

# Predicting, Detecting, and Responding to Insider Attacks

By E. Eugene Schultz – ISSA member, Los Angeles, USA chapter

**Insider attacks are one of the least understood and inadequately mitigated security-related risks facing organizations today. This article identifies patterns and behaviors that will help security professionals defend against the malicious insider.**

Insider attacks are one of the least understood and inadequately mitigated security-related risks facing organizations today. The fact that the term “insider” is in and of itself elusive only makes this issue and its many dimensions more difficult to truly understand. Although employees of an organization are unquestionably insiders, whether (or perhaps better said, the degree to which) the term “insider” applies to consultants, contractors, vendors, third-party business partners, and others is usually to at least some degree ambiguous. Many believe that anyone who is granted access to an organization’s computing resources and is, hence, given an account and password is an insider. If this definition is taken literally, however, users who are granted access to a subscription website are insiders, even if they have never worked for the organization that operates the website.

Regardless of difficulties in defining exactly what an insider is, very few information security professionals debate the meaning of the term “insider threat.” This term refers not only to malicious actions by employees and others considered to be insiders, but also to a much wider set of potentially undesirable risks, e.g., such as when a CD or USB drive is inserted without any knowledge of its contents, or when an email attachment from an unknown sender is opened, causing massive destruction of data.

The meaning of the term “insider attack” or “insider misuse” also varies considerably. Tugular and Spafford<sup>1</sup> have defined malicious insiders as individuals who are capable of using a computing system at an assigned privilege level, but who use the system in a way that bypasses or exceeds this level,

thereby violating their organization’s information security policy. Schultz and Shumway<sup>2</sup> define an insider attack as “the intentional misuse of computer systems by users who are authorized to access those systems and networks.” Einwechter<sup>3</sup> views an internal attacker as an authorized user who instead of meeting assigned responsibilities subverts system access to take advantage of it in some way.

## Predicting insider attacks

Ideally, organizations should be able to predict and then circumvent insider attacks, much in the same way that clandestine government agencies are often able to identify and thwart the actions of potential spies within their ranks by profiling them. Various predictive models of insider behavior, many of which are empirically-based, have been proposed. These models describe numerous kinds of insider attack-related motives, behaviors, and symptoms, many of which are subtle and thus easily overlooked<sup>4</sup> including:

- 1. Personality traits** – Several personality traits, namely introversion and depression, weakness in dealing with stress or conflict, and frustration with the workplace environment, are often found in individuals who commit insider attacks.
- 2. Verbal behavior** – Those who engage in malicious insider acts often tip off others that they intend to do so.

2 E.E. Schultz & R. Shumway, *Incident response: A strategic guide for system and network security breaches*. (Indianapolis: New Riders, 2001) p. 189.

3 N. Einwechter, “Preventing and detecting insider attacks using IDS,” on-line document (2002), <http://online.securityfocus.com/infocus/1558>.

4 E.E. Schultz, “A framework for understanding and predicting insider attacks,” *Computers and Security*, 21 (6) (2002), pp. 526-531.

1 T. Tugular & E.H. Spafford, “A framework for characterization of insider computer misuse,” unpublished paper, (Purdue University, 1997).

3. **Across system usage patterns** – Correlated usage patterns are patterns of computer usage (e.g., login attempts against accounts for which a user does not have authorization to access) across different computing systems. Although the actions might not seem significant on any one system, a pattern across numerous systems can reveal malicious purpose on the part of perpetrators.
4. **Negative work environments** – Negative/hostile work environments are also often associated with an increase in the likelihood of insider attacks occurring.
5. **Preparatory behavior** – Preparatory behavior (e.g., reconnaissance activity) can indicate that an insider attack is imminent.
6. **Meaningful errors** – Insiders typically make errors when they launch attacks, and some of these errors often indicate malicious intent before the insiders are successful in engaging in malicious actions. An attacker may, for example, mistype one or more letters of a malicious command such as `rm -rf / *.*` on a Unix host.<sup>5</sup> Identifying meaningful errors in near real time can enable security and other staff to quickly stave off impending attacks by taking defensive actions such as killing a perpetrator's remote sessions.
7. **Deliberate markers** – Insiders may leave small, intentional signs (e.g., hostile comments in one or more system configuration files) that they have attacked a system or device simply to make a "statement." Identifying such markers as early as possible can lead to the ability to thwart further attacks.

Additionally, motives for insider attacks can generally be characterized as hostility, the desire for revenge, and/or greed. In a study jointly conducted by CERT/CC and the U.S. Secret Service, 57 percent of the insiders who committed IT sabotage were disgruntled, 84 percent were motivated by revenge, and 92 percent of all of inside attackers attacked shortly after a negative work-related incident.<sup>6</sup> Insiders who have committed IT sabotage tend to be in technical positions (86 percent of insiders who committed attacks held technical positions) and also have had privileged system access (90 percent).

Given the growing body of knowledge produced by insider attack research, it is now more than ever possible to construct accurate profiles for individuals to allow prediction and special monitoring of insider attack-related behavior. Serious legal and privacy concerns have impeded using this kind of profiling, however. Clandestine government agencies are the major exception to the rule.

5 In case you are not sure what that does: `rm -rf / *.*` causes the file system in a Unix or Linux host to be recursively, starting with the root directory and working all the way down, deleted – a total wipeout of host.

6 D.M. Cappelli, A.G. Desai, A.P. Moore, T.J. Shimeall, E.A. Weaver, & B.J. Willke, "Management and education of the risk of insider threat (MERIT): Mitigating the risk of sabotage to employers' information, systems, or networks," (2007) [www.cert.org/archive/pdf/merit.pdf](http://www.cert.org/archive/pdf/merit.pdf).

---

**If organizations constructed profiles for their employees, consultants and contractors, insider attacks would be considerably easier to identify.**

---

## Detecting insider attacks

Detecting insider attacks is usually an elusive and difficult task until the loss, damage, and/or disruption from the attack becomes obvious. Insider attacks are launched by individuals who are generally authorized to access an organization's computing resources, so much of their behavior in the course of an attack will superficially appear to be normal. A simple access to a file might constitute a very serious insider attack, but intrusion detection systems (IDS) and intrusion prevention systems (IPS) generally do not issue alerts when this type of event occurs because the overwhelming preponderance of such alerts would constitute false alarms. As stated previously, if organizations constructed profiles for their employees, consultants and contractors, insider attacks would be considerably easier to identify, but most organizations hesitate to use such profiles.

The ability of IDS to identify externally-initiated attacks has gotten considerably better over the last five years or so, but the same is not true for internally-initiated attacks. Not only are many behaviors of internal attackers superficially innocuous, but externally-initiated attacks generally involve network traffic that goes through not only devices such as routers and switches that serve as intermediate network hops, but also devices (including IDS and firewalls) located within external gateways as well as numerous devices within an organization's internal network. This increases the likelihood that at least one device would detect externally-originated malicious traffic. In contrast, many internally-initiated attacks occur without any network traffic whatsoever being generated.

Despite limitations in the ability today's IDS to identify and report insider attacks, several reasonably effective intrusion detection methods exist.

## Tripwire tools

Tripwire tools (available as commercial as well as open source products) spot and report changes made in files, directories, and registry keys and values in association with an insider's efforts to install rootkits, back doors, or other types of malicious code. An example is a new entry in the `etc/init.d/` directory in a Debian Linux host that starts a new service that establishes a listener on an ephemeral port every time the compromised host boots. Tripwire tools are designed to detect and report this type of change as well as a wide range of others.

## Anomaly detection IDS

An anomaly-based IDS can look for and report deviations from known usage patterns (date/time, CPU and memory utilization, and so on) for each user. In one case, a successful insider attack against a Solaris host was detected and reported because an anomaly-based IDS identified behavior that did not fit the profile of a user who mostly used the host for email access and Web browsing. When someone logged into this user's account late at night and then started using the C compiler to compile a program, the IDS determined that this usage pattern deviated too far from the user's profile and thus sent an alert.

## Data extrusion detection

Data extrusion detection methods look for and report the presence of keywords such as "proprietary," "business initiative," and "engineering research" in file transfers, IM sessions, postings to social networking sites, and so on. Although many methods of evading detection based on this logic exist, data extrusion-based methods are capable of identifying insider attacks in which the motive is to steal critical and/or sensitive information. Similar detection methods allow critical files to be marked such that if anyone other than users (e.g., data owners) who are exempted from analysis downloads or attempts to download them, an alert is issued.

One of the best ways to detect insider attacks is to use security information and event management (SIEM) tools that aggregate log data and then perform event correlation analysis.<sup>7</sup> Audit data from a wide variety of sources – firewalls, network IDSs, servers, VPNs, and others – are conducive to pattern matching and identification needed to discover insider attacks. As opposed to detecting most external attacks, detecting internal attacks often depends on the discovery and analysis of numerous small, subtle clues, something for which event correlation algorithms excel. Consider the following examples:

1. The fact that a file on a particular host has been accessed may not in and of itself be very noteworthy. However, the fact that a copy of that file was soon afterwards sent as an email attachment would increase the level of suspicion in connection with the first action. If the address of the recipient is someone not associated with a particular organization, the level of suspicion would be ever greater. An effective SIEM tool can easily identify and report this series of events (file download => emailing file as an attachment => email is sent to external address). Conventional IDS are in contrast not designed to detect chains of events such as the one in this example.
2. An effective SIEM tool will readily identify and report other types of subtle indications of insider attacks such as failed login attempts on a variety of hosts within an organization's network that are well-spaced over time. Attempting

only one login on one host, then attempting another login on another a half hour later, and then attempting still another login on yet another host a half hour afterwards is a strategy that malicious insiders frequently use to evade detection capabilities. Conventional IDSs will spot only the individual login attempts, and given that failed logins occur all the time, the login failures in this hypothetical example would not in general be particularly noticeable. In contrast, effective SIEM tools are designed to detect and report the relationship between the distributed failed login attempts initiated from a single host.

## Responding to insider attacks

Because no risk that information security professionals face is as great as insider-related risk, every effort to contain an insider incident as soon as it is discovered must be expended. In this vein, identifying every malicious action that the perpetrator has performed up to the point of discovery is imperative in that different variations of insider attacks dictate particular courses of action. For example, containment of an attack in which malware has been installed in victim systems will be quite different from containment of an attack in which no malware was used. Additionally, different organizations have different business drivers and risk appetites, both of which very much affect the particular incident response strategy that each organization should use. Still, effective incident response for insider attacks will normally require following some or all of the following recommendations:<sup>8</sup>

- First and foremost, understand that you as an information security professional have not been empowered to carry on an investigation of one or more persons who work for your organization. Human relations (HR) and/or legal (but not information security) generally has the authority to conduct personnel-related investigations. In most cases, your job will be to serve HR and/or legal by collecting and then turning over to others information about the nature of the accesses the potential perpetrator(s) gained and by containing the incident until it is over; and then eradicating the cause of the incident and recovering any compromised systems and information.
- Always keep your management, as well as HR, legal, and possibly also public relations, advised of the status of the insider incident you are handling. Document (e.g., in a notebook, PDA, or through some other means) everything you find that is relevant to the incident as well as your conversations with others. Ensure that all evidence that you and others in your response team collect is gathered and handled in accordance with established forensics practices.
- Through inspecting audit log output and other means, identify every access path available to all personnel

7 A. Chuvakin, "Log analysis vs. insider attacks," *ISSA Journal*, November 2007, pp. 36 – 39.

8 Your organization should have high level provisions concerning handling insider attacks in its information security policy and should also have detailed procedures for responding to such incidents.

who may have been involved in perpetrating an insider attack. Be ready to promptly close each one if doing so becomes necessary. Meanwhile, if such personnel are still allowed to access computing systems, greatly increase the amount of audit data obtained concerning their activity and inspect this output promptly and regularly.

- Look for evidence of a much more extensive incident than what you and your team currently believe exists. Many insider incidents involve a greater number of highly coordinated perpetrators than was originally believed. Event correlation tools can serve a critical role in the effort to identify patterns that indicate complicity among multiple perpetrators.
- Perform a thorough, aggressive malware detection effort; do not quit until you are very sure that all malware has been found and removed from victim hosts (after, of course, you have made and then transferred custody of forensics copies of this code). Malware left in victim systems represents one of the most severe and lingering threats to your organization's systems and data.
- Consider using honey pot servers or a honey net to discover the nature of attacks that are occurring and who is perpetrating them. These technologies are by no means a first line of defense, but they can produce valuable information about the nature of attacks and can also deflect attacks that might otherwise be directed at critical servers.
- Once an internal attacker has been fired,<sup>9</sup> left your organization, or has been found guilty and disciplined, perform the following additional actions:
  - In the case of attackers who have been terminated or have left your organization, at first disable and in time (e.g., after 90 days) delete every account assigned to the perpetrator(s) as well as any other means of access that has been granted.
  - Change passwords of other accounts that have been used without authorization. If you are not sure whether a particular account has been used in this manner, it is better for security's sake to change the password anyway.
  - Examine the integrity of important information and binaries to which the perpetrator has had access and if unauthorized changes have been made, restore from a known, clean backup.
  - Continue to look for backdoors and other types of malware that the perpetrator may have left behind in victim systems.
- After the incident is over, conduct follow-up activity in which you and your team evaluate what did and did

not go well in responding to the incident, how many hours of effort were involved, and other critical aspects of the incident response effort. Incorporate the details of the incident, actions taken, recommendations for changes needed, and an estimate of the cost of responding to the incident in a report to be given to senior management within one or two weeks after the incident is over. Revise your team's incident response procedures as needed.

## Conclusion

To say that predicting, detecting, and responding to insider attacks are immense challenges is a gross understatement. Advances in research on insider attacks have increased understanding in these areas, enabling persons who predict, detect and respond to insider attacks in real-world settings to be more proficient. Despite all the new knowledge as well as a variety of administrative, procedural, and technical controls that can be deployed, however, defending against insider attacks remains as difficult a struggle as it was decades ago. The most trusted individuals in an organization can turn out to be the organization's worst enemies. Personnel screening performed not only when people apply for employment, but also regularly afterwards is one of the most effective preventative control measures. Quickly detecting and efficiently responding to insider attacks are necessary additional components in a successful strategy to mitigate insider attack-related risk to an acceptable level.

## About the Author

*Eugene Schultz, Ph.D., CISM, CISSP, is the CTO and CISO at High Tower Software, a company that develops security information and event management software. He has written/co-written five books and over 120 published papers and is an associate editor of three information security journals. He is a member of the ISSA Hall of Fame and can be reached at [eeschultz@sbcglobal.net](mailto:eeschultz@sbcglobal.net).*



<sup>9</sup> It is unwise to fire a suspected inside attacker if your organization wants cooperation from that person.