ETSI TR 104 016 V1.1.1 (2024-10)

CYBER; Quantum-Safe Cryptography (QSC);
A Repeatable Framework for Quantum-Safe Migrations

Reference
DTR/CYBER-QSC-0024

Keywords
cybersecurity, framework, migration,
Quantum Safe Cryptography

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from the
ETSI Search & Browse Standards application.

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format on ETSI deliver.

Users should be aware that the present document may be revised or have its status changed,
this information is available in the Milestones listing.

If you find errors in the present document, please send your comments to
the relevant service listed under Committee Support Staff.

If you find a security vulnerability in the present document, please report it through our
Coordinated Vulnerability Disclosure (CVD) program.

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.
No recommendation as to products and services or vendors is made or should be implied.
No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.
In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.
The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2024.
All rights reserved.

Contents

List of figures

List of tables

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™** and **LTE™** are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**® and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This Technical Report (TR) has been produced by ETSI Technical Committee Cyber Security (CYBER).

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Introduction

Migrating an enterprise's information security systems and infrastructures from a quantum-vulnerable state to a Fully Quantum-Safe Cryptographic State (FQSCS) can be enormously complicated. Depending on the enterprise, the quantum-safe migration can take a decade or longer to complete. Moreover, to successfully migrate the entirety of an enterprise to a FQSCS, it is reasonable to expect a whole-of-enterprise commitment will be required. According to the current best-estimates, the likelihood that a quantum computer capable of breaking RSA-2048 (within 24 hours) emerges within the next ten years is materially high [i.1]. From a risk management perspective, this implies that enterprises should already have started their quantum-safe migration planning.

Unfortunately, there are several reasons for why many enterprises have not yet begun their quantum-safe migration planning. The enormity of the migration itself can act as a disincentive to begin. Budgets and internal resources need to be allocated, which can be difficult to do when there is a lack of internal expertise, governance-level buy-in, or because the timeline for the emergence of a Cryptographically Relevant Quantum Computer (CRQC) cannot be precisely estimated. These issues are compounded when enterprises are working under fixed budgets and are routinely faced with significant security threats including ransomware attacks and zero-day exploits. Moreover, even for enterprises that wish to begin their quantum-safe migration planning now, they often have difficulties deciding where and how to begin.

Various standards development organizations, government agencies, and industry members have published guides and frameworks - with more likely under development - that give enterprises actionable recommendations on formulating and executing their quantum-safe migration plans. Examples include ETSI TC CYBER WG QSC's "Migration strategies and recommendations for Quantum Safe schemes" [i.2] and the "Preparing for Post-Quantum Cryptography" roadmap and infographic by the United States' Department of Homeland Security (DHS); created in conjunction with the National Institute of Standards and Technology (NIST) [i.3]. In late 2022, the Accredited Standards Committee X9 Inc. (ASC X9) published a broadly scoped report on the subject through their Quantum Computing Risk Study Group [i.4]. More recently, a collaboration of cryptographic research groups from the Netherlands published a migration handbook, fundamentally designed around [i.2], defining various migration personae and urgency levels, providing methods to diagnose an enterprise's persona and urgency level, and giving concrete recommendations for preparing, designing, and executing a quantum-safe migration [i.5]. Each of these documents provide their own insight and recommendations for performing a quantum-safe migration. They have commonalities, but they also have differences.

Any framework for a quantum-safe migration is going to be, by nature, not a one-size-fits-all solution. Frameworks are designed to be flexible, can be scoped and tailored as needed, and usually do not attempt to address every possible detail. It is the responsibility of the entity implementing the framework to address any gaps. For quantum-safe migrations it is simply not feasible to create a highly detailed document addressing every single consideration in a way that is directly applicable to an arbitrary enterprise. Therefore, enterprises are encouraged to leverage several resources to create the most appropriate migration strategy for themselves. The present document aims to complement existing guidance by proposing a more detailed methodology for determining the order in which to migrate an enterprise's assets, including selecting the solutions the assets are migrated to, and recommending an iterative risk-based methodology for performing those migrations.

By defining a prioritized order in which to migrate the enterprise's assets and by taking guidance from several reputable sources, the enterprise can become well-positioned to plan and execute their quantum-safe migration. The framework described herein follows a divide and conquer-type strategy by creating individual migration plans for the discrete structural elements of the enterprise and refining those plans through various analyses. Moreover, the framework is iterative, in the sense that it can be re-run each year (or after whichever length of time the enterprise prefers). There are several reasons for taking an iterative approach. The primary reason is because it is generally not feasible to migrate all an enterprise's assets in one step. This can be due to numerous causes, such as budget constraints, system dependencies, requirements for system availability, technological limitations, supply chain limitations, the need for certification or validation, lack of standardization, constraints on employees' time and expertise, and so on.

The result of this framework is a sustainable approach for executing an enterprise-wide quantum-safe migration over time. Although not discussed in detail within the present document, much of the information gathered and analysed through this framework can additionally serve as inputs to other processes within the enterprise, including but not limited to the more general change, risk, and vendor management programs.

1 Scope

The present document describes a repeatable divide and conquer-style framework for migrating, in a prioritized order, an enterprise's information security assets from quantum-vulnerable states to quantum-safe states. First, the approach gives recommendations for partitioning the enterprise into discrete elements. Following, through various analyses within and between the elements of the partition, a methodology is described for establishing quantum-safe migration plans for each of those partition elements.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] Global Risk Institute: "2023 Quantum Threat Timeline Report".

[i.2] ETSI TR 103 619: "CYBER; Migration strategies and recommendations to Quantum Safe schemes".

[i.3] Department of Homeland Security: "Preparing for Post-Quantum Cryptography: Infographic".

[i.4] ASC X9 IR-F01-2022: "Quantum Computing Risks to the Financial Services Industry".

[i.5] T. Attema, J. Duarte, V. Dunning, M. Lequesne, W. van der Schoot, and M. Stevens: "The PQC Migration Handbook", 2023.

[i.6] ETSI GR QSC 004: "Quantum-Safe Cryptography; Quantum-Safe threat assessment".

[i.7] ETSI TR 103 967: "CYBER; Quantum-Safe Cryptography (QSC); Impact of Quantum Computing on Symmetric Cryptography".

[i.8] World Economic Forum: "Quantum Readiness Toolkit: Building a Quantum-Secure Economy".

[i.9] IETF RFC 8446: "The Transport Layer Security (TLS) Protocol Version 1.3", August 2018.

[i.10] Recommendation ITU-T X.509: "Information technology - Open Systems Interconnection - The Directory: Public key and attribute certificate frameworks".

[i.11] NIST SP 800-207: "Zero Trust Architecture".

[i.12] National Cyber Security Centre: "Next steps in preparing for post-quantum cryptography".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

asset: resource controlled by the enterprise as a result of past events and from which future economic benefits are expected to flow to the enterprise

asset migration: act of changing the cryptography of an asset

asset migration planning: process of planning an asset migration

cryptographic asset: non-human asset that performs cryptographic operations

cryptographically protected asset: asset which has cryptographic operations performed on it

dependency cycle: cycle in a dependency digraph, occurs when an asset is directly or indirectly dependent on itself

dependency digraph: directed graph whose nodes correspond to assets and whose edges indicate dependencies between assets

fully migrated asset: asset which has been migrated to the end-state identified in its asset migration plan

quantum-safe migration: act of migrating an enterprise's cryptographic and cryptographically protected assets to quantum-safe states

quantum-safe state: state of an asset wherein the cryptography used by or on it is quantum safe

migration conflict: any situation wherein a department's asset cannot be migrated according to the order suggested by its Department Migration Priority Report

migration interval: period wherein the migration plans of a given framework iteration are performed

migration period: period starting from the enterprise's initial asset migration and lasting until all assets in the enterprise asset inventory have been fully migrated

NOTE: The enterprise's migration period includes every iteration of this framework.

migration priority vector: ordered vector of migration priority levels of the nodes of a component of a department's dependency digraph

migration requirement: requirements for migrating an asset. Includes technical and non-technical considerations

3.2 Symbols

For the purposes of the present document, the following symbols apply:

D_i	The i^{th} department of the enterprise partition
M	The total number of departments in the enterprise partition
$a_{i,j}$	The j^{th} asset of the i^{th} department
G_i	The dependency digraph for department D_i
G_i^k	The k^{th} component of G_i
$a_{i,j}^*$	The fully migrated version of $a_{i,j}$
$a_{i,j}'$	The backwards compatible migrated version of $a_{i,j}$
p_i^k	The migration priority vector for G_i^k
X	The "shelf-life" variable in Mosca's *XYZ* Theorem
Y	The "migration time" variable in Mosca's *XYZ* Theorem
Z	The "collapse time" variable in Mosca's *XYZ* Theorem
$X_{i,j}$	The "shelf-life" variable for $a_{i,j}$
$Y_{i,j}$	The "migration time" variable" for $a_{i,j}$

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AES	Advanced Encryption Standard
AI	Artificial Intelligence
API	Application Programming Interface
CA	Certificate Authority
CBOM	Cryptography Bill Of Materials
CMDB	Configuration Management DataBase
COTS	Commercial Off-The-Shelf
CRM	Customer Relationship Management
CRQC	Cryptographically Relevant Quantum Computer
CVE	Common Vulnerabilities and Exposures
DHS	Department of Homeland Security
EOL	End-Of-Life
EOS	End-Of-Support
FQSCS	Fully Quantum-Safe Cryptographic State
IT	Information Technology
ITAM	Information Technology Asset Management
KEM	Key Encapsulation Mechanism
KPI	Key Performance Indicator
MAC	Message Authentication Code
MFA	Multifactor Authentication
NCSC	National Cyber Security Centre
OT	Operational Technology
PEP	Policy Enforcement Point
PII	Personally Identifiable Information
PSK	Pre-Shared Key
RSA	Rivest Shamir Adleman
SBOM	Software Bill Of Materials
SLA	Service Level Agreement
TLS	Transport Layer Security
VPN	Virtual Private Network
ZT	Zero Trust
ZTA	Zero Trust Architecture

4 Framework summary

Table 1 provides a high-level summary of the framework. The outputs of each Step shown in Table 1 are examples only, additional information can be included as desired. Regardless, the exact outputs and their content and formatting are left to the discretion of the framework implementor.

Table 1: Migration framework summary

Step	Purpose	Output
1) Identify Enterprise Architecture	To conceptually partition the enterprise into distinct components to enable migration planning for each component.	Enterprise Partition
2) Generate Asset Inventories	To produce a complete list of assets within each component of the Enterprise Partition.	Enterprise Asset Inventory
3) Dependency Analysis	To examine internal dependencies among the assets of each component of the Enterprise Partition.	Enterprise Dependency Digraph
4) Vulnerability Analysis	To examine the vulnerabilities of each asset and to collect an initial list of potential mitigating solutions.	Enterprise Vulnerability Report
5) Cross-Department Analysis	To augment the dependency analyses from Step 3 by considering asset dependencies between different components of the Enterprise Partition.	Enterprise Cross-Analysis Report
6) Migration Requirements Analysis	To produce initial requirements for migrating each asset.	Enterprise Migration Requirements Report

Step	Purpose	Output
7) Migration Risk Analysis	To perform a risk analysis for each asset and to select candidate solutions for which to migrate each asset.	Enterprise Migration Risk Report
8) Initial Priority Analysis	To compute a risk-based migration priority level for each asset, assuming the assets are migrated to the solutions identified in Step 7.	Enterprise Migration Priority Report
9) Generate Migration Plans	To construct migration plans for each component of the Enterprise Partition by identifying and resolving any issues preventing assets from being migrated to the solutions identified in Step 7 and in the order computed in Step 8.	Enterprise Migration Plan
10) Execute Migration Plans	To execute the migration plans (for the current migration interval) constructed in Step 9 for each component of the Enterprise Partition.	Enterprise Migration Status Report
11) Prepare for Next Iteration	To compile lessons-learned from the completed migration interval, note relevant changes and events which occurred during the migration interval, and to otherwise prepare for the next iteration of the framework.	Enterprise Migration Status Report

5 Preliminaries

5.1 Motivation

Primarily, the present document is motivated by the need for enterprises to address the cybersecurity threat posed by CRQCs. However, quantum computers are not the only motivating factor for enterprise cryptographic migrations. Often, enterprises do not have a good understanding of where or how they consume cryptography. Further, many enterprises lack the internal resources to evaluate the cryptography they are aware of (e.g. in terms of security models, algorithm parameters, modes of operations, appropriate primitives, and other configuration considerations). Largely, this can be attributed to the relative stability of cryptographic algorithms as compared to other technologies deployed throughout a modern enterprise. Moreover, most cybersecurity incidents are not due to attacks against cryptographic algorithms and protocols themselves. Together, these factors have historically created a disincentive for enterprises to rigorously inventory, monitor, and assess their cryptographic usage.

Unfortunately, as threat actors become more sophisticated, as Artificial Intelligence (AI) becomes more powerful, as equipment and technologies age, and as quantum computing advances, the potential likelihood and impact of cryptographic attacks only increase. If enterprises continue to ignore their cryptography, or take it for granted, then the consequences can be significant.

Other benefits can be realized from performing a quantum-safe migration besides increased security and peace of mind. For example, it demonstrates that the enterprise takes security seriously, not only for itself, but for its customers, clients, and partners as well. In an era of evermore frequent cyberattacks, such actions can yield competitive advantages against enterprises who do not adopt next-generation security.

Many organizations are already performing Information Technology (IT) modernization activities, either for similar reasons as described above, or due to changing regulatory or other compliance requirements. Efficiencies can be gained by performing a quantum-safe migration concurrently with those other IT modernization efforts (see clause 5.4), such as the adoption of Zero Trust (ZT). Finally, by planning and executing a quantum-safe migration in the manner described herein, technology switching costs can be reduced or amortized.

5.2 Background

The present document does not describe how or why classical cryptography can be vulnerable to quantum-aided attacks. For example, no description is provided herein for Shor's or Grover's Algorithms, no commentary is provided on the quantum security of classical cryptosystems, and no recommendations are given on symmetric key lengths or hash function outputs. A primer on such information can be found in [i.4] and [i.6]. A detailed analysis of quantum computing's impact on the security of various symmetric algorithms and primitives can be found in [i.7].

Although the present document describes a framework for an enterprise cryptographic migration, it is important to understand that a successful enterprise quantum-safe migration will require the input of various stakeholders, not just that of cryptographers. Different Steps of this framework can require specialized expertise to perform, such as compliance, operations, risk management, procurement, IT, Operational Technology (OT), and so on. It is the responsibility of the enterprise implementing this framework to ensure that appropriate personnel are assigned to each Step.

The present document makes a distinction between migration and migration planning. Asset migration is the act of changing the cryptography of an asset (i.e. the cryptography performed by the asset, or the cryptography performed on the asset). Migration planning is the process of planning some number of asset migrations; the planning efforts done before the migrations are performed. Quantum-safe migrations and quantum-safe migration planning refer to asset migrations and migration planning where the cryptography the assets are migrated to, or are planned to be migrated to, is quantum safe.

NOTE: Implicitly, asset migration includes wholly replacing assets instead of only applying patches or other updates. Concretely, if an asset is replaced with a solution that performs or consumes different cryptography, then that asset is also said to be migrated. Moreover, asset migration does not include the decommissioning of an asset. That is, if an asset has reached the end of its lifecycle, and is not replaced with a new solution, then that asset has not been migrated. However, decommissioning of assets can be a valid action within a migration plan.

An asset is fully migrated when it has been migrated to the end-state identified in its migration plan. For example, if an asset was planned to be migrated from signing and verifying RSA-2048 signatures to RSA-4096 signatures, then the asset is fully migrated when it can sign and verify RSA-4096 signatures. If the asset were instead upgraded to RSA-3072 (as an intermediary solution), then the asset would have been migrated, but not fully migrated. In both cases, the migration would not be quantum safe.

The critical reason for distinguishing between asset migration and migration planning is that the migration planning comes before the actual migration. Unfortunately, enterprises often delay their migration planning because they feel they do not need to begin their asset migrations for some time, or because asset migrations are interpreted to be lower priority than other activities. However, a risk-based determination of when an asset should be migrated to a quantum-safe state cannot reasonably be made until at least some of the migration planning has been performed. Hence, while certain enterprises can indeed delay their asset migrations for some time, the present document strongly encourages all enterprises to begin their migration planning as soon as possible.

5.3 Assumptions

While the present document includes discussion of both technical and non-technical aspects of quantum-safe migrations, certain aspects are assumed to have been addressed prior to an enterprise's implementation of this framework.

It is assumed that some entity within the enterprise has been given ownership of the quantum-safe migration and the quantum-safe migration planning, be it an individual person, a committee, or some other group. It is further assumed that all necessary awareness and knowledge raising has been done, that personnel are sufficiently aware of the quantum computing threat to cryptography, adequate training has been provided, and that there is governance-level support for the quantum-safe migration efforts. Recommendations for aligning the governance structure to the quantum-safe migration can be found in [i.8]. Moreover, in a quantum-safe migration, assets which are not cryptographic nor cryptographically protected can require changes as well (see clause 6.2). For example, human assets can require re-training or upskilling on new quantum-safe equipment, technologies, and processes. Such considerations, while of critical importance, are outside the scope of the present document.

The framework is described herein as a series of linear Steps. In practice, some of these Steps can be combined or performed in an alternative order, depending on the specific needs of the implementing enterprise. Indeed, it is expected that this framework be modified, tailored, and scoped for the specific needs of the implementing enterprise. That is, a general framework is described, and it is expected that profiles of this framework be used in practice rather than the general framework itself.

The framework is designed to be iterated. However, the Steps are described from the perspective of a first iteration. The Step descriptions implicitly assume that it is the first time the Step has been performed. In practice, if a Step has been performed previously, then the output of that Step from the previous iteration can be updated rather than recreated from scratch, to avoid unnecessary duplication of efforts.

Certain inputs and outputs from the various Steps of this framework can be sensitive in nature, and it is expected that the implementing enterprise takes appropriate actions to protect that information. For example, by classifying and labelling the data and using controls to limit who can access, read, or modify the data.

Some Steps of this framework require communication and data sharing across components of the enterprise. It is assumed that appropriate controls are in place, that data is shared according to the enterprise's policies, and that information is only accessed, readable, or modifiable by appropriate entities.

5.4 Program, plan, and process alignment

The framework described in the present document can benefit heavily from programs, plans, and other processes already being carried out within the enterprise. For example, if the enterprise maintains vendor, risk, or change management programs, these programs can be helpful resources for completing various steps of this framework. Similarly, if the enterprise has certain roadmap plans (e.g. for technological updates, enterprise structural change, or other IT modernization initiatives) then efficiencies can be gained by aligning the migration plans developed herein to those already existing plans. Conversely, it is possible that certain Steps of this framework can be used to augment the enterprise's current programs, plans, and processes. These points are reiterated several times throughout the present document but are emphasized here.

6 Framework description

6.1 Step 1: Identify enterprise architecture

Input:

- None if this is the first framework iteration
- Else, the Enterprise Migration Status Report from Step 11 of the previous framework iteration

An enterprise can comprise various departments, divisions, and other structural units. If the entire enterprise is to be migrated to a quantum-safe state, then necessarily each of the units that compose the enterprise are to be migrated as well. Therefore, to enable a divide and conquer-style approach to the quantum-safe migration, this framework partitions the enterprise into discrete structural units and gives recommendations for planning a quantum-safe migration within each of those units, where those migration plans are supplemented by cross-analyses with other enterprise units (e.g. by considering dependencies between enterprise units, such as shared assets and workflows).

For simplicity, the present document assumes the structural units selected are the departments of the enterprise. However, the enterprise may select another type of partition if desired, such as by networks or physical geographies. Regardless of the units selected, the language of "departments" is used throughout the present document.

NOTE: Although the present document attempts to be agnostic to the way an Enterprise Partition is constructed, no guarantees can be made of the suitability of the framework, as presented herein, to an arbitrary Enterprise Partition. If the enterprise is partitioned into units other than departments, special care should be taken to ensure the framework Steps are suitably modified, if required, to fit the chosen Enterprise Partition.

Hence, the first step of the framework is to identify the structure of the enterprise and assign a label to each department identified. The i^{th} department is denoted D_i, the total number of departments is denoted M, and the set of resulting departments if referred to as an Enterprise Partition. An Enterprise Partition can be represented graphically in a chart. It can also be helpful to produce an Organization Chart of the enterprise, depicting key personnel in each department and their roles.

EXAMPLE 1: The enterprise identifies 6 departments: Legal, Administrative, Sales and Marketing, Finance, Research and Development, and Manufacturing. The enterprise assigns these departments the labels of D_1, D_2, D_3, D_4, D_5, and D_6, respectively. Here, $M = 6$.

EXAMPLE 2: The enterprise identifies the same 6 departments as above but decides to sub-divide Sales and Marketing into two distinct units, one for Sales and another for Marketing. In this case, the enterprise can assign the label D_1 to Legal, D_2 to Administrative, D_3 to Sales, D_4 to Marketing, D_5 to Finance, D_6 to Research and Development, and D_7 to Manufacturing. Here, $M = 7$.

Instead of performing a quantum-safe migration of the entire enterprise, it is possible to apply this framework to a proper subset of all departments (i.e. to N departments, where $1 \leq N < M$). However, special care should be taken to ensure that non-migrated departments are not negatively affected by the migration of the other departments. For example, multiple departments can share network resources. By making quantum-safe updates to a network, the operations of a non-migrated department can be disrupted.

Output:

- An Enterprise Partition

6.2 Step 2: Generate asset inventories

Input:

- An Enterprise Partition

Repeat for each department in the Enterprise Partition.

In business generally, an asset can be defined as a resource controlled by the enterprise as a result of past events and from which future economic benefits are expected to flow to the enterprise. This definition encompasses both tangible and intangible assets (e.g. hardware vs electronic data), as well as human and non-human assets (e.g. an employee vs. a piece of Property Plant and Equipment). For the purposes of a quantum-safe migration, the present framework only considers non-human assets that perform cryptographic operations, or which have cryptographic operations performed on them. Such assets are referred to as cryptographic assets and cryptographically protected assets, respectively.

An example of a cryptographic asset is a TLS server [i.9], and an example of a cryptographically protected asset is encrypted data. An X.509 digital certificate [i.10] can be considered as a cryptographically protected asset, and any public and private key pair corresponding to that certificate can be considered cryptographic assets. Notably, some assets can be both cryptographic and cryptographically protected. The present document includes recommendations for migrating both cryptographic and cryptographically protected assets.

EXAMPLE: A Policy Enforcement Point (PEP) in a Zero Trust Architecture (ZTA) [i.11] can cryptographically authenticate an access request to a database whose columns are encrypted. The user can then use their secret information to decrypt the database columns and make changes to the plaintext data. Here, both the PEP and the user's device can be considered cryptographic assets, whereas the encrypted data in the database are cryptographically protected assets.

In Step 2, the department performs an inventory of its assets. Details on how an asset inventory is obtained is outside the scope of the present document. Recommendations for building a cryptographic asset inventory can be found in [i.2], [i.4], and [i.5]. However, if the enterprise maintains things such as Software Bills Of Materials (SBOMs) or Cryptography Bills Of Materials (CBOMs), these can aid in the asset inventorying process, as well as in later Steps of this framework (such as Steps 3, 6, and 8). The resulting inventory should distinguish between assets the department wholly owns and assets it shares with other departments, but both kinds should be included in the resulting inventory. The inventory should also note when an asset is cryptographic, cryptographically protected, or both. If available, each asset in the inventory should include a label for its respective system or data classification level. As mentioned in clause 5.3, due to the potentially sensitive nature of the inventoried assets, care should be taken to ensure they are appropriately represented in the inventory, and that the inventory itself is suitably controlled.

For each department, a Department Asset Inventory is produced. The collection of all Department Asset Inventories is referred to as an Enterprise Asset Inventory.

It is helpful to order and uniquely label the assets for each department. For the remainder of the present document, the notation $a_{i,j}$ is used to denote the j^{th} asset of department D_i, where the choice of asset order is left to the department.

Although $a_{i,j}$ serves as a label of the j^{th} asset of the i^{th} department, the asset inventory can include more than just identifying information about the assets. For example, $a_{i,j}$ can include any additional information about the asset collected during the inventorying phase, such as information about any systems the asset is a part of or resides within, information on the asset's manufacturer, or information about the owner of the asset. This additional information can be helpful when performing later Steps of this framework. As the migration planning will most likely require communication and collaboration with asset vendors and suppliers, it can be helpful to include information about vendors and suppliers in this Step, if such information is not already managed through some other process (such as a vendor management program).

At the discretion of the department, certain assets can be excluded from the Department Asset Inventory. For example, if an asset is scheduled to be decommissioned, or will otherwise reach the end of its lifecycle and not be replaced (migrated) during the current iteration of this framework, then it can be reasonable to exclude the asset from the inventory. Another example could include assets whose values are expected to reach sufficiently close to zero during the current framework iteration. If new assets are added to the department after the completion of Step 2 but before the next iteration of this framework, then those assets should be added to the department asset inventory immediately. If new assets are added, then a partial reanalysis of each completed Step should be done to understand and record any impacts caused by the introduction of the new asset.

The decommissioning of an asset can be handled through the department's regular change management processes. This includes managing and resolving any dependencies between the to-be-decommissioned asset and any other assets of the department. However, the decommissioning of an asset can have second-order effects on other assets. If such plans are made, they should be considered during the migration planning process.

Assets which have already been migrated in previous iterations of this framework should still be included in the Department Asset Inventory, as they can still have dependencies with non-migrated assets which can impact the migration planning.

Finally, the ongoing maintenance of an Enterprise Asset Inventory can be useful for the enterprise for regular change and risk management activities, as well as for enacting future cryptographic migrations.

Output:

- An Enterprise Asset Inventory

6.3 Step 3: Dependency analysis

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory

Repeat for each department in the Enterprise Partition.

If an asset is migrated, then interoperability can be lost between that asset and assets with which it has dependencies. For example, if an authentication system is migrated to understand how to process certain quantum-safe signatures and can no longer verify RSA signatures, then any credentials signed using RSA can no longer be authenticated by that system. Understanding such dependencies between assets is critical for developing a quantum-safe migration plan.

A distinction is made between direct dependencies and indirect dependencies. For example, System B is directly dependent on System A if System B takes as input some output of System A, or where the operation of System B directly relies on the operation of System A. System B is indirectly dependent on System A if there exists at least one intermediate system, say System C, separating Systems A and B. There are other ways in which direct or indirect dependencies can exist, further examples are given below.

EXAMPLE 1: (Direct dependency) A Root Certificate Authority (CA) signs and issues a public key certificate to an Intermediate CA. If the Root CA's signing certificate is revoked due to a quantum-capable attacker recovering the associated private signing key, then the Intermediate CA's signing certificate immediately becomes untrusted as well. In this case, the Intermediate CA certificate is directly dependent on the Root CA certificate.

EXAMPLE 2: (Indirect dependency) As in the above, a Root CA issues a signing certificate to an Intermediate CA. Now, the Intermediate CA issues a certificate to one of the enterprise's TLS servers. Users communicating to that TLS server will use the server's certificate to authenticate the server and to establish TLS sessions. The authentication of the server is partly done by verifying the Intermediate CA's signature on the server's certificate and verifying the Root CA's signature on the Intermediate CA's certificate. Here, if the Root CA's certificate is no longer trusted, then the user cannot establish trust with the TLS server. In this way, the TLS server's certificate has an indirect dependency with the Root CA certificate.

EXAMPLE 3: A software signing server issues a software update to a firewall, where that firewall separates the department's Customer Relationship Management (CRM) system from the rest of the network. Because customer data is sensitive information, the CRM requires user authentication before allowing access. Therefore, for a user to access the CRM, they need to be granted access by the firewall as well as have their credentials authenticated by the CRM. Here, the firewall is directly dependent on the software signing server, the CRM is directly dependent on the firewall, and the CRM is indirectly dependent on the software signing server. Observe that in this situation, the firewall is unable to receive new updates if the software signing server is brought offline. However, not being able to receive an update does not necessarily stop the firewall from working. Thus, the operation (availability) of the CRM is not necessarily affected.

In example 3, the software signing server likely does not belong to the enterprise. Meaning, the server is likely not included within the Enterprise Asset Inventory. This illustrates that enterprise-owned assets can have dependencies with third-party assets outside of the enterprise's direct control. Even though the enterprise does not have the direct ability to migrate such assets, it can be helpful to note such dependencies during this Step, if feasible. Third-party dependencies are explicitly considered in Step 5.

In Step 3, for each asset in the Department Asset Inventory, the department compiles a list of all other assets in that inventory the asset is directly dependent on. The aggregated results can be transformed into a directed graph whose nodes are the department's assets and whose directed edges show the dependencies between assets. Such a graph is called a Department Dependency Digraph, denoted G_i for department D_i. The collection of all Department Dependency Digraphs is referred to as an Enterprise Dependency Digraph. An example dependency digraph is shown in Figure 1.

Dependency digraphs are only one possible way to track dependencies among assets. They are used within the present document to simplify the description of the migration planning process. In practice, an enterprise can already have in place a method to detect and track dependencies between assets. Further, the enterprise can choose to use a method other than dependency digraphs, if desired. For example, relevant information about asset dependencies can possibly be gathered from Configuration Management Databases (CMDBs), IT Asset Management (ITAM) systems, cryptographic discovery tools, or from other established processes relating to IT risk, or change, management.

NOTE 1: A dependency digraph can be represented visually, but in practice the department can have too many assets and dependencies to effectively visualize. Whichever way the data is represented, the present document refers to the data as a dependency digraph.

Importantly, a dependency digraph can have separate components. I.e. there can exist assets within the same department that have no direct or indirect dependencies between them. In such cases the updates to assets in one component are independent of those in other components. Consequently, the digraph components can be migrated independently of each other. The department should identify each component of its dependency digraph. The k^{th} component of department D_i's dependency digraph is denoted G_i^k, so that $G_i = \{G_i^k\}_k$. The ordering of the digraph components is left to the discretion of the department.

NOTE 2: It is possible that some of the department's assets can have direct or indirect dependencies with assets of other departments. For example, if an asset is shared between departments, or due to certain workflows between departments. Such information is also important for migration planning but is considered in Step 5 of this framework.

As shown in Example 3 above, if two assets belong to the same digraph component but one is not directly dependent on the other, then updates to one of the assets does not affect the other asset. In Figure 1, asset $a_{i,8}$ is indirectly dependent on $a_{i,1}$. Although the two assets are connected by $a_{i,4}$, it is not necessarily the case that migrating $a_{i,1}$ disrupts the operation of $a_{i,8}$.

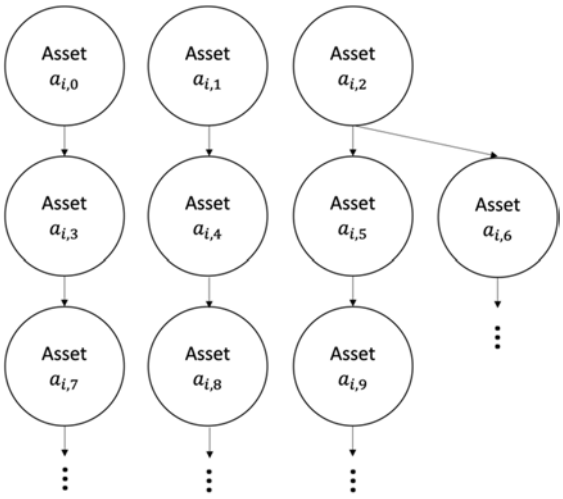

Figure 1: Example dependency digraph for department D_i

In Figure 1, an arrow pointing to a node indicates that the corresponding asset is dependent on the asset corresponding to the node at the tail of the arrow. For example, asset $a_{i,6}$ is (directly) dependent on asset $a_{i,2}$. If there is no path (ignoring the arrow directions) from node $a_{i,0}$ to $a_{i,1}$, then asset $a_{i,0}$ and $a_{i,1}$ belong to separate components of the digraph. Consequently, the corresponding assets in either component can be migrated independently of those in the other component.

In the simplest presentation, an arrow between digraph nodes serves as a binary indicator for whether one of the nodes (assets) is directly dependent on the other. In practice, additional information about the nature of the dependencies can be included in the digraph. This is similar to how $a_{i,j}$ serves as a generic label for the j^{th} asset of the i^{th} department, but additional information can be included if desired, as discussed in Step 1.

Output:

- An Enterprise Dependency Digraph

6.4 Step 4: Vulnerability analysis

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Dependency Digraph

Repeat for each department in the Enterprise Partition.

In Step 4, the department produces a collection of reports describing the cryptographic characteristics and vulnerabilities of each asset identified in its Department Asset Inventory together with a list of potential solutions for mitigating the vulnerabilities. For each asset, an Asset Vulnerability Report is produced. The collection of all a department's Asset Vulnerability Reports is called a Department Vulnerability Report. The collection of all Department Vulnerability Reports is referred to as the Enterprise Vulnerability Report.

Although the precise content and formatting of each report is not defined in the present document, some recommendations are given below. Notably, some of the information relating to the cryptographic characteristics of the assets can have been included in the Department Asset Inventory, as discussed in Step 2.

EXAMPLE 1: If a system produces digital signatures for software updates, then characteristics of the signature algorithms used, including sub-routines and parameters, can be included in the Asset Vulnerability Report for that system (asset). If the asset is a TLS-based Virtual Private Network (VPN), then the cryptographic information recorded can include information on the TLS versions and configurations.

EXAMPLE 2: If a system produces digital signatures for software updates, then a vulnerability can be the forging of signatures for illegitimate software updates. If the asset is a TLS-based VPN, then the vulnerability can be in the use of classical cryptographic algorithms for key establishment or authentication.

NOTE 1: Although the present document is primarily concerned with addressing the quantum-vulnerabilities of assets, non-quantum vulnerabilities can also be included in this analysis and the corresponding migration planning.

Table 2 provides example questions and considerations for assessing the cryptographic characteristics and vulnerabilities of an asset.

Table 2: Considerations for assessing vulnerabilities

Questions	Considerations
What is the basic purpose of the cryptographic operations performed by or on the asset?	• Confidentiality: – E.g. data encryption. • Authentication: – E.g. via passwords, Multifactor Authentication (MFA), or verification of other credentials. • Key establishment: – E.g. using a Key Encapsulation Mechanism (KEM), a key transport or exchange algorithm, or a technique for distributing Pre-Shared Keys (PSKs). • Integrity check: – E.g. verifying hash fingerprints. • A combination of the above: – Digital signatures or Message Authentication Codes (MACs).
More specifically, what is the use case or application for the cryptography used by or on the asset?	• For encrypting data-in-motion: – E.g. between an internal server and employee devices. • For encrypting data-at-rest: – E.g. a database, or sensitive documents. • For controlling access to specific resources: – E.g. a database, or sensitive documents. • For signing or issuing credentials: – E.g. for users or devices. • For verifying data integrity: – E.g. for logs, communications, or certain documents (such as legal agreements). • For data origin authentication: – E.g. between an internal server and employee devices.
Are the cryptographic operations symmetric or asymmetric?	• Are the asymmetric operations known to be quantum-vulnerable? – E.g. to Shor's Algorithm. • Are the symmetric primitives known to be vulnerable to classical cryptanalysis? – E.g. SHA-1 or MD5, or blockcipher modes of operations such as ECB or CBC. • Are the symmetric primitives believed to offer sufficient security in the face of Grover's Algorithm? – E.g. if using 128-bit hash outputs, does the enterprise believe the security level provided is adequate for the asset's intended purpose?
Are the cryptographic operations employed directly, or are they part of a higher-level protocol?	• Is AES-256 being used to directly encrypt data, or is a symmetric key first established through a protocol such as TLS? – Who has provided the relevant protocol code, cryptographic libraries, or supporting hardware?

Questions	Considerations
Are the results of the cryptographic operations static or ephemeral?	• If the operation establishes a shared key, is that key used for a single session, or is it expected to be used for an arbitrary number of sessions? How are the keys managed?
For the algorithms, protocols, or libraries used by the asset, what versions and what parameters are employed?	• Does the asset use obsolete, deprecated, or outdated versions of protocols, libraries, etc.? • If so, what are the reasons for the use of non-current versions? • Are the cryptographic primitives, parameters, or configurations up to date with current recommendations? • Does the enterprise have a change management program, and if so, is it being applied to the asset?
Are there Common Vulnerabilities and Exposures (CVEs), or similar issues, known to be associated to the asset, or any other hardware or software component of the system it resides in?	• Have the vulnerabilities already been recorded and mitigated? • If not, why not?
For any algorithm or protocol implementations, and other relevant libraries, to what extent have they been analysed and tested for security?	• Have implementations undergone rigorous quality assurance or other testing, such as static or dynamic code analysis, formal verification, side channel analysis, and so on? • If so, who performed the tests, when, and what were the results? • Is there any other relevant information about these tests which can be recorded?
Is there other supporting information which can be included?	• What are the permitted cipher suites or key types? • What are the relevant cryptographic parameters, primitives, or sub-routines? • What are the relevant sizes of keys, ciphertexts, or signatures permitted or used by the asset? • What are the asset's requirements for bandwidth, latency, power consumption, packet sizes, memory, storage, and so on?
Do the cryptographic operations employ hardware acceleration, or have other hardware dependencies?	• If a machine offers hardware acceleration for a hash function, what is the quantum-security of that hash function? Is an adequate level of quantum-security provided?
Does the cryptographic functionality come from a third-party or an open-source repository, or is it implemented directly by the enterprise?	• Do the providers of the cryptographic libraries have a good reputation for quickly addressing (critical) vulnerabilities? • Are the open-source libraries still supported and maintained by the original provider, or are they maintained by the enterprise? • For libraries implemented by the enterprise, who has the responsibility or capability of maintaining the library?

The final component of the Asset Vulnerability Report is a summary of the potential solutions to mitigate the identified vulnerabilities. It is possible that multiple solutions are available to mitigate a given vulnerability, and that different solutions mitigate the vulnerabilities to different extents (e.g. decommissioning an asset eliminates all its vulnerabilities but can yield unwanted second-order consequences). The enterprise should make a list of every reasonable solution per vulnerability. The solution eventually selected can affect the order in which other assets are migrated. In this Step, it is more important to identify solutions than to assess the feasibility or practicality of obtaining or implementing those solutions. In Steps 7 and 9, the enterprise will analyse additional information to determine an appropriate solution for each asset. For recommendations on quantum-safe cryptographic solutions, and other guidance, the reader can find the 2023 whitepaper "Next Steps in preparing for post-quantum cryptography" by the National Cyber Security Centre (NCSC) [i.12] helpful.

The present document recommends that for each asset a solution is chosen which mitigates the associated risk to a level acceptable by the enterprise. Unfortunately, this is not always possible. For example, some assets can rely on cryptographic techniques which do not currently have suitable quantum-safe equivalents. If such situations cannot be reasonably avoided (e.g. by switching to a different technology or decommissioning the asset), alternative solutions can be required, such as obtaining some form of insurance, or simply accepting the risk and taking no action. Such options can be included in the associated Asset Vulnerability Report, if desired. These options are further discussed in clause 6.7, but are mentioned here for completeness.

It is expected that the migration of many assets will largely be done by the asset's manufacturer or supplier. For example, the enterprise can acquire, install, or patch their assets, but the new assets, updated versions of assets, or asset software updates can come from a third party.

Although the present document cannot provide specific solutions for each asset, three generic strategies for asset migration are described below. First, two definitions are introduced. A migration period is taken to be the period starting from the enterprise's initial asset migration and lasting until all assets in the Enterprise Asset Inventory have been fully migrated. The enterprise's migration period includes every iteration of this framework. A migration interval is the period between iterations of this framework. That is, the enterprise's current migration interval is the period wherein the migration plans of the current framework iteration are performed (Step 10). The reason for this distinction is that asset migrations are performed one migration interval at a time. Some assets can be migrated at later intervals than others, whereas all assets are migrated within the migration period.

NOTE 2: It is possible for an enterprise to iterate (some version of) this framework indefinitely, such as by incorporating it into the enterprise's regular risk or change management programs. If so, then the migration period as defined above is also indefinite.

There are at least three distinct types of approaches for migrating an asset: backwards compatible migrations, parallel migrations, and pure migrations. Each type has their own advantages and disadvantages depending on the situation. The three approaches are briefly described below. The reader is made aware that the three approaches described below is not intended to be an exhaustive list of all possible approaches.

Pure Migration: Where the asset is directly fully migrated to the desired end-state. In terms of a quantum-safe migration, a pure migration is where the classical cryptography of an asset is entirely replaced by quantum-safe cryptography. An advantage of this approach is that it eliminates the quantum vulnerabilities of the assets and presents no threat of future downgrade attack. A disadvantage of this approach is that any non-migrated dependent assets will no longer be able to interoperate with the updated asset.

Parallel Migration: Where a separate system is installed in parallel with the current system. In this case, dependent migrated assets can utilize the new system, and non-migrated assets can continue to use the legacy system. Advantages of this approach are that it can create a clean separation between legacy and migrated assets while reducing the potential of downgrade attacks in the future. Moreover, a parallel approach does not necessarily have to be done in a single step; a parallel system can be built over time. Possible disadvantages of this approach include increased costs (time, compute, money, etc.) to simultaneously support the parallel systems.

NOTE 3: There are types of hybrid approaches, sometimes called composite approaches, which are not strictly speaking parallel migration approaches, as that term is defined above. These approaches can be thought of as parallel systems designed as single systems. For simplicity, these composite approaches are included as parallel approaches for the purposes of this framework. However, depending on set policies and system configurations, these approaches can also be considered as backwards compatible approaches.

Backwards Compatible Migration: A type of hybrid approach which maintains backwards compatibility with non-migrated assets. In terms of a quantum-safe migration, non-migrated dependent assets can still utilize the legacy, classical cryptography and protocols, whereas migrated dependent assets can use the quantum-safe components of the hybrid system. An advantage of this approach is the maintenance of interoperability and non-disruption of dependent assets. A disadvantage of this approach is that it can potentially permit downgrade attacks in the future if the legacy components can still be utilized. In this way, backwards compatible solutions can be a convenient intermediary solution, but likely not an acceptable end-state. The backwards compatible migration approach can be useful for systems that are too complex to migrate in a pure fashion, or where building a parallel system is not a viable option, due to costs or other resource constraints.

These three approaches are further examined below.

Consider the simplified digraph component show in Figure 2. Here, $a_{i,1}$ is directly dependent on $a_{i,0}$, and both $a_{i,2}$ and $a_{i,3}$ are directly dependent on $a_{i,1}$.

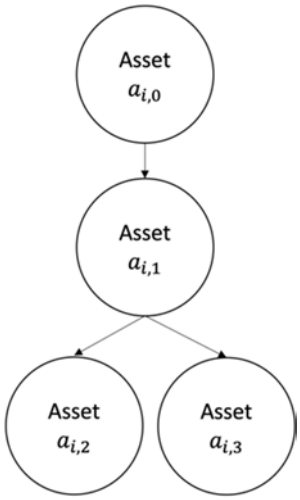

Figure 2: A simplified dependency digraph component for D_i

Pure Migration Approaches

If $a_{i,0}$ is migrated before $a_{i,1}$, then the dependency $a_{i,1}$ has on $a_{i,0}$ can be broken. Meaning that the operation of $a_{i,1}$ can be disrupted. If $a_{i,1}$ cannot operate, then the operations of assets $a_{i,2}$ and $a_{i,3}$ can also be interrupted.

If $a_{i,2}$ is updated before $a_{i,1}$, then the dependency $a_{i,2}$ has on $a_{i,1}$ can be broken. In this case, there are not necessarily any effects on assets $a_{i,0}$ or $a_{i,3}$.

The conclusion is that a sequential pure migration approach-from the top-down, bottom-up, or middle out-has a risk of breaking interoperability and causing significant disruption. This issue can be avoided by performing a simultaneous migration of all assets in the dependency digraph component. While feasible in some situations, a simultaneous update approach will likely be prohibitively difficult and costly in many real-world situations.

Parallel Migration Approaches

As seen above, if any one asset is purely migrated to a quantum-safe state, then interoperability can be broken. The consequences can include unacceptable disruptions to work-flows and business operations. A parallel migration approach avoids this issue because the legacy solution (e.g. classical cryptography) is still in use and available to any dependent asset requiring it.

A parallel approach does not necessarily have to be done in a single step; a parallel system can be built over time. For example, still considering Figure 2, a quantum-safe version of $a_{i,1}$ can be installed even without a quantum-safe version of $a_{i,2}$. This approach can be costly, as more systems are supported simultaneously and the quantum-safe version of $a_{i,1}$ can be of limited use until quantum-safe versions of the dependent assets are available. However, the cost can be amortized.

In Figure 3, a parallel system is created over three migration intervals. Time $t = 0$ represents the beginning of the first migration interval, and the graph components between the vertical hatched lines depict the state of the parallel system by the end of the indicated migration interval. Here, $a^*_{i,j}$ denotes the migrated (quantum-safe) version of asset $a_{i,j}$. In this example, the legacy (classical) version of the digraph component can be decommissioned after the second interval.

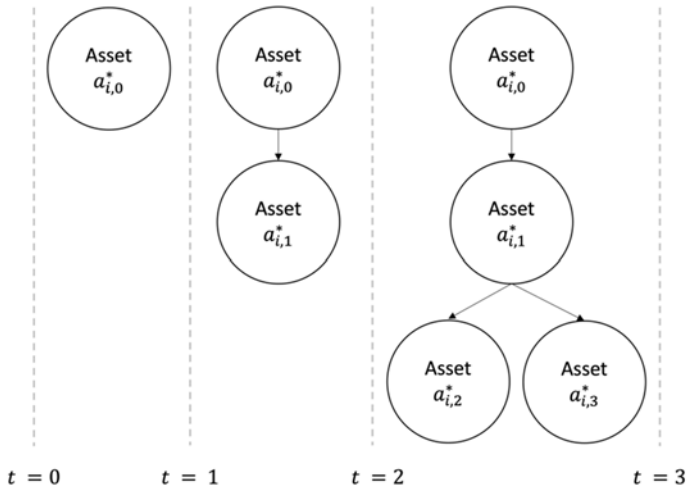

Figure 3: Parallel system installed over three migration intervals

A disadvantage of the phased parallel approach is that before the entire digraph component is migrated, the quantum-vulnerable versions of the assets are still in operation; meaning that the quantum-vulnerabilities have not been mitigated and will not be mitigated until the entire parallel graph component is operational and the legacy component decommissioned. Further, the maintenance of parallel systems can incur significant overhead costs.

Backwards Compatible Migration Approaches

A backwards compatible approach is one which maintains interoperability with non-migrated dependent assets while simultaneously offering updated (quantum-safe) functionality to purely migrated dependent assets. For example, suppose that $a_{i,1}$ has been given a backwards compatible update, $a_{i,2}$ has not been migrated, and $a_{i,3}$ has been fully migrated. In this case, the dependency $a_{i,2}$ has on $a_{i,1}$ is undisturbed, and $a_{i,2}$ can continue to operate as normal. At the same time, $a^*_{i,3}$ (the migrated version of asset $a_{i,3}$) can rely on the updated features of $a_{i,1}$. This situation is depicted in Figure 4, where $a'_{i,1}$ denotes the backwards compatible migrated version of asset $a_{i,1}$.

The backwards compatible approach shares some of the disadvantages with the parallel approach. In particular, the vulnerabilities of the assets are not entirely mitigated. Moreover, this approach is potentially susceptible to downgrade attacks in the future if the legacy components are still in operation. This approach can be less costly in terms of equipment than the parallel approach, as it requires less duplication of assets. However, this approach can be more costly than the parallel approach in terms of the administration and project management activities to handle the mixture of system capabilities. An advantage this approach shares with the parallel approach is that it enables a stepwise migration over time while maintaining the functionality of assets which have not yet been migrated.

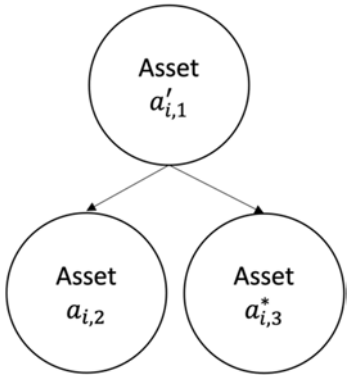

Figure 4: Example backwards compatible migration

The three migration approaches described above are not equally suitable or feasible for every asset. For example, a TLS server can be migrated in a hybrid fashion (e.g. composite or backwards compatible) with relative ease compared to a physical system with cryptographic keys burned into its hardware.

Output:

- An Enterprise Vulnerability Report

6.5 Step 5: Cross-department analysis

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Dependency Digraph
- An Enterprise Vulnerability Report

Repeat for each department in the Enterprise Partition.

In the previous Steps, assets have only been considered within the context of a single department. In practice, assets within one department can have implications for other departments. The goal of Step 5 is to gather relevant information about the asset dependencies between enterprise departments to supplement the information gathered about individual departments in the previous Steps.

For each asset in the department's Department Asset Inventory, a cross-analysis is performed between that asset and the assets of each other Department Asset Inventory in the Enterprise Asset Inventory. The result for the department is a set of Asset Cross-Analysis Reports; one report for each other department. The collection of a department's Asset Cross-Analysis Reports is called a Department Cross-Analysis Report. The collection of all Department Cross-Analysis Reports is referred to as an Enterprise Cross-Analysis Report.

There is a natural redundancy within many of the reports which can be used to simplify the dependency compilation process. For example, consider two departments, D_i and D_j. For each asset in D_i, the Asset Cross-Analysis Report of D_i will contain cross-dependency information about each asset in D_j. This information can be helpful when compiling the Asset Cross-Analysis Reports for D_j. The asset dependency digraphs can also be helpful in mapping out the cross-department dependencies. For example, in Step 2 it was recommended that the asset inventories include, but distinguish between, assets wholly owned by a department and those it shares with other departments. Consequently, when two different departments share an asset, the corresponding dependency digraphs will share nodes. By comparing the digraphs, useful information about department cross-dependencies can be gathered. It is also possible that an asset of one department can have indirect dependencies with assets of a second department, even if the asset is not a member of the second department's asset inventory. Again, a comparison of dependency digraphs can be helpful in detecting and understanding such situations.

In addition to a dependency analysis between departments, an Asset Cross-Analysis Report can include information on external dependencies, if any exist. That is, if an asset of department D_i has dependencies with systems external to the enterprise, then that information can also be included in the Asset Cross-Analysis Report of D_i. Such dependencies can arise from collaborations with, services offered to, or services received from third parties.

Table 3 gives further considerations for examining cross-department relationships.

Table 3: Cross-department considerations

Context between departments	Considerations
The same, or sufficiently similar, asset exists in, or is used by, multiple departments.	• It is possible that the cost of migrating can be reduced if done for both departments simultaneously. • If the asset is shared by multiple departments, special care should be taken to avoid business disruptions, or a loss of interoperability caused by performing updates within one department only, or one department first.

Context between departments	Considerations
Indirect dependencies exist between assets of different departments.	• Migrating an asset in D_i can inadvertently disrupt the functionality of assets within D_j. Such possibilities should be documented and understood before the migration is executed, as much as is reasonable.
Different assets performing similar functions exist in different departments, and the enterprise already has roadmap plans to reconcile the differences.	• If the enterprise already has technological changes planned, then those plans should be noted and accounted for in the migration plan: – E.g. the enterprise plans to either replace one asset with the other or replace both with a third asset.
Planned changes to enterprise architecture	• Departments can be merged, divided, added, or spun-out from the main enterprise. New departments can be added through mergers or acquisitions: – The quantum-safe migration strategy should be reconciled insofar as possible with the other relevant plans and roadmaps of the enterprise.
Workflow dependencies between departments	• If the output of department D_i serves as input to department D_j, issues can arise if the output is cryptographically protected by processes of D_i that D_j is not upgraded to handle: – E.g. digitally signed data is sent from D_i to D_j, but D_i is unable to verify the quantum-safe signature on the data.
Enterprise architectural relationships	• Two separate departments are embedded within the same higher-level structure of the enterprise, and their budgets, or other resources, are joined: – Recall from Step 1 that an implementation of this framework can define "departments" arbitrarily.
Staffing and skills dependencies	• If an employee is cross assigned to multiple departments, they can have relevant insights into how the departments can be migrated.

As mentioned in Step 1, it is possible to apply this framework to a proper subset of the enterprise's departments. If such an approach is taken, special care should be taken to limit any issues or conflicts that arise from excluding departments from the cross-analysis.

The allocation of resources for performing quantum-safe migrations can be done at the department level. However, this is not always going to be the case. For example, the Board of Directors can approve a single budget for the enterprise-wide quantum-safe migration. It is also possible that the enterprise has established a quantum-safe migration steering committee which oversees the quantum-safe migration of the entire enterprise, and consequently has limited resources to allocate to any specific department. In either of these cases, individual departments can have to compete for migration resources.

Output:

- An Enterprise Cross-Analysis Report

6.6 Step 6: Migration requirements analysis

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Vulnerability Report

- An Enterprise Dependency Digraph
- An Enterprise Cross-Analysis Report

Repeat for each department in the Enterprise Partition.

If an enterprise is using a particular asset, then there is a reason for it. Possibly several. Consequently, an asset typically cannot be migrated arbitrarily without causing disruption or other negative effects. To effectively migrate an asset, the enterprise should have a thorough understanding of the reasons the asset is in use and of the changes required before the asset can be migrated.

Step 6 seeks to define initial sets of requirements for migrating each asset. For each of the department's assets, the results of the analysis are included in an *Asset Migration Requirements Report*. The collection of all such analyses for a department is called a Department Migration Requirements Report. The collection of all Department Migration Requirements Reports is referred to as the Enterprise Migration Requirements Report.

Many of the migration requirements can be discerned from the outputs of Steps 2 through 5, such as asset dependencies and application use cases. Table 4 gives examples of additional questions the enterprise can ask during the migration requirements gathering process.

Table 4: Migration requirements gathering questions

Questions	Examples
Why is the asset currently in use? Why is this asset used instead of other options with similar functionalities?	• It is required by law, regulation, contractual agreement, or enterprise policy. • The asset offers specific functionality not reasonably available elsewhere. • Use of the asset is necessitated by the use of other assets. • There is a cost benefit to using this asset over alternative assets. • Due to some other convenience such as employees having pre-existing knowledge or experience with the asset.
What are the technical requirements for the cryptography used by, or on, the asset?	• What are the requirements for bandwidth, latency, key sizes, ciphertext or signature sizes, memory, storage, power consumption, or other relevant metrics? • What are the reasons for these limitations? – Due to currently used hardware, protocol specifications, system dependencies, third-party dependencies, inherent characteristics of the application, etc.
What are the Key Performance Indicators (KPIs) or other relevant metrics relating to the asset?	• How often is the asset used? • What are requirements for system availability? • What is the cost of ownership of the asset? • If the asset was acquired to solve a specific problem, does that problem still exist? Is the asset still required to address that problem?
What is the expected lifetime of the asset?	• Different technical requirements can be had for assets with different expected lifetimes: – An asset that requires protection for 50 years can have drastically different requirements than one requiring protection for only a few hours.
Are there other relevant IT or business migration projects currently being planned or executed within the department?	• Are systems already planned to be upgraded in some way? Is the department expecting to undergo a structural change? – How do such plans impact the asset migration requirements? – How can the plans be reconciled?

Questions	Examples
Does migrating the asset require changes to policy, or updates to existing standards, laws, or regulations?	• Which entities or bodies oversee maintaining the laws, policies, regulations, or standards? • How much influence does the enterprise have in making the relevant updates? • What are the bottlenecks to updating enterprise policy? – The content of the new policy cannot be decided until further research and analysis is conducted, or until after certain other assets are migrated. – Updating policy now means some assets will be non-compliant to the new policy, creating excessive difficulties in asset and policy management.
For each potential solution identified in the corresponding Asset Vulnerability Report, what are the bottlenecks to acquiring or implementing that solution?	• The solution is not yet commercially available due to lack of relevant standards or compliance validation and certification. • As research is ongoing, there is insufficient confidence in the capability or security of the identified solution. • The solution can only be implemented by a specific vendor, but that vendor cannot yet perform the updates due to reliance on other elements of its supply chain, or other industrial, regulatory, or logistical constraints.
For each potential solution identified in the corresponding Asset Vulnerability Report, what other requirements can be identified for acquiring or implementing that solution?	• Some solutions can be more difficult or costly to implement than others or require a different scope of changes be made. • If an initial feasibility assessment can be made for each identified solution, then the department should consider including that information in the Department Migration Requirements Report.

Not every solution can be directly implemented by the enterprise. Some can only be implemented by the supplier or manufacturer of the asset. In such cases, the implementation of the solutions can require coordinating with the relevant third party.

Output:

- An Enterprise Migration Requirements Report

6.7 Step 7: Department migration risk analysis

Input:

- An Enterprise Partition
- An Enterprise Vulnerability Report
- An Enterprise Cross-Analysis Report
- An Enterprise Migration Requirements Report

Repeat for each department in the Enterprise Partition.

The goal of Step 7 is to, for each asset, select specific solutions to the vulnerabilities identified in Step 4 from among the potential solutions identified in the corresponding Asset Vulnerability Reports.

The Asset Vulnerability Reports generated in Step 4 contain a summary of the potential solutions to the asset vulnerabilities also identified in Step 4. In Step 7, the department selects a specific solution for each vulnerability and begins the migration planning process under the assumption that solution will be implemented. The chosen solution may later be changed to another depending on the results of Step 9. For each of the department's assets, the results of this Step's analysis are included in an Asset Migration Risk Report. The collection of all the department's Asset Migration Risk Reports is called a Department Migration Risk Report. The collection of all Department Migration Risk Reports is referred to as an Enterprise Migration Risk Report.

There are 4 ways in which an entity can react to a given risk. The risk can be mitigated, accepted, transferred, or avoided. Risk mitigation is when an action is taken which reduces the overall level of risk (e.g. by implementing a security control, updating a system, moving to a stronger cryptographic algorithm, etc.). Risk acceptance is when the entity intentionally takes no action to alter the risk level (i.e. where the entity decides that the risk level is tolerable). Risk transference is when the overall risk is distributed among different entities, thereby reducing the level of risk faced by any one entity (e.g. through insurance). And risk avoidance is when the activity which incurs risk is entirely avoided (i.e., where the entity takes on no risk by simply not engaging in the risk-inducing activity).

For each asset in the Department Asset Inventory, the associated risk is assessed, and a risk reaction selected. For the purposes of a quantum-safe migration, the present document strongly recommends that a (quantum-safe) risk-mitigating solution is selected for each asset, if possible.

NOTE: Risk avoidance is generally not possible for assets already owned or operated by the enterprise; the risk already exists as the risk-inducing activity has already been engaged in. Decommissioning an asset does reduce the risk to zero (ignoring any second-order consequences) but is a risk mitigation action rather than an avoidance action.

When selecting a solution for a given asset, there are two broad types of considerations to be made. Crudely, these can be called technical considerations and business considerations. Intuitively, technical considerations include many of the items addressed in previous Steps of this framework as well as in the risk analysis discussed in the present Step. Business considerations can include the governance, operational, managerial, and administrative considerations not captured by technical analyses. The reader is cautioned that these are not rigorous definitions, nor are they necessarily mutually exclusive of each other. They are nebulous terms and are introduced here for illustrative purposes only.

A technical analysis can shed light on the level and nature of a risk, and the costs and barriers to different risk reactions. Such an analysis can be used to inform and support the risk reaction selected by the enterprise. In practice, it can be the case that business considerations carry more weight in the decision-making process than do technical considerations. That is, the final choice of risk reaction can ultimately be a "business decision". However, that business decision should be informed by rigorous technical analysis whenever possible.

In what follows, methods are discussed for assessing the risk levels of assets and for selecting appropriate mitigating solutions, or other risk reactions.

A key step for selecting an appropriate solution for a given asset is to assess, as much as is possible, the risks associated with that asset. Information security risk assessment is a mature and sophisticated field, and as such, the present document cannot responsibly give a complete treatment of the topic. Depending on its specific needs, the enterprise is encouraged to supplement the guidance given herein with their own preferred methods of risk analysis, possibly including tools and frameworks already in use within the enterprise.

Although methodologies can differ, information security risk typically has two components: the impact of a successful vulnerability exploit and the probability that an attempted exploit will be successful. There are different options for estimating these components and for deriving a risk from them. For example, the enterprise can use a quantitative approach by assigning specific numerical values to the risk components and multiplying those values to produce a numerical risk estimate, i.e. $\mathrm{risk} = (\mathrm{impact}) \times (\mathrm{probability})$. Another option is to use a qualitative approach by assigning levels, or grades, to the risk components, such as Low, Medium, and High (or something more granular, if desired). In the qualitative approach, a vulnerability with, for example, Low impact and Low probability can be interpreted as being lower risk than one with Low impact and Medium probability, or High impact and High probability, etc.

More generally, risk analysis can include a third component: the expected frequency of the event. In this case, risk can be numerically estimated as $\mathrm{risk} = (\mathrm{impact}) \times (\mathrm{probability}) \times (\mathrm{frequency})$. Notably, qualitative levels become difficult to compare when considering a third component. The inclusion of a frequency component can be useful when the enterprise cannot reasonably influence the frequency of the event. For example, if the threat under consideration is an earthquake, tornado, or other natural disaster. The present document assumes that if a cryptographic vulnerability is successfully exploited, then that vulnerability will be remediated in a timely manner. Hence, individual cryptographic vulnerabilities are exploited at most once.

Table 5 gives example questions and considerations for assessing the impact of successful exploits of vulnerabilities. The issues discussed in Table 5 are not strictly limited to quantum computer-aided attacks. Rather, they can be considered another dimension of the enterprise's typical risk assessment and analysis processes. The emergence of a CRQC does not necessarily change the enterprise's risk appetite, but it can change the probabilities of security incidents occurring. Thereby impacting the results of a risk assessment. It is possible that much of the information required for assessing the impact of successful exploits has already been analysed within the enterprise's existing risk management program.

Table 5: Considerations for assessing impacts of exploits

Questions	Consequences
Are there potential reputational harms?	• Sensitive communications are decrypted and leaked to the public, causing embarrassment: – The enterprise is now seen as one that does not take information security seriously, jeopardizing future business. – Other enterprises which have already performed their quantum-safe migrations are now seen as more reputable and reliable.
Are there potential legal or other liabilities?	• A successful quantum-aided attack brings a system offline, where the enterprise is obligated by Service Level Agreements (SLAs) or other contracts to maintain a certain amount of availability (e.g. 99,99 % uptime): – The results can include financial penalties, loss of trust, reputational harm, and loss of future business.
What are the expected costs of disruptions to business operations?	• A critical system is disrupted due to a quantum-aided attack, preventing the enterprise from conducting mission-critical operations until the system can be restored.
What are the other business costs of the exploit?	• A trade secret was disclosed and made publicly available, thereby eroding competitive advantage.
What are the expected effects to the asset?	• Is recovery expected to be possible, or will the asset require destruction and replacement? • If the vulnerability is documented as a CVE, or similar issue, how extensively has it been exploited elsewhere?
What is the asset value?	• A general rule is to not spend more protecting an asset than the asset is worth. Hence, both the value of the asset (possibly as a function of time) and the estimated cost of the solution should be considered.
What are the expected effects to directly and indirectly dependent assets?	• Can the damage be reasonably limited to the exploited asset, or does a successful exploit give the attacker access to other assets and resources? – E.g. through privilege escalation attacks, credential compromise, remote (arbitrary) code execution, etc.
What are the expected effects to other departments?	• Including directly and indirectly dependent assets, disruptions to workflows and business operations, etc.

The next step in the risk analysis is to estimate the probabilities of successful exploits. Table 6 provides some considerations for estimating exploit probabilities. Like the impact factor, it can be difficult to accurately estimate a vulnerability's probability of being exploited. Probability estimates typically require making certain assumptions, such as the resources available to a threat actor, completeness of knowledge about the nature of the vulnerability, and that the enterprise's security controls are implemented and working as expected. Another similarity to the impact factor is that the enterprise can choose to use either a quantitative or a qualitative approach; they can assign numerical probabilities to each exploit or assign each exploit probability a grade such as Low, Medium, or High.

Regarding quantum vulnerabilities, part of the migration priority assessments will include estimating when a quantum computer with sufficient capabilities to exploit the vulnerabilities will be available to threat actors. Such estimates are deferred to Step 8. For the current Step, it is important to estimate what the risks would be if the quantum-vulnerable assets are not migrated before CRQCs become available. In Step 8, questions such as "how long will asset migration take?" and "when will a threat actor have sufficient quantum capabilities to attempt exploits?" are considered. In other words, the enterprise (tentatively) identifies appropriate solutions to mitigate risks in this Step and uses that information along with other timeline considerations in Step 8 to determine the appropriate order in which to perform the asset migrations. In Step 9, non-risk factors are considered and used to update the planned asset migration order.

Again, the considerations given herein can be supplemented, or replaced entirely, by the enterprise's preferred risk assessment methodology.

Table 6: Considerations for assessing exploit success probabilities

Questions	Considerations
What is the asset value to an attacker?	• The value of an asset is not necessarily the same for the enterprise as for a threat actor. Generally, the more valuable an asset is to a threat actor, the more resources they are willing to dedicate to an attempted attack, increasing their success probability.
What are the relevant mathematical and cryptanalytical estimates?	• If the exploit is mathematical in nature, such as the application of Grover's or Shor's Algorithms, then the success probability can be estimated by understanding the algorithm parameters and the associated costs of running the cryptanalytic algorithms.
Is the asset protected by defence in depth? What are the logical attack vectors?	• Can the asset be attacked directly, or can it only be attacked after certain other assets are compromised? – Is the asset protected by a defence-in-depth paradigm? – Is the asset protected under some kind of n-of-m, split-key, or two-person control? – Is the asset in an air-gaped room, protected with strict access controls? – Can the asset be accessed through a public internet-facing Application Programming Interface (API)?
What are the expected origins for attacks?	• From where can the asset be attacked? – E.g. over the public internet, from the corporate network, or from within a certain geographical or physical boundary?
How many attackers are expected to be required for success?	• Can the attack reasonably be executed by a single entity, or are multiple attackers expected?

Once the risks have been estimated, the final part of Step 7 is to select the solutions for each asset.

Table 7 gives some recommendations for how an enterprise can select an appropriate solution for each asset. The considerations in Table 7 should be supplemented by an analysis of the aggregate information obtained from the previous steps of this framework, such as the Asset Vulnerability Reports, Asset Cross-Analysis Reports, and Asset Migration Requirements Reports. Further, it is worth repeating that the following recommendations are examples only and are not considered to be exhaustive. There are potentially many other factors the enterprise can consider when making the determination of which solutions to migrate their assets to.

Table 7: Considerations for selecting solutions

Questions	Considerations
How many vulnerabilities does the asset have?	• It is possible that a combination of solutions is required to mitigate, to an acceptable level, the estimated risks.
What is the acceptable level of risk for the asset, and which solutions reduce the risk to at, or below, that level?	• Just because a potential solution mitigates a risk does not mean that it mitigates the risk enough. The selected solution should reduce the risk to an acceptable level: – E.g. a solution that reduces the risk from High to Medium can be insufficient according to the enterprise's risk tolerance. • There are likely to be cost or performance trade-offs incurred by using solutions that excessively reduce the risk: – E.g. a new cryptographic algorithm (or parameter set) offering 512-bits of post-quantum security where 256-bits is the minimum level required by enterprise policy.
How do the solutions fit in with the other technological roadmaps of the enterprise?	• A solution can use a security approach that is technically sufficient but conflicts with other goals of the enterprise. In general, the migration planning should be reconciled as much as possible with other plans the enterprise has: – E.g. the enterprise wants to move away from certain types of technologies in favour of something else. • The decommissioning of an asset can have second order effects on other assets. If there are plans for such changes, they should be considered during migration planning.
Are the candidate solutions available?	• Are the solutions obtainable as Commercial Off-The-Shelf (COTS) products? • Are proprietary solutions necessary, such as modifications to enterprise-owned code? – Are resources available to make the required changes? • If the solution is still under development, how complex is the relevant supply chain, and is the enterprise aware of the roadmap plans of the suppliers along the chain?
How many distinct solutions are required?	• If the same solution can be used for multiple assets, then it can be more effective to use that same solution repeatedly, rather than use a collection of distinct solutions: – E.g. in terms of costs to acquire multiple solutions, and the skills and knowledge required to implement and maintain the solutions.
Can the enterprise estimate and compare the different lifetime costs of acquiring, implementing, and maintaining the candidate solutions?	• If the solution is a simple policy change (requiring modifications to behaviours, processes, or workflows), then that approach can be more efficient than a technical solution (which can require significant testing, maintenance costs, training, etc., over the solution's lifetime). Of course, depending on the context, the opposite can be true.

Output:

- An Enterprise Migration Risk Report

6.8 Step 8: Initial priority analysis

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Vulnerability Report
- An Enterprise Dependency Digraph
- An Enterprise Migration Requirements Report
- An Enterprise Migration Risk Report

Repeat for each department in the Enterprise Partition.

The goal of Step 8 is to derive an initial migration priority order for the department's assets. The asset migration priorities are derived from the analyses performed in earlier Steps and a combination of three estimates computed in this Step: the required operational lifetimes of the cryptographic assets (or the required protection lifetimes, in the case of cryptographically protected assets), the lengths of time required to migrate the assets to the solutions selected in Step 7, and the lengths of time until quantum computers capable of exploiting the vulnerabilities identified in Step 4 emerge (where applicable).

The migration priority level of an asset is a real number, and the enterprise has discretion in how the value is calculated. Recommendations for calculating migration priority levels are given later in the present clause. The outputs of Step 8 are formulated in terms of the dependency digraphs. The reason for this is that the concrete migration plans for each department will be constructed, in Step 9, from the migration plans made for the separate, independent, digraph components. An implementation of this framework can choose to formulate the outputs of Step 8 in an alternative way, if desired. The purpose of using the dependency digraph terminology is to emphasize that even within a department, certain sets of assets can be migrated without impacting the other assets of the department. These sets of assets are exactly those corresponding to the nodes of the separate dependency digraph components. To create migration plans for all the assets of the department, it suffices to create migration plans for each of the components of the department's dependency digraph. Such plans are created in Step 9.

Recall from Step 3 that G_i is the dependency digraph of department D_i and that G_i^k is the k^{th} component of D_i's dependency digraph. Now, the migration priority vector of G_i^k, denoted p_i^k, is the ordered vector of migration priority levels of the nodes of G_i^k.

EXAMPLE: If G_i^k has three nodes, $\{a_{i,1}^k, a_{i,2}^k, a_{i,3}^k\}$, ordered in some way, and the nodes have respective migration priorities levels of 1, 3, and 5, then the migration priority vector of G_i^k is the 3-dimensional vector $p_i^k = (1, 3, 5)$.

The collection of all migration priority vectors for G_i is called a Department Migration Priority Report. The collection of all Department Migration Priority Reports is referred to as the Enterprise Migration Priority Report.

As mentioned above, the migration priority levels are partly determined by three separate time estimates. These estimates are formalized in Mosca's XYZ Theorem, due to Michele Mosca, described below.

Let X denote the time, in years, for which the cryptographically protected asset requires protection or for which the cryptographic asset is required to be operational. Also referred to as the "shelf-life".

Let Y denote the time, in years, required to fully migrate the asset. Also referred to as the "migration time".

Let Z denote the time, in years, until a Cryptographically Relevant Quantum Computer becomes operational. Also referred to as the "threat timeline" or the "collapse time".

By comparing the sum of the X and Y variables to the Z variable, one obtains an estimate for a dimension of quantum risk not captured in Step 7. Concretely, Mosca's XYZ Theorem states that if $X + Y > Z$, then it is expected that a CRQC will be able to compromise the asset before the asset has been migrated to a quantum-safe state, or while the asset is still required to be protected or operational. Conversely, if $X + Y < Z$, then it is expected that the asset can be fully migrated, or will be decommissioned, before the advent of a CRQC. To err on the side of caution, it is recommended to treat the case of equality the same as the $X + Y > Z$ case.

NOTE 1: The above presentation is a slight generalization of the original theorem. In typical presentations of Mosca's XYZ Theorem-such as in [i.1] - the X variable is defined strictly in terms of how long certain data is required to be protected by the enterprise. The Y variable is typically presented as the time required to migrate the system (the cryptographic asset) that protects the data (the cryptographically protected asset).

Mosca's XYZ Theorem pertains to the protection of assets from CRQCs, it does not consider the protection of assets against threats unrelated to quantum computing. Hence, Mosca's XYZ Theorem, as stated above, cannot be used to compute migration priority levels for assets with only non-quantum computing related vulnerabilities. If an enterprise chooses to use this framework to mitigate non-quantum computing related risks, an alternative method should be used to compute the migration priority levels for those assets. For example, the department can estimate the X and Y variables following the guidance given in the present clause but can replace the Z variable with some other estimate, such as the time it is expected to take a threat actor to exploit the vulnerability.

If this framework is used in the above way, it is recommended the enterprise derive any "non-quantum Z variable" estimates in consultation with their risk management experts and program. To keep the presentation as simple as possible, the remainder of the present document assumes the Z variable is defined in terms of years until a CRQC becomes operational.

It is not always straightforward to estimate any of the three variables of Mosca's XYZ Theorem. Each estimate requires unique considerations, possibly including both technical and business considerations. Further, separate X and Y estimates should be done for each asset. A single Z variable estimate can be used for all assets in the enterprise.

NOTE 2: Different Z values can be used for different assets. If two vulnerabilities require sufficiently different quantum computing resources to exploit, it can take different lengths of time to build machines capable of exploiting the different vulnerabilities. For example, it is expected that a quantum computer capable of breaking RSA-2048 will emerge before a quantum computer capable of breaking RSA-4096. However, a risk-averse assumption is that the underlying quantum computing technology will be somewhat scalable. Hence, it is unclear if there will be much of a practical difference in the true Z values for different assets. Moreover, estimating different Z values per asset can complicate the migration analysis and planning.

The remainder of the present document assumes a single Z value is used for all assets in the Enterprise Asset Inventory. However, the enterprise can choose to use multiple Z value estimates, if desired. Regardless, as they are subject to change over time, each $X, Y,$ and Z estimate should be revised during every iteration of this framework.

The present document assumes that the $X, Y,$ and Z variables are estimated in terms of years. This assumption does not imply the estimates are necessarily whole numbers. The enterprise may estimate the variables to any level of specificity they choose. The enterprise is cautioned against using a level of specificity not commensurate with their confidence in the accuracy of the estimates, as doing so can create unacceptably small margins of error for the execution of the migration plans. For example, as 0,001 years is roughly equal to 8,76 hours, it can be inappropriate to estimate Y variables to three decimal places. However, it can be reasonable to estimate X variables to three decimal places, for example if the required protection period for an asset is known to end at a specific time and date. The enterprise should use their best judgment when deciding on the number of significant digits to use in these estimates.

For clarity, the X and Y value estimates for asset $a_{i,j}$ are denoted as $X_{i,j}$ and $Y_{i,j}$, respectively. Recommendations for calculating these estimates are given below. If the enterprise decides to use separate Z values for different assets, the notation $Z_{i,j}$ is recommended. Else, the enterprise is recommended to derive their Z-value estimate from a reputable, expert-backed source, such as the Global Risk Institute's annual Quantum Threat Timeline Report [i.1]. For example, in the 2023 version of the report, more than two-thirds of the respondents agreed that the likelihood of a CRQC emerging within the next 10 years capable of breaking RSA-2048 within 24 hours is at least 5 %. Nearly half of the respondents said the likelihood is at least 50 %. From an enterprise risk perspective, even a 5 % likelihood can be unacceptably high. Hence, taking Z to be 10 years (from the time of publication of [i.1]) can be reasonable.

X-Value Analysis

The lifespan of an asset is not always easy to determine and can be subject to change for various reasons. Hardware can be used for an indefinite time and is often not replaced (or considered for replacement) until it breaks down or some other hard-to-predict factor necessitates change. Some assets will have a pre-determined retirement date, but for the ones that do not, best guesses can be needed.

The protection requirements for cryptographically protected assets are often determined by law or legal agreements, but such laws or agreements can change (e.g. extending the required protection period or updating the technical protection requirements). Similarly, the protection requirements and lifetime of a cryptographically protected asset can be changed or extended. For example, suppose the enterprise has Personally Identifiable Information (PII) of a customer, which is required by the enterprise to offer the customer some service. Laws can require that PII be protected in a certain way and for a certain period. At the end of the pre-determined period, the enterprise may be required to securely destroy the data. However, if the customer can renew their agreement with the enterprise, then the protection period can be extended. In this way, the true shelf-life of the PII becomes difficult to determine.

Another dimension of an asset's shelf-life is the asset's value to the enterprise. The value of an asset to the enterprise is not necessarily constant over time. Many factors can influence the value of an asset over time, such as technological innovations, changes to business operations or processes, updates to standards or regulations, asset deprecation, and customer, vendor, or supply chain changes. If an asset's value is expected to lower over time, then an X-value estimate can be partly derived from the estimated time until the value of the asset reaches 0, or sufficiently close to 0.

Table 8 provides further example questions the enterprise can use to help estimate an asset's X value. However, the reader is made aware that the estimation of X values can ultimately be a business decision.

Table 8: Questions for estimating X values

Questions	Impact on X Value
If the asset is a cryptographic asset, what is the shelf-life of the data it protects? How do dependencies with other cryptographic assets impact the shelf-life?	• The length of time for which a cryptographic asset is required to be operational can depend, in part, on dependencies with other assets, or the cryptographically protected asset(s) they protect: – E.g. if the purpose of a (set of) cryptographic asset(s) is to protect specific data with known X-value estimates, those estimates can influence the X value of the cryptographic asset(s).
If the asset is a cryptographic asset, has the vendor or manufacturer set End-of-Life (EOL) or End-of-Support dates (EOS)?	• If a cryptographic asset cannot be feasibly maintained after a known date, it can be reasonable to schedule the decommissioning or replacement of that asset based on that date: – Although such assets can be excluded from the Enterprise Asset Inventory generated in Step 2, it can be helpful to check again at this point, as new EOL or EOS dates can be announced after the inventory compilation. • If a cryptographic asset is expected to still be operational even after the EOL or EOS dates, then disruptions to that asset can have greater impact after those dates than before: – E.g. the risk can increase if the asset cannot be patched or repaired as effectively without third-party support.
If the asset is a cryptographic asset, does it (or the system it resides in) have a set decommissioning date or hardware refresh cycle?	• Although such assets can be excluded from the Enterprise Asset Inventory generated in Step 2, it can be helpful to check again at this point, as new sunset or refresh dates can be determined after the inventory compilation: – Again, the decommissioning of an asset can be handled through regular change management processes (as mentioned in Step 2).

Questions	Impact on X Value
If the asset is a cryptographically protected asset, why does it require protection?	• Due to laws, policies, or regulations. • Due to legal agreements with third parties. • Because it is the enterprise's proprietary information. • To promote or ensure workflows or other business operations: – E.g. the corruption of data (in motion, in use, or at rest) can have serious impacts to business operations. • For Goodwill reasons, such as strong cybersecurity practices being a competitive advantage for the enterprise.
Besides EOL or EOS dates, are there other considerations to be made around warranties, support agreements, or supply chain requirements?	• The level or cost of support changes after a set date or after some other predefined conditions: – E.g. it can be undesirable to continue using the asset after warranties or support contracts end; resulting in a decreased shelf-life.
Does the asset require the same level and types of protections throughout its protection period?	• If the value of the asset changes over time, or the classification level of a data asset changes, it is possible that the level of cryptographic protection afforded it also can be changed over time. • If an asset's value is expected to decrease over time, there can be less urgency to migrate the asset.
Considerations for policies, procedures, contractual agreements, laws, and regulations for cryptographically protected assets	• The shelf-life of an information asset can be heavily influenced by laws, policies, regulations, or contractual agreements. • If it is expected that relevant laws, policies, regulations, or agreements will change in the future, then the estimated shelf-life of the asset, or the cyber systems protecting the asset, can be affected: – Technical changes to types of protections, such as cryptographic algorithm standards, can mean new cyber systems are needed to meet the new requirements. – Changes to the required protection period can likewise demand changes to the systems and processes currently protecting the information.
Who are the users of the asset?	• If the asset is used by specific employees, or by those in specific roles, then it is possible that the asset is no longer required if those employees leave the enterprise, or if the description of their roles change: – Again, while such assets can be excluded from the Enterprise Asset Inventory generated in Step 2, it can be helpful to check again at this point, due to changes that occurred between the inventory compilation and now.

Y-Value Analysis

Recalling from the above, the Y value of an asset is the time required to fully migrate the asset. Implicitly, this assumes that a final state for the asset is known. I.e. that a solution to migrate the asset to has been identified. Hence, the Y-value estimates can be made based on the solutions chosen in Step 7 and recorded in the Asset Migration Risk Reports.

Table 9 provides some recommendations for determining the Y values of assets. Importantly, the recommendations given in Table 9 consider assets in isolation and do not consider how the Y values are affected by asset dependencies. In practice, it is likely that some assets cannot be migrated until some other assets in their dependency digraphs have been migrated. Although the recommendations in Table 9 do include considerations for when an asset's migration can be initiated, considerations for the impacts of dependencies can be difficult to make until at least some initial, risk-based, migration order has been determined. Such dependency considerations are made in Step 9.

Table 9: Questions for estimating Y values

Questions	Impact on Y Value
How does the state of standards affect how soon the migration can be started?	• If the solution involves a cryptographic algorithm which has not yet been standardized, then the migration can possibly be begun using pre-standard draft specifications. However, this approach comes with risk. Namely, that the final standards are non-interoperable with the draft specifications and that the implementations will need to be updated once the final specifications are available. • There can be delays to migrating assets to new cryptographic standards if algorithm certification and validation are required but are not yet available or if the time required to receive certification and validation is difficult to estimate.
What resources are required, and when will the resources be available, to migrate the asset?	• Migrating systems can be technically complicated and resource intensive. The enterprise should estimate the resource costs of migrating to a given solution, determine availability, and use this information in the Y-value estimate.
What technical expertise is required to migrate the asset? Is that expertise currently available?	• If the number of people with sufficient expertise to perform the asset migration is limited, then this can increase the time taken for migration. • If additional expertise is required, then considerations for budget cycles to acquire that expertise can lengthen the Y-value estimate.
For cryptographic assets, does the supplier or manufacturer have plans to release a quantum-safe version of the asset? What are those expected timelines?	• Some migrations will not be created or implemented by the enterprise, rather by product suppliers, vendors, or systems integrators. In such cases, the enterprise should communicate with the relevant parties to learn their roadmap plans and coordinate accordingly.
Has a risk-transference solution been selected for the asset?	• If some form of insurance has been identified as the appropriate solution for the asset, then the Y value can be taken to be the estimated time until an appropriate policy and provider are identified, plus the time until the policy comes into effect.
If the asset has no vulnerabilities or if estimated risk is already acceptably low, then the asset can be considered to already be fully migrated.	• If the asset does not need to be migrated, has already been migrated, or if a risk-acceptance option has been selected for the asset, then the Y value can be taken as 0: – E.g. if this is not the enterprise's first iteration of this framework, then some assets have already been fully migrated.

Now that the department has estimates for the various $X, Y,$ and Z values, they can compute the migration priority levels of the assets. The following is one possible way to compute the migration priority levels. However, it is entirely the enterprise's decision as to the precise method they use.

For asset $a_{i,j}$, the migration priority level can be computed as $X_{i,j} + Y_{i,j} - Z$. An alternative, but related, method for calculating migration priority levels can be found in [i.5].

Mosca's XYZ Theorem states that if the sum of the X and Y values for a given asset is greater than the Z value, then the asset is at risk of quantum-aided attack. Specifically, if $X + Y > Z$, then a CRQC can exploit the quantum vulnerabilities of the asset while the asset still requires cryptographic protections, or before the quantum-safe migration of that asset is completed. Stated differently, if $X + Y - Z > 0$, then the asset is at risk of quantum-aided attack. Further, the larger the value on the left-hand side of the above inequality, the greater the risk to the corresponding asset. In other words, the greater the risk to the asset, the higher the asset's migration priority level.

EXAMPLE: If asset $a_{i,j}$ has a migration priority level of 1,2, then $X_{i,j} + Y_{i,j} - Z = 1{,}2$, and in particular, $X_{i,j} + Y_{i,j} > Z$. According to Mosca's XYZ Theorem, this asset is at risk of quantum-aided attack. In general, assets with a non-negative migration priority level are at risk of quantum-aided attack.

To provide a margin of error for the migration, the enterprise can consider adjusting their value estimates to compensate for possible delays in migration, or for CRQCs emerging sooner than estimated.

The Migration Priority Reports are not the final plans for how and when the assets will be migrated. The Migration Priority Reports are used as inputs into Step 9, wherein the final migration plans for the current iteration of the framework are developed.

Output:

- An Enterprise Migration Priority Report

6.9 Step 9: Department migration planning

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Dependency Digraph
- An Enterprise Cross-Analysis Report
- An Enterprise Migration Risk Report
- An Enterprise Migration Priority Report

Repeat for each department in the Enterprise Partition.

The migration priority levels computed in Step 8 imply an asset migration order. Namely, that the assets are migrated in descending order of migration priority level. Unfortunately, this migration order is still somewhat idealized and can have practical limitations. To facilitate a risk-based prioritization order, the analysis in Step 8 intentionally ignored how dependencies between assets can affect the migration order or when the migrations can be initiated. That is, the calculation of an asset's Y value only considered things such as the time until a solution is available and the estimated time to complete an asset's migration. The analysis did not consider situations, for example, where an asset's migration cannot be initiated until that asset's dependent assets have been fully migrated. Such constraints are considered in this Step.

Generically, the present document refers to a migration conflict as any situation wherein an asset cannot be migrated according to the order suggested by the analysis of Step 8. Just because the risk-based analysis of Step 8 suggests that assets be migrated in a certain order does not imply that the assets can practically be migrated in that order. Hence, the goal of Step 9 is to identify and address migration conflicts, and to ultimately design practical and executable migration plans for each department.

For each component of the department's Dependency Digraph, the department produces a Digraph Component Migration Plan. The collection of all a department's Digraph Component Migration Plans is called a Department Migration Plan. The collection of all Department Migration Plans is referred to as an Enterprise Migration Plan.

To detect migration conflicts, the department should examine the migration priority order suggested by Step 8 and analyse the practical consequences of following that order. This analysis should be supplemented by the outputs of previous steps of this framework, such as the dependency digraphs and Department Cross-Analysis Reports.

Migration conflicts can be caused by things such as the need to preserve functionality of an asset, preserve interoperability between assets, due to constraints on solution availability, or due to conflicts between departments. Several theoretical examples are described below. Following, recommendations are provided for resolving migration conflicts. The reader should be aware that other types of migration conflicts can exist besides those described herein, and that the recommendations provided are not necessarily applicable in all situations.

Migration Conflict Examples

A dependency cycle occurs when an asset is directly or indirectly dependent on itself. As a simple example, suppose that asset $a_{i,3}$ is dependent on asset $a_{i,2}$, asset $a_{i,2}$ is dependent on asset $a_{i,1}$, and asset $a_{i,1}$ is dependent on asset $a_{i,3}$. Such a dependency cycle is shown in Figure 5. An example of a dependency cycle is a pair of cross-signed certificates.

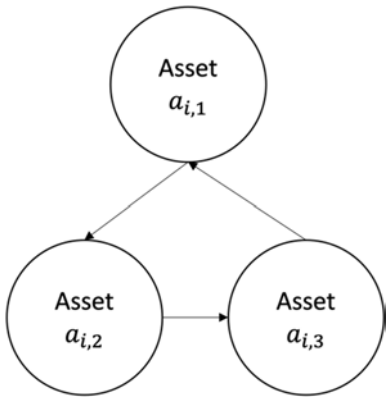

Figure 5: Dependency cycle

Dependency cycles need not be limited to two or three assets, they can possibly be much longer. Further, the longer a dependency cycle is, the less obvious it can be to detect. A computer-aided search for dependency cycles can be helpful when formulating the migration plan.

The potential issue with dependency cycles is that migrating a single asset in a cycle can cause a cascading disruption to the other assets in the cycle.

EXAMPLE 1: In Figure 5, if asset $a_{i,1}$ is migrated first, then interoperability can be broken between $a_{i,1}$ and $a_{i,2}$ and between $a_{i,1}$ and $a_{i,3}$. Because $a_{i,2}$ is dependent on $a_{i,1}$, it is possible that $a_{i,2}$ can no longer function as required after the migration of $a_{i,1}$. If $a_{i,2}$ is no longer functioning, then because $a_{i,3}$ is dependent on $a_{i,2}$ the functionality of $a_{i,3}$ can also be lost. Finally, even if interoperability is not broken between $a_{i,1}$ and $a_{i,3}$, if $a_{i,3}$ loses functionality, then $a_{i,1}$ can also lose functionality. The result is that all assets in the cycle are at risk of becoming non-functional due to the migration of a single asset.

NOTE: Migrating an asset does not necessarily break interoperability between that asset and its dependent assets. The impacts to interoperability and functionality depend on the nature of the assets and the migrated-to solution.

Even without the presence of dependency cycles, the order in which assets in a dependency digraph component are migrated can have significant impacts on interoperability and functionality of other assets.

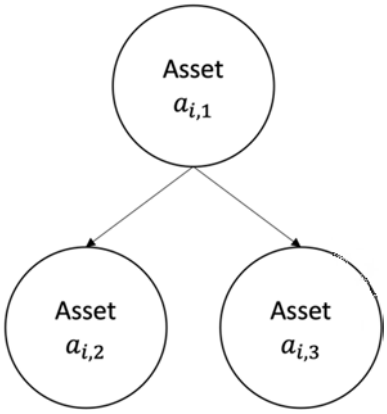

Figure 6: Multiple dependencies

EXAMPLE 2: In Figure 6, assets $a_{i,2}$ and $a_{i,3}$ are independent of each other. Consequently, the migration of $a_{i,2}$ has no impact on the functionality of $a_{i,3}$, and vice versa. However, as both are dependent on $a_{i,1}$, if $a_{i,1}$ is migrated first then the functionality of both other assets can be affected.

One take-away from the above example is that the more dependencies (direct or indirect) an asset has, the greater the impact can be to the rest of the assets in the dependency digraph when that asset is migrated. This is the second type of migration conflict: when an asset has a higher migration priority level than some of its dependent assets, but where migrating that asset before its dependent assets causes unacceptable disruption or is otherwise impractical.

A third type of migration conflict can occur due to certain solutions not yet being available. The Y-value analysis of Step 8 considered situations where a desired solution is not immediately available, where there are factors outside the enterprise's control which delay when the asset's migration can be initiated. Similar scenarios are considered in Section 3.1.1 of [i.5]. In these situations, it can make sense to begin the migrations of lower priority assets instead of taking no action while waiting for higher-priority solutions to become available.

A fourth type of migration conflict is one between departments or third parties. The migration priority level calculated for an asset in Step 8 is ideally equal for any department which shares the asset. That is, if an asset is shared by multiple departments, say D_i and D_j, then the migration priority level D_i calculated for the asset should equal that calculated by D_j. Hence, the asset can ideally be migrated at the same time for both departments. However, this is not necessarily the case in practice. For example, it is possible that when D_i is ready to migrate the asset, D_j is not ready. Perhaps D_j still has other assets to migrate before it can accept a migration of the shared asset, there are budget or resource limitations on the part of D_j, or other reasons causing an asynchronicity between the two departments. The result is a situation where an asset is desired to be migrated simultaneously for multiple departments, but where a simultaneous migration is not feasible. Similar situations can occur with external dependencies with third parties.

To summarize, potential causes of migration conflicts include:

- the need to avoid disruptions to interoperability or asset functionality;
- waiting periods until certain solutions are available, or until the migration of certain assets can be initiated;
- differing constraints and priorities between departments which share assets; and
- differing constraints and priorities between the department and third parties.

Approaches to Resolving Migration Conflicts

Step 4 described three distinct types of approaches for migrating an asset (backwards compatible migrations, parallel migrations, and pure migrations) as well as how the three approaches can impact dependencies. Table 10 summarizes some recommendations for using the three migration approaches to resolve migration conflicts. Table 10 also provides examples not using the three migration approaches, such as by altering the migration priority levels, or coordination and synchronization between departments. In all cases, the enterprise has the option to select an entirely different solution for an asset than was identified in Step 7.

The Y values calculated in Step 8 for each asset were partly based on the specific end-state solutions selected for the assets in Step 7. If a different solution is selected for an asset, then the migration priority level for that asset will likely be changed as well. Moreover, the new solution can mitigate the vulnerabilities identified in Step 4 to a different extent than the solution selected in Step 7. For example, if a parallel migration approach is selected to resolve a migration conflict (where a fully quantum-safe solution was selected in Step 7), then the quantum-vulnerabilities of the legacy system will not be mitigated at all. The vulnerabilities will persist until the parallel system is fully implemented, and the legacy system is decommissioned. The enterprise should use their own best judgement when deciding how to resolve migration conflicts.

Even without the presence of a migration conflict, it is possible for an alternative solution to be selected for an asset during this Step, different from the solution selected in Step 7. Similarly, an alternative migration order for a digraph component than the one identified in Step 8 can be selected in this Step, even without the presence of a migration conflict. This framework has attempted to provide robust technical methods for assessing when and how assets should be migrated. However, due to the natural complexities of enterprises and of cryptographic migrations, the final migration plans will likely be the results of business decisions rather than purely technical, risk-based, decisions.

Table 10: Recommendations for resolving migration conflicts

Migration Conflict	Possible Remediations
Disruption of interoperability	• Parallel migrations. • Backwards compatible solutions. • Pure asset migration. • Raise or lower migration priority level(s). See note.
Waiting periods until an asset's migration can be initiated	• Begin migrating lower-priority assets until the migration of higher-priority assets can be initiated.
Cross-departmental or external conflicts	• Parallel migrations. • Backwards compatible solutions. • Merging the Migrations Plans of affected departments and rerunning the framework analysis. • Raise or lower migration priority level(s) to synchronize priority orders. See note.
NOTE: Lowering the migration priority level of an asset can incur risk. The lower the migration priority level, the longer it can be before that asset is fully migrated. If the enterprise decides to lower the migration priority level of an asset to resolve migration conflicts, then care should be taken to ensure the new level of risk is within the risk appetite of the enterprise.	

Once all migration conflicts have been identified and resolutions to them selected, the formal migration plans are constructed. If no migration conflicts are identified, then the assets can be migrated in the order suggested by the migration priority vectors and to the solutions selected in Step 7.

A migration plan (of any type) should include all available information required to execute the plan. Additional supporting or clarifying information can also be included. Example suggestions for items to include in the three types of migration plans are given in Table 11 below.

Table 11: Recommendations for constructing migration plans

Type of Migration Plan	Example Content
Digraph Component Migration Plan	• Information related to solution acquisition or development. • Information on the planned mechanisms or processes for asset migrations. • Estimated X and Y values. • Migration priority vectors. • For each asset, if the identified solution is an intermediary solution. • The migration interval in which each asset is planned to be migrated. • Associated roles and responsibilities. • Any other supporting information helpful for the successful migration of each asset.

Type of Migration Plan	Example Content
Department Migration Plan	• Summary information of each Digraph Component Migration Plan. • The migration intervals in which each digraph component is expected to be (fully) migrated. • Associated roles and responsibilities. • Information on commonalities between different Digraph Component Migration Plans: – E.g. noting commonalities between different Digraph Component Migration Plans can be helpful in avoiding duplication of efforts, can streamline solution acquisition cycles, improve resource allocation planning, etc. • Any other supporting information helpful for the successful migration of the department.
Enterprise Migration Plan	• Summary information of each Department Migration Plan. • The migration intervals in which each department. is expected to be (fully) migrated. • Associated roles and responsibilities. • Information on commonalities between different Department Migration Plans: – E.g. noting commonalities between different Department Migration Plans can be helpful in avoiding duplication of efforts, can streamline solution acquisition cycles, improve resource allocation planning, etc. • Any other supporting information helpful for the successful migration of the enterprise.

Although every asset in the Enterprise Asset Inventory will be assigned a migration priority level, it is expected that some (perhaps many) assets will not be migrated, or planned to be migrated, during the current migration interval. It is possible that due to resource constraints and other priorities, even assets with high migration priority levels will not be migrated during the current migration interval. Moreover, it is possible that assets which are planned to be migrated during the current migration interval fail to be migrated due to unforeseen circumstances or other issues. Finally, the *Y*-value analysis of Step 8 asked how long it will take to fully migrate the asset to the desired end state. Some of the approaches to resolving migration conflicts described in this Step can result in the asset reaching a non-fully migrated state. These are some reasons for iterating the framework; further reasons are discussed in Step 11. Regardless, it should be noted in the relevant Department Migration Plan if an asset is not planned to be, or expected to be, migrated during the current migration interval.

Output:

- An Enterprise Migration Plan

6.10 Step 10: Execute migration plans

Input:

- An Enterprise Partition
- An Enterprise Asset Inventory
- An Enterprise Migration Plan

Repeat for each department in the Enterprise Partition.

The goal of Step 10 is straightforward: execute each Department Migration Plan over the current migration interval.

At the end of the migration interval the department produces an Asset Migration Status Report for each of its assets. The collection of all a department's Asset Migration Status Reports is called a Department Migration Status Report. The collection of all Department Migration Status Reports is referred to as an Enterprise Migration Status Report.

The purpose of the migration status reports is to record and track important occurrences during the migration interval and to provide information for updating the migration plans in the following iteration of the framework. For example, in a real-world scenario, it is likely that the migrations will not all proceed precisely as planned. Unforeseen circumstances, accidents, disruptions, errors, and miscalculations can all contribute to flaws in the execution of migration plans. Moreover, in the face of such issues, the department can alter their migration plans ad hoc. Such events should be recorded in the relevant Status Report.

Table 12 gives examples of items which can be recorded in an Asset Migration Status Report.

Table 12: Asset Migration Status Report considerations

Asset Migration Status Report	Considerations
Was the asset fully migrated to the solution identified in Step 7?	• Yes - the asset was fully migrated to the solution identified in Step 7. • No - the migration is in-progress and was not completed during the intended migration interval. • No - the asset was fully migrated to an alternative end-state solution than was identified in Step 7. • No - the asset was migrated to an alternative intermediary solution than was identified in Step 7. • No - the asset is being migrated to an alternative intermediary solution than was identified in Step 7, but the migration was not completed during the intended migration interval. • No - the migration was halted because the identified solution was discovered to be insufficient, a superior solution has not yet been identified. • No - the migration was halted because the identified solution was discovered to be insufficient, a superior solution was identified, but migration has not yet begun. • No - the asset was not planned to be migrated during this migration interval. • No - other
Were there any important changes to the asset during the migration interval, aside from migration considerations?	• Were vendor-supplied updates made to the asset which altered the viability of the Asset Migration Plan? • Did the asset experience any errors, failures, or other factors causing it to be decommissioned or replaced? • Did the asset's dependencies change during the migration interval?
Were there any significant changes to the vulnerabilities, risks, or solutions for the asset?	• For any asset, new vulnerabilities can emerge, risk components can change: – E.g. zero-days or other exploits, and improved cryptanalysis making certain attacks easier to execute. • The identified solutions can also experience change: – E.g. changes to algorithm (draft) specifications, improved cryptanalysis, changes to parameters or cryptographic primitives. • Were there any unexpected delays in obtaining solutions? • Did new solutions emerge, which were not included in the Step 4 analysis, and which are potentially more desirable than currently planned-for solutions? • Did any other events occur which significantly changed the estimated X, Y or Z values for the asset?
Other important notes or observations?	• Comments about the experience of migrating the asset can be helpful when formulating or updating future migration plans: – E.g. lessons learned.

Table 13 gives examples of additional items which can be recorded in a Department Migration Status Report.

Table 13: Department Migration Status Report considerations

Department Migration Status Report	Considerations
Did the department achieve its migration goals for the migration interval?	• Which assets, or digraph components, have been migrated as planned? Which have not been migrated as planned?
Were there any important, unexpected, changes to the department during the migration interval?	• The migration plans can be affected by changes to personnel, the department structure, the department's mission or mandate, and so on.
Did the department's asset inventory change unexpectedly during the migration interval?	• The department's inventory of assets is not necessarily fixed, assets can be added or removed. Making note of such changes here can simplify parts of the next iteration of this framework.
Other important notes or observations?	• Comments about the experience of migrating the department can be helpful when formulating or updating future migration plans: – E.g. lessons learned.

Finally, Table 14 gives examples of additional items which can be recorded in the Enterprise Migration Status Report.

Table 14: Enterprise Migration Status Report considerations

Asset Migration Status Report	Considerations
Were there any important, unexpected, changes to the Enterprise during the migration interval?	• The structure of the enterprise can be changed due to things such as mergers and acquisitions, the combining of departments, the creation of new departments, or the dissolution or closure of departments. • The migration plans can be affected by changes to key personnel, including whoever is responsible for overseeing the various migrations.
Other important notes or observations?	• Comments about the experience of migrating the enterprise can be helpful when formulating or updating future migration plans. – E.g. lessons learned.

Output:

- An Enterprise Migration Status Report

6.11 Step 11: Prepare for next iteration

Input:

- An Enterprise Migration Status Report

The goal of Step 11 is to decide whether the framework will be iterated, and if so, to make appropriate preparations for the coming iteration.

The output of Step 11 is different from the outputs of the other framework Steps. Step 11 produces no original reports. However, action is still taken by the enterprise within this Step, as described below. The Migration Status Reports generated in Step 10 are repeated as the output of Step 11 for convenience, as those reports can be used as input to Step 1, in the case of iteration.

There are many reasons for why the enterprise can decide to iterate this framework. Examples include, incomplete assets migrations, changes to department or enterprise structures, changes to department asset inventories, or due to other information included in the Migration Status Reports. Moreover, it is possible that the enterprise will iterate the framework indefinitely, and eventually incorporate it as a normal component of their risk management, change management, or other enterprise programs.

It is often beneficial to first review, update, and optimize a framework or process before it is repeated. Hence, before iterating the framework, the enterprise should spend some time examining their processes, incorporate lessons-learned from the previous iterations, revisit their underlying assumptions, and make any appropriate adjustments before the next iteration.

Output:

- An Enterprise Migration Status Report.

History

Document history		
V1.1.1	October 2024	Publication