Version 1.0, May 2024

Technology assurance

The NCSC's Technology Assurance activities provide a means to gain confidence in the cyber security of the services and technologies on which the UK relies.

Principles: Product design and functionality

6 principles which describe security functionality intended to defend against the most common techniques used by cyber attackers.

Like all others, quantum security products must follow these principles, as their goal – security – is the same. Users should not be expected or required to have specific quantum knowledge, nor to understand the quantum operation of the device.

Furthermore, it is extremely likely that quantum communications will always be integrated with conventional communications (even within a product or system), so these principles should apply to both aspects, separately but also when together in a system. Quantum-specific comments presented here pertain to the quantum hardware/layer, but with the understanding that the principles will also apply to all conventional hardware/layers, including e.g. key management.

In publishing these principles – on Product development, Product design and functionality (this document) and Through life – so they can be implemented, it is assumed that suppliers of all security products (non-quantum and quantum), whether UK-based or not, will have to provide whatever is needed to evidence that the principles have been followed. Therefore, for all the application and deployment scenarios where historically GCHQ/NCSC would have overseen, or provided, the direct assessment and assurance for their use in the UK, in the future developers of all security products and services will need to provide arguments that they meet assurance claims that underpin the principles, backed by evidence.

Note: This provision of evidence would seem to be a new challenge for companies (particularly non-UK) wishing to supply security products to UK markets or non-UK companies wishing to participate in the supply chains (see later) of UK companies producing security products in the UK. It will be interesting to monitor the response of these companies to this new approach from NCSC (effectively now defining principles and outsourcing the assurance, rather than undertaking or directly overseeing it).

It's essential that products implement the security functionality needed to mitigate the cyber threats they will face in use.

The *Product Design and Functionality principles* are intended to:

- Assist product vendors, designers and developers in making security-related decisions as they take a product from concept to installation and use.
- Help risk owners to gain confidence that a technology solution mitigates the specific threats which they expect it to face.

These principles provide a framework that is not only relevant to cyber security products, but any product which must be resistant to cyber attack, even though its main purpose may not be cyber security related.

Cyber attacks

The *Product Design and Functionality principles* describe security functionality which is intended to defend against the most common techniques used by cyber attackers. They include measures to:

- Protect sensitive data in transit, and at rest on the product
- Maintain the secure operation of the product
- Enable malicious activity to be detected and acted upon

These fundamental principles apply across the assurance spectrum, from defending against the most basic commodity-level threats to countering the elevated threat from extremely capable, and motivated, threat actors.

Attacking the specifically quantum aspects of a quantum security product to extract information (as opposed to simply rendering it inoperable, by denying service) will likely require some level of quantum expertise. Product vendors, designers and developers should therefore make the worst-case assumption, that threat actors possess the same expertise and capability as they do. Note that this already holds for quantum security proofs, where threat actors are assumed to possess any form of technology that operates according to quantum laws, even if this technology does not yet exist. Some weaker quantum security analyses have been made, where bounds or limitations are placed on the quantum capabilities of the threat actor. These analyses may still have significant practical application and value, if the bounds or limitations are consistent with current quantum technology capabilities.

Implementing the principles

How the principles are met in practice, and the strength of any protective measures, is expected to vary according to the anticipated level of threat. The amount of confidence needed will also vary according to risk appetite.

The principles provide a framework against which the design and functionality of the product can be analysed. For each principle in this collection we describe the underlying security issue which needs to be addressed and give a series of example measures which could be used.

There are 6 principles

1. 1

Usability of the product

2. 2

Restrict access to authorised users

3. 3

Protect sensitive data when in transit

4. 4

Protect against unauthorised access and modification

5. 5

Protect against compromise from connected technology

6. 6

Security events should be logged and monitored

These principles detail the areas to be considered when assessing whether the security functionality of a product is sufficient for the level of threat it faces.

Guidance on additional mitigations for elevated threat scenarios can be found in the NCSC Design guidelines for High Assurance products.

1. Usability of the product

Making secure operation the natural choice.

Any product which must be resistant to cyber attack should be designed to naturally promote safe and secure use. As far as possible, security functionality should not interfere with day-to-day operation, remaining easily accessible but unobtrusive.

Security functions and interfaces should be intuitive for the people who use them. Support should be available to help ensure the product is configured and used in the way intended.

By making secure operation the preferred choice for users, these design goals will help to ensure that a product remains as secure as possible throughout its life.

Very often, if security doesn't work for people, it doesn't work at all. For instance, if a product is clumsy to use, people will find a lower friction way to get their task done. If a product is difficult to configure, then a mistake in its set up could lead to a lack of security functionality further down the line.

These usability challenges can potentially bypass the controls that have been put in place to keep them, and the systems in which they are working, safe.

Quantum expertise should not be required to configure, use and (as required) reconfigure products. Or if it is, this expertise needs to be from an assured expert or the product producer/supplier. Basically, users should see no operational difference between quantum and non-quantum security products, so this usability principle applies equally to both.

Examples of defensive measures

- The product should have a focus on the human-centred design qualities of efficiency, effectiveness, user satisfaction, inclusivity, and accessibility. Education, training, and configuration support should be available to people installing and using the product.
- Usability testing of the product should be carried out by a representative sample of people, covering potential roles, tasks, constraints, and situational contexts. Product designers should take time to understand who their likely users will be.
- Product performance should be predictable for users. If performance is not good enough, there will be pressure for people to find alternative ways to get their task done, potentially bypassing security controls.
- A process for logging usage patterns and reporting pain points should be in place to identify potential sources of vulnerability and enable continuous improvement.
- The user should be made aware when the product is insecurely configured, and reverting to a secure default state should be easy. The product should provide support, and clear ways to recover, if users have made an error. Any feedback from the product to the user should be clear and meaningful.
- Measures are put in place, where possible, to prevent or make scenarios that facilitate malicious use or lead to security incidents less likely or dangerous. These could include identifying and controlling for unintended interactions with the product, or moving dangerous actions away from commonly used functions.

2. Only authorised users should have access to data and functionality

Keep access permissions to the minimum necessary.

A key principle in security is that users should only have access to data and functionality that is necessary to support their legitimate aims. ‘Access’ can mean both physical access to a device, and remote access to services and functions provided by the product.

Products that may be subject to cyber attack need to be managed and configured. Access to the management and configuration functions of a device should be regarded as a privileged role, restricted to authorised users and implemented securely. Access should then be logged and monitored accordingly (see Principle 5).

An attacker who can gain access to the management of a product can affect its security and functionality, and compromise sensitive data. If regular users also have access to low level functionality beyond that needed for their role, there is unnecessary potential for them to gain access to more sensitive functions.

This limited access principle also applies to quantum security products. See also the added measure below, with regard to specifically quantum aspects of the products. Users should not be able to expand, or introduce new, quantum side channels[1] in these quantum aspects, either deliberately or inadvertently.

[1] For more details and terminology, refer to the “Introduction_Quantum Assurance” document.

Example defensive measures

- The product should support role-based authentication and access control. Access to data and functionality is defined by the role. All users should be issued with unique, but usable, authentication credentials before their first access to the system.
- All requests for access should be authenticated before being granted, so that users are only given access to the data and functionality to which their role entitles them. Authentication mechanisms might range from simple usernames / passwords to large certificate-based trust architectures, depending on the complexity and security requirements of the system.
- Only authorised and authenticated administrators should have access to the management interface - it should be unavailable to all others. Management can be remote or local, and if either mode is not required, it should be disabled.
- Privileged functions should afford access to the minimum amount of sensitive user data necessary. The purpose of management and configuration is to support the operation of the device, not give access to all data.
- Authentication credentials should be generated and managed securely. Management could involve technical considerations, such as not storing passwords in plain text, or procedural, such as controlling distribution of physical access tokens and applying time limits to credential validity. Default credentials, such as those used during manufacturing must be removed before the product is operational.
- You should consider how to protect sensitive user data from other users, either from attempts to bypass or undermine security functions, or through inadvertent implementation errors (see Principle 4).
    - Users should not need to reconfigure the (for them) hidden quantum aspects of a product, so they should not have access to these. For example, a user should not be able to increase the intensity of a weak coherent pulse source in a QKD system, mistakenly thinking they could increase the key rate, but in fact rendering the product insecure.

3. Protect sensitive data in transit

When sending sensitive data across any network it must be protected against eavesdropping and tampering.

Users and product developers need to be confident that whenever sensitive data is in transit, it is protected against eavesdropping and tampering. This is true regardless of the type of connection: it could be a physical (wired) connection across a network, a wireless or Bluetooth connection between devices, or a broadcast radio frequency transmission. The mechanisms used should protect both communications over public (untrusted) networks and within private (trusted) networks.

Current quantum security products support the protection of sensitive data in transit either by providing keys (QKD) for the encryption of sensitive data, and/or random numbers (from QRNGs) to support these processes. The specific remarks below refer to QKD and QRNG products. New quantum security products offering new functionality will need separate consideration when these products come to market.

An adversary who can intercept a communication may seek to gain an advantage in a number of ways.

For QKD, these interventions would be with respect to the establishment of key material, and the comments inserted reflect this. For the sensitive data transmission itself, all the non-quantum comments and measures still apply.

- They may want to extract sensitive data directly.
    - Extraction of key material as it is being established is not possible for QKD systems that operate within the assumptions that underpin their security proof. This is the quantum advantage of QKD.
- They may want to modify the communication in order to masquerade as a legitimate user and send malicious messages.
    - QKD requires separate authentication (e.g. pre-shared secret material, or a conventional cryptographic mechanism) to prevent such attacks.
- They may look to replay previously transmitted data to cause a disruptive effect.
    - This cannot undermine the QKD security proof.
- They may seek to prevent data reaching its intended recipient, causing a denial of service.
    - QKD cannot operate if the quantum channel is broken, denying service. (Refer to the redundancy defensive measure below.)

Protective mechanisms aim to defend against these attacker objectives in two main ways:

1. Preventing an adversary from intercepting data in the first place
2. Preventing loss of confidentiality or integrity if it is intercepted

Example defensive measures

- Strong cryptography should be used to establish trusted connections, ensuring sensitive data only goes where it is intended to. This will also provide confidentiality and integrity protection. Standardised algorithms and transport protocols provide the mechanisms to do this effectively, and to detect a range of attacks.
    - QKD products can provide assured keys, to support this defensive measure.
- Where content is encrypted, encryption should happen at source, and decrypted only at the final destination, not *en route*. This ensures that an attacker intercepting data in transit cannot learn its content.
    - At present QKD has a distance limitation, so long-distance QKD relies on trusted nodes and suitable key management to provide keys over longer distances. Once keys are shared over long distance, there is no need to decrypt and re-encrypt sensitive data at trusted QKD nodes. Future quantum security products (repeaters, etc.) are being developed to overcome the current need for trusted QKD nodes.
- Cryptographic mechanisms rely on secret values – keys – that should be unpredictable by an adversary. Strong random number generators should be used to generate keys, and there should be appropriate processes for distributing, managing and storing keys in a secure manner, throughout their lifetime.
    - QKD systems, operating securely, share strong random numbers as keys. QRNGs can also support this measure by providing strong randomness that is assured to be unique.
- Data in transit is less likely to be at risk from an adversary if it is hard to identify. Use of standardised, widely used protocols can help with this for electronically transmitted data. Unusual regions or patterns of use of the radio frequency spectrum should be avoided for data transmitted over-the-air.
- Where availability of communications is a critical requirement, you should consider building in redundancy to the system, so that if one connection is unavailable, data can be transmitted through an alternative route.
    - Service can be denied for QKD by breaking of the quantum channel. A network topology providing multiple quantum channel routes can mitigate this, along with key distribution in advance, when QKD service is available.

4. Maintain the integrity of a product and any sensitive data held on it

Ensuring the product is resilient to attempts to change its behaviour.

Adversaries may want to gain access to sensitive data, either to compromise the user or to aid development of future attacks. They may look to modify software, firmware or hardware to alter the operation of the product, or to enable a persistent presence.

Product designers should ensure that there is appropriate identification of, and protection for all sensitive data in the product - both user data and device-sensitive data. They should also ensure that mechanisms exist to protect against physical modification, and to give confidence in the integrity of the product and the components it relies upon.

Example defensive measures

- Sensitive user data and device-specific data, should be clearly identified during design and, where possible, be separated from non-sensitive data, for example, in separate memory, or separate locations within file systems. This enables data protection mechanisms to be well targeted.
- If sensitive data needs to be persistent, apply appropriate confidentiality and integrity mechanisms. Where sensitive data is updated, update mechanisms should provide authentication. This gives the user and the product developer confidence that their personal and proprietary information is well protected.
- It is important to minimise the amount of data that is potentially accessible to an adversary. For user data, information should not be retained when it is no longer required. For device-specific data, the developer should limit the amount of information available to someone scanning or probing the device.
- Verify the integrity of software and hardware components during start-up and operation, and through product updates. Where the device relies on external components (e.g. unique cables or peripherals), these should also be verified prior to use. This provides confidence that the device remains in a trusted state throughout its lifetime.
- Where they are available, use the built-in security features of components within a product. Many security-focussed components provide protections for memory contents, or mechanisms to aid with separation of sensitive and non-sensitive data.
- Incorporate methods to detect and respond to attempts at physical compromise of the product. These methods could be procedural, but can also include passive or active anti-tamper technologies. A layered approach, comprising a few such approaches, provides defence-in-depth.
    - All the defensive measures identified above also apply to quantum security products. However, particular attention should be given to the introduction of new quantum side channels[1] by attacks that seek to modify the behaviour of the quantum hardware. Resilience against these is desirable and where this not possible, detection and deployment of countermeasures provides defence.
    - For QKD systems, attention should also be given to the integrity of the separate authentication mechanism.

[1] For more details and terminology, refer to the “Introduction_Quantum Assurance” document.

5. Protect against compromise from connected technology

Ensuring connectivity can be achieved without compromising security.

Handling connectivity to external devices, or networks, is a critical security function for most systems. This can range from the pairing of a Bluetooth headset with a mobile phone, to the dynamic interconnections of a large scale enterprise architecture across a global network.

While these connections are essential, they have the potential to create additional areas of focus for an adversary. If an attacker can gain control over a connected device, and if they can use such a device to send malicious data through your system, the have the potential to compromise security functionality.

Control over how connections are authorised and managed, what data is allowed to pass through them, and the way your product protects against exploitation, can help manage this risk.

Attacks may also come from supposedly trusted connections if another device on your network has been compromised, or there is an insider threat. So, measures to protect against compromise from connected technology should be considered even when you are only connecting inside your local network.

Example defensive measures

- You should only establish connections with devices, systems or networks you can trust. Use standardised cryptographic methods for device authentication and management of connection sessions, wherever possible. This means that you can be confident that whoever you are communicating with is who they claim to be.
- Ensure you have a method for revoking access granted to connected devices. If an attacker is able to control a compromised device, you want to be able to stop them from establishing a trusted connection with your product.
- You should control the processing or forwarding of data. This means checking that data reaching your product has an expected format and that you are able to discard anything that does not match expectations. When forwarding data to other devices, you should apply technical controls to establish trust and protect against malign content.
- Products should be designed so that an attacker cannot easily exploit any vulnerability even if they can establish an initial presence. As described in Principle 2, restoring and rebooting from a trusted, known state at power-up make establishing persistent presence hard, and strong physical separation between sensitive and non-sensitive data area can make accessing sensitive data harder.
    - All the defensive measures identified above also apply to quantum security products. However, particular attention should be given to the conventional technologies connected to the quantum parts. The approach to quantum

security proofs targets composability, to allow for connections of multiple quantum technologies.

6. Security events should be logged and monitored

Keeping your eyes open for possible attacker activity.

Logging and auditing of events that indicate changes in the secure operation of a device can help highlight possible attacker activity, provide early warning of compromise and offer a foundation for analysis in response an any security incidents.

Effective logging will not always deter an adversary, but as long as appropriate monitoring and auditing of logs is in place, it does increase the likelihood that their activity will be exposed. This, in turn may decrease their potential appetite for attack. Logging should be informed by the threat that a product is likely to be exposed to.

Security logs should be treated like other sensitive data, and protected accordingly. Compromise of logs can enable adversaries to gain insights into security configurations, but also enable malicious activity to go undetected.

Example defensive measures

- Logs are most useful when their purpose is clearly defined. You should choose the events you log based on the threats that might apply to a device, and consider categorising events by level of importance, based on the level of action to be taken in response to them.
- Access to management functions that allow modification of the security configuration of a device may be of particular concern. For these, or other events critical to security, you should consider triggering alerts for more immediate attention.
- Access to logs should be a privileged function, limited to those in administrator roles, and disabled by default. Logs should be protected through mechanisms to detect unexpected modification and alerting in response to unauthorised changes.
- Logs should persist for the lifetime of the product. This may mean you need to periodically back them up, or export them securely. Remote logging can help avoid on-device storage limitations over long periods of time. Developers should consider how to enable modelling and measurement of typical logging behaviour to help users choose appropriate storage and handling.
- Logs are only useful if there is an appropriate strategy for monitoring and auditing, to identify abnormal behaviour. Structured logs make auditing easier, whether this is manual or automated. Where auditing is manual, critical events should be highlighted and easy to identify, reducing cognitive burden on the auditor.
- Auditing is most effective if suspicious behaviour patterns can readily be recognised. For automated monitoring, models of typical suspicious behaviour are needed. A system-

wide monitoring and auditing system may be helpful, and all monitoring systems should provide meaningful and actionable alerting in a way that is useful to the user or system owner.

- All the defensive measures identified above also apply to quantum security products. There will also be additional quantum security parameters that can be logged and audited. For example, a QKD system can operate for a range of QBER (quantum bit error rate), but a sudden change in QBER even within the acceptable operating range could be a signature of attacker activity.