Understanding the insider threat & how to mitigate it

Prepared by:
Katy Winterborn, senior security consultant

Table of contents

1. Introduction

This paper is intended to give a high-level view on the insider threat for those looking to implement a defensive programme. It considers the types of attack that may take place and some of the common weaknesses that aid insider attacks. It also covers some of the policies and controls that can be implemented to detect, deter or defend against the insider threat. This paper is intended to be a summary, however, the final section details further reading and resources that provide more in-depth information.

It is a widely held belief that the vast majority of threats to businesses are from outside attackers, with the stereotypical view of hackers trying to make money through crime. Additionally, while working on-site during penetration test engagements it is often perceived that issues found during build reviews and internal infrastructure security tests are not a real cause for concern. This is because of mitigations in place that prevent external attackers from gaining access to internal resources.

The problem with this viewpoint is that it does not consider the threat from a malicious insider. There is a tendency to trust staff once they have been hired and have passed any policy-based background checks or vetting processes, and to only consider external attack vectors. As described in a report by Computer Economics [1], fewer than 40 per cent of those surveyed viewed any type of malicious insider incident as a major threat. However, research by IBM has shown that up to 60 per cent of attacks against businesses come from those with legitimate access. In addition, three quarters of those were by individuals with malicious intent [2], whether from employees, contractors or service staff in the form of an ‘evil maid’-style attack [3]. Very few employees join an organisation with malicious intent, however, circumstances either within the workplace or outside can adversely change an employee’s behavior or viewpoint.

This is a large problem. Insiders are likely to have extensive knowledge on how systems and businesses work, allowing them to be more targeted in their activities and have more time to carry out their attacks as they have a legitimate reason for accessing the property or systems. In addition, they do not need to devote time to bypassing external security controls, which might otherwise slow down or deter external attackers.

While the main issue raised in discussions around the insider threat is loss of confidentiality, often through theft or accessing restricted systems, it is also the case that the integrity of systems can be compromised. This could be through fraudulent transactions or data manipulation as well as the availability of systems, through sabotage or other deliberate damage.

The remainder of this paper will consider the three common types of attack carried out by malicious insiders, as well as common weaknesses and areas that are often targeted and some mitigation strategies for those weaknesses.

2. Types of insider attack

There is no one definition for an attack by an insider. They are carried out in many different ways for a large variety of reasons, affecting both IT and physical systems. This makes categorising insider attacks difficult.

For the purposes of this paper the definition of an insider will be taken as an individual who has legitimate access to systems or premises and uses this privileged position from which to launch an attack. Note that this definition is not restricted to purely physical access as remote workers with access to corporate resources, via VPN for example, can also be considered insiders. In addition, the term insider does not just refer to employees, but to anyone with legitimate access. This could be, but is not limited to, contractors, suppliers, consultants etc.

Although categorisation is a difficult task, attacks can broadly be broken down into three types: sabotage, theft and fraud. The attacks were first described by CERT in its wide-ranging study on insider threat [4]. An interesting aspect of this research was that attackers are not as technically competent as may be assumed. A large number of attacks were found to be via simple vectors that did not require a large amount of skill or specialist knowledge. An early study into the insider threat in the banking and finance sector [5] highlighted this explicitly, however, it was a common theme throughout the majority of investigations studied.

Note also that while potential motivations for attacks are documented in this paper, the research carried out by CERT showed no pattern in those that could be considered malicious insiders. They covered every demographic, level of skill and position in the management chain [6].

There is no one definition for an attack by an insider, they are carried out in many different ways for a large variety of reasons

2.1 Sabotage

Sabotage can be considered damage to systems or physical property and is often carried out by disgruntled employees in revenge attacks, for example after a dismissal or when they feel they have been unfairly treated.

Sabotage attacks are not necessarily conducted while the employee is still working for company, instead they may be set up so that they retain access after leaving or have set a timer to trigger an event.

One of the most well-known examples of sabotage is a 'logic bomb' [7]. This is where an attacker will leave a programme hidden on the network to be triggered by a certain event; this may be a given date or specific activity on the system. Once this event occurs, the programme will execute, often causing massive amounts of damage, for example the deletion or corruption of client data from a database. One example of a logic bomb causing widespread damage was seen in South Korea in 2013, which caused the widespread destruction of hard drives at banking and broadcasting companies [8].

Sabotage may also include physical damage. As was the case with an employee who spent three years damaging systems by spraying them with 'Cillit Bang' after being denied a pay rise, causing widespread damage and system failures [9].

2.2 Theft

Theft is often the first thing that comes to mind when considering risks to an enterprise and this is of particular concern when a company is dealing with personally identifiable information (PII). Theft can target intellectual property, such as source code, client lists and proprietary research or may target more tangible items such as company property or even money. A Washington Post article suggested that up to 60 per cent of departing employees may be stealing company data [10].

The motivation behind theft is not always straightforward but often stems from a desire for financial gain. This may be because the attacker believes they are owed something, are struggling financially or even trying to gain work with a competitor and want to take something of value with them.

In other cases theft may be an act of retaliation, similar in motivation to the sabotage examples described above.

The most famous example of employee theft in recent times is the case of Edward Snowden. Snowden was a US government contractor who stole classified information and leaked it to the media [11]. While it appears that, in this case, his actions were for ideological reasons, cases such as this could conceivably be due to disgruntlement or the desire for monetary gain.

It is worth noting that theft may involve third parties rather than being for personal advantage, for example aiding a friend or family member or in cases of blackmail. This is also true for fraud and, to a lesser extent, sabotage, however, it is most commonly associated with theft.

2.3 Fraud

Fraud often involves circumventing controls or business processes. In a more technical attack, this could involve changing source code in order to benefit an attacker in some way, most often for financial gain.

Examples of fraudulent activity may include managers signing-off their own transactions in order to hide unauthorised transactions, exploiting logic errors in software to bypass controls or the infamous 'salami attack' in which the difference from rounding errors can be rerouted to a controlled account without detection [12].

The motivation behind fraud is very similar to theft in that it is often for personal gain. One additional area that may lead to activity of this kind is to cover up mistakes or failures. As was the case of an investor who manipulated transactions so he would appear to be the star salesman, when in fact he was losing the company vast sums of money [13].

"One additional area that may lead to activity of this kind is to cover up mistakes or failures. As was the case of an investor who manipulated transactions so he would appear to be the star salesman, when in fact he was losing the company vast sums of money."

3. Risks

As mentioned previously, there is no obvious pattern which is guaranteed to indicate insider activity or the type of person most likely to commit an offense. This section will attempt to highlight some of the risk areas, however, this should not be taken as definitive proof of wrongdoing.

This section has been organised into three areas: personnel, technical and business. However, there is some overlap between the three sections.

Insider cases are complex and individual indicators cannot be used to detect attacks in isolation. Instead, a holistic approach must be taken, reviewing all areas, in order to mitigate risks.

3.1 Personnel

Personnel issues are wide-ranging and complex and care should be taken to avoid accusations or over-monitoring of individuals with no evidence as this may contravene their human rights and result in legal action. An important area to consider to promote a good security culture is the attitudes of those in senior positions and the language chosen to address issues of concern. This may include, for example, not calling people suspects or perpetrators but using softer language such as ‘behaviour of concern’.

A report released by CPNI [14] highlighted a number of areas that may indicate a reason for concern, including personality traits, lifestyle issues and workplace behaviours. It should be noted that although these areas were only of interest if they were frequent or severe in nature, one point of interest from the studies CERT conducted into insider threat was that offences were often preceded by a deterioration in general behavior and a number of minor offences. This included worsening relationships with colleagues, minor reprimands from management or a seeming lack of engagement [17].

Financial trouble has been seen in a number of cases of malicious insider threats, particularly in cases of fraud and theft. This is of particular concern if the individual is attempting to keep their issues hidden, maintain an extravagant lifestyle that they cannot afford or should they have issues with drink, drugs or gambling.

Other personal circumstances, such as family problems and health issues, may cause financial issues and may lead to desperate measures on the part of the individual.

Another factor often seen is a lack of job satisfaction or feeling unfairly treated. This may stem from being passed over for promotion, receiving insufficient pay, experiencing

office favoritism or even feeling a sense of entitlement. In this case, the offender may believe the company owes them something and they are not doing anything wrong, merely taking what is rightfully theirs. An example of this was seen in the case of an employee passed over for promotion who deleted files from a number of machines the day before a new finance director started in the company [15].

The most extreme cases of dissatisfaction often occur when roles are terminated, as this can leave the individual in difficult circumstances with nothing to lose. Employees often have a notice period of at least one month and if they retain the same rights and access to the system there is a potential risk that they can use this time to commit an offense.

Bad or absent management is often another factor in malicious insider cases. However, this is more of a risk in the case of remote workers, who often have little to no contact with people in the company and can be neglected. A feeling of solitude and lack of engagement can lead to the impression (and often the reality) that no one is watching and offences will go undetected [16].

A final area of consideration is those who commit offences just because they can. This could simply be for the challenge of it or because they spot an area of weakness in a system. This is better known in traditional, external hacking attacks, however, it should still be considered an internal threat and it should be noted that insiders are often better acquainted with systems and in a better position to spot potential weaknesses.

It is also of note that loyalty, satisfaction and personal circumstances change over time. An individual may have no intention and no motivation to commit an offence but this does not give any guarantees for future behavior or situations. A study by CPNI found that “over 75 per cent of insider acts were carried out by staff who had no malicious intent when joining the organisation but whose loyalties changed after recruitment.” [18]

It is important to note that one area cannot be addressed in isolation, in order to effectively mitigate the insider threat

3.2 Technical

While technical issues do not directly cause insider attacks, they can often make them easier, opening up avenues of attack for non-skilled individuals. If technical systems were secured it would make attacks substantially more difficult.

One of the most prevalent security issues relates to password management, both in complexity and password sharing. If passwords can be easily guessed or are shared among multiple users the system loses confidentiality, integrity and nonrepudiation. This means that attackers can not only access and change information they are not supposed to but also make it appear as if someone else (the logged on user) was responsible. This is of particular concern in the case of administrative passwords or those used to elevate privileges.

Alongside password sharing, a robust, role-based access control system should be implemented, ensuring that users only have access to systems and resources they need, denying and logging requests to other resources.

Unmitigated vulnerabilities in the system are another area to consider. If patches are not applied or other vulnerabilities exist, but are unknown or ignored, a skilled attacker could exploit these to gain access to a system they are not supposed to or gain elevated privileges.

Sometimes, while you may not be able to prevent an attack, it is possible to detect and mitigate one swiftly. However, in order for this to be possible an effective logging and monitoring system needs to be in place and adequately tuned. It is often the case that monitoring is missing altogether or has been left at the default settings, which can lead to a high false positive or false negative rate.

Monitoring is also not effective unless someone is actually checking logs and alerts. If it does exist within the organisation, the person responsible for this role often has to do it in addition to other responsibilities, leaving little time for triage.

In addition, if someone is in a development role, a significant risk is that they may keep the source code on his or her machine and it may be the only copy. This is a particular risk in small teams. If this is the case and the code is lost, it could set projects back or cause them to fail altogether.

3.3 Business

As with the technical risks, business issues do not directly cause attacks. However, having the appropriate controls in place makes it substantially harder to commit an offence and ensures that the correct procedures are in place in order to deal with such issues.

Lack of oversight is one potential area for concern and this can take two forms. In the first, the individual does not have adequate supervision and is free to perform malicious activities unchecked. In the second, the malicious insider is in a position of

responsibility, meaning they can ask a subordinate to perform an action on their behalf and then act as the authoriser.

In some cases, particularly if an individual has been in a role for a long time, they can become a single point of failure. In this case, they alone have access and rights to certain systems, or manage entire code bases themselves. Should they choose to damage systems or change all the passwords, for example, this could cause a significant amount of distruption. Failure to ensure that there are second copies of all business-critical data, including code bases, leaves an organisation at risk of attack. If data is securely backed up, even if the local copy is deleted or corrupted, they can be recovered.

Many companies do not have a secure termination or user deprovisioning policy. When a contract is terminated, retaining access beyond what is necessary leaves systems and data at risk of theft or sabotage. Alongside this, there have been examples where an ex-employee's access was not revoked at all, meaning they could access systems and premises after their employment had ended [19]. Identity and access management solutions (IAM) can be critical to helping mitigate against the insider threat but, as discussed in an article on the subject by Forbes [20], identity and access management solutions (IAM) can be critical to helping mitigate against the insider threat but are often lacking in organisations.

Privilege creep is a term for a user retaining access to systems and data from a previous role when it should have been removed and gradually accruing more and more access rights as part of a new position. The more access a user has, the more opportunity they have to commit an offence. An example of this can be seen in a case study given by CERT [19] where an employee retained access to payroll data after switching roles within the company. The employee provided confidential data to an associate of theirs starting up a new firm and this then cost the employer over $1 million in damages.

Senior support can be one of the most critical mitigations for insider threat, but it is often overlooked. In businesses that fail to support security at board level can lack the necessary controls to prevent insider attacks [21]. Studies have shown that malicious insiders often display signs that have been picked up by colleagues, however, if there is no safe mechanism to report these concerns to management, they often go undisclosed. Care must be taken to avoid false accusations in these situations, so any whistleblowing strategy must be carefully considered [22].

The CERT study into insider threats showed that those committing offences were unaware of, or did not consider, the consequences of their actions. In other cases, they were not aware of the monitoring solutions or other controls that would detect their activity.

4. Mitigation

In this section, potential mitigation options will be considered, however, this list is not exhaustive and only gives a flavour of the areas to address.

When considering any security strategy, no single control will be sufficient to prevent all threats. A good approach is to adopt a defence-in-depth strategy and implement a variety of measures. This ensures that if one mitigation fails, the attacker does not gain access to everything.

However, it is possible to go overboard with a security policy and so care must be taken to ensure that the measures put in place do not create a culture of mistrust and suspicion or prevent people from being able to do their work effectively. This could then have the opposite effect to that intended.

4.1 Personnel

One of the most efficient personnel controls is staff vetting. The level of vetting will depend on the risk and security assurance requirements of data to be protected. This will highlight any areas of concern with existing or prospective employees if done at the right level. Vetting should not be considered as one-off activity.

Circumstances can change, so an appropriate level of continued vetting throughout employment may be necessary. As an example for government-cleared staff at the highest level, full vetting takes place every five years and for a lower level of clearance it is every ten years. This gives an indication of how vetting is carried out for those handling classified material, however, each organisation will need to decide on the appropriate level of vetting as part of a risk assessment.

As noted previously, individuals who become malicious insiders often display a change in general behaviour and may be reprimanded for minor offences. In many cases this is noted by colleagues or happens in separate instances. A secure procedure for reporting concerning behaviour can, if handled correctly, act as an early warning system for potential attacks. Care must be taken when implementing any kind of whistleblowing scheme, such that both the employee reporting the behavior and the employee being reported are protected from unjustified backlash. Any process implemented must be transparent and detailed in company policy to prevent any accusations of constructive dismissal or bullying.

Management engagement from both direct management and those more senior can provide a framework to help protect against malicious insiders. If an employee feels they have been unfairly treated, having a more senior member of staff that you are able

to raise any concerns with may help reduce dissatisfaction. This is particularly important in the case of remote workers who may be isolated and have little direct contact with others in the company. Regular meetings, whether in person or on the phone, can provide a sense of belonging in the team or business. In addition, areas of concern or specific incidents should be dealt with by a line manager, preferably via a phone call to provide personal contact, rather than an impersonal email or even a message from IT departments.

4.2 Technical

As discussion in the previous section, technical issues do not, in themselves, cause the insider threat, but weak technical controls can enable attacks to succeed or exacerbate the consequences.

A robust and enforceable password policy is necessary to adequately protect an environment. Passwords should be sufficiently strong and follow recommended guidelines. In addition, password sharing should be discouraged and each user who requires access to a system should have their own logon (this has the added advantage of enabling accountability) and no user should be permitted to reveal their logon details.

Looking beyond passwords may help to defend against password sharing and misuse issues. Two-factor authentication (2FA) is one method. Users must supply a password as well as another form of identifier such as a randomised token, for example RSA tokens or a code that is supplied to a device and in the user's possession, such as via text or email.

Biometrics is another area that is becoming more prevalent, although there are a number of issues depending on the biometric mode in use and its underlying technology. It may also be harder to commit fraudulent actions with biometrics as they can be difficult to spoof. CPNI has included recommendations for biometric systems in its good practice guide to ongoing personnel security [24].

"A good approach is to adopt a defence-in-depth strategy and implement a variety of measures, so if one mitigation fails the attacker does not gain access to everything."

A secure environment should be applied to all machines via group policy (if available) and the build of all servers and workstations should be hardened to the appropriate level. This makes circumventing controls or accessing forbidden data much less likely.

All systems in the network should undergo regular security tests and issues raised should be considered from the point of view of a malicious insider. Much of the effort in defending systems and networks goes on the outer boundary, assuming attackers will be external to the business. As mentioned in the introduction this is a false assumption - results from internal infrastructure assessments, build reviews and other internal checks should be given an appropriate priority.

In order to protect all important data, backups should take place at regular intervals. These should be to a secure, off-site location that doesn't allow access without prior authorisation. In addition to this, for projects involving source code, the use of a version control repository, such as GitHub [25] or Apache subversion (more commonly known as SVN [26]) should be mandated. This means that regular snapshots of the code will be taken, making it more difficult for an insider to destroy the project they are working on. If these solutions are implemented from the beginning of a project and introduced to new team members as mandatory ways of working it provides an easier means for handling users who do not comply with the policy at a later date.

Implementing a logging and/or monitoring system will give visibility of what is going on in the network as well as on individual systems and will provide valuable evidence in the case of an actual attack. For a monitoring solution to be effective an analyst must be regularly checking alerts and output, otherwise signs of an attack may be missed. To detect anomalous behaviour there should be a baseline for what constitutes 'normal' on any given system or network so activity deviating from this can be investigated. From a user perspective, pop-ups that warn users monitoring is in place when they attempt to access a restricted area are often an effective deterrent and can be implemented as part of a monitoring solution.

For companies developing software internally, it is recommended that they undergo a thorough code review prior to deployment. If done by someone outside of the development team who is skilled in secure coding, it will remove security vulnerabilities that may otherwise go undetected. The other advantage is that it ensures nobody on the development team can insert a back door or deliberate defect in the code that they can exploit.

Data Loss Prevention (DLP) solutions are another safeguard that may give an indication of malicious activity inside a company. This can work by fingerprinting data and analysing it as it passes the boundary of a company, for example leaving the company network. Solutions may also analyse traffic volumes for indications of mass data exfiltration or prevent large data transfers via removable media.

4.3 Business

As with technical controls, business controls are implemented to discourage attackers or make malicious activities seem like a less appealing option. The difference with business controls is that technical controls are likely to be as effective at defending against both internal and external threats. Business strategies are much more focused on those already in the environment and are vital for countering the insider threat.

Separation of privileges/duty should be enforced as much as possible, meaning that no one person should be responsible for both taking an action and authorising an action. In addition to this, the two-person rule should be implemented where appropriate. This means that two separate users will need to enter distinct credentials in order for certain, business critical functions to take place. These actions give a measure of oversight and ensure that at least one extra person is involved in the process. This does not protect against collusion or coercion but gives a measure of security against a single attacker and acts as a deterrent for those considering carrying out an offense.

It is important to ensure that there is no single point of failure anywhere within the organisation and that no one employee is allowed to amass a critical amount of data that they alone are responsible for. This can relate to things like source code, database access, process knowledge or any other business critical functionality. In extreme cases, this may mean putting policies in place to prevent employees working in certain combinations of areas of the business throughout a career, in order to prevent the possibility of amassing the required knowledge and experience to carry out an attack.

The principle of least privilege should be applied to all users - this means that every employee should only be allowed the level of access they require to perform their role and nothing more. As administrative rights give the greatest amount of flexibility on any system special care should be taken to keep the list of users with this role to a minimum. However, it also applies to data access and system access, for example databases containing PII and any areas containing sensitive information an employee does not need to do their job.

Ensure that there is no single point of failure anywhere within the organisation

Policies should be in place for removing access rights when a user changes role and access that is no longer required should be removed as soon as is possible. This is of particular importance when an employee is leaving the company and may be working a notice period. It is important to consider the risk of allowing this individual continued access versus what they need to continue working. There are a number of good practice guides available and CPNI provides a large amount of advice, however, a quick overview can be found easily online [27].

It is important that policies are meaningful and not creating the wrong culture. For example, if the password policy says that you should never share your log-in details but the IT department routinely asks for log-in credentials to provide technical support, there is a risk of teaching employees that sometimes it is ok for them to ignore policies. This then increases the chances of them doing it when they are being socially engineered.

An important measure in combatting the insider threat is the joining up of different functional areas within a business. For example, if the physical security team become aware of someone coming into the office outside of normal business hours, the HR team have had reports of the person making threats about stealing the customer database and the IT team have seen several attempts at downloading large amounts of data then this information should be combined and highlighted as a credible risk, rather than it being discovered in hindsight.

Senior support and sponsorship is vital for making any security policy viable. A board level member of staff should be responsible for security, taking into account both internal and external attackers. As part of this responsibility, planning and auditing activities should take place which consider all systems and processes, where the threats are and how to mitigate them. Doing so leads to a much more robust and implementable security strategy.

This should also extend to creating a good security culture. Any policies put in place should be seen to apply to all staff members at all levels of seniority. For example, if there is a pass wearing policy, the CEO must wear their pass or risk undermining the policy. This issue and recommendations for implementation are discussed in detail in the CPNI guide to the Holistic Management of Employee Risk (HoMER) [28].

As noted by CERT, a large proportion of malicious insiders did not realise the consequences of what they were doing and were unaware of any monitoring [23]. Running a user awareness campaign may make potential attackers think twice about what actions they are taking. However, care must be taken to ensure technical details are not revealed as this may allow an attacker to bypass controls. In addition a transparent security policy makes staff aware of investigatory procedures, following this policy allows any necessary investigations to be conducted without treating staff members unfairly.

5. Conclusion

The threat from malicious insiders is a growing concern to many companies but is often overlooked when considering where best to focus security. Insiders act for a vast number of reasons but by considering potential motivations, any weak points in the system can be identified and addressed.

In order to create the best chance of mitigating the insider threat, it is worth identifying the high risk roles in an organisation and determining the right measures to put in place to reduce the chances of an individual causing damage to the company, either deliberately or by accident.

Alongside this, organisations should identify their ‘crown jewels’ and protect them. This may involve not only putting controls in place to stop attackers accessing them, but also reducing the number of individuals with legitimate access so that only those who genuinely require access have it.

There are a number of potential actions that can be taken to address the personnel, technical and business risks faced and in order for a system of controls to be successful it must consider all areas.

Traditional penetration testing will highlight technical risks, but will not be as effective at discovering procedural or business holes. When used in conjunction with business audits and risk assessments, security assurance activities may discover more procedural and business controls that may be lacking and thus needing remediation.

It is possible to perform scenario-based testing around the insider threat which may be the most effective solution. This is where a consultant has the rights of a specific insider group, for example management, technical staff or analysts, and attempts to perform a specific action. This may include exfiltration of sensitive data or planting a (harmless/simulated) ‘logic bomb’ for example. The end report then highlights all issues discovered across all areas and makes recommendations for maximising assurance around insider threats in a holistic way.

6. Further reading & resources

6.1 Further reading

This paper has provided a brief overview of some of the threats from a malicious insider, potential areas for concern and ideas for mitigation strategies. However, this cannot be considered exhaustive. A much more in-depth study into the subject has been completed by both CERT and CPNI and both have extensive resources available online.

The HoMER is guidance produced by CPNI that is aimed at board members and risk managers and gives a framework for assessing and mitigating the risk from the insider threat:
https://www.cpni.gov.uk/system/files/documents/62/53/Holistic-Management-of-Employee-Risk-HoMER-Guidance.pdf

CPNI has also produced a short summary on the nature of insider risk and some of the factors to be aware of. This document provides a useful, concise summary:
https://www.cpni.gov.uk/system/files/documents/63/29/insider-data-collection-study-report-of-main-findings.pdf

CERT has been working on an in-depth study into the insider threat for over a decade and has produced a number of papers examining how the threat differs in different sectors. For a very short, one-page set of guidelines the insider threat best practices page can be used:
https://www.cert.org/insider-threat/best-practices/index.cfm

For a more in-depth analysis of the insider threat, along with a list of 20 best practices to adopt, CERT has released a 'Common Sense Guide to Mitigating Insider Threats':
http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=484738

For a permanent offline reference, similar information along with numerous case studies the 'CERT Guide to Insider Threats' by Dawn Cappelli et al. might be considered.

The full list of resources available from CPNI can be found here:
https://www.cpni.gov.uk/reducing-insider-risk

The full list of resources available from CERT can be found here:
https://www.cert.org/insider-threat/index.cfm

6.2 Available tools

There are a number of tools aimed at detecting and mitigating the insider threat available freely online. There are also a number of commercial solutions available. A selection of the tools are listed here for reference, however, these have not been tested in-depth. Each comes with its own benefits and limitations and should be tested and assessed for suitability before use. Some of the solutions must be considered in line with the end-user's human rights and notice should be given to end-users that they should have no expectation of privacy in a corporate network.

Scout is a proprietary tool that is designed to proactively detect threats and monitors communications for risk indicators:
https://www.strozfriedberg.com/press-release/stroz-friedberg-announces-insider-threat-detection-tool-scout/

Merit is a training simulator developed by CERT that places users in various business situations from which they must make decisions related to insider actions. From this the impact is simulated and users can view the potential outcomes from their actions:
https://www.cert.org/insider-threat/research/merit-interactive.cfm

A research paper from CERT describes the use of anti-plagiarism algorithms to detect the insider threat. The premise is that attackers may be using webmail or other encrypted, unexamined exfiltration vectors. The solution involves building a database of known, sensitive data. However, if this approach was to be implemented, care must be taken to protect this database as it is in itself an attractive target. A simpler solution may be to block webmail access:
http://resources.sei.cmu.edu/asset_files/TechnicalNote/2013_004_001_64688.pdf

In combination with the previous paper, CERT has developed a tool for inserting tags into sensitive documents, known as Tagger:
http://resources.sei.cmu.edu/asset_files/TechnicalNote/2013_004_001_40234.pdf

CERT has developed an ontology for describing indicators of potentially malicious activity. This is an incredibly detailed document and implementation may be difficult, however, it does describe a standardised approach that may be beneficial once it is implemented and running:
http://resources.sei.cmu.edu/asset_files/TechnicalReport/2016_005_001_454627.pdf

CPNI has developed a personnel security maturity model that may be used to assess the current state of an organisation with respect to the insider threat:
https://www.cpni.gov.uk/system/files/documents/c3/69/CPNI-personnel-security-maturity-model.pdf

CERT also provide a whitepaper on analytics that may be used as indicators for early detection of the insider threat:
http://resources.sei.cmu.edu/asset_files/WhitePaper/2015_019_001_451069.pdf

7. References

[1] http://www.computereconomics.com/article.cfm?id=1537
[2] https://hbr.org/2016/09/the-biggest-cybersecurity-threats-are-inside-your-company
[3] http://theinvisiblethings.blogspot.co.uk/2009/10/evil-maid-goes-after-truecrypt.html
[4] http://resources.sei.cmu.edu/asset_files/TechnicalReport/2016_005_001_484758.pdf, page 23
[5] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52247.pdf , page 9
[6] http://www.dtic.mil/dtic/tr/fulltext/u2/a441249.pdf, page 22
[7] http://www.sans.edu/cyber-research/security-laboratory/article/log-bmb-trp-door
[8] https://www.wired.com/2013/03/logic-bomb-south-korea-attack/
[9] http://www.telegraph.co.uk/news/uknews/crime/9845363/Disgruntled-employee-spends-three-years-destroying-work-computers-with-Cillit-Bang.html
[10] http://www.washingtonpost.com/wp-dyn/content/article/2009/02/26/AR2009022601821.html
[11] https://www.theguardian.com/world/2014/feb/01/edward-snowden-intelligence-leak-nsa-contractor-extract
[12] http://security.stackexchange.com/questions/76070/what-is-a-salami-attack
[13] Fraud Case 3, Page 265, The Cert Guide to Insider Threats, Cappelli, Moore, Trzeciak
[14] https://www.cpni.gov.uk/system/files/documents/63/29/insider-data-collection-study-report-of-main-findings.pdf Page 11
[15] http://resources.sei.cmu.edu/asset_files/SpecialReport/2005_003_001_51946.pdf page 3
[16] https://www.cpni.gov.uk/system/files/documents/d0/d2/ongoing-personnel-security-a-good-practice-guide-edition-3.pdf Page 12
https://www.cpni.gov.uk/system/files/documents/af/05/personnel-security-in-remote-working-a-good-practice-guide.pdf
[17] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52266.pdf, appendix c
[18] https://www.cpni.gov.uk/ongoing-personnel-security
[19] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52266.pdf
[20] https://www.forbes.com/sites/benkerschberg/2011/12/07/data-security-and-identity-access-management/#1d7f9bec31de
[21] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52266.pdf Page 10
[22] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52266.pdf page 37
[23] http://resources.sei.cmu.edu/asset_files/WhitePaper/2008_019_001_52266.pdf page 25
[24] https://www.cpni.gov.uk/system/files/documents/d0/d2/ongoing-personnel-security-a-good-practice-guide-edition-3.pdf page 15
[25] https://github.com/
[26] https://subversion.apache.org/
[27] http://searchsecurity.techtarget.com/answer/Enterprise-user-de-provisioning-best-practices-How-to-efficiently-revoke-access
[28] https://www.cpni.gov.uk/system/files/documents/62/53/Holistic-Management-of-Employee-Risk-HoMER-Guidance.pdf page 10 onwards

8. About NCC Group

NCC Group is a global expert in cyber security and risk mitigation, working with businesses to protect their brand, value and reputation against the ever-evolving threat landscape.

With our knowledge, experience and global footprint, we are best placed to help businesses identify, assess, mitigate & respond to the risks they face.

We are passionate about making the Internet safer and revolutionising the way in which organisations think about cyber security.

Headquartered in Manchester, UK, with over 35 offices across the world, NCC Group employs more than 2,000 people and is a trusted advisor to 15,000 clients worldwide.