HAPKIDO

Deliverable 1.1

A method for assessing the risks of quantum computing on public key infrastructures

Team work-package 1

Tom Klunder – TU Delft

Nitesh Bharosa – TU Delft

Yoram Meijaard – TNO

Marieke Klaver – TNO

Contents

List of abbreviations

BIA	Business Impact Analysis
CA	Certificate Authority
CARAF	Crypto Agility Risk Assessment Framework
CIA	Confidentiality, Integrity, Availability
DH	Diffie-Hellman key exchange method
ECC	Elliptic Curve Cryptography
eIDAS	Electronic Identification And Trust Services
GRNV	Integrated Risk Analysis National Security
HAPKIDO	Hybrid Approach for quantum-safe Public Key Infrastructure Development for Organisations
PKI	Public Key Infrastructure
Post quantum computing	Post Quantum Cryptography
RA	Risk Assessment
Risk	Likelihood*impact
SECRAM	Security Risk Assessment Methodology
SRA	Societal Risk Assessment
SWOT	Strengths, Weaknesses, Opportunities, Threats
TSP	Trust Service Provider
Quantum computing	Quantum Computing
QTSP	Qualified Trust Service Provider

Summary

Why do we need a method for assessing the societal risks posed by quantum computing on public key infrastructures?

Many sectors like government, banking and telecom today rely on digitalized processes. Examples include filing taxes or paying for online purchases. These kinds of digitalized processes heavily depend on cryptography. Cryptography is one of the ways information and communication processes are secured by employing encryption. It prevents eavesdroppers from listening in and ensures that the received data is trustworthy. In addition, cryptography enables trusted online transactions such as e-commerce and digital document signing, for instance, when signing a lease contract or bank loan. In other words, cryptography is the basis for our digital society.

Many large-scale cryptographic systems are organized using a Public Key Infrastructure (PKI) model. PKI is a combination of software, hardware, roles, guidelines and processes. One of the PKI processes is Key Management. One of the mechanisms to distribute keys (the public portion) are digital certificates. In essence, PKI is the infrastructure needed to manage and distribute public/private key pairs. Software applications use these key pairs. The usage of PKI for creating digital trust is widespread and interwoven in our society. It is extensively used in organisational identity management, internet security (e.g. website & server), secure email, VPNs & intranets, software updates, the Internet of Things, healthcare, finance, and critical infrastructures.

Accordingly, it is not an exaggeration to say that PKI based cryptography lays the foundation for trust in our digital world. At the core of PKI is asymmetric cryptography. One level deeper, at the foundation of conventional asymmetric cryptography lie two mathematical problems. The hardness of these problems guarantees the security of the cryptographic schemes that use them. Reasoning back, the hardness of the two mentioned mathematical problems is what most digital trust systems, by using PKI, depend on. In other words, the main strategy for information security is to use mathematical problems that – in practice – cannot be solved with the currently available computing power. But what if these mathematical problems are not so hard to solve anymore? If that is the case, the cryptographic schemes cannot guarantee information security, i.e. confidentiality, integrity and availability. And if information security cannot be guaranteed anymore, there is no trust.

Such a scenario is becoming realistic due to the rise of quantum computing. Because of the fundamental differences in the way of computing compared to classical computing, quantum computing can apply Shor's algorithm and solve two classically hard mathematical problems with ease. As quantum computers are still being developed around the world, a quantum computer to powerful enough to crack today's cryptographic standards does not yet exist. However, experts estimate the chance that it is developed before 2030 small, but realistic. Some researchers suggest that it is just a matter of getting over the most difficult parts. Developing one qubit was the culmination of decades of work, then figuring out how to double the one qubit into two would be extremely difficult, but once that is done and we know how to double qubits, there is suddenly not so far from two to five, and from five to fifteen. Some countries share a clear ambition about this. Denmark's first fully functional generally applicable quantum computer will be available in 2034. This is the objective of the ambitious Novo Nordisk Foundation Quantum Computing Programme that is being launched in collaboration with the University of Copenhagen[1].

[1] https://nbi.ku.dk/english/news/news22/major-investment-for-developing-denmarks-first-fully-functional-quantum-computer/

The Dutch General Intelligence and Security Service (AIVD) follows this vision. The quantum threat, meaning the looming ability to crack today's cryptographic standards and thereby break digital trust using quantum computing, is real. It is expected that the quantum threat materialising will have devastating impact to society. However, it is not described in literature what exactly the consequences for society will be, aside from relatively general statements about areas impacted. Rather, most studies focus on the technical impacts. There is a need to know more about the potential societal impact, so (1) it becomes clear how pressing and far reaching the quantum threat can be and (2) help stakeholders in collaborating and transitioning towards quantum safe PKI systems.

What: the research objective

The objective of this study as part of the HAPKIDO project is to develop and test a method for assessing the societal risks of quantum computing, particularly focussing on domains that use PKI systems. An important distinction to make in potential societal risks is those that result from PKI failing because of the quantum threat and those that result from mitigation and migration efforts. This research is limited to the former. As such, potential solutions to the quantum threat such as quantum key distribution, post-quantum cryptography, and hybrid cryptography are out of scope.

With the objective and the scope of the research set, the following research question is defined: *What is a suitable method for assessing the potential societal risks of the quantum threat to PKI systems?*

How: design science research

While there are several societal and risks assessments methods in literature, there is no ready to use method tailored to assess the societal risks posed by quantum computing on PKI systems. Therefore we chose to develop a tailored societal risks assessment (SRA) method based on the design science research approach. We do this via three design science research cycles that include literature review, co-creating and iterating versions of the artefact, and multiple evaluation workshops with stakeholders throughout the design process. The idea of co-creation from Design Thinking is applied by engaging in discussions during workshop sessions with intended users. The SRA borrows concepts and insights from existing methods, including the organisational impact score matrix, the confidentiality-integrity-availability structure from SecRAM and parts of the underlying method of the Netherlands' National Risk Assessment. The SRA was completed after six workshops with actors performing different roles within a PKI operating in the Netherlands. Actors that apply the method can do so to explore multiple risks categories, including risks to the assessing organisation, risks to society that are within the sphere of influence of the assessing organisation and the risks to society that arise from the presence of the quantum threat.

Results

The resulting SRA consists of six steps: (1) determining the scope, (2) identifying the threats & vulnerabilities, (3) identifying the assets, (4) assessing the risks, (5) assessing the likelihood, and lastly (6) synthesis. Each step is an interactive process that produces the knowledge necessary to arrive at the outcome: an overview of risks to the organisation and to society within the sphere of influence of the assessing organisation. After six workshops, the proposed method was evaluated to be useful, yet more guidance would be needed in order to fully apply the method in practice. This document describes the method. The results of the application of the method on PKI systems are discussed in a separate HAPKIDO deliverable (D1.2).

1 Introduction

1.1 Background

Quantum computing refers to a rapidly developing set of technologies with unprecedented potential for tackling computational problems that conventional computing cannot. It promises to provide great computational power if harnessed, facilitating scientific breakthroughs in many fields (Gill et al., 2021; Vermaas, 2017). However, while it is expected to provide benefits in cybersecurity, it is simultaneously expected to threaten the security of our digital society (Raban & Hauptman, 2018). This may seem paradoxical, but it is not.

Currently, most online communication is secured by leveraging cryptography that is hard to break. The industry standard cryptographic schemes can theoretically be broken, but this takes thousands of years in practise. After decades of trying, nobody has found a way to do so in a reasonable time. However, whilst it is not capable yet, quantum computing is fully expected to break the current standard cryptography quickly. This significantly threatens the security of our digital society. We will explain why this *quantum threat* is an issue.

Part of the cryptographic standards is public key infrastructure (PKI). PKI is "a combination of software, hardware, roles, guidelines and procedures that are required to manage keys as digital certificates" (Bharosa et al., 2015). Essentially, it is a way to put cryptography as a technology to use and create digital trust. This will be further detailed in Section 2. PKI plays a crucial role in digital society as we know it. It allows online authentication so that only you can access your social media account[2], secure communication of sensitive data such as between medical professionals, and legally binding electronic signatures or financial transactions. PKI usage is very much interwoven in many if not most of our online activities, especially those requiring high trust.

1.2 Unknown societal risks

More specifically, this research investigates the consequences of the quantum threat on PKI *to society*. The technological consequences (i.e., public key cryptography standards easily being broken) are explained above, but in what ways this will impact society is less clear. The term risk is more suitable, as it implies on the one hand an element that we would like to protect and on the other hand uncertainty. This applies as digital trust and its benefits are the things of value that we would like to protect and there is great uncertainty surrounding the development of the quantum threat: a large enough quantum computer to break current public key cryptography standards.

The societal risks are not clearly known, which increases the vagueness of the problem at hand. A vague problem causes uncertainty in solving it. There is a need to know more about the societal risks, so (1) it becomes clear how pressing the matter of the quantum threat is and (2) where to focus activity to mitigate its impact.

We can accomplish the first task by creating a shared picture of what the societal risks are in cooperation with the relevant actors who are involved with PKI. It is when we do this we start to see the problem perceptions of these actors align. The alignment of problem perceptions bolsters the actor's commitment to solving the matter through cooperation (Bruijn & Heuvelhof, 2008). This is particularly necessary regarding the quantum threat to PKI, as PKI inherently involves many actors with differing roles in the trust chain.

[2] Dismissing the possibility that somebody else knows/guesses your login credentials

The second benefit of knowing the societal risks is that it helps stakeholders to prepare for the transition to a solution. Logically, by knowing what specific instance of PKI causes the worst societal risks, you know where to prioritise in mitigating the quantum threat. These two points clearly show the societal relevance of knowing the societal risks of the quantum threat.

1.3 Societal Risk Assessment

There is little research on the societal risks of the quantum threat to PKI, which makes this research all the more necessary. Additionally, there is no developed method available to assess the societal risk. This research aims to provide a method to assess these risks: a Societal Risk Assessment (SRA). In other words, the objective of this research is to develop and test a method for assessing the societal risks of quantum computing, particularly in domains that use PKI systems.

In the previous section, the societal relevance of knowing the societal risks of the quantum threat is established. Consequentially, providing a sound method to uncover the societal risks is of societal relevance as well.

1.4 Scope

PKI is a broad subject, therefore some scoping is required. This report is part of NWO project HAPKIDO work package 1, and we made several scoping choices accordingly. First, we focus on PKI systems in the Netherlands. Even though this research focusses on the Netherlands, we expect that findings can be transferred to other EU and non-EU countries since most countries use the same public key cryptography standards.

Secondly, this work-package focusses on PKI as opposed to public key cryptography in general. More specifically, this research limits itself to PKIs with a hierarchical trust model, which is a very widely adopted paradigm (Amadori et al., 2022).

Thirdly, we consider only qualified trust services as defined by the European eIDAS regulation[3]. The eIDAS regulation is being adopted throughout the EU standardising trust services. Market parties expect the need for higher levels of trust, and thus qualified trust services, to keep growing. As the quantum threat is a threat of the future and qualified trust services are expected to become the solidified standard, the scoping choice makes sense. Additionally, qualified trust services have the strictest requirements, providing the most trust, and therefore are likely to have the largest societal impacts if that trust is broken. These are provided by Qualified Trust Service Providers (QTSPs).

Fourth and finally, the scope excludes self-signed CAs used for applications within an organisation. It is assumed that these internal applications are more easily adapted to deal with the quantum threat, as there will be less governance struggles because there is no large actor network involved. Therefore, they are less relevant to finding the most important societal risks.

1.5 Research question

With the objective and the scope of the research setting, the following research question was defined: **What is a suitable method to assess the potential societal risks of the quantum threat to PKI systems? (RQ1).** Answering this question is the focus of this report (HAPKIDO Deliverable 1.1).

Note that, there is a second deliverable for work package 1 (deliverable 1.2), which focuses on the application of the SRA method proposed in this document. Hence, deliverable 1.2 focusses on the question 'What are the potential societal risks of the quantum threat to PKI in relation to societal

[3] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2014.257.01.0073.01.ENG

effects in the sectors of government (public services), banking and telecommunication in the Netherlands?' Therefore, this report does not discuss the application of the developed SRA method in practice (this is done in D1.2). Other works packages in HAPKIDO focus on the requirements for QS architectures, possible migration and governance strategies.

1.6 Reading guide

This document proceeds as follows. Section two elaborates on the research approach followed to develop the SRA method. Section three presents the results of the literature review. Section four presents the final version of the SRA method based on the steps included in the method. Section five reflects on the evaluation results for the SRA. Section six concludes with a discussion and recommendations. The appendix provides more detail on the results of the workshops conducted.

2 Research approach

2.1 Design Science as a research approach

The design science research approach as described by Hevner et al. (2004) was used to perform this study for three reasons:

1. Literature does not provide for a ready-to-use SRA method for assessing the risks of quantum computing on PKI systems. Nonetheless, the literature does provides a plethora of risks assessment theories, instruments and techniques that can be useful for developing the SRA method we seek.
2. For the SRA method to be useful, we need to find out what the user group for the method looks like and what concerns and constraints they have. Hence, we need a research approach that allows for interaction with representatives from the prospective user group.
3. Method engineering is not a first time right process; you need several rounds of design, development, application, and evaluation before the method satisfies the requirements. The design science research approach provides the necessary flexibility for doing so.

Design Science research promotes three cycles. These are the relevance cycle, the rigor cycle, and the design cycle. These are depicted in Figure 1.

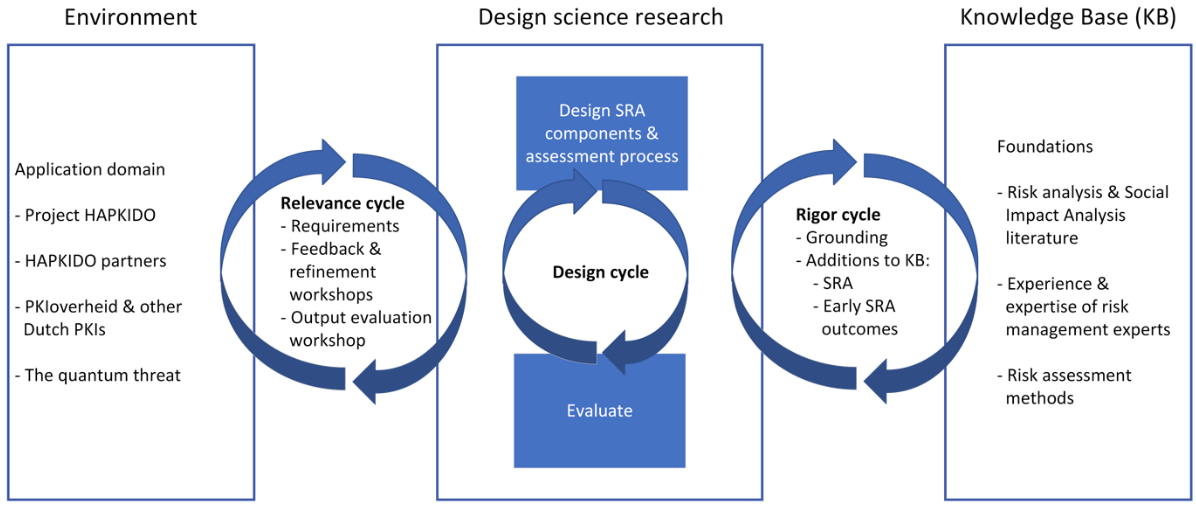
Environment
Design science research
Knowledge Base (KB)
Application domain
- Project HAPKIDO
- HAPKIDO partners
- PKIoverheid & other Dutch PKIs
- The quantum threat
Relevance cycle
- Requirements
- Feedback & refinement workshops
- Output evaluation workshop
Design SRA components & assessment process
Design cycle
Evaluate
Rigor cycle
- Grounding
- Additions to KB:
- SRA
- Early SRA outcomes
Foundations
- Risk analysis & Social Impact Analysis literature
- Experience & expertise of risk management experts
- Risk assessment methods

Figure 1: Cycles in design science research (Adapted from Hevner, 2007)

The relevance cycle connects the designing activities with the application environment of the research project, so that relevancy of the research is ensured. Usually, design science research spawns from problems or opportunities in an application environment. When the need for an artefact arises, it initiates the design process. In the case of this research, there is a need for an SRA that emerges from the context of our digital society that is vulnerable to the quantum threat. After the initialisation, the relevancy cycle serves to keep the artefact being designed connected to the contextual reality. This happens throughout the design process so that the artefact will be well adjusted to the contextual needs and thus useful.

The rigour cycle draws existing knowledge from the knowledge base to fuel the design process. This knowledge base consists of expertise, pre-existing artefacts, and theories. Then the lessons learnt and artefacts produced during the design process add to the knowledge base. The design science researcher should explore the state-of-the-art in the application domain, so that the design science research may provide scientific knowledge contributions rather than routine design. In this research, it is necessary to learn from previous risk assessment methods and theories to come to an effective SRA.

The last of the three cycles is the design cycle. This cycle is at the core of design science research. It iterates between generating new design alternatives, evaluating them, and considering which way to proceed with the artefact. To support these activities, the design cycle takes input from both the relevance and the rigor cycles as described above. In turn, it provides input to the other two cycles as well. The intermediary artefact is used in the relevance cycle for field-testing and its results steer the design cycle. The artefact is based on lessons from the knowledge base and adds lessons obtained during the design process to the knowledge base in the form of experiences from field testing, extensions of theories, and new artefacts.

2.2 Literature review

In this research, the rigor cycle from Hevner's model (2007) partially takes form as a literature review. The literature review was a continuous process looking into several knowledge domains. The Scopus and Google Scholar databases were used to search for relevant literature. The snowballing method is applied to find newer and older publications. The results of the literature review is presented in section three.

The first part of the literature review is part of the relevance cycle. As described in section 2.1, the design science research is initiated by the need for an artefact that stems from a problem in the application context. In this case, the problem is the lack of insight into the societal risks posed by the quantum threat. To better understand the application context and thus increase the relevancy of the SRA, the current cryptographic standards and how they are threatened by quantum computing are explored. Additionally, an attempt is made to explore literature on the societal risks of the quantum threat, or rather the lack thereof.

Then, the literature review investigates ways in which the (societal) impact of technology are assessed. To do so, literature on risk analysis and on societal impact assessment is consulted. From this literature, lessons are drawn to support the structure and building blocks of the SRA. This is an apparent expression of the rigor cycle.

After a first exploration of the potential methods that may serve as inspiration to the SRA, a draft version of the SRA was made. This very first version served as a stepping-stone to later iterations based on feedback from field experts. This feedback was gathered in two types of workshops: expert workshops and user workshops. Both are discussed next.

2.3 Expert workshops

This study builds on six workshops with experts. Here experts refer to the team working on Work Package 1. The experts have deep knowledge of one or more of the following topics:

1. Public key infrastructures
2. Trust services
3. Risk assessment methods
4. (post)Quantum computing

The four workshops took place between September 2021 and March 2022. The workshops included between 2-6 experts, and took about 90 minutes each. In each workshop, participants were asked to reflect on how the available version of the SRA could be further improved in order to satisfy the requirements (see section 4.1).

2.4 User workshops

The second kind of workshop used in this study focusses on getting feedback from potential future users of the SRA method. These includes representatives from one of the HAPKIDO application domains: government, banking and telecom. Specifically, users include persons active in providing or making use of trust services, enabled by PKI. Each workshop was video-recorded in order to analyse improvement suggestions after the workshop. The following table provides an overview of the expert workshops.

Table 1: Overview of user workshops

User Workshop	SRA version	PKI domain	Participants	Date
1	0.1.	Government	PKI manager at PA, manager at TSP, PKI researcher	02-11-2021
2	0.2.	Banking	M2M communication program lead at a large bank	13-12-2021
3	0.3.	Telecom	PKI manager at QTSP, compliance officer at QTSP, technical expert at QTSP	12-01-2022
4	0.4.	Government	M2M communication program lead at tax authority, risk management expert	26-01-2022
5	0.5	Government	PKI manager at PA, cyber security scientist	24-02-2022
6	0.9	Government	PKI manager at PA, cyber security scientist, Researcher	13-04-2022

The workshops took approximately 90 minutes each. During these workshops, input is gathered on how to satisfy the requirements for the SRA method (see section 4.1). The workshops also serve to implement DT ideas of co-creation. They do so by placing extra emphasis on frequent contact with the intended user-base, engaging in discussions, and establishing agreed upon design requirements.

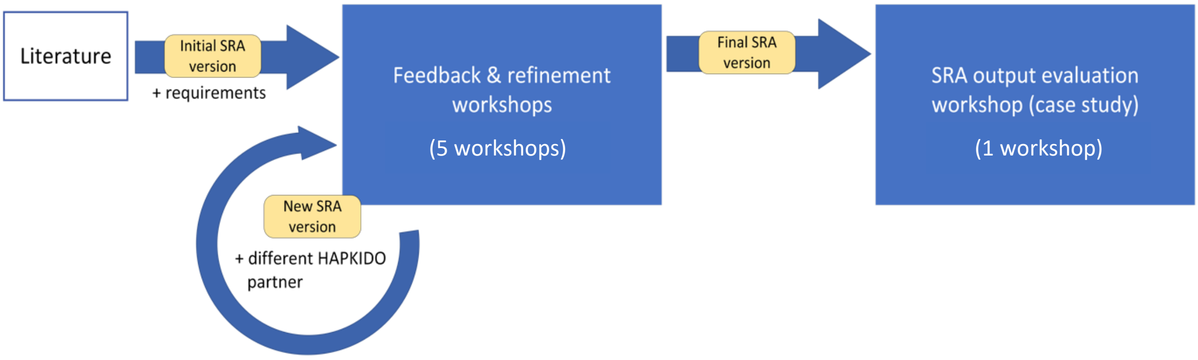
Literature
Initial SRA version
+ requirements
Feedback & refinement workshops
(5 workshops)
New SRA version
+ different HAPKIDO partner
Final SRA version
SRA output evaluation workshop (case study)
(1 workshop)

Figure 2: Schematic overview of workshops

Every workshop served as an opportunity for evaluation of the artefact in the process of being designed. Leveraging the SWOT analysis technique, after every workshop the artefact is refined. Using this technique allows for bolstering the strong aspects of the artefact and improving the weak aspects. This constant evaluation and iteration of adjustments resembles the design cycle.

2.5 Evaluation

After every workshop, feedback was structured in SWOT analyses by listening back the recordings of the workshops. These SWOT analyses are documented in Appendix C. In the last workshop, the participants were asked to collaboratively fill in a SWOT analysis themselves. The results of all SWOT analyses are discussed in section 5 (Evaluation).

Based on the executed SWOT analysis, the SRA method has been updated and simplified. In particular, the abstraction level has been raised slightly to make the SRA less complex. The SRA method presented in section 4 is the simplified SRA method.

3 Literature review

This section describes the relevant literature found and how it relates to this research. Firstly, it will describe the current cryptographic landscape and how quantum computing can have an impact. Then, the scarce research on the societal impact of the quantum threat is highlighted. Next, the literature is analysed. All lessons drawn from the literature that are directly implemented in the SRA are summarised in Table 2.

3.1 Current cryptographic standards

The security of our digital society is largely dependent on cryptography. Here we need to consider symmetric and asymmetric cryptography. Both are used to encrypt our data, from highly confidential state secrets to stored grocery lists. When using symmetric cryptography, data is encrypting with a shared secret, or a key, that renders data to be unreadable for those without the shared key. Later, using the same key, the process can be reversed or decrypted, so that the original readable data is recovered. The original goal of cryptography is to ensure none other than the intended recipient can read the message. This is called confidentiality. There is only one currently known cryptographic cipher that is perfectly secure, meaning that it is theoretically impossible to break it. This is the Vernam cipher. This sounds like the perfect way to secure information, but there are two drawbacks.

The first drawback is key size. Using the Vernam cipher requires a key that is the same size as the data. This means that sending a confidential file of 200MB requires sending the file plus a 200MB sized key, effectively doubling the necessary capacity of electronic communication. This is extremely inefficient and thus not practical for day-to-day use. To solve this issue, many new ciphers were invented, leveraging computational problems. These are mathematical problems that cannot realistically be solved by conventional computers. Even if you were to use huge amounts of the computational power available today, breaking such ciphers would still take hundreds or thousands of years. This means that even though the ciphers are not theoretically secure, they are practically secure.

The second drawback of the Vernam cipher is that it is a form of symmetric cryptography. The result is that when communicating, it requires both parties to know the key. This means that they need a way to securely communicate about the shared keys, before being able to securely exchange other types of data. This is called the key distribution problem. One solution to this problem is asymmetric cryptography.

Asymmetric cryptography is a form of cryptography that uses key pairs, consisting of a private and a public key. Here, we can use two people, Alice and Bob, as examples. Alice can encrypt a message to Bob with his public key, which can now only be decrypted with Bob's private key. This is also called public key cryptography. Public key cryptography is a necessity for secure internet standards such as TLS and PGP. The most popular form of public key cryptography is RSA, which relies on the hardness of prime factorisation. Other popular public key cryptography schemes are Diffie-Hellman (DH) and Elliptic Curve Cryptography (ECC). These rely on the hardness of the discrete logarithm problem. In case of both problems, computation becomes infeasible once the parameters (i.e., key size) are large enough.

Aside encryption, public key cryptography can be used for signing. Doing so proves that a certain message is written by the owner of the private key that belongs to a certain public key. Alice uses her private key to create a signature based on a message. Bob can then use Alice's public key and the signature to verify that it is Alice who signed the message. More information on cryptography and its workings can be found in the book Cryptography Made Simple by Smart (2016).

Observant readers may have noticed that public key cryptography requires knowledge beforehand as well. After all, how can you be sure that an insecurely communicated public key is in fact the public key of the intended recipient? There may not be a need for a shared secret, but there is a need for trust in the public key. This is where Public Key Infrastructure comes into play.

3.2 Public Key Infrastructure

There is a need for trust that a certain public key belongs to a certain party. In the current digital landscape, this trust is provided by Public Key Infrastructure (PKI). PKI is "a combination of software, hardware, roles, guidelines and procedures that are required to manage keys as digital certificates" (Bharosa et al., 2015). PKI is a system that brings together actors, using institutions, and cryptographic technology to achieve digital trust. At the basis of PKI lies assumed trust in a trusted third party (TTP). This TTP signs the certificate (i.e., the public key) of Bob. Alice can now verify that the certificate presented by Bob is signed by the TTP. As Alice trusts the TTP, she can trust that the certificate presented by Bob is in fact Bob's. In this scenario, the TTP fulfils the role of root Certificate Authority (CA). There may be another party that is trusted by the root CA to sign certificates, which sign Bob's certificate. In this scenario, the other party is an intermediary CA. The concept presented in this scenario is the chain of trust. It is crucial that CAs are strict in their policies that dictate which certificates are signed for whom, as otherwise trust is lost.

There are many parties active as root CA. Their certificates often come pre-installed on consumer devices and in web browsers. The collection of pre-installed certificates on a device or in an application is called the root store.

There are several functions that need to be fulfilled in a PKI by several roles. These include certificate usage, certificate revocation, certificate generation, and setting the rules for all of these.

Aside from confidentiality, there are other purposes that are enabled by PKI. These are the integrity of data, authentication of persons, organisations, and systems, and non-repudiation. Altogether they serve to enable trustworthy and legally binding electronic transactions. This in turn enables business and public administration to be conducted online, heavily improving the efficiency of society. This is evident from the massive adoption of PKI and how interwoven it is in society. It is extensively used in organisational identity management, internet security (website & server), secure email, VPNs & intranets, software updates, the Internet of Things, healthcare, finance, and critical infrastructures (Mulholland et al., 2017).

3.3 The quantum threat

Now that the current cryptographic landscape is clear, the vulnerability to quantum computing can be explored. As mentioned, the present-day digital trust relies on public key cryptography, which relies on the hardness of the prime factorisation and discrete logarithm problems. These two mathematical problems are hard for conventional computing. However, quantum computing presents a fundamentally different way of computing, which allows Shor’s algorithm to be applied in order to break or solve the mathematical problems used for encryption. Shor’s algorithm can efficiently solve both the prime factorisation and discrete logarithm problems. In turn, current public key cryptography provides no security, as it is easily broken in theory. For now, there is no quantum computer available with enough processing power to attack current public key cryptography standards. However, technological strides have been made over the past years which make the advent of such a quantum computer in the future very likely (Skosana & Tame, 2021).

When exactly the efforts of the scientific community will come to fruition is anybody’s guess. There are many different views among experts of how long such endeavours will take. However, experts generally agree that it will happen eventually. In a report on the expected timeline of quantum computing, Mosca & Piani (2021) found that more than half of the inquired experts (n=44) thought that such a quantum computer is unlikely to exist before 2030. On the other hand, a quarter thought that it was 50% likely or more. The near consensus found in the report is that the threat is likely to occur before 2040. 36/44 respondents indicated they ought the threat likely, very likely, or extremely likely to occur before 2050. The Dutch General Intelligence and Security Service (AIVD) follow this vision and advises parties that have an information security need to act accordingly (AIVD, 2021).

Overall, there is a small but realistic chance that by 2030, and otherwise by 2040 or 2050, the foundation of digital trust will be no more. As illustrated before, our society and economy depend on reliable digital financial transactions and confidential communications, which cease to exist without trust. Simply put, disaster would strike (De Wolf, 2017).

3.4 Post-Quantum Cryptography & Project HAPKIDO

The vulnerability of conventional digital infrastructure to the quantum threat is a worldwide problem. In the United States, NIST is developing standards expected to be final by 2024 for new public key cryptography schemes resistant to quantum computing. To transition to quantum-safe PKI systems in the Netherlands, the research project HAPKIDO is running. Solutions like Post-Quantum Cryptography, Quantum Key Distribution, and Hybrid Cryptography offer promising prospects (Amadori et al., 2022).

Yet, before we can look to future solutions, the current issue must be clear. The consequences of the quantum threat could be extreme, but it is unknown how these consequences could play out in more detail. This research is part of project HAPKIDO, investigating the consequences of quantum-unsafe PKI in the scenario where a large quantum computer is in existence.

3.5 The societal impact of quantum computing

The three articles of De Wolf (2017), Raban & Hauptman (2018), and Vermaas (2017) are concerned with the impact of quantum computing in general. In 2017, Vermaas wrote an article about the societal impact of quantum computing. In this article, they describe the potential positive impact of quantum computing through advancements in science. These advancements are expected to happen because of the increased simulation power quantum computing is expected to hold. They also describe the fact that Shor's algorithm is expected to break current cryptographic standards which are used to secure financial transactions and governmental and company secrets. Furthermore, they describe the changes in cryptographic implementation that are expected to happen to counter the quantum threat. However, this article serves more as a call to societal debate and suggestions on how to do so, rather than an explicit review of the consequences felt by society if the quantum threat to cryptography is not mitigated. In the same issue, De Wolf (2017) elaborates on why the impacts mentioned by Vermaas will happen. However, they also refrain from describing on a detailed level what the societal impact of failing cryptography are. Raban & Hauptman (2018) mention the expected impact of quantum computing on cybersecurity as part of a larger literature review looking into the future of cybersecurity. Again, there is not more detail other than that it is likely to benefit the offensive capabilities in cyber-attacks. They do however mention the defensive capabilities of quantum computing in that it offers ways to make communication *perfectly secure*.

Then there are the studies of Joseph et al. (2022), Mavroeidis et al. (2018), Mulholland et al. (2017), and Yunakovsky et al. (2021) that focus on the quantum threat to cryptography. Of these, Mavroeidis et al. (2018) and Joseph et al. (2022) remain rather general in their mentions of the societal impact of breaking asymmetric cryptography by quantum computers. Additionally, most of their analysis of consequences is rather technical and has little to do with the societal aspects of the technological implementations. Mulholland et al. (2017) and Yunakovsky et al. (2021) are more specific in their discussions of the impacts of the quantum threat. However, both lack the connection to the consequences on a societal level. All of these studies consistently place focus on the technological impacts of the quantum threat. While the technological aspect is important, as it is where societal impact stems from, it should not be forgotten that the real value of technology lies in its implementations in society. This is something that the SRA aims to shed light on.

Lindsay (2020) places quantum technologies in an international intelligence context and brings a new perspective to societal impact. On the one hand, it may be end of privacy, because of quantum computing breaking cryptography. On the other hand, it may be end of intelligence, because quantum computing can enhance secure communication to a new level, making it even more difficult or impossible to decrypt intercepted messages. Perhaps reality will lie somewhere in-between. Alarmists may warn for one or the other end of the spectrum, but the practical and human sociotechnical reality in which technology lives renders their arguments void. There are simply too many factors that reduce the effectiveness of cyber offence and defence, also in the face of quantum computing. Even in the case of a powerful quantum computer being realised before quantum-safe solutions are in place, human organisation and strategic interaction will prevent major shifts in geopolitical balance.

Although no earth-shattering geopolitical weight changes are expected, there is still value in finding the societal risks on a lower (i.e., national, organisational, and individual) level. Not adopting quantum safe technologies will leave societies vulnerable. Lindsay (2020) says that the importance of PKI in our societies should not be understated. It is a pillar in online banking, software updates, electronic medical records, virtual private networks (VPN) and intranets, remote maintenance on industrial machinery, and the monitoring and control of industrial operations, military & governmental access to secure facilities and top-secret information. These applications are confirmed in the other studies in this section. This once more confirms the necessity of this research. Nevertheless, Lindsay's work points out that we should not be fearing the end of the world with the dawn of quantum computing and this notion should be kept in mind when performing the SRA.

Though not specifically investigating the societal risks of the quantum threat, there are two methods designed to deal with the quantum threat to cryptography. Mosca's XYZ is a simple model to assess the timeline of the quantum threat and provide guidance on when action should be taken (Mosca & Mulholland, 2017). It works by subtracting the shelf life of certain data and the time it takes to implement a quantum-safe version of a system from the time it takes until the threat is (estimated to be) actualised. These numbers are of course not meant to be known accurately, but the model does provide a way to think about the threat and the time there is to act. This model only says something about vulnerability, rather than impact.

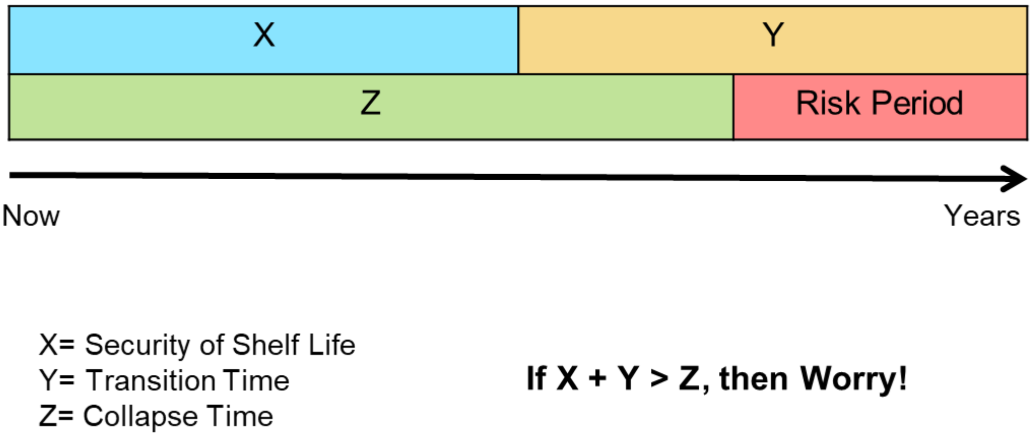

Figure 3: Mosca's Theorem (adapted from Mosca & Mulholland, 2017)

Another more extensive method is the Cryptography Agility Risk Assessment Framework (CARAF). While it is made not to be specific to the quantum threat, it was designed with the quantum threat as the main example. It includes a component in which the assessor estimates the impact of an expected future threat. However, this is only done on an organisational level and does not reflect the potential influence of the organisation on society. This is where the SRA can fill the gap in the knowledge base.

3.6 Social / Societal Impact Assessment

When investigating the ways in which societal impact of technology are assessed, the field of Social Impact Assessment (SIA) presented itself. SIA is concerned with identifying and managing social issues arising from planned interventions (Esteves et al., 2012). Quantum computing as a new technology bound to cause disruptions in people's lives seems a good fit for the field. However, in the case of the quantum threat as the object of analysis, there are two issues. Both issues stem from the origins of SIA being project appraisal. The first issue is that SIA is designed to assess the social impacts of a project, meaning the impacts on humans and their interactions, rather than the broader socio-technical impacts. The broader perspective is necessary in face of the quantum threat, because PKI is a prime example of a system in which technology, institutions, and actors coincide.

The second issue is that SIA is founded on the assumption that the analysed development is planned and controllable (Esteves et al., 2012). More specifically, SIA is centred on the implementation of an artefact. This is a problem for the case of the quantum threat, which is not so much an artefact, but rather an imposed threat to security enabled by technological development. Here, the technological development cannot be logically controlled or constrained, as it is sought after all over the world.

In 2015, Wadhwa et al. proposed a SIA-based methodology adapted for security research and security measure implementation. They propose use of the term *Societal* Impact Assessment, indicating the inclusion of impact on social as well as natural and artificial systems. This approach solves the first issue discussed above, but the methodology presented by Wadhwa et al. is nonetheless prone to the second issue.

The usage of the term 'societal' and its meaning is a useful takeaway from the SIA field. Moreover, knowing why SIA is not a good fit for analysing the quantum threat brings us closer to what is. A change in perspective that highlights the fact that the quantum threat inevitable rather than one possibility is necessary. This is where the concept of risk comes in.

3.7 Risk analysis

While there are many different interpretations of risk in literature (Aven, 2018; Joosten & Smulders, 2014), they share the same two components: values at stake and uncertainties (Aven, 2018). In case of the quantum threat, the many benefits of digital trust and our reliance on it as a society fit in the first component. As for the second component, the big unknown is when quantum computing will be advanced enough to break current cryptographic standards. This illustrates that when discussing the potential societal impact of the quantum threat to PKI, risk (to society) is a fitting concept. Inspired by Wadhwa et al. (2015) and Aven (2018), this research defines societal risk as the values at stake in human, natural, and artefactual systems; their interactions and the accompanying uncertainty. For the sake of ease, values at stake will be termed impact in the rest of this research.

The separation of risk into the two components of impact and uncertainty is apparent in the SRA. Step 3 establishes the expected impact while step 4 deals with the uncertainty. Both components are part of the eventual risks in step 5.

Before continuing to describe how to find the societal risks, some background on risk as a scientific subject is provided. In 2018, Aven wrote an influential article calling for recognition of Risk Analysis as a distinct science. For Aven, risk analysis encompasses "risk assessment, risk characterization, risk communication, risk management, and policy relating to risk, in the context of risks that are a concern for individuals, public- and private-sector organizations, and society at a local, regional, national, or global level". Risk analysis is often considered supplemental to other fields. But by comparing risk analysis to statistics, Aven makes the case that risk analysis is in fact a field of its own. Risk analysis and statistics both have limited explanatory power but help greatly in dealing with uncertainties. Knowledge generation in both fields is often highly related to other fields in practical application. For example, statistics are used to inform public policy in curbing a pandemic and risk analysis aids the safe design of a chemical plant. To further conceptualise this, Aven distinguishes between two types of knowledge generation in risk analysis:

(1) Type A: knowledge related to an activity in the real world, such as the examples given above.
(2) Type B: knowledge on concepts, theories, frameworks, approaches, principles, methods, and models to understand, assess, characterize, communicate, and (in a broad sense) manage risk.

This ties in with the design science research cycles by Hevner (2007). Type A knowledge generation can be linked to the relevance cycle, whereas type B knowledge generation can be linked to the rigor cycle. The design cycle, residing in the middle (see figure), helps to generate both type A and type B knowledge. In this research, the answer to RQ1 is knowledge of type B and the answer to RQ2 is of type A.

One approach to generate type A knowledge recognised by Aven et al. (2018) is predictive analysis. This approach suits the purpose of this research well, as quantum computing and its threat to PKI are future and no similar risks have materialised yet to draw from. When adopting the predictive analysis approach, the following questions are central:

- What will happen if a specific activity is realized?
- What might go wrong?
- Why and how might it go wrong?
- What are the consequences?
- How bad is it?
- What will happen if we (do not) intervene?
- How soon, with what consequences?
- What do we know; what do we not know?
- What are the uncertainties and likelihoods?

The results of a method aiming to uncover the societal risks of the quantum threat should answer these questions.

3.8 Information security risk assessment

To assess the societal risk, a Societal Risk Assessment method (SRA) is necessary. There are many risk assessment (RA) methods available for information security risk. Shamala et al. (2013) describe the shared needs for information in their Information Security Risk Assessment framework. The quantum threat is very much related to information security, as the encryption broken by quantum computing serves it as its main purpose. Another comparative study of information security RA methods was done in 2018 by Wangen et al. They identified three main activities in information security RA: risk identification, risk estimation, and risk evaluation.

The first typically consists of threat, outcome, asset, and vulnerability identification. After these have been identified, the values that go with them are estimated in the risk estimation phase. Lastly, in the risk evaluation phase, the risks are compared and prioritised to finalise the risk assessment process. This provides a basic structure that is proven in practise. Additional to credibility, adopting this basic structure in the SRA promotes two of the design requirements. First, the established information security RA methods were made to be used widely. Basing the SRA on their structure helps the SRA be *usable*. Second, by adopting a structure that is industry standard, the SRA becomes recognisable. This helps the assessor to understand the use of the SRA quicker, making the SRA more *transferable*. The basic structure is adopted as follows: The first phase, risk identification, as recognised by Wangen et al. (2018), is represented by steps 1 and 2 in the SRA. The second phase, risk estimation, is represented by steps 3 and 4. Lastly, the third phase, risk evaluation, is represented by step 5.

Usually in information security RA, threat identification is based on previously identified critical assets (Shamala et al., 2013). CARAF teaches us that when assessing risk related to specific cryptographic vulnerabilities, such as in case of the quantum threat, it is sensible to swap the order. It makes more sense to identify the assets based on the specific cryptographic vulnerability and threat that exploits this vulnerability (Ma et al., 2021). That is why in the SRA, the threats and vulnerabilities are first identified in step 1 and after that, the relevant assets are identified in step 2.

Another lesson from CARAF is that in case of the quantum threat, estimating a probability to capture the uncertainty aspect of risk is not logical. They say that "lack of information about incidents is particularly challenging in the context of crypto agility as the goal is to model the risk of an event that has not yet materialized. For example, let's consider the threat of quantum computing to current cryptosystems. It is not meaningful to consider the frequency of exposure to quantum computing" (Ma et al., 2021, p. 5). Another way to deal with the limited knowledge available on the uncertainty of the quantum threat, is to use Mosca's XZY. This is a simple numerical model which helps to indicate how urgent the risk is. A slightly adapted version is used in step 4 of the SRA.

Wangen et al. (2018) point out that typically, ISRA is tends to prioritise technical over organisational aspects. Only one of the eleven compared methods identifies business processes as assets to the organisation. In case of assessing the societal risks of the quantum threat, this is clearly a weakness. It is necessary to widen the scope to be socio-technical, rather than solely technical. Therefore, the SRA incorporates business processes as assets in step 2.

After the comparisons of Shamala et al. (2013) and Wangen et al. (2018), we can conclude that traditional information security RA builds from the bottom up. The risk identification phase is characterised by reasoning from parts (e.g., assets, threats, vulnerabilities, outcomes) to complete scenarios. This allows for a fine-grained approach that considers many separate but related components. Adopting this structure appeases the *highly granular* requirement. Starting at this broken down level, the assessor works towards the consequences on a higher, more general level. However, this higher level is limited to the scope of the organisation in traditional information security RA. When aiming to find societal risks, the scope of the organisation needs to be surpassed.

A tried-and-true method that operates on the national level of analysis and assesses risks to national security is the Integrated Risk Analysis National Security (GRNV) (Analistennetwerk Nationale Veiligheid, 2019). This method has no risk identification phase (or risk evaluation phase). It assumes threat scenarios and initiates the risk estimation phase, in which impacts and likelihood are estimated.

The SRA adds value by combining both approaches above. Leveraging lessons from information security RA practise, risk identification and risk estimation are performed. One information security RA method, SecRAM, is particularly useful for its easy-to-use organisational impact estimation (Le Fevre et al., 2017). The organisational impact estimation from SecRAM is used in step 3.2.1 of the SRA. The SRA borrows the organisational impact areas, the organisational impact score matrix, and the confidentiality-integrity-availability structure from SecRAM. Then, to surpass the organisational scope and provide a societal perspective, the GRNV is used in step 3.2.3 of the SRA.

The confidentiality-integrity-availability structure from SecRAM is also applied throughout the rest of the SRA. It uses the three well-established information security properties to categorise different types of information security breaches and the accompanying risks. The properties are defined as follows according to ISO/IEC 27000:2018 (ISO, 2018):

Confidentiality:	property that information is not made available or disclosed to unauthorized individuals, entities, or processes
Integrity:	property of accuracy and completeness
Availability:	property of being accessible and usable on demand by an authorized entity

By imagining that one of the properties cannot be guaranteed for a certain asset, threat scenarios are imagined. For example, what does it mean when medical dossiers can unintendedly be seen by third parties (i.e., a breach of confidentiality)? Or when digital financial transactions can be edited by skilled hackers (i.e., a breach of integrity)? Or when pressing charges digitally is not possible anymore (i.e., a breach of availability)? Adopting this structure in the SRA helps the assessor to think in a more granular way.

Table 2: Lessons from literature as input for the SRA

Method/framework	Lesson	Author(s)
Social Impact Analysis	Risk rather than impact is the right concept for discussing the expected potential impact of the quantum threat. Risk assessment rather than social impact analysis is field that is suited to the goals of this research.	Esteves et al., 2012
Risk analysis		Aven, 2018
Core subjects of risk analysis		Aven et al., 2018
Risk analysis	Risk is composed of impact and uncertainty. These are separately estimated in steps 3 and 4 and combined in step 5 of the SRA.	Aven, 2018
Societal Impact Analysis for security research	When studying societal impact, one should include natural and artefactual systems aside from social systems. This is taken up in my definition of societal risk.	Wadhwa et al., 2015
Information security RA comparisons	Reasoning bottom-up from threats, vulnerabilities, and assets to higher level threat scenarios gives room for a granular approach. The SRA employs this idea throughout its general structure.	Shamala et al., 2013
		Wangen et al., 2018
Information security RA completeness framework	Three phases of risk identification, risk estimation, and risk evaluation should be (and are) present in the SRA, making it more *usable* and *transferable*.	Wangen et al., 2018
Information security RA completeness framework	Business processes as assets are a useful way to broaden the RA scope to be less technical. This is applied in step 2 of the SRA.	Wangen et al., 2018
CARAF	Asset identification should happen based on the identified threats, rather than the other way around. This is present in the order of steps 1 and 2 in the SRA.	Ma et al., 2021
	In case of the quantum threat, it makes sense to break with conventional information security RA methods and estimate the uncertainty by using Mosca's XYZ. This is applied in step 4 of the SRA.	Ma et al., 2021
SecRAM	The organisational impact estimation of SecRAM is very straightforward. It is used in step 3.2.1 of the SRA.	Le Fevre et al., 2017
	The confidentiality-integrity-availability structure is useful to facilitate granular thinking. It is applied in steps 3, 4, and 5 of the SRA.	Le Fevre et al., 2017
GRNV	The GRNV provides a way to include a societal impact perspective. It is used in step 3.2.3 of the SRA.	Analistennetwerk Nationale Veiligheid, 2019

4 Result: the final SRA method

In this chapter the final version of the Societal Risk Assessment method (SRA) resulting from this research is presented. The previous versions are available upon request. Some of the elements are taken from existing risk assessment methodology, risk analysis and SIA theory. The section is organised as follows. First, the general purpose of the SRA is laid out. Then, an outline of the method and its components is presented. After that, each individual step is discussed in-depth. Lastly, the application context is explained.

4.1 Purpose and requirements

The main purpose of this method is to capture the most important risks to society that arise from the development of quantum computing in relation to PKI related to an organisation. In other words, the developed method allows an organisation to assess the societal impact that would occur when quantum computing breaks the public key cryptography used by the PKI of that organisation. While the method is focussed around an organisation, the impact under consideration is indeed impact on society at large.

Therefore, the method is engineered to have the assessing organisation adopt multiple viewpoints and a broad perspective on risk. This desire for a broad perspective has led to the choice to design for use by parties with differing roles in the PKI chain. This means that a PA, CA, intermediary CA, end-user, or any other party reliant on PKI usage can make use of the method.

In line with the goal, this method must be useful when a party involved with PKI wishes to assess their societal risks considering quantum computing. It helps the assessing organisation to achieve a clear understanding of their societal risks, prioritise risks to be treated first, and prepare for transition to quantum safe PKI.

Accordingly, we formulated the following requirements (R) for the SRA method:

- R1. Usability – experts across various PKI domains should be able to use this method.
- R2. Self-explanatory – domain experts should be able to use this method by themselves, without external guidance.
- R3. Relevance – the method must include organisational, national, and individual perspectives on risks.
- R4. High granularity – Method must facilitate the detailed and specific description of assets and associated risks.

Our literature review reveals that there are no risks assessment methods available that satisfy all of the requirements mentioned above. Previous risk assessment methods fall short when it comes to providing a holistic view of the risks of quantum computing to PKI. SecRAM provides a solid general approach to cyber security related risk covering many types of impact. Therefore, it provides a solid base for the SRA. However, it is insufficient in two areas. Firstly, it limits its view to consequences to the assessing organisation. This means it is not adequate for assessing the broad notion of societal risk. To broaden the scope and get closer to societal as opposed to organisational risk, the GRNV is introduced. This is a method to assess risks to Dutch national security. The GRNV is designed by the Analyst Network National Security (ANV), consisting of prominent knowledge institutes on national security. By making use of the perspective the GRNV provides, the SRA is better able to assess societal risk.

Moreover, SecRAM departs from what an organisation would like to protect (assets) and then derives vulnerabilities to these assets from existing technologies. This is useful when assessing general cyber risk, but not when there is a need to analyse specific future vulnerabilities. Quantum computing both creates very specific vulnerabilities to specific assets and is a future technology. Therefore, in this case, a different approach is required. This approach is taken by CARAF. It is specifically designed to deal with newly developing technological threats to encryption that create very specific vulnerabilities. Borrowing ideas from CARAF suits the SRA to the quantum threat.

Note that this SRA-methodology has been expanded and evolved during the design process. In particular, this final version of the SRA method has chosen a higher abstraction-level of the SRA-methodology. From its SWOT analysis it became clear that the specificity of the method would be of little practical use, due to its dependence on data that is likely to be absent. This problem has been addressed in the current version.

4.2 Component overview

To provide a clear overview of the SRA components, IDEF0 modelling is used. This is a suitable tool to model methodologies (Presley & Liles, 1998). The SRA broken down into six steps can be found in Figure 4. The activity and purpose of every step can be found in *Table 3*. Each of these steps is thoroughly elaborated upon in the next segment.

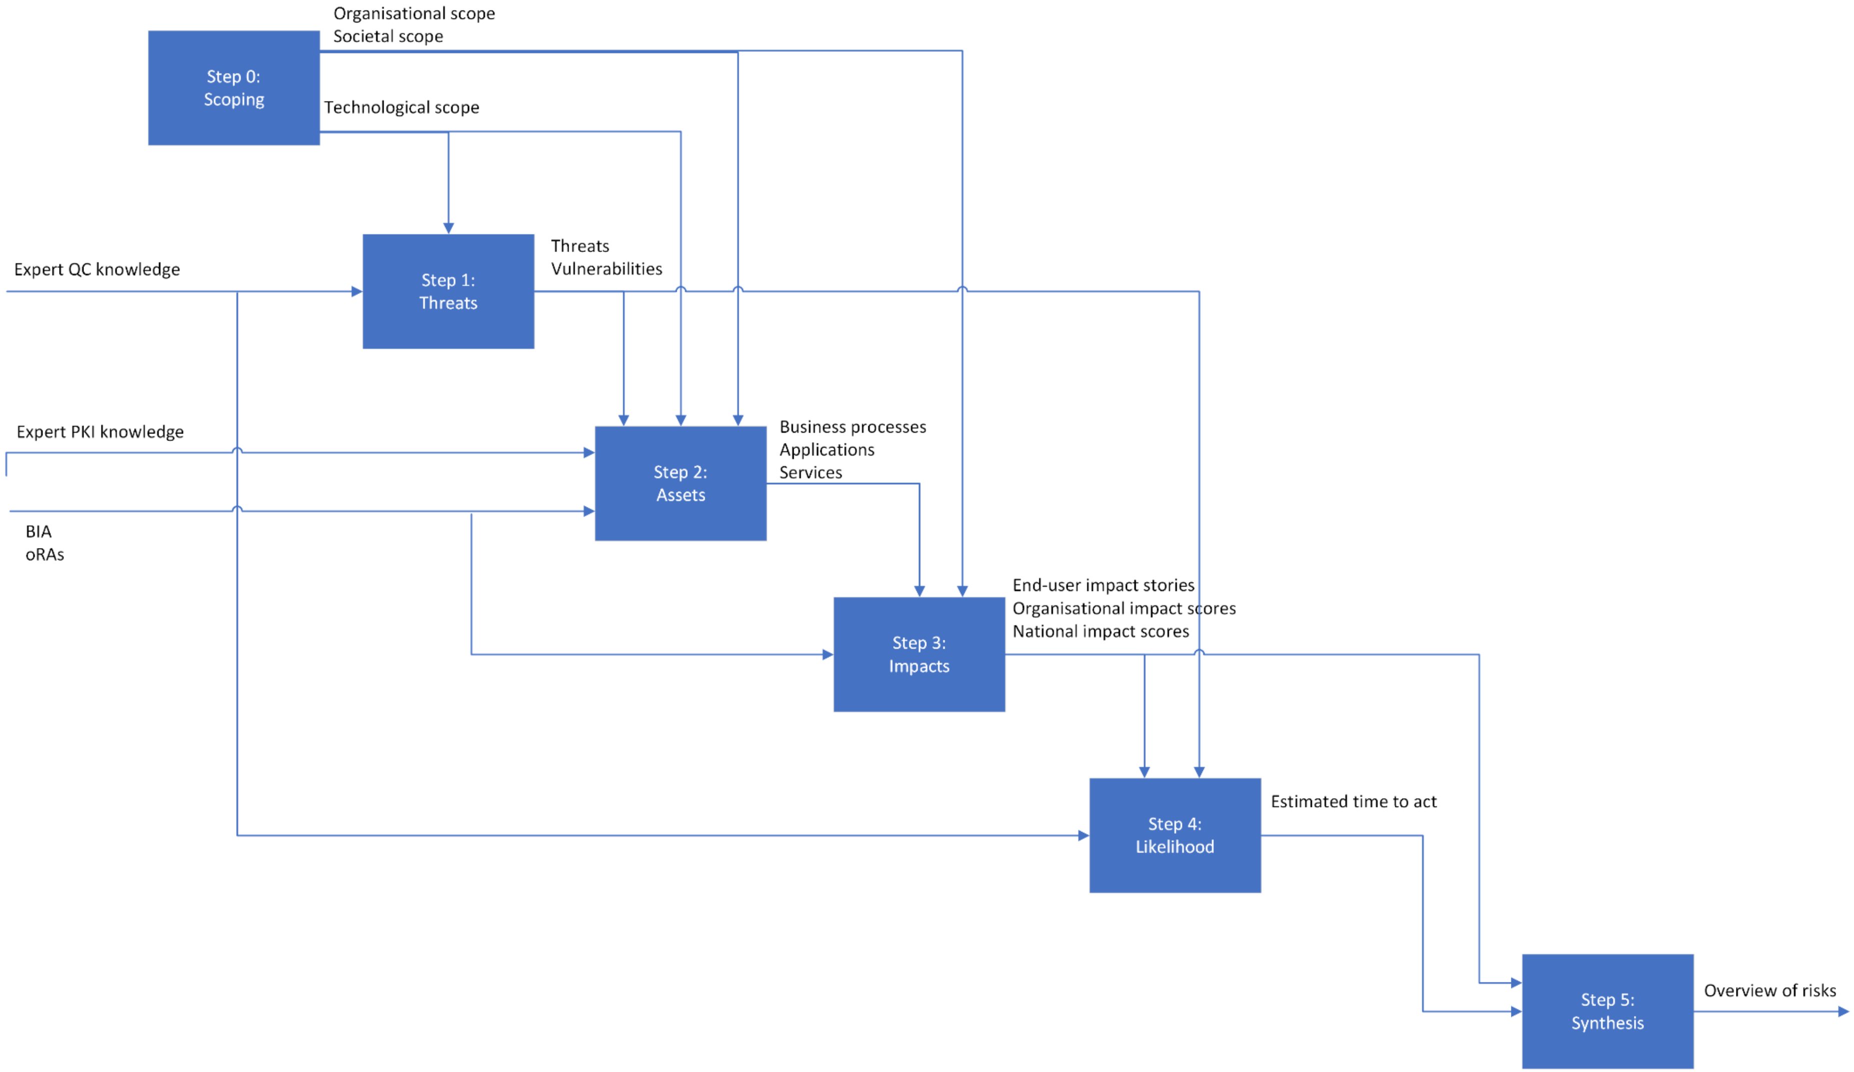

Figure 4: Overview of the SRA method using the IDEF0 modelling notation

Table 3: Overview of SRA steps

Step	Activity	Purpose
0: Determine scope	To determine the technological, organisational, and societal influence scope.	To ensure a focussed and highly relevant assessment of risks and avoid wandering.
1: Identify threats	To identify what threats are in scope and what vulnerabilities are exploited by these threats.	To get an idea of what there is to be defended against, so that we may find what to defend.
2: Identify assets	To list the relevant business processes, related PKI applications, and their dependent services.	To have an overview of what is to be defended, so that we may find what is at stake.
3: Assess impact	To assess the potential impacts of the threats on the assets from an organisational and a societal perspective.	To give an idea of what is at stake when threats materialise.
4: Assess urgency	To review available expert judgement of the urgency of the threat.	To provide an understanding of the timescale in which certain impact are to occur.
5: Synthesise	To combine the risk components from the previous steps and rank threat scenarios according to their need to be mitigated.	To generate an overview of the societal risks that the quantum threat brings from the point of view of a single organisation.

As can be seen in Figure 4, the output of step 0 is three different scopes to be used in the rest of the SRA: technological, organisational, and societal influence. The technological scope is used to delineate steps 1 and 2, and the organisational and societal influence scopes are used to delineate steps 2 and 3. In step 1, the technological scope is used to inform which threats and vulnerabilities are relevant in the SRA. The threats and vulnerabilities are taken from expert quantum computing knowledge. In step 2, the technological scope is used again to inform which business processes, applications, and services are relevant. These are the three different types of assets used in the SRA. The organisational scope is applied to find the relevant business processes and depending on the assessing organisation also the relevant applications. The applications may also fall within the societal influence scope, as do the services provided. Then, in step 3, the three types of assets are used to find how their compromise could impact the assessing organisation and society. Both the organisational and societal influence scopes are used in this regard. Both the assets and the impacts are assessed by processing a business impact analysis, other previously completed risk assessments, and expert PKI knowledge. Step 4 takes each threat and estimates the corresponding time to act. This estimation is based on expert knowledge of quantum computing and PKI. Now that the organisational and societal impacts and the estimated time per threat is known, they can be combined in step 5 to create an overview of the risks.

4.3 Step 0: Determine the scope

Scoping is important in risk management. A scope is seen as "an area which requires a specific type of attention." Good scoping ensures that people can focus their attention on that which matters and avoids discussions that stray from the point (Joosten & Smulders, 2014). Hence, at the start of the SRA, the scope is determined.

An important limitation to the scope of the SRA is technological. In this research, it is only sensible to limit the scope to quantum computing and its threat to the asymmetric cryptographic basis of PKI. This means that only those things related to quantum computing and PKI should be considered in the SRA. The technological scope could be further specified, depending on the assessing organisation. For example, symmetric cryptography or a distinction between data at-rest, in-transit, or in-use could be considered. Hence, in this step, the assessor is asked to specify the technological scope.

Moreover, it is important that there is a clear distinction between the scope of the organisation and the scope of the influence the organisation has on society. The distinction is illustrated in the following figure.

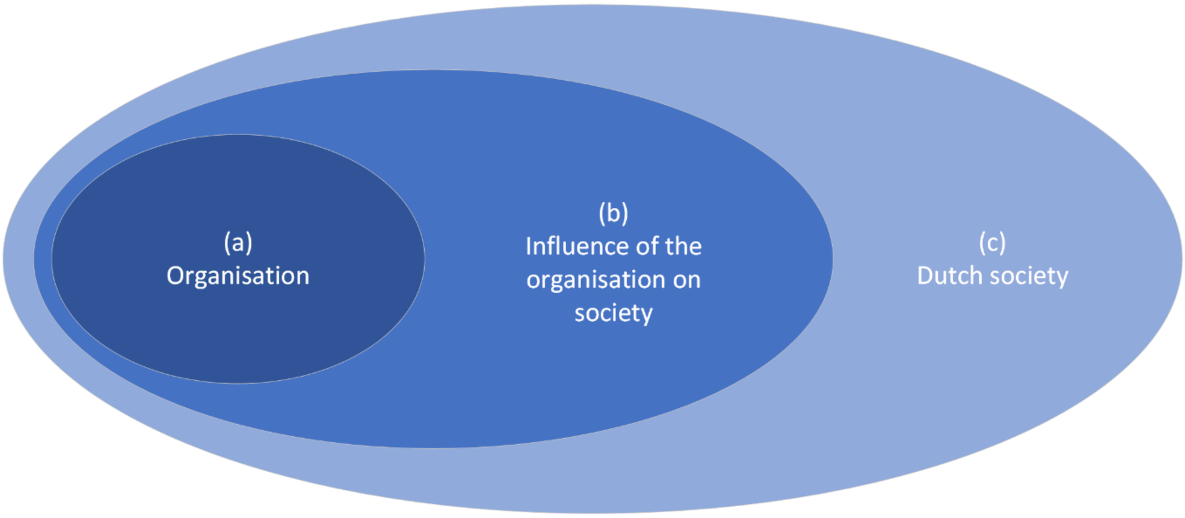

Figure 5: Scoping organisational societal influence

The world of PKI is socio-technical and complex, involving many actors in varying roles that depend on one another to enable services to society. Limiting the scope of analysis to a sole organisation (a) does not do right by this complexity and interdependency. The societal impact cannot be investigated if the assessing organisation does not look further than its own boundaries. On the other hand, broadening the scope to encompass society in its entirety (c) takes away from the specificity to the assessing organisation and thus the usefulness of the SRA to the organisation. As one can imagine, an analysis of how to reduce energy consumption in the Netherlands in general does not provide many actionable focus points to isolating one's house. Therefore, during the SRA, the assessor should walk the line between the scope of the organisation and its influence on society (b).

The steps of the SRA are designed to accommodate this. In step 0, the assessor is asked to define an organisational scope and a societal influence scope. As an example, the organisational scope could be limited to a specific branch of the organisation that facilitates or makes use of PKI. The societal influence scope could be geographically limited to the Netherlands. Both scopes are used in the following steps of the SRA. Any scoping decisions can be filled in a table like Table 4.

While the scopes considered may seem to be relatively set in stone, there is room for the assessor to further specify wherever deemed necessary. This helps to make the SRA more relevant and thus useful to the assessing organisation. Additionally, this step serves as an exercise to create a better shared understanding among the members of the assessing team of what exactly is to be assessed. This benefits the focus of the assessment and thus the quality of the output.

Table 4: Scoping

	Scope
Technological	The technological scope is limited to - risks that emerge from quantum computing (quantum computing) - risks related to PKI - …
Organisational	…
Societal influence	…
Explanation:	

4.4 Step 1: Identify threats

In step 1, the following question is central: “What to defend against?”. This goes against common practice. It is common practice to first identify what is of value and then see what could potentially endanger that. This order of work is present in ISO3100x and risk assessments based on those standards (ISO, 2018). It makes perfect sense when attempting to find what risks in general an organisation is exposed to. However, as CARAF points out, when investigating a specific technology that poses a threat, it makes more sense to depart from those threats (Ma et al., 2021). The reason is that you may otherwise include many elements in your analysis that are not vulnerable to the specific technology you want to investigate. That is why in the SRA, assessing the threats and vulnerabilities comes first.

The most obvious threat is quantum computing being used to break the security of data real-time, enabling an attacker to spy upon, modify, and interrupt data in-transit, in-use, and at-rest. This threat can be realised once a large enough quantum computer is available. Another threat is that of a store now, decrypt later attack. Such an attack is performed by capturing encrypted data and decrypting it once a powerful enough quantum computer is available. This attack is only relevant to data that needs to remain confidential for longer than the time it takes until a powerful enough quantum computer is available. For example, encrypted medical records that need to remain confidential for 20 years are captured and can be decrypted in 15 years. Such an attack is less powerful than the type described first, but the risk of a store now, decrypt later attack is more immediate. The two threats described are considered the main general threats, but there may be more relevant threats to specific organisations.

The threats are related to vulnerabilities. Vulnerabilities are weaknesses of systems in place that can be exploited to materialise threats. In case of the quantum threat, the reliance of cryptographic security on the factorisation or discrete log problem is a vulnerability. In practice, this means that any reliance on RSA, DH, or ECC is a vulnerability. Another vulnerability could be the use of non-doubled symmetric key lengths (because of Grover's algorithm), but only if symmetric cryptography is within the technological scope of the SRA.

During this step, the assessor establishes which threats and vulnerabilities are present and links each threat to a vulnerability. The results can be filled in a table like Table 5. The identified vulnerabilities can act as input for step 2.

Table 5: Threats and vulnerabilities

Threat	Vulnerability
quantum computing breaking asymmetric cryptography real-time	The use of any cryptographic solution dependent on the discrete log or factorisation problems (e.g., RSA, ECC, DH), which can be broken by quantum computing (i.e., virtually all asymmetric cryptography in place)
'Store now, decrypt later' tactic capturing sensitive data in-transit protected by PKI	
…	…
Explanation:	

4.5 Step 2: Identify assets

Now that we know what is to be defended against from step 1, the next question is "What to defend?". Finding the answer to this question is necessary to reveal what is at stake. It will be essential in determining the potential impacts of the threats. According to SecRAM, an asset is an "[element] in the system that [has] value for the achievement of business objectives or [an] element that [supports] the existence of the business objectives".

The only assets that are relevant are the ones vulnerable to the threats in scope. Hence, the identified vulnerabilities from step 1 are used to inform the asset identification in step 2.

The SRA defines three different types of assets: business processes, applications, and services. The first, business processes, are "a set of logically related tasks performed to achieve a defined business outcome" (Davenport & Short, 1990). These lie within the organisational scope defined in step 0. Keep in mind that any business process, or other asset, included in the SRA should be within the technological scope. This means that every business process should at least have some dependency on or relation to PKI.

After the relevant business processes, the applications are identified. Applications are applications of PKI technology that are reliant on one of the business processes. Applications can be within the organisational scope or only in the societal influence scope. This depends on whether running the full application is a responsibility of the organisation, or the organisation is only responsible for facilitating a part of the application. For example, a CA giving out certificates that are supposed to be used by businesses to automatically authenticate their tax statements is not responsible for the functioning of the application verifying the messages. However, they are responsible for the validity of the certificates used by this application. In this case, the application would be outside of the organisational scope of the CA, but in their societal influence scope.

Lastly, the services dependent on the applications need to be determined. A service is "The execution of information processes provided by an organisation" (Bharosa et al., 2015). Applications in and of themselves are not of value to society. The services they facilitate create societal value and thus the services are a key factor in determining the societal impact if the threats materialise.

All the business processes, applications, and services can be organised in a table such as *Table 6*. Finding any of the assets can be assisted by making use of a pre-existing Business Impact Analysis. Pre-existing risk assessments might be of use as well.

Table 6: Assets

Business process	**Application**	**Service**
Organisational	Org. / Soc. influence	Societal influence
Explanation:		

4.6 Step 3: Assess impacts

It is important to note the difference in the impacts of a breach and the impacts of a mitigating measure. In this step, the assessor should stick to the impacts of a breach.

Following the definition of risk, impact is an integral part. Therefore, the relevant potential impacts and their gravity need to be assessed.

All impacts are systematically assessed by examining relevant assets.

4.7 Step 3.1: Assess organisational impacts

When assessing the impact from an organisational perspective, it is important to keep the organisational scope in mind. This means that the business processes defined in step 2 are the assets to be considered, as they fall within the organisational scope. In assessing the organisational impact, SecRAM provides guidance. As in SecRAM, every asset is judged on a compromised information security property. SecRAM prescribes seven different impact areas to be considered when assessing the impact. These are personnel, capacity, performance, economic, branding, regulatory, and environment. Every combination of business process, CIA property, and impact area is given an impact score from A to E. These scores can be filled in

Table 7. To determine which score should be given, a score card such as *Table 14* can be used. The highest score of each row should be copied to the 'total' column.

Table 7: Organisational impact

Business process	Personnel	Capacity	Performance	Economic	Branding	Regulatory	Environment	Total
Explanation:								

4.8 Step 3.2: Assess societal impacts

After the organisational impacts are assessed, the assessor moves on to the societal impacts. For this part, the societal influence scope is leading instead of the organisational scope (Figure 5). Assessing the organisational impact is done for every business process. Every business process is associated with at least one, but perhaps multiple services. To assess the potential societal impacts, the services are the assets to judge. The societal impact of each business process is decided by the societal impact of its related services.

4.8.1 Step 3.2.1: Assess end-user impacts

Assessing the societal impacts per business process is done in two steps. First, the end-user perspective is adopted. The assessor estimates what impact the end-user of a service of an application would feel in case of a compromise. This produces a few sentences of written text for every application, as can be filled in *Table 8*. Taking this step is necessary to force the assessor to think about the impact that a compromise within the scope of their organisation may have on individuals outside of that scope. It is easier to think of how an individual may be impacted than to come up with ideas how the whole of society may be impacted. This intermediary step serves as a bridge to help think of impacts for the whole of society.

Table 8: End-user impact

Business process	
Application	**Consequences for end-users in case of compromise**

4.8.2 Step 3.2.2: Assess national impacts

The next step in assessing the societal impact is taking on a national perspective and giving corresponding impact scores. To do so, the GRNV[4] is consulted. It is specifically designed to assess disrupting impact to society on a national level from a security point of view. This method is used to inform the national security strategy of the Dutch government (*TNO*, n.d.).

Similar to the way the organisational impact is scored, the societal impact is scored. The impact areas in *Table 7* taken from SecRAM are replaced by impact areas from GRNV. Some impact areas from GRNV are deemed more relevant than others. Although not all impact areas will be applicable, all should be considered so that no kind of impact will go unnoticed. The full list of impact areas with the more relevant ones in bold can be found in *Table 9*.

Table 9: Societal impact areas

National security concern	**Impact area**
1. Territorial security	1.1 Violation of the integrity of (Dutch) soil
	1.2 Violation of the integrity of the international position of the Netherlands
	1.3 Violation of the integrity of cyberspace
	1.4 Violation of the integrity of allied soil
2. Physical security	2.1 Deaths
	2.2 Seriously injured and chronically sick
	2.3 Lack of basic needs
3. Economic security	3.1 Costs
	3.2 Degradation of the vitality of the Dutch economy
4. Ecological security	4.1 Long-lasting damage to nature and the environment
5. Social and political stability	5.1 Disruption of day-to-day life
	5.2 Degradation of the democratic rule of law
	5.3 Societal unrest
6. International rule of law	6.1 Degradation of the norms of state sovereignty, peaceful co-existence, and peaceful conflict resolution
	6.2 Degradation of the working, legitimacy or compliance with international treaties and norms concerning human rights
	6.3 Degradation of a rule-based international financial-economic order
	6.4 Degradation of the effectiveness and legitimacy of multilateral institutions

[4] Geïntegreerde Risicoanalyse Nationale Veiligheid: https://www.nctv.nl/documenten/publicaties/2019/6/07/geintegreerde-risicoanalyse-nationale-veiligheid

To determine the scores for each impact area, the GRNV is applied. For every impact area, the GRNV describes how to determine the (societal) impact score. The threat scenario that is considered is that a certain type of compromise occurs affecting all services of all applications dependent on a certain business process. This is repeated for every business process. The scores are taken up in *Table 10*. To assess the impacts, deep knowledge about the business processes is necessary. Additionally, organisations can make use of business impact analyses and previously conducted risk assessments.

Table 10: Societal impact

	1.2	**1.3**	**3.1**	**3.2**	**5.1**	**5.2**	**…**	
Business process	**International position**	**Cyberspace**	**Costs**	**Economy**	**Day-to-day life**	**Rule of law**	**…**	**Total**
Explanation:								

4.9 Step 4: Assess urgency

As defined earlier, societal risk has two parts. The previous steps have dealt with the value at stake. This step deals with the uncertainty. The quantum threat is a future threat of which there has not been a materialised instance.

However, following Mosca's XYZ model summarized in figure 3 (Mosca & Mulholland, 2017), an estimation of urgency can be deduced. The variables of Mosca's model are defined as follows:

X = longest shelf life of data within the scope of the business process

Y = time it will take to replace all vulnerable cryptography used within the scope of the business process to quantum safe cryptography

Z = estimated time until a powerful enough quantum computer is available to execute a threat

Urgency serves as an abstraction of these three variables. Urgency should be read as 'how quickly will a business process be vulnerable to a threat'. For example, a process depended on long-term confidentiality of data needs to be addressed with higher urgency compared to a process depended on ephemeral data. Similarly, a process that cannot easily replace its vulnerable cryptography has a higher urgency, compared to a crypto-agile process.

The estimated urgency can be extreme, high, middle or low, see *Table 11.*

Table 11: Likelihood assessment

Business Process	Threat	Urgency
		Extreme/High/middle/low
		Extreme/High/middle/low
Explanation:		

4.10 Step 5: Synthesis

In this step, the insights from the previous steps are combined. This results in a priority list of business processes to secure first. The high scoring business processes can be traced back in previous steps, to see where the high scores originate from (table 13). This gives insight in the internal dependencies of the SRA.

The results from all previous steps are combined in a single table. The organisational and national impact are taken from Steps 3.1 and 3.2.2 respectively. The theoretical time to act (t) is taken from Step 4. Lastly, using the conversion matrix, the risks can be ranked (see table 14). Taking the highest of the two impact scores and the theoretical time to act, the matrix ranks the risk for each threat to each business process.

Table 12: Compounded impact for business processes

Business process	**Threat**	**Organisational impact**	**National impact**	**Urgency**
A…..	Real-time access			
	Store now, decrypt later			
B….	Real-time access			
	Store now, decrypt later			

It should be noted that the conversion matrix can be adjusted according to the risk appetite of the assessing organisation. This is illustrated in the following table

Table 13: Urgency and impact

urgency ⇒ ⇓ impact	low	middle	high	extreme
A	Acceptable	Acceptable	Moderate	Moderate
B	Acceptable	Acceptable	Moderate	Severe
C	Acceptable	Moderate	Severe	Severe
D	Moderate	Severe	Severe	Critical
E	Severe	Severe	Critical	Critical

Table 14: Organisational impact score card

Impact Areas	A No impact	B Minor	C Severe	D Critical	E Catastrophic
Personnel	No injuries	Minor injuries	Severe injuries	Multiple severe injuries	Fatalities
Capacity	No capacity loss	Loss of up to 10% capacity	Loss of 10% - 30% capacity	Loss of 30% - 60% capacity	Loss of 60% - 100% capacity
Performance	No quality abuse	Minor system quality abuse	Severe quality abuse that makes systems partially inoperable	Major quality abuse that makes major system inoperable	Major quality abuse that makes multiple major systems inoperable
Economic	No impact	Minor loss of income	Large loss of income	Serious loss of income	Bankruptcy or loss of all income
Branding	No impact	Minor complaints	Complaints and local attention	National attention	Government & international attention
Regulatory	No impact	Minor regulatory infraction	Multiple minor regulatory infractions	Major regulatory infraction	Multiple major regulatory infractions
Environment	Insignificant	Short term impact on environment	Severe pollution with noticeable impact on environment	Severe pollution with long term impact on environment	Widespread or catastrophic impact on environment

4.11 SRA output

An important part of the SRA is how the output is handled afterwards. The SRA is not complete without a report detailing the particularities encountered during the assessment and what they mean for the results and any potential next steps in mitigation. This report should be like a management summary for the SRA complete with conclusions and recommendations.

4.12 Application context

The SRA is a method, a tool that needs to be applied. Without proper application, it cannot be effective. Imagine a hammer gripped by the head instead of the handle, thrusted instead of swung, or held by a person not strong enough. All of these ways of handling the tool render its use not as effective as it can be. This is why the application context matters. Just like the hammer, the SRA benefits from proper application. This starts with who is intended to use to SRA.

4.12.1 Who should apply the SRA?

As mentioned in 4.1, the SRA is intended to be used by a wide variety of actors. But, the SRA is not meant to be used by just any- and everybody. The organisation from whose viewpoint the SRA is applied is called the assessing organisation. The assessing organisation should in some way be reliant on PKI. This may be through activities in trust service provision in varying roles such as policy authority, registration authority, root CA, or intermediary CA. It could also be through the use of leaf certificates to gain access to trust services or being reliant on the use of leaf certificates by other parties. The assessing organisation should also be interested in knowing their societal influence and/or doing business in a socially responsible manner. These parties can use the knowledge gained from the SRA to take action or changing their policies on becoming quantum safe.

With this user group in mind, the SRA was designed. This resulted in the choice to limit the scope to the viewpoint of the assessing organisation and their influence on society, rather than the composite influence on society of multiple organisations. This choice was made because of three reasons. First, it provides an incentive for organisation to perform the SRA. It allows organisations to get insights relevant to them. They get to see their risks and their influence on society. Second, the assessment will be less complicated as no multiple viewpoints clash and clutter the analysis. Third, it is easier to organise a session within an organisation than with people from multiple organisations.

However, there is a drawback to this scoping choice. There is less detailed oversight on the actual societal risks. The chained nature of hierarchical trust present in PKI creates dependencies outside of the view of the assessing organisation. For example, Logius as the policy authority of PKIoverheid is not aware of all exact use cases of all certificates under their umbrella infrastructure. This information resides with parties lower in the trust chain that manage applications dependent on PKIoverheid certificates.

To overcome this drawback, research projects (e.g., HAPKIDO) or other public-private collaborations can combine insights of SRAs done by multiple organisations. This way, knowledge from the entire trust chain, top to bottom, and in multiple sectors can be leveraged and the societal risks can be accurately and widely assessed. It is the task of such projects and collaborations to decide how much detail needs to be known in order to paint an adequate picture of the societal risk.

Apart from which organisation is to perform the SRA, it also matters who inside the organisation is present during the actual performance of the SRA. As this is a more practical question, it is discussed in the next section.

4.12.2 How to apply the SRA?

As an organisation, there are a multitude of ways in which tools such as the SRA can be applied. In the case of the SRA, we first take a look at what kind of people should be involved during the SRA. This group of people is called the assessor. Ideally, there is a mix of expertise present. Risk assessment in general thrives on diversity of expertise. In case of the SRA, it is advised to have at least one expert present on the following knowledge areas: risk management, (senior) management, technology (i.e., PKI), potentially affected business processes, and compliance. Additionally, while in-depth knowledge of quantum computing is not necessary, it is strongly recommended to have some shared knowledge on how quantum computing affects PKI. In-house experts are necessary, as (tacit) knowledge of the assessing organisation is necessary. Yet, not all experts need to be in-house. External experts can be a great addition. It should be noted that the composition of the assessor is based on expert consultation and is not rigorously researched.

Another aspect of the way in which the SRA is applied, is the depth of analysis. Because of the high granularity requirement, the SRA facilitates a deep analysis. However, per case it is applied to, the required depth of analysis may vary. It is the task of the assessor to attain the right depth of analysis for the case. This is done in multiple steps of the SRA. Steps 0, 1, particularly 2, but also 3 are where the depth is decided. In step 0, the scoping can be described in a very broad and general way, or it can be made very specific. The same goes for the threats and vulnerabilities in step 1. In step 2, the assets can be categorised in a fine-grained way, or in brush strokes. For example, a business process can be described as broad as “safeguarding digital trust in our PKI” or as narrow as “assisting our customers of service X with certificate renewal”. Similarly, an application could be Digipoort as a whole, or a specific part of Digipoort[5]. As for the dependent services, one could be described as tax reporting in general or reporting a specific kind of tax by a specific kind of party. Taking a deep and granular approach may benefit the output of the SRA, but it takes longer to do the assessment. Of course, at a certain point, it may take much more time to be more specific, while it is marginally beneficial to the output. Hence, the trade-off between granularity/depth and time to complete the assessment.

So, the time it takes to perform the SRA is dependent on the depth of analysis. Ideally, the SRA should be completed in a single session, so that the assessment remains focussed. Comparable RAs take a single session of between four to six hours to be completed. Based on the speed of the user workshops and the evaluation session, this seems realistic. However, this has not been properly tested. The time necessary is also dependent on the efficiency of the session. This is reliant on proper guidance of the assessor.

Proper guidance of the assessor is important for the efficiency of the session and the quality of the SRA output. For example, improper guidance can lead to vague scoping. Here, participants will keep adding information to the discussion which may not be as relevant to the assessment as they think. It could lead to participants talking about differing concepts without realising that they are. This muddies the analysis and takes up time and mental space. This problem in common in risk management practise. Good guidance can curb such and other behaviours distracting from a focussed assessment. There are several ways in which guidance can be offered. One of which is the provision of a manual. Another option is a tutorial video or workshop to be followed beforehand. This is a good way to more actively show the purpose and workings of the SRA. To add to this, an exemplary case can provide a means to make the abstractions of the SRA concrete. These three options can be provided in conjunction to maximise their effectiveness. Lastly, a session leader

[5] This example is drawn from HAPKIDO and is further elaborated in Deliverable 1.2.

should be appointed to guide the assessment. This leader should be well-aware of the guidance materials and preferable have pre-existing knowledge of risk management. The guidance options are directly related to the transferable requirement.

A more practical concern is tooling. For this research, the use of Word documents and online collaborative word processing documents was sufficient. However, this was at times awkward because of the need for flexible table editing and tables being cut off due to page sizes. A better workflow would be enabled if a tool specific for the purpose of applying the SRA was made. This could be a web-tool which runs client-side to avoid issues with sensitive data, as it would be easily accessible to users. On the other hand, because very sensitive information about business vulnerabilities is being handled, users may prefer to stay away from a browser communicating with the open internet altogether. In this case, a portable open source program would be the better choice.

5 Evaluation of the SRA

The SRA is evaluated iteratively using SWOT analyses. After every workshop, feedback was structured in SWOT analyses by listening back to the recordings of the workshops. These SWOT analyses are documented in Appendix B. In the output evaluation workshop, the participants were asked to collaboratively fill in a SWOT analysis themselves. The results of all SWOT analyses have been combined in

. In the following section, the SRA is evaluated per requirement.

Table 15: Combined SWOT analysis of the SRA

Strengths	**Weaknesses**
This method addresses the recognised need for insight into societal risks because of the quantum threat to PKI. The method feels familiar to those with some experience in risk assessment. The multiple perspectives on impact assessment. The step-by-step nature enables 'peeling off' the different layers and provides granularity and insight in dependencies.	The method does not distinguish between threat actors. The method does not include mitigation risks. The method requires a lot of knowledge on business processes. Information with the required level of detail can be hard to find. Risk assessment is a tough process for business owners. Difficult to pull mitigation strategies from the method.
Opportunities	**Threats**
Documents such as a Business Impact Analyses and previously done risk assessments can be valuable input. Present the SRA as a self-assessment and combine the results with recommendations/action pathways.	The urgency is very hard to properly assess. Organisational impact: capacity & performance may be differently interpreted by different assessing parties, depending on their industry. Parties are not likely to share sensitive information about their perceived vulnerabilities such as a Business Impact Analysis and previously done risk assessments. The method can be perceived as lengthy and complex. There is lack of clarity about the need for extensive insights.

5.1 Evaluation per requirement

Each strength, weakness, opportunity, and threat is related to one or more requirements. These are discussed below per requirement, starting with the *usable* requirement.

R1. **Usable** – Experts across domains should be able to use this method.

To be clear, this requirement is about how well users across domains are able to apply the method, given that they have no problems understanding how the method is supposed to be applied. The following weaknesses and threats identified by the participants of the user workshops indicate that the SRA is not an assessment that can be quickly executed. It is clear that a lot of knowledge and thought is required. This emphasises the need for a wide variety of experts during the assessment. Although the SRA may be perceived as lengthy, this is consider to be a necessary evil. It is simply quite a complex topic, which means it requires some time to figure out. There is an opportunity that may be leveraged to improve the usability of the SRA, by providing access to clear and already processed information ready to be used in the SRA. Documents that are often in existence by good business practice can offer valuable input. These documents are business impact analyses and previously done risk assessments as part of business continuity planning.

W The method requires a lot of knowledge on business processes.

W Risk assessment is a tough process for business owners.

W Information with the required level of detail can be hard to find.

T The urgency is very hard to properly assess

T The method can be perceived as lengthy and complex.

O Documents such as a Business Impact Analyses and previously done risk assessments can be valuable input

The following threat and opportunity are linked. The problem that parties are not willing to share sensitive information may be alleviated by having the SRA be a self-assessment in which the sensitive information can be kept to the assessing organisation. By combining the results of the SRA with recommendations/action pathways, the assessing organisation can effectively use the results without ever needing external help. This is something to be further researched by combining work package 1 with other work packages from project HAPKIDO.

T Parties are not likely to share sensitive information about their perceived vulnerabilities such as a Business Impact Analysis and previously done risk assessments

O Present the SRA as a self-assessment and combine the results with recommendations/action pathways

In conclusion, the combined SWOT points towards some problems with usability in the sense that it may be hard for some users and requires effort. This may be helped somewhat in future versions of the SRA. On the other hand, it will always require effort to do the SRA, as the topic itself remains complex. Therefore, it is recommended to create a trustable atmosphere whilst executing the SRA in which all participating parties are willing to share sensitive information and cooperatively assess the quantum threat their organisations. In terms of usability across domains, no weaknesses or threats were identified. This seems to be a good thing.

R2. **Transferable** – Domain risk experts should be able to understand and apply this method

This requirement is about how well users understand how to apply the SRA. As it turns out, the use of existing information security RA structures and building blocks has paid off. Because of the familiarity, the method is more easily understood. However, it can still seem complex and lacking clarity about why the granularity is necessary. These threats can be taken on by applying the advice on guidance presented in section 4.4.

S The method feels familiar to those with some experience in risk assessment.

T The method can be perceived as lengthy and complex.

T There is lack of clarity about the need for extensive insights.

One identified weakness was about the potential for some concepts to be understood differently by differing assessing organisations. It should be noted that this may not be an issue per se. In the end, risk assessment is subjective. The meaning of the scores and filled in answers is what the assessor assigns to them. The results of the SRA are still valid, as they are relevant to the assessing organisation according to the assessor.

W Organisational impact: capacity & performance may be differently interpreted by different assessing parties, depending on their industry.

Summarising, the transferability seems to be not perfect, but not too bad either. With some additional work on guidance, the transferability can be improved. It is recommended for the assessors to provide clear terminology to improve transferability.

R3. **Relevant** – Method must include organisational, national, and individual perspectives

Both strengths identified below can be linked to the relevancy of the SRA. Many of the participants indicated their interest in the SRA and that the goals of the method are relevant. Moreover, the multi-perspective approach is seen as a plus.

S This method addresses the recognised need for insight into societal risks because of the quantum threat to PKI.

S The multiple perspectives on impact assessment.

However, the limited scope of the SRA focussing on the risks of quantum un-safe PKI in the quantum era was less well-received. Some participants would rather have a method that includes the risks involved with mitigation of the quantum threat. Moreover, some participants would have preferred to couple the SRA to potential mitigation strategies. Both these points, whilst very relevant indeed, are outside of the scope of work package 1 of project HAPKIDO, for which this SRA was developed. This should be taken as a sign that participants are in fact interested in seeing the risks that come with handling the quantum threat as well and thus it should be seen as a stimulus for the rest of project HAPKIDO.

W The method does not include mitigation risks.

W Difficult to pull mitigation strategies from the method.

In short, the SRA is relevant. Where it falls short is because of the explicit scoping choice to fit in project HAPKIDO’s work package 1. This should be taken up in the other work packages of HAPKIDO.

R4. **Highly granular** – Method must facilitate granular description of assets and associated risks

The strength identified signals that the SRA is deemed to be granular and at the right level of detail to be useful. One slight comment on that is that there is no distinction between threat actors. This allows for less granularity of the threat scenarios. Yet, this was a deliberate choice, as most participants did not view the distinction to add much value to the assessment, but it does complicate the assessment.

S The step-by-step nature enables 'peeling off' the different layers and provides granularity and insight in dependencies.

W The method does not distinguish between threat actors

One weakness that might take away from the granularity of the assessment is that highly granular information can be hard to find. Unfortunately, this is something that cannot be dealt with from the SRA development side. The information resides with the users and after all, the output of any method can only be as good as its input.

W Information with the required level of detail can be hard to find

To conclude, the SRA is a granular method. This benefits the results. Not always can the SRA be implemented as granular as the method allows, as highly detailed information may be hard or even impossible to find.

5.2 Concluding remarks

Overall, the combined SWOT reflects two main beliefs about the SRA among experts. Firstly, the SRA serves a relevant purpose, is of value, and achieves its goal. The second belief is that the SRA is somewhat complex, requires deep knowledge of business processes which might be hard to find, and asks for experience with risk assessment. To help ease the problems that come with the second belief, guidance should be offered and business impact analyses and previously completed risk assessments can be used as input. However, these input documents may be hard to gain access to as well, as they are often confidential.

Finally, note that during the finalisation of the SRA, the method has been brought to a higher abstraction level. That is, some complexity has been reduced to improve usability. The assessment of urgency in particular has been significantly reduced in complexity. Whilst this reduces the granularity of the assessment, it substantially improves the SRA as a whole.

6 Conclusions and recommendations

6.1 Conclusions

The goal of WP 1 is to design, develop and test a method for assessing the societal risks of quantum computing, in particular focussing on domains that use PKI systems. After six workshops and design iterations, the SRA presented in section is the end-result. Here, each component, inspired by ISRA and adapted with a societal perspective, is described.

This answers the primary question of WP 1: What is a suitable method to assess the potential societal impact of the quantum threat to PKI systems?

The resulting SRA is found to be of value by experts and prospective users. The prospective users suggest that the method serves a relevant purpose, and achieves its goal of generating insight into the societal impact of the quantum threat. However, applying the SRA properly is a complex process that requires knowledge about the method and business processes. Moreover, it may prove difficult to obtain the detailed information to reach the desired level of granularity. While the output of any method is only as good as its input, it remains useful to take note of the effort required to achieve a detailed overview of risks and their origins. After validating the case study of PKI government and its PA, workshop participants indicated that assistance in applying the method is useful and necessary.

Information security risks within an organisational context can impact society. This report has provided more insight into how these risks can be structurally assessed, by providing a method to do so. Additionally, the research serves as a stepping-stone to assess the societal risks of the quantum threat to PKI.

6.2 Recommendations

Within the context of HAPKIDO, this method can be applied to help answer another research question of WP1: What are the societal risks of a quantum-unsafe PKI in the quantum era? This question will be answered in Deliverable 1.2. The SRA can be applied multiple times in cooperation with a representative selection of stakeholders. Then, the results can be combined to gain practical insight into the societal risks of the quantum threat to PKI in the Netherlands. To further hone the SRA for this purpose, researchers could investigate whether the organisational impact matrix and the national impact scoring should be adapted specifically to the case context. Additionally, this research left PKIs not managed by QTSPs out of scope. For the purpose of WP1, it might be valuable to consider PKIs that are managed by the organisations that use them, such as the Ministry of Justice and Safety. Another important factor to consider is how trust loss in multiple domains can cause cascading or multiplicative impacts.

Concerning the other work packages of HAPKIDO, it should be noted that the SRA does not consider all types of risk relevant to HAPKIDO. The SRA is designed to assess the risks in case nothing is done to thwart the quantum threat. Essentially, it assesses the risks of failing PKI. However, there are two other types of risk very relevant to HAPKIDO. The first is mitigation risk. This is the risk resulting from taking mitigating measures, such as downscaling PKI usage and using other means to the same end. The second risk type is a form of mitigation risk: migration risk. This is the risk that arise from migrating to a quantum safe PKI. These include not knowing where all leaf certificates are, implementation mistakes, and hardware restrictions at the end-user level.

As a recommendation for speeding up the migration process on an organisational level, the CARAF can be used.

Aside from project HAPKIDO and the quantum threat to PKI, further research could be done to see if the SRA is appropriate in general information security contexts. Perhaps with some adaptation, it can provide a way to assess the societal risks of new technologies that present a threat, just as quantum computing does.

Another avenue to explore is the fitness of the SRA for mitigation and migration risk mentioned above. This was out of scope for this research, but should be investigated, as it helps to further make the outcomes actionable. By knowing the risk of inaction, the assessor knows where to prioritise asset-wise when addressing the quantum threat. By knowing the mitigation and migration risk, the assessor can draw on those to form action/migration plans.

7 References

Algemene Inlichtingen- en Veiligheidsdienst. (2021). Bereid je voor op de dreiging van quantum computers. https://www.aivd.nl/binaries/aivd_nl/documenten/publicaties/2021/09/23/bereid-je-voor-op-de-dreiging-van-quantumcomputers/Brochure+Dreiging+Quantumcomputers%2C+webversie+september+2021.pdf

Amadori, A., Duarte, J. D., & Spini, G. (2022). Literature Overview of Public-Key Infrastructures, with Focus on Quantum-Safe Variants Deliverable 4.1, HAPKIDO Project. TNO.

Analistennetwerk Nationale Veiligheid. (2019). Leidraad risicobeoordeling Geïntegreerde risicoanalyse Nationale Veiligheid. https://www.rivm.nl/sites/default/files/2019-10/Leidraad%20Risicobeoordeling%202019.pdf

Aven, T. (2018). An Emerging New Risk Analysis Science: Foundations and Implications. Risk Analysis, 38(5), 876–888. https://doi.org/10.1111/risa.12899

Aven, T., Andersen, H. B., Cox, T., Droguett, E. L., Greenberg, M., Guikema, S., Kröger, W., McComas, K., Renn, O., Thompson, K. M., & Zio, E. (2018). Core Subjects of Risk Analysis (p. 7). Society for Risk Analysis.

Bharosa, N., Wijk, R. van, Winne, N. de, Janssen, Marijn, Luitjens, S., & Veld, P. (2015). Challenging the chain: Governing the automated exchange and processing of business information.

Davenport, T. H., & Short, J. E. (1990). The New Industrial Engineering: Information Technology and Business Process Redesign. MIT Sloan Management Review. https://sloanreview.mit.edu/article/the-new-industrial-engineering-information-technology-and-business-process-redesign/

de Wolf, R. (2017). The potential impact of quantum computers on society. Ethics and Information Technology, 19(4), 271–276. https://doi.org/10.1007/s10676-017-9439-z

Esteves, A. M., Franks, D., & Vanclay, F. (2012). Social impact assessment: The state of the art. Impact Assessment and Project Appraisal, 30(1), 34–42. https://doi.org/10.1080/14615517.2012.660356

Gregor, S. (2006). The Nature of Theory in Information Systems. MIS Quarterly, 30(3), 611–642. https://doi.org/10.2307/25148742

Grimes, R.A., 2019. Cryptography apocalypse: preparing for the day when quantum computing breaks today's crypto. Wiley, 1st Edition

Hevner, A. (2007). A Three Cycle View of Design Science Research. Scandinavian Journal of Information Systems, 19(2). https://aisel.aisnet.org/sjis/vol19/iss2/4

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. MIS Quarterly, 28(1), 75–105. https://doi.org/10.2307/25148625

International Organization for Standardization [ISO]. (2018). Information technology—Security techniques—Information security management systems—Overview and vocabulary (ISO/IEC 27000:2018). https://www.iso.org/obp/ui/#iso:std:iso-iec:27000:ed-5:v1:en

Joosten, R., & Smulders, A. (2014). How to successfully manage risks in hyperconnected value networks (p. 44). TNO.

Joseph, D., Misoczki, R., Manzano, M., Tricot, J., Pinuaga, F. D., Lacombe, O., Leichenauer, S., Hidary, J., Venables, P., & Hansen, R. (2022). Transitioning organizations to post-quantum cryptography. Nature, 605(7909), 237–243. https://doi.org/10.1038/s41586-022-04623-2

Le Fevre, M., Gölz, B., Flohr, R., Stelkens-Kobsch, T., & Verhoogt, T. (2017). SecRAM 2.0: Security Risk Assessment methodology for SESAR 2020. SESAR. https://www.sesarju.eu/sites/default/files/documents/transversal/SESAR%202020%20-%20Security%20Reference%20Material%20Guidance.pdf

Lindsay, J. R. (2020). Demystifying the Quantum Threat: Infrastructure, Institutions, and Intelligence Advantage. Security Studies, 29(2), 335–361. https://doi.org/10.1080/09636412.2020.1722853

Ma, C., Colon, L., Dera, J., Rashidi, B., & Garg, V. (2021). CARAF: Crypto Agility Risk Assessment Framework. Journal of Cybersecurity, 7(1), tyab013. https://doi.org/10.1093/cybsec/tyab013

Mavroeidis, V., Vishi, K., Zych, M. D., & Jøsang, A. (2018). The Impact of Quantum Computing on Present Cryptography. International Journal of Advanced Computer Science and Applications (Ijacsa), 9(3), Article 3. https://doi.org/10.14569/IJACSA.2018.090354

Mosca, D. M., & Piani, D. M. (2021). Quantum Threat Timeline Report 2020 (p. 52). Global Risk Institute. https://globalriskinstitute.org/publications/quantum-threat-timeline-report-2020/

Mosca, M., & Mulholland, J. (2017). A Methodology for Quantum Risk Assessment (p. 6) [Whitepaper]. Global Risk Institute.

Mulholland, J., Mosca, M., & Braun, J. (2017). The Day the Cryptography Dies. IEEE Security Privacy, 15(4), 14–21. https://doi.org/10.1109/MSP.2017.3151325

Presley, A., & Liles, D. (1998). The Use of IDEF0 for the Design and Specification of Methodologies.

Raban, Y., & Hauptman, A. (2018). Foresight of cyber security threat drivers and affecting technologies. Foresight, 20(4), 353–363.

Shamala, P., Ahmad, R., & Yusoff, M. (2013). A conceptual framework of info structure for information security risk assessment (ISRA). Journal of Information Security and Applications, 18(1), 45–52.

Smart, N. P. (2016). Cryptography made simple. Springer. https://link.springer.com/book/10.1007/978-3-319-21936-3

Stake, R. (2003). Case Studies. In N. K. Denzin & Y. S. Lincoln (Eds.), Strategies of qualitative inquiry (2nd ed). Sage.

TNO. (n.d.). Analistennetwerk Nationale Veiligheid (ANV). TNO. Retrieved 16 March 2022, from https://www.tno.nl/nl/aandachtsgebieden/defensie-veiligheid/roadmaps/nationale-veiligheid/crisisbeheersing-nieuwe-uitdagingen-nieuwe-kansen/analistennetwerk-nationale-veiligheid/

Vermaas, P. E. (2017). The societal impact of the emerging quantum technologies: A renewed urgency to make quantum theory understandable. Ethics and Information Technology, 19(4). https://doi.org/10.1007/s10676-017-9429-1

Wadhwa, K., Barnard-Wills, D., & Wright, D. (2015). The state of the art in societal impact assessment for security research. Science and Public Policy, 42(3), 339–354. https://doi.org/10.1093/scipol/scu046

Wangen, G., Hallstensen, C., & Snekkenes, E. (2018). A framework for estimating information security risk assessment method completeness. International Journal of Information Security, 17(6), 681–699. https://doi.org/10.1007/s10207-017-0382-0

Yunakovsky, S. E., Kot, M., Pozhar, N., Nabokov, D., Kudinov, M., Guglya, A., Kiktenko, E. O., Kolycheva, E., Borisov, A., & Fedorov, A. K. (2021). Towards security recommendations for public-key infrastructures for production environments in the post-quantum era. EPJ Quantum Technology, 8(1), 1–19. https://doi.org/10.1140/epjqt/s40507-021-00104-z

Appendix A – Menti meter results

Menti-meter results of user workshop 6 (13-04-2022) with four participants.

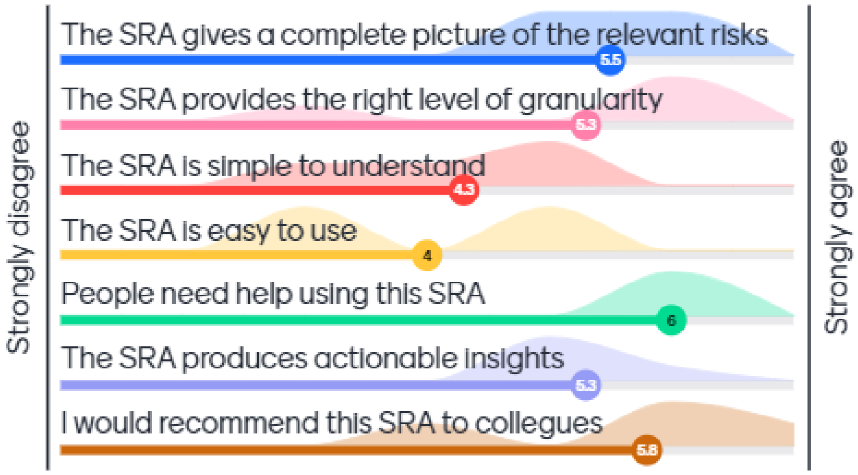

Figure 6: Menti-meter results workshop 6 (n=4)

Appendix B – Results from the various workshops

Workshop 1

Date	02-11-2021
SRA version	0.1
Participants	PKI manager at PA, Microsoft, manager at TSP, cyber security of critical infrastructures scientist

Strengths	Weaknesses
Threat capacities are clear and recognised Necessity is recognised The trust supply side sees the potential of the method	Does not include cyber criminals as a threat actor Does not recognise that anyone may eventually acquire access to quantum computing [o]
Opportunities	**Threats**
Make use of existing Business Impact Analyses [r] Invest in storytelling [a] Present the three impact levels visually	Skipping national impact areas may leave some impacts unidentified [b]

SRA update v0.1 -> v0.2		
Choice		Rationale
a	Add a text box beneath each table to be filled in, that is used to add context to the choices made in the tables. Applies to all steps.	The audience from workshop 1 gave the feedback that storytelling is important to convey the message and reach the goals of this SRA. Moreover, it allows for elaboration on the choices made during the application process of the SRA.
b	Include all national impact criteria from the GRNV instead of a selection	The audience of workshop 1 indicated that there were relevant impacts outside of the initially selected impact areas. There is a risk of not capturing all relevant impacts when leaving out impact areas.
c	Add cybercriminals to the threat actors	Cybercriminals are potential threat actors as well. These were initially overlooked.

Workshop 2

Date	13-12-2021
SRA version	0.2
Participants	M2M communication program lead at a large bank

Strengths	Weaknesses
There are many different applications/solutions dependent on PKI which are captured by the asset identification step (e.g. all customer interfaces; My business, my data; BIV; SSC) Organisational impact matrix is intuitive and invites discussion Very recognisable method Different levels of analysis make sense	It does not matter so much what threat actor attacks for the societal impact[o] The terms friendly and hostile state actor are not nuanced enough to reflect reality[d,o]
Opportunities	**Threats**
Couple likelihood of scenarios with the threat actors[o] Add a middle layer state actor category[d]	The likelihood is very hard to properly identify

SRA update v0.2 -> v0.3		
Choice		Rationale
d	Refrain from adding extra middle layer state actor category	This is not likely to increase the quality of the output. It will not capture the nuance as it is still too limited and will also weaken the focus of the analysis.

Relevant notes
The bank has firmly committed to digitalisation and the use of certificates. The entire modern way of banking is reliant on certificates. In the worst case we will have to fall back to a ‘paper society’.
The large bank is planning to employ a QTSP to provide QES technology for its digital trust needs.
For various processes that the bank fulfils, several different types of certificates are used that rely on differing PKI structures.
When all Dutch electronic financial transaction traffic is interrupted for half a day, all of the Netherlands will be bankrupt.
There is no other sector that has more societal impact than the financial sector when it is interrupted, even healthcare.

Workshop 3

Date	12-01-2022
SRA version	0.3
Participants	PKI manager at QTSP, compliance officer at QTSP, technical expert at QTSP

Strengths	Weaknesses
Making an inventory of PKI applications and peeling off each layer of impact is a good approach The inclusion of multiple perspectives.	No distinction between compromise and mitigation risks [h]
Opportunities	**Threats**
Have the outcome of the RA be a prioritisation of assets to make quantum safe first [g] Leave room for a broad perspective of risk and cascading impact of trust loss Mind the PKI use-case as it affects the impacts [e] Include the risks that come with transitioning to QS PKI [f] Use the term 'applications' to indicate assets that make use of PKI [e]	The term 'assets' can be interpreted too broadly [e] The absence of a distinction between compromise and mitigation risks can muddy the discussion [h]

SRA update v0.3 -> v0.4		
Choice		Rationale
e	Change the term 'asset' to 'application'	'Application' better reflects the use of PKI rather than the technical components behind it
f	Explicitly keep the risks of transitioning/migrating to QS PKI out of scope	This best reflects the original research question of HAPKIDO WP1 upon which this research is based
g	Have the last step be a prioritisation of risks	Not all systems can be simultaneously updates to be QS, therefore it is important to identify the biggest risks so they can be mitigated first
h	Scope the RA to risks of compromise and keep mitigation measures and their risks out of scope	It is better fitting for this research and HAPKIDO WP1 to focus on the risks that emerge from not acting. Additionally, the different types of compromise cover potential mitigation risks for a large part.
t	Have the national impact judged per business process	By combining the trust loss of all affected applications, the assessor is free to speculate about a scenario used for the national impact assessment including cascading loss of trust

Relevant notes
The level of trust demanded by the market is expected to keep increasing. As the demand for a higher level of trust increases, so does the demand for qualified trust services, as qualification is a means to ensure a high level of trust.
Some assets are less sensitive / critical than others and it is not realistic to make all systems quantum safe at once.

Workshop 4

Date	26-01-2022
SRA version	0.4
Participants	M2M communication program lead at tax authority, risk management expert

Strengths	Weaknesses
Organisational impact assessment feels familiar, invites discussion, and is workable Individual impact assessment is workable	Friendly state actors are not really a clear concept, as they can still employ hostile activities [o] The term ‘applications’ is ambiguous as it can mean a technical component using PKI as well as a business process in which PKI is applied [k] Lack of specific scoping can muddy the discussion and decrease the quality of the outcomes [l m n]
Opportunities	Threats
Define state actors by their motivation to spy or disrupt [o] The assessor can decide the level of analysis when deciding on the assets: applications can be subdivided in domains per application in order to group the dependent services [i] Similarly, applications can be grouped into application groups [i] Give room for differences in the levels of confidentiality of information [j]	The list of dependent services per application can get quite long [i] Organisational impact: capacity & performance may be differently interpreted by different assessing parties, depending on their industry [s]

SRA update v0.4 -> v0.5		
Choice		**Rationale**
i	Make explicit that it is the duty of the assessor to decide the proper level of analysis for the assets	Different situations call for different levels of analysis. The assessor is able to see the specific situation and can choose the appropriate level of analysis, rather than the method specifying a specific level of analysis for all situations.
j	Abstain from explicitly incorporating levels of confidentiality	Difference in levels of confidentiality is reflected in impacts. The impact of highly confidential information leaking will be higher. In the end, the impact resulting from the level of confidentiality is what counts. Additionally, it is reflected in the time to act when considering a store now, decrypt later attack.
k	Separate applications and business processes when making an inventory of assets	The two concepts need to be separated as it can cause confusion to only use 'applications'. This allows for the assessor to link the two together and get a clear picture of what is discussed.
l	Add an extra step beforehand for specifying the scopes used in the RA	Lack of specific scoping can hinder the analysis and decrease the quality of output
m	Introduce three different scopes: technical, organisational, and societal influence	These three scopes should keep the discussions focussed. The technical scope limits the discussion to the relevant technologies
n	Put the organisational impact assessment before the identification of services to society	By swapping these steps, the assessor stays within the organisational scope before taking a broader societal perspective. The aim is to reduce confusion because of scope switching.
o	Change the threat assessment to identifying threats and vulnerabilities that can be exploited by the threats, leaving out threat actors	For this type of risk assessment, the threat actor does not really matter. In the end, the attacks and following impacts are all the same.
s	Abstain from further defining the impact areas capacity & performance in the organisational impact	The room for interpretation in these impact areas is a necessary evil. This way, the SRA is more widely applicable as different parties in the trust chain may have different kinds of capacity.

Relevant notes
By using PKI, the tax authority gathers and sends data from and to many kinds of actors for various business processes.
Realising that there is a threat might be the biggest obstacle in transitioning to a QS version of SBR.
Going back to paper is not a viable option.

Right now, within the organisation there is not enough thought on how risks will play out nationally in a broad context. It is necessary to place risks in a broad perspective.

Workshop 5

Date	24-02-2022
SRA version	0.5
Participants	PKI manager at PA, cyber security scientist

Strengths	Weaknesses
The goal this method sets out to achieve is very relevant	Switching between making an inventory of assets and assessing the impacts was confusing [p] Terms are not clearly defined [q]
Opportunities	**Threats**
Documents such as a Business Impact Analyses and previously done risk assessments can be valuable input [r]	Parties are not likely to share sensitive information about their perceived vulnerabilities such as a Business Impact Analysis and previously done risk assessments [r]

SRA update v0.5 -> v1.0		
Choice		Rationale
p	Swap back the steps making an inventory of dependent services and organisational impact assessment	Switching between making an inventory of assets and assessing the impacts turned out to be confusing. Keeping them separate and switching between the organisational and societal influence scope seems better.
q	Give a clear definition for all terms used in the template	Clear definitions avoid confusion and improve the quality and speed of the analysis
r	Make explicit what input documents might be of use	Increased quality of input means increased quality of output. As most serious organisations should have useful input documents available, it is wise to make use of them.
s	Remove the CIA component in the SRA	We have remove the CIA component in the SRA in order to reduce complexity.

Relevant notes
Any method that helps taking steps in dealing with the quantum threat to PKI is welcome.
The method can be used making use of confidential input documents but can also be used without these documents to arrive at more general output.
Most serious organisations should have a Business Continuity Plan, including a Business Impact Analysis and risk assessments.