PA Consulting Group

National Protective Security Authority

HOLISTIC MANAGEMENT OF EMPLOYEE RISK (HoMER)

Next

Foreword by the Information Commissioner

The Information Commissioner's Office (ICO) has been involved with the HoMER project since its planning stages. It has been of particular interest to us because - on the one hand - we are concerned with the information rights of employees subject to workplace monitoring and - on the other - we have an interest in the security of personal data held by organisations. We recognise that keeping data secure inevitably involves making sure employees do not abuse their access to it - indeed the Data Protection Act (DPA) requires organisations to take steps to ensure the reliability of any employees who have access to personal data.

Managing employee risk has become a critical issue for organisations, for which a fine balance is required between treating employees fairly and ethically, and ensuring comprehensive data security. This guidance is positive, pragmatic and business-led. It will help organisations to achieve a step-change in their security and to strengthen their future position in the eyes of their people, customers and other relevant stakeholders.

Too often we have seen workplace monitoring schemes that are cloaked in unnecessary secrecy and that lack proper governance and senior-level accountability. The operation of such schemes can contribute to the negative and damaging behaviour that HoMER's principles are intended to prevent. We have therefore been impressed with HoMER's emphasis on the transparency and governance that must accompany any effective workplace monitoring scheme - this is very much along the lines of the ICO's 'Employment Practices Data Protection Code'.

We recognise that in some circumstances the procedural detail of a scheme - its precise methodology and scheduling - must be kept secret. However, the advantages of being as open as possible with employees about the operation of a workplace monitoring scheme are obvious.

Christopher Graham

Information Commissioner

August 2012

Introduction

Holistic Management of Employee Risk (HoMER) is new guidance to help you manage the risk of employees' behaviour damaging your business. Employee risk is defined as counterproductive behaviour, whether inadvertent, negligent or malicious, that can cause harm to an organisation.

The guidance sets out:

- Principles, policies, procedures and examples of good practice which help manage the risk of counterproductive behaviour in the workplace;
- Ways to strengthen compliance with legal and regulatory frameworks;
- A framework to help improve trust amongst employees, customers and shareholders.

The holistic use of targeted security measures and interventions (e.g. information, personnel and physical) will help you spot high-risk workplace behaviour and reduce the threat of employees carrying out malicious attacks.

This guidance is for board members and the managers of risk in your organisation.

Produced by NPSA in collaboration with PA Consulting Group

Introduction: key points

A risk-based approach: focus on assets and those who can access them. Understand where risks lie and allow your organisation to concentrate its investment in protective measures.

Holistic people management: seek input from across your organisation to address the risk effectively.

A strong security culture: make your values explicit and build a security culture where employees are committed to safeguarding security.

Single accountable ownership of people risk: ensure that all functions with a responsibility for people risk report to a senior, single accountable owner.

Legality and transparency: transparent policies and informed employees will gain trust and buy-in from stakeholders and strengthen compliance.

Find the sections most relevant to you

This guidance is arranged according to its relevance to the board and top level management, the person accountable for people risk and the implementation team.

This interactive PDF allows you to navigate easily to the chapters of most relevance to you. Access each section by clicking on the titles below. Use the forward and backward arrows to navigate within each chapter.

Board and top level management
Introduction and section 1

Single accountable owner of people risk
Sections 1 – 5

Implementation team
Sections 2 – 5 and technical appendix

The guidance is structured as follows:

Vision and leadership: the importance of sound and engaged leadership, corporate governance and transparent policies in managing people risk and strengthening compliance;

Assess: the importance of adopting a demonstrable risk-based approach and of having reliable asset, access and identity management;

Protect: the importance of developing compliant policies and procedures for protective monitoring;

Respond: the importance of advance 'response preparation' to an employee incident and guidance on how respond in ways most likely to minimise damage and retain stakeholder trust;

Recover: effective steps for recovery following an incident while maximising lessons learned in order to improve security.

Sources and types of threat and the process of managing people risk

Sources of threat

Criminals
(inc competitors)

Foreign
governments

'Hacktivists'

Human error

Insiders
(self-motivated or manipulated)

Terrorists

Types of threat

- Espionage
- Seizing control
- Theft, fraud and extortion
- Forensic attack & reconnaissance
- Harassment and coercion
- Disruption / denial of service
- Sabotage

Manage people risk effectively and strengthen your business

Employee risk (‘people risk’) is defined as high-risk behaviour, whether inadvertent, negligent or malicious, that can harm an organisation. It includes risk from employees (including contract or agency staff who have authorised access to your premises and/or assets) to the organisation or to other employees.

Occasionally members of staff betray the trust placed in them. This can have catastrophic results.

In the face of increasing technological complexity, threats from cyber crime, the high-impact fallout from incidents and a toughening regulatory framework, **effective management of people risk can reduce risk across your organisation and strengthen its competitive position.**

Case history 1: Indicators for people risk

Ms G was a 40-something, middle-ranking employee who had worked for the same organisation all her adult life. Her line manager was 10 years her junior and seen by the organisation as something of a rising star. In her annual appraisal, the line manager described Ms G as ‘change resistant’ and ‘an average performer’. Ms G complained to her senior manager, non-confrontational Mr A. She argued that her line manager had given her no credit for the quality of her work, or for the depth of experience she brought to it. She described his attitude to her as being ‘close to sexist and ageist bullying’. After some discussion, Mr A said that although he recognised the important contribution she made, he could not uphold her complaint. He urged her to improve her relationship with her line manager and to broaden her knowledge of the organisation to try to gain a better understanding of how it was modernising and expanding.

In the ensuing weeks, access control data recorded that Ms G had changed her pattern of work, was working longer hours and was often the last person to leave the office. IT logs captured the fact that Ms G was frequently accessing corporate data that had little to do with her job and many of her searches appeared to be random and unrelated. Strictly speaking none of these searches broke corporate policy but the accumulated data would be of significant value to a competitor.

What are the indicators for people risk in this case history?

Would your organisation identify these indicators?

What conclusions can you derive from the change in behaviour?

Vision and leadership

This section introduces the high-level principles behind HoMER. It provides guidance on setting the vision for your organisation and discusses some of the benefits that a holistic approach can bring in addition to enhanced protection from people risk.

Section one: key points

The HoMER vision

- A positive, clear security culture based on trust.
- Well-targeted security measures.

Leadership and culture

- Senior buy-in is a pre-requisite for implementing HoMER.
- Decide on and develop a clear organisational security culture.

Corporate governance

- Appoint a senior single accountable owner of people risk.
- Take a holistic approach , drawing on information from across the organisation.
- Manage people risk as a corporate risk.

Transparent, ethical policies

- All principles, policies and procedures in HoMER need to be transparent and accessible.
- Take an ethical approach, proportionate to the risk, to gain support.

Compliance

- Always act in accordance with the law and regulatory frameworks.
- Move beyond compliance and establish what risks matter most for your organisation.

The HoMER vision: a holistic, risk-based approach

HoMER aims to prevent damaging employee behaviour in a sensitive, ethical and transparent way. Central to it, is that information about people risk should be brought together under a single point of accountability. This is the holistic approach, but it can only work if employees respect and understand the measures being taken to protect the organisation and its people.

A holistic approach to managing people risk will help your organisation to have:

- an effective security culture, and
- employees with values which support this culture.

A risk-based approach

Organisations which invest in security without maintaining an up-to-date personnel security risk assessment may find they have invested in the wrong place or that they could have achieved the same results more cost effectively.

Leadership and culture

An organisation's leaders determine its culture and values. These cultures and values should be clearly expressed in transparent policies which can be accepted, followed and understood by all employees, from the board down.

Build a strong security culture

In a strong security culture, employees display an intuitive awareness of risk, security and trust in a way that attracts the respect of colleagues, the admiration of regulators and the on-going trust of customers.

Evidence of a strong security culture can be observed if all employees make a contribution to managing the problem. As well as complying with policy themselves employees should be empowered to challenge colleagues who do not.

A simple example of the impact that senior leadership can have on security within an organisation is in the wearing of passes. If leaders do not wear their passes it sets a bad example and weakens the security culture of the organisation.

Studies show that most employees are unwilling to challenge colleagues who are not complying with policy. They are prepared to report breaches 'up the chain', but won't step in themselves. In some circumstances, the response to a report may be too late to prevent a security problem becoming an incident. Organisations should identify those policies which they want employees to actively enforce. They should provide the necessary training and support to enable them to do so in a non-confrontational manner which does not destroy the trust between colleagues.

Top tip: Evaluate your security culture

NPSA has developed a tool to help your organisation measure its security culture, diagnose problems and implement changes to organisational culture.

To find out more about the Security Culture Review and Evaluation Tool (SeCURE 2) you should contact NPSA-enquiries@npsa.gov.uk

Case history 2: HoMER in practice

A UK technology organisation employs people across 20 countries. It believes its most valuable assets are its people; its culture; its reputation as a consistently high-performing and trusted supplier; its market share; its customer base; the confidence of its shareholders and its Intellectual Property. It has data centres in three countries.

In an environment where legal, industry and regulatory codes of practice are inconsistent and dynamic across the globe, it prizes the ethical approach that underpins its corporate values and transparent business culture. This helps drive the right behaviours in its people, enhances overall business performance and ensures compliance.

Its internal regulatory framework is simple, accessible to all and is promoted as its central guide to 'always doing the right thing'.

People risk is 'owned' by the global director of security who has several functional leads reporting to him: Operational Risk, Information and Technology Security and Personnel Security. He brings these together through frequent, close and formal review of the 'top 20' corporate risks, which include risks presented by the organisation's people, to ensure a coherent risk picture is available and being managed. Where these touch on other corporate functions, such as HR and IT, he takes action to align or engage them.

How many of your organisation's top risks relate to people?

How does your organisation identify, assess and manage people risk?

Is there clear accountability for the management of people risk in your organisation?

Corporate governance

Appoint a board member as single accountable owner of people risk

One of the key principles of HoMER is to appoint a board member as the single accountable owner of people risk. In most organisations the responsibility for people risk is spread across different people in different functions. Appointing a single accountable owner helps your organisation ensure that people risk is assessed and mitigated effectively.

Take a holistic approach

Without a holistic approach, silos develop and indicators for an incident may be missed because they occur in separate silos. To break down these silos, the single accountable owner should have a broad team accountable to them from across the entire organisation.

Manage people risk as a corporate risk and apply the same principles of oversight

The major people risks can be comparable to other significant corporate risks and require the attention of the board. It is for this reason that the single accountable owner of people risk must be a senior member of the organisation and report to the CEO.

As people risks are comparable in size to other corporate risks, organisations should apply the same principles of oversight. This should include:

- ensuring the Audit Committee is fully informed of the organisation's use of HoMER and has approved it
- ensuring that the Audit Committee regularly reviews people risk and takes an active interest in its status
- ensure that people risk is the explicit responsibility of a non-executive director to balance the executive discharge of this responsibility.

Build transparent, ethical policies

A transparent approach

Communicate principles, polices and procedures clearly. Ensure policies are easy to understand and follow. If they are difficult or inaccessible, employees will not comply and breaches or ‘work-arounds’ will create security weaknesses.

Consult with unions or other employee representation as part of the process.

Be prepared to enforce your policies. To be effective, organisations must react to non-compliance, irrespective of the outcome. Non-compliance can easily become a habit, resulting in lack of respect for policies in general.

An ethical approach

It is vital for an organisation to be clear about its ethical stance, since this will underpin its internal culture and ultimately its approach to everything it does.

An important part of HoMER is the guidance on protective monitoring. An organisation’s ethical position and beliefs will determine whether it applies protective monitoring and to what extent, as well as its approach.

Beyond compliance

In today's increasingly regulated world, all organisations need to comply with mandatory requirements, as enacted in the law. Compliance with these external regulatory frameworks is necessary. However, it is not sufficient because:

- The legal and regulatory framework is dynamic, it changes with time and it varies from region to region.
- Interpretation can be subjective and one expert's view may be different from another's.

For these reasons the principles in HoMER advocate an approach founded on an organisation's own principles, ensuring that the measures it takes are appropriate, proportionate and reasonable and that it adopts a holistic risk-based approach throughout.

Top tip: Information on legal issues

Find more guidance on legal issues on the Information Commissioner's Office website:

Data Protection and Freedom of Information advice – ICO

Case history 3: A transparent approach?

Ms D has worked in a Critical National Infrastructure (CNI) sector as an office manager for three months. From the start she made a point of getting to know the company's policies, especially the security policies, telling her colleagues that the rules wouldn't be there if they weren't important.

She recalls from her contract of employment that employees were subject to monitoring. However, she hasn't seen or heard any further company reference to monitoring since she joined.

Her friend in the company's network services department mentions to her that he is currently working on the installation of new automated monitoring software. He has been impressed at how configurable it is and disappointed that it was being installed 'out of the box'.

One of the company's security policies states that any transfer of data from the company networks via removable media must have prior authorisation. Ms D is not aware of the process for such transfers but has seen colleagues insert USB sticks and CDs into their company laptops. She thinks they may have been downloading company material. While she's not suggesting that they are doing anything wrong, she's concerned that they don't appear to follow any process to obtain permission first.

Is the application of protective monitoring transparent in this case history? How could it be improved?

Is the application of protective monitoring legal in this case history?

Assess

This section explains how to identify risks by focusing on assets across the organisation rather than area by area. This will ensure that nothing is missed. Once you understand what you are trying to protect you can identify the vulnerabilities and threats that lead to the risks and decide on countermeasures.

Risks that arise from people can often be mitigated by doing simple things well.

Section two: key points

Apply a risk-based approach to strengthen compliance and focus investment

Use a risk management framework that works for your organisation

- Undertake regular risk assessments for your key assets, including personnel security risk assessments.
- Identify countermeasures. Start with those that are already available.

Security breaches, what they say about the security culture and countermeasures

Legal requirements

Relevant UK legislation and codes of practice

Making the business case

- The business case is the responsibility of senior leaders.
- Consider all elements of the risk assessment, including safety and compliance.

Apply a risk-based approach to strengthen compliance and focus investment

Start by completing a Personnel Security Risk Assessment

Measures should be based on an assessment. Risk assessment is commonly accepted as good business practice. There can also be a regulatory expectation that organisations will complete such assessments frequently.

Ensure that risk assessment is holistic

When assessments are carried out in different areas of the business, i.e. in silos, risks (and simple measures for managing them) can be missed.

The single accountable owner of people risk will be responsible for ensuring that risks from people across the entire organisation are identified, managed and reduced to an acceptable level.

Once your organisation has identified its people risks it can make an informed decision on mitigation. This could involve a change in process, an investment in new software or hardware, or a decision to tolerate a particular risk.

By taking a risk-based approach and making informed decisions your organisation will make protection more cost-effective, by using existing processes where possible and only investing when it makes business sense. There are many instances of organisations making uninformed investment in protection that provides only cosmetic assurance. One example is identity and access management systems where significant investment delivers a poor return if employees are allowed access long after they have resigned or retired.

Use a risk management framework that works for your organisation

Choose and apply a risk management method which suits your circumstances. Any method you use should include:

- an assessment of the main vulnerabilities and threats to the assets your organisation values most
- an assessment of the likelihood and impact of the risk – this will allow you to prioritise the risks
- identification of possible countermeasures
- an assessment of the controls already in place to mitigate against the risks, and identification of residual risks
- a plan that identifies countermeasures to reduce the residual risks to an acceptable level.

All of the information gathered through a risk management process should be stored securely on a risk register.

Too often risk management is used as a 'tick in the box' for compliance. This is contrary to its primary value as a stimulus for action. Proactive risk management means making regular assessments as the environment changes. In addition, risk assessments should be undertaken in response to internal and external events.

Larger organisations will need to align risk registers across the various sources and owners of the risks. Senior executives can then consider people risk within the context of the organisation's wider risks and take coordinated action. HoMER provides a framework for managing people risk during this process.

Top tip: Risk management

An organisation should:

- have an understanding of its valued assets
- understand who has access to them
- identify what risks those with access could pose to those assets.

Further information on risk assessments can be found on the NPSA website: Personnel Security Risk Assessment

Case history 4: Risk assessment

Acme plc. is a successful high-tech organisation serving the energy sector. It produces components for industrial control systems from its own designs and employs 150 people. Four weeks ago Ms G joined the organisation as the new CEO. Her first impressions were of an organisation with brilliant technical people who have given a lead in the market. However, she's concerned that the focus on technology may have led to complacency in some of the organisation's approaches to its administration and security.

Four specific examples have led to her concerns:

1. On a visit to the organisation's manufacturing facility, she was briefed on what was described as 'the inventory control system'. This kept a record of all components valued at more than £500, as well as finished stock. In answer to her question, there was no knowledge of any separate inventory of information or data assets.
2. On a tour of one of the organisation's design centres, when she asked the manager how the designs were protected, he said they were given the same protection as any of the organisation's documents and material. The networks were all properly firewalled and a market-leading antivirus (AV) software was used. In response to one of her questions on a new design, the presenter logged in to 'his old account from his previous job' to show her the relevant images. Asked when he had left his previous role, he said 'about three months ago' and 'it was very useful still to be able to access these drawings'.
3. Visiting HR, she was told that the organisation had three databases of employees, one of roles, one which showed who was responsible for which assets, and the payroll. Asked which one was used as the master, the HR Director paused and said that was a good question and that he thought it was probably the payroll, and suggested it was really a question for the CIO.
4. Visiting a data centre, she saw the same access control procedures (swipe card in and out) applied to her as were used across the rest of the organisation. She was particularly interested in the size of the media library that she and a colleague walked through on their way out of the data centre.

Example risk assessment

The following is an example of a risk assessment from the case history: Risk assessment.

Assets	Vulnerabilities	Threats	Likelihood 1 = least likely, 5 = most likely	Impact	Countermeasures
Assets under £500	• No record of lower value assets	• Theft / loss can go unnoticed • No way to audit these assets	4	• Financial losses over time • Low or medium to high when aggregated	Extend inventory control system to all assets, search employees on leaving factory at end of day
Information or data assets	• No asset register for information or data	• Unauthorised disclosure or loss • Inadequate protection of sensitive assets	4	• Potential reputational damage • Valuable documents are lost, such as IP or confidential personal or financial information	Maintain information asset register listing all information assets and their confidentiality, integrity and availability ratings
Information assets	• Access rights not corresponding with job role	• Lack of segregation of duties • Inappropriate access privileges	5	• Abuse of privileges leading to loss of information, fraud, malicious activity	Joiners, movers and leavers process implemented and regularly reviewed
Staff records	• No single view of HR information • HR view that CIO is responsible	• Ghost employees, failure to remove access right for leavers • Inappropriate elevated privileges	5	• Fraud, abuse of privileges	Implement a single HR database to provide a single view of staff on payroll, what assets they can view and position held
Media library	• Unnecessary access to data held there	• Theft / loss • Disclosure of data to unauthorised personnel	2	• Inability to perform appropriate backups • Data disclosure could lead to reputational damage	Limit access solely to those needing it for legitimate work purposes

Security breaches, what they say about the security culture and countermeasures

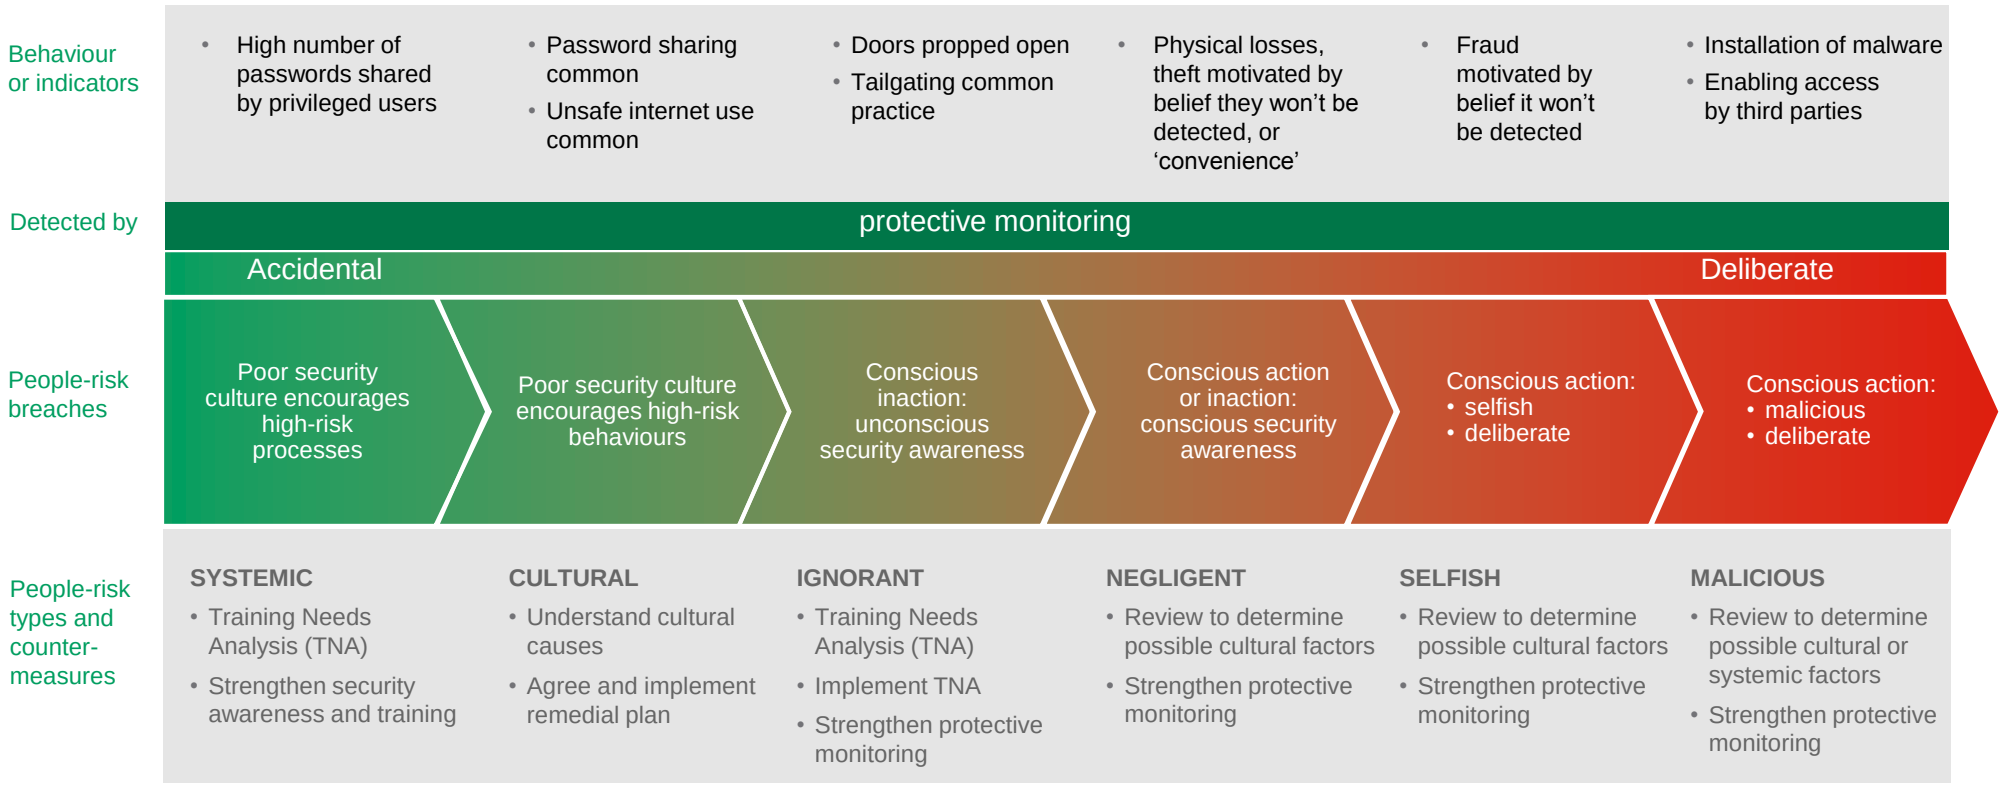

Legal requirements

The countermeasures which your organisation decides to implement may include some form of protective monitoring. All monitoring is subject to legal and regulatory controls. Your approach should be lawful and transparent. Transparency can be instrumental in allaying suspicion and gaining acceptance and support from regulators, employees, unions and employee representatives.

This guidance does not provide a definitive list of laws which organisations must obey but the main UK legislation is listed on the following page.

If your organisation is carrying out any form of monitoring in the workplace in the UK, you should obtain legal advice and consult relevant websites such as the Information Commissioner's Office (ICO) to ensure that monitoring is in line with legislation and guidance. If your organisation operates outside the UK, you should seek legal advice on the requirements that apply in the relevant countries.

More information on relevant legislation

NPSA provides a list of relevant legislation when implementing monitoring in On-going personnel security – NPSA Website

The Information Commissioner's Office also provides guidance on legislation

Relevant UK legislation and codes of practice (August 2012)

Data Protection Act 1998 (DPA): almost all forms of monitoring will involve the collection of personal data. The DPA places responsibilities on organisations to ensure that such personal data is collected lawfully and processed in a fair and proper way.

Human Rights Act 1998: article 8 of the Act provides for the right to respect for private and family life. Individuals' Article 8 rights extend to the workplace.

Employment Practices Data Protection Code: the Code is issued by the Information Commissioner and is intended to help employers comply with the DPA, and to encourage them to adopt good practice. Part 3 of the Code addresses monitoring in the workplace.

Telecommunications (Lawful Business Practice) (Interception of Communications) Regulations 2000: these regulations provide for certain circumstances in which intrusive techniques such as the interception of communications can be used in the business context.

CCTV code of practice (revised edition 2008): this code is issued by the Information Commissioner and helps organisations to comply with the law when using closed-circuit television to carry out monitoring.

Public Interest Disclosure Act (1998): explains the protection given to 'whistle-blowers' and defines the terms 'Qualifying Disclosures' (the process of making disclosures about malpractice), and 'Protected Disclosures' (disclosing to the right person and in the right way).

Regulation of Investigatory Powers Act 2000 (RIPA): RIPA regulates the use of intrusive surveillance and investigation techniques, including the interception of communications.

The legal and regulatory frameworks are dynamic so organisations should obtain legal advice.

Making the business case

A basic risk assessment will include the risk, its impact and the probability of the risk occurring. A simple way to assess the cost effectiveness of the proposed control is to consider the revenue costs or the cost of the man hours required to respond and recover to the incident should it occur. If we assume that probability will sit somewhere along the scale of 0 to 1, then multiplying the cost by the probability gives the weighted risk, which can then be compared to the cost of the control.

However, data needed to calculate the probability of and impact from people risk incidents are often subjective and not widely shared. This may be because organisations which suffer incidents fear a loss of employee, customer, partner and shareholder confidence if an incident becomes public.

Return on investment (ROI) should not be the only consideration. Research shows that the common corporate hierarchy of decision making, in order of importance, is:

- Safety
- Compliance
- ROI.

Senior leaders and executives are responsible for making safety and compliance investment decisions. They are also responsible for setting their organisation's corporate culture.

Senior leadership needs to be clear about the main reasons for investing in mitigating against people risks in the context of their business. Do not be distracted by assessments of ROI since these tend to be least useful.

Protect

This section provides guidance on protective countermeasures to people risk. It focuses on the methods for collecting, aggregating and analysing people-related data in a transparent, ethical and legal manner.

Protective monitoring is one of the most powerful tools for mitigating against people risk. While it can provide insight into many aspects of the organisation, its focus in HoMER is how to support the management of people risk.

Before deciding which measures to take to protect against the risk from an employee you should use an appropriate risk framework to:

- understand valued assets
- understand who has access to valued assets
- identify what risks those with access could pose to those assets.

Once your organisation understands these risks it can make informed decisions on how to protect itself.

Section three: key points

A holistic approach

- Simple controls can prevent and deter people risk.

Access controls

- Every role should have a clear list of access privileges.
- Additional privileges should be controlled and when employees change status, access privileges should be reviewed.

Identity and access management

- Have a clear, up-to-date master database of employees and minimise the number of databases.
- Changes in the master database of employees must be reflected promptly on secondary databases.
- Have a clearly defined policy and process for verifying identity.

Protective monitoring

- Protective monitoring does not necessarily require new investment.
- Start by improving existing processes.
- Involve HR in identifying people risk.

An ethical, legal and transparent protective monitoring capability

- Sensitivities.
- Appropriate handling of data.

A roadmap to a basic capability

A proactive protective monitoring capability

- The principles of protective monitoring.
- Identify incidents that might occur instead of simply reacting to incidents when they have occurred.
- Analyse past data to identify incidents in the future.

A holistic approach

Many incidents involving people risk can be identified in advance and prevented if simple processes and procedures are in place and the organisation is alert to the indicators. This section provides guidance on reducing people risk through simple processes such as access controls and identity management.

Many organisations already apply some form of monitoring. This is usually network monitoring of technical systems. While network monitoring has a value it rarely reveals the whole picture, often increases false positives and, on its own, may result in important events being missed.

A holistic approach which takes input from various functions such as IT, HR and Physical Security rather than a silo approach where different parts of the organisation manage their own risks, means that your organisation has a more complete picture of risk, can identify important events and decrease the number of false positives.

By following the guidance in this chapter your organisation can reduce the opportunity for employees to harm it.

Access controls

Access should be controlled rigorously to minimise the risk from counter-productive workplace behaviour. Posts carry different levels of access and therefore different levels of risk. To assess the risk associated with a post:

- consider the value of the asset to which the holder has access
- identify the vulnerability of the asset
- assess the likelihood of an incident and its impact.

Of course, some members of staff need privileged access (e.g. HR or IT systems administrators). Such access makes monitoring all the more important. To reduce the risk:

- every role should have a clear list of access privileges. Additional privileges should be controlled by a clear policy
- when employees change roles, access privileges should be reviewed
- review employees' access privileges when they tender their resignation
- ensure everyone knows that the organisation reserves the right (at any time) to change or restrict access or authority on granting privileged access.

Identity and access management

Multiple databases often exist to manage access to different assets, and frequently there is no synchronisation across them and no master database. This can lead to inappropriate access being granted when an employee changes role or residual access to organisation-sensitive information/physical space when an employee leaves.

Your organisation can mitigate against inappropriate access, reduce administration costs and manage employee data more effectively by implementing the following controls:

- have a master database of employees, usually the payroll database, which is the most likely to be kept up to date
- minimise the number of databases. Having too many can lead to cumbersome and expensive business processes and increased risk
- changes in the master database should be reflected in the secondary databases. If an employee is leaving the organisation then there will be a notice to payroll to stop payment; this should be reflected in the access databases
- have a clearly defined policy and process for verifying identity upon joining the organisation and when moving within the organisation.

Case history 4: Access management

In April 2009, Mr A, an auditor at a US utility organisation resigned. That night he used his (still-active) electronic key card to gain access to the company's secure facilities.

Having previously gained the access codes, he used the computers of two senior executives to make and then approve three electronic transfers totalling $9million to an account in Qatar before fleeing to Canada.

The circumstances of the attack indicate that normal security measures were in place at the utility organisation. The site required swipe access and, in order to make and approve transactions, Mr A needed to access two separate machines situated in different buildings. Despite this Mr A, as an employee, was able to circumvent these measures.

Which roles have privileged access in your organisation?

How does your organisation control addition and removal of access? How long does this process take?

How does your organisation manage the increased risk that roles requiring privileged access present?

Protective monitoring

In this section we consider the use of protective monitoring to help identify the indicators of people risk. Protective monitoring in this context is defined as the collection of information for investigation and in some cases, compliance purposes. It can be for reactive or proactive risk-based investigation.

This approach must have the informed support of its employees under appropriate corporate governance. This does not mean investing new money in expensive tools; it means having the right people, policies, processes and structures in place to enable this capability to be effective.

The holistic approach

Many organisations carry out monitoring but it is usually performed in silos and thus misses indicators of people risk. By taking a holistic approach, (i.e. taking information from different parts of the business - IT, HR, line management, physical security and so on) organisations can gain a more complete picture, identify significant events and decrease the number of 'false positives'.

Of course, input from HR can include highly sensitive information which requires careful handling and a transparent and ethical approach.
For this reason, it is worth considering appointing the HR Director to be the senior single accountable owner of people risk.

HR practice and line management

Implementing good HR practice and effective line management are essential in mitigating people risk. Traditionally many aspects of identifying and managing people risk are left to line management and HR. Much of the time this works well because most people do not pose exceptional risk. However, in the case of those who do, leaving shared responsibility with line management and the HR function means significant indicators of increased risk are more likely to be missed. In taking a more holistic approach your organisation is more likely to identify threats early.

A holistic approach to protective monitoring

Research from the US shows that most insider acts come from employees who, until only 3 - 4 weeks previously, had been perfectly loyal and committed. The change from a trustworthy employee into someone who feels disgruntled and motivated to damage the organisation can be triggered by a negative work place experience which alters the employee's expectations of the organisation. At that stage the employee often displays social cues of dissatisfaction which are visible to their colleagues and manager. The employee may also behave differently in terms of the information, systems and sites or assets they access and retrieve.

Looked at in isolation these anomalies could be signs of factors unrelated to the organisation such as issues in the employee's personal life. However they may be signs of counter-productive workplace behaviour. Too often these are overlooked until it is too late. Taking a holistic approach to the collecting, aggregating and analysing of behavioural data can help prevent an organisation from reaching an incorrect conclusion, avoid the risk of unlawful discrimination or harassment and allow for the early identification of undesirable behaviours.

For further information on Insiders see www.npsa.gov.uk

Common indicators for people risk

- **Change of working pattern**
- **Conflicts at work**
- **Decline in performance**
- **Drug or alcohol abuse**
- **Aggressive behaviour**
- **Mood swings**
- **Missed promotions**
- **Debt**
- **Unexplained wealth**

An ethical, legal and transparent protective monitoring capability

Sensitivities of protective monitoring

Almost all forms of monitoring will involve the collection of personal data. The Data Protection Act places responsibilities on organisations to ensure that such personal data is collected lawfully and processed in a fair and proper way.

In an office context, it is intrusive for an organisation to routinely access 'reasons for absence' for protective monitoring. However, in a hospital context, where the 'reasons for absence' could have a direct impact on an employee's ability to provide care, it may be decided that such data should be routinely captured and monitored. Another example where protective monitoring will require routine access to personal data is when it is important to know about specific medical conditions which may affect an employee's ability to operate machinery or drive.

The subjective and sensitive nature of this topic means that there is no right answer: each decision needs careful consideration.

Appropriate handling of protective monitoring data

An organisation's protective monitoring team will have access to highly sensitive information including people-related data and substantive organisational data. This means that the protective monitoring team has some of the most privileged access in an organisation and must be trusted to a degree associated with this access.

Information held by the protective monitoring team will provide exceptional information about the workings of the organisation. When aggregated, this data may be among the organisation's most valuable assets and will need appropriate protection.

It is essential that all personal data held by the protective monitoring team are respected and properly protected in accordance with legislation.

Implementing an ethical, legal and transparent protective monitoring capability

Monitoring should be both transparent and unpredictable

Employees being monitored should be aware of and consent to, the monitoring. The process of informing them should be as transparent as possible: policy is a good mechanism for doing this. Other mechanisms could include a prompt on a log-in screen. This example has the benefit of being both a frequent reminder of monitoring and a means of obtaining consent to it.

For monitoring to be transparent, employees need clear guidance on the organisation's privacy policy. This must include, but not be limited to, what is and is not acceptable behaviour, from the use of corporate IT and social media to information handling policies.

Due to the dynamic nature of the law and the employment relationship, organisations should review protective monitoring policies regularly to ensure they are up to date with good commercial practice and legally compliant.

It is essential to ensure transparency about the existence of protective monitoring but details of the monitoring should not be disclosed as this could allow it to be circumvented. To be effective the capability needs to be unpredictable.

As an example, protective monitoring can be used to verify the creation of system administrator accounts. This can provide an early indicator of people risk as it is a common indicator of a potential threat. A monthly and random review of new accounts can reduce this risk.

Risks and fraudulent behaviour are often exposed by fellow employees or workers. It is important to have a whistle-blowing policy that has the confidence of all employees. This will underpin any other protective monitoring systems and policies which are in place.

Top tip: Monitoring notification

Avoid using a repetitive, negative or bureaucratic approach to letting employees know that monitoring is taking place. Find a constructive way which supports a positive security culture.

Principles of protective monitoring

- The protective monitoring team may not define its own remit or go beyond the remit defined by the organisation. This prevents the team members ‘snooping’ or conducting ‘fishing trips’.
- Only designated individuals may task the protective monitoring team. This will typically be specific staff from Legal and HR. Others wishing to task the team must do so through designated individuals.
- The people in the protective monitoring team are in roles that require a very high level of trust from the organisation. A risk assessment should identify these roles as high risk and suitable countermeasures should be put in place (e.g. extra screening).
- The protective monitoring team may only undertake tasks which are within their policy-defined limits. They are permitted to report, but not to judge or intervene themselves.
- The protective monitoring team may not be aware of additional information that may change the significance of a piece of apparently high risk behaviour.
- The data used in protective monitoring needs to be processed and protected to maintain integrity. Also, since the majority of high-risk anomalies are likely to be innocent, visibility and knowledge of such anomalies needs to be restricted.
- Application accounting logs should be captured as soon as possible after they have been created and immediately passed through a one-way gate to a dedicated protected database. A separate copy of the logs should be saved to verify the integrity of the log data used for protective monitoring. The separate copies also ensure that the data retain evidential quality if needed to support disciplinary action against an employee.
- For data protection compliance purposes organisations should minimise data processing. The protective monitoring team's access to personal data should be restricted to the lowest level possible proportionate with the successful performance of the task.
- In larger organisations network monitoring should be carried out by a different team from the one monitoring people risk. The two teams are looking for different types of event and each type of monitoring is more effective when carried out by a dedicated team.

Roadmap to a basic protective monitoring capability

Some organisations may only need a very low level of capability and be able to implement holistic employee risk management gradually by following the roadmap below:

Base policies on transparency	Appoint a single accountable owner of people risk	Train and develop good managers	Ensure effective access controls	Apply the three prerequisites	Appoint a trusted protective monitoring team	Collect, aggregate and analyse HoMER data	Apply a policy-based communications plan
• Be transparent • Develop and apply transparent policies • Gain informed consent to protective monitoring • Gain employee support	• Ensure a joined-up approach among all risk owners	• Line managers • HR • Corporate	• Physical • Information • Privileged • Other	• Manage assets • Manage identity and access • Manage risks	• Control responses through policy and management	• Access sensitive data only in accordance with policy	• Training and awareness programmes • Brief on risks and cases, appropriately protected • Deter opportunistic crime

A proactive protective monitoring capability

Most organisations carry out protective monitoring reactively, i.e. they react to an incident after it has taken place by looking at the logs for information.

This approach has inherent limitations because the event has already taken place. Perhaps more significantly, it fails to recognise the value of adopting a data-based, proactive approach that enables early identification of, and mitigation against risks.

One reason why organisations take a reactive approach is the perception that proactive monitoring is associated with snooping. Nevertheless there is a clear and valuable place for proactive monitoring to identify and mitigate against new risks

Build a proactive monitoring capability

Identify and weigh your assets according to value or sensitivity to provide a starting point for monitoring what happens to them, when, how, and by whom.

Identify, analyse and split out your risks to show where the key people risks lie, for example:

- does the latest analysis of risk to assets from malware identify where technical risk is also a people risk?
- spear-phishing is an email fraud attempt that seeks unauthorised access to confidential data. Even people in the most technically sophisticated organisations fall victim to it. How recently has your organisation reminded employees of the threat from spear-phishing in a way that brings it to life? How regularly do you use penetration testing by an independent professional firm to strengthen your defences, your assurance and reduce your people risk?
- tailgating can provide the risk of access to an organisation's assets, such as server rooms or other controlled access rooms.
- apply an informed, risk-based and disciplined approach that controls individual access to any asset at any time.

A risk-based, policy-led protective monitoring capability

To be both proportionate and effective, protective monitoring should focus on risks, not individuals. You can start to build a proactive protective monitoring capability for your organisation by asking the following questions:

- What events have affected this asset, over a given time period, and in what way?
- Is there anything unusual or apparently anomalous about these events?
- If so, how do these unusual or apparently anomalous events compare with a previous period or a different asset?
- What volume of anomalies is associated with this application/ action/asset/environmental event? What questions does this information trigger?

The above questions focus on the organisation's assets, not its people, however the answers to the questions will help identify anomalous behaviour by individual employees.

Research in the USA has shown that employees are more likely to commit acts such as intellectual property theft when leaving an organisation. By applying a simple risk-based 'flagging system' which does not target individuals, organisations can implement ethical yet effective monitoring. The 'flag' designates employee status, e.g. 'in employment', 'resigned'. 'leaving' or 'retired'. This allows the organisation to implement a policy whereby anyone flagged as 'leaving' may be subject to a random review covering the previous 30/60/90 days. This can prevent or deter individuals from stealing intellectual property prior to leaving an organisation.

Review past events to predict and prevent future incidents

The principle is simple and risk-based. By reviewing retrospective behaviour, for example answering the four basic questions posed in the previous section (against time, volume, access and employee status), your organisation can inform and focus the monitoring to help predict increased risks in future.

The benefits include:

- Demonstrably risk-focused protective monitoring.
- Reduced protective monitoring of irrelevant subjects.
- Detecting riskier behaviour and responding appropriately.
- Deterrence of ‘opportunistic crime’.
- A strengthening of the security culture.

Case history 5: Identifying high-risk behaviour

Mr X, a sales representative, found his remuneration package steadily decreasing, due in part to his being a 'chronically unhappy employee'. Upset by his reduced status within the company, Mr X told a colleague that if he left the organisation, he would take all his business accounts and 20% of the company's other accounts with him. He accepted an offer of employment with a new company, but falsely informed his supervisor that he had turned the position down. Over a period of one month, while waiting for the new job to start, Mr X forwarded more than 100 confidential email messages to his home computer. Several of these messages contained passwords for accessing secure areas of his employer's network and proprietary information on a computer program that he had helped build (and which he believed he rightfully owned).

The passwords allowed him to view confidential customer information to which he did not have access during his normal course of duties. His goal was to take this information to his new employer. Mr X deleted the contents of his hard drive to cover his tracks and then submitted his resignation.

What indicators of high-risk behaviour could have been observed?

Does your organisation have the capability to identify these types of behaviour?

What information would be required by a protective monitoring team in your organisation? Is there an easy way to access this information?

Relevant data and integration

In most organisations, digital data relevant to people risk can be found in various software application logs which record employees' actions. These are useful for showing the normal working patterns.

Useful applications include:

- physical entry/exit logs, with a focus on time and access to physical spaces
- log-on, log-off logs, with a focus on time and user credentials
- email application logs
- database application logs.

This guidance does not aim to prescribe which data is most useful as most is of some use. Access to the data improves understanding of 'normal' and 'abnormal' behaviour. Linking the different sources and anomalies builds a comprehensive risk picture which you can act upon.

The diagram illustrates this by setting normal behaviour inside the boundaries. The red dots illustrate anomalies which occur outside the normal boundaries. These feeds are then integrated and, when merged, illustrate that the combined feed presents a higher risk profile warranting further investigation.

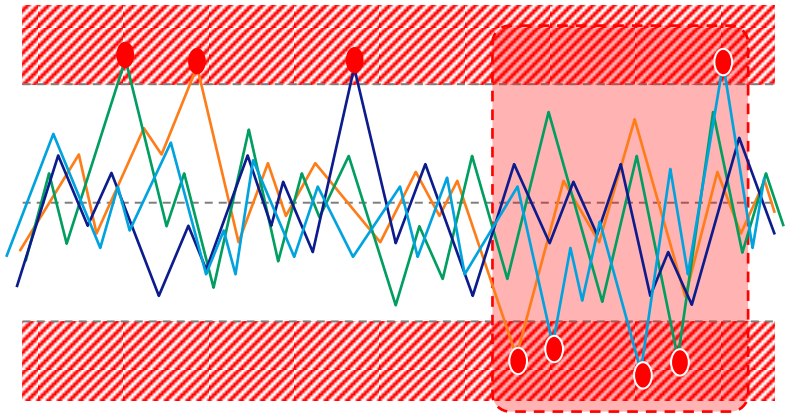

Integration of protective monitoring to ensure a holistic approach

The diagram illustrates the flow of information feeds from various parts of the business to the single point of accountability.

These feeds deliver data on observed anomalies. These are then collated at the level of the tasking team. The tasking team is a group of designated individuals (most likely from the legal and HR functions). All tasking for the protective monitoring team should be done through the tasking team.

In larger organisations there may be another team that collates the feeds and performs analysis to identify indicators of people risk. Risk is likely to be reviewed by a larger number of stakeholders. In smaller organisations this might take the form of a meeting between HR, legal and security, chaired by the single accountable owner of people risk.

These meetings should examine anomalies raised by the different protective monitoring teams, cases of concern and possible future action. Future action could include more protective monitoring to investigate cases of concern or refinements of protective monitoring criteria.

In smaller organisations where there are fewer feeds, less data and one or two people cover all the roles, the same principles should be applied.

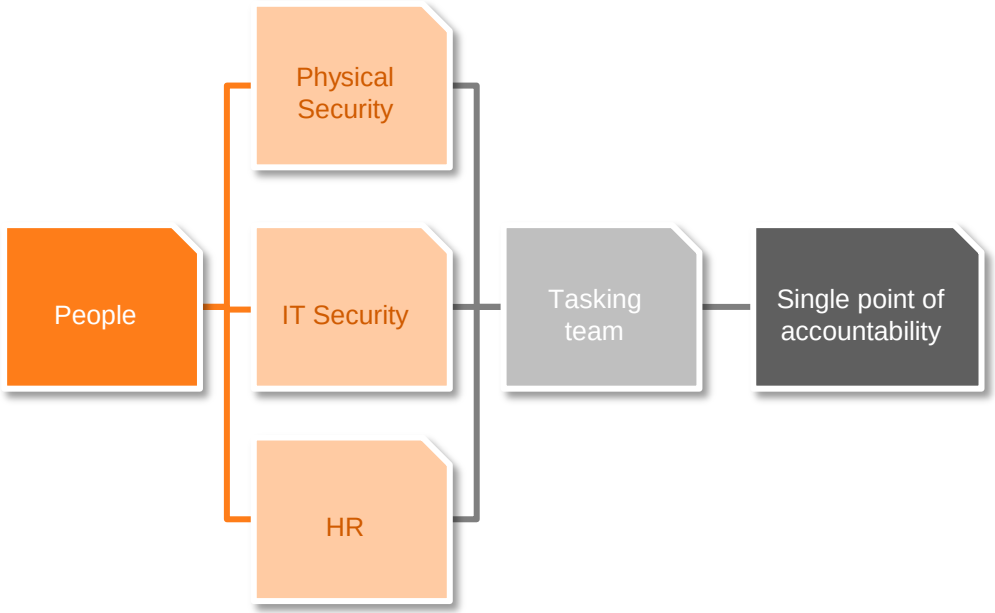

Building an integrated protective monitoring capability

Next steps

To build an integrated protective monitoring capability for the holistic management of employee risk, there are four simple steps an organisation can take:

- encourage people to do the right thing
- monitor normal and unusual behaviour
- determine the meaning of anomalous indicators
- respond appropriately.

These steps are expanded on the following page.

Protective monitoring business process cycle

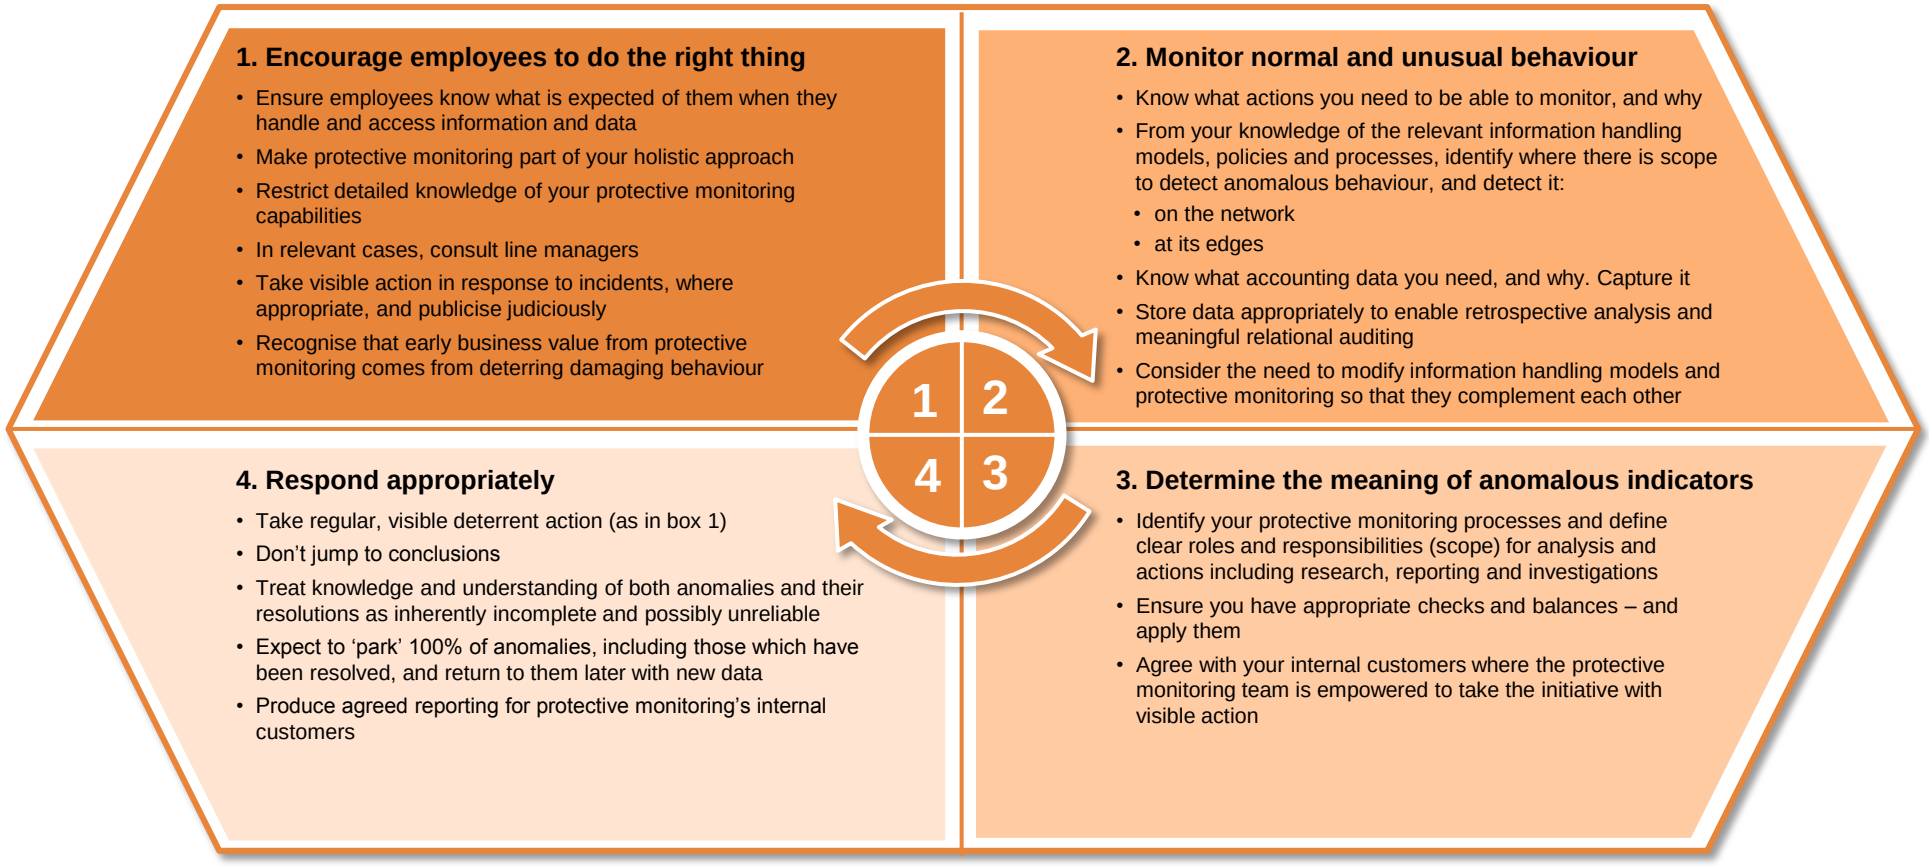

Respond

This section provides guidance on how to respond to potential and actual incidents by setting out the immediate steps that need to be followed prior to any possible investigation.

It covers the principles and policies that should be in place before an incident occurs and outlines the actions that your organisation should take in response to an incident.

It then focuses on the operational response, detailing the range of events that can lead to an incident and suggesting appropriate responses. This is because it is at an operational level that an organisation will typically discover or confirm high-risk behaviours and have to deal with them.

Incidents involving people risk often pose policy and governance challenges. These challenges are easier to address and resolve if an organisation has developed a contingency plan.

Section four: key points

Identifying and responding to incidents

- The importance of pre-planning.
- Types of incident.

Protective monitoring triggers

- Protective monitoring and risk management should allow an organisation to respond to indicators before an incident occurs.
- Most conclusions on employee behaviour are conditional and only partially informed.
- Decide which data feeds should be monitored routinely and which require a risk threshold for access.

Considerations before responding

- Choose appropriate responses on a case-by-case basis.

Response policy

- Act in accordance with the law, regulatory requirements and relevant codes of practice as well as the values and traditions of the organisation.
- Ensure your response is controlled, measured, proportionate, justifiable, trusted, timely, disciplined and informed.

Managing the response

- Have a clear plan that includes responsibilities and the chain of command for incident response.
- Only intervene if you are likely to achieve something useful.
- Draw up policy-based rules of engagement.

Operational response states

- Understand what is normal in order to be able to identify the abnormal.

Communication during an incident

- Actions need to be coordinated and effectively communicated
- Restrict knowledge of an event and your response to those who need access.
- Take care in communicating an event. Ensure the details disclosed meet your obligations to protect personal data and respect privacy.

Identifying and responding to incidents

An incident is an event which causes damage or has the potential to cause damage. It can often be difficult to define exactly.

Types of incident

When defining a major and minor incident, consider:

- damage to a core asset
- damage to reputation, affecting market, partner and shareholder confidence
- damage to employee morale and confidence
- time required to recover from the incident
- potential media interest in the incident.

From the perspective of data protection regulation, the seriousness of an incident is judged by the effect on the individuals concerned or harmed and the number of people affected.

The importance of pre-planning

In the event of an incident your organisation will be better placed to respond if pre-planning has taken place.

Pre-planning involves:

- setting out the policy, roles and responsibilities for managing a response
- testing and reviewing the response.

This approach is in line with data protection law and the law of negligence which require organisations to have contingency plans in place for security incidents.

Managing incident response effectively may well enhance your organisation's reputation.

Protective monitoring triggers

A proactive approach to protective monitoring and managing risk should allow your organisation to respond to indicators before an incident occurs.

Your organisation should carry out a risk assessment to identify the threshold that will trigger a response from the protective monitoring team.

Three types of trigger could merit a response or further investigation:

An event: an occurrence recorded by monitoring systems, such as an entry log. This will most often trigger no intervention.

An anomaly: an occurrence, such as an entry log, with a feature such as time/volume indicating that the event is different from 'normal'. In most cases an anomaly does not indicate a security threat and should be viewed with the presumption of innocence. It will usually trigger 'no active response'. In some cases, however, it will prove to be significant, especially when correlated with other relevant data feeds. In these cases action will be required.

An incident: typically an occurrence with immediate and clear adverse security implications such as an attempt to gain and exploit unauthorised access.

Case history 6: Line management concerns

Example 1: A line manager, Ms E, has concerns about the punctuality of Mr L and thinks that it might be a sign of poor performance. Ms E does not want to upset Mr L so requests sight of the access control logs showing what times of day Mr L arrives for work. Ms E also asks for CCTV footage to prove that Mr L was working at his desk during his time in the building.

Example 2: Another line manager, Mr H, suspects that Mr O has shared his login details with a colleague and that the system is being accessed while Mr O is not at work. Mr H asks to see CCTV coverage of Mr O's desk to see whether he is at his desk when the system is being accessed with his login.

Is access justified for each of the line managers?

How should the protective monitoring team respond?

Considerations before responding

Employees who cause incidents may not be malicious in intent. They could be the victims of manipulation by others (e.g. social engineering), they may be ignorant of the organisation's security policies, or they could be bending rules 'to get the job done'.

Even in the case of a disaffected employee causing deliberate damage, dismissing the person without subsequently understanding how s/he came to be disaffected and was allowed to continue with that attitude uncorrected is not good practice.

When considering appropriate action to take (up to and including dismissal) your organisation should take account of any investigatory and disciplinary procedures and be mindful of employment legislation which gives protection to employees against dismissal.

Investigate causes and motives

To understand what led to an event or incident, you will need to identify cause (direct and indirect) and motive. What you find will have a bearing on your response.

Human error: training and other support, such as supervision and coaching will normally resolve this category.

Example: an employee rushing to complete work to a tight deadline mistakenly sends an email to an external addressee who should not receive it. If spotted quickly, this can normally be easily remedied. Ideally the individual spots their own mistake, reports it and takes action to obtain the recipient's assurance that the email has been immediately deleted.

Negligence: this implies a level of carelessness where disciplinary action will normally be appropriate.

Example: an employee leaves a briefcase containing sensitive material in a public place allowing unauthorised access and potential damage.

Deliberate act: theft and sabotage are criminal offences and will require careful, legally aware and informed operational decision making and handling. Even the initial interview of the individuals involved may need to be done by the police since, if the appropriate steps are not followed, statements may be inadmissible, thereby jeopardising future legal proceedings. All evidence must be handled carefully in a way that maintains its integrity.

Response policy

Your organisation should have an incident management plan for major incidents, including those that involve people risk. An effective incident management plan must be supported by a transparent and accessible response policy. It is important that your organisation follows its own policies and procedures (which may include disciplinary procedures, suspension and investigation procedures).

For all organisations the approach to policy response needs to be holistic and should address:

- **Bringing information together to enable a holistic view.** The specifics will vary depending on the organisation. In smaller organisations it may simply be a meeting between the head of HR and the head of security, where the latter is likely to manage the protective monitoring team. In large organisations this might be a regular meeting of the various functions that are involved in managing risk to review cases of concern.
- **Actions to be taken if the actual or potential damage from the incident could affect the organisation's future value, including reputation**. This should include preparing an internal and external communications plan. All public-sector organisations' policies will need to fulfil any obligations to report high-impact incidents to the Cabinet Office. Commercial organisations should consider whether they have a legal obligation to report matters to the relevant policy or regulatory authority, e.g. the Information Commissioner's Office or the Financial Services Authority.
- **Actions to be taken if data belonging to others, such as customer credit card details, have been or may be affected**. You may want to inform those who may be affected early to ensure they are kept updated and reassured by your handling of the incident.

The accountable officer will have to decide on the appropriate action, including considering whether the police should be informed with a view to taking criminal action. Actions must be consistent with disciplinary policies which often contain rules relating to suspension as well as investigation procedures.

Response policy

Each organisation's response plan to incidents will be different and will be shaped by the risks identified during the Assess stage. It may be necessary to include public relations and marketing in the response policy. If an incident occurs, how should these functions respond to external enquiries? It is quite likely that people outside the organisation will become aware of the incident.

Your organisation's policies should consider the role of line managers in dealing with incidents and potentially high-risk employees. Line managers should be able to exercise their judgement when responding to these challenges. Guidance, training and coaching should be available to help line managers make these judgements.

Guidance should cover a range of cases, such as:

- low-risk indicators which gradually build in frequency
- a sudden high-risk event that is a clear cause for concern.

In the first case, indicators will typically only be visible to the monitoring team. They will normally be set aside but may be reviewed periodically by the single accountable owner of people risk. In the second case, a sudden high risk event, for which there is no obvious explanation, may require an early consultation with the individual's line manager.

Managing the response

Your organisation should have a clear **incident response plan** which includes responsibilities and chain of command. It should set out what is required from senior leadership. Alternative responsible persons may need to be appointed for cases where those with responsibility for ensuring the correct running of an investigation are themselves implicated or even under investigation.

Ensure that there is a clear understanding of the stance your organisation will take in the event of an incident. Will the organisation cut off access to limit damage, or continue to allow access to build an actionable case, perhaps ensuring that forensic copies of logs are acquired? What takes priority? Damage limitation, damage assessment or evidence building? Without clear guidance, confusion in the response team could cause counterproductive behaviour and even result in a situation where the correct action cannot be taken. For example, evidence needed for internal discipline or referral to law enforcement agencies may not be properly gathered.

There might be misleading indicators and those to whom the protective monitoring team is accountable must be prepared to tolerate this rather than accusing the auditors/investigators of continually jumping to the wrong conclusions. Some indicators will eventually prove to be true but this may not become evident until new information is added.

Your organisation should draw up contingency plans which address all types of event and levels of impact, and test them so that it has confidence they will work. Testing serves as on-the-job training and is essential. Contingency plans should include actions for everyone who has a role to play in responding. For a minor event, this may be just the operational protective monitoring team and the board member accountable for people risk. A major incident could involve anyone from the board downwards.

The Information Commissioner and the Financial Services Authority regard the absence of contingency planning as an aggravating feature when they assess whether an organisation has complied with its legal obligations for security and confidentiality.

Guidance for incident response

Your organisation should draw up guidance setting out how everyone involved in responding to an incident should carry out their responsibilities. This guidance would cover the following points:

- Treat individuals as good and loyal employees. Most interventions will discover reassuringly positive findings. ‘Intervention’ means direct engagement with an employee whose behaviour is causing concern, or with their line manager.
- Control any communication coming from the protective monitoring team. Start by consulting within the protective monitoring team and management. The reason for controlling communication is that the protective monitoring team is dealing with sensitive employee data.
- Designate someone to consult with the wider organisation beyond the protective monitoring team in the event that this is needed. This person will normally be in a management or supervisory role on the protective monitoring team.
- Have a clear procedure for intervention including guidance on who should respond, with what response, to whom, how, when and why, with what potential outcomes in mind.
- Develop an interview strategy before you speak to an individual about potentially harmful behaviour that has come to your attention and consider what information you need to disclose. For example, if an anomaly has come to your attention through protective monitoring, you may find another source of correlation for the anomaly that would be better to use in your conversation. This would also have the benefit of protecting details of the specific nature of the monitoring.

Consult NPSA's Guide to on-going personnel security

Intervention must be carefully considered and controlled

Only intervene if you are likely to achieve something useful.

Keep the initial review anonymised as far as possible so that the reviewer does not know the identity of the person under review. If you need to identify a person after that point, because, for example, intervention is needed this can be done. However, a policy framework will need to describe the conditions under which removing anonymity is permitted.

Whilst all indicators from the data need to be examined, experience suggests that most will not give rise to concern. Since most anomalies observable in log data will not be indicators of a security threat but rather will indicate normal behaviour for the employee concerned, they should not generate any intervention. This is an important first principle and should help set the expectations and approach of the protective monitoring team and its management. The protective monitoring team needs to be open-minded, detailed and thorough.

Operational response states

To be effective, protective monitoring needs to provide sufficient coverage of behaviour to allow the protective monitoring team to be able to recognise what is normal. Once ‘normal’ is established and understood the team can recognise anomalies.

Since there are anomalies in everyone’s normal work pattern care is needed to ensure ‘normal anomalies’ do not trigger an inappropriate response. Such anomalies would not, in isolation, be expected to trigger an active response such as intervening with the line manager or employee. Rather, it should begin to form a baseline of variations that become the ‘norm’ for that employee.

Operational response states are outlined in the table on the following page.

For a large organisation, response state 1 or 2 will be normal. An obviously damaging incident will trip the state into response state 4. However, there will also be examples of a gradual increase of anomalies and apparently adverse events that should trigger a general proactive response (for example raise security awareness in an area) or a focused proactive response (for example an investigation to determine whether an employee is causing harm to the organisation).

Protective monitoring teams will observe a range of activity including ‘no activity’.

Operational response states

Response state	Action
Response state 0 No adverse events detected	• Protect the discovery – even state 0 is sensitive since the state observed through the data may not reflect reality • Expect this response state – this will be normal for most people most of the time • Train the monitoring team • The protective monitoring team needs to continue to believe there are adverse events 'out there' • Review historic data – for example, on a sampling basis covering high-risk assets, looking for anomalies
Response state 1 Anomaly or tip-off	• Protect the discovery • Review the data • Due diligence – look for meaningful correlations • Park – where the meaning is unclear
Response state 2 Adverse occurrence	• Protect the discovery • Review the data • Due diligence – look for meaningful correlations with other relevant data feeds • Report to the tasking team – where the data justifies escalation and possible subsequent intervention is required. • Park – where on examination there is no reason to escalate to customers who may be affected
Response state 3 A major incident	• Protect the discovery – if it is not yet public knowledge and until a decision is reached on how it will be communicated • Report to the single accountable owner • Requires immediate action (if you have time, produce an actionable implementable plan within a useful timeframe) • Requires deliberate action over more time (more considered approach, involving wider consultation and deliberation) • May require a combination of immediate and deliberate action
Response state 4 After an incident	• Capture immediate learning points • Identify improvements • Implement improvements

Communication during an incident

Developments during an on-going incident need to be communicated appropriately, therefore you should devise and implement a communications policy and plan.

The policy does not need to be long or complicated. It should simply put in place appropriate guidance and processes.

An example of a simple communications plan:

Stakeholder	Message	Last contacted	Next contact	Action
Senior Management				
Employees				
Customers				
Regulators				
Media				

Recover

This section shows how your organisation can recover from an incident that involves people risk, ensuring capability is restored and the organisation strengthened.

It outlines principles regarding investigation, dealing with the findings of an investigation and how an organisation might learn from an incident and evolve while communicating its actions effectively. Lessons learned from an incident should feed into future risk assessments.

Section five: key points

Investigate

- Appoint someone to lead the investigation.
- Be impartial, regardless of who the instructing party is.
- Identify all sources.
- Seek evidence from a range of different sources.
- Protect original information or data once it is gathered.
- Ensure findings are verified.
- Tailor for the intended audience.
- Maintain an audit trail.
- Carry out regular risk assessments.

Actions to take

- Immediate actions.
- Longer-term changes.
- Consider whether internal or external action is appropriate.

Learn

- Learn lessons from each incident, both those generated within the organisation and incidents with external sources.

Evolve

- Embed the lessons you learn into future risk assessments.
- Anticipate technological and regulatory trends.

Communicate

- The importance of a communications plan.

Emerge stronger

- Feed ‘lessons learned’ into the risk assessment process to improve the holistic management of people risk.

Investigate

A number of factors will influence how an investigation of an incident is carried out. Your organisation may have its own internal investigator or investigative team. There are third parties who provide investigative services or, where there is a probable infringement of the law, a law enforcement agency may be involved.

The following principles should be applied:

- Appoint someone to lead the investigation. This person will be responsible for the findings and ensuring that the investigation is carried out in accordance with the law.
- Be impartial, regardless of who the instructing party is. The investigation must always be carried out independently and show no bias.
- Interview relevant people, where appropriate, to identify all sources.
- Evidence should be kept secure in case it is necessary for legal proceedings.
- Seek evidence from a range of different sources to ensure you produce the best picture of events.
- Original information or data should only be accessed by trained individuals. They must handle the data in a way that maintains its evidential integrity.
- Information gathered as part of the investigation should not be changed.
- Ensure your findings are verified.
- Ensure that the output from the findings and recommendations is relevant for the intended audience.
- Maintain an audit trail of the activities carried out in the investigation.
- Carry out regular risk assessments. Due to the sensitive nature of an investigation it is important to be aware of and manage risk.

Actions to take

Immediate actions

The incident might have revealed specific weaknesses in defences such as physical access controls, data and systems access controls and protective monitoring. These types of weakness are normally the easiest to address because they are concrete, often the least expensive to put right and it is easy to measure when they have been corrected.

Longer-term changes

Simple examples include new data feeds to address identified gaps, additional resources or new monitoring capabilities and tools. These changes might not only require new investment but also training, up-skilling and, in the case of new data feeds, negotiation with the data asset owners to obtain the data, whilst ensuring that you are legally compliant. The cost and effort will need to be balanced against the risk mitigated by taking the action.

The human factors that contributed to the incident can be more challenging to address. One contributing factor might be a weak security culture. Another might be an ineffective response culture or team that ignored protective monitoring reports that detected the incident precursors and failed to take timely action that could have prevented the incident.

Trust and confidence

Immediately following an incident people's confidence in the organisation might be shaken. Unless there has been prompt, effective communication, people may have learned about the incident through gossip or the media. Communications tailored to the audience should reassure and show that the organisation is being transparent and is in control both of the incident and of the response.

Agree a problem statement

A problem statement sets out the gap between the desired state and the actual state, and is based on the facts established by the incident investigation. **Its aim is to help the investigator to go beyond accepting a tactical explanation of an incident. For example, a briefcase of documents left on a train could be explained by 'this was just an accident'. Or, the investigator could look deeper for underlying causes which could include a failing in policy, process or culture. A clear problem statement can help an organisation understand and address any significant underlying causes.**

The investigator should take the facts and analyse them for causes, with a view to determining the problems. This needs to be done transparently and thoroughly. The investigator should take care not to rush to judgement and remain alert for non-obvious, underlying causes. Once the investigator has drawn up the problem statement, the senior owner of people risk should certify that s/he is happy with the investigation, the analysis and the problem statement (since the senior owner of risk is ultimately accountable for ensuring that the necessary recovery action is identified and taken).

Actions to take

Decide what to do about it

The senior owner of people risk should retain supervisory accountability for implementation which will usually be delegated to functional managers.

To decide on appropriate recovery actions, the senior owner of people risk may need to consult with others, such as lawyers who may be advising on legal issues or HR who may be advising on employee issues.

Decisions will usually be made on a case-by-case basis, but the two main options are to deal with an incident internally or externally.

Internal

Your organisation should still consider whether to inform regulatory bodies at an early stage even if you are dealing with the incident internally. This method ensures that the details of an incident can be controlled and is usually more timely and cost effective. It is essential to ensure that any action involves Legal and HR and is performed within the legal framework set out under UK employment, privacy and data protection law for dealing with investigations, disciplinary and grievances. You should also consider whether any whistleblowing legislation applies.

It is important that any action is proportionate and appropriate. Even in cases such as intellectual property theft where the stolen asset is recovered it may still require additional action. In many cases taking internal action can lead to disciplinary proceedings including dismissal.

External

External action is likely to involve a law enforcement agency, regulatory body and possible prosecution. Your organisation must decide whether external action is a viable option before undertaking an investigation. This is because there are specific requirements for maintaining the forensic integrity of evidence. The exact process for external action or prosecution differs across the UK and internationally. Therefore, before pursuing this route, your organisation should ask an appropriately qualified legal advisor to identify the legal implications and ensure that the legal process is properly managed as external action can be costly, time consuming and may not return the desired result.

Top tip – Disciplinary action

Guidance on good practice and guidance on discipline and grievances can be found on the Chartered Institute of Personnel and Development website.

Case history 7: Recovering from an incident

One night, Mr S, who worked for a large telecommunications company embarked on a series of acts that he believed would enable him to 'save the day' and impress his new supervisor. First, Mr S used a contractor's badge to gain unauthorised physical access to the company's Network Operations Centre (NOC), where all of the computers were left logged in with system administrator access. He then used those computers to bring down the system that provided address information for emergency 911 calls based on the incoming phone number for emergency services. Before leaving the NOC, he stole the backup tapes for the system. On leaving the NOC, Mr S proceeded to the backup offsite storage facility, where he gained unauthorised physical access to that facility by once again using the contractor's badge, and then stole additional system backup tapes. In all, 55 tapes were stolen. Fortunately, most areas affected by the disabled central system were able to switch over to regional 911 systems; however, some areas had no 911 backup capabilities. These actions caused over $200,000 damage to the company.

How would your organisation discover what has happened in this incident?

What action would your organisation take to recover from this incident?

Learn

By establishing processes that allow it to learn from each incident, your organisation will be able to improve its management of employee risk continuously. Even as an incident is unfolding, it is important to note points that need remedial attention and pass them to someone beyond the investigation team to implement.

Top tip: Capturing lessons

Designate someone to keep a log of all aspects of incident response. The log can be used during the post-incident review to enable lessons to be identified and learned, and necessary improvements to be made.

The log will make it easier to conduct a constructive learning review of successes and failures in both readiness and response. The learning review should address:

- What went well/didn't go well?
- What was expected/unexpected?
- Why? Should it have been anticipated?
- What could you control/not control?
- How could you have done it better?

What specific improvements do you need to make so that you are better prepared in future to prevent and respond to a similar incident?

Where disciplinary proceedings are still under way any log should avoid making judgements on the personal responsibility or liability of employees involved.

Evolve

Your organisation does not have to suffer an incident to evolve: you can learn from the experiences of others. The most effective organisations use their awareness of the external environment to anticipate new vulnerabilities, threats and risks.

As technologies change they bring new benefits but also new risks. Your organisation should be anticipating these environmental changes for business reasons but it is also encouraged to do so in the Data Protection Act.

Anticipate regulatory trends

As society's tolerance for data loss and invasion of privacy decreases, new guidance and codes of practice are developed by industry, government and regulators to prevent or deter such breaches. The consequence of these regulations can drastically change a risk assessment. New fines for data loss will certainly change the impact of an incident. It is good practice that the implications of such changes are understood early and embedded into policy and practice by the time regulations come into effect.

UK organisations should look at both UK and EU legislation for regulatory trends. Remain vigilant to changes that could affect your business. International organisations should be aware of what they are obliged to do and can do by law as well as best practice in the countries in which they operate.

Example

A 'personal data breach', is proposed in the EU data protection law, which is a *…breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data.*

Organisations will need to adjust to detect and address these kinds of incidents.

The Information Commissioner and the Financial Services Authority both build upon prior regulatory cases when considering enforcement action, creating a regulatory form of legal precedent.

Cloud computing

How should organisations manage the risk from employees to their data held in the cloud? In many cases, while the data storage service is outsourced, the risk remains with the client organisation since the risk cannot be outsourced. In such cases, your organisation should consider seeking contractual guarantees that your service provider will address them when the contract comes up for renewal and put in place provisions that enable the service provider to take an active role in mitigating the risk through appropriate correlation of data.

Communicate

Your organisation should plan a communications strategy in advance and control communications to ensure that the right messages are reaching the right people post-incident. There are three options. Which you choose will depend on your assessment of the pros and cons:

- say nothing
- communicate internally
- communicate externally.

A well-drafted and timely internal communication about an incident will:

- help correct or prevent rumours
- inform and build confidence
- deter others from acting in the same or similar way.

Guard against providing information that will allow employees involved in the incident to tip off any accomplices, who might then destroy useful evidence.

Good external communication can actually improve stakeholder confidence and shareholder value by showing that your organisation has effective response management and recovery processes in place.

Top tip: Lines to take

As part of your communications plan for an incident, agree, in advance, lines to take with both internal and external audiences on a range of scenarios.

Emerge stronger

By learning from each incident and feeding any ‘lessons learned’ into the risk assessment process your organisation can continually improve its holistic management of employee risk. By applying this approach you will:

- mitigate against identified people risks
- remain compliant with dynamic regulatory requirements
- anticipate regulatory trends
- anticipate the expectations of employees and wider stakeholders
- reflect, strengthen and contribute to the security culture of the organisation.

Appendices

Case history 1: Indicators for people risk

(Click question to return to case history)

A

What are the indicators for people risk in this case history?

Indicators for people risk for Ms G are:

- poor performance review
- demonstrated she was disgruntled
- working hours anomalies
- access of data unrelated to their role
- access of valuable information.

Find more information on indicators for people risk on the CERT website

Would your organisation identify these indicators for people risk?

Most organisations can identify these indicators although may not consider them as high risk. Normally the indicators are identified in their respective silo and not brought together. This guidance recommends taking a holistic approach, for example in this case taking input from the line manager, Physical Security and IT.

What conclusions can you derived from the change in behaviour?

There are many conclusions that can be drawn from this change in behaviour. Two examples are:

- Ms G had taken Mr A's advice to heart, was working harder to impress Mr B and broadening her knowledge of the organisation.
- The interview with Mr A had been the tipping point into disaffection and she was acquiring as much confidential information as possible with a view to taking it to a rival employer.

It is important to consider both conclusions but also to presume innocence as a majority of cases like this prove to be innocent behaviour.

Case history 2: HoMER in practice

(Click question to return to case history)

How many of your organisation's top risks relate to people?

Every organisation is encouraged to perform risk assessments and risk management and most organisations will have a corporate risk register. However, many organisations do not associate these corporate risks with people when, in reality, they often are.

How does your organisation identify, assess and manage people risk?

Many organisations identify risks, though may not associate them with people. It is common that these risks will have many different accountable stakeholders. These risks are normally addressed in silos, which results in an incoherent approach. In this guidance we advocate taking a holistic approach that addresses all the risks that relate to people and addressing them in an integrated manner.

Is there clear accountability for the management of people risk in your organisation?

If identified, the elements of people risk are often spread across various functions depending on who the risk is to. This normally involves HR, IT and Physical Security acting in silos which results in an incoherent approach. However in this guidance we consider who the risk is from. We recommend a single point of accountability to ensure a holistic approach.

Case history 3: A transparent approach?

(Click question to return to case history)

A

Is the application of protective monitoring transparent? How could it be improved?

Although the monitoring is detailed in the contract this is not sufficient and the employee needs to be better informed. This can be done by:

- documenting monitoring in policy
- informing employees when and where they are being monitored for example through the use of physical signs
- informing employees when they log on to the network.

Is the application of protective monitoring legal in 'a transparent approach'?

Based on the evidence in the case history the organisation is not carrying monitoring out legally:

- because the monitoring activities have not been properly explained to the employees the organisation could be accused of unfair data processing
- the organisation should have an internal regulatory framework that requires:
  - a written policy framework for monitoring
  - an impact assessment before monitoring commences
  - a supervisory function, so that monitoring decisions and activities are subject to periodic and specific reassessment.

Case history 4: Access management

(Click question to return to case history)

Which roles have privileged access in your organisation?

Most organisations do not have a holistic single view of privileged access. By following the guidance in Access controls an organisation should have a defined set of access for each role.

How do you control addition and removal of access? How long does this process take?

Every organisation should have a documented process for managing access. This process needs to be fit for the purposes of the business and should ensure that changes are implemented effectively and in a timely manner.

How does your organisation manage the increased risk that roles that require privileged access present?

Regular risk assessments should be carried out for those with privileged access. Section 2 'Assess' shows how to assess these risks.

Case history 5: Identifying high-risk behaviour

(Click question to return to case history)

A

What indicators of high-risk behaviour could have been observed?

The following are indications of high risk behaviour that could have been identified:

- Reports that he was ‘a chronically unhappy employee’.
- Disgruntlement communicated to his co-worker.
- His communication regarding a potential job offer.
- High volumes of data sent to an external address.
- High levels of access to secure areas of the network.

Does your organisation have the capability to identify these types of behaviour?

The following sources are normally accessible but may require the appropriate policy in place:

- The line manager could have identified that the employee was disgruntled. This could have been fed to those responsible for people risk.
- His co-worker could have reported his intent to those responsible for people risk.
- IT could have identified the high use of privileged access.
- IT could also identified the high volume of data leaving the organisation for an external account.

What information would be required by a protective monitoring team in your organisation? Is there an easy way to access this information?

These sources are usually easily accessible. Reports on the employee’s behaviour will require appropriate policy in place. A protective monitoring team may be allowed access to network application logs for the purpose of monitoring provided the employee is made aware that such monitoring may be carried out and has consented to it. This can be done simply through a screen prompt on log on to the organisation’s IT systems.

Case history 6: Line management concerns

(Click question to return to case history)

Is access justified for each of the line managers?

Access in the first is not justified as the line manager is not using a means that is proportionate.

In the second case the risk is greater and monitoring may be justified but this will vary from business to business.

In neither case, should the line manager be given access to the logged data, since this is inherently sensitive, and the use of logged data even by the monitoring team always needs careful consideration.

How should the protective monitoring team respond?

The protective monitoring team should review times when the employee had logged-in, only, perhaps on a sampling basis, to verify either the guilt or innocence of the employee.

Case history 7: Recovering from an incident

(Click question to return to case history)

How would your organisation discover what has happened in this incident?

It is quite likely that the details of this incident will not be evident on first examination. A detailed investigation would be required to understand what had happened and why.

The principles in the Respond and Recover chapters should be followed and, since there was a theft of data tapes, the organisation should consider external action.

What action would your organisation take to recover from this incident?

Action to recover should include deciding on what route should be taken, internal or external.

Given that the employee acted on good intent this should be taken into account. Action should be taken to learn lessons and integrate these into the organisation.

Technical appendix

Technical appendix

This section looks at technology which can be deployed to develop an integrated employee risk monitoring system that draws upon observations in the physical, information and personnel worlds to enable an informed risk assessment to be made about employees and others with inside access to an organisation.

This section is informed by a survey of available products – either specific security products or products that can be used for this purpose. There are a large number of technology providers with the capability to gather and analyse information to form employee risk assessments; the survey of available products informing this appendix is therefore not exhaustive. The technology described here should be readily available to anyone seeking to deliver a holistic monitoring system; however, it may not represent the cutting-edge of available tools, particularly as these products mature and evolve at a rapid pace.

Each organisation needs to approach employee risk according to its needs. Not all organisations need the range of technology outlined here, nor will they necessarily need to implement employee risk monitoring across the entirety of the enterprise.

Top Tip: Protective monitoring guidance

For government departments, more guidance can be found in the CESG Good Practice Guide 13

A pragmatic, technical approach to holistic monitoring

Ideally, a holistic monitoring system should be capable of taking feeds from physical security systems, security-relevant events from IT systems and personnel reports and amalgamate these information streams into a single database, supplemented and enriched by organisational context information (organisational structure, sensitive data, sensitive locations, etc.). A set of tools for data presentation (visualisation), data analysis and alert generation helps towards the interpretation of any specific alerts based on current activity and the overall employee risk assessment.

In summary, the high-level technical capabilities required from such a system include:

- Identifying and alerting policy breaches or other known high-risk events in near-real time (i.e. within seconds or minutes of the event revealing the policy breach occurring);
- Identifying actions over a longer period of time which indicate an increasing risk to the organisation;
- Monitoring trends in policy breaches to inform the need for corrective action;
- Adjusting assessed risk of individuals and assets according to observed events and alerting when risk thresholds are breached (in line with Personnel Security Risk Assessment;)
- Supporting investigation of suspected or actual incidents.

To take advantage of **all** available information, the system should ideally be able to:

- Aggregate data from multiple real-time and near-real time sources (IT systems and physical security systems);
- Use background information on areas of concern to the business such as inherently high-risk roles, sensitive information and physical assets and adjust risk reports as these change (probably slowly) over time;
- Make use of contextual personnel information such as organisational structure and absence (planned, unplanned and unused sickness and vacation) information for individuals;
- Use sensitive personnel information available from performance reports, anonymous telephone and e-mail hotlines, and regular background security assessments to adjust risk assessments.

The system also needs to protect the integrity and confidentiality of all information it processes, stores and produces.

A pragmatic, technical approach to holistic monitoring

A zone-based structure for a monitoring system

The rationale for dividing the system into the three zones is in part driven by the capability of available tools, and in part by the need to protect the most sensitive information. Structuring the system into these three broad, independent, zones also makes it easier for an organisation to incrementally upgrade the system, taking advantage of new capabilities as they become available. The suggested zones are:

zone 1: collection of data from real-time sources and background organisational data.

zone 2: aggregation of data and real-time alerting; long-term historical data storage, analysis and investigation support.

zone 3: final risk assessment based on sensitive personnel data and the output from the two previous areas.

The sensitivity of data increases with each zone, therefore security controls are required at the boundaries of the zones. Overall, the monitoring system should be considered to belong to a different security domain from the systems being monitored, the zones forming sub-domains within that overall security domain.

The overall zone structure and functionality is shown in the diagram opposite.

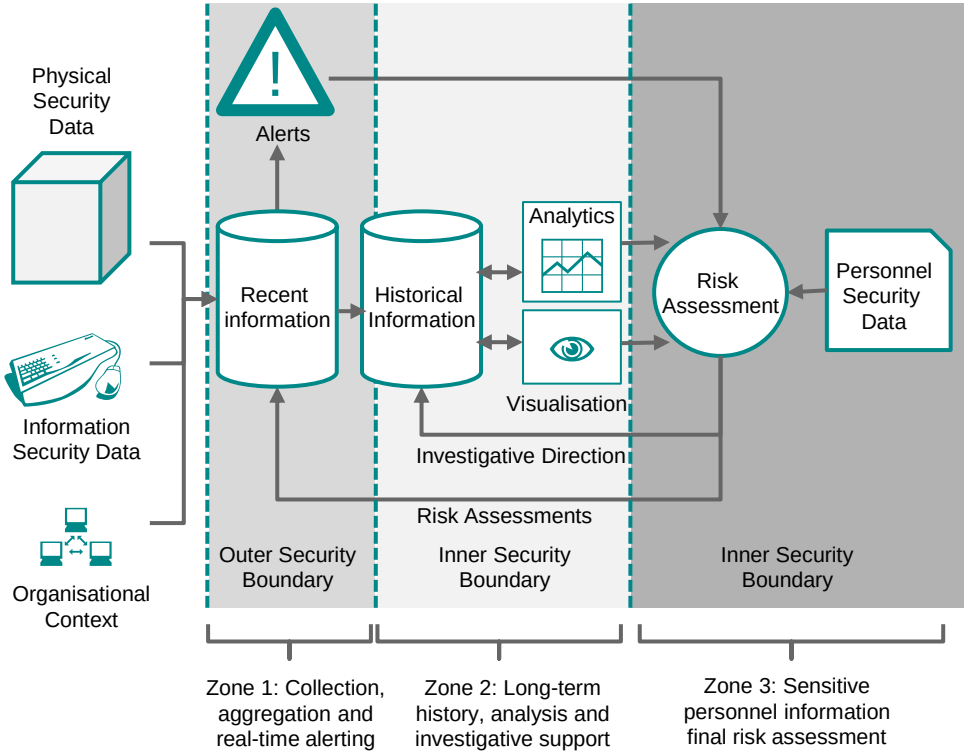

A pragmatic, technical approach to holistic monitoring

Ideally, only a single zone would be required that amalgamates the functions of the three zones, however currently-available technology is not yet capable of allowing this architecture to be implemented.

The division between zone 1 and zone 2 is mainly driven by the capability of available tools. Currently, there is a distinction between tools which have good capability in collecting, aggregating, correlating and alerting on near-real time events (i.e. those events received within seconds or less of the triggering activity) and those which have good capability in rich analysis of long-term historical data, although this distinction is becoming less important as both the real-time and the historical analysis tools are converging in capability. However, even if a tool were available with sufficient capability in real-time and historical analysis, there is still an argument that investigation support, which typically requires historical analysis of data as well as continuing monitoring of current events, should be segregated to a small need-to-know team of investigative analysts.

This zoned structure enables the feed of data from IT systems and physical security systems into a repository of recent information (zone 1), enriched with contextual organisational information. Data is correlated and rules run to detect patterns of high-risk activity, producing alerts. Data is also held in a historical database for long-term trend investigation and analysis (zone 2) and to provide data for directed investigations. The separation of recent and long-term data allows trend analysis and retrospective investigations to take place without impacting the performance of real-time alerting.

A pragmatic, technical approach to holistic monitoring

A separate store for long-term data also allows the use of a different data model that may be more suited to retrospective analysis tools and tasks than a data model designed for performance of correlation rules. The personnel information used to make the final risk assessment is some of the most sensitive data to be handled in the aggregation and analysis. For this reason, the assessment of the overall individual risk score should be performed within another security compartment (zone 3) in the system so the rationale for the final score is visible to a minimum number of people.

Many commercially available tools are capable of calculating a risk score for staff and assets based on events received and correlation rules. They are not yet sophisticated enough, however, to score risks directly from indicators in personnel management reports. For this reason, the final risk assessment must be formed from personnel security data combined with data based on events from IT systems and physical security systems. The resulting risk assessment can be fed back in to the information and physical event aggregation and analysis to be used in the generation of alerts, prioritising rule events that are associated with high-risk individuals or roles and allowing the revised risk assessments to pass into the historical database for inclusion in trend analysis.

The process for setting and reviewing risk scores needs to take account of possible feedback loops. Individuals can trigger further protective monitoring which will then find further alerts. These events can result in a higher risk assessment and further events. This feedback loop can result in false results and an organisation should mitigate against an un-checked feedback loop.

Some tools and techniques for monitoring employee risk overlap with those used for monitoring for signs of technical and physical intrusion. While this presents an opportunity to use the same (or similar) tools and techniques for both purposes, it can lead to a blurring of the personnel risk assessment mission into one of intrusion detection. Projects specifically setting out to create a personnel risk assessment capability need to be wary of drifting away from this objective and towards physical or electronic intrusion detection. While monitoring for the signs of external intrusions into an organisation is also a worthy (and necessary) objective, there is a danger that the primary focus on detecting the risk from employees may be lost. It is possible, however, that an existing capability in monitoring and managing external (physical or electronic) intrusion could form the basis from which a personnel risk assessment system can be developed as there is an overlap in the capabilities and data required.

A pragmatic, technical approach to holistic monitoring

Proving capability through a pilot system

A pilot system can be of great value in setting achievable expectations and in uncovering implementation issues early. Experience from the pilot can be used to develop the business case for a full-scale implementation. All tools and technologies have limitations. Identifying these limitations early on and adapting designs to minimise the impact of these limitations is essential if the HoMER implementation is not to be derailed or constrained in its abilities at some point in its lifetime.

Starting the project with a pilot to explore the limitations of available tools and technologies can prevent unpleasant and expensive surprises during or post implementation of the full system and help ensure the project starts with a good selection of tools and technologies. Identifying limitations early on means that a full implementation can proceed with more confidence in the tools, technology and high-level design and with more confidence in the benefits that can be realised.

A pilot should attempt to identify the limitations of tools to inform the design and set realistic expectations, and should seek to find limitations relating to at least the following features:

Real-time event processing – what is the maximum rate of events that the tools can receive without losing data? What is the limiting factor (processor, network, disk, software, etc.)?

Handling of delayed data – particularly for real-time analysis, how does the tool cope with data sources which may be delayed (perhaps due to events that have to be fetched in batches)? Does this cause analysis to proceed at the pace of the 'slowest' data feed? Do delays cause alerts to be missed or false-positives to be generated?

Volume of events – how do the tools perform as the database grows in size? Can real-time alerting rules keep pace as the volume of events grows? Does retrospective analysis take too long to scan large databases?

Flexibility in data received – how hard is it to add new data sources? Can new data formats be added easily? Can the data model be extended to include new data fields easily?

Repeatability of results – do queries and alerts return consistent results as the overall volume of data increases?

A pragmatic, technical approach to holistic monitoring

- **Complexity of alerts/queries** – how complex can the alerts and analytical queries become before performance and reliability start to suffer?
- **Correlation over time** – particularly with real-time alerting that attempts to correlate event A with event B, what is the practical limit on the time difference between event A and event B occurring?

A useful pilot system can be developed based on one-time extracts of security event data and evaluation of open-source tools, commercial tools that may be available with an evaluation license, or with any suitable tools already available within the organisation.

Using a one-time extract of security event data allows for tool comparison using the same data set to see which tool gives the best results. Extracting the data will not reveal some of the issues that may be encountered with real-time data provision, but will certainly reveal any difficulties in interpretation of data that may be encountered. Data should be extracted from some related systems (to test correlation ability) and from systems that use diverse technologies (particularly some physical and IT systems should be included). Ideally, two to three months' worth of data would be available for the pilot to allow for monthly trends in data to be identified, with the same time-period covered by all sources.

Part of the difficulty in running a pilot is that you may be uncertain whether there are any incidents or anomalies to be detected in the pilot data set, and uncertain whether anything the pilot detects represents a real security incident. To help increase confidence in the ability of the pilot to detect relevant behaviour, data representing security incidents, counter-productive activities and anomalous behaviour can be seeded in the pilot data, either before collection by deliberate actions on the systems that logs will be taken from or by editing/inserting crafted events to ensure that there are some known anomalies to be found. Seeded incidents should be based on a range of the assessed risks in order to gain confidence in the degree to which a full-scale system could meet the needs of the organisation.

Performance under high-volumes of data can be investigated by developing tools to generate additional data using the original data as a template (perhaps as simply as duplicating the data with different time-stamps). Real-time performance can be simulated by building test-beds to replay the live data in real-time or faster.

A pragmatic, technical approach to holistic monitoring

Starting to investigate the value of the data by investigating historical analysis capability (zone 2) rather than real-time analysis capability may be more valuable. The historical analysis will give some feeling for the level of policy breaches and trend indicators, and will also help point to areas that should be able to be detected in real-time.

Involving all stakeholders (physical and IT security, personnel, line managers, legal, etc.), in the pilot, particularly in evaluating the analytical capability of the tools, is vital to ensuring that the tools can produce relevant and valuable results. Staff with audit and investigative roles (or those likely to have such responsibility) should be as hands-on as possible in the development of the pilot to allow full investigation of the capabilities of various tools and to allow the tool capabilities to prompt questions the stakeholders may wish to pose on a real system.

Organisations wishing to develop a pilot should recognise that training may be needed on the tools which are being evaluated. Remember that training courses do not make instant experts! Consultancy support from the tool vendor can provide the 'instant expertise' necessary for a successful pilot, but can also mask some of the difficulties that an organisation may experience in developing their own capability to use the tool or technique in question. Consultancy support should therefore supplement, not replace, configuration and use of the pilot by those who will have that responsibility after the pilot.

Detailed capabilities required and availability of tools

This section provides more detail on the capabilities required to collect, aggregate and analyse information for the holistic analysis of anomalous behaviour. The emphasis here is on the automated assistance required for real-time and historical analysis (zones 1 and 2) where automation is key to success. Final risk assessment (zone 3) based on sensitive personnel information is less reliant on complex IT support.

Analysis and aggregation of physical and information security data, zone 1

To perform analysis and aggregation of data from information and physical security systems, the following capabilities are needed:

- generation and assured collection of accurate and timely log data from sources
- normalising event data into a common format and semantics
- representations of organisational structure (part of the organisational context)
- the ability to update the sensitivity of assets (physical and information) and the resultant assessed risk for staff (part of the organisational context)
- calculation of new data from the raw data received
- correlation between event data received
- classification of assets and staff based on event correlations (e.g. to update a member of staff as ‘high risk’ or an asset as ‘under attack’ based on rules)

Generation and collection of log data

Device operating systems typically generate logs for a wide range of individual events, but these events are generally intended to allow system managers to monitor the status of the device to detect problems and indicators of future problems. Security events generated by device operating systems typically report successful and unsuccessful attempts to access controlled resources (user accounts, files, commands, etc.) and this can reveal technical attacks or an insider attempting to access resources for which they are not authorised. Security events are also generated by normal system activities and activities by users – the production of a security event does not necessarily mean that a security incident has occurred, merely that a security-relevant event has occurred.

Most of the security events produced by a system will be the result of normal activity, the analysis of these events is necessary to determine whether they are indicating that activity that would cause concern is taking place.

Detailed capabilities required and availability of tools

Physical security systems typically generate events that are similar to IT system events, and it is relatively easy for products designed to cope with IT security alerts to aggregate and analyse physical security events as well, provided that technical interface issues can be dealt with. The exception is CCTV data, which will may need to be handled separately due to the specialist nature of video software (even if a digital feed is available) and referred to as required, based on time-stamps from other events if the CCTV system is not capable of automatically generating and forwarding alerts.

Systems vary in the logs they generate. For bespoke systems, it is important to include requirements to reveal user actions and the results/consequences of those actions.

To supplement IT application and operating system log generation, it is possible to install monitoring applications that will capture user application usage in a more direct fashion, either in the form of keystroke capture or of regular screen captures. These capabilities typically record continually, but may have the ability to retain only the information for a defined period around a detected policy violation (e.g. 30 minutes before and after a suspicious event is detected) to minimise the amount of data that has to be stored. Screen capture is similar to CCTV data from physical security systems and it can provide strong contextual evidence of the existence of an anomalous action. It can also provide details of the user's actions for evidential purposes. Care should be taken if such direct recording capabilities are to be deployed, as these may capture highly sensitive and/or personal information (e.g. passwords) and the intrusive nature of such capability may not be acceptable.

Events need to be captured in a secure location as fast as possible to prevent malicious or accidental alteration. However, this may be difficult when collecting logs from a widespread enterprise with limited network bandwidth. It may be necessary to provide local secure storage and send information in batches or to send summary data and pull details only when required. Collecting logs from mobile devices such as laptops that are often used while disconnected from any network can prove difficult.

While planning the events that will be captured and designing the ways in which those events will be captured, the legal implications of monitoring discussed in Section 2: Assess should be considered.

Detailed capabilities required and availability of tools

Modern physical security systems which use IP networks to transmit event data need to be handled with care, as the sensors for these are often placed in non-secure locations (e.g. a surveillance or access control or device on the outer perimeter accessible to the public) that could present an opportunity for an IT attack on networked assets.

Generation and collection of log data – available capability

Security Information and Event Management (SIEM) tools provide mechanisms for collecting logs from devices in near real-time, and also typically provide a level of assurance on delivery of logs and authenticity of source and destination. In some cases the SIEM tools can be configured to cache log information in a local area and send summary information for analysis to reduce network load.

SIEM tools often provide the ability to customise log collection agents to pull information from non-standard locations that may be used by applications. These customisation capabilities are likely to be important for acquiring data from physical security systems. SIEM tools may also store event data in proprietary, often compressed, formats that facilitate analysis by that tool. The use of proprietary (not publicly documented) data formats for event logs may present some issues if log data needs to be retained for several years as this could effectively lock-in the specific SIEM tool and its development roadmap.

There are a number of log collection tools available independently of SIEM products. Some of these are designed to allow analysis tools to plug-in to the dataset directly to perform batch analysis of the logs.

Detailed capabilities required and availability of tools

Event data normalisation

Unfortunately there is no universal format or syntax for reporting events. Different operating systems and applications produce events in their own unique form. Even if a common logging format is used, e.g. the syslog standard, the meaning of the data is likely to be unique to the system or expressed in a unique fashion. Normalising event data into a common format with standardised meaning prior to analysis makes the analysis task more efficient. Many SIEM tools provide normalisation for common operating system and application event data, and usually provide a means of scripting the conversion of non-standard event data into the normalised form, although there is a risk that standardised normalisation may lead to loss of some data in the original event (e.g. flags that are rarely useful may be ignored).

Event data normalisation – available capability

Most SIEM vendors provide conversion of common event data into a standard form, and usually provide a means to script the normalisation of non-standard event data. Since event data formats, meanings, and so on, can change with patches and new versions, use of SIEM-provided normalisation tools can mean that an update to a new operating system or application version has to wait until either the SIEM vendor releases a compatible normalisation function or a custom normaliser script can be developed.

Event correlation

Correlation is performed using pre-set rules (or signatures) designed to detect known sequences of events. Correlation can be simple ‘If *A* and *B* occur within *n* seconds’ type rules to complex signatures of threat activities. More complex rules can be built up using standard Boolean AND, OR, NOT type statements to describe the sequence of events being looked for, typically with some time-window in which the events must occur. Composite events can typically be created from rule matches and fed into further rules to create complex and flexible signatures.

Detailed capabilities required and availability of tools

Correlation rules tend to be triggered from an event collected by the system. In some cases, rules need to be created which include the absence of an event (e.g. a user accesses a desktop system without first having entered the building in which the desktop is located, potentially indicating the use of ‘borrowed’ credentials).

This can require work to turn the absence of an event (in this case the building entry event) into something that can be checked. For example, the ability to add users who have entered a building to a list of active users (and remove them from the list when they leave the building) provides a record that can be checked against. A correlation rule can then be written that to alert to desktop accesses when a user is not on the active list, delivering the desired result of alerting when a user accesses a system without entering the building.

Correlation rules become more powerful if the monitoring system has the ability to enrich events, either by creating sets of users or assets based on observed activity, or by adding calculated values (e.g. average and total number of file download events). The ability to enrich events with calculated or rule-generated information is a significant asset to the detection of pre-defined event sequences.

Event correlation – available capability

Correlation of events to match a pre-defined sequence is a basic feature of SIEM tools. The complexity of the rules that can be generated vary, as can the ease of rule creation, management
and predictability of results.

Some SIEM tools are limited in the time over which events can be correlated as events often have to be held in working memory for the period of time of the correlation. Most SIEM tools can correlate events that occur over minutes, but as the correlation window extends to hours or days fewer SIEM tools can cope.

Detailed capabilities required and availability of tools

Calculating new data

The ability to calculate new data from events allows significantly more powerful correlation rules to be created. Mathematical operations can be used to create new data that can be referred to in rules, for example to count the number of events or calculate statistics.

Data can also be created in response to some observed event or sequence of events, for example, to create sets of ‘active’ or ‘inactive’ users, but may also be used to indicate ‘threatened’ assets, for example by noting an asset that has been the subject of low-level probes, to escalate the response should further events involving that asset be detected.

The ability to calculate threat or risk scores for users and assets based on observed events during a period can be useful in prioritising responses to alerts and in detecting ‘slow burning’ attacks or increasing sensitivity to assets or users associated with reconnaissance activities.

Calculating new data – available capability

Most SIEM tools and enterprise event management tools have the ability to perform some degree of data calculation, the results of which can be referred to in analysis rules.

SIEM tools often have the ability to add or remove assets, users, or external entities to various sets or watch lists based on observed events (correlation rules matched). The sophistication of calculated data varies between tools.

As with correlation events, the ’window’ of events over which data can be calculated can vary, and the ability to hold the results over a long period can also vary.

Detailed capabilities required and availability of tools

Classifying entities

The ability to classify assets, users and external entities that will be the subject or cause of events can be used to develop watch-lists of high-value assets and high-risk users or external entities. Watch lists can be created manually (entering a set of high-risk assets identified through the risk analysis and assessment process or a set of high-risk users identified from personnel security reports or line-management feedback) or can be generated automatically by the system based on observed events (promoting users observed performing minor suspicious actions to a watch list of high-risk users). These watch-lists can then be used to prioritise alerts or to select which rules will be activated.

Hierarchical classification can be used to represent relationships between assets and relationships between individuals. Modelling an organisation's structure (classifying users into roles, business units, etc.) allows observed actions to be placed in the context of the individual's role in the organisation (for example, a member of the personnel department accessing proprietary engineering designs would probably be more significant than a member of the engineering team accessing the same document). This information can also be used to develop profiles of expected behaviour – for example, analysing observed behaviour within peer groups to find common and exceptional patterns. Richer organisational information can be used to refine this analysis, for example rank, length of employment, etc. can also be used as peer groups for comparison (behaviour of senior ranks may be significantly different to junior ranks, even within the same business unit – behaviour of new hires may be different to the behaviour of those with more time in service).

Classifying entities – available capability

Most SIEM tools have some form of watch list capability that can be populated with information such as high-value assets or high-risk users. Some have the ability to classify on-the-fly, for example using rules to generate risk or threat scores and effectively adjust threat/risk classifications automatically based on behaviour observed by the system.

The ability to classify users based on organisational structure is found in some SIEM tools that can read role and other information from a user identity management system or directory of users (typically used to manage authentication and authorisation). It should be noted, however, that the automated population of role/business unit classifications from a user management system reduces the effort required to keep information in the monitoring system up to date, but does not replace the need to ensure that useful and accurate information is present at the source.

Detailed capabilities required and availability of tools

Analysis of historical event data, zone 2

The capabilities described in this section are considered to be desirable in a holistic monitoring system. These extend the basic monitoring capabilities in areas such as detailed understanding of observed behaviour, dynamic adaptation of monitoring and improving the analysis and initial triage/investigation of detected incidents. The range of tools considered here goes beyond the basic log management and event integration/correlation systems focussed on in the essential capabilities section, and begins to introduce techniques and capabilities that have yet to be widely adopted in mainstream tools. While the use of these capabilities can be powerful additions to the overall monitoring ability, it should be noted that these can place more demands on the knowledge and skills of the analysts using the capabilities.

Data visualisation

The ability to generate and manipulate graphical representations of events, users, assets and the relationships between them has several valuable uses in the holistic monitoring process. Firstly, good data visualisation has a role in understanding the data to be analysed and the behaviour they represent, both up-front prior to implementation and as a regular check (or after implementation of a new system or following business change) that actual behaviour and analytics are aligned. This understanding can be used to help analysts identify patterns (either normal or anomalous) in the data that can be used in detection rules. Visualisation techniques also help to identify problems in data quality, for example missing or corrupted events, and can be used to confirm that event data is producing the expected information. The human brain is adapted to detecting patterns in complex visual data-sets, and can analyse complex visual scenes for patterns and anomalies in a single glance.

Secondly, data visualisation techniques are useful in retrospective analysis of data, allowing analysts to spot patterns or anomalies that are only apparent over a long period of time or over a broad data set. Data visualisation allows analysts to become pattern recognition and anomaly detection systems themselves to spot patterns that would not be apparent from analysing test-based alerts or would only be apparent over a long time period.

Detailed capabilities required and availability of tools

Thirdly, visualisation techniques can be a useful way of presenting a situational overview of security incidents and other events. Many monitoring tools provide a scrolling list of alerts as the primary user interface, and it can be easy to lose oneself in the detail, thereby overlooking the overall threat picture. Having a display that shows an overview picture of activity can be an effective way of spotting broader trends and maintaining perspective.

Fourthly, visualisation can be a useful technique for validating alerts during the triage stage produced from monitoring systems. Visualisation can place an alert in context with other alerts, with events that did not generate alerts and can also contextualise the alert with time, either presenting a time-line view or an animation (or sequence of pictures in time) of events in a context. Visualisation can also help explain why an alert was generated by showing the sequence of events and rule matches. Lastly, visualisation can help with the investigation of incidents, allowing investigators to gain an understanding of the incident in context.

Visualisation techniques are many and varied. However, there are some particular styles of visualisation that have been shown to be consistently useful in investigating and detecting security incidents. Standard graphs and charts used in many spread sheet applications are good at displaying simple data sets in a way that can show trends, anomalies and comparative data. Tree maps can be useful to reveal the structure of large data sets. Network maps, showing a stylised view of a network and allowing the events and/or volumes of traffic flow are common techniques. Link charts are more generalised forms of network maps, allowing relationships between entities to be displayed using a variety of styles and algorithms, and can be highly effective in revealing previously un-noticed patterns or anomalies. Several data dimensions can often be shown on link charts using node size and shape, line thickness and style, and of course colour. Overlaying security event information on other data, for example geographic (map) data to highlight location or more abstract representations of organisational structure or information categories can create rich information sets.

Static views can be informative and useful, however visualisations become powerful investigative and discovery tools when they are made ‘live’ and allow the user to navigate through data sets by re-arranging and filtering data, calling up details on aspects of the picture or focussing on different areas of the picture.

Detailed capabilities required and availability of tools

Data visualisation – available capability

SIEM tools commonly use data visualisation for management overview reporting, explaining the events and rules causing alerts, and showing network views of events and alerts. Generalised link display charts are also available in some SIEMs.

There are also a range of commercial and open-source tools dedicated to producing visualisations. Some of these provide a range of pre-defined visualisations that can be customised (similar to the graph/chart capabilities of spread sheets or report generation tools), others are general-purpose data visualisation tools that resemble specialised programming environments rather than end-user tools. Specialised security event visualisation tools are also available and have the ability to interface to commonly used log capture systems. These have the advantage that users often publish the source code for visualisations created in these tools that they have found to be particularly useful, demonstrate a particular technique, or just produce aesthetically pleasing results.

General-purpose data mining tools (commercial and open-source) often have sophisticated data visualisation capabilities to enable the initial understanding of a data set, for use as an analytic tool and to display the results of analysis.

Investigative tools designed for law-enforcement and security agencies often have good visualisation tools built in for relationship link graphing and timeline analysis. These tools are naturally designed for the exploration of large data-sets, often containing information from multiple sources. They offer useful capability for investigating and uncovering previously unknown relationships in data acquired from multiple sources, for example, file and database access logs, building accesses, telephone logs and e-mail records. Information discovered through these investigative tools can be used directly to create incidents and to inform real-time analysis through the discovery of new patterns of suspicious or normal activity, or by adjusting the threat/risk levels associated with individuals or information assets. Needless to say, these tools also have great value in investigating incidents detected through monitoring or from anonymous sources such as employee hotlines and e-mail systems.

Social network analysis tools are designed for graphing and analysing the relationships between entities. However, these tools tend to have interfaces designed for academics studying this field, and the user is often expected to understand the algorithms offered, how to use them and how to tune their parameters.

Detailed capabilities required and availability of tools

Statistical analysis

The ability to subject large data sets of events to statistical analysis can be a way of identifying patterns, trends and anomalies in the data that can either directly reveal security-relevant incidents or provide valuable information to help construct rules for detecting incidents. Powerful statistical analysis tool sets rely heavily on employing staff able to understand the mathematics involved and able to correctly apply the correct statistical techniques for the data.

Statistical analysis – available capability

SIEM tools have the ability to perform basic statistical analysis to enrich data prior to analysis by rules and to produce trend reports. More sophisticated analysis is possible in commercial and open source statistical and mathematical packages and statistics-oriented data analysis tools.

General purpose data mining tools also often have statistical analysis tools able to perform basic operations, regression analysis and (in some cases) Bayesian analysis. Bayesian analysis can be a powerful technique for assessing data where there is a high degree of uncertainty to assess the probability of some conclusion from this uncertain data. Automated clustering techniques using statistical approaches (eg K-Means) are also common in these tools.

Machine learning

Machine learning typically requires a training data set, i.e. a set of events that contain known security incident patterns. In the case of insider threats or advanced technical attacks, having such a data set available for training is highly unlikely, but machine learning techniques can still be used if one assumes that the data set contains a baseline of ‘normal’ data (the system ‘learns’ normality rather than the data associated with known incidents).

Some approaches to machine learning, such as neural networks, can be opaque, in that it is not always easy or even possible to determine exactly what the system has ‘learnt’ or why it has ‘learnt’ what it has. Rule induction algorithms are relatively open, as they typically produce if-then-else style rules that can be studied by subject matter experts to verify that the induced rules are ‘sensible’ and have not included features that cannot be relevant to the desired result.

Bayesian learning systems can bridge the world of statistical analysis and machine learning, and in recent years have been successfully applied to a range of data prediction problems.

Detailed capabilities required and availability of tools

Machine learning – available capability

General-purpose data mining tools typically come with a range of machine learning algorithms and tools for testing their reliability. Open source machine learning software is available, and there are plenty of books available on the subject. However, these tend to be more academic than practical.

Automated anomaly detection

Ideally, security monitoring would automatically and accurately identify security-relevant events without manual intervention or lengthy rule creation processes. Given enough data, patterns of normal behaviour should emerge and deviations can be treated as alerts. While this ideal has yet to be found, there has been good progress in the area of automatic anomaly detection.

Tools with these capabilities can reduce the burden of creating and maintaining detection rules by identifying patterns of events over time (repeated sequences of events or repeated sequences of similar events) and presenting these to analysts, allowing rules to be generated when observed behaviour deviates from the expected pattern. The focus of these techniques has historically been on the detection of zero-day malware and novel technical attacks, but is now beginning to turn to the insider threat.

As with automated machine learning, there are cautions to be observed. Commercial tool vendors may not be forthcoming with the details of the techniques used for anomaly detection, let alone the actual algorithms used, which may mean that there are weaknesses or limitations in the detection mechanism of which the user is not aware.

Automated anomaly detection – available capability

Tools now exist to analyse data sets for patterns and propose (or automatically create) rules to alert when an anomaly is detected. As mentioned above, some tools have developed these specifically for detection of anomalous behaviour; others have generalised techniques developed for detecting zero-day technical attacks (signature-free malware or intrusion detection).

Detailed capabilities required and availability of tools

Final risk assessment, zone 3

The final risk assessment performed in zone 3 requires access to alerts and reports from real-time and historical analysis, as well as sensitive personnel reports, including background security check information and regular performance and security reports. Investigation case management tools may also be valuable to track the progress of investigations and to assist in providing assessors with an overview of the information provided from real-time and historical analysis.

The capability in this area is likely to be closely scrutinised as part of the overall governance, ethical and legal oversight of the monitoring programme, however other aspects of the overall capability, particularly the analysis capability in zone 2.

In practice, any of these areas could be the starting point, and a manager or colleague could come forward with concerns about someone's behaviour which then triggers a review of Information Systems (IS) and physical records.

Available capability

Investigative case management tools are available which track investigation progress and key indicators/evidence. These tools can provide additional assurance that the knowledge of, existence and content of investigations is carefully controlled by providing a workflow and document management capability that can be configured to ensure access is only possible on a need-to-know basis while providing overview capability so managers can track and manage the overall investigative workload and the progress of investigations. These tools can provide valuable supporting capability in responding to security incidents and subsequent investigations.

Case management tools typically provide support for managing the documentation produced during the investigation (either to capture evidence or notes made on the progress of the investigation) by collecting these into a 'case' that can then be managed as a single entity and to link related cases together. Case management tools typically also provide audit trails of the actions performed during an investigation and accesses to case files, which can help provide assurance of the integrity of an investigation.

Detailed capabilities required and availability of tools

Case management tools often provide a capability for managing the investigative processes and workflow itself, for example providing alerts to people working on the case when new information about that case is available and for allocating cases to investigative staff to manage the workload. The workflow/workload management capabilities can also provide useful management information, allowing managers to assess the total case workload, produce reports on the type and number of cases, provide reports on the number of cases by affected business area or geography, and ensure that investigators are assigned efficiently according to the current workload.

Some case management tools are sold as part of an integrated suite that may also provide incident management and visualisation tools. Some case management tool vendors also produce versions of their case management and investigation tools customised to deal with particular industries or types of investigation, for example financial fraud, theft, etc. These may provide specialised processes, reports, etc. to assist in compliance with the requirements of specific regulations or regulatory bodies.

Key published research findings

References

Common Sense Guide to Prevention and Detection of Insider Threats 3rd Edition – Version 3.1, Dawn Cappelli , Andrew Moore, Randall Trzeciak, Timothy J. Shimeall, January 2009. Information from this book has been referred to on pages 33, 41, 43 and 67.

The 'Big Picture' of Insider IT Sabotage Across U.S. Critical Infrastructures, Andrew P. Moore, Dawn M. Cappelli, Randall F. Trzeciak, May 2008

Useful links

NPSA Website

Data Protection and Freedom of Information advice – ICO

The CERT Program

The Chartered Institute of Personnel and Development

CESG Homepage

Governance and risk management – HM Treasury

The Institute of Risk Management

UK Cyber Security Strategy – Cabinet Office

Institute of Information Security Professionals

Association of Chief Police Officers

Information Assurance Advisory Council

Financial Reporting Council

Acknowledgements

PA Consulting Group and NPSA are grateful for the comments and suggestions received from Field Fisher Waterhouse LLP.

About the authors

This document was produced jointly by PA Consulting Group and NPSA.

National Protective Security Authority
Central Support
PO Box 60628
London
SW1P 9HA

Email: NPSA-enquiries@npsa.gov.uk

Web: www.npsa.gov.uk

PA Consulting Group
123 Buckingham Palace Road
London
SW1W 9SR

Tel: +44 20 7730 9000

Fax: +44 20 7333 5050

Email: info@paconsulting.com

Web: www.paconsulting.com

PA offices worldwide

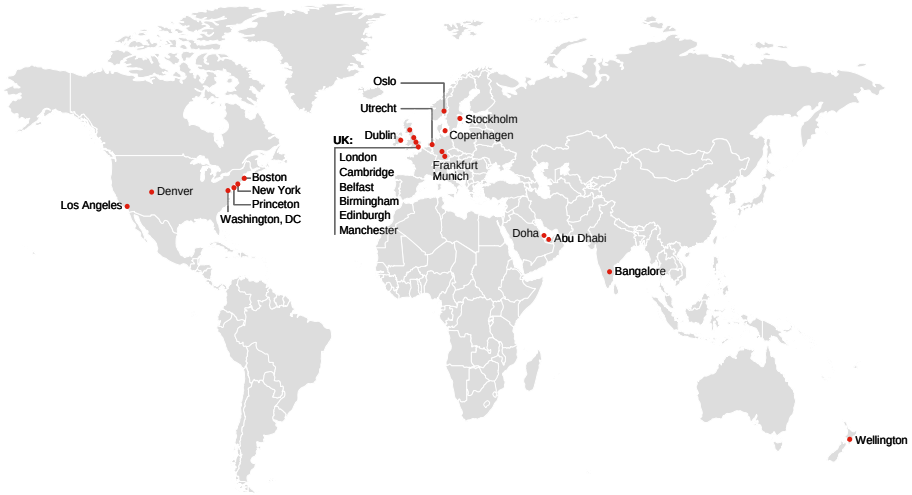

At PA Consulting Group, we **transform** the performance of organisations.

We put together teams from many disciplines and backgrounds to tackle the most complex problems facing our clients, working with leaders and their staff to turn around organisations in the private and public sectors. Clients call on us when they want:

an **innovative** solution: counter-intuitive thinking and groundbreaking solutions

a highly **responsive** approach: we listen, and then we act decisively and quickly

delivery of hard results: we get the job done, often trouble-shooting where previous initiatives have failed.

We are an independent, employee-owned firm of talented individuals, operating from offices across the world, in Europe, North America, Middle East, Asia and Oceania. We have won numerous awards for delivering complex and highly innovative assignments, run on of the most successful venture programmes in our industry, have technology development capability that few firms can match, deep expertise across key industries and government, and a unique breadth of skills from strategy to IT to HR to applied technology.

• defence • energy • financial services • government and public services • life sciences and healthcare • manufacturing
• postal services • retail • telecommunications • transportation

---

• strategic management • innovation and technology • IT • operational improvement • human resources • complex programme delivery

Delivering business transformation

Corporate headquarters
123 Buckingham Palace Road
London SW1W 9SR
United Kingdom
Tel: +44 20 7730 9000

www.paconsulting.com