Technical security for e-infrastructures

24 November 2014

“Technical security for e-Infrastructures ”

© Jisc
Published under the CC BY 4.0 licence
creativecommons.org/licenses/by/4.0/

Contents and introduction

E-infrastructures are large computer systems with considerable processing and storage capacity and in some cases, holding valuable or sensitive data.

They are therefore likely to be attractive targets for attackers with a wide range of motivations. However, to support international research, e-infrastructures must be accessible to users located anywhere on the Internet. In many cases users will upload and run their own software or virtual machines and exchange large volumes of data over high-speed networks. Operators of e-infrastructures are therefore challenged both to provide the open and flexible computing platform that is inherent to the e-infrastructure concept and to protect against the consequences of attacks on that platform over the Internet. To help them, the e-infrastructure model offers many different ways to implement security controls. This paper reviews the security measures used by e-infrastructures against a widely-used model – the Cyber-Security Council's Top 20 Controls – to assess what is being done and where improvements may be possible.

The Cyber-Security Council model

The Cyber-Security Council's Top 20 controls[1] are widely recognised as a useful framework for discussing how to protect organisations and systems that are connected to the Internet.

The controls are not intended to be auditable, but provide reassurance to operators and users that measures have been taken to address the most prevalent threats. Furthermore, the Top 20 controls only cover threats arising from connection to the Internet: other threats, such as misuse of systems by authorised users, may require different controls. It is recognised that particular circumstances may involve both different threats and different approaches to dealing with them[2]. Organisations should use information risk assessments to determine whether each control objective is relevant and cost effective to the system being considered. Where organisations choose not to implement one of the suggested controls, either because the threat does not apply or because it has been mitigated in another way, this should be noted and periodically reviewed.

Applying this process to a typical e-Infrastructure one significant difference is immediately apparent. The Top 20 Controls presume that the organisation to be protected is defined by a "network" connected to the outside world at a point that can be protected by firewalls and other networking devices. The perimeter between "inside" and "outside" is principally controlled by the configuration of those networking devices. The Controls therefore divide into one group covering the establishment and management of the perimeter and a second group that apply within that perimeter: for example control CSC1 says "Actively manage (inventory, track, and correct) all hardware devices **on the network**". "Outside" the perimeter (in other words "not on the network") is considered to be beyond the organisation's influence so no controls can be applied there.

By contrast, in the e-infrastructure case users may be globally distributed anywhere on the Internet; processing and storage devices may themselves be in different network domains, connected across a shared public network. The "perimeter" that separates controlled and uncontrolled zones cannot therefore be defined as the boundary between two different networks. Instead the main technical perimeter of an e-infrastructure falls between its administrators, who may themselves be distributed across multiple locations rather than being on a shared network, and its users. Users, who will mostly be outside the technical perimeter, can be subject to some controls through policies on how the infrastructure may be used. Controls such as awareness and participation in incident response can be effective outside the technical perimeter, however, it is not possible to use ability to connect to a particular network to distinguish authorised users, who are subject to controls, from others who are not.

This discussion therefore first considers the measures that e-infrastructures should take to establish a technical perimeter between the zone containing people and systems that have administrative rights and the zone containing users who do not. This involves both the Top 20's "perimeter" controls and additional technical options likely to be available to e-infrastructure designers. Then we consider how the Top 20's "internal" controls apply to users and devices in the administrative zone; then which "internal" controls can be applied to all authorised e-infrastructure users, even if they are outside the technical perimeter.

[1]

1 counciloncybersecurity.org/critical-controls/
2 tripwire.com/state-of-security/featured/threat-mitigation-and-the-20-critical-security-controls-with-tony-sager

Implementing the perimeter

Whereas on a traditional network a firewall marks the boundary between the zone that the organisation controls and the zone it does not, the situation of an e-Infrastructure is much more complex.

Users, who may be anywhere on the Internet, will often provide their own software to run on the infrastructure; some infrastructures may even allow users to install their own virtual machines and operating systems. Although there may be policies regulating what is installed, the e-infrastructure operator has less ability to apply preventive controls to this software or virtual machines. Different types of perimeter control are therefore needed to manage the risk that such installations may pose to the rest of the e-infrastructure. Risk can be reduced either by preventing security problems or by containing their impact (or both): for example user code with a higher risk of vulnerabilities might be run inside a sandbox that reduces the scope for security problems to spread to other systems or users.

Compared to the network-based model of perimeter security, the e-infrastructure model and technologies provide a rich set of controls that can be used, often in combination, to maintain risks at an acceptable level. Thus even if one approach is unsuitable for particular circumstances there should be substitutes available that can achieve an equivalent risk reduction. Designing a perimeter for an e-infrastructure is likely to involve selecting measures from a menu of possible controls at different levels in the infrastructure stack, for example, separation may be achieved at physical, network, operating system or software layers. It may also be possible to compensate between preventive and responsive controls, as in the sandbox example above. It is unlikely that all controls will be appropriate for all circumstances: the essential thing is to select a consistent group of controls that deliver the required protection against threats for each security zone.

This section first considers some of the controls likely to be available to e-infrastructure designers, both from the Top 20 set and elsewhere. This is followed by examples of how these may be combined into a consistent perimeter.

» **CSC10 Secure Configurations for Network Devices such as Firewalls, Routers and Switches:** Although e-infrastructures cannot rely solely on network devices to establish the perimeter, these still perform an important function and must be configured and operated securely to prevent unauthorised modifications. For example access to the infrastructure's management interfaces and services should still be protected by network device settings

» **CSC11 Limitation and Control of Network Ports, Protocols and Services:** Some e-infrastructures only support pre-selected communications protocols: others may need to allow users to install or implement their own. Where users are allowed to open new ports or protocols, the operators should ensure (and test, for example by internal and external scanning) that the port ranges available to users are not used by any system services, and vice versa. This complements the separation between user and system privileges within the infrastructure components. Control of ports and protocols – which may be implemented within e-infrastructure components, network devices, or both – may also restrict the ability of malware and other intrusions to propagate within e-infrastructures, even if one program or process can be compromised

- **CSC12 Controlled Use of Administrative Privileges:** E-infrastructures should use group and process permissions to establish perimeters both between user and administrative zones and between individual users and groups. If it is necessary to grant administrative privileges to user-provided software or operating systems then other controls – such as port restrictions, sandboxes, air gaps and intense monitoring – must be used to establish a perimeter around these privileged applications and reduce their capacity to damage others

- **CSC13 Boundary Defence:** As discussed above, the trust boundaries of e-infrastructures are more likely to fall between different processing activities and datasets than between different networks. Where trust levels are similar (but not necessarily identical), permissions and access control mechanisms should provide sufficient separation; for more sensitive applications it may be necessary to provide separate processing and storage facilities, supported by secure locations, management and access controls

- **CSC19 Secure Network Engineering:** The requirement that e-infrastructures be accessible to users anywhere on the Internet challenges traditional security architectures that rely on a network perimeter. However, network designs can still contribute to security by providing appropriate separation between different users of the infrastructure (in particular between administrative and user tasks). A proposed network design – the Science DMZ – that separates the e-infrastructure from both the Internet and the local network of its host site(s) is discussed below

- **CSC6 Application Software Security:** Where users of the e-infrastructure provide their own software, controls within the software may contribute to maintaining the infrastructure's perimeter. For example, it may be a requirement, before granting relaxation of port/protocol controls (CSC11) or increased access to privileges (CSC12), that software be audited or follow approved design methods to reduce the risk that it will be compromised. Improving the quality of research software is an explicit goal of the UK's e-infrastructure policy; a Community of Research Software Engineers[3] has been established to develop and disseminate good practice for the profession

- **CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers:** As for software above, where e-infrastructure users are able to install their own virtual machines or operating systems, audit or other approval processes may be required before these are granted access to ports, protocols or privileges

- **Virtual Machine (VM) Sandboxes:** Many e-infrastructures provide their users with virtual machines rather than access to the underlying hardware. Virtual machine configurations provide an additional layer of fine-grained, per-user or even per-process control over many of the Top 20 controls. For example, VM configuration might limit the range of ports available to send and receive communications (CSC11/19), the permissions available both within and between VMs (CSC12) and the amount of system resources a VM can consume

- **Monitoring and alerting:** Most e-infrastructures perform detailed monitoring of the activity of processes for accounting and performance purposes; many can also raise alerts when unexpected behaviour occurs. These will often allow operators to respond quickly, in some cases even automatically, to security incidents by suspending or terminating a process or component that appears to have become a threat

[1]

[3] rse.ac.uk

- **VM Freezing:** The ability of a Virtual Machine platform to temporarily "freeze" the operation of a VM may be preferable to terminating a suspect process, since if the action is found to be legitimate the VM's operation can be resumed without losing information. It also lets a virtual machine be saved and subjected to detailed forensic investigation should this be necessary

- **Physical separation:** The controls provided by e-infrastructures are in addition to traditional approaches using physical separation of hardware and logical separation of communications. Where the sensitivity of work done on an e-infrastructure requires it, separation can provide additional assurance. However, this inevitably comes at the cost of some decreased flexibility and increased inconvenience for users, since the barriers to gaining access to and using these high assurance infrastructures must also be raised to avoid them becoming the weak link in the infrastructure's security. Face-to-face identification and two-factor authentication are likely to be required for access to patient data, but are likely to be unsuitable for e-infrastructures used by citizen scientists and researchers on less sensitive data sets

- **Change control**: Where an e-infrastructure relies on a particular configuration, feature or software to maintain part of the perimeter, a change control process should be used to ensure that it is not changed unexpectedly or unintentionally. Change monitoring can also be used more widely to detect unexpected changes that may indicate a security problem

Examples of composite perimeters

There are many ways to use these controls to establish a perimeter between administrative and user zones on an e-infrastructure. The choice will depend on the technical approach taken by the infrastructure and the threats to which it, and the data and processing it contains, are exposed. Some e-infrastructures will need to establish additional perimeters between different groups of users of the same underlying systems. The following sections describe, as illustrations, some of the approaches that have been taken. Others are possible. In each case, prompt detection of incidents and an effective response to them is likely to be an important aspect of the security of the e-infrastructure, its users and data, and the networks and systems around them.

E-infrastructure security

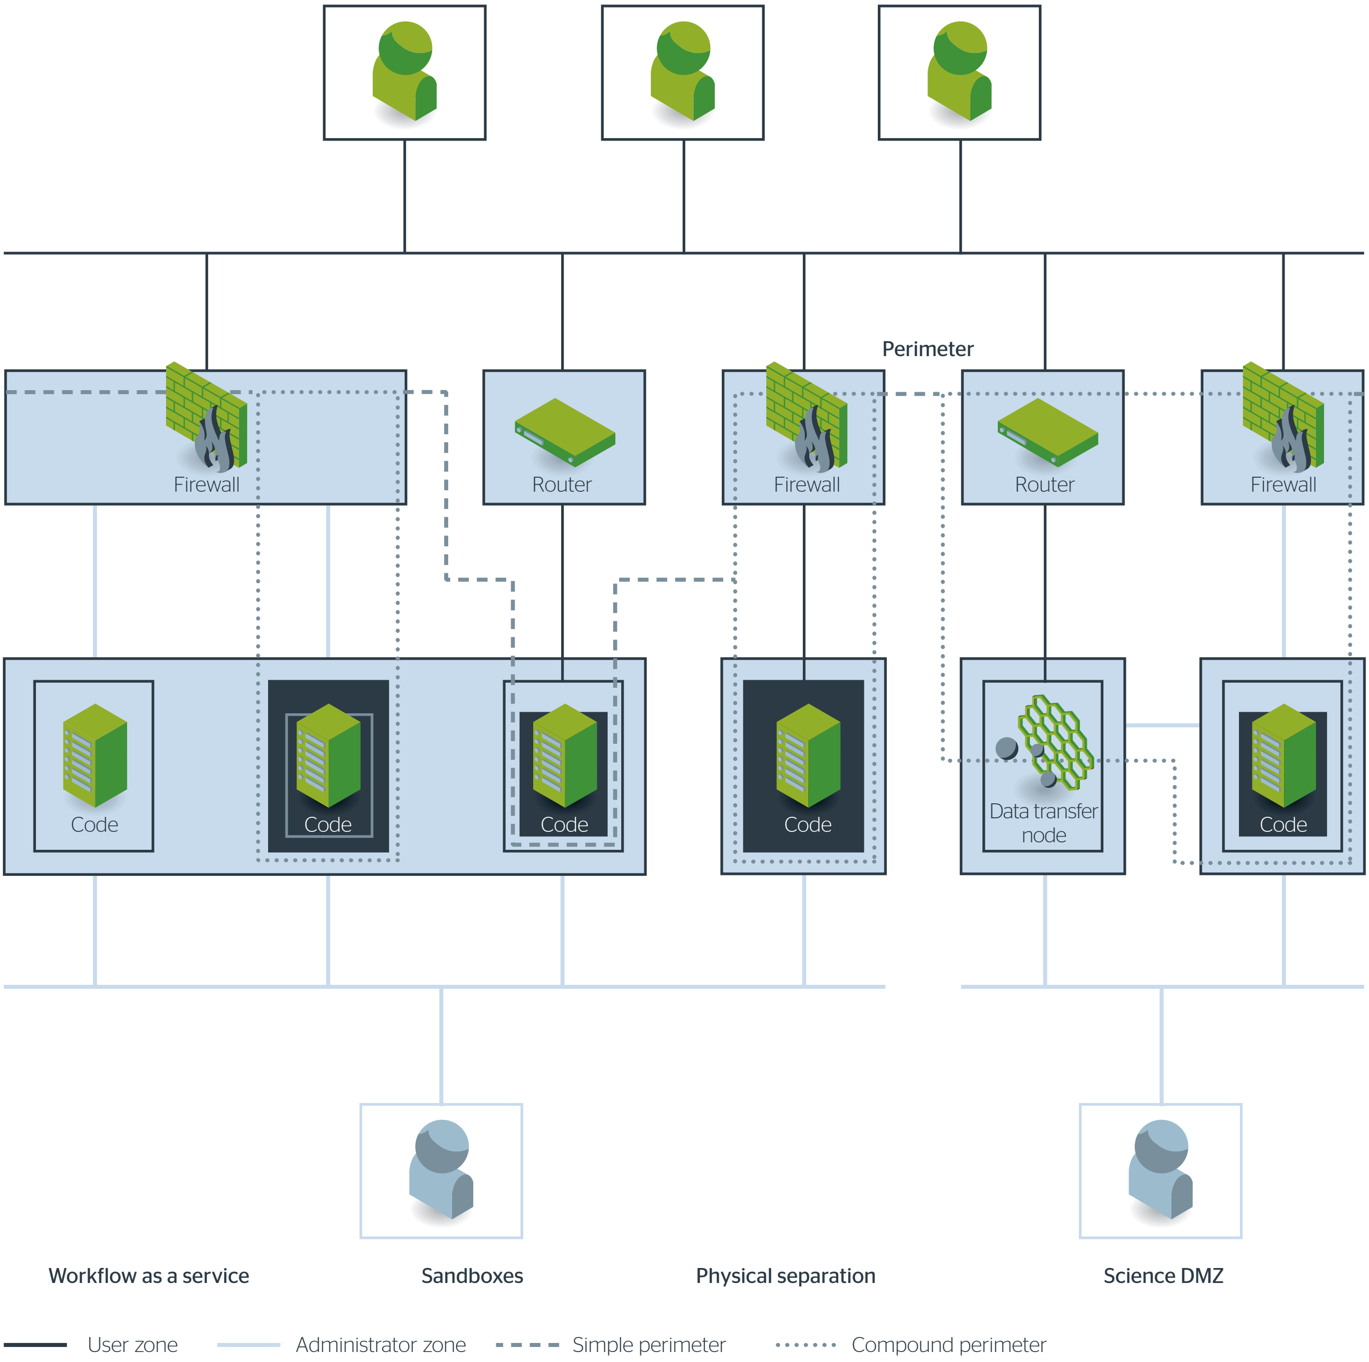

Workflow as a Service

Some e-infrastructures provide services to users at a very high level. Rather than a command line interface where users can install and run software, these provide pre-selected libraries of software that users can combine into workflows (perhaps using and modifying workflows developed by others) to analyse and display data. These services typically offer a web interface, so a security design like that of a traditional outward facing web server may well be appropriate. Only a limited range of inbound protocols from users needs to be supported (perhaps just HTTP and HTTPS); others can be blocked by a router or firewall. The operating system and web server processes must be securely configured and managed with the web server having minimal privileges; pre-selecting software components lets them be checked for security weaknesses. Access for administrators can be provided through a different protocol such as an authenticated VPN that is only accessible from specified IP addresses or ranges.

The perimeter in this familiar configuration is comprised principally of CSC11 Limitation and Control of Network Ports, Protocols and Services, CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers, and CSC6 Application Software Security.

Sandboxes

Where an e-infrastructure lets users install their own software or operating systems there are two main concerns: that the installation may itself damage the security of the infrastructure (for example by consuming excessive resources) and that it may have vulnerabilities that let others attack the infrastructure or its users. User installations are likely to require segregation on all sides: from the administration of the infrastructure, from other users, and from the external network. This is normally done by placing the user installation in a “sandbox”. E-infrastructures provide many different controls that can be combined to establish suitable sandbox perimeters; different requirements may suit a different balance between these controls.

For example:

» If software runs in an unprivileged user account, normal operating system file and process permissions may provide sufficient protection within the infrastructure. Since any required network connections will be known in advance these ports can be configured into an external router or firewall; the open ports range should be scanned to ensure no other (user or system) processes offer ports in the same range

» Where user-installed software requires privileged access (“root” or other super-user privileges) controls within the operating system are less effective. Instead, processes may be run within a virtual machine container that can only access dedicated file systems, not shared with other users, and does not let the software receive connections from the external network

» In a few cases users may wish to set up their own networks of virtual machines running on different e-infrastructure services and communicating across the networks that connect them. As well as sandboxes surrounding each endpoint, these applications may need to run their own security, logging and incident response processes if details of process and network activity are not visible to the infrastructure operators

Sandboxes are likely to be established using layered CSC controls, typically involving complementary application of CSC11 Limitation and Control of Network Ports, Protocols and Services, CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers, CSC6 Application Software Security, CSC12 Controlled Use of Administrative Privileges, and CSC13 Boundary Defence.

Physical separation

Some e-infrastructures may handle sufficiently sensitive information that additional security precautions are required. For example, some types of data may involve policy and ethical requirements on the individuals granted access and physical security requirements on the locations from where they access it. These situations are likely to require a perimeter that extends around both the client and server ends of the connection: if there are requirements that users be vetted or sign agreements before obtaining access and that access be from physically secure locations then equivalent precautions will be needed for the server location and those who may have administrative access.

This perimeter also needs to be maintained in the arrangements for digital storage and transfer: transfers over networks should use end-to-end encrypted tunnels, with securely configured and managed computer and networking equipment at both client and server ends. If the server is used only for storage then it may be possible to keep information in encrypted form with decryption keys held elsewhere. If, however, the information needs to be processed in unencrypted form then it will be hard to achieve a sufficiently secure perimeter between it and other, less sensitive, uses of the same equipment. Physically separate hardware is likely to be required in these cases.

Science DMZ

For e-infrastructures requiring extreme high-speed transfers of data across wide area networks, complex firewalls may create too great a risk of packet loss and other performance problems. The "Science DMZ" network architecture that has been used to address this problem shows how a composite perimeter can provide the necessary performance and security.

The key innovation in the Science DMZ[4] is to separate data transfer and data processing. Dedicated Data Transfer Nodes (DTNs) are placed outside the organisational firewall as endpoints for high-speed data transfers. These DTNs run only the software and protocols that are needed for this transfer function. Since the protocols are well defined, ACLs on a high-performance router can block any other protocol between the DTN and the wide area network. DTNs run only their single-purpose software, so can be designed and tested to reduce the risk of compromise through the limited range of open ports.

Alongside the DTN, connected to it by dedicated network links, are the infrastructure's storage and processing systems. The DTN places received data on the storage system, making it accessible to the processing system. These two systems, which may need to run a wider range of software, can be protected from external networks using a firewall since their communication rates are much lower than the DTN's. Storage and processing systems still need to be securely configured and maintained in case attacks from local or wide-area networks are able to pass through the firewall.

The local organisation's network can be connected to the DTN, storage and processing systems through the organisation's normal perimeter firewall, since the low latency of these short-distance connections should not produce the same performance issues as using a firewall on a higher-latency WAN. Instruments, general purpose computers and other equipment whose security cannot be proactively managed should also be connected through a firewall into the Science DMZ.
The Secure DMZ model thus constructs a perimeter primarily using CSC11 Limitation and Control of Network Ports, Protocols and Services, CSC19 Secure Network Engineering, CSC6 Application Software Security, CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers, together with Physical Separation of networking functions and traffic flows.

[1]

[4] fasterdata.es.net/science-dmz/science-dmz-security

Within the perimeter

With a perimeter established between the users' and administrators' zones of access to the e-infrastructure, most of the Top 20's "internal" controls can be applied within the administrator zone as a matter of good operational security.

In most cases administrators will access the e-infrastructure components over some form of private network connection, though this may be implemented in different ways, including Virtual LANs, Virtual Private Networks, or physically separate network segments. Whatever technology is used for this private connection, it should be taken as the "network" to which the Top 20 controls are applied.

The controls may need minor adjustments where the e-infrastructure or its administration are distributed across multiple locations, for example where several organisations provide servers or data storage to the infrastructure. In this case there may be no single network or organisation that connects all administrators; central control of security will instead be through the agreed policies that those managing individual infrastructure components are required to follow. This may result, for example, in separate registers of authorised devices (CSC1) at each location rather than a single central register, not least because a device (or user) authorised to act as an administrator at one location may well not be an authorised administrator at another.

» **CSC1 Inventory of Authorized and Unauthorized Devices:** All devices connected to each administrative network should be known. New devices have to be added to the local inventory and network configuration before they can send or receive traffic on the administrative networks

» **CSC2 Inventory of Authorised and Unauthorised Software:** Any device permitted to connect to the administrative network should have a hardened 'management workstation' configuration

» **CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers:** As above, any device on the administrative network should have a managed, hardened configuration

» **CSC4 Continuous Vulnerability Assessment and Remediation:** Devices authorised to connect to the administrative network should be supported by vulnerability assessment and remediation processes

» **CSC5 Malware Defences:** Devices authorised to connect to the administrative network should run relevant and up to date malware defences, including safe user behaviour

- **CSC6 Application Software Security:** Software applications used within the administrative zone should be selected, configured and maintained for security. Application software security is also a control used in establishing the perimeter, so is described further above

- **CSC7 Wireless Access Control:** It appears unlikely that there should be any need to provide wireless access to the administrative network

- **CSC8 Data Recovery Capability:** Information and device configuration for the administrative network should be securely backed up. Note that the use of virtual machines can provide particularly efficient data recovery

- **CSC9:** User control, see next section

- **CSC10, 11, 12, 13:** Perimeter controls, see previous section

- **CSC14 Maintenance, Monitoring and Analysis of Audit Logs:** Audit logs from devices and software in the administrative zone should be collected and monitored for security events

- **CSC15 Controlled Access Based on the Need to Know:** Access to the administrative network and zone and information about the zones should be limited to those responsible for managing the e-infrastructure

- **CSC16 Account Monitoring and Control:** Accounts for access to the administrative network and tools should be actively managed, being created and deleted as required

- **CSC17, 18:** User control, see next section

- **CSC19:** Perimeter control, see previous section

- **CSC20:** User control, see next section

Outside the perimeter

The Top 20 controls, being designed for a network perimeter surrounding a single organisation, treat all people and systems outside the perimeter as outwith the organisation's control.

This is not true of e-infrastructures where users outside the technical perimeter are still subject to the control of those who operate the infrastructure and are responsible for the data used on it. The following controls are applicable to these users:

» **CSC9 Security Skills Assessment and Appropriate Training to Fill Gaps:** This is widely recognised as a significant challenge. Much good work is being done, for example, in developing the skills of those who operate e-infrastructures, those who write code for them,[5] and in research areas that handle sensitive information. However, establishing a consistent formal framework that addresses the needs of everyone from infrastructure operators to researchers and data providers will be a long-term activity

» **CSC17 Data Protection:** E-infrastructures handle such a wide range of data of different sensitivities (from astronomy to medicine and beyond) that a standard approach to data protection is unlikely to be appropriate. Instead infrastructures should provide appropriate tools for researchers to manage the security of their data in accordance with contractual and ethical requirements

» **CSC18 Incident Response and Management:** The openness and flexibility of e-infrastructures would be severely limited by disproportionate preventive measures. The ability to detect incidents quickly and deal with them effectively is therefore an important compensating control. Infrastructures should provide researchers and operators with monitoring and alerting tools to detect unexpected behaviour by systems and processes. Technical and policy measures allowing prompt suspension of particular programs, users, systems or sites are essential to limit the impact and spread of security problems. Incidents may also be detected by local and national network providers so these should be included in incident detection and response processes. The EGI project's work on incident response is being disseminated and developed through the Security for Collaborating Infrastructures[6] activity

» **CSC20 Penetration Tests and Red Team Exercises:** These may be a useful way to discover vulnerabilities, though on an open e-infrastructure it may be difficult to define "success". For infrastructures handling sensitive data penetration tests may already be a requirement

In addition, the following controls, though mainly applied within the perimeter, may also be suitable for external user processes and software.

- **CSC3 Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers:** E-infrastructures that provide their own operating system, virtual machine, or software installations should also maintain secure configurations for those. Some e-infrastructures are considering auditing virtual machines provided by their users
- **CSC4 Continuous Vulnerability Assessment and Remediation:** E-infrastructures should apply vulnerability assessment and remediation to the operating systems, virtual machines, or software interfaces they provide to users. Some also provide users with information about vulnerabilities in commonly used software and limit access to the infrastructure if serious vulnerabilities are not patched promptly
- **CSC5 Malware Defenses:** Tools to monitor inappropriate propagation of jobs running on the e-infrastructure may also be able to detect signs of malware propagation
- **CSC14 Maintenance, Monitoring and Analysis of Audit Logs:** Audit logs may also be used for accounting and analysis of user code behaviour, including security events. EGI run a series of security challenges to ensure their logging, monitoring and analysis processes are effective
- **CSC15 Controlled Access Based on the Need to Know:** User data may be subject to different policy requirements according to the type of information being processed. This may drive different technical and procedural security controls, for example medical data may require separate hardware with stronger approval and authentication checks of both researchers and operators
- **CSC16 Account Monitoring and Control:** Tools used by researchers to administer groups of collaborators may provide account lifecycle controls. Where the e-infrastructure uses home sites for federated authentication, this will benefit from account lifecycle management by those home organisations

[1]

5 For example: software.ac.uk
6 eugridpma.org/sci

Conclusion

Although e-infrastructures appear very different from the organisational networks for which the Cyber-Security Council controls were designed, the controls nonetheless provide a useful framework for planning technical security.

The model of security zones separated by perimeters is still relevant but, unlike traditional organisational architectures, the perimeter of an e-infrastructure is not a point on a network. Instead the perimeter runs through different devices controlled by the e-infrastructure – including routers, firewalls, physical and virtual machines – increasing the range of controls available and the opportunities to implement layered security measures.

Many security controls are already used by e-infrastructures: the challenge is to ensure they are consistent with one another and with infrastructure users' requirements and expectations for access and security. As with any security architecture, the overall approach to security and choice of particular measures should be guided by an assessment of risk. The annex to this paper suggests some threat scenarios that can be used to identify risks and controls relevant to individual e-infrastructures. In particular, risk assessment should inform the balance between preventive and reactive approaches: infrastructures that give their users freedom to upload and run software need a quick and effective response when there are indications of a security problem. The option to temporarily suspend operation without losing existing work should make it easier to adopt a precautionary approach to incident detection and response.

E-infrastructure technologies such as virtualisation are also used by campus services creating opportunities to share expertise and experience. Information about general and specific threats and mitigations should be exchanged with security and incident response teams at campus, network and national levels.

As e-infrastructures serve increasingly diverse communities they are likely to need to provide users with more information about the risks that were considered, the security approach adopted and the types of mitigation measures chosen. This will help users decide which infrastructure options are most suitable for their data and application: some will need strong preventive controls whereas others may be suited to a more reactive approach. Matching the security expectations of e-infrastructure users and providers will be essential to avoid security disasters.

Annex – Threat scenarios

One approach to designing security systems is to consider them from the point of view of potential attackers.

This can be particularly helpful in ensuring that a group of security controls provides a consistent level of protection: sufficient to reduce risks to an acceptable level without making legitimate access to the service intolerably onerous. Attackers will normally concentrate on the weakest point in any security scheme since that provides them with the easiest way to achieve their objective. Defenders should also try to identify those weak points and either bring them up to the required level or else supplement them with layered controls to improve the overall protection.

Considering likely attackers can also help determine the appropriate level of security for the e-infrastructure. If the service or information is likely to attract well-resourced, well-motivated attackers prepared to spend considerable time and/or money to achieve a specific goal, then an equivalent amount of effort should be spent on defending it.

Finally, role-playing attack scenarios can also be helpful in planning how to respond to an incident and identifying the tools and resources that are likely to be needed to detect and mitigate it.

The following scenarios suggest different motives for attacking an e-infrastructure and some of the approaches that attackers might adopt. All are based on real incidents involving research computing, though not necessarily (as far as we know) e-infrastructures. They are intended as a starting point for discussions and do not cover all possibilities: new forms of attack on networked systems are continually being discovered.

Example attack scenarios

- An authorised user of an e-infrastructure decides to use its CPU power to mine digital currency (e.g. Bitcoins) for personal gain. The required software is freely available for operating systems used by e-infrastructures, and little additional network traffic is required

- A group of digital 'pirates' notice that the e-infrastructure's storage capacity and high bandwidth network connection would make it ideal for distributing their collection of unlicensed software and movies. Although the e-infrastructure itself is secure, they find a vulnerability in a scientist's open source code that lets them install and run a server program to distribute the files using peer-to-peer or encrypted web protocols

- A would-be penetration tester decides that compromising a major e-infrastructure will enhance her CV. She calls a local administrator and says she is a PhD student who urgently needs to complete some work but can't find her login details after a recent move. Having been 'reminded' of her credentials she proves her success by posting confidential information obtained from the system on Pastebin

- A hacktivist hears that an e-infrastructure is being used for research that is contrary to his ethical beliefs. He considers using a botnet to launch a distributed denial of service attack against the service, though he would prefer to find a like-minded colleague to modify the program code - so that the research results could later be publicly discredited

- A company discovers that a competitor is collaborating with a university to develop advanced simulations of its next product. If it could obtain key information about the results it could improve its own product and get to market first. According to the university website, the professor leading the research has recently retired, so the competitor creates a Gmail account in her name and mails one of the researchers asking how the work is progressing [note that for some types of product, the attacker could be a nation state seeking to start up its own industry]

- An authorised user has written his own program code and tested it with small parameter values on his local computer. When he recompiles it with larger parameters on an e-infrastructure it continually creates new processes, consuming all the system's memory and CPU resources

For further information in this area, please visit:
community.ja.net/groups/uk-e-infrastructure-security-access-management-wg

Share our vision to make the UK the most digitally advanced education and research nation in the world

jisc.ac.uk

Jisc
One Castlepark
Tower Hill
Bristol, BS2 0JA
0203 697 5800
info@jisc.ac.uk