

Cyber Essentials

Requirements for IT Infrastructure

NCSC Publication
6th February 17

QG Adaption
Publication 6th July 17

Document No.
BIS 14/696/1.2

Requirements for IT Infrastructure

Specifying the technical controls required for IT infrastructure, for assessment under the Cyber Essentials scheme.

We specify the requirements under five technical control themes:

- firewalls
- secure configuration
- user access control
- malware protection
- patch management

As a Cyber Essentials scheme Applicant, you must ensure that your organisation meets all the requirements. You may also be required to supply various forms of evidence before your chosen Certification Body can award certification at the level you seek. Proceed as follows:

1. Establish the **boundary of scope** for your organisation, and determine **what is in scope within this boundary**.
2. Review each of the five **technical control themes** and the **controls they embody as requirements**.
3. Take steps as necessary to **ensure that your organisation meets every requirement**, throughout the scope you have determined.

What's new

We've reworked this information so that it is easier to understand and use. If you're familiar with the earlier PDF version you'll notice lots of changes in sense and structure.

Many statements are now clearer, to reduce ambiguity and potential to misinterpret requirements. Otherwise, there are few **technical** changes in this release. The main changes are:

- Made the title more specific, to help differentiate this document from other requirements documents we may produce in future (to deal with cloud services, for example).
- Clarified how to determine the scope and which devices are within the scope.

- Added requirement for authentication in services that allow Internet-based users to access data which must be protected.
- Removed requirement for regular password changes in Internet-facing services. Replaced with a choice of responses to deal with repeated failed authentication attempts.
- Added content to cover the use of certificate-based application whitelisting or sandboxing to defeat malware.
- Refined requirements for patching, to be more flexible about devices in scope yet also more specific.

Definitions

- **Software** includes operating systems, commercial off-the-shelf applications, plugins, interpreters, scripts, libraries, network software and firmware.
- **Devices** includes all types of hosts, networking equipment, servers, networks and end-user equipment such as desktop computers, laptop computers, tablets and mobile phones (smartphones) — whether physical or virtual.
- **Applicant** means the organisation seeking certification, or sometimes the individual acting as the main point of contact, depending on context.

1. Business Scope

Overview of the scope

Assessment and certification can cover the whole of the Applicant's IT infrastructure, or a sub-set. Either way, the boundary of the scope must be clearly defined in terms of the business unit managing it, the network boundary and physical location. The scope must be agreed between the Applicant and the Certification Body before assessment begins.

We strongly recommend that the scope should include the whole IT infrastructure if possible, to achieve the best protection.

The requirements apply to all the devices and software that are within this boundary and that meet the conditions below:

- accept incoming network connections from untrusted Internet-connected hosts
- establish user-initiated outbound connections to arbitrary devices via the Internet
- control the flow of data between any of the above devices and the Internet

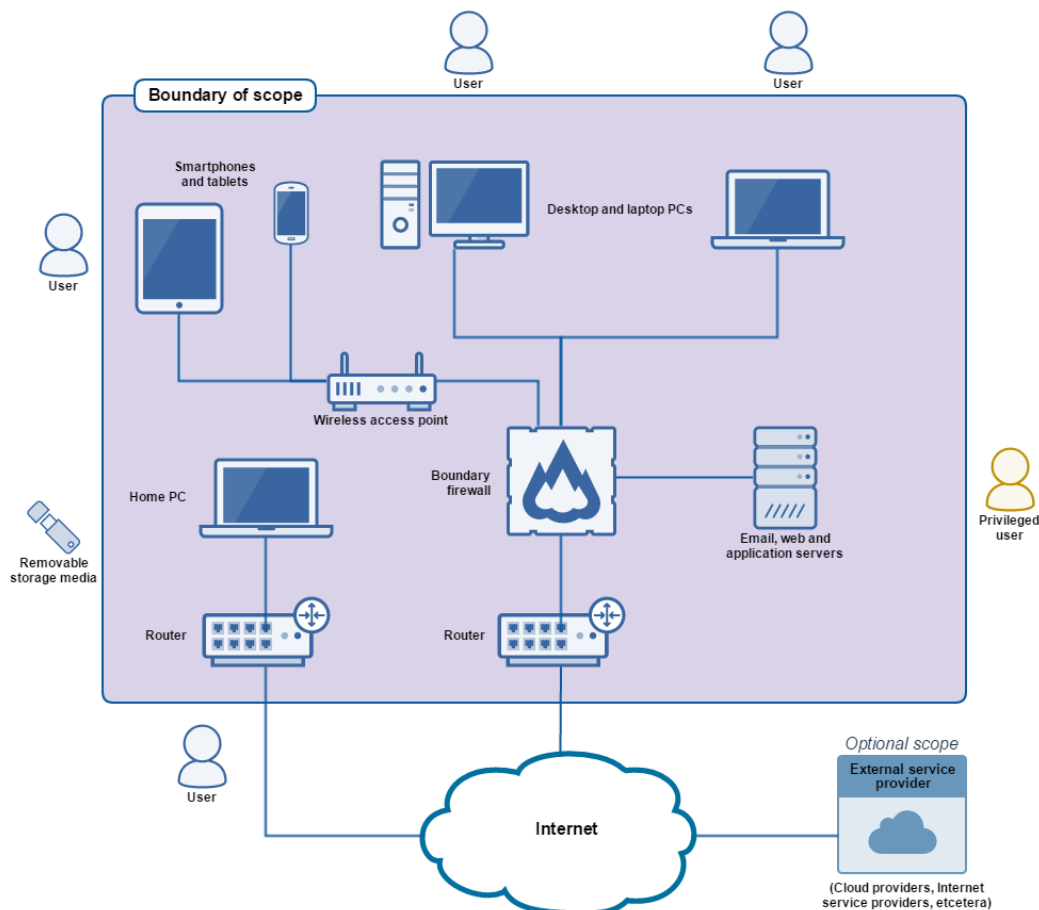


Figure 1: Scope of the requirements for IT infrastructure.

Bring your own device (BYOD)

In addition to mobile or remote devices owned by the organisation, user-owned devices which access organisational data or services are **in scope**.

Traditionally, user devices were managed through centralised administration, ensuring consistency across the organisation. In such cases, certification of the security controls is straightforward as there will be a standard build or reference to assess.

BYOD complicates matters, as users are given more freedom to ‘customise’ their experience making consistent implementation of the controls more challenging.

Wireless devices

Wireless devices (including wireless access points) are:

- **in scope** if they can communicate with other devices via the Internet
- **not in scope** if it is not possible for an attacker to attack directly from the Internet (the Cyber Essentials scheme is not concerned with attacks that can only be launched from within the signal range of the wireless device)

Externally managed services — cloud

If it is practicable for the Applicant to apply the requirements to its cloud services then it should include these services within the boundary of scope.

Example

Acme Corporation has procured infrastructure as a service (IaaS) from a cloud service provider. Acme has control of the operating systems on the infrastructure, and so it is able to apply the requirements. Acme will therefore include this service in its scope.

At present, software as a service (SaaS) and platform as a service (PaaS) are **not in scope** — the current requirements can not be mapped against them.

Externally managed services — other

Where the Applicant is using other externally managed services (such as remote administration) it may not be possible for the Applicant to meet all the requirements directly. The Applicant may **choose** whether or not to include these services within the boundary of scope, according to feasibility.

If included then the Applicant must be able to attest that the requirements that are outside of the Applicant's control are being adequately met by the service provider. Existing evidence may be considered (such as that provided through PCI certification of a cloud service, and ISO 27001 certifications that cover an appropriate scope).

Web applications

Commercial web applications created by development companies (rather than in-house developers) and which are publicly accessible from the Internet are **in scope** by default.

Bespoke and custom components of web applications are **not in scope**. The primary mitigation against vulnerabilities in such applications is robust development and testing in line with commercial best practices, such as the [Open Web Application Security Project \(OWASP\)](#) standards.

2. Password-based authentication

Password-based authentication (Question 2.1, 2.2, 2.3, 4.1 & 4.3)

The Applicant must make good use of the technical controls available to it on password-protected systems. As much as is reasonably practicable, technical controls and policies must shift the burden away from individual users and reduce reliance on them knowing and using good practices.

Users are still expected to pick sensible passwords.

For password-based authentication in Internet-facing services the Applicant must:

- protect against brute-force password guessing, by using at least one of the following methods:
 - lock accounts after **no more** than 10 unsuccessful attempts
 - limit the number of guesses allowed in a specified time period to **no more** than 10 guesses within 5 minutes

For password-based authentication in Internet-facing and non-internet facing services the Applicant must:

- set a **minimum** password length of at least 8 characters
- **not** set a maximum password length
- change passwords promptly when the Applicant knows or suspects they have been compromised

- have a password policy that tells users:
 - how to avoid choosing obvious passwords (such as those based on easily-discoverable information like the name of a favourite pet)
 - not to choose common passwords — this could be implemented by technical means, using a password blacklist
 - not to use the same password anywhere else, at work or at home
 - where and how they may record passwords to store and retrieve them securely — for example, in a sealed envelope in a secure cupboard
 - if they may use password management software — if so, which software and how
 - which passwords they really must memorise and not record anywhere

The Applicant is **not** required to:

- enforce regular password expiry for any account (we actually advise against this — for more information see [The problems with forcing regular password expiry](#))
- enforce password complexity requirements

3. Firewalls

Applies to: boundary firewalls; desktop computers; laptop computers; routers; servers.

Objective

Ensure that only safe and necessary network services can be accessed from the Internet.

Introduction

All devices run network services, which create some form of communication with other devices and services. By restricting access to these services, you reduce your exposure to attacks. This can be achieved using firewalls and equivalent network devices.

A boundary firewall is a network device which can restrict the inbound and outbound network traffic to services on its network of computers and mobile devices. It can help protect against cyber attacks by implementing restrictions, known as 'firewall rules', which can allow or block traffic according to its source, destination and type of communication protocol.

Alternatively, a host-based firewall may be configured on a device. This works in the same way as a boundary firewall but only protects the single device on which it is configured. This approach can provide for more tailored rules and means that the rules apply to the device wherever it is used. However, this increases the administrative overhead of managing firewall rules.

Requirements under this technical control theme

Every device that is in scope must be protected by a correctly configured firewall (or equivalent network device) **(Question 3.1 & 3.2)**

For all firewalls (or equivalent network devices), the Applicant organisation must routinely:

- change any default administrative password to an alternative that is difficult to guess (see [Password-based authentication](#)) — or disable remote administrative access entirely **(Question 3.3)**
- prevent access to the administrative interface (used to manage firewall configuration) from the Internet **(Question 3.4)**, unless there is a clear and documented business need and the interface is protected by one of the following controls:
 - a second authentication factor, such as a one-time token
 - an IP whitelist that limits access to a small range of trusted addresses
- block unauthenticated inbound connections by default **(Question 3.5)**

- ensure inbound firewall rules are approved and documented by an authorised individual; the business need must be included in the documentation (**Question 3.6**)
- remove or disable permissive firewall rules quickly, when they are no longer needed (**Question 3.7**)
- use a host-based firewall on devices which are used on untrusted networks, such as public Wi-Fi hotspots (**Question 3.2**)

4. Secure Configuration

Applies to: email, web, and application servers; desktop computers; laptop computers; tablets; mobile phones; firewalls; routers.

Objective

Ensure that computers and network devices are properly configured to:

- reduce the level of inherent vulnerabilities
- provide only the services required to fulfil their role

Introduction

Computers and network devices are not always secure in their default configurations. Standard, out-of-the-box configurations often include one or more weak points such as:

- an administrative account with a predetermined, publicly known default password
- pre-enabled but unnecessary user accounts (sometimes with special access privileges)
- pre-installed but unnecessary applications or services

Default installations of computers and network devices can provide cyber attackers with a variety of opportunities to gain unauthorised access to an organisation's sensitive information — often with ease.

By applying some simple technical controls when installing computers and network devices you can minimise inherent vulnerabilities and increase protection against common types of cyber attack.

Requirements under this technical control theme

Computers and network devices

The Applicant must be active in its management of computers and network devices. It must routinely:

- change any default or guessable account passwords to something non-obvious
(Question 4.1 & 4.3)

- remove and disable unnecessary user accounts (such as guest accounts and administrative accounts that won't be used) **(Question 4.2)**
- remove or disable unnecessary software (including applications, system utilities and network services) **(Question 4.4)**
- disable any auto-run feature which allows file execution without user authorisation (such as when they are downloaded from the Internet) **(Question 4.5)**
- authenticate users before allowing Internet-based access to commercially or personally sensitive data, or data which is critical to the running of the organisation **(Question 4.6)**

5. User Access Control

Applies to: email, web and application servers; desktop computers; laptop computers; tablets; mobile phones.

Objective

Ensure user accounts:

- are assigned to authorised individuals only
- provide access to only those applications, computers and networks actually required for the user to perform their role

Introduction

Every active user account in your organisation facilitates access to devices and applications, and to sensitive business information. By ensuring that only authorised individuals have user accounts, and that they are granted only as much access as they need to perform their role, you reduce the risk of information being stolen or damaged. **(Question 5.1)**

Compared to normal user accounts, accounts with special access privileges have enhanced access to devices, applications and information. When such accounts are compromised, their greater freedoms can be exploited to facilitate large-scale corruption of information, disruption to business processes and unauthorised access to other devices in the organisation.

'Administrative accounts' are especially highly privileged, for example. Such accounts typically allow:

- execution of software that has the ability to make significant and security relevant changes to the operating system
- changes to the operating system for some or all users
- creation of new accounts and allocation of their privileges

All types of Administrator will have such accounts, including Domain Administrators and Local Administrators. **(Question 5.6)**

Now consider that if a user opens a malicious URL or email attachment, any associated malware is typically executed with the privilege level of the account that user is currently operating. Clearly, you must take special care over the allocation and use of privileged accounts.

Example

Jody is logged in with an administrative account. If Jody opens a malicious URL or email attachment, any associated malware is likely to acquire administrative privileges.

Unfortunately, this is exactly what happens. Using Jody's administrative privileges, a type of malware known as ransomware encrypts all of the data on the network and then demands a ransom.

The ransomware was able to encrypt far more data than would have been possible with standard user privileges, making the problem that much more serious.

Requirements under this technical control theme

The Applicant must be in control of its user accounts and the access privileges granted to each user account. It must also understand how user accounts authenticate and control the strength of that authentication. This means the Applicant must:

- have a user account creation and approval process (**Question 5.2**)
- authenticate users before granting access to applications or devices, using unique credentials (see [Password-based authentication](#)) (**Question 5.3**)
- remove or disable user accounts when no longer required (when a user leaves the organisation or after a defined period of account inactivity, for example) (**Question 5.4**)
- implement two-factor authentication, where available (**Question 5.5**)
- use administrative accounts to perform administrative activities **only** (no emailing, web browsing or other standard user activities that may expose administrative privileges to avoidable risks) (**Question 5.6**)
- remove or disable special access privileges when no longer required (when a member of staff changes role, for example) (**Question 5.7**)

6. Malware Protection

Applies to: desktop computers; laptop computers; tablets; mobile phones.

Objective

Restrict execution of known malware and untrusted software, to prevent harmful code from causing damage or accessing sensitive data.

Introduction

The execution of software downloaded from the Internet can expose a device to malware infection.

Malware, such as computer viruses, worms and spyware, is software that has been written and distributed deliberately to perform malicious actions. Potential sources of malware infection include malicious email attachments, downloads (including those from application stores), and direct installation of unauthorised software.

If a system is infected with malware, your organisation is likely to suffer from problems like malfunctioning systems, data loss, or onward infection that goes unseen until it causes harm elsewhere.

You can largely avoid the potential for harm from malware by:

- detecting and disabling malware before it causes harm (anti-malware)
- executing only software that you know to be worthy of trust (whitelisting)
- executing untrusted software in an environment that controls access to other data (sandboxing)

Example

Acme Corporation implements code signing alongside a rule that allows only vetted applications from the device application store to execute on devices. Unsigned and unapproved applications will not run on devices.

The fact that users can only install trusted (whitelisted) applications leads to a reduced risk of malware infection.

Requirements under this technical control theme

The Applicant must implement a malware protection mechanism on all devices that are in scope. For each such device, the Applicant must use at least one of the three mechanisms listed below:

Anti-malware software

- The software (and all associated malware signature files) must be kept up to date, with signature files updated at least daily. This may be achieved through automated updates, or with a centrally managed deployment. **(Question 6.1.1)**
- The software must be configured to scan files automatically upon access. This includes when files are downloaded and opened, and when they are accessed from a network folder. **(Question 6.1.2)**
- The software must scan web pages automatically when they accessed through a web browser (whether by other software or by the browser itself). **(Question 6.1.3)**
- The software must prevent connections to malicious websites on the Internet (by means of blacklisting, for example) — unless there is a clear, documented business need and the Applicant understands and accepts the associated risk. **(Question 6.1.4)**

Application whitelisting

- Only approved applications, restricted by code signing, are allowed to execute on devices. **(Question 6.2.1)** The Applicant must:
 - actively approve such applications before deploying them to devices **(Question 6.2.2)**
 - maintain a current list of approved applications **(Question 6.2.3)**
- Users must not be able to install any application that is unsigned or has an invalid signature. **(Question 6.2.4)**

Application sandboxing

- All code of unknown origin must be run within a 'sandbox' that prevents access to other resources unless permission is explicitly granted by the user. **(Question 6.3)**
This includes:
 - other sandboxed applications
 - data stores, such as those holding documents and photos
 - sensitive peripherals, such as the camera, microphone and GPS
 - local network access

7. Patch Management

Applies to: web, email and application servers; desktop computers; laptop computers; tablets; mobile phones; firewalls; routers.

Objective

Ensure that devices and software are not vulnerable to known security issues for which fixes are available.

Introduction

Any device that runs software can contain security flaws, known as 'vulnerabilities'.

Vulnerabilities are regularly discovered in all sorts of software. Once discovered, malicious individuals or groups move quickly to misuse (or 'exploit') vulnerabilities to attack computers and networks in organisations with these weaknesses.

Product vendors provide fixes for vulnerabilities identified in products that they still support, in the form of software updates known as 'patches'. Patches may be made available to customers immediately or on a regular release schedule (perhaps monthly).

Product vendors do not generally release patches for products they no longer support — not even to fix vulnerabilities.

Requirements under this technical control theme

The Applicant must keep all its [software](#) up to date. Software must be:

- licensed and supported (**Question 7.1**)
- removed from devices when no longer supported (**Question 7.2**)
- patched within 14 days of an update being released, where the patch fixes a vulnerability with a severity the product vendor describes as 'critical' or 'high risk'* (**Question 7.3**)

*If the vendor uses different terms to describe the severity of vulnerabilities, see the precise definition in the [Common Vulnerability Scoring System \(CVSS\)](#).

For the purposes of the Cyber Essentials scheme, 'critical' or 'high risk' vulnerabilities are those with the following values:

- attack vector: **network** only
- attack complexity: **low** only
- privileges required: **none** only

- user interaction: **none** only
- exploit code maturity: **functional** or **high**
- report confidence: **confirmed** or **high**

Some vendors release patches for multiple issues with differing severity levels as a single update. If such an update covers any 'critical' or 'high risk' issues then it must be installed within **14 days**.

8. Approval

It is important that in all organisations, senior management buy in to Cyber Security of the organisation, therefore it is mandatory that a Board level officer (or equivalent) of the organisation has approved the information given to a Certification Body.