

Computer security meets KM—knowing when to worry

By Judith Lamont, KMWorld senior writer

Critical knowledge assets must be protected from external and internal threats, but the flood of alerts and alarms soon becomes a knowledge management issue itself. How can that data be interpreted and turned into actionable knowledge?

The good news is that security techniques are becoming more sophisticated and proactive, but the bad news is that enterprises must now manage an enormous amount of security-related information. About a dozen companies have jumped into this market, which is referred to as security event management (SEM). The goal is to centralize input from a variety of devices and software products that range from routers and firewalls to virus and intrusion detection software. Otherwise the data becomes worthless because security administrators do not have time to review it, much less respond to the threats.

Industry analysts forecast that this need will translate into a booming market. [IDC](#) projects that vendor revenues for software that correlates security events will increase at an average annual compound rate of 61.3% for the years 2000 to 2005, to \$404 million. The rate is higher than the expected growth rate for security management software in general (37.4%), a market that IDC believes will reach \$1.7 billion per year by 2005. A growing number of companies are finding this potential market to be very attractive.

[IBM's](#) IBM's line of security management products includes the [IBM Tivoli Risk Manager](#), which is designed to integrate and correlate alerts from many products into one console. From that centralized vantage point, a security administrator can view information that helps determine the source and severity of attacks. "One of our customers was getting 300,000 alerts per day," says Arvind Krishna, VP of security products, Tivoli Software, IBM. "What can they do when they have 17 consoles and all are on red alert?"

The customer—a large global bank that uses several levels of firewalls, intrusion detection systems from multiple vendors and other security products—eventually deployed IBM Tivoli Risk Manager. Now those 300,000 events have been reduced to 30 meaningful security alerts, which enabled the security administration team to focus on real security threats, reducing the time and cost of investigating alerts that do not represent a threat.

Tivoli Risk Manager allows the administrator to see an integrated picture of the alerts and to determine, for example, whether multiple alerts have the same underlying cause. The solution then becomes much simpler because the administrator may have to respond to only one event. Behind the scenes, a great deal of knowledge must be embedded in the software. Tivoli Risk Manager must know that a certain kind of alert has particular implications, or that a second alert looks different on the surface but maps to the same kind of underlying cause. The software can also offer solutions—sometimes a precise fix and sometimes a pointer to a URL or other guidance. In addition, it can incorporate new information as users gain experience with problems or patterns of attacks. Tivoli Risk Manager is updated on the Tivoli Web site on a regular basis for routine information and on an ad hoc basis for critical input.

According to Krishna, one of the biggest sources of security problems is the widespread practice of leaving security until the last stage of application development. "Often, a project is mature before security is considered," he maintains, "and the result is a mix of projects that are administered independently, without sharing best practices, data or other common elements."

He advises the use of a common security architecture and a roadmap to coordinate security across different projects. Although Tivoli Risk Manager includes some of its own security technologies such as intrusion detection systems and wireless security vulnerability scanners, a key benefit is its interoperability with a wide range of security systems from other manufacturers. "One of our strengths," says Krishna, "is

that we help companies leverage what they already have in place.”

Curing the pain

[ArcSight](#) was designed to integrate business objectives, security policies and procedures, and operations to provide a holistic view of the security environment. The company was founded two years ago after extensive research to discover areas of security where large organizations were “feeling the most pain.” ArcSight correlates and analyzes data from disparate sources in real time and summarizes it on a centralized console. Analytics are embedded in the product but can be expanded over time by the security staff to include new rules or patterns of attacks. More than 100 preconfigured reports are available, which also can be extended by adding reports specific to a customer’s security environment.

“ArcSight can import industry-standard information,” says Larry Lunetta, VP of marketing at ArcSight, “and connect in real time with events that are flowing through the system.” Security administrators can drill down to obtain more detailed information after an alert appears. If the system identifies a buffer overrun (which can deliver malicious code to an operating system), for example, the administrator can click on a cell in the display grid, read comments from the monitoring organization and seek company policy to guide a response. Lunetta cites the product’s rapid deployment as one of its biggest competitive advantages, with implementation typically completed in less than a week. The company’s top four markets are financial services, manufacturers, government agencies and ISP service providers.

“Organizations mature gradually in their approach to security,” Lunetta points out. “At first, they are reactive and unfocused, but eventually a corporate mandate emerges that recognizes security as a critical business process.” At that point, security measures are implemented but do not necessarily provide a meaningful picture. Additional sources of security are added that expand the amount of security data, such as information from applications and servers. “Finally, companies begin to take a proactive role by using security management technology,” says Lunetta, “so that the efficiency of talented but overworked security professionals increases, and their knowledge is leveraged effectively.”

The ArcSight solution consists of three components. ArcSight SmartAgents collect alarms and alerts from security devices and systems, and the ArcSight Manager normalizes the data. The normalization process translates information encoded in various formats into a unified view that is then presented in the ArcSight Console. A Console Dashboard provides real-time display for the metrics likely to be accessed most frequently. ArcSight reports that companies using its solution are detecting and resolving issues more quickly, and spending less time on false alarms.

Introduced as a commercial product in May, CyberWolf software from [CyberWolf Technologies](#) has been used by agencies in the federal government for several years. The software was initially developed by [Mountain Wave](#) through a research project at the [Defense Advanced Research Projects Agency](#) (DARPA), a central research and development organization for the Department of Defense, and was refined through several Small Business Innovation Research (SBIR) grants. CyberWolf matches patterns of events against a set of 3,000 patterns built into the software that could indicate attacks. A forms-based input system makes it easier to record new information and make it part of a permanent knowledgebase.

“One of our strengths,” says chairman and co-founder Juanita Koilpillai, “is the ability to easily capture knowledge from analysts.” Koilpillai notes that the technology has been used for a number of years by government agencies that are concerned with security issues including those in the intelligence community.

In addition to improving security, use of centralized monitoring helps alleviate critical labor shortages. One of CyberWolf’s government customers reduced its group of security employees from nine to two, and was able to reassign the remaining staff to other security tasks. It typically reduces hundreds or thousands of alerts to five or 10 per day that require human intervention. Right now CyberWolf is a rules-based system, but Koilpillai says that the company is planning to include other artificial intelligence approaches such as

neural network technology in the future.

CyberWolf is being used at the Federal Emergency Management Administration (FEMA, fema.gov) to monitor firewalls, routers, authentication servers and other devices. Other government clients include the [Air Force Research Labs](#), which is using the software to weed out alarms that result from equipment failure rather than real attacks, and the [Naval Sea Systems Command](#), which is using it to monitor security on workstations.

Comparing software products in this large and growing market can be a challenge. "Products vary in how well they carry out the normalization function," says Michael Rasmussen, director of research and information security at the [Giga Information Group](#). "That's an important evaluation criteria because it affects how well the correlation across events can be done." He advises organizations to consider how well the back end will hold up under real-world circumstances as opposed to pilot testing. In addition, some products create a better audit trail than others, by providing digital time stamping, for example, which can be useful if a cybercrime is prosecuted.

A patch in time

[BigFix](#) also operates throughout the network, but rather than detecting attacks, it monitors and repairs computer vulnerabilities on individual machines. Quick to install (it can be up and running in a day), BigFix can operate in standalone mode or work in conjunction with security event management software. As part of its service, BigFix researches alerts issued by authorities, such as [CERT](#) or [Microsoft TechNet](#), and then codifies the information into individual "Fixlet" messages that specifically define each vulnerability.

Fixlet messages are then gathered by the BigFix server, which evaluates them against the network. Because BigFix conducts a detailed analysis of each computer, it knows whether a certain patch needs to be installed. Patches are lines of code added to a program for the purpose of eliminating vulnerabilities or fixing "bugs." Although the present focus of Fixlets is on security, BigFix can apply patches for any type of problem, including bugs, upgrades or functionality enhancements. Since characteristics of the machine can be queried and retrieved, some companies use the system for asset management.

"The diagnostics allow an examination at a very granular level of the computer, including applications, versions, dlls or configurations," says Scott Texeira, director of business development for the BigFix Enterprise Suite products. "Companies can also author their own Fixlet messages to customize solutions unique to their network or policies."

Organizations can pay a high price for failing to install patches. Cybersecurity czar Richard Clarke estimated that the Nimda worm cost \$2 billion to repair, yet it was a known vulnerability in Microsoft products for which a patch had been available for months.

In the Department of Statistics at [Stanford University](#), BigFix is used to keep systems up to date. "Before we began using BigFix, our security staff had to go to every computer to install patches," says Balasubramanian Narasimhan, a senior research scientist who oversees the 30-machine network. Users couldn't install the patches on their own machines because updating Windows security requires administrator's privileges. Narasimhan and his staff found the process cumbersome and time-consuming. BigFix automatically installs the patches from a centralized console. "Now we update every week," says Narasimhan, who is pleased with the new method.

On the horizon

As enterprises extend their use of wireless computing, concerns about its vulnerabilities have increased, and new technologies are emerging in response. [Meetinghouse Data Communications](#), which was launched in 1988 to develop networking communications software, now provides a range of

authentication and security products including wireless security. Its AEGIS client and server software is an 802.1x-based (the IEEE protocol for Local Area Network port-based authentication) solution, offering end-to-end user authentication for wireless LANs in both enterprise and public access networks. (IEEE is the [Institute of Electrical and Electronics Engineers](#)). The authentication allows users to log on securely at wireless "hot spots" in public locations such as airports, as well as private locations such as corporate headquarters.

The Wireless Equivalent Privacy (WEP) protocol initially created by the IEEE 802.11 Working Group was intended to make the hop between the radio transceiver and the wireless computer as secure as a wired line. However, it uses relatively weak encryption. "It is distressingly easy to eavesdrop on WEP-secured wireless traffic," says Paul Goransson, president and founder of Meetinghouse. "Using 802.1x to dynamically pass a new key to the encryption hardware at frequent intervals makes the readily available algorithms for cracking the WEP encryption much more difficult to use."

[Secure Computing](#), which offers authentication, firewall and Internet filtering products, worked with [3Com](#) to develop an innovative product called Embedded Firewall. Developed under a DARPA research project, Embedded Firewall resides in each network interface card (NIC). The NIC firewalls protect individual computers even if perimeter security is broken. Embedded Firewall prevents the host from using the NIC to monitor traffic or "sniff" passwords. Introduced in February 2002 and available through 3Com distribution channels, Embedded Firewall has attracted the attention of the U.S. Navy, which is preparing to purchase it in bulk. With only a modest increment in cost over a standard NIC, Embedded Firewall offers protection to back-end servers by limiting the type of messages that can be sent.

"Policies are created and pushed to every card," says Paul DeBernardi, director of product marketing at Secure Computing. "By controlling each desktop computer, Embedded Firewall addresses the issue of insider attacks, which has been a difficult problem to solve."

Along with technology, attitudes are also changing. At [RCG Information Technology](#), Rachele McLure, national director for solutions, encourages clients to think about security early in the application development cycle. RCG IT is an information technology professional services provider, supporting such clients as [Reader's Digest](#) with database development and the [City of San Diego](#) with IT project management guidance.

"Some of our clients are now beginning to build in security testing as part of the application development cycle," says McLure, "particularly in the financial services and pharmaceutical industries." However, she believes that the prevailing "firewall mentality" of blocking outsiders needs revision, pointing out that not everything should be blocked in the same way for different groups of users. In addition, many enterprises still do not implement adequate security measures. McLure sees education as a key component of improvement, providing users with the information they need to understand the methods and importance of computer security.

Judith Lamont is a research analyst with Zentek Corp., e-mail jlamont@sprintmail.com.

KMWorld July/August 2002, Volume 11, Issue 7