

Protecting Sensitive Information: Workshop Report

This report summarises the conclusions of a workshop organised by Janet CSIRT in March 2013 to discuss with invited participants how sensitive research and other information is protected in universities, and how existing best practice could be developed and expanded. We are very grateful to all those who contributed to the workshop and to this report.

1.Introduction: what to protect?

Universities and colleges process many different types of information on their computers and networks. Research or consultancy may mean handling information with commercial value (e.g. in chemistry, engineering, or research on advertising); research on offenders, victims or patterns of movement may involve information that could affect safety; staff may be acting as expert witnesses or performing forensic analysis that can affect the justice system; health research can affect patients and all those who may rely on treatments in future; credit card and bank details can have serious financial consequences; CCTV, telephony, staff and student records can all affect privacy, safety and the operation of the organisation.

Taking appropriate care of all this information is critical to the organisation. Inaccurate information about teaching, research, staff or students, can disrupt the organisation's business; losing confidential or sensitive information will affect the confidence of funders and partners when they next look for a safe place to obtain advice; failure to protect financial or personal data can result in significant fines. All of these damage the whole organisation, not just the area where the problem arose. A reputation as an unsafe place for information will be hard to overcome.

This presents a challenge to the organisation as it must protect information against unauthorised inspection, use or modification while ensuring that it is available to those who need to use it. This balance will be different for different information and different parts of the organisation. Universities and colleges are particularly unsuited to a one-size-fits-all approach, since their business may well require them to provide facilities both for highly sensitive information and student-owned laptops. Different security measures will be needed in different areas, with clear divisions when individuals move between them.

In assessing the appropriate measures to protect information it is often useful to consider scenarios: how could an activist, a disgruntled student, or a competitor obtain access to this information to do damage, and what would the impact of an accident be? Attackers will often use networks as part of their attack, but these will rarely be the only, or even most important, part. Most attacks now involve deceiving humans (by persuading them to give out information in a phone call, respond to an e-mail or visit a website), then use the resulting access to do further damage. Activists and competitors, in particular, may use precisely targeted messages to deceive their victim. Only a combination of policy, awareness, behaviour and technology can defend against these types of attack.

Universities and colleges already use technical tools to protect themselves and their information from security problems. The challenges are to identify areas where these may be particularly required, and to expand the known best practices of responsibility, behaviour and technology to meet their particular requirements.

2. Standards

Grant-awarding bodies and other external bodies are increasingly referring to security standards when setting conditions for handling their data. This means both that organisations need to show the required conformance to the standard in order to receive a particular grant/award/contract, and that failure to meet the requirements of the standard may put future awards or contracts at risk. Three commonly-referenced security standards are:

- **ISO/IEC 27001** Requirements for Information Security Management Systems (ISMS). This is a high-level standard, setting out how organisations should go about identifying risks and appropriate measures to control them. It does not prescribe particular risks or measures, though various subsidiary standards (e.g. ISO/IEC 27002, 27011, 27015) suggest lists of measures that may need to be considered in particular settings. Although it is possible to have an ISMS formally accredited against ISO/IEC 27001, most external bodies currently require conformance to the standard, which can be assessed internally;
- **PCI DSS** Compliance with the Payment Card Industry Data Security Standard is a requirement for anyone handling payment card data. The standard has detailed requirements for both organisational and technical security measures;
- The NHS Information Governance Toolkit (**IGTK**) is usually referenced when handling health data. It falls between ISO27001 and PCI DSS in level of detail: mostly requiring that organisation have policies/processes/contracts covering relevant information security areas, rather than specifying detailed measures.

All three standards contain the idea of scope: those people, information and systems “in scope” are subject to the security requirements of a particular standard, those “out of scope” are not. An area in scope for a particular standard (such as an area handling payment card data) must be protected against threats leaking in and will have strong perimeter security to divide it from out of scope areas. This perimeter also ensures that the measures needed to satisfy the requirements inside the scope do not have to be applied to the whole organisation. In the security literature, such areas are often described as “security zones”.

The variety of activities within a university makes it very unlikely that a single set of security

requirements will cover the whole organisation. Instead there will be a number of different in-scope environments within the organisation, each relating to a given standard, with its own specific requirements and policies to address them.

3. Architectures

The key to making security usable is to identify the right boundaries between security zones. Within each zone, a consistent set of security rules apply to technology, behaviour and systems governing access to the zone. In a different zone, a different set of rules may apply to the information, services and people within that zone. Boundaries between zones therefore act both to protect information and services from threats outside the zone, and also to limit how broadly that zone's technical and behavioural rules need to be applied. For example, if accessing particular information requires individuals to satisfy prior background checks, it would be highly inconvenient and expensive for this to apply to the whole organisation. Much better to define a small security zone around the information so only those with access to that zone need to undergo the checks.

Historically, networks were sometimes seen as containing only two zones – “inside” and “outside” the organisation. Since “outside” was considered to be more dangerous, the boundary between the two was strongly enforced by membership, rules and firewalls. Once inside the boundary, however, every person, system and information was subject to the same rules. This was sometimes known as the “igloo” model – a hard shell with a soft interior.

It is not obvious that that model was ever appropriate for universities; however, the range of uses of their networks now makes its assumptions even more obviously false. Parts of the organisation will often need to allow students, staff or visitors to connect their own computers (and those users will expect this to be at least as easy as using wifi in a hotel): other parts will be handling sensitive information with contractual or legal requirements for how it is protected. The result now looks more like a pomegranate: there is still a perimeter/skin, because universities still have rules that do not apply to the Internet (though some partner organisations may appear within the perimeter), but within the organisation there are many more security zones, at least some of which will have significantly different rules and strongly enforced boundaries. Some of those zones may be sufficiently different from the norm to require their own de-militarised zones (DMZs) where information flows can be controlled and security threats isolated. Most of these internal zones will have higher security protection, but some (for example zones supporting guest machines) may be lower, and so may have a boundary designed to protect the rest of the organisation **from** them.

Zones are likely to have a technical implementation, but it is at least as important that the change of rules is matched by a change of behaviour by authorised individuals working “inside” the security zone. For example, technology cannot prevent someone talking about what they have seen on a computer screen or taking home a paper document. Only the understanding of the individual that they are in a specially protected zone and must behave appropriately can protect the information and systems it contains. Maintaining this discipline may be particularly hard when different security zones can be accessed from the same place: having to go to a different (locked) location may actually be a useful subconscious reminder that different rules apply here. Computer interfaces should aim to provide similar instinctive cues, without making use unreasonably inconvenient. Above all, individuals need to feel that security zones exist for their own benefit – to protect their data, contract, research or career –

and are not imposed arbitrarily by the organisation.

4.Tools

A wide range of technical tools is already used in universities to protect sensitive information. However, these can only be effective if accompanied by suitable **policies, awareness** and **behaviour**. There is little point in protecting information on computers if users then read it or disclose it in mobile phone calls in public places. Indeed, technical protection that information users perceive as unnecessary can increase the risks to information, by encouraging users to implement 'work-arounds' such as putting information on portable devices, home computers or consumer cloud services.

Physical separation: some information may be so sensitive that it can only be processed in physically secure locations. Secure data centres, locked machine room cabinets and offices only accessible to authorised users are all possible options. However, they must be accompanied by matching restrictions on access to those locations by guests and cleaners, and restrictions on the transfer of information and equipment into and out of secure areas.

VLANs, ACLs & Firewalls: network configurations can be an effective way to create technical security zones. Placing computers in the same zone into their own Virtual LAN (VLAN) prevents others seeing their network traffic; router ACLs can then be used to allow simple traffic to flow between VLANs with firewalls providing more advanced traffic management features.

Encrypted VPNs: if the computers in a security zone are in different locations then encrypted Virtual Private Networks (VPNs) can be used to carry traffic between separated parts of the zone across untrusted networks. For example if servers are placed in a remote data centre an encrypted VPN can connect them to PCs in a local security zone.

Virtualisation: virtualised servers can be much easier to create and securely destroy than physical ones, and standard secure configurations can be stored and reused. Although there have been concerns in the past about the possibility of information leaking between virtual servers, technical measures to prevent this seem sufficient to satisfy many security requirements.

Device configuration management: as with servers, standard secure configurations for client devices (PCs, laptops or mobile) may be required in secure zones. Universities often use device management for open-access terminal rooms and standard desktops. For secure zones the processes for testing and adding new software need to be sufficiently fast to deal with new user requirements. The IEEE 802.1X protocol can be used on both wireless and wired connections to place devices into appropriate zones based on their type (e.g. laptop, smartphone, printer), status (e.g. systems without up to date patches and anti-virus can be temporarily assigned to a quarantine network until they have updated themselves), or who is logged in. For mobile devices this can be combined with automated wireless network configuration.

Portable device encryption: portable devices (including USB drives, smartphones and laptops) risk being lost or stolen, and need to protect the information and access credentials they may contain. Encryption of these devices should be regarded as normal, as indicated by the monetary penalties being levied by the Information Commissioner on organisations that

lose unencrypted laptops.

Access control: users need to prove their claimed identity before being granted access to non-public systems or information. Static passwords are likely to remain the most common way to do this, but need to be protected by behaviour and technology against common attacks such as network sniffing and phishing. For access to sensitive information and systems, access control methods that avoid re-usable credentials, such as two-factor authentication, should be considered.

Disposal of assets: equipment that has contained sensitive information needs at least to be securely deleted before it is disposed of. For more sensitive information, and for devices that cannot be reliably wiped, physical destruction may be required. Far too much sensitive information has ended up on disks sold on eBay.

5.Examples

a.Commercial information: engineering

Several universities have developed specialist facilities to support research and manufacturing work. External partners are essential to this success as the key is not just to research new technology but to build it. However these partners may be commercial competitors, so secure perimeters within the facility, as well as between the facility and the rest of the organisation, are required. Depending on the partner and project, there may also be national security concerns. Within the facility each partner, and indeed individual projects, will define the level of security required. For some the lowest common denominator of security offered by the facility (building access) and University infrastructure (policies, account control) are sufficient. For others additional controls have been put in place, for example completely standalone IT equipment with no connection to the facility's network. By clearly defining the risk presented by each project competitors are able to share facilities that would be cost-prohibitive to develop individually.

b.Government information: modelling and forecasting

When using financial, governmental or climate data for undertaking detailed modelling and forecasting, the key component for protecting the information involved is often data integrity rather than confidentiality. Assurance must be provided that the figures used are timely, accurate and have not been altered in any undocumented or uncontrolled manner. Doubts about data integrity in such projects can lead to perfunctory dismissal of the findings and a loss of reputation for the host institution. Methods of providing such assurance include deploying ACLs to ensure only appropriate people can change data, irrespective of their ability to read it (for instance, by using the Biba Integrity Model), mandatory document version control, and clear access and alteration audit trails.

c.Financial information

The Payment Card Industry Security Standards Council set strict security standards ^[1]for anyone handling credit card details, covering the security of computers, networks and people;

processes for managing vulnerabilities, change control, monitoring and testing security are also required. Where individuals have access to large numbers of card details background checks may be required for criminal convictions and other relevant issues. To achieve the required technical standards, credit card processing is often done on dedicated PCs (virtualised systems may also be permitted if the underlying hypervisor is sufficiently secured) with encrypted local networks or point-to-point connections connecting them to card terminals. Access to these systems from other networks must be restricted by router and/or firewall configurations that are subject to effective change control. Card data may be transmitted to payment service providers using dedicated communications lines or encrypted VPNs over the public Internet. Physical access must also be controlled, for example by ensuring that visitors can be identified and their presence recorded.

d.Security information

Groups dealing with the security of networks and systems, and handling incidents relating to them, use information that could be used to cause serious harm to those networks and systems and to their users. They also need to communicate with system and network managers within the organisation and with external incident response teams such as Janet CSIRT. Security teams normally have their own servers and networks firewalled off from the main university network – these may also contain internal firewalls to allow activities such as testing security patches to be done without risking the security of the rest of the network. Clients and servers are securely configured and managed; access to them may well require two-factor authentication. Communication between team members and between teams is encrypted using VPNs and encrypted e-mail. Staff working on security incidents may be subject to special procedures or contract terms. Physical access to offices may be restricted, for example by requiring cleaning to take place when team members are present.

e.Personal information

Research on health data about individual patients is subject to NHS security requirements as well as being regulated as sensitive personal data under the *Data Protection Act 1998*. Protecting this information is likely to involve securing servers, clients and networks, as well as clear policies for those with access. One project has its own dedicated servers both to store and process the data, as well as its own authentication, directory and web servers. These are located in a university data centre providing a secure environment with power and air conditioning, physical security and access control and monitoring; the health systems are in separate cabinets with their own lock. To reduce the risk of data being placed on mobile computers, staff are provided with managed desktop computers in their secure offices. These connect via the main university network to the firewall protecting access to the servers. For out of office use there is a pool of laptops with tightly secured settings: users are prevented from installing software on these. E-mail, provided by the university's standard outsourced service, may be used to discuss the project but staff are prohibited from using it for any research data.

6.Protection as a Service

As the examples above illustrate, at present there is a tendency to create security zones individually around particular information or services, using local resources and configuration.

This may be appropriate for a very small number of zones, but is likely to lead to increased costs and inconvenience in the longer term. Security benefits from considerable economies of scale (through automation, facilities, and skills) so universities should aim to develop standard approaches that can be re-used. A standard approach should also result in new zones being much quicker to set up, once the security responsibilities and requirements have been identified. A standard approach is also less likely to suffer from incompatibility problems when individuals need to work across more than one zone.

A vision for “protection as a service” would be for a Principal Investigator, Finance Director or other risk owner to be able to request a security zone around information or services that need particular protection. A central service could then generate appropriate policies for them to adopt, and provide advice on tailoring them; once the policy was agreed, the appropriate configurations of network devices, (virtual) servers and end-user systems would be made automatically. Such a system would make protecting information efficient and effective for information users, risk owners and their organisations.

7. Conclusions

Several different factors may require universities to establish secure enclaves to protect particular information and services. These include the requirements of funders and research partners, external regulation, and protecting the university’s reputation. However universities also need to be able to provide general-purpose networks and services that are easy for staff, students and visitors to access. Current examples indicate that it is possible to provide appropriate levels of security within a university environment, provided that sensitive information and services are identified and security zones with specific policy, behavioural and technical measures defined around them. At present this is largely done on an individual basis for each enclave, an approach that is unlikely to be scalable in either cost or effort. Technologies such as virtual networks and servers allow secure enclaves to be provisioned more quickly, consistently and efficiently; universities need to develop systems and policies that combine these with efficient identification of requirements and appropriate policies.

Source URL: <https://community-stg.jisc.ac.uk/library/janet-services-documentation/protecting-sensitive-information-workshop-report>

Links

[1] https://www.pcisecuritystandards.org/documents/pci_ssc_quick_guide.pdf