Realtime publishers

Monitoring, Detecting and Preventing Insider Fraud and Abuse

Dan Sullivan

sponsored by

Attachmate

Copyright Statement

Chapter 2: Technical Barriers to Monitoring, Detecting, and Investigating Insider Fraud

Businesses and other organizations have long used security controls to protect both physical and information assets. Buildings are protected with locked doors, surveillance monitors, and guards. Access to information assets is controlled with authentication and authorization systems, log monitoring, and vulnerability management. These are all well-developed methods for keeping out those who should not be in. They are not as useful when threats originate with those who have been granted access to the physical and information assets of an organization.

Insiders, such as employees, contractors, consultants, and business partners, are typically granted access to applications and data they need to do a particular job. They walk through the front door in the morning without much notice from guards, they move about the building with the wave of a badge before a magnetic card reader, and they work with enterprise applications throughout the day. A combination of trust in employees and other outsiders coupled with verification through monitoring and auditing can mitigate some risk of insider abuse but not all.

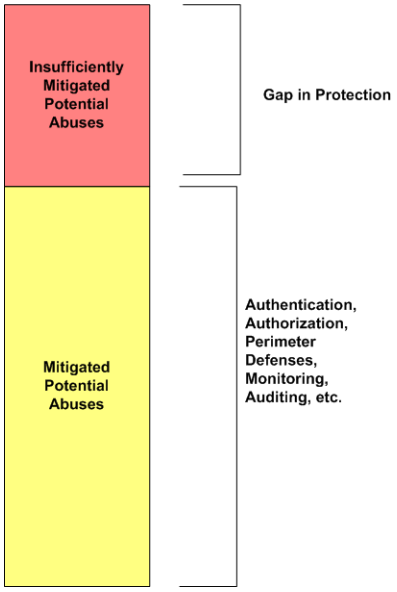

Figure 2.1: Commonly used security controls do not adequately mitigate all risks—particularly the risk of insider abuse.

This chapter examines the limitations of common security controls when it comes to controlling insider abuse, with special attention to the types of abuse that can occur when insiders have legitimate access to enterprise information systems. We examine the problem of insider abuse by discussing

- Special challenges with insider abuse
- Examples of insider abuse
- Five key challenges to detecting insider abuse

As we will see, there is a persistent challenge to preventing insider abuse because of insiders' access to applications coupled with detailed knowledge of internal operations.

Special Challenges with Insider Abuse

Security is about protecting assets. In many cases, this means preventing access to assets by those who would steal or otherwise damage those assets. Homeowners lock their doors when they leave their houses. Drivers lock their cars. We do not want strangers or outsiders to enter our homes or take our property; however, we might not think twice before lending our car to a friend or letting a friend stay in our house. Physical access controls such as door locks are designed to prevent those we do not trust from accessing our valuables. People we do trust, we give access to our things and may even share information about how we use those things.

Similar to the ways we protect personal property, we want to protect our business assets. We grant access to people we trust, such as employees, contractors, and business partners. Most of the time, trusting employees and other insiders is not a problem. People, for the most part, abide by a social contract and "play by the rules." Unfortunately, it does not take a large number of unscrupulous individuals to commit sizeable fraud, theft, or other damage.

When we think about security and asset protection with respect to insiders, we have to think in terms of probability and impact. If one in 1000 employees steals from their employer, and a company employs 10,000 individuals, 10 of those are likely to abuse the trust the company has placed in them. What statistics and probability cannot tell us is which 10 of the 10,000 will be threats. What the statistics do tell us is sobering with regard to impact. According to the Association of Certified Fraud Examiners, companies lose 5% of annual revenue to fraud, with each instance of fraud averaging $160,000—but nearly 25% of fraud cases are valued at more than $1,000,000. It might be hard to imagine how so many instances of fraud are possible until we realize the median duration of an instance of fraud is 18 months (Source: Association of Certified Fraud Examiners, 2010 Report to the Nations on Occupational Fraud and Abuse).

Pre-screening candidates can help eliminate some potential abusers, but we cannot depend entirely on that to avoid fraud and abuse. Screening procedures may miss someone who is intent on becoming an employee with the intention of committing fraud. This type of individual may invest considerable time and effort learning how to avoid detection by pre-employment screens. Existing employees may become potential abusers long after they are hired. Some may find it difficult to manage personal finance and turn to insider fraud to relieve a financial burden. Another employee may become dissatisfied with supervisors or have problems with coworkers that lead him to seek revenge for a perceived wrongdoing. Whatever the motivation, some employees may attempt to commit fraud and abuse. Our security control planning has to account for these threats as well as threats that emerge from outside the organization.

Insiders posses three characteristics that make them especially difficult to control:

- Legitimate access to resources
- Insider knowledge
- Potential ability to tamper with security controls

When a single individual possess all three characteristics, it is especially difficult to prevent fraud and abuse using only commonly deployed security controls designed to prevent theft and fraud by outsiders.

Legitimate Access to Resources

Employees and other insiders are routinely granted both logical and physical access to business assets. Such access creates unwanted opportunities to explore or even to commit fraud or tamper with such assets. This access also allows insiders to perform legitimate tasks.

Logical Access to Applications and Data Resources

Employees need applications and data to do their jobs. For example, a financial analyst needs current and historical financial transaction data. She may need access to reporting systems and repositories of consolidated data that combine financial and operational data in order to perform operational efficiency studies. Insiders often need access to multiple types of data, and therefore multiple sources of data. For example, the financial analyst mentioned might require access to data sources for:

- Account structure
- Accounts payable and receivables
- Payroll summary data
- Inventory levels over time

These data sources might span multiple applications with each application using different underlying structures. When the underlying data sources are frequently used together, it is often reasonable to build an application that provides services integrating or federating the various data sources (see Figure 2.2).

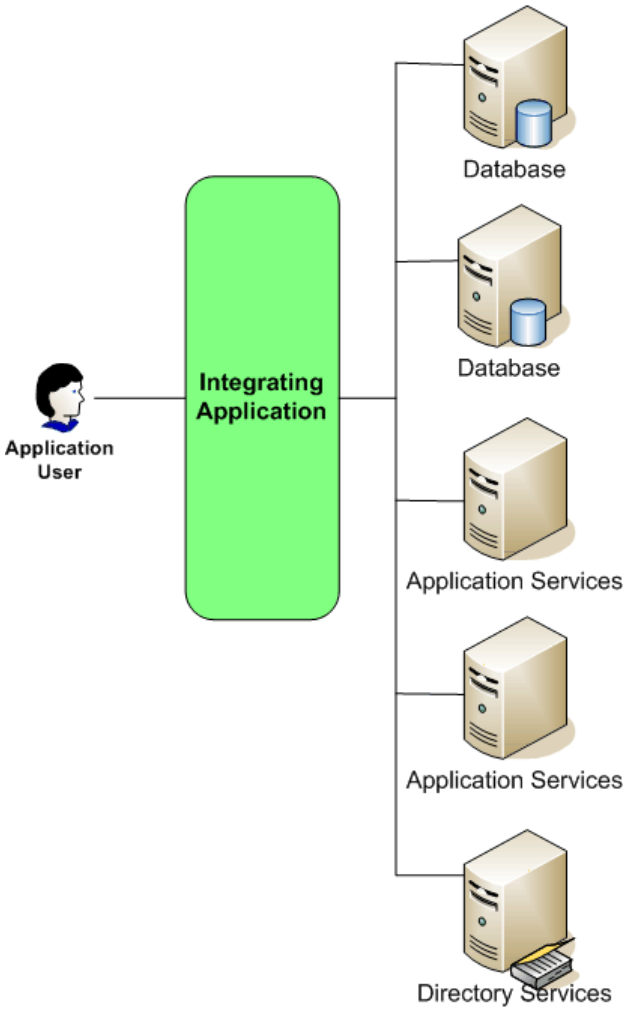

Figure 2.2: Integrating applications provides a single point of access to multiple application services and databases, which can be an effective means for committing fraud or abuse.

Common security controls, such as activity monitoring and audit logs are helpful in that they provide information about activities and events on a single component, such as a database or application server. Their scope, however, is limited to a single component and fraud and abuse often entails activities across multiple components. This type of single component monitoring, known as single channel monitoring, is inadequate to protect against fraud that occurs using multiple channels.

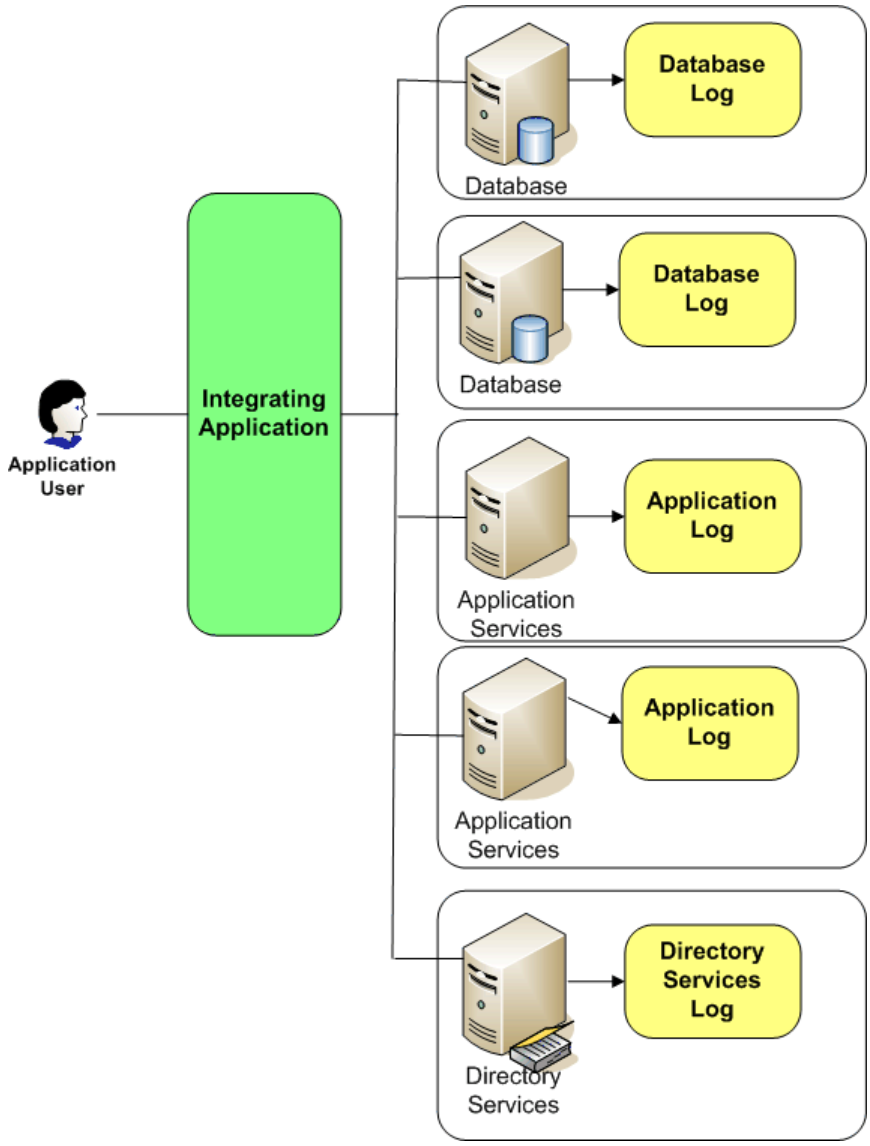

Figure 2.3: Integrated applications often use component-based logging and monitoring mechanisms, which do not give a comprehensive view of application-level activity.

Physical Access

In addition to access to multiple applications and data sources, employees and other trusted insiders have physical access to buildings and equipment. Employees often work late in order to stay on schedule or clear a backlog so that someone who is apparently working late might never raise suspicion. Combine logical access to applications and the ability to move about an office with relatively little supervision or observation, and you have an ideal setup to experiment with methods of using applications for abusive and fraudulent acts.

Insider Knowledge

Logical and physical access to assets is just one of the advantages insiders have when it comes to committing fraud and abuse; they also have detailed insider knowledge about business processes and workflows, potential collaborators, and potential ability to tamper with controls.

Insider Knowledge about Business Processes

Employees who work in a particular area long enough will accumulate knowledge about complex processes, their vulnerabilities, and potentially how to exploit them. Examples include:

- How to use supervisor overrides for some transactions that do not perform normal data consistency checks
- The fact that full monitoring is not enabled during maintenance windows when minimal services are available using replicated data on a standby server
- Thresholds used to determine when a transaction triggers automated checks or possible manual review
- Delays between the timestamp for events in one application and the timestamps that result from those events as they are posted to other related systems or applications
- Application details, such as the version of the database or the patch level of an application server, which can be used to find exploits to known vulnerabilities in those systems

As these examples show, insider knowledge is quite diverse and includes knowledge of application features, operational details, and exploitable characteristics of multi-system processes.

Operational knowledge about how applications and data are secured and audited is some of the most valuable insider knowledge. If someone understood the authentication and authorization rules for a system, that person could avoid triggering events that flag attempts at improper access. If the probability that a transaction is selected for auditing is directly proportional to the amount of the transaction, an insider could opt to commit fraud through a large number of small transactions rather than one or two sizeable transactions.

Colluders

It is important to remember that insiders can collude to commit fraud. Businesses and other organizations often separate duties to reduce the risk that a single person could abuse a business process. The person responsible for generating invoices is not the same person responsible for managing accounts payable. Insiders know other insiders. They may know others with particular privileges or access to application functions that are necessary to complete a fraudulent operation.

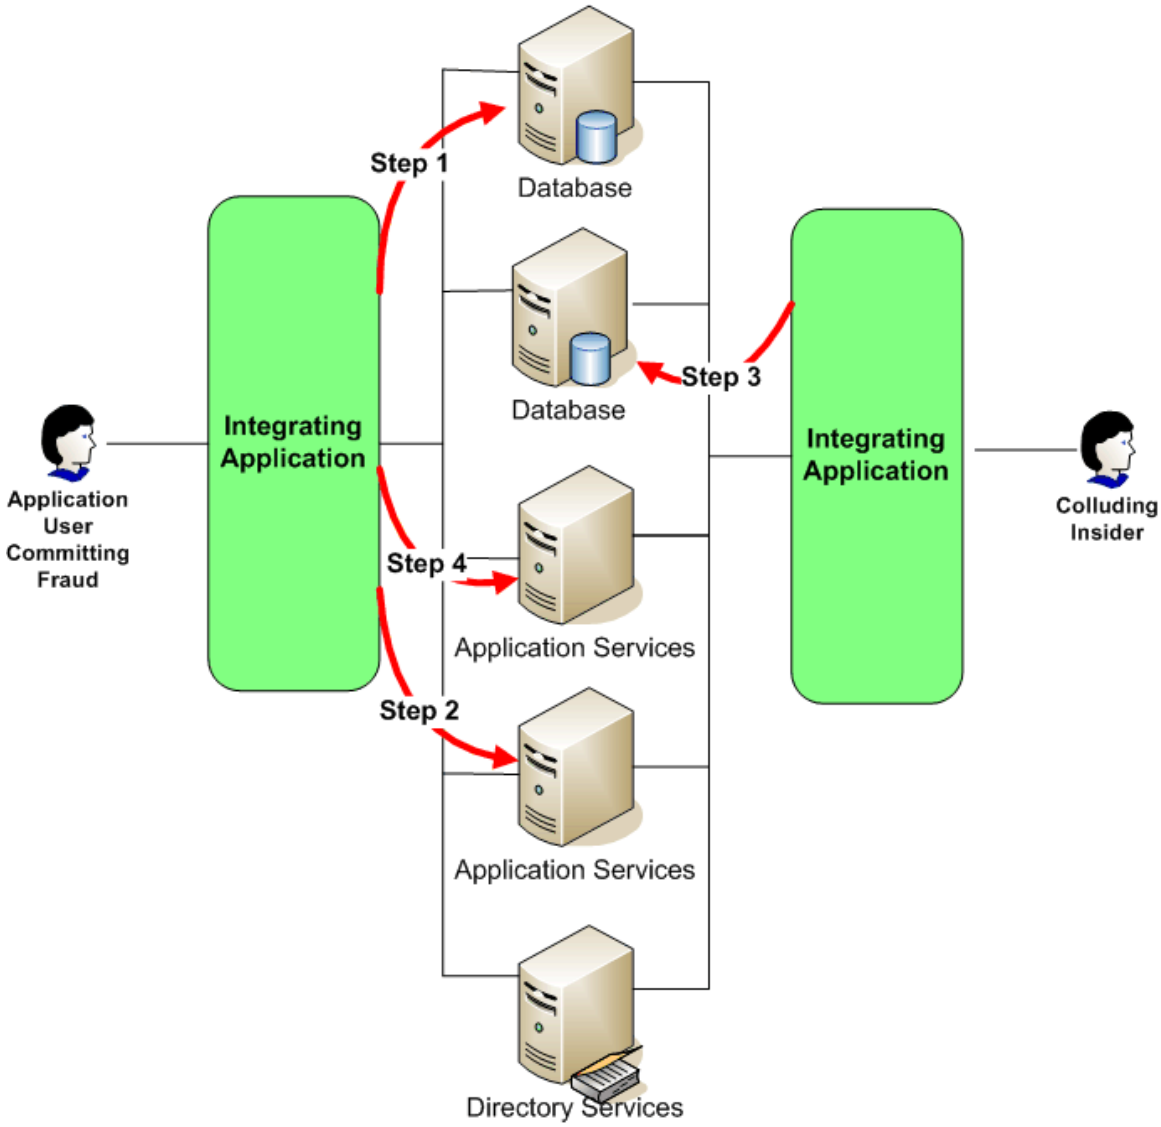

Figure 2.4: Insiders can work together to use enterprise applications to commit fraud and abuse that cannot be detected using conventional security controls.

Potential Tampering with Controls

Insiders can exploit improperly configured security controls to their advantage. At the most basic level, abusive insiders might find another employee who is careless about leaving passwords written down on sticky notes and attached to their workstation monitor or using easy-to-guess passwords. Poor password security plus a little social engineering on the part of the unscrupulous insider, and he will have access to applications and data via another employee's compromised account.

If an abusive insider were to gain access to an account with elevated privileges, the consequence could be more damaging. For example, an employee in the process of committing fraud who gains elevated privileges could:

- Alter application logs
- Change audit control records
- Modify access control privileges on other accounts
- Tamper with physical controls, such as changing physical access privileges to restricted parts of the facility

Insiders clearly have advantages over external attackers. Insiders are trusted and move about the physical facilities with greater access than non-insiders. Insiders have access to applications that can be used to commit fraud, whereas an outsider would first have to expend considerable effort to collect information about the types of applications in use, their configurations, access controls, and so on. Insiders get paid to learn and use that kind of knowledge. Furthermore, insiders understand internal business processes. Over time, they collect detailed knowledge about how applications work, including vulnerabilities that could be exploited for fraudulent purposes.

The special challenges businesses face in preventing insider abuse stem largely from their knowledge of and access to applications that can be used to commit fraud. Countermeasures to detect insider abuse and fraud must be designed to detect the type of activity an insider would perform in the course of an insider crime—not just the types of activities an outsider would attempt. Another way to understand the technical challenges of detecting and preventing insider abuse is to consider the following hypothetical example.

Example Scenario of Financial Theft

Businesses and other organizations face a number of types of threats from malicious insiders, including financial theft, intellectual property theft, and exposure of private and confidential information. Of these, financial theft is probably the most widely applicable, so it will be the focus in this section. The goal at this point is to highlight elements of financial fraud that are relevant to prevention and detection. This is not a step-by-step guide to perpetrating fraud using a particular enterprise application but a high-level outline of the issues a defrauder must consider and methods for dealing with those issues.

Means and Motive

We begin with an employee, Bob, who has become disgruntled with his job. His performance has been fair but there have been some incidents that required review. The employee is reportedly having family difficulties as well, including financial concerns. The employee has been with the company for 10 years and has access to several applications used for financial management, order fulfillment, and inventory. Over the years, the employee has worked in three different departments and so has accumulated knowledge of business processes in those areas. He also has colleagues in each area with whom he maintains friendly relations.

This description shows someone with both the means and the motive to commit fraud. The motive is driven by a combination of dissatisfaction with the employee's current work situation and by external factors, including financial and other stresses at home. This description probably fits a large number of employees who would never commit or even seriously consider committing fraud. If all we had to do was watch for employees who fit this profile to prevent fraud, we would not need the more sophisticated analytical tools that are required. The problem is that the pool of insiders that fit this or other relevant profiles is large, and the profile is insufficient to determine who will actually commit fraud. Profiling in this way produces too many false positives (that is, employees identified as committing fraud although they are not); it also does not identify everyone who commits fraud, leading to false negatives (that is, employees identified as not committing fraud but who actually do commit fraud).

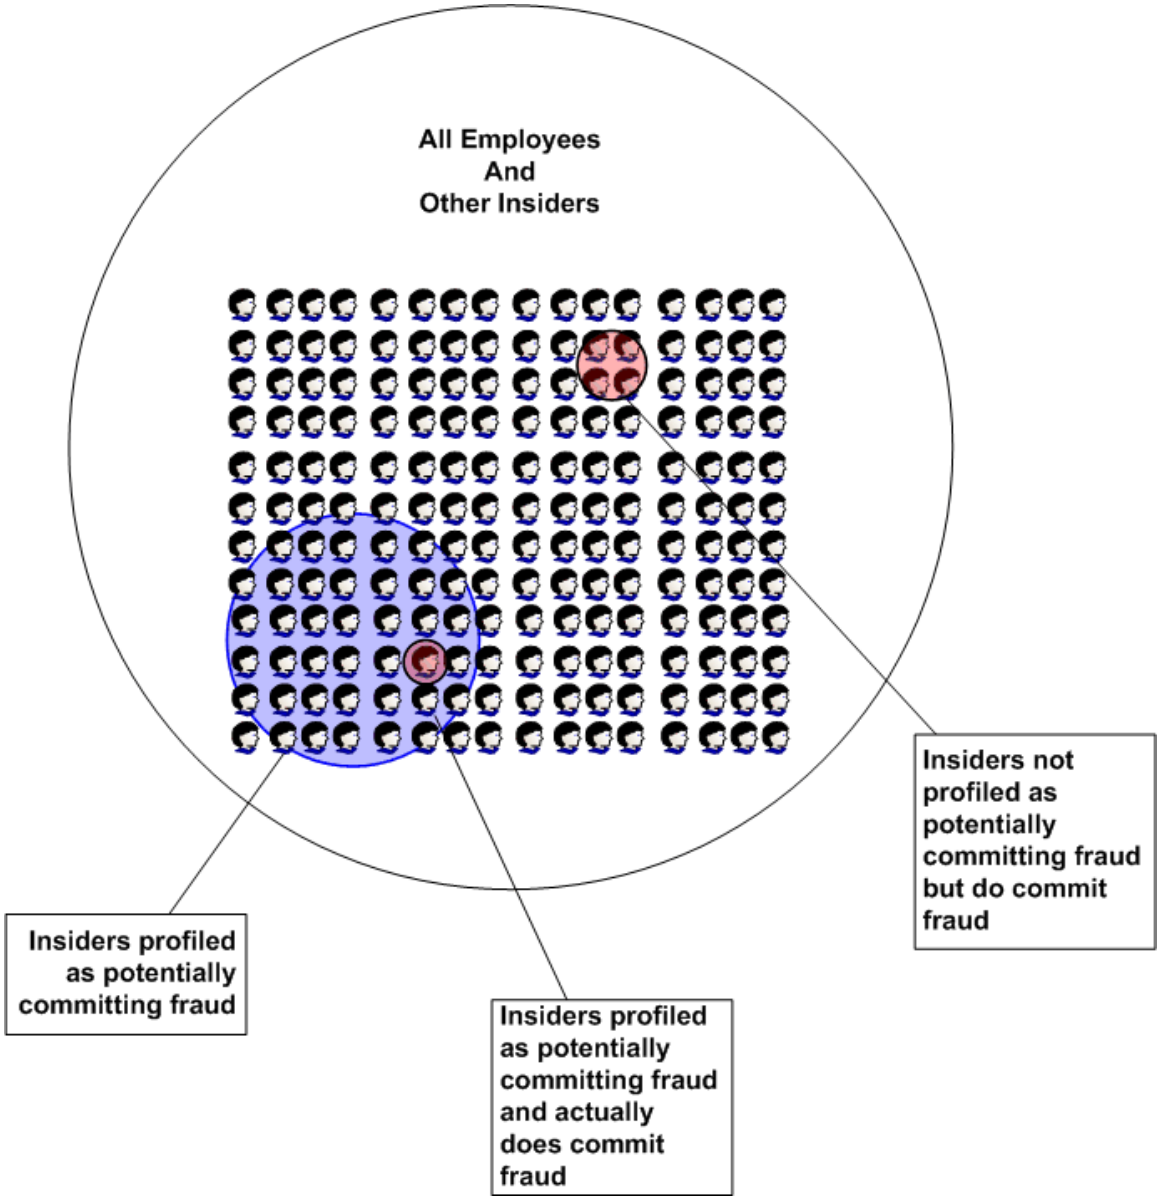

Figure 2.5: Profiling based on motive and means is an insufficient means to detect and prevent fraud.

Access to Applications

Bob, as noted earlier, has been with the company for a decade and knows his way around the business. In particular he has:

- Knowledge of account structures
- Knowledge of account receivables and payable applications
- Access to electronic funds transfer services

With knowledge of account structures, Bob can devise a scheme for creating a fake accounts payable record, which he will use to make payments to himself. Bob plans to set up a business checking account using a phony business name to receive funds, which he will withdraw at a later time. Of course, it would be easy to trace that account back to Bob, so he needs to ensure that no one would go looking too deeply into that account.

In Bob's mind, he has an advantage here: he worked in the finance department for 4 years and has a good understanding of the accounts payables system. He knows there are two processes for setting up payables accounts: the conventional method, which requires a number of due diligence steps to verify the vendor is a legitimate business and is providing goods or services to the company, and the "rapid pay" process, which is used for small vendors who bill infrequently and for small sums. The "rapid pay" process exposes the company to greater risk, but the cost of executing the conventional process for all small vendors is too costly to warrant it.

To further control the costs of dealing with a large number of small payments, small vendors, such as Bob's front company, are paid using electronic funds transfers. Bob is familiar with this process too from his days in the finance department, and he suspects there should be no problem with that payment method once he has established his fake payable account.

Slow, Methodical Observation

Bob feels he has a plan in place. His fake company has established a bank account. He has gone over the steps in his mind several times and feels like he has a foolproof plan, but Bob is patient and cautious and decides he need to observe the accounts payable procedures more closely with his plan in mind. In particular, Bob is concerned with:

- Detecting weaknesses in the application monitoring process
- Verifying he can actually execute the sequence of steps he has planned without triggering existing single-monitoring systems
- Testing his ability to make small fraudulent transactions without detection

Bob has made friends with several IT professionals that support the finance applications, and he uses these relationships to cull information about monitoring processes. Of course, he does this indirectly, casting his questions in terms of application performance, an interest in new reports to help him better manage his budget, and so on. Also, he never asks too many questions of one person in order to avoid suspicion. The IT professionals might disclose bits of useful information similar to victims of social engineering attacks—someone earns and builds their trust and then exploits it. The IT professionals are not colluders in any kind of conventional sense, but they do fill holes in Bob's knowledge of system functionality.

Attempts to Avoid Detection

Bob has worked out his plan and understands, he believes, what needs to be done to execute it. Part of the plan is making sure he avoids detection because even with the best of plans, unexpected events can occur that could lead another employee or auditor to suspect fraud. Bob needs to cover his tracks at several levels.

Covering Tracks with Documentation

For starters, he creates fake documentation about his bogus company's relationship to his employer. These include paperwork typically used for the kinds of business transactions Bob is supposedly carrying out: contracts, service orders, and invoices. Bob has access to the accounts payable system, so he is able to scan documents and add them to the bogus company's records. He does this in the unlikely case his account is selected by an auditor for further review. A more likely case is that someone running a quality control check to identify accounts that are not in compliance with policy will flag Bob's fake company account if proper documentation is not in place. Another possibility is that without proper initiation paperwork and subsequent transactions, the account might be flagged as inactive, which would prompt further review.

Creating a New Account

Bob does not want to draw attention to his malicious activities by using his own accounts, so he decides to create an account using the identity of a former employee. He can access application administrator accounts because he spent some time working with administrators during an upgrade of the application and learned of a shared administrator account. The password was easily guessed as is often the case with shared passwords. Bob decides to use a former employee's identity to deflect any suspicion should the account be discovered; with a bit of luck administrators would assume the account had not been properly deleted when the former employee left the company.

Keeping Transactions Small

Pay yourself $10 million and someone is sure to notice; pay yourself $50 dollars and chances are it will not draw much attention.

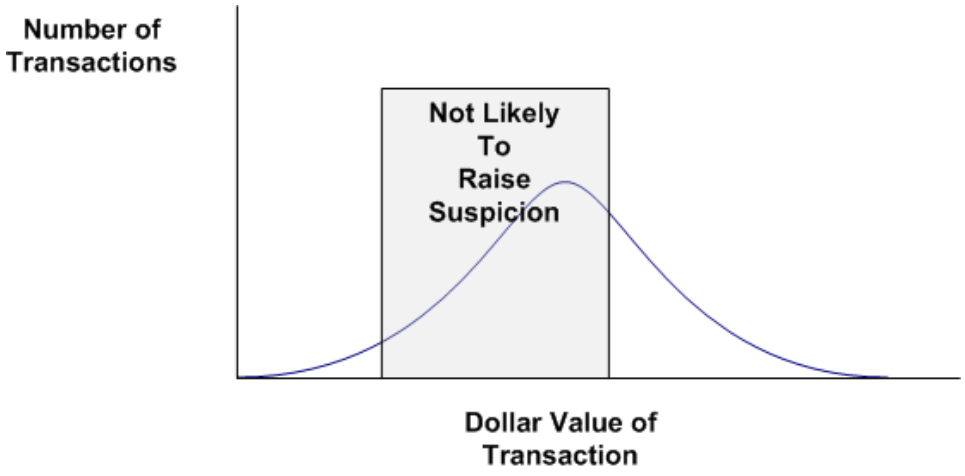

Figure 2.6: Transactions that look like most other transactions are not likely to stand out or raise suspicion based on the dollar value of the transaction.

Bob is familiar with the distribution of the dollar value of transactions in the accounts payable systems, and keeps his fraudulent payments in the range of the most common but toward the smaller amount of the range so as not to trigger amount-based controls.

Finally, Bob's philosophy is to know his enemy and what they are up to, so he occasionally logs in to the shared administrator account to see what types of single channel monitoring reports are running. One of the advantages of the shared administrator account, at least for Bob, is that his occasional use does not raise suspicion among the administrators. If they were to notice the last login time and it was not their last log in, they would likely assume it was one of the other administrators.

Putting Together the Fraudulent Pieces

In this scenario, an insider is able to devise and implement a scheme to defraud his employer. His insider knowledge, acquired over years of service in a number of departments, is invaluable. He understands business processes and has noticed weaknesses over the years. He has access to enterprise applications, which he uses to perpetrate his fraud. Since he has been an employee for so long, he has developed collegial relations with other employees, which he exploits to his advantage. Finally, he knows how to reduce the likelihood of being caught through the usual monitoring and auditing procedures. His countermeasures are not guaranteed to avoid detection, but they are sufficient to warrant the potential payoff from fraud.

All of this takes place in an organization with typical security controls designed to keep outsiders out. It also recognizes that even with the best intentions, poor security practices, such as shared accounts and easily guessed passwords, can occur.

In the first section of this chapter, we consider some of the special challenges with regards to preventing insider abuse, and we have just seen how those challenges can play out in the case of financial fraud. In the next section, we will summarize the key challenges to detecting insider abuse.

5 Key Challenges to Detecting Insider Abuse

At this point, it is probably clear that detecting insider abuse poses exceptional challenges, and common security practices are insufficient to detect such abuse. What is the solution? Controls to prevent insider abuse will have to address five key challenges to detection:

- Traditional access controls are insufficient to prevent potentially abusive access
- Insiders can collect data from multiple systems
- Insiders can perform malicious activities over an extended period of time
- Insiders can tamper with logs and other audit controls
- It is difficult to distinguish malicious from legitimate transactions

Any solution that attempts to addresses insider abuse will have to meet each of these challenges.

Insufficient Traditional Controls

Traditional security controls include:

- Perimeter defenses, such as firewalls
- Access controls, such as authentication and authorizations
- Encryption, such as disk encryption and virtual private networks (VPNs)
- Vulnerability scanning and patch management

Each of these controls assumes there are two sets of users: those who should have access to an application or data and those who should not. Once a user is deemed trustworthy, these controls are no longer relevant.

For example, an employee with a desktop workstation connected to an internal local area network (LAN) is unaffected by firewalls. Users who need an application to perform their jobs are given usernames and passwords (or other authentication mechanisms), so access controls can block functions unrelated to a user's job but they still have access to authorized functions. Encryption works well in preventing eavesdropping but is of little use when an employee has legitimate access to encryption/decryption keys. Vulnerability scanning and patch management help reduce the chance that an attacker can exploit a vulnerability in an application. Insiders already have access to enterprise applications, so exploiting bugs may actually be more work than using legitimate functions in fraudulent ways. Additional security controls are needed to detect and block insider fraud and abuse.

Insiders Can Collect Data from Multiple Systems

Application designers are well versed in creating systems that meet some set of requirements but no more. This reduces the business functions and data exposed through a single application, which is sometimes an advantage and sometimes a disadvantage. The fact that functions and data are limited means someone with access to the system can only do so much, and this promotes security. It is sometimes a disadvantage if applications become silos of functions and employees need access to multiple systems to perform a single business process. This is not uncommon: insiders have access to multiple systems with different functions.

From a monitoring perspective this means that monitoring a single application is not enough. We need to monitor multiple applications and look for patterns indicative of abuse that span multiple systems.

Insiders Can Perform Malicious Activities Over Extended Periods of Time

Insiders can use time to mask their activities. For example, an insider in the early stages of planning fraud might run reports or create fraudulent transactions and then wait to see if anyone notices. If the actions are detected, the insider gains knowledge about monitoring practices; if they are not detected, the insider is similarly rewarded with knowledge about monitoring, or lack thereof. In some cases, insiders can move even more quickly. For example, in a major case of fraud in the United Kingdom, a temporary employee in the social housing sector created a bogus company and submitted invoices for more than £2 million in merely 3 weeks (See "The Internal Betrayal: A CIFAS Report on Beating the Growing Threat of Staff Fraud," August 2010).

As a general principle, if a malicious insider were committing a series of steps and extended the time over which those steps executed so that they fall outside the usual monitoring window, then the fraudulent activity may not be detected. For example, if weekly reports analyze the past 2 weeks for suspicious activity, a perpetrator would need only to space out activities over 3 or more weeks to reduce the chance of detection.

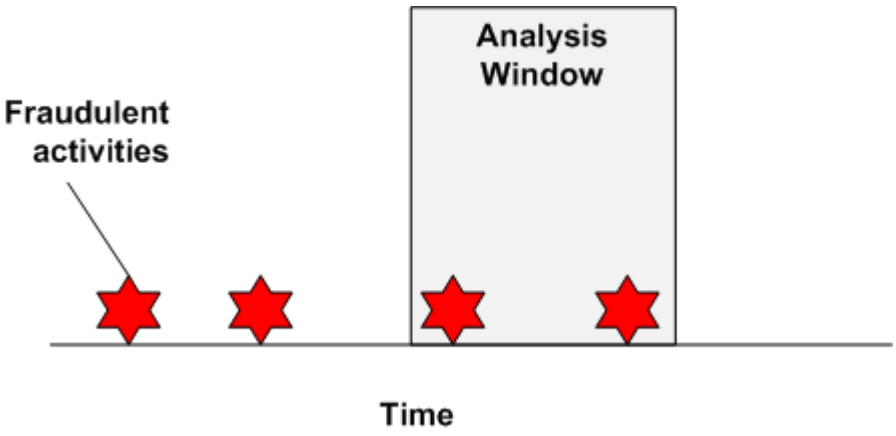

Figure 2.7: Fraudulent activities can be spread over extended periods of time, making it more difficult to detect sequences of events that are indicative of fraud.

Insiders Can Tamper with Logs and Other Audit Controls

Insiders might gain access to privileged accounts, either through malicious means, such as those described in the previous section, or because they have been granted elevated privileges in order to do their jobs. One of the challenges in protecting applications and data is that administrators are effectively granted the "keys to the kingdom." Although some key infrastructure providers, such as relational database vendors, address this situation with restrictions on privileged users, we will always have the case where some users are allowed to do more than others. With that comes the risk that privileged users will employ their privileges to either commit fraudulent activity directly and/or cover it up after the fact.

Difficult to Distinguish Malicious from Legitimate Transactions

Fraudulent transactions do not carry markers identifying themselves as illegitimate. Insiders can use their knowledge of the range and frequency of transaction amounts and types to design transactions that blend in with legitimate transactions. In systems with a large number of transactions, it is especially difficult to find small numbers of fraudulent transactions unless we have more information than what is contained in a transaction. For example, names and amounts may not indicate fraud, but the way a transaction was entered, the other events that preceded and followed the transaction, and other information that provide a context for the transaction can provide valuable indicators of potential fraud and abuse. Furthermore, baseline measures of the number and types of transactions performed by others in the same department or with the same role in the organization can be used to identify unusual activity. A teller that performs two to three times the average number of a particular type of transaction warrants some investigation because this is an indicator of potential fraud.

Summary

Detecting insider abuse is challenging. Insiders have detailed knowledge about business processes as well as legitimate access to applications that can be used to perpetrate fraud. Insiders can leverage their knowledge about weaknesses in security practices and monitoring procedures. Conventional security controls, such as perimeter controls, access controls, and encryption are not sufficient to address these challenges. Fortunately, techniques exist for monitoring application activity in ways that can detect anomalous and suspicious activity. Those will be the topic of the next chapter.