

Good practice in information handling in schools

Audit logging and incident handling

A guide for staff and contractors tasked with implementing data security

Contents

1 Introduction	4
2 The need for audit logging	4
2.1 Which devices do we audit?	5
2.1.1 Routers and firewall devices	5
2.1.2 Windows servers	5
2.1.3 Linux servers	7
2.1.4 Connected devices	7
3 Planning a basic audit logging infrastructure	7
3.1 Basic audit/logging infrastructure	7
3.2 Implementing a basic audit log infrastructure	8
3.2.1 Audit and logging tools	9
3.3 Barriers to implementing audit logging	10
4 Security event auditing	10
4.1 What is security event auditing?	10
4.1.1 Identifying vulnerabilities	11
4.1.2 Identifying suspicious activities, break-in attempts and security breaches	11
4.2 What are the principal audit data sources?	11
4.2.1 Security configuration snapshots	11
4.2.2 Event logs	12
4.3 What is event logging?	12
4.3.1 What kinds of events are logged?	12
4.4 Why use event logging?	12
4.5 International standards for event logging	13
4.6 The limitations of event logging	14
5 Monitoring strategy	14
5.2 Detection	15
5.3 Response and notification	15
5.4 Damage assessment	16
5.5 Event anticipation	16
5.6 Corrective resolution support	16

6 Defining requirements for security incident response	18
6.1 Goals for monitoring with respect to incident response preparation.....	18
6.2 Detection requirements	18
6.2.1 Insider threat detection requirements	18
6.4 Compliance monitoring requirements.....	18
6.5 Response requirements	19
6.6 Resource classification.....	19
6.7 Platform coverage requirements	20
6.8 Audit source requirements	21
6.9 Corrective resolution requirements	21
7 Good practice for audit logging	22
7.1 Prerequisites	22
7.2 Archiving strategies	22
7.3 Rolling over the archive from online to permanent storage	24
7.3.1 Protecting data integrity	24
7.4 Disk space.....	24
7.5 Audit data management	25
8 Building an effective security incident response capability	25
8.1 Management commitment.....	26
8.2 The resolution team	26
8.3 Communications plan.....	26
8.4 Documentation	27
8.5 Awareness campaign	27

1 Introduction

All educational organisations must have a rigorous approach to data security. Much of the data held by schools and local authorities is sensitive, and protecting it is a legal obligation.

As well as encrypting data, controlling access and safely destroying sensitive data, schools and local authorities must keep audit logs to provide evidence of accidental or deliberate security breaches. These could include loss of protected data or breach of an acceptable use policy, for example.

This document outlines why schools need to produce a policy and procedures for audit and event logging so they are able to detect and respond to security incidents. It is aimed at network administrators and other staff (or contractors) who have a responsibility for audit and data security.

The document contains information and guidance on:

- planning and implementing a basic audit logging infrastructure
- security event auditing – collecting ICT system events and reviewing their impact
- devising a monitoring strategy
- defining the actions that you need your logging system to detect
- good practice in audit logging
- planning ahead so you can respond effectively to a breach of data security.

2 The need for audit logging

The report, Data Handling Procedures in Government requires public sector organisations to log audit and event data. Collecting logs is a critical part of providing a safe and secure ICT infrastructure for educational environments. Adopting ISO 27001/27002 will also make an organisation compliant with Data Handling Procedures in Government.

Schools will be required to configure a system that consolidates all of the data recorded by the information management systems, learning platforms and portals.

A typical network arrangement for an educational organisation has a number of items of hardware that schools or local authorities should collect audit logs from, including:

- hardware and software based firewalls
- web servers

- domain controllers
- learning platform web-servers
- web portal, database, query and index servers
- mail servers and web mail servers
- file servers
- routers
- networked PCs and other connected devices.

Audit logs should be collected from each of the above, and held for the length of time stated in the relevant local audit policy (section 7 contains information to help you in drawing up your own policy). Logging produces large amounts of data, so schools do not need to retain it indefinitely. However, specific security events should be archived and retained at evidential quality for seven years.

2.1 Which devices do we audit?

The rapid growth of ICT in schools, combined with the current regulations, requires educational organisations to collect audit data on an enterprise-wide basis. This will include physical devices, network and security devices, hosts, databases and the wide range of commercial and bespoke applications.

Log collection infrastructures must therefore be capable of meeting the needs of distributed heterogeneous networks whilst delivering secure and evidential quality audit log collection.

You can find further information on archiving data in section 7.

2.1.1 Routers and firewall devices

All school or local authority routers should be able to stream logging data to a central audit server. The central server keeps the data organised. For example, routers can stream log data to a designated audit and compliance server.

2.1.2 Windows servers

The Windows servers that record access, transfer data and record throughput include (but are not necessarily limited to):

- Internet Security and Acceleration (ISA) firewall servers
- Portal servers
 - Web front end servers
 - Database servers
 - Indexing servers
 - Query servers

- MIS server
- School central/ domain controller servers.

The native log events that are available come from:

- applications: information on actions performed within applications as programmed by the software vendor(s)
- security devices: information about user and processes with regard to security policy
- system level events: information about processes as related to core system functions such as DHCP, netlogon and other specialist services and events.

Additional logs will have to be enabled based on the specific hardware configuration located in each school. Internet Security and Acceleration (ISA) servers are able to monitor their own usage through using tools bundled within their distribution (see figure 1).

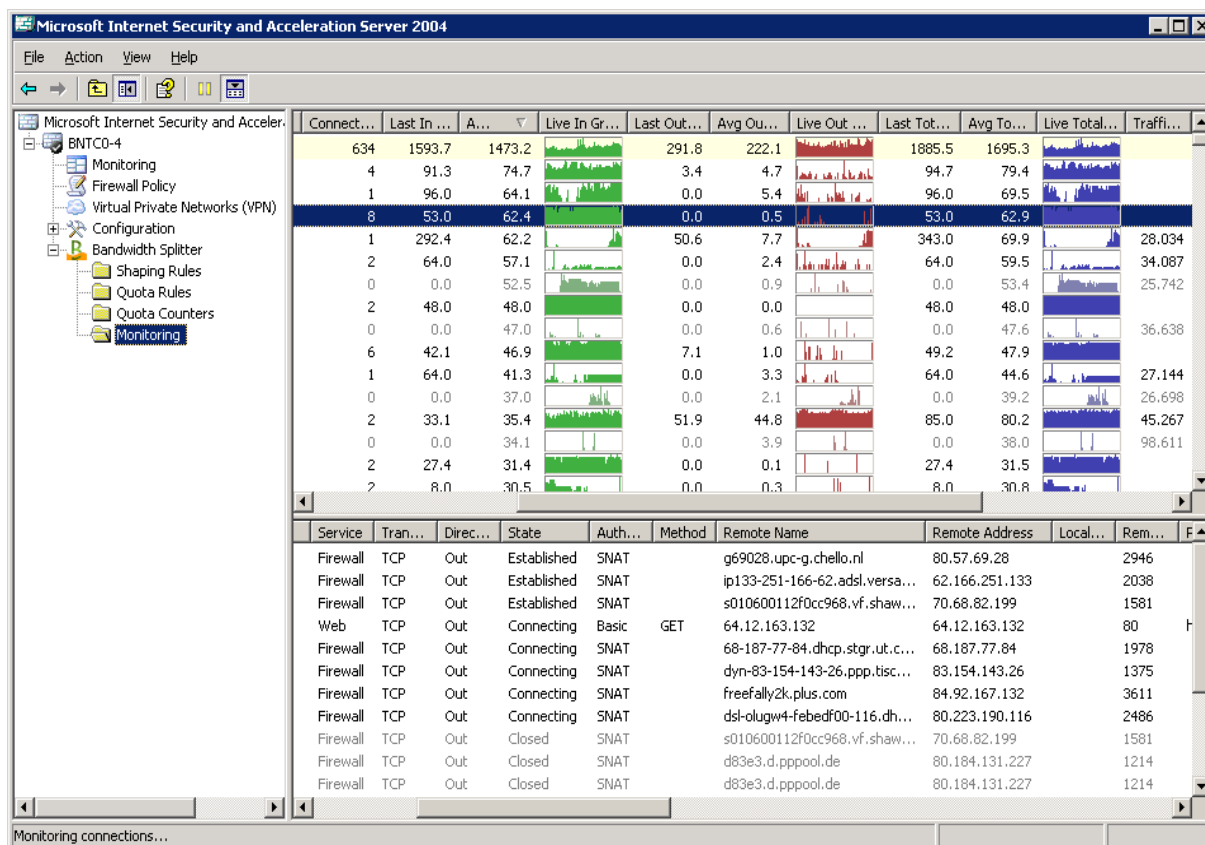


Figure 1: Screenshot of an Internet Security and Acceleration (ISA) server usage monitoring tool

2.1.3 Linux servers

The Linux servers record access, data transfer and throughput to:

- Learning Platform servers
 - web servers
 - database servers
 - audit servers.

2.1.4 Connected devices

These devices include all devices that may be attached within a school network that may need to access the facilities in question.

3 Planning a basic audit logging infrastructure

The information below is based on industry good practice and is intended as a guide. You need to define your own policies and procedures to meet your own requirements.

There are three ways to monitor systems for security breaches:

- Network level TCP/IP
- Server and application
- Process-specific.

3.1 Basic audit/logging infrastructure

Currently, the majority of schools and local authorities do not have an audit logging policy or consolidated auditing infrastructure. Most simply log the firewalls supporting the internet and email. Figure 2 shows the infrastructure that organisations will need to implement to provide basic audit and event logging. In this infrastructure organisations will need to manually consolidate log data.

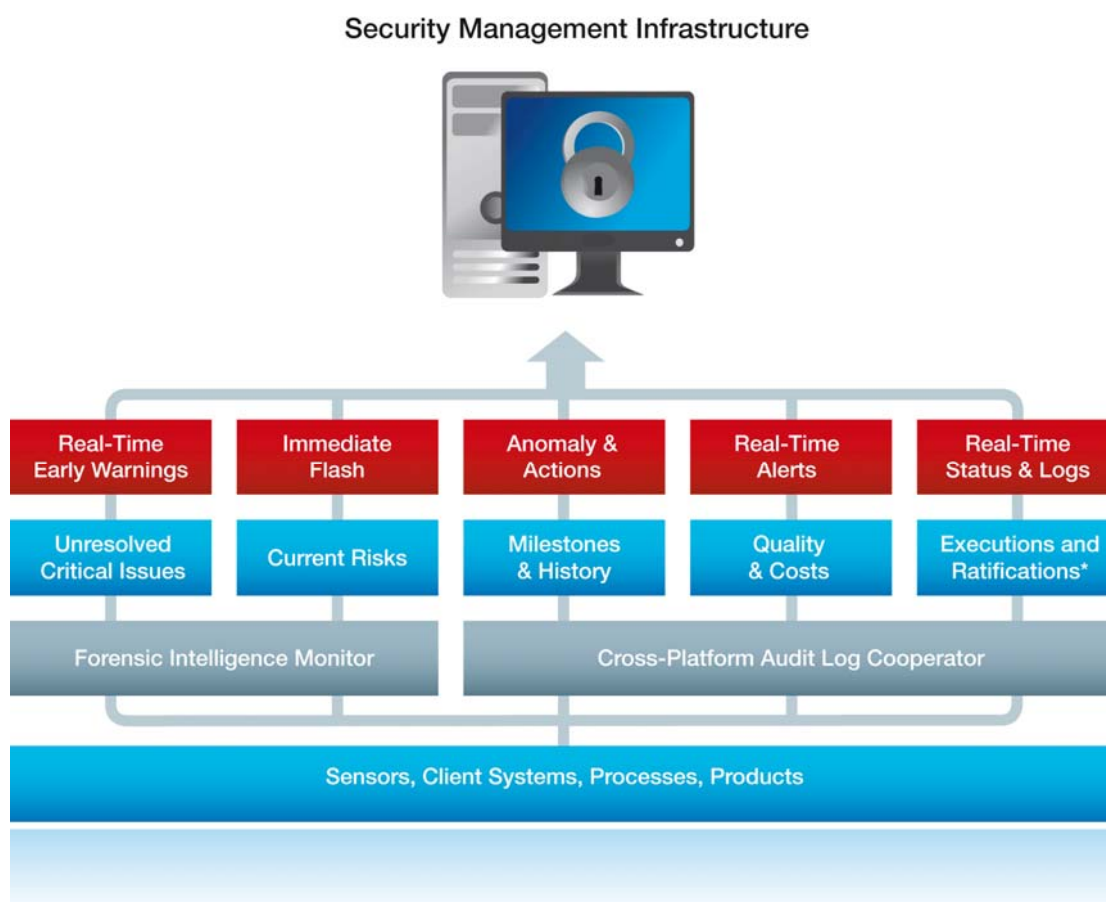


Figure 2: A basic audit and event logging infrastructure

The basic infrastructure shown in Figure 2 requires organisations to install extra hardware and software to provide storage and archiving that is of evidential quality. The extra hardware and software collects data from selected servers or appliances.

Raw log data is difficult to analyse and interpret correctly. In particular, it is difficult to understand how different logs relate to each other, and advanced ICT and statistical skills are needed. It is recommended that you consider additional audit log reduction tools to help with analysis.

3.2 Implementing a basic audit log infrastructure

The first steps necessary to implement the basic infrastructure include:

- Completing an inventory of critical systems (including those with data that is IL2–Protect and above) and determining what auditing or logging functions are turned on, where their data is written, the format, who owns the system and who obtains access to that data.

- Compiling a report that summarises this data and focus on the amount of data produced to determine network bandwidth, data storage requirements and recording format.
- Acquiring the necessary servers, hubs, network attached storage and firewalls to build a secure DMZ for these items to be installed. It may be necessary to build multiple audit/logging devices due to the distributed nature of the school's infrastructure.
- Nominating staffing who have responsibility for operating this security solution, including the information that is to be reported, archival processes and procedures for resolving discoveries and remediation requirements.

The main obstacles to implementing a basic audit log infrastructure are the time needed, concerns from staff about logging and access to log data and the extra storage space needed to hold logs. Schools need to consider all these issues before implementing audit logging.

3.2.1 Audit and logging tools

Schools should give technicians tools to carry out cross-platform query and reporting. This will be complicated by the need to manually consolidate the event logs. The formats will not be natively normalised, so cross-platform scripting may need to be developed. Staff will need to be appropriately skilled in order to organise and interpret the voluminous data and discover the hidden exploits. Initially, there may be significant gaps in the data sets and operators may not be able to discover certain types of exploits as they will not have tools to assist them.

A number of open source and commercial software products are available to help users collect and analyse audit and log data. A small selection of the tools available is presented here (Becta has not conducted formal evaluation of these and does not favour any particular tool).

Event Log Explorer

<http://www.eventlogxp.com/>

Event Analyst

<http://www.doriansoft.com/totalsolution/index.htm>

Exchange Log Analyzer

<http://www.mechanicalminds.com/site/ela>

Snare (Open Source)

<http://www.intersectalliance.com/>

Advanced Log Analyser

<http://www.abacre.com/ala/>

StealthWatch

<http://www.lancope.com/>

CA Spectrum

<http://www.ca.com/gb/products/product.aspx?id=7832>

WizRule

<http://www.wizsoft.com/>

Dragon by Enterasys

<http://www.enterasys.com/products/advanced%2Dsecurity%2Dapps/>

Not all the tools listed need extensive technical knowledge of the target system. These tools can communicate with the target system in its native language, perform the necessary transformations, and provide an integrated report, regardless of the operator's ability to write a query. These tools reduce the need for users to have technical knowledge, training and ICT forensic experience.

3.3 Barriers to implementing audit logging

Implementation projects have often failed for the following reasons:

- Lack of support from senior management and lack of directives mandating that all involved parties actively support the tasks
- Lack of management enforcement of the adopted policies
- Issues relating to data ownership and network bandwidth availability
- Inadequate funding to upgrade servers owned by others
- Lack of skilled staff and the ability to dedicate the required time, in spite of this being of the highest priority.

4 Security event auditing

4.1 What is security event auditing?

Security event auditing is the process of collecting ICT system events and reviewing the impact of these against security policy. Context is gained by linking the series of events with users' actions.

The process of auditing determines compliance or non-compliance to established security policies and procedures. Regular audits are recommended as part of running an efficient and well-controlled ICT operation.

4.1.1 Identifying vulnerabilities

Reviewing the security configuration will indicate whether it is set in accordance with a defined baseline guided by the security policy. Examples of the security configuration include security level, audit policy, password characteristics, registry, file/directory permissions, user accounts, groups, rights, privileges, and network configuration.

4.1.2 Identifying suspicious activities, break-in attempts and security breaches

The purpose of the audit is to identify any events that indicate suspicious activity. These would include, for example, users who repeatedly failed to log on; who logged on at unusual times; login from unknown remote systems; users who failed to open files or folders due to insufficient permissions; unusual use of admin privileges; and users who have repeatedly attempted to access system services, but failed due to insufficient privileges.

It is likely that these events are the result of a normal system activity – mistyping of log-on passwords or accidentally trying to open a files without appropriate permission are common mistakes. What is suspicious is the number of repeat attempts that are made. Only familiarity with the routine activity within your environment can determine whether any of the events warrant deeper investigation. It is therefore important to periodically scan audit logs even when no alert has been raised.

You are also looking for a possible intrusion from the outside, and need to find out whether an internal user has performed an unauthorised activity that violates the security policy. In other words, you are checking whether the security of a system has been breached and, if it has, determining exactly what has happened and where the first point of breach occurred.

To be able to do this, the computer system must have an event-logging facility to record the occurrence of significant events in the first place. The more sophisticated the event logging, the sooner you will detect an unauthorised activity.

4.2 What are the principal audit data sources?

4.2.1 Security configuration snapshots

In order to identify system vulnerabilities, the relevant security information will need to be collected and reviewed. This will include security level, audit policy, password characteristics, registry, file/ directory permissions, user accounts, groups, rights, privileges, and network configuration. This is the first type of audit data source, often referred to as ‘security configuration snapshots’ when captured and stored with a time stamp.

4.2.2 Event logs

In order to identify suspicious activities, break-in attempts and security breaches, you will need to collect and review all the significant events recorded by the event-logging facility. Note that the auditing parameters need to be appropriately configured in the first place, so that the event logging facility records all the significant events you want to monitor. This is the second type of audit data source, often referred to as event logs. On Windows, these are known as system, application or event logs, and syslogs on Linux, Mac and Unix.

4.3 What is event logging?

Event logging is the process of noting the occurrence of a significant event and recording it in a persistent medium. Each event is recorded in a log. Each record written to the log is referenced by date and time.

4.3.1 What kinds of events are logged?

In a security event log, you are likely to see security relevant actions such as:

- users logging on and off the system
- changes made to system security and user privileges
- attempts to access file, directories, printers, and other system objects that are under audit control.

These types of events inform ICT network managers and/or dedicated security managers how users and processes are attempting to access and use the system. The security log therefore contains a persistent record of 'who is doing what, when and where from'. Such an audit trail would give an indication of repeated attempts to illegally log on to a remote computer, gain access to secured files, or install unauthorised software. On a more mundane level, it would simply point to someone who seems to be printing out a lot of unnecessary documents, or spending too much time on the internet.

The auditing policy determines which events are recorded in a log. Specific system and user events are pre-selected by the security administrator based on how the system is configured and for what it is used.

4.4 Why use event logging?

It is impossible to guarantee that any computer system is 100 per cent secure. There will always be security flaws that can be exploited, and unfortunately the greatest risks to security are the users themselves.

If illegal entry and access to protected data cannot be prevented, such problems at least need to be recorded and tracked in an event log for the purpose of revealing the security flaws in the systems and possibly identifying those (humans or computers) that have exploited these flaws. This is a requirement of Data Handling Procedures in Government .

To summarise, event logs provide the information required to identify attempted attacks, to investigate what happened when an incident has occurred, and to potentially provide evidence in support of an investigation.

4.5 International standards for event logging

The best practice standard for measuring and performing event logging is based upon the International Common Criteria, ISO 15408. Education systems will generally require a minimum security level based upon the classes defined in ISO 15408. The Common Criteria security evaluation class requires that a computing system must have an auditing mechanism with the following minimum capabilities:

- The system has the ability to record all security-related events that occur on the system in the form of audit records
- The system provides a way for the audit records to be reviewed by the system administrators
- The auditing software and logs must be protected by the operating system from unauthorised access and modification, and access must be limited to authorised system administrators
- A mechanism must exist that allows the selection of security events to be audited
- The system must be able to audit individual users.

The Common Criteria mandate that each audit log record must contain the following information:

- Date and time of the event's occurrence
- Unique ID of the user creating the event
- Type of event
- Success or failure of the event
- Origin of the event (user, system, terminal, etc.)
- Name of the object accessed (a system file, piece of data or computing process)
- Description of any modifications made to security databases.

For these requirements to be met, the logging system must be able to monitor both its own activities and those of all local and remote users connected to the system. It

must also be able to report events pre-selected by the system administrators to one or more central structures. The logs and the logging mechanism must be guarded using the highest sensitivity level possible for system objects.

On protected systems, it is possible to track a user's access of files and directories, printers, network volumes and shares, and attempts to modify any security aspects of the system such as changing file permissions, adding a user account or privilege, or disabling auditing. All such actions may be recorded as events in a security audit log.

4.6 The limitations of event logging

All of this information does not necessarily indicate the specific events to log, or how to classify the events. For example, is a user logon and logoff a single event with two states, or two separate events? The audit events recognised by a system depend entirely on the capabilities and sophistication of the system's components and security mechanisms.

Changes to the security configuration should be recorded in the event logs. However, inherent weaknesses are observed in most event logging facilities. It may not be possible to record all user and process actions in sufficient detail. Because of the deficiencies in the event logs themselves, it becomes necessary to rely additionally on security configuration snapshots that identify vulnerabilities.

The following examples may help to illustrate the problem.

Example 1

If a security administrator changes the value for minimum password length (as part of the password policy settings), the security event log will record that a policy change has occurred, but not record what exact change was made. To be able to detect such a change would require a snapshot of the policy settings before and after the change was made.

Example 2

The default logging mechanisms capture commands that were executed but not always the parameters associated with those commands. If the access permission on sensitive files was changed temporarily for a few hours during the working day to gain inadvertent access to secured files, and if the snapshot of the file system is only collected at the end of each working day, the fact that permission on those files was changed and reset to its original value after a few hours will not be detected.

5 Monitoring strategy

Monitoring provides a significant deterrent effect, since people who know that their actions may be monitored are less likely to breach regulations. Deterrence can be observed as a significant shift in user behaviour based on the introduction of a monitoring system. It is not necessary to monitor every computer in the network to provide the desired effect.

Where users are not deterred, monitoring also provides support for corrective resolution as the data gathered is a crucial evidential tool. Evidence of unauthorised activity produced by the monitoring system is usually sufficient to take administrative action against an offender. This level of response, in conjunction with an awareness campaign about the monitoring, is usually sufficient to make the most obstinate user (or administrator) stop unauthorised activity.

5.2 Detection

Detection in a broad context means the identification of activities of note. Noteworthy activities may not necessarily be considered misuse or an intrusion. For example, detecting accesses to a mission-critical file may be misuse if the user account accessing the file is unauthorised to do so. However, simply counting the number of unique individuals that access a mission critical file can indicate that critical data is too widely available within a given school, local authority or other institution. An example would be the number of individuals accessing protected data and attempting to download the entire database. This is a noteworthy activity and may be tracked by a monitoring system.

5.3 Response and notification

Some monitoring systems are able to react after detecting misuse. These reactions may be automated or manual and include both local and remote actions and notifications. The notifications or alerts are similar to network-based monitoring systems and usually include:

- Pager
- SNMP Trap
- On-screen
- Audible
- Email

Once they have detected misuse, most commercial products include the following response actions:

- Log off user
- Shut down system
- Disable account

- Execute local script

5.4 Damage assessment

If an event and data protection loss has been positively identified, the first question to be addressed is 'What was the extent of the damage?'

Typically, files must be pulled from storage and hundreds of staff-hours might be spent looking through the data to determine the extent of damage. Additional hours must be spent on analysing databases and file systems in search of unauthorised modifications. If your threat was unauthorised disclosure, then only archived event log data will be of any value. If you have no archive then you will not be able to determine the extent of damage.

A key part of monitoring is maintaining an archive of information that can be mined using data forensics tools. An event log archive helps answer the following questions:

- Which computers were accessed?
- Which sensitive files were accessed and/or modified?
- What methods and tools were used to gain access?
- Was the individual working alone?
- How long have the events been going on?

5.5 Event anticipation

Many events are characterised by a set of preliminary activities before the actual loss is incurred. For example, a user who is looking for sensitive data with malicious intent may start a systematic search of sensitive systems before finding the critical data. If this pattern of systematic search can be identified then the user can be locked out before the system is compromised.

Many breaches will, however, be accidental. An example of this is when a teacher downloads data from the school information management system. Generally, this is protected data at IL2 or IL3. Frequently, they take this data with them by saving it on a USB portable drive. This action is typical for a single class or two. Doing this for the entire database should be an exception and require authorisation by the Information Asset Owner. This type of activity would be detectable by the monitoring system. Once the detection is made, the user can be warned whilst still in the process of generating the report or creating the file.

5.6 Corrective resolution support

Server tools can provide data to support corrective resolution – that is, action which follows detection of an infringement by a specific individual. This data includes

access patterns to files and computers with specific dates and times. This data may not be sufficient by itself to lead to a conviction but in conjunction with other evidence, it can indicate specific computer activities.

Note, however, that there are rules that control the admissibility of monitoring data as evidence in a court of law. Forensic evidence such as logs must pass several tests including chain of custody and integrity tests. This means that you have to be able to reasonably protect the data from modification between the time it is collected and the time it is submitted in court as evidence.

Computer evidence is admitted relatively easily. Taking all this into account suggests that the only data that can sufficiently be protected is the original raw event log files if they were created in a secure manner. Server and application based monitoring can provide the secure collection and storage mechanisms so that the data may be admitted as evidence in court.

You may need to employ the assistance of a forensic analyst to explain the actual log events and the relationship between system events. A critical aspect of any audit/log collection system is accurate time stamps of these events. Network time synchronisation is a priority, or you will need to calculate the time difference for each event across the various systems. Without consistent time stamps (such as network time domain) it is almost impossible to provide the required assurances for court evidence.

The term 'corrective resolution' refers to action(s) that may be taken against individuals as the consequence of their culpability in the loss of data, a breach of security or a breach of security policy. Corrective resolution may cover exclusion, suspension, restriction of privilege, restriction of use, termination of contract, prosecution or any other measure required to be enacted. Severity of corrective resolution must be determined based on the context of the incident and within the judgement of the relevant authority.

6 Defining requirements for security incident response

6.1 Goals for monitoring with respect to incident response preparation

It is good practice to have a statement of intent that covers both the type and reason for monitoring, such as:

“We are going to use monitoring to identify threatening behaviours including inappropriate use of resources or access to sensitive data, and attempts to circumvent policy across our network so that we can focus our limited resources on measurable threats.”

Monitoring of this sort should be made clear in any staff handbook, and in any acceptable use policy for staff, together with the response that the school will take in cases of misuse.

6.2 Detection requirements

Stipulating your detection requirements is your most fundamental need. Whether using an automated system or manual methods for detection, you need to ascertain what you want to look for in the data so that you can gather the proper events.

Using your statement of intent, define the actions that you would like to detect. Start at a high level and move into more detailed requirements as the process develops.

6.2.1 Insider threat detection requirements

Requirements for insider threat detection cover the actions of authenticated users. The methods to detect insider misuse usually focus on trend and behavioural analysis. By observing access patterns to critical data and systems, an investigator can detect suspicious behaviours before damage can occur. The requirements are likely to be as described below.

Requirement: Monitoring shall be used to track access patterns to critical data (that is, protected data) and systems.

Requirement: Monitoring shall be used to detect common vulnerabilities used to gain privilege as an authenticated user.

(Note: Unauthorised use of privileges is a general class of detection requirement serviced by server and application based detection mechanisms.)

6.4 Compliance monitoring requirements

Compliance monitoring ensures that people and processes are following the security policy guidelines. Policies are established to provide a verifiable level of protection in a network. Compliance monitoring can be provided through behavioural monitoring or static configuration analysis. To establish these requirements, consider which of your security policies, if ignored, would result in the most significant losses.

Requirement: The audit system shall detect when people or systems are not following the rules set out in the audit/logging policy.

(Note: This is a specific ISO 27001 compliance detection requirement.)

6.5 Response requirements

These are the requirements directly associated with incident response. Response requirements often start out being very ambitious (such as automatically logging out users and shutting down access to school systems) but as the security co-ordinators begin to understand the risks associated with automated responses, they can relax these requirements significantly.

Most response requirements should be associated with escalation procedures designed to make the system most effective. Once again, these requirements should relate to your statement of intent.

If your primary reason for using logging is to record breaches, then your response requirements may focus on escalating administrative actions as an investigation proceeds. If your primary reason is real-time analysis, then you need to carefully consider the timing requirements for manual response and balance automated responses against the risks they pose.

Requirement: When a user falls under suspicion during the course of an investigation, the user's account shall be disabled on the approval of the ICT network manager or site security manager. The account shall remain disabled until authorised for re-enablement by the appropriate party.

(Note: This requirement relates to escalation procedures and prevents continued misuse during an investigation. However, there may be circumstances where the goal is to not alert the individual under review and allow continued access for surveillance, so this requirement could be restated. This will be dependent upon the specific circumstances involved.)

6.6 Resource classification

Not everything can be monitored, so you want to be able to focus analysis on critical assets and systems containing sensitive data. You will need to differentiate critical data from non-critical data, and sensitive data from non-sensitive data. The process of defining data is known as resource classification. Definitions must be established for identifying, controlling and handling the different classifications. These include information elements such as resource, assigns, grants, requirements, and recovery. The information populated in these categories establishes the requirements for each resource in each classification level. These are defined as:

- **Resource** – This is the definition of each resource. Define these as necessary to match the critical assets in your school. For example, all Human Resource data files, protected directories, website content files, and so on.
- **Assigns** – This is the person who assigns the classification to the resource, for example, the Senior Information Risk Owner (SIRO), the department manager, and so on. It's important to control data classification carefully and commonly between all users who share that data in different organisations. Data controls cost time and money so you need to be able to balance security with access requirements.
- **Grants** – This is the person who grants access to the resource (for protected data this is usually the Information Asset Owner – IAO). This is a very different role from the assignment role. The assignment role requires decision-making abilities and a higher-level view while the granting role is more about following instructions. The significance of the granting role depends on the exact nature of the control requirements.
- **Requirements** – These are the control requirements. They cover as broad a range as necessary to protect the data at its classification including who can access the data and how exceptions are handled. For example, the restriction 'Only persons in the school leadership team may open and read summary information management data' could allow as an exception 'Teachers may access personally identifying data for individuals in their school with permission from the granting authority'.
- **Recovery** – If data is important enough to be controlled then it should have recovery and back-up requirements.

6.7 Platform coverage requirements

Platform coverage requirements describe the functions and systems that you want to monitor. This includes network protocols, operating systems, computers and applications.

It is important to note that analysis tools are not compatible with all network hardware and software, so you have to know what you're going to monitor in order to select the right tools. You will also need to describe the target elements – the elements that you wish to monitor – in order to establish the proper operational requirements. The prioritisation process should focus on critical network elements at risk and stipulate requirements such as those below.

Requirement: Monitoring shall be required on designated platforms (Windows, Linux, for example).

Requirement: Monitoring shall be deployed to sensitive applications first, and then migrated to support infrastructures and end user applications.

(Note: This requirement sets a priority order for the deployment and relative importance of the various systems.)

6.8 Audit source requirements

Once you have defined your platforms, you need to define which audit sources you need to monitor on those platforms. Many servers have multiple operating system and application logs. The different logs have different levels of information and security associated with them. This requirement will enable you to enumerate the different audit sources that your monitoring procedures will need to cover.

Requirement: The logging software shall have the ability to monitor the various logs from operating systems, routers, wireless access points, firewalls and applications (school information management system and learning platform).

6.9 Corrective resolution requirements

As you will be using data collected during monitoring to enforce potential action against individuals, you need to take care to protect this as evidence. This includes selecting appropriate audit sources and tools that protect data sufficiently for admissibility in court. Many of the requirements for corrective resolution are process-related and these requirements establish the level of information that needs to be available at various points during the investigation, or for reference by the relevant authorities.

Requirement: When a user falls under suspicion during the course of an investigation, all raw data shall be stored for evidence.

(Note: Evidentiary requirements place constraints on the quantity and quality of data as well as how it is handled. For further information, refer to ISO 27001 and the CSIA Incident Handling documentation , amongst others.)

7 Good practice for audit logging

This section suggests good practice for logging under the Windows, Mac and Linux environments. It is not intended to replace the specific guidance provided by the operating system manufacturers.

7.1 Prerequisites

Most importantly, you need to determine what you should collect, and how frequently you should collect the security configuration snapshots and the event logs.

Your monitoring strategy should determine what security configuration baseline you need to collect and how frequently. These snapshots can be used to determine whether the underlying system configuration has been changed, forming the baseline for future security audits with changes being tracked and assessed.

Your monitoring strategy should also determine what events you would need to monitor and hence, how frequently you should collect the event logs. It may be that you employ sophisticated tools and techniques to review only certain events online in real-time.

7.2 Archiving strategies

There are several archiving strategies from which you can choose. The archiving strategy is based on when logs are archived (when the audit log is copied to another location locally or to storage media and the log is cleared to gather more data) and when the logs are centralised (moved to a central archiving location so that while disk space is freed up on the target machine, the data is still available for online immediate reference). Below are some methodologies and their associated implications:

Strategy 1: Archive logs once an hour. Centralise logs once a day.

- The advantage with this strategy is that the raw audit logs will be small enough to sift through at a later date. In addition, the live data file will not need to have as much disk space available.
- The disadvantages include the fact that smaller file sizes mean more files to archive, and leaving them on the local machine longer gives an attacker

a chance to modify or delete them before they are centralised to a secure location.

Strategy 2: Archive logs once an hour. Centralise immediately.

- Here, the log sizes are small enough to be analysed effectively and the immediate centralisation means less time on the target machine to be at risk of compromise.
- The disadvantage of this strategy is that the immediate centralisation will force a spike in network traffic each hour rather than a previously scheduled time when network bandwidth is available.

Strategy 3: Archive logs once a day. Centralise once a week.

- This is the most risky from the perspective of log integrity but the easiest on resources because centralising is less frequent. This represents the least intensive in operational effort and is commonly chosen for this reason.

To select an archiving strategy for your school, you will need to balance the availability of resource against the requirements to collect logs in a timely and secure fashion.

How long these logs are kept is dependent on a local policy that you determine, the system(s) being monitored and whether protected data is involved. The points below represent good practice for log management:

- Your audit/log policy should state which systems are logged and the retention periods for each system.
- If you collect the audit logs from your system for a month and during that month no incidents have occurred, you could archive the data offline and retain it for one academic year (recorded in one-month intervals). If an event happened in a subsequent period, it would be necessary to go back to check if any previous pattern existed or if this was an isolated event. In such cases the archive schedule would likely provide an acceptable and relevant data set. This should be outlined in the audit logging policy.
- If a breach of the acceptable use policy had occurred, you would need to extract all the related data surrounding that particular incident and create a case file. As you would not know at the outset whether the breach might result in a court case, you should archive this set of records for a period of not less than seven years. This is a common practice similar to email retention in regulated industries.

Note that the seven-year retention times only apply to logging of actual events (breaches of the acceptable use policy). The regulatory policy is not intended to collect every bit of log data and retain it for extended periods.

7.3 Rolling over the archive from online to permanent storage

To carry out an investigation you will need to be able to access the data you have archived. Some of this data will be kept online for easy access and some will be stored on permanent media. Again, these requirements will depend on your school or local authority requirements or recent incidents that affect the risk profile within your institution.

A very common practice is to keep 30, 60, or 90 days online. Then, every 30 days or so, depending on your requirements, put two copies of the last 30 days onto separate permanent storage media and store them in a separate secure locations.

7.3.1 Protecting data integrity

Your data storage should be a secure media such as an optical Write-Once Read-Many (WORM) drive that can ensure the integrity of the data. Your aim, in the case of data being used in evidence, is to be able to say that the data was created by a secure data source, copied very quickly to a secure central server, and then put on indelible media so you can demonstrate that the data was not subsequently changed. The surety with which you can say these things is dependent on how quickly you copy it, what mechanisms you use, and what process you use.

7.4 Disk space

You will be required to set aside disk space for both local collection on the target machine and archiving disk space requirements for the online data. To estimate the amount of disk space necessary on the target system you will need to calculate the amount of data created between each collection period. The collection period is defined by your archiving strategy. If you archive logs once an hour and centralise logs once a day, then you need to define how much data may be created in a day before the logs can be centralised and the disk space is made available for more logs.

You should account for additional risks when determining how much disk space is required for storing event logs. For example, if a server is suddenly unable to send data to a logging device it must be able to store sufficient data until communication is restored and data is resynchronised, or until a process can be enacted to protect the data manually.

To calculate the online disk space required, multiply average daily log size by the number of logs that you will keep online and the duration of the storage requirement:

$(\text{Number of log files} * \text{Average Daily Log Size} * \text{Number of Days}) = \text{Archived Logs}$

Warning: If you do not have enough disk space set available on the target machines then disk space will run out and may cause your target server to stop operating and possibly shut down.

7.5 Audit data management

Audit data management is the practice of securely collecting and archiving audit (event) data for the purposes of analysis, investigation, and corrective resolution. Data management also addresses the issue of collecting audit data in a way that does not affect the operation of your devices and network. It is important that the data is collected securely so that it can withstand cross-examination if it is ever necessary for a case to be taken to court.

Data should be gathered in a secure manner that is tamper resistant. It is critical that the analysis methods are incapable of changing or altering the integrity of the data.

An effective audit policy is one that gathers just the right amount of data such that you do not gather so much data that you suffer performance problems and not so little data that events cannot be proved beyond doubt.

8 Building an effective security incident response capability

Data Handling Procedures in Government stipulates that public sector organisations should establish a security incident response capability. Depending on the severity of the security incident, it may also take the form of disaster recovery. Depending upon the size of your organisation and the specific skills of your staff, an effective security incident response may necessitate engagement with suppliers, your local authority or your regional broadband consortium.

Your policy for incident response should be tailored to both the above and to your specific site and operational requirements.

A prerequisite for an effective security incident response is the detection of the security incident in the first place. Thereafter, the effectiveness of the response team should be measured based on the extent of damage that resulted from each incident. The sooner the incident is contained, the lower the risk of financial loss or data compromise.

Good practice highlights the following components for the successful resolution of an incident:

- Management commitment, in terms of human resources, budget and priority
- A resolution team of technical and legal experts
- Primary responsible person for each incident
- Communications plan, including escalation procedures and interfaces with inter-departmental and law enforcement agencies □ Plan of action for rapid resolution
- Plan of action for non-recurrence
- Knowledge base of past security incidents, including steps taken for resolution and non-recurrence
- Awareness campaign

It is recognised that not all resources may be available and further guidance and support may be necessary to facilitate all the components outlined above.

8.1 Management commitment

Within the educational environment, strong emphasis is required from all responsible individuals in managing and responding to potential threats and risks to ICT. Senior representatives from all departments should establish a formal steering group to analyse the current education system threats. It should cover the following functions and activities at a high level: infrastructure, networks, systems and database administration, applications development, security, audit, legal and compliance.

8.2 The resolution team

The steering group should define the composition of the team that will address the technical and legal aspects of the issues at hand and expedite the solution.

8.3 Communications plan

Any activity as part of incident response is sensitive in nature. Improper publicity may harm the reputation of the institution. What is needed is a formal plan for when a security incident occurs.

The plan can be constructed by envisaging, provisioning and rehearsing for the events that occur when a security incident arises. Personnel contact lists, escalation procedures, operational responsibilities and guidelines should be formulated in advance and rigorously adhered to. Changes in plan should be avoided unless the situation leaves no other alternative.

Underpinning all of the above is a process of communicating your procedures to all relevant members of staff in schools and local authorities, associated parties and law enforcement agencies, to ensure:

- rapid resolution
- non-recurrence
- limitation of damage
- protection of all users.

8.4 Documentation

All processes enacted during security incidents should be recorded and archived securely to ensure that complete evidential quality is maintained. Coupled with the recorded evidence of the event, a debriefing operation should be conducted to evaluate the performance of the team and effectiveness of resolution(s).

Subsequent to the incident, debrief material should be analysed to ascertain whether changes in infrastructures, operational practices or policies are required. Suggested improvements may then be fed back into the processes and procedures to enhance the response.

8.5 Awareness campaign

An awareness campaign must be developed that identifies the security incident team, contact points or agencies required to be informed on discovery of a suspected event. This should be disseminated to any party that may have cause to use such a system. This should include (but not be limited to) school staff, support contractors, local authority employees, parents, governors and (where appropriate) children.