

Insider Threat *Deep Dive*



Combating the enemy within

Sponsored by



How to thwart employee cybercrime

To guard against the insider threat, trust your detection and prevention countermeasures, not your instincts

By Roger A. Grimes

THE MOST DIFFICULT PART about insider threats is that they are almost always a surprise. The stereotype of the sullen malcontent who pilfers data after hours can blind management to other possibilities. Just as often, it's an ambitious, well-liked employee who turns out to have a hidden criminal agenda.

Unmonitored, any insider with the appropriate privileges is a potential threat. That's why managers who rely on personal instincts alone put themselves and their companies at risk. A clever criminal can lull the boss into believing nothing is amiss. Systems designed to monitor the network for patterns of criminal or destructive behavior are much harder to fool.

Detecting and preventing insider threats is its own discipline, with its own dynamics. The good news is that if you put these countermeasures in place, you can reduce the threat dramatically. Before delving into the details of detection and prevention, however, it pays to get an idea of the scope of the problem.

THE SIZE AND SHAPE OF THE PROBLEM

How big is the insider threat, really? Many sources have long cited an estimate that insiders are responsible for 80 percent of malicious attacks. Unfortunately, no publicly available statistics back up that statement – in fact, they indicate a much lower percentage. And yet, for reasons that will become apparent, that 80 percent figure is probably pretty accurate.

The [2009 CSI Computer Crime survey](#), probably one of the most respected reports covering insider threats, says insiders are responsible for 43 percent of malicious attacks, with 25 percent of respondents indicating that more than 60 percent of their losses were

due to nonmalicious actions by insiders.

Another respected source, the [2009 Verizon Data Breach Report](#), pegs data breaches due to insider attacks at only 20 percent. Why the big difference? For one thing, the Verizon report states that trusted partners are responsible for another 32 percent of data breaches – and that insiders are responsible for three times more breached records per incident than either external attackers or partners. So a weighted average of the damage due to trusted insiders moves that 20 percent much closer to 60 percent.

Add users with nonmalicious intent, and the percentage of insiders involved (wittingly or unwittingly) in a malicious attack is probably at least 80 percent. These days, most malicious attacks employ socially engineered Trojans, where someone is tricked into installing malware. Therefore, if you include unwitting accomplices, insiders are directly responsible for almost every attack that succeeds in breaching outer defenses using spam or phishing.

According to the [Microsoft Security Intelligence Report 8](#), almost 80 percent of all the successful malware threats detected and cleaned in the last quarter of 2009 required human complicity (such as installing Trojans, spyware, or downloaders). And the Microsoft report doesn't include the complete impact of spam, phishing, accidental mistakes, intentional maliciousness, and misconfigurations. The threat may have originated with an external attacker, but it would not have been successful without an end-user to play the goat.

The Verizon report indicates that end-users and IT administrators are by far the biggest culprits in internal data breaches, with regular end-users just nudging past admins. And a growing percentage of trusted partner attacks can be attributed to the actions of partners' administrators.



The same report goes on to say that two-thirds of data breaches attributed to insiders were intentional. A whopping 83 of all attacks, whether internal or external in origin, were simple in nature. The most telling statistic of all: 87 percent of attacks could have been avoided through simple or intermediate controls.

DETECTING INSIDER ATTACKS

Because insiders are usually trusted and have authorized access to internal systems, it's often hard to detect when they are performing unauthorized activities. But like minor criminals who get greedy, most insiders who get caught did something that would be considered aberrant in the course of their normal duties. An array of monitoring, scanning, and other techniques can help you determine if something unpleasant is afoot.

LOOK FOR UNUSUAL TRAFFIC PATTERNS

Most data breaches involve moving large amounts of data between unauthorized locations. If you aren't already doing so, find and use a tool that monitors network traffic patterns.

For best results, you'll need to spend weeks or months determining normal traffic baselines. Which computers normally talk to which computers? What is the normal traffic pattern? How much data is typically transferred, and in what amount of time? How many workstations talk to other workstations? Most shouldn't. The normal data or file server doesn't talk to a whole bunch of other servers, like a network management or anti-virus update server might. In any environment, you will find a normal flow of data between various machines. What a good detection system does is alert you to aberrant patterns and amounts.

Establishing good baselines is hard. Many computers have periods of low and heavy use depending on the system and its use lifecycle. For example, a patching server will experience its heaviest use when patches are released and pushed, whereas an accounting server typically reaches peak activity at the end of every month.

Still, take the time to create meaningful baselines and create alerts for the unusual events. However, always be prepared for the aberrant yet "normal" events that occur on any network. One big company event or news event (such as a natural disaster) can blow your baselines out

of the water -- but when your alerts go off, investigate the cause, even if the reason may seem obvious. If you're getting too many false alerts, you need to reset your baselines. The best designs minimize false positives and ferret out the malicious insiders near the beginning of their attacks.

REVIEW EVENT LOGS

Event logs and the content they collect frequently contain signs of an intruder's attack, but in many enterprises, logs are neither collected nor analyzed. Setting up a system to collect and analyze log file information is not trivial, but it has become vital to the health of any modern organization (see InfoWorld's [Log Analysis Deep Dive](#)).

As with a network monitoring tool, you will need to fine tune your event log management system to alert you only to incidents that require further investigation.

What type of information should trigger alerts? It depends on the environment and system, but certainly unexpected and excessive log-ons, attempted log-ons to old account names, unusual activity times, and a sudden appearance of previously unrecorded critical errors should raise alarms in any event log management system.

RECORD DATA ACCESSES

Companies that are most successful at detecting trusted insiders conducting unauthorized activities use systems that track data access -- and record who is responsible for each instance of data access.

Ultimately, all computer security defenses are about protecting data, and nothing provides a clearer data trail than a record of who accessed what. For example, many hospitals have discovered employees accessing the health records of acquaintances or celebrities. To avoid similar snooping, police departments often track who runs background checks on whom. Companies with these types of tracking abilities often educate employees about the feature in order to "keep honest people honest."

USE DATA LEAK DETECTION/ PREVENTION PRODUCTS

Internal attackers often focus on stealing confidential data. Data leak detection and prevention tools are designed to flag potential data leaks. Products can be



implemented as software on each managed device or on the network looking for emitting data streams. No product can provide perfect protection, but many do a capable job of bringing suspicious patterns of data activity to light.

SCAN FOR HACKING TOOLS

Many insiders download and use well-known hacking tools to begin their explorations and compromises. Most antivirus scanners detect the most popular tools. Make sure your anti-malware scan has these capabilities and that they are enabled. If hacking tools are found, you have a choice: Confront the employee immediately or enable detailed tracking to try and determine the employee's intent and target.

USE INTRUSION DETECTION/ PREVENTION SYSTEMS

Intrusion detection or prevention systems detect insider attacks that use known exploits. The key to detecting and preventing insider attacks is to place these systems on managed hosts and in the middle of internal data streams. They can work in conjunction with network traffic baselines to create incident alerts for aberrant events. Unfortunately, many insider attacks do not involve exploits that would trigger these systems (such as regular log-ons using authorized credentials), so be aware of that significant limitation.

LOOK FOR FUNKY FILES

Attackers stealing data often create very large compressed files or use archive formats that are not used in the normal course of business. Look for very large data files or unexpected file extensions appearing in unusual places.

CONDUCT RANDOM AUDITS

Companies worried about internal attacks should conduct random audits on internal employees. Of course, you should also focus on high-risk employees, contractors, and partners (where permissible). Audits should review local event logs for unusual log-on activity and look for hacking tools and unexpected files. If employees are given company laptops or home computers, random checks should look for unauthorized company

data. If employees use portable storage media, the media should be inspected.

DEPLOY HONEY POTS OR RED HERRING DATA

A honey pot is any computing device used for unauthorized access detection. Take a couple of computers you're getting ready to throw away, place them on the network in attractive locations, and configure them to send alerts if anyone tries to connect to them (recording originating IP address information and log-on credentials tried).

You will need to spend a few hours for a couple of days filtering out the false positive log-on attempts (from patching servers, anti-virus update servers, and so on). But after all the filtering is completed, no other computer or person should try to access the system. Honey pots are fake systems and no one should try to contact them.

You may want to go as far as installing specialized honey pot software (such as Honeyd or Kfsensor) to give the honey pot a particular persona (a Web server, SQL server, e-mail server, etc.) and to gain additional detection functionality that you might not have with a regular computer. Honey pot systems should be prevented from connecting to the Internet or any other system in your environment to minimize security risk. Although some types of honey pots allow intruders to log on and gain full access to the system, the real point of a honey pot is simply to detect unauthorized log-on attempts.

Don't underestimate the value of a few well-placed honey pots. A honey pot is a cheap, low-noise (that is, it generates few false positives) early warning system. Honey pot systems are a great way to detect malicious activity from previously undetected malware and intruders. Trusted insiders who are exploring the network beyond their authorized level of authority often run afoul of honey pots.

Knowledge, use, and installation of a honey pot should be kept to the smallest group possible. For obvious reasons, you don't want insiders to be aware of them. You can even keep honey pot systems from the knowledge of the incident response team and other administrators. If someone asks how the unauthorized activity was detected, people in the know can claim it was found using a normal intrusion detection system, router log, or system event log.



If your enterprise is hesitant to deploy full honey pot systems, consider sprinkling “red herring” data around your environment. Red herring data is fake data that should not appear anywhere else in the environment. Some companies create fake user records with unique names, and then scan for copies of these fake records outside the original systems. Another good technique is to rename your administrator and root accounts and immediately track any log-on attempts using the original names that no longer exist.

PREVENTING INSIDER ATTACKS

Detection is great; prevention is better. A good detection system can provide early warning of unauthorized activities and function as a prevention tool. The following ideas, tools, and techniques will help you prevent attacks before they occur.

START WITH BACKGROUND CHECKS

Companies concerned about internal attacks should perform employee background checks. As listed above, half the internal attacks involve IT administrators, but the other half involves regular employees. Employee background checks should cover former employers as well as state and federal criminal checks. Many internal attacks are committed by repeat offenders who have signed forms allowing background checks that were delayed or never performed.

EDUCATE WITH POLICIES

New employees should be made to sign acceptable use policies that outline what is and isn't allowed. Most acceptable use policies should clearly state that accessing unauthorized systems, processing data in unauthorized locations, or possessing hacking tools are terminable offenses. Acceptable use policies should educate the user about good password usage, not sharing passwords, locking unattended workstations, physical security protections, data encryption requirements, and other policies designed to safeguard the enterprise's data and assets.

IMPLEMENT A SEPARATION-OF-EMPLOYMENT PROCESS

Many data breaches are performed by fired employees. Every enterprise should have a very specific and

strict set of processes that occur when an employee is no longer employed. This should include removing the person's access to physical buildings, collecting all company assets, changing all passwords (including any passwords to other accounts they may have learned), and disabling previously enabled network access. Don't forget to terminate remote access methods as well.

CHANGE CONTROLS WHEN EMPLOYEES CHANGE POSITIONS

Most companies are great at giving an existing employee new permissions and access when they move to a new position, but are not as good at removing the old permissions and accesses that are no longer needed. Often, an employee moving to a new position is asked to continue to assist with the old position until a replacement is up to speed, but then admins forget to remove the old access rights. All companies should have clear policies that stipulate how access and permissions are handled when an employee changes positions. Any passwords previously known to the moved employee should be changed if they are no longer needed.

CREATE THIRD-PARTY AGREEMENTS

Trusted partners are responsible for a significant portion of data breaches. All partners, contractors, and third parties (including clients who have access to your systems) should be required to sign an acceptable use policy describing what is and isn't allowed. At the very least, all systems used to access your data should follow the usual best practices: up-to-date anti-virus programs, enabled host firewalls, secure configurations, fully patched operating systems and applications, and so on. Third parties should be required to use encryption during remote accesses and to store all confidential data at rest. Some companies go so far as to require that third parties allow random audits. Third-party clients found to be out of compliance should face warnings or disciplinary actions. Repeated incidents should result in a termination of access.

ENFORCE SEPARATION OF DUTIES

“Separation of duty” is a long-held accounting control. Policy controls are put in place and enforced to prevent a single employee from conducting an action that



could prove significantly harmful to the company. For example, an employee that approves payroll amounts should not also sign or print the checks.

In the IT world, separation of duties should also be deployed where reasonable. Here are some examples:

- Account and group creation should require manager approval.
- Highly privileged accounts should be separated from the user's regular account.
- IT administrators' regular accounts should not be able to access the encrypted data of other accounts without an explicit sign-off.
- Highly privileged account creation should be approved and performed by people who will not use the accounts.
- Some companies go so far as to require two people to enter a compound password for a highly privileged account, with no single employee knowing the whole password.

As you can see, the idea is require two or more employees to collaborate on very sensitive roles and actions, making it less likely that unauthorized activities will occur. Also, IT auditors love to hear that IT departments understand the concept of separation of duties and are more likely to have confidence in who they are auditing.

USE LEAST-PRIVILEGE ACCESS CONTROL

Implementing least-privilege access control is one of the best ways to prevent insider attacks. If attackers and their tools can't access data, they can't steal it. Access control means configuring file and folder permissions, as well as separating users from unneeded networks and computers.

Access control should be granted by the application or data owner, with each required user and their least privilege access determined by the most knowledgeable person involved (often not the application or data owner). All access control should be audited periodically, with the owner asked to reaffirm the list of members and their access. This process should be automated where possible. If you fail to manage access control over the lifecycle of applications and data, you will surely suffer unneeded access in the future, without knowing one way or the other who really needs the access they enjoy.

USE ROLE-BASED ACCESS CONTROL

Role-based access control (RBAC) is the process of providing least privilege access control determined by the authorized actions and role of the user needing the access. In RBAC systems, least privilege permissions are granted to role-based groups instead of user accounts or departmental groups. At the very least, RBAC systems help administrators implement least privilege access control by forcing IT employees to think about the bare minimum permissions needed for the employee to perform his or her job.

Advanced RBAC systems go even further and allow employees access only during specific tasks. For example, in a traditional HR system, all HR employees might be given full control to the HR system and database. In an advanced RBAC HR system, there are many roles or groups, each with only the permissions needed to do a specific job or task (HR administrator, benefits administrator, payroll supervisor, data entry clerk, and so on). For example, when a data entry clerk enters payroll records, the clerk gets write access to the payroll database. When the clerk runs a report, the permission reverts to read-only. And when the data entry clerk is not logged on to the system, he or she has zero access to the databases behind the system (something absolutely not true in non-RBAC systems).

AUTOMATE WORKFLOWS

Data isolation can also be performed by automating workflows that require interaction with data. For example, hiring a new employee could literally cause a half-dozen employees to interact with a dozen databases.

In a case like this, it's best to isolate all the involved users in a tightly controlled user interface that never allows the involved employees to access the actual databases at all. Instead, each involved employee is sent simple questions to answer, in the form of an e-mail or word processing document, for example. The responses manipulate the involved databases and automate other workflows to other involved employees.

To extend the HR example, once HR enters the new employee's information, a confirmation e-mail is sent to the new employee's boss to confirm the hiring details. That confirmation automates the process of putting the employee in all the necessary security groups. A good



automated workflow minimizes mistakes and minimizes employee access to the database behind the scenes.

ISOLATE DOMAINS

In most networks, most workstations don't need to connect to other workstations. Most servers don't need to connect to most other servers (although they may need to connect to a few). If users don't need access to a particular computer or network, don't give it to them. Domain isolation gives users access only to the computers and application necessary to perform their jobs. If users can't access it, they can't hack it (as easily). Domain isolation can be set up using access controls, routers, firewalls, IPSec, network access control, and so on. Domain isolation is just access control taken to the network level. For example, my company has tens of thousands of servers. When I connect, I can connect to exactly 15. I don't know what else exists. I can't ping the other computers, and I certainly can't easily hack them.

USE IDENTITY MANAGEMENT SYSTEMS

Identity management systems help manage the user account lifecycle, from the initial creation (provisioning) to deletion (deprovisioning) and everything in between. They help automate necessary workflows and isolate administrators from the databases behind the scenes. Besides significantly reducing data entry errors and identity synchronization across multiple systems, identity management systems can reduce the security risk of excessive permissions and left-behind, stale accounts.

USE DATA LEAK PREVENTION

Data leak protection systems are uniquely qualified to prevent employees from accessing or transferring unauthorized, monitored data. Host- and network-based systems have proven their value to many companies and stopped many security incidents. The key is to configure the data leak protection system to recognize unauthorized data and to minimize the number of false positives.

ENCRYPT SENSITIVE DATA

Lost storage media is involved in a large number of reported data loss claims. Internal attackers often copy company data to portable storage media and take them off premises. Require and implement encryption to protect sensitive data, both in transit and at rest. Require encryption on storage media, such as tape, portable hard drives, and USB keys. Many systems today can require that all portable media be encrypted in order to store data on them and to prevent the accessing of that data on nonsystem computers.

HARDEN WORKSTATIONS AND SERVERS

Many compromises happen due to poor default security or misconfigurations. Ensure that all computers have been configured and secured using industry-accepted best practices. Most OS vendors configure their products to be reasonably secure by default, and several trusted entities offer security configuration baselines (such as www.nist.org or www.cisecurity.org), which can be used as starting points for security evaluation.

IMPLEMENT CONFIGURATION AND CHANGE MANAGEMENT

Once you have least-privilege access controls and secure default configurations, you need to make sure computers don't get modified or "drift" to less secure states. All computers should be periodically audited to ensure they remain in the required secure state and require a change control process (and monitoring) to prevent unauthorized modifications.

Companies that follow the detection and prevention recommendations discussed here can vastly decrease the risk of internal attacks. You may not be in a position to implement all of the systems and processes we've covered, but if you don't take at least some of these countermeasures, you leave the inside door wide open to those who would betray your trust or do unintentional harm to your network and your business.