

WHITE PAPER

Combating Insider Threats: The Application-Level User Behavior Tracking Approach

Sponsored by: Intellinx

Dan Yachin

April 2006

IDC OPINION

The growing awareness of the insider threat, and the recognition that security breaches by internal, trusted users are at least as risky as malicious outsiders, is pushing organizations to take action. Fueled by regulatory requirements, the market for various security solutions that can help detect and mitigate risks associated with the insider threat is growing rapidly.

At the same time, as organizations are shifting focus to the insider threat they realize that new security approaches and products are required to tackle different aspects of the problem, which for the most part have not been addressed by existing solutions.

One such emerging approach is application-level user behavior tracking, which allows organizations to detect fraud and other misconduct by insiders by tracking user activities in corporate business applications. This approach enables the tracking of authorized user access to corporate data that normally does not leave any traces, such as queries and other read-only actions that can be misused for personal gain – for example, selling sensitive customer information. By proactively detecting suspicious behavior at the application level, instant alerts can be generated, and immediate action can be triggered to suspend the suspect user until further investigation is made.

Application-level user behavior tracking can serve as a complement to existing Identity and Access Management solutions, which enable the management of users' access rights to sensitive information but lack the ability to control the use of information once access has been granted.

METHODOLOGY

IDC developed this white paper using a combination of existing market studies and direct, in-depth, primary research. To gain insight into internal threat mitigation technologies and approaches, and to learn about Intellinx' solutions, IDC interviewed the company's team on the issues of technology, product offerings, competitive landscape, and go-to-market strategy.

IN THIS WHITE PAPER

This IDC white paper looks at the growing problem of the insider threat. It discusses the regulatory and other drivers that fuel the demand for new security solutions in this area and outlines different approaches towards mitigating associated risks, including the emerging application-level user behavior tracking approach.

SITUATION OVERVIEW

Introduction

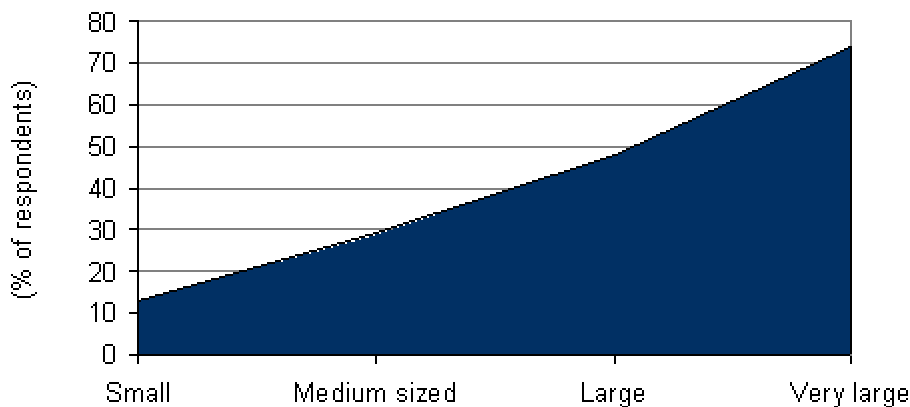
For years, organizations have been focusing their information security efforts to protect themselves against external threats, posed by the growing exposure to the Internet. Deploying an expanding array of solutions such as firewall, antivirus, antispam, intrusion detection/prevention, and anti-spyware, most organizations have built solid walls to protect their perimeters. Today, many of them are realizing that these defenses for the most part cannot help them deal with a different type of threat, which can no longer be seen as less risky – the insider threat.

IDC estimates that internal sources are responsible for over 60% of all security breaches. From deliberately stealing or destroying sensitive corporate data to falling victim to hackers exploiting unaware insiders to undertake different types of attacks, damage created by insiders may be costly.

IDC's survey conducted in 2005 shows that the larger the organization is, the more it is vulnerable to the insider threat (see Figure 1 below). Accordingly, the larger the organization is, the more it is concerned with internal threats (see Figure 2).

FIGURE 1

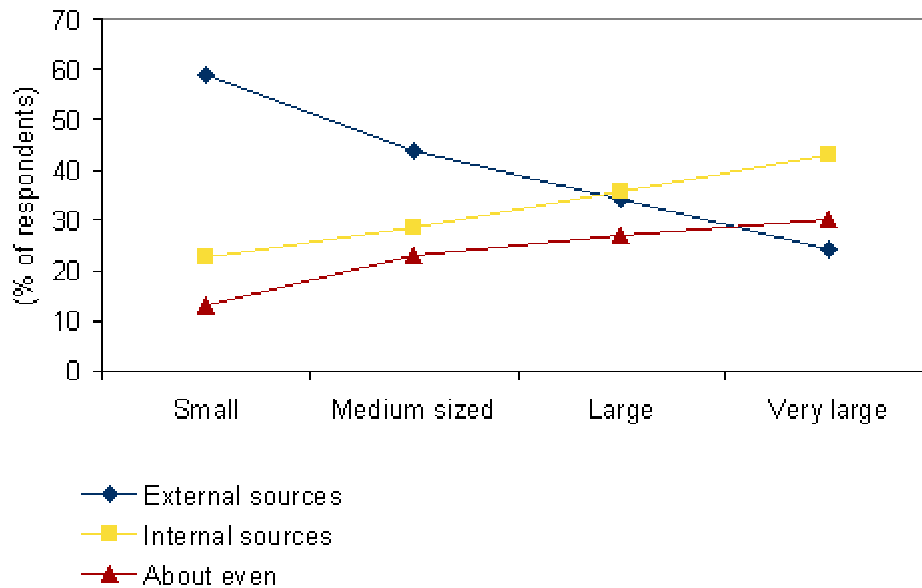
Internal Security Violations by Company Size



Source: IDC, 2006

FIGURE 2

Internal Versus External Security Threats to Enterprise Security by Company Size



Source: IDC, 2006

According to the 2005 CSI/FBI Computer Crime and Security Survey, 56% of respondents reported security incidents originated from the inside. Although this reflects a decrease from 66% in 2004, the survey indicates that the drop may be due to the increase in the "don't know" response, which makes interpretation ambiguous. The survey estimated the losses caused by insider net abuse at more than \$6.8 billion. Although this category is ranked 5th in the dollar amount losses by type chart, it is likely that internal sources are behind some portion of other higher-ranked incident types such as theft of proprietary info (\$30.9 billion).

Another indication of the extent of the insider threat was provided by a study published by the Association of Certified Fraud Examiners (ACFE) in 2004. According to the study, almost 80% of the fraudulent activity takes place on companies' premises and during normal business hours. Company employees are responsible for about 60% of this activity, while authorized partners or contractors carry out about an additional 20%. The study estimates that the typical U.S. organization loses 6% of its annual revenues to fraud (with total GDP loss of \$660 billion).

In addition, the 2005 "Insider Threat Study: Illicit Cyber Activity in the Banking & Finance Sector", conducted by the U.S. Secret Service National Threat Assessment Center and the CERT Coordination Center of Carnegie Mellon University's Software Engineering Institute, highlighted interesting findings about the nature of insider threats. Among them:

- ☑ Most incidents required little technical sophistication. In 87% of cases studied, the insiders employed simple, legitimate user commands to carry out the incidents. Only 23% of the insiders were employed in technical positions, with 17% of the insiders possessing system administrator/root access within the organization.

- ☒ In 78% of the incidents, the insiders were authorized users with active computer accounts at the time of the incident. In 43% of the cases, the insider used his or her own username and password to carry out the incident.
- ☒ The motive and goal for most insiders studied was the prospect of financial gain (both 81%); followed by revenge (23%); dissatisfaction with the company management, culture, or policies (15%); and desire for respect (15%).
- ☒ Insiders ranged from 18 to 59 years of age. Less than half (42%) of the insiders were female. A total of 83% of the insider threat cases involved attacks that took place physically from within the insider's organization. In 70% of the cases, the incidents took place during normal working hours.

Dealing with the insider threat has become more challenging in recent years as organizations now provide internal network access to a widening scope of users, including remote employees, partners, customers, subcontractors, consultants, and others. Because these sources are considered trusted, they are permitted access to sensitive corporate information. In this situation, organizations are becoming increasingly exposed to insider threats such as resource misuse, privacy violations, destruction of critical data, proprietary information loss, fraud, planting of logic bombs, and others.

The Impact of Regulations on Insider Threat Mitigation

Until recently, most organizations have neglected to protect themselves against the insider threat. Nonetheless, in the last few years, awareness of this issue has been increasing, especially in light of new information-intensive regulations (such as the Health Insurance Portability and Accountability Act, the Sarbanes-Oxley Act of 2002, the Gramm-Leach Bliley Act, and Basel II) that require organizations to protect the integrity of customer and employee personal information and corporate digital assets. As non-compliance with the expanding set of those regulations may carry criminal and/or civil penalties that can lead to the prosecution of individual executives or substantial fines, organizations are increasingly looking to implement solutions and practices that will help them comply.

Technology plays a significant role here, providing organizations with different means to meet regulatory demands such as maintaining information integrity, preventing unauthorized access, securing and storing communication records for future investigations, and so on. In fact, those requirements have become a major driver for greater spending on security, storage, backup, disaster recovery, information life cycle management, and other related IT product groups.

Dealing with the insider threat is one of the main security challenges that need to be addressed in the context of regulatory compliance. The fact that many of the abovementioned compliance-driven security requirements can be addressed by different insider threat mitigation solutions serves as a key driver for this market.

The following sections describe the main industry and government information-intensive regulations that are driving the adoption of insider threat mitigation solutions.

Sarbanes-Oxley Act (SOX)

The Sarbanes-Oxley Act of 2002 defines new requirements regarding the financial management of publicly traded companies. It is aimed at ensuring the integrity and the accuracy of reporting and preventing accounting errors and wrongdoings that may affect a company's shareholders and the general public. SOX lays responsibility on CEOs and CFOs, who must certify that their companies' financial reports are complete and do not contain any inaccurate or misleading statements. Non-compliance may lead to fines of up to \$5 million for individuals and up to \$25 million for entities, and prison sentences of up to 20 years.

Section 404 of the act requires companies to "provide reasonable assurance regarding prevention or timely detection of unauthorized acquisition, use or disposition of the registrant's assets that could have a material effect on the financial statements". The broad category of "assets" includes digital assets such as source code, trade secrets, M&A information, patient records, and any other sensitive information the unauthorized disclosure of which may have a negative impact on the company's stock price and its financial performance. Thus organizations are required to closely monitor the usage of those assets and be able to detect such events in real-time or near real-time.

Health Insurance Portability and Accountability Act (HIPAA)

The Health Insurance Portability and Accountability Act Privacy Rule defines administrative, physical, and technical safeguards for Covered Entities (CE), which include standards for maintaining the privacy of Electronic Protected Health Information (EPHI). These standards deal with several requirements such as the implementation of policies and processes on issues like assigning and controlling access to EPHI; reporting incidents; keeping tracks of EPHI moving in, out and within CE; and securing the transmission of EPHI over networks. Non-compliance with the Security Rule requirements may carry criminal penalties of up to \$250,000 in fines and jail time of up to 10 years.

Gramm-Leach Bliley Act (GLBA)

The GLBA Safeguard Rule requires all financial institutions to "develop, implement and maintain a comprehensive written information security program that contains administrative, technical and physical safeguards" to protect customer "Non-Public Information" (e.g., account numbers and details, social security numbers, credit card numbers, and so on). It mandates different requirements for safeguarding NPI, including the establishment of access controls of IT systems on which NPI is stored; encryption of electronic records; and monitoring of systems in order to detect intrusion attempts and attacks. Non-compliance with the Safeguard Rule may carry severe penalties in fines and prison terms of up to five years for individuals.

European Union Data Protection Directive

The Data Protection Directive, which outlines principals for protecting the privacy of individuals' Personal Identifiable Information of individuals (PII) on information systems, was adopted by the EU in 1995. The Directive requires each EU member state to pass local legislation that meets the different privacy protection principals that are an outcome of the OECD guidelines of 1980. Among those, Security Safeguards Principle 11 requires that "personal data should be protected by reasonable security safeguards against such risks as loss or unauthorized access, destruction, use, modification or disclosure of data". The Directive also classifies certain types of

personal data such as financial, health, and so on that requires additional safeguards given its sensitive nature.

Basel II Accord

The New Basel Capital Accord, due to become effective in most OECD countries by the end of 2006, requires financial institutions to calculate credit, market, and operational risks, to ensure they have enough capital reserves to cover risk exposures.

Although the Accord does not discuss information security measures directly, as operational risk is defined as "direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events", security breaches are considered as such and should be handled accordingly. In that respect, protecting the integrity of sensitive and customer private information and preventing its unauthorized distribution is aligned with operational risk management as the Accord requires. Internal Fraud is mentioned in the Accord appendix as one of the types of operational risks.

Japan's Personal Information Protection Act (PIPA)

Japan's PIPA was passed to "protect individuals' rights and welfare while preserving the usefulness of personal information. The Act's intent is to set out a policy for handling personal information, and measures for protecting it. Among other things, it requires companies maintaining personal information to protect its security and prevent unauthorized disclosure, loss, or destruction thereof; and supervise employees who handle it.

High-Profile Insider Abuse Incidents

The growing awareness of the insider threat was catalyzed by a series of incidents in which internal users were abusing their access to sensitive or private information for committing crimes or damaging their employers' business. The following highlights several high-profile examples of such incidents.

- ☒ In 2001 FBI agent Robert Hanssen was convicted of spying for the former Soviet Union and Russia after using his access to the bureau's computers to steal classified information for 15 years.
- ☒ In December 2002 a UBS PaineWebber systems administrator was charged with using a "logic bomb" to cause more than \$3 million in damage to the company's computer network. His failed plan was to drive down the company's stock with detonation of the logic bomb.
- ☒ In September 2004, a former helpdesk employee at Teledata Communications pleaded guilty to a scheme to steal and sell 30,000 consumer credit reports of the company's customers.
- ☒ In February 2005 a former AOL employee pleaded guilty to stealing and selling 92 million email addresses of 30 million AOL customers to known spammers.
- ☒ In April 2005 the Japanese police arrested a contract employee at an NTT DoCoMo affiliate for leaking personal information of 24,600 of the company's mobile phone subscribers in an alleged attempt to demonstrate system vulnerability.

- ☒ In April 2005 the New Jersey police arrested 9 people for what may have been the biggest bank security breach in U.S. history. Suspects in the insider fraud ring include bank employees at Wachovia, Bank of America, Commerce Bank, and PNC Bank. According to police investigators, the suspects — none of whom were IT workers — created a database containing stolen information from 676,000 accounts, using names and social security numbers. The bank employees would normally conduct between 40 and 50 searches of customer bank accounts in the course of a normal day. While the scam was underway, that figure jumped as high as 500 searches per day. After retrieving the account data, the bank employees allegedly either printed out screen shots containing the information or wrote it down by hand. The information was then sold to more than 40 law firms and collection agencies by a shell company.

Addressing the Insider Threat

The growing awareness of the insider threat has been pushing the adoption of different security solutions such as client antivirus, personal firewall, host-based intrusion detection, host-based vulnerability assessment, file encryption, database security, storage security, and other products. Most of these solutions, however, are aimed at protecting corporate networks against different external attacks designed to exploit insider vulnerabilities.

Given the elusive nature of internally generated threats, there is no one product that can address all possible scenarios. Hence, organizations usually employ different types of products for this purpose, each of which tackles different aspects of the problem. For example, many organizations use Identity and Access Management (IAM) solutions to control users' access to resources, implementing such technologies as single sign-on, user provisioning, authentication (including PKI), and directory services.

While IAM solutions are efficient for controlling information access, for the most part they cannot prevent authorized users from doing unauthorized actions with the data once access is granted. Given that, organizations are increasingly looking for other solutions in order to monitor actual data traffic to detect unauthorized delivery and usage of sensitive information.

In this regard, a new class of Information Leakage Detection and Prevention (ILD&P) solutions emerged over the last few years to provide dedicated response to this need by tracking the digital transfer of sensitive data. This category includes network-based products, which in most cases use sniffers installed next to corporate firewalls to monitor outbound network traffic via email, Web, IM, P2P, and other channels; and desktop-based solutions that use agents to monitor and control user activities, e.g., attaching a file to an email or IM message, printing, copying files to USB devices, burning to CDs, and so on.

Given the multiple channels and ways from which information can leak, a multi-layered ILD&P approach that combines network and desktop solutions could provide a comprehensive coverage of information leakage channels. However, organizations should still be aware that malicious insiders, for example, or disgruntled employees — if "motivated" enough — may eventually find a way around these security measures. For example, once gained access to sensitive data through an application that displays it on screen, a malicious insider can copy the data from the screen and write it on paper or use a cellular phone camera to photograph it.

Another type of threat that is not fully addressed by ILD&P solutions is fraud committed by authorized users that use legitimate commands to manipulate applications and data for financial gain or other reasons.

In addition, system administrators, database administrators, application programmers, and other privileged users pose a unique threat since they have both high-level access rights and the technical knowledge how to access and manipulate sensitive data. Monitoring these users is a challenge even for organizations that maintain application logs and other controls over regular end users.

Application-Level User Behavior Tracking

To address the above scenario of deliberate internal malicious acts, a different approach, which can either complement or in some cases serve as an alternative to IAM and/or ILD&P, is based on tracking user behavior on corporate applications. Thus, it can be referred as "application-level user behavior tracking".

Under this approach, internal users' activities in corporate applications are monitored in order to identify any type of abnormal behavior, as defined by corporate rules, which may indicate malicious activity is underway. Such abnormal behavior can be detected based on users' frequency of access to sensitive data, resources accessed, specific customer profiles, login time (i.e., after work hours) and location, session duration, and so on.

Tracking user activities in business applications allows for mitigating insider threats before they are "actualized". In other words, unlike ILD&P solutions, for example, which monitor actual outbound traffic to identify unauthorized delivery of recognized or alleged sensitive information, the application-level user behavior tracking approach is aimed at providing early alarm on insider activities that may eventually result in information leakage incidents or other harms.

For this approach to be effective, it should address the following requirements:

- Real-time inspection of user activities in corporate applications
- Comprehensive application coverage across heterogeneous IT environments, including legacy, client/server and Web applications
- Monitor all corporate users, including privileged users (e.g., DBAs, systems administrators, programmers)
- Provide post-event analysis for forensic purposes
- Operate in full transparency to end users without compromising performance
- Allow exception monitoring and handling based on corporate defined policies, which may include, for example, automatic suspension of a suspicious user

Securing Business Processes and the Legacy Challenge

One key driver for using application-level user behavior tracking is the difficulty of controlling application usage in complex environments. Today, many organizations are building IAM and other security mechanisms into their critical business applications. While this approach might work for small organizations, managing and

integrating security into each of the dozens applications typically used by large corporations is an administrative burden that translates to significant costs.

In addition, many large organizations are using legacy applications for mission-critical business processes, and over the last years have Web-enabled them to fit to today's open environments. However, given their complexity and the common lack of know-how required for modifying applications that were developed in previous decades, the need for controlling the usage of legacy applications is in many cases compromised. The old IT adage of "if it works, don't touch it" is in many cases a risky reality, as these legacy systems may become a security Achilles heel of corporate business processes that often span across multiple applications, including client-server, Web and legacy applications. To secure these processes, organizations are increasingly looking for automated, centrally managed solutions that provide the ability to centrally define use and access policies to all applications, and identify violations or suspicious activities as they occur.

INTELLINX USER BEHAVIOR TRACKING SOLUTIONS

Company Overview

Intellinx Ltd. was founded in 2005 as a spin-off from Sabratec, a developer of legacy integration solutions, which was acquired by Software AG in January last year. Founded in 1997, Sabratec deployed its products at more than 300 customers worldwide. In 2003, Sabratec started developing the Intellinx product for the purpose of real-time monitoring of business processes and user behavior on corporate applications that run in legacy environments. Following the acquisition of Sabratec, which included only the legacy integration division, the Intellinx division was transferred to the new company and the product was extended to support other operating environments.

Product Offering

The Intellinx product was designed to detect and act upon internal fraud by tracking user behavior patterns, provide an audit trail for investigation of suspicious events, and achieve regulatory compliance. To accomplish that, the product consists of three main layers:

- ☒ Record/Replay: All user activities in the internal business applications across the enterprise are recorded to allow internal auditors to visually replay them screen-by-screen, keystroke-by-keystroke. All end users across the enterprise are recorded including regular end users as well as privileged users.
- ☒ Content analysis: The content of the recorded screens is analyzed to identify screen headers and field captions and values. The information is stored in the system's database, allowing the querying of application usage data, e.g., which users accessed which customer accounts in a given timeframe.
- ☒ Business rules engine – this layer is used for tracking user behavior patterns and generating alerts on exceptions in real-time. The alert can be sent either as email or SMS or can initiate a pre-defined process in another system using protocols such as MQ, Web services or SNMP, for example, for suspending a suspicious user.

The above layers are based on a non-invasive technology, which applies an agent-less approach for intercepting network traffic between the server and the client. This is done using the Intellinx sensor (see Figure 3 below) – a network sniffer connected passively through a mirror port or tap device to the switch that connects clients to the host. The sensor captures all packets flowing to and from the host and passes the packets to the session analyzer through a queue.

The patent-pending session analyzer matches each packet to the session it belongs to, and stores each session's related packets in a manner that enables the reconstruction of a complete user session, including all screens and keystrokes, and replaying it upon demand.

The reconstructed screens are passed to the event analyzer, which in addition to content analysis is used for tracking end-user behavior patterns and generating alerts in real-time in cases of abnormal actions, according to predefined indicators. For example, as bank clerks typically search for customer account information by account number, frequent searches by account name can be defined as indicating malicious activity that need to be reported. Other examples for alert triggering may include repeated access to accounts after work hours, excessive usage of specific transaction types, excessive access to high profile accounts, and so on.

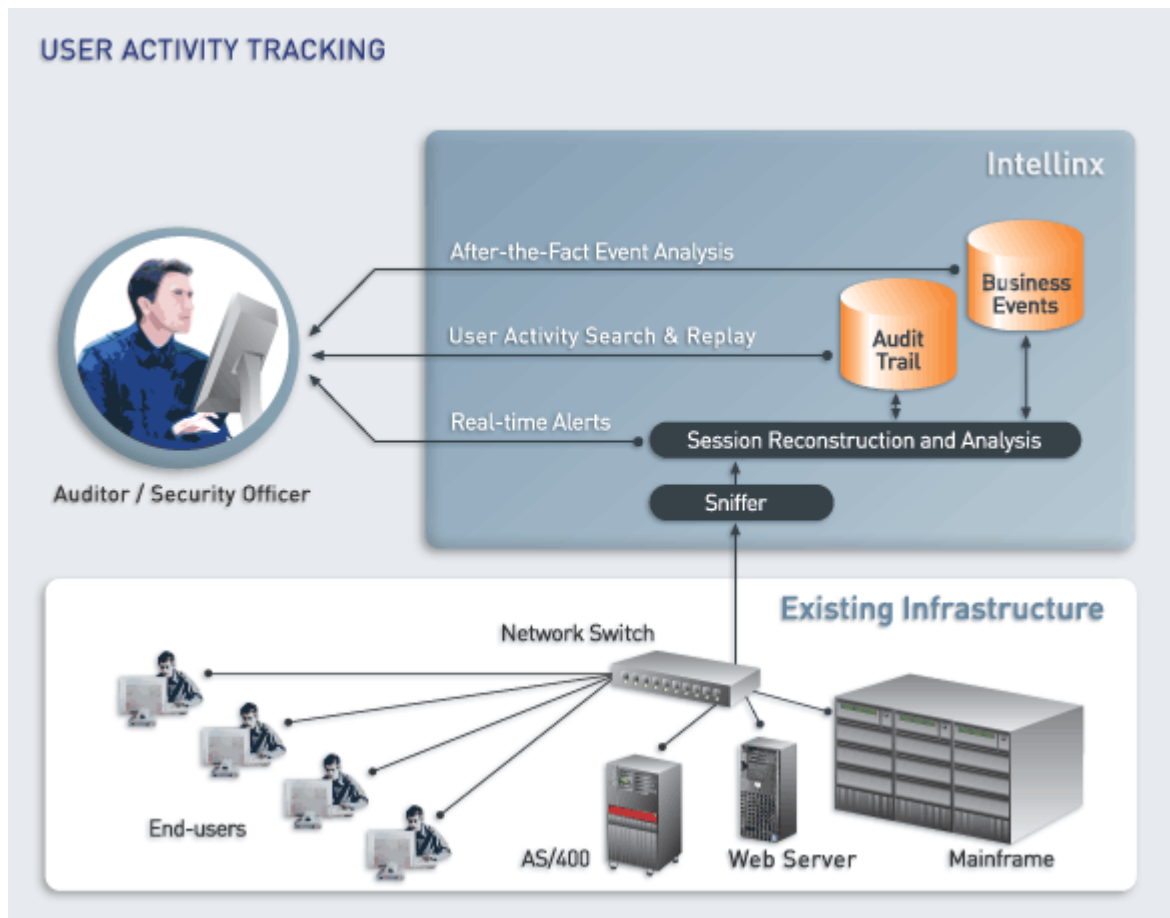
Maintaining a complete log of user activities on corporate application and data is highly important for forensics. While most organizations are keeping system logs for this purpose, in most cases these logs provide only transaction-level information, e.g., which users performed which transactions and when; or provide information on update transactions only. To tackle this issue, logs stored by the Intellinx system provide information on the exact customer data, for example, which was accessed by the user, and the exact actions that were made with it. In addition, the system provides information on read-only actions that due to the frequency of their occurrence are neglected from the log in most cases to avoid overhead. By including this information, the system allows to protect the integrity and privacy of sensitive customer information and mitigate the risk of malicious activities by authorized users.

As the Intellinx product stores the intercepted network traffic only, rather than screen images, the amount of disk space required is significantly reduced. In legacy systems, only the delta change between one screen and the next is transmitted. According to the company, recording a full workday of a single user requires about 50-60K in average. Recorded data is encrypted and digitally signed, allowing it to be used as court admissible evidence, if necessary.

The Intellinx product runs on a separate server (running Linux, Unix, or Windows), with no need to install or change anything in the host or clients' software or hardware, hence there is no impact on the system or network performance. The passive nature of the product architecture enables a simple implementation process, which does not require integration with existing systems. Once installed, the system can immediately start recording screen communications, allowing internal auditors to search and replay user activities.

FIGURE 3

Intellinx Monitoring Solution Architecture



Source: Intellinx, 2006

Building on the company's expertise in legacy systems, the first version of the Intellinx product supported 3270 and 5250 protocols over TCP/IP and SNA. The next version, currently available, added support for client-server architectures. Support for mainframe and AS/400 servers is also provided, covering such protocols as LU0, LU6.2 and MQ. The product's next version, due in Q2 2006, will add HTTP and FTP support.

The heterogeneous support allows Intellinx to provide large organizations with comprehensive application-level user behavior tracking, covering composite business processes that span over various applications and platforms. For example, it enables auditors to search for all the users who accessed a specific patient file across multiple applications on mainframe, iSeries, and Web-based systems.

CHALLENGES/OPPORTUNITIES

In light of the growing awareness of the insider threat, the market opportunity for products such as Intellinx could be significant. But many small start-ups and larger, established vendors are pursuing this opportunity and competition might be fierce.

Given this situation, Intellinx should highlight those specific product capabilities that are not met by most solutions in this area. For example, the ability to monitor legacy applications without changing code lines could be compelling for large organizations, especially in the financial services market where they are highly utilized. In addition, providing a comprehensive coverage of corporate applications across heterogeneous environments using an agent-less solution, which does not affect network or server performance, could be favored by organizations that are reluctant to engage in resource-intensive client deployments. The non-intrusive implementation could also appeal to organizations that are concerned about compromising the privacy of their employees.

As mentioned above, no one solution can provide complete coverage of internal threats. Given that, Intellinx should be seeking partnerships with other vendors operating in this space. In this regard, as the company's product is a natural complement to IAM solutions, working with IAM vendors could make much sense. Partnering with ILD&P solution providers to complement their offerings and create comprehensive internal threat mitigation package is an option that should also be considered.

Technology-wise, Intellinx should focus its development efforts on building automatic prevention capabilities into the product. Although the non-intrusive approach is highly valued by companies that are concerned of violating their employees' privacy, doing prevention in clear-cut cases, as defined by corporate policies, may be required by customers. Additional development efforts should revolve around helping organizations identify normal usage patterns that can serve as a baseline for detecting abnormal behavior. Adding automatically-created baselines to the Intellinx business rules engine could improve the product's detection capabilities and response time.

CONCLUSION

Driven by the growing awareness of the insider threat, demand for security solutions for protecting corporate sensitive digital assets from internal attacks is on the rise. In order to effectively mitigate such risks as data loss, theft, or misuse by trusted insiders, new security approaches are being increasingly required.

One such approach is provided by Intellinx. The company's product allows tracking user activities at the application level in heterogeneous environments, including legacy systems, in real time. Unlike other solutions in this space that provide a "last line" defense by monitoring outbound network traffic or user desktop activity associated with sensitive files, the Intellinx solution tracks suspicious user behavior patterns at the application level and provides early alerts of potential insider threats. Thus, it helps organizations reduce fraud losses, improve internal audit effectiveness, and achieve compliance with government and industry information-intensive regulations.

As the market for insider threat mitigation is rapidly evolving and competition is heating up, the company must establish a clear differentiation over the competition

and find the right partners and channels in order to capitalize on the market opportunity.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2006 IDC. Reproduction without written permission is completely forbidden.