BEYOND blue

NEWSLETTER
SEPTEMBER 2022

OPERATING RESILIENTLY

THE VIRTUES OF BUILDING OPERATING MODELS AROUND RESILIENCE

The operating models of financial services firms in the years leading up to the COVID-19 pandemic were in a state of constant flux and modification as a result of evolving services driven by technological advancement, and increasingly complex regulatory regimes. Following the financial crisis in 2008, the international regulatory community implemented considerable reforms and measures aimed at reducing the risk to the public finances. In the UK, the implementation of a comprehensive bank resolution regime and the passing of the Financial Services (Banking Reform) Act 2013 introduced a new found focus on ensuring that firms and banking systems in aggregate are financially resilient. Adequately capitalised, with sufficient liquidity buffers, and enhanced recovery and resolution mechanisms aimed at lessening the impact of a firm's financial failure. However, despite creating a system more resilient to financial vulnerabilities, organisations in the post-pandemic era also face a pressing need to both modernise their operating model to fit a more digital, decentralised and data-driven environment. In doing so they also encounter regulators who are increasingly concerned about Operational Resilience and the systemic implications of an increasingly hyperconnected financial sector.

An operating model represents organisational DNA. It articulates how parts of an organisation work together to deliver a strategic vision whilst also creating value. There are several ways in which the different components of this framework (i.e., people, processes, technology and governance) could be configured to form a unique blueprint which fits the specific needs of a firm. The structure can be based on a series of value chains or 'layers' helping organisational leaders depict how these different parts of their organisation interact to deliver on their ambition of having a truly resilient set of important business services.

There are varying academic and theoretical takes on what a good operating model and framework looks like. Large and medium scale businesses usually have a mature view of how their operating model complements their business needs. But are these models truly geared to manage an organisation's journey through shifting risk conditions, all whilst embedding operational resilience processes and meeting resilience expectations? Have the financial resilience gains of the 2010s resulted in a robust readiness for systemic-level disruptions to critical operations and core business lines? Operational resilience is not a goal which once achieved can be forgotten - it must be continually managed. Hence, to assess their present and future preparedness, firms need to develop an operating model which is more than just a set of plans, programs, processes and people. It should be a model that's composed of these independent, yet interrelated components but these elements must be unified by a focus on achieving a common goal. Namely, ensuring the resilient provision of important business services. Moreover, a commitment from the top of the business downwards to the prioritisation of the resilience of important business services is paramount to ensuring progress is made against resilience objectives.

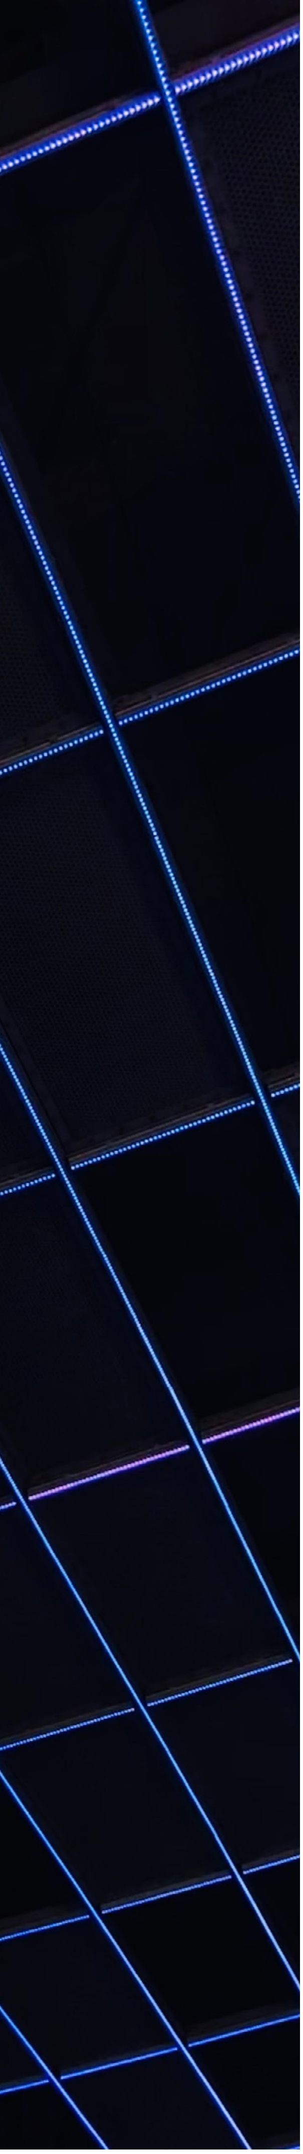

Linking the Operational Resilience agenda with a proactive and resilience-driven approach to operating model design is key in taking early, considered and comprehensive action. Whilst it's up to each firm to determine what their operating model framework looks like, the following are a few crucial principles that will underpin its success.

- Define strategic intent for operational resilience with the buy-in of senior stakeholders across the organisation.
- Embody a "Resilience First" culture, with a view to ensuring the customer journeys of important business services remain viable even during disruptions.
- Leadership should drive strategic investment decisions such that they address not only known and existing operational resilience deficiencies but also 'grey swan' events (i.e., risks which may seem improbable but are nevertheless conceivable).
- Converge and leverage frameworks already embedded within the organisation such as security and technology risk management, third-party oversight, business continuity and incident management to drive an overall resilience framework.
- Explore opportunities to adapt existing governance structures, risk forums, etc., to integrate operational resilience performance and risk management. This ensures resilience risks are visible and are assessed in conjunction with other types of business risks (e.g., operational risks, financial risks, credit risks).
- Consider the pros and cons of centralised vs. decentralised approaches to embedding your operational resilience framework. It's essential that organisations fully define the roles and responsibilities of the Three Lines of Defence, as well as aligning the incentives of senior managers.
- Empower the business to draw on skills sets from multiple disciplines whilst also breaking down silos between business-line risk and control teams. Try to build a holistic knowledge and expect that tackling resilience issues will require cross-disciplinary working.
- Develop a data-driven understanding of capabilities that directly maps to and evidences the resiliency of the firm and its important business services. Through evidence-based reporting, gain thematic insights around the maturity of capabilities and ensure management information supports timely decision making.
- Embed effective change management processes to include "Resilience by Design" whilst addressing any areas where resilience may contradict or override other disciplines e.g., security or technology management.

Whilst designing an effective operating model around operational resilience will be a ongoing project for most firms, it will ensure that they are more proficient at assessing, analysing and addressing any resilience implications of future changes to the overall business. An effective framework will also pave the way for firms to be able to conduct model-based testing on operational failure scenarios.

A focus on safety (e.g., prevent, protect, detect) in the context of operational resilience and security risk management might appear prudent for many organisations, but it ignores the necessary actions to build the ability of the organisation to absorb and recovery from a unexpected jolt or disruption. Having a strategic resilience mindset is key to that. Prioritizing a resilience-oriented operating model which embeds and builds resilience is likely to be more impactful than focusing on individual risks or concerns.

POST-QUANTUM CRYPTOGRAPHY

WHAT YOU NEED TO KNOW AND WHAT YOU SHOULD BE DOING ABOUT IT

In recent months, you may have noticed more and more people getting worked up about post-quantum cryptography. To those with an aversion to all things mathematical, this can induce a state equal parts nausea and apathy. To those with any understanding of the substantial challenges that must be overcome to develop reliable quantum computers, thinking about these problems can simply seem premature. But the recent approval by the US National Institute of Standards and Technology (NIST) of a set of new quantum-secure encryption algorithms has served as a reminder that the threat posed by quantum computers is both serious and increasingly pressing. With that in mind, it's worth taking a moment to reflect on what the problem is, when UK financial institutions (UKFIs) need to start taking action and what early actions to take.

WHAT'S THE PROBLEM WITH CRYPTOGRAPHY?

Classical cryptography refers to the kinds of encryption which are hard to crack by digital computers (i.e., computers which are only capable of manipulating 1s and 0s). This includes all the well-known and widely used symmetric and asymmetric encryption ciphers (e.g., AES, DSA, Diffie-Helman, RSA, SHA) and the protocols which govern the use of these ciphers within networks. Asymmetric ciphers and their deployment in widely used protocols like Transport Layer Security (TLS) are particularly important. The reason for this is that they make it possible to do certain things than with the use of symmetric ciphers. In particular, they provide a relatively straightforward way for two parties to establish secret keys over an insecure channel and, by extension, for ensuring the authenticity and integrity of any communications between both parties.

To do this, asymmetric ciphers require a function which is easy to perform in one direction but computationally very difficult to reverse. One of the most widely used asymmetric ciphers called "RSA encryption" does this by relying on the simple fact that it is easy to calculate the product N of two prime numbers p and q but it can be very hard to find p and q given only N. In mathematical parlance, it is very hard to calculate the prime factorisation of N. The problem is that we now know that there exist algorithms for quantum computers which enable them to find the prime factorization of any positive integer with relatively little difficulty, even if that integer is very large.[1] The upshot of this is that whereas it would take all of the digital computers in the entire world thousands of millions of years to crack a modest RSA encryption, it would take a quantum computer a matter of seconds.

What makes this particularly chastening is that TLS, which relies heavily on this sort of cipher, is used to secure almost all web connections. This includes not only connections between systems internal to your network and those external to it, but also in much of the network traffic internal to your network. There is the possibility that anything within your networks, from critical middleware infrastructure to your video conferencing software, is using TLS.

[1] Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. In Proc. 35th Ann. Symp. on Foundations of Computer Science (FOCS '94) 124–134 (IEEE, 1994).

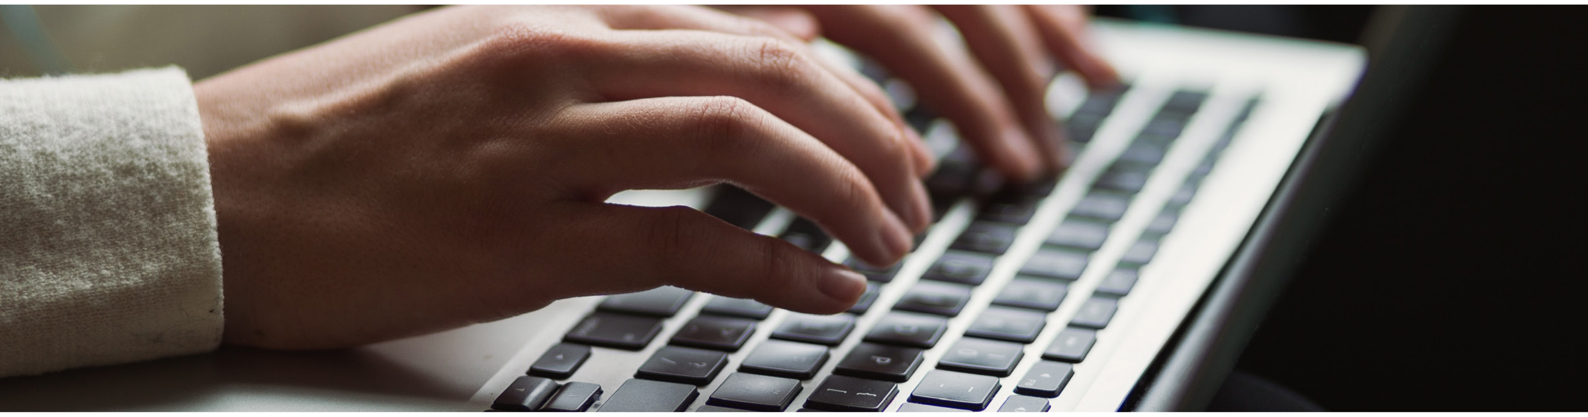

WHAT CAN WE BE DOING TO PREPARE?

Fortunately, we still have time left before the old ciphers can no longer be relied upon. Recent research suggests that reliable, error-corrected quantum computers capable of cracking ciphers like 2048-bit RSA encryption will likely be in the wild by 2036. So, there is clearly no need to take any immediate remedial action. However, it is important to note that the eventual deprecation of these widely used ciphers will be require a substantial response. It will be a significant challenge for large, complex organisations to migrate away from public key cryptography seeing as it is likely embedded in hundreds of different applications and platforms throughout their networks. They will also need to move away from certain low-latency symmetric ciphers to those which involve much larger keys which take more computational power to process.

For large UKFIs, this begs the question "What, if anything, should we being doing now to prepare?" We think there are two key steps that UKFIs can take without delay to prepare the stage for the shift to post-quantum cryptography.

CLARIFY HOW AND WHERE YOU ARE USING CRYPTOGRAPHY

Given the widespread deployment of encryption required across complex enterprise networks, there is a high likelihood that most UKFIs have only limited visibility on their cryptographic estates. Network security teams, however well-resourced or managed, will often struggle to maintain a map of this vast and diffuse array of deployments. Moreover, a general anxiety amongst non-specialists surrounding the complexity and mathematical nature of cryptography can mean that even those responsible for managing applications which use encryption often lack a full understanding of which protocols and ciphers are being deployed and to what end.

This fragmentary nature of enterprise knowledge about the use of encryption could become more of an issue. If UKFIs are to be able to effectively manage the risk posed by the inevitable quantum insecurity of classical modern cryptographic techniques, it is critical that they be able to exhaustively identify the most acute vulnerabilities, such as any dependencies on asymmetric ciphers for key exchanges and digital signatures. Without adequate visibility over the cryptographic estate, it is unlikely this will be achievable.

Fortunately, however, developing this more robust knowledge ecosystem is something which UKFIs can start doing now. Don't wait to be driven by regulatory pressures. This will take too long. Task your security teams to risk-assess the landscape for use of encryption technologies.

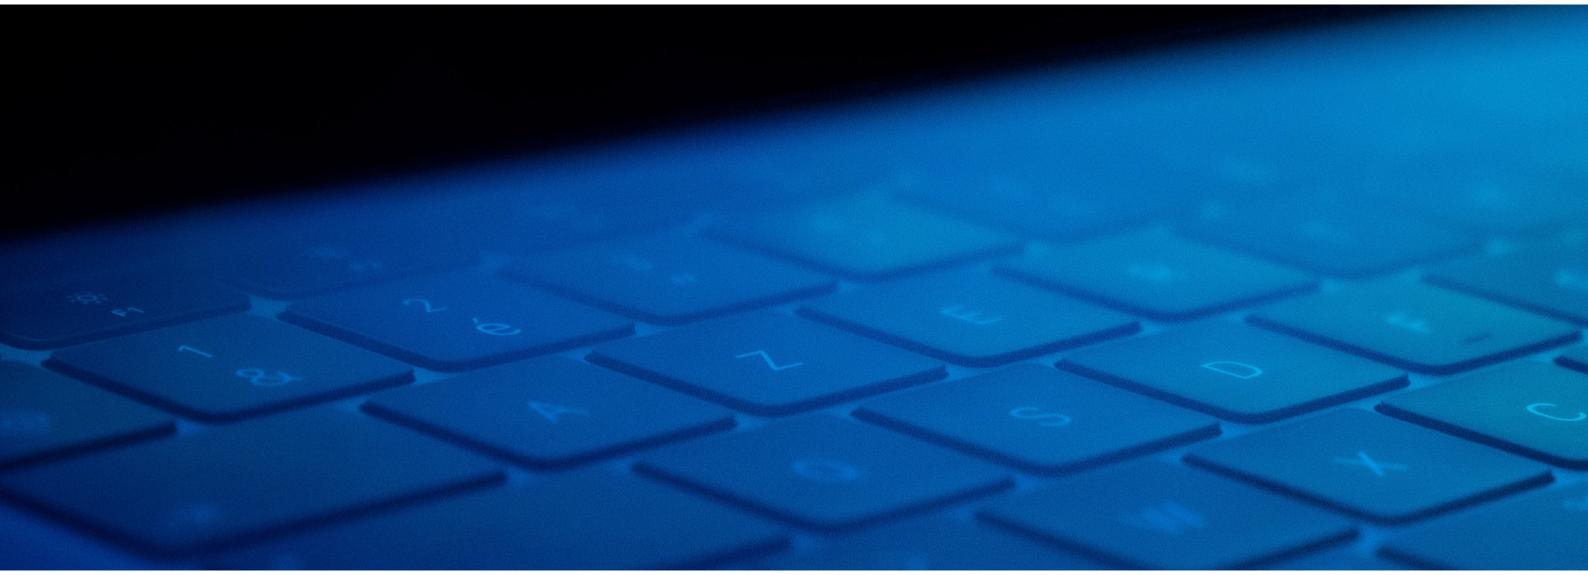

WHERE POSSIBLE, MODULARISE YOUR NETWORK ARCHITECTURE WHICH REQUIRES ENCRYPTION

One option for UKFIs when confronted with the challenge posed by quantum computers is to adopt the commercially available, first-to-market post-quantum cryptographic solutions. However, for most, this is not necessary. These solutions are expensive, high-latency and (given the lack of any existing error-corrected quantum computers) cannot yet be tested. Moreover, most of them were developed in advance of the publication of the NIST standards and, whilst ingenious, are not designed with them in mind. Whilst this latter complaint is likely to soon become obsolete, it is important to remember that the NIST standards are themselves untested in a live environment and are relatively immature. Much more time and work will be needed before we can be confident that they will provide the security we require of them.

UKFIs are much better placed to use the time available to them now to ensure that they are able to switch out their classical cryptographic solutions for post-quantum ones when necessary. The key to making this approach to remediation possible is to ensure that the parts of your network architecture which handle cryptographic encryption and decryption are modular. In other words, you should seek to ensure that the hardware, firmware and software completing cryptographic workloads in your network is both interoperable with your existing systems but also self-contained. This will make it much easier for you to retrofit your network with the relevant post-quantum solutions when they reached the requisite levels of maturity and affordability required, and when the risks posed by quantum computers become more acute.

It will also enable UKFIs to respond to another challenge posed by post-quantum cryptography: the fact there is no one-size-fits-all form of post-quantum algorithm that will enjoy the cross functional applicability of classical ciphers like RSA or Diffie Helman. Ensuring your cryptographic estate is modular in the manner above will enable your security and application teams to pick the post-quantum encryption algorithm best suited to their needs when the time comes.

So, the approach is one of cautious, strategic preparedness. Quantum computers are liable to transform the world of digital finance. One of the ways in which they will do this is to necessitate the widespread adoption of sophisticated Quantum-computing resistant ciphers and protocols. But all of this is still some way off. For now, it will suffice to take stock, prepare the ground for change and await the coming Quantum revolution.

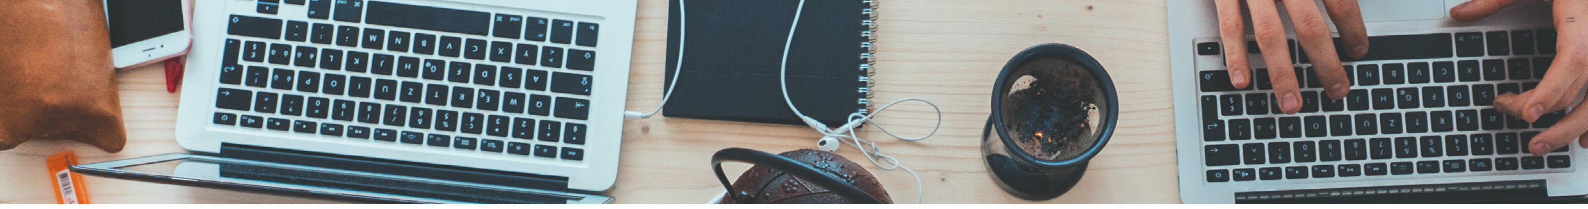

WHAT'S NEXT FOR CRITICAL THIRD PARTIES?

On 21 July the latest resilience related discussion paper landed, this time focused on critical third parties (CTPs). This comes 3 and a half years after the operational resilience discussion paper from the FCA, PRA and Bank of England. The 2018 discussion paper reignited discussions about where the responsibility for the oversight of Financial Market Infrastructure (FMI) and systemically important third parties resided. Was it with the financial institutions (FIs) who used their services or with the regulators? The latest discussion paper sees the regulators concede that "no single firm or FMI can adequately monitor or manage the systemic risks that certain third parties pose of the supervisory authorities' objectives, including UK financial stability, market integrity and consumer protection." For many FIs, this will have been music to their ears. The discussion paper goes on to propose that suitably systemic third parties could receive a formal designation of a CTP by Her Majesty's Treasury. That said, the proposal is not to bring CTPs in their entirety into the regulatory ringfence, just the material services that they provide to the sector. After all, whilst not explicitly mentioned in the discussion paper, the most obvious candidates for CTP designation may be the top Cloud Service Providers. The likes of Microsoft Azure, Google Cloud Platform and Amazon Web Services are unlikely to look favourably on the prospect of regulation.

Whilst we will seek to analyse the ins-and-outs of the latest discussion paper in more detail elsewhere, for now, we would note that the paper raises the following critical question:

"Who has the responsibility for overseeing the resilience of these CTPs in the short term?"

Looking at the timeline for the original Operational Resilience discussion paper, nearly 6 and a half years will have passed between its release in December 2018 and the regulatory compliance deadline of March 2025. Applying that same timeline to the CTP discussion paper, it may be the start of 2029 before the proposals set out in the discussion paper demand compliance. This window of ambiguity may push firms to question the minimum level of oversight and, by extension, compliance with the Operational Resilience Policy that they can get away with. This would have the perverse effect of cultivating in UK firms the exact opposite of the mindset the Policy was setting out to instil. Unfortunately for FIs, if a significant incident were to happen at a CTP that breached one or more impact tolerances, the regulators would likely end up knocking at their door with questions about the level of oversight the FI had over the third party's operational resilience and, perhaps most importantly, any trade-offs that were made between the operational efficiencies the third party provided and the transparency they offered into their resilience.