

Management of Networks to Ensure Security

John C. Hoag, Ph.D. and David L. Post
McClure School of Information and Telecommunication Systems
Ohio University, Athens Ohio
{John.Hoag.1, David.L.Post.1} @ ohiou.edu

Abstract

At least one view of network management integrates both network monitoring and resource allocation based on a maintaining system state. While standards-based network management has provided the telemetry to interpret security events, advances in security have come in the form of tools and techniques. This paper explores the potential use of system state, as well as integrated network monitoring and operations support systems, to advance the field of information security.

1. Introduction

The primary objectives of network management have been achieved, albeit often lacking in automation and efficiency. Events are trapped and correlated; circuits are provisioned, tested, and ultimately billed. Interesting advances in network management have sought to integrate information systems and databases, and the trajectory of such systems is to model both network elements and their status

In this paper, we will attempt to assess network security practices and artifacts against the emerging state-oriented systems for network management. In order, we discuss the objectives of network security; the emerging capabilities of telecommunications network management; a converged model for network management and security state; and conclude with an assessment of this model.

Network operations has long been an open-loop process with many manual tasks including fault intervention. Network management in

telecommunications has, until recently, involved separate subsystems for realtime network monitoring and business/operations transaction; in the electric industry, the stability of the transmission grid is largely due to integration, monitoring, and automated control.¹ We may characterize the advances in telecommunications network management as adjuncts to implementation of new infrastructure like MPLS or through trends in software practice such as open source and COTS (commercial off-the-shelf software). Interested parties, both providers and enterprise customers, participate in standards-setting for objects and interfaces.

We have termed our intellectual contribution to the topic of telecommunication network management “State Based Network Management,” embracing a semantic web approach and pursuing open source implementations.²

The strategic goals of network security have shown their permanence. U.S. Law defines information security to have three attributes: integrity, confidentiality, and availability.³ The U.S. Navy describes its objectives as follows: to protect information systems and services from accidental or intentional disclosure, modification, destruction, and denial of service.⁴ The tactical objectives for network security have shifted over time, however. Antivirus vendor Symantec states that two important trends are that attacks on software have now exceeded attacks on network perimeters, and that web vulnerabilities now account for over half of all software attacks.⁵

Achievements in the science of security have focused on systems such as servers and application suites and their architectures. Bellovin pessimistically notes that the security attributes of these systems are not amenable to metrics and are unlikely to be improved measurably.⁶ Some time-invariant metrics are the time-to-compromise a system and the time-to-exploit a vulnerability. Moreover, per Symantec, the risk associated with software vulnerability *is eliminated* when patches are installed.

An interpretation of the trends listed above is that software quality is becoming the driver for system security. The traditional view has been that software source code becomes more reliable through testing and execution experience, although this concept had eluded the assumptions necessary for modeling.⁷ A more current and more accepted notion is that of a “relative attack surface” composed of discrete vectors such as: open sockets, enabled guest accounts, network address access control lists, etc.⁸ Rather than looking at source code, a Relative Attack Surface Quotient (RASQ) provides the means for comparing a system against, say, Windows 2000.

2. Background to Security

U.S. authorities, based on recent law, prescribe an information security approach that is risk-based, anticipatory, standards-based, and traceable through security policies. The starting point for this approach is the appreciation of risk, defined as magnitude times probability of occurrence. This approach emphasizes planning: developing policies, anticipating threats, and planning responses consistent with external standards.

The U.S. National Institute of Standards and Technology (NIST) is tasked by law with development of security standards and implementation guidelines, for which its Computer Security Resource Center (CSRC) has produced the “800 Series” of publications. Elements of the 800-100 Information Security Handbook include: security governance; lifecycle, capital, and contingency planning; performance measurement; certification, incident response, etc.⁹ The NIST guidelines specify an executive role for the head of information security, with control over security aspects of all systems in development or production.

The U.S. Department of Defense has funded similar efforts at Carnegie Mellon to advance software engineering, emergency response, and survivability/information assurance. The Computer Emergency Response Team (CERT), whose most visible role is to

assess security activity and issue alerts, also has developed an excellent tool for security management named OCTAVE, Operationally Critical Threat, Asset, and Vulnerability Evaluation. As a strategic planning and evaluation tool, OCTAVE focuses on systemic risks rather than technology evaluation.¹⁰ The OCTAVE Method involves the following three phases, repeated endlessly:

- Catalog assets, threats, practices, organizational vulnerabilities, and other security requirements
- Identify key components for technical vulnerability
- Assess risks to critical assets and plan a protection strategy, plans, and actions

Some of the first security oriented systems developed were firewalls. These systems exist as either stateless or stateful (the former giving way to the latter). In a stateless system, the firewall cannot distinguish between traffic that is established by a trusted source and that which is not. The most well known example of this problem is a standard FTP connection which needs to connect to a random high level port. In this case, a protected source initiating a FTP session from behind the firewall would have their legitimate transaction denied. Stateful firewalls allow for connection tracking and would associate the random high level traffic with a trusted internal source.

The next stage in security development came with the advent of Intrusion Detection (IDS). One such system is the open source project known as SNORT (and developed commercially by Sourcefire). This product, now the de facto standard for intrusion prevention, was initially developed in November 1998 (ten years to the month after the release of the Morris Worm) by Martin Roesch as a packet sniffer. Limited rule based detection features were added in January of 1999 and the turn of the century brought numerous community developed modules that expanded the package allowing for logging to a database, SNMP support, and Intrusion Prevention (IPS) capabilities (adjusting the rules on a Cisco PIX firewall for example) to name a few.¹¹

Since that time, the fundamental way in which SNORT and most other industry available options work has not changed. The system is rules based—sniffing packets for strings that match a defined rule set and taking action accordingly. These actions could be as mundane as quietly logging for later review, issuing an alert/notification/generating a ticket, or result in traffic termination. Rules and actions are a compilation of both mainstream product updates and user forked custom rule sets.¹²

Gartner Group recently released its "hype cycle" assessment of security tool adoption, noting both the maturity of each product but the tendency to overstate expectations of each. Also, this assessment indicates that some products will not reach maturing before the need for them dissipates.¹³ Two such products are contrasted below. Intrusion Detection/Prevention Systems are reaching a "plateau of productivity," while Network Behavior Analysis and Anomaly Detection Systems are approaching their "peak of inflated expectations."

Intrusion Detection/Prevention Systems (IDS/IPS) go beyond typical firewall Access Control List (ACL) functionality. These systems can look deeper than the source/destination of a packet and allow for decisions to be made based on packet contents. Cisco defines Intrusion Detection as "A security service that monitors and analyzes system events for the purpose of finding, and providing real-time or near real-time warning of, attempts to access system resources in an unauthorized manner."¹⁴ Whereas Intrusion Prevention Systems take detection one step further and will attempt to mitigate the intrusion in addition to issuing an alert. Many systems available today identify themselves as one, the other, or a mix of both—while the functional line dividing the two is somewhat blurred. These systems can be deployed on either a network or host basis (giving way to NIDS/HIDS style acronyms) and perform their function while either promiscuously monitoring the network or serving as an inline gateway.

Recognizing when a network service or segment is exhibiting abnormal behavior is also an important element of network security. To this end Anomaly Detection Systems (ADS) have been developed to monitor utilization of network elements and issue an alert when the defined metrics are violated. Because an ADS utilizes pattern analysis, it can often alert operators to a possible problem with the network (such as symptoms exhibited from a DDOS attack) before it would otherwise be noticed by other systems that simply check for availability or end users.

Continued convergence of monitoring technologies is expected. The combination of IDS/IPS solutions with those that provide anomaly detection (monitoring traffic and reacting to deviations from learned expectations) is well underway.¹⁵ Ultimately, the efficacy of the network security provided by these systems is only the sum of its parts. The greater the systems knowledge of the network and its components, the better able it is to respond. In fact, with the inclusion of a dashboard style view of current alerts and network health, these systems often resemble a network management system.

We presume the next evolution of security development to revolve around a system that—based on previous vulnerability knowledge and full awareness of its own health—can diagnose itself. The limits of the current systems revolve around the need for human update/analysis/intervention. The intelligence to react to day zero threats autonomously should be the next stage of development. A nigh self-aware central control engine that compares network faults, security sensor data, network traffic history, and previous threat reactions could yield a system that can diagnose itself, remedy if possible, patch around (using its relations with an OSS), or—ultimately—request outside intervention. This type of system broadens the definition of "self-healing" and brings us to realization of systems which are more capable of defending themselves from day zero threats.

3. Background to Network Management

Broadly speaking, there are two forms of network management: real-time monitoring, characterized by IP in-band standards-based SNMP messaging; and operations/business support systems, composed of inventory, order entry, provisioning, testing, fault management, and billing systems. Note that the former category can apply both to enterprise and provider networks, both LAN and WAN, and that the latter category is primarily the realm of transport providers. Given that nearly ten thousand enterprises with fifty thousand endpoints utilize MPLS, for instance, and that VPNs are common for remote access even in private TDM WANs, we conclude that large, interesting networks require components of both categories of network management systems.

Network monitoring is performed by commercial products such as HP Openview that must be implemented uniquely for each organization. This involves manual setup, for large enterprises, of not only the network topology, but also the rules for filtering alarms and correlating information and, optionally, the presentation of coherent views to operators. At its most sophisticated level, this NMS can create trouble ticket records in a companion business system upon determination of an error; recent advice recommends against this, however, asserting up to a 50% error rate in automated trouble ticketing.¹⁶ NMS development is complex, requiring both incremental “tuning,” expertise working with the specific reporting nuances per each interface, and a commitment to maintenance.

While a market for network monitoring software had developed – software purchases over \$5 billion annually with three firms accounting for nearly half¹⁷ – operations and business support systems are largely proprietary. Gallen estimated that providers combine to spend \$30 billion annually on OSS/BSS systems, which may be neither integrated within the firm nor compatible with peers or trading partners.¹⁸ Software vendors and network providers have collaborated through the Telemanagement Forum (TMForum) to

develop interface standards for a “next generation” OSS (NGOSS), a process that has to date yielded few ratified documents or demonstrations. Proof-of-concept was established in 2005 by Gallen’s project using open source code for OSS modules. An interesting component of the OpenOSS initiative is the use of web services and XML documents for interfaces.¹⁹

A primary focus of this research project has been the design of a system, State-Based Network Management (SBNM), to integrate NMS and OSS activities. This activity has determined that semantic reasoning can be the key to efficient interoperability among OSS, NMS, and requirements modules. As evidenced by open source versions of the first two modules, web services can be developed as user-less “black box” transaction portals. Effective transactions between interconnected systems depend on shared ontologies (or translation tools). Leverage at the ontology level, in our design, may take place at three interfaces:

- Between customer requirements and technical specifications
- Between customer specifications and alternative implementations
- Among different ways to describe location

The goal of this system is to allocate and reallocate resources from providers’ inventories in order to fulfill requirements. Its notion of system state encompasses requirements, provisioned network elements, their current status, and remaining items in provider inventory. In these terms, one goal of the system is to maintain system state after a perturbation.

4. A Converged Model

The challenge in this paper is to investigate whether corresponding definitions of state, dynamic equilibrium, and transaction processing make sense in a security context. Moreover, we are concerned to what extent ontologies advance system interoperability in security.

For small providers or enterprises, a single staff may have responsibilities for all aspects of operations including security. In a larger setting, a separate and distinct security organization may exist to lead, for example, the activities outlined in OCTAVE – primarily at a policy level. Given that the execution of policies is an operations task, industry analyst Gartner has called for the convergence of IT security and operational management processes and tools.²⁰

Recognizing that endpoint (i.e., workstation) software is becoming the greatest source of risk, and that timely application of patches is a known effective remedy, configuration management (CM) is potentially the first task for converged network and security management. The scope of converged CM should include both endpoint software and communication node configuration. Tasks for both categories of CM have been implemented through modest scripting, i.e., through login profiles and tftp, for instance.

The set of configuration data can be seen as a goal state to be maintained, with automated CM tools as means to maintain state. Configuration data, though, is static by contrast to incident response.

The notion of state is much harder to define in the dynamic case – and also harder to enforce. We define dynamic state as the set of permitted flows between endpoints, which ultimately requires knowledge of applications and characterization of their use. In theory, this is knowable, but in practice these kinds of data are elusive. In policy terms, we might state that practice is to deny all traffic except that which has been registered with security; moreover, an extension to this policy might be to discard registered traffic in excess of its prescribed rate.

The state-based security model could also encompass the collective status of every instance of a security policy in action. State-altering transactions (add, modify, delete) probably come from either a CM subsystem as well as policy records maintenance. These definitions

imply methods for acquiring state information; actions such as policy enforcement and incident response are left unresolved with no intent herein for automation.

A new genre of software applications known as Security Information and Event Management (SIEM) is appearing largely to fulfill compliance reporting as dictated by law. Among leading systems, one has evolved from a web filtering product (Network Intelligence and Websense) and another is an extension to the correlation suite of an NMS system (CA-Unicenter Network System Manager). While each of these will provide enhanced visibility (e.g., dashboards) to operators, neither is based on application (i.e., traffic) requirements or maintains a complete model of desired state.

5. Conclusions and future work

The concept of policy-oriented and application-oriented system state aligns well with the goals of information security. The notion of a central repository of state information appears workable, and such a system would be valuable even in an open loop configuration (i.e. no automatic incident response or policy enforcement). This functionality is consistent with the network management concept introduced in SBNM.

Two factors make implementation of this approach impractical, and one will be very hard to overcome. The lesser factor is that the necessary applications and traffic data to form the model is not typically maintained by network operations staffs; the extent of granularity available may be to the level of address and port. An organization must be committed to life-cycle design based on requirements in order for this task to succeed.

The stronger factor working against state-based security management is the variation in threats and vulnerabilities. While some trends may never reverse, the nature and origin of threats change constantly. Regardless of whether an organization embraces a system-oriented

risk-based or component-oriented technology-based approach to security management, the results are individual security policies. Policy information must be translated into the language of system state, which further must be traced to telemetry devices in the field. Unlike network management, which has only sets of communication nodes to monitor, security management must consider static and dynamic data for all endpoints.

This paper provides three clear contributions: to extend the concept of state-based network management to security management; to introduce a definition of security system state; and to clarify the scope of the new genre of commercial security management tools.

¹ Hoag, John, "Telecommunications Systems Failure: The Electricity Blackout of August 14, 2003," presented at the 12th International Conference on Communication Systems Modeling, 2004.

² Hoag, John and C. Gunderson, "State-Based Network Infrastructure Allocation," in Proceedings of IEEE MILCOM, 2005.

³ Federal Information Security Management Act of 2002, U.S.C. Chapter 35, Title 44, Section 3541 et seq.

⁴ Automated Information Systems Security Guidelines, U.S. Department of the Navy, accessed from the Naval Postgraduate School, August 29, 2006.

⁵ Symantec Internet Security Threat Report XI, May, 2006.

⁶ Bellovin, S.M., "On the Brittleness of Software and the Infeasibility of Software Metrics," *IEEE Security and Privacy* 4:4, 2006.

⁷ Musa, J.D., *Software Reliability*, McGraw-Hill, 1987.

⁸ Howard, M., Pincus, J., and Wing, J., "Measuring Relative Attack Surfaces," Computer Science Technical Report TR04-102, Pittsburg: Carnegie Mellon University, 2004.

⁹ Information Security Handbook, U.S. National Institute of Standards and Technology Computer Security Division Draft Special Publication 800-100, June, 2006.

¹⁰ OCTAVE Criteria, Version 2.0, Software Engineering Institute Computer Emergency Response Team, Pittsburgh: 2001.

¹¹ Roesch, M. (Speaker). (2005). *The Story of Snort: Past, Present and Future* (Audio Recording).

Retrieved August 30, 2006, from Help Net Security:
<http://www.net-security.org/article.php?id=860>

¹² Caswell, B., & Hewlett, J. (2006). *Snort Users Manual* (2.6.0)., Retrieved August 30, 2006 from Snort.org:

http://www.snort.org/docs/snort_manual/2.6/snort_manual.pdf

¹³ Wheatman, V., et al, "Hype Cycle for Information Security, 2006," Gartner Group Research ID G00139428, July 10, 2006.

¹⁴ http://www.cisco.com/en/US/products/sw/secursw/ps2113/products_user_guide_chapter09186a008007ead5.html

¹⁵ Sourcefire Network Security. *Real-time Network Awareness (RNA) Sensor*

Retrieved August 30, 2006 from Sourcefire.com:

<http://www.sourcefire.com/products/rna.html>

¹⁶ Rosales, M., "Business Class SLAs for VOIP Services," *Business Communications Review*, July 2006.

¹⁷ "HP Ranked No. 1 in Worldwide Distributed Systems Management Software," Reuters UK, August 17, 2006.

¹⁸ Gallen, C. and J. Reeve, "Investigating the Feasibility of Open Development of Operations Support Systems," in Proceedings of the 9th IFIP/IEEE International Conference on Integrated Network Management, 2005.

¹⁹ "Open Operation Support Systems", British Telecommunications PLC, 2005.

²⁰ Nicoleff, M., and J. Girard, "IT Security and Operational Management Must Converge," Gartner Research Publication G00124711, November 8, 2004.

John C. Hoag is Assistant Professor in the McClure School of Information and Telecommunication Systems at Ohio University, Athens, OH. He earned his Ph.D. in Industrial and Systems Engineering at The Ohio State University. He is a Senior Member of the IEEE and a member of the IEEE SMC and Communication Societies.

David L. Post is a graduate student in the Master of Communication Technology Policy program in the McClure School of Information and Telecommunication Systems at Ohio University, where he also earned his Bachelors degree