Transitioning organizations to post-quantum cryptography

https://doi.org/10.1038/s41586-022-04623-2

Received: 18 May 2021

Accepted: 8 March 2022

Published online: 11 May 2022

Check for updates

David Joseph[1]✉, Rafael Misoczki[2], Marc Manzano[1], Joe Tricot[1], Fernando Dominguez Pinuaga[1], Olivier Lacombe[2], Stefan Leichenauer[1], Jack Hidary[1], Phil Venables[2] & Royal Hansen[2]

Quantum computers are expected to break modern public key cryptography owing to Shor's algorithm. As a result, these cryptosystems need to be replaced by quantum-resistant algorithms, also known as post-quantum cryptography (PQC) algorithms. The PQC research field has flourished over the past two decades, leading to the creation of a large variety of algorithms that are expected to be resistant to quantum attacks. These PQC algorithms are being selected and standardized by several standardization bodies. However, even with the guidance from these important efforts, the danger is not gone: there are billions of old and new devices that need to transition to the PQC suite of algorithms, leading to a multidecade transition process that has to account for aspects such as security, algorithm performance, ease of secure implementation, compliance and more. Here we present an organizational perspective of the PQC transition. We discuss transition timelines, leading strategies to protect systems against quantum attacks, and approaches for combining pre-quantum cryptography with PQC to minimize transition risks. We suggest standards to start experimenting with now and provide a series of other recommendations to allow organizations to achieve a smooth and timely PQC transition.

In the past few decades, the field of cryptography has developed from an obscure set of rudimentary scrambling techniques into a mature, formal science. Along with better cryptographic techniques, a set of cryptanalysis techniques has arisen. One of these cryptanalysis techniques is related to quantum computers and threatens the foundations of the security guarantees that cryptography strives to offer[1] (see a review[2] for a comprehensive overview of the post-quantum cryptography (PQC) field).

The adoption of such post-quantum cryptographic techniques constitutes a challenge in itself. In this Perspective, we present an application-focused perspective of the transition process to protect organizations (including businesses, government departments and non-profit organizations) from quantum threats. Our perspective is derived from extensive discussions across security teams within Alphabet, and substantially agrees with established best practices in the information security and cryptography communities. The scale of the challenges faced by our colleagues and the pressing timeline within which they must be confronted lead us to believe that now is an opportune moment to open the discussion to a wider array of stakeholders in business, government and other organizations.

We present a set of actionable recommendations to organizations: from outlining the reasons why they should craft a robust strategy to start the migration to post-quantum cryptosystems now and increase awareness and understanding of PQC, to an analysis of the computational resources these new cryptosystems will require. We believe that taking critical steps now will be beneficial to reduce the future shortcomings of rushing through poorly planned countermeasures down the road. We intend this document to be relevant to a wide audience, and particularly to those in industry and government.

Post-quantum cryptography

In general terms, cryptography is the study of mathematical techniques to enforce policies on information. These policies broadly specify who is allowed to send, read and edit digital information. Some common uses include security against eavesdroppers, enforcing read and write access to data, and message authentication. All these techniques have something in common: they depend on the intractability of certain mathematical problems. To ensure that a cryptosystem is secure, there is a need to show that breaking such a cryptosystem is at least as hard as solving some mathematical problem considered intractable to anyone who does not possess knowledge of some piece of secret information, henceforth known as a key. Implementation errors aside, the hardness of this problem is the core security guarantee of the cryptosystem, and if the hardness is refuted by a cryptanalysis technique, then the cryptosystem is considered broken.

The quantum threat to traditional cryptography

The key mathematical techniques underpinning today's cryptosystems are closely related, and are based on the integer factorization problem and the discrete logarithm problem. The cryptosystem's security relies on the hardness of solving these problems. In 1994, the mathematician

[1]SandboxAQ, Palo Alto, CA, USA. [2]Google, Mountain View, CA, USA. ✉e-mail: david.joseph@sandboxaq.com

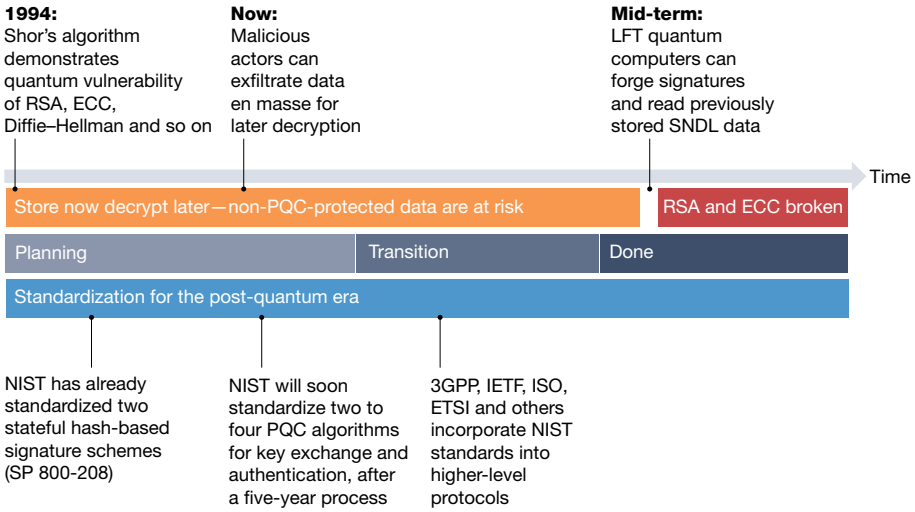

Fig. 1 | Post-quantum cryptography timeline. The three timelines can be thought of as: the threat to cryptography (top), the steps organizations should pass through during the migration (middle) and the process of standardization (bottom), which is led by multinational standards bodies.

Peter Shor devised a quantum algorithm that promised an exponential speed-up for factoring integers and finding discrete logarithms[1] over non-quantum algorithms, which in theory allows a quantum computer to crack the majority of the currently used public key cryptosystems. That is, many of our present cryptosystems will be broken when sufficiently large and fault-tolerant (LFT) quantum computers are built.

Quantum computers exist today, but they are highly rudimentary and imperfect machines and a great deal of technology evolution is needed to achieve wide application. The roadblocks for quantum computing lie mostly in creating high-precision hardware. Even with qubits that can execute basic operations with 0.1% error rates, over an entire system these errors propagate and grow exponentially, limiting the size of a useful quantum computer. Each extra qubit doubles the power of a quantum computer, and so when Google AI Quantum announced quantum supremacy in late 2019[3] their experiment was performed on a processor of only 53 qubits. The number of noisy qubits required to break RSA-2048—where RSA (Rivest–Shamir–Adleman) is the cryptosystem and 2048 is the most commonly used parameter set—is estimated to be around 20 million[4]. We argue below why action is urgent despite the engineering challenges this development implies.

Consequently, new cryptographic primitives are required to maintain the security of communication and information storage in the face of quantum threats. These cryptographic algorithms are known as post-quantum cryptography, and are based on mathematical problems that are believed to be quantum resistant. Although there exist quantum-based cryptographic techniques that are secure against quantum computers (see refs. [5–7] for relevant discussions), a substantial advantage of PQC over any quantum alternatives is that PQC schemes can be plugged into any conventional communication infrastructure or contemporary devices.

PQC transition timeline

This Perspective makes a set of recommendations to organizations about the process and timeline by which the PQC transition should take place, summarizes the landscape of the field, maps out the standardization timelines and compiles a list of resources for stakeholders. Figure 1 depicts a timeline of important PQC-related events that lie ahead. This timeline is composed of three parallel sequences of events and is not to scale.

The top red timeline in Fig. 1 captures the two most important quantum threats and when they become of critical importance. The first, known as a store-now-decrypt-later (SNDL) attack, is already an active threat. It corresponds to adversaries capturing valuable encrypted information now, storing it and decrypting it later once LFT quantum computers are available. The SNDL attack assumes that this information remains valuable in the future. The second quantum threat refers to the capability of breaking RSA and elliptic curve cryptography (ECC), the two most widespread public key algorithms for encrypting information today that can be broken with Shor's algorithm. This would allow adversaries to forge RSA and ECC digital signatures and pose risks to systems that rely on them, such as secure web browsing[8], zero trust architectures[9] and cryptocurrencies[10].

The middle grey timeline in Fig. 1 depicts the two actions required by organizations in transitioning to PQC. The first regards the strategic planning and technological experimentation for this transition, whereas the second regards the effective adoption of PQC in production systems. We emphasize that the strategic planning phase must be completed well before LFT quantum computers are able to effectively attack RSA and ECC (that is, a process whose start should not be further delayed).

Finally, the bottom blue timeline in Fig. 1 concerns the standardization processes organized by relevant government and industrial bodies, with particular focus on the National Institute of Standards and Technology (NIST) PQC process to determine the fundamental security of proposed PQC candidates.

Recommendations to organizations regarding strategy and timelines

Given its low cost, ease of integration into current infrastructure plus the whole set of cryptographic features, there has been a natural convergence of standardization bodies and organizations towards PQC. As such, we recommend that organizations interested in protecting their systems and users against quantum attacks should adopt PQC (over quantum cryptography) as their main quantum protection strategy.

For those organizations that have not started integrating PQC in their systems or even planning for it, we highly recommend starting their efforts now. Those organizations and enterprises with sensitive data with time value exceeding five years should consider PQC immediately. The SNDL attack is already practicable, so in this context, such organizations are already late and at increasing risk.

We recommend that, despite the natural pressure from industry, the standardization bodies should have a clear focus on creating standards with a security-first mindset. A standard that is published quickly but that lacks clarity brings about many more problems than a standard that is published later.

We recommend that organizations should explicitly express their interest in the PQC standards to motivate standardization bodies to allocate resources to this effort.

Timing the migration to PQC

When talking about quantum attacks, it is natural to question when the PQC transition needs to start. Given that LFT quantum computers are not available yet, in this section we provide arguments reinforcing why starting the PQC transition now is important.

Store now decrypt later

The SNDL attack poses a threat to information that is encrypted now using quantum-vulnerable cryptography. Such encrypted data, which are often transmitted over the public internet infrastructure, can be collected, stored indefinitely and then decrypted in the future once the adversary has access to a LFT quantum computer. In some situations, this is not a major concern. However, there are important trade secrets, medical records, national security documents and more that have multidecade shelf lives and must remain confidential for an extended period of time. For this reason, the SNDL attack is one of the most important arguments to not delay starting the transition any further.

Far-horizon projects

Another reason that PQC is of immediate importance concerns projects that are being designed and planned now but have long lifespans (for example, of multiple decades). Vehicles are a good example of this[11]—many cars, planes, trains and ships in production now are expected to be in service in up to 20 years' or even 30 years' time. In some cases, they will contain modules where one cryptosystem can be swapped out for another, but this is not true for all modules. This is particularly so where application-specific hardware is used to implement cryptography and remains immutable for the lifetime of the product.

Critical national infrastructure projects are another example where high availability is essential (with some applications requiring 99.999% availability, or 6 minutes of downtime per year[12]), and upgrading the cryptography software or hardware represents an unacceptable cost.

Cryptography transition takes time

History has shown that a cryptography transition takes a considerable amount of time. ECC was proposed back in the 1980s and, despite the fact it is much more efficient (space and speed wise, depending on parametrizations) than RSA[13–16], it took over two decades to finally gain widespread adoption. Hash functions are also another example of cryptographic tools that took a long time from their inception until gaining some adoption. For example, the NIST SHA-3 competition was announced back in 2007, its winner was announced in 2012[17] and still in 2021 SHA-3 has not seen widespread adoption. Therefore, cryptography transitions (even much simpler ones than the PQC transition) commonly take several years, or even decades. The PQC transition is more complex, given the fact many of the approaches are relatively new and that the performance of many candidates is considerably worse than current algorithms.

Clearer path for PQC standards adoption

The fourth reason to start this transition now is the series of announcements about PQC standardization. In 2019, NIST published the stateful hash-based signatures standard[18–20], and in 2020 it announced the third-round finalists[21] (and alternative track) of their PQC competition. As a result, the cryptography community now has a clearer direction of which primitives are likely to form the backbone of the PQC suite of standardized cryptosystems. Moving early allows time for the ironing out of bugs, training the workforce and preparing adequately for what will be a lengthy process.

Recommendations to organizations regarding the urgency to act

We recommend that organizations interested in protecting their systems and users against quantum attacks should start, at least, planning their PQC transition strategy now. The SNDL attack highlights the fact that most organizations are already late.

We caution against a sense of complacency that may develop from viewing this as 'just another cryptography transition'. This migration covers a wider and more complex scope than previous transitions. As such, more planning, time and resources should be allocated to this migration than for other past migrations.

PQC standardization

There are a number of standardization bodies working on standardization processes of PQC. These efforts are being led by the US-based NIST, the International Standards Organization (ISO), the Internet Engineering Task Force (IETF) and the European Telecommunications Standards Institute (ETSI). Each one of these processes is at a different stage and covering different PQC schemes.

Standardization of stateful hash-based signatures

Stateful hash-based signatures (HBS) are digital signature schemes whose security relies solely on the security of hash functions. This represents an advantage when compared with other digital signature schemes, from both quantitative and qualitative perspectives. From the quantitative perspective, HBS security relies only on the security of hash functions, whereas other signature schemes rely on the security of hash functions plus some other presumably hard problems. Recall that the fewer security assumptions the better for any cryptosystem. From the qualitative perspective, hash functions are among the most studied topics in cryptology, which means that their security properties are well understood, including their expected resistance against quantum attacks[22] (assuming appropriate digest sizes are used).

The statefulness property means that the signer needs to keep track of a state in-between signature generation. In practice, the state is an increasing counter. Reuse of the same state would break the security of the system. This poses deployment challenges for some applications, but not for all. For example, code-signing applications seem to be very suitable for stateful HBS schemes as the signer is on the server side, and thus it should be able to manage the state properly. There are stateless HBS schemes too[23], but those are comparatively less efficient than their stateful counterparts.

Given their optimal security properties and acceptable performance metrics (usually less efficient than RSA and ECC but not by a great margin), stateful HBS have already been (or soon will be) standardized by multiple standardization bodies, being the first PQC standards available for widespread adoption.

The IETF has published stateful HBS request-for-comments (RFCs), which are often adopted by industry players as informal standards. The IETF HBS standards cover the Leighton–Micali signature (LMS) scheme and eXtended Merkle signature scheme (XMSS)[18,19], and their multitree variants.

NIST is running two standardization efforts related to PQC, one of which is discussed in 'The NIST PQC project' below, whereas the other (completed) was focused specifically on stateful HBS[20] and has already finished: NIST standardized the same schemes for which the IETF published RFCs, namely the XMSS and LMS schemes, and their multitree variants. Finally, the ISO and the IEC have also been very active in this field. After a successful study period focused on the stateful HBS topic, ISO/IEC JTC1 SC27 Work Group 2 has begun work on the first ISO PQC standard draft, ISO/IEC 14888-4 on stateful HBS algorithms.

The NIST PQC project

NIST has been very active in shaping cybersecurity best practices for quite some time. For example, the widely used advanced encryption

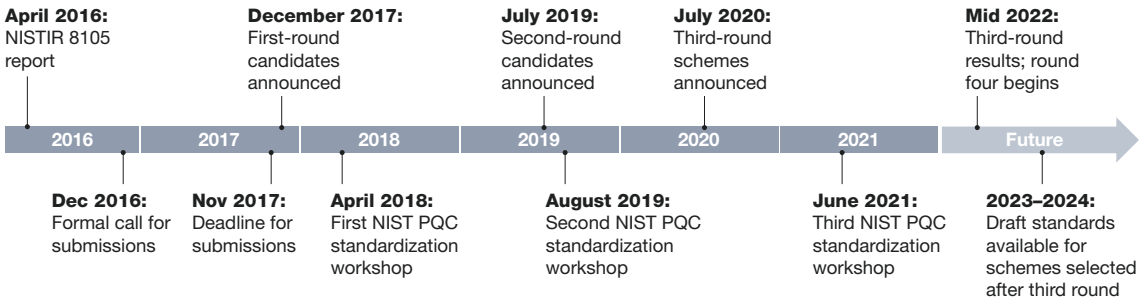

Fig. 2 | NIST post-quantum cryptography process timeline. The notable events during the course of the NIST PQC standardization process are shown, from its inception in 2016 to the present day. This process is the longest and most comprehensive study into PQC conducted thus far.

standard (AES) algorithm[24] and, more recently, the SHA-3 algorithm[17] are examples of these initiatives. In 2016, with the quantum threat looming, NIST launched a process to standardize public key PQC algorithms. Figure 2 depicts the timeline of events for this process.

Since the beginning of the process, NIST acknowledged that this particular process would be substantially more complex than the processes for AES and SHA-3[25]. One reason for this is that the requirements for public key cryptography and digital signatures are more complex than symmetric cryptography. Another is the sheer breadth of proposed solutions that research has provided. Comparing such varied approaches brings unique challenges, such as weighing up security, key sizes, latency, bandwidth and ease of secure implementation.

The process considers two cryptographic functionalities: stateless digital signature, and asymmetric encryption and key encapsulation mechanisms. The evaluation criteria for this process focus on security first and foremost, then on the practical considerations of efficiency and performance, and as a last priority consider other factors such as intellectual property claims and ease of secure implementation[21]. Parameter sets for five security levels ranging from the equivalent to conducting exhaustive key search on AES 128 (that is, level I) to AES 256 (that is, level V) are analysed, which allows cryptosystems from different families to be roughly compared with one another.

To narrow down the field of 82 submissions initially received, NIST considered the security evaluations that were provided along with the submissions, external security analyses, as well as internal cryptanalysis performed by NIST's own researchers. After just over a year, NIST announced 26 algorithms that would proceed to the second round of the process. In July 2020, NIST announced the 15 candidates that would proceed to the third round[21]. Of these 15 candidates, 7 were classified as 'finalists' (4 asymmetric encryption or key encapsulation mechanisms (KEMs) and 3 stateless signature schemes) and 8 were classified as 'alternatives' (5 asymmetric encryption or KEMs and 3 stateless signature schemes). At the end of the third round, it is expected that NIST will standardize a few finalist schemes, and will continue to consider alternative candidates for future standardization in an eventual fourth round. The intention to keep alternative candidates in the process could be explained by several reasons, including achieving diversity of primitives, suitability to special use cases, and more. NIST aims to release results of the third round by mid 2022, with the final standards taking up to another two years.

During its standardization process, NIST has disclosed benchmark results to illustrate potential performance gaps between the candidates. The charts in Fig. 3 show these differences. Here we make a few observations. Isogenies are extremely space efficient with small public keys and ciphertexts, but suffer poor speed performance. For lattices, the unstructured variants are considered the most conservative approach, and enjoy more confidence from the crypto community regarding security than their structured counterparts. Similarly the McEliece cryptosystem[26] is considered more conservative than other more recent code-based systems. By contrast, the structured variants perform better in all metrics at the cost of at least one additional security assumption. In summary, these two classes (structured versus unstructured code and lattice schemes) represent a trade-off between security and efficiency. NIST's decision to keep representatives from both categories for its third round seems prudent, as it gives time to the community to determine where the line should be drawn between efficiency and security.

Other PQC-related standardization efforts

The IETF, which was responsible for crafting the transport layer security (TLS) protocol[27,28] that is used extensively for secure web browsing[29,30], has several ongoing efforts to integrate post-quantum primitives in different protocols, for instance in TLS[31] and the internet key exchange (IKE) standard[32]. The intention is to combine RSA- and ECC-based and PQC schemes, providing a stepping stone towards PQC without risking naked vulnerabilities that are inherent with relatively new cryptography. The 3rd Generation Partnership Project (3GPP) have started initial discussions on the topic but will probably wait for the standards to be published by NIST before proposing new wireless encryption protocols. At ISO, SD8 of JTC 1/SC 27 offers information on a range of PQC algorithms.

ETSI is another organization that is taking an active role in the standardization of quantum-resistant communication technologies, with working groups on quantum key distribution (QKD) and PQC, the latter having published a standard for quantum-safe key exchange[33]. In 2015, ETSI released a white paper[34] analysing some of the most promising post-quantum cryptography algorithms and discussing the main challenges for this transition.

In a separate non-standardization project, NIST and the National Cybersecurity Centre of Excellence are working to stimulate development of tools, playbooks and proofs of concept to ease migration[35].

In China, the Chinese Association for Cryptographic Research completed a short competition held over a period of months in 2019 to quickly settle on a small number of algorithms for standardization. The first prize was awarded to the lightweight authenticated encryption cipher (LAC) scheme[36,37] in the key exchange category, a cryptosystem that made it to the second round of the NIST process[38] but not the third, owing to a number of successive attacks.

Recommendations to organizations regarding PQC standardization

Stateful HBS is a technology already standardized by multiple standardization bodies (for example, NIST). Despite the need to implement a robust state-management mechanism, this technology has an outstanding benefit: its security guarantees, which are based on minimal assumptions. Therefore, organizations that need to transition to PQC in applications amenable to state management (such as any software code signing application) should consider HBS as a potential solution.

The NIST PQC project is close to the end of the third round, and standards for the algorithms selected are expected to be released no

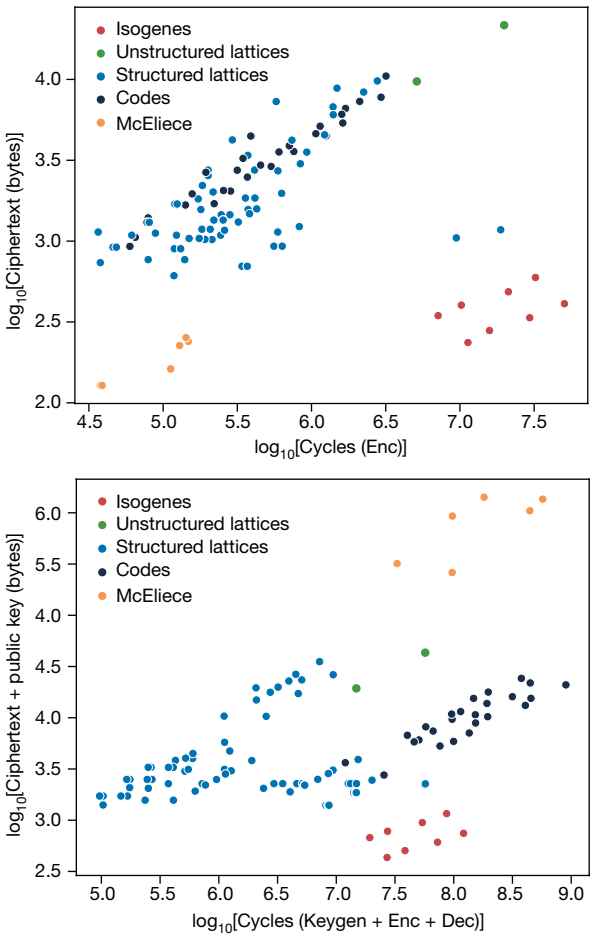

Fig. 3 | NIST post-quantum cryptography algorithm performance. The top plot shows the sizes of ciphertext and public keys for key exchange algorithms, where similar algorithms are grouped by colour. The bottom plot shows the ciphertext and public key sizes versus the computation required. A key generator (Keygen) is used to generate a public and private key, encapsulation (Enc) encrypts a secret using the public key, and decapsulation (Dec) decrypts the secret using the private key, and the secret is then securely shared.

later than 2024. We recommend that organizations monitoring this process should start experimenting now with the finalist and alternative candidates. This will be important to minimize the transition time once the NIST PQC standards are published.

General recommendations

At first glance, transitioning cryptographic algorithms may look like a simple task and similar to any other algorithmic replacement: the old generation of (quantum vulnerable) algorithms are replaced by the new generation of (quantum resistant) algorithms. Unfortunately, this task is anything but simple given the fact that adversaries can (and will) exploit any insecure node at any given time to trigger devastating attacks. To complement the topic-specific recommendations given throughout this Perspective, we now provide a set of generic recommendations that should further help organizations to transition to PQC in a manner that minimizes security risks, ensures a shorter transition time and optimizes costs.

Crypto-agility

To develop a holistic approach towards infrastructural security in the face of the post-quantum migration, organizations must take steps towards crypto-agility. This is because changing cryptographic algorithms, or layering them with existing cryptographic algorithms, requires more than the direct work itself. There are also differences in key sizes, encrypted file sizes and signature lengths. This latter set of attributes will probably have wider ramifications to application and infrastructure software, protocol specifications and application programme interfaces, and the standards that define them. Adapting infrastructure to accommodate such considerations will be a considerable part of the work of migrating to PQC[39]. In addition, despite the best efforts of the new algorithm authors and evaluators, there is potential for some degree of ongoing change, in algorithms, modes of operations or specific parameters that may in turn affect the wider system configuration.

In preparing to implement these changes, organizations should plan for crypto-agility, specifically to adopt abstraction layers on centrally managed toolkits and services that minimize the effort on any subsequent changes. Similarly, the implementation of such tooling should also include capabilities for the application or infrastructure users of such tooling to cope with adjusted data formats and sizes. Explicitly, in initiating a crypto-agility programme, organizations should: implement a centrally provided set of cryptographic libraries and services that abstract algorithms in use from application and infrastructure teams; and identify data field and size dependencies, and adjust surrounding databases, datastores, protocols and other software that assumes current fixed field sizes.

Concerning standards bodies and regulators, additionally: crypto-agility should be embedded in any standards that are currently being developed, for example, 6G must be inherently crypto-agile and PQC compatible; and industry-specific regulators across critical infrastructure sectors should urgently start planning for sectoral coordination to reduce systemic risk.

Prioritization strategy

The first thing organizations need to carefully pay attention to in order to ensure a successful transition is prioritization. This refers to the task of identifying where the PQC transition is needed first. This is important because the workforce able to perform this task is highly specialized and usually scarce. Consequently, deploying a transition strategy that does not take into account the main security bottlenecks will probably consume all the resources available without necessarily protecting systems and users.

Regarding this prioritization of efforts, first we need to identify the cryptographic schemes that are at highest risk. In the context of immediate need for confidentiality, key exchange algorithms are at highest risk. This is because the outcome of a key exchange procedure can be captured to be broken later (SNDL attack). Digital signatures, however, require an online adversary (that is, the adversary needs to be able to forge signatures at the time of signing). Meanwhile some systems are hard to update but do not absolutely require immediate quantum-resistant confidentiality, such as vehicular communications. In these cases, secure digital signatures should be a first step, which can then be used to push updates to key exchange algorithms at a later date.

Hybrid algorithms

Rather than replacing existing algorithms with comparatively less-studied post-quantum alternatives, the scientific community came up with a simple and effective approach. This approach consists of combining both a traditional algorithm and a post-quantum algorithm into a single mechanism. If done correctly, the overall system's security is lower bounded by the stronger of the two cryptosystems composing the hybrid system. In other words, even if the PQC algorithm is subsequently identified as flawed, the security offered by the classical scheme is still guaranteed. In this way, security is only potentially increased in this transition, never decreased.

To combine the key exchange algorithms, one uses each algorithm (one PQC and one classical) to generate a single shared secret. Then,

these two secrets could be combined to produce a single symmetric key. In this way, an adversary willing to attack the system will necessarily need to break both classical and PQC schemes. One question remaining is how to combine these two shared secrets. The main approaches are: to concatenate the shared secrets, either before or after passing through a key derivation function (KDF); XOR the shared secrets.

This is a topic that is gaining increasing traction among the academic community[40,41], although the KDF option seems to lead to the most conservative approach with minimal cost. Moreover, there are already proposals being drafted to combine digital signature schemes in a hybrid manner[42].

Certification requirements

For compliance reasons, many users of cryptography are required to adhere to the Federal Information Processing Standard (FIPS) defined by NIST. Switching out traditional public key cryptosystems for post-quantum alternatives would be a problem if this meant loss of FIPS certification. For example, federal agencies purchasing cryptography-based security systems must ensure a FIPS 140-2 certificate exists, otherwise it is ineligible. Fortunately this has been provisioned for, and the 'hybrid' approach mentioned above allows practitioners to maintain their FIPS 140-2 rating, owing to the fact that one is only increasing security, and not replacing it[43]. All that is required is that at least one of the schemes used in hybrid mode is FIPS 140-2 approved.

PQC resources

Part of planning for the transition is education on both the underlying theory of the hard problems providing the security for these new schemes, and getting to grips with software implementations to analyse their integration with companies' network infrastructure. Good theoretical resources include a lattice cryptography review[44] to better understand primitives pertaining to lattice cryptography, and the book *Post-Quantum Cryptography*[45], providing a broader explanation of the field, more relevant to those looking to understand schemes in the alternative track for the third round.

There are a number of places from which to obtain software implementations of the NIST schemes. The most comprehensive repository of these is Liboqs[46], of the Open Quantum Safe project. Other resources include BoringSSL[47] and Tink[48], which implement NTRU-HRSS and X25519 in hybrid mode (together called CECPQ2). For highly optimized implementations, consider SUPERCOP[49], which has become the de facto benchmarking tool for PQC schemes. Crypto hardware for PQC is less readily available, and is an area that seems to require more research moving forward. One can find relevant and up-to-date material at a number of conferences with varying focuses and levels of technicality. To name a few conferences on the International Association of Cryptological Research (IACR) calendar, there is Post-Quantum Cryptography, Real World Cryptography, Public-Key Cryptography, and Cryptographic Hardware and Embedded Systems.

Summary

In 2021, more than half of experts surveyed believed that the probability of an LFT quantum computer breaking integer factorization and discrete logarithm-based cryptography within 15 years was greater than 50% (ref. [50]), and so there is little time to transition to PQC, especially considering that even today's private data can be compromised by tomorrow's quantum computers owing to the SNDL attack, and that cryptographic hardware is being deployed now that is expected to remain in the field for many decades.

Owing to the diversity of schemes, the need to instantly switch from one algorithm to another in the event of a successful attack, and the requirement for increasing connectedness between systems, crypto-agility must be front of mind when designing new protocols.

Work on the transition does not have to wait for full NIST standards, as hybrid cryptography allows practitioners to safely deploy quantum-resistant schemes without compromising current security levels. The US government released a memorandum[51] on transitioning to quantum-resistant cryptographic protocols in early 2022, setting a strong example to both public and private organizations globally, indicating that preparatory work for the transition should begin as soon as possible. The steps that organizations can be taking now to mitigate against the quantum computing threat include: building a crypto inventory of where all public key cryptography exists within their infrastructure and products; creating a roadmap for enacting the transition once standards have been published; experimenting with different families of PQC algorithms, to measure performance, investigate different approaches (for example, hybrid), assess interdependencies and so on; and ensuring that systems are crypto-agile, that is, ready to transition to PQC with minimal cost and time.

None of these tasks is a trivial matter, and they can all be started now, helping to responsibly work towards a safer future for organizations, systems and users.

Data availability

The datasets analysed in the report are available from SUPERCOP at https://bench.cr.yp.to/supercop.html. Source data are provided with this paper.

1. Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. In *Proc. 35th Annual Symposium on Foundations of Computer Science* 124–134 (Soc. Industr. Appl. Math., 1994).
   Shor's quantum algorithm demonstrated how to factorize large integers in polynomial time, which is an exponential speed-up over the best classical algorithms.
2. Bernstein, D. J. & Lange, T. Post-quantum cryptography. *Nature* **549**, 188–194 (2017).
3. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
4. Gidney, C. & Ekerå, M. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum* **5**, 433 (2021).
   Gidney and Ekerå describe the resources required to implement Shor's algorithm to break today's standard cryptography, assuming noisy qubits.
5. Bennett, C. H. & Brassard, G. Quantum cryptography: public key distribution and coin tossing. *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing* 175–179 (1984).
6. Alagic, G. et al. Computational security of quantum encryption. In *International Conference on Information Theoretic Security* 47–71 (Springer, 2016).
7. Barnum, H., Crepeau, C., Gottesman, D., Smith, A. & Tapp, A. Authentication of quantum messages. In *Proc. 43rd Annual IEEE Symposium on Foundations of Computer Science* 449–458 (IEEE, 2002).
8. Paquin, C., Stebila, D. & Tamvada, G. Benchmarking post-quantum cryptography in TLS. In *International Conference on Post-Quantum Cryptography* 72–91 (Springer, 2020).
9. Rose, S., Borchert, O., Mitchell, S. & Connelly, S. *Zero Trust Architecture* (NIST, 2020); https://csrc.nist.gov/publications/detail/sp/800-207/final
10. Kearney, J. J. & Perez-Delgado, C. A. Vulnerability of blockchain technologies to quantum attacks. *Array* **10**, 100065 (2021).
11. Lemke, K., Paar, C. & Wolf, M. *Embedded Security in Cars* (Springer, 2006).
12. Anderson, R. & Fuloria, S. Security economics and critical national infrastructure. In *Economics of Information Security and Privacy* 55–66 (Springer, 2010).
13. Gura, N., Patel, A., Wander, A., Eberle, H. & Shantz, S. C. Comparing elliptic curve cryptography and RSA on 8-bit CPUs. In *International Workshop on Cryptographic Hardware and Embedded Systems* 119–132 (Springer, 2004).
14. Rivest, R. L., Shamir, A. & Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* **21**, 120–126 (1978).
15. Miller, V. S. Use of elliptic curves in cryptography. In *Conference on the Theory and Application of Cryptographic Techniques* 417–426 (Springer, 1985).
16. Koblitz, N. Elliptic curve cryptosystems. *Math. Comput.* **48**, 203–209 (1987).
17. Chang, S. et al. *Third-Round Report of the SHA-3 Cryptographic Hash Algorithm Competition* NISTIR 7896 (NIST, 2012).
18. Hülsing, A., Butin, D., Gazdag, S.-L., Rijneveld, J. & Mohaisen, A. XMSS: eXtended Merkle signature scheme. RFC 8391 (2018); https://datatracker.ietf.org/doc/html/rfc8391
19. McGrew, D., Curcio, M. & Fluhrer, S. Leighton-Micali hash-based signatures. RFC 8554 (2019); https://datatracker.ietf.org/doc/html/rfc8554
20. Cooper, D. A. et al. *Recommendation for Stateful Hash-based Signature Schemes* NIST Special Publication 800-208 (NIST, 2020); https://csrc.nist.gov/publications/detail/sp/800-208/final
21. Alagic, G. et al. *Status Report on the Second Round of the NIST Post-quantum Cryptography Standardization Process* (US Department of Commerce, NIST, 2020); https://csrc.nist.gov/publications/detail/nistir/8309/final

This report describes NIST's findings after evaluation of the second round, and explains the motivation for selecting the seven finalist schemes as well as the eight alternative track schemes for evaluation in the third round.
22. Gheorghiu, V. & Mosca, M. Benchmarking the quantum cryptanalysis of symmetric, public-key and hash-based cryptographic schemes. Preprint at https://arxiv.org/abs/1902.02332 (2019).
23. Bernstein, D. J. et al. SPHINCS: practical stateless hash-based signatures. In *Proc. EUROCRYPT* Vol. 9056 368–397 (Springer, 2015).
24. Nechvatal, J. et al. Report on the development of the advanced encryption standard (AES). *J. Res. Natl Inst. Stand. Technol.* **106,** 511–577 (2001).
25. Chen, L. et al. *Report on Post-quantum Cryptography* (NIST, 2016); https://csrc.nist.gov/publications/detail/nistir/8105/final
26. McEliece, R. J. A public-key cryptosystem based on algebraic coding theory. Jet Propulsion Laboratory, Pasadena. DSN Progress Reports **4244**, 114–116 (1978).
27. Dierks, T. & Allen, C. The TLS protocol version 1.0. RFC 2246 (1999); https://www.ietf.org/rfc/rfc2246.txt
28. Rescorla, E. & Dierks, T. The transport layer security (TLS) protocol version 1.3. RFC 8446 (2018); https://datatracker.ietf.org/doc/html/rfc8446
29. Rescorla, E. & Schiffman, A. The secure hypertext transfer protocol. RFC 2660 (1999); https://datatracker.ietf.org/doc/html/rfc2660
30. Holz, R., Amann, J., Mehani, O., Wachs, M. & Kaafar, M. A. TLS in the wild: an Internet-wide analysis of TLS-based protocols for electronic communication. *Proceedings of the Network and Distributed System Security Symposium* (NDSS) (2016).
31. Steblia, D., Fluhrer, S. & Gueron, S. *Hybrid Key Exchange in TLS 1.3* (IETF, 2020); https://tools.ietf.org/id/draft-stebila-tls-hybrid-design-03.html
32. Tjhai, C. et al. *Multiple Key Exchanges in IKEv2* (IETF, 2021); https://www.ietf.org/archive/id/draft-ietf-ipsecme-ikev2-multiple-ke-03.txt
33. *CYBER; Quantum-Safe Hybrid Key Exchanges* ETSI TS 103 744, (ETSI, 2020); https://www.etsi.org/deliver/etsi_ts/103700_103799/103744/01.01.01_60/ts_103744v010101p.pdf
34. *Quantum Safe Cryptography and Security; An Introduction, Benefits, Enablers and Challenges* White Paper No. 8 (ETSI, 2015); https://www.etsi.org/technologies/quantum-safe-cryptography
35. Barker, W., Souppaya, M. & Newhouse, W. *Migration to Post-Quantum Cryptography* (NIST & CSRC, 2021); https://csrc.nist.gov/publications/detail/white-paper/2021/08/04/migration-to-post-quantum-cryptography/final
36. Lu, X. et al. LAC: practical ring-LWE based public-key encryption with byte-level modulus. *IACR Cryptol. ePrint Arch.* **2018**, 1009 (2018).
37. Announcement of nation-wide cryptographic algorithm design competition result. *Chinese Association for Cryptology Research* https://www.cacrnet.org.cn/site/content/854.html (2021).
38. Alagic, G. et al. *Status Report on the First Round of the NIST Post-Quantum Cryptography Standardization Process* (NIST, 2019); https://www.nist.gov/publications/status-report-first-round-nist-post-quantum-cryptography-standardization-process
39. Ott, D. et al. Identifying research challenges in post quantum cryptography migration and cryptographic agility. Preprint at https://arxiv.org/abs/1909.07353 (2019).
40. Bindel, N., Brendel, J., Fischlin, M., Goncalves, B. & Stebila, D. Hybrid key encapsulation mechanisms and authenticated key exchange. In *International Conference on Post-Quantum Cryptography* 206–226 (Springer, 2019).
41. Crockett, E., Paquin, C. & Stebila, D. Prototyping post-quantum and hybrid key exchange and authentication in TLS and SSH. *IACR Cryptol. ePrint Arch.* **2019**, 858 (2019). **Implementations of NIST round two PQC algorithms in TLS, providing insightful data on which algorithms are likely to be performant enough for widespread use and which will suffer severe performance issues**.
42. Ounsworth, M. & Pala, M. *Composite Signatures For Use In Internet PKI* (IETF, 2021); https://www.ietf.org/archive/id/draft-ounsworth-pq-composite-sigs-05.txt
43. Barker, E., Chen, L. & Davis, R. *Recommendation for Key-Derivation Methods in Key-Establishment Schemes* (NIST, 2020); https://www.nist.gov/publications/recommendation-key-derivation-methods-key-establishment-schemes
44. Peikert, C. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.* **10**, 283–424 (2016).
45. Bernstein, D. J., Buchmann, J. & Dahmen, E. *Post-Quantum Cryptography* (Springer, 2009).
46. Stebila, D. & Mosca, M. Post-quantum key exchange for the internet and the open quantum safe project. In *International Conference on Selected Areas in Cryptography* 14–37 (Springer, 2016).
47. Langley, A. BoringSSL. *GitHub* https://github.com/google/boringssl (2020).
48. Duong, T. Tink. *GitHub* https://github.com/google/tink (2020).
49. Bernstein, D. J. & Lange, T. SUPERCOP: system for unified performance evaluation related to cryptographic operations and primitives (VAMPIRE Lab, 2018); https://bench.cr.yp.to/supercop.html
50. Mosca, M. & Piani, M. *Quantum Threat Timeline* (Global Risk Institute, 2021); https://globalriskinstitute.org/publications/2021-quantum-threat-timeline-report/
51. Memorandum on Improving the Cybersecurity of National Security, Department of Defense, and Intelligence Community Systems. *The White House* https://www.whitehouse.gov/briefing-room/presidential-actions/2022/01/19/memorandum-on-improving-the-cybersecurity-of-national-security-department-of-defense-and-intelligence-community-systems/ (2022).

Author contributions D.J., R.M. and M.M. drafted the paper and provided technical expertise. J.T., F.D.P., O.L., P.V. and S.L. participated in extensive discussions, providing business and organizational perspectives and edits, and J.H. and R.H. drove the project from an executive level, helping to gather resources, provide direction and edit the manuscript. A substantial part of this paper was written while all the authors were a part of Alphabet.

Competing interests The authors declare no competing interests.

Additional information
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41586-022-04623-2.
Correspondence and requests for materials should be addressed to David Joseph.
Peer review information *Nature* thanks Tanja Lange and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.
Reprints and permissions information is available at http://www.nature.com/reprints.
Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.