TOPIC

Migration to post-quantum cryptography

Mastercard R&D white paper

Table of contents

Executive summary

The promise of unprecedented quantum computational power is both a risk and an opportunity for the financial industry.

Quantum computing, a transformative technology with the potential to outperform classical computers, is advancing. Governments and private sectors worldwide have invested billions in its development, signalling a shared belief in its ability to reshape industries.

While quantum computing offers enormous potential in fields like pharmaceuticals, logistics, and material sciences, it threatens the cryptographic foundations that secure financial systems today.

Public-key cryptographic methods help establish digital trust in most digital infrastructures, including our financial systems, which are primarily comprised of parties communicating across insecure or untrusted channels. Today's public-key cryptosystems, built mostly on the cryptosystems of RSA, named after its inventors Rivest, Shamir, and Adleman, as specified in Section 6 of NIST Special Publication 800-56B, and Elliptic Curve Cryptography (ECC) as detailed in NIST Special Publication 800-186, are foundational in securing transactions and sensitive data.

Sufficiently advanced quantum computers would be able to completely break these cryptographic algorithms and associated protocols, rendering financial institutions vulnerable to data breaches, financial losses, and reputational damage. While the timeline for the building of a Cryptographically Relevant Quantum Computer (CRQC) is uncertain, we believe the urgency to act is clear. A reactive cybersecurity approach against quantum threats no longer suffices. Financial organizations must plan to adopt quantum-safe practices to mitigate the risks.

It is our position investing in quantum-safe technologies that can ensure security when large-scale quantum computers become widely available. Predominantly, there are two categories of quantum-safe security technologies, namely, Post-Quantum Cryptography (PQC) and Quantum Key Distribution (QKD). Quantum migration refers to the process of

transitioning from classical cryptographic systems to quantum-safe alternatives such as PQC and QKD. Standardisation bodies and government agencies across the world have been actively working on PQC and QKD standards to guide the industry in transitioning to quantum-safe security systems.

Financial institutions should begin exploring these alternatives now to make their systems and infrastructure resilient against quantum threats. It is our view that this migration is pivotal in maintaining the integrity of encrypted communications, securing payment systems, and protecting sensitive customer data. The journey is replete with not just technological challenges but also operational, regulatory, and strategic complexities. As quantum computing edges closer to reality, financial institutions can start to prepare now, implementing proactive quantum-resistant strategies that align with evolving standards and regulatory frameworks.

This whitepaper highlights the significance of quantum migration for the financial sector, offering insights into the challenges and necessary steps to ensure a timely, smooth, and safe transition. In this paper, we separate the hype from the practical reality and provide an evidence-based assessment and analysis on the presumed threat.

We argue that early adopters of quantum migration today will be best positioned to protect their assets and maintain resilience in the face of future threats.

Objectives

The key topics addressed in this paper

1. ### Understanding the quantum threat

   Although large-scale quantum computers capable of breaking cryptographic systems do not exist yet, it is critical to address developments that financial institutions should know to assess the urgency of proactive preparations and the likely timeline for the preparation to be enacted.

2. ### Exploring quantum-resistant cryptographic alternatives

   This section will delve into available alternatives that are designed to be resistant to quantum attacks. It will discuss standard algorithms, the role of government agencies, standardisation bodies, and regulatory frameworks in pushing for adoption, and how much progress has been made in promoting these alternatives for widespread use in the financial sector.

3. ### Mandates and regulations around the world on quantum-safe migration

   We focus on directives and mandates by national or regional cybersecurity authorities (e.g. NIST) on approaches and timeline to migrate. Industry-specific initiatives, especially those concerning the financial service industry, will also be mentioned.

4. ### Evaluating quantum-resistant cryptographic readiness

   Are the alternative quantum-resistant cryptographic schemes ready for deployment in today's financial applications, infrastructure, and broader digital ecosystems? This section will provide an in-depth analysis of the maturity and applicability of quantum-safe algorithms and their compatibility with current systems.

5. ### Migration pathways to quantum-safe cryptography

   How should financial institutions go about migrating their existing systems to quantum-safe cryptographic standards? This section will outline the steps that organizations should take, from evaluating current infrastructure to deploying new tools and resources.

PART I

Understanding the quantum threat

Quantum computing has long been hyped as the next big thing in technology. Promises to range from solving climate change to revolutionizing the pharmaceutical industry. The potential of quantum computing is starting to feel more real. While the short-term hype might well be overblown, the long-term potential cannot be ignored or overlooked. If a risk cannot be avoided, then decisions must be made regarding mitigations, revolving around when and how much to spend. The quantum threat is to communication infrastructures, and like all infrastructures, change cannot be enacted quickly. The timescale to mitigating against quantum computers threat will be years not months.

One pressing concern is the impact on public-key cryptography, the system that underpins nearly all our secure digital communications, from banking transactions to encrypted emails. The security of public-key cryptography relies on mathematical problems, such as factoring large numbers or solving discrete logarithms, that classical, that is, digital non-quantum, computers find extremely difficult to do. This difficulty forms the basis of digital security. When large-scale quantum computers arrive, this entire foundation could be swept away.

The theoretical breakthrough to enable this revolution came in 1994 when *Peter Shor* announced algorithms that can efficiently solve these difficult mathematical problems using quantum computers. Shor's Algorithm, in theory, allows quantum computers to break cryptographic algorithms like RSA and ECC in a fraction of the time it would take classical computers. To put it in perspective: while today's best classical algorithms for factoring would take a single supercomputer millions of years to break RSA-2048, a sufficiently powerful quantum computer could do the job in just hours. It was estimated by *Gidney+Ekerå* in 2019, provided that some assumptions are met, that a 2048-bit RSA key could be recovered by an attacker in just 8 hours using a quantum computer with around 20 million qubits. Gidney 2025 has published a revised estimate using 900,000 qubits for 4.63 days to achieve the same result. Both numbers of qubits, however, are huge when contrasted with the largest publicly announced quantum computers of today, which have a few hundred physical qubits at most, with no clear route to scaling to 900,000 qubits never mind 20 million qubits.

Using the analogy of a meteor on course to strike the earth, we know that a CRQC is far away. The closing speed, however, is uncertain. A CRQC is no closer than 10 years away and more likely to be 20 years away at least.

Quantum computers attack two cryptographic services, namely digital signatures and the management of encryption keys. The latter is primarily used for confidentiality.

A reasonable strategy to counter the quantum threat to digital signatures is to set alarms to identify a latest action date. One can, for example, define that a quantum computer with ten thousand qubits would indicate 10 years to CRQC. A successful construction of such a machine would then trigger the execution of a pre-prepared 10 year-long action plan.

Managing the threat to confidentiality is more nuanced. It requires entities to assess the time value of confidential data versus the likelihood that an adversary could commercially deploy a quantum computer to recover historic data. The next section considers such a scenario. The key observation here is that organizations cannot afford to take a reactive approach and must at least build a plan, especially important when there is not much clarity on how the threat would evolve. Planning is essential and can be undertaken now. Furthermore, adoption of low-cost mitigations should be prioritised as soon as practical.

Governments and private companies are pouring billions of dollars into quantum research. While we are not yet at the point where quantum computers can break cryptography, there is evidence to heed expert warnings that we are getting closer.

As Steve Brierley put it…

"The short-term hype is a bit high, but the long-term hype is nowhere near enough."

Quoted in an article on quantum computing spotlight *The Race to Find Quantum Computing's Sweet Spot* written by Michael Brooks in *Nature 25th May 2023* (pp. S1 to S3 DOI:10.1038/d41586-023-01692-9).

The danger is not immediate, and the timeline is unclear, causing uncertainties in the security community. Organizations cannot afford to take a reactive approach here, especially when there is still a lack of clarity on how the threat would evolve. In any case, we believe the warning signs are clear, and it is high time we shore our defences up.

What makes this even more concerning is the fact that current public-key infrastructure (PKI) is everywhere. It secures the internet and, hence, everything that flows through it. If our PKI is broken, then the consequences could be catastrophic. Sensitive information would be exposed, financial systems compromised, and the digital backbone of entire industries undermined. One threat to the security of today's digital systems is closely tied to the anticipated arrival of large-scale quantum computers. It is often called the *Harvest Now, Decrypt Later (HNDL)* attack.

Harvest Now Decrypt Later (HNDL)

The HNDL attack paradigm involves malicious actors intercepting and storing sensitive encrypted communications where the symmetric encryption keys were distributed by using the RSA or ECC key exchange mechanism, *e.g.* by using classical Transport Layer Security (TLS). The intent of HNDL is to break the RSA/ECC part in the future using a CRQC to reveal the symmetric secret keys and then, decrypt the data. HNDL attacks are not new, and industries have been dealing with this possibility ever since single Data Encryption System (DES) and RSA-1024 or weaker key exchange mechanisms were widely deployed. All data which was historically encrypted with a single length DES can now be cheaply decrypted by attackers had they bothered to intercept and store it. Data encrypted by a symmetric encryption key which was protected by an RSA-1024-bit key could similarly be feasibly attacked today by classical methods.

These historic HNDL opportunities are probably being exploited by governments, but it does not seem that significant. Indeed, hardly any corporate damage has resulted from the decryption of historic data encrypted by using single-length DES. Notwithstanding this prior HNDL history, active decisions are now required for data with long-term shelf life such as medical records, state secrets, personally identifiable information (PII), property records, and investment holdings in the context of the quantum threat. Large organizations that retain vast amounts of sensitive data need to be mindful of the risks.

Official documents, including the U.S. Quantum Computing Cybersecurity Preparedness Act of 2022 and the Netherland's PQC Migration Handbook, frequently cite HNDL as a critical threat to national security and long-term data confidentiality. Numerous whitepapers, opinion pieces, and cybersecurity experts have also echoed the dangers posed by HNDL. While there are dissenting voices questioning the practicality or existence of such attacks, the consensus highlights the risk quantum computers present to encrypted data stored today.

According to the FS-ISAC PQC Working Group's 2023 report, high-profile data breaches in the financial sector have resulted in the theft of encrypted data. The report warns that, once sufficiently powerful quantum computers become available, any RSA/ECC mediated key establishment could be broken to reveal the symmetric encryption key, enabling the decryption of this stolen data, compromising customer confidentiality and security. It must be noted that, if the symmetric data encryption key was not exchanged using RSA or ECC, then the data would remain safe since known attacks by quantum computers cannot effectively break symmetric algorithms.

Despite widespread concerns about HNDL attacks, their relevance to specific sectors, such as financial services, vary and is context dependent. Let us pick EMVCo-based smart card transactions as an example. The impact of HNDL, be it in a contact or wireless

environment, is inconsequential. During a payment at a retail store, the transaction between the card and the terminal is secured by dynamic cryptograms, which are unique to each transaction and become invalid after use. Even if an attacker managed to intercept the communication during this brief interaction using advanced equipment, the captured data quickly loses value because the cryptogram cannot be reused for future payments.

Additionally, regardless of whether wireless payments are done using the contact or the near-field communication (NFC) protocol, harvesting requires physical proximity. Doing this is possible in theory but incurs a confounding complication for large-scale data collection attempts. Extending the EMV transaction example further, RSA is not used to establish symmetric keys and, hence, HNDL cannot be applied to recover symmetric keys to break a number of subsequent encryptions.

Clearly HNDL is more relevant in scenarios where transaction data is stored long-term as opposed to used transiently. Institutions need to perform a risk analysis in order to select and prioritise quantum-safe encryption mitigations for back-end systems to protect assets such as credit histories, investment records, and loan documents, which are often retained for decades. This means ensuring data travels along data routes that are physically protected or protected by symmetric encryption keys that have been set up by a manual technique or by RSA/ECC techniques that are not available for attackers to intercept. OpenSSL 3.5 standards now support a quantum resistant hybrid KEM, as do some variants of TLS 1.3, these provide cryptographic protection against the quantum threat.

HNDL cannot be ignored. A crucial part, therefore, of preparing for quantum computer attacks would be to carry out a threat analysis of HNDL now. For a start, organisations need to understand and assess the time value of their organisational data as of now and as it ages. This is a valuable exercise in its own right, independent of the specific threat from quantum computers.

The second step in the analysis would be to understand the threat actors who could act against the organisation, along with their motivations and their resource capabilities.

FOR EXAMPLE:

WHO? Criminals, governments, or both.

WHY? For monetary gain or strategic dominance.

HOW? What resources they might wish to bring to bear.

It has been claimed that governments do harvest and store colossal amounts of encrypted (and plaintext) data in the hope of one day it will become intelligible to them in some way. As possible evidence of this governmental hunger for data, in a 2016 incident reported in a paper by Demchak and Shavitt, internet traffic originating from Canada and intended for South Korea was mysteriously rerouted through China on several occasions, raising concerns about interception and long-term data harvesting. Other incidents have been reported in the OECD Digital Economy Papers No. 330 on Routing Security dated October 2022.

The business case attractiveness to the threat actors of obtaining data today, next year, and in 20 years' time is crucial in eventually arriving at a risk score regarding HNDL. Let us put ourselves in the shoes of would be HNDL attackers. We are presented with a selection problem, namely, which data intercepts should we keep?

By definition this data is encrypted and cannot be evaluated.
For any '***harvested communication that can turn valuable in the future***', there are myriads of communications that are useless. The default behaviour, assuming ignorance as to which communications would be valuable, is to collect and store everything.

The immense resources required to indiscriminately store, and curate petabytes of encrypted traffic are so huge, they would be beyond a criminal enterprise. This leads many to presume that it would be primarily nations with long-term strategic espionage and intelligence goals that are both motivated and capable of carrying out effective HNDL attacks. It is therefore very plausible that HNDL is primarily a governmental issue instead of a viable criminal strategy having very short term and sharp "*business*" objectives, say, requiring a return on investment multiplier of at least ten times and payback within a year. A strategy that relies on criminal access to a quantum computer in 20 years' time simply will not meet these criminal business requirements. Storing and curating huge amounts of data for decades, because of the selection problem mentioned earlier, will be an expensive upfront cost. This makes it very unattractive from a discounted cash point of view.

When the day the criminals have access to a CRQC comes, which constitutes a problem in itself, more costs would have to be borne. Breaking an RSA or ECC key using a quantum computer will incur significant cost per break, roughly estimated to be at least thousands of dollars per key and possibly millions in power consumption. If we combine this logic with an assessment made by Adi Shamir at The Cryptographers' Panel at RSA Conference 2023 that 99% of all encrypted messages are junk, we can deduce a hundred breaks with the associated wasted cost will result in one message that was once certainly interesting being decrypted. The message itself stands a good chance of already being "timed out" upon successful decryption. When assessing the probability and, hence, the risk of criminally motivated HNDL, organisations need to make judgements on criminal rationality.

Summarising the points made thus far, data hungry criminals can choose between a highly expensive and speculative route to commit remunerative crimes in 20 years' time by relying on access to a technology where access will be strictly controlled or by bribing someone today or by buying a zero day vulnerability for, say, one million dollars to infect thousands of companies and achieving a return today.

If an organisation is a nation state target, the sad hypothesis would be that penetration has occurred and the best damage limitation mitigations would at least be to exercise timely patch management, rigorous background checks on employees, and airgap its IT systems.

Quantum resource estimation

When will there be quantum computers capable of destroying our current public key infrastructure? More concretely, what timeline can we reasonably infer for a quantum computer capable of breaking, say, RSA-2048 or ECDSA based on curve P-256 to become available? Estimates widely vary, depending on whom we ask and what their interests in quantum computing are.

Given the information available in the open literature, we can conservatively infer that current quantum computers in existence have no more than a few hundred or one thousand or so physical qubits under control. In their December 2020 roadmap, IonQ declared that by 2028 it would have quantum computers of 1024 algorithmic qubits, which is defined as the largest number of effectively perfect qubits for a typical quantum program. The updated picture on IonQ quantum computers, as of March 2025, stands at commercially available **36 physical qubit** with **2 quantum gate fidelity at 99.6%**. Information on the development and expected number of qubits from IBM Quantum can be publicly followed online at https://www.ibm.com/quantum/technology.

By 2029, for example, IBM hopes to have a 200-qubit processor with modular error-control capabilities. Despite different technological approaches tested by makers of quantum computers, the numbers of qubits that we are seeing in public releases are quite similar, give or take some small factors.

Let us consider the state-of-the-art for quantum resource estimation to implement Shor's algorithm, which breaks ECC and RSA, and Grover's algorithm, which, in theory, could reduce attack costs on symmetric algorithms. In practice, however, Grover's algorithm is currently still 'outperformed' by classical computers. Most symmetric cryptosystems' immunity from best-known classical attack points to immunity to attack using Grover's algorithm as the next table indicates. Taking practical considerations into account, there are significant challenges to implementing these quantum algorithms. To be cryptographically relevant, the algorithms require a large number of logical qubits. Fault-tolerant computation necessitates quantum error control, which introduces significant overhead in both the number of physical qubits and the runtime.

Somewhat outdated but one of the clearest technical summaries is given in Table 4.1 of National Academies of Sciences, Engineering, and Medicine. 2019. *Quantum Computing: Progress and Prospects*. https://doi.org/10.17226/25196. We reproduce the table here for convenience.

TABLE 4.1 Literature-Reported Estimates of Quantum Resilience for Current Cryptosystems, under Various Assumptions of Error Rates and Error-Correcting Codes

Cryptosystem	Category	Key Size	Security Parameter	Quantum Algorithm Expected to Defeat Cryptosystem	# Logical Qubits Required	# Physical Qubits Required[a]	Time Required to Break System[b]	Quantum-Resilient Replacement Strategies
AES-GCM[c]	Symmetric encryption	128 192 256	128 192 256	Grover's algorithm	2,953 4,449 6,681	4.61×10^6 1.68×10^7 3.36×10^7	2.61×10^{12} years 1.97×10^{22} years 2.29×10^{32} years	
RSA[d]	Asymmetric encryption	1024 2048 4096	80 112 128	Shor's algorithm	2,050 4,098 8,194	8.05×10^6 8.56×10^6 1.12×10^7	3.58 hours 28.63 hours 229 hours	Move to NIST-selected PQC algorithm when available
ECC Discrete-log problem[e-g]	Asymmetric encryption	256 384 521	128 192 256	Shor's algorithm	2,330 3,484 4,719	8.56×10^6 9.05×10^6 1.13×10^6	10.5 hours 37.67 hours 55 hours	Move to NIST-selected PQC algorithm when available
SHA256[h]	Bitcoin mining	N/A	72	Grover's Algorithm	2,403	2.23×10^6	1.8×10^4 years	
PBKDF2 with 10,000 iterations[i]	Password hashing	N/A	66	Grover's algorithm	2,403	2.23×10^6	2.3×10^7 years	Move away from password-based authentication

Since 2019, there have been some progress in reducing the number of physical or logical qubits required to break ECC discrete log problem based on specific popular curves, but the improvement does not yet alter the general picture significantly enough. Hence, the general recommendation of retaining the use of currently recommended key lengths for

symmetric-key schemes and begin transitioning to the new standards for asymmetric-key schemes, selecting suitable key lengths and recommended security components for the use cases, remains valid.

An update has been recently given by *V. Gheorghiu* and *M. Mosca* in Quantum resource estimation for large scale quantum algorithms (DOI: https://doi.org/10.1016/j.future.2024.107480).

They analyse the security of symmetric schemes and hash functions against quantum adversaries. The assumed error control relies on the quantum surface codes and braiding techniques. We note that surface codes form a subfamily of the more general and powerful quantum stabilizer codes. The former type of codes is currently the most implementable. The general picture may change significantly if we can implement concatenated stabilizer codes down to the desired fault-tolerant layer. The quantum security parameters, based on the assumptions of using state-of-the-art algorithms and fault-tolerance methods, for symmetric and hash-based cryptographic schemes are summarized in Table 1 of the above-mentioned paper relating to the costs of a Grover attack on AES.

Table 1
Quantum security parameter (qs) for the AES family of ciphers, SHA family of hash functions, and Bitcoin, assuming a conservative physical error rate per gate $p_g = 10^{-4}$.

Name	qs
AES-128	106
AES-192	139
AES-256	172
SHA-256	166
SHA3–256	167
Bitcoin's PoW	75

Grover's algorithm asymptotically square roots the difficulty of a classical unstructured search. Naively, this means halving the security level of some symmetric cryptosystems. Gheorgiu and Mosca' s paper, however, illustrates that a search of size 2^{128} classical trials does not reduce to 2^{64} quantum trials, which is the asymptotic complexity, but rather to 2^{106} quantum trials. Highlighting the authors' assertion that '*the constants in the complexity matter*', the consequence of this analysis is that to perform the 2^{106} quantum trials that break one AES-128 key in one year would require 2^{80} quantum computers working in parallel, which is simply impossible.

Specific to widely deployed elliptic curve cryptosystems, the paper supplies their respective space-time trade-offs. Here is an example on the popular P256 curve. To break ECC based on this curve in roughly 24 hours requires approximately 67.7 million physical qubits. The estimate here is larger than the one presented by a team from Microsoft Research in Asiacrypt 2017 (see https://arxiv.org/abs/1706.06752) in which the overhead for error control was not considered.

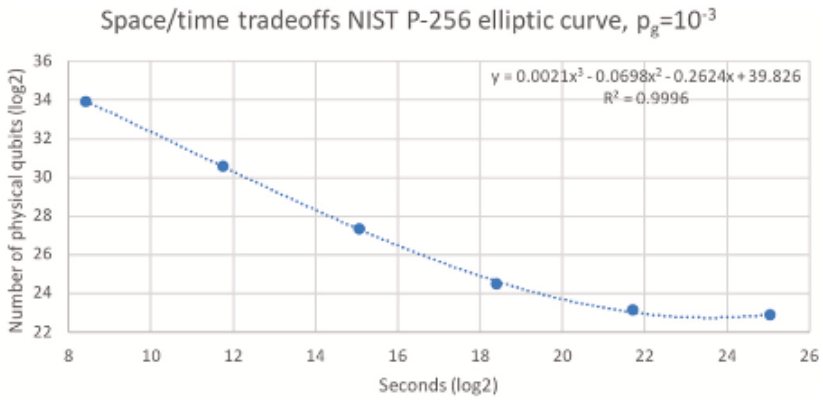

Fig. 49. NIST P-256 elliptic curve space/time tradeoffs with physical error rate per gate $p_g = 10^{-3}$. The scale is logarithmic (base 2). Approximately $y(16.3987) \approx 6.77 \times 10^7$ physical qubits are required to break the scheme in one day (24 h). The number of T gates in the circuit is 8.82×10^{11}, the corresponding number of logical qubits is 2330, and the total number of surface code cycles is 1.72×10^{14}. The classical security parameter is 128 bits.

The next table provides a snapshot on popular ECCs. One may be tempted to buy time by moving to curves P-384 and P-521 in the interim. It is reasonable, however, to assume that once we can scale up to several million qubits from the present hundred or thousand qubits, going to tens of million qubits would be easier engineering-wise.

Table 2
The total physical footprint (nq) required to break the ECC schemes in 24 h, together with the required number of T gates (Tc), the corresponding number of surface code cycles (scc), and the corresponding classical security parameter (s). We assume a very conservative physical error rate per gate $p_g = 10^{-3}$, more likely to be achievable by the first generations of fault-tolerant quantum computers.

Name	nq	Tc	scc	s
P-160	1.81×10^7	2.08×10^{11}	4.05×10^{13}	80
P-192	3.37×10^7	3.71×10^{11}	7.23×10^{13}	96
P-224	4.91×10^7	5.90×10^{11}	1.15×10^{14}	112
P-256	6.77×10^7	8.82×10^{11}	1.72×10^{14}	128
P-384	2.27×10^8	3.16×10^{12}	6.17×10^{14}	192
P-521	6.06×10^8	7.92×10^{12}	1.56×10^{15}	260

For Shor's algorithm on various RSA modules, the work of Gidney and Ekera in **https://arxiv.org/pdf/1905.09749.pdf has been superseded by Gidney** https://arxiv.org/abs/2505.15917.

The efficiency of Shor's algorithm over the current best classical factorisation algorithm, which is the General Number Field Sieve (GNFS), is spectacular. For more on the GNFS,

one can consult an excellent exposition by Carl Pomerance titled *A Tale of Two Sieves* in the December 1996 edition of Notices of the American Mathematical Society. Factoring a 2048-bit in the RSA cryptosystem would take one classical supercomputer thousands of years. Gidney and Ekera stated that one needs roughly 20 million physical qubits, provided that the error rate can be kept at 10-3 to factor 2048 in 8 hours, the new estimate by Gidney reduces the spacetime volume by approximately 35%, using the same assumptions as in the 2019 paper the new resource estimate is 900,000 noisy qubits for 4.63 days as opposed to 20 million qubits for 8 hours. Adjusting the error rate would impact the other design parameters and, hence, the performance. Here lies the catch. Realising quantum fault tolerance is hard. Quantum bits are inherently noisy and come with differing physical characteristics. Factoring or finding discrete logarithm using Shor's algorithm surely require many physical qubit interactions, which bring us to a large quantum layout problem (QLP) and back to quantum error-control.

Other significant references for resource estimations include the energy requirement estimates undertaken by Parker and Vermeer in

https://arxiv.org/pdf/2304.14344.pdf in this paper and a high-level cryptanalysis comparison of Ekera and Gartner in https://arxiv.org/pdf/2405.14381.

On the theoretical side, there have been algorithmic breakthroughs that claim to reduce the physical to logical qubit ratio by a few orders of magnitude. A recent report with a 2 order of magnitude reduction can be found in Bravyi, S., Cross, A.W., Gambetta, J.M. *et al.* High-threshold and low-overhead fault-tolerant quantum memory. *Nature* **627**, 778–782 (2024). While more scrutiny is required to validate this claim, acceleration in the reduction ratio can drastically alter the landscape by shortening time-to-quantum and, hence, heightening the urgency to migrate.

In a paper published in Nature in August 2024, the Google team presented a quantum processor called Willow. It doubled the number of qubits from 53 reported in 2019 to 105 in 2024. The coherence time increases to 1 second. Given the challenging engineering context, this improvement is fantastic. Seen as a part of the big picture, however, this is not something to get too excited about. The fact that it has taken this long to reach this stage highlights the difficulty.

A rough extrapolation reveals that, at the current rate of progress, Google in five years has also doubled the number of gate operations performed from 20 to 40. Since approximately 20 million qubits are required to break RSA-2048, at the current rate of progress we need about 90 years to get there, ignoring the scaling limit of Google's

current technology. Assuming the need for 2.1 billion gate operations to break RSA 2048, we will reach the target in about 125 years. A single physical qubit is stable for 30-60 millionths of a second, inferred from 49 qubits performing like 1 logical qubit. To break RSA 2048, qubits must be stable for 8 hours, with 20 million qubits that is, 28,800 seconds or stable for 400,000 seconds with 900,000 qubits.

Sam Jaques of University of Waterloo has charted a landscape of quantum computing, emphasizing the connection of number of qubits and error rate. The chart in its latest update for 2024 (https://sam-jaques.appspot.com/quantum_landscape_2024) is reproduced here for convenience. One needs to keep in mind that the chart is drawn in log-log scale in visualizing the gap between where we are and breaking RSA. The limited progress from 2021 to 2024 points to the enormous challenges in this field and the large gap to utility.

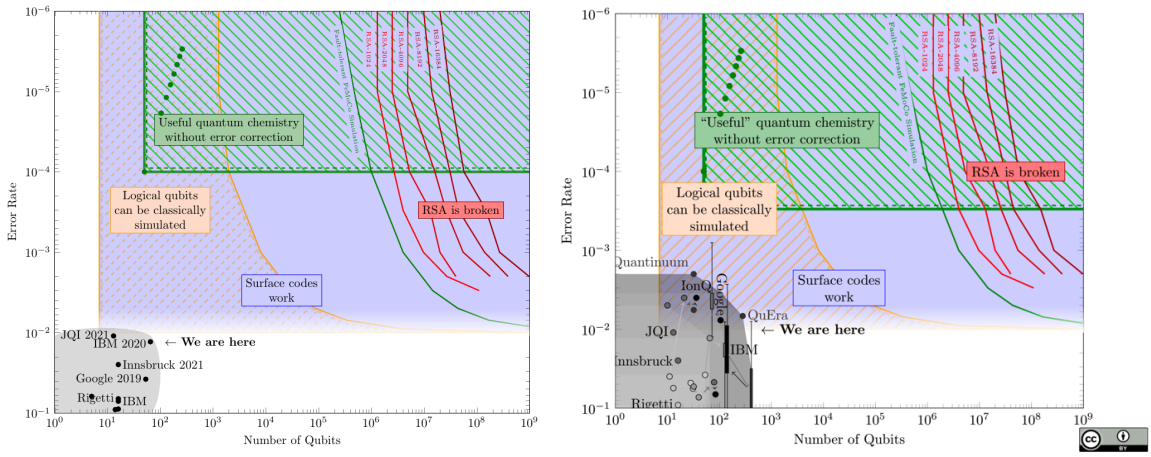

Quantum computers need excellent support systems, in particular, a cooling system to ensure normal operation. In superconducting quantum computing, dilution refrigerators cool all components that control and measure the system and the states of the qubits. Scaling up the number of qubits requires cooling equipment that can be mass produced and deployed modularly. This is another massive engineering challenge. Setting aside the costs, it is not so clear by how many times bigger the refrigerators need to get before we see significant leaps as this depends on their modularity and integration overhead. A recent study on the status of quantum computer development released by BSI Germany concludes that quantum computing is "steadily progressing towards cryptanalytic relevance". Improved fault-tolerant execution has been achieved with surface coding on superconducting systems and with colour coding on ion-based systems. The study states that, conservatively, cryptanalytic relevance will **not** be achieved within the next 16 years, that is by around 2040, unless major leaps occur before then.

Estimating the strength of current cryptographic schemes to resist realistic quantum attacks is a moving target. It depends on a variety of parameters. These include, among others, fault-tolerant quantum error correction, novel circuit optimization techniques at the logical level, new cryptanalysis tools, and improved quantum algorithms. Monitoring this space should be a priority and stresses the importance of preparing for migration to quantum-resistant systems.

In summary, our current quantum technologies as of 2025 are far from being cryptographically relevant. Quantum computers are difficult to build. On the other hand, it is hard to predict the future. Breakthroughs may be imminent. Governments and large enterprises are supporting many top researchers and engineers on the quest to build large-scale quantum computers. Improvements continue to appear in the open literature. Another interesting development is the integration of quantum computers and classical supercomputers for pre- or post-processing of computational tasks carried out in the quantum processors. IBM, for example, has put this agenda forward and called it quantum-centric supercomputing in this article.
https://www.ibm.com/quantum/blog/supercomputing-24

It is likely that some components of quantum technologies may or will be treated as trade secrets or matters of national security. Therefore, significant developments may well be kept confidential. It seems unlikely, however, that classified research is far and qualitatively ahead, given that some of the largest commercial players would need to justify their spending and keep their shareholders happy, for example by announcing milestones almost as soon as they have been vetted.

We hold that the new post-quantum algorithms and associated protocols are an improvement on today's practices which, however, remain entirely fit for current purposes. ECC, in particular, is very performant indeed. Even if cryptographically relevant quantum computers turn out to be impossible to engineer over time, these new protocols, especially their hybrid variants, may become the norm best practice and any performance penalties would become sustainable.

PART II

Transition to quantum safe systems

Risk management

It is often argued that "***Risk can never be eliminated, only managed***". Assets in the form of information and its systems face continuous risks from unauthorized disclosure, alteration, as well as manipulation of all kinds. Data at rest can be protected by a combination of robust physical security and logical cryptographic security. Since it is harder to protect data in transit with robust physical security, protection naturally relies more heavily on cryptography. It is even harder to comprehensively protect data in use. Most users prefer convenience and flexibility over stringent security measures.

There are many risk analysis models, including *MITRE, STRIDE, PASTA, DREAD*, and the HNDL analysis above. They typically take as input the probability of some negative events occurring. The models then assess if a malicious agent is engineering the events, typically based on some assumptions on the rewards that the malicious agent stands to gain.

Once cryptographically relevant quantum computers become available, one must assume their availability to the adversaries. For many governments, risks of attacks against critical information systems and data should be mitigated even if the costs could be a significant fraction of their GDPs. National security reasons often take precedence over economic considerations, particularly for governmental agencies with large resources. In the commercial world, estimating how much investment in technological protections as a form of insurance an enterprise would be willing to take against quantum risks may be more nuanced. Many businesses would readily accept some risks that would be unacceptable to major government bodies. In short, risk appetite dictates the insurance costs one is willing to accept. The timing to take the plunge can be markedly different across jurisdictions and industries. The same holds for the timeline and priorities to migrate digital assets.

Broadly speaking, there are two prominent choices for symmetric key establishment that offer security against quantum attacks. These are **quantum key distribution** (QKD) and **post-quantum cryptography** (PQC).

QKD is a physical approach that securely generates and distributes random bits as symmetric encryption keys between two parties, who are commonly referred to as "Alice" and "Bob", typically by using quantum optics. When certain assumptions are met, QKD comes with an information-theoretic security guarantee. It cannot be broken, even by quantum computers. Recent advances in QKD have resulted in improved key generation

rates, long-distance quantum communications, and even satellite-based network deployments. While QKD requires specialized quantum infrastructure, there is much active research to enable its use globally.

On the algorithmic side, PQC offers a more readily deployable solution for most organizations. Here, the term PQC refers to new public-key cryptographic protocols that are designed to be secure against quantum attacks but can run on classical computers. In 2016, the National Institute of Standards and Technology (NIST) initiated a competition to standardise quantum-resistant algorithms for key exchange mechanism (KEM) and digital signature algorithm (DSA). In July 2022, NIST announced the first batch of the winning algorithms, and on 13 August 2024, it published three standard PQC algorithms after an extensive, eight-year public review. These are ML-KEM (FIPS 203) for quantum resistant key-exchange, based on CRYSTALS-Kyber, and two digital signature schemes, namely ML-DSA (FIPS 204) and SLH-DSA (FIPS 205). Another scheme named FN-DSA, originally known as Falcon, has also been chosen to be standardised. The final version detailing its standard specifications is expected to be published soon. In March 2025, NIST made public the choice of standardising HQC (Hamming Quasi-Cyclic) as the second KEM. The official standard and technical specifications are expected to be final by early 2027.

We provide a brief overview below of PQC and QKD to understand their respective roles.

PQC as a quantum-safe alternative

Unlike traditional cryptographic algorithms, PQC algorithms are developed to resist attacks by both large-scale quantum computers and classical ones. They explore various mathematical problems that are believed to be practically impossible to solve even by quantum computers to use as a security foundation. These problems may come, for instance, from lattices, algebraic codes, and multivariate quadratic polynomials. They give rise to new quantum-safe public key cryptosystems. Some of them require larger key and signature sizes compared to RSA and ECC. Fortunately, a good number of quantum-safe algorithms perform competitively or even better when the parameters are chosen judiciously and implemented cleverly.

SOME BRANCHES OF PQC ALGORITHMS INCLUDE

- ### Lattice-based cryptography

  Lattice-based cryptography has become the most promising approach that offers high performance and quantum resistance. Their security relies on two main problems that are provably hard. These are the shortest vector problem (SVP) and the learning with errors (LWE) problem. As the name suggests, SVP involves finding the shortest non-zero vector in an algebraic lattice. In LWE, one mathematically hides the true structure of the keys through linear transformation and further scrambles the message by deliberately adding errors. To boost efficiency, the polynomial ring variant of LWE, abbreviated to R-LWE, is often used. Unlike factorization or discrete logarithm problems, no known quantum algorithm can solve SVP or LWE efficiently, making them an excellent foundation for post-quantum cryptography. CRYSTALS-Kyber and CRYSTALS-Dilithium, the names of the original proposals that eventually evolved, respectively, into ML-KEM and ML-DSA, are R-LWE-based schemes that provides excellent balance between security, performance, and key sizes. The third lattice-based scheme which has been selected to be standardised is Falcon. It provides smaller signatures with higher efficiency, ensuring strong quantum resistance for authentication and integrity.

- ### Code-based cryptography

  Code-based cryptography leverages error-correcting codes, such as binary Goppa codes or codes with certain cyclic properties. Such codes, originally designed for information fidelity, can secure communication by keeping decoding functions secret while sharing only disguised encoding functions. This protects the plaintext by mapping it to a scrambled codeword that can only be decoded with the secret function. The core security foundation relies on the hard problem of syndrome decoding on random-looking codes, making it resistant to both classical and quantum attacks. Introduced in the McEliece cryptosystem in 1978, this approach offers fast encryption but requires large key sizes. Despite this drawback, its long-standing resistance to cryptanalysis has earned it significant trust. Three code-based schemes, namely Classical MCEliece, BIKE, and HQC, were the candidates that advanced to NIST's fourth round of PQC standardisation for KEMs. HQC has very recently been chosen to be standardised. The other two candidates have been ruled out of contention.

For digital signing, code-based candidates face limitations due to inefficient signing, large keys, and cryptanalytic vulnerabilities, making them less favourable for standardisation in terms of performance. Two code-based schemes, namely CROSS and LESS, made it to the second round of the additional digital signature PQC standardisation process.

- Hash-based cryptography

  Hash-based cryptography provides quantum-safe signature schemes that rely on the collision-resistance of hash functions. Earlier one-time signature schemes faced usability challenges, prompting the introduction of Merkle trees to generate multiple signatures from a single key. Modern schemes like eXtended Merkle Signature Scheme (XMSS) and Leighton-Micali Signature Scheme (LMS) have improved efficiency, reduced key sizes, and added forward secrecy. Both are now NIST-approved standards as specified in NIST SP 800-208. They are also recognized as ISO standards ISO/IEC 14888-4:2022 for digital signatures, affirming their global acceptance.

  The main strength of hash-based cryptography is its algorithmic agility. If a hash function becomes insecure, switching to a secure one remedies the issue. However, traditional schemes like XMSS and LMS are stateful, requiring careful tracking of key usage in practice. To address this, SPHINCS+, a stateless hash-based signature scheme, has been standardised as SLH-DSA (FIPS 205). Its design is simple and assumption-free. It offers strong long-term security, albeit having larger signature sizes. Despite this trade-off, hash-based cryptography remains a trusted and resilient option for post-quantum security.

There are at least two other mathematical domains that can provide hard problems for quantum computers to solve. **Multivariate and Isogeny-based schemes** are also being considered in the standardisation process. Proposed KEM candidates from them, however, have not gained sufficient favour from the security community beyond the realm of academic research. In terms of digital signature, one isogeny-based candidate and four multivariate-based candidates have made it to the second round of additional PQC standardisation process.

PQC offers a significant advantage over QKD. The former does not typically require an extensive upgrade to existing hardware infrastructures. PQC can deliver quantum-resistant security on classical communication channels, making it compatible with a wide range of devices, from low-cost microcontrollers to high-performance servers and dedicated hardware security modules (HSMs).

PQC's versatility to provide both confidentialities, for instance, as defined in FIPS 203, and integrity/authenticity, for examples, as specified in FIPS 204 and 205, turns PQC into a direct replacement for RSA and ECC. PQC can in fact be used to enhance the security of QKD setups by providing additional authentication and integrity guarantees for the classical communication channel.

As PQC standards are now available, there has been a heightened push for migrating a range of applications to support PQC. While National Security Agency of USA requires federal agencies to migrate to NIST-standard PQC algorithms within the designated timeline, European agencies like British Standards Institute (BSI) and Agence Nationale

de la Sécurité des Systèmes d'Information (ANSSI) promote hybrid use of classical cryptography and PQC. The share of PQC support in internet communications has been increasing rapidly. By March 2024 nearly two percent of connections with Cloudflare over TLS 1.3 were secured with PQC. The ratio hit double digits percentile by the end of 2024. Popular communication platforms, such as Apple *iMessage*, *Signal*, and *Zoom*, already support PQC. Web browsers such as Google *Chrome* began supporting hybrid X25519+Kyber for most outbound connections in 2024. Now with the standards officially out, we can expect wider and more uniform adoption across the internet.

PQC has its fair share of challenges, some of which we now outline.

- Computational complexity and overheads

  PQC algorithms can be computationally more intensive than RSA and ECC, leading to a performance degradation in certain setups. This is evident in real time applications with strict latency deadlines. The larger key sizes also contribute to the overhead. More memory is required to store the keys, spanning several kilobytes to megabytes, which is challenging in low-end embedded devices. Further, transmitting large keys can drain the battery of wireless devices faster as the radio frequency components typically remain active for extended periods. The ephemeral nature of PQC, where new key pairs are generated for each session, further contributes to the overhead.

- Future quantum attacks

  The PQC standards are built on algorithms that have undergone extensive evaluation and scrutiny in terms of security. Although PQC algorithms are designed to be resilient against known quantum attacks, this may not hold against future quantum algorithms. As a precautionary measure, NIST is still in the process of standardising alternative PQC algorithms based on different mathematical foundations as backup options and for variety.

- QKD as a quantum-safe alternative

  Before the widespread adoption of public key cryptography, point-to-point solutions such as couriers carrying encrypted messages or tamper-resistant envelopes were the main methods for secure key exchange. They were slow and cumbersome. QKD represents a modern take on the point-to-point approach, offering security based on the principles of quantum mechanics.

  QKD is a mechanism for secure symmetric key establishment, it does not provide a digital signature mechanism.

  A QKD protocol does **not** require a quantum computer while offering a level of security that is qualitatively different from conventional cryptographic schemes. It provides **information-theoretic security** instead of computational security. QKD is theoretically

secure, even against adversaries with unlimited computational power, not limited to quantum computers.

The basic principle of QKD can be explained through the operation of a QKD link, which is a point-to-point connection between Alice and Bob. The link consists of a quantum channel and a classical channel. Alice generates a sequence of random bits, encodes it into non-orthogonal quantum states, and sends it down the quantum channel. Bob then measures the received quantum states to get a bit string that correlates with Alice's. They use the classical channel to check for the correlation. High correlation indicates minimal eavesdropping, allowing them to distil a shared symmetric key. The security of QKD is based on the quantum principle that measuring quantum states disturbs them. This disturbance can be detected, allowing legitimate users to measure how much information the eavesdropper Eve has gained. If the level of eavesdropping exceeds a predefined threshold, then the communication is aborted.

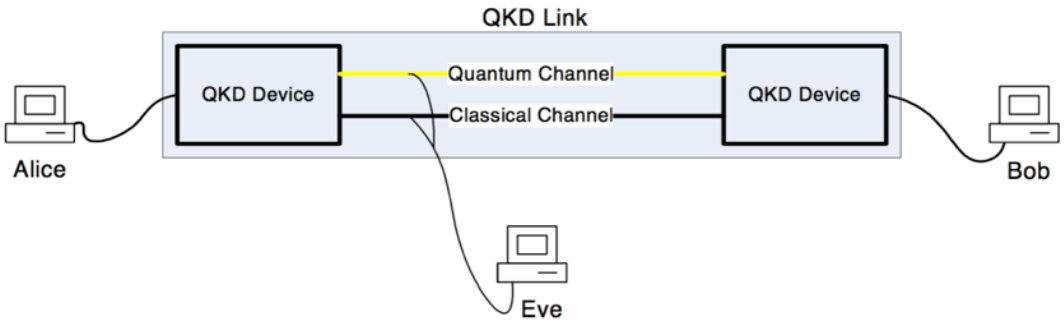

Limitations of QKD

There are notable limitations to QKD, pointing to the need for PQC for a holistic protection against quantum threats.

Security strengths and availability concerns

QKD boasts robust security against eavesdropping by making any attempt at interception detectable. This sensitivity, while safeguarding the integrity of the random strings being exchanged, introduces severe availability issues. Specifically, Eve can effectively cause a denial of service by persistently disturbing the quantum channel, triggering loops of detection and subsequent shutdown of the transmission. Moreover, differentiating between a genuine channel degradation caused by technical faults or environmental conditions and deliberate eavesdropping attempts complicates the reliability and practical deployment of QKD systems. This ambiguity can hinder the effectiveness of QKD in critical communication channels requiring near perfect reliability.

QKD still requires quantum-resistant cryptography for authentication

QKD cannot operate in isolation to guarantee unconditional security. It requires both a public quantum channel and an authenticated classical channel. Without authentication on the classical channel, man-in-the-middle attacks could easily compromise the keys being exchanged. Therefore, QKD must be technically integrated into an existing security infrastructure. Authentication can be achieved by using either a private key cryptography, via a message authentication code (MAC) with symmetric keys, or a public key cryptography, via digital signatures, which can be based on classical or post-quantum schemes. This authentication process is independent of the keys being exchanged over the quantum channel.

Transporting QKD key in a quantum secure manner is challenging

While the QKD keys themselves are generated in a quantum-safe manner, they must be transmitted to applications or devices that need to consume them. This must be done through classical communication channels, either by electrical signals through copper, optical signals through optical fiber or wireless signals electromagnetically. The classical channels need to rely on cryptographic algorithms that are themselves quantum secure. This brings us back to the essential role of PQC.

The above two issues imply that QKD systems themselves need PQC algorithms to transport the QKD keys safely to the endpoints before the keys can be used for secure

communication. Research into QKD is foundational towards realizing a quantum internet. QKD in combination with One-Time Pad (OTP) encryption not only provides a high level of security but also introduces the benefit of deniability. No unauthorised party can ever know what was communicated by linking plaintext to ciphertext in a unique way. There remains a lot of work to be done before QKD can be widely adopted in terms of scalability, performance, and costs.

While there are no formal mandates or advisories from standardisation bodies or government agencies specifically promoting QKD, significant efforts are underway to weave QKD into existing infrastructures. In April of 2023, European Telecommunications Standards Institute (ETSI) released the first **protection profile** for QKD for the purpose of common criteria assurance.

Next, we will explore various practical deployments of QKD and discuss its role in enhancing the existing security measures within these infrastructures.

Practical deployments of QKD

QKD is increasingly being deployed in specific application use-cases. A notable example is the Quantum Secure Metro Network (QSMN), a joint project by Toshiba and BT Group, which connects multiple sites in London using QKD to secure data transmitted over fiber optic cables. HSBC became the first bank to join this commercial trial, using QKD to protect sensitive information, including financial transactions. The trial currently secures communications between HSBC's global headquarters and a data centre 62 km away, using BT's infrastructure, Toshiba's quantum technology, and AWS Edge Compute Services. This practical deployment highlights how QKD can be scaled across multiple customers without significant changes to the existing network, providing a template for future quantum-secure networks.

Singapore has also launched its own quantum-safe initiative, called the National Quantum-Safe Network Plus (NQSN+). It aims at deploying QKD across the entire island nation. Building on a decade of quantum research, this network enables businesses to access quantum-safe solutions. Singtel and SPTel, in collaboration with SpeQtral, are leading efforts to establish nationwide quantum-safe networks that can integrate both QKD and PQC. Singapore's approach showcases the potential for global interoperability. These deployments are paving the way for more widespread adoption of QKD to protect data in industries ranging from finance to national security.

Augmenting security infrastructure with QKD

QKD is currently too expensive for widespread deployment due to its specialized hardware requirements. It has been deployed as an add-on option for specific high-security applications, often complementing existing infrastructure already protected with PQC. This layered approach introduces *cryptographic agility*, allowing organizations to switch between or combine different cryptographic methods. By integrating QKD where feasible and relying on PQC more broadly, government bodies and large businesses can strike a balance between cost-efficiency and quantum resistance.

In our view, QKD is best suited for scenarios where highly sensitive information requires the strongest possible encryption. Its cost and infrastructure demands may be justifiable in niche, high-security environments. Such an environment, however, typically demands high availability, which remains problematic for QKD.

Here are some use cases that have been explored:

1. **Intra bank transfers**: QKD can secure the exchange of encryption keys for intra-bank transactions involving large sums or high-value assets, providing additional protection against potential eavesdropping or man-in-the-middle attacks.
2. **Government communications**: QKD can safeguard confidential communications between government agencies, embassies, or intelligence services to prevent interception.
3. **Data centres and cloud infrastructure**: QKD can secure communication channels between critical data centres in which large volumes of sensitive information, such as financial or healthcare records, flow.
4. **Stock exchanges and financial trading platforms**: For high-frequency trading platforms or stock exchanges, QKD ensures that critical financial data is protected from being intercepted and manipulated.

In these scenarios, QKD acts as a high-security add-on, complementing the security of existing systems built with PQC. This hybrid approach introduces agility, allowing critical infrastructures to adopt QKD where justified, while relying on PQC for broader and more cost-effective protection.

PART III

Mandates and regulations around the world on quantum-safe migration

A key milestone in the transition was the official publication of three PQC standards by NIST in August 2024. (https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards)

This marked a significant step towards establishing a foundation for quantum-resistant cryptographic algorithms. With four PQC standards, three of them already final and one in advanced drafting, plus another already chosen to be standardised as well as available QKD technologies, the focus seems to be shifting towards migrating to a quantum-safe future based on regulatory mandates rather than concerns over HNDL or debates on time-to-quantum-disaster. This migration is likely to be gradual and take decades as billions of applications and devices are being upgraded. The transition to quantum-resistant cryptography is a significant undertaking. Given the vast number of software and hardware systems that need to be updated, and the diverse range of technologies involved, this migration is likely to be a lengthy and complex process.

To guide organizations through this transition, cybersecurity agencies and regulatory bodies have been providing recommendations and, in some cases, issuing mandates. These guidelines recommend planning a migration strategy that ensures service continuity and compliance while addressing economic concerns. The strongest motivation and driver towards PQC migration are security mandates. Let us look at examples of these mandates.

In the context of PQC migration, two primary strategies have emerged. These are direct and hybrid transitions. The direct transition approach replaces existing cryptographic systems with quantum-resistant algorithms. This method is straightforward but requires confidence in the maturity and security of new PQC methods. It tends to be more costly and disruptive. The hybrid transition strategy combines currently deployed cryptographic mechanisms with quantum-resistant ones. The idea is to provide layered security that benefits from the strengths of both approaches while mitigating their individual weaknesses. This allows for a more flexible and phased integration as quantum-resistant technologies continue to evolve and being evaluated. Those opting for the hybrid approach will not take the direct one.

The Quantum computing cybersecurity preparedness act mandates U.S. Federal Agencies to transition to PQC. The goals are to build proactive defence against quantum

threats and to maintain national security. Augmenting this, the NSA's Commercial National Security Algorithm Suite 2.0 (CNSA) specifies requirements and compliance timelines for their National Security Systems to adopt quantum-resistant algorithms. ML-KEM and ML-DSA are required for public key cryptography while AES-256 and SHA-384/512 are mandated for symmetric key cryptography. For software and firmware updates, XMSS and LMS are deemed sufficient. Starting in 2025, all new software and firmware must support signing with the listed algorithms. By 2030, all deployed software and firmware should be fully transitioned to the PQC algorithms. Similar timelines apply to web browsers, networking equipment, operating systems, and other components. Ultimately, the goal is for all applications to adhere to CNSA Suite 2.0 by default by 2033. In terms of security, the mandate is to use the strongest level specified in NIST official documents. While focused on the military and the intelligence communities, these mandates also inform quantum-safe practices across critical sectors, including financial institutions.

European institutions have also published recommendations. Compared with the US directives, there is more flexibility in terms of acceptable algorithms. The timeline is not as definitive. In particular, the BSI's technical guideline, given in BSI TR-02102-1 dated 31 January 2025, stresses the inevitability of quantum computers and the urgent need for post-quantum cryptography adoption, distinctively advocating for a hybrid approach that combines classical and post-quantum schemes for robust security. This strategy aims for cryptographic agility to ensure easy updates as new threats emerge. In contrast, the U.S. CNSA 2.0 guideline does not explicitly endorse a hybrid approach, focusing instead on transitioning directly to quantum-resistant algorithms. This highlights a strategic difference between the transition approaches.

The French ANSSI's strategy for post-quantum cryptography (PQC) migration leans heavily towards a hybrid approach, blending established cryptography with the new PQC algorithms. ANSSI advocates for the development of cryptographic products that smoothen transitions between standards, which deemed crucial given the evolving landscape of quantum computing. ANSSI's conservative yet proactive approach to PQC migration strives to balance immediate security needs with long-term objectives. Systems must not only meet current security demands but are also being prepared for future advancements in cryptography.

The National Cyber Security Centre (NCSC) of the UK has also issued recommendations. They suggest ML-KEM for key establishment, ML-DSA for digital signatures, and SLH-DSA, XMSS, or LMS for firmware and software signing. For AES, NCSC recommends AES-128 and SHA-256 as relatively safe options while acknowledging the need for continued vigilance and future transitions.

We have also seen industry-specific initiatives, including in telecommunication and defence technology sectors. Closely related to our industry, the European Cybercrime Centre of Europol created the Quantum Safe Financial Forum (QSFF) in 2024. The forum is an effort to address the transition to PQC across the financial sector, with Europe as the focus, to share best practices and coordinate actions. Behind the creation of the forum is an acknowledgement that migrating to a quantum-safe approach will be a complex project that will require dedicated resources. Industry peers, the public sector, and academia would benefit from coming together to identify both opportunities and challenges in advance.

In summary, the global consensus is shifting towards quantum-resistant cryptography. Other national agencies are expected to release their recommendations soon. Organizations worldwide are being urged to adopt recommended algorithms and strategies to ensure the long-term security of their digital assets in the face of emerging quantum threats.

Except for national security systems in several countries, for which the mandate is to directly adopt PQC standard algorithms with the highest security Level 5, the general agreement converges towards the hybrid approach. There are relatively minor differences in the specific recommendations in terms of the minimally required security levels. We can reasonably expect that these differences will lessen either by clearer future directives or by industry convention. We reproduce Table 1 from NIST IR 8547 on Transition to Post-Quantum Cryptography Standards for ease of reference regarding security levels.

TABLE I: Security Categories in NIST PQC standards.

Security	Attack Type	Example
Level 1	Key search on a block cipher with a 128-bit key	AES-128
Level 2	Collision search on a 256-bit hash function	SHA-256
Level 3	Key search on a block cipher with a 192-bit key	AES-192
Level 4	Collision search on a 384-bit hash function	SHA3-384
Level 5	Key search on a block cipher with a 256-bit key	AES-256

Government agencies on PQC vs QKD as quantum-safe alternatives

BSI, NCSC, Swedish Armed Forces, and the Netherland's NLNCSA carried a cautious tone on QKD. They are quite critical regarding the practicality and maturity of QKD solutions. These agencies underscore the current limitations of QKD that we have listed earlier. QKD is practical only for niche applications and does not yet offer a comprehensive

security solution at the level required for national security systems. The agencies push for PQC as a wider and ready quantum-safe alternative. The following is a quote from the conclusion of the joint position statement...

"However, due to current limitations, QKD is only practical for niche use cases and cannot replace classical key agreement schemes in most scenarios. Additionally, QKD is not yet fully mature from a security standpoint. Given the urgency to move away from quantum-vulnerable public-key cryptography, the primary focus should be on migrating to post-quantum cryptography or adopting symmetric keying solutions."

Similarly, the NSA advises that...

"... the technology (QKD) involved is of significant scientific interest, but it only addresses some security threats and it requires significant engineering modifications to NSS communications systems."

QKD is not yet seen as a practical security solution to protect national security information.

Most standardisation bodies and government agencies around the world advocate for migrating to PQC or adopting enhanced symmetric keying solutions to protect against quantum threats. This perspective aligns with the urgent need to develop quantum-safe alternatives that are more universally applicable, economical, and secure than current QKD technologies.

Quantum Security: QKD or PQC

	Security Objective	Infrastructure and Cost	Coverage	Standardization	Direct Applications
QKD	Confidentiality	*New Optical Hardware* High Cost	Distance Limit ~150KM Only Point to Point Low Bandwidth	Not backed by Standardization Bodies as Quantum Safe Solution[1]	Limited Critical Point to Point Comm. Links
PQC	Confidentiality, Integrity, Authenticity, Non-repudiation	*No Hardware Changes (Only Plugin)* *Low Cost*	No Such Limitation	NSA (US), BSI (Germany), ANSSI (France), NIST (US), NCSC (UK), ISO	Can be used to secure all devices (HW/SW)

QKD cannot be standalone, but can come as an *additional layer* of security along with PQC[1].

PQC acts as standalone solution, that provides holistic quantum security.

1. Position Paper on QKD by ANSSI, GIS, BSI, SAF: https://cyber.gouv.fr/sites/default/files/document/Quantum_Key_Distribution_Position_Paper.pdf

PART IV

Performance of PQC algorithms

A key consideration during the PQC standardisation process is performance. Benchmarking among the candidate algorithms in Round 3 greatly influenced the decision to choose the four new standard algorithms as detailed in the status report published as NIST IR 8413. Similarly in Round 4 for the decision to select HQC as explained in NIST IR 8545. Earlier, we have touched upon the larger key and signature sizes. Concrete benchmarking studies are available and more are being carried out on diverse platforms and use cases. The respective official standard documents for ML-KEM and ML-DSA provide the following information.

Table 3. Sizes (in bytes) of keys and ciphertexts of ML-KEM

	encapsulation key	decapsulation key	ciphertext	shared secret key
ML-KEM-512	800	1632	768	32
ML-KEM-768	1184	2400	1088	32
ML-KEM-1024	1568	3168	1568	32

Table 2. Sizes (in bytes) of keys and signatures of ML-DSA

	Private Key	Public Key	Signature Size
ML-DSA-44	2560	1312	2420
ML-DSA-65	4032	1952	3309
ML-DSA-87	4896	2592	4627

For signature schemes, continuously updated and extensive performance benchmarking, including schemes not chosen as standards, is made available online by PQShield. We see that the key and signature sizes are typically in the kilobytes, which are much larger than those of RSA and ECC cryptosystems. In perspective, however, the resource requirements fall well within the capabilities of most current computing devices. In terms of run time, the standard PQC algorithms are competitive, often performing much better than the widely deployed RSA-OAEP and ECCs, even when carried out inside a trusted execution environment such as in an Intel SGX-capable processor.

When compared with the X25519 key exchange protocol, which has security Level 1, we have the following simulation results on a Desktop PC with 10th generation Intel i5 processor and 32GB of RAM running Ubuntu 20.04, equipped with SGX.

TABLE III: Performance comparison of key exchange algorithms. TPS denotes transactions per second.

Scheme	Untrusted		Trusted		Ratio
	Time (μs)	TPS	Time (μs)	TPS	(in %)
X25519	**141**	**7 078**	**172**	**5 806**	**82**
Keygen	35	28 521	62	16 026	56
Encaps	70	14 268	86	11 691	82
Decaps	35	28 187	43	23 272	83
Kyber512	**27**	**36 971**	**98**	**10 173**	**28**
Keygen	8	121 597	62	16 136	13
Encaps	11	92 055	35	28 521	31
Decaps	8	124 983	14	69 166	55
Kyber768	**42**	**23 830**	**134**	**7 460**	**31**
Keygen	13	76 785	82	12 202	16
Encaps	16	60 686	46	21 561	36
Decaps	13	79 962	20	50 980	64
Kyber1024	**57**	**17 407**	**174**	**5 762**	**33**
Keygen	18	57 104	101	9 881	17
Encaps	22	44 664	61	16 436	37
Decaps	18	57 102	25	39 394	69
ML-KEM-512	**24**	**41 284**	**92**	**10 895**	**26**
Keygen	8	124 394	58	17 229	14
Encaps	8	119 190	33	30 416	26
Decaps	8	129 356	14	70 465	54
ML-KEM-768	**38**	**26 571**	**126**	**7 919**	**30**
Keygen	13	78 471	77	12 937	16
Encaps	13	77 716	43	231 084	30
Decaps	12	82 674	19	52 093	63
ML-KEM-1024	**53**	**19 020**	**164**	**6 090**	**32**
Keygen	17	57 362	96	10 385	18
Encaps	18	55 435	57	17 671	32
Decaps	17	58 368	25	40 148	69

The three variants of Kyber and ML-KEM have, respectively, security Levels 1, 3, and 5. Across the board, ML-KEM variants have higher number of transactions per second than their corresponding Kyber variants. The speed of ML-KEM is impressive. Its most secure standard variant is still 2.7 times as fast as X25519. Inside the SGX secure enclave, the performance of Kyber and ML-KEM drops to about 30 percent of its performance in the unprotected memory. The drop is not as significant for X25519 as, inside the secure enclave, it still performs at 86 percent of its performance outside. Even if one opts for ML-KEM-1024 executed inside the secure enclave over X25519 in the normal memory space, the speed ratio is still a decent six over seven.

On the same platform, we get the following simulation results on signature schemes from which one can derive performance comparison in a manner like the one for KEMs. Both Ed25519 and Secp256k1 provide security Level 1. The three variants of Dilithium and ML-DSA have, respectively, security Levels 2, 3, and 5.

TABLE IV: Performance comparison of digital signature algorithms. TPS denotes transactions per second.

Scheme	Untrusted		Trusted		Ratio
	Time (μs)	TPS	Time (μs)	TPS	(in %)
Ed25519	**176**	**5 673**	**204**	**4 901**	**86**
Keygen	37	27 229	64	15 517	57
Sign	35	28 638	45	22 090	77
Verify	104	9 580	114	8 791	92
Secp256k1	**139**	**7 220**	**150**	**6 672**	**92**
Keygen	73	13 662	87	11 485	84
Sign	31	31 945	37	26 768	84
Verify	34	29 160	39	25 395	87
Dilithium-2	**116**	**8 614**	**230**	**4 353**	**51**
Keygen	26	38 272	98	10 248	27
Sign	66	15 184	114	8 765	58
Verify	25	40 262	32	31 186	77
Dilithium-3	**192**	**5 198**	**348**	**2 873**	**55**
Keygen	44	22 797	143	6 993	31
Sign	108	9 289	169	5 905	64
Verify	41	24 124	50	20 021	83
Dilithium-5	**272**	**3 681**	**469**	**2 131**	**58**
Keygen	70	14 221	192	5 218	37
Sign	135	7 417	216	4 632	62
Verify	67	14 909	76	13 076	88
ML-DSA-44	**118**	**8 454**	**231**	**4 334**	**51**
Keygen	26	37 884	98	10 184	27
Sign	67	14 908	114	8 745	59
Verify	25	39 688	32	31 166	79
ML-DSA-65	**194**	**5 158**	**349**	**2 865**	**56**
Keygen	44	22 682	143	6 973	31
Sign	109	9 216	170	5 887	64
Verify	42	23 882	50	20 011	84
ML-DSA-87	**268**	**3 731**	**471**	**2 125**	**57**
Keygen	70	14 357	192	5 205	36
Sign	134	7 474	217	4 618	62
Verify	65	15 313	77	13 061	85

Post-quantum TLS

Pure and hybrid implementations

TLS is foundational to internet security. It protects data in transit across networks by encrypting connections among web servers and clients. It is crucial in preventing eavesdropping and tampering. In short, TLS is the backbone protocol for secure web communications.

Research to integrate PQC algorithms highlights efforts to fortify this essential security structure against quantum threats. Performance evaluations reveal manageable efficiency trade-offs that do not disrupt business. At the same time, we benefit from enhanced security provided by PQC algorithms. Experiments with PQC-enhanced TLS confirmed slight increases in handshake times and data overhead due to larger key and signature sizes. While challenges exist in bandwidth-constrained environments, the overall feasibility for broad application in existing devices is high in ensuring continued confidentiality and integrity of data. In the transition journey, adopting PQC hybrid TLS 1.3 is an economical entry level step that can already be taken.

The following data on implementation and performance of post-quantum TLS are taken from Appendix A of an excellent survey article by Nouri Alnahawi, Johannes Müller, Jan Oupický, and Alexander Wiesmaier, A Comprehensive Survey on Post-Quantum TLS. IACR Communications in Cryptology, vol. 1, no. 2, Jul 08, 2024, DOI: 10.62056/ahee0iuc. The test machines have typical CPU and RAM. So are their communication's round trip time and bandwidth. The helper scripts that can be used to replicate the experiments are available in this github repository.

The trade-off between improved security and performance reduction is reflected in the *slowdown coefficient*, which is the ratio between the average number of connections per second in the post-quantum schemes and that of X25519 for KEMs or Ed25519 for signature schemes. One can see from the next two figures, reproduced from the survey paper, that the values involving Kyber, Dilithium, SPHINCS+, and Falcon are well within our tolerance to absorb. We have stated earlier that the respective standard algorithms ML-KEM, ML-DSA, SLH-DSA have slightly better performance than their predecessors Kyber, Dilithium, and SPHINCS+. We should add that, guided by the specifications in the standard documents, significant optimization can be expected to come from dedicated cryptographic accelerators being built for specific platforms.

The following figure presents the TLS slowdown coefficients for individual algorithms. For Falcon and Dilithium, the variants are indicated by their security levels in the parentheses. For example, L3 means the variant with Security Level 3. For SPHINCS+, s and f denote the short and full variants, respectively, with SHA-2 as the chosen hash function component. The curve for the ECDSA is indicated to be P-256.

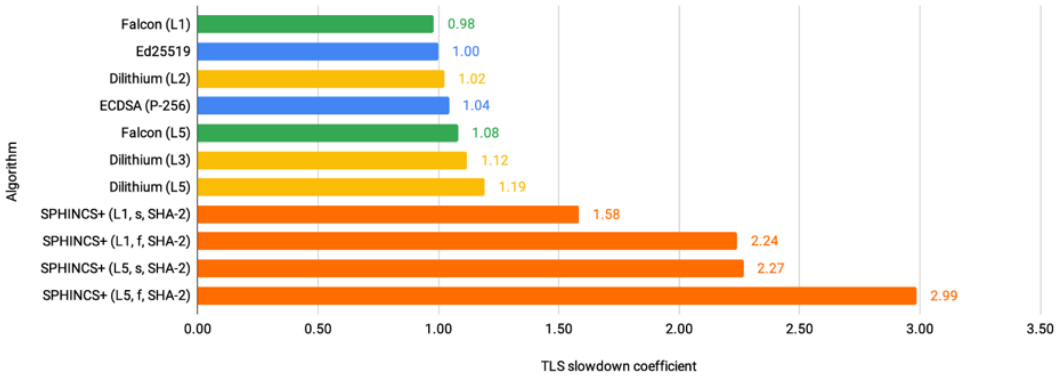

Figure 3: TLS slowdown coefficient for purely post-quantum signatures.

The next figure shows the slowdown coefficient for hybrid algorithms, that is, classical ECDSA and PQC digital signature algorithms working in tandem. The benchmark is taken against Ed25519.

Figure 4: TLS slowdown coefficient for hybrid post-quantum signatures.

Going to a more complex implementation setup, one can consult A Performance Evaluation of IPsec with Post-Quantum Cryptography by Seungyeon Bae *et al.* DOI: 10.1007/978-3-031-29371-9_13 to be sufficiently convinced of the readiness of PQC algorithms for deployment.

A little discussed aspect of replacing RSA/ECC by quantum resistant analogues is the impact on constrained computational devices, *e.g.*, smartcards and IoT devices. These devices are widespread. The US government Personal Identity Verification (PIV) program, for instance, uses RSA-enabled smart cards. For EMV chip payment cards, the primary security service of protecting customer accounts from non "lost/stolen" card fraud is based on symmetric algorithms currently unaffected by the quantum threats. It is important to note then the core fraud protection for EMV chip cardholders, as of today, is based on symmetric algorithms resistant to both classical and quantum computer attacks. However, risk management information exchanged between the card and the terminal is protected by RSA and the integrity of this process is possibly at risk from attacks using a CRQC. EMVCo is clearly aware of this and the situation is fluid, since combining chip cards and quantum resistance is still a work in progress.

Idemia has performed some experiments updating EMV chip card protocols to use a quantum resistant public key algorithm instead of RSA in the context of EMV payment and, ignoring the considerable data communication burden, they found that the computation in the smart card chip, which is an M-3 Cortex processor, consumed almost all of the typical transaction time budgets in current use cases in https://www.gsa.gov/technology/it-contract-vehicles-and-purchasing-programs/federal-credentialing-services

A report posted in the website of Banque de France titled "Observatoire De La Sécurité Des Moyens de Paiement Rapport Annuel 2023" estimated the extra resource implications for chip cards if post-quantum algorithms were to be supported. The report notes that, depending on which quantum algorithm is chosen, 4 to 6 times more RAM chip memory will be required as compared to the RAM used today. Unlike the analysis above which demonstrates that computational devices richer than IoT/Chip cards do not require hardware upgrades to support post-quantum algorithms in terms of performance, this is not the case for IoT/chip card devices. Some constrained computational devices may not be able to be economically upgraded to replace RSA/ECC with post-quantum analogues and may, instead, move to security based on symmetric algorithms.

PART V

PQC migration

In the financial sector, full-scale PQC migration is not yet imminent but preparing and strategizing for migration to quantum-secure technologies is an important task. For the banking and financial industry, the security of critical financial transaction data typically relies on **symmetric key cryptography**, which is assessed to remain quantum-secure, more than on a public key infrastructure.

In the finance industry, using public key cryptosystem for key management is a means to an end. Sharing keys is not a typical business objective or deliverable. In many countries, key distribution is manually done within and between major banks. Nearly all banks have some legacy manual key distribution mechanisms. Manual key component in the context of hub-and-spokes with 'n' parties scales very well. A trusted third party can play the role of the hub with banks acting as spokes. Public key infrastructure has much larger roles to play in the point-to-point networks.

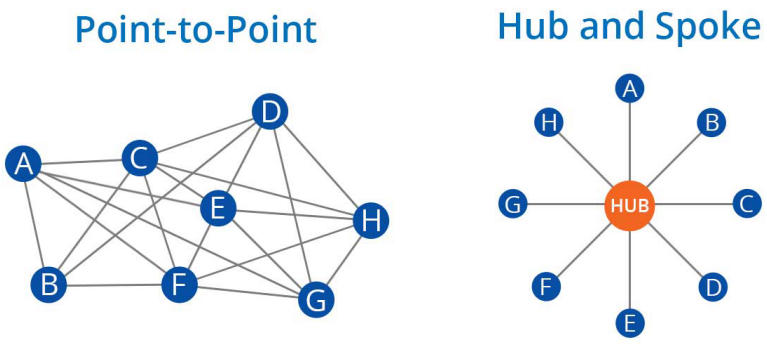

The status of digital signature in the finance industry is materially different. Digital signature techniques are foundational for many business processes. These include reporting to regulators, maintaining auditable business commitments to third parties, and protecting the integrity of agreements that last decades in some cases. Migrating digital signatures is the real problem and more relevant than countering HNDL attack.

The use of quantum-resistant signature algorithms aims to functionally replace what is being done today with RSA/ECC. However, we believe that careful thought must be given to any migration process to quantum- resistant signature schemes. Utmost care must be taken on how to re-confirm existing long term business commitments, how to advise relying parties on the application of any quantum-resistant mechanism, and how to ensure that the associated public key can be recognised as such by the committing party.

In short, it is crucial to confirm that all relying parties are aware of changes in the digital signature methodology.

Ripping out and replacing an infrastructure to mitigate deficiencies tends to be more expensive than adding to it. One can compare it to in-filling rotting cast iron pipes with plastic pipes instead of digging out and replacing them. In fact, analogous infilling processes already exist for digital signatures. These processes enable the validation of signatures beyond the lifetime of the underlying public-private key pair and algorithm combination. By design, they deal with exactly the type of contingency event that cryptographically relevant quantum computers would present.

One such process is known as **qualified time stamping** and is used in conjunction with archiving. The idea is simple. A recognised body is trusted to keep a public-private key pair or algorithm combination that is believed to be quantum-resilient for, say, at least five years into the future. This body reliably supplies an overlay, called a time stamp digital signature, using the latest cryptographic techniques, onto an existing electronic document, its digital signature, and the public key of the "signing" party. The trusted body carries out a due diligence to confirm that the existing digital signature can be currently trusted and applies its time-stamped digital signature. The lifetime of the validity of the underlying business agreement is now extended by the time stamp and the relying parties can act on that knowledge safely. The parties are now vulnerable to the time stamping process becoming vulnerable. However, the time stamping process is designed to be crypto-agile. If circumstances require it, users can obtain a new time stamp with renewed techniques and, hence, extend the lifetime even further. Let us say that, if within the next two years the original contracting party's private key has become compromised, then the time stamp still asserts the validity of the time-stamped data. Better still, the time stamping process can be repeated as often as one likes.

One key issue for the finance industry is to select an appropriate quantum resistant digital signature scheme and, based on that, build an appropriate structure for public key certification. Given that business processes involving digital signatures are typically long term, the execution time of the signature process and the size of the signature data are likely not crucial concerns.

Setting aside regulations and mandates, assessing the situation purely from the current state of quantum technologies, the risk incurred by delaying migration by several years while waiting for PQC standards to be battle-tested in real deployments is reasonably low. We believe that what every financial institution should invest in right now, regardless of the timeline for PQC migration that it deems to be ideal, is on cryptographic inventory tools. This investment is mentioned in **all** guidelines and mandates that have been made public, regardless of sectors and levels. A report from FS-ISAC on infrastructure inventory

can give a general idea of what to check. Commercial products that generate and display cryptography bill of materials (CBOM) have started to hit the market. Examples include QVision from PQStation, SandboxAQ's AQtive Guard, IBM Guardian Quantum Safe, and CipherInsights from QuantumXChange.

It is important to know, independently of any quantum threat, what protocols are effectively securing the flows of data in our networks, where are public key certificates located and used, which processes use which symmetric keys and which public key pairs, and where is the most valuable data located and protected. In short, the task is to build an information and associated protection asset register. This investment in cryptographic inventory is a win-win. When the time comes for a full-fledged migration to PQC or other cryptographic standards in the future, we already know where we are, where we want to be, and the gaps between the two. In the meantime, prudence dictates that we continue to closely watch the development of quantum computing and the supply chains of quantum-secure products and solutions.

Concluding remarks on quantum-secure migration[1]

The financial industry has been self-regulating in terms of its use of cryptography and in the past has executed many cryptographic migrations; they have all been triggered by tangible widely accepted evidence points which have galvanised suppliers to build new product and stimulated customers to buy those products with their improved security. For example, the migration from Single DES to 2 key Triple-DES was triggered by Michael Wiener's paper "*Efficient DES Key Search*", this paper described a circuit, how to build it and presented a cost, based on facts, for breaking DES and began the move to 2-key triple-DES. The financial industry as part of critical digital systems needs a clearer risk analysis akin to the Michael Wiener paper to galvanise progress on a quantum resistant digital signature approach which will inevitably be hugely disruptive and very expensive: it needs to be the right approach executed at the right time.

AES-128 is considered secure for decades against both classical and quantum computer attack and, consequently, remains fit for purpose for the financial industry. National security agencies must minimise residual risk in protecting information assets that may need to be protected for at least 100 years. For such government data, transitioning to AES-256 provides a security margin against developments, however unlikely given current knowledge, which we cannot yet foresee across such a huge timespan.

[1] The views and conclusions expressed in this document are those of the authors and do not necessarily reflect the official policy or position of any affiliated organization. This material is provided for informational purposes only and does not constitute legal, financial, or professional advice. Readers are encouraged to conduct their own analysis and form their own opinions based on the information presented.

Economic and environmental considerations should influence policies. Most payment transactions stay well protected by current systems and are under no immediate threat. The industry's communication security workhorse is the TLS protocol since data in transit is typically protected by TLS-type protocols. A quick win against the quantum threat is to upgrade to a hybrid TLS scheme built on ECC plus ML-KEM as soon as economically sensible. This can be used to perform internal symmetric key management and to protect communications in general. With this enhancement, data that is at rest can be secured by symmetric cryptography where the symmetric keys are established without relying on just ECC or RSA.

It appears then that the HNDL threat against RSA/ECC symmetric key establishment is being tackled by initiatives such as OpenSSL 3.5 and the enhancements to TLS 1.3, and the way forward is clear. However, for quantum resistant digital signature there is currently a lack of clarity and indeed a state of flux exists since NIST are running a search for new quantum resistant digital signature techniques, the current standardised digital signature techniques do not fit all environments sufficiently well.

As regards digital quantum resistant digital signature, beyond analysing use of RSA/ECC and assessing the performance features of the current quantum resistant digital signature standards and given the immaturity of the supply chain, there is little further concrete action that can be taken.

WHITE PAPER CONTRIBUTORS

Mastercard contributors	
John Beric	Vice President, Product development
Rob Byrne	Vice President, Software engineering
Bruno Chagas	Lead Data Scientist
Steve Flinter	Distinguished Engineer

NTU Singapore and PQStation Pte. Ltd. contributors	
Prasanna Ravi	Digital Trust Centre
Anupam Chattopadhyay	College of Computing and Data
Martianus Frederic Ezerman	School of Physical and Mathematical Sciences
Shivam Bhasin	Temasek Labs