Security and privacy
September 2006

Stopping insider attacks: how organizations can protect their sensitive information.

Contents

Introduction

More business transactions occur electronically every year, and organizations are retaining a growing volume of sensitive data. For many organizations, data has become an invaluable asset—the lifeblood of their operations. Access to this data is available to an expanding user base, including employees, business partners, suppliers and customers. IT infrastructures are more extensive, more complex, more distributed—and more accessible.

This interconnectedness affords many benefits for businesses, government agencies and consumers alike—but it also potentially introduces a great deal of risk. The more access points an organization maintains, the greater the possibility of compromised systems or data theft. And the stakes are high—especially as repositories of private information expand. Companies and government agencies must answer to stringent regulatory requirements and protect intellectual capital from competitors or subversive political entities. And consumers are typically most concerned with the potential for identity theft and other privacy violations.

Growing public concern over the security and privacy of personal data has placed many companies and government agencies in the spotlight—and countries around the world are developing regulations designed to support confidentiality. United Kingdom laws, such as the Data Protection Act of 1998, have evolved in recent years to guard against fraud and identity theft. In the United States, the state of California passed its Security Breach Notification Law requiring companies (regardless of geographic location) to disclose data fraud incidents to California state citizens.[1] Since then, 23 other states have passed their own disclosure laws. And Canada plans to adopt similar measures. These recent legislative actions, as well as frequent media reports of security breaches, have made the prevalence of such threats clear.

Highlights

In its 2005 global business security index report, IBM identifies a burgeoning trend toward small, targeted attacks—rather than sweeping global threats such as worms, spam, viruses and other malware.[2] Insider attacks, in particular, can be a significant threat to the security and privacy of data. This white paper seeks to provide the reader with a better understanding of the issue of insider attacks and offers suggestions that may help organizations mitigate their risk.

The growing threat of insider attacks

Strong perimeter defenses can block external threats effectively, but provide only part of the protection organizations need.

Organizations have spent decades counterbalancing more open, connected and distributed networks with stronger defenses against intrusion, including firewalls, anti-virus software, biometrics and identity access badges. These measures have made the business world more effective at blocking threats from the outside, and made it increasingly difficult for would-be hackers or viruses to penetrate systems. Such technologies, however, are largely passive in approach and designed to thwart only *un*authorized access. They provide only a first line of defense.

Though often overshadowed by attacks from the outside, the risk of insider threats is nevertheless a pressing concern for practically every organization.

"The Global State of Information Security 2005," a study published by PricewaterhouseCoopers and *CIO*, showed that 33 percent of information security attacks originated from internal employees, while 28 percent came from ex-employees and partners.[3] Though certainly not forgotten by businesses, government agencies and industry analysts, the risk of security breaches from within is often overshadowed by more dramatic intrusions such as denial-of-service attacks, widespread viruses, or outright thefts of intellectual or financial capital.

Highlights

"Dishonest insiders" can exploit an organization's vulnerabilities to commit identity fraud and expose confidential information—for personal gain or as part of a larger crime ring.

A list kept by the Privacy Rights Clearinghouse shows hundreds of data breaches reported in the United States alone since February of 2005.

The prioritization of external threats over internal dangers is misguided, however, and can leave gaping vulnerabilities in an organization's defenses. The back-office processing and customer support firm HSBC Electronic Data Processing (India) Private reported in June, 2006, that an employee, as part of a larger theft ring, accessed customer debit card information and used it to defraud 20 UK customers of US$425,000.[4] The incident represents only one of many in the past few years. A list kept by the Privacy Rights Clearinghouse in the U.S. shows hundreds of data breaches since the highly visible ChoicePoint breach in February of 2005.[5] Many of these breaches originated from within the organization, by what the list refers to as "dishonest insiders." Consider the following examples found on the Privacy Rights Clearinghouse list:

- *At one full-service securities firm, a former employee illegitimately accessed more than 100 customer records.*
- *A hotel's systems were compromised by either a dishonest insider or a hacker, exposing 55,000 records—names, addresses, credit card details, Social Security numbers, driver's license numbers and bank account data.*
- *"A dishonest insider or possibly malicious software" accessed the systems of an Internet billing company, exposing names, phone numbers, addresses, e-mail addresses, Internet Protocol (IP) addresses, login names and passwords, credit card types and purchase amounts online.*
- *At an insurance company, an employee accessed confidential data, including names, Social Security numbers, birth dates and addresses on foreclosure properties, and used the information for her own personal gain.*

Several other cases were associated with stolen laptops, lost backup tapes or unauthorized setup or use of accounts.[5] While not strictly attributable to "insider attacks," these incidents can put an organization at risk in the same way—by exploiting authorized channels to bypass perimeter defenses and, thus, escape detection.

Highlights

Because employees carry valid authorization and are privy to the organization's vulnerabilities, insider attacks can be more difficult to detect than external penetration attempts.

According to a recent study, the average fraud scheme continues undetected for 18 months.

Undetected attacks can cause serious harm, including legal liability for compromised data, loss of competitive position and disrupted business operations.

Your organization at risk: understanding the stakes

Attacks from the inside carry the potential for significant damage that can rival or even exceed the damage caused by external forces. As an integral and trusted member of the organization, the perpetrator carries valid authorization and typically enjoys relatively unchallenged presence and movement within the organization's IT infrastructure. The attacks typically target specific information and exploit established entry points or obscure vulnerabilities. In many respects, insider attacks can be more difficult to detect than penetration attempts from the outside.

In its "2006 ACFE Report to the Nation on Occupational Fraud and Abuse," the Association of Certified Fraud Examiners (ACFE) indicates that most cases of general fraud—asset misappropriation, corruption or fraudulent statements—are exposed by accident or through tips given by employees, suggesting a widespread lack of effective monitoring and oversight capabilities. The report also finds that the average fraud scheme continues for 18 months before being detected.[6]

Internal attacks that continue undetected can cause serious harm to an organization. Perhaps most significant, they can expose the personal information of customers or employees. A breach of this kind—whether it is identity theft, inappropriate use of data or the sale of sensitive information—can leave an organization legally liable for associated damages and subject to regulatory fines. In addition, a company's competitive position could suffer if an insider uses intellectual property or trade secrets for unauthorized purposes. Attacks may also be designed to extort money or damage an organization's reputation. If they lead to IT downtime or damaged systems, they can also disrupt business operations and reduce the value of IT investments.

With so much at stake, it is becoming increasingly important to address the threat of insider attacks—*before* they occur.

Highlights

Distributed, global work environments and rapidly changing business conditions require a balance between end user accessibility and data protection.

Protecting against attacks from the inside requires greater sophistication and granularity on the part of security systems.

There are four basic elements that can provide the sophistication needed to help prevent insider attacks.

Building greater sophistication into security measures

Today's distributed environments and rapidly changing business conditions (such as mergers and acquisitions, layoffs and global sourcing) make for a wide geographic distribution of users, a system of multiple entry points and the potential for disgruntled employees. As a result, today's organizations carry greater risk of insider attacks. Every organization must adopt a strategy that can help manage that risk effectively, striking a balance between end user accessibility and protection against security breaches.

When watching for insider attacks (as opposed to external threats), the security question changes from "Is the access authorized?" to "Is the behavior acceptable?" Whereas the former question asks for a simple yes-or-no answer at a single point in time, the latter question addresses much more complexity. A user's behavior encompasses all events in a given session from beginning to end, and involves long-term patterns and subtle variations. Answering this question requires more sophistication and granularity on the part of the security systems. As we intend to show in the next few pages, there are four basic elements—behavioral analysis, integrated security components, automatic response and an iterative modeling process—that, as part of a comprehensive approach to the threat of insider attacks, can help provide this level of security sophistication.

Behavioral analysis

The key to thwarting an insider attack lies in understanding the range of normal behavior in a given business process and pinpointing behavior that deviates from the norm. Thus, one of the first steps must involve policy making—the definition of parameters for acceptable behavior within a peer group. These parameters will serve as the baseline for comparative analysis, so it is important to establish user profiles based on historical data or concrete

Highlights

experience—not just business expectations that may or may not be realistic. (Parameters set too wide can let some dangerous behavior go unnoticed, while too-narrow parameters are prone to trigger a flood of false positives.) As user roles change, organizations should update the profiles accordingly.

Security systems should automatically monitor the online activities of authorized users, detect abnormal behavior and even help to prevent potential misuse.

Using the parameters as a baseline for comparative analysis, security systems should automatically monitor every aspect of the online activities of authorized users—from the beginning to the end of each session. Not only should systems have the ability to identify abnormal behavior through comparative analysis, they should also be able to predict and even help to prevent potential misuse—by responding immediately to certain trigger events. Systems should monitor variables such as:

- *Initial connection—date and time of logon, IP addresses involved and connection frequency*
- *Data access—requests for data, organized according to specific type*
- *Application usage—frequency and duration*
- *Overall usage—total session time and overall data usage requests.*

Behavioral analysis can help pinpoint small deviations and unusual patterns in high-traffic, dynamic work environments.

Behavioral analysis can be indispensable in high-traffic, dynamic work environments such as call centers where customer information can be vulnerable to fraud or misuse. The employees who work in these centers have extensive access to customer records, but they can be expected to access a predictable number of them during the workday. For example, if it has been determined through historical analysis that each agent in a certain call center typically accesses 10 to 15 records per day, it may be reasonable to investigate an agent who accesses 30 or more. Likewise, an organization may deem a situation suspicious if an agent views information that is not normally required for customer interactions. Only through ongoing, controlled behavioral analysis can an organization identify these deviations.

Highlights

Security elements should interact seamlessly—in real time—to enable thorough analysis and quick response to potential threats.

Integrated security components

Many organizations have at least some of the security elements needed to protect against malicious internal attacks: authentication systems, asset tracking software, device and Internet usage monitoring capabilities, and other tools. It is critical, however, for these pieces to interact as seamlessly as possible. Indeed, one reason organizations find it difficult to detect insider attacks is the time it takes to analyze a vast amount of data coming from a wide array of devices, entry points and user accounts.

Organizations need to enable communication, correlation and analysis at a granular level among a wide range of security components, including authentication gateways, physical security systems, asset management tools, network monitoring capabilities and Web security platforms. These systems should communicate in real time so the organization can react quickly before data can be used for illegitimate purposes—and potentially even predict and prevent malicious attacks.

The systems an organization puts in place to monitor user behavior should also be designed to simplify monitoring and pattern detection tasks for administrators. Administrators should be able to access a central console that compiles messages and events from systems that monitor everything from network devices to application usage. Manually reviewing historical logs and searching for complex relationships across systems can divert too much effort away from activities of higher value and priority.

Effective pattern detection depends on the ability to correlate messages and events from different monitoring systems across the IT environment.

Consider how much more powerful an organization's pattern detection capabilities can become when events are correlated across the IT environment. For example, an organization may run a sensitive application that generally should not be accessed remotely. If an employee logs on to that application without having passed through physical access points, such as a badge reader or an

Highlights

onsite workstation, an integrated system can immediately identify the behavior as unusual and potentially harmful. Without this automatic, real-time correlation, the remote access may not be detected quickly enough. A delay of even a few hours can provide an ample window of opportunity for a would-be attacker. Or, as another example, a credit card call center may field several customer complaints of erroneous billing over a period of weeks. An administrator with employee access records can quickly detect deviant behavioral patterns within the same time period.

Automatic response

Organizations need to recognize and respond to deviations from normal behavior as quickly as possible. Relying only on human detection and response may not suffice, especially if an attack occurs during non-business hours.

The security systems themselves must be capable of responding immediately to unacceptable user behavior.

To prevent or mitigate damage, the systems themselves must be capable of acting immediately in response to unacceptable behavior. Once the behavior departs from the standard beyond a certain threshold, for example, the system should deny access to a requested application or data resource. This near-immediate response allows time for network administrators to receive an alert, analyze the patterns and choose an appropriate course of action. And the network administrator should not have to maintain deep security expertise to interpret the data or determine the next steps. The security systems should automatically suggest a range of relevant responses that are based on the latest research and insight into security threats. In addition, the systems should be capable of sorting through false positives. An alert system that simply passes information along without basic levels of analysis fails to add value to the monitoring process.

Automatic denial of access can thwart attacks before they occur—and give network administrators the opportunity to determine a suitable course of action.

Highlights

To stay a step ahead of evolving security threats, organizations must continuously revise and enhance their security efforts.

Self-tuning systems should react appropriately and intelligently to dynamic business conditions—without human intervention.

Iterative modeling process

No matter how much an organization prepares for today's security threats, the risks continue to evolve. Employees come and go. IT infrastructures grow and incorporate new technologies that can introduce unforeseen vulnerabilities. To keep sensitive data protected, organizations must work continuously to remain a step ahead of potential attacks. Security systems should play a significant role in these ongoing efforts.

It is important not to limit detection systems to narrow, specific rules, because the range of valid behavior shifts over time. Instead, organizations should institute self-tuning systems that can react appropriately and intelligently to dynamic business conditions—without the need for a full redefinition of rules. For example, a call center may frequently alter the average length of a call to reach certain cost or customer satisfaction objectives. Or a marketing campaign might require agents to access data that is not normally needed. Without the ability to adapt dynamically to these kinds of changes, the security systems may inundate administrators with false positives—thus, reducing the value of the alerts. At the same time, the systems need thresholds that are sensitive enough to detect subtle deviations within large samples of behavioral data. Striking a balance between the two extremes can be done only through an iterative modeling process, wherein monitoring systems can learn the organization's natural rhythms and sort through several overlapping layers of acceptable behavior.

Highlights

Organizations must be prepared to fend off attacks wherever they originate—even as the boundaries between organizations, partners, users and customers blur.

Conclusion

Recent events have shown that organizations across industries cannot afford to continue to ignore the potential for insider attacks. As organizations grow, they employ workforces that are increasingly spread across geographies; they implement systems that are more heterogeneous, more complex and more connected; they retain more confidential data; and they are subject to changing regulatory requirements.

Traditional boundaries between organizations, partners, users and customers have become blurred, making security policies more difficult to define and enforce. Organizations must be prepared to fend off attacks wherever they originate—by addressing vulnerabilities in precisely this gap between traditional business and the open, distributed organizations of today and the future.

For more information

As a leader in the security and privacy space, IBM Global Services can help you learn more about the threat of insider attacks and consider potential ways to address it. Our IBM Information Security Framework, designed to provide a methodical and efficient approach to key security issues, can help organizations address their evolving threats, risks and business demands as they relate to data protection and overall security. In addition, the IBM Center for Business Optimization can provide unique insights into security and privacy issues, and can help organizations create effective strategies through advanced mathematical research, business performance management, business intelligence systems, software and deep computing.

For more information, contact your IBM sales representative or send an e-mail to:

IBM Center for Business Optimization
Toby Cook, Associate Partner, IBM Center for Business Optimization—toby.cook@us.ibm.com.

IBM Information Security Framework
Michel Bobillier, Global Offering Executive, IBM Security and Privacy Services—bobillier@ch.ibm.com

Or visit:
ibm.com/services

IBM Global Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
09-06

---

1 California Security Breach Information Act (S.B. 1386), enacted July 1, 2003; http://info.sen.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.html

2 IBM global business security index report, 2005.

3 Scott Berinato (with Research Editor Lorraine Cosgrove Ware), "The Global State of Information Security 2005," September 15, 2005, published by PricewaterhouseCoopers and CIO; http://www.cio.com/archive/091505/global.html

4 "A Chronology of Data Breaches Reported Since the ChoicePoint Incident," Privacy Rights Clearinghouse; August 5, 2006, used with permission of the Privacy Rights Clearinghouse, www.privacyrights.org

5 John Ribeiro, "HSBC claims customer fraud in Indian services center," Network World (IDG NewsService), June 27, 2006; http://www.networkworld.com/news/2006/062706-hsbc-claims-customer-fraud-in.html

6 "2006 ACFE Report to the Nation on Occupational Fraud and Abuse," Association of Certified Fraud Examiners; http://www.acfe.com/fraud/report.asp

GSW00316-USEN-00