assertions are typically digitally signed objects obtained from a trusted source by a protocol such as Security Assertion Markup Language (SAML). Because the local police department and the DSI are not co-located and the nature of the data is sensitive, secure network interconnectivity (e.g., a VPN) must be employed to ensure the confidentiality of data passed between the two enterprise networks.

Step 4: In order to provide the information requested by the officer, the DSI may, depending on the type and sensitivity of information, need more electronic credentialing information. The officer may not explicitly know about this information (such as access privilege specifications meaningful to DSI), but this credentialing information is stored in the CS computer system managed by the police department.

While verifying the officer's credentials, the DSI server may also need to verify digital signatures – perhaps the police department's digital signature on the officer's authentication information or her credentials. The digital signature, supported by a public key infrastructure (PKI), is a tool that can be used to establish electronic trust among the officer, the police department network, and the DSI system.

Step 5: Once the DSI's AA is convinced that the officer has been properly authenticated and carries the valid and required credentials, the DSI is confident that the officer is an authorized party and can transmit the requested intelligence information to the officer.

Application-to-application: In an application-to-application use-case, there is generally no exchange of identity information for individual users. Rather, it is the computer systems of two or more organizations that exchange information. When the DSI shares information with the police department network, the DSI must trust that the police department will treat this information within agreed-upon security guidelines and provide the information only to authorized users. It is

Figure 2. The concept for user-to-application identity and authentication

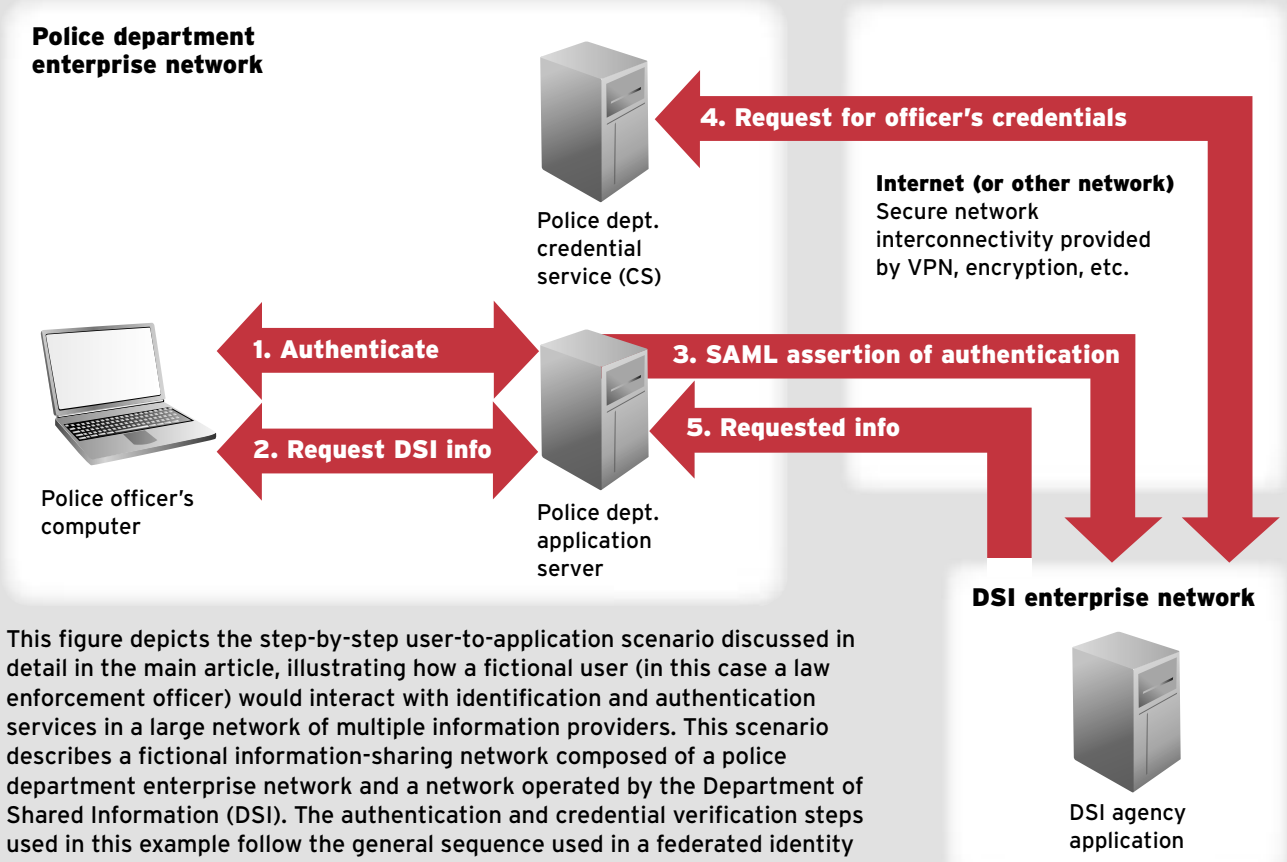

This figure depicts the step-by-step user-to-application scenario discussed in detail in the main article, illustrating how a fictional user (in this case a law enforcement officer) would interact with identification and authentication services in a large network of multiple information providers. This scenario describes a fictional information-sharing network composed of a police department enterprise network and a network operated by the Department of Shared Information (DSI). The authentication and credential verification steps used in this example follow the general sequence used in a federated identity architecture and support information sharing on a large scale.

Public key infrastructures

A public key infrastructure (PKI) is an aggregate of technology, people, policies, and procedures that work collectively to provide a managed trust framework and to facilitate the secure transfer and validation of information between two or more parties. A PKI can establish relationships with other PKIs and work together to create a federated PKI model supporting inter-organization managed trusts agreements.

Encryption

The technology at the heart of public key infrastructures is public key cryptography (also known as asymmetric cryptography), a modern branch of cryptography that uses mathematical processes, or algorithms, to transform information to/from plaintext and ciphertext (encrypt and decrypt information).

In asymmetric cryptography, pairs of cryptographic keys – a public key and a private key – are used to encrypt and decrypt data; information encrypted with one key can only be decrypted with the other, and vice versa. Typically, asymmetric key pairs are issued to users. From this pair, the public key is made known to everyone, and the private key is made known (or accessible) only to the user. While the private and public keys are mathematically related, it is not computationally feasible to compute the private key from the public key. This allows a private message to be sent to any party by encrypting it with his or her public key.

Asymmetric cryptography algorithms are not as efficient for bulk data encryption as symmetric cryptography algorithms. Symmetric encryption algorithms employ a single, shared secret key to encrypt and decrypt data. In many implementations, asymmetric cryptography is used to distribute the symmetric (also known as "session") key. In this case, the sender would encrypt a message with a symmetric key and encrypt the symmetric key with the recipient's public key. The message recipient could then decrypt the symmetric key with his or her private key, and then decrypt the message with the symmetric key.

Digital signature

In addition to encrypting information, public key cryptography is used to create digital signatures. A digital signature in the electronic world has the same purpose as a handwritten signature in the paper world – it provides non-repudiation. Through a paper-based handwriting analysis or an electronic public key cryptography operation, the signer of a document cannot convincingly deny having signed the document. The digital signature process adds another level of integrity by ensuring that the contents of the digitally signed document have not been changed. The "signing" process involves encrypting a digest of the document with his or her private key. (A digest is produced by processing the document through a one-way hash function. The hashing process is such that it is extremely difficult to create a different document that would produce the same digest.) The digital signature verification process requires the verifier of the document to decrypt the digest with the signer's public key (verifies the signer), and then match this digest to his or her own computed digest of the document (verifies document integrity).

Key functions of public key infrastructures

Both PKI applications described above, encryption and digital signature, require the distribution of trusted public keys. For example, a person verifying a digital signature needs to be absolutely sure that they are using the signer's public key for the verification process. To establish this trust, public keys are published in digital certificates, which contain the user's public key along with other information, such as name, issuing authority, and expiration date, and are digitally signed by a Certificate Authority (CA), described below. To maintain the trust in these digital certificates, PKIs carry out several important functions, including identity proofing, issuing and revoking certificates, and protecting private keys.

- *Identity proofing:* The process for obtaining certificates must include an identity proofing component, which is generally handled by a Registration Authority (RA) that is responsible for the identification and authentication of certificate subscribers before certificates are issued. The RA does not sign or issue the certificates. Once the RA has established an individual's identity, the certificate request is forwarded to the CA for certificate generation.
- *Issuing and revoking certificates:* The CA serves as a trusted third-party that certifies the identity of its subscribers, issues digital certificates to these users, and manages the lifecycle of public key certifications, from issuance through to revocation and/or expiration.

To issue a certificate, the CA validates the certificate request (proof of possession through a challenge-response protocol) and creates and signs the certificate using the defined X.509 certificate profile. [X.509 is a public key infrastructure standard, defined by the International Telecommunication Union (ITU-T), that specifies formats for public key certificates.] The subscriber is then notified that the certificate has successfully been created. Authenticity of the published digital certificates is established by digitally signing each certificate with the CA's private key. Relying parties can verify the authenticity of a digital certificate using the CA's public key.

If a certificate's private key is compromised, the subscriber contacts the CA. After the identity has been validated, the CA publishes the associated certificate serial number to a certificate revocation list (CRL), which is a list of the serial numbers of certificates that have been revoked or

are no longer valid. As part of the digital signature verification process, a recipient of a signed message should check the CRL to determine if the certificate is still valid.

The CA is also responsible for publishing a list of all certificate serial numbers, along with the revocation information in the CRLs, to a repository service, which stores and makes publicly available the subject/subscriber public verification certificates, encryption certificates, and revocations lists. The repository supports the retrieval of X.509 PKI information. Typically, CAs publish PKI information to repositories using Lightweight Directory Access Protocol (LDAP). The repository is responsible for the availability of such information, and can be designed to replicate to multiple nodes supporting the defined security and availability requirements.

The CRL is used for validating certificate status. The CRL may contain end entity (user), certification authority, and cross certificates. A CRL will provide the certificate's serial number (unique within the CA security domain), revocation date, and revocation reason code. Another way to validate certificate status is via the Online Certificate Status Protocol (OCSP), a protocol defined by the Internet Engineering Task Force (IETF) RFC 2560, which allows a client/relying party to request the status of a certificate without the need to download a CRL. Because OCSP typically provides more timely revocation information than is possible with CRLs, this protocol is used in industries that participate in high-value and/or highly sensitive transactions (e.g., the banking and financial industries). OCSP can also be leveraged in mobile environments where bandwidth is at a premium and client-side processing power is constrained.

• *Protecting private keys:* Private keys must be secured from malicious compromise by using a layered security posture that includes physical, logical, and personnel protections. Key protection applies to all levels within the PKI, to include CA private key protection, application/device private key protection, and subscriber private key protection. The NIST FIPS 140-2 standard defines four levels of cryptographic module implementations, ranging from software implementations to physical security requirements to include tamper-evident, tamper-resistant, multi-level role-based authentication, and identity-based authentication.

PKI enablement

Once a PKI has been established, applications, operating systems, and devices must be configured to take advantage of the security capabilities enabled by the PKI. This configuration requires the integration of certificates for authentication, encryption, and digital signatures. Many applications have this capability built in; others require an external application to support this functionality. PKIs can provide support for a variety of business applications and platforms, including: web-based, legacy, customer relationship management, and enterprise resource planning applications; Windows, UNIX, and LINUX platforms; and personal digital assistants, virtual private networks, and mobile/smart phones.

likely that the DSI does not know about the police network's users and therefore must trust the police network to manage and authenticate users in a manner that is consistent with the DSI's security policies.

While the DSI may have little or no knowledge of users' identities, many information-sharing applications have a definitive requirement to track access to information on a transaction-by-transaction basis. The DSI may later go back to the police department and ask: "Remember that information we gave you yesterday? We would now like a list of the individuals who had access to it, and if and when they did so." This requirement places an auditing responsibility on the police network – a responsibility that it may be required to assume as a condition of DSI sharing its information.

When the number of applications exchanging information is small, it is relatively easy for the DSI to keep track of external requests and exchanges. However, the DSI is always required to authenticate that requests come from trusted systems. The DSI will most likely accomplish that through digital signature verification or asymmetric key-based authentication. To accomplish this verification and authentication, the police network and the DSI must either participate in the same public key infrastructure (PKI) or in separate PKIs that can interoperate and exchange compatible digital certificates (i.e., cross-certify).

User-to-user: In a user-to-user case, there may be limited involvement from an applications server, and communications takes place directly between users. For example, a common information-exchange application, instant messaging, operates to a large extent on a peer-to-peer basis. In this case, an individual relying party must be able to retrieve a subscriber's digital credentials directly from the credential service (CS). To accomplish this, the CS must be able to accept individual, user-sourced requests for credentials, verify the authenticity of the requests, and provide the necessary

Federated identity management

Strong authentication and identity management practices are vital for government information-sharing networks, and lack of strong practices can significantly hamper national security efforts.

Weak authentication controls could enable malicious users, for instance, to access "breeder documents," or documents that are used to obtain other documents for identity – such as drivers' licenses, social security cards, and birth certificates. With these documents, long-term identities can be established and maintained.

Authentication and identity processes, of course, rely on the protection afforded by public key infrastructures, and can also be strengthened through the use of such technologies as multi-factor authentication, which is the combination of "something you know (e.g., a password), something you have (e.g., a smart card), and something you are (e.g., a fingerprint or retinal scan). Multi-factor authentication is becoming more prevalent as the various associated technologies become more cost-effective.

In addition, a new paradigm for authentication is emerging for large-scale information-sharing networks. In recent years, a new concept, called federated identity management, has risen out of the emergence of distributed web services that have enabled the potential for large-scale e-commerce and e-government applications.

Federated identity management involves multiple entities, such as different government departments, organizations, and companies, entering into an agreement – a federation – to employ common authentication and identity management practices and to recognize acceptable identity credentials, thereby establishing mutual trust in each others' authentication capabilities.

In this way, a person's electronic identity can be moved easily and transparently across the various network elements (platforms, devices, databases), applications, and network boundaries of federation members, allowing subscribers to use a single sign-on and enabling governments and businesses to extend their security perimeters to trusted partners.

A federated identity architecture, then, provides a trust infrastructure that allows a user's identity information to be managed separately from the domain that provides access to the application.

Federations offer many benefits to their members and users, including cost economies realized through the acceptance of credentials issued by authorized credential service providers (CSPs), which eliminates the need for local identity enrollment or management services.

The authentication transparency afforded by the federation also improves the user's experience. Similar in concept to "single sign-on" services used within the traditional distributed networking environments, identity credentials are passed transparently between sites within the federated environment. Once an identity is authenticated and the user has been enrolled and issued a credential, that credential is accepted by any participant within the federation.

Figure 2 in the main article illustrates how a federated identity management architecture would function in the context of a large government information-sharing network. This vision is consistent with the NIST Electronic Authentication Guideline, and shows how users interact with identity services in a large network with other information services.

Methods to enroll users in federations vary widely, but are usually formalized. In some cases, users are directed to an approved CSP that validates the user's identity and provides a credential, such as a PIN/password, digital certificate, or smart card. Other federations may allow users to be automatically enrolled simply as a part of another process. The level of confidence in presented identity credentials and the subsequently afforded transitive trust among federation participants relies almost exclusively on the standardization of strong identity management processes.

To this end, the U.S. Government has adopted comprehensive identity management guidelines within its E-Authentication Framework and the NIST PIV standards (see sidebar, page 56). Under these guidelines, the General Services Administration's E-Authentication Portal (EAP) was developed to meet the U.S. President's Management Agenda of providing online access for citizens and businesses to interact with government. The EAP enables the sharing of identity credentials across autonomous government agency domains. The EAP accommodates both assertion-based credentials (PIN/password) and digital certificate-based authentication technologies. Transactions can begin at either the U.S. Government's official web portal (First.gov), with the CSP or directly at the agency's application. The EAP eventually will be technically interoperable with a wide variety of identity architectures and will support emerging specifications and technologies.

Indeed, several identity management frameworks and technology standards are beginning to take shape. For instance:

• The Liberty Alliance Project is an alliance of more than 150 companies and nonprofit and government organizations seeking to develop open standards for federated identity management. The core of the Liberty Alliance standard is an emerging protocol called Security Assertion Markup Language (SAML), which was defined by the Organization for the Advancement of Structured Information Standards (OASIS) Security Services Technical Committee. SAML is an XML-based framework and is used to pass security assertions over the Internet protocol HTTP. SAML enables such transactions as passing authentication and digital credentialing information. SAML release 2.0, with a significant

increase in function, was recently released.

• WS-Federation: Originally developed by Microsoft and IBM, but now primarily supported by Microsoft and a small group of technology vendors, the WS-Federation supports a competing protocol for identity federation that does not use SAML. There is wide speculation that there will be some convergence between these standards and those prescribed by the Liberty Alliance.

• Shibboleth: A project of Internet2/MACE, Shibboleth is a Federation framework that supports the higher education community. In this framework, campuses serve as Identity Providers, providing attribute assertions about the user to the service provider's (usually the application owner's) site. Also standards-based, Shibboleth uses OpenSAML, as developed by the OASIS Security Services Technical Committee. As with the Liberty Alliance standards, it is expected that Shibboleth will migrate support to SAML 2.0.

credentials to support secure information exchange. These credentials will include the necessary public cryptographic keys to establish secure communications between peers. For example, the credentials may include a digital certificate for the users that allows them to establish a session key and exchange information securely through a symmetric encryption algorithm.

Looking to the future

Over the next five to ten years, as various industry players continue to move technologies and standards forward, the "to-be" model for secure information-sharing networks will move closer to reality. In this model – as depicted in Phase Three of Nortel's Network Security Architecture – federations in both government and commerce will become more commonplace, and trusted information-sharing networks widely implemented. In fact, Nortel Government Solutions is working with its government clients to define some of the key elements of this vision, including:

• *Standards for digital credentials in each government community of interest:* Public safety agencies, healthcare officials, and homeland security officials will all have XML-formatted digital credentials that establish their identity, organizational affiliation, and authorities. This credential will be trustable and honored by all peer organizations for the purpose of authorizing access to shared information repositories.

• *A standardized federated identity architecture:* Identity servers and authorized credential service providers will be distributed throughout each information-sharing community of interest. These entities will provide such federation-wide functions as credential issuance and single sign-on.

• *Interoperable PKIs:* To enable high-assurance confidentiality and integrity, all government PKIs will be interoperable, capable of exchanging digital certificates and managing certificate information through mechanisms such as certificate revocation lists.

• *Standardized MOUs:* When two organizations agree to share information, they typically put in place a memorandum of understanding (MOU) that specifies, among other items, the security policies and practices that they agree to follow. To facilitate these agreements, a government-wide standardized MOU will exist that goes beyond one-to-one sharing between two organizations and is extensible to many-to-many sharing situations. The MOU will reference standardized evaluation and rating methodologies. For example, it would be extremely helpful if an organization that owned information participated in a rating system that allowed the following statement to be made: "We are a security level 7 organization. We will share our least-sensitive information with organizations that qualify at the level of 4 through 6. We will share our most sensitive and more valuable information with organizations that qualify at level 7 and above." Several possible sources from organizations such as the National Institute of Standards and Technologies (NIST) currently exist for developing such a rating system.

In realizing this vision, Nortel Government Solutions has broad experience in implementing many of the information-assurance technologies described in this article, and is actively engaged in promoting and implementing secure information sharing both horizontally and vertically for its government clients. ■

Thomas Casey *is Manager, PKI Engineering Department, Nortel Government Solutions.*
Dr. Alan Harbitter *is Chief Technology Officer, Nortel Government Solutions.*
Margaret Leary *is Manager, Security Policy Department, Nortel Government Solutions.*
Ian Martin *is Vice-President, Information Assurance and Engineering Division, Nortel Government Solutions.*

The new security frontier: Moving toward open, trust-based networks

An interview with Rod Wallace

Throughout the past decade, Nortel's advanced research team has played a key role in shaping the technology direction for network security, working to embed carrier-grade, robust, and cost-effective security solutions into every layer of the network. Rod Wallace, in his role as leader of Nortel's advanced security technologies and now in his new position driving security-related professional services for the company's enterprise and carrier customers, shares his perspectives on how network-driven security capabilities will be fundamental to unlocking new value propositions for converged multimedia networks going forward – and vital for realizing the open, secure, trust-based communications environments envisioned for the future.

From your perspective, where does the industry stand at this moment in the evolution of network security?

Over the past several years, network security has progressed through several important evolutions. First, we were focused on securing the physical points in the network – protecting the infrastructure and traffic, creating a secure path over which to communicate, securing the bits of information being transferred across that path, and hardening all the hardware and software used along the way. Then the focus was put on securing the different applications being used.

Now, we're working to create networks that protect and defend themselves autonomically – a critical development that will help ease the burden on end users. There is a wide gap between the number of people who have properly secured systems and those who don't – a gap that exists because too much onus rests on the person at home to invest their time and money, as well as to gain the knowledge needed to properly install the software protections, firewalls, and other security measures. We need to remove much of that burden from end users, by enabling service providers to take care of the security needs with autonomic networks (see page 7).

Once that is done, we can start to add other capabilities, such as authentication and identity management, on a truly network-wide basis. What we're seeing is the application, over time, of the fundamental buildings blocks that will ultimately lead to a virtual world that is trusted, where individuals and businesses will have the confidence that the network is a safe and secure place where they can share personal information and conduct business.

When this happens, new business models that forge online trust relationships will emerge, ubiquitous e-commerce will take off, and digital multimedia content will become even richer. Essential to this vision, though, is the ability for networks to establish electronic trust.

How do you establish electronic trust?

Trust has historically been something associated with people. You either trust people or you don't. The concept of electronic trust is similar: you either trust the network, device, or application to do what it is supposed to do, or you don't. And you need to trust that users and devices on the network are who they say they are. The electronic version of trust is enabled by a raft of different technologies and policies, among them multifactor authentication, en-

Rod Wallace, Leader of Security Solutions and Services, in one of Nortel's advanced research labs, where designers are working on innovative new security technologies.

cryption and digital signatures, identity management, and information rights management, to name a few, along with others that are needed to secure all layers of the physical infrastructure.

While the concept of electronic trust has been discussed for decades, advanced technologies are getting us much closer to creating reliable trust capabilities in the network: we are starting to reach the point of having a trustable environment combined with techniques to make trustable individuals.

How far away are we from seeing trusted communications environments?

Early forms of trust-based networks are already emerging in a few pockets of the industry.

Online auction systems are great examples of trust-based networks in the consumer space. How do you know whether you can trust a vendor? These systems have a self-regulating trust mechanism. You can see how many stars a person has behind his or her name, and you can read customer feedback online. People are willing to buy from somebody around the world simply because they have four stars behind their name, because they see those stars as a figure of merit.

These online auction systems, while they may represent rudimentary trust environments, are extremely effective. They are scaleable and instantaneous, and offer buyers a dynamic, global, and virtual shopping experience, with real-time pricing. Billions of dollars of sales are being transacted simply on the basis of how many stars a person has behind their name.

In fact, I recently read an article where an eBay executive said that the company's online auction was selling toys at the rate of 85 per minute during the weeks leading up to the past Christmas season. The success of eBay shows how eager millions of people are to participate in a global, virtual marketplace, and how a trusted network can pave the way for highly ad hoc, highly dynamic, and effective business and information-sharing relationships.

Similar trust-based environments are emerging in the service provider space. For instance, NTT DoCoMo – the largest wireless service provider in Japan – operates an interactive content portal that brings together a large ecosystem of applications and content providers with whom DoCoMo has established trusted relationships. Through this portal, subscribers can purchase everything from ring tones to books, and the transaction is billed to their wireless account. DoCoMo consolidates the billing, and in exchange, garners a part of the revenue from each sale.

These are the kinds of trust-based networks that are emerging, but they could be considered fairly rudimentary, in the sense that they are closed and proprietary systems – called "walled gardens" – where content and services are offered exclusively to subscribers.

For the vision of universal trusted networks to become a reality, these networks need to evolve to be open and standards-based.

As they do, large and open federations of businesses and organizations will emerge, with formal agreements among member organizations to establish common authentication and identity management infrastructures that enable them to trust in each other's infrastructures. In this way, consumers and entire communities will have secure access to the Internet-based infrastructure used by these federations for content sharing and e-commerce.

You could say that creating trusted network environments that are universally accessible is the industry's "grand challenge." It's a challenge that comes with interesting technology, business, policy, and standards hurdles that need to be addressed by all players in our industry – from chip vendors, mobile device manufacturers, equipment vendors, and network service providers to content owners, financial intermediaries, and international standards bodies.

From a technology perspective, what is required to begin addressing that grand challenge?

Nortel and others in our industry have already put a lot of effort and innovation into creating the spectrum of security technologies needed to protect networks and information, and we are continuing to do so. In fact, much of the fundamental technology that we need to get started is already beginning to roll out.

Our challenge, if you will, is to take these security technologies and make them work across the entire converged network.

Identity management, for example, is one of the technologies considered essential to realizing trust-based networks. Identity management systems, along with associated authentication techniques, refer to a broad area that deals with identifying individuals in a system (such as in a country, a network, or an enterprise) and using their secure credentials to seamlessly access the services, applications, and content they want and need. With an identity management solution, the network can store and manage a single identity profile for each individual and transfer that identity across different network domains.

Identity management systems themselves are not new and Nortel has long been a leader in this area. Nortel's own enterprise network, in fact, has had in place an internally developed identity management system since before the industry really began using this terminology.

Our Norpass system addresses several key security and usability requirements across the enterprise, including: enforcement of strong password rules, including expiration; password synchronization across multiple application domains (allowing use of common user IDs and passwords even among non-integrated applications); and single

sign-on across all of our web-based applications. This continues to save us a significant amount in direct IT costs and dramatically improves employee productivity and security. Employees need to sign on only once, typically at the start of the day, and are logged into a suite of Norpass-compliant applications, including email and market research databases. Norpass has driven a tremendous reduction in the number of employee log-ons, passwords, and different sets of credentials in use and has significantly simplified the user experience, resulting in far fewer calls being made to the IT help line to reset forgotten passwords.

Norpass is also used as the basis for access to the internal network. As an example, Nortel employees and partners who are working outside the office (e.g., at home, in a hotel, or wherever Internet access is available) use their same Norpass user ID and password to log-on remotely and securely. This capability leverages these users' single identity profiles in conjunction with encryption provided by Nortel's VPN Router and VPN Gateway solutions.

Now, imagine the benefits for users – and for operators – of further integrating identity management capabilities into the network. We could connect users, along with their credentials, not only to all types of communications services, but also to multiple different networks or federations of networks – whether wired or wireless, public or private.

With such a federated identity management capability, the network could manage my identity and credentials as I move from place to place, and from network to network, using any number of different devices. Wherever and whenever I log on to any network – whether I access the Internet through my DSL modem at home or even from a friend's desktop computer, or through the corporate network from my laptop or cellphone while traveling – the network would recognize me and be able to verify my identity. And my identity would roam with me as I move about this virtual, connected world, regardless of the application I'm using. The network would know what information I am allowed, or entitled, to access. I could even customize my online experience by setting up multiple identities – me at work, me at home, me as the hockey coach – depending on a specific context or domain.

As we begin to embed technologies like these directly into the core network, we will be able to address the values that people really care about: Am I interacting with people that I trust? Are people that I don't trust disallowed from being part of this transaction? Am I comfortable sharing my information with this person or that organization? Is this a person with whom I want to do business?

For end users, this next generation of identity management can make their lives simpler – improving their quality of experience – by allowing them to customize their online experience and establish the terms for how they interact with others, and vice versa.

To help make this a reality, our advanced research teams are working in this area, integrating identity management into all parts of the network.

Networked identity management is one solution that helps establish trusted networks. What are some other technology requirements?

Perhaps our biggest challenge is making the trusted network one that is very easy to use. This is an area where Nortel is very well-positioned. To make the network simple to use and reliable, you have to be a master of many disciplines, and that is precisely Nortel's heritage.

Think of the basic telephone – it is undoubtedly the gold standard for "simple to use." When did you ever need to read a manual to learn how to use the phone? A lot of the network security capabilities that I referred to earlier are, to varying degrees, possible today, but it would take a computer programmer to figure out how to use them. How do I get my mother to use it? How do we push all the complexity back into the network, so that it just works, invisible to the user, much like the way that cellular roaming works today?

In cellular networks, you turn on your handset and it automatically goes to the network and checks whether that person or device is allowed to make a phone call on that mobile network. As you roam into another provider's network, the device and network are both working diligently behind the scene to confirm that the two service providers have an agreement in place that lets this user roam. The user isn't aware of these activities. The network just works, and it is instantaneous.

This trusted roaming capability is an area where Nortel has been a technology leader for decades, and we're now bringing that technology to bear more generally through networking. In fact, wireless networks are leading the way in many aspects of network security – including identity management, simplicity of use, and trusted infrastructures. Much of our technology leadership in this area can easily be leveraged. For instance, in the identity management space, we are building on our industry-leading HLR/HSS (Home Location Register/Home Subscriber Server) capabilities.

Another capability essential to creating trusted networks is knowledge of real-time networking – understanding the performance requirements of end-to-end networks and knowing how to integrate security across all the different network domains, elements, and applications in a seamless manner. It isn't just enough to be really good at security. To create world-leading, secure, and reliable communications solutions, you must be the best in both security and in networking. Here again, Nortel is well-positioned. One of our central value propositions as a company has been built around our leadership in designing and building real-time communications that are secure, trusted, reliable, and

Key functions of network-wide identity management

Network-wide identity management (IdM) enables the network to seamlessly manage a person's electronic identity and credentials as that person moves from place to place and across different access domains, and uses any number of different devices and applications.

In order for IdM to operate seamlessly across networks, five key functions must interwork across every domain related to the network – from users and devices, through the network elements themselves, to all services, applications, and content – and regardless of access technology and device type.

The five key functions are:

• *authentication* – reliably verifying the identity of a party (user, device, or network element). Authentication is usually a prerequisite to authorization, since entitlements depend upon the identity of the user. User authentication is based on one or more credentials, such as a password (something a user knows), a smart card (something a user has), or a fingerprint (something a user is).

• *authorization* – reliably controlling access to resources as well as enforcing entitlements, considering end-user identities and policies;

• *accounting* – reliably charging and enabling unified end-user billing for service consumption;

• *auditing* – reliably logging, storing, and making available relevant end-user information and transaction records, to comply with current regulations such as Sarbanes-Oxley; and

• *user management* – the process used by end-users to reliably update specific aspects of their user profiles and preferences (considering end-user roles, time of day, location, etc.). Users can also be operators in a service provider organization who need to manage such administrative activities as updating end-users' service level agreements (e.g., bronze, silver, or gold service) and security policies.

Enabling consistent IdM in network solutions is a powerful element for enabling new levels of security. The network's ability to track a user's identity and associate it with all of that user's actions and individual policies, reinforces the real-life concept of trust and accountability for one's actions in the virtual network environment. For instance, a user who might be inclined to engage in malicious or otherwise unethical behavior is less likely to do so if his or her actions can be easily traced to their identity instead of to an anonymous IP address. Today's networks, however, are largely unable to trace actions against user IDs. Moreover, the means for authenticating users are typically prone to impersonation and identity theft.

Network-wide IdM brings several other benefits, including single or reduced sign-on for end users (one password); simplification, lower costs, and fewer help-desk calls for enterprises; more finely grained security across a broader and virtual user base (including employees, partners, contractors, suppliers, and customers), as well as improved business effectiveness with emerging federated systems.

Federated IdM is a logical extension of IdM, whereby multiple different entities, such as enterprises and service providers, form a "circle of trust" and agree to use common authentication and identity management practices, as well as a common set of policies and procedures to enforce confidentiality and privacy of end-user information. Federated IdM would enable users to conduct secure and trusted transactions among different companies or service providers by means of a single or reduced log-on.

Nortel is currently incorporating advanced network-based identity management capabilities into its solutions.

simple to use.

At Nortel, we're drawing from our expertise in all these areas, as well as our deep understanding of network security, to explore some of the exciting new services and capabilities that future trusted networks will enable.

What are some of these exploratory areas?

Our advanced research teams are looking into a broad spectrum of potential opportunities with a number of initiatives and exploratory projects.

One of these areas is IPTV. Nortel announced an end-to-end IPTV solution that brings our industry-leading SIP-based multimedia communication technology into the television experience. With this solution, IPTV subscribers can use their televisions to communicate and interact with their friends and family through a variety of media, such as voice, instant messaging, video, and picture sharing. The solution also works with wireless devices such as PDAs and cellphones.

Moreover, as the 3G IP Multimedia Subsystem (IMS) architecture becomes well established over the next few years, the IPTV solution will allow content to be received and delivered across wireless and wireline boundaries.

Securing television on an IP network is critical for enabling these new capabilities. Our broadband solutions team is working to integrate security into the network, and address such security challenges as the need to protect user privacy and ensure that content cannot be pirated, among others. In fact, working with its IPTV ecosystem of third-party developers, our development team has completed detailed integration and testing of its IPTV solution with

third-party products and capabilities to integrate security into the network core and enable the secure transmission of information to a set-top box.

A secure, trusted IPTV environment will make possible a range of new applications that will break down the barriers between traditionally separate communications environments, and they will fundamentally change the user experience. For instance, with a secure IPTV solution, videoconferencing and video-on-demand – which today are separate environments that use different devices (the telephone and the television) – can be integrated. A group of teenagers, while in their respective homes, could all watch the same television program as if they were in one room, sharing comments and laughter in real-time while also sending instant messages to each other. Or, a family just returning from, say, their child's soccer tournament, could send photos or a video of the winning game to their grandparents living across the country, and they could all watch the slide show or video together, in real-time, on their respective televisions while engaging in a group conversation.

We're also exploring many interesting future opportunities to extend the types of applications that service providers could provide over their legacy infrastructures. One of these applications is health-related services. With trusted networks, people could hook up their heart-rate monitors to their set-top box and transmit that information in real-time to their doctor's office or to a distant hospital. The technical challenge here lies in ensuring that the patient's privacy is protected – that the information is not intercepted and it reaches the intended recipient.

The cable provider space represents equally fertile ground for innovation. With a trusted cable infrastructure, it would become possible for providers to achieve their holy grail of releasing new Hollywood movies simultaneously with the cinemas. For that to happen, though, the network must be able to prevent piracy of content by the subscribers themselves – an interesting technical challenge.

We've been discussing how security needs to be embedded in the network. But what about end-user devices themselves? What is Nortel's strategy to protect mobile devices from viruses?

Fortunately, there have been only a few reported viruses affecting mobile devices, but we expect such events to increase as data-enabled mobile devices become more commonplace. Here, network service providers can play a big role in protecting mobile users from attacks. We believe that providers should have technology in their networks to intercept viruses and other threats before they are sent to the cellphone. The network should actually take care of most of the security on behalf of the user.

We're actively working to make this happen. For instance, we recently announced an agreement with Websense, Inc. to develop an innovative URL filtering solution that will filter out unwanted web content – from viruses to content inappropriate for children. This solution combines Websense's security and filtering expertise with our end-to-end packet networking leadership (see page 48).

Even so, cellphones do need some basic capability to protect themselves. Because mobile devices can talk directly to other devices without a network being involved, through such technologies as Bluetooth, they will need to have the ability to connect only with trusted devices. We are in discussions with some of the key mobile device players, in both the handset and the software areas, looking at the technologies that will be necessary in the future. In addition, we will continue to leverage the mobile handset knowledge and technology expertise of LG Electronics, through our joint venture relationship.

New generations of mobile handsets are also now emerging with greater levels of security built in. For instance, handsets are just coming to market now that meet the Digital Rights Management standard published by the Open Mobile Alliance (OMA). These handsets are designed to govern and protect the use of new types of mobile content – short movie clips, for instance – that can now be downloaded to cellphones. [The OMA was established in 2002 and has grown to more than 300 member companies, including wireless operators, device and network vendors, IT companies, and content providers.]

Nevertheless, it makes a lot of sense to put security capabilities directly into the network. If you had to rely on having security software on all end devices, it would become the user's responsibility to make sure that all his or her end devices had the same security settings and the right security software. By putting security in the constant part of the network (the network itself) rather than in the changing part (the devices), we can achieve improved and more consistent security, as well as a more satisfying and productive user experience.

It sounds like the industry is moving quickly to put the pieces in place for trusted networks.

It is happening very quickly, because security has become a top-of-mind issue with anyone who builds and operates networks or offers network-based services. That's one of the reasons why it's a very exciting time to be in the network security business. It's a very dynamic and technical field, and because we're still in what I'd characterize as a first-generation state, there is a lot of room for innovation. There's a real sense of excitement, both in our own advanced research labs and throughout the security industry. The solutions that we're working on will have a huge impact, not just on the market, but on the nature of how people, communities, commercial entities, and governments will communicate. ■

Rod Wallace *is Leader of Security Solutions and Services.*

Newsbriefs

Nortel demonstrates world's first integrated data encryption for 10-Gbit/s optical networks

by Kim Roberts

Nortel is the first in the world to demonstrate integrated data packet encryption for high-speed 10-Gbit/s optical networks – an advanced technology that could revolutionize the way networks are secured, providing a global-scale approach to security that is always-on and transparent to users.

The technology has the potential to make the integration of encryption security a standard component of tomorrow's transport networks. Encrypting packets for transmission over the network provides protection when the data has not been encrypted at source, or provides an additional layer of protection when it has.

The ability to use real-time encryption to secure the transport of large amounts of sensitive or confidential information has applications in many industries, such as media distribution in the entertainment industry, global scientific collaboration using grid computing, and real-time data back-up for the financial industry. The technology will address the growing concern about protecting such information as it becomes more distributed across multiple networks through virtual private networks (VPNs) and applications such as grid computing.

Security over these networks is a concern because, contrary to what many believe, fiber is relatively easy to tap. Clip-on couplers are widely available that can simultaneously inject or extract light from the fiber at any intermediate point between the source and destination. These couplers are used for legitimate purposes, providing a convenient and quick solution for connecting talk sets, power meters, and light sources by installers or designers. However, anyone with the appropriate technical skill and equipment can also employ them to tap into fiber – in a non-destructive and almost undetectable way – to intercept sensitive communications. Other more sophisticated methods of intercept, such as etching, also exist and are even harder to detect from the terminals, but they harm the fiber.

Nortel's 10-Gbit/s integrated encryption technology uses the 256-bit Advanced Encryption Standard (AES), a globally accepted standard for encryption adopted by the U.S. National Institute of Standards and Technologies (NIST) as an improvement and replacement for its predecessor, the Data Encryption Standard (DES). Because Nortel has implemented this technology in hardware at high bit rates, latency is almost non-existent. Delay is just 400 nanoseconds at each end, compared to some competing implementations that experience delays of 100-200 microseconds – 250 to 500 times greater. Nortel's hardware-based technology implementation also delivers full 10-Gbit/s throughput, unlike other processor-based methods of encryption such as IPsec and Secure Socket Layer (SSL) that send overhead information along with the transmission, significantly lowering the throughput of short packets.

Nortel demonstrated the world's first successful prototype of integrated data encryption for 10-Gbit/s optical networks at the iGrid event in

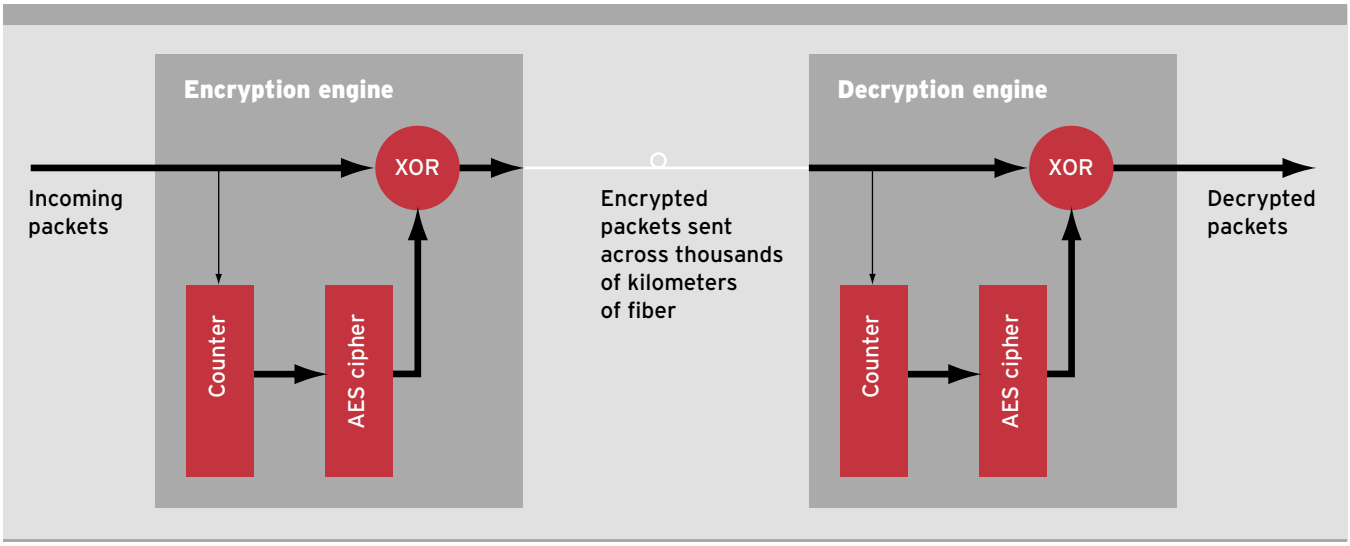

San Diego in September 2005. The demonstration was carried out in collaboration with leading global research institutions – including SURFnet (which connects universities, research centers, and scientific libraries in the Netherlands to one another and to other networks in Europe and the rest of the world), SARA Computing & Networking Services (also based in the Netherlands), CANARIE (a non-profit corporation funded by Industry Canada to facilitate the development and use of next-generation research networks and the applications and services that run on them), the International Center for Advanced Internet Research at Northwestern University, and the Electronic Visualization Laboratory of the University of Illinois.

The live demonstration transmitted real-time encrypted data from an electronic visualization application over a 10-Gbit/s SONET network that spanned thousands of kilometers of fiber and six network operators to the iGrid show floor in San Diego, where the data was viewed and manipulated on an integrated 55-screen, 100-million pixel video display. The hardware for 256-bit AES encryption at 10-Gbit/s line speeds was integrated into a standard Nortel Optical Multiservice Edge (OME) 6500 switch.

By encrypting the application data payload before packaging it into a SONET envelope for transport, the encrypted traffic can travel across today's standards-based SONET networks to be decrypted on the other side. As shown in the diagram, each incoming packet is separated into 128-bit words, which are sent to the exclusive OR (XOR) bitwise logic function. The counter is incremented for each word that arrives in that packet connection. The counter value for each 128-bit word is then encrypted by using the AES cipher algorithm and the secret 256-bit key that corresponds to the connection to which that packet belongs. Each bit in the resulting apparently random word is then XORed with each bit in the word of the data itself, and the result is transmitted. At the receiving end, the counter is incremented for each word of the packets in that connection and synchronized to the same value as the counter on the sending side. Each 128-bit counter value is then encrypted using the same AES algorithm and secret 256-bit key, and XORed again with the word of the packet. This operation decrypts the data.

The integration of encryption functionality into the Nortel optical products offers the potential to fuel the growth of optical networks by:

- providing data privacy and protection against denial of service attacks, enhancing the security of customer data and the network;
- reducing network complexity by eliminating the need for an external encryption device;
- lowering operational expenses through reduced power and space requirements; and
- providing higher network reliability through the inherent carrier-grade attributes of the optical switch. ■

Kim Roberts *is Technical Advisor, Systems and Architecture, for Optics Research.*

NORTEL

The photograph on page 62 was taken by Mike Pinder.