CMORG
CROSS MARKET OPERATIONAL RESILIENCE GROUP

Collaborative PQC Readiness Engagement

VERSION 1.0 | OCTOBER 2026 | TLP CLEAR

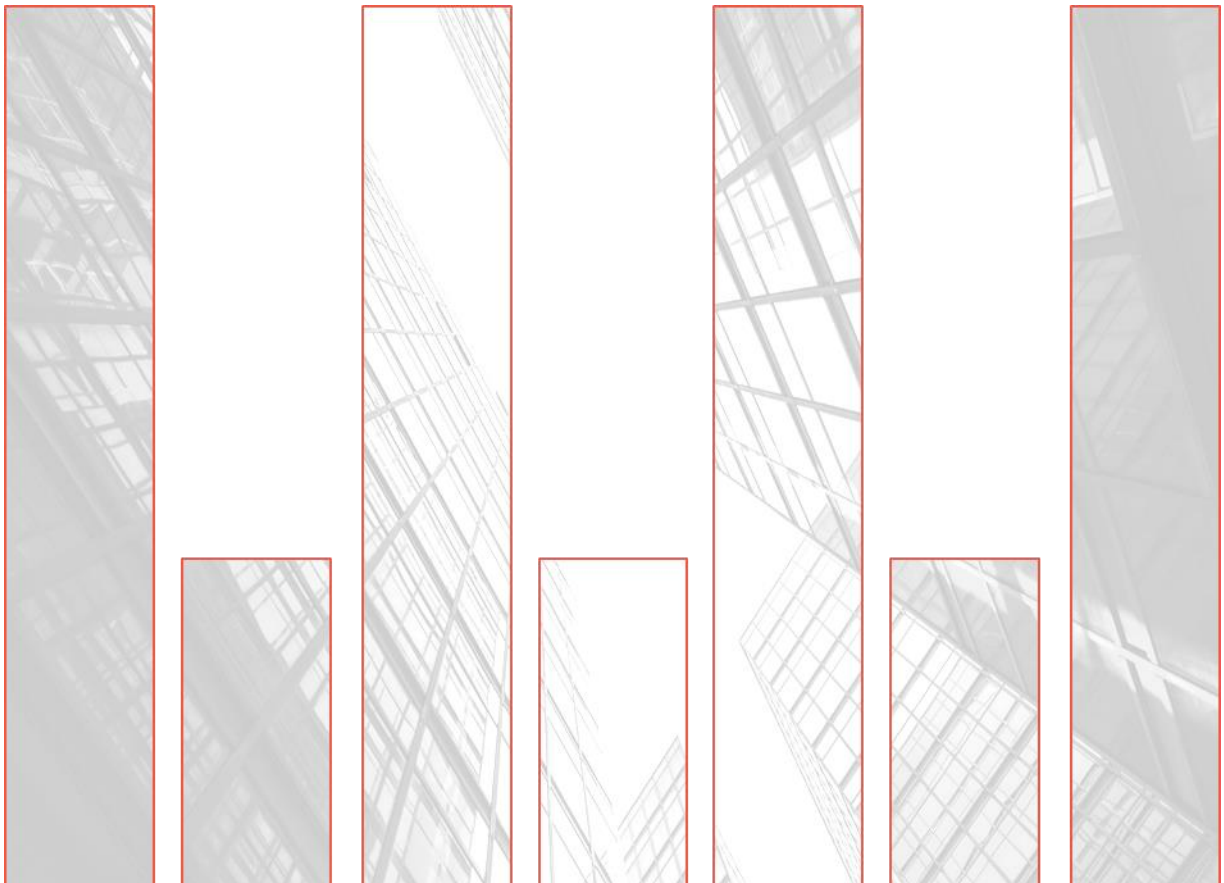

Contents

Introduction **2**

1 Initial engagement **6**

1.1 Engagement trigger 6
1.2 Scope definition 6
1.3 Audience and participation 7
1.4 Information governance 7

2 Data collection and analysis **8**

2.1 Assessment framing 8
2.2 Assessment standardisation 8
2.3 Proportionate evidence 9
2.4 Confidence and evidence handling 9
2.5 Dependencies and constraints 10

3 Evidence review **11**

3.1 Challenge and validation 11
3.2 Shared perspective 11
3.3 Sector relevance 12
3.4 Outcome and follow-up 13

4 Reporting and playback **14**

4.1 Findings and recommendations 14
4.2 Playback and prioritisation 15
4.3 Collective learning and reuse 15

Pilot outcomes and next steps **17**

Appendix A: Terms and definitions **21**

Appendix B: Key reference material **22**

CMORG-endorsed capabilities (including good practice guidance, response frameworks and contingency tools) have been developed collectively by industry to support the operational resilience of the UK financial sector. The financial authorities support the development of these capabilities and collective efforts to improve sector resilience. However, their use is voluntary, and they do not constitute regulatory rules or supervisory expectations; as such, they may not necessarily represent formal endorsement by the authorities.

Introduction

Post-quantum cryptography (PQC) presents a sector-wide migration challenge that is both technically complex and operationally non-trivial. While individual firms will manage their own cryptographic exposures, material dependencies sit outside firm boundaries, particularly where multiple organisations rely on shared infrastructure, common service providers, and widely deployed products. As a result, achieving a secure and effective transition to PQC requires more than isolated activity by individual organisations. It requires coordinated engagement, shared problem definition, and a practical means of surfacing dependencies and constraints early.

The **PQC Taskforce Sector Coordination Pilot** was established under CMORG to design and test a structured, repeatable approach for collective engagement with priority common infrastructure providers. The pilot is intended to generate a coherent view of technical feasibility, ecosystem dependencies, implementation constraints, readiness constraints, governance considerations and "no regret" preparatory actions that could support future migration planning. The pilot is intentionally time-boxed and discovery-oriented. It is not intended to design or deliver full solutions, but to make explicit what can and cannot yet be answered, and to provide actionable outputs that can be reused across the sector.

For the purposes of this pilot, "readiness" refers to the extent to which an organisation, service, product or dependency is able to **operate securely following the transition to post-quantum cryptography**, taking account of any relevant technical, operational and dependency-related constraints. Experience from previous cryptographic migrations suggests that readiness challenges often extend beyond the technology itself. Changes that appear straightforward from a central infrastructure perspective may depend on the readiness of software suppliers, service users, managed service providers, governance processes and broader operational ecosystems. As a result, understanding and managing dependencies may be as important as understanding the underlying cryptographic change itself.

Purpose and audience

This report describes the objectives, scope, and early outputs of the pilot, and is intended to support a shared understanding of the pilot's purpose, methods, and emerging lessons. The report is aimed at:

- **senior decision makers and those with operational accountability** for PQC readiness, cryptographic risk management, and technology dependency governance.
- **CMORG and CIO Forum members**, who require a coherent and evidence-based view of where sector coordination is most valuable and where constraints currently sit.

Scope and approach

The pilot adopted a **twin-track approach**, designed to deliver a short-term practical discovery with a willing pilot organisation (in this case, Pay.UK and VocaLink), while developing a repeatable engagement model for broader use. Across both tracks, the pilot emphasised:

- a focus on **learning and transparency**, rather than implementation;
- explicit articulation of **constraints and uncertainty**; and
- production of outputs that are **reusable**, including repeatable questions and practical considerations that organisations can apply as part of their own third-party engagements.

The two tracks are intentionally complementary. Track A provides practical experience and evidence from a real-world dependency, while Track B seeks to capture and generalise those lessons into a repeatable engagement model that can be applied across other common dependencies.

Bacs TLS discovery pilot (Track A)

Track A focused on improving collective understanding of the PQC readiness for a **deliberately limited subset of the Bacs ecosystem**, selected as a high-value case study due to broad sector usage. The in-scope focus was intentionally constrained to be deliverable within a three-month window through two lenses:

- **Bacs services in scope**: internet-facing connections for Bacstel-IP (BIP) and the Payment Services Website (PSW), deemed as the most critical connections for payments. Other payment schemes or deeper analysis of additional Bacs channels (e.g., ETS/STS) were agreed as out of scope.

- **Cryptographic focus**: The pilot was limited to the Transport Layer Security (TLS) aspect of the selected services, given the readily available TLS migration path, and reusability of findings across broad range of TLS use cases across the financial sector. Specifically, implementation of hybrid configurations which combine classical algorithms (such as X25519) with quantum-resistant mechanisms (e.g. ML-KEM768) were agreed as the most practical and currently available focus area for near-term progress. Broader PQC topics where standards and product support are less mature (e.g., certificate formats for PQC authentication/signing, full PKI redesign, HSM migration, or signature mechanism uplift), except where they emerged as relevant dependencies or "known unknowns".

The approach for Track A was anchored on Pay.UK's interaction model for Bacs, with a focus on identification of dependencies, feasibility constraints, client ecosystem considerations, and the steps required to support a transition path to a hybrid PQC-capable TLS v1.3, rather than implementing change. Steps included:

- identifying and reaffirming TLS termination points and relevant middleboxes under scheme control;
- reviewing readiness of TLS stacks and dependencies needed for TLS 1.3 and hybrid ML-KEM support;
- articulating plausible migration paths for the client ecosystem (including software suppliers and end-user environments), recognising that client readiness may be the primary constraint; and
- making explicit the operational and implementation risks that would need to be tested and managed (e.g., interoperability, performance impacts, packetisation issues, and rollback considerations).

Collaborative readiness engagement methodology (Track B)

In parallel, Track B developed a structured, repeatable methodology to enable market participants to engage collectively with a shared dependency (e.g., a common provider or market infrastructure) to:

- determine an evidence-based view of quantum readiness in line with an agreed scope;
- understand systemic implications, constraints, and prioritisation needs; and
- produce reusable engagement artefacts (question sets, design principles, reporting structure) that can be applied to other dependencies over time.

Track B was designed to be capability-oriented rather than technology-specific, enabling comparability across diverse environments while avoiding unnecessary exposure of sensitive implementation detail. Leveraging similar approaches used in previous collaborative efforts under CMORG, Track B followed a four-stage engagement process as depicted in **Figure 1**.

Figure 1. Approach workflow for Track B

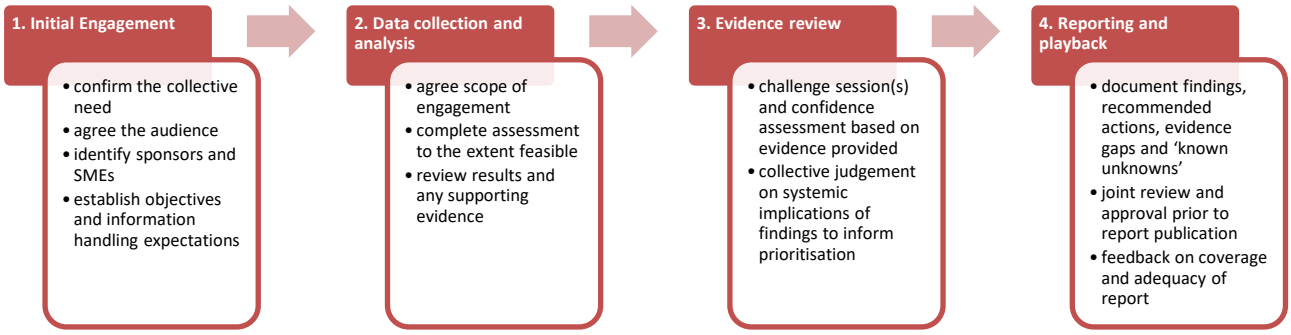

The guidance sets out **16 principles** spanning four phases which are elaborated on in the rest of this document.

Phase	#	Principle	Outcome
1. Initial engagement	1.1	**Engagement trigger**	Collective engagement is initiated where there is a shared need to evaluate or address the PQC readiness of a common dependency.
	1.2	**Scope definition**	The initial scope of engagement is defined clearly, including whether the focus is on an organisation as a whole, one or more services, products or capabilities, or a deliberately constrained subset of these.
	1.3	**Audience and participation**	The engagement establishes an appropriately defined collective audience and the key participants needed to support the work.
	1.4	**Information governance**	Information-handling expectations are established early, including the intended audience, level of detail, any tiered reporting, and how outputs will be shared or controlled.
2. Data collection and analysis	2.1	**Assessment framing**	Data collection begins from a clearly defined assessment boundary, using the service or dependency in scope as the primary organising frame.
	2.2	**Assessment standardisation**	The assessment uses a minimum common question set but remains flexible enough to reflect different dependency types and maturity levels.
	2.3	**Proportionate evidence**	Evidence gathering is proportionate and decision-useful, while making uncertainty and partial completion visible.
	2.4	**Confidence and evidence handling**	Confidence is grounded in the quality of supporting evidence, with sensitive detail abstracted where necessary.
	2.5	**Dependencies and constraints**	Dependencies, blockers, downstream constraints, and transition options are surfaced explicitly.
3. Evidence review	3.1	**Challenge and validation**	Findings are subject to structured challenge and validation to determine whether the readiness position can reasonably be relied upon.
	3.2	**Shared perspective**	Evidence collected during the assessment is reviewed and interpreted to establish a shared understanding of significance, materiality, and residual uncertainty.
	3.3	**Sector relevance**	Findings are assessed for broader relevance beyond the immediate engagement, including whether they indicate common patterns, shared dependencies, or wider coordination challenges.
	3.4	**Outcome and follow-up**	The review identifies matters requiring further evidence, remediation, escalation, or wider sector action, and establishes what can reasonably progress next.
4. Reporting and playback	4.1	**Findings and recommendations**	Engagement outcomes are communicated in a clear, structured format that explains what was learned, any resulting recommendations, and the rationale supporting future activity.
	4.2	**Playback and prioritisation**	Engagement outcomes support a common understanding of priorities, enabling participating organisations to determine where attention, coordination and future activity may be most beneficial.
	4.3	**Collective learning and reuse**	Engagement outputs are retained, shared and reused in a way that supports cumulative learning while respecting agreed information-handling expectations.

Limitations and key considerations

What this is:

- A time-boxed, exploratory pilot intended to surface reusable engagement artefacts, blockers, dependencies, and feasibility constraints for sector coordination on PQC readiness.
- A mechanism to improve shared understanding of where PQC uplift is currently feasible, where it is constrained, and what "no regret" preparatory actions can be advanced.

What this isn't:

- A full implementation programme, prescriptive technical architecture, or sector-wide migration mandate.
- A complete PQC readiness evaluation or migration plan for Bacs, firms, or the wider ecosystem; the pilot focuses on a limited slice of the overall PQC challenge.

Key considerations include:

- **Dependency-driven feasibility**: client, supplier and ecosystem readiness may dominate constraints, and successful migration may depend as much on governance, coordination, testing and operational adoption as on the availability of cryptographic solutions themselves.
- **Security vs. availability trade-offs**: migration paths must anticipate transitional operation, interoperability risks, and practical controls to manage disruption.
- **Information handling and sensitivity**: outputs should be framed to enable collective learning without requiring disclosure of unnecessary sensitive implementation detail.
- **Repeatable but adaptable model design**: the engagement model should provide a consistent structure that can be reused across shared dependencies over time, while remaining flexible enough to accommodate different technologies, levels of maturity, and patterns of sector coordination.

1 Initial engagement

Initial engagement is the entry point for any collective readiness exercise and is intended to establish whether a coordinated engagement is justified, who should be involved, and the practical conditions under which the work can proceed. This phase establishes the basis for collaboration, confirming the shared dependency or use case in scope, identifying the relevant participating organisations, and setting the expectations, sponsorship, and handling arrangements needed before any substantive analysis begins.

1.1 Engagement trigger

Principle 1.1: Collective engagement is initiated where there is a shared need to evaluate or address the PQC readiness of a common dependency.

The trigger should be a **common sectoral interest** in understanding or improving readiness for a dependency that affects multiple participants, rather than a purely firm-specific internal priority, although the same engagement approach may also be used by individual financial institutions in engaging the third parties on which they depend. The model should support both sector-led requests for engagement and provider-led offers to participate. It should therefore accommodate both a 'pull' dynamic, where the sector seeks answers from a shared dependency, and a 'push' dynamic, where a provider wishes to engage collectively with its client or user community. Where a provider initiates engagement proactively, the process may benefit from an existing willingness to share information, discuss readiness challenges, and explore potential migration considerations with affected stakeholders.

The purpose, intended benefits and broad approach should be made explicit at the outset so that the engagement is understood as a **coordinated readiness exercise** rather than a compliance process. This helps to set expectations early, clarify the value of participation, and reduce the risk that the exercise is interpreted as supervisory, contractual or audit activity.

Collective engagement may be most effective where there is sufficient interest from affected organisations to demonstrate that the dependency represents a meaningful shared concern rather than an isolated request. The level of participation required may vary by dependency type, criticality and concentration, but the **perceived legitimacy of an engagement** may influence the willingness of the target organisation to participate constructively.

1.2 Scope definition

Principle 1.2: The initial scope of engagement is defined clearly, including whether the focus is on an organisation as a whole, one or more services, products or capabilities, or a deliberately constrained subset of these.

Early scoping should make clear what is in scope and out of scope for the current cycle of work, recognising that near-term engagements may need to be intentionally narrow to target areas where practical progress is currently feasible.

Early engagements may need to adopt **deliberately constrained scopes** to remain practical and achievable. Narrowly defined assessments can generate useful learning, validation and reusable artefacts while avoiding unnecessary complexity. Broader organisational readiness questions may not always produce decision-useful outputs where readiness varies materially between individual services, products or capabilities, which may then be addressed through subsequent engagement cycles.

A **standardised scoping template** may help establish a consistent assessment boundary, identify relevant participants, document assumptions and constraints, and improve consistency between future engagements.

1.3 Audience and participation

Principle 1.3: The engagement establishes an appropriately defined collective audience and the key participants needed to support the work.

The audience should represent the firms, sector bodies and other stakeholders with a **legitimate interest in the outcome**. It should be large enough to reflect the shared dependency and support collective value but sufficiently bounded to keep the engagement manageable and to align with agreed information-handling expectations. The engagement should establish which organisations wish to participate, helping to ensure that the collective audience remains both relevant and manageable. Executive sponsors and subject matter experts should be identified on all relevant sides to authorise participation, agree engagement scope, resource assessment activities and support effective discussion.

Effective participation may require representation from **multiple functional perspectives**, including technology, architecture, operational resilience, operational risk, supplier management and business ownership. The right mix of participants may vary depending on the dependency and objectives of the engagement.

Roles should be clear at the outset and **documented as part of the agreed scope**, including those accountable for sponsorship and decision-making, those responsible for providing or reviewing evidence, and the intended consumers of any resulting outputs.

1.4 Information governance

Principle 1.4: Information-handling expectations are established early, including the intended audience, level of detail, any tiered reporting, and how outputs will be shared or controlled.

This should be agreed before substantive exchanges begin so that participants understand what may be shared more broadly, what remains restricted, and whether an intermediary or central handling mechanism is needed.

The willingness of participants to share information may depend heavily upon confidence that appropriate handling arrangements are in place. Information governance should therefore be viewed not solely as an administrative activity, but as an enabling condition for effective engagement.

The model should recognise that different information may need to be communicated at different levels of detail. In practice, this may require summary reporting, detailed technical annexes, trusted-reviewer arrangements, or other mechanisms that allow collective learning without unnecessary disclosure of sensitive information.

2 Data collection and analysis

Data collection and analysis is the stage at which the engagement moves from initial framing into evidence gathering. Its purpose is to translate the agreed scope into a structured view of what is known, what can be evidenced, and where the key dependencies, constraints, and uncertainties sit. This phase aims to gather enough information to support collective understanding and prioritisation, while remaining proportionate and practical for the organisation in scope. In other words, it is not intended to produce an exhaustive technical inventory in every case, but to collect sufficient information to identify the most material issues, establish where confidence is high or low, and make explicit any gaps that would need to be resolved later.

2.1 Assessment framing

Principle 2.1: Data collection begins from a clearly defined assessment boundary, using the service or dependency in scope as the primary organising frame.

The evidence-gathering phase should translate the agreed engagement scope into a **practical assessment boundary**, making clear what is in and out of scope for the current cycle of work. In most cases, the service, product, platform, or shared dependency should provide the primary frame, because PQC readiness may vary materially across different services within the same organisation.

The term **assessment boundary** should be understood as the specific organisational, service, product, technology or capability scope against which readiness is being evaluated. The agreed assessment boundary should be documented as part of the engagement scope to ensure consistency of interpretation throughout the assessment and review process. Boundaries should be explicit enough to make clear what evidence is relevant, what dependencies are in scope, and which questions do not need to be answered during the current engagement cycle. For example, an assessment may focus on a single internet-facing service rather than the organisation as a whole, recognising that capability maturity and PQC readiness may differ materially across different products and service lines. Defining this boundary clearly at the outset helps avoid unnecessary complexity, supports proportionate evidence gathering, and provides greater consistency between engagements. However, where relevant, broader organisational readiness activities may provide useful context for the assessment but should not replace service- or dependency-specific evaluation where readiness may vary materially across different parts of the organisation.

Domain- or capability-based prompts can then be used underneath that boundary as **modular assessment components**, allowing only those questions relevant to the service, technology, capability or dependency under review to be selected. This helps avoid unnecessary information gathering, reduces burden on the organisation being assessed, and improves the likelihood of obtaining actionable responses that are directly relevant to the agreed scope.

2.2 Assessment standardisation

Principle 2.2: The assessment uses a minimum common question set but remains flexible enough to reflect different dependency types and maturity levels.

A **structured assessment artefact, questionnaire, or prompt catalogue** should be used to support consistency, comparability, and repeatability across engagements. The engagement model is intended to remain assessment-methodology agnostic. Organisations may choose to leverage existing readiness assessment approaches, maturity models or sector frameworks where these are already established, provided they support the agreed assessment scope and objectives (examples cited in **Appendix B**).

A **baseline question set** should provide a common baseline that can be reused consistently across engagements, enabling comparability and reducing repeated effort for both assessors and assessed organisations. Additional modules may then be applied according to the nature of the dependency, service model, technology type, or maturity of the ecosystem being assessed. For example, a common baseline assessment might include questions relating to governance, cryptographic visibility, inventory, migration planning and dependency management, with additional modules applied only where particular technologies, protocols or service types are relevant. This approach seeks to balance standardisation and flexibility by avoiding excessive tailoring while recognising that not every assessment requires the same depth or coverage.

The objective is not to force every assessment into an identical structure, but to ensure that **equivalent readiness questions** are asked and interpreted in a broadly consistent manner. Organisations participating in multiple engagements should be able to recognise common concepts, terminology and evidence requirements, reducing repeated effort and improving comparability of outcomes across different dependencies and service providers. It should be **reusable across engagements**, but adaptable where the organisation, service model, technology type, or maturity of available standards means that different levels of detail are realistic.

2.3 Proportionate evidence

Principle 2.3: Evidence gathering is proportionate and decision-useful, while making uncertainty and partial completion visible.

The aim should be to obtain sufficient information to understand readiness, constraints, dependencies and systemic implications **within a practical timeframe**. The process should remain workable where an organisation can describe a migration path or dependency landscape only in broad terms, or cannot yet quantify implementation feasibility, effort, or timelines. In practice, readiness assessments are unlikely to begin with complete information. Organisations may still be developing inventories, dependency mappings or migration plans, and some questions may not yet have definitive answers. Assessments may also identify areas where visibility of downstream customers, software suppliers or service providers is limited. The purpose of the assessment is therefore not to eliminate uncertainty entirely, but to understand its nature, significance and potential impact on future decision-making. In those cases, the limitation should be **recorded explicitly as an evidence gap, assumption, or unresolved position**, so that incomplete information does not prevent progress but remains visible in later confidence assessment and reporting.

The principle also recognises that readiness assessments are intended to **support practical decision-making** rather than complete certainty. Organisations may therefore need to proceed on the basis of partial evidence, provided that limitations are understood, documented and revisited where appropriate. The emphasis should remain on obtaining sufficient information to inform action, rather than delaying progress until every uncertainty has been resolved.

2.4 Confidence and evidence handling

Principle 2.4: Confidence is grounded in the quality of supporting evidence, with sensitive detail abstracted where necessary.

The process should **distinguish** between asserted positions, evidenced positions, and unresolved positions. Different forms of evidence may justify different levels of confidence, and unresolved positions should not automatically be interpreted as negative findings where supporting information is not yet available. Where possible, readiness claims should be **supported by** artefacts, technical evidence, architecture information, test results, supplier statements, or other corroborating material. Confidence assessment should distinguish

between different forms of evidence and recognise that not all supporting material provides equivalent assurance. At one end of the spectrum, readiness positions may rely solely on verbal or written assertions. Increasing confidence may be obtained through supporting artefacts, technical evidence, documentation, cryptographic inventories, component bills of materials, testing outputs, independent assessments or formal certification mechanisms where available. A **structured confidence taxonomy** may therefore help ensure that evidence is interpreted consistently across engagements and that confidence levels remain comparable between different assessments.

Confidence assessment may also need to consider the **temporal context** in which evidence is reviewed. A readiness position that appears reasonable early in a migration journey may warrant greater scrutiny if little progress is evident as agreed milestones approach. Confidence should therefore be informed not only by the quality of evidence provided, but also by whether the readiness trajectory appears consistent with relevant implementation timelines, dependency constraints, and planning assumptions.

Consistent with the information-handling arrangements agreed at the outset of the engagement, some technical or commercially sensitive detail **may not be suitable for collective disclosure**. In those cases, evidence may be abstracted into confidence judgements, validated summaries, maturity assessments, or other derived outputs that support collective understanding without unnecessarily exposing commercially sensitive, security-sensitive, or technically sensitive detail. Where direct disclosure is not feasible, organisations may also choose to rely upon trusted-third-party validation mechanisms capable of reviewing sensitive evidence and providing confidence assessments on behalf of the wider collective.

2.5 Dependencies and constraints

Principle 2.5: Dependencies, blockers, downstream constraints, and transition options are surfaced explicitly.

The evidence-gathering phase should identify **constraints, unresolved questions, downstream dependencies, implementation blockers, and ecosystem factors** that affect the readiness position. This should include client ecosystem and supplier dependency information where available, while recognising that the organisation in scope **may not have full visibility** of customer environments, internal effort requirements, or adoption timelines. The readiness position of a dependency may be influenced by actors beyond the direct control of the organisation being assessed, including customers, suppliers, software vendors, service providers and other parties within the broader ecosystem

Where plausible transition options or migration pathways are already known, they should be captured at a high level, not to define the solution, but to **clarify feasibility, sequencing constraints, and issues** requiring later review. The purpose of capturing these pathways is not to prescribe implementation decisions, but to improve understanding of sequencing requirements, prerequisite activities and practical feasibility.

Dependency analysis should remain proportionate and risk-based. While organisations should seek to understand material downstream dependencies where practical, the model should **avoid creating an expectation** that every engagement will achieve exhaustive visibility across multiple layers of a complex supply chain. In many cases, visibility will decrease with each additional layer of dependency, and pursuing complete understanding may become impractical or lead to analysis paralysis. The objective is therefore to identify material constraints, concentration risks and readiness dependencies that may affect decision-making, while recognising that some uncertainty will remain. Where complete visibility cannot be achieved, the output should make those limitations explicit and enable organisations to make informed business judgements based on the information available.

3 Evidence review

Evidence review is the stage at which the collective moves from information gathering into shared interpretation and challenge. Their purpose is not simply to replay the evidence received, but to test understanding, assess the degree of confidence that can be placed in the position presented, and identify the implications that matter most from a collective perspective. This phase should function as a structured sense-making exercise: surfacing where participants agree, where uncertainty remains, and where the findings indicate the need for prioritised follow-up, deeper engagement, or broader sector action.

3.1 Challenge and validation

Principle 3.1: Findings are subject to structured challenge and validation to determine whether the readiness position can reasonably be relied upon.

Challenge and validation activities **may take multiple forms** depending on the nature, maturity and complexity of the engagement. These may include collective challenge workshops, specialist technical review, review of supporting artefacts, trusted-third-party validation, independent assessment, or externally produced assurance material where available. The objective is not to repeat the assessment itself, but to test whether the **conclusions presented are sufficiently well supported**, whether important assumptions or gaps remain, and whether the resulting position is decision-useful for the agreed audience.

The review should remain **collaborative in tone and proportionate in practice**. Given the relative novelty of PQC readiness work, the process should avoid implying that conclusions are fixed or punitive where the underlying technology, standards, implementation pathways or evidence base may continue to evolve. Instead, the review should provide a structured opportunity to test interpretation, identify areas requiring clarification, and establish whether further evidence or follow-up engagement is needed.

Where appropriate, the review stage may **bring in additional perspectives** that were not involved in the earlier assessment activity, such as senior resilience, operational risk, architecture, technology, supplier management or business-owner representatives. This can help test whether the findings are understood consistently across relevant functional perspectives and whether the operational implications of any readiness gaps have been adequately considered.

The process should be sufficiently structured to support **repeatable and evidence-based outcomes**, while avoiding an overly rigid or audit-like framing. Consistency matters because different engagements should not reach materially different conclusions merely because challenge activities were conducted in different ways. This may require common criteria, minimum review prompts, documented rationale, and clarity on the level of depth expected for the service, product, platform or dependency in scope.

3.2 Shared perspective

Principle 3.2: Evidence collected during the assessment is reviewed and interpreted to establish a shared understanding of significance, materiality, and residual uncertainty.

The purpose of this phase is not to reassess evidence quality, which is addressed through Principle 2.4, but to interpret the assessment outputs and **determine what they mean in practice**. Review activities should seek to develop a shared understanding between the organisation being assessed and participating stakeholders regarding the significance of findings, the implications of identified constraints, and the extent of any remaining uncertainty.

The process should recognise that **differences in interpretation** may arise for several reasons. These may include differences in organisational maturity, incomplete evidence, varying assumptions about what "ready" means, or different views on the depth at which PQC readiness should be assessed. For example, an organisation may be able to evidence PQC support at one layer of a solution while unresolved questions remain about underlying firmware, update mechanisms, signing processes, interoperability, or downstream dependencies. These differences should be discussed openly and resolved where possible through reference to agreed criteria and supporting evidence.

Where differences cannot be resolved, the output should **record the divergence clearly** rather than forcing artificial consensus. This should include the nature of the difference, the evidence or assumptions underpinning each view, and any implications for confidence, follow-up action or future reassessment. The emphasis should remain on building a shared understanding for the benefit of the wider migration challenge, rather than positioning the process as a pass/fail judgement on the organisation being reviewed.

The engagement model should also recognise that some interpretive differences may have commercial, contractual or operational implications outside the review process itself. For example, unresolved gaps may affect how participating organisations view substitutability, customer impact, service continuity, or their ability to influence a dependency through contractual or sector-level channels. These implications should be signposted where relevant, but the engagement model should **avoid prescribing a single response** because the appropriate action will depend on context, criticality, available levers and the nature of the dependency.

3.3 Sector relevance

Principle 3.3: Findings are assessed for broader relevance beyond the immediate engagement, including whether they indicate common patterns, shared dependencies, or wider coordination challenges.

The purpose is not for the collective to determine whether an issue is systemically important in a formal sense. Rather, participants should consider whether findings appear unique to the organisation under review or whether they may have **broader implications** for organisations with similar dependencies, technologies, suppliers, service models or migration challenges.

Issues of potential wider relevance should be **highlighted for consideration** by relevant sector bodies or governance forums. This may include common implementation blockers, recurring evidence gaps, dependencies that appear across multiple engagements, areas where readiness is constrained by the client ecosystem, or findings that suggest a need for further sector coordination.

The review should also consider how reusable learning can be extracted without disclosing sensitive implementation detail or identifiable information about the organisation assessed. This may require abstraction of findings into common themes, maturity indicators, anonymised lessons, or aggregated observations. Such outputs should be designed to support wider sector learning while respecting the information-handling expectations agreed at the outset of the engagement.

Where the number of engagements increases, manual consolidation of lessons may become impractical. The model should therefore signpost the potential need for a controlled mechanism to retain, aggregate and query lessons across engagements, subject to appropriate governance, non-identifiability controls, and participant confidence in the party operating the mechanism. Any such mechanism would need to address the risk that sensitive or identifying information could be inferred indirectly, including from the nature of gaps, technologies or dependencies described.

3.4 Outcome and follow-up

Principle 3.4: The review identifies matters requiring further evidence, remediation, escalation, or wider sector action, and establishes what can reasonably progress next.

The review should **distinguish** between findings that support immediate action and those that require further clarification, additional evidence, reassessment or follow-up engagement. Outputs may include recommended remediation activities, areas for additional evidence collection, proposed future engagement, escalation of significant issues, or identification of "no regret" preparatory actions that can proceed even where the full readiness picture remains incomplete. Whilst the review may identify potential actions, remediation activities or areas requiring further engagement, the responsibility for determining and implementing those actions remains outside the engagement itself.

The process should **allow progress to continue where evidence is partial**, provided that limitations are made explicit. In many cases, organisations may not yet have complete visibility of their own downstream dependencies, client ecosystem constraints, supplier readiness, or future implementation path. The output should therefore identify what is known, what remains uncertain, and what can still reasonably be progressed without waiting for every dependency or assumption to be resolved.

Findings should be prioritised in a way that reflects their **significance, urgency and dependency implications**. This may require a structured severity or RAG-style approach to distinguish matters requiring immediate attention from those that can be tracked, revisited or addressed through normal planning cycles. The model should avoid treating all gaps as equally important, because doing so would reduce the usefulness of the final output and make prioritisation harder for sponsors and governance bodies.

The review should also recognise that different remediation paths may create **different longer-term readiness positions**. A narrow upgrade path may address an immediate transition requirement but leave limited crypto-agility for future algorithm change, whereas more fundamental re-architecture may place an organisation in a stronger position over time. Where relevant, the output should distinguish between short-term compliance or migration steps and broader maturity considerations, including the extent to which a dependency can adapt to future cryptographic change.

Finally, the review should provide enough rationale to support **progress tracking over time**. PQC readiness is expected to evolve as standards, products, implementation guidance and organisational plans mature. The output should therefore identify which positions may require reassessment, when known gaps should be revisited, and how updated evidence should feed back into the collective understanding.

4 Reporting and playback

Reporting and playback is the stage at which the engagement is translated into a usable output for the agreed audience. Its purpose is to consolidate the outcomes of the previous phases, communicate what has been learned, and support practical decisions about next steps, prioritisation and future use of findings. This phase should not reopen the assessment or evidence review but should present the outputs in a form that is accessible, proportionate and sufficiently traceable to support confidence in how conclusions and recommendations have been reached.

4.1 Findings and recommendations

Principle 4.1: Engagement outcomes are communicated in a clear, structured format that explains what was learned, any resulting recommendations, and the rationale supporting future activity.

Reporting should **consolidate the outputs of the engagement** rather than reproduce the underlying evidence, assessment artefacts or review discussions. The report should provide a clear account of the key observations, areas of shared understanding, recommended actions or next steps, and the assumptions, dependencies, limitations and evidence gaps that are relevant to those outputs. It should be possible for readers to understand why a recommendation has been made, what it is based on, and how firmly it is supported by the evidence and review activity undertaken in earlier phases.

The report should **distinguish clearly** between what was observed, how those observations were interpreted, and what actions or options may be beneficial as a result. Recommendations should derive from the outcomes established during the evidence review phase, rather than introducing new analysis during reporting. Where there are multiple possible ways to respond to a finding, the output may present options rather than a single preferred path, provided the rationale, trade-offs and relevant uncertainties are made clear.

Evidence gaps, assumptions, confidence considerations and known unknowns should **remain visible** so that recommendations are understood in context. Where the evidence base is incomplete, or where confidence depends on assumptions, this should be explained transparently without undermining the usefulness of the report. The objective is to support informed decision-making, not to imply a level of certainty that the engagement cannot support.

Reporting should also be designed for the needs of the agreed audience. Different stakeholders may require different levels of detail, ranging from concise executive summaries, scorecards or briefing material through to more detailed technical annexes or supporting outputs. A **common reporting structure and format** can help readers interpret outputs consistently across engagements, while still allowing the level of detail and format to be adapted for senior, operational, technical or sector-level audiences.

Where helpful, reporting may also need to explain the wider context or complexity behind the findings in a way that is accessible to non-specialist readers. Visible remedial activity may **mask underlying complexity** associated with a broader set of dependencies, unresolved standards, legacy constraints, and ecosystem coordination issues. Reporting should therefore help senior audiences understand why apparently narrow findings may point to larger readiness or coordination challenges, without overloading the output with unnecessary technical detail.

4.2 Playback and prioritisation

Principle 4.2: Engagement outcomes support a common understanding of priorities, enabling participating organisations to determine where attention, coordination and future activity may be most beneficial.

Playback should focus on **communicating outcomes** and developing a **common understanding of priorities**, rather than repeating the assessment or evidence review activity. It should help participants understand which issues appear most important, where dependencies or constraints may warrant additional coordination, and where further evidence gathering, follow-up engagement or future review may be beneficial. This is particularly important where the evidence review has identified multiple findings that cannot all be progressed with equal urgency.

This engagement approach recognises that priorities may look different depending on perspective. An issue may be a local priority for an individual organisation, a priority for the organisation or service being assessed, or a broader priority for the collective where common dependencies, shared constraints or repeated findings emerge across multiple engagements. Playback should therefore **distinguish** between organisation-specific observations, common implementation challenges, and matters that may benefit from broader sector discussion or coordination.

The prioritisation process should not assume that the engagement model itself can direct or require action. Where significant findings are identified, the report may articulate recommended actions, priority areas or collective expectations, but the mechanism for progressing those actions may **depend on other levers** such as contractual arrangements, commercial influence, governance ownership, regulatory expectations or sector coordination mechanisms. The engagement should therefore help clarify what needs attention and why, while recognising that ownership of action may sit outside the engagement itself.

Prioritisation should also take account of both engagement-specific outputs and wider external signals. For example, national authority guidance, emerging standards, technology developments, or findings from other engagements may affect whether a constraint should receive greater attention. As PQC readiness work scales, **consistent reporting** may help compare issues across engagements and identify where collective effort is best directed.

The purpose of playback is not to prescribe a single course of action for all participants, but to **help stakeholders make informed decisions** about planning, coordination and future work. In practice, this may include identifying issues that individual firms can manage through their own programmes, issues where organisations are "all in it together" because unilateral action is not sufficient, and issues where collective attention may help accelerate or unblock progress.

4.3 Collective learning and reuse

Principle 4.3: Engagement outputs are retained, shared and reused in a way that supports cumulative learning while respecting agreed information-handling expectations.

Reporting should operate within the **governance and information-sharing arrangements** established during the initial engagement phase. Detailed evidence, sensitive implementation information, or organisation-specific findings may need to remain within the agreed audience for the engagement. However, the model should also seek to extract learning that can be reused more widely where this can be done safely and appropriately.

Outputs should support **reuse across future engagements** wherever practical. This may include reusable assessment questions, common findings, recurring dependency themes, maturity observations, implementation considerations, or recommended approaches that could help other organisations facing similar readiness challenges. Reuse should reduce duplication, improve consistency, and help the sector build a stronger understanding of common PQC migration constraints over time.

The model should distinguish between information that is specific to an organisation or provider, and information that can support broader learning. Where detailed findings cannot be shared more widely, they may be **abstracted** into themes, lessons, maturity indicators, confidence statements, or other derived outputs that provide useful insight without exposing sensitive, proprietary or security-relevant detail. This engagement approach acknowledges the tension between PQC capability and proprietary advantage, while also recognising that shared approaches for addressing common problems may be valuable to a wider audience.

Reusable learning may relate both to findings and recommendations. For example, a specific provider issue may only be relevant to those directly dependent on that provider, whereas a common implementation challenge, migration approach, testing model, or solution pattern may be relevant across a wider set of organisations. Where an engagement identifies **lessons with broader applicability**, those should be captured in a form that can be reused, subject to agreed handling constraints.

Over time, reporting should **support comparison between engagements** and provide a reference point for tracking how readiness positions evolve over time. This does not require every engagement to disclose the same level of detail, but it does require sufficient consistency in concepts, terminology and output structure to allow themes and patterns to be recognised. The model should therefore support the **gradual accumulation of sector knowledge** regarding readiness challenges, common dependencies, migration constraints and emerging good practice.

Pilot outcomes and next steps

The pilot has provided both a proposed framework for collaborative PQC readiness engagement and practical insights from its application to a shared sector dependency. The following sections summarise the key outcomes from the pilot, opportunities for further development of the approach, and wider considerations identified during the exercise.

Pilot outcomes

The pilot has demonstrated that a structured, collaborative approach to PQC readiness engagement is **theoretically possible** of producing useful sector-wide insights. Across both tracks, the pilot achieved its primary objective of developing a repeatable engagement model while generating a deeper understanding of the readiness challenges associated with a shared dependency.

From a methodology perspective (Track B), the pilot designed a four-stage engagement process covering initial engagement, data collection, evidence review, and reporting. The exercise **supported the position** that collective assessment can be conducted in a manner that supports shared understanding, challenge, and prioritisation without requiring the disclosure of unnecessary sensitive information. The resulting principles, engagement structure and supporting artefacts provide a foundation that can now be applied, refined and tested through future engagements. While the methodology was informed and iteratively refined through the Bacs pilot (Track A), further application from the outset of a future engagement would provide an opportunity to assess its effectiveness under operational conditions.

From a technical perspective, the pilot found **no immediate barriers** to enabling hybrid TLS capability within the subset of Bacs services examined. However, the assessment also demonstrated that readiness is influenced less by the central infrastructure itself and more by the wider ecosystem of client organisations, software suppliers, bureaux, managed service providers, and downstream dependencies that interact with it. The pilot therefore reinforced the importance of viewing PQC readiness as an ecosystem challenge rather than solely a technology upgrade exercise.[1]

Taken together, the findings suggest that collaborative readiness engagements can provide a **valuable mechanism** for improving sector understanding of PQC migration risks and dependencies, whilst also identifying areas where broader coordination or supporting capabilities may be required.

Next steps

The pilot has provided useful insight into both collaborative readiness assessment and the challenges associated with understanding the readiness of shared dependencies. More importantly, it has **highlighted the scale and complexity** of the broader PQC transition challenge facing the financial sector. Successful transition is likely to require not only technical change, but also coordinated approaches to readiness assessment, dependency visibility, prioritisation, knowledge sharing, migration planning and ecosystem engagement. The pilot should therefore be viewed as an **early contribution towards a much larger body of work** required to support sector-wide transition.

[1] Detailed insights into the BACS Discovery Exercise (Track A) findings and recommendations are not included in the TLP CLEAR version of this report, and are only available to CMORG, CIO Forum, and PQC Taskforce members via request to the CMORG PMO.

Future activity may focus on developing the capabilities, approaches and coordination mechanisms needed to support collective understanding of readiness and migration risk across the sector. This includes:

- improving sector awareness and communication of PQC developments;
- establishing practical approaches for readiness assessment and collective visibility of readiness;
- identifying and promoting reusable migration approaches, implementation guidance and "no regret" activities;
- developing opportunities to build practical experience through collaborative testing and experimentation; and
- improving the coordination, prioritisation and sequencing of migration activity across shared dependencies.

Many of these considerations are explored further in the sectoral challenges set out below. The outputs generated through this pilot will contribute to these wider efforts. The collaborative readiness engagement approach provides an initial basis for future readiness assessments, whilst the findings from the Bacs TLS discovery provide an example of how shared dependencies can be examined to improve understanding of migration feasibility, ecosystem constraints and readiness considerations.

Sectoral challenges to be addressed

As part of the pilot, the taskforce identified several broader challenges that sit outside the immediate scope of the engagement model itself but will influence the effectiveness of sector-wide PQC readiness and transition activity.

Challenge	Description	Why it matters	Capability area(s)
Scaling ecosystem readiness assessment by end 2028	• Assessing the PQC readiness of the organisations, products, services and suppliers that comprise the financial sector ecosystem represents a substantial undertaking. • Without greater standardisation and centralisation, organisations may face a growing volume of overlapping assessment requests, creating inefficiency for both assessors and assessed parties. • Practical mechanisms may be needed to support assessment at scale whilst balancing consistency, proportionality and information handling considerations.	• The sector may struggle to establish a sufficiently comprehensive understanding of migration exposure within available timeframes if assessment activity remains fragmented across individual organisations.	• Readiness assessment
Limited time to establish a common PQC assessment approach	• The sector may require a sufficiently common approach to readiness assessment, evidence evaluation and confidence determination if findings are to be interpreted consistently across organisations and dependencies. • However, the time available before large-scale readiness activity is required limits the sector's ability to develop a bespoke methodology. • Consideration may therefore be needed as to whether existing assessment approaches can be adopted, adapted or recommended in order to accelerate readiness activity whilst maintaining sufficient consistency and comparability.	• Delays in agreeing assessment approaches may reduce the time available for practical readiness evaluation and migration planning, and risks pushing out transition timelines beyond target dates. • Conversely, inconsistent approaches may make findings difficult to compare, aggregate or prioritise.	• Readiness assessment

Challenge	Description	Why it matters	Capability area(s)
Establishing and maintaining a collective visibility of readiness	• Sector-wide situational awareness may require mechanisms to maintain visibility of readiness positions, blockers, dependencies and emerging risks. • However, such visibility may be difficult to achieve where readiness information is distributed across organisations, constrained by information-sharing considerations, or assessed using differing methodologies and confidence criteria. • Questions also remain regarding how readiness information could be sourced, consolidated, maintained and interpreted consistently across the sector, including whether a trusted mechanism could be established to aggregate, maintain and communicate readiness information in a manner that is both decision-useful and consistent with information-handling expectations.	• Without a sufficiently coherent picture of readiness, it may be difficult to identify systemic constraints, prioritise collective action, track progress, or understand where migration activity is most likely to encounter delays.	• Sector communications • Readiness assessment
Collective knowledge retention and secure sharing	• Reusable learning, implementation experience, migration approaches, dependency insights, technical findings and lessons identified through sector activity are likely to accumulate over time. • Mechanisms may therefore be required to retain, share and reuse knowledge whilst respecting agreed information-handling expectations and participant confidence. • Consideration may also be needed as to whether a central repository, controlled information-sharing environment or other mechanism could support cumulative sector learning.	• Failure to capture and reuse learning may increase duplication of effort, reduce consistency of understanding and slow sector progress.	• Sector communications
Developing practical migration experience and reusable approaches	• The sector may benefit from opportunities to identify, test and refine reusable migration approaches, implementation patterns and "no regret" activities. • Practical experience may help reduce uncertainty regarding migration feasibility, implementation effort, ecosystem dependencies and operational constraints, whilst supporting the development of reusable approaches that can be applied across multiple organisations and dependencies. • There is currently limited practical experience of collaborative PQC migration activity across shared dependencies and common infrastructure, creating uncertainty regarding how migration approaches will perform under real-world operational conditions.	• Shared implementation experience may accelerate learning, reduce repeated effort across organisations and improve confidence in migration planning and execution. • Without opportunities for practical testing and experimentation, organisations may be required to make migration decisions with limited evidence regarding feasibility, interoperability, ecosystem impacts or operational constraints.	• Effective practices • Testing and experimentation
Dependency mapping and migration visibility	• Effective transition planning requires improved understanding of shared dependencies, critical paths and migration interdependencies across the sector. • While individual readiness assessments may provide insight into the state of particular organisations, services or suppliers, additional	• Migration sequencing and prioritisation may be difficult where dependency relationships, critical paths and migration	• Migration coordination

Challenge	Description	Why it matters	Capability area(s)
	visibility may be needed to understand how migrations interact, where sequencing constraints exist, and which dependencies may influence the pace or feasibility of broader transition activity. • Questions remain regarding how such information could be sourced, maintained and shared appropriately.	interdependencies are not well understood. • Delays affecting a small number of strategically important dependencies could have disproportionate implications for wider sector transition activity.	
Prioritisation and coordination of migration activity	• The scale of the transition raises questions regarding which dependencies should be prioritised for assessment or engagement, how collective effort should be directed, and how migration activity should be sequenced across interconnected ecosystems, typically where the sectoral exposures are most material. • As assessment and migration activity expands across the sector, consideration may be needed as to how duplicate effort is avoided, how visibility of ongoing activity is maintained, and how a coherent and executable view of sector migration priorities can be developed (i.e. a sectoral 'air traffic control' function). • Different organisations may view priorities differently based on their own risk appetites, contractual relationships, business dependencies and migration timelines, making it difficult to establish and maintain a shared view of sector priorities.	• Coordination challenges may lead to duplication of effort, inconsistent priorities or fragmented migration activity. • Without sufficient coordination, organisations may optimise for local priorities whilst overlooking wider dependency relationships and sector-wide constraints.	• Migration coordination
Governance, influence and sustaining momentum	• Collective assessments may identify blockers, dependencies or readiness concerns, but the mechanisms available to influence outcomes, coordinate action and maintain progress over time are less clear. • Consideration may be needed regarding ownership, escalation pathways, reassessment triggers and ongoing governance. • Questions also remain regarding what interventions, incentives or coordination mechanisms may be available where a strategically important dependency is unwilling to engage, materially lags collective need, or fails to address issues identified through the assessment process. • The pilot also highlighted the challenge of maintaining visibility of progress following an assessment. • Whilst assessments may identify recommended actions, evidence gaps or areas requiring follow-up, questions remain regarding how progress should be tracked over time, when reassessment may be appropriate, and what mechanisms are available where an assessed organisation fails to address issues identified through the engagement.	• Long-term progress may be difficult to sustain without clear coordination, accountability and mechanisms for collective action. • Even where readiness concerns are identified, the sector may have limited ability to influence outcomes if ownership, escalation routes, communication mechanisms, progress-tracking arrangements and collective expectations are not sufficiently defined.	• Migration coordination • Sector communications

Appendix A: Terms and definitions

Term	Definition
cryptographic (crypto) agility	The capabilities needed to replace and adapt cryptographic algorithms in protocols, applications, software, hardware, firmware, and infrastructures while preserving security and ongoing operations. *(Source: NIST CSWP 39-upd1)*
post-quantum cryptography (PQC)	Cryptographic algorithms designed to be secure against attack by both classical and quantum computers; in practice, this refers to encryption and digital signature algorithms intended to resist attacks from cryptanalytically relevant quantum computers. *(Source: NIST)*
quantum algorithm	An algorithm in which one or more steps are quantum operations, or which is implemented on a quantum computer using quantum-mechanical effects such as superposition and entanglement. *(Source: ISO/IEC 4879:2024)*
quantum computing	Computing that uses quantum-mechanical phenomena to represent and process information, typically by operating on quantum bits (qubits) that can exploit properties such as superposition and entanglement. *(Source: ISO/IEC 4879:2024)*
quantum cryptography	The application of quantum-mechanical principles to cryptographic tasks, most commonly to establish or distribute cryptographic keys with security properties grounded in the laws of physics rather than computational hardness assumptions. *(Source: ITU)*

Appendix B: Key reference material

Cross Market Operational Resilience Group (CMORG)

- Guidance for Post-Quantum Cryptography (Oct 25)

UK National Cyber Security Centre (NCSC)

- Preparing for Quantum-Safe cryptography (Nov 20)
- Next steps in preparing for post-quantum cryptography (Aug 24)
- Timelines for migration to post-quantum cryptography (Mar 25)
- Quantum networking technologies (Aug 25)

Internet Engineering Task Force (IETF)

- Post-quantum hybrid ECDHE-MLKEM Key Agreement for TLSv1.3 (May 25)
- RFC 9794: Terminology for Post-Quantum Traditional Hybrid schemes (Jun 25)
- Guidance for migration to Post-Quantum Cryptography (Jan 26)

European Telecommunications Standards Institute (ETSI)

- ETSI TR 104 016 v1.1.1: A Repeatable Framework for Quantum-Safe Migrations (Oct 24)

National Institute of Science and Technology (NIST)

- SP 1800-37: Addressing Visibility Challenges with TLS 1.3 within the Enterprise (Feb 25)

Public Key Infrastructure Consortium (PKIC)

- PQC Maturity Model (PQCMM) v1.0.1 (Jun 26)

CyberSecurity NonProfit (CSNP)

- Quantum Readiness Assurance Maturity Model (QRAMM) (Feb 25)