7

Information management

This chapter addresses some of the considerations involved in designing an information management scheme and making it operate in practice. It forms part of Stage 2 – Planning, assessment and evaluation in the Toolkit Route map.

Key topics

- **The benefits of having an information management scheme**
- **The components of an information management scheme**
- **Tips for creating and using a workable and appropriate scheme**

7.1 What is an information management scheme?

An information management scheme provides a framework within which information can be identified, its security requirements determined and instructions given to those who may handle it. Although it may be tempting to have the information management scheme echo the full complexity of an educational organisation, this is not desirable and should not be necessary. A complex scheme is too easy to misunderstand and mistakes could expose the organisation to significant risks. The aim should be to have the simplest scheme that will satisfy the organisation's requirements: identifying this is likely to involve a series of iterations between theoretical and practical considerations. The most effective schemes take a pragmatic approach that can be understood by all their users.

An information management scheme, in a simple form, can be made up as follows:

- A classification scheme: a list of classifications with definitions to allow people to consistently classify information.
- A labelling scheme: a way for documents and other information to be visibly associated with a classification.
- Handling rules: information on how to use and protect information with each of the defined classifications.
- A process which explains how to use the above three documents (e.g. how to decide who is responsible for classifying a given item of information).

As usual, these documents should be supported by a policy statement, and endorsed by top management. The statement specifies the scope of the information management scheme; who is responsible for maintaining and controlling it; and what sanctions should apply in the case of non-compliance (sanctions can often be handled via normal disciplinary processes).

The organisation should appoint a suitable role(s) to develop this scheme (see Chapter 8, Roles and competencies), and ensure that the scheme is tested and approved (see Chapter 2, Information security governance).

7.2 Classification

In our daily lives we tend to see a huge number of attempts to mark information with a classification – "confidential", "personal", "commercial in confidence", "private", "off-the-record", and even "classified". In addition, there are formal schemes such as the Information Sharing Traffic Light Protocol[2] and the UK Government's Security Classification scheme.

When designing an information management scheme for an organisation, it may seem prudent to implement all of the above classifications, if not more. However, a scheme that is too complicated will produce confusion, non-compliance and other unintended consequences – e.g. either information being seen when it shouldn't be, or not disclosed when it needs to be. A scheme that is too simple carries the same risks, as it forces people to either over- or under-classify. It should be noted that the UK Government's new scheme only has three classifications above the base level of unclassified – "Official", "Secret" and "Top Secret", of which the top level may never be encountered in most branches of the Government.

Classifications must apply to information, not to the particular form it is in: it makes little sense to say that a printed copy of a document must not be left on a desk, if computers with access to the same information are left logged in when unattended. As information changes format, it must experience a consistent level of protection.

7.3 How many levels?

To develop a classification scheme, the organisation should decide how few different sets of information handling rules it needs to use. In many situations, particularly given the need for consistency in handling across different formats of information, that turns out to be surprisingly small. Many organisations in the educational sector have settled on classification schemes with three or four levels, despite initially expecting that they would need more.

If the initial attempt at designing a scheme does produce a large number of classifications, the organisation should check for mixed or inconsistent treatments for different formats of information. This is likely to undermine both the actual effectiveness and the credibility of the scheme.

The classification levels chosen should be compatible with the classification structure implied by the Freedom of Information Act 2000 (FoIA). The Act effectively groups information into three classes with regard to confidentiality:

- information that is routinely published
- information that is disclosed subject to a public interest test
- information that is not disclosed.

There is no advantage in sub-dividing the first of these classifications, since anyone can obtain the information merely by asking for it. There may be some point in sub-dividing the third (and possibly the second) if there are clear divisions within the FoIA categories. For example "does not leave the building" and "viewable from outside" might be different sub-classes of "not disclosed". A classification scheme whose breakpoints do not match those of FoIA may be both confusing and liable to error. Note: this also applies to the Data Protection Act 1998.

Business Impact Levels (BILs), as used by the UK Government, are a way to formalise the assessment of risk. However the seven columns and significant detail in their Impact Level tables are likely to be too complex for practical information classification schemes. Noting the Business Impact Levels that are likely to match the organisation's own information classifications may, however, be a useful check, especially if the organisation will be expected to engage with BILs in its interactions with other organisations, e.g. funding bodies.

Classifications should take account the information's requirements for integrity and availability, as well as confidentiality.

7.3.1 A note on special cases

Encryption keys must be handled with extra care, as they constitute security controls in themselves, not just information. Equally, information on security controls, such as plans for a security system, lists of roles with privileged access, maps of sensitive areas, results from penetration testing, and risk treatment plans, are special cases where the risk may be intrinsically higher.

7.3.2 Naming the classifications

Once the organisation has agreed the number and definition of its classifications, they need to be given easy-to-recognise names. Classification schemes are harder to use if names don't form an obvious sequence – is "confidential" more or less restricting than "personal", for example? If a classification scheme does have a sequence, then its names should make that sequence obvious: something like "non-sensitive", "sensitive" and "highly sensitive", or "low integrity", "medium integrity" and "high integrity" (e.g. research data) works well.

Classification instructions/definitions should be clear and easy to follow, so that information owners can quickly classify their information, and have appropriate cues during information creation (or at receipt) to remind them to classify information. Guidance for information owners may also include policy and regulatory requirements that apply to particular kinds of information, and which require them to be given a particular classification. For example, personal information is likely to have an elevated classification.

Be careful if borrowing terms from classification schemes in other sectors. If the organisation decides to use a classification entitled "Secret", for example, then that will mean something very specific to anyone who has worked with Government documents, and that meaning, and related handling rules, probably are not consistent with those defined by the organisation.

7.4 Labelling

It is good practice, and may be a requirement, to label information with its classification. Different formats of information will need to be labelled in different ways: for digital documents or e-mails the label should be in a standard place in the digital content; for paper it should be on the file or envelope (double envelopes may be required if the actual classification needs to be protected); for on-line systems the label may need to be on a login page if it's not possible to put it on every screen.

The important thing is that the label is understandable (this is why the classifications are created first), and visible to all readers, even those who only skim the start of a message. Everyone who sees information must know how to handle it.

It may be necessary to consider information as being of two types: structured (e.g. files in a database or CMS) and unstructured files (anything ad hoc, e.g. in a private file system, email, or in a notebook). Structured information will be much easier to label than unstructured information, so it may be necessary to consider how information of value is being managed in general, in order to make labelling and handling it more feasible.

One method for labelling information, which is simple but very effective, is to specify that everything in a particular system, or environment, is automatically of a particular classification. This approach requires there to be a verification process at the point where information is introduced into the system, to make sure that information with a higher classification is not entered into the system, and at the point of data extraction, to ensure that it is labelled and handled effectively outside the system.

While the classification label, along with a handling scheme, defines how information should be handled, labelling related to confidentiality can also be used to indicate who should handle the information. Here the most important thing may well be that those who are not entitled to see the information should be able to immediately recognise that fact, return the information and report a security breach. Labels also need to make clear to those who are entitled to see the information who they may share it with. Provided that labels meet these twin requirements of being immediately clear to both authorised and unauthorised recipients, they can be relatively flexible. For example, information might be labelled with the department(s) where it should be used, or with the name of a project, event or function.

The "how" and "who" labels may appear together, for example as "SENSITIVE:Finance". "SENSITIVE:Finance" and "SENSITIVE:Physics". They must require the same handling rules in all departments, otherwise the security of the department's information may be breached by accident.

7.5 Handling

Each classification of information should have its own set of rules for how that information should be handled. Although many information management schemes concentrate on the confidentiality of information, rules should also address the organisation's requirements for integrity and availability. These too, must be consistent across different formats of information: making a written note of information from a conversation or phone call, and ensuring that work is not left on a single laptop or memory stick, both protect the availability of information. Ensuring that only authorised individuals can alter information, whether it is on paper or digital form, protects its integrity.

Information handling rules will probably have emerged during the development of the classification levels (especially if the advice above has been followed), but it is recommended that they be revisited after the initial decision on classifications, to ensure that they are appropriate, clear, and provide consistent risk management.

Each classification should relate to unique rules for how information with that classification is handled: if two different classification levels impose the same rules, information owners are likely to be confused about which classification they should apply, and users are less likely to understand how they should handle the information. A useful test for consistency is to consider which format of information a determined attacker would find it easiest to gain unauthorised access to: do they need to hack central servers or can they just hang around in the coffee room? With consistent handling rules, the difficulty (or ease) of unauthorised access should be about the same for all formats.

Once the organisation has established and agreed a consistent set of handling rules, it should look at how current processes require information to be used, to identify any inconsistencies. For example, if tender documents have been given the highest classification, but have to be sent to external assessors for review, then a "does not leave the building" rule will not work and the classification, and the handling rules, should be reviewed and revised as necessary.

Organisations should therefore expect to make a series of adjustment to classifications and rules as inconsistencies with either the organisation's risk or operational requirements are discovered. The goal should be a classification and rules that satisfy both.

7.6 Example handling scheme

Here is an example of a three-tier approach, which focuses exclusively on protecting the confidentiality of information. If integrity and availability are of particular interest (e.g. if the information in question is likely to be viewed on a website), then the handling rules should be extended to protect these attributes, in addition to confidentiality.

Table 3 - Three tier information handling scheme

	Classification 1 (no concern)	Classification 2 (slightly unsettled)	Classification 4 (genuinely scared)
Store, process and transmit - where	Anywhere	Premises of organisation or trusted third party (can take work home with minor precautions)	High security location (can't take work home unless you live in a bunker)
Store, process and transmit - how	Any method allowed	Only approved methods (e.g. encrypted, or via registered post)	Storage only in bunker. Processing with formal approval on high security systems. Hand transport by security personnel only. Face to face discussions only.

7.7 Documenting the scheme

Once an information management scheme has been designed, it must be documented, for the benefit of both information owners (who will be marking information with the relevant classification), and information users (who need to understand how to handle material with each classification).

One way to document handling rules, and to highlight the need for consistency across formats, is to start with the high-level risk the information needs to be protected against, then list the measures to be taken for each classification level and each format of information. For example:

Table 4 - Example documentation for handling rules

Risk	Information should not be seen/heard by unauthorised people	
	Sensitive	Don't leave papers lying around; lock your screen when you leave it; don't have conversations or phone calls in public places.
	Highly sensitive	Keep paper under lock and key; password-protect individual files; have conversations only in private offices.

7.7.1 Asset inventories

It is no longer required by ISO/IEC 27001 that an asset inventory be created or maintained. If the organisation, as a result of a risk assessment, decides that one would be appropriate, perhaps in certain areas, then this should be managed just as any other control, providing additional focus to the information management system.

7.7.2 The information management policy and process

An information handling policy and process are required that describe when and how the information management scheme should be applied. They should not only explain the classifications and labelling rules, but provide a process for classification- who should do it, how, and when? How should it be audited and verified to ensure that it is being applied consistently? The process should also address how to deal with situations where there are differences of opinion, and where information is found to have been incorrectly classified /labelled /handled (e.g. what steps should be taken to identify potential incidents).

A particular characteristic of educational organisations is that information may have different classifications and different points in its lifecycle. Both research data and exam results change their requirements for confidentiality and availability after publication, for example. The information management process therefore needs to be able to handle these time-related aspects. Lifecycles, like classifications, are best identified when the collection or creation of information is planned, whether the information is destined for publication, archiving or destruction. The Jisc model Records Retention Schedules may be a useful starting point.

7.8 Information management as part of an organisation's ISMS

An Information Management Scheme is actually a set of controls. In order to ensure that information classifications are applied appropriately, and that information is handled suitably, it is important to handle the controls in the context of the wider ISMS (also see Chapter 6, Controls). Here is an example of how this could work.

1. Roles are chosen to take responsibility for classifying, labelling and handling information (see Chapter 8, Roles and competencies) – this may include the information owner.
2. Decisions are made as to how the information management scheme is maintained and how compliance with it is measured (see Chapter 10, Measurement).
3. Top management are regularly informed of how the scheme is working, and whether there is anything they need to be aware of, or make decisions on (see Chapter 2, Information security governance).
4. Staff are trained in the content and use of the scheme (see Chapter 9, Awareness raising).
5. Problems with the implementation and running of the scheme are identified and addressed (see Chapter 11, When things go wrong: nonconformities and incidents).
6. Opportunities to make the scheme work better are identified, assessed and implemented if appropriate (see Chapter 12, Continual improvement).

Summary

- An Information Management scheme should comprise: a classification scheme, a labelling scheme, handling rules and processes to define how these all interact
- A classification scheme describes how information should be classified
- A handling scheme describes how information given a particular classification should be treated
- Don't have more classification levels than necessary, or practical
- Labelling can indicate both the classification and who should (and should not) see the information
- Handling rules must provide consistent protection across different media

[2]http://www.terena.org/activities/tf-csirt/publications/ISTLP-v1.1.pdf

Resources

Information Classification Scheme – University of York

Development of an Information Classification and Handling Policy – Cardiff University, case study

Information Classification and Handling Policy – Cardiff University

University Guidance on Classification of Information - University of Oxford

Reading list

UK Government information classification scheme
www.ucisa.ac.uk/ismt31
www.gov.uk/government/uploads/system/uploads/attachment_data/file/251480/Government-Security-Classifications-April-2014.pdf

CESG, Business Impact Level Tables
www.ucisa.ac.uk/ismt32
www.cesg.gov.uk/publications/Documents/business_impact_tables.pdf

JISC records retention schedule
www.ucisa.ac.uk/ismt33
http://bcs.jiscinfonet.ac.uk/