

Leiden University

ICT in Business and the Public Sector

Auditing Digital Sovereignty: Design and Evaluation of an Operational Reversibility Framework for IT Auditors

Name:	Zaynab Musawi
Student ID:	s4566564
Date:	03/07/2026
1st supervisor:	Prof. dr. ir. J.M.W. Visser
2nd supervisor:	Drs. J.B. Kruiswijk
Company supervisors:	N. Ali & J. Rijksen-Hartman

MASTER'S THESIS

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Einsteinweg 55
2333 CC Leiden
The Netherlands

Abstract

Background: European organizations are increasingly dependent on a small number of non-EU cloud providers, creating structural risks to operational continuity, data control, and institutional knowledge. Regulatory frameworks, including NIS2, DORA, and the GDPR, impose resilience and continuity obligations that require organizations to maintain the capability to exit a cloud provider without catastrophic loss, a concept referred to as *operational reversibility* which is a core dimension of digital sovereignty. While European frameworks such as the DICTU Toetsingsinstrument, the EC Cloud Sovereignty Framework, and EuroStack define what sovereign cloud capability looks like at a policy level, none provides IT auditors with a structured, independently verifiable instrument for assessing whether that capability actually exists within a client organization. The gap between a normative definition of digital sovereignty and an auditable criterion for it remains unaddressed.

Aim: This study designs and evaluates the Digital Sovereignty Audit Framework (DSAF), an audit instrument that operationalizes the operational reversibility dimension of digital sovereignty into five control areas, each comprising structured audit questions, cumulative design maturity criteria (M1 to M5), and independent operating effectiveness criteria (E1 to E4), for use by IT auditors in cloud-dependent organizations.

Method: The DSAF was developed and validated using Design Science Research (DSR). Design requirements were derived from a comparative gap analysis of five existing European sovereignty frameworks. The artifact was evaluated across two expert interview rounds with IT audit and Technology Risk professionals from a single Big Four firm (Round 1: n=7; Round 2: n=2), using a TAM-based twelve-item Likert questionnaire across four constructs (Perceived Usefulness, Perceived Ease of Use, Social Norm, and Behavioral Intention), combined with reflexive thematic analysis of the qualitative interview data.

Results: Round 1 produced a differentiated TAM profile across the four constructs. Perceived Usefulness was strong, with all seven respondents unanimously agreeing that the framework surfaces cloud exit risks a standard audit would miss and that it improves the completeness of audit conclusions, though the item assessing whether the framework provides sufficient structure for a defensible audit judgment was the most divided in the data set. Perceived Ease of Use was the lowest-scoring construct overall, with the distinction between design maturity and operating effectiveness producing the most interpretive ambiguity. Social norms showed a striking divergence: respondents near unanimously agreed that formal adoption by a leading professional body would increase their own willingness to use the framework. Behavioral intention remained high despite the moderate ease of use scores, indicating that usability concerns did not translate into reduced adoption intent. Five qualitative themes emerged from the thematic analysis, establishing that practical applicability determines adoption, that professional judgement must be preserved rather than replaced by rigid compliance criteria, and that institutional endorsement is the primary driver for adoption in the IT audit profession. Four concrete problems were identified and addressed in Version 2; the most significant was the translation to spreadsheet-based workbook format, identified by six of seven respondents as the format most compatible with existing audit practice. Round 2 confirmed that all four Version 2 adjustments were perceived as resolving the problems they targeted, while surfacing two new implementation-specific findings concerning the technical enforcement of cumulative scoring logic and the design and operating effectiveness hierarchy.

Conclusion: The DSAF provides IT auditors with a structured, evidence-grounded instrument for assessing operational reversibility, filling a gap that no existing European sovereignty framework addresses at the level of practical audit application. The findings indicate that formal endorsement by professional bodies represents the most significant driver for adoption, beyond further refinement of the framework itself. Future work should address the technical enforcement of the scoring hierarchy within the workbook, develop realistic testing proxies for operating effectiveness compatible with standard audit engagement constraints, and evaluate the framework's applicability across sectors and organizational sizes beyond the Big Four context in which it was validated.

Acknowledgements

This thesis would not have been possible without the support of several people to whom I owe sincere gratitude.

First, I thank my two academic supervisors, Prof. dr. ir. J.M.W. Visser and Drs. B. Kruiswijk, for their guidance, feedback, and encouragement at every stage of the process. Their insights helped me sharpen both the conceptual depth and the scientific precision of this work.

I would also like to thank my two company supervisors, N. Ali and J. Rijksen-Hartman, for their practical support and their confidence in this project. Their involvement made the empirical part of this research possible.

I am especially grateful to all nine respondents who generously gave their time and expertise to evaluate the framework. Their honest and constructive feedback was invaluable and directly shaped the final artifact. I would also like to thank the many professionals and peers I encountered throughout this period who were willing to discuss, question, and think along with me about digital sovereignty and cloud governance. Those conversations, however informal, contributed more to this work than they may realize.

Finally, and most importantly, I want to thank my parents and my little sister. This thesis was written during a period that was not easy for our family, and there were moments when I could not be there as much as I wanted to be. The fact that my parents and sister never once made me feel guilty for that and instead continued to encourage me every step of the way means more to me than I can put into words. This thesis is as much theirs as it is mine.

Contents

Abstract . . . 1
Acknowledgements . . . 1

1 Introduction **4**
1.1 Problem statement . . . 4
1.2 Research questions . . . 5
1.3 Research scope . . . 6
1.4 Overview of the thesis . . . 6

2 Background and related work **7**
2.1 Introduction to IT auditing . . . 7
2.1.1 Financial audit . . . 7
2.1.2 IT audit . . . 7
2.1.3 Internal versus external IT auditing . . . 8
2.1.4 What gets audited? . . . 9
2.1.5 Why IT audits are performed . . . 9
2.1.6 The IT audit process . . . 10
2.2 Cloud Computing fundamentals . . . 11
2.2.1 Definition . . . 11
2.2.2 Cloud service models . . . 11
2.2.3 Cloud deployment models . . . 12
2.2.4 Shared responsibility models . . . 12
2.2.5 Cloud dependencies and risk factors . . . 13
2.2.6 Cloud governance and security standards . . . 13
2.3 Digital sovereignty . . . 14
2.3.1 Digital sovereignty . . . 14
2.3.2 Digital autonomy . . . 15
2.4 Existing cloud sovereignty frameworks . . . 15
2.4.1 European Commission Cloud Sovereignty Framework . . . 16
2.4.2 EuroStack . . . 18
2.4.3 The EuroStack Buy European Procurement Framework . . . 19
2.4.4 The DICTU Toetsingsinstrument Soevereiniteit Clouddiensten . . . 21
2.4.5 The SUSE Cloud Sovereignty Framework Self-Assessment . . . 25
2.4.6 Comparative evaluation . . . 26
2.5 IT Audit and assurance frameworks . . . 31
2.5.1 COBIT . . . 31
2.5.2 ISO/IEC 27001 . . . 31
2.5.3 Service Organizational Control (SOC) . . . 32
2.5.4 Cloud Control Matrix (CCM) . . . 32

3 Method **34**
3.1 Research paradigm and overall strategy . . . 34
3.2 The Technology Acceptance Model (TAM) . . . 34
3.3 Research design . . . 35
3.4 Sampling strategy . . . 35
3.5 Data collection instrument . . . 36
3.5.1 Interview protocol . . . 36

3.5.2 Transcription procedure . . . 37
3.5.3 Likert instrument . . . 37
3.6 Data analysis . . . 37
3.6.1 Quantitative analysis of Likert data . . . 37
3.6.2 Qualitative analysis: Thematic analysis . . . 37
3.7 Ethical consideration . . . 38
3.8 Second evaluation round . . . 38
3.9 Research quality criteria . . . 39
3.10 Synthesis . . . 39

4 Design **40**
4.1 Design requirements and artifact architecture . . . 41
4.1.1 Design requirements derived from the comparative gap analysis . . . 41
4.1.2 Artifact architecture . . . 43
4.2 Generic structure of a control area . . . 44
4.2.1 Sovereignty Objective, Risk, and Sovereignty Link . . . 44
4.2.2 Audit questions . . . 46
4.2.3 Part A: Design Maturity (M1-M5) . . . 46
4.2.4 Part B: Operating Effectiveness (E1-E4) . . . 47
4.2.5 Aggregation and decision logic . . . 47
4.3 The five control areas . . . 48
4.3.1 Control area 1: Exit plan and contractual framework . . . 49
4.3.2 Control area 2: Data portability and technical exportability . . . 50
4.3.3 Control area 3: Knowledge transfer and operational documentation . . . 52
4.3.4 Control area 4: Continuity of operations during migration . . . 54
4.3.5 Control area 5: Testing programme and periodic evaluation . . . 55
4.4 Design rationale . . . 57
4.4.1 Scoring logic within a control area: cumulative maturity ladder . . . 57
4.4.2 Aggregation across control areas: average with weakest-link signaling . . . 58
4.4.3 Decomposition into control areas: five distinct controls . . . 59
4.4.4 Quantitative indicators: IaC coverage as a verified, principle-based criterion . . . 60

5 Evaluation **61**
5.1 Results: TAM (Round 1) . . . 61
5.1.1 Results by TAM construct . . . 61
5.1.2 Perceived Usefulness (PU) . . . 64
5.1.3 Perceived Ease of Use (PEOU) . . . 64
5.1.4 Social Norm (SN) . . . 65
5.1.5 Behavioral Intention (BI) . . . 65
5.2 Themes from the interviews . . . 66
5.2.1 Theme 1: The framework provides structure and improves audit quality . . . 66
5.2.2 Theme 2: Practical applicability determines adoption . . . 66
5.2.3 Theme 3: The framework is valuable but requires methodological maturation . . . 66
5.2.4 Theme 4: Digital sovereignty is gaining relevance, but adoption is shaped by external drivers . . . 67
5.2.5 Theme 5: Auditing digital sovereignty requires professional judgment, not rigid compliance criteria . . . 67
5.3 Concrete problems identified (Round 1) . . . 67
5.4 Concrete improvements (Version 2) . . . 68
5.5 Results: Round 2 (confirmatory evaluation) . . . 69
5.5.1 Qualitative findings . . . 70
5.6 Comparison between rounds . . . 72
5.7 Conclusion . . . 74
5.8 Limitations . . . 74

6 Discussion **75**
6.1 Findings in relation to the literature . . . 75
6.1.1 The DSAF fills a real and specific gap, not a general one . . . 75
6.1.2 The two-dimensional scoring model . . . 76

6.1.3 Adoption is institutionally driven, not individually driven . . . 76
6.1.4 Professional judgement cannot be replaced, but it can be anchored . . . 76
6.2 Position within the broader digital sovereignty landscape . . . 77
6.2.1 Operational reversibility is one sub-dimension, not a complete treatment . . . 77
6.2.2 Cryptographic access control is the most urgent next subdimension . . . 77
6.3 Reflection on the absence of a real-world case study . . . 78
6.3.1 Why expert evaluation is sufficient at this stage . . . 78
6.3.2 What a real-world application would look like . . . 79
6.4 Limitations . . . 79

7 Conclusions **80**
7.1 Answers to the research questions . . . 80
7.1.1 SQ1: What is digital sovereignty, and which key elements characterize it in organizational contexts? . . . 80
7.1.2 SQ2: What aspects of digital sovereignty are not adequately addressed by existing IT audit and assurance frameworks? . . . 81
7.1.3 SQ3: Which audit criteria can be used to translate digital sovereignty into an auditable concept for IT audit practice? . . . 81
7.1.4 SQ4: How can these audit criteria be operationalized into audit procedures, required evidence, and sovereignty assurance levels? . . . 81
7.1.5 SQ5: How do IT audit and Technology Risk professionals evaluate the usability and practical relevance of the proposed framework? . . . 81
7.1.6 Main research question . . . 82
7.2 Contributions . . . 82
7.2.1 Scientific contribution . . . 82
7.2.2 Practical contribution . . . 82
7.3 Future work . . . 82

A Interview guide **88**

B DSAF Audit Control Form **102**

C The full Atlas.ti code list **116**

Chapter 1

Introduction

This chapter introduces the research by outlining the problem that motivated it, the objective it pursues, and the questions it seeks to answer. It concludes with a delineation of the research scope and an overview of the thesis structure.

1.1 Problem statement

Organizations are increasingly reliant on cloud service providers for mission-critical operations and data-intensive workloads. Worldwide end-user spending on public cloud services reached $595.7 billion in 2024 and is forecast to grow to $723.4 billion in 2025, with 90% of organizations projected to adopt hybrid cloud models by 2027 [28]. While compliance with existing security and regulatory controls can often be demonstrated, such reliance may introduce long-term structural dependencies that constrain organizational autonomy and strategic decision-making power while reducing transparency over how data is processed, where it is stored, and under whose legal jurisdiction it falls [43]. These dependencies are not always visible through traditional IT control assessments [32], yet they can significantly affect an organization's ability to govern, adapt, or exit cloud-based arrangements. Moreover, these dependencies are often obscured by selective risk representations [21] and narrowly scoped assurances provided by cloud service providers and policy documents. While formal compliance statements may emphasize specific safeguards or legal limitations, they frequently fail to address broader systemic exposures such as residual jurisdictional reach, provider discretion, and asymmetric power relations [25]. As a result, organizations and auditors may overestimate actual control and autonomy based on incomplete or misleading information.

Digital sovereignty has therefore emerged as a prominent concept in European policy and academic debates, referring broadly to the ability of organizations and states to retain control over their data, digital infrastructure, and decision-making in a globalized digital ecosystem [25]. However, despite its growing relevance, there is limited consensus on how digital sovereignty can be operationally defined and measured within organizational assurance contexts [46]. As a result, sovereignty remains largely implicit in IT audit practice: audits typically assess whether controls exist and operate effectively, but do not explicitly evaluate whether organizations retain sufficient autonomy, legal control, transparency, and exit feasibility in cloud-dependent environments [32].

This gap is becoming increasingly urgent due to the rising geopolitical volatility and jurisdictional risks that affect cross-border data processing and cloud services [21]. A frequently cited concern is extraterritorial access to data under foreign legislation such as the United States CLOUD Act, which allows U.S. authorities, under specific legal conditions, to compel certain providers to disclose data within their "possession, custody, or control," even when this data is stored outside the United States [55]. These legal mechanisms create uncertainty for organizations operating under European data protection and public accountability regimes, particularly in sectors handling sensitive or public-interest data.

A second risk is that organizations can lose access to cloud services without warning and without having made that decision. This can happen because a provider changes its policies, because governments impose trade restrictions, or because of geopolitical developments entirely outside the organization's control. In 2025, for example, Anthropic was forced to disable access to two of its AI model series for users in

certain countries following U.S. export control measures [4], leaving dependent organizations without access overnight. This shows that cloud dependency creates two distinct risks: a foreign government can access your data, or you can simply lose access to the service itself. Both risks are expressions of the same underlying problem, and this study addresses both.

At the European level, digital sovereignty has recently shifted from a high-level strategic notion to a more concrete assessment concept. This is exemplified by the European Commission's Cloud Sovereignty Framework, which defines sovereignty objectives and introduces assurance levels and scoring mechanisms in procurement contexts [16]. However, such policy-level frameworks are not designed as IT audit methods and do not directly translate into audit objectives, assessment criteria, or maturity models used in practice.

From an academic perspective, this research addresses a clear gap between high-level digital sovereignty concepts and their application in operational assurance and IT audit practice. Existing research on digital sovereignty and strategic digital dependencies focuses primarily on policy, governance, and legal dimensions [25, 45, 11], offering limited guidance on how these concepts can be transformed into auditable constructs at organizational levels. By proposing an auditable conceptualization of digital sovereignty and a structured assessment framework tailored to IT audit contexts, this study contributes both theoretically and practically. The expected outcome is a validated, audit-oriented framework that enables IT auditors to systematically assess the operational reversibility dimension of digital sovereignty in cloud-dependent organizations, providing both a proof of concept and a reusable template for extending the assessment to additional sovereignty dimensions in future work.

1.2 Research questions

Digital sovereignty has emerged as a central concern in European digital policy and cloud governance, particularly in relation to organizational dependence on external cloud service providers. While the concept is increasingly translated into policy objectives and procurement requirements, there is currently no established method within IT audit practice to assess digital sovereignty in a structured and auditable manner. As a result, sovereignty-related risks remain largely implicit in audits, despite their strategic and operational significance.

This goal is addressed through the following main research question:

RQ: *How can IT auditors assess digital sovereignty in organizations using cloud services in a structured and auditable manner?*

To provide a comprehensive and well-founded answer to the main research question, it is decomposed into a set of sub-questions. Each sub-question addresses a specific aspect of the research problem and contributes to the development and evaluation of the proposed audit framework. Together, the sub-questions form a coherent research path from conceptual foundations to practical application.

The identified sub-questions are as follows:

SQ1: What is digital sovereignty and which key elements characterize digital sovereignty in organizational contexts?

SQ2: What aspects of digital sovereignty are not adequately addressed by existing IT audit and assurance frameworks?

SQ3: Which audit criteria can be used to translate digital sovereignty into an auditable concept for IT audit practice?

SQ4: How can these audit criteria be operationalized into audit procedures, required evidence, and sovereignty assurance levels?

SQ5: How do IT audit and Technology Risk professionals evaluate the usability and practical relevance of the proposed digital sovereignty assessment framework?

To answer these research questions, we will design audit criteria not for all aspects of digital sovereignty, but as a first step for one of its prime aspects: operational reversibility, defined as the ability of an organization to exit its current cloud provider without catastrophic loss of service continuity, data, or institutional knowledge. This focus was chosen because operational reversibility is the dimension most

directly accessible through existing IT audit methodology and because it represents the most significant unaddressed gap in current European sovereignty frameworks, as the comparative analysis in chapter 2 demonstrates. This scoping decision allows the study to produce a complete, validated instrument rather than a superficial treatment of a broader topic.

The research approach follows a logical progression. SQ1 establishes the conceptual and policy foundations of digital sovereignty. SQ2 identifies the gap between these high-level concepts and current IT audit practice. Based on these insights, SQ3 and SQ4 focus on designing an auditable framework that translates sovereignty into concrete assessment criteria. Finally, SQ5 evaluates the proposed framework through expert validation to assess its practical applicability and relevance in real-world IT audit contexts.

1.3 Research scope

This research focuses on digital sovereignty within the context of IT audits and cloud-dependent organizations. Several deliberate scoping decisions shape this focus, each motivated by the need to produce an instrument that is both practically and methodologically tractable within the bounds of a single study.

The scope is restricted to cloud-dependent organizations because digital sovereignty risks, while relevant at the macro level of European industrial policy, only become auditable at the organizational level where concrete evidence exists: exit plans, contracts, data export configurations, and continuity procedures. Organizations without meaningful cloud dependency do not face the same structural lock-in risks and therefore fall outside the relevant problem space.

The scope is further restricted to IT audit and assurance practice, rather than legal or financial advisory, because operational reversibility is the dimension of digital sovereignty most directly assessable through existing IT audit methodology. Assessing data sovereignty in a legal sense requires jurisdictional expertise; assessing economic sovereignty requires financial modeling. Assessing whether an organization can exit its cloud provider without catastrophic loss of continuity, data, or institutional knowledge requires evidence-based evaluation of controls, processes, and documentation, which is what IT auditors are trained to do.

Finally, this study does not seek to assess digital sovereignty at the macro level of national or European infrastructure policy, nor does it compare specific cloud vendors or evaluate geopolitical risk at the state level. These are important questions, but they fall outside the competence of an IT audit engagement and beyond the scope of a single master's thesis.

1.4 Overview of the thesis

Chapter 2 - Related work: This chapter reviews the existing landscape of European digital sovereignty frameworks and IT audit standards. It conducts a comparative gap analysis of five frameworks, identifying the structural limitations that motivated the design of the Digital Sovereignty Assessment Framework (DSAF).

Chapter 3 - Method: This chapter described the Design Science Research (DSR) methodology underpinning the study, including the TAM-based evaluation instrument, the sampling strategy, the data collection procedure, and the analytical methods used for both quantitative and qualitative data.

Chapter 4 - Design: This chapter documents the DSAF in its second, revised iteration. It presents the five control areas, the generic structure of each control area, and the design rationale for the four principal design decisions.

Chapter 5 - Evaluation: This chapter reports the results of two rounds of expert evaluation. Round 1 involved seven IT audit and Technology Risk professionals and produced five qualitative themes and five concrete design improvements. Round 2 confirmed these improvements and surfaced two new implementation-specific findings.

Chapter 6 - Discussion: This chapter reflects on the findings in relation to the broader literature on digital sovereignty, IT audit methodology, and Design DSR.

Chapter 7 - Conclusion: This chapter answers the research questions, summarizes the contributions of the study, and presents recommendations for future research.

Chapter 2

Background and related work

This chapter provides the necessary background and an overview of existing academic and practical literature relevant to this research. It sets the foundation for understanding the context and significance of applying a digital sovereignty assessment method.

2.1 Introduction to IT auditing

The increased reliance of organizations on IT has made IT audits important prerequisites for reliable financial audits. Understanding what an IT audit is, what it covers, and how it is performed in practice is essential for this thesis, because the framework developed in chapter 4 is designed to be used by IT auditors in their existing engagements, without requiring a new professional paradigm.

2.1.1 Financial audit

A financial audit is an independent, systematic examination of an organization's financial statements. The purpose is to determine whether these statements are free from material misstatement and prepared in accordance with the relevant reporting framework (such as IFRS or GAAP [33]). Financial auditors evaluate accounting records, internal controls, and financial processes to form an audit opinion, which enhances the confidence of stakeholders such as shareholders, regulators, investors, and the public.

Financial audits follow structured standards issued by recognized bodies such as the International Auditing and Assurance Standards Board (IAASB). One of the most relevant standards is ISA 315, which requires auditors to understand the entity and its internal control environment in order to identify and assess risks of material misstatement [33]. Traditionally, this process focuses on evaluating financial transactions, journal entries, reconciliations, and other accounting-related evidence.

Over the past decades, however, the financial audit has become inseparable from IT. Modern financial reporting relies heavily on digital systems such as Enterprise Research Platforms (ERP), automated workflows, and electronic data processing [32]. These systems directly influence the completeness, accuracy, and integrity of financial information. For this reason, financial auditors increasingly depend on specialized IT auditors to assess the reliability of the digital systems that support financial reporting.

2.1.2 IT audit

An Information Technology (IT) audit focuses on evaluating the governance, risk management, and control structures embedded within an organization's IT environment [27]. Its purpose is to determine whether IT systems safeguard assets, maintain data integrity, support organizational objectives, and operate effectively and securely [27].

Within the context of a financial statement audit, IT audit procedures aim to assess whether IT-dependent controls and automated processes can be relied upon. Since financial transactions are often initiated, processed, and recorded within IT systems, weaknesses in IT control may lead directly to material misstatements [33].

Frameworks such as COBIT 2019 [32], and the COSO Internal Control Framework [10] provide structured approaches for evaluating IT control environments. Additionally, ISA 315 explicitly requires auditors to understand the entity's IT environment, including general IT controls, automated application controls, and IT-dependent manual controls [33].

IT controls are generally categorized into General IT Controls (GITCs), IT Application Controls (ITACs), and IT Dependent Manual Controls (ITDMs) [33].

General IT Controls (GITCs) GITCs provide the foundational control environment that all other controls depend on. If GITCs fail, the reliability of both automated and manual controls built on top of them is undermined [48]. GITCs typically cover four areas. **Logical access control** ensures that only authorized users can access systems and data, consistent with the principle of least privilege [32]. **Change management** ensures that system changes are authorized, tested, approved, and properly implemented before going live. **IT operations** covers job scheduling, incident management, backup procedures, and system monitoring. **Program development control** governs the design and implementation of new systems or major modifications to existing ones.

IT Application Controls (ITACs) ITACs are fully automated controls embedded within specific applications that execute at the transaction level without human intervention [33]. A concrete example is an automated three-way match in a procurement system that blocks payment if the purchase order, delivery receipt, and invoice do not align. Because ITACs are automated and consistent, auditors can place greater reliance on them, reducing the need for manual transaction testing [33]. However, the reliability of ITACs depends entirely on the effectiveness of the underlying GITCs. If change management controls are weak, an ITAC that appears to function correctly today may have been altered without authorization. For this reason, auditors test GITCs before relying on automated application controls [48].

IT Dependent Manual Controls (ITDMs) ITDMs are manual controls that depend on IT-generated output to function. The control itself is performed by a person, but only because an IT system surfaces the information or trigger that makes the control possible [33]. A concrete example is a financial controller who reviews a system-generated exception report each week and manually investigates flagged transactions. The review is human, but if the underlying system produces incorrect output, the manual control fails even if the person performs it diligently. In the context of cloud exit capability, ITDMs are particularly relevant. Many of the controls assessed by the artifact in chapter 4 are ITDMs: an exit plan that a CIO approves annually based on system-generated dependency reports, a key-person analysis that a manager updates using a skills matrix produced by HR systems, or a rollback decision during migration that a team lead makes based on real-time monitoring alerts. Auditors assessing these controls must recognize that the quality of the IT-generated input is as important as the quality of the human judgment applied to it [33].

2.1.3 Internal versus external IT auditing

Internal IT audit

Internal IT auditing is conducted by employees with auditing responsibility within the organization and is commonly positioned as a governance mechanism that supports management and the board. Internal audit focuses on whether internal controls and governance processes are designed and operated in a way that supports organizational objectives, manages risk, and enables compliance. Internal audit is often motivated by improvement and operational assurance, and it typically leads to recommendations that management can implement to reduce risk or increase efficiency [27].

Internal audit independence is functional rather than absolute. Internal auditors are expected to be objective, but they remain part of the organization. This means that internal auditors may, in practice, give primacy to supporting the organization's strategic objectives rather than providing fully independent assurance. This makes internal audits especially useful for continuous monitoring, iterative improvements, and embedding control awareness into daily practices [53].

External IT audit

External IT auditing is performed by auditors who are independent from the audited organization. An external audit is commonly driven by external accountability requirements such as financial reporting

obligations, regulatory mandates, or certification needs. An external audit emphasizes demonstrable assurance to third parties. The purpose is not primary improvement but credible verification that stated requirements are met [27].

External audits are also commonly commissioned by service organizations, such as cloud providers or software suppliers, to demonstrate to their clients that their operations are well-controlled and that they deliver reliable services. A cloud provider that undergoes a SOC 2 audit, for example, can provide its customers with an independent assurance report confirming that specified controls over security, availability, and confidentiality were in place and operating effectively during the audit period [2].

In financial statement audits, external auditors must understand the IT environment and its role in producing information used in financial reporting. ISA 315 explicitly requires auditors to assess risks of misstatements arising from the use of IT, including automated processing and system-generated information [33]. Because financial information is increasingly produced, processed, and stored within IT systems, the reliability of financial reporting depends directly on the reliability of those systems. IT audit work is therefore structurally important even in engagements whose primary objective is financial, since a financial auditor who does not understand the IT environment cannot reliably assess whether the numbers they are auditing are complete and accurate.

2.1.4 What gets audited?

IT auditors do not only audit systems in isolation. They assess an audit universe that spans technologies, processes, controls, and organizational capabilities, all of which together determine whether IT audit supports organizational objectives and manages risk effectively [27]. This includes IT governance structures, IT processes such as change management and continuity planning, IT assets such as applications, databases, and cloud services, and the control mechanisms that govern them.

What matters for this thesis is that cloud exit capability falls directly within this audit universe. Whether an organization can exit a cloud provider without catastrophic loss involves IT governance (who owns the exit plan and is accountable for its maintenance), IT processes (migration procedures, rollback mechanisms, and data export workflow), IT assets (data portability and infrastructure configuration), and control mechanisms (contractual exit rights and testing programmes). Each of these is an established IT audit object. Yet despite this, no existing audit framework currently provides auditors with structured criteria, evidence requirements, and maturity levels to assess cloud exit capability in a systematic and defensible manner. What this means in practice depends on which controls within this universe are actually relevant to cloud exit, and that question is not straightforward.

Of the controls that fall within the audit universe described above, not all are equally relevant to cloud exit capability. A control is any mechanism that reduces the likelihood or impact of an undesired outcome, whether a policy, a technical enforcement rule, or a workflow step [27]. In the context of cloud dependency, the relevant controls are not primarily cybersecurity or access management controls. They are the planning, documentation, portability, continuity, and testing controls that determine whether an organization can actually execute an exit when needed. The absence of these controls represents a category of IT risk that is currently implicit in audit practice rather than explicitly assessed, and it is this absence that motivated the framework developed later in this thesis. The following sections examine whether existing European sovereignty frameworks address this gap and identify precisely where and why they fall short.

2.1.5 Why IT audits are performed

Laws and regulations

Aside from requirements that ensure an IT audit is part of financial audits, as discussed above, regulations around cybersecurity and data protection create independent legal obligations for organizations to demonstrate that their IT controls are effective. In the EU, the NIS2 Directive explicitly establishes a broad cybersecurity risk management and accountability regime across critical sectors, and it shapes what organizations must implement and what auditors must be able to evaluate [19].

Certifications and standards as formalized baselines

Certification standards create a clear audit baseline. ISO/IEC 27001 is particularly important because it defines requirements for an Information Security Management System (ISMS) and is widely used as a certification target and as an audit reference model [39]. The 2022 revision is the currently published version and is frequently used in cloud and outsourcing contexts as a governance anchor for security assurance [37].

In practice, certification also influences auditability. It forces organizations to document control objectives, assign responsibilities, and maintain evidence trails. Those features are exactly what auditors require in order to form evidence-based conclusions [27].

Operational effectiveness and continuous improvement

Not all audits are external forces. Many organizations audit IT internally to improve reliability, reduce incidents, strengthen service management, and demonstrate maturity. This logic aligns with continuous improvement models such as PDCA, where audits are part of organizational learning cycles rather than only compliance exercises [27].

2.1.6 The IT audit process

Although audit methodologies may differ across firms, the IT audit process generally follows a structured risk-oriented sequence aligned with ISA 315 and ISA 330.

IT understanding and scoping

The initial phase involves obtaining an in-depth understanding of the entity's IT environment. Auditors evaluate system architecture, data flows, IT governance structures, tasks or processes outsourced to service providers, and cybersecurity risk exposure. This understanding enables identification of relevant IT systems that support financially significant processes [33]. Risk assessment procedures include walkthroughs, inquiries, system documentation reviews, and evaluation of prior-year findings. The outcome of this phase is a defined audit scope targeting systems and controls relevant to the audit objective. In financial statement audits, this objective is the reliability of financial reporting. In non-financial audits, such as cybersecurity assessments or the cloud exit capability assessment that this thesis develops, the same understanding and scoping logic applies: the auditor must understand the IT environment, identify what is relevant to the objective at hand, and define a scope that is both complete and proportionate.

Artifact collection

Following scoping, auditors obtain documentation and evidence supporting the design and implementation of IT controls. This includes for example:

- IT policies and procedures
- Access control listings
- Change management logs
- System configuration settings
- Incident reports
- Service Organization Control (SOC 1) reports for outsourced IT providers

Documentation serves as preliminary evidence for assessing whether controls are appropriately designed to mitigate identified risks [2].

Control testing

Control testing generally consists of two components:

1. **Test of Design (TOD)**: Auditors assess whether controls are suitably designed and have been implemented. A control is properly designed if it effectively addresses a specific risk of material misstatement. Implementation is verified by inspecting documentation and observing control execution [33].

2. **Test of Operating Effectiveness (TOE)**: If reliance on a control is planned or sometimes discretionary, auditors perform tests of operating effectiveness to determine whether the control operated consistently throughout the audit period. Testing methods may include inspection of logs, reperformance, and inquiry, combined with corroboration and sampling techniques [48]. Sample sizes are determined based on risk assessments, control frequency, and desired assurance levels. Controls addressing higher-risk areas typically require larger sample sizes and more rigorous testing procedures [33].

Client clarification and remediation

During and after testing, auditors communicate preliminary findings to management. Where deficiencies are identified, management may provide additional evidence or implement remediation measures. Depending on timing severity, auditors may retest controls post-remediation [33].

Deficiencies are classified according to severity (control deficiency, significant deficiency, or material weakness), consistent with professional standards [48].

Reporting and completion

The final stage involves evaluating the aggregate impact of identified deficiencies on audit risk and financial statement reliability. IT control deficiencies may affect the auditor's reliance on controls and increase the need for substantive testing. Auditors communicate significant deficiencies and material weaknesses to those charged with governance in accordance with ISA 265 [33]. In integrated audits (e.g., under SOC), opinions may also be expressed on internal control over financial reporting [48].

2.2 Cloud Computing fundamentals

Understanding cloud computing is essential for this thesis because the sovereignty risks that the framework addresses, vendor lock-in, loss of exit capability, and reduced organizational control arise directly from how cloud services are structured, delivered, and governed. This section explains the relevant concepts and establishes the technical and governance context for the framework developed in chapter 4.

2.2.1 Definition

Cloud computing is typically defined as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or with little to no interaction with a service provider [41].

ISO/IEC 17788 provides a standardization perspective on the same concept, which frames cloud computing as an approach to enabling network access to a scalable and elastic pool of shareable physical or virtual resources and provides a shared vocabulary for cloud concepts used across governance and assurance contexts [35].

In academic treatments, cloud computing is frequently described as a paradigm that delivers computing services over the internet using virtualization and distributed architectures, while emphasizing both efficiency gains and a distinct security and risk profile [35].

Cloud adoption is therefore not only a technical migration choice but also a governance decision that changes control boundaries, accountability structures, and the locus of authority over infrastructure and data [5].

2.2.2 Cloud service models

The cloud service taxonomy is typically structured around three primary service models, which represent increasing levels of abstraction and provider management of the computing stack [41].

Infrastructure as a Service (IaaS) IaaS provides fundamental computing resources such as processing, storage, and networking, where the consumer can deploy and run arbitrary software while retaining significant responsibility for operating systems and applications [41].

Platform as a Service (PaaS) PaaS provides the capability to deploy consumer-created or acquired applications onto a provider-managed platform, reducing the consumer's responsibility for underlying infrastructure while increasing dependency on provider runtime services and APIs [41].

Software as a Service (SaaS) SaaS provides access to a provider-hosted application where the consumer primarily configures usage and access, while the provider manages the application and underlying infrastructure [41].

Academic discussions often frame these service models as "pay-per-use" service delivery forms that shift cost structures and scalability capabilities while simultaneously changing the organization's control over security configuration and operational decision-making [41].

Which service model an organization uses determines how much control it retains over its own infrastructure and how difficult exit becomes. In IaaS environments, organizations manage their own operating systems and applications, making migration technically complex but feasible. In PaaS and SaaS environments, increasing provider management reduces that control, deepening dependency and raising the stakes for exit capability assessment.

2.2.3 Cloud deployment models

Cloud deployment models describe how cloud services are organized based on control and sharing of physical or virtual resources, which directly affects governance, jurisdictional exposure, and dependency risk [35].

The widely used deployment model taxonomy includes private cloud, public cloud, community cloud, and hybrid cloud [41].

"Public cloud" refers to a deployment model where cloud services are potentially available to any customer, and resources are controlled by the cloud service provider, which typically increases dependency on the cloud provider for external operational and legal conditions [35].

"Private cloud" refers to cloud infrastructure provisioned for exclusive use by a single organization, which may increase control but does not automatically eliminate third-party dependencies where managed services or co-location are involved [41].

Hybrid cloud refers to a composition of two or more distinct cloud infrastructures that remain unique entities but are bound together, which introduces interoperability and data movement challenges that are directly relevant for portability and sovereignty objectives [41].

Community cloud refers to cloud infrastructure provisioned for exclusive use by a specific community of organizations with shared concerns, making it particularly relevant for regulated sectors that attempt to embed shared compliance and governance requirements into the deployment model [41].

For cloud exit assessment purposes, the deployment model matters because it shapes what an organization actually controls and therefore what it can take with it when leaving a provider. A private cloud organization typically has more direct access to its infrastructure and data. A public cloud organization must rely on the provider's export mechanisms and contractual exit provisions, which are precisely what the frameworks control areas in chapter 4 are designed to assess.

2.2.4 Shared responsibility models

A central governance concept in cloud computing is the shared responsibility model, which describes how security and control responsibilities are divided between the cloud service provider and the cloud customer [42].

The shared responsibility model is practically important because it clarifies that adopting cloud services does not transfer accountability to the provider but reallocates operational responsibilities across organizational boundaries [42]. The customer remains accountable for outcomes even when the provider is responsible for execution.

Professional guidance highlights that responsibility allocation differs by deployment model and service model, with customers generally retaining more responsibility in IaaS and less in SaaS, while providers retain responsibility for underlying infrastructure operations [32].

From a security assurance perspective, the shared responsibility model implies that effective governance requires explicit delineation of control ownership and evidence sources, because relevant controls may be implemented by the provider, the customer, or both.

2.2.5 Cloud dependencies and risk factors

Cloud computing introduces distinctive risk factors because it is characterized by multi-tenancy, virtualization, and shared resource pooling, which create potential isolation failures if segmentation mechanisms fail [5].

Cloud environments may increase the risk of "loss of governance," where organizations reduce direct operational control over assets and must rely on provider transparency, contractual rights, and assurance reports [21].

Vendor lock-in is a recurring dependency risk in cloud adoption. Proprietary interfaces, data formats, and managed service entanglements can limit the portability of data and workloads, making it difficult for organizations to migrate to alternative providers or bring services back in-house [21].

Cloud security risk assessments also identify management interface compromise as a critical category because cloud control planes concentrate administrative power and can become high-impact attack targets [21].

Academic overviews of cloud risks emphasize that confidentiality, integrity, and availability concerns are shaped not only by technical controls but also by organizational dependencies such as provider governance, subcontracting chains, and geographic distribution of data processing [21].

ENISA also emphasized that cloud adoption creates both security risks and security benefits, because providers can achieve economies of scale in security investment while simultaneously concentrating systemic risk and shared platforms [21].

These dependency dynamics are directly relevant for digital sovereignty because sovereignty concerns are frequently expressed through control over data location, legal exposure, substitutability of providers, and the feasibility of exit or migration [35].

For the purpose of this thesis, the most relevant risk factors are those that directly affect an organization's ability to exit a cloud provider: data portability constraints, dependency on provider-specific tooling, absence of documented exit procedures, and the concentration of operational knowledge in provider systems or personnel. These are precisely the risks that the framework's control areas are designed to surface and assess.

2.2.6 Cloud governance and security standards

Cloud governance and security are commonly structured through internationally recognized standards and control frameworks that define baseline expectations for security management, assurance, and accountability [7]. The Cloud Security Alliance Cloud Controls Matrix (CCM) is specifically designed as a cloud-focused control framework that supports assessment of cloud service security across governance, risk management, compliance, and technical control domains [7].

ISO/IEC 27017 provides cloud-specific guidance for information security controls applicable to both cloud service providers and cloud customers, thereby explicitly supporting shared responsibility and control allocation [36].

ISO/IEC 27018 specifies control objectives and guidelines for protecting personally identifiable information in public cloud environments, which is particularly relevant where sovereignty concerns are expressed through privacy, lawful processing, and data handling constraints [38].

At the EU level, ENISA's work on the European Cybersecurity Certification Scheme for Cloud Services (EUCS) represents an effort to harmonize cybersecurity certification expectations for cloud services, even though the scheme has been developed through drafts and policy debate over the inclusion of sovereignty-related requirements [21].

Notably, none of these standards directly addresses operational reversibility or cloud exit capability, as an assessable dimension [36, 38]. ISO 27017 and ISO 27018 focus on security controls during operation,

not on what happens when an organization needs to leave its provider [36, 38]. The EUCS addresses sovereignty at the provider certification level, not at the organization audit level [21].

2.3 Digital sovereignty

Understanding the conceptual foundations of digital sovereignty is essential for developing structured approaches to assessing sovereignty risks in digital environments. The concept has gained increasing attention in both academic research and policy debates, particularly in response to growing dependence of states and organizations on global digital infrastructures operated by a small number of international technology providers [11].

Despite its prominence in contemporary digital policy discourse, digital sovereignty remains an evolving and contested concept, with different interpretations emerging across political science, information systems, and technology governance literature [45].

In addition to digital sovereignty, the closely related concept of digital autonomy frequently appears in academic and policy discussions on technological independence and strategic control over digital infrastructures [45]. While the two concepts are often used interchangeably in public discourse, academic literature typically distinguishes between them in terms of governance authority and operational capability.

Establishing clear conceptual definitions of both digital sovereignty and digital autonomy is therefore necessary in order to develop a coherent analytical framework for evaluating sovereignty-related risks in organizations that rely on cloud computing services.

2.3.1 Digital sovereignty

The concept of digital sovereignty originates from the broader political and legal notion of sovereignty, which traditionally refers to the authority of a state to exercise control over its territory, institutions, and governance without external interference [45]. Historically, sovereignty has been associated with territorial authority and the ability of states to establish and enforce rules within their jurisdiction [26]. However, the rapid expansion of digital technologies and globally interconnected communication networks has fundamentally challenged the territorial foundations of sovereignty.

Digital infrastructures such as cloud computing platforms, global data networks, and distributed computing architectures operate across national borders, often under the control of private corporations rather than state institutions. As a result, traditional mechanisms of state authority over infrastructure and information flows have become increasingly difficult to apply in digital environments [11].

The concept of digital sovereignty emerged as a response to these developments and reflects attempts to re-establish forms of authority and control in digital ecosystems that are characterized by technological interdependence and cross-border governance challenges [45].

In academic literature, digital sovereignty is generally understood as the capacity of states, organizations, or individuals to exercise meaningful control over digital infrastructures, technological development, and data governance within their sphere of authority [50].

Several scholars emphasize that digital sovereignty involves the ability to determine how digital technologies are developed, deployed, and governed, as well as the ability to regulate digital activities according to locally defined legal and societal norms [45, 11, 50].

From this perspective, digital sovereignty encompasses not only control over technological infrastructures but also governance authority over digital ecosystems, including the regulation of digital platforms, the management of cross-border data flows, and the establishment of cybersecurity frameworks [45].

Other interpretations emphasize the importance of data governance as a core dimension of digital sovereignty, highlighting the ability of actors to control the storage, processing, and transfer of digital information [57].

Across these interpretations, digital sovereignty can therefore be conceptualized as a multidimensional construct that includes technological independence, governance authority, and control over digital resources. Scholars commonly identify several key dimensions of digital sovereignty, including infrastructural sovereignty, data sovereignty, technological sovereignty, and regulatory sovereignty [11].

- Infrastructural sovereignty refers to control over the physical and virtual infrastructures that enable digital communication and computing services.
- Data sovereignty refers to the ability to control how data is collected, stored, processed, and transferred across jurisdictions.
- Technological sovereignty refers to the capacity to develop and operate digital technologies independently from external actors.
- Regulatory sovereignty refers to the authority to define and enforce legal frameworks governing digital activities within a specific jurisdiction.

Taken together, these dimensions illustrate that digital sovereignty extends beyond purely technical considerations and includes governance, economic, and geopolitical aspects of digital ecosystems.

2.3.2 Digital autonomy

Closely related to digital sovereignty is the notion of digital autonomy. Where digital sovereignty concerns authority and governance in a legal and institutional sense, digital autonomy is best understood as its operational dimension: the practical capability of an organization or state to act independently within digital ecosystems without excessive reliance on external technological actors [50].

Digital autonomy is frequently discussed in the context of industrial policy and the development of domestic technological ecosystems that reduce dependence on foreign technology providers [45]. In European policy debates, it is framed as a strategic objective aimed at strengthening the capacity of European institutions to develop and operate critical digital infrastructures [9].

Consequently, digital autonomy and digital sovereignty should not be understood as complementary elements within a broader governance framework. Digital sovereignty focuses primarily on authority, governance, and regulatory control over digital environments, while digital autonomy emphasizes the technological and operational capabilities required to exercise that authority effectively.

An important conceptual distinction must be acknowledged here. In the strict political and legal sense, sovereignty refers to the authority of a state over its territory and governance, implying a legal and institutional basis for control [54]. Autonomy, by contrast, refers to the practical capability of an actor, whether a state or an organization, to make independent decisions and determine its own actions without excessive dependence on external parties [54].

This distinction matters for this research. The framework developed in this study assesses whether organizations can exit their cloud provider without catastrophic loss of continuity, data, or operational knowledge. That is primarily an organizational capability question, which places it closer to digital autonomy than to digital sovereignty in the state-level, legal sense.

Nevertheless, the term *digital sovereignty* is used throughout this thesis for two reasons. First, the European policy framework that forms the comparative basis of this study explicitly applies the term to organizations and procurement contexts, not only to states. The term has already been operationalized at the organizational level in European policy practice, as using consistent terminology makes the contribution of this study traceable to those frameworks. Second, organizational autonomy and state sovereignty are not independent. An organization that cannot exit its cloud provider is an organization whose operational continuity is subject to decisions made by a foreign corporation operating under foreign law. That dependence has direct implications for the sovereignty of the state in which the organizations operate, which is precisely the concern that European digital sovereignty policy is responding to. Operational reversibility is therefore best understood as the organizational expression of a broader sovereignty objective: the point at which an abstract policy goal becomes something that can be assessed, evidenced, and improved.

2.4 Existing cloud sovereignty frameworks

Digital sovereignty frameworks have been developed in response to a widening gap between Europe's regulatory ambitions and its continued dependence on external providers for cloud infrastructure, software ecosystems, data processing capacity, and digital innovation [51].

The frameworks in this field operate at different levels of intervention, and that distinction matters for this research. EuroStack and the EC-CSF operate primarily at the level of European industrial strategy and supranational policy: they define what digital sovereignty should look like at a continental scale. The Buy European Framework operates at the level of procurement decisions: it helps organizations choose sovereign providers. DICTU operates at the level of national operationalization: it translates European policy into assessable criteria for Dutch government procurement. None of these frameworks operates at the level where an IT auditor works: the level of an individual organization's controls, evidence, and audit conclusions. This gap, between defining sovereignty at a policy level and assessing it at an organizational audit level, is the central problem this thesis addresses.

2.4.1 European Commission Cloud Sovereignty Framework

Origin and institutional context

The European Commission Cloud Sovereignty Framework (hereafter EC-CSF), published by the Directorate-General for Digital Services in version 1.2.1 (October 2025) [16], forms the most authoritative supranational assessment framework currently available for the systematic assessment of cloud services on sovereignty criteria [16]. Its provenance lies in a convergence of regulatory and political developments: the European Cybersecurity Certification Scheme (EUCS) developed by the European Union Agency for Cybersecurity (ENISA); the Gaia-X policy rules and trust framework; and CIGREF's Trusted Cloud Referential (v2) [16]. The EC-CSF explicitly acknowledges these antecedents and situates itself within the broader landscape of national cloud sovereignty strategies, citing France's "Cloud de Confiance" and Germany's "Souveräner Cloud" as precedents [16].

The institutional significance of the EC-CSF is considerable. As a document produced by the European Commission's Directorate-General for Digital Services and intended for use in public procurement procedures, it carries quasi-regulatory weight [16]. It operationalizes sovereignty as a procurement criterion, that is, as a dimension of contract award that can be assessed, scored, and used to differentiate between bidders. This marks a shift from treating digital sovereignty as a normative aspiration to treating it as a verifiable, auditable characteristic of a cloud service offering.

Conceptual structure

The EC-CSF is structured around eight Sovereignty Objectives (SOV-1 through SOV-8), which together define the scope of what sovereignty means in the context of cloud service procurement [16]. These objectives cover a broad set of sovereignty-related concerns, including strategic control, legal jurisdiction, operational independence, technological autonomy, supply-chain governance, and regulatory compliance. The framework, therefore, conceptualizes digital sovereignty as a multidimensional construct rather than a single technical property. Table 2.1 presents the eight Sovereignty Objectives as defined in the EC-CSF.

Scoring methodology and SEAL model

A distinctive characteristic of EC-CSF is the introduction of a structured scoring methodology that produces an overall Sovereignty Score based on weighted contributions from each Sovereignty Objective [16]. Scores for individual objectives are normalized and combined using predefined weights that reflect policy priorities. For example, supply-chain sovereignty is assigned a higher weight than environmental sustainability, reflecting the perceived importance of dependency risks in cloud environments [16].

The framework introduces Sovereign Effective Assurance Levels (SEAL) as mechanisms to express the degree to which sovereignty objectives are effectively achieved [16]. The SEAL concept represents a structured interpretation of sovereignty assurance, linking assessment questions to observable evidence and documented controls. Rather than representing organizational maturity, SEAL levels (table 2.2) indicate the strength and reliability of sovereignty-related assurances. This distinction is important because sovereignty assessments are intended to evaluate verifiable controls and guarantees rather than organizational development stages.

The SEAL model represents a significant step towards making digital sovereignty auditable. The framework explicitly connects sovereignty objectives to assessment questions and evidence expectations, allowing sovereignty claims to be verified through documented controls and technical configurations [16]. This structure aligns closely with established IT audit methodologies in which audit conclusions are derived from evidence-based evaluation of controls.

Table 2.1: EC-CSF Sovereignty Objectives (SOV-1 through SOV-8) as defined in the European Commission Cloud Sovereignty Framework v1.2.1 [16].

Criterion / Dimension	**Description / Key Criteria**
SOV-1: Strategic Sovereignty	Ownership stability, EU legal/financial ecosystem anchoring, governance influence, alignment with EU industrial strategy, resistance to service suspension.
SOV-2: Legal & Jurisdictional Sovereignty	Governing legal system, exposure to extraterritorial laws (US CLOUD Act, FISA 702), channels for foreign access compulsion, IP jurisdiction.
SOV-3 Data & AI Sovereignty	Customer-exclusive cryptographic control, data access auditability, irreversible deletion, strict EU-only storage/processing, EU-controlled AI model governance
SOV-4: Operational Sovereignty	Portability and exit planning, EU-operator capability without non-EU vendor, EU talent pool, exclusive EU legal operational support
SOV-5: Supply Chain Sovereignty	Geographic origin of hardware/firmware, software architecture jurisdiction, non-EU dependency degree, full supply chain auditability
SOV-6: Technology Sovereignty	Open API/protocols, open licensing and source code access, architectural transparency, EU independence in high-performance computing
SOV-7: Security & Compliance Sovereignty	EU/international certifications (ISO, ENISA, NIS2, DORA, GDPR), EU-only SOC operations, independent EU audit capacity
SOV-8: Environmental Sustainability	Energy efficiency (PUE), circular economy hardware practices, transparent carbon water reporting, renewable energy sourcing

Each of the eight Sovereignty Objectives is assessed against the SEAL taxonomy, and a minimum SEAL level is specified per objective for each procurement procedure. Bids that do not meet the minimum SEAL threshold across all objectives are rejected. This binary gate function introduces important structural parallels with other frameworks discussed in this chapter: the distinction between threshold-based pass/fail criteria and graded scoring criteria.

The overall sovereignty score is computed as a weighted sum of normalized per-objective scores. The weights reflect the Commission's assessment of risk distribution; Supply Chain Sovereignty (SOV-5) carries the highest weight at 20%, followed by Strategic Sovereignty (SOV-1) and Operational Sovereignty (SOV-4) at 15% each and Technology Sovereignty (SOV-6) at 15%. Security and Compliance (SOV-7) and Legal and Jurisdictional (SOV-2) receive lower weights of 10% each. The framework's rationale is that procurement procedures already incorporate significant legal safeguards through standard contractual mechanisms. Environmental Sovereignty (SOV-8) is weighted at 5%, reflecting its secondary but growing role in procurement criteria.

Institutional interpretation of digital sovereignty

The EC-CSF reflects a specifically European interpretation of digital sovereignty that emphasizes strategic autonomy, regulatory control, and technological independence. Within this interpretation, sovereignty is closely associated with the ability of European organizations and institutions to maintain control over data, infrastructure, and operational decision-making despite the reliance on external service providers [16].

This interpretation is consistent with broader European policy developments in which digital sovereignty is framed as a prerequisite for economic resilience and regulatory effectiveness. Sovereignty is therefore treated as a practical governance objective rather than an abstract political concept. The framework operationalizes this policy perspective by translating sovereignty into observable characteristics such as jurisdictional control, technical interoperability, and auditability of cloud operations [16].

Table 2.2: EC-CSF Sovereign Effective Assurance Levels (SEAL-0 through SEAL-4), defining the sovereignty classification scale used in the European Commission Cloud Sovereignty Framework [16].

Criterion / Dimension	**Description / Key Criteria**
SEAL-0: No Sovereignty	Service entirely under non-EU control; governed in non-EU jurisdictions.
SEAL-1: Jurisdictional Sovereignty	EU law formally applies, but with limited practical enforceability, a non-EU third-party retains exclusive control.
SEAL-2: Data sovereignty	EU law applicable and enforceable; material non-EU dependencies remain; indirect non-EU control.
SEAL-3: Digital resilience	EU law enforceable, EU actors exercise meaningful but not full influence, and marginal non-EU control.
SEAL-4: Full digital sovereignty	Complete EU control over technology and operations, subject exclusively to EU law; no critical non-EU dependencies.

Table 2.3: EC-CSF objective weights

Sovereignty Objectives	**Weight in scoring**
SOV-1: Strategic Sovereignty	15%
SOV-2: Legal & Jurisdictional Sovereignty	10%
SOV-3 Data & AI Sovereignty	10%
SOV-4: Operational Sovereignty	15%
SOV-5: Supply Chain Sovereignty	20%
SOV-6: Technology Sovereignty	15%
SOV-7: Security & Compliance Sovereignty	10%
SOV-8: Environmental Sustainability	5%
Total	100%

2.4.2 EuroStack

Origin and institutional context

The EuroStack initiative represents a qualitatively different category of framework from the EC-CSF. Where the latter is a procedural instrument embedded in procurement law, EuroStack is a macro-level industrial policy proposition: a mission-driven call for the construction of a comprehensive, sovereign European digital infrastructure ecosystem. The term "EuroStack" was coined by economist Glen Weyl in 2024 and rapidly catalyzed a coalition of technologists, academics, and industry representatives across Europe [24].

The flagship policy report, authored by Bria, Timmers and Gemoets and commissioned by the Bertelsmann Stiftung with support from the Center of European Policy Studies (CEPS) [24]. Stiftung Mercator, and UCL Institute for Innovation and Public Purpose (UCL IIPP), forms the most academically grounded articulation of the initiative [24]. The document has been widely circulated among European technology policy stakeholders and serves as a reference point for the broader EuroStack legislative advocacy [24]. A companion open letter to European Commission President von der Leyen, signed by a broad coalition of European businesses and associations in March 2025 [22], further institutionalizes the initiative's political ambitions.

The significance of EuroStack lies in its explicit framing of digital dependency as a structural rather than merely technical problem. The report observes that more than 80% of core digital technologies used in Europe are imported, that 70% of foundational AI models used globally originate in the United States, and that European companies account for just 7% of global research spending on software and the internet (Bria, 2025). The figures situate the sovereignty dilemma within a broader political economy of digital colonialism, echoing the work of scholars who have argued that the EU risks becoming a "profitable consumer market" for non-EU platform firms rather than a hub of technological creation [50].

Conceptual structure: The seven-layer technology stack

Conceptually, the EuroStack framework is organized around the metaphor of a technology stack, a layered architecture that integrates raw materials and hardware at the infrastructure layer through networks.

IoT and devices, cloud infrastructure, software platforms, and data and AI at the application layer. The framework argues that sovereignty interventions must operate across all layers simultaneously, as vulnerabilities at any layer propagate upward. This systemic, layer-based conceptualization resonates with Floridi's multi-dimensional understanding of digital sovereignty [26], which distinguishes control at the physical code and information layers.

The EuroStack framework is grounded in seven guiding principles, which together constitute its normative architecture:

Table 2.4: EuroStack sovereignty dimensions

Criterion / Dimension	**Description / Key Criteria**
Sovereignty and Security	Digital infrastructure must be under EU legal jurisdiction and operational protection, insulating European public and private actors from extraterritorial legal coercion.
De-proprietarization and Interoperability	A federated, open-source technology stack that prevents vendor lock-in and ensures workload portability across providers and jurisdictions.
Sustainability	Energy-efficient climate-aligned design principles across all infrastructure layers.
Data as a common good	Data as a shared resource governed by transparent, rights-protective mechanisms that balance innovation with citizen privacy.
Decentralized Sovereign Infrastructure	Combing edge computing with centralized systems to maximize both operational efficiency and jurisdictional sovereignty.
Inclusive Governance	Harmonized EU regulation and accountability mechanisms that balance resilience with democratic autonomy.
Strong democracy & Compliance Sovereignty	Digital technologies designed to serve and reinforce democratic governance structures, not to undermine them.

The Open Letter: Immediate goals and policy demands

The accompanying Open Letter (March 2025) translates the EuroStack report's strategic vision into four immediate policy demands addressed to the European Commission [22]. These demands are worth examining carefully, as they reflect the initiative's theory of change.

First, the Open Letter calls for the creation of demand through procurement obligations, specifically a formal "Buy European" requirement for the public sector [22]. This demand anticipates and motivates the Buy European Procurement Framework examined in section 2.4.3 of this chapter. Second, it advocates for viable supply through a "pooling and federating" approach: aggregating existing dispersed European digital assets and solutions into coherent, commercially available offerings based on open standards and interoperability [22]. Third, it calls for the prioritization of services with high adoption potential, particularly those satisfying cybersecurity objectives, with key performance indicators defined at the project level based on market outcomes [22]. Fourth, it requests the development of harmonized certification requirements for sovereign cloud services, specifically augmenting the EUCS to include a High+ sovereignty level or finding an adequate alternative if the EUCS process remains potentially blocked [22].

2.4.3 The EuroStack Buy European Procurement Framework

Origin and institutional context

The EuroStack Industry Initiative's "A proposed framework for a 'Buy European' Regulation of Strategic Digital Procurement" (September 2025) [23] is the most legally focused of the four frameworks examined. Authored by multiple supporters of the EuroStack Industry Initiative and subject to legal review from procurement specialists, it is explicitly positioned as a proposal for binding EU regulation, the most direct and uniform instrument of EU law under Article 288 TFEU [23]. The document explicitly disclaims any lobbying affiliation, presenting itself as a volunteer effort by unpaid technologists and economists motivated by shared concern about European digital dependency.

Its intellectual and political context is the perceived inadequacy of existing EU procurement law to ad-

dress digital sovereignty concerns. The document opens by characterizing the EU's default procurement posture as "buy from wherever," as a principle of non-discrimination that, in the digital sector, has systematically advantaged non-EU hyperscalers at the expense of European providers. In contrast, the United States maintains a "Buy American Act" default that privileges domestic procurement, and China effectively mandates local suppliers. The Buy European Framework proposes to reverse this asymmetry for strategic digital technologies through a new EU regulation invoking the national security exception of the WTO Government Procurement Agreement (GPA) [23].

The five dimensions of a sovereign European Provider

The Buy European framework defines "Sovereign European Provider" as a technical qualification, not a badge of trust or a moral judgment, but a set of objective, auditable criteria that a provider either meets or does not meet. This framing reflects a deliberate rejection of vague marketing labels and responds to the documented risk of "sovereignty washing," the practice of labeling non-European services as "European" through subsidiary structures of contractual arrangements that offer no genuine protection against extraterritorial legal coercion.

Table 2.5 presents the five dimensions and their key criteria.

Table 2.5: The five hierarchically ordered dimensions of a Sovereign European Provider as defined in the EuroStack Buy European Framework for Strategic Digital Procurement [23]. Dimensions I through IV are mandatory requirements; Dimension V functions as a key differentiator in award scoring.

Criterion / Dimension	**Description / Key Criteria**
I. Jurisdiction and Governance (Mandatory Pass/Fail)	Ultimate parent, entity incorporated and headquartered within EU/EEA/EFTA; majority voting rights held by European entities/citizens; no non-EU blocking minority stake; no structural exposure to CLOUD Act, FISA 702, or non-EU export control on core IP.
II. Technical Sovereignty (Core Guarantee)	Open standards and predominantly OSI-approved open-source software for core components; source code available for inspection (public or European escrow); service fully documented and automatable for third-party redeployment; no punitive data egress fees.
III. Operational Sovereignty (Core Guarantee)	Physical infrastructure and control plane exclusively within European territory; 100% of privileged access personnel resident, employed, hardware/software dependency mitigation.
IV. Data Sovereignty (Core Guarantee)	All data (primary, backups, logs, metadata) stored and processed exclusively in EU territory; verifiable technical measures (confidential computing or customer-exclusive key management) making provider access cryptographically impossible; and there is a contractual obligation to legally challenge and notify customer of non-EU data access requests.
V. Economic Sovereignty (Key Differentiator)	>50% of global R&D expenditure and personnel for core technology located in EU; no punitive data egress fees or tied selling; active partnerships with European SMEs and meaningful open-source contributions.

The three-level procurement strategy

A central methodological contribution of the Buy European framework is its three-level hierarchy to apply sovereignty criteria in public procurement procedures. This hierarchy maps the five dimensions to distinct stages of the procurement evaluation process, providing a clear procedural architecture for contracting authorities.

Level-1 Mandatory Prerequisite (Pass/Fail): Dimension I (Jurisdiction and Governance) functions as a binary eligibility gate at the outset of the evaluation process. Any provider failing to demonstrate full compliance, specifically an ultimate parent entity incorporated and headquartered in EU/EEA/EFTA, free from non-EU de jure and de facto control, and structurally insulated from non-EU extraterritorial laws, is disqualified before technical evaluation begins. This is grounded in the legal logic that a provider

subject to non-EU legal authority is categorically incapable of guaranteeing a sovereign service, regardless of its technical capabilities.

Level-2 Core Guarantees (Technical specifications and mandatory minimum standards): Dimensions II, III, and IV are evaluated through a combination of mandatory technical specifications (non-negotiable minimum requirements) and qualitative award criteria. The principle of proportionality governs the specification of minimum levels: a contract involving sensitive national security data would demand the highest possible rating across all three dimensions, while a contract for citizen health records might permit a higher minimum on Data Sovereignty than on Technical Sovereignty.

Level-3 Key differentiator (Award Criteria): Dimension V (Economic Sovereignty) functions as a differentiator among providers that have successfully passed Levels 1 and 2. Criteria such as location of R&D expenditures, commitment to European SME partnerships, and open-source contributions create an incentive structure that rewards providers genuinely invested in the European digital ecosystem, not only those meeting minimum compliance requirements.

Legal foundation

A unique contribution of the Buy European framework is its detailed legal analysis of how a "Buy European" mandate could be structured to comply with international trade law. The document identifies Article III of the WTO Government Procurement Agreement (the "Essential Security Interests" exception) as the legal foundation for a proportionate, targeted digital procurement preference. The argument proceeds as follows: public sector digital infrastructure and the sensitive data it contains constitute a national security interest; Europe's exposure to the extraterritorial reach of foreign laws through non-EU cloud dependencies forms a direct and verifiable security risk; therefore, a procurement preference for structurally sovereign providers is a legitimate exercise of sovereign security prerogative, not an act of economic discrimination [23]. The framework also grounds this argument in Article 346 TFEU, which affirms member states' right to take measures necessary to protect their essential security interests

2.4.4 The DICTU Toetsingsinstrument Soevereiniteit Clouddiensten

The DICTU Toetsingsinstrument Soevereiniteit Clouddiensten (Assessment Instrument for Cloud Service Sovereignty), published by the Dienst ICT Uitvoering (DICTU) in version 1.0.1 (January 2026), represents the most operationally specific of the four frameworks examined here. DICTU is the ICT implementation agency of the Dutch central government (Rijksoverheid), providing shared digital infrastructure and services to multiple Dutch ministries and governmental bodies. Its assessment instrument was developed in direct response on non-European cloud providers, a concern crystallized by the Microsoft ICC incident, in which Microsoft's role in implementing Trump administration sanctions against the International Criminal Court was described in the document itself as "a serious wake-up call" [15].

The DICTU instrument is explicitly grounded in and cross-references both the EuroStack framework and the EC-CSF. The document acknowledges that the five-dimensional structure draws primarily on EuroStack, while the SEAL and Sovereignty Score concepts from the EC-CSF confirm and validate its independent analytical choices. This institutional intertextuality is analytically significant: DICTU functions as a national-level operationalization of European-level frameworks, adapted to the specific legal and organizational context of Dutch central government procurement.

Conceptual structure

The DICTU framework organizes its sovereignty assessment across five dimensions (table 2.6), each structured around a central question, a set of criteria, and five levels of sovereignty for each criterion. Crucially, DICTU departs from EuroStack's original five-dimension model by replacing the Economic Sovereignty dimension with a "Mens" (Human/Personnel) dimension, reflecting the specific operational reality of governmental IT service provision, where access privileges, personnel, and workforce control are of primary concern.

Table 2.6: Overview of the five sovereignty dimensions within the DICTU framework and their respective key assessment criteria

Criterion / Dimension	**Description / Key Criteria**
Juridisch (Legal)	Ultimate legal accountability: establishment (vestiging) of the ultimate parent entity in EU/EEA/EFTA; ownership and control (zeggenschap) free from non-EU influence; applicable law insulated from extraterritorial reach (CLOUD Act, FISA 702)
Data & AI	Data residency (all data, backups, metadata, and logs stored and processed exclusively in EU territory); technical access security (cryptographic isolation, HYOK + confidential computing); legal access security (mandatory legal challenge of non-EU data requests, MLAT referral).
Technologie (Technology)	Interoperability and portability (all services built on most widely used open standards); open-source (core components built on OSI-approved licenses); architectural transparency (source code auditable via escrow or public access); operational reversibility (third-party redeployment capability, exit plan).
Operationeel (Operational)	Infrastructure and control plane (physical infrastructure and management plane exclusively in EU territory); personnel (100% of privileged access personnel resident in EU, employment by EU entity, operating from within EU territory); supply chain resilience (auditable non-EU hardware/software dependency mitigation strategy).
Mens (Human/Personnel)	Who has access to systems and data; visibility and control over individuals working on digital infrastructure; personnel screening, nationality considerations, and access management under EU legal frameworks.

Table 2.7: DICTU Toetsingsinstrument Sovereignty assessment criteria and maturity levels — Part I (Legal, Data & AI, Technology)

Dimension	**Code**	**Criterion**	**Guiding question**	**Level 1 (lowest)**	**Level 5 (highest)**
Juridisch (Legal)	1.1	Establishment (Vestiging)	Is the Ultimate Parent Entity legally incorporated and headquartered within the EU, EEA, or EFTA?	No EU establishment	Fully EU-established with Dutch local presence
	1.2	Ownership and control (Zeggenschap)	Is the provider free from direct or indirect control by non-EU entities, both formally and substantively?	Full non-EU control	All voting rights and influence are exclusively European
	1.3	Applicable law (Toepasselijk recht)	Is the provider legally insulated from extraterritorial non-EU legislation such as the US CLOUD Act or FISA 702?	Non-EU jurisdiction	EU jurisdiction (binary; no intermediate levels)

Continued on next page

Continued from previous page

Dimension	Code	Criterion	Guiding question	Level 1 (lowest)	Level 5 (highest)
Data & AI	2.1	Data residency	Are all customer data assets stored and processed exclusively within EU/EEA/EFTA territory?	No EU storage guarantee	All data exclusively EU; national or on-premise option available
	2.2	Technical access security	Are technical measures in place to ensure unencrypted customer data is inaccessible to anyone, including the provider itself?	No technical protection; provider has full access	HYOK plus confidential computing with cryptographic isolation
	2.3	Legal access protection	Is the provider contractually obligated to challenge non-EU data access requests and notify customers?	No obligation; compliance without resistance	Mandatory legal challenge with MLAT referral
Technologie (Technology)	3.1	Interoperability and portability	Is the service built on open standards enabling provider switching or multi-cloud operation?	Full lock-in; proprietary only	All services based on widely accepted open standards
	3.2	Open source	Does the service use open-source software for its core components?	Fully proprietary	Fully open-source with active upstream contribution
	3.3	Software transparency	Is the source code of core technology publicly available for inspection or via a recognized European escrow?	Black box; no transparency	Independently audited supply chain; reproducible software
	3.4	Operational reversibility	Can a qualified third party operate the service using available documentation and design?	Non-transferable	Fully transferable and independently tested

Table 2.8: Sovereignty assessment criteria and maturity levels — Part II (Operational, Human/Personnel)

Dimension	Code	Criterion	Guiding question	Level 1 (lowest)	Level 5 (highest)
Operationeel (Operational)	4.1	EU infrastructure and control plane	Is the complete infrastructure and management control plane located and operated within EU/EEA/EFTA?	Infrastructure and control plane fully outside EU	Fully EU-based; periodically audited by EU partner
	4.2	Exclusively European personnel	Are all privileged access activities performed exclusively by EU personnel working from Europe?	Non-EU privileged access	100% EU personnel; periodically audited
	4.3	Supply chain resilience	Does the provider have a demonstrable and audited long-term supply chain strategy?	No strategy	Audited strategy with guaranteed continuity
Mens (Human/Personnel)	5.1	Training, certification and screening	Is personnel with access to critical systems screened, certified and independently verifiable?	No screening or training	Periodically audited by a recognized European body
	5.2	Capacity and skills	Does the provider have sufficient internal EU capacity to autonomously operate a sovereign cloud?	Fully dependent on non-EU capacity	Leading EU sovereign cloud innovator with local (NL) capacity

Assessment methodology

The core methodological innovation of the DICTU instrument is its five-level ordinal scale, applied consistently across all thirteen criteria. Each level describes a specific, observable state of the provider, moving from a condition of complete non-EU dependency (Level 1) to full sovereignty that is independently verified by a recognized European partner (Level 5). This progression is analytically important: the highest assurance level is not merely a matter of technical capability but requires third-party audit confirmation. This "audit-contingent Level 5" design appears in three criteria: EU Infrastructure and Control Plane (4.1), Exclusively European Personnel (4.2), and Training, Certification, and Screening (5.1), and forms the instrument's most direct engagement with independent assurance practice.

Two criteria warrant specific methodological attention because they illustrate the instrument's analytical depth. Criterion 2.2 (Technological Access Security) provides a five-level taxonomy of encryption and key management models ranging from Platform Managed Keys (Level 2 - provider can technically access data if legally compelled) to Hold Your Own Key (HYOK) plus Confidential Computing (Level 5 - provider is cryptographically incapable of accessing unencrypted data even during processing). This taxonomy directly operationalized the CLOUD Act exposure assessment: only levels 4 and 5 provide genuine protection against a legally compelled decryption order, because at these levels the provider lacks the technical means to comply. Criterion 1.3 (Applicable law) is treated as strictly binary: either EU jurisdiction applies (Level 5) or it does not (Level 1), with no intermediate states, reflecting DICTU's view that there is no meaningful partial state between EU and non-EU legal authority over cloud service delivery.

Regulatory and policy grounding

A distinctive feature of the DICTU instrument is its explicit regulatory embedding. The instrument situates each dimension within the applicable EU and Dutch legal framework: GDPR and the CLOUD Act conflict for the Legal dimension; NIS2, DORA, and the EU AI Act for the Data & AI and Operational dimension; and the Baseline Informatiebeveiliging Overheid (BIO) and the Implementatiekader Risicoafweging Cloud Gebruik (2022) at the national policy level. This multi-level regulatory grounding makes the instrument directly actionable for Dutch public sector organizations operating within the Rijksbrede cloud governance architecture. The document also explicitly identifies the political impasse surrounding the EUCS High+ level as a governance gap, positioning both EuroStack and the EC-CSF as the institutional responses that motivated and contextualized the DICTU instrument itself.

2.4.5 The SUSE Cloud Sovereignty Framework Self-Assessment

Origin and institutional context

The SUSE Cloud Sovereignty Framework Self-Assessment (hereafter SUSE CFSA), launched on 29 January 2026, forms a qualitatively distinct type of instrument compared to the four frameworks analyzed in preceding sections. Where the EC-CSF, EuroStack, DICTU, and the Buy European framework are primarily normative or policy instruments, produced by public institutions, think tanks, and civil society actors to define what sovereignty means and how it should be governed, the SUSE CFSA is a vendor-developed, technology-operationalized self-assessment tool designed to help organizations measure their current sovereignty posture against the EC-CSF and identify remediation pathways [52].

SUSE (Software and Linux AG) is a German-headquartered, globally operating enterprise open-source software company, best known for SUSE Linux Enterprise Server and the Rancher Prime Kubernetes management platform. Its positioning as a European-headquartered open-source vendor gives it a particular strategic stake in the digital sovereignty agenda: it presents itself as the structural alternative to non-EU hyperscalers. The SUSE CFSA is therefore simultaneously an assessment instrument and a market positioning tool. It creates measurable evidence of a "sovereignty gap" and channels remediation toward SUSE products and its European partner ecosystem [52]. This dual function is important to keep in mind when evaluating its outputs: the scoring is derived from the EC-CSF, but the remediation road map reflects a vendor perspective.

Conceptual structure and assessment architecture

Unlike the other frameworks in this analysis, the SUSE CFSA introduces no independent sovereignty taxonomy. It adopts the EC-CSF's architecture wholesale, the eight Sovereignty Objectives (SOV-1 through SOV-8) and the five-level SEAL scale (0-4), and operationalizes them into a 32-question structured questionnaire (SUSE, 2026). The questions are distributed across the eight SOV domains, weighted identically to the EC-CSF, and assessed by the organization itself on the SEAL scale. Table 2.9 presents the complete SOV structure as implemented in the SUSE CFSA, including an IT auditor perspective on the verifiability of each objective.

A methodologically important feature of the SUSE CFSA - one that distinguishes it from a simple generic self-rating exercise, is that each of the 32 questions carries its own contextualized description of what every SEAL level (0-4) means for that specific question [52]. Rather than asking an organization to self-assign a generic SEAL level and interpret its meaning independently, the tool provides tailored level descriptions anchored in the technical and operational reality of each assessed dimension. This contextualization significantly reduces response ambiguity and brings the scoring closer to the precision of the DICTU instrument's five-level per criterion scale. To illustrate: for a question under SOV-4 Operational Sovereignty concerning the location and governance of the management control plane, the question-specific SEAL descriptions translate the abstract scale into concrete operational states, SEAL-0 indicating that the control plane is entirely outside the EU and managed by a non-EU party; SEAL-2 indicating that infrastructure resides in the EU but the control plane is partially externally managed; and SEAL-4 indicating that the control plane is fully EU-located, autonomously operated and free from non-EU dependencies [52].

Similarly, SUSE has articulated an Application Trust Hierarchy that maps these SEAL levels to specific organizational contexts: SEAL-1 for regular applications where EU legal compliance is the primary concern; SEAL-2 for business-critical applications requiring verifiable control over data location; SEAL-3 for

mission-critical systems such as banking or healthcare platforms; and SEAL-4 for national critical infrastructure such as payments systems or electricity grid management [52]. This question-level specificity partially mitigates the self-assessment limitation, because the respondent is guided toward the technically accurate answer rather than an impressionistic one. The fundamental verification gap nonetheless remains: the accuracy of the selected level is not independently confirmed, and no documentary evidence is required to substantiate it.

Table 2.9: Sovereignty objectives and scope of assessment

Code	**Sovereignty Objective**	**Weight (%)**	**Scope of Assessment**
SOV-1	Strategic Sovereignty	15	EU ownership, governance stability, resistance to non-EU control or service suspension.
SOV-2	Legal Jurisdictional	10	Governing legal system, CLOUD Act/FISA 702 exposure, contractual enforceability under EU law.
SOV-3	Data & AI Sovereignty	10	Customer-exclusive cryptographic control, data residency, AI model governance under EU control.
SOV-4	Operational Sovereignty	15	Control plane location, EU-resident privileged personnel, migration and exit feasibility.
SOV-5	Supply Chain Sovereignty	20	Hardware origin, firmware provenance, software supply-chain transparency, non-EU vendor dependency.
SOV-6	Technology Sovereignty	15	Open APIs, open-source adoption, source code transparency, architectural auditability.
SOV-7	Security & Compliance	10	EU certifications (ISO 27001, EUCS), NIS2/DORA/GDPR compliance, EU-based SOC operations.
SOV-8	Environmental Sovereignty	5	Energy efficiency (PUE), renewable energy sourcing, hardware lifecycle and circular economy.

The policy-to-technology gap

The SUSE CSFA's most substantive analytical contribution is its explicit identification of the "policy-to-technology gap": the structural disconnect between what sovereignty regulations require and what the actual technical stack delivers [52]. This insight is directly relevant for IT audit practice. It establishes that a sovereignty assessment cannot be conducted solely to the contractual or policy level; reviewing a DPA and data residency commitments is necessary but insufficient. An effective sovereignty assessment must also engage with the technical infrastructure layer: the operating system, virtualization platform, key management architecture, container orchestration, and firmware provenance [52]. This is precisely the layer on which self-assessment tools systematically underperform, because organizations frequently lack visibility into what they do not know about their own stack.

2.4.6 Comparative evaluation

Rationale for a weighted criteria evaluation

The five frameworks analyzed in this chapter were each developed for a specific institutional purpose: the EC-CSF as a procurement evaluation standard, EuroStack as an industrial strategy proposition, DICTU as a national operational assessment instrument, the Buy European Framework as a legal procurement mechanism, and the SUSE CFSA as a commercial self-assessment tool. Presenting them as comparable in a flat confrontation matrix risks obscuring the differences that matter most for this research: not what each framework covers in the abstract, but how well each framework serves as a building block for a purpose-built IT audit instrument for digital sovereignty assessment in cloud-dependent organizations.

This section applies a structured comparative evaluation in which each framework is assessed against six criteria derived from the functional requirements of IT audit practice [32]. The criteria and their

relative weights reflect the priorities identified in the preceding literature review: evidence standard and scoring granularity carry the highest weights because these are the dimensions most directly relevant to producing a defensible audit conclusion, while proportionality carries the lowest weight because it functions as a secondary calibration property rather than a foundational requirement.

Six criteria are applied, each grounded in professional audit practice:

1. **Dimensional coverage:** does the framework address all material sovereignty dimensions relevant to cloud audit practice?
2. **Evidence standard:** Does the framework specify or require the collection of independently verifiable documentary or technical evidence, rather than self-assertion?
3. **Scoring granularity:** Does the framework provide a sufficiently granular measurement scale to support gap analysis, target-setting, and before-and-after comparison?
4. **Regulatory grounding:** Is the framework explicitly anchored in applicable EU and national legal and regulatory obligations?
5. **Auditability:** Are the framework's criteria sufficiently precise and observable to support reproducible, independent assessment by different auditors?
6. **Proportionality:** Does the framework support risk-based calibration of assessment depth to the sensitivity and criticality of the assessed service?

The evaluation was conducted by the author through systematic analysis of each framework's primary source documentation. For each criterion, a score of 1 (not addressed), 2 (partially addressed), 3 (substantially addressed), or 4 (fully addressed) was assigned based on whether the framework explicitly defines, requires, or operationalizes the relevant property. A score of 4 requires that the criterion is not only present in principle but is operationalized in a form that an IT auditor could directly apply without additional interpretation. Scores were assigned conservatively: where a framework addresses a criterion implicitly or only at a high level of abstraction, a score of 2 rather than 3 was assigned. The scoring rationale for each framework and criterion is documented in the narrative that follows the table.

Table 2.10 presents the weighted evaluation across all five frameworks.

Table 2.10: Comparative audit scorecard across European sovereignty frameworks

Audit criterion	**Weight (%)**	**EC-CSF**	**EuroStack**	**DICTU**	**Buy European**	**SUSE CFSA**
Score = 1 (Poor) to 4 (Excellent)		Score / Weight	Score / Weight	Score / Weight	Score / Weight	Score / Weight
(1) Dimensional coverage: does the framework address all material sovereignty dimensions relevant to cloud audit practice?	20	4 / 0.80	3 / 0.60	4 / 0.80	3 / 0.60	4 / 0.80
(2) Evidence standard: does the framework specify or require independently verifiable documentary or technical evidence, rather than self-assertion?	25	4 / 1.00	1 / 0.25	4 / 1.00	4 / 1.00	1 / 0.25
(3) Scoring granularity: does the framework provide a sufficiently granular measurement scale to support comparison?	20	3 / 0.60	1 / 0.20	4 / 0.80	2 / 0.40	3 / 0.60
(4) Regulatory grounding: is the framework anchored in applicable EU and national legal obligations?	15	4 / 0.60	2 / 0.30	4 / 0.60	3 / 0.45	2 / 0.30
(5) Auditability: are criteria sufficiently precise to support reproducible independent assessment?	15	3 / 0.45	1 / 0.15	3 / 0.45	4 / 0.60	2 / 0.30
(6) Proportionality: does the framework support risk-based calibration of assessment depth?	5	3 / 0.15	1 / 0.05	3 / 0.15	4 / 0.20	2 / 0.10
Total weighted score	100	3.60 / 4.00	1.55 / 4.00	3.80 / 4.00	3.25 / 4.00	2.35 / 4.00

Interpretation of the evaluation results

The evaluation results reveal a clear two-tier structure. The DICTU instrument (3.80/4.00) and the EC-CSF (3.60/4.00) score substantially above the remaining three frameworks on the six audit-relevant criteria. Both achieve the maximum score on the evidence standard, reflecting their requirement for independently verifiable documentary or technical evidence rather than self-assertion. DICTU is unique in achieving the maximum score on scoring granularity through its five-level per-criterion maturity scale, while the EC-CSF achieves the highest dimensional coverage through its eight Sovereignty Objectives. Together, these two frameworks come closest to meeting the requirements of an IT audit context, though neither was designed specifically for that purpose: DICTU was developed for Dutch government procurement assessment, and the EC-CSF is for supranational policy benchmarking.

The Buy European framework (3.25/4.00) scores in the secondary range. Its highest scores are on auditability and proportionality, reflecting the precision and risk-calibration of its three-level procurement hierarchy. Its low score on scoring granularity reflects its binary procurement orientation: it was designed to filter providers, not to produce a diagnostic profile of an organization's exit capability.

The SUSE CFSA (2.35/4.00) scores below the audit-utility threshold on the most critical criterion, evidence standard, because its outputs are entirely self-reported and unverified [52]. However, its methodological contribution lies in demonstrating how EC-CSF objectives can be operationalized into a structured question-and-response format, even if the resulting outputs lack the independence required for assurance purposes.

EuroStack (1.55/4.00) scores lowest because it is a strategic policy document, not an assessment instrument. It provides no scoring mechanism, no evidence requirements, and no per-criterion criteria. Its value lies in the political-economic context it supplies for understanding why cloud sovereignty matters at the European level.

Gap analysis: What the existing frameworks do not provide

Even the two highest-scoring frameworks leave four material gaps that the DSAF must fill. Together, these gaps constitute the academic justification for the artifact's existence as a DSR contribution [31].

Gap 1: No integrated audit protocol with defined evidence requirements Neither the EC-CSF nor DICTU specifies the audit procedures an IT auditor must execute to verify a sovereignty claim, nor does either document define the minimum evidence set required to substantiate a given score. The EC-CSF specifies what to assess (contributing factors per SOV objective), but not how it leaves evidence collection entirely to the contracting authority's discretion [16]. DICTU defines sovereignty levels with precision but does not specify what documentary or technical evidence must be obtained to place a provider at Level 3 versus Level 4 [15]. This gap is the DSAF's primary design objective: every audit question in the DSAF protocol is accompanied by an explicit audit procedure and a defined minimum evidence set.

Gap 2: No professional independence standard None of the five frameworks is designed to be applied by an independent IT auditor operating under professional standards such as ISAE 3402, ISAE 3000, or equivalent national audit standards. The EC-CSF and Buy European Framework assume a contracting authority assessing a tenderer in a procurement context, not an audit context. DICTU is an internal assessment instrument for a governmental agency that evaluates its own cloud suppliers. The SUSE CFSA is explicitly a self-assessment.

Gap 3: No cross-sector proportionality mechanism Both the EC-CSF and DICTU implicitly assume a public sector context in which the highest sovereignty levels are the default aspiration. The EC-CSF permits contracting authorities to set minimum SEAL thresholds per procurement procedure but does not provide a structured methodology for calibrating those thresholds to the specific data sensitivity, operational criticality, and regulatory obligations of the assessed service.

Gap 4: No integrated reporting standard None of the existing frameworks specifies how sovereignty assessment results should be communicated to management, boards, regulators, or procurement authorities in a format consistent with professional audit reporting standards. The SUSE CFSA produces a PDF report, but it is vendor-branded and covers only self-assessed findings.

These four gaps collectively define the problem space that the framework is designed to occupy. They confirm that the existing landscape, while rich in normative frameworks and policy instruments, lacks

a single integrated instrument that IT auditors can apply systematically, independently, and verifiably to produce credible sovereignty assurance outputs. The following chapter presents the framework's methodology as the direct response to this documented gap.

The five frameworks examined in this chapter, the EC-CSF, EuroStack, DICTU, the Buy European Procurement Framework, and the SUSE CFSA, do not represent competing visions of digital sovereignty. They are complementary instruments operating at different levels of abstraction, institutional authority, and practical application. Together, they reveal a converging cross-institutional consensus around the core dimensions of cloud sovereignty, while exhibiting meaningful divergences in scope, scoring methodology, assessment direction, and the role accorded to independent verification.

Three structural insights emerge from this synthesis. First, jurisdictional sovereignty is universally treated as the foundational prerequisite, the condition without which all other sovereignty controls are legally unenforceable [26]. Whether operationalized as a pass/fail gate [15, 23]; a minimum SEAL threshold [16], or a strategic prerequisite, every framework identifies the Ultimate Parent Entity's EU domicile and freedom from non-EU legal coercion as the sine qua non of sovereign cloud procurement. Second, all frameworks share a documented understanding that data sovereignty, operational control, and technological openness are independent. The "sovereign prison" concept [23] identifies the failure modes that arise when one dimension is addressed while others are neglected. Third, all frameworks consistently treat technological controls, cryptographic isolation, open standards, and reproducible builds as superior to contractual promises that can be overridden by foreign law [15].

The most significant divergence concerns the role of independent verification. The EC-CSF, DICTU, and Buy European framework all require or assume third-party evidence [16, 15, 23]. The SUSE CFSA is explicitly self-reported [52]. EuroStack does not prescribe an evidence standard [24]. This divergence defines the gap that a purpose-built IT audit framework must fill: it must combine the dimensional comprehensiveness of the EC-CSF and DICTU with the auditability requirements of professional assurance practice, producing outputs credible to regulators, procurement authorities, and boards. Table 2.11 provides a systematic confrontation of all five frameworks across fifteen analytical dimensions.

Table 2.11: Comparative analytical dimensions across European sovereignty frameworks

Analytical dimension	**EC-CSF**	**EuroStack**	**DICTU**	**Buy European**	**SUSE CFSA**
Institutional type	EU Commission (supranational regulatory)	Academic / civil society (industrial strategy)	Dutch government agency (national operational)	Industry initiative (procurement-law proposal)	Commercial vendor (technology self-assessment tool)
Primary purpose	Procurement evaluation standard	Industrial policy and investment strategy	Cloud service assessment instrument	Binding "Buy European" regulatory proposal	Organizational self-assessment and gap analysis
Assessment direction	Provider assessment	Ecosystem / market level	Provider assessment	Provider assessment	Organization assesses itself
Legal & jurisdictional dimension	Full: SOV-1 and SOV-2; weighted 25%	Full: Dimension I (mandatory prerequisite)	Full: three criteria; hard pass/fail gate	Full: Dimension I pass/fail; GPA legal foundation	Partial: SOV-1 and SOV-2; self-reported only
Data sovereignty dimension	Full: SOV-3; residency, AI governance, encryption	Full: Dimension IV; EU residency plus HYOK	Full: three criteria; five-level encryption taxonomy	Full: Dimension IV; confidential computing required	Partial: SOV-3; self-assessed with limited depth
Operational sovereignty dimension	Full: SOV-4; control plane and personnel	Full: Dimension III; EU infrastructure and staff	Three criteria; audit-contingent Level 5	Full: Dimension III; 100% EU personnel required	Partial: SOV-4; self-reported; control plane not verified
Technology sovereignty dimension	Full: SOV-6; open standards and auditability	Full: Dimension II; open source and reversibility	Full: four criteria; transparency and reversibility	Full: Dimension II; open source and escrow	Partial: SOV-6; self-assessed; no code review
Supply chain dimension	Full: SOV-5; highest weight (20%)	Full: Dimension III; hardware and software chain	Partial: single criterion; strategy focus only	Full: Dimension III; SBOM and audit rights	Partial: SOV-5; highest weight but self-reported
Human & personnel dimension	Partial: included operationally in SOV-4	Partial: EU personnel requirement in Dimension III	Full: dedicated dimension with two criteria	Partial: EU personnel included in Dimension III	Absent: not explicitly assessed
Economic / ecosystem dimension	Partial: SOV-1 strategic sovereignty	Full: Dimension V; R&D and SME ecosystem	Absent: replaced by human dimension	Full: Dimension V; R&D and open-source contribution	Partial: SOV-1; self-assessed only
Security & compliance dimension	Full: SOV-7; certifications, SOC, and audit rights	Partial: implicit via open source and governance	Absent: not a standalone dimension	Partial: implicit via EU jurisdiction and audit rights	Partial: SOV-7; certification check; self-assessed
Environmental sustainability	Full: SOV-8; 5% weight	Partial: sustainability principle; not scored	Absent: out of scope	Absent: out of scope	Partial: SOV-8; 5% weight; self-assessed
Scoring / measurement model	SEAL 0–4 with weighted sovereignty score	Qualitative framework; no formal scoring	Five-level ordinal scale; no aggregate score	Pass/fail gate plus proportionality scoring	SEAL 0–4 with weighted scores (mirrors EC-CSF)
Independence verifiability	Full: contracting authority assesses on evidence	N/A: strategy document; no evidence requirement	Full: Level 5 requires independent EU audit	Full: corporate registry and legal attestation	Absent: fully self-reported; no independent check
IT audit utility	Full: primary reference standard; evidence-based	Partial: strategic context only; no audit protocol	Full: best operational instrument for gap analysis	Partial: procurement focus and legal framework	Partial: rapid screening only; not suitable for assurance

2.5 IT Audit and assurance frameworks

As organizations increasingly depend on complex information systems, interconnected digital services, and cloud-based infrastructures, structured governance and assurance frameworks provide a systematic foundation for evaluating whether technological controls effectively mitigate operational, security, and compliance risks [14]. These frameworks enable organizations to implement consistent governance structures and standardized control mechanisms that ensure information systems operate in accordance with organizational objectives and regulatory expectations [32]. In addition, they provide auditors and stakeholders with a structured methodology for assessing whether information systems maintain data integrity, confidentiality, and availability, which are fundamental pillars of secure digital operations [30].

Within this context, IT auditing functions as an essential mechanism for evaluating whether organizations implement appropriate governance processes, risk management practices, and internal controls that collectively safeguard digital assets and ensure the reliability of information systems [32]. By applying standardized audit criteria and assurance methodologies, IT audit frameworks contribute to organizational transparency and accountability while strengthening stakeholder confidence in the security and resilience of digital infrastructures [14].

In practice, IT audits involve the systematic examination of governance structures, technical controls, operational procedures, and risk management practices that collectively shape how organizations manage their information systems and digital resources [32]. Assurance frameworks play a central role in structuring these assessments by defining clear control objectives, audit criteria, and evaluation methodologies that guide auditors during the audit process [2]. By providing standardized benchmarks and structured evaluation models, these frameworks allow auditors to determine whether implemented controls effectively mitigate identified risks and comply with relevant regulations and best practices [32]. Furthermore, IT assurance frameworks support organizations in demonstrating accountability toward regulators, business partners, and customers by providing verifiable evidence that appropriate security and governance mechanisms are in place [14]. As digital ecosystems continue to expand in scale and complexity, the importance of structured assurance frameworks has increased significantly, particularly in sectors that rely heavily on cloud computing, outsourcing, and interconnected digital services [7].

2.5.1 COBIT

One of the most influential frameworks in the field of IT governance and assurance is COBIT 2019, developed by the Information Systems Audit and Control Association (ISACA) as a comprehensive governance and management framework for enterprise information and technology [32]. COBIT 2019 provides organizations with an integrated model for aligning IT governance with organizational strategy, allowing management to ensure that technological resources support business objectives while simultaneously managing associated risks [32]. The framework organizes IT governance into a set of governance and management objectives that collectively address strategic alignment, value delivery, risk management, resource optimization, and performance monitoring within information systems [32]

By structuring governance practices around these objectives, COBIT facilitates the systematic evaluation and improvement of IT governance processes and organizational maturity levels [14]. A significant strength of COBIT lies in its holistic approach to governance, which integrates strategic decision-making with operational management practices and enables organizations to assess the effectiveness of their IT governance structures and identify opportunities for continuous improvements [14]. Despite these strengths, the framework focuses primarily on governance effectiveness within the organizational boundary and focuses on aligning IT with business objectives rather than explicitly addressing broader geopolitical or jurisdictional dimensions related to digital sovereignty [15].

2.5.2 ISO/IEC 27001

Another widely adopted framework for information security assurance is the international standard ISO/IEC 27001, which defines the requirements for establishing, implementing, maintaining, and continuously improving an information security management system (ISMS) within an organization [39]. The core principle of ISO/IEC 27001 is the implementation of a risk-based approach to information security management, whereby organizations systematically identify, assess, and mitigate information security risks through the implementation of appropriate security controls [39]. These controls are further elaborated in ISO/IEC 27002, which provides detailed guidance on the implementation of security

practices across domains such as asset management, access control, cryptography, incident management, and supplier relationships [39].

Through certification and periodic external audits, ISO/IEC 27001 allows organizations to demonstrate compliance with internationally recognized information security standards and thus strengthen trust among customers, regulators, and business partners [58]. The global recognition and structured methodology of ISO/IEC 27001 represent significant advantages, as they allow organizations across different industries and jurisdictions to adopt a consistent approach to managing information security risks [39]. Nevertheless, while the standard provides extensive guidance on security governance and risk management, its main focus remains on organizational security practices rather than on broader issues related to jurisdictional authority, data sovereignty, or geopolitical control over digital infrastructure [15].

2.5.3 Service Organizational Control (SOC)

In addition to governance and security standards, third-party assurance frameworks play an important role in evaluating the internal control environments of service providers, particularly within cloud computing and outsourced IT services [2]. One of the most widely used assurance mechanisms in this context is the Service Organization Control (SOC) reporting framework developed by the American Institute of Certified Public Accountants (AICPA) [2]. SOC 2 reports specifically evaluate whether service organizations implement effective controls based on the Trust Service Criteria, which encompass the principles of security, availability, processing integrity, confidentiality, and privacy [2].

Independent auditors assess whether service provider systems and processes are designed and operated effectively to meet these criteria and protect client data [2]. SOC 2 reports, therefore, provide organizations that rely on external service providers with assurance that appropriate internal controls exist to manage operational security risks associated with outsourced IT services [2]. This form of assurance is particularly important in cloud environments, where organizations frequently depend on third-party providers for critical infrastructure and data processing services [7]. However, despite the value of SOC 2 reporting in evaluating service provider controls and operational security practices, the framework focuses primarily on internal control effectiveness rather than addressing broader concerns related to national jurisdiction, digital autonomy, or sovereignty over data and infrastructure [15].

2.5.4 Cloud Control Matrix (CCM)

The rapid expansion of cloud computing has also led to the development of specialized assurance frameworks designed specifically to address cloud-related security challenges, among which the Cloud Security Alliance's Cloud Controls Matrix (CCM) represents one of the most prominent examples [7]. The Cloud Control Matrix provides a comprehensive set of security controls that are tailored to the unique risks associated with cloud computing environments, enabling organizations to systematically evaluate the security capabilities of cloud service providers [7]. These controls are organized into multiple domains, including identity and access management, infrastructure and virtualization security, data protection, application security, and compliance management [7].

By aligning cloud-specific security controls with industry standards and regulatory requirements, the CCM facilitates standardized security assessments across cloud providers and supports transparency in cloud governance practices [7]. In addition, the framework enables organizations to compare provider capabilities and evaluate whether cloud services meet required security and compliance standards before adoption [7].

Despite its comprehensive approach to cloud security controls, the Cloud Controls Matrix largely concentrates on technical and organizational safeguards, while broader questions concerning jurisdictional control, geopolitical dependency, and digital sovereignty remain outside its central scope [15].

Taken together, COBIT 2019, ISO/IEC 27001, SOC 2, and the CSA Cloud Controls Matrix illustrate the diverse set of instruments currently used to govern, audit, and evaluate digital infrastructures [32, 39, 2, 7]. These frameworks collectively enable organizations to structure internal controls, manage information security risks, and demonstrate accountability toward regulators, partners, and stakeholders in increasingly complex digital environments [14]. They also provide auditors with standardized methodologies for evaluating governance structures, technical safeguards, and operational processes in a consistent and transparent manner [2].

However, while these frameworks effectively address issues related to cybersecurity, risk management, and governance within organizational contexts, most traditional IT audit frameworks were not originally designed to address the strategic and geopolitical dimensions of digital sovereignty [15]. As concerns about jurisdictional control, cross-border data flows, technological dependency, and geopolitical influence over digital infrastructures continue to grow, the limitations of traditional IT audit frameworks in addressing sovereignty-related risks become increasingly evident [15]. This observation highlights the need for further academic research into how existing IT governance and assurance frameworks can be extended or complemented to incorporate considerations of digital sovereignty within modern digital ecosystems [15].

Chapter 3

Method

This chapter outlines the methodological architecture underpinning the study, articulating the research approach, research design, data collection procedures, and analytical strategies that govern the empirical inquiry. The research is situated within DSR paradigm and employs a mixed-method validation strategy grounded in the Technology Acceptance Model (TAM). The chapter proceeds from the overarching research strategy through sampling rationale, instrumentation, data collection procedures, and analytical methods, concluding with a reflection on the research quality criteria that underpin the credibility of the findings.

3.1 Research paradigm and overall strategy

The study adopts DSR as its overarching methodological paradigm. DSR, as conceptualized by Hevner, et al. [31], is premised on the iterative construction and rigorous evaluation of the artifacts, purposive objects designed to solve identified problems in their organizational or technical context. Unlike purely explanatory research traditions that seek to describe or predict phenomena, DSR is fundamentally prescriptive: its central output is a designed artifact whose validity is assessed against the criterion of utility in the problem domain.

The present study operationalizes DSR across three interconnected cycles, each corresponding to a distinct phase of knowledge production. The first cycle, the relevance cycle, grounds the artifact design in the practitioner problem space through an extensive review of extant frameworks, including the DICTU Toetsingsinstrument Soevereiniteit Clouddiensten (2026) [15], the EC Cloud Sovereignty Framework (2025) [16], and the EuroStack Buy European Framework (2025) [23]. The second cycle, the design cycle, covers the construction of the Digital Sovereignty Audit Framework, specifically the Operational Reversibility control instrument, drawing on the conceptual vocabulary and normative criteria established in the prior literature. The third cycle, the rigor cycle, forms the empirical evaluation of the artifact through expert elicitation, and it is this cycle that the present chapter addresses.

This methodological architecture positions the artifact not as a static deliverable but as an evolving object subject to iterative refinement in response to expert critique, a theoretical commitment that is operationalized through the validation study described below.

3.2 The Technology Acceptance Model (TAM)

The validation of the artifact is structured around the Technology Acceptance Model (TAM), originally proposed by Davis (1989) [13] and subsequently extended by Venkatesh and Davis (2000) [56] in the TAM2 formulation. TAM posits that user adoption of information technology or a tool is determined by two primary cognitive antecedents: Perceived Usefulness (PU), defined as the degree to which a user believes that employing the tool will enhance their job performance, and Perceived Ease of Use (PEOU), defined as the degree to which use of the tool is expected to be free of cognitive effort [13].

The present study extends the classical TAM specification in two respects. First, Social Norm (SN) is incorporated as a third antecedent construct, drawing on the TAM2 extension [56] and the Theory of

Planned Behavior [1]. In the professional audit context, SN captures the degree to which respondents perceive that normative peer expectations, professional body endorsements, and organizational mandates would support or impede framework adoption. This extension is theoretically warranted given that IT audit practice is embedded in a densely institutional environment, one in which NOREA guidance, Big Four methodology standards, and regulatory expectations structurally condition practitioner behavior. Second, Behavioral Intention (BI) is measured as a dependent construct, capturing the prospective likelihood that respondents would deploy or recommend the framework in a real engagement [56].

The juxtaposition of these four constructs (PU, PEOU, SN, BI) provides a multidimensional lens through which the artifact's perceived utility, usability, social legitimacy, and adaptability can be assessed simultaneously and compared across respondents differing in seniority, sector exposure, and prior familiarity with sovereignty concepts.

3.3 Research design

The instruments were structured across nine sequential sections, progressing from respondent profiling and sovereignty context-setting through framework review, construct-specific open questioning, Likert statement elicitation, and format preference. This sequencing was deliberate: by situating the sovereignty primer and framework review before the TAM items, the instrument ensured that respondents evaluated the artifact from a position of informed engagement rather than abstract speculation, a procedural decision that enhances the ecological validity of the quantitative scores.

3.4 Sampling strategy

This research was conducted in collaboration with EY Netherlands, which served as the host organization for the study. Respondents were recruited through the professional network of the company supervisors at EY Netherlands, targeting IT audit and Technology Risk professionals with direct experience in cloud governance and IT assurance engagements.

The study employs purposive sampling, targeting IT audit and Technology Risk professionals with demonstrable experience in cloud-dependent client environments. The selection criteria required respondents to hold a current or recent role in IT audit, Technology Risk advisory, or assurance practice and to have engaged with cloud-dependent client environments in a professional capacity. Familiarity with the sovereignty framework was explicitly not required, as the instrument incorporated a standardized sovereignty primer to establish a common interpretive baseline.

Seven respondents were recruited, spanning four seniority levels: Staff/Assistant (n=2), Senior (n=2), Manager (n=2), Senior Manager (n=1). All respondents were affiliated with Big Four professional services firms operating in the Netherlands. The deliberate seniority spread serves an analytical purpose: it enables the examination of whether TAM construct scores, qualitative critique, and design improvement suggestions vary systematically across the practitioner career trajectory, a dimension of analysis that is particularly pertinent in a professional knowledge domain where methodological expectations differ sharply between junior auditors, who evaluate usability primarily, and senior practitioners, who inspect structural validity and theoretical consistency.

The sample size of seven exceeds the saturation threshold of four to six respondents identified by Guest, Bunce, and Johnson [29] for homogeneous expert populations. This criterion, empirically validated through interview data demonstrating thematic recurrence across the final interviews with minimal emergence of novel codes, provides the justificatory basis for the sample size decision.

Table 3.1: Overview of respondents (Round 1)

Respondent	Role	Experience
R01	Senior, Technology Risk Advisory	5 years (2 in TR)
R02	Manager, Cyber & Privacy	8 years
R03	Senior, Technology Risk Advisory	3 years
R04	Staff/Assistant, IT Audit	1.5 years
R05	Staff/Assistant, IT Audit & Advisory	2.5 years
R06	Manager, ITO, Assurance Advisory	8 years
R07	Senior Manager, IT Audit	9 years

Table 3.2: Overview of respondents (Round 2)

Respondent	Role	Experience
R08	Manager, ITO, Assurance Advisory (also participated in Round 1 as R06)	8 years
R09	Partner, Technology Risk Advisory & Cyber Security Advisory	12 years

3.5 Data collection instrument

Data were collected using two complementary instruments: a semi-structured interview protocol and a Likert-based questionnaire. Both are described below.

3.5.1 Interview protocol

Data were collected through semi-structured expert interviews, each lasting approximately sixty minutes and conducted in Dutch. The interview guide was structured across nine sections: (1) consent and respondent profiling; (2) sovereignty context-setting via a standardized primer; (3) framework content review across the five control areas; (4-7) construct-specific open questioning for PU, PEOU, SN, and BI; (8) format preference; (9) closing remarks. The complete first-round interview guide is provided in Appendix A.

The sovereignty primer, delivered verbally in sections two and three, was designed to bridge the knowledge gap between respondents unfamiliar with digital sovereignty theory and the conceptual vocabulary embedded in the framework. It framed operational reversibility as a fundamentally audit-analytical question: whether an organization could exit a cloud provider relationship without catastrophic loss of service continuity, data, or operational knowledge, thereby anchoring respondent evaluation in their existing professional schemas rather than requiring sovereignty expertise.

The framework document was distributed to respondents forty-eight hours before the interview, accompanied by a briefing note requesting familiarization reading without formal assessment. This pre-sessional preparation protocol was adopted to ensure that Likert ratings reflect considered professional judgment rather than first-impression reactions, a methodological precaution that improves the reliability of the quantitative data.

At the start of each interview, respondents were asked to confirm whether they had reviewed the framework document prior to the session. All seven Round 1 respondents confirmed they had read the framework before the interview, though the depth of review varied, some respondents indicated a thorough reading, while others described a familiarization reading as intended by the briefing note. This variation is consistent with the pre-sessional preparation protocol's stated objective of ensuring considered rather than first-impression responses and is noted as a contextual factor in interpreting the Likert scores.

3.5.2 Transcription procedure

Audio recordings of all seven interviews were transcribed using Faster Whisper, an open-source automatic speech recognition model based on the Whisper architecture [49], run entirely in a local computing environment without transmission of audio data to any external server or cloud-based service. The decision to process recordings locally rather than through a cloud-based transcription service was made on both ethical and substantive grounds. Ethically, all respondents had consented to audio recording solely for transcription of this research, with the assurance that recordings would remain accessible only to the researcher; uploading audio to an external provider would have exceeded the scope of that consent. Substantively, given that this study concerns the risks of cloud dependency and data sovereignty, transcribing sensitive professional interviews via a cloud provider would have been inconsistent with the data governance principles the research itself is designed to evaluate.

Each automatically generated transcript was reviewed and corrected manually against the original recording to address transcription errors, particularly in Dutch technical terminology and proper nouns. The correct Dutch-language transcripts were then imported into Atlas.ti for qualitative coding.

3.5.3 Likert instrument

The quantitative component comprised twelve five-point Likert items operationalizing the four TAM constructs. Four items operationalized PU (PU1-PU4), three operationalized PEOU (PEOU1-PEOU3), three operationalized SN (SN1-SN3), and two operationalized BI (BI1-BI2). Items were adapted from Davis [13] for PU and PEOU, and from Venkatesh and Davis [56] and Ajzen [1] for SN and BI, with domain-specific reformulation to ensure content validity within the IT audit context.

The five-point response scale ranged from 1 (strongly disagree) to 5 (strongly agree). To minimize agreement bias, items were introduced after the open qualitative questioning for each construct, ensuring that respondents had articulated their substantive views before committing to a numerical rating. Items were administered verbally with the response scale visible to the respondent throughout.

3.6 Data analysis

3.6.1 Quantitative analysis of Likert data

Likert scales were analyzed descriptively. For each item, the mean and standard deviation across all seven respondents were computed. Construct-level means were derived by averaging item scores within each construct. An overall TAM score was calculated as the mean across all twelve items. Interpretive thresholds follow established conventions in TAM-based technology evaluation research [13]: a construct mean of 4.0 or above is interpreted as a positive indicator of perceived acceptance; a mean between 3.0 and 3.9 signals moderate acceptance requiring qualitative investigation; and a mean below 3.0 forms a material concern warranting substantive redesign.

Standard deviation is treated as both a measure of dispersion and an analytically informative indicator of respondent divergence. Items exhibiting high standard deviation are flagged for cross-referencing with qualitative data to identify the structural or conceptual sources of intra-respondent disagreement. This approach compares quantitative variance with qualitative critique and operationalizes the sequential explanatory logic of the mixed-methods design.

3.6.2 Qualitative analysis: Thematic analysis

The qualitative interview data were interpreted using reflective thematic analysis as developed by Braun and Clarke [6]. This approach was selected in preference to more constrained qualitative methods, such as content analysis or grounded theory, because it permits the analyst to move between deductively motivated and inductively generated codes without requiring strict theoretical saturation of categorical exclusivity. In the present study, the deductive level of analysis is provided by the TAM construct architecture, which structures the initial coding schema; the inductive level is generated from the emergent substantive content of respondent critique.

The analysis proceeded through six phases: In familiarization, the researcher read all seven transcripts in full, noting preliminary observations and recording analytical memos. In the second phase, initial code generation, every meaningful unit of text was assigned a first-order code capturing its semantic content

in language proximate to the respondent's own formulation. Dutch-language transcripts were coded in the original language; all codes were assigned English-language Atlas.ti labels to enable cross-respondent comparison on integration with the quantitative construct schema. In the third phase, theme searching, first-order codes were grouped into candidate themes by identifying semantic proximity, recurrence across respondents, and relevance to the four TAM constructs. In the fourth phase, theme reviewing, candidate themes were interrogated against the full dataset to assess their evidential grounding: a theme was retained as analytically robust if supported by data from three or more respondents. In the fifth phase, theme definition, each retained theme was assigned a name and a concise analytical description articulating its content and boundaries. In the sixth phase write-up, themes were integrated with quantitative findings to produce the results and discussion presented in the subsequent chapters.

Atlas.ti (version 24) was used as the qualitative data management platform, facilitating systematic code application, code co-occurrence analysis, and the development of analytical memos. The two-layer coding architecture, deductive TAM construct codes applied in parallel with inductive first-order content codes, enabled the identification of both construct-level patterns and specific design improvement themes that go beyond the TAM framework.

3.7 Ethical consideration

All respondents provided informed consent before data collection. The briefing note distributed forty-eight hours before the interview specified the research purpose, the voluntary nature of participation, and the intended use of data solely for academic purposes. Participants consented explicitly to audio recording for transcription purposes, with the assurance that recordings would not be shared through the use of pseudonyms and would be accessible only to the researcher and thesis supervisor.

3.8 Second evaluation round

While the first evaluation round, described above, was designed to generate problems and design implications from a broad expert sample, the second evaluation round served a confirmatory rather than exploratory purpose. Its objective was to establish whether the design adjustments made to the artifact in direct response to the first-round findings, documented in full in chapter 4, were perceived by IT audit professionals as resolving the specific problems they targeted.

The second round was conducted against the Version 2 artifact and involved two respondents, coded R08 and R09. R08 is a manager who had also participated in the first evaluation round, which allowed a direct comparison between the Version 1 and Version 2 artifacts and provided insight into whether specific, previously voiced concerns had been addressed in the respondent's own assessment. R09 is a partner who had not been exposed to either version of the framework before this session, which allowed the second round to also capture an independent first impression of the Version 2 artifact rather than only a recollection-based comparison. The sample size of two is appropriate to the confirmatory aim of this round and should not be read against the saturation logic used to justify the first-round sample in section 3.6: rather than seeking thematic saturation across a heterogeneous population, the second round targeted specific, previously identified design adjustments with a small number of respondents capable of assessing them in depth.

The second round reused the same data collection logic as the first. The twelve-item TAM questionnaire described in section 3.5.3 was administered unchanged, allowing scores to be compared directly against the first-round results on the same items and the same five-point scale. The questionnaire was again followed by a semi-structured interview, conducted and transcribed in Dutch and translated into English for analysis and reporting. Unlike the first-round interview guide, the second-round interview guide was structured around the specific design adjustments made in Version 2, pairing each adjustment with the first-round finding it was intended to address and prompting respondents to confirm, qualify, or reject whether that adjustment had resolved the underlying concern.

Qualitative analysis of the second-round transcripts followed the same reflexive thematic analysis procedure described in section 3.6.2, using the same Atlas.ti code environment as the first round. Codes from the first round were reused where the same underlying concept recurred, and new codes were created where second-round respondents raised considerations specific to the Version 2 implementation that

had not been, and in some cases could not have been, identified in the first round, since those were considerations of workbook features that did not exist in Version 1.

3.9 Research quality criteria

The quality of the mixed-methods study is evaluated against criteria appropriate to each methodological strand. For the quantitative component, internal consistency across construct items is assessed through consideration of convergent validity (the degree to which items operationalizing the same construct produce consistent ratings across respondents). For the qualitative component, quality is evaluated against the criteria of credibility, transferability, and confirmability as articulated by Lincoln and Guba [40].

Credibility is enhanced through member-checking procedures: key findings were verified against verbatim transcript data, and all qualitative interpretations are grounded in direct citations. Dependability was supported by the use of a standardized protocol that ensured procedural consistency across all seven sessions. Confirmability is established through the use of Atlas.ti as an auditable coding environment, in which each code is traceable to a specific verbatim transcript fragment, and through the provision of a complete code list as an appendix. Transferability, which in qualitative research cannot be claimed but only facilitated, is addressed through a rich description of the respondent sample, the interview procedure, and the analytical strategy, enabling readers to assess the applicability of findings to analogous professional contexts.

3.10 Synthesis

The methodological architecture assembled in this chapter is not a collection of independent decisions but a coherent structure in which each choice reinforces the others. DSR provides the rationale for expert-validated artifact construction; TAM provides the theoretical vocabulary for operationalizing expert judgment; the sequential explanatory mixed-methods design ensures that numerical signs are always interpretable in light of substantive professional reasoning; and reflexive thematic analysis ensures that the richness of expert critique is not reduced to aggregate scores but preserved as a source of actionable design intelligence.

The broader intellectual implication of this architecture is that artifact-based IT audit research need not choose between precision and relevance. By embedding a theoretically grounded acceptance model within a design science evaluation cycle, and by subjecting both quantitative and qualitative data streams to systematic analytical procedures, the study demonstrates that practitioner validation can be conducted with the same methodological precision demanded of any empirical inquiry while remaining directly responsive to the design challenges of a nascent, practically urgent domain.

Chapter 4

Design

The preceding chapters established two foundations for the design work presented here. Chapter 2 identified a structural gap in the European cloud sovereignty landscape: while frameworks such as DICTU's Toetsingsinstrument, the EC Cloud Sovereignty Framework, and EuroStack collectively cover the relevant sovereignty dimensions, none of them provides IT auditors with the combination of scoring granularity, independently verifiable evidence requirements, and a gated aggregation mechanism needed to produce a defensible, evidence-based audit conclusion on operational reversibility. Chapter 3 described the DSR methodology through which a purpose-built artifact is constructed to fill that gap and then evaluated against the criterion of utility in the professional audit context.

It documents the DSAF. A distinction is necessary here. The DSAF as a whole is conceived as an audit framework: a structured approach for IT auditors to assess digital sovereignty in cloud-dependent organizations. Within that framework, individual control instruments address specific sovereignty dimensions. This chapter presents the first such instrument, covering the operational reversibility dimension, meaning the ability of an organization to exit its cloud provider without catastrophic loss.

This chapter presents the artifact. It documents the DSAF in its second, revised iteration. Version 1 of the DSAF was the instrument presented to and evaluated by the seven Round 1 expert respondents (R01 to R07) described in chapter 3. Version 2, documented here, is the result of incorporating the design implications of that evaluation into the artifact before it was submitted to the second, confirmatory evaluation round. A significant change between versions concerns the implementation format. Version 1 was presented as a document-based instrument. Version 2 is implemented as an Excel workbook. This change was directly motivated by the Round 1 evaluation findings: six of seven respondents indicated that an Excel format was significantly more compatible with their existing audit practice than a document-based instrument, and that format compatibility was a precondition for practical adoption rather than a secondary preference.

The distinction matters for reading this chapter correctly: the design choices explained below are not hypotheses awaiting validation but decisions that have already been through one cycle of expert critique and have been revised accordingly. Where a choice changed between versions, the reason is stated explicitly. Where it did not change, that stability itself is part of the design record.

In the interest of scientific transparency and practical accessibility, the complete Version 2 workbook is made publicly available as an open-source resource. It can be downloaded at `https://zaynab-m.github.io/DSAF/`. Researchers seeking to extend the DSAF to additional sovereignty dimensions, and practitioners wishing to apply the instrument in their own audit engagements, are encouraged to use the workbook as a starting point. The workbook as documented in this chapter and the downloadable version are identical.

The central question this chapter answers is, given the gap identified in chapter 2 and the design requirements it generated, what is an effective design that an IT auditor can actually use to assess whether an organization can exit its cloud provider without catastrophic loss of service continuity, data, or institutional knowledge? The answer is developed in four steps. Section 4.1 states the design requirements derived from the gap analysis and presents the overall artifact architecture. Section 4.2 explains the generic structure that is instantiated directly across five control areas that together cover the full scope

of operational reversibility as an auditable property. Section 4.3 works through each of the five control areas in turn, documenting what, where, how, and why. Section 4.4 closes with the design rationale for the four decisions that required the most deliberation, including the alternatives that were considered and the reasons why they were rejected.

4.1 Design requirements and artifact architecture

This section translates the gap analysis of section 2.4 into a concrete instrument design.

4.1.1 Design requirements derived from the comparative gap analysis

The design requirements for the DSAF are not stated as design preferences; they follow directly from the comparative analysis of existing European sovereignty frameworks conducted in section 2.4. The table 2.10 scored five existing frameworks (the European Commission's Cloud Sovereignty Framework (EC-CSF), EuroStack, the DICTU Toetsingsinstrument, the EuroStack Buy European Framework for strategic procurement, and the SUSE Cloud Sovereignty Framework Self-Assessment (SUSE CFSA)) against six criteria considered material to audit practice, each weighted by its relevance to an IT audit use case.

Three patterns in table 2.10 directly shaped the DSAF's design requirements. First, DICTU (3.80/4.00) and EC-CSF (3.60/4.00) score highest overall and, specifically, achieve the maximum score (4/4) on criterion (2), evidence standard. Both require independently verifiable documentary or technical evidence rather than self-assertion. By contrast, EuroStack and the SUSE CFSA score only 1/4 on this criterion: EuroStack is a strategic policy document with no evidentiary apparatus, and the SUSE CFSA, despite being a structured self-assessment tool, produces entirely self-reported, unverified outputs. This gap (high-scoring frameworks demand verifiable evidence) motivates DR4 (verifiable quantitative indicators) and DR5 (audit work-paper structure with an 'evidence to provide' field per criterion).

Second, DICTU is the only framework to score the maximum on criterion (3) scoring granularity (4/4), while EuroStack scores 1/4 and the Buy European Framework scores 2/4. DICTU achieves this through its five-level (M1-M5) maturity scale for criterion 3.4 [15]; EuroStack and Buy European Framework, by contrast, are oriented toward binary or non-binary procurement gates rather than graded diagnostic profiles. This gap motivates DR2 and DR3: the DSAF adopts DICTU's graded M1-M5 structure (rather than EuroStack's or Buy European coarser logic) specifically because granularity scores as the differentiator between an audit-grade instrument and a procurement-screening instrument.

Third, the Buy European framework is the only framework to score the maximum on criterion (5) auditability (4/4), despite only scoring only 2/4 on scoring granularity; its contribution is structural rather than diagnostic: a binding jurisdictional gate that filters providers pass/fail before any graded assessment occurs [23]. This is the structural origin of DR1 and, more specifically, of the weakest-link gating logic in table 2.10: the DSAF reuses the Buy European framework's pass/fail gate principle but applies it not as a jurisdictional pre-screen but as the C1 'exit plan exists at all' prerequisite that gates the remainder of the assessment.

Section 4.1 lists the resulting design requirements together with the corresponding artifact component and the primary source from which each requirement was derived: the DICTU and EC-CSF sources for DR1, DR2, and DR6 (dimensional coverage and regulatory grounding, criteria 1 and 4 of table 2.10); the evidence-standard gap for DR4 and DR5 (criterion 2); the granularity gap for DR3 (criterion 3); and the auditability/proportionality gap for DR7 and the gating logic (criteria 5 and 6). This table operationalizes the requirements-to-design mapping recommended in Peffers et al. (2008) [44] and corresponds to Hevner et al.'s (2004) [31] second guideline, problem relevance: every design decision in this chapter can be traced to a requirement in this table, and every requirement can be traced to a specific score gap in table 2.10.

Table 4.1: Design Requirements and Corresponding artifact Components

#	**Design requirement**	**artifact component**	**Primary source / Table gap**
DR1	The framework must produce an auditable conclusion on whether an organization can exit its current cloud provider without catastrophic loss of continuity, data, or institutional knowledge.	Five control areas (C1-C5), each with audit questions and a documented finding (Appendix B).	DICTU criterion 3.4 [15]; EuroStack criterion 2.3 [23]; Table 2.10 criterion (1) dimensional coverage — DICTU/EC-CSF score 4/4
DR2	Maturity must be assessed on two independent dimensions: whether a control is designed adequately, and whether it operates as designed.	Part A (Design Maturity, M1-M5) and Part B (Operating Effectiveness, E1-E4) per control area.	DICTU v1.0.1, five sovereignty levels for criterion 3.4 [15]; EuroStack criterion 2.3 [23]; Table 2.10 criterion (3) scoring granularity — DICTU 4/4 vs EuroStack 1/4
DR3	The framework must be usable by IT auditors without requiring them to make a graded binary judgment on an inherently graded judgment.	Average maturity score with a weakest-link signaling guideline, rather than a binary Effective/Not Effective verdict per control.	Professional-judgment principle; DICTU v.1.0.1 [15]; Table 2.10 criterion (3) granularity gap
DR4	Quantitative indicators used within the framework (e.g., Infrastructure-as-Code coverage) must be verifiable and resistant to assertion-based reporting without becoming rigid pass/fail gates that displace professional judgment.	IaC coverage thresholds (50%/80%) embedded as cumulative criteria within the maturity ladder, paired with verification procedures (spot-checks, functional deployment tests).	EuroStack (2025) [23], criterion 2.1; Table 2.10 criterion (2) evidence standard — DICTU/EC-CSF 4/4 vs EuroStack/SUSE CFSA 1/4
DR5	Evidence requirements per control must mirror the structure of a standard audit workpaper, so the framework integrates into existing IT audit methodology without requiring auditors to learn a new evidentiary paradigm.	Each control area follows: Sovereignty objective → Risk → Sovereignty link → Audit questions → Part A → Part B → Evidence to provide → Reached (Appendix B)	ISAE 3402 [34] / SOC 2 [3] work paper structure; Table 2.10 criterion (2) evidence standard

#	Design requirement	artifact component	Primary source / Table gap
DR6	The framework must address not only contractual and technical exit conditions but also the continuity of operations during the exit itself and the periodic testing of exit capability, without which a documented exit capability remains theoretical.	Control areas C4 (Continuity of operations during migration) and C5 (Testing programme and periodic evaluation)	NIS2 Art. 21 [19]; DORA Art. 11 [18]; ISO 22301 §8.5 [37]; Table 2.10 criterion (4) regulatory grounding — DICTU/EC-CSF 4/4
DR7	The framework must be replicable: an auditor unfamiliar with its construction must be able to apply it consistently from its documentation alone, without recourse to the designer.	Fixed generic structure (section 4.2) applied identically across all five control areas; explicit mapping tables linking audit questions to maturity and effectiveness levels (section 4.3); full evidence specification in Appendix B	Hevner et al. [31], guidelines 3 and 7; Table 2.10 criteria (5) auditability en (6) proportionality — Buy European 4/4 en 4/4

4.1.2 Artifact architecture

The architecture of the DSAF is a direct structural response to the two largest unresolved gaps identified in table 2.10: scoring granularity (criterion 3, where the spread between DICTU's 4/4 and EuroStack's 1/4 was the widest of any criterion) and auditability (criterion 5, where the Buy European framework's 4/4 came from a binary gate rather than a graded scale, and DICTU and the EC-CSF (despite their highest overall scores) only reached 3/4. No single existing framework combines a high-granularity grade scale (DICTU's strength) with a high-auditability gating mechanism (Buy European's strength) within a single per-criterion structure that also carries its own evidence specification (criterion 2). The DSAF's architecture is designed to combine these three properties (granularity, gating, and evidence) within a unified structure. This combination is illustrated in fig. 4.1.

Figure 4.1 presents the artifact as a process flow because this structure makes explicit both the sequential logic an auditor follows and the repeating unit that applies identically across all five control areas, while preserving the three architectural properties the design requirements demand: granularity in the per-control scoring, an evidence specification per step, and a gating aggregation at the end.

The flow is organized across four steps. Step 1 is a one-time orientation phase in which the auditor reads the Instructions tab, which contains the Step-by-Step Workflow and the explanation of how to score Part A versus Part B, and the Definitions tab, which provides a shared glossary of terms such as RTO, IaC, and tabletop exercise. This orientation layer addresses the replicability requirement of DR7: an auditor unfamiliar with the framework's construction must be able to apply it consistently from the documentation alone.

Step 2 is the assessment phase and forms the core repeating unit of the process flow. It is executed identically five times, once for each control area in sequence from C1 to C5. Each iteration moves through five sub-steps: (2a) reviewing the sovereignty objective, risk statement, and sovereignty link, which anchors the assessment in its normative authority and addresses criterion (4), regulatory grounding, of table 2.10; (2b) answering the fixed set of audit questions, which forms the evidence specification that EuroStack and SUSE CFSA lacked and addresses criterion (2), evidence standard; (2c) scoring Part A, Design Maturity, by selecting the highest M1 to M5 level whose cumulative criteria are all satisfied, which operationalizes the graded granularity that only DICTU achieved at 4/4 on criterion (3); (2d) scoring Part B, Operating Effectiveness, by selecting the highest E1 to E4 level, which is the DSAF's own addition since none of the five frameworks in table 2.10 addressed demonstrated operation as a

separate axis from design intent, justified further in section 4.2.4; and (2e) recording the Responsible Role table together with the documented finding and rationale for the assigned scores.

Each completed iteration of Step 2 produces one score pair per control area: an M-level (M1 to M5) reflecting the maturity of the control's design, and an E-level (E1 to E4) reflecting whether that design actually operates as intended in practice. These two scores are always recorded independently and are never averaged together, because a well-designed control that is not operating (high M, low E) represents fundamentally different risk than a poorly designed control that somehow produces adequate outcomes (low M, high E). Together with the documented finding and rationale, this score pair gives the auditor a precise, two-dimensional picture of each control area before moving to the next.

Step 3 is the aggregation phase, executed once after all five control areas have been assessed. Throughout Step 2, the auditor selects the reached M-level and E-level for each control area using a drop-down in the corresponding control sheet. Because the Version 2 artifact is implemented as an Excel workbook, these selections are automatically carried through to the Summary tab without manual re-entry, which eliminates transcription errors and ensures that the summary always reflects the scores as assigned. The summary tab then computes the average Design Maturity and average Operating Effectiveness automatically across all five control areas. The Decision Guidelines, shown in table 4.3, flag weak links: like the Buy European framework's pass/fail gate, they allow a single control area to determine the overall outcome regardless of the average, but unlike Buy European's binary gate, this happens only after the full M1 to M5 and E1 to E4 assessment has been completed for every control area before the Decision Guidelines are applied. This means the gating logic and the granularity scoring coexist: a single critically weak control area can still determine the overall conclusion, but the detailed scores for all other control areas are preserved and visible in the audit record. In the Buy European Framework, the gate replaces granular scoring entirely. In the DSAF, the gate is applied after granular scoring is complete, so no detail is lost in producing the overall conclusion.

Step 4 is the overall audit conclusion, in which the auditor documents the outcome of the assessment as an advisory finding, supported by the findings, recommendations, and follow-up actions recorded per control area throughout Step 2.

This combination of a graded per-control assessment, a fixed evidence specification per step, a responsible-role and rationale record per control, and a gating aggregation across all five controls is what no single framework in table 2.10 provided in one instrument, and it is the principal architectural contribution operationalized in fig. 4.1.

4.2 Generic structure of a control area

Each of the five control areas in the DSAF instantiates the same five-column structure shown in fig. 4.1. This section explains each in turn: what it is, what role it plays in the audit procedure, and why it was designed this way and from which source it derives.

4.2.1 Sovereignty Objective, Risk, and Sovereignty Link

Each control area opens with three linked elements. The sovereignty objective is a single sentence stating, in auditable language, what must be verified; for example, for C1, 'verify that a formally approved, contractually supported exit plan exists, enabling an orderly transition to an alternative provider within a defined Recovery Time Objective'. The risk statement describes the consequences if this objective is not met, expressed in terms of the operational or contractual harm that would follow, not in terms of the control itself, so that the risk remains intelligible to non-technical stakeholders (e.g., legal and the audit committee) reading the audit report. The sovereignty link then maps the control explicitly to its originating criterion: the DICTU Toetsingsinstrument criterion [15], the corresponding EuroStack Buy European Framework criterion [23], and, where applicable, EU regulatory provisions (GDPR [20], DORA [18], and ISO 22301 [37]).

This three-part opening exists because of DR1 and DR5: an auditor reviewing the resulting finding in an engagement file must be able to reconstruct the reasoning behind each score independently, without needing to contact the person who performed the assessment, 'What exactly was checked? Why does it matter? And on what authority is this check based?' Separating objective, risk, and link into distinct elements, rather than combining them into a single narrative paragraph, as in Version 1 of the framework,

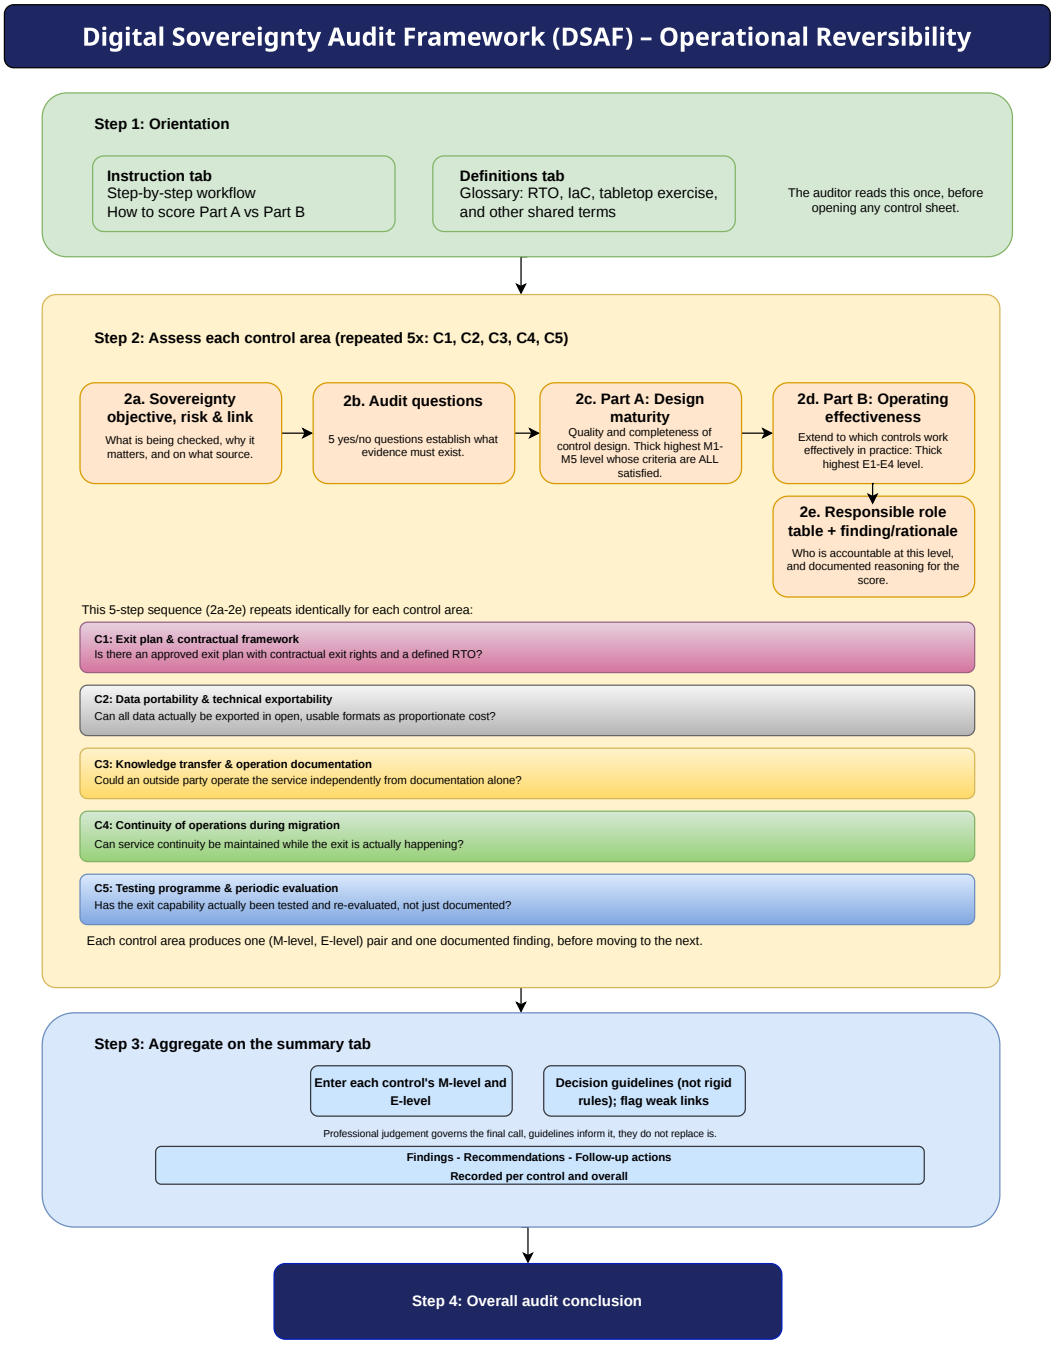


Figure 4.1: Process flow of the DSAF Operational Reversibility instrument (Version 2). M = Design Maturity (M1 - M5); E = Operating Effectiveness (E1 - E4). The instrument is available at `https://zaynab-m.github.io/DSAF/`

makes each element independently citable in an audit report: a finding can reference the sovereignty link directly (e.g., 'see DICTU 3.4, EuroStack 2.3') without re-deriving the risk argument each time.

4.2.2 Audit questions

Each control area contains exactly five closed (yes/no) audit questions (fig. 4.1). These questions establish what evidence must exist, for example, 'Does a formally approved, version-controlled exit plan document exist for this cloud service,' independently of how complete, current, or verified that evidence is. The completeness, currency, and verification requirements are specified separately in the Part A maturity criteria.

Two design choices are embedded in this column. First, the questions are deliberately closed (yes/no) rather than open, because the framework is intended to integrate into a standard audit work paper (DR5), where closed questions with a documented answer and evidence reference are the established format (ISAE 3402) [47]. Second, the number of questions is fixed at five per control area across all five control areas, producing a 25-question instrument; this uniformity was chosen so that the instrument's length and structure are predictable for an auditor moving between control areas, supporting the replicability requirement (DR7). The audit questions remain stable across maturity levels by design so that the same evidence-gathering step (answering the 25 questions) always precedes the maturity determination, regardless of which level is ultimately reached.

4.2.3 Part A: Design Maturity (M1-M5)

Part A (fig. 4.1) assesses design effectiveness: would the control, if it operated exactly as documented, mitigate the underlying risk? Table 4.2 reproduces the five generic maturity levels verbatim from the DICTU Toetsingsinstrument v1.0.1 (2026, p.20) [15], which defines five sovereignty levels for criterion 3.4 in generic form, ranging from 'no reversibility' M1 to 'full operational reversibility' (M5).

Table 4.2: Overview of Maturity Levels and Audit Conclusions

Level	Designation	DICTU definition	Audit conclusion
M1	No reversibility	Service is not transferable; continuity upon provider failure is not guaranteed.	Not effective - critical finding
M2	Limited reversibility	Documentation is available to a limited extent; transfer to a third party is complex and not feasible.	Not effective - significant finding
M3	Contractual reversibility	Service largely transferable; documentation and design available at a high level to enable third-party operation over time.	Partially effective
M4	Active reversibility	Service fully transferable; documentation and design sufficiently detailed to enable independent third-party operation.	Effective
M5	Full operational reversibility	Service fully transferable; documentation and design sufficiently detailed and tested for independent third-party operation.	Fully Effective

For each control area, the DSAF operationalizes these five generic levels into concrete, cumulative criteria specific to that control's evidence: a level can only be reached if all criteria of the preceding level are also satisfied (full criteria per control area: Appendix B). This cumulative design is the subject of section 4.4.1, where it is justified against two rejected alternatives. Two points are stated here directly because they recur throughout section 4.3: first, the M1-M5 scale and its audit-conclusion mapping (the right-hand column of table 4.2 is identical across all five control areas, so a finding of 'M2' always carries the same conclusion label ('Not effective - significant finding') regardless of which control area it applies to. This consistency allows the M_{avg} aggregate in section 4.4.2 to be meaningfully computed across control areas. Second, M5 (across all control areas) always adds a testing or independent verification requirement on top of the M4 criteria (e.g., the documentation independence test in C3, the tabletop-with-rollback in C4);

this reflects the DICTU (2026) wording that M5 requires documentation and design to be 'sufficiently detailed and tested,' and it is the structural reason why M5 is addressed separately from M1-M4 in the design-rationale discussion testing.

4.2.4 Part B: Operating Effectiveness (E1-E4)

Part B (fig. 4.1) assesses operating effectiveness: did the designed control actually operate as intended during the audit period? This is independent of Part A; a control can be well designed (M4) yet never exercised in practice (E2), or only partially designed (M3) yet operated consistently within that limited scope. M-scale and the E-scale, therefore, form a two-dimensional space (5 x 4 = 20 possible (M, E) combinations per control area), and a single combined score (e.g., 'M4/E2') is more informative than either dimension alone.

The four-level E1-E4 scale (Not effective, Partially Effective, Largely Effective, Fully Effective) was adopted from the design-versus-operating effectiveness distinction conventional in ISAE 3402 assurance engagements [34], a deliberate choice to align the framework's second dimension with a scale that IT auditors already use in other engagements, again in support of DR5 and DR7. A feature that recurs by design across all five control areas is that E4 always requires demonstrated execution (a completed export, an executed rollback, or a live deployment from documentation) rather than the existence of a plan or a tabletop discussion of intent; this is what operationalizes the word 'tested' from the DICTU (2026) M5 definition [15] on the operating effectiveness side of the framework.

4.2.5 Aggregation and decision logic

After each control area (each row of fig. 4.1) has been assigned an (M, E) pair, the framework computes an average design maturity (M_{avg}) and an average operating effectiveness (E_{avg}) across the five control areas and combines these with the decision guidelines in table 4.3 to reach the single overall audit conclusion shown at the base of fig. 4.1. Version 1 of the artifact stated this logic as fixed decision rules; Version 2, documented here, restates the same five conditions as guidelines using deliberately non-mechanical language ('generally,' 'usually,' 'absent specific countervailing evidence'). The reasoning behind this change is set out in section 4.4.2.

Table 4.3: Decision Guidelines for Overall Audit Conclusions

Guideline	Condition	Guidance
1	Control 1 at M1 or E1	Treat Control 1 (the exit plan) as a prerequisite for the others. If it sits at M1 or E1, the overall conclusion should generally be Not Effective — strong scores elsewhere cannot compensate. Document any departure from this.
2	Any control at M1 + E1	If any single control is at M1 and E1, either design or operating evidence are effectively absent for that area. This generally points to Not Effective overall, unless the auditor can justify why the gap is not material to reversibility.
3	All M2+/E2+, but Control 5 at M1 or E1	Where controls are broadly adequate (M2+/E2+) but the testing programme (Control 5) lags at M1 or E1, reversibility is designed but not yet confirmed in practice. Partially Effective is usually appropriate; consider recommending testing as the priority action.
4	All controls M3+/E3+	If all five controls reach at least M3 in design and E3 in operation, exit capability is both designed to a contractual standard and operationally demonstrated. The auditor can generally conclude Effective, absent specific countervailing evidence.
5	All controls M5/E4	Reserve Fully Effective for cases where every control reaches M5 and E4, maximum operational reversibility (DICTU Level 5). Expect this to be supported by tested evidence across all controls.

The averages (M_{avg}, E_{avg}) satisfy DR3 by preserving the graded information from all five control areas in a single comparable figure, useful for trend reporting across repeated audits of the same organization. The decision guidelines then encode two dependency relationships that an average alone cannot express: Control 1 (the exit plan itself) is a logical precondition for the other four controls to be meaningful; an organization cannot have a tested migration continuity plan (C4) for an exit it has not planned, and Control 5 (the testing program) verifies whether the capability described by C1-C4 has actually been demonstrated. Guidelines 1 and 2 encode the first relationship; Guideline 3 encodes the second. This 'average plus weakest-link guidelines' design, and the alternatives considered for it (including an earlier version of the framework that used a binary per-control verdict) are discussed in section 4.4.2. The guidelines are framed as guidance for professional judgment rather than as a mechanical checklist: where the auditor finds the evidence ambiguous or the organizational situation atypical, a documented rationale is required and may justify deviation from the guideline-based outcome, consistent with the professional judgment principle stated in DICTU v1.0.1 [15].

4.3 The five control areas

This section instantiates the generic structure of section 4.2 for each of the five control areas. For each control area, the discussion opens with what the control area is and why it exists, followed by its five audit questions, and then works through each Part A maturity level (M1-M5) and each Part B effectiveness level (E1-E4) in turn. At each level, the text states (quoting or closely paraphrasing Appendix B) what specifically must be evidenced to reach that level and explains why the threshold is placed where it is. This level-by-level structure is intended to support replication (DR7): an auditor reading this section should be able to determine, for any piece of evidence encountered in the field, which maturity or effectiveness level it does or does not satisfy, and why. Each control area closes with a short synthesis identifying its most distinctive design feature. Full cumulative criteria and evidence requirements are provided in Appendix B.

4.3.1 Control area 1: Exit plan and contractual framework

C1 verifies that the organization processes a formally approved, contractually supported exit plan with a defined Recovery Time Objective (RTO) for the cloud service under review. It is written as the first control area and as the gating control in the aggregation logic (section 4.2.5), because an exit plan is the precondition for everything the other four control areas assess: C2 (data portability) is only relevant if an exit is planned to begin with; C3 (knowledge transfer and operational documentation) presupposes that there is a transition to document for, since the quality of exit documentation only matters when someone will actually use it to take over operations; C4 (continuity during migration) presupposes that a migration has been scoped; and C5 (testing programme) presupposes that there is a programme to test. The risk this control area addresses is stated in direct terms in Appendix B: without a plan, 'vendor lock-in is not a risk in this scenario but a certain outcome' in the event of provider insolvency, geopolitical sanctions, or acquisition by a non-EU entity. The sovereignty link is DICTU criterion 3.4 [15] and EuroStack criterion 2.3 [23], the latter stating directly that 'a contractually binding exit plan is mandatory' [23].

Audit questions. The five audit questions for C1 are not five independent checks; together, they test whether an exit plan exists as a credible artifact rather than as a label. Questions 1.1 and 1.2 establish the document itself: whether a version-controlled plan exists at all and whether the contract underneath actually grants the rights (data export, transition assistance, and no punitive egress fees) that would let the plan be executed. A plan without these contractual rights is, in practice, a description of an exit the organization has no legal means to carry out [23]. Question 1.3 adds a time dimension that the first two questions cannot capture on their own (a formally approved RTO) because a plan with no defined recovery time objective gives the auditor no way to judge whether the plan is fast enough to matter. Questions 1.4 and 1.5 then test whether the plan is actually usable rather than merely approved: 1.4 checks that at least one credible, EU-domiciled alternative has been identified (so the plan has somewhere to go), and 1.5 checks that a named owner and review cycle exist (so the plan does not silently go out of date). The five questions are deliberately ordered this way (existence, then contractual substance, then timing, then destination, then maintenance) because this is also the order in which the Part A criteria in Appendix B accumulate from M1 to M5.

What must be shown for each maturity level (Part A)

M1 - No Reversibility. Appendix B specifies that M1 applies if 'no exit plan exists in any form,' 'no contractual exit rights in the service agreement,' and 'no consideration has been given to what happens if the provider fails.' Critically, the evidence requirement for M1 is framed as an absence-by-default: 'absence of a named exit plan document confirms M1. If no evidence is produced this level applies by default.' This means M1 is not something the auditor needs positive evidence for. It is the result of the organization failing to produce evidence for 1.1 and 1.2. This design choice (discussed further in section 4.2.4) prevents an organization from arguing its way out of M1 through narrative explanation in the absence of a document.

M2 - Limited Reversibility. To move beyond M1, an exit plan document must exist in some form (draft, informal note, or process document). However, Appendix B is explicit that the mere existence of such a document does not by itself indicate adequacy: 'the plan is significantly outdated (more than 24 months without review) or lacks formal approval,' and contractual exit rights are 'absent or contain terms that make exit economically prohibitive.' M2 is therefore the level at which an organization has 'thought about' exit but has not formalized or maintained that thinking. The auditor guidance notes that 'outdated or unapproved plans confirm M2' as a ceiling, not a pass.

M3 - Contractual Reversibility. M3 is the first level at which the exit plan is treated as a credible artifact: it must be 'formally approved by an appropriate authority (minimum: IT Director or CIO),' the contract must show 'explicit exit rights, including data export rights, no punitive egress fees, and minimum transition assistance obligation,' and the RTO 'is defined (though may not have been tested or validated).' This is the level DICTU [15] describes generically as the service being 'largely transferable' with 'documentation and design available at a high level.' Note the explicit qualifier 'though may not yet be tested,' which is what keeps M3 distinct from M4 and M5: at M3 the plan is approved and current, but its workability has not been demonstrated.

M4 - Active Reversibility. M4 adds operational specificity to the M3 plan. The exit plan must now be 'detailed and current (reviewed within 12 months), containing migration phases, phase-gate criteria, parallel operational periods, data synchronization approaches, and rollback trigger conditions.' This is the point at which audit questions 1.4 and 1.5 become decisive, requiring 'at least two credible alternatives

identified with compatibility assessment; at least one EU-domiciled alternative' and a 'named owner by role; formal review cycle defined and evidenced.' Additionally, the RTO must have been 'validated against the business impact analysis of the relevant cloud service, i.e., the RTO is no longer just a number the CIO approved (M3), but a number that has been checked against what the business can actually tolerate.

M5 - Full Operational Reversibility. M5 requires that the M4 plan 'has been validated through at least a structured tabletop exercise simulating an unplanned exit trigger (provider insolvency, sanctions, or non-EU acquisition),' that 'at least one alternative provider is pre-qualified or has a contract on standby, ' and that the exit plan and RTO have been 'reviewed or independently assessed within 12 months.' The wording 'unplanned exit trigger' is deliberate and recurs at M5 across multiple control areas (see C4 and C5): a tabletop exercise that only rehearses a voluntary, planned migration does not satisfy M5, because it does not test the scenario the control objective is actually concerned with, an exit forced by circumstances outside the organization's control.

What must be shown for operating effectiveness (Part B)

E1 - Not effective. The auditor guidance for E1 is explicit about its triggering condition regardless of the Part A score: 'if the exit plan has not been reviewed in 24 months or key stakeholders are unaware of it, rate E1 regardless of maturity score.' The audit procedure includes interviewing the exit plan owner to check 'can they explain the plan without looking at it?.' Operating effectiveness is partly about whether the plan lives in institutional memory, not only on a server.

E2 - Partially effective. E2 requires the plan reviewed within 24 months, contractual exit rights confirmed as current, RTO approved by the CIO, and alternative providers identified (i.e., the M3 criteria, demonstrating it as still current. The auditor guidance flags this as 'the most common state: 'plan is maintained but has never been tested. Document as an observation and recommend a tabletop exercise as a priority action.' E2 is therefore not a weak score in absolute terms; it describes a normal, defensible position that nonetheless falls short of demonstrated capability.

E3 - Largely effective. E3 requires a tabletop exercise within 24 months whose 'scenario included an unplanned trigger,' with findings 'addressed and documented.' The guidance is explicit that 'planned migration only is insufficient for E3,' mirroring the M5 wording above; the operating effectiveness side of the framework applies the same 'unplanned trigger' filter to whatever evidence of testing is presented.

E4 - Fully effective. E4 requires the exit plan to have been 'tested via live exercise or partial migration,' an alternative provider 'pre-qualified or contracted on standby,' and all prior exercise findings closed. The guidance draws the same proof of execution line seen at M5: 'An organization that has only conducted tabletop exercises but no live test is at most E3.' E4 is therefore reserved for organizations that have moved from rehearsal to execution.

What this means for C1 in practice.

The distinctive feature of C1, evident across both Part A and Part B, is the recurring 'unplanned trigger'/'live execution' threshold that separates M4/E3 from M5/E4: an organization can have an approved, detailed, reviewed exit plan with named owners and identified alternatives (M4) and have rehearsed it via tabletop (E3) and still not qualify for M5/E4 unless that rehearsal specifically simulated an involuntary exit and unless at least one element has been executed live (a pre-qualified standby provider or a partial migration). This design choice (the same demonstrated execution threshold recurring at the top of both scales) is what the framework uses to operationalize the word 'tested' in DICTU's M5 definition ('sufficiently detailed and tested' [15]) in a way that is checkable against a specific artifact (a tabletop report, naming the trigger scenario) rather than against a general impression of preparedness.

4.3.2 Control area 2: Data portability and technical exportability

C2 verifies that all data (primary data, backups, configurations, metadata, and logs) can be exported at any time, in open, nonproprietary formats, at proportionate cost, with the export procedures themselves documented and tested. Where C1 establishes an exit is contractually possible, C2 establishes that it is technically possible: a contract granting exit rights is of limited value if the data itself cannot leave the platform in a usable form. Appendix B frames this as the 'data gravity gap.' A situation in which proprietary formats, excessive egress fees, and missing export tooling make data technically immobile regardless of contractual permission. EuroStack (2025) describes punitive egress fees specifically as

'strategic barriers that prevent a competitive European market' [24]. The sovereignty link is DICTU criterion 3.1 (Interoperability and Portability) [15], EuroStack criterion 2.1 [23], and GDPR Art. 20 (the right to data portability) [17].

Audit questions. C2's five-question test portability comes from two directions that are easy to conflate but evidently distinct: whether data can leave the current environment (export) and whether it can be used once it has left (import). Questions 2.1 through 2.4 establish exportability, open nonproprietary formats for primary data (2.1); a contract that does not financially punish the act of exporting (2.2); infrastructure configuration held outside the provider's own systems (2.3); and the same portability standard extended to backups, not only primary data (2.4). Question 2.5, by contrast, explicitly requires 'import validation in an alternative environment,' not export alone. This control area combined export and import questions deliberately, rather than treating them as two separate controls, because an export that cannot be read elsewhere provides no real portability: EuroStack [24] frames this as a 'data gravity trap' in which data is technically copyable but practically immobile, and DICTU's own M5 definition for this criterion requires documentation and design to be 'sufficiently detailed and tested' [15]. A requirement that is only meaningful if testing covers both ends of the transfer. Splitting export and import into separate control areas would also have broken the cumulative logic described in section 4.2.3: an organization could score well on a stand-alone 'export' control while having never confirmed that anything it exported was usable, which is precisely the gap question 2.5 is designed to close.

What must be shown at each maturity level (Part A)

M1 - No Portability: M1 is defined by three cumulative absences: 'data stored exclusively in proprietary provider formats with no open export mechanism'; 'egress fees are prohibitive, or the organization has not assessed the termination cost'; and 'no IaC exists; infrastructure configurations are undocumented or stored only in provider consoles.' As with C1's M1, the evidence requirements are framed around absence: the organization cannot 'provide any export format specification.' DICTU (2026) describes this generically as a service that is 'fully proprietary' where 'migration' of 'data and workloads is not possible without significant effort or cost. [15]'

M2 - Limited Portability: M2 requires that export is 'possible for at least some data types but requires significant provider involvement or reformatting effort,' that 'some core data types remain in proprietary formats,' and that egress fees 'have been reviewed but not modeled for a full exit scenario.' The evidence at this level is a 'data inventory showing which data types have open export and which do not.' In other words, M2 is reached once the organization can at least describe the shape of its portability gap, even if it has not closed it.

M3 - Open Standards Partial: M3 is the level at which Infrastructure-as-Code coverage first becomes a quantitative criterion: 'IaC covers at least 50% of production infrastructure components in organization-controlled repositories.' In addition, 'all primary data types have a documented open format export option,' the contract 'confirms no punitive egress fees at termination,' and 'backup export formats have been assessed, even if not yet tested.' The 50% figure functions as a midpoint marker. By M3, a majority of the infrastructure configuration is no longer trapped in provider consoles, but the framework does not yet require this to be demonstrated through an actual report.

M4 - All Open Standards: M4 raises the IaC threshold to 'exceeds 80% of all production infrastructure components' and critically requires that this figure is not merely reported but evidenced: the corresponding evidence entry calls for an 'IaC coverage report confirming more than 80% (spot-check by auditor recommended).' M4 also requires an 'export test conducted for all primary data types within 24 months, confirming successful export without provider tooling,' that 'backup and DR replicas are confirmed to use the same open format and portability standard as primary data,' and that 'data egress cost has been modeled for a realistic full-exit scenario, confirming proportionality'. The explicit 'spot-check by auditor recommended' instruction at M4 is the textual anchor for the verification requirement discussed in section 4.4.4.

M5 - Maximum Portability: M5 requires a 'full end-to-end export test: all data types, backups, metadata, IaC configurations, and logs' and, distinctly, 'import validation in an alternative environment that confirms no data loss and readability without providing tooling.' The distinction between M4 and M5 is therefore the distinction between exporting data and ensuring that the export is usable elsewhere. M4 stops at 'successful export'; M5 requires 'import validation ... not just export success.' M5 additionally requires that results are 'independently reviewed or included in internal audit scope.'

What must be shown for operating effectiveness (Part B)

E1 - Not Effective: E1 applies if 'export has never been tested for any data type,' the organization 'cannot identify the data format in use,' and 'egress fees at termination have not been reviewed.' The audit procedure includes an 'attempted walkthrough of export for one minor data type.' A practical, hands-on check rather than a documentation review. The guidance states: 'if no export test has been conducted and the organization cannot explain what format their data is in, rate E1 regardless of contractual claims.'

E2 - Partially Effective: E2 requires primary data types 'confirmed in open format' and IaC repository ownership confirmed, but allows that 'some export tests were conducted, but coverage is incomplete.' The guidance flags a specific recurring gap: 'partial test coverage does not constitute demonstrated portability. Metadata and logs are frequently the untested gap.' Naming the categories of data most likely to be overlooked when an organization claims export readiness.

E3 - Largely Effective: E3 requires all primary data types tested for export, export duration validated against RTO, and backup portability confirmed. The audit procedure for E3 includes a concrete readability test: 'attempt to open or parse the exported file using a non-provider tool.' The guidance is explicit that 'E3 requires that exported data has been validated as usable, not only that export completed. A file that exports but cannot be opened in a standard tool is a finding.'

E4 - Fully Effective: E4 requires the full end-to-end export test (all types, backups, IaC, logs), import validation in an alternative environment, and independent review of portability controls. As at M5, the differentiator from E3 is import, not export: 'E4 requires import validation in a target environment. 'Export ran successfully' without import validation at most E3.

What this means for C2 in practice.

Two design features recur throughout C2 and distinguish it from C1. First, the IaC coverage percentage (50% at M3, 80% at M4) is the only quantitative threshold in the framework that is paired with an explicit instruction to verify it ('spot-check by auditor recommended') rather than accept it as reported. This is what section 4.4.4 generalizes into a design principle for quantitative indicators across the framework. Second, the export/import distinction recurs at the top of both scales (M4 vs M5, E3 vs E4) in the same way the planned/unplanned trigger distinction recurs in C1: at M4/E3 the organization has demonstrated that data can leave the environment; at M5/E4 it has additionally demonstrated that the data is usable once it has left, a higher and qualitatively different bar than simply producing an export file.

4.3.3 Control area 3: Knowledge transfer and operational documentation

C3 verifies that the cloud environment is documented in sufficient detail for a competent but previously uninvolved third party to assume full operational control within the exit RTO, without requiring knowledge transfer from current personnel. C1 and C2 address contractual and technical lock-in; C3 addresses a third, often overlooked form: lock-in through people. Appendix B states the risk directly: 'tacit knowledge held by a small number of key employees, or embedded in provider-specific tooling, creates human lock-in as dangerous technical lock-in. When these people leave, or the provider's specialist knowledge is withdrawn at termination, the organization loses operational capability regardless of data portability. The sovereignty link is DICTU criterion 3.4 [15], especially levels 4 and 5, which describe documentation 'sufficiently detailed for a third party to operate the service independently' [15]; EuroStack criterion 2.3 [23]; and DICTU §5.5.2 on capacity and skills.

Audit questions. C3 deliberately combines two evidentiary strands that address the same underlying risk from different angles: written documentation (architecture records, runbooks) and the concentration of undocumented knowledge in specific individuals. Questions 3.1, 3.2, and 3.4 test the written strand. Current architecture documentation covering topology, data flows, and security (3.1); infrastructure configuration that is substantially automated and verifiably functional rather than manually maintained (3.2); and runbooks detailed enough for an unfamiliar engineer to follow (3.4). Question 3.3 tests the human strand directly: whether a key-person dependency analysis exists and whether identified dependencies have a remediation plan. These two strands are kept within a single control area, rather than split into a documentation control and a separate succession-planning control, because the underlying risk (that operational knowledge of the cloud environment is not actually transferable) can be produced by either gap on its own: an organization with an excellent runbook but two irreplaceable engineers is no more reversible than one with well-distributed knowledge and no written documentation. Question

3.5 then tests both strands together: it requires a party external to the current team and explicitly not an auditor of the documentation to execute a procedure using only what 3.1-3.4 produced. This fifth question is what DICTU's source definition for this criterion calls the difference between documentation that is 'sufficiently detailed' and documentation that is 'sufficiently detailed and tested' [15].

What must be shown for maturity level (Part A).

M1 - No Documentation: M1 applies if 'no current architecture documentation exists,' 'operations depend entirely on knowledge held by one or two individuals or on provider specialist assistance, ' and 'IaC coverage is below 20%; most infrastructure is manually configured.' As with C1 and C2, the criterion is framed as an evidentiary absence: 'if the organization cannot produce an architecture document reviewed within 36 months, M1 applies.'

M2 - Limited Documentation: M2 requires that 'some documentation exists but is significantly outdated (not reviewed in 24 or more months) or covers only parts of the environment,' that 'multiple key person dependencies exist without a remediation plan,' and that 'IaC coverage is between 20 and 50%'. The evidence at this level is described as 'key person analysis showing unaddressed single-person dependencies.' Meaning that at M2, the organization may already know where its dependencies are but has not acted on that knowledge.

M3 - Basic Documentation: At M3, the 'architecture document is current (reviewed within 24 months) and covers main components,' 'basic operational run books exist for primary procedures,' 'IaC coverage exceeds 50%', and 'key persons have been identified, and at least partial cross-training is underway.' M3 therefore represents documentation that is current and broadly scoped but not yet comprehensive. 'Main components' and 'primary procedures,' not the full operational surface.

M4 - Comprehensive Documentation: M4 raises the currency to 'reviewed within 12 months' and the scope to 'comprehensive, covering all components including integrations, security controls, and recovery procedures.' IaC coverage must 'exceed 80%', and (as in C2) the evidence requirement specified that a 'recent deployment from IaC confirms it is functional,' i.e., coverage is demonstrated by use, not asserted by percentage alone. Run books must cover 'all routine and non-routine procedures,' and 'no unaddressed critical single-person dependencies remain.' M4 is the level at which C3's three sub-concerns (architecture documentation, IaC, and key-person risk) are each individually resolved.

M5 - Tested Documentation: M5 adds a single, qualitatively different requirement on top of M4: 'documentation tested by a party outside the current team: a competent engineer unfamiliar with the environment successfully executed at least one deployment or operational procedure using only the available documentation.' The evidence specifications are unusually precise about what 'outside the current team' means: 'tester name and role (must not be an author), data, procedure executed, outcome,' and confirmation that the tester had no prior knowledge of the specific environment.' Nothing in M1-M4 is undone if this test is absent, but M5 cannot be reached without it, regardless of how thorough the M4 documentation is judged to be.

What must be shown for operating effectiveness (Part B) E1 - Not Effective: E1 applies where architecture documentation is 'significantly outdated or absent,' key-person dependencies are 'unaddressed,' and IaC is 'below 50%'. The audit procedure includes a specific practical test: 'ask a senior team member who is not the primary cloud engineer to locate and describe the data flow for a named application. Inability to do so without the primary engineer is a clear E1 indicator.' This operationalizes the key-person risk as something the auditor can observe directly during fieldwork, not only infer from an organization chart.

E2 - Partially Effective: E2 requires documentation to be current, IaC present 'with known coverage,' and key persons identified with remediation in progress. The guidance repeats the verification point from C2 in stronger terms: 'IaC coverage claims must be verified by actually opening the repository and checking, not accepted on assertion.

E3 - Largely Effective: E3 requires IaC above 80% confirmed (by sampling five components against the infrastructure repository), run books covering all criteria procedures (verified by selecting three non-routine procedures and checking they exist, are complete, and are current), and no remaining critical single-person dependencies. The audit procedure includes a test of documentation completeness that foreshadows the M5/E4 independence test: 'select one infrastructure component and walk through the IaC deployment with a team member who is not the author. Can they execute it? If not, E2 at best.'

E4 - Fully Effective: E4 requires the documentation test from M5 to have been performed and recorded, with the deployment 'successfully executed using only documentation' and the tester's lack of prior knowledge confirmed. The guidance names this as 'the defining criterion for E4: the person who executed the test must not be the person who wrote the documentation. Verify this explicitly.'

What this means for C3 in practice. C3 is structured around three sub-concerns (architecture documentation, IaC coverage, and key-person dependency) that progress largely in parallel through M1-M4 (each addition at a given level applies to all three) and then coverage on a single criterion at M5/E4: an independence test performed by someone who is neither a key person nor a documentation author. This is the most concrete instantiation, across the five control areas, of the 'M5 = independent verification' pattern introduced in section 4.2.3: where C1's M5 requires a person, and the audit evidence specifically records who that person is and confirms they are not the author, which is the most direct operationalization of 'tested' across the entire framework.

4.3.4 Control area 4: Continuity of operations during migration

C4 verifies that the exit plan contains a migration continuity plan addressing migration phases, parallel-operation periods, data synchronization, service degradation tolerance, a rollback procedure with trigger criteria, and stakeholder communication. C1-C3 establish that an exit is contractually possible (C1), technically possible (C2), and operable by a third party (C3); C4 addresses what happens during the exit itself. Appendix B frames the risk as a second-order one: 'even when the exit plan, portability, and documentation are all adequate, the transition period is a high-risk window: services are partially operational, data is split across old and new environments, and teams must manage both simultaneously. Without a migration continuity plan, exit can trigger a service availability crisis that forces the organization to abort the migration.' In other words, C4 exists because an organization can score well on C1-C3 and still be unable to execute an exit safely. The sovereignty link is DICTU domain 4 (Operational) [15], NIS2 Art. 21 (business continuity) [19], DORA Art.11 (ICT continuity policy) [18], and ISO 22301 §8.5 [37].

Audit questions. C4's five questions move through the migration window in the order an organization would actually experience it, rather than testing five unrelated properties. Question 4.1 established the structural skeleton (phases, parallel operation, and phase-question,-gate criteria) without which the remaining four questions would have nothing to attach to. Question 4.2 then tests the mechanism that keeps the two environments consistent while both are live (synchronization and integrity), and question 4.3 tests the safety mechanism for when that consistency breaks down: a rollback procedure with explicit triggers, not merely an intention to roll back if needed. Question 4.4 addresses a risk that is easy to overlook once rollback is in place (that the service may legitimately degrade for a bounded period even when nothing has gone wrong) and ties that tolerance to existing SLA, NIS2, or DORA obligations so it is not set arbitrarily. Question 4.5 closes the set with the stakeholder dimension: a migration itself remaining unmanaged. ISO 22301 frames the underlying concern as maintaining continuity 'during' disruption, not only planning for it in advance [37].

What must be shown at each maturity level (Part A)

M1 - No Continuity Plan: M1 applies if 'the exit plan contains no continuity section for the migration period,' 'no rollback procedure is defined,' and 'service degradation tolerance is undefined.' The evidence requirement is again an absence: 'absence of a continuity section in the exit plan confirms M1 for this control.' Note that M1 here can co-exist with a strong C1 score; an organization can have an approved exit plan (C1) that simply does not address the migration period at all.

M2 - Continuity Acknowledged: M2 requires that 'migration continuity is acknowledged as a risk in the exit plan, but no detailed plan exists,' that 'rollback is mentioned but not procedurally defined (no trigger criteria, no execution steps),' and that 'service degradation is not formally assessed.' The evidence is a 'high-level rollback mention without procedure details.' M2 captures the situation where the organization has written down the continuity requirements without having worked out what it would actually do.

M3 - Basic Plan Defined: M3 requires that 'migration phases are defined with a phase diagram,' a 'parallel operation period is specified,' the 'data synchronization protocol is described at a high-level,' and the 'rollback procedure is documented with trigger criteria, even if not yet tested.' M3 is the point at which the plan becomes procedurally concrete (phases, a diagram, and named trigger criteria), but as the explicit qualifier states, without any test history.

M4 - Detailed and Validated Plan: M4 requires the synchronization protocol to specify 'frequency, integrity check method, dual write, or freeze period'; the service degradation tolerance to be 'formally defined per service and validated against SLAs'; the rollback RTO to be 'approved by CIO or equivalent'; and the communication plan to be 'completed, identifying all stakeholders, channels, and escalation paths, and reviewed by all named parties.' Each of the four high-level elements introduced at M3 (phases, sync, rollback/degradation/communication) is individually detailed and formally approved at 4, but as in M3, none of it has yet been exercised.

M5 - Tested Continuity Plan: M5 requires that the migration continuity plan has been 'validated in a structured exercise; (minimum: tabletop with an unplanned trigger scenario),' that the 'rollback procedure was explicitly tested in the exercise (rollback was activated as part of the scenario),' and that 'exercise findings are documented and remediated.' The evidence specification distinguishes between a scenario that merely mentions rollback and one that activates it: 'confirmation that rollback was activated and executed in the exercise' is listed as a separate evidence item from the exercise report itself. The same activation/discussion distinction seen at M5 in C1.

What must be shown for operating effectiveness (Part B)

E1 - Not Effective: E1 applies if there is 'no continuity section in exit plan,' 'no rollback procedure,' and 'service degradation tolerance not defined.' The guidance states the structural consequence directly: 'if the exit plan has no continuity section, rate E1 for Control 4. This automatically limits the overall conclusion to mostly partially effective.' An explicit cross-reference from the per-control rating to the aggregation logic of table 4.3.

E2 - Partially Effective: E2 requires the continuity plan to be reviewed within 24 months, the rollback procedure to be documented, and an SLA cross-check to be performed. As with C1's E2, the guidance frames this as the modal outcome: 'plan exists and is maintained but never tested. The most common state. Priority recommendations: schedule a tabletop exercise with an unplanned scenario.'

E3 - Largely Effective: E3 requires a tabletop exercise with an unplanned exit trigger in which 'rollback was explicitly tested, not merely discussed,' and findings were remediated. The guidance draws the activation/discussion line explicitly: 'an exercise were rollback was discussed but not executed is E2.'

E4 - Fully Effective: E4 requires a 'partial migration test conducted with service continuity maintained throughout,' rollback 'executed and confirmed functional' with 'service resuming without data loss, ' and the communication plan 'tested with actual stakeholders.' The guidance is unusually direct about the gap between E3 and E4: E4 requires actual partial migration execution. A tabletop simulation, however thorough, does not constitute E4 for this control.

What this means for C4 in practice. C4 is the control area where the framework states most explicitly that a single missing element (the continuity section) caps the entire audit conclusion, independent of M_{avg} and E_{avg}. This is the textual basis for the discussion in section 4.2.5 of why Guideline 3 in table 4.3 is necessary in addition to an average. C4 also reproduces, for rollback specifically, the activation-versus-discussion distinction that recurs throughout the framework's M5/E4 levels: an exit plan can name rollback, document its criteria, and have those criteria approved by the CTO (M4), and a tabletop exercise can even mention rollback as part of its scenario (E2), but only an exercise in which rollback was actually activated or a partial migration in which it was actually executed reaches M5/E4.

4.3.5 Control area 5: Testing programme and periodic evaluation

C5 verifies that exit capability is subject to a structured, periodic test programme, and that findings from testing are tracked and reported to senior management. Where C1-C4 each assess one part of an exit capability (plan, data, documentation, and continuity arrangements), C5 assesses whether that capability, as a whole, has actually been exercised. Appendix B states the risk in terms of assurance rather than operational harm: 'exit plans and documentation that are created but never tested provide false assurance. Technology, provider terms, and organizational context change continuously. The DICTU level 5 explicitly requires documentation that is 'tested' [15]. Without periodic testing, the organization cannot demonstrate that its exit capability is real rather than theoretical. This framing (C5 as the control that converts the other four from 'designed' to 'demonstrated') is the basis for its gating role in Guideline 3 table 4.3 and is discussed further in section 4.4.2. The sovereignty link is the DICTU criterion 3.4, level 5 ('sufficiently detailed and tested') [15], EuroStack criterion 2.3 (tested deployment capability) [23], and ISO 22301 §8.5 (BCM test programme) [37].

Audit questions. C5 combines governance-oriented evidence and technical-verification evidence within a single control area, because the framework treats 'has this been tested' as a single question that cannot be answered by either type of evidence alone. Question 5.1 establishes that testing is programmatic rather than incidental (a documented programme with defined test types, frequency, and success criteria) before any individual test is examined. Questions 5.2 and 5.3 then test the two specific exercise types the rest of the framework relies on: the tabletop exercise referenced at M5 in C1 and C4 (5.2, with its 'unplanned exit trigger' requirement deliberately mirroring the M5 wording used elsewhere), and the data export test referenced at M5 in C2 (5.3, again requiring import validation, not only export). Rather than re-deriving these tests separately for each control area, C5 collects the evidence that they occurred at all and were repeated. Question 5.4 adds a dimension that the other four questions do not cover (whether the programme reacts to material change, not only to the calendar), and question 5.5 closes the set with an accountability mechanism (a tracked findings register reported to senior management) without which a test programme can run indefinitely without producing remediation. DICTU's own definition of its highest maturity level requires documentation that is 'sufficiently detailed and tested' [15]; C5 is the control area that applies the evidence for the second half of that requirement across the framework as a whole.

What must be shown at each maturity level (Part A)

M1 - No Testing: M1 applies if 'exit capability has never been tested in any form,' 'no test programme exists,' and 'no findings register is maintained.' Across all five control areas, the criterion is framed as an absence: 'absence of any test records confirms M1. If no evidence is produced, this level applies by default.'

M2 - Ad Hoc Testing: M2 again requires, first, that the organization has progressed beyond M1 (the M1 absence conditions no longer apply), and specifically, that 'some informal testing or discussion of exit scenarios has occurred,' but 'no structured programme, no documented findings, no remediation tracking.' The evidence is described as 'email records or meeting notes mentioning exit scenario discussion,' with the explicit caveat that 'informal discussion without documented outcomes confirms M2 at most.'

M3 - Annual Tabletop: M3 requires that 'a formal test programme document exist specifying test types and frequency,' that an 'annual tabletop exercise was conducted within the past 12 months' whose 'scenario' included at least one unplanned exit 'trigger,' and that 'exercise findings are documented (data export test not yet conducted).' The bracketed qualifier is part of the criterion itself: M3 is explicitly defined as the level at which the tabletop has happened, but the export test ('C2's empirical equivalent) has not.

M4 - Full Test Programme: M4 requires a 'data export test conducted within 24 months with import validation' (closing the gap left open at M3), 'test findings tracked in findings register with owners and due dates,' findings 'reported to senior management at least annually,' and a 'trigger-based review procedure in place.' M4 is therefore the level at which both major test types referenced elsewhere in the framework (the tabletop C1, C4, C5 themselves) and the export test (C2)) have each been performed at least once, and the resulting findings governed (tracked, owned, and reported).

M5 - Comprehensively Tested: M5 adds a third test type on top of M4's tabletop-plus-export-test: 'partial migration test (or equivalent highest-level exercise) conducted in addition to the tabletop and export test.' In further requires that 'all prior-period findings are closed or have an accepted risk.' M5 for C5 is thus the point at which the framework's own three escalating test types (tabletop, export test, and partial migration) have all been completed at least once, and an external party (not the team being assessed) has reviewed the programme itself.

What must be shown for operating effectiveness (Part B)

E1 - Not Effective: E1 applies if there is 'no evidence of any exit-related testing in the past 24 months,' 'no test programme document, ' and 'no findings register.' The guidance is unconditional: 'if no testing has been conducted in the past 24 months. E1 applies regardless of how the design maturity is scored.' An explicit statement that Part A and Part B can diverge sharply for C5 in particular.

E2 - Partially Effective: E2 requires an annual tabletop exercise conducted and findings documented, but 'remediation is incomplete or data export test not yet conducted.' The guidance adds a check on the tabletop's content, not just its occurrence: 'confirm the tabletop scenario included an unplanned trigger. A planned migration session labeled as a tabletop does not meet this requirement.

E3 - Largely Effective: E3 requires both the tabletop and the data export test to be completed within 24 months, with all findings from both closed and annual management reporting evidenced. The guidance explains why both are required: 'E3 requires both test types: tabletop for governance and decision-making and data export test for technical validation. One without the other is E2 at most.'

E4 - Fully Effective: E4 requires the full programme (tabletop, export test, and partial migration test), plus independent oversight and all prior findings closed. The guidance states the consequence of an untested migration directly: 'an organization that has never tested an actual migration cannot claim E4 regardless of how thorough its tabletop exercises are.'

What this means for C5 in practice

C5 is structured around three escalating test types: tabletop (governance-oriented, tests decision-making and documentation), data export test (technical, tests C2's portability claims empirically), and partial migration test (operational, tests the actual execution of an exit step). Each of which also appears, in a control-specific form, as the M5/E4 criterion of C1, C2, and C4, respectively. C5's own M4/M5 and E3/E4 transitions are therefore best read as a summary checklist of whether those other controls' top-level tests have actually occurred across the framework as a whole. This is also why C5's gating role in Guideline 3 table 4.3 is structurally different from C1's role in Guidelines 1 and 2: C1 gates because it is a logical precondition for the other controls to be meaningful section 4.2.5, whereas C5 gates because it is the control area in which the evidence for 'has this actually been tested' across the whole framework is collected and reported to senior management.

4.4 Design rationale

4.4.1 Scoring logic within a control area: cumulative maturity ladder

Design Problem. Within a single control area, how should the five maturity levels (M1-M5) relate to one another? The criteria attached to each level could be designed as independent sets (a control could, in principle, satisfy the M4 criteria without satisfying the M2 criteria) or as a cumulative ladder in which each level presupposes all criteria of the levels below it.

Alternative 1 (independent (non-cumulative) criteria per level. Each maturity level would be defined by its own self-contained set of criteria, and the auditor would select the highest level whose criteria are fully met, irrespective of whether lower-level criteria are also met. This alternative was rejected because operational reversibility is, by construction, an additive capability: the DICTU definitions themselves describe each level as a refinement of the level below (e.g., 'active reversibility' presupposes 'contractual reversibility') [15]. Independent criteria sets would permit logically inconsistent outcomes that do not correspond to any meaningful real-world configuration and would undermine the audit conclusion's face validity. For example, a service satisfying the documentation-testing criterion of M5 while failing the basic exit-rights criterion of M3.

Alternative 2: weighted scoring model (CMMI-style capability levels with weighted sub-practices). Following the Capability Maturity Model Integration (CMMI) convention [8], each maturity level could be reached through a weighted combination of sub-practices, where partial satisfaction of higher-level practices could compensate for incomplete satisfaction of lower-level ones, producing a continuous (e.g., 0-100) score subsequently mapped onto M1-M5 bands [8]. This alternative was rejected for two reasons. First, it requires assigning numerical weights to qualitatively heterogeneous criteria (e.g., 'an exit plan exists' versus 'the exit plan has been reviewed within 24 months'), and the weighting itself would constitute an additional, empirically unvalidated design decision. Validating both the maturity criteria and a weighting scheme simultaneously would be required. a disproportionate research effort: the weights themselves would need to be derived empirically, for example, through a Delphi study with IT audit experts, before the weighted instrument could be meaningfully evaluated for utility. This is a separate research programme rather than a design decision that can be made within a single study. Second, a continuous score reintroduces the compensability problem of Alternative 1 in a less transparent form: an auditor reading 'score 72/100, mapped to M4' cannot inspect which underlying criteria drove that score without consulting the weighting table separately.

Alternative 3 (adit question completing percentage). Maturity could be derived directly from the proportion of 'yes' answers to five audit questions per control area (e.g., 3 out of 5 = 60% = some intermediate level). This alternative was rejected because the audit questions are designed to establish

the existence of evidence, not its currency, completeness, or verification states. Three 'yes' answers to questions 1.1-1.3 (a plan exists, exit rights exist, RTO defined) describe a materially different situation depending on whether the plan was reviewed last month or four years ago. A percentage-of-question metric would discard exactly the temporal and verification dimensions that the cumulative criteria in Appendix B are designed to capture.

Chosen design and justification. The cumulative ladder (section 4.2.3, instantiated per control area in section 4.3) was retained. Each maturity level's criteria explicitly include 'all criteria of the preceding level, plus' the new requirements for that level (Appendix B). This design satisfies DR2 (graded but logically consistent assessment), reflects the DICTU source definitions without requiring an additional translation layer (DR5), and keeps the relationship between a maturity level and the underlying evidence directly inspectable. An auditor disputing an M3 conclusion can point to the specific unmet M4 criterion quoted in section 4.3, which is a property the weighted and percentage-based alternatives do not offer.

4.4.2 Aggregation across control areas: average with weakest-link signaling

Design problem. Once each of the five control areas has an M-score and an E-score, how should these five-by-two scores be combined into a single overall audit conclusion? The earliest version of the framework, developed before expert validation, derived a binary 'Effective'/'Not Effective' verdict directly per control area and required all five control areas to be 'Effective' for an overall positive conclusion.

Alternative 1 (binary per control verdict, unanimous requirement (original design)). Under this design, each control area would be classified as 'Effective' or 'Not Effective' (collapsing the M1-M5 and E1-E4 scales into a single binary judgment per control), and the overall conclusion would require all five to be 'Effective.' This is the framework's initial design but was rejected after reflection on the professional judgment principle articulated in DICTU v.1.0.1 (2026, p. 2) [15]: collapsing a five-point maturity scale into a binary verdict at the control level discards exactly the graded information that the M1-M5 and E1-E4 scales were designed to capture and does so before that information reaches the overall conclusion. The point at which it is most useful for prioritizing remediation.

Alternative 2 (unweighted average without gating). The overall conclusion could be derived purely from M_{avg} and E_{avg} (the arithmetic mean across the five controls), without any additional rules. This alternative was rejected because it permits full compensation between control areas: an organization could score M5 on Controls 2-5 and M1 on Control 1 ('no exit plan at all') and still obtain an average of M4, despite having, by definition, no verified path to exit at all. Given that C1 (exit plan) is a logical precondition for the other four controls being exercised in a real exit (an organization cannot meaningfully test a migration continuity plan (C4) or a testing programme (C5) for an exit it has not planned). An unweighted average misrepresents the dependency structure of the underlying risk.

Alternative 3 (fixed per control weights (weighted average)). A weighted average could assign, for example, a higher weight to C1 and C5 than C2-C4, reflecting their gating role, while still producing a single continuous score. This was considered a refinement of Alternative 2 but rejected for the same reason as the weighted scoring model in section 4.4.1: assigning specific numerical weights (e.g., C1 = 30%, others = 17.5% each) is itself a design decision requiring empirical justification that was outside the scope of validating the underlying maturity criteria. A weighted average still permits compensation in degree, only a different exchange rate, which does not resolve the underlying objection that a near-absent exit plan (C1-M1) is not a 'partial' problem that can be offset by strength elsewhere.

Chosen design and justification. The framework combines an unweighted average (M_{avg}, E_{avg}), which preserves the graded information requested in DR3 and provides a single comparable figure for trend reporting across repeated audits, with the small set of weakest-link 'Decision Guidelines' in table 4.3, which encode the gating role of C1 and C5 without requiring a numerical weighting scheme. This design was explicitly revised twice during the research process. First, the original binary per-control design (Alternative 1) was replaced with the average-plus-gating design in response to the professional-judgment principle in the DICTU source material; this revised design (not the original binary one) is what was presented to, and discussed with, the seven experts (R01-R07) in the first round of evaluation reported in chapter 5. Second, in direct response to feedback from that first round, the five conditions were restated from fixed 'Decision Rules' into 'Decision Guidelines' using explicitly non-mechanical wording ('generally,' 'usually') to make the requirement for professional judgment, rather than mechanical rule-following, unambiguous in the artifact's own language. This second revision is part of the Version

2 artifact documented throughout this chapter and is the version submitted to the second, confirmatory round of evaluation, also reported in chapter 5.

4.4.3 Decomposition into control areas: five distinct controls

Design problem. The DICTU criterion 3.4 ('Operational Reversibility') is, in its source form, a single criterion with five generic maturity levels [15]. The design problem was whether to operationalize this single criterion as one composite control or to decompose it into multiple control areas, and if the latter, how many and along which boundaries.

Alternative 1 (a single composite control for criterion 3.4). DICTU 3.4 could be operationalized directly as one control, with a single set of audit questions and a single M1-M5 / E1-E4 assessment. This alternative was rejected because the source definitions for 3.4 reference at least three analytically distinct evidentiary domains (contractual exit rights, technical data portability, and operation documentation), each governed by different evidence types (legal/contractual documents, technical export tests, and architecture and run book documentation) and, in practice, falling within the remit of different organizational functions (legal/procurement, infrastructure engineering, and operations) [15]. A single composite control would either require an auditor to assess all of them simultaneously against one undifferentiated maturity scale (reducing the diagnostic value of the resulting finding) or would implicitly decompose during fieldwork in any case, in which case the decomposition should be made explicit in the design rather than left to ad hoc practice (in tension with DR5).

Alternative 2 (four control areas (merging exit planning and migration continuity)). Controls 1 (Exit Plan and Contractual Exit Framework) and 4 (Continuity of Operations During Migration) could be merged into a single 'Exit and Continuity Planning' control on the grounds that both concern planning artifacts produced before an exit occurs. This was considered but rejected because the underlying risks differ in kind and timing: C1 concerns whether an exit is contractually and strategically possible at all (a precondition that can be assessed independently of any specific migration), whereas C4 concerns whether service continuity can be maintained during the execution of a migration that has already been deemed possible. Merging them would obscure a plausible case observed during the design process and discussed in chapter 5: an organization with a strong contractual exit plan (C1) but weak continuity during transition (C4). The Decision Guidelines in table 4.3 are designed to highlight this distinction.

Alternative 3 (six control areas (splitting documentation from key-person dependency)). Control 3 (Knowledge Transfer and Operational Documentation) addresses both written documentation (architecture documents and run books) and key-person dependency (concentration of tactical knowledge in specific individuals). These could be split into two separate control areas, since an organization could, in principle, have excellent documentation while still being critically dependent on one or two individuals, or vice versa. This was considered but rejected on the grounds of evidentiary overlap rather than conceptual distinctness: the evidence used to assess key-person dependency (cross-training records, run book usage, and the independence test at M5) is the same evidence used to assess documentation quality, and the M1-M5 cumulative criteria for C3 (section 4.3.3) already treat key-person remediation as one of several jointly-required criteria at each level rather than as an independent axis. Splitting the control would have introduced a sixth control area whose audit questions would substantially duplicate those of the fifth, without a corresponding gain in diagnostic resolution proportional to the added length of the audit instrument. A consideration directly relevant to the practical usability concerns (perceived ease of use) examined through TAM [13] in chapter 5.

Chosen design and justification. The five-control-area decomposition (C1-C5, section 4.3) was retained as the version presented for expert validation. The boundary between C1 and C4 separates 'is an exit possible' from 'can continuity be maintained while existing,' and the boundary around C5 separates 'has the capability been designed' (C1-C4) from 'has the capability been demonstrated' (C5), a distinction that the aggregation logic in section 4.4.2 depends on directly. Five control areas, each with five audit questions, also produce a 25-question instrument of a length considered, during design, to be compatible with a single audit fieldwork session. A practical constraint that is examined empirically through the perceived ease of use construct in chapter 5.

4.4.4 Quantitative indicators: IaC coverage as a verified, principle-based criterion

Design problem. Several DICTU and EuroStack source criteria reference Infrastructure-as-Code (IaC) coverage as an indicator of technical portability and documentation maturity [15] [23]. The design problem was how to incorporate a quantitative indicator of this kind into a framework whose primary mode of operation is qualitative, evidence-based professional judgment.

Alternative 1 (IaC coverage as a hard pass/fail gate). A coverage percentage (e.g., 80%) could function as a binary gate: below the threshold, the relevant maturity level is unreachable regardless of other evidence; at or above it, the criterion is satisfied. This alternative was rejected for two reasons. First, coverage percentages reported by organizations are frequently self-reported and difficult to verify independently within a single audit engagement, so a hard gate based on an unverified figure would create a false sense of precision (DR4). Second, a hard numeric gate conflicts with the professional judgment principle in DICTU v.1.0.1 (2026, p. 2) [15] and risks producing 'cliff-edge' findings (a result of 79% being treated as categorically different from 81%) that do not correspond to a meaningful difference in actual reversibility risk.

Alternative 2 (no quantitative indicator; fully qualitative assessment). IaC coverage could be omitted entirely, with the relevant maturity criteria relying only on qualitative descriptions (e.g., 'most infrastructure is defined as code'). This option was rejected because it eliminates the concrete, source-grounded indicator that provides auditors a tangible starting point for fieldwork. Furthermore, experience during the design process indicated that purely qualitative formulations of this criterion are more prone to organizations overstating their position than formulations that reference a specific, directly inspectable artifact, namely the IaC repository.

Chosen design and justification. The 50% (M3) and 80% (M4) coverage figures (section 4.3.2) were retained as one of several jointly-required cumulative criteria at each level, paired explicitly with a verification procedure: the auditor guidance in Appendix B states that coverage claims must be checked against the repository directly (e.g., by spot-checking a sample of components), and that an unverified percentage does not satisfy this criterion. This design treats the percentage as an orientation point that anchors the auditor's judgment to a recognizable, source-grounded benchmark (satisfying DR4's verifiability requirement) without making the overall maturity level a direct, automatic function of that single number (avoiding the cliff-edge problem of Alternative 1). The M3/M4 criteria walkthroughs in section 4.3.2 and section 4.3.3 illustrate the intended operations of this design choice.

The four design choices discussed in this section were finalized before the first round of expert evaluation. Where the interviews produced evidence bearing on these specific choices, for example, on the practicality of the verification procedures referenced in section 4.4.4, or on the reported perceived appropriateness of the weakest-link guidance in section 4.4.2. This is reported in the first-round evaluation in chapter 5 and is reflected directly in the Version 2 artifact documented throughout this chapter. Chapter 5 then reports the second, confirmatory round of evaluation of this Version 2 artifact.

Chapter 5

Evaluation

Consistent with the build and evaluate cycle central to DSR [31, 44], the DSAF artifact described in chapter 4 was evaluated across two rounds. The first round involved seven IT audit and Technology Risk professionals from a single Big Four firm, coded R01 to R07, with professional experience ranging from several years to over two decades in IT audit, Technology Risk, and cloud governance roles. The second, confirmatory round was conducted against the resulting Version 2 artifact and involved two additional respondents, coded R08 and R09. R08 is a manager who had also participated in the first round, which allows a direct within-respondent comparison between the two versions of the artifact. R09 is a partner who had not previously been exposed to the framework, which allows the second round to also capture an independent first impression rather than only a confirmation of remembered improvements. Both rounds used a sequential explanatory mixed-methods design [12]. A structured questionnaire based on the Technology Acceptance Model, hereafter TAM [13, 56], was administered first, followed immediately by a semi-structured interview in which respondents explained the reasoning behind their scores.

The questionnaire covered four TAM constructs with twelve items in total, namely Perceived Usefulness with four items, Perceived Ease of Use with three items, Social Norm with three items, and Behavioral Intention with two items, each rated on a five-point scale. All interviews were conducted and transcribed in Dutch and coded in Atlas.ti, after which the transcripts were analyzed using Braun and Clarke's reflexive thematic analysis [6]. Every quotation reproduced in this chapter is therefore a translation from the original Dutch transcript into English, prepared for this thesis and presented alongside the source language reasoning that produced it. The full Atlas.ti code list is provided in Appendix C, and every quotation used in this chapter is cross-referenced there by respondent code and code label.

This chapter is structured to move from quantitative results to the qualitative themes that explain them, to the specific problems and design responses those themes motivated, and finally to a second, confirmatory evaluation round conducted against the resulting Version 2 artifact (documented in full in chapter 4). Section 5.1 interprets the Round 1 TAM results. Section 5.2 reports the five themes produced by Round 1 thematic analysis. Section 5.3 summarizes the concrete problems identified in Round 1. Section 5.4 reports the concrete improvements made in response. Section 5.5 reports the results of the second evaluation round. Section 5.6 compares the two rounds directly. Section 5.7 concludes the chapter and section 5.8 discusses its limitations.

5.1 Results: TAM (Round 1)

This section reports the quantitative results from the twelve-item TAM questionnaire administered in Round 1. Results are reported per construct, followed by an interpretation of the patterns across constructs taken together.

5.1.1 Results by TAM construct

Table 5.1 reports the mean and standard deviation for all twelve TAM items across the seven Round 1 respondents. Scores of four and above are treated as broadly positive, a score of approximately three as

moderate, and scores below three, or items combining a moderate mean with a higher standard deviation, are treated as signaling a genuine concern that warrants qualitative investigation.

Table 5.1: Overview of TAM Statements, Respondent Ratings, and Averages (n=7). Cell shading indicates the interpretive category of the average score: green (AVG $>$ 4.0) indicates broad positive agreement; orange (AVG 3.0-3.9) indicates a moderate or mixed response that warrants qualitative investigation; red (AVG $<$ 3.0) indicates a predominantly negative response. Color coding applies to the AVG column only.

#	**Statement**	**R01**	**R02**	**R03**	**R04**	**R05**	**R06**	**R07**	**AVG**	**STDEV**
PU1	This framework facilitates the identification of cloud exit risks that may otherwise be overlooked or insufficiently emphasized in a standard audit context.	4	4	4	4	4	4	4	4.00	0.00
PU2	The use of this framework would enhance the quality and completeness of my audit conclusions regarding vendor lock-in and exit capability.	5	5	5	4	4	5	3.5	4.67	0.52
PU3	This framework provides sufficient structure to develop a defensible, evidence-based audit judgment on operational reversibility.	2	2	5	3	4	5	4	3.57	1.27
PU4	This framework adds value beyond what existing audit standards, such as ISO 27001 or NIS2, already require me to assess.	4	3	4	3	3	4	4	3.57	0.53
PEOU1	The audit questions in this framework are clearly and unambiguously formulated, enabling me to understand precisely what is required to be verified.	3.5	3	4	4	5	4	2	3.67	1.03
PEOU2	The distinction between design maturity (M1–M5) and operational effectiveness (E1–E4) is easy to understand and apply.	4.5	2	3	4	5	4	3	3.50	1.05
PEOU3	The evidence requirements defined at each maturity level are sufficiently specific to clearly indicate what should be requested from a client organization.	3	3	3	5	4	5	3	3.71	0.95
SN1	Individuals whose opinions I value within my professional environment would support the use of a structured sovereignty assessment framework such as this in cloud audits.	5	4	5	5	5	5	4	4.71	0.49
SN2	I believe that the use of this type of framework is aligned with what is expected of a professional IT auditor within the current regulatory environment (e.g., NIS2, DORA, GDPR).	3	2	3	3	4	4	4	3.29	0.76
SN3	If leading professional bodies or prominent organizations within my field were to adopt this framework, I would be more inclined to use it myself.	5	4	5	5	5	5	5	4.86	0.38
BI1	I would use this framework, or a refined version thereof, in a future cloud sovereignty or vendor lock-in audit engagement.	4	4	4	5	5	5	4	4.43	0.53
BI2	I would recommend this framework to a colleague conducting an assessment of cloud exit or operational reversibility risks.	4	4	5	5	5	5	4	4.57	0.53

5.1.2 Perceived Usefulness (PU)

PU1 produced the most unanimous score in the entire dataset. All seven respondents converged on a mean of 4.00 with zero variance, agreeing without qualification that the framework surfaces cloud exit risks that a standard audit would otherwise miss or insufficiently emphasize. PU2 scored even higher, with a mean of 4.67 and a standard deviation of 0.52, making it the second highest PU item and indicating broad agreement that the framework improves the completeness of audit conclusions specifically, rather then only their breadth. R01, for instance, separated the two judgments clearly by rating PU2 a full point above PU1, because the completeness of the conclusion was the more decisively improved property (code: framework quality of audit conclusions).

PU3 broke this pattern by producing the most divided score in the entire dataset. It was both the lowest-scoring PU item and, with a standard deviation of 1.27, the most divided item in the entire questionnaire. The spread is the more informative signal here, since it shows that respondents did not converge on a moderate view but instead split between two substantively different positions. R01 and R02 scored PU3 at 2, explicitly tying their score to a perceived confusion between the framework's substantive criteria and the procedural guidance needed to apply them (code: insufficient structure for defensible audit opinion, framework versus guidance confusion). R01 explained the rating directly: "that is why I find that the framework, in its current form, does not yet provide sufficient structure for a fully defensible audit judgment, and I give it a 2." R03, by contrast, scored the same item a 5, but on a materially narrower reading of the question. His explanation, "then I would give a 5 for that context," referred specifically to the framework applied to one single, already identified cloud provider, not to the general case that R01 and R02 have in mind (code: sufficient structure when applied to one specific cloud provider). The PU3 split is, therefore, not a simple disagreement about quality but evidence that the question was being answered against two different mental models of what sufficient structure requires.

PU4 scored similarly low to PU3 on average, at 3.57, but with a far lower standard deviation of 0.53, suggesting a consistently held rather than divided view. R03 articulated the underlying reasoning precisely: "ISO 27001 does not really address sovereignty. So I would say it does add value" (code: added value beyond ISO 27001 and NIS2). The value of the framework is acknowledged here, but several respondents qualified it as not yet sharply differentiated within the document itself.

- PU1 and PU2 together confirm that the framework identifies cloud exit risk in a way standard audits overlook and provides sufficient structure to translate those risks into a defensible audit conclusion, two properties that respondents agreed are absent from existing audit practice.
- PU3 is the most divided item in the dataset, reflecting two different readings of what sufficient structure requires, namely general applicability as opposed to a single, bounded engagement.
- PU4 is consistently moderate, reflecting real but not yet clearly differentiated added value over ISO 27001 and NIS2.

5.1.3 Perceived Ease of Use (PEOU)

PEOU produced the lowest mean of the four TAM constructs. PEOU2, the item asking whether the distinction between design maturity, referred to as Part A, and operating effectiveness, referred to as Part B, is easy to understand and apply, was its lowest and most divided item, with a mean of 3.50 and a standard deviation of 1.05. R02 scored this item a 2 and gave one of the sharpest critiques in the dataset: "I do not find it logical that you score the same control on two dimensions" (code: two dimensions on on control is illogical, creating false security in board reporting), a concern he linked explicitly to the risk that splitting one control into two scores could create a false sense of assurance once the result reaches the board level reporting. R03 scored the same item a 3, taking a narrower position. The distinction was visible in principle, but in his words, "you can see the difference, but it is not consistently applicable everywhere" (code: distinction visible but not consistently applicable), noting that it worked better for process-type control than for static, document-based ones.

Not every respondent shared this difficulty. R04 recognized the Part A and Part B split immediately as a structure already familiar from his own audit practice: "we also work with tests of design and tests of effectiveness. With design, you check whether the process exists, and with effectiveness, whether it actually works" (code: design versus operating effectiveness). This divergence, in which the distinction was confusing for R02 and R03 but immediately intuitive for R04, suggests that the difficulty lies not with the underlying design versus operation logic itself, which has a direct analogue in standard IT audit

methodology, but with how explicitly the framework states the logic itself before an auditor reaches the scoring step.

PEOU1 and PEOU3 scored moderately higher, at 3.67 and 3.71, respectively, but with comparable dispersion, at standard deviations of 1.03 and 0.95, suggesting that the uncertainty was not confined to the Part A and Part B split but extended more broadly to how the form communicates what is expected of the auditor before scoring begins. R03's assessment of PEOU1 captured this directly: "you do know what you need to do, but the coherence between elements needs improvement."

- PEOU2 is the construct's lowest and most divided item, splitting respondents between those who found Part A and Part B logic confusing and those, notably R04, who recognized it from existing design versus effectiveness audit practice.
- PEOU1 and PEOU2 are moderate and comparably dispersed, pointing to a broader clarity gap in how the form sequences what is expected of the auditor.

5.1.4 Social Norm (SN)

SN3 produced the highest score in the entire dataset, with a mean of 4.86 and a standard deviation of 0.38. Nearly every respondent indicated that formal adoption by a leading professional body would make them more inclined to use the framework themselves. SN1 followed closely, with a mean of 4.71 and a standard deviation of 0.49. R06 explained the underlying mechanism plainly: "it is simply much easier when a professional organization sets the norm for good practice. If they say this will be used, implementation becomes easier" (code: leading organizations adopting the framework would increase adoption likelihood). The mechanism described is institutional rather than purely reputational. In a professional environment as closely structured by recognized frameworks as a Dutch IT audit, an unendorsed instrument carries a higher adoption cost than an endorsed instrument does not, namely the need to individually justify why it is worth using.

SN2 is the construct's clear outlier and the lowest-scoring item of all twelve, with a mean of 3.29 and a standard deviation of 0.76. Respondents were considerably less convinced that using a framework of this kind is already expected of them under current regulation, as distinct from being something their professional peers would generally welcome. R02 captured the uncertainty concisely when asked whether current regulation already demands this kind of assessment, responding, "I don't know, because that usually isn't required" (code: current regulation does not yet require sovereignty assessment). R03 offered a more structured version of the same hesitation, scoring SN2 a 3 because an auditor must remain flexible because not every regulatory mandate specifically calls for this particular framework (code: auditor must remain flexible; not all regulatory requirements demand this framework).

- SN1 and SN3 are the two highest scores in the dataset, indicating strong professional and peer support for the framework and a clear, institutionally driven adoption pathway through formal endorsement.
- SN2 is the lowest overall, reflecting that current regulation, including NIS2, DORA, and GDPR, is not yet perceived as explicitly requiring this type of sovereignty assessment, independent of whether peers would support its use.

5.1.5 Behavioral Intention (BI)

Both BI items scored highly with identical, comparatively low dispersion, at 4.43 and 4.57, respectively, both with a standard deviation of 0.53, indicating that the structural and presentational concerns visible in PU3 and PEOU2 did not translate into weaker intention to use or recommend the framework. R03 scored BI2 points above BI1, explaining that he would recommend the framework specifically as reference and comparative material even where he might not yet apply it as a complete, stand-alone instrument in every engagement (code: would recommend as reference and comparative material), a more conditional but still clearly positive form of intention than a simple top score would suggest.

Read together with PEOU, this pattern points to a decoupling between perceived ease of use and behavioral intention that departs from the relationship TAM would predict in a purely individually driven adoption context [13, 56]. Behavioral intention remained high despite a moderate and divided PEOU construct. Section 5.2 returns to this decoupling in more depth, since it is one of the clearest threads connecting the quantitative results to the qualitative themes.

- BI1 and BI2 are consistently high and tightly clustered, showing that usability concerns did not weaken respondents' stated intention to use or recommend the framework.

5.2 Themes from the interviews

Reflexive thematic analysis of the seven Round 1 interviews produced five themes. Each is reported below together with its supporting Atlas.ti codes and illustrative quotations. As stated in section 5.0, every quotation has been translated from the original Dutch into English for presentation in this chapter; the full set of codes underlying each theme is provided in Appendix C.

5.2.1 Theme 1: The framework provides structure and improves audit quality

All seven respondents affirmed, without qualification, that operational reversibility is an audit-relevant concern poorly served by existing standards and that a structured instrument of this kind fills a genuine methodological gap. This theme is the qualitative counterpart to the high-low variance PU1 and PU2 scores reported in section 5.1.2 where respondents diverged not on whether the framework adds value but on how completely the current form realized that value in practice, which is the subject of Themes 2, 3, and 5 below.

5.2.2 Theme 2: Practical applicability determines adoption

Respondents consistently separated two judgments that a simple usefulness score would otherwise blend together, namely whether the framework's underlying logic is sound and whether its current presentation is immediately workable in a live engagement without prior familiarity. R02 expressed this distinction directly when discussing whether he would recommend the framework to a colleague, conditioning his answer on revisions rather than withholding it: "yes, if the tweaks are made" (code: would recommend to colleague conditional on tweaks). R04 offered the clearest single articulation of the pattern, treating the gap between liking the framework and being able to apply it unaided as a one-off onboarding cost rather than a structural flaw.

This theme also surfaced a constructive, practice-grounded suggestion for how operating effectiveness specifically should be tested. R04 pointed out that a full provider migration cannot reasonably be demanded as audit evidence, stating, "you cannot ask a client to just switch to another provider," and proposed bounded, realistic proxies instead: "you could ask the organization to show that its data can migrate, to run a test with a backup, and to demonstrate how long that takes" (code: testing exit capability, critique of practical feasibility). The same respondent argued more generally that the framework's guidance on how to test remains underdeveloped relative to its criteria on what to assess: "it could be more practical in how operating effectiveness is tested" (code: realistic evidence requirements).

5.2.3 Theme 3: The framework is valuable but requires methodological maturation

This theme captures the recurring conflation, visible quantitatively in PU3's high variance, between the framework's substantive criteria, namely how to assess it, which several respondents linked directly to inconsistent interpretation between auditors. A related but conceptually distinct strand within this theme concerns the framework's underlying maturity logic. R04 endorsed the cumulative structure of the maturity levels as directly consistent with existing audit reasoning: "yes, I find that logical. This resembles how we also work in audits. You look at what must minimally be present and keep building on that" (code: cumulative maturity model). The same respondent separately distinguished, unprompted, between an emergency recovery plan and a genuine exit plan: "an emergency plan is about recovering within the same environment, but an exit plan is about switching to a different provider, which is considerably more complex" (code: emergency plan versus exit plan, low maturity in practice), noting that most organizations he encounters have only the former.

A further, more conceptual concern in this theme was raised regarding the maturity scale's wording itself. DICTU's maturity levels are named after sovereignty outcomes, for example, 'contractual reversibility' or 'full operational reversibility.' Respondents expected process maturity language instead, such as 'identified,' 'managed,' or 'optimized.' This mismatch between outcome-oriented labels and process-oriented

expectations created interpretive friction during the evaluation. R04's own distinction between emergency and exit plans is itself evidence for this concern in practice. His own organization audit approach, by his own account, "currently looks at operational continuity but not at knowledge, contracts, and alternative providers" (code: gap in existing methodology, audit approach without framework), confirming that the gap the framework targets is real, while simultaneously illustrating why its terminology needs to map cleanly onto auditors' existing mental models in order to be applied consistently.

5.2.4 Theme 4: Digital sovereignty is gaining relevance, but adoption is shaped by external drivers

Respondents indicated that their willingness to adopt the framework depends primarily on external endorsement, such as recognition by a professional body like NOREA or an explicit regulatory expectation, rather than on the framework's own qualities. This is consistent with the high SN1 and SN3 scores combined with the comparatively low SN2 score reported in section 5.1.4. R01 situated this in a broader trajectory, observing that digital sovereignty is currently at a turning point, moving from a niche topic to something discussed more broadly in organizations in a manner comparable to how cybersecurity developed earlier, before linking adoption directly to formal recognition: "in practice we use many frameworks from organizations like NOREA. If such a framework is formally endorsed, that directly increases the likelihood it will actually be used" (code: leading organizations adopting a framework would increase adoption likelihood). R06's explanation of why endorsement matters so much functionally mirrors R01's reasoning, namely that formal recognition removes the burden of individually justifying the framework's value before it can be used.

5.2.5 Theme 5: Auditing digital sovereignty requires professional judgment, not rigid compliance criteria

Respondents valued structured assessment criteria but resisted evidence requirements that read as mandatory checklists, which directly explains the moderate PEOU construct mean in section 5.1.3. R02 articulated the underlying principle most directly: "as an auditor you have a norm, a standard. It does not matter how an organization meets that standard, whether through documentation, through processes, or through tooling. That may differ" (code: principle-based norm preferred over prescriptive evidence). R05 connected the same concern to a concrete risk of superficial compliance: "the difficulty with these kind of questions is that they quickly become a tick-box exercise. You ask whether they have an exit plan, and they say yes, but perhaps that is one page long and does not really mean much" (code: risk of checkbox compliance).

The underlying argument across this theme is not that structured criteria are unwelcome, since several respondents also valued specific evidence requirements, but that criteria should anchor professional judgment rather than substitute for it, which is consistent with established perspectives on the fundamental nature of judgment in audit practice [47].

5.3 Concrete problems identified (Round 1)

Drawing on the TAM results and the thematic findings together, five concrete problems were identified in the Version 1 artifact, summarized in table 5.2 together with their primary quantitative and qualitative evidence and the corresponding Atlas.ti codes, listed in full in Appendix C.In brief: the framework combined substantive criteria with procedural guidance, making consistent application difficult; the Part A and Part B distinction was not self-evident on first exposure; testing guidance for operating effectiveness was underdeveloped; the binary overall conclusion discarded the M and E scores; and the document-based format was incompatible with how practitioners structure their audit work.

Table 5.2: Overview of Identified Problems, Evidence, and Coding

#	**Problem**	**Evidence and codes**
1	The framework conflates assessment criteria and application guidance, which leads to inconsistent interpretation among auditors.	PU3 (mean 3.57, SD 1.27); Theme 3 (R01, R02, R04); codes: insufficient structure for defensible audit opinion; gap in existing methodology
2	The distinction between Part A (design maturity) and Part B (operating effectiveness) is not self-evident on first exposure to the form, despite mapping onto existing tests of design and tests of effectiveness audit practice for some respondents.	PEOU2 (mean 3.50, SD 1.05); Theme 3 (R02, R03, R04); codes: two dimensions on one control is illogical; design versus operational effectiveness
3	Guidance on how to test operating effectiveness in practice is underdeveloped relative to the criteria stating what should be assessed.	Theme 2 (R04); code: testing exit capability, critique of practical feasibility
4	A binary Effective or Not Effective conclusion collapses graded M1 to M5 and E1 to E4 evidence into a single judgment before that information reaches the overall conclusion.	Theme 3, raised in interview discussion of PU3 (R02); code: insufficient structure for defensible audit opinion
5	The instrument format (a multi-page PDF or Word document) does not integrate into standard audit workflow and requires manual re-entry of scores into any summary or audit file, creating transcription errors and reducing the likelihood of practical adoption.	Theme 2 and Theme 4, raised in interview discussion of (R01, R02, R03, R04, R05, R07, six of seven respondents); BI1 (mean 4.43); code: Excel format preferred, practical applicability determines adoption

5.4 Concrete improvements (Version 2)

Table 5.3 maps each problem in section 5.3 to the corresponding design adjustment made in Version 2 of the artifact, documented in full in chapter 4.

One problem identified in Round 1 does not have a corresponding design adjustment in Version 2: Problem 3, the underdeveloped testing guidance for operating effectiveness. R04's proposed proxies, specifically a data migration test and backup restoration test with documented timing against the RTO, were considered as candidate additions to the Part B evidence guidance but were not incorporated. This decision was deliberate: adding specific testing procedures requires empirical validation of whether those procedures produce reliable and consistent evidence across different organizational contexts, which fall outside the scope of a single design iteration. Problem 3 is therefore carried forward as an open item and is addressed as a future work recommendation in chapter 7.

Table 5.3: Overview of Problems and Version 2 Design Adjustments

#	**Problem**	**Design adjustment (Version 2)**
1	Conflation of criteria and procedural guidance.	A dedicated Step-by-Step Workflow was added to the Instruction tab, and a Definitions and glossary sheet was added, both within the existing workbook rather than as a separate guidance document.
2	Part A and Part B distinction is not self-evident.	Explicit workflow steps were added stating what each Part assesses, with an explicit statement that Part A and Part B scores are never averaged together.
3	Guidance on how to test operating effectiveness in practice is underdeveloped relative to the criteria stating what should be assessed.	No adjustments made
4	Binary concluding collapses graded evidence.	The overall conclusion logic was restated as the average Design Maturity and Operating Effectiveness across all five controls, combined with the weakest link Decision Guidelines rather than a binary control verdict. The underlying Decision Rules were also reframed as Decision Guidelines using non-mechanical language such as generally and usually.
5	Instrument format does not integrate into audit workflow.	The artifact was rebuilt as an Excel workbook with per-control dropdown selections, automatic score carry-through to the Summary tab, and a rationale and notes field per control area. This addresses both the workflow integration concern and the transcription-error risk identified in Round 1.
6	No indication of accountability within the client organization per maturity level (R01, R04).	Responsible Role tables were added to each control area, specifying which organizational role is accountable for evidence at each M and E level.
7	M2 criterion wording based on absence rather than progression, creating interpretive friction (R01, R06, R07).	The M2 criterion was reworded across all five control areas to state that 'the organization has progressed beyond M1 (the M1 absence conditions no longer apply)' rather than describing M2 through what is absent.

Not all improvements made in Version 2 are direct responses to the four numbered problems in table 5.2. Two additional changes were made in response to thematic signals that did not resolve into a single discrete problem but nonetheless carried sufficient evidential weight across multiple respondents to justify a design adjustment:

1. Evidence requirements throughout the form were reframed from a fixed list to 'Suggested Evidence to Provide,' responding to Theme 5 and the concern raised by R02, R06, and R07 that mandatory evidence lists risk producing checkbox compliance rather than professional judgment.
2. The M2 maturity criterion was reworded across all five control areas from an absence-based formulation to interpretive friction identified by R01 and R02.

5.5 Results: Round 2 (confirmatory evaluation)

The second round was conducted with two respondents against the Version 2 workbook described in chapter 4. R08 is a manager who had also participated in Round 1, allowing a direct within-respondent

comparison. R09 is a partner who had not been exposed to the framework before this session, providing an independent first impression. As in Round 1, the TAM questionnaire was administered first, and the semi-structured interview followed immediately afterward.

The fact that R09 had no prior exposure to the framework raises a methodological consideration that is worth acknowledging. First impressions of a new instrument may not accurately reflect the experience of an auditor who has completed a training session or applied the framework in a real engagement. Aspects that appear counterintuitive at first glance, such as the Part A and Part B distinction, may become natural after a brief orientation, as evidenced by R04's immediate recognition of the design-versus-effectiveness logic from its own ISAE 3402 practice. Concerns that are not immediately visible may surface after deeper use, particularly implementation constraints that only become apparent in a real client engagement with time pressure and limited access to documentation. The Round 2 evaluation therefore captures early adoption perception rather than experienced practitioner judgment, and the PEOU scores in particular should be interpreted with this temporal dimension in mind.

Table 5.4: Respondent Ratings per Item (Round 2)

Construct item	**R08 (returning)**	**R09 (new)**
PU1	5	4
PU2	5	5
PU3	5	5
PU4	5	3
PEOU1	5	5
PEOU2	5	4
PEOU3	5	5
SN1	5	5
SN2	5	3
SN3	5	4
BI1	5	4
BI2	5	5

R08 rated all twelve items at the maximum score of 5, producing a ceiling effect with no remaining variance to interpret quantitatively. This pattern is interpreted qualitatively in section 5.5 and is treated with caution in section 5.8, since R08's prior involvement in Round 1 introduces a plausible source of confirmation bias that a wholly independent respondent would not carry. R09's scores are more differentiated and therefore qualitatively richer. Three items stand out. PU4, concerning added value beyond existing standards such as ISO 27001 or NIS2, scored a 3, the same moderate result recorded for this item in Round 1, with R09 explaining the score in terms close to those used by Round 1 respondents, namely that other frameworks serve a different purpose rather than that this framework lacks value. SN2, concerning alignment with current regulatory expectations, also scored a 3 for the same underlying reason given in Round 1, namely the absence of an explicit regulatory mandate for this type of assessment. PEOU2, the item that produced the lowest and most divided Round 1 score, improved to a 4 for R09, who explained the remaining gap as the connection between the two dimensions, meaning design and operational effectiveness, that could still be stated more clearly.

5.5.1 Qualitative findings

The Round 2 interviews confirmed that the four Version 2 design adjustments reported in section 5.4 were perceived as resolving the problems they target, while also surfacing two new findings concerning the technical implementation of the Version 2 workbook that had not been identified in Round 1.

Confirmation of the Excel format and workflow

Both respondents independently and spontaneously welcomed the shift from a document-based instrument to an Excel workbook before any question specifically asked about the format. R08 stated that the workbook "looks really good. It makes it considerably easier and more applicable and also prevent

human error" (code: Excel format preferred). R09 connected the same preference directly to existing audit practice: "in practice, we work a lot with Excel, and the notes you have added fit that precisely" (code: Excel format preferred, notes and rationale field essential for audit defensibility).

The Step-by-Step Workflow added to the Instructions tab, the design adjustment made in response to Problem 1 in table 5.2, was also confirmed by both respondents as addressing the original framework versus guidance conflation. R08 described the M2 wording revision, in which the criterion now states explicitly that the organization has progressed beyond M1, as "very clear" (code: M2 reworded as progressed beyond M1). R09 confirmed the same point independently while reviewing the Definitions sheet, describing the explanation of the relationship between design maturity and operating effectiveness as clearer than before and explicitly crediting the Definitions and the workflow together for that clarity.

Confirmation of suggested evidence wording

The change from a fixed evidence list to suggested evidence to provide, made in response to Theme 5 rather than to one of the four numbered Round 1 problems, was endorsed by both respondents. R08 explained that the wording "gives direction without saying it has to be exactly this document" (code: principle-based norm preferred over prescriptive evidence). R09 went further and proposed an additional refinement not present in Version 2, namely an explicit disclaimer stating that the listed evidence is a minimum suggestion and that the auditor may request additional evidence based on professional judgment, a suggestion he summarized by remarking that auditors tend to use disclaimers (code: principle-based norm preferred over prescriptive evidence, new sub-code: explicit professional judgment disclaimer suggested).

New finding: workflow ordering of the Responsible Role table

R08 identified a sequencing problem in the Step-by-Step Workflow that had not been raised in Round 1. Step 5 of the workflow instructs the auditor to document the conclusion and identify ownership using the Responsible Role table, which is presented after the audit questions, Part A, and Part B steps. R08 pointed out that this ordering is logically inverted, since the individuals named in the Responsible Role table are typically the same individuals an auditor would need to approach earlier in order to collect evidence in the first place: "you already need those people for evidence at step 3 and step 4" (code: workflow ordering critique, responsible roles too late). This is treated as a new, Version 2-specific finding rather than a confirmation of a prior problem, since the Responsible Role table itself was a Version 2 addition not present in Version 1 and therefore could not have been raised in Round 1.

New finding: design maturity and operating effectiveness hierarchy not enforced

R09 raised a methodologically significant concern that had not surfaced in Round 1, namely that the Version 2 workbook does not technically prevent an Operating Effectiveness score from exceeding the corresponding Design Maturity score for the same control, which he argued is a logical impossibility: "operating effectiveness can currently be higher than design maturity, and that is not correct, since you must first design something before you can execute it" (code: design cannot be lower than operating effectiveness, logically hierarchy not enforced). R09 proposed two possible remedies, namely, enforcing the hierarchy technically within Excel, for example, through data validation rules that block an Operating Effectiveness entry exceeding the corresponding Design Maturity entry, or stating the constraint explicitly in the Instructions tab if a technical control is not implemented.

New finding: cumulative maturity logic not technically enforced

A closely related concern raised independently by R09 is that the cumulative principle underlying the M1 to M5 scale, confirmed as conceptually sound by R04 in Round 1, is not technically enforced in the Version 2 workbook. R09 observed that an auditor can currently select M3 in the drop-down without the workbook requiring confirmation that the M1 and M2 criteria were also satisfied: "you can currently select level 3 without checking levels 1 and 2, but that is not how it should work" (code: cumulative logic is not technically enforced). This finding extends rather than repeats Theme 3 from Round 1, since R04's Round 1 endorsement concerned the conceptual soundness of cumulative scoring, while R09's Round 2 finding concerns the absence of technical safeguards ensuring that the conceptual principle is actually applied during scoring.

Other observations

R08 also questioned the precision of the term 'live exercise' as defined in the Definitions sheet, where it is described as non-simulated execution of exit or migration steps. He noted an apparent contradiction in this wording, observing that if an exercise is genuinely non-simulated, it is difficult to see in what sense it remains an exercise rather than an actual migration (code: terminology ambiguity, live exercise definition). R09 raised a broader positioning concern, recommending that the introduction state explicitly that operational reversibility is one component of the wider cloud sovereignty domain rather than a complete treatment of it (code: framework needs explicit scope positioning within broader sovereignty domain) and separately recommended that the framework be positioned as an advisory tool rather than a formal assurance instrument (code: advisory positioning, not assurance), a positioning also consistent with Theme 4 from Round 1.

5.6 Comparison between rounds

Table 5.5 compares the four Round 1 problems and their Version 2 adjustments against the Round 2 evidence collected for each and adds the two new findings that emerged specifically from the Round 2 evaluation of the Version 2 implementation.

Table 5.5: Overview of Evaluation Rounds, Adjustments, and Evidence

#	**Round 1 problem**	**Version 2 adjustment**	**Round 2 evidence**
1	Framework and guidance conflated.	Step-by-Step Workflow and Definitions sheet was added.	Confirmed resolved by both R08 and R09; both independently credited the workflow and Definitions sheet for added clarity.
2	Part A and Part B definitions unclear.	Explicit workflow steps stating what each part assesses; never averaged.	Confirmed positively by R08, who specifically praised the requirement to document rationale rather than apply a mechanical rule.
3	Guidance on how to test operating effectiveness in practice is underdeveloped relative to the criteria stating what should be assessed.	Not resolved in Version 2. R04's proposed proxies, namely a data migration test and a backup restoration test with documented timing, are candidate additions to the Part B evidence guidance. These are carried forward as recommendations for future revision in chapter 6.	Not directly probed in Round 2. Neither R08 nor R09 independently raised this concern, but the absence of a complaint does not constitute confirmation that the underlying gap was resolved. The item remains open.
4	Binary conclusions collapse graded evidence.	Decision rules reframed as Decision Guidelines; weakest-link logic with documented departure.	Confirmed positively by R08, who specifically praised the requirement to document rationale rather than apply a mechanical rule.
5	New in Round 2	Not applicable	Workflow sequencing of the Responsible Role table places role identification after evidence collection steps, when the same roles are needed earlier (R08).
6	New in Round 2	Not applicable	Design Maturity and Operating Effectiveness hierarchy is not technically enforced, allowing a logically inconsistent score combination (R09).
7	New in Round 2	Not applicable	Cumulative maturity logic (M1 through M5) is not technically enforced in the drop-down selection mechanism (R09).

Two broader patterns follow from this comparison. First, all four Version 2 design adjustments directly probed in Round 2 were confirmed effective by at least one respondent, and none were reported to have introduced a new problem of their own. This supports treating the Version 1 to Version 2 documented in chapter 4 as having addressed the concerns it was designed to address, at least within the bounds of a two-respondent confirmatory sample. Second, both new findings from Round 2, namely the Responsible Role sequencing issue and the unenforced design and operating effectiveness hierarchy together with the unenforced cumulative logic, are specific to the Excel implementation of Version 2 rather than to the underlying conceptual design described in chapter 4. They were not and could not have been raised in Round 1, since the workbook format, the Respondent Role table, and the drop-down-based scoring mechanism did not exist in Version 1. This suggests that moving the artifact into a more usable, structured format surfaced a distinct category of findings, namely implementation-level enforcement gaps, that a document-based Version 1 could not have exposed regardless of how thoroughly its conceptual content was reviewed.

5.7 Conclusion

This chapter reported two rounds of expert evaluation of the DSAF Operational Reversibility artifact. Round 1, involving seven IT audit and Technology Risk professionals, established strong support for the framework's underlying value, reflected a near-unanimous PU1 and PU2 scores and captured an unconditional endorsement in Theme 1, alongside four concrete problems concentrated in the framework's presentation rather than its substance, namely the conflation of criteria and guidance, the unclear Part A and Part B distinction, underdeveloped testing guidance, and a binary overall conclusion that discarded graded evidence. Four corresponding design adjustments were made and documented in chapter 4, moving the artifact from a document-based instrument to an Excel workbook with an explicit workflow, a glossary, suggested rather than mandatory evidence, and guideline-based rather than rule-based conclusions.

Round 2, involving one returning and one new respondent, confirmed that these four adjustments were perceived as resolving the problems they targeted while surfacing two new, implementation-specific findings concerning the technical enforcement of the design and operating effectiveness hierarchy and of the cumulative maturity logic, together with a workflow sequencing issue concerning the Responsible Role table. Taken together, the two rounds demonstrate both that the artifact's core conceptual design, established in chapter 4, has held up under independent evaluation across two cycles and that translating that design into a usable instrument introduced new, narrower technical considerations that a purely conceptual review could not have surfaced. These considerations are carried forward as recommendations for a future revision in chapter 6.

5.8 Limitations

Several limitations qualify the findings reported in this chapter. The Round 1 sample of seven respondents and the Round 2 sample of two respondents are both small, and neither is large enough to support statistical generalization beyond the specific context of a single Big Four operating in the Dutch market from which all nine respondents were drawn. The findings should accordingly be read as analytically rich, expert-informed evidence of perceived usefulness, ease of use, and intention to adopt, rather than as a population-level estimate of these constructs.

Within round 2 specifically, R08's involvement in both rounds introduces a plausible source of confirmation bias. A respondent who is aware that a second version was produced in direct response to feedback, including feedback that the respondent may have given, has an incentive, whether conscious or not, to perceive and report improvement. R08's uniform maximum score across all twelve TAM items, discussed in section 5.1.2, is consistent with this risk and is one reason this chapter places more interpretive weight on R08's specific, concrete qualitative observations, such as the Responsible Role sequencing issue, than on his quantitative scores in isolation. R09's more differentiated score and independent first exposure to the framework partially offset this limitation but cannot substitute for a confirmatory round conducted entirely with respondents who had no prior involvement with the version of the artifact.

The translation of all interview material from Dutch into English for presentation in this chapter introduces an additional source of interpretive distance between the original respondent statements and the analysis presented here. While the translations were prepared with care to preserve the respondent's intended meaning, translation is not a neutral or fully reversible process, and subtle connotations presented in the original Dutch phrasing may not be perfectly recoverable in English. Finally, both evaluation rounds relied on self-reported intention to use rather than observation of the artifact in an actual audit engagement. Behavioral intentions scores, while consistently high across both rounds, measure stated rather than demonstrated adoption, and the gap between the two is a well-documented limitation of TAM-based evaluation more generally [13, 56].

Chapter 6

Discussion

This chapter interprets and connects the findings reported in chapter 4 and chapter 5, places them in the context of existing literature and reflects on what they mean for research and practice.

6.1 Findings in relation to the literature

The findings of this study can be read at three levels: what they say about the DSAF as a design artifact, what they say about how IT audit professionals relate to digital sovereignty concerns, and what they say about the conditions under which a new audit instrument enters practice.

6.1.1 The DSAF fills a real and specific gap, not a general one

The comparative analysis in chapter 2 identified a specific structural gap: no existing European sovereignty framework combines independently verifiable evidence requirements, granular maturity scoring, and an aggregation mechanism that produces a defensible overall conclusion in a format compatible with IT audit practice. The near-unanimous PU1 score (mean 4.00, SD 0.00) confirms that this gap is real from the perspective of practitioners: every respondent agreed that the DSAF surfaces cloud exit risks that a standard audit would otherwise miss. This is not a trivial finding. It means that the problem identified through the literature was not a theoretical construction but a gap that experienced professionals recognize from their own engagements.

What is equally important is understanding why the gap existed in the first place. DICTU, the highest-scoring framework in chapter 2 at 3.80/4.00, provides a rigorous five-level maturity definition for operational reversibility in criterion 3.4. But DICTU was developed as a procurement assessment instrument for the Dutch government, not as an audit instrument for IT assurance engagements. Its definitions tell an organization what sovereignty level it should aim for; they do not tell an auditor what evidence to request, how to verify that evidence independently, how to score two independently observable properties of a control (its design and its operations), or how to aggregate findings across multiple control areas into a single auditable conclusion. These are audit-specific requirements that DICTU was never designed to meet. The DSAF's contribution is, therefore, not to replace or improve DICTU, but to translate DICTU's normative definitions into an instrument that serves a fundamentally different professional purpose.

The moderate PU4 score (mean 3.57, SD 0.53) reflects this same distinction from a different angle. Respondents were uncertain about whether the DSAF adds value beyond what ISO 27001 or NIS2 already require. But ISO 27001 and NIS2 address security and resilience at the operational level; neither requires an assessment of whether an organization can exit its cloud provider. R03 articulated this precisely when he noted that ISO 27001 does not address sovereignty and concluded that the DSAF therefore does add value. The moderate PU4 score reflects a communication gap rather than a substantive overlap: the DSAF's scope differentiation from existing standards needs to be stated more prominently within the instrument itself.

6.1.2 The two-dimensional scoring model

PEOU2 (mean 3.50, SD 1.05) was the lowest and most divided individual score in the entire evaluation. It asked whether the distinction between Design Maturity (Part A) and Operating Effectiveness (Part B) is easy to understand and apply. The variation in scores reflects differences in prior exposure to this logic, not disagreement about whether the distinction itself is valid. R04, who was already familiar with the test-of-design and test-of-effectiveness distinction from ISAE 3402 practice, found Part A and Part B immediately intuitive and rated PEOU2 substantially higher. R02 and R03, who approached the framework without that reference point, found it initially confusing. The problem is therefore one of framing, not the concept: the instrument in Version 1 asked users to apply a distinction without explicitly connecting it to the audit logic they already use. The Step-by-Step Workflow and Definitions tab added in Version 2 resolved this for both Round 2 respondents, and R09's PEOU2 score improved from the Round 1 average to 4. This confirms that the solution is presentational rather than structural, and that as IT audit practice becomes more familiar with two-dimensional maturity scoring in sovereignty contexts, this barrier will continue to decrease.

6.1.3 Adoption is institutionally driven, not individually driven

The TAM evaluation produces a finding that is counterintuitive from a standard TAM perspective but entirely consistent with the sociology of professional practice: Behavioral Intention remained high (BI1: 4.43, BI2: 4.57) despite Perceived Ease of Use being the lowest-scoring construct. In the standard TAM model, ease of use is expected to influence intention to use, particularly for voluntary adoption of tools that are not yet mandated [13, 56]. The decoupling observed here contradicts this prediction and requires an explanation rooted in the specific professional context.

The explanation lies in the social norm findings taken together. SN3 (mean 4.86, SD 0.38) is the highest score in the entire dataset: almost all respondents indicated that formal adoption by a leading professional body would make them more inclined to use the framework. SN2 (mean 3.29, SD 0.76) is the lowest score overall: respondents did not feel that current regulation already requires this type of assessment. Reading these two findings together reveals the adoption dynamic at work. IT audit is a profession in which individual practitioners do not independently decide which frameworks to use in client engagements. Those decisions are shaped by professional bodies such as NOREA, by Big Four methodology committees, and by regulatory expectations. A practitioner who finds the DSAF slightly difficult to use in a first engagement will still use it if their professional environment expects them to. On the other hand, a practitioner who finds it straightforward will still not deploy it if it carries no institutional recognition. This is why PEOU did not predict BI in this study: the adoption decision does not primarily sit with the individual.

This finding has a direct implication for the DSAF's deployment pathway that goes beyond what further design refinement can achieve. The gap between SN2 and SN3 shows clearly that the DSAF's primary barrier is not its quality but its institutional status. The most impactful next step is submission to NOREA or an equivalent body for review and potential endorsement, not a Version 3, with additional technical improvements. This mirrors the trajectory of COBIT, which became the standard for IT governance audit not because individual auditors evaluated it as superior to alternatives, but because ISACA's institutional authority made it the expected approach. The same logic applies here, and the SN3 score confirms that the professional willingness is already in place.

The SN2 finding also carries a signal for regulators. NIS2 Article 21, DORA Article 11, and GDPR Article 20 collectively create obligations that the DSAF operationalizes, but practitioners do not yet experience those provisions as an explicit mandate for operational reversibility assessment. The translation of existing regulatory obligations into specific audit expectations has not yet happened, and until it does, adoption will depend on professional endorsement rather than regulatory compulsion.

6.1.4 Professional judgement cannot be replaced, but it can be anchored

Theme 5 from the qualitative analysis surfaces a tension that is well established in the audit literature: the tension between standardization and judgment. Power [47] argues that an audit is always at risk of becoming a ritual of verification rather than a substantive assessment of risk, precisely because the drive toward standardization tends to produce checklists that organizations learn to satisfy without genuinely addressing the underlying concern. R05's observation that a one-page exit plan will generate a 'yes'

answer to the question 'does exit plan exist' describes exactly this risk in the sovereignty context.

The Version 2 response, reframing evidence requirements as suggestions and rules as guidelines with non-mechanical language, is directly consistent with the principle-based standard-setting approach used in ISAE 3402 and SOC 2, where control objectives specify what must be achieved without prescribing how, leaving evidence determination to professional judgment [34]. The evaluation confirmed this was the right direction: R08 described the suggested evidence wording as giving direction without dictating form. What the consistently high BI scores confirm, read alongside Theme 5, is that practitioners do not want less structure. They want structure that functions as an anchor for their judgment rather than a substitute for it. The DSAF's Version 2 design achieves this balance, and the evaluation evidence supports treating it as stable on this dimension.

6.2 Position within the broader digital sovereignty landscape

6.2.1 Operational reversibility is one sub-dimension, not a complete treatment

This study deliberately scoped itself to operational reversibility, defined as the ability of an organization to exit its current cloud provider without catastrophic loss of service continuity, data, or institutional knowledge. This was the right scoping decision for a first instrument in this domain: it is the sub dimension most directly assessable through existing IT audit methodology, it has the largest gap between normative definition and practical assessment tool, and it is auditable without requiring legal or financial expertise beyond the standard IT audit skill set. But it is one sub-dimension among several, and a strong DSAF score does not imply sovereignty across all dimensions.

The DSAF should be positioned explicitly as the first module of a broader digital sovereignty audit programme rather than as a standalone sovereignty assessment. Three additional modules identified in the comparative analysis of chapter 2, would be needed. A cryptographic access control module would assess who holds the encryption keys protecting cloud-stored data and whether the provider can technically access that data when legally compelled, grounded in DICTU criterion 2.2 and its five-level taxonomy from platform-managed keys to HYOK with confidential computing [15]. A data residency and jurisdictional compliance module would assess whether data is processed and stored within EU jurisdiction and whether contractual and technical measures are sufficient to limit extraterritorial exposure, grounded in DICTU criterion 2.1 [15] and EC-CSF SOV-2 [16]. A personnel independence module would assess whether the organization has sufficient internal EU-based capacity and skills to manage its cloud environment without dependence on provider-specific knowledge or non-EU personnel, grounded in DICTU domain 5 [15] and EC-CSF SOV-1 [16].

Together, these four modules would constitute a complete digital sovereignty audit programme, covering the dimensions that existing European frameworks identify as material but that no single existing instrument currently operationalizes for IT audit practice. The cryptographic access control module is identified as the most urgent next step, as it is fully within IT audit competence, directly grounded in all five frameworks examined in chapter 2, and addresses a risk that organizations face regardless of whether they have planned an exit.

6.2.2 Cryptographic access control is the most urgent next subdimension

Consistent with the reasoning that motivated the choice of operational reversibility, the next subdimension should be technically assessable through IT audit methodology without requiring legal or financial expertise, and it should address a risk that organizations face regardless of whether they have planned an exit. Cryptographic access control, specifically the question of who holds the encryption keys that protect cloud-stored data and whether the provider can technically access that data when legally compelled to do so, meets both criteria and is directly grounded in the frameworks examined in chapter 2.

All five frameworks in the comparative analysis treat cryptographic control as a distinct, independently assessable dimension. The EC-CSF addresses it under SOV-3 (Data and AI Sovereignty), which requires customer-exclusive cryptographic control and data access auditability. DICTU criterion 2.2 (Technical Access Security) defines a five-level taxonomy specifically for this purpose, ranging from Platform Managed Keys at level 2, where the provider can technically access data if legally compelled, to HYOK (Hold

Your Own Key) plus Confidential Computing at level 5, where the provider is cryptographically prevented from accessing the data regardless of any legal demand. The Buy European Framework requires verifiable technical measures, including confidential computing or customer-exclusive key management as a core guarantee in dimension IV. EuroStack identifies cryptographic isolation as central to technical sovereignty. This cross-framework convergence is not coincidental: it reflects a shared recognition that contractual data sovereignty claims are only as strong as the technical enforcement underneath them. A provider can contractually promise not to access data and still be legally compelled to hand it over if the organization does not hold its own keys.

This is directly relevant to IT audit for two reasons. First, key management architecture is already a standard IT audit object. Auditors assess encryption controls in SOC 2, ISO 27001, and ISAE 3402 engagements as part of access control and cryptography domains. Extending this to assess whether an organization has implemented customer-managed keys, whether key management systems are operated exclusively within EU infrastructure, and whether confidential computing isolates workloads from provider access does not require new professional competencies. It requires applying existing cryptography audit skills to a sovereignty-specific question. Second, this dimension fills a gap that the DSAF explicitly leaves open. An organization can achieve M5 across all five DSAF control areas, meaning it has a tested exit plan, exportable data, comprehensive documentation, continuity provisions, and a periodic testing programme, and still have its data technically accessible to its provider throughout. The DSAF establishes that the organization can leave; a cryptographic access control instrument would establish that the provider cannot look inside while the organization is still there. Together, the two instruments would address both sides of the cloud dependency risk.

Building a cryptographic access control module in the same Excel workbook, using the same M1 to M5 and E1 to E4 structure, would allow organizations to use both instruments within a single engagement and produce a combined score across operational reversibility and cryptographic sovereignty. The DICTU criterion 2.2 taxonomy already provides the five-level maturity foundation for such an instrument, just as criterion 3.4 provided the foundation for the DSAF, which means the design work required to produce a second module is considerably smaller than what this study required for the first.

6.3 Reflection on the absence of a real-world case study

This study evaluated the DSAF through expert interviews and a TAM-based questionnaire. It did not apply the framework in a real audit engagement. This section addresses that methodological choice directly.

6.3.1 Why expert evaluation is sufficient at this stage

The decision not to include a case study reflects both practical constraints and a deliberate methodological position. Applying a new framework in a real audit engagement requires client access and engagement scope agreements that are not available within the scope of a master's thesis. But beyond the practical constraint, expert evaluation is also the methodologically correct choice for the stage of artifact maturity this study represents.

Hevner et al. (2004) identify expert evaluation as one of the primary methods for assessing artifact utility at early stages of development, alongside field testing and simulation [31]. Expert evaluation is not a substitute for field testing, but it is the correct method when the primary objective is to establish that an artifact is both useful and usable before investing in full deployment. A field study conducted before the conceptual and presentational issues identified in Round 1 were resolved would have produced data about a weaker version of the instrument. The iterative DSR cycle used here produces a better-validated Version 2 than a single-round field study of Version 1 would have.

This view was also expressed directly by a Round 2 respondent with extensive client-facing audit experience. They noted that in professional practice, review and endorsement of a framework by qualified IT audit professionals is generally considered sufficient evidence of practical applicability. An audit framework does not need to have been deployed in a client engagement to be considered ready for use; it needs to have been assessed by people who do those engagements. That is precisely what this study's evaluation design achieved.

6.3.2 What a real-world application would look like

If the DSAF were applied in an actual engagement, the four-step workflow described in chapter 4 would structure the fieldwork. The auditor would first orient using the Instructions and Definitions tabs, then work through each of the five control areas, interviewing the relevant role from the Responsible Role table, requesting the suggested evidence at each maturity level, and scoring Part A and Part B independently. The Summary tab would aggregate scores automatically and flag weak links through the Decision Guidelines.

A full assessment would likely require between three and seven audit days of fieldwork, depending on the organization's cloud complexity and documentation maturity. The five control areas involve different organizational stakeholders: C1 requires legal and senior management, C2 requires infrastructure engineering, C3 requires operations and architecture teams, and C4 and C5 require risk management and the audit committee. Based on what respondents described about their current clients, the most likely initial finding for most organizations is a score around M2 to M3 across most control areas. R04 noted that most organizations he encounters have a basic continuity plan but no genuine exit plan, have never tested data portability in a non-provider environment, and have not conducted a tabletop exercise for an unplanned exit scenario. This means the DSAF would, in practice, surface risks that organizations are genuinely not aware of, which is exactly what the high PU1 and PU2 scores predict.

6.4 Limitations

Three limitations qualify the conclusion of this study.

First, the sample is small and homogeneous. Nine respondents across two rounds, all from a single Big Four firm in the Netherlands, does not support statistical generalization. The findings represent the perceptions of experienced IT audit professionals in a specific professional and national context. The saturation threshold cited for Round 1 applies within this homogeneous expert group, not across the broader population [29]. Whether smaller firms, public sector audit bodies, or auditors in other European countries would evaluate the DSAF differently remains an open empirical question.

Second, the study measured stated rather than demonstrated adoption. Behavioral Intention captures what respondents say they would do, not what they actually do when facing client constraints, time pressure, and engagement scope limitations [13, 56]. The gap between expressed intention and actual behavior is a recognized limitation of TAM-based evaluation and cannot be resolved without longitudinal observation of real engagements [13, 56].

Third, Round 2 includes one respondent, R08, who also participated in Round 1 and was therefore aware that Version 2 was built in direct response to his earlier feedback. This creates a plausible source of confirmation bias. R08's uniform maximum score of 5 across all twelve TAM items is consistent with this risk, which is why this chapter places greater interpretive weight on the specific qualitative observation than on his quantitative scores taken in isolation.

Chapter 7

Conclusions

This chapter brings the thesis to a close. It answers the main research question and each sub-question in turn, summarizes the contributions of the study to research and practice, and identifies the most important directions for future work.

7.1 Answers to the research questions

The main research question asked: How can IT auditors assess digital sovereignty in organizations using cloud services in a structured and auditable manner? The five sub-questions provided the path to answering it.

7.1.1 SQ1: What is digital sovereignty, and which key elements characterize it in organizational contexts?

"Digital sovereignty" refers broadly to the ability of organizations and states to retain control over their data, digital infrastructure, and decision-making in a globalized digital ecosystem [45]. The academic literature identifies four key dimensions [11]: infrastructural sovereignty, meaning control over the physical and virtual infrastructures that enable digital communication; data sovereignty, meaning the ability to control how data is collected, stored, processed, and transferred across jurisdictions; technological sovereignty, meaning the capacity to develop and operate digital technologies independently from external actors; and regulatory sovereignty, meaning the authority to define and enforce legal frameworks governing digital activity within a specific jurisdiction. These dimensions show that digital sovereignty extends beyond technical considerations and includes governance, economic, and geopolitical aspects.

At the policy level, European frameworks operationalize these dimensions differently. The EC-CSF structures its assessment around eight Sovereignty Objectives (SOV-1 through SOV-8), ranging from Strategic Sovereignty and Legal and Jurisdictional Sovereignty through Supply Chain Sovereignty, Technology Sovereignty, and Environmental Sustainability. DICTU organizes its assessment across five domains: Juridical, Data and AI, Technology, Operational, and Human/Personnel, with the Technology domain containing criterion 3.4 on Operational Reversibility that forms the direct foundation of this study. The EuroStack framework is grounded in seven guiding principles, including sovereignty and security, deproprietarization and interoperability, and inclusive governance. The Buy European Framework defines five hierarchically ordered dimensions from Jurisdiction and Governance as a mandatory pass/fail gate through to Economic Sovereignty as a key differentiator.

Closely related is "digital autonomy", which refers to the practical operational capability to act independently within digital ecosystems without excessive reliance on external providers [50]. Where "sovereignty" describes authority and governance, "autonomy" describes the operational capacity to exercise that authority. Operational reversibility, the focus of this thesis, is one specific and auditable expression of digital autonomy: the ability to exit a cloud provider relationship without catastrophic loss of service continuity, data, or institutional knowledge.

7.1.2 SQ2: What aspects of digital sovereignty are not adequately addressed by existing IT audit and assurance frameworks?

A comparative analysis of five existing European sovereignty frameworks, the EC Cloud Sovereignty Framework, EuroStack, the DICTU Toetsingsinstrument, the EuroStack Buy European Framework, and the SUSE CFSA, against six audit-relevant criteria showed that no single framework provides IT auditors with the combination of scoring granularity (criterion 3), independently verifiable evidence requirements (criterion 2), and a gated aggregation mechanism (criterion 5) needed to produce a defensible audit conclusion. DICTU scored highest overall (3.80/4.00) and provided the strongest maturity model, but was developed for procurement assessment rather than IT audit. None of the five frameworks assessed demonstrated operation as a dimension independent from design intent. The gap is specific: the dimension most underserved is operational reversibility, the ability of an organization to exit its current cloud provider without catastrophic loss, which is auditable through standard IT audit methodology without requiring legal or financial expertise.

7.1.3 SQ3: Which audit criteria can be used to translate digital sovereignty into an auditable concept for IT audit practice?

Seven design requirements were derived from the gap analysis, grounded in the DICTU source definitions and the audit-specific criteria identified in the comparative scorecard. The most critical were the following: the framework must produce an auditable conclusion on exit capability (DR1); maturity must be assessed on two independent dimensions (DR2); the instrument must not force auditors into binary judgments on inherently graded capability (DR3); quantitative indicators must be verifiable and principle-based rather than rigid gates (DR4); the instrument must integrate into standard audit workpaper structure without requiring a new evidentiary paradigm (DR5); the framework must address not only contractual and technical exit conditions but also continuity of operations during exit and periodic testing of exit capability (DR6); and the framework must be replicable by an auditor unfamiliar with its construction (DR7). Together, these requirements translate the normative concept of operational reversibility into an auditable construct with concrete criteria for what evidence to request and how to evaluate it.

7.1.4 SQ4: How can these audit criteria be operationalized into audit procedures, required evidence, and sovereignty assurance levels?

The DSAF operationalizes the design requirements across five control areas, each covering a distinct dimension of operational reversibility: C1 (Exit plan and contractual framework), C2 (Data portability and technical exportability), C3 (Knowledge transfer and operational documentation), C4 (Continuity of operations during migration), and C5 (Testing programme and periodic evaluation). Each control area contains five closed audit questions, a cumulative Design Maturity scale (M1 to M5) derived from DICTU criterion 3.4, and an independent Operating Effectiveness scale (E1 to E4) derived from the Test of Design and Test of Effectiveness logic established in ISAE 3402 and SOC 2 [33]. The five control areas are aggregated through an average Design Maturity and Operating Effectiveness score, combined with Decision Guidelines that apply weakest-link logic to ensure that a critically weak control area can determine the overall conclusion without discarding the graded information from the other four. The instrument is implemented as an Excel workbook with dropdown-based scoring, automatic score carry-through to a Summary tab, and a rationale field per control area to support audit defensibility.

7.1.5 SQ5: How do IT audit and Technology Risk professionals evaluate the usability and practical relevance of the proposed framework?

Two rounds of expert evaluation were conducted. Round 1 involved seven IT audit and Technology Risk professionals from Big Four firms, who evaluated Version 1 of the DSAF using a twelve-item TAM questionnaire and a semi-structured interview. Perceived Usefulness was strong, with unanimous agreement that the framework surfaces cloud exit risks that a standard audit misses (PU1: mean 4.00, SD 0.00) and improves audit conclusion completeness (PU2: mean 4.67, SD 0.52). Four concrete problems were identified, leading to five improvements in Version 2, the most significant of which was the transition to an Excel workbook format preferred by six of seven respondents. Round 2, involving two respondents evaluating Version 2, confirmed that all four improvements resolved the problems they targeted. Social

Norm produced the sharpest finding of the entire evaluation: SN3 (mean 4.86, SD 0.38, the highest score in the dataset) showed near-unanimous agreement that formal adoption by a professional body such as NOREA would drive adoption, while SN2 (mean 3.29, SD 0.76, the lowest score in the dataset) showed that current regulation is not yet perceived as explicitly mandating this type of assessment. Behavioral Intention remained high despite moderate ease of use scores (BI1: 4.43, BI2: 4.57), confirming that adoption in the IT audit profession is institutionally driven rather than individually driven [13, 56].

7.1.6 Main research question

IT auditors can assess digital sovereignty in cloud-dependent organizations in a structured and auditable manner by applying the DSAF: a five-control-area instrument that translates the operational reversibility dimension of digital sovereignty into cumulative design maturity criteria, independent operating effectiveness criteria, and a weakest-link aggregation mechanism implemented in an Excel workbook compatible with standard IT audit practice. The DSAF does not cover all dimensions of digital sovereignty; it covers the one dimension that is most directly assessable through existing IT audit methodology and that has the largest gap between normative definition and practical audit tools. Extending the programme to additional dimensions, beginning with cryptographic access control as the next most urgent and technically auditable subdimension, is identified as the primary direction for future work.

7.2 Contributions

This study contributes to both the academic literature on digital sovereignty and the professional practice of IT auditing. The scientific and practical contributions are discussed in turn.

7.2.1 Scientific contribution

This study makes three contributions to the research literature. First, it produces the first structured, evidence-based IT audit instrument for assessing operational reversibility as a dimension of digital sovereignty. Existing research on digital sovereignty focuses primarily on policy, governance, and legal dimensions; this study demonstrates that at least one dimension of digital sovereignty can be operationalized into auditable criteria using standard DSR methodology and evaluated through expert elicitation with TAM-based instrumentation.

Second, the study contributes an empirically grounded gap analysis of five major European sovereignty frameworks against six audit-relevant criteria, producing a comparative scorecard that identifies where and why existing frameworks fall short for IT audit purposes. This analysis is a reusable resource for future research designing audit instruments for other sovereignty dimensions.

Third, the evaluation findings extend the TAM literature in a professional audit context by demonstrating that Behavioral Intention can remain high despite moderate Perceived Ease of Use when Social Norm functions as the primary adoption driver. This decoupling suggests that TAM's standard prediction about ease of use and intention requires qualification in highly institutionalized professional environments where adoption decisions are shaped by professional bodies and regulatory expectations rather than individual usability assessments.

7.2.2 Practical contribution

For IT audit practice, the DSAF provides a ready-to-use instrument that can be integrated into existing cloud audit programmes without requiring new professional competencies or a new evidentiary paradigm. The Excel workbook format, confirmed as the preferred format by six of seven Round 1 respondents, is compatible with existing audit tooling and supports the documentation of professional judgment through per-control rationale fields. The instrument operationalizes obligations that NIS2, DORA, and the GDPR already impose on organizations but that no existing audit framework has yet translated into structured assessment criteria, giving IT auditors a defensible basis for addressing cloud exit capability in their engagements.

7.3 Future work

Three directions for future work follow directly from this study.

The most immediate priority is the technical enforcement of the scoring logic within the Excel workbook. Two implementation gaps identified in Round 2 remain unresolved: the workbook does not technically prevent an Operating Effectiveness score from exceeding the corresponding Design Maturity score for the same control area, and the cumulative maturity logic is not enforced through the dropdown selection mechanism. Both can be addressed through Excel data validation rules and would eliminate a category of logically inconsistent score combinations that the current version permits.

The second priority is the development of a cryptographic access control module as the next instrument in the digital sovereignty audit programme. As established in chapter 6, cryptographic access control, specifically who holds the encryption keys that protect cloud-stored data and whether the provider can technically access that data when legally compelled is the sub-dimension most directly grounded in the frameworks examined in chapter 2, most urgent for organizations currently exposed to CLOUD Act jurisdiction, and most compatible with existing IT audit competencies. DICTU criterion 2.2 already provides a five-level maturity taxonomy that could serve as the foundation for such an instrument in the same way that DICTU criterion 3.4 served as the foundation for the DSAF.

The third priority is a real-world case study applying the DSAF in an actual audit engagement. The evaluation conducted in this study established that the instrument is perceived as useful and applicable by practitioners who assess frameworks of this type professionally. A field study would provide the next level of evidence: whether the instrument produces consistent findings across auditors, what the typical score profile looks like for organizations that have not previously assessed their cloud exit capability, and what practical implementation constraints arise in client engagements that laboratory evaluation cannot surface. Such a study would also provide the empirical basis for submitting the DSAF to NOREA or an equivalent professional body for formal endorsement, which the Social Norm findings identify as the single most impactful step toward widespread adoption.

Bibliography

[1] Icek Ajzen. The theory of planned behavior: Frequently asked questions. *Human behavior and emerging technologies*, 2(4):314–324, 2020.

[2] American Institute of Certified Public Accountants. Trust services criteria for security, availability, processing integrity, confidentiality, and privacy. Technical report, American Institute of Certified Public Accountants, New York, NY, 2023. URL `https://www.aicpa-cima.com/`.

[3] American Institute of Certified Public Accountants (AICPA). SOC 2® reporting on an examination of controls at a service organization relevant to security, availability, processing integrity, confidentiality, or privacy. Technical report, AICPA, New York, NY, 2022. URL `https://www.aicpa-cima.com/cpe-learning/publication/soc-2-reporting-on-an-examination-of-controls-at-a-service-organization-relevant-to-security-availability-processing-integrity-confidentiality-or-privacy`. Trust Services Criteria.

[4] Anthropic. Important update on Fable and Mythos access, 2025. URL `https://www.anthropic.com/news/fable-mythos-access`.

[5] Daniele Catteddu. Cloud computing: benefits, risks and recommendations for information security. In *Iberic Web Application Security Conference*, pages 17–17. Springer, 2009.

[6] Victoria Clarke and Virginia Braun. Thematic analysis. *The journal of positive psychology*, 12(3): 297–298, 2017.

[7] Cloud Security Alliance. Cloud Controls Matrix (CCM). Technical report, Cloud Security Alliance, 2026. URL `https://cloudsecurityalliance.org/research/cloud-controls-matrix`.

[8] CMMI Product Team. CMMI for development, version 2.0: Improving processes for developing better products and services. Technical report, CMMI Institute, Pittsburgh, PA, 2018. URL `https://cmmiinstitute.com/cmmi/dev`. In partnership with ISACA.

[9] European Commission. Shaping europe's digital future. Communication, European Commission, Luxembourg, February 2020. URL `https://commission.europa.eu/document/download/84c05739-547a-4b86-9564-76e834dc7a49_en?filename=communication-shaping-europes-digital-future-feb2020_en.pdf`. Manuscript completed in February 2020. PDF ISBN 978-92-76-16363-3, doi:10.2759/091014.

[10] Committee of Sponsoring Organizations of the Treadway Commission. Internal control — Integrated framework. Technical report, Committee of Sponsoring Organizations of the Treadway Commission, New York, NY, 2013. URL `https://www.coso.org/internal-control`.

[11] Stephane Couture and Sophie Toupin. What does the notion of "sovereignty" mean when referring to the digital? *New media & society*, 21(10):2305–2322, 2019.

[12] John W. Creswell and Vicki L. Plano Clark. *Designing and conducting mixed methods research*. SAGE Publications, Thousand Oaks, CA, 3rd edition, 2018.

[13] Fred D Davis, Richard P Bagozzi, and Paul R Warshaw. Technology acceptance model. *J Manag Sci*, 35(8):982–1003, 1989.

[14] Steven De Haes. Enterprise governance of information technology. 2009.

[15] Dienst ICT Uitvoering (DICTU). Toetsingsinstrument soevereiniteit clouddiensten. Rapport, Ministerie van Economische Zaken en Klimaat, January 2026. URL https://www.dictu.nl/sites/default/files/bestanden/website/DICTU%20Toetsingsinstrument%20Soevereiniteit%20Clouddiensten%20v1.0.1.pdf. Toetsingskader voor het beoordelen van digitale soevereiniteit van clouddiensten binnen de Nederlandse overheid.

[16] European Commission, Directorate-General for Digital Services. Cloud sovereignty framework, version 1.2.1. PDF document, Oct 2025. URL https://commission.europa.eu/document/download/09579818-64a6-4dd5-9577-446ab6219113_en. European Commission policy framework document; accessed 2026-01-22.

[17] European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), 2016. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679.

[18] European Union. Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011 (Text with EEA relevance), 2022. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022R2554.

[19] European Union. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive), 2022. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32022L2555.

[20] European Union. Regulation (EU) 2025/2518 of the European Parliament and of the Council of 26 November 2025 laying down additional procedural rules on the enforcement of Regulation (EU) 2016/679 (Text with EEA relevance), 2025. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32025R2518.

[21] European Union Agency for Cybersecurity (ENISA). ENISA threat landscape 2023. Technical report, Publications Office of the European Union, Athens, Greece, 2023. URL https://www.enisa.europa.eu/publications/enisa-threat-landscape-2023. Accessed: 2026-01-22.

[22] EuroStack Coalition. Open letter to european commission president ursula von der leyen, March 2025. URL https://eurostack.eu. Signed by a broad coalition of European businesses and associations.

[23] EuroStack Industry Initiative. A proposed framework for a "buy european" regulation of strategic digital procurement. Technical report, EuroStack, September 2025. URL https://eurostack.eu/a-proposed-framework-for-buy-european-regulation/. Position paper proposing a legally enforceable framework for defining and procuring sovereign European digital providers.

[24] EuroStack Initiative. The eurostack: A european alternative for digital sovereignty. Technical report, EuroStack, 2025. URL https://eurostack.eu/the-white-paper/. White paper outlining a strategic vision for European digital sovereignty across AI, cloud, open source, and critical digital infrastructure.

[25] Henry Farrell and Abraham L Newman. Weaponized interdependence: How global economic networks shape state coercion. *International security*, 44(1):42–79, 2019.

[26] Luciano Floridi. The fight for digital sovereignty: What it is, and why it matters, especially for the eu. *Philosophy & Technology*, 33, 08 2020. doi: 10.1007/s13347-020-00423-6.

[27] Stephen D Gantz. *The basics of IT audit: purposes, processes, and practical information*. Elsevier, 2013.

[28] Gartner, Inc. Gartner forecasts worldwide public cloud end-user spending to total $723 billion in 2025. Technical report, Gartner, Inc., November 2024. URL https://www.gartner.com/en/newsroom/press-releases/2024-11-19-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-total-723-billion-dollars-in-2025.

[29] Greg Guest, Arwen Bunce, and Laura Johnson. How many interviews are enough? an experiment with data saturation and variability. *Field methods*, 18(1):59–82, 2006.

[30] James A Hall. *Information technology auditing*. South-Western College Publishing, 2015.

[31] Alan R Hevner, Salvatore T March, Jinsoo Park, and Sudha Ram. Design science in information systems research1. *MIS quarterly*, 28(1):75–106, 2004.

[32] Information Systems Audit and Control Association (ISACA). *COBIT 2019 Framework: Governance and Management Objectives*. ISACA, Schaumburg, IL, USA, 2018. ISBN 978-1-60420-728-6. URL `https://www.isaca.org/resources/cobit`. COBIT® 2019 Framework.

[33] International Auditing and Assurance Assurance Board. International standard on auditing 315 (Revised 2019): Identifying and assessing the risks of material misstatement. Technical report, International Federation of Accountants, New York, NY, 2020. URL `https://www.iaasb.org/publications/isa-315-revised-2019-identifying-and-assessing-risks-material-misstatement`.

[34] International Auditing and Assurance Standards Board (IAASB). International standard on assurance engagements (ISAE) 3402: Assurance reports on controls at a service organization. Technical report, International Federation of Accountants (IFAC), New York, NY, 2011. URL `https://www.iaasb.org/`.

[35] International Organization for Standardization. Information technology — Cloud computing — Overview and vocabulary. Technical Report ISO/IEC 17788:2014, International Organization for Standardization, Geneva, Switzerland, 2014. URL `https://www.iso.org/standard/60544.html`.

[36] International Organization for Standardization. Information technology — Security techniques — Code of practice for information security controls based on ISO/IEC 27002 for cloud services. Technical Report ISO/IEC 27017:2015, International Organization for Standardization, Geneva, Switzerland, 2015. URL `https://www.iso.org/standard/43757.html`.

[37] International Organization for Standardization. *ISO 22301:2019: Security and resilience — Business continuity management systems — Requirements*. ISO, Geneva, Switzerland, 2nd edition, 2019.

[38] International Organization for Standardization. Information technology — Security techniques — Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors. Technical Report ISO/IEC 27018:2019, International Organization for Standardization, Geneva, Switzerland, 2019. URL `https://www.iso.org/standard/76037.html`.

[39] International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). ISO/IEC 27001:2022 – Information security, cybersecurity and privacy protection — Information security management systems — Requirements. International Standard, 2022. URL `https://www.iso.org/standard/27001`. ISO/IEC 27001:2022, Edition 3; published by ISO/IEC.

[40] Yvonna S Lincoln and Egon G Guba. *Naturalistic inquiry*. sage, 1985.

[41] Peter Mell, Tim Grance, et al. The nist definition of cloud computing. 2011.

[42] National Cyber Security Centre. Cloud security guidance. Technical report, National Cyber Security Centre, London, United Kingdom, 2022. URL `https://www.ncsc.gov.uk/collection/cloud`.

[43] Organisation for Economic Co-operation and Development (OECD). Digital security risk management. Online, 2026. URL `https://www.oecd.org/en/topics/digital-security-risk-management.html`. Accessed: 2026-01-22.

[44] Ken Peffers, Tuure Tuunanen, Marcus A Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of management information systems*, 24(3):45–77, 2007.

[45] Julia Pohle. *Digital sovereignty. A new key concept of digital policy in Germany and Europe*. Berlin: Konrad-Adenauer-Stiftung, 2020.

[46] Julia Pohle and Thorsten Thiel. Digital sovereignty. *Internet Policy Review*, 9(4), December 2020. ISSN 2197-6775. doi: 10.14763/2020.4.1532. URL `http://dx.doi.org/10.14763/2020.4.1532`.

[47] Michael Power. *The audit society: Rituals of verification*. OUP Oxford, 1997.

[48] Public Company Accounting Oversight Board. Auditing standards. Technical report, Public Company Accounting Oversight Board, Washington, DC, 2023. URL https://pcaobus.org/oversight/standards/auditing-standards.

[49] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202 of *Proceedings of Machine Learning Research*, pages 28492–28518. PMLR, 2023. URL https://proceedings.mlr.press/v202/radford23a.html.

[50] Huw Roberts. Digital sovereignty and artificial intelligence: a normative approach. *Ethics and information technology*, 26(4):70, 2024.

[51] Haroon Sheikh. European digital sovereignty: A layered approach. *Digital Society*, 1(3):25, 2022.

[52] SUSE. Cloud sovereignty framework self assessment, January 2026. URL https://www.suse.com/cloud-sovereignty-framework-assessment/. Online zelfevaluatie-instrument gebaseerd op het EU Cloud Sovereignty Framework (SEAL-niveaus).

[53] The Institute of Internal Auditors. Global internal audit standards. Technical report, The Institute of Internal Auditors, Lake Mary, FL, 2024. URL https://www.theiia.org/en/standards/2024-standards/global-internal-audit-standards/.

[54] Tweede Kamer der Staten-Generaal. Digitale soevereiniteit Rijksoverheid en besluitvorming bij migraties naar buitenlandse IT-diensten. Technical Report Kamerbrief 2025D06452, Tweede Kamer der Staten-Generaal, 2025. URL https://debatdirect.tweedekamer.nl/2025-02-19/digitale-zaken/suze-groenewegzaal/digitale-soevereiniteit-rijksoverheid-en-besluitvorming-bij-migraties-naar-buitenlandse-it-diensten-14-00.

[55] U.S. Congress. H.r. 4943 – clarifying lawful overseas use of data act (cloud act), 115th congress. Legislative Text, Congress.gov, Feb 2018. URL https://www.congress.gov/bill/115th-congress/house-bill/4943. Introduced in the House of Representatives, 115th Congress.

[56] Viswanath Venkatesh and Fred Davis. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46:186–204, 02 2000. doi: 10.1287/mnsc.46.2.186.11926.

[57] Lukas von Ditfurth. The european union's pursuit of digital sovereignty through legislation. *J. Intell. Prop. Info. Tech. & Elec. Com. L.*, 16:286, 2025.

[58] Basie Von Solms and Rossouw Van Niekerk. From information security to cyber security. *Computers & Security*, 38:97–102, 2013. doi: 10.1016/j.cose.2013.04.004. URL https://doi.org/10.1016/j.cose.2013.04.004.

Appendix A

Interview guide

This appendix contains the interview guide used in the first evaluation round (Round 1) of this study. The guide was administered to seven IT audit and Technology Risk professionals (R01 to R07) recruited through EY Netherlands. Each interview followed the same structure to ensure comparability across respondents, while allowing sufficient flexibility for respondents to elaborate on points of professional interests.

The guide is organized across nine sections: consent and respondent profiling, a sovereignty primer to establish shared vocabulary, a framework content review across the five control areas, construct-specific open questioning for each of the four TAM constructs (Perceived Usefulness, Perceived Ease of Use, Social Norm, and Behavioral Intention), a format preference section, and closing remarks. The Likert questionnaire was administered after the open qualitative questioning for each construct, as described in section 3.5.3. Interviews were conducted in Dutch and lasted approximately sixty minutes.

Expert interview guide

Technology Acceptance Model Validation

Digital Sovereignty Audit Framework

Operational Reversibility Control Form

1. Purpose and structure of this guide

This guide supports the validation of the Operational Reversibility control form, one component of the Digital Sovereignty Audit Framework developed as part of this master thesis. The guide is used in five structured interviews with IT audit and Technology Risk professionals. Its purpose is to assess how these professionals evaluate the practical relevance and usability of the framework. The data collected through these interviews address Sub-question 5 of the research: how do IT audit and Technology Risk professionals evaluate the usability and practical relevance of the proposed digital sovereignty assessment framework?

The session has 9 parts, each with a defined purpose and time allocation.

Part	Duration	Content	Purpose
1	5 min	Consent and professional background	Establish respondent profile; confirm eligibility
2	3 min	Sovereignty context	Provide sufficient context to evaluate the framework; bridge knowledge gap
3	20 min	Framework review	Content validity: completeness, audit questions, evidence requirements
4	10 min	Perceived Usefulness questions	Qualitative depth on audit value and risk coverage (TAM: PU)
5	10 min	Perceived Ease of Use questions	Clarity, formulation, and practical handleability (TAM: PEOU)
6	5 min	Social Norm	Peer expectations and professional endorsement
7	5 min	Behavioral Intention	Adoption intent
8	2 min	Format Preference	Format preference
9	5 min	Closing	Final remarks and administrative closure

2. Covering Note to Respondent

Send this note 48 hours before the session together with the framework document. Ask the respondent to read through it and note initial reactions. Do not ask them to complete or formally assess the document in advance.

Dear [Name],

Thank you for agreeing to participate in this interview as part of my master's thesis research at Leiden University. The session will take approximately 60 minutes and will take place on [date] via [platform/location].

My research focuses on developing a structured audit framework that enables IT auditors to assess digital sovereignty in cloud-dependent organizations. Specifically, I have developed a control form for the Operational Reversibility dimension, which assesses whether an organization can exit its current cloud provider arrangement without catastrophic loss of continuity or data. I am now in the validation phase and would like your professional judgment on whether the framework is practically useful and easy to apply.

Attached you will find the Operational Reversibility Control Form. Please read through it before our session. You do not need to complete it or form a final view. Simply familiarize yourself with the structure and note any initial reactions or questions.

The session consists of a short explanation of the sovereignty context relevant to the framework, a structured discussion based on your audit experience, and a brief questionnaire at the end. You do not need to be a specialist in digital sovereignty. I am interested in your judgment as an IT audit or Technology Risk professional.

All responses will be anonymized and used solely for academic research. With your permission, I would like to record the session for transcription purposes. The recording will not be shared outside this research.

Kind regards,

Zaynab Musawi
Leiden University

3. Part 1: Introduction and professional background (5 minutes)

Interviewer script

"Thank you for making time for this session. My name is Zaynab Musawi and I am conducting this interview as part of my master's thesis research at Leiden University. I am developing a structured audit framework for assessing digital sovereignty in cloud-dependent organizations. Today I am asking for your professional evaluation of one specific component: the Operational Reversibility control form. I am interested in your judgment as an IT audit or Technology Risk professional, not in testing your knowledge of sovereignty theory. I will provide the relevant context as we go. The session will take approximately 60 minutes. Do you consent to being recorded?"

Background questions

#	Question	Response
B1	What is your current role and what is the primary focus of your work in relation to IT audit or Technology Risk?	
B2	Approximately how many years of experience do you have in IT audit, Technology Risk, or a related discipline?	
B3	Have you been involved in IT audits or risk assessments at cloud-dependent organizations? Roughly how many such engagements in the past three years?	
B4	Have you encountered cloud-related vendor lock-in, data residency, or exit planning issues in your audit work?	
B5	Are you familiar with the DICTU Toetsingsinstrument, the EC Cloud Sovereignty Framework, or similar documents? If yes, in what context?	

4. Part 2: Sovereignty context (2 minutes)

Script

“Digital sovereignty, in the context of this research, refers to the degree to which an organization maintains genuine control over its own digital infrastructure and data when using cloud services. The practical problem is this: when an organization moves its systems to a cloud provider, it often loses control in ways it does not immediately notice. It may not be able to switch providers because its data is in a proprietary format. It may not know whether its data is legally accessible by non-EU authorities under laws such as the US CLOUD Act. And if the provider fails or is acquired, the organization may not have the documentation or technical capability to keep its services running with a different provider.

The specific dimension I am asking you to evaluate today is called Operational Reversibility, sometimes referred to as exit capability. The central audit question is: if this organization decided tomorrow that it needed to leave this cloud provider, could it actually do that without a catastrophic loss of service, data, or operational knowledge? This is fundamentally an audit question. It is about whether the right controls are in place and actually work: an exit plan, data portability, documentation, continuity during migration, and periodic testing of exit scenarios.

The framework you have read translates this question into a structured set of audit controls with five maturity levels (M1 to M5) and four operating effectiveness levels (E1 to E4). I would like your assessment of whether this is a useful and practical tool for IT auditors.”

5. Part 3: Framework review (20 minutes)

These questions address the content validity of the framework.

#	Question	Notes
Q1	Looking at the five control areas in this form (1 through 5), do you think they collectively cover the most important aspects of cloud exit capability from an audit perspective? Is anything material missing?	
Q2	For control 1 (Exit Plan and Contractual Framework), look at the five audit questions. Are these the right questions to ask a client organization in an audit context? Are they formulated clearly and unambiguously?	
Q3	The framework uses a cumulative maturity model from M1 to M5, where each higher level requires all criteria of lower levels plus new ones. Does this logic work from an audit perspective, or does it create problems in practice?	
Q4	The form distinguishes between design maturity (Part A) and operating effectiveness (Part B) for each control area. Is this distinction useful and meaningful for a real IT audit engagement, or is it unnecessarily complex?	
Q5	At each maturity level, the form specifies what evidence the organization must provide. Are these evidence requirements realistic? Would a typical client organization be able to produce this evidence within a standard engagement?	

6. Part 4: Perceived Usefulness (10 minutes)

TAM construct: Perceived Usefulness. These questions assess whether the respondent believes the framework adds genuine audit value. Questions are framed from an audit practice perspective to accommodate the knowledge gap.

#	Question	Notes
PU-Q1	Without a framework like this one, how would you currently approach assessing a client's cloud exit capability? What would your existing methodology give you, and what would it miss?	
PU-Q2	If you found that a client scored M1 or M2 on control 1 (no formal exit plan), how useful would that finding be for your audit report and for the client? Would it be a finding that management takes seriously?	
PU-Q3	Does this framework identify risks that are not already covered by the IT audit standards you currently use, such as ISO 27001, or your firm's standard cloud audit methodology?	
PU-Q4	Is this framework primarily useful for compliance engagements or for risk-based audits? Which type of engagement do you see as its primary use case?	

Interviewer script:
I am now going to ask you to rate eight statements about the framework you have reviewed today. Please respond based on your experience and impression from this session. Use a scale from 1 (Strongly Disagree) to 5 (Strongly Agree). There are no right or wrong answers.

Scale:

1 = Strongly Disagree

2 = Disagree

3 = Neutral

4 = Agree

5 = Strongly Agree

Construct 1: Perceived Usefulness (PU)

Adapted from Davis (1989). Assesses the degree to which the respondent believes using this framework would enhance their professional performance in cloud sovereignty audit engagements.

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
PU1	This framework would help me identify cloud exit risks that I would otherwise miss or underweight in a standard audit.	☐	☐	☐	☐	☐

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
PU2	Using this framework would improve the quality and completeness of my audit conclusions on vendor lock-in and exit capability.	□	□	□	□	□
PU3	This framework provides sufficient structure to produce a defensible, evidence-based audit opinion on operational reversibility.	□	□	□	□	□
PU4	This framework adds value beyond what existing audit standards such as ISO 27001, or NIS2 already require me to assess.	□	□	□	□	□

7. Part 5: Perceived Ease of Use (10 minutes)

TAM construct: Perceived Ease of Use. Focus on practical handleability: clarity of audit questions, realism of evidence requirements, and accessibility of language for auditors without deep sovereignty expertise.

#	Question	Notes
PEOU-Q1	Without specific training in this framework, how much effort did it take you to understand what each control area is asking you to do? Was the structure intuitive?	
PEOU-Q2	Look at the maturity level descriptions for control 1. Are the criteria for each level formulated clearly enough that two different auditors reading them would reach the same assessment for the same client?	
PEOU-Q3	Are the evidence requirements at each maturity level specific enough that you would know exactly what to request from a client organization? Or are any too vague?	
PEOU-Q4	Is the language in this form accessible to an IT auditor who does not have deep expertise in cloud architecture or digital sovereignty theory?	

Interviewer script:
I am now going to ask you to rate eight statements about the framework you have reviewed today. Please respond based on your experience and impression from this session. Use a scale from 1 (Strongly Disagree) to 5 (Strongly Agree). There are no right or wrong answers.

Scale:

1 = Strongly Disagree

2 = Disagree

3 = Neutral

4 = Agree

5 = Strongly Agree

Construct 2: Perceived Ease of Use (PEOU)

Adapted from Davis (1989). Assesses the degree to which the respondent believes using this framework would be free of cognitive effort.

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
PEOU1	The audit questions in this form are clearly and unambiguously formulated. I know exactly what I am being asked to verify.	☐	☐	☐	☐	☐

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
PEOU2	The distinction between design maturity (M1 to M5) and operating effectiveness (E1 to E4) is easy to understand and apply.	□	□	□	□	□
PEOU3	The evidence requirements at each maturity level are specific enough that I know exactly what to request from a client organization.	□	□	□	□	□

8. Part 6: Social Norm (10 minutes)

These questions assess how the framework…

#	Question	Notes
SN-Q1	In your expertise, do partners, managers, or senior colleagues at your firm or organization talk about digital sovereignty or cloud exit capability as an audit topic? Or is it still considered niche?	
SN-Q2	If a leading professional body (such as the NBA, NOREA, or a Big 4 firm) were to formally endorse or adopt a framework like this as part of standard cloud audit methodology, would that influence your likelihood of using it?	
SN-Q3	Do your clients, or the organization you audit, would welcome or resist this type of assessment? Would they see it as adding value or as additional compliance burden?	

Interviewer script:
I am now going to ask you to rate eight statements about the framework you have reviewed today. Please respond based on your experience and impression from this session. Use a scale from 1 (Strongly Disagree) to 5 (Strongly Agree). There are no right or wrong answers.

Scale:

1 = Strongly Disagree

2 = Disagree

3 = Neutral

4 = Agree

5 = Strongly Agree

Construct 3: Behavioral Intention to Use (BI)

Adapted from Davis, Bagozzi and Warshaw (1989). Assesses the degree to which the respondent intends to use the framework in future professional engagements.

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
SN1	People whose opinion I value in my professional environment would support the use of a structured sovereignty assessment framework like this in cloud audits.	□	□	□	□	□
SN2	II believe that using this type of framework would be consistent with what is expected of a professional IT	□	□	□	□	□

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
	auditor in today's regulatory environment (NIS2, DORA, GDPR).					
SN3	If leading professional bodies or prominent firms in my field adopted this framework, I would be more inclined to use it myself.					

9. Part 7: Behavioral Intention (5 minutes)

These questions assess how the framework fits with existing audit standards and current practice, and identify practical integration barriers.

#	Question	Notes
BI-Q1	Based on your experience with this form today, would you consider using it, or a defined version, in a future cloud-related audit engagement? What conditions would need to be met for you to actually use it?	
BI-Q2	Would you recommend this framework to a colleague conducting a cloud exit or vendor lock-in risk assessment? What would be your primary reason for recommending or not recommending it?	

Interviewer script:
I am now going to ask you to rate eight statements about the framework you have reviewed today. Please respond based on your experience and impression from this session. Use a scale from 1 (Strongly Disagree) to 5 (Strongly Agree). There are no right or wrong answers.

Scale:

1 = Strongly Disagree

2 = Disagree

3 = Neutral

4 = Agree

5 = Strongly Agree

Construct 3: Behavioral Intention to Use (BI)

Adapted from Davis, Bagozzi and Warshaw (1989). Assesses the degree to which the respondent intends to use the framework in future professional engagements.

#	Statement	1 Strongly Disagree	2 Disagree	3 Neutral	4 Agree	5 Strongly Agree
BI1	I would use this framework, or a refined version of it, in a future cloud sovereignty or vendor lock-in audit engagement.	□	□	□	□	□
BI2	I would recommend this framework to a colleague conducting a cloud exit or operational reversibility risk assessment.	□	□	□	□	□

10. Part 8: Format preference question (1 minute)

FP-1

In what format would you ideally want to receive and use a framework like this in practice? Please think about your actual workflow when conducting an IT audit engagement.

A. Word or PDF document: printable, can be used offline, familiar format for audit workpapers.
B. Excel spreadsheet: scoring cells, automatic calculations, integrates into existing workpaper templates.
C. Digital form or canvas: easy to fill in on screen, auto-saves responses
D. Interactive web tool or dashboard: online shareable, with client, visual scoring overview
E. Other (please specify):

11. Part 9: Closing (5 minutes)

Thank the respondent for their time and contribution. Confirm how results will be used.

Closing script:

Thank you very much for your time and your candid responses. Your input is directly useful for refining the framework before the final version is submitted as part of my thesis. The findings from all five interviews will be reported in aggregate and anonymized form. You will not be identified by name or organization in any publication.

If any further questions come to mind after the session, please feel free to contact me at z.musawi@umail.leidenuniv.com. I will share a summary of the overall findings with you once the analysis is complete, if you would like.

Appendix B

DSAF Audit Control Form

This appendix contains the complete DSAF Operational Reversibility control instrument in its Version 2 form. It documents all five control areas (C1 through C5), including the sovereignty objective, sovereignty link, audit questions, Design Maturity criteria (M1 to M5), Operating Effectiveness criteria (E1 to E4), suggested evidence per maturity level, and the Responsible Role table for each control area.

The instrument is also available as a downloadable spreadsheet workbook at `https://zaynab-m.github.io/DSAF/`, which is the primary format intended for use in audit engagements. The version reproduced here serves as a static reference for readers of this thesis.

Digital Sovereignty Audit Framework (DSAF) — Operational Reversibility

Technology Domain: Operational Reversibility | Exit Strategy and Reversibility | Version 1.1

How to Use This Form

This form can be used to assess whether an organization can exit its current cloud provider arrangement without catastrophic loss of continuity, data, or institutional knowledge. The guiding question from the DICTU Toetsingsinstrument (2026, criterion 3.4) is: 'Can a third party take over the service upon provider failure or contract termination, based on available documentation and design?' The EuroStack framework (2025, Criterion 2.3) adds: 'A contractually binding exit plan is mandatory.'

The form covers all five control areas. The audit questions (the numbered Yes/No questions) are the starting point: they define what the auditor must verify. All five control areas are assessed on two dimensions: design maturity (whether the control is designed to achieve its objective) and operating effectiveness (whether the designed control actually operated as intended during the audit period).

Part A: Design Maturity (M1–M5)	Part B: Operating Effectiveness (E1–E4)
Assesses whether the control is designed to achieve its objective. The question is: if this control operates exactly as written, would it mitigate the risk?	Assesses whether the designed control operated as intended during the audit period. The question is: did the control work in practice?

Framework Origin

The five maturity levels in this form correspond directly to the five DICTU sovereignty levels for Criterion 3.4. The definitions below are reproduced verbatim. Source: DICTU Toetsingsinstrument Soevereiniteit Clouddiensten v1.0.1 (January 2026), p. 20.

Although presented here in generic form, the definitions of these levels have been further specified per norm and subdimension in the remainder of this document, providing additional context to make them more relevant and applicable to each specific assessment area.

The sovereignty dimensions, control objectives, and structure of this framework are derived from DICTU, relevant ISO standards, the EuroStack Buy European Framework, and broader sovereignty frameworks. These sources have been consolidated and operationalized into auditable control language tailored to IT auditing in the context of operational reversibility.

Quantitative tresholds (e.g. Infrastructure-as-Code (IaC) coverage, review intervals) are design choices that make the quantitative DICTU maturity levels measurable.

Step-by-Step Workflow

Step 1	Read the control area. For each of the five control areas (C1-C5), read the Risk, Objective, and Sovereignty Link at the top of the sheet before starting the audit questions, to understand what is being tested and why.
Step 2	Answer the Audit Questions. Work through the numbered Yes/No questions for that control area. These define the minimum evidence the auditor must request and review; they do not yet determine a score.
Step 3	Score Part A (Design Maturity, M1-M5). Starting from M1, check whether ALL criteria at each level are satisfied, moving upward. Tick the HIGHEST level for which every cumulative criterion is met. Evidence listed is suggested, not mandatory; document your rationale if alternative evidence was used.
Step 4	Score Part B (Operating Effectiveness, E1-E4). Independently of Part A, assess whether the control has actually been demonstrated in practice (tested, exercised, executed), not only designed. Part A and Part B scores are never averaged together.
Step 5	Document the conclusion and ownership. Complete the Maturity/Effectiveness Level Assigned fields and the Rationale/Finding box, and identify the responsible role for follow-up using the Responsible Role table on each control sheet.
Step 6	Repeat for all five control areas, then complete the Summary tab. Enter the assigned Design Maturity (M1-M5) and Operating Effectiveness (E1-E4) level for each control area; the average scores and decision-rule outcome are calculated automatically.
Step 7	Record the overall conclusion. Use the Decision Rules on the Summary tab as guidance, document the basis for the overall conclusion, and list priority recommendations with named owners and deadlines.

Note: a level applies only if every cumulative criterion at that level is satisfied. If evidence is missing, ambiguous, or the situation is atypical, apply professional judgement and document the rationale rather than forcing a score.

Part A: Design Maturity (M1–M5)

Level	Designation	Description
M1	**No Reversibility**	No reversibility: The service is not transferable; continuity in the event of provider failure is not guaranteed.
M2	**Limited Reversibility**	Limited reversibility: Documentation is available to a limited extent, but transferring to a third party is complex and not fully feasible.
M3	**Contractual Reversibility**	Contractual reversibility: The service is largely transferable; documentation and design are available at a high level to enable a third party to independently operate the service over time.
M4	**Active Reversibility**	Active reversibility: The service is fully transferable; documentation and design are sufficiently detailed to enable a third party to independently operate the service.
M5	**Full Operational Reversibility**	Full operational reversibility: The service is fully transferable; documentation and design are sufficiently detailed and tested to enable a third party to independently operate the service.

Part B: Operational Effectiveness (E1–E4)

Level	Designation	Description
E1	**Not Effective**	Control has never been tested or operated
E2	**Partially Effective**	Control exists but untested or incomplete
E3	**Largely Effective**	Control demonstrated with minor gaps
E4	**Fully Effective**	Control fully demonstrated and independently verified

Definitions and Key Terms

Glossary of terms used in the Digital Sovereignty Audit Framework: Operational Reversibility

Framework & Sources	
Digital Sovereignty Audit Framework (DSAF)	The structured audit framework used in this workbook to assess whether an organisation can exit a cloud provider without catastrophic loss of continuity, data, or institutional knowledge. It covers five control areas, each scored on design and operating effectiveness.
Operational Reversibility	The capacity to transfer a cloud service to another provider (or bring it back in-house) in an orderly way, so continuity is maintained on provider failure or contract termination. The central theme this framework assesses.
Digital Sovereignty	The degree of control an organisation retains over its data, systems, and technology choices, including the ability to change provider without undue dependence. 'Technology Sovereignty' is the specific DICTU dimension (3.4) on reversibility.
DICTU	The Dutch government's shared ICT service organisation. Its assessment instrument for cloud sovereignty provides the maturity levels and criteria on which this framework is based.
DICTU Toetsingsinstrument	The DICTU 'Toetsingsinstrument Soevereiniteit Clouddiensten' (assessment instrument for cloud-service sovereignty). Criterion 3.4 and its five sovereignty levels underpin the design-maturity scale used here.
EuroStack (Buy European Framework)	A European framework promoting digital autonomy and use of European infrastructure. Its Criterion 2.3 ('a contractually binding exit plan is mandatory') is referenced as a sovereignty requirement in this framework.
Sovereignty link	The field at the top of each control sheet that maps the control to the underlying sovereignty requirement (e.g. DICTU 3.4, EuroStack 2.3), showing why the control matters.

Assessment Structure	
Control area	One of the five themes assessed (C1–C5): Exit Plan, Data Portability, Knowledge Transfer, Continuity of Operations, and Testing Programme. Each has its own sheet.
Audit question	A numbered Yes/No question (e.g. 1.1, 1.2) defining what the auditor must verify for a control. The starting point of the assessment.
Design Effectiveness (Maturity)	Whether a control is designed to achieve its objective: if it operated exactly as written, would it mitigate the risk? Scored on the M1–M5 maturity scale (Part A).
Operating Effectiveness	Whether the designed control actually operated as intended during the audit period: did it work in practice? Scored on the E1–E4 scale (Part B).
Maturity level (M1–M5)	The design-effectiveness rating. Levels are cumulative: a higher level requires all lower-level criteria to be met. The auditor ticks the highest fully satisfied level.
Effectiveness level (E1–E4)	The operating-effectiveness rating, from Not Effective (E1) to Fully Effective (E4), based on evidence that the control operated during the period.
Cumulative	A scoring principle: a given level can only be assigned if all criteria of every lower level are also satisfied.

Maturity Levels (Design, Part A)	
M1: No Reversibility	The service is not transferable; continuity on provider failure is not guaranteed.
M2: Limited Reversibility	Documentation exists to a limited extent, but transfer to a third party is complex and not fully feasible.
M3: Contractual Reversibility	The service is largely transferable; documentation and design are available at a high level for a third party to operate it over time.
M4: Active Reversibility	The service is fully transferable; documentation and design are detailed enough for a third party to operate it independently.
M5: Full Operational Reversibility	As M4, and the documentation and design have been tested, confirming a third party can independently operate the service.

Effectiveness Levels (Operating, Part B)	

E1: Not Effective	The control has never been tested or operated; evidence of operation is absent.
E2: Partially Effective	The control exists but is untested or incomplete in operation.
E3: Largely Effective	The control has been demonstrated to operate, with only minor gaps.
E4: Fully Effective	The control is fully demonstrated and independently verified.
Exit & Continuity	
Exit plan	A formally approved, contractually supported plan enabling an orderly transition to an alternative provider within a defined RTO. The prerequisite control (C1).
Vendor lock-in	Dependence on a single provider that makes leaving prohibitively difficult or costly. Without an exit plan it is a certain outcome rather than a risk.
RTO (Recovery Time Objective)	The maximum acceptable time to restore the service with an alternative provider in an exit scenario; should be formally defined and approved.
Rollback	A defined procedure to revert to the original provider or state if a migration fails, including trigger criteria and execution steps.
Parallel operation period	A phase during migration in which old and new environments run simultaneously to safeguard continuity before cut-over.
Business Impact Analysis (BIA)	An assessment of how critical a service is to the organisation, used to validate that the RTO matches the service's importance.
Service Level Agreement (SLA)	The contracted service performance levels; cross-checked to ensure continuity and degradation tolerance are maintained during migration.
Service degradation tolerance	The level of reduced service the organisation can accept during a migration, defined per service and validated against SLAs.
Data & Technical Portability	
Data portability	The ability to export data and workloads to another environment without significant effort, cost, or loss, ideally in open formats.
Open standard / open format	A non-proprietary, publicly documented data or interface format that a third party can read and process without the original provider's tooling.
Egress fees	Charges levied by a provider for extracting data out of its platform; punitive egress fees can make exit economically prohibitive.
Infrastructure as Code (IaC)	Infrastructure defined and managed through version-controlled code (e.g. Terraform), enabling reproducible deployment elsewhere. Coverage % indicates how much of production is captured.
Runbook	Documented step-by-step operational procedures enabling someone other than the original engineer to run or recover the service.
Key person dependency	A situation where operation relies on knowledge held by one or few individuals, creating risk if they become unavailable.
Disaster Recovery (DR)	Arrangements (including backups and replicas) for restoring service after a major failure; their portability and format are assessed for exit readiness.
Import validation	Confirming that exported data can actually be loaded and used in an alternative environment without loss, not merely that the export completed.
Testing & Assurance	
Tabletop exercise	A structured, discussion-based simulation of an exit scenario to test the plan without live migration; must include an unplanned trigger to count.
Unplanned (exit) trigger	An involuntary exit cause (e.g. provider insolvency, sanctions, or non-EU acquisition) that the test scenario must include; a voluntary planned migration alone is insufficient.
Partial migration test	A live test that actually migrates part of the service to an alternative environment while maintaining continuity, providing proof of execution.
Live exercise	Any real (non-simulated) execution of exit or migration steps, demonstrating capability beyond intent.
Findings register	A log of issues identified in tests or exercises, with owners and due dates, used to track remediation to closure.
Independent oversight / independence test	Verification that testing or documentation was executed or reviewed by someone who did not author it; the defining criterion for the highest levels.
Roles & Documents	

CIO (Chief Information Officer)	The senior executive who, in this framework, is the minimum authority for approving the exit plan and the RTO.
IT Director	An alternative minimum approval authority for the exit plan where a CIO role is not present.
Data Processing Agreement (DPA)	The contractual agreement governing data handling with the provider; often where exit rights, export rights, and egress terms are recorded.

Control 1: Exit plan and contractual exit framework

Risk	Without a formal exit plan the organization is structurally dependent on the current provider. In the event of forced exit, insolvency, geopolitical sanctions, acquisition by a non-EU entity, the absence of a plan results in operational crisis, uncontrollable data loss, or contractual imprisonment. Vendor lock-in is not a risk in this scenario but a certain outcome.
Objective	Verify that a formally approved, contractually supported exit plan exists enabling an orderly transition to an alternative provider within a defined Recovery Time Objective (RTO).
Sovereignty link	Technology Sovereignty (DICTU 3.4) | EuroStack Criterion 2.3: 'A contractually binding exit plan is mandatory' | Guiding question: 'Can a third party take over the service upon provider failure or contract termination?'

Audit Questions

#	Audit Question	Yes/No	Notes
1,1	Does a formally approved, version-controlled exit plan document exist for this cloud service?		
1,2	Does the cloud service contract contain explicit exit rights, including data export rights, provider transition assistance obligation, and absence of punitive egress fees at termination?		
1,3	Is a Recovery Time Objective (RTO) for the exit scenario formally defined and approved by the CIO or equivalent authority?		
1,4	Does the exit plan identify at least two credible alternative providers (including at least one EU-domiciled alternative) with a documented compatibility assessment?		
1,5	Does the exit plan have a named owner (by role), a defined review cycle, and evidence of at least one review within the past 24 months?		

Part A: Design Maturity (M1–M5)

Instructions: Tick (✓) the HIGHEST level at which ALL criteria are fully satisfied. Levels are CUMULATIVE.

Level		Definition	Criteria (cumulative)	Suggested Evidence to Provide	Reached (✓)	Notes
M1	No Reversibility	Service not transferable; continuity upon provider failure not guaranteed.	1. No exit plan exists in any form. 2. No contractual exit rights in the service agreement. 3. No consideration has been given to what happens if the provider fails.	→ *Absence of a named exit plan document confirms M1. If no evidence is produced, this level applies by default.*		
M2	Limited Reversibility	Documentation is limited; transfer to a third party is complex and not fully possible.	1. Cumulative: the organization has progressed beyond M1 (the M1 absence conditions no longer apply), PLUS: 2. An exit plan document exists in some form (draft, informal note, or process document). 3. The plan is significantly outdated (more than 24 months without review) or lacks formal approval. 4. Contractual exit rights are absent or contain terms that make exit economically prohibitive.	→ *Exit plan document (draft or otherwise); note: outdated or unapproved plans confirm M2.* → *Contract text showing absence of exit rights or presence of prohibitive exit fees.*		
M3	Contractual Reversibility	Service largely transferable; high-level documentation available for a third party to operate the service in due course.	1. Cumulative: all M2 criteria satisfied, PLUS: 2. Exit plan is formally approved by an appropriate authority (minimum: IT Director or CIO). 3. Contract contains explicit exit rights including data export right, no punitive egress fees, and minimum transition assistance obligation. 4. RTO for the exit scenario is defined (though may not yet be tested or validated).	→ *Exit plan with approval record, dated within 24 months.* → *Contract or DPA section confirming explicit exit rights and no punitive egress fees.* → *RTO definition document or section in exit plan approved by CIO.*		
M4	Active Reversibility	Service fully transferable; documentation sufficiently detailed for a third party to operate the service independently.	1. Cumulative: all M3 criteria satisfied, PLUS: 2. Exit plan is detailed and current (reviewed within 12 months), containing: migration phases, phase-gate criteria, parallel operation period, data synchronization approach, and rollback trigger conditions. 3. At least two credible alternatives identified with compatibility assessment; at least one EU-domiciled alternative. 4. Named owner by role; formal review cycle defined and evidenced. 5. RTO has been validated against the business impact analysis of the relevant cloud service.	→ *Exit plan with review date within 12 months, phase diagram, and gate criteria.* → *Alternative provider assessment document with EU jurisdiction confirmation.* → *Business Impact Analysis linking RTO to criticality of service.* → *Owner assignment and review cycle record.*		
M5	Full Operational Reversibility	Service fully transferable; documentation sufficiently detailed and tested for a third party to operate independently.	1. Cumulative: all M4 criteria satisfied, PLUS: 2. Exit plan has been validated through at least a structured tabletop exercise simulating an unplanned exit trigger (provider insolvency, sanctions, or non-EU acquisition). 3. At least one alternative provider is pre-qualified or has a contract on standby. 4. Exit plan and RTO reviewed or independently assessed within 12 months.	→ *Tabletop exercise report: date, scenario (must include unplanned trigger), participants, findings, remediation.* → *Pre-qualification documentation or standby contract with alternative EU provider.* → *Independent review or internal audit report on exit plan adequacy.*		
DESIGN MATURITY CONCLUSION		**Maturity Level Assigned:**		Rationale / Finding:		

Part B: Operational Effectiveness (E1–E4)

Instructions: Assess whether the designed control operated as intended during the audit period.

	Level	What the Organization Must Demonstrate	Audit Procedure	Audit Guidance	Reached (✓)	Notes
E1	Not Effective	1. Exit plan exists on paper only. 2. Plan has never been operationally referenced or tested. 3. Contractual exit rights unverified.	1. Inspect exit plan currency and formal approval. 2. Interview exit plan owner: can they explain the plan without looking at it? 3. Check contract for exit rights clause.	*If the exit plan has not been reviewed in 24 months or key stakeholders are unaware of it, rate E1 regardless of maturity score.*		
E2	Partially Effective	1. Exit plan reviewed within 24 months. 2. Contractual exit rights confirmed current. 3. RTO approved by CIO. 4. Alternative providers identified.	1. Review evidence of last plan update and approval. 2. Confirm contractual exit rights are still valid. 3. Interview exit plan owner on RTO rationale and plan awareness.	*Plan is maintained but has never been tested. This is the most common state. Document as observation and recommend tabletop exercise as priority action.*		
E3	Largely Effective	1. Exit plan reviewed within 12 months. 2. Tabletop exercise conducted within 24 months. 3. Findings from exercise addressed and documented.	1. Review tabletop exercise report: scope, scenario, participants, findings. 2. Verify findings register and confirm all items are closed or on schedule. 3. Confirm scenario included an unplanned trigger; planned migration only is insufficient for E3.	*E3 requires evidence that the plan has been stress-tested, not only maintained. Unplanned triggers must feature in the scenario.*		

E4	Fully Effective	1. Exit plan tested via live exercise or partial migration. 2. Alternative provider pre-qualified or contracted on standby. 3. All prior exercise findings closed.	1. Review live exercise documentation or partial migration test evidence. 2. Verify pre-qualification or standby contract with alternative EU provider. 3. Confirm all findings from prior exercises are closed.	*E4 requires demonstrated exit capability, proof of execution, not intent. An organization that has only conducted tabletop exercises but no live test is at most E3.*		
OPERATING EFFECTIVENESS CONCLUSION		**Effectiveness Level Assigned:**		Rationale / Finding:		

Responsible Roles	
Design Maturity Level	**Responsible Role (typically accountable for evidence at this level)**
M1–M2	No named owner exists at this level (criteria reflect absence). Escalate to IT Director / CIO to initiate an exit plan.
M3	IT Director or CIO: approves the exit plan and the RTO (minimum approval authority per criteria).
M4	Named exit plan owner (by role) + CIO: owner maintains the plan and review cycle; CIO validates RTO against business impact analysis.
M5	Named exit plan owner + CIO + Internal Audit: exercise participants per the tabletop scenario; Internal Audit performs the independent review.

Operating Effectiveness Level	**Responsible Role (typically accountable for evidence at this level)**
E1	Exit plan owner (if named): interviewed to assess awareness; escalate to CIO if plan is unreviewed beyond 24 months.
E2	Exit plan owner: accountable for plan currency and RTO awareness.
E3	Exit plan owner + tabletop exercise participants (IT, Legal, Compliance as applicable): accountable for findings remediation.
E4	Exit plan owner + CIO: CIO confirms pre-qualification/standby contract; owner closes prior findings.

Control 2: Data Portability and Technical Exportability

Risk	If data cannot be exported in a usable format, exit is technically impossible even when contractually permitted. Proprietary formats, excessive egress fees, and missing export tooling create a 'data gravity trap'. EuroStack (2025) identifies punitive egress fees as 'strategic barriers that prevent a competitive European market.'
Objective	Verify that all data, primary data, backups, configurations, metadata, and logs, can be exported at any time in open, non-proprietary formats at proportionate cost, and that export procedures are documented and tested.
Sovereignty link	Technology Sovereignty (DICTU 3.1, Interoperability & Portability) | EuroStack Criterion 2.1 (data portability, no punitive egress fees) | GDPR Art. 20 (personal data portability)

Audit Questions

#	Audit Question	Yes/No	Notes
2,1	Can the organization demonstrate that all primary production data types can be exported in open, non-proprietary formats (e.g. CSV, JSON, Parquet, XML, OVA/OVF) without requiring provider-specific tooling to read the output?		
2,2	Does the cloud service contract explicitly prohibit punitive egress fees at contract termination, and does the cost modelling for exporting current data volumes at termination confirm fees are proportionate to actual transfer costs?		
2,3	Are Infrastructure-as-Code (IaC) artefacts for all production infrastructure stored in an organization-controlled version control repository, not solely in provider-managed systems?		
2,4	Are backup copies and disaster recovery replicas exportable with the same open format and portability standard as primary data, and is this verified through separate backup export testing?		
2,5	Has the organization conducted an end-to-end export test, including import validation in an alternative environment, for all primary data types within the past 24 months?		

Part A: Design Maturity (M1–M5)

Instructions: Tick (✓) the HIGHEST level at which ALL criteria are fully satisfied. Levels are CUMULATIVE.

Level		Definition	Criteria (cumulative)	Suggested Evidence to Provide	Reached (✓)	Notes
M1	No Portability	Service fully proprietary; migration of data and workloads not possible without significant effort or cost.	1. Data stored exclusively in proprietary provider formats with no open export mechanism. 2. Egress fees are prohibitive, or the organization has not assessed the termination cost. 3. No IaC exists; infrastructure configurations are undocumented or stored only in provider consoles.	→ *Provider documentation showing only proprietary formats, or the organization cannot provide any export format specification.*		
M2	Limited Portability	Partial openness; migration is limited and requires significant effort.	1. Cumulative: the organization has progressed beyond M1 (the M1 absence conditions no longer apply), PLUS: 2. Export is possible for at least some data types but requires significant provider involvement or re-formatting effort. 3. Some core data types remain in proprietary formats. 4. Egress fees have been reviewed but not modelled for a full exit scenario.	→ *Contract or DPA section listing available export formats.* → *Data inventory showing which data types have open export and which do not.*		
M3	Open Standards Partial	Open standards widely used but not for all components; portability partially limited.	1. Cumulative: M2 satisfied, PLUS: 2. All primary data types have a documented open format export option. 3. Contract confirms no punitive egress fees at termination. 4. IaC covers at least 50% of production infrastructure components in organization-controlled repositories. 5. Backup export format has been assessed, even if not yet tested.	→ *Export format specification per data type (contract or technical documentation).* → *Contract clause confirming absence of punitive egress fees.* → *IaC repository location and ownership evidence with coverage percentage estimate.*		
M4	All Open Standards	All components use open standards; high portability, not yet fully optimized.	1. Cumulative: M3 satisfied, PLUS: 2. IaC coverage exceeds 80% of all production infrastructure components. 3. Export test conducted for all primary data types within 24 months, confirming successful export without provider tooling. 4. Backup and DR replicas confirmed to use the same open format and portability standard as primary data. 5. Data egress cost modelled for a realistic full-exit scenario confirming proportionality.	→ *IaC coverage report confirming more than 80% (spot-check by auditor recommended).* → *Export test results per data type, dated within 24 months.* → *Backup export test evidence or configuration confirming same standard.* → *Egress cost model for full exit scenario.*		
M5	Maximum Portability	Most widely used open standards; maximized portability; migration easy and cost-efficient.	1. Cumulative: M4 satisfied, PLUS: 2. Full end-to-end export test completed: all data types, backups, metadata, IaC configurations, and logs. 3. Import validation in an alternative environment confirms no data loss and readability without provider tooling. 4. Results independently reviewed or included in internal audit scope.	→ *Full export test report including all data types, backups, configs, logs.* → *Import validation evidence in alternative environment (not just export success).* → *Independent review record or internal audit finding on portability controls.*		
DESIGN MATURITY CONCLUSION		**Maturity Level Assigned:**		Rationale / Finding:		

Part B: Operational Effectiveness (E1–E4)

Instructions: Assess whether the designed control operated as intended during the audit period.

Level		What the Organization Must Demonstrate	Audit Procedure	Audit Guidance	Reached (✓)	Notes
E1	Not Effective	1. Export has never been tested for any data type. 2. Organization cannot identify data format in use. 3. Egress fees at termination have not been reviewed.	1. Request list of data types and their export formats. 2. Attempt walkthrough of export for one minor data type. 3. Review contract for egress fee structure.	*If no export test has been conducted and the organization cannot explain what format their data is in, rate E1 regardless of contractual claims.*		
E2	Partially Effective	1. Primary data types confirmed in open format. 2. Some export tests conducted but coverage is incomplete. 3. IaC repository ownership confirmed.	1. Review available export test results and note which types are missing. 2. Confirm IaC repository is organization-controlled. 3. Review contract egress fee clause.	*Partial test coverage does not constitute demonstrated portability. Metadata and logs are frequently the untested gap.*		
E3	Largely Effective	1. All primary data types tested for export. 2. Export duration validated against RTO. 3. Backup portability confirmed.	1. Review export test results for all data types. 2. Verify duration is within RTO. 3. Perform readability test: attempt to open or parse one exported file using a non-provider tool.	*E3 requires that exported data has been validated as usable, not only that export completed. A file that exports but cannot be opened in a standard tool is a finding.*		
E4	Fully Effective	1. Full end-to-end export test: all types, backups, IaC, logs. 2. Import validation in alternative environment confirms no data loss. 3. Independent review of portability controls.	1. Review full export test report. 2. Confirm import validation in alternative environment (must include at least one critical data type imported and validated). 3. Review independent assessment scope.	*E4 requires import validation in a target environment. 'Export ran successfully' without import validation is at most E3.*		
EFFECTIVENESS CONCLUSION		**Effectiveness Level Assigned:**		Rationale / Finding:		

Responsible Roles

Design Maturity Level	Responsible Role (typically accountable for evidence at this level)
M1–M2	No named owner exists at this level. Escalate to Infrastructure/Platform Lead to assess export and IaC posture.
M3	Infrastructure/Platform Lead: owns IaC repository and data export format documentation (>50% coverage).
M4	Infrastructure/Platform Lead + Data Owner: Infrastructure Lead delivers IaC coverage >80% and export tests; Data Owner confirms backup/DR parity.
M5	Infrastructure/Platform Lead + Internal Audit: Infrastructure Lead executes full export/import test; Internal Audit reviews results.

Operating Effectiveness Level	Responsible Role (typically accountable for evidence at this level)
E1	Infrastructure/Platform Lead: accountable for identifying data formats in use.
E2	Infrastructure/Platform Lead: accountable for closing untested data types (commonly metadata/logs).
E3	Infrastructure/Platform Lead: accountable for RTO-aligned export duration and readability testing.
E4	Infrastructure/Platform Lead + Internal Audit: Internal Audit reviews import validation evidence.

Control 3: Knowledge Transfer and Operational Documentation

Risk	Tacit knowledge held by a small number of key employees or embedded in provider-specific tooling creates human lock-in as dangerous as technical lock-in. When these people leave, or the provider's specialist knowledge is withdrawn at termination, the organization loses operational capability regardless of data portability.
Objective	Verify that the cloud environment is documented in sufficient detail for a competent but previously uninvolved third party to assume full operational control within the exit RTO, without requiring knowledge transfer from provider personnel.
Sovereignty link	Technology Sovereignty (DICTU 3.4 levels 4–5: documentation 'sufficiently detailed for a third party to operate the service independently') | EuroStack Criterion 2.3 | DICTU 5.5.2 (capacity and skills)

Audit Questions

#	Audit Question	Yes/No	Notes
3,1	Is the system architecture documented in a current, version-controlled document (reviewed within 12 months) covering: network topology, data flows, integration points, component dependencies, and security architecture, stored in an organization-controlled system?		
3,2	Does Infrastructure-as-Code (IaC) cover at least 80% of all production infrastructure components, with artefacts stored in an organization-managed repository and a recent IaC-based deployment confirming the IaC is functional (not merely present)?		
3,3	Has the organization conducted a key person dependency analysis, identified all cloud functions relying on fewer than two individuals, and established a documented remediation plan for each identified dependency?		
3,4	Are operational runbooks available for all routine and non-routine procedures (deployment, scaling, backup, recovery, incident response) in sufficient detail for a competent engineer unfamiliar with the environment to execute them?		
3,5	Has the documentation been tested by a party external to the current team, confirming that a competent engineer unfamiliar with the environment can successfully execute a deployment or operational procedure using only the available documentation?		

Part A: Design Maturity (M1–M5)

Instructions: Tick (✓) the HIGHEST level at which ALL criteria are fully satisfied. Levels are CUMULATIVE.

Level		Definition	Criteria (cumulative)	Suggested Evidence to Provide	Reached (✓)	Notes
M1	No Documentation	Service not transferable; continuity not guaranteed	1. No current architecture documentation exists. 2. Operations depend entirely on knowledge held by one or two individuals or on provider specialist assistance. 3. IaC coverage below 20%; most infrastructure is manually configured.	→ *Absence of documentation confirms M1. If the organization cannot produce an architecture document reviewed within 36 months, M1 applies.*		
M2	Limited Documentation	Documentation limited; transfer to third party complex and not fully possible.	1. Cumulative: the organization has progressed beyond M1 (the M1 absence conditions no longer apply), PLUS: 2. Some documentation exists but is significantly outdated (not reviewed in 24 or more months) or covers only parts of the environment. 3. Multiple key person dependencies exist without a remediation plan. 4. IaC coverage between 20 and 50%.	→ *Architecture document (note: outdated or incomplete documents confirm M2 ceiling).* → *Key person analysis showing unaddressed single-person dependencies.*		
M3	Basic Documentation	Service largely transferable; high-level documentation available.	1. Cumulative: M2 satisfied, PLUS: 2. Architecture document is current (reviewed within 24 months) and covers main components. 3. Basic operational runbooks exist for primary procedures. 4. IaC coverage exceeds 50%. 5. Key persons have been identified and at least partial cross-training is underway.	→ *Architecture document with review date within 24 months.* → *Runbook list for primary procedures.* → *IaC repository with coverage estimate above 50%.* → *Key person analysis with cross-training progress record.*		
M4	Sufficient Documentation	Service fully transferable; documentation sufficiently detailed for a third party to operate independently.	1. Cumulative: M3 satisfied, PLUS: 2. Architecture document current (reviewed within 12 months) and comprehensive, covering all components including integrations, security controls, and recovery procedures. 3. IaC coverage exceeds 80%; recent deployment from IaC confirms it is functional. 4. All routine and non-routine procedures are in runbooks. 5. No unaddressed critical single-person dependencies remain.	→ *Architecture document review date within 12 months; comprehensive scope confirmed.* → *IaC coverage evidence above 80%; recent IaC-based deployment log as proof of function.* → *Runbook index covering all procedures (routine, incident, and disaster recovery).* → *Key person analysis with all critical dependencies addressed.*		
M5	Comprehensive Documentation	Service fully transferable; documentation sufficiently detailed and tested.	1. Cumulative: M4 satisfied, PLUS: 2. Documentation tested by a party outside the current team: a competent engineer unfamiliar with the environment successfully executed at least one deployment or operational procedure using only the available documentation. 3. Test result is documented with: date, tester (not an author of the documentation), procedure executed, and outcome.	→ *Documentation test report: tester name and role (must not be an author), date, procedure executed, outcome.* → *Confirmation that tester had no prior knowledge of the specific environment.* → *IaC deployment log from the test run.*		
DESIGN MATURITY CONCLUSION		**Maturity Level Assigned:**		Rationale / Finding:		

Part B: Operational Effectiveness (E1–E4)

Instructions: Assess whether the designed control operated as intended during the audit period.

Level		What the Organization Must Demonstrate	Audit Procedure	Audit Guidance	Reached (✓)	Notes
E1	Not Effective	1. Architecture documentation significantly outdated or absent. 2. Key person dependencies unaddressed. 3. IaC below 50%.	1. Inspect architecture document version date. 2. Request IaC repository and spot-check coverage. 3. Ask two non-primary cloud engineers to explain the architecture; if they cannot, documentation is insufficient.	*Practical test: ask a senior team member who is not the primary cloud engineer to locate and describe the data flow for a named application. Inability to do so without the primary engineer is a clear E1 indicator.*		
E2	Partially Effective	1. Documentation current. 2. IaC present with known coverage. 3. Key persons identified with remediation in progress.	1. Review IaC coverage via repository spot-check; do not accept stated coverage without verification. 2. Verify architecture document review date. 3. Review key person analysis and cross-training records for progress.	*IaC coverage claims must be verified by actually opening the repository and checking, not accepted on assertion. 'We have Terraform scripts' is not evidence of 80% coverage.*		
E3	Largely Effective	1. IaC above 80% confirmed. 2. Runbooks cover all critical procedures. 3. No critical single-person dependencies remaining.	1. Verify IaC coverage against infrastructure inventory (sample five components). 2. Select three non-routine procedures and verify runbooks exist, are complete, and current. 3. Confirm key person analysis shows no open critical dependencies.	*Test documentation completeness: select one infrastructure component and walk through the IaC deployment with a team member who is not the author. Can they execute it? If not, E2 at best.*		
E4	Fully Effective	1. Documentation tested by external or non-authoring party. 2. Deployment successfully executed using only documentation. 3. Test result documented with independence confirmed.	1. Review documentation test report, confirm tester is not an author. 2. Verify deployment log from the test run. 3. Confirm tester had no prior knowledge of the specific environment.	*The independence test is the defining criterion for E4: the person who executed the test must not be the person who wrote the documentation. Verify this explicitly.*		
EFFECTIVENESS CONCLUSION		**Effectiveness Level Assigned:**		Rationale / Finding:		

Responsible Roles

Design Maturity Level	Responsible Role (typically accountable for evidence at this level)
M1–M2	No remediation owner exists at this level. Escalate to Operations Lead to initiate documentation and key-person remediation.
M3	Operations Lead: owns architecture documentation currency and cross-training progress for key persons.
M4	Operations Lead + Engineering Manager: Operations Lead maintains comprehensive documentation and runbooks; Engineering Manager closes critical single-person dependencies.
M5	Operations Lead + independent/non-authoring tester: tester (not a documentation author) executes the independence test; Operations Lead records the outcome.

Operating Effectiveness Level	Responsible Role (typically accountable for evidence at this level)
E1	Operations Lead: accountable for architecture currency; a non-primary engineer is interviewed for the practical test.
E2	Operations Lead: accountable for verified (not asserted) IaC coverage and cross-training records.
E3	Operations Lead: accountable for runbook completeness and closure of critical dependencies (sample-verified).
E4	Independent/non-authoring tester + Operations Lead: tester executes the deployment; Operations Lead confirms tester independence.

Control 4: Continuity of Operations During Migration

Risk	Even when exit plan, portability, and documentation are all adequate, the transition period is a high-risk window: services are partially operational, data is split across old and new environments, and teams must manage both simultaneously. Without a migration continuity plan, exit can trigger a service availability crisis that forces the organization to abort the migration.
Objective	Verify that the exit plan contains a migration continuity plan addressing: migration phases, parallel operation periods, data synchronization protocol, service degradation tolerance, rollback procedure with trigger criteria, and stakeholder communication.
Sovereignty link	Operational Sovereignty (DICTU 4 Operational) | NIS2 Art. 21 (business continuity) | DORA Art. 11 (ICT continuity policy) | ISO 22301 §8.5

Audit Questions

#	Audit Question	Yes/No	Notes
4,1	Does the exit plan contain a dedicated migration continuity section with: defined migration phases, services operational at each phase, parallel operation period, data synchronization protocol, and explicit phase-gate criteria for cut-over?		
4,2	Is there a documented data synchronization and integrity protocol specifying: synchronization frequency during migration, integrity check method, change-freeze or dual-write period, and a post-cut-over validation procedure?		
4,3	Does the continuity plan contain a documented rollback procedure with explicit trigger criteria (conditions that activate rollback), step-by-step execution instructions, and an approved rollback RTO?		
4,4	Has the service degradation tolerance for each cloud-hosted service been defined (e.g. read-only mode acceptable for x hours) and validated against contractual SLA obligations and any applicable NIS2 or DORA requirements?		
4,5	Does the exit plan include a stakeholder communication plan for the migration period covering: internal stakeholders (IT, Legal, Risk, Board), external stakeholders (customers, supervisors), communication frequency, and escalation path for migration failures?		

Part A: Design Maturity (M1–M5)

Instructions: Tick (✓) the HIGHEST level at which ALL criteria are fully satisfied. Levels are CUMULATIVE.

Level		Definition	Criteria (cumulative)	Suggested Evidence to Provide	Reached (✓)	Notes
M1	No Continuity Plan	No reversibility; continuity not guaranteed.	1. Exit plan contains no continuity section for the migration period. 2. No rollback procedure defined. 3. Service degradation tolerance undefined.	→ *Absence of continuity section in exit plan confirms M1 for this control.*		
M2	Continuity Acknowledged	Limited reversibility; transfer is complex.	1. Cumulative: the organization has progressed beyond M1 (the M1 absence conditions no longer apply), PLUS: 2. Migration continuity is acknowledged as a risk in the exit plan but no detailed plan exists. 3. Rollback is mentioned but not procedurally defined (no trigger criteria, no execution steps). 4. Service degradation not formally assessed.	→ *Exit plan section acknowledging continuity risk.* → *High-level rollback mention without procedure details.*		
M3	Basic Plan Defined	Service largely transferable; high-level documentation available.	1. Cumulative: M2 satisfied, PLUS: 2. Migration phases are defined with a phase diagram. 3. Parallel operation period is specified. 4. Data synchronization protocol described at high level. 5. Rollback procedure documented with trigger criteria, even if not yet tested.	→ *Phase diagram in exit plan.* → *Parallel operation period specification.* → *Data synchronization protocol (can be at high level for M3).* → *Rollback procedure with trigger criteria.*		
M4	Detailed and Validated Plan	Service fully transferable; documentation sufficiently detailed.	1. Cumulative: M3 satisfied, PLUS: 2. Data synchronization protocol is detailed (specifying frequency, integrity check method, dual-write or freeze period). 3. Service degradation tolerance formally defined per service and validated against SLAs. 4. Rollback RTO approved by CIO or equivalent. 5. Communication plan completed identifying all stakeholders, channels, and escalation path.	→ *Detailed synchronization protocol with integrity check procedure.* → *Service degradation tolerance document cross-referenced to SLA schedules.* → *Rollback RTO with CIO approval record.* → *Communication plan with stakeholder list reviewed by all named parties.*		
M5	Tested Continuity Plan	Service fully transferable; documentation sufficiently detailed and tested.	1. Cumulative: M4 satisfied, PLUS: 2. Migration continuity plan validated in a structured exercise; minimum: tabletop with an unplanned trigger scenario. 3. Rollback procedure explicitly tested in the exercise (rollback was activated as part of the scenario). 4. Exercise findings documented and remediated.	→ *Exercise report with unplanned exit trigger scenario (not only voluntary migration).* → *Confirmation that rollback was activated and executed in the exercise.* → *Findings register with all exercise findings closed or on schedule.*		
DESIGN MATURITY CONCLUSION		**Maturity Level Assigned:**		Rationale / Finding:		

Part B: Operational Effectiveness (E1–E4)

Instructions: Assess whether the designed control operated as intended during the audit period.

Level		What the Organization Must Demonstrate	Audit Procedure	Audit Guidance	Reached (✓)	Notes
E1	Not Effective	1. No continuity section in exit plan. 2. No rollback procedure. 3. Service degradation tolerance not defined.	1. Review exit plan for continuity section. 2. Verify rollback procedure exists. 3. Check SLA schedules for degradation tolerance.	*If exit plan has no continuity section, rate E1 for Control 4. This automatically limits the overall conclusion to at most partially effective.*		
E2	Partially Effective	1. Continuity plan reviewed within 24 months. 2. Rollback procedure documented. 3. SLA cross-check performed.	1. Review continuity section currency. 2. Interview continuity plan owner on rollback trigger logic. 3. Verify SLA cross-check evidence.	*Plan exists and maintained but never tested. The most common state. Priority recommendation: schedule a tabletop exercise with an unplanned scenario.*		
E3	Largely Effective	1. Tabletop exercise conducted with unplanned exit trigger. 2. Rollback explicitly tested in exercise. 3. Findings remediated.	1. Review exercise report; confirm scenario included unplanned trigger. 2. Verify rollback was activated (not merely discussed) in exercise. 3. Confirm findings closure.	*E3 requires that rollback was actually activated and tested in the exercise, not only that the scenario mentions rollback as an option. An exercise where rollback was discussed but not executed is E2.*		
E4	Fully Effective	1. Partial migration test conducted with service continuity maintained throughout. 2. Rollback executed and confirmed functional. 3. Communication plan tested with actual stakeholders.	1. Review partial migration test documentation. 2. Verify rollback was executed and services resumed without data loss. 3. Confirm stakeholder communication test evidence.	*E4 requires actual partial migration execution. A tabletop simulation, however thorough, does not constitute E4 for this control.*		
EFFECTIVENESS CONCLUSION		**Effectiveness Level Assigned:**		Rationale / Finding:		

Responsible Roles

Design Maturity Level	Responsible Role (typically accountable for evidence at this level)
M1–M2	No continuity owner exists at this level. Escalate to Exit Plan Owner / CIO to initiate a migration continuity section.
M3	Exit Plan Owner + Infrastructure Lead: define migration phases, parallel-operation period, and rollback trigger criteria.
M4	Exit Plan Owner + CIO: CIO approves rollback RTO and degradation tolerance vs. SLA; Owner completes the stakeholder communication plan.
M5	Exit Plan Owner + exercise participants (IT, Legal, Risk, Board as applicable): exercise rollback under an unplanned trigger scenario.

Operating Effectiveness Level	Responsible Role (typically accountable for evidence at this level)

E1	Exit Plan Owner: accountable for the existence of a continuity section; absence caps the overall conclusion (Rule 1 / Summary tab).
E2	Exit Plan Owner: accountable for continuity plan currency and SLA cross-check.
E3	Exit Plan Owner + exercise participants: accountable for rollback activation (not only discussion) during the exercise.
E4	Exit Plan Owner + Operations Lead: jointly accountable for executing and confirming a partial migration test.

Control 5: Testing Programme and Periodic Evaluation

Risk	Exit plans and documentation that are created but never tested provide false assurance. Technology, provider terms, and organizational context change continuously. DICTU level 5 explicitly requires documentation that is 'tested'. Without periodic testing, the organization cannot demonstrate that its exit capability is real rather than theoretical.
Objective	Verify that exit capability is subject to a structured, periodic test programme, and that findings from testing result in documented, tracked remediation activities reported to senior management.
Sovereignty link	Technology Sovereignty (DICTU 3.4 level 5: 'sufficiently detailed AND TESTED') | EuroStack Criterion 2.3 (tested deployment capability) | ISO 22301 §8.5 (BCM test programme)

Audit Questions

#	Audit Question	Yes/No	Notes
5,1	Does a documented exit capability test programme exist, specifying: test types (tabletop, data export test, partial migration), test frequency (minimum annual for critical services), responsible parties, scope, and success criteria?		
5,2	Has an annual tabletop exercise specifically testing exit scenarios been conducted, attended by IT, Legal, Compliance, and senior management, with a scenario that includes at least one unplanned exit trigger (provider insolvency, geopolitical sanctions, or non-EU acquisition)?		
5,3	Has a complete or representative data export test, including import validation in an alternative environment, been conducted within the past 24 months, with results documented including data integrity confirmation?		
5,4	Is there a procedure for trigger-based review of the exit plan upon material changes (new service onboarded, contract renewal, major data volume increase, regulatory change, geopolitical development affecting the provider), and has at least one trigger-based review taken place in the past 24 months if triggers have occurred?		
5,5	Is a findings register maintained for all exit-related testing, with named owners and due dates, and are findings reported to senior management or the Audit Committee at least annually?		

Part A: Design Maturity (M1–M5)

Instructions: Tick (✓) the HIGHEST level at which ALL criteria are fully satisfied. Levels are CUMULATIVE.

Level		Definition	Criteria (cumulative)	Suggested Evidence to Provide	Reached (✓)	Notes
M1	**No Testing**	Service not transferable; continuity not guaranteed.	1. Exit capability has never been tested in any form. 2. No test programme exists. 3. No findings register maintained.	→ *Absence of any test records confirms M1. If no evidence is produced, this level applies by default.*		
M2	**Ad Hoc Testing**	Documentation is limited; transfer to a third party is complex and not fully possible.	1. Cumulative: the organization has progressed beyond M1 (the M1 absence conditions no longer apply), PLUS: 2. Some informal testing or discussion of exit scenarios has occurred. 3. No structured programme, no documented findings, no remediation tracking.	→ *Email records or meeting notes mentioning exit scenario discussion (note: informal discussion without documented outcomes confirms M2 at most).*		
M3	**Annual Tabletop**	Service largely transferable; high-level documentation available.	1. Cumulative: M2 satisfied, PLUS: 2. A formal test programme document exists specifying test types and frequency. 3. Annual tabletop exercise conducted within the past 12 months. 4. Tabletop scenario included at least one unplanned exit trigger. 5. Exercise findings are documented (data export test not yet conducted).	→ *Test programme document with frequency and test types.* → *Tabletop exercise report: date, scenario, attendance list, findings.* → *Confirmation that scenario included an unplanned trigger (not only voluntary migration).*		
M4	**Full Test Programme**	Service fully transferable; documentation sufficiently detailed.	1. Cumulative: M3 satisfied, PLUS: 2. Data export test conducted within 24 months with import validation. 3. Test findings tracked in findings register with owners and due dates. 4. Findings reported to senior management at least annually. 5. Trigger-based review procedure in place.	→ *Data export test report with import validation evidence.* → *Findings register showing open and closed items from all test types.* → *Management report or Audit Committee minutes referencing exit capability test results.* → *Trigger-based review procedure document.*		
M5	**Comprehensively Tested**	Service fully transferable; documentation sufficiently detailed and tested.	1. Cumulative: M4 satisfied, PLUS: 2. Partial migration test (or equivalent highest-level exercise) conducted in addition to tabletop and export test. 3. Test programme subject to independent oversight (internal audit or external review). 4. All prior-period findings are closed or have an accepted risk.	→ *Partial migration test report or equivalent live exercise evidence.* → *Independent oversight record (internal audit scope or external review engagement).* → *Findings register showing all prior-period items closed or risk-accepted with approval.*		
DESIGN MATURITY CONCLUSION		**Maturity Level Assigned:**		Rationale / Finding:		

Part B: Operational Effectiveness (E1–E4)

Instructions: Assess whether the designed control operated as intended during the audit period.

Level		What the Organization Must Demonstrate	Audit Procedure	Audit Guidance	Reached (✓)	Notes
E1	**Not Effective**	1. No evidence of any exit-related testing in the past 24 months. 2. No test programme document. 3. No findings register.	1. Request test records for exit scenarios over 24 months. 2. If none produced, rate E1. 3. Verify whether test programme document exists.	*If no testing has been conducted in the past 24 months, E1 applies regardless of how the design maturity is scored.*		
E2	**Partially Effective**	1. Annual tabletop conducted. 2. Findings documented. 3. Remediation incomplete or data export test not yet conducted.	1. Review most recent tabletop report: date, scenario, participants, findings. 2. Check findings register for open items. 3. Verify data export test status.	*Confirm the tabletop scenario included an unplanned trigger. A planned migration session labelled as a tabletop exercise does not meet this requirement.*		
E3	**Largely Effective**	1. Tabletop and data export test completed within 24 months. 2. All findings from both tests closed. 3. Annual management reporting evidenced.	1. Review data export test results including import validation. 2. Confirm all tabletop and export test findings are addressed. 3. Review management report or Audit Committee minutes.	*E3 requires both test types: tabletop for governance and decision-making, and data export test for technical validation. One without the other is E2 at most.*		
E4	**Fully Effective**	1. Full programme: tabletop, export test, and partial migration test. 2. Independent oversight of test programme. 3. All prior findings closed.	1. Review partial migration test documentation. 2. Confirm independent oversight (internal audit scope or external review). 3. Verify all prior-period findings are closed.	*E4 requires a partial migration test and independent oversight. An organization that has never tested an actual migration cannot claim E4 regardless of how thorough its tabletop exercises are.*		
EFFECTIVENESS CONCLUSION		**Effectiveness Level Assigned:**		Rationale / Finding:		

Responsible Roles (Testing Programme and Periodic Evaluation)

Design Maturity Level	Responsible Role (typically accountable for evidence at this level)
M1–M2	No test programme owner exists at this level. Escalate to CIO / Audit Committee to mandate a test programme.
M3	Test Programme Owner: schedules and reports the annual tabletop exercise (with unplanned trigger) within the past 12 months.
M4	Test Programme Owner + Infrastructure Lead: Owner maintains the findings register and trigger-based review procedure; Infrastructure Lead delivers the data export test.
M5	Test Programme Owner + Internal Audit: Internal Audit provides independent oversight of the full test programme (tabletop, export test, partial migration).

Operating Effectiveness Level	Responsible Role (typically accountable for evidence at this level)
E1	Test Programme Owner: accountable for the existence of any test record within 24 months.
E2	Test Programme Owner: accountable for tabletop scenario design (must include an unplanned trigger) and findings tracking.

E3	Test Programme Owner + Infrastructure Lead: jointly accountable for completing both tabletop and export test with closed findings.
E4	Test Programme Owner + Internal Audit: Internal Audit confirms independent oversight; Owner confirms all prior findings closed.

Overall Audit Conclusion

The overall conclusion is based on the individual conclusions across the five control areas. The assessor applies professional judgement. The decision rules below provide guidance, but do not replace judgement in cases where the situation is atypical.

#	Control Area	Design Maturity	Operating Effectiveness	Key Finding	Recommendations
1	Exit Plan and Contractual Framework				
2	Data Portability and Technical Exportability				
3	Knowledge Transfer and Operational Documentation				
4	Continuity of Operations During Migration				
5	Testing Programme and Periodic Evaluation				

Scoring

Average Design Maturity:	#DIV/0!	Average Operating Effectiveness:	#DIV/0!

Decision Guidelines

These are guidelines to support the auditor's professional judgement, not mechanical pass/fail rules. Where a situation does not fit neatly into one category, the auditor applies judgement and documents the rationale for the conclusion reached.

Guideline 1	Control 1 at M1 or E1	Treat Control 1 (the exit plan) as a prerequisite for the others. If it sits at M1 or E1, the overall conclusion should generally be Not Effective — strong scores elsewhere cannot compensate. Document any deviation from this.
Guideline 2	Any control at M1 + E1	If any single control is at M1 and E1, both design and operating evidence are effectively absent for that area. This generally points to Not Effective overall, unless the auditor can justify why the gap is not material to reversibility.
Guideline 3	All M2+/E2+, but Control 5 at M1 or E1	Where controls are broadly adequate (M2+/E2+) but the testing programme (Control 5) lags at M1 or E1, reversibility is designed but not yet confirmed in practice. Partially Effective is usually appropriate; consider recommending testing as the priority action.
Guideline 4	All controls M3+/E3+	If all five controls reach at least M3 in design and E3 in operation, exit capability is both designed to a contractual standard and operationally demonstrated. The auditor can generally conclude Effective, absent specific countervailing evidence.
Guideline 5	All controls M5/E4	Reserve Fully Effective for cases where every control reaches M5 and E4, maximum operational reversibility (DICTU Level 5). Expect this to be supported by tested evidence across all controls.

Overall Conclusion

Basis for overall conclusion:	
Priority recommendations for management:	
Required follow-up actions (owner and deadline):	

Appendix C

The full Atlas.ti code list

This appendix contains the complete code list produced in Atlas.ti during the qualitative analysis of the Round 1 and Round 2 interview transcripts. Codes are organized by the TAM construct to which they were primarily assigned (Perceived Usefulness, Perceived Ease of Use, Social Norm, and Behavioral Intention), followed by inductively generated codes that emerged from the data and were not directly derivable from the TAM framework. Each code is listed with its definition and the theme to which it was applied. The full untranslated Dutch transcripts are available from the author upon reasonable request, subject to the anonymization commitments made to respondents as described in section 3.7.

Code	Theme
PU: existing standards insufficient	Theme 1: The framework provides structure and improves audit quality
PU: director liability relevance	Theme 1: The framework provides structure and improves audit quality
PEOU: intuitive on first reading	Theme 1: The framework provides structure and improves audit quality
PEOU: language accessible	Theme 1: The framework provides structure and improves audit quality
BI: client insight as recommendation reason	Theme 1: The framework provides structure and improves audit quality
IMP: evidence requirements mostly adequate	Theme 1: The framework provides structure and improves audit quality
PU: finding relevance depends on client risk perception	Theme 1: The framework provides structure and improves audit quality
PU: identifies missed risks	Theme 1: The framework provides structure and improves audit quality
PEOU: framework fits daily audit practice	Theme 1: The framework provides structure and improves audit quality
PU: framework quality of audit conclusions	Theme 1: The framework provides structure and improves audit quality
PEOU: framework understandable without training	Theme 1: The framework provides structure and improves audit quality
PU: framework provides more complete picture	Theme 1: The framework provides structure and improves audit quality
PU: insufficient enforcement of evidence collection	Theme 1: The framework provides structure and improves audit quality
PU: exit plan distinct from emergency plan; different scope	Theme 1: The framework provides structure and improves audit quality
PEOU: strong table structure level definition criteria evidence	Theme 1: The framework provides structure and improves audit quality
PEOU: audit questions clear	Theme 1: The framework provides structure and improves audit quality
PEOU: emphasis on tangible evidence not assertion	Theme 1: The framework provides structure and improves audit quality
PEOU: evidence requirements specific and clear	Theme 1: The framework provides structure and improves audit quality
PU: finding relevant due to stakeholder reporting obligation	Theme 1: The framework provides structure and improves audit quality
BI: framework prevents ad hoc assessment and missed elements	Theme 1: The framework provides structure and improves audit quality
BI: framework prevents manipulation of audit scope	Theme 1: The framework provides structure and improves audit quality
PU: existing approach relies on ad hoc document collection without structure	Theme 1: The framework provides structure and improves audit quality
PU: added value beyond ISO 27001 and NIS2	Theme 1: The framework provides structure and improves audit quality
PU: growing client awareness of sovereignty risks drives framework relevance	Theme 1: The framework provides structure and improves audit quality
BI: baseline measurement and risk assessment as use cases	Theme 1: The framework provides structure and improves audit quality
BI: facilitates conversation about desired maturity target	Theme 1: The framework provides structure and improves audit quality
BI: advisory mode preferred	Theme 2: Practical applicability determines adoption
BI: condition is an improved structure	Theme 2: Practical applicability determines adoption
IMP: Excel format preferred, practical applicability determines adoption	Theme 2: Practical applicability determines adoption
BI: condition — maturity model revision needed	Theme 2: Practical applicability determines adoption
BI: structure value as adoption motivation	Theme 2: Practical applicability determines adoption
PU: primary use case risk-based not compliance	Theme 2: Practical applicability determines adoption
BI: use dependent on specific engagement type	Theme 2: Practical applicability determines adoption
BI: would recommend as reference and comparative material	Theme 2: Practical applicability determines adoption
IMP: visual dashboard with colour-coded scoring per maturity level	Theme 2: Practical applicability determines adoption
IMP: add automatic score calculation to working format	Theme 2: Practical applicability determines adoption
IMP: add detail tab alongside summary overview	Theme 2: Practical applicability determines adoption
BI: framework as starting point for own audit templates	Theme 2: Practical applicability determines adoption
PU: primary use case advisory maturity assessment	Theme 2: Practical applicability determines adoption
BI: practical applicability is decisive for actual adoption	Theme 2: Practical applicability determines adoption
BI: future adoption potential via audit management tools and systems	Theme 2: Practical applicability determines adoption
IMP: merge design and operating into single Excel table with columns	Theme 2: Practical applicability determines adoption
BI: would recommend to colleague conditional on tweaks	Theme 2: Practical applicability determines adoption
PU: insufficient structure for defensible audit opinion	Theme 3: The framework is valuable but requires methodological maturation
PEOU: illogical question order	Theme 3: The framework is valuable but requires methodological maturation
PEOU: need for explicit formulation	Theme 3: The framework is valuable but requires methodological maturation
IMP: reorder questions outside-in	Theme 3: The framework is valuable but requires methodological maturation
IMP: exit scenario question unclear	Theme 3: The framework is valuable but requires methodological maturation
IMP: add guidance document	Theme 3: The framework is valuable but requires methodological maturation
PU: missing organisational recovery standpoint	Theme 3: The framework is valuable but requires methodological maturation
PEOU: criteria clarity insufficient for consistent assessment	Theme 3: The framework is valuable but requires methodological maturation
PEOU: two dimensions on one control is illogical	Theme 3: The framework is valuable but requires methodological maturation
PEOU: creating false security in board reporting	Theme 3: The framework is valuable but requires methodological maturation
PEOU: SURF approach as alternative model	Theme 3: The framework is valuable but requires methodological maturation
IMP: maturity level and recovery time are independent variables	Theme 3: The framework is valuable but requires methodological maturation
IMP: cloud business model contradicts internal knowledge retention requirement	Theme 3: The framework is valuable but requires methodological maturation
IMP: control 1 needs organisational standpoint definition	Theme 3: The framework is valuable but requires methodological maturation
PU: risk stratification by provider criticality missing	Theme 3: The framework is valuable but requires methodological maturation
PEOU: concepts overlap — exit plan and tabletop test mixed	Theme 3: The framework is valuable but requires methodological maturation
PEOU: score 4 — audit questions clear but coherence needs improvement	Theme 3: The framework is valuable but requires methodological maturation
PEOU: distinction visible but not consistently applicable	Theme 3: The framework is valuable but requires methodological maturation
IMP: add provider criticality assessment before framework application	Theme 3: The framework is valuable but requires methodological maturation
IMP: operating effectiveness not always testable for document-based controls	Theme 3: The framework is valuable but requires methodological maturation
IMP: clarify distinction between exit plan document and tabletop test activity	Theme 3: The framework is valuable but requires methodological maturation
PEOU: need for concrete test scenarios for operating effectiveness	Theme 3: The framework is valuable but requires methodological maturation
PEOU: specific time thresholds need justification	Theme 3: The framework is valuable but requires methodological maturation
PEOU: design failure makes operating testing redundant	Theme 3: The framework is valuable but requires methodological maturation
IMP: add question about roles and responsibilities	Theme 3: The framework is valuable but requires methodological maturation
IMP: follow-up actions are management responsibility not auditor task	Theme 3: The framework is valuable but requires methodological maturation
PEOU: missing glossary for technical terms	Theme 3: The framework is valuable but requires methodological maturation
PEOU: operating effectiveness not immediately testable for newly implemented contr	Theme 3: The framework is valuable but requires methodological maturation
IMP: add post-migration control — can organisation operate normally after switching	Theme 3: The framework is valuable but requires methodological maturation
IMP: fix inconsistent use of confirmed absent vs confirmed satisfied	Theme 3: The framework is valuable but requires methodological maturation
IMP: add not yet testable operating effectiveness judgement option	Theme 3: The framework is valuable but requires methodological maturation
IMP: missing glossary for technical terms	Theme 3: The framework is valuable but requires methodological maturation

PEOU: confusion about single vs multiple exit plan objects	Theme 3: The framework is valuable but requires methodological maturation
IMP: clarify that one object is assessed against maturity level criteria	Theme 3: The framework is valuable but requires methodological maturation
PU: cumulative maturity model	Theme 3: The framework is valuable but requires methodological maturation
PU: low maturity in practice	Theme 3: The framework is valuable but requires methodological maturation
PEOU: testing exit capability	Theme 3: The framework is valuable but requires methodological maturation
PEOU: critique of practical feasibility	Theme 3: The framework is valuable but requires methodological maturation
IMP: realistic evidence requirements	Theme 3: The framework is valuable but requires methodological maturation
PU: emergency plan versus exit plan	Theme 3: The framework is valuable but requires methodological maturation
PU: gap in existing methodology	Theme 3: The framework is valuable but requires methodological maturation
PU: audit approach without framework	Theme 3: The framework is valuable but requires methodological maturation
PEOU: framework versus guidance confusion	Theme 3: The framework is valuable but requires methodological maturation
PU: sufficient structure when applied to one specific cloud provider	Theme 3: The framework is valuable but requires methodological maturation
PEOU: design versus operating effectiveness	Theme 3: The framework is valuable but requires methodological maturation
SN: sovereignty becoming mainstream	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: analogy with prior emerging topics	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: client misconception about sovereignty	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: client maturity too low for formal audit	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: align with regulatory direction	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: risk quantification as precondition	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: sovereignty still niche in audit practice	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: client adoption depends on risk awareness	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: sovereignty not yet mainstream in audit engagements	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: organisation sees opportunity in sovereignty services	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: broad acceptance lowers adoption threshold	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: professional environment supportive of frameworks	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: sovereignty outside scope of financial audit	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: broad openness to frameworks within team	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: organisational mandate as primary adoption driver	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: US cloud providers and Cloud Act as sovereignty awareness driver	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: NIS2 and geopolitical situation as drivers of sovereignty awareness	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: current regulation does not yet require sovereignty assessment	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: Cloud Act as concrete sovereignty risk in client conversations	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: sovereign cloud development as internal strategic priority	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: growing attention but not yet embedded in standard audit practice	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: professional body as norm-setter for adoption	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: leading organizations adopting a framework would increase adoption likelihood	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
SN: auditor must remain flexible; not all regulatory requirements demand this framew	Theme 4: Digital sovereignty is gaining relevance but adoption is shaped by external drivers
PEOU: evidence requirements too prescriptive	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PEOU: context-dependency of audit questions by organisational scale	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
IMP: flexible evidence requirements necessary	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PEOU: IaC threshold too prescriptive	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PEOU: principle-based norm preferred over prescriptive evidence	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PU: risk of checkbox compliance	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PEOU: specific evidence requirements may cause blind spots	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PEOU: subjectivity in qualitative criteria requires professional judgement	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
IMP: framework should support evidence description not just binary answers	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
PU: framework is compliance-driven rather than risk-driven	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter
IMP: reformulate closed questions as open norms for auditor to verify	Theme 5: Auditing digital sovereignty requires professional judgement not rigid compliance criter