Mitigating cyber threat from malicious insiders[1]

Authors
Jason Anthony Smith, MSc (Royal Holloway, 2014)
William Rothwell, Abatis (UK) Ltd

Overview

Insider threat is one of the hardest problems in cyber defence. "Traditional" defence and more advanced intelligence-led defence against advanced persistent threats (APT) have been ineffective against malicious insiders. In this article we look at an insider attack as a sequence of phases and examine the characteristics of each phase. We also describe a practical 10-step programme for mitigating malicious insider threat.

Introduction

Employees have a degree of trust invested in them by their employer and have likely been granted physical and logical access to the organisation's IT systems and information in order to fulfil their in-role duties. Most of these employees are thankfully honest, however there exists the threat that some will abuse their insider position and commit crimes like fraud, theft or IT sabotage against their organisations.

Probably the most well-known insider crime was committed by US Army soldier Bradley Manning who, in his role of intelligence analyst, abused legitimately granted access to classified databases and, through the WikiLeaks organisation, released the largest set of classified documents ever leaked to the public. This included 250,000 U.S. diplomatic cables and 500,000 other reports relating to the Iraq and Afghan wars. In 2013 he was convicted under the espionage act and sentenced to 35 years confinement.

When a roundtable of US government agencies compiled a list of "*the hardest and most critical challenges in INFOSEC research that must be addressed for the development and deployment of trustworthy systems for the U.S. Government",* insider threat was ranked as the number two hardest problem behind 'global scale identity management.'

[1] This article is to be published online by Computer Weekly as part of the 2015 Royal Holloway info security thesis series. The full MSc thesis is published on the ISG's website.

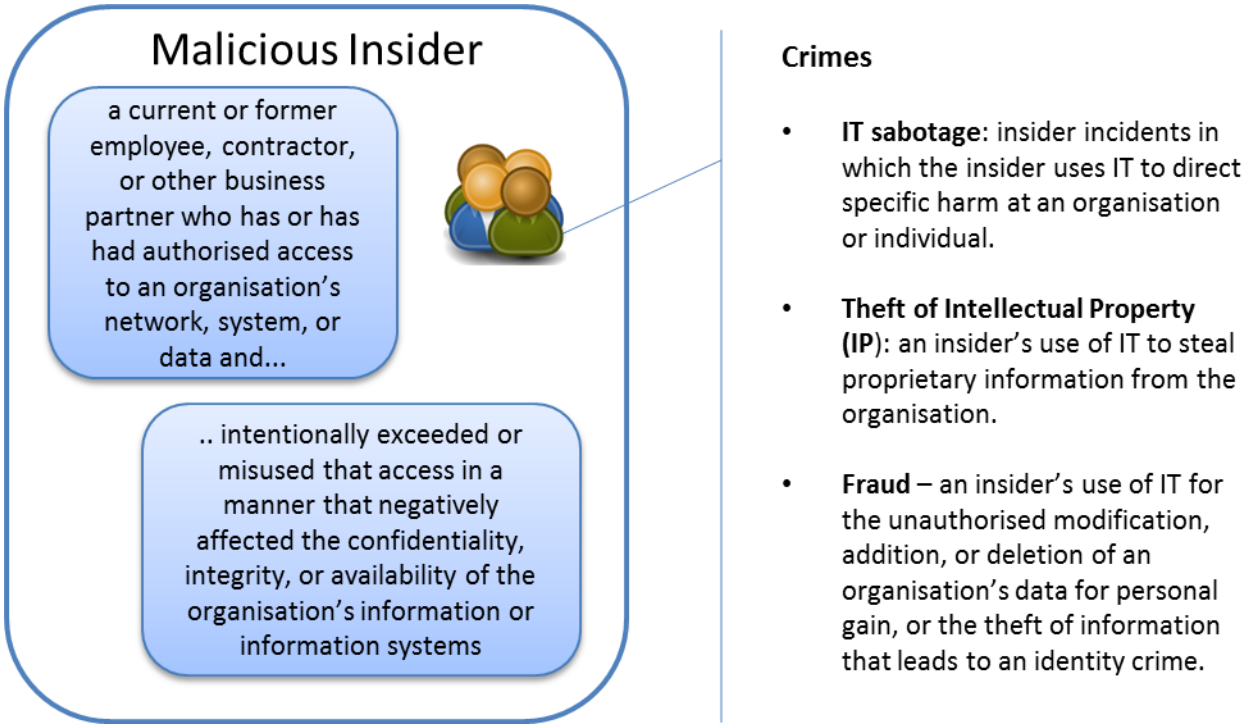

Figure 1. Malicious insider definition

Traditional cyber defence

Most implementations of information security management systems, such as the widely adopted ISO17799, are traditionally focused on keeping external attackers at bay by defending the digital perimeter, identity and access management, testing policy compliance, and removing vulnerabilities from systems. In concert with creating a positive security culture, relieving time and work-load pressure, awareness training, and providing usable security tools and processes, these programmes can also alleviate threats from non-malicious insiders who typically cause incidents through user-error and negligence. There may even be pre-employment screening processes and Data Loss Prevention (DLP) deployed, however these programmes remain largely ineffective against malicious insiders.

Advanced intelligence-led cyber defence

The current digital threat landscape now includes well-resourced sophisticated attackers perpetrating targeted dynamic attacks over long periods of time. These so-called advanced persistent threat (APT) attacks often make use of zero-day exploits, social engineering and other adaptive techniques to get remote hands-on keyboard before ultimately exfiltrating data or completing some other objective.

Defending against APT has required organisations to implement a non-traditional approach to defence often using predictive and diagnostic analytics of security intelligence data which has typically been gathered from specialist companies, government agencies, public sources and through private sharing agreements with trusted partners. These approaches to APT defence is unfortunately also ineffective for defence against malicious insider threat.

Malicious insiders

The sophistication of human beings makes tackling insider threats a difficult problem which requires much more than technology alone. A combination of techniques from the sociological and the socio-technical domains will be required too.

People exhibit polymorphous behaviour in that they tend to hide their true emotions from their peers and their behaviour tends to be self-modifying in response to changes in the environment. For example, an insider could wait for an opportune moment to access a system and steal intellectual property and would be sophisticated enough to continually monitor and evaluate the risk of being caught versus the reward of accomplishing the theft.

Personality is a dynamic and organised set of characteristics possessed by a person that uniquely influences his or her cognitions, motivations, and behaviours in various situations. It is this intrinsic personality, combined with extrinsic temporal and dynamic factors such as motivation and other external stimuli that form the basis for indicators that may correlate with a higher risk of insider threat.

Intrinsic factors (personality) + **Extrinsic factors** (motivation/events/etc) = **Indicator** of higher risk

Emotional distress, disappointment, frustration, disgruntlement, introversion, perceived entitlement, lack of empathy and other traits have been discovered to correlate with an increase in the risk of an insider performing illegal acts or being vulnerable to recruitment or manipulation by others. Strong indicators of insider threat risk include disgruntlement, difficulty accepting feedback, anger management, disregard for authority, low performance, stress and confrontational behaviour; whilst weaker indicators include personal issues, self-centredness, dependability and absenteeism.

A personal or work related event could be the trigger (an external stimuli) which causes an insider to decide to commit a crime against the organisation that trusts them. This could include being fired, financial pressure, family or marital issues, coercion or recruitment by a third party, or many other events or changes in circumstance. For example, people become affiliated with various organisations of like-minded individuals with similar beliefs, ideologies or causes. These could include religious sects, gangs, clubs, political parties, radical and extremist groups etc. Peer pressure from other members of these groups to commit a crime could cause the trigger point to be reached for a particular individual with a particular personality at a particular time in their life when they were vulnerable.

In addition, when an insider crime is being prepared for or is in progress, it's reasonable to assume there may be observable actions that could give rise to suspicion. For example, the US Defense Security Service regard certain behaviours as reportable including:

- Keeping classified materials in an unauthorised location
- Removing classification markings from documents
- Attempting to conceal foreign travel
- Sudden repayment of large debts

- Repeated or un-required work outside of normal duty hours

The key realisation is that a complex interaction of many factors lead to the malicious insider threat being realised with no single one of the factors being a decisive predictor in its own right. Rather, each extra indicator adds weight to the possibility that an insider is going rogue.

The phases of a malicious insider attack

An attack can be considered as a sequence of phases, each of which may or may not necessarily be present and which may also overlap with each other during some attacks. Using a phase based model of defence and countermeasure strategy was described in a paper by Lockheed Martin in 2010. In contrast to traditional network defence strategies, which focus upon the vulnerability element of risk, this approach focusses on the threat element of risk.

For each moment along the time sequence of the attack there is an opportunity to design countermeasures to monitor and mitigate the threat. This sequence of opportunities is called the 'kill-chain'.

At each phase of the kill-chain, observable data will likely be generated and can be monitored. It may be generated by technical actions taken on systems and electronically captured, or as a result of behaviours and physical actions associated with the insider which may be observed by other people. One or more observable data points might be considered to be an indicator that the threat level of a particular insider is increasing which could lead to a particular intervention or course of action being pursued. The data could be used not just to indicate that an attack is in progress or has succeeded (through diagnostic analytics) but also to warn that the probability of an attack is rising (predictive analytics).

As described earlier, threat mitigation can also be directed during each phase of the kill-chain. Specifically, countermeasures can be implemented to 'deter' insiders from attempting a particular course of action, 'deny' attempts to execute a particular course of action, 'disrupt' a particular course of action after it has begun, 'degrade' a particular course of action to be less effective or efficient, 'deceive' the insider so that they are disadvantaged, or 'defeat' the insider such that all future possible courses of action by the insider are of no consequence to the organisation.

Countermeasures:

- Deter
- Deny
- Disrupt
- Degrade
- Deceive
- Defeat

A characteristics of each phase in a six-phase kill chain model for malicious insider threat is described below and shown graphically in Figure 2.

1. **Non-Insider phase**
    - People start off as non-insiders to the organisation, progress to becoming an insider by being granted access, and then commit a crime which causes them to be labelled as a malicious insider.

- People may have joined an organisation deliberately intending to the commit the insider crime, or may decide to do so after joining. This choice may have been because of recruitment or coercion by others or it may be self-initiated.
- People have intrinsic characteristics and traits which make up their personality. They also have different social and cultural backgrounds, religious and political beliefs.
- People may have criminal records and other historic records of past unwanted behaviour in their background.
- People have families and friends and also may have associations with many other groups of likeminded individuals or individuals desiring similar actions or outcomes. These associations and these people can influence them.

2. **Normal insider phase**
   - A key transition state from non-insider to insider happens when the person is granted access by the organisation as an employee, contractor or business partner.
   - The mental states of people, whilst not directly observable, are able to be inferred to some degree of accuracy by people interacting with or observing the individuals behaviour and demeanour.
   - In the course of performing their role, the insider will perform physical actions and technical actions which are observable and will be the basis for the establishment of 'normality'.
   - There may be a very short or no normal insider phase, particularly if the insider joined the organisation with the deliberate intent of committing a crime. It may be possible to infer normal technical actions for a role from what peers in the same role do.
3. **Tipping point phase**
   - At some point in time there is a tipping point, before which the person is not intending to commit a crime and after which that person has the intention to commit the crime.
   - The tipping point may have been caused by one or more precipitating events, or triggers, in the person's professional or personal life.
   - Events happen in people's personal and professional life continuously which may influence them and change their motivations and outlook. Some of these events will be obvious to the organisation, such as giving the person notice of redundancy, and others will be less visible.
   - Leading up to the tipping point there may have been a progressive deterioration of trustworthiness of the individual towards the tipping point which resulted in observable behaviour changes.
4. **Crime preparation phase**
   - Following the tipping point, in order to accomplish the mission of achieving the crime, the person may need to plan and execute one or more preparatory actions.
   - Actions and behaviours after the tipping point may be different to those actions and behaviours which were considered normal for the individual before the tipping point. For example, the individual may start hording information or working strange hours.
5. **Crime execution phase**

- The actual criminal act may happen after the person has, or should have had, their access rights removed, for example, after they are dismissed from the organisation.
- There may be opportunities to deter or prevent the person from committing a crime right up until the crime is actually committed, both before and after the tipping point.
- The crime may not be an instantaneous or one-time event. Instead it may execute over a long period of time or be repeated numerous times by the same actor.
- Deterrence, detection and prevention may be achieved using technology, process or people controls.

6. **Organisation recovery phase**
   - As well as response and recovery from the crime, this phase should give opportunities to perform root cause analysis and feed lessons learned back into the defence system.
   - Depending on the nature of the crime, the organisation may not ever fully recover.

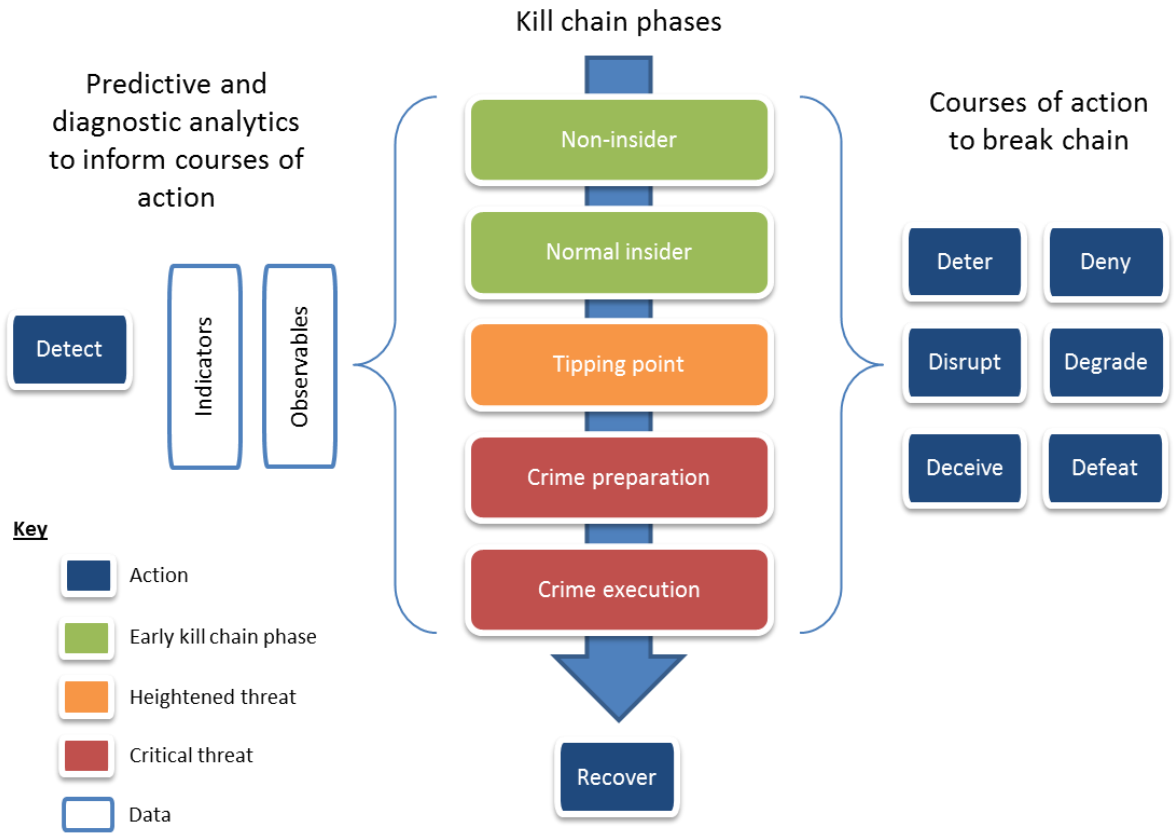

Figure 2. Malicious insider cyber threat kill-chain model

A practical 10-step programme for mitigating malicious insider threat

Implementing a programme to effectively mitigate malicious insider threat can be complicated, costly and take a long time – but organisations still need to make a start. The paragraphs below describe a programme which incrementally builds organisational capability in 10 practical steps.

1. **Recognition** In the first instance, the threat to the organisation from malicious insiders must be recognised by someone, perhaps the Chief Information Security Officer (CISO), who is capable of championing its importance to senior executives and establishing support to build a mitigation programme.
2. **Baseline security** A pre-requisite for a mitigation programme focussed on malicious insider threat is a well-functioning traditional Information Security Management System (ISMS) which maintains a baseline security control environment for the organisation. This will serve to protect the organisation from a broad range of threats and will play a role in reducing risk from malicious and non-malicious insiders.
3. **Incident response** Extend existing incident response processes to properly handle incidents caused by malicious insiders and to respond to escalations by staff who have suspicions of wrong-doing by others.
4. **Communication and awareness** Establish, communicate and enforce acceptable use policies which sufficiently set out the expectations and rules around using the organisations information and systems. Introduce insider threat awareness and training programmes to begin to educate personnel to recognise the threat and respond to it.
5. **Identify critical assets** A well-run traditional ISMS should already have established the critical assets of the organisation and this should form the starting point for prioritising increased protection against the threat from malicious insiders. In particular understand the location and flow of sensitive or valuable. The systems processing this data should be considered critical as well as any systems which play a critical role in the execution of mission critical business processes
6. **Access control** Ensure that identity and access management processes are operating efficiently. This means that people are granted least privilege entitlement for access to systems; segregation of duties are checked and enforced; and access changes are made quickly following personnel joining, moving roles within, or leaving the organisation. Pay particular attention to those users with elevated privileges, and as with other measures, in the first instance focus on access which is to critical assets.
7. **Vetting** Introduce vetting processes prior to on-boarding new personnel and negotiate contractual clauses into master services agreements with business partners for a similar level of screening to take place. Focus in the first instance on those people who will be granted access to critical assets.
8. **Data Loss Prevention** Consider Data Loss Prevention (DLP) and Information Rights Management (IRM) implementations, especially if exfiltration of sensitive data has been identified as a key risk for the organisation. Consider the introduction of honey-tokens and configure DLP to detect and respond to their discovery in unauthorised locations.
9. **Monitoring** Ensure Security Information and Event Monitoring (SIEM) logging and event analysis and correlation processes are properly functioning and establish a baseline of normal network traffic patterns and system usage.

10. **Data analytics** Extend the scope of monitoring, logging and auditing to harvest richer data relevant to insider threat including behavioural data and personal event data from HR. Implement data analytics capabilities for prediction and diagnosis of insider incidents. Take particular care and apply real resources to ensure the monitoring teams are not overwhelmed by increasingly larger volumes of data each day and the expected benefits get lost in the noise.

Concluding remarks

Imagine the difficulty of defending against a sophisticated well-resourced, perhaps state-sponsored, attacker whom was deliberately placed in the organisation to target specific information for which their role would typically give them access. Whilst it may be prohibitively expensive to defend against this degree of attack, there is certainly a great deal that can be done at the opposite end of the spectrum. For example, many people might think it is normal, or an unspoken right, to take customer lists, or other intellectual property (IP), when they leave an organisation. Significantly reducing this threat may be as simple as rolling out a well thought through awareness campaign which explains the wrongs of IP theft.

Biographies

Jason A Smith has worked in IT and Information Security for more than 25 years and is currently the Chief Security Architect at BP. He was previously the CISO for BP's global trading businesses and has held senior risk management roles in many industries including energy, financial, manufacturing and telecommunications. He gained an MSc with distinction in Information Security from Royal Holloway, University of London.

A qualified security professional, *William Rothwell* is an experienced security practitioner with over 25 years in the information security field with clients including government, military, finance and commercial organizations. He specializes in security architecture and design, crypto application and system intrusion defence research. In 2004 he founded Abatis Limited, which provides a unique, patent protected, non-signature based solution to protect computers from malware infection and hacking intrusions.