DARKTRACE

WHITE PAPER

# Detecting the Threat Within

## The Real Challenge of Insider Threat

# Introduction

Bradley Manning, the soldier convicted in July 2013 of violations of the U.S. Espionage Act, and Edward Snowden, the former contractor to both the CIA and NSA, have recently made headline news. The damage of insider attacks may not be a new issue, but the sheer volume of data that they were able to extricate while holding positions of trust, and the profile of the organisations that they stole from, make these individuals particularly interesting.

The cases have created a wave of interest throughout the IT security world and beyond. The fact that even the most secure establishments (and, by implication, networks) could be breached by those with motivation and technical know-how sends a clear signal to those who have been entrusted to protect the security of an enterprise. If it can happen in the U.S. intelligence and defence community, with its associated level of personal and physical security, it can happen anywhere.

Access to data – and lots of it – has become an everyday norm. The ubiquitous nature of high-speed data connectivity, coupled with the ever-growing capacity of portable storage devices, means that an individual with intent can steal, corrupt and/or destroy vast amounts of data with relative impunity.

Some would argue that the commoditisation of data is one of the greatest benefits of the expansive growth of the internet. It has also become a growing risk for those whose businesses rely on that information, and companies are increasingly aware that their intellectual and commercial valuables are on the line, day after day.

This paper addresses the issue of insider threat by highlighting current industry best practises, but more importantly by describing why those steps are no longer sufficient to meaningfully reduce the risk of serious damage – whether data corruption, information theft or disruption to business activities.

The overwhelming narrative of a burgeoning cyber security sector has been about external threat: without doubt, those threats remain the most significant in terms of scale of activity. Insider attacks happen less frequently, but the threat of them presents a real and potentially far more disruptive danger to both commercial and government activities. This paper shows that an insider threat does not have to be catastrophic – as long as it is identified quickly, and processes exist to intervene. It will suggest that a new approach to the problem, based on monitoring and understanding subtle changes in human behaviour within the network, is required if we are to get a handle on the most potentially devastating attack vector of all.

> **"Mr Snowden learned something critical about the NSA's culture: while the organization built enormously high electronic barriers to keep out foreign invaders, it had rudimentary protections against insiders."**
>
> New York Times

## Defining Insider Threat: An Impossible Task?

When we think about an insider, we tend to think of a shady employee in our midst, with login details to the corporate systems and a building pass. In the internet age, the insider now includes anyone who has access to the digital organisation, whether as a basic or privileged user, broadening insider liability out across a company's entire supply chain, temporary worker base, subcontractors and other related parties with which they are connected.

The Centre for the Protection of National Infrastructure (CPNI) conducted an extensive study on insider activity in 2009, identifying a number of common themes among the individuals and organisations involved, which was updated in 2013, reflecting the concern that this topic continues to engender. The definition of insider that was used in the report is "a person who exploits, or has the intention to exploit, their legitimate access to an organisation's assets for unauthorised purposes".

It is interesting to note how the definition of 'insider' has changed. The emphasis is placed on the fact that the authorised individual has the ability to access data, rather than on their position relative to the company – inside or outside.

The CPNI study identified various types of insider activity, the most frequent types being unauthorised disclosure of sensitive information and process corruption. While a majority of insider attacks are carried out by permanent staff, a significant minority involved contractors and agency or temporary staff. 59% of insider attacks had a duration of over six months, with 11% lasting for more than five years. 60% of cases involved individuals who had worked for the organisation for less than 5 years.

Such findings are insightful, and highlight the disparity between different types of insider threats and their methods. Insiders can be classified into five categories, which also include the individual who is unaware that they have exposed their company to risk.

| | |
|---|---|
| Malicious | Someone who deliberately exploits their access to do reputational and/or financial harm to the organisation |
| Opportunist | Someone who steals data for his own gain, often on leaving the company. Motivated by self-interest, rather than malice against the company |
| Negligent | Someone who circumvents security and puts data at risk in the name of convenience |
| Unwilling | Someone who has been coerced by an outsider to provide support to a fraud or cyber-attack |
| Unwitting | Someone who has unwittingly exposed the company to risk, often through infected IT |

All of the types above are not limited to company employees, but include anyone who has access to the data that you wish to defend. The fundamental question must be what are 'insiders' inside of, exactly? The fact is that ideas and intellectual property are routinely stolen from companies, yet insider activity mostly goes undetected until the damage is done.

This paper will address the fundamental contradiction of the 'insider' in a world of blurry boundaries, and examine why old ideas of perimeters – of our businesses, networks and information systems – have become degraded and irrelevant in today's digital age.

# Getting the Basics Right: Security Best Practises

The threat landscape has steadily grown in complexity, resource and sophistication. However, while advanced attacks are on the rise, a significant proportion of security incidents could be avoided or minimised by adopting basic security measures. These measures can be classed in three categories, People, Network and Data Management.

## People

### 1. Basic background checks

Many internal attacks are committed by repeat offenders. Organisations are advised to perform background checks on employees, including former employers where suspicion is high.

### 2. Acceptable use policy

Acceptable use policies are helpful to establish the behaviour expected from employees. They should sign an agreement that clearly states that accessing unauthorised data is a serious offence, and which educates them about good password practises and physical security protections.

### 3. Effective HR termination procedures

Every company should have a standardised termination process when an employee leaves the company, which should include the prompt termination of the person's access to company buildings and digital networks.

### 4. Access rights associated with role / function

As an employee changes roles within an organisation, companies should formally reassess his or her access rights to ensure that permissions and access do not simply accumulate. Passwords that were previously known to the moved employee should be changed.

### 5. Supply chain & third-party agreements

A significant number of data breaches are performed by trusted third parties. All third parties should sign an acceptable use policy, and all programs that they use should observe good security hygiene procedures, such as having up-to-date anti-virus protection and firewalls, etc.

### 6. Separation of role and functions

Policy controls should be put in place that prevent a single employee from abusing their power and significantly damaging the company. For example, highly-privileged accounts should be separated from the user's regular account.

## Network Management

### 7. Effective domain management

If users do not need access to a particular network or computer, don't give it to them. Not all workstations need to be connected to other workstations.

### 8. Least-privilege access control

Access control should be granted by the application to the data owner, and all control should be audited periodically.

### 9. Role-based access control (RBAC)

RBAC provides least-privilege control determined by authorised actions and the role of the user. In RBAC systems, least-privileged permissions are granted to role-based groups instead of user accounts or departmental groups.

## Data Management

### 10. Data leak prevention

Effective monitoring systems can help prevent against the access or transfer of unauthorised data. However, it is necessary to configure them to recognise unauthorised data, and they must be highly tuned to avoid a high number of false positives. More importantly, somebody needs to be accountable for monitoring and acting on the information identified from these systems.

### 11. Encryption

Implement encryption in order to protect sensitive data both in transit and at rest, including on storage media such as portable hard drives and USB keys.

### 12. Design security into your system from the outset

Poor default security or misconfigurations cause many security incidents. It is important to ensure all computers have been configured using industry-accepted best practises.

### 13. Configuration and change management

Ensure that computers do not get modified and become less secure. All computers should be periodically audited and require a control process that prevents unauthorised modifications.

### 14. Understand your workflows

Data isolation can be performed by automating workflows that require interaction with data. A good automated workflow minimises mistakes and access to the underlying database.

**"When you look at the sheer volume of the attackers, it really shows that certainly an organisation is going to have more outsiders than insiders, no matter what. Just with the sheer number of possible actors, that's going to be the case forever. But that doesn't negate the fact that insiders can do damage."**

Suzanne Widup, senior analyst, Verizon RISK team

## Mind the gap: threat in today's environment

The detection and prevention recommendations listed above can help protect companies against a significant portion of internal attacks, and reduce the likelihood that employees tempted to betray your trust have the opportunity to do so. As such, they are very useful steps that companies are advised to follow.

That said, it is clear that these guidelines only go so far in dealing with the problem of cyber threats, namely because:

**a) They're not always practical**

Not all of these best practises are implementable in every network or organisation, and it is virtually unheard of to implement all fourteen best practises. However sensible these guidelines are, they remain an ideal standard, which do not always bear up to the reality of running a modern business.

**b) They don't go far enough**

Even if you did manage to implement all 14 best practises, your troubles would not be over. Trying to second-guess how the next insider threat will manifest itself is impossible, as there are always ways around traditional defences. Something more is needed that does not rely on policies or static barriers.

Let us examine in more detail why traditional best practises fall short and ultimately fail to protect organisations against the advanced, persistent threats that plague the landscape today.

## You can't define the perimeter

At its heart, the traditional approach to security falls down due to its fundamental dependence on making a distinction between 'inside' and 'outside', both in terms of information and users.

The world has moved on in the past ten years, to the point where perimeters that

used to demarcate one network from another network have become difficult to define and more difficult still to police.

Supply chains are particularly difficult to manage: on the one hand, they represent convenience, and on the other, they bring complexity and vulnerability to the corporate network. M&A activity is another factor that carries a double impact of growth and expansion, but also of liability to unknown risks. While financial due diligence is usually very thorough, network topography and control points are sometimes overlooked, and two completely different networks are simply plugged together – with unfortunate results.

And of course as enterprises have evolved to support web-based business services, e-commerce, internet hosting and mobile devices, cloud computing has vastly expanded the attacker's surface of opportunity. The cloud perimeter does not really exist at all – it is off-premises and ever-changing. Multi-tenanted data is practical but means that you have no control over the infrastructure, and makes it difficult to draw down on specific front-line network configurations, since they must be suitable for the entire cloud tenant base.

The reality – the 'new normal' – is that the information systems and networks that we have progressively built and elaborated are like dark, labyrinthine underworlds, whirring with frenetic activity: they perform amazing feats, but are also riddled with holes and vulnerabilities, at once porous and dynamic.

> **"Systems are too complex and perimeter defences are too porous to treat everything inside the network perimeter as trusted"**
>
> Hugh Thompson, computing security professor, Columbia University

## Who's there? And where's my data?

When the perimeters of our information networks lose their meaning, how do you even know who is inside and who is outside?

Given the dependency of modern organisations on large and interconnected global networks, data has become easy to target. Modern working habits, such as the use of personal devices at work, remote access and uncontrolled physical spaces, mean that segmenting and protecting different sets of data has become increasingly difficult, and often impossible. Data no longer sits static on a corporate server, but constantly flows within and outside of an organisation.

The dynamic nature of information flows means that it is difficult to put our finger on where our most valuable assets actually are. Information is the lifeblood of our organisations, but companies are struggling to identify where the really critical parts of it are. While physical asset management and stock control are normal and taken for granted in product management, the digital world has proven far more difficult to handle. Unlike physical assets, a digital file can be in many places simultaneously, and its movements and duplications are very difficult to track.

Given the sheer volumes of data that we produce and pass around amongst ourselves from day to day, we cannot possibly expect to secure all of it at the same level of protection. Nor would we want to. Not all information is born equal: it is not a disaster if a company's lunch menu gets leaked, but the theft of its manufacturing blueprints would be. Less important data should benefit from greater flexibility, whereas our most valuable assets should be dealt with more carefully.

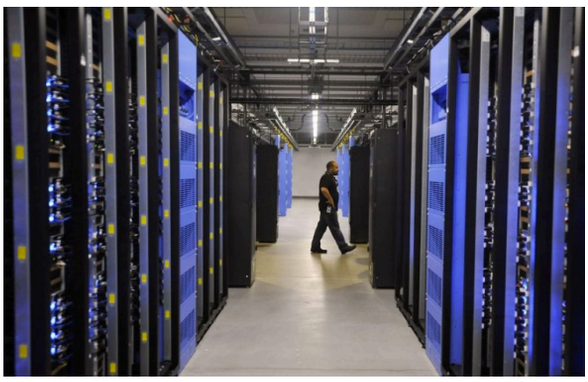

*A data warehouse: Do you know who this person is?*

The difficulty in tracking data has opened up opportunities for people with sufficient inroads into an organisation to inflict serious damage, whether for self-interest or in the service of ideological ideals. However, these people are not necessarily on your payroll – so who are they?

The CPNI study on insiders showed a large range of different insiders, acting from different motivations. While the largest proportion of insiders were motivated by financial gain (47%), the study highlighted a number of other common motivations, such as ideology (20%), a desire for recognition (14%) and loyalty to friends, family or country (14%). General disaffection with the employing organisation was also found to be a significant contributing factor in many cases. With such a range of personalities and motivations involved, it is extremely challenging to identify high-risk users – particularly without alienating your workforce.

Workforces are global in most large organisations, with offices worldwide and an extensive supplier network. Global enterprises typically subcontract various business functions to third parties abroad, as well as employ temporary as well as permanent staff. All these standard trappings of the modern global business mean corporate networks have expanded to adapt to a global and flexible user base.

The rigorous application of 'need to know' and segregation of duties principles, and the strict enforcement of security policies and practises, such as locking computer terminals after use, clear-desk rules and pass access to secure areas will only get you so far, faced with a determined adversary – especially when we don't even know who or where that adversary is.

## Overwhelmed and Under-served

The Holy Grail of the battle against insider threat is to be able to spot incidents as they are developing and then engage early enough – not after the data has left the firewall and the damage is done.

While advances have been made in terms of the data that is available for analysis of network activity, with log management and security information and event management (SIEM), the volume of data produced is often overwhelming. An average firewall alone can produce more than 500,000 messages every day.

**"Trusted insiders with the intent to do harm can exploit their access to compromise vast amounts of sensitive and classified information as part of a personal ideology or at the direction of a foreign government… The unauthorized disclosure of this information to state adversaries, nonstate activists, or other entities will continue to pose a critical threat"**

James Clapper, U.S. Director of National Intelligence, 2014

Multiple studies by SANS have shown that organisations rely heavily on log management and SIEM platforms that are unable to handle the deluge of data fed into them. Nor can threat analysts handle the data coming out of them either. Such tools can help companies investigate and learn more about past threat actors, but – if left unharnessed – they are stuck in a reactive role. Billions of logs producing millions of incidents is of zero use for the threat analyst looking to spot trends, glean intelligence from the data and proactively prevent serious insider threat.

## Inaction

The security profession therefore relies on automation of detection, but requires solutions that do not simply send a list of security events to be chased down and turn out to be no more than slightly unexpected user or network behaviour.

False positives are often the Achilles Heel of machine-learning systems: intrusion detection systems that label normal network traffic as malicious, valid email messages blocked by anti-virus systems, unexploitable software defects that are flagged up by security analysis systems, etc. Large volumes of such false positive alerts make it very difficult for IT security managers to focus on the most severe events, which gives way to a lack of action taken in response to all that data.

A recent SANS study on security analytics found that another major challenge was obtaining a better overview picture of the entire infrastructure. Without any overall insight into information systems, and without knowing which data assets you most want to protect, the collection of data risks becoming a meaningless process, which overwhelms rather than illuminates the security team.

## Overzealous Security Policies

Whilst good basic security hygiene, such as updating anti-virus software and changing passwords regularly, is important, organisations should be wary of overly restrictive policies. According to a recent report, almost half of employees that admit to breaking security policies do so because they were an impediment to their job. 40% believed that if they breached the policy, it would go undetected. It is evident that security policies are not sufficiently effective as deterrents to uncompliant behaviour. They also encourage a reporting culture that makes no distinction between a time-constrained or careless employee who means no harm, and an insider intent on damage.

## Behavioural Problem, Behavioural Solution

We have discussed the ways in which security practise can be significantly improved, in the first part of this paper. Secondly we have examined why these basic steps – even in the rare scenarios where they might be fully implementable – fail to adequately protect corporations and organisations from serious cyber threats.

Despite a healthy debate about good cyber practise, with the support of business confederations and governments, which have helped bring the issue into the boardrooms, we still see large companies fall victim to malicious attacks from one month to the next. The cyber solutions industry needs to fundamentally refresh its approach if it is to properly address this new age of threat, where the attacker's speed of evolution and the scale of the operation far exceeds those of the defender.

At its core, every insider threat is directed by a human being, and can be boiled down to human behaviour, as manifested on computer systems. Behaviour is therefore key to understanding threat, and not just as a means of investigating an attacker retrospectively. Behaviour refers to the way in which one acts or conducts oneself – for humans, it is not static or standardised across the species, but is by nature dynamic and changeable. Humans are incredibly good at adapting themselves to different situations and contexts, and this is exactly what an attacker does too.

The ability to see and understand different types of behaviour, as they evolve and adapt, is the only way that it is possible to detect real threats to a high degree of accuracy. In turn, it is the only way that threat analysts and security managers can focus on true threat and intervene in the attack kill chain before real damage is done or information is stolen.

## The Darktrace Approach

Darktrace takes a novel approach to the problem of cyber threat, which takes as its starting point the presence of threats on all information systems. However well patched the network is or well trained your staff may be, there is always a risk of compromise – and the first step to a resilient cyber strategy is an acknowledgement of this reality.

Given that assumption of ever-present threat, it is unnecessary and indeed unhelpful to make false distinctions between inside and outside. It is precisely the breakdown of network barriers and the blurring of our perimeters that has led to the complexified threat landscape that we face. Darktrace uniquely responds to ever-evolving threats with adaptive, ever-evolving technology. It does not sit at the network border, or rely on sets of rules that are outdated as soon as the threat has changed its tactics.

Based on complex, probabilistic mathematics, Behavioural Cyber Defence is a new category of cyber technology that passively sees all network interactions and events and self-learns to build dynamic models of the normal behaviour of each user and machine, and the enterprise as a whole. As more information is gathered and contexts change, these models adapt themselves to each new set of circumstances, and are continually updated. This evolving picture of network activity and normal behaviour means that Darktrace is uniquely able to spot abnormal behaviour, as it begins to manifest itself in real time. The technology is able to form a highly compelling picture of threat activity by correlating multiple subtle shifts in behaviour, and sends alerts based on its understanding and judgement.

## Patterns of Life

Darktrace Cyber Intelligence Platform (DCIP) is the flagship cyber defence software solution pioneering Behavioural Cyber Defence. The solution has been specifically designed and developed for the threat within faced by large enterprises and government bodies at risk from state-sponsored espionage, organised criminal groups and 'insiders'.

Using network flow data, DCIP builds unique 'pattern of life' models for every network entity, human and machine, consisting of roughly 250 dimensions, such as reboots, protocol changes, and external connections. For example, if an HR employee logs on to the network from an unknown overseas IP address, DCIP will make a judgement about the user and the device. It knows that it is unusual for this particular member of staff to log on remotely, and the device in question is a laptop, but very rarely logs on from outside the enterprise. DCIP would attribute a higher threat level to this activity than if it were a legal advisor, for whom travel was usual, and mobile device use common.

The power of this approach lies in its rejection of the binary rules and signatures of traditional models. The Darktrace platform does not need to make a rule for every possible transgression that might catch an attacker. Instead, it understands the changeable nature of human and machine behaviour, and detects subtle anomalies that are truly indicative of threat.

The real power of this approach is that there are no preconceived rules about threat behaviour, or definitions about how the threat might manifest itself. Employees do not have to predetermine how an attack might take place, or where it might come from. Rather, once a pattern of life is developed, DCIP is able to detect the subtle outliers to the pattern of life, and take a probabilistic view of whether these indicators are indicative of threat. Unlike traditional endpoint defences, an attack does not need to have been seen before in order to be detected.

> **"Our approach of using cutting-edge mathematics to make sense of weak indicators buried in a sea of data is absolutely critical, and allows you to engage much earlier in potentially devastating situations. If your current critical success factor is how fast you can tidy up, you're missing the point."**
>
> Andrew France OBE, CEO, Darktrace

This ability to self-learn and detect normality in order to spot true anomalies is ground-breaking, allowing organisations of all sizes to understand the behaviour of users and machines on their network at both an individual and group level.

## Conclusion

When Edward Snowden singlehandedly stole classified documents from one of the most secretive and well-defended organisations in the world, a rising wave of suspense came crashing down within the information security industry. The uncomfortable reality that our key information systems are insecure can no longer be avoided, and while insider threats may be rare, they have greater potential to be deadly.

'Could this happen to me?' is a question that keeps many security officers up at night. Insider incidents do not need to be of the scale of a Snowden or Manning to have a significant negative impact on your business and reputation. In spite of the introduction of enhanced security measures and better sharing of industry best practises, insider threat remains an unsolved problem that threatens the very survival of many businesses.

To solve this problem, we need to reset our mentality around cyber security. While organisations are rightly encouraged to make various improvements to their security policies in order to reduce their exposure to risk, persistent adversaries have proven, time and again, that they can persevere and triumph against legacy defences.

Networks today are too complex, people are too variable and data is too dynamic to rely on perimeter and rules-based approaches to cyber security. Instead, organisations need an adaptive immune system to keep them healthy. Darktrace's Behavioural Cyber Defence technology is based on mathematical innovation and plays the same critical, self-learning defensive role within the security infrastructure that the immune system does within the body. It accepts the reality of threat all around, and performs highly complex analysis of its environment, based on small and subtle indicators, to mathematically identify emerging patterns of abnormality.

When boundaries break down and perimeters fail, the definition of the insider starts to lose its relevance. Whether from inside or outside, threat actors are becoming harder and harder to spot because, one way or another, they are acting from within. Only the most sophisticated technologies that understand threat behaviours at the deepest level are capable of defending organisations in a world where everyone is an insider.

## About Darktrace

Darktrace is the world leader in Behavioural Cyber Defence technology. Based on pioneering Bayesian mathematics developed at the University of Cambridge, Darktrace's unique approach helps organisations to defend against insider threat and advanced, persistent attackers within the network, by detecting new attack vectors as they emerge. Darktrace's self-learning platform works out normal and abnormal behaviour within an organisation in real time, in order to detect anomalous and threatening activity. Darktrace is made up of world-class cyber intelligence experts and mathematicians. The company is head-quartered in Cambridge, UK, with offices in London, New York, Paris and Milan.

T: +44 (0) 1223 350653 E: info@darktrace.com www.darktrace.com