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: Through-life

5 principles for maintaining security through all stages of a product's lifecycle.

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, and Through-life (this document) – 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).

Against a backdrop of rapid technological advancement and a forever evolving threat environment, continuous management of device security, throughout its lifetime, is vital.

The Through-Life principles are intended to:

- Assist vendors, designers and developers to make decisions about maintaining product security through life
- Help risk owners to gain confidence that a technology solution mitigates the specific threats which they expect it to face through all stages of a product's lifecycle

About these principles

This guide introduces five principles of through life product security. These principles are intended to draw out considerations for maintaining security through all stages of a product's lifecycle: from conception, design and development, testing, initial release, in-field updates, decommissioning and disposal.

These principles complement the secure development and product design and functionality principles that also form part of this collection. 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 five base through life principles:

1. **1**

Enable people to manage their risks

2. **2**

Protect the product from unauthorised access or modification

3. **3**

Monitor services used by a product

4. **4**

Review and improve the security offered by a product

5. **5**

Be prepared to respond to external events

The need for through life security

Sources of vulnerability

There are many ways in which a product can become insecure.

Features which are intrinsic to a product or its implementation may become increasingly susceptible to attack over time. Similarly, vulnerabilities overlooked during development may only emerge once a product is in widespread use.

For quantum security products, degradation of quantum hardware components away from their original specifications could increase susceptibility to attack. Emergence of new quantum side channels[1] during deployment and use would have to be handled in the field, or by recall.

[1]For more details and terminology, refer to the "Quantum Assurance Introductory Foreword".

The environment in which a product is developed, built, tested and maintained may be targeted as a means to discover confidential information, or introduce weaknesses through surreptitious modifications.

Particularly for quantum security products, the physical environment where these are developed, built and tested needs to be secure, as quantum security relies on physical aspects of the hardware.

External events, outside the control of a product developer, also have the potential to impact on the security of a product.

As long as quantum security products continue to satisfy the conditions specified in their security proof, there is immunity to external events. For example, development of any new conventional or quantum technologies will not provide new tools to hack quantum security.

Other parties in the supply chain may be compromised, or components from third parties incorporated into a product may be found to be vulnerable, or become obsolete. Services on which a product depends also need to be considered, as any compromise of these could have a security impact for a product and its users.

Nature of mitigations

Mitigations must address a range of risks, many of which are not solely related to the technical features of a product. So, it's necessary to consider all potential avenues of compromise.

It is important then, that users of a product are provided with access to sufficient information and capability to help them protect themselves and manage any security risks to their information.

1. Enable people to manage their risks

Equip users with the tools and information needed to mitigate any residual risk.

Technical defensive measures designed into a product may have limitations, and their effectiveness may degrade over time. Developers should ensure that advice on these limitations, along with associated mitigations and improvement, is readily available and actively advertised to both risk owners and the people using the product.

The people using the product should have easy access to a range of product support options throughout its lifetime. They should be notified promptly of exploitable vulnerabilities so that they can protect themselves and others, and when a product is updated, they should be advised of the security implications.

Users of quantum security products cannot be expected to have quantum knowledge, so easy access to support options is important. In preference to service or function simply being

denied, vulnerable quantum secure products should default to conventional security (even if not proven) whilst vulnerabilities are addressed.

When a user passes a product on to another party, they should be confident that their sensitive data is not exposed.

Example defensive measures

- User manuals should be clearly written, and updated as new features are added to a product, highlighting how to improve security. User training should be offered regularly on the security aspects of installation, configuration, maintenance and use, drawing attention to common mistakes to avoid, along with an explanation of security implications.
    - For quantum security products, users without quantum expertise (likely most) should not be undertaking any hardware installation, configuration and maintenance that requires quantum knowledge. To them, these products should simply be security products, handled and used as per non-quantum products.
- Where a product, or a service it depends on, has become vulnerable to compromise, security alerts should be issued immediately to affected users. These should clearly detail mitigating actions they should take. An effective product support service, with contact details accurately maintained, can help users get support in resolving security issues, as well as reporting defects.
- Making the process for updating a product simple and painless maximises the likelihood that users will adopt security enhancements. Whenever a product update includes new features, those features should be configured in the most secure state by default, so that the user is in control of activating any feature that increases the attack surface or otherwise puts them at greater risk.
    - Hardware updates of quantum components will require appropriate technical support. Conventional software updates should proceed as per those for non-quantum security products.
- End-of-life dates should be clearly published to users, so that they know when security updates will no longer be offered for a product. When a product is known to offer insufficient security, it should be retired.
    - End-of-life for quantum security products due to insufficient security will correspond to the hardware no longer being able to satisfy the security proof assumptions, or the product becoming vulnerable to new attacks (e.g. based on the emergence of new quantum side channels) that the vendor cannot or chooses not to address. A security proof does not degrade over time against its original assumptions. However, if a new vulnerability (e.g. quantum side channel) is discovered, it may be desirable to modify the assumptions and correspondingly update the security proof, potentially with the further addition of a physical countermeasure. If this route is not available in a timely fashion, the alternative is to devise and implement a protective measure against the

attack. If neither of these approaches are possible, a vendor may have to switch to a completely different protocol, or even withdraw the product.

- Users should be able to clear all their sensitive information and configuration details from a product. This means that the risks to their data are minimised when a product is returned for repair, transferred to a new owner, or disposed of.

2. Protect the product from unauthorised access

All assets used to design and build the product should be protected from unauthorised modification.

Management of access to products throughout their lifecycle reduces opportunities for malicious interference.

This means that data, software, systems and component parts should be protected from compromise during all stages of development, build, test, use and maintenance. The integrity of tools used to build a product is as important as the integrity of any data that supports development.

Particularly for quantum security products, the physical environment where these are developed, built and tested needs to be secure, as quantum security relies on physical aspects of the hardware.

Unauthorised or unexpected changes to a product can be detrimental to the security it offers, and it is much easier to gain confidence in products when they function consistently, conforming to known build specifications that only change when new versions are formally approved for release.

Example defensive measures

- Using controlled processes for incorporating changes into both production releases and the tooling used to build a product provides confidence that the product can be built consistently over time. These processes should take into account the relative levels of trust you have in development, build and deployment systems and associated data sources.
- Access to data and systems should be effectively controlled, using the principle of least privilege, and access should be removed for those who no longer need it. Different individuals have different levels of influence across a product's lifecycle. Measures to manage the insider threat should be considered alongside threats from outside the organisation.
- Long-term sensitive values held in a product, such as cryptographic secrets, should be protected from exposure to unauthorised parties (see principles 1 and 2 of product design and functionality). Sensitive parameters that are generated through the build process and

need to be retained to support ongoing development should also be appropriately protected.

  - Quantum keys, random numbers or other sensitive values need similar protection, after their generation and before their (single) use.

- Your processes should minimise opportunities for attackers to gain knowledge of weaknesses in your product. Details of weaknesses revealed during development should be restricted to those who have a need to know. Opportunities to discover additional weaknesses can arise from exposure of detailed technical product information, so you should limit this only to what is strictly necessary during sales and marketing.

  - Quantum side channels in quantum security products need suitable mitigation when discovered. Simply hiding these and hoping attackers don't find them would undermine the whole purpose of quantum security. When quantum side channels are discovered, they should be promptly addressed and published. Root causes should be analysed to determine if they are endemic and ensure that similar side channels are avoided in the future.

3. Monitor services used by a product

Maintain awareness of the security status of all services used by the product.

A product may be dependent on other services to operate, with some being critical to maintaining security. For some services, availability will be paramount. In other cases, services may increase the attack surface of the product, especially as more becomes known about the service and its connection to a product.

Some services – for example, web hosting and cloud storage - may be offered by established global service providers, while others might be provided by the same organisation that develops the product. Outsourcing service provision can offer many advantages but it is important to ensure that outsourced services respect the security needs and expectations of users throughout the lifetime of a product.

This principle and the associated defensive measures apply equally to conventional aspects of quantum security products that utilise conventional services. In addition, it can be envisaged that future device-independent quantum security products might use external quantum services, such as third-party-supplied entanglement. In such scenarios, this monitoring principle is already defined to be part of the protocol. Entanglement can be certified, for example by the quantum product(s) testing for violation of a Bell inequality, without requiring knowledge of the third-party source. Protocols utilising such entanglement therefore specify that it is monitored, in order to verify and maintain the security.

Example defensive measures

- Continuous monitoring of the security status of all services on which a product depends means you can respond promptly if new risks are identified. “Right of audit” clauses in contracts, exercised periodically, can help you maintain confidence in a service providers’ security posture.
- You should have processes to help you maintain awareness of planned or actual changes to services or service providers. These could include internal changes in their technology or the details of their service provision, but also external factors such as the geographical location of their operations, or use of new subcontractors.
- Accreditations or certifications from relevant bodies, maintained throughout the lifetime of your product, can help you gain confidence in the security of services your product depends on. If necessary, you can supplement this by commissioning your own security assessments.
- You should require service providers to give immediate notification in the event that the security of their service is compromised. This will enable you to assess the impact for your product, and implement mitigations to protect users promptly.

4. Review and improve the security offered by a product

Put mechanisms in place to ensure product security keeps pace with external developments.

Activities conducted during development can give users and developers confidence in the security of a product when it is launched. However, the nature of threats, and of the state-of-the-art available to attackers looking to exploit a product will evolve through its lifetime. So, security measures should be continuously reassessed over time against the prevailing threat environment, and developers will need mechanisms in place that enable them to update and upgrade their products.

Ideally, you will be able to identify and mitigate security vulnerabilities before an attacker can exploit them. Product developers can gain information on vulnerabilities from a range of sources:

- their own analysis of known defects
- reports from users
- work by security researchers
- ongoing testing

Importantly, mitigations do not have to be reactive. You may have advance notice that a particular technology which you depend on is coming out of support. You can use this time to develop alternatives.

This principle applies to quantum side channels[1] in quantum secure products, and so the defensive measures identified below can be adapted for side-channel vulnerabilities.

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

Example defensive measures

- Performing security-focussed testing and verification regularly, informed by the threats you expect your product to be exposed to, can help ensure that security issues do not recur once they have been encountered for a first time. You should also provide mechanisms for easy disclosure of vulnerabilities, from users and security researchers, so that mitigations can be developed before flaws can be exploited.
- Studying user interactions with a product can help you learn where it is not being used as expected, to the detriment of security, and will enable you to implement changes that lead to better security outcomes. Good usage monitoring (see product design principles) can help you actively identify attempts to bypass or subvert security measures.
- Maintaining and regularly reviewing a log of all defects in a product will enable developers to identify security implications and prioritise remediation. Following a clear roadmap for security improvements will help ensure the product continues to offer the best possible protection as it is developed over time.
- Product updates that address security issues should be made available promptly, to minimise the window of opportunity for an attacker. Investing in additional capacity at design time means that improvements can be implemented without being constrained by performance or storage limits.

5. Be prepared to respond to external events

Plan for failures, vulnerabilities, and supply chain changes.

Not everything that affects the security offered by a product is within the control of developers, users, or the product itself. Having well-rehearsed plans in place to deal with unexpected events that have security implications will help manage such risks.

External events could include vulnerabilities being released publicly without warning, or technical breakthroughs that the product has not been designed to defend against. They can also arise from changes that affect suppliers and services the product depends on.

Quantum security products should be immune to future quantum technical breakthroughs. Any new quantum technology that might emerge in the future will not undermine an existing

security proof. Therefore, other forms of external event (e.g. new quantum side channels) should be the focus in applying this principle to quantum products.

Hardware components may become hard to source, support for software may end, third-party supplies may suffer a security incident, or manufacturing premises may suffer unexpected damage. All these things have direct implications for the ability of a developer to maintain the security offered by the product.

Vendors need to consider that replacing a component with one that works by a different principle may bring new (potentially yet-to-be discovered) side-channels.

Example defensive measures

- Maintain a register of all components supplied by a third party, including those that are free or open source, and routinely check for vulnerabilities that affect them. Vulnerability news feeds and threat intelligence services can help you maintain awareness of trends in cyber security and develop plans to address any relevant risks
- An obsolescence management plan, based on the known, supported life of components, should keep an up-to-date record of options that can be pursued to replace them. You should take care to ensure that you are acquiring replacements from reputable sources.
- Third party suppliers should only be allowed to retain the information they need, and have access to your systems that is necessary for them to fulfil their role. Accurate records of third-party information holdings can help you assess the impact if they are compromised.
- Business continuity plans should be updated regularly, and you should ensure you follow the measures defined in them. For example, creating secure offsite and offline back-ups will help with disaster recovery, but only if they are updated frequently and the restoration process is tested regularly.
- You should be ready to respond to vulnerabilities in tool chains by assessing the potential impact and take mitigating actions, including installation of updates or patches. It is sensible to plan for tool chain elements becoming obsolete, which may result in loss of functionality or, if cloud-based, complete unexpected withdrawal.

    - In extreme cases, external events could render quantum security products inoperable. Examples include the complete removal or breaking of the quantum channel, or the cessation of the supply of third-party entanglement. In these extreme cases, (quantum) service is denied, so quantum security products should default to conventional security, rather than cease to function entirely.