

Leiden University

ICT in Business and the Public Sector

Designing follow-up governance for ICT-related recommendations: a case study at SSC-ICT

Name:	Justin Simon de Weert
Student ID:	s3075400
Date:	01/07/2026
1st supervisor:	Prof. Dr. Ir. J.M.W. Visser
2nd supervisor:	T.D. Offerman MSc.
Company supervisor:	H.J. Bouwknegt, A. Roos, M. Westzaan

MASTER'S THESIS

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

Abstract

Background: Government organisations depend on ICT to execute policy and deliver services. Audits, evaluations, and advisory bodies therefore produce a steady stream of ICT-related recommendations. At SSC-ICT, the shared service provider of the Dutch central government, audit findings are tracked in a structured way, but strategic and advisory recommendations are not. As a result, they can stay open for years, lose there relevance or disappear without an explicit decision.

Aim: This thesis designs and evaluates a follow-up and governance method that helps SSC-ICT, and possibly similar organisations, to keep ICT-related recommendations visible and to steer them over time.

Method: The study follows a design science research approach, combined with a qualitative case study at SSC-ICT. Eight interviews were held inside the organisation and four with producers of recommendations (Algemene Rekenkamer, ACICT, CIO Rijk, and the Vereniging van Rekenkamers). From the interviews and the literature, 25 design requirements were derived, using the Volere template. The artefact built on these requirements is a recommendation register with a three-layer status model, based on the "effect chain" of follow-up, implementation and impact. The artefact is demonstrated on five historical cases from an inventory of 126 recommendations, and validated in a workshop with seven stakeholders, using the Method Evaluation Model (MEM).

Results: The artefact could hold recommendations from all four follow-up patterns found in the organisation, including abstract recommendations that need decomposition and recommendations that depend on other organisations. In the workshop, perceived usefulness, intention to use and perceived ease of use were measured and scored above the midpoint. The participants saw the value mainly in the shared overview and the explicit ownership. Their main doubt was the administrative burden, and whether owners would keep the register up to date.

Conclusion: The follow-up of ICT-related recommendations can be made more structured by treating each recommendation as an object that is followed over time, with a clear owner, a status in three layers, and a review and closure in a decision body that already exists and meets regularly. The demonstration and the workshop show that such a register fits the SSC-ICT context and that the stakeholders would use it. Because the method is tested before real implementation, its effect on the follow-up practice is not shown yet, and the method only works if the register is kept up to date. Testing it on real recommendations inside the organisation is the main next step.

Acknowledgements

This thesis could not have been written without the guidance of Prof. Dr. Ir. Joost Visser and Tyron Offerman MSc. Not only were they of great value for their academic insight and their experience from practice, but most of all they were very pleasant to work with.

I have known both of them since my bachelor. I followed several of Tyron's classes back then, and they were among the more engaging lectures I had during my studies. Afterwards I was a student assistant for his classes for two and a half years, and we always worked together very pleasantly. Joost supervised my bachelor thesis as well, and his interest and clear passion for the public sector kept my motivation high throughout this thesis.

I also want to thank Harm-Jan Bouwknegt, Arie Roos, and Monique Westzaan for their guidance at SSC-ICT, and above all for making me feel welcome in such a pleasant and open organisation.

Contents

Abstract . . . 1
Acknowledgements . . . 1

1 Introduction **3**
1.1 Problem statement . . . 3
1.2 Research questions . . . 4
1.3 Overview of the thesis . . . 5

2 Background and related work **6**
2.1 Public-sector ICT and policy implementation . . . 6
2.2 Policy recommendations, audits, and advisory reports in government . . . 7
2.3 Governing ICT-related recommendations under uncertainty . . . 9
2.4 Ex ante anticipation vs. ex post accountability . . . 10
2.5 Fragmentation of governance across reports, time, and decision arenas . . . 11
2.6 Temporal decay, obsolescence, and learning failure . . . 13
2.7 The effect chain of advisory research: a lens from Dutch Rekenkamer practice . . . 15
2.8 The Dutch context . . . 17

3 Methodology **19**
3.1 Research Design . . . 19
3.2 Case context and scope . . . 20
3.3 Execution of DSRM . . . 20
3.4 Data sources and collection . . . 21
3.4.1 Document analysis . . . 21
3.4.2 Semi-structured interviews . . . 21
3.4.3 Data analysis and synthesis approach . . . 21
3.4.4 Evaluation criteria . . . 22
3.4.5 Ethics and data handling . . . 23
3.5 Scope and intended generalisability . . . 23

4 Requirements synthesis **24**
4.1 Interview analysis . . . 24
4.1.1 Theme coding overview . . . 25
4.1.2 Cross-cutting observations from the interviews . . . 25
4.2 Processmodeling . . . 31
4.2.1 Formalised audit follow-up . . . 31
4.2.2 Programmatic follow-up via taskforce . . . 32
4.2.3 Strategic follow-up without structural embedding . . . 33
4.2.4 Chain-dependence . . . 34
4.3 External perspectives on recommendation follow-up . . . 35
4.3.1 Synthesis of producer-side observations . . . 41
4.4 Synthesis: design requirements for a follow-up and governance method . . . 42
4.4.1 Design problem and contribution argument . . . 42
4.4.2 Approach to requirements synthesis . . . 42
4.4.3 Project drivers and constraints . . . 44
4.4.4 Functional requirements . . . 46
4.4.5 Non-functional requirements . . . 52
4.4.6 Traceability matrix . . . 56

5 Artefact design **57**
5.1 Recommendation register . . . 57

5.1.1 Schema . . . 57
5.2 Status model . . . 59
5.2.1 The receiver-side reading of the effect chain . . . 60
5.2.2 Status fields and value sets . . . 61
5.3 Intake . . . 62
5.4 Review and closure . . . 63
5.5 Learning loop . . . 65
5.6 Tool mapping . . . 65
5.7 Governance cycle . . . 66
5.8 Generalisability . . . 67

6 Demonstration **68**
6.1 Demonstration definition . . . 68
6.1.1 Inventory and case selection . . . 68
6.1.2 Case A - baseline audit finding (Interpretation 1) . . . 71
6.1.3 Case B - abstract recommendation requiring decomposition (Interpretation 1) . . . 72
6.1.4 Case C - strategic advisory recommendation (Interpretation 3) . . . 73
6.1.5 Case D - programmatic follow-up over a long horizon (Interpretation 2) . . . 74
6.1.6 Case E - research-report recommendation that decays over time (Interpretation 3) . . . 75
6.1.7 What the demonstration shows . . . 77

7 Validation **79**
7.1 Why the MEM and not a software-usability test . . . 79
7.2 Workshop design . . . 79
7.3 Analysis approach . . . 80
7.4 Results . . . 81
7.4.1 Questionnaire . . . 81
7.4.2 Observation . . . 82
7.4.3 Open questions . . . 82
7.5 Interpretation against the requirements . . . 83
7.6 Limits of the validation . . . 86

8 Discussion **87**
8.1 Key findings . . . 87
8.2 Adoption in context: a deepening interview . . . 88
8.3 Practical implications . . . 89
8.4 Limitations . . . 90

9 Conclusions **91**
9.1 Answers to the research questions . . . 91
9.2 Contribution to theory . . . 92
9.3 Future work . . . 92

A Exploratory interview protocol (receiver side) **97**

B Producer-side interview protocol (advisory and oversight bodies) **100**

C Validation questionnaire (Method Evaluation Model) **103**

D Code-to-theme inventory **105**

E Recommendation inventory **108**
E.1 Source documents . . . 109
E.2 Recommendations . . . 109

F Workshop materials **121**

G Glossary **126**
G.1 Management structure at SSC-ICT . . . 126
G.2 Institutional terms . . . 126
G.3 SSC-ICT and governance terms . . . 127

Chapter 1

Introduction

Government organisations increasingly depend on information and communication technologies (ICT) to execute policy, coordinate work, and deliver services. As digital infrastructure becomes much more embedded in public administration, audits, evaluations, advisory reports, and review bodies produce a growing number of (ICT-related) recommendations intended to improve governance, reduce risk, and ensure execution. These recommendations are meant to improve the governmental organisation, provide corrective action, and support organisational learning across cycles of advice and implementation.

This thesis examines that problem in the context of SSC-ICT, the shared service provider for multiple ministries within the Dutch central government. It is classified as an executive agency ("agentschap")[1] within the Dutch governmental organisation. Within such an environment that manages many of the IT infrastructure of the Dutch government, it is of the utmost importance that it takes up recommendations in a proactive manner, as its influence concerns not only technical issues, but also governance arrangements, implementation capacity, interdependencies and longer term organisational behaviour. This makes the follow-up of recommendations more complex than a simple question of whether an action has been formally completed. Recommendations may compete for attention, depend on multiple actors, change in urgency over time or lose relevance as technologies, risks, and policy contexts evolve.

With this background in mind, this thesis focuses on the design and evaluation of a structured follow-up and governance method for recommendations. The study aims to support monitoring, prioritisation, and steering over time, so that recommendations from different audits, evaluations and advisory reports can be made more visible and manageable within SSC-ICT and similar government organisations.

1.1 Problem statement

In the Dutch public sector, advisory and oversight reports are frequently acknowledged, however, the translation into action and follow-up is often lacking. This problem is especially relevant in ICT-intensive settings, where implementation often depends on shared infrastructures, multiple governance layers, scarce capacity and changing technological conditions. In such contexts, recommendations do not automatically become concrete implementation actions once they are accepted in principle.

Within SSC-ICT, this challenge appears in a specific way. Audit findings are tracked fairly consistently, but the more strategic, cross-cutting, or policy-driven recommendations are not: they often have no fixed follow-up process, no clear owner, and no fixed moment to review them. As a result, recommendations can stay open for a long time, get reinterpreted along the way, lose relevance, or disappear into scattered governance routines. It then becomes hard to keep an overview of what is still open, why it is still open, and what should be done next.

The problem is not only administrative. When recommendations are followed up inconsistently, unresolved actions pile up, weak spots in governance stay open, and changes are delayed. In the public sector this does not only slow projects down. Over time it also touches accountability, security, continuity, and the ability to keep up with new policy and technology. This weighs heavier for SSC-ICT than for a single

[1] Dutch terms and institutional abbreviations used in this thesis are listed in Appendix G.

organisation, because its recommendations often concern IT services and infrastructure used by several ministries, so the effects reach beyond one unit.

This thesis also addresses an academic gap. ICT governance, digitalisation, audits, evaluations, and project success factors have each been studied on their own. What has been studied far less is how a public organisation governs the follow-up of many different ICT-related recommendations over time, when those recommendations come from different sources, differ in scope and urgency, and have to be steered across organisational borders. This thesis looks at that follow-up as a long-term organisational process, and not as a single project or report.

To address this problem, this thesis develops and evaluates a structured follow-up and governance method for ICT-related recommendations. The intended contribution is twofold. Practically, the study aims to provide SSC-ICT with a usable artefact that supports more coherent monitoring, prioritisation, and steering of recommendations over time. Academically, it aims to contribute to understanding how organisational processes, roles, and decision structures shape the translation of advisory inputs into implementation action in complex public-sector environments.

1.2 Research questions

How can a structured follow-up and governance method be designed and evaluated to support government organisations such as SSC-ICT in monitoring, prioritising, and steering the implementation of ICT-related recommendations over time?

RQ1 How are ICT-related recommendations from audits, evaluations, and advisory reports currently followed up and governed within SSC-ICT?

RQ2 What limitations and challenges do stakeholders experience in the current follow-up and governance of ICT-related recommendations within SSC-ICT?

RQ3 What requirements for an effective follow-up and governance method for ICT-related recommendations can be derived from academic literature and organisational practice?

RQ4 How can these requirements be translated into a structured follow-up and governance method for ICT-related recommendations within SSC-ICT?

RQ5 To what extent does the proposed method support stakeholders at SSC-ICT in terms of perceived ease of use, perceived usefulness, and intention to use after application on a case?

These questions reflect the design-science oriented nature of the study. The research does not only aim to analyse an existing governance problem, but also to develop a practically useful and theoretically grounded artefact. Accordingly, the study combines problem analysis, requirements analysis, artefact design, and evaluation within a design science research approach.

1.3 Overview of the thesis

The rest of his thesis follows the steps of the design science process. ‘

Chapter 2 gives the background and related work. It reviews the literature on public-sector ICT, on audits and advisory reports, on governing recommendations under uncertainty, on the fragmentation of follow-up across reports and decision arenas and on temporal decay. It also introduces the effect chain of follow-up, implementation, and impact and describes the Dutch context in which SSC-ICT operates.

Chapter 3 explains the methodology used. The design science approach, the case study at SSC-ICT, the data sources and the way there are analysed, the evaluation criteria, and the handling of ethics and data.

Chapter 4 reports the empirical work and turns it into requirements. The receiver-side interviews and the themes that come from them, the four process interpretations of how recommendations are currently followed up, the producer-side interviews, and the set of design requirements that follows from these inputs and literature.

Chapter 5 describes the artefact, namely the recommendation register with its three-layer status model, the intake, the review and closure, the learning loop, the mapping of existing tools and the governance cycle, and it separates the general elements from the case-specific ones.

Chapter 6 demonstrates the artefact by applying it tot historical recommendation cases from the SSC-ICT context, to show that the register can hold the recommendations the organisation receives.

Chapter 7 validates the artefact in a stakeholder workshop, using a questionnaire based on the Method Evaluation Model, and reports the results on perceived ease of use, perceived usefulness, and intention to use.

Chapter 8 answers the research questions, gives contribution to theory, and describes the future work.

Chapter 2

Background and related work

This chapter gives the background for the study. It reviews literature on public-sector ICT, recommendation follow-up, governance under uncertainty, and the Dutch governmental context in order to clarify the problem space and review concepts relevant for the artefact design.

2.1 Public-sector ICT and policy implementation

Information and Communication Technologies (ICT) have become increasingly embedded in public administration, extending beyond support for routine administrative work into broader governmental operations and policy implementation. Early uses of ICT in the public sector focused primarily on automation, data storage, and internal efficiency, such as document archiving and transaction processing. However, in recent decades, ICT has increasingly been integrated into the core governance processes, influencing coordination, decision-making, and the organisation of public services [21].

Liu and Yuan describe in a broad literature review covering developed and developing-countries that public-sector ICT has moved from separate, stand-alone tools to connected systems that work throughout the organisation, where technology is linked to changes in rules, structures, and how policies are implemented [21]. Their longitudinal review suggests that the literature increasingly links ICT not only to office automation, but also to inter-governmental communication, service delivery, citizen participation, and policy and decision-making.

Because of this, researchers no longer look only at efficiency. They also study how ICT changes the way public services create value. Cordella and Paletti conceptually distinguish between a manufacturing logic and a service logic of ICT-enabled public service production [9]. In traditional "manufacturing logic", ICT is used to standardise processes, improve control, and improve efficiency within bureaucratic organisations. In contrast, service logic highlights more open forms of value creation, in which ICT can enable co-production with external actors such as citizens, firms, and civil society organisations [9]. From this view, ICT changes not only how services are provided, but also who takes part in putting policies into practice and how value is created.

In a conceptual paper, Misuraca and Viscusi argue that ICT-enabled innovation in the public sector should be analysed in relation to governance processes and policy-making mechanisms, rather than as a purely technical matter [25]. Their framework suggests that ICT-enabled innovation can reshape governance relationships, participation, openness, and inclusion, especially in networked governance settings. They also note that limited uptake and diffusion of ICT-enabled innovation in the public sector often reflects organisational fragmentation, institutionalised distrust, and misaligned incentives rather than purely technical shortcomings. This shows that ICT should be studied as a built-in part of how policies are carried out, not as a separate tool used only after decisions are taken.

The role of ICT in policy implementation is further elaborated by Di Giulio and Vecchi in a theory-building study with Italian e-government cases, who conceptualise ICT as an "ordering force" that can restructure organisational interdependencies and governance modes [14]. Their study suggests that different ICTs require different intra- and inter-organisational interdependencies and therefore call for different governance arrangements and implementation strategies. This affects how people coordinate, who has authority, and how systems are put into practice. As a result, using ICT directly changes how policies turn into everyday work, and it shapes discretion, accountability, and learning inside and between public organisations.

What these contributions have in common is that they do not treat ICT as an administrative support tool only. They treat it as a part of how policy is organised and put into practice. ICT shapes who has authority, how coordination is done, and how well public organisations can carry policy decisions through. For this reason, policy implementation in modern public administration cannot be understood without including ICT in the picture.

2.2 Policy recommendations, audits, and advisory reports in government

Governments rely extensively on audits, evaluations, and advisory reports as instruments to monitor performance, ensure accountability, and guide organisational and policy change. In the context of public-sector ICT, these instruments can play a particularly prominent role. For example, in the Indonesian Ministry of Finance, large ICT investments, operational dependence on information systems, and the growing relevance of IT governance increased the importance of IT audit practices [31]. Therefore, ICT governance audits, project reviews, and advisory reviews are commonly used to see if ICT resources are being managed in line with the organisation's goals, legal rules, and standard governance practices.

In public-sector ICT, ICT audits identify deficiencies, risks, and inefficiencies and provide recommendations for improvement, particularly with regard to security, compliance, and operational effectiveness [31]. In ICT governance contexts, reviews and improvement efforts may also be structured through recognised frameworks and governance models. In the Indonesian Ministry of Finance case, IT governance audits were explicitly structured through COBIT 4.1 and COBIT 5 [31], while the Finnish central-government study by Lappi, Aaltonen, and Kujala shows how project governance and portfolio management can be used to connect ICT projects to higher-level digitalisation objectives. Ideally, such recommendations are not treated as a one-time compliance task, but as inputs for organisational learning and improvement over time.

Formal follow-up mechanisms may include action plans, timelines, reporting obligations, and periodic monitoring or review. In the Tanzanian case discussed by Innocent and Kwesigabo, timely implementation is explicitly linked to management engagement, resource allocation, communication, and subsequent monitoring of implementation [15]. Their case study found that where management support is strong, recommendations are more likely to be translated into concrete actions, supported by resource allocation and structured monitoring [15]. In the Finnish central-government context, structured project governance and portfolio management were used to connect ICT projects with higher-level digitalisation strategy across organisational levels, thereby improving alignment and oversight [19].

Dain and Rahmat study this in the Malaysian public-sector audit context. They find that the follow-up of recommendations does not only depend on a formal action or review process. It also depends on the accountability mechanisms around the implementation [11]. In their case these mechanisms are the follow-up audit, the Audit Committee, the accountability index, and the Auditor General's Dashboard. The follow-up audit checks again whether earlier recommendations are acted upon. When a recommendation stays open, it can be reported again. The Audit Committee reviews the recommendations, decides on corrective and preventive actions, and monitors who is responsible. The accountability index is a rating mechanism. It scores the compliance and compares the performance between organisations. The Auditor General's Dashboard shows the implementation status of audit issues to the public. It uses a traffic-light logic to show whether an action is taken or not [11]. Their survey of 226 public-sector auditees found that the accountability index, the Audit Committee, the Auditor General's Dashboard, auditors, and follow-up audit were associated with actual implementation, while media, the Audit Committee, and the accountability index were associated with intention to implement [11]. Although this is a non-ICT audit context in Malaysia, it is relevant here because it shows why correlations of this kind matter: recommendation follow-up is shaped not only by formal plans, but also by whether recommendations are made visible, compared, revisited, and reinforced through concrete accountability arrangements.

However, the formal adoption of ICT governance processes does not automatically guarantee effective implementation in practice. In a Brazilian state-company pilot, Canedo et al. show that implementation depended strongly on process maturity, coordination between actors, and understanding of the processes involved [6]. When these conditions are absent, follow-up usually depends on ad hoc actions, personal effort, or short-term project setups, and not on fixed governance routines. This makes it fragile, as the recommendation depends on people instead of a process.

Capacity and resources also matter. ICT-related recommendations may require structural changes, new competencies, or investments that exceed organisational capabilities, which can contribute to delayed or partial implementation [15, 6]. In addition, public organisations often deal with clashing priorities, political pressure, and strict budget cycles, which makes steady follow-up even harder.

Another reason is that the recommendations may not really align with how things work in practice. Formal frameworks do not by themselves ensure workable implementation. In the Brazilian pilot studied by Canedo et al., processes initially believed to be mature and implemented proved, on closer examination, to be incomplete or outdated, while Lappi, Aaltonen, and Kujala identify perceived gaps between strategic digitalisation goals and operational ICT project governance in the Finnish central-government context [19, 6].

Finally, there is also a temporal element to recommendation follow-up. ICT environments evolve rapidly, while audit recommendations are issued at a specific moment in time. Innocent and Kwesigabo focus primarily on timely execution [15], which leaves the longer-term question of continued relevance under changing technological and organisational conditions largely unaddressed. This could logically contribute to situations in which recommendations are formally "open" for years, even when their original assumptions are no longer valid.

This broader follow-up problem is also visible in the Dutch policy-advice context, where cabinet responses to advisory reports often acknowledge the issue and list actions already underway, while concrete steps for following up specific recommendations remain limited [18].

What the studies above show is that audits, evaluations, and advisory reports are important instruments in public-sector governance. Their impact, however, is limited. The follow-up is fragmented, capacity is constrained, the recommendations do not always fit the local context, and the longer-term relevance is rarely revisited. These limitations are especially visible in ICT-heavy settings.

2.3 Governing ICT-related recommendations under uncertainty

In the public sector, ICT recommendations are usually written as concrete advice for action. They suggest changes in systems, processes, governance structures, or capabilities, which we assume will help to improve performance, lower risks, or create more public value. This kind of framing reflects what Orlikowski and Iacono call the "tool view of technology", in which technology is treated as a relatively unproblematic resource with specified capabilities [28]. In the public-sector reform literature, Cordella and Bonina argue that this kind of view has often informed ICT reform thinking by treating ICT as a tool for a relatively predetermined process of change, assessed mainly through managerial and economic lenses [8]. However, the core challenge for governing such recommendations is that their effects are rarely fully knowable at the moment they are issued, and they may remain difficult to assess even after implementation. This uncertainty is not accidental. ICT-enabled reforms are shaped by wider political and organisational contexts [8]. They also often create dependencies between units and organisations that make implementation more complex [14]. As a result, their effects are difficult to predict in advance, particularly because technologies and policy priorities continue to evolve [13].

A first source of uncertainty concerns the nature of the "outcomes" that the recommendations implicitly aim to achieve. Public-sector ICT initiatives are often assessed through managerial lenses (for example, efficiency and cost reduction), but such assessments can be too narrow for reforms whose consequences extend to legitimacy, fairness, and trust. From the public value perspective developed by Cordella and Bonina, ICT-enabled public sector reforms involve multiple, partly contested goals that cannot be reduced to a single value-for-money measure [8]. Because these objectives reflect collective expectations and political priorities, they are complex, context-dependent, and not fully predefined. In practice, this implies that the "impact" of implementing a recommendation can remain ambiguous: the same intervention can improve one dimension of public value while creating tensions in another [8].

A second source of uncertainty is whether a recommendation is feasible. It is also a question how something that is technically possible can be turned into a measure that organisations and political actors can really implement. A digital government initiative can be technically achievable, but it can still not fit the organisational structures, the resources, or the political constraints. Dawes reviews e-governance in U.S. state and local government. The study shows that e-governance develops in an environment of changing risks and opportunities. New policy issues keep adding to an already complex set of concerns, and the policy landscape stays uneven and incomplete [13]. Implementers must therefore read a recommendation inside policy requirements that change, for example on privacy, security, and oversight, and inside changing expectations about cross-boundary information use. So a recommendation is not a static task. Initiatives that are technically possible may not be organisationally feasible or politically or socially desirable, and this can change when the policy and governance context around them changes [13].

A third type of uncertainty stems from the ways in which ICT structurally transforms governance itself. In a theory-building study illustrated with four Italian e-government cases, Di Giulio and Vecchi conceptualise ICTs as "ordering forces" that can restructure power and authority when introduced into a system of relations [14]. They further argue that different ICTs generate different intra- and inter-organisational interdependencies, and that policy actors must learn how to govern them and activate learning processes in response to failures [14]. For recommendation governance, this suggests that follow-up should include mechanisms for learning and adaptation, rather than being treated only as a box-ticking compliance exercise.

A fourth source of uncertainty concerns path dependence and infrastructural constraints, particularly in public administrations characterised by legacy systems and fragmented IT landscapes. In a multi-case analysis of digital transformation in European public administration, Irani et al. show that legacy systems create multiple layers of complexity for digital transformation, including interoperability and integration issues arising from systems that have been built incrementally in isolation, frequently linked to specific statutory roles and changing policy frameworks. For recommendation follow-up, this means that a recommendation can be reasonable in theory, but still hard to carry out in different kinds of infrastructures, and that some side effects only become clear when we start with the implementation. Their analysis also shows that for a transformation to work well, we must understand in detail not only the systems that change, but also the policies these systems support and the broader administrative context in which they function. This further supports the view that the follow-up of ICT recommendations needs to be managed with clear consideration of dependencies and context, since these elements consistently

determine what can be implemented and what outcomes can reasonably be expected [16].

Finally, uncertainty also arises from gaps in the existing evidence base and from the constraints of current methods used to measure ICT-enabled governance innovation. In a conceptual paper based on a broad literature review, Misuraca and Viscusi argue that empirical evidence on the effects of ICT-enabled innovation on public governance remains limited, and that clearer and better-tested indicators are still needed [25]. They note that there is not yet evidence of direct intervening effects on governance systems, indirect effects on institutional settings, or clear causal connections to policy outcomes and socio-economic impacts [25]. In such a setting, governance cannot rely only on linear implementation-to-impact assumptions. This suggests that recommendation governance should be approached as an iterative and longitudinal process that monitors upcoming effects and supports organisational learning over time.

What we see in these sources is that governing ICT-related recommendations is more than checking whether the proposed actions were done. ICT-based reforms can create more than one type of public value at the same time. They develop in policy and risk situations that keep changing. They also depend on existing infrastructure and on choices that were made earlier. As a result, both the prediction before implementation (ex ante) and the assessment afterwards (ex post) come with a level of uncertainty. This creates a governance problem. If recommendations are treated as one-off deliverables, organisations either close them too early, without knowing what the longer-term consequences are, or they keep them open for years because no clear impact can be shown. A follow-up approach that fits these insights should therefore allow for ongoing steering under uncertainty: it should make context and dependencies explicit, and it should include periodic review moments, instead of assuming that implementation and impact can be fully predicted in advance.

2.4 Ex ante anticipation vs. ex post accountability

Within the e-government literature, research on evaluation has long been described as relatively limited, while ex post evaluation is discussed less frequently than ex ante assessment, despite its relevance for learning and correction over time [17]. In the literature on developing countries specifically, Choi et al. similarly note that most reports focus on ex ante evaluations and that ex post evaluations are rarely discussed [7].

Ex ante evaluation is defined as a forward-looking assessment carried out before the implementation, with the purpose of mapping the consequences and giving the option to compare possible alternative solutions. In Jukić, Vintar, and Benčina, ex ante logic is linked to predicting costs, benefits, effects, opportunities, and challenges before an intervention, while ex post evaluation takes place after implementation. Ex post evaluation focuses on identifying actual effects (for example, organisational, economic or social results) and also offers feedback that can be used to improve future ex ante assessments [17].

This connection between anticipation and accountability is quite logical when positioned in the broader policy-cycle view. Jukić, Vintar, and Benčina place ex ante evaluation in the phase where alternatives are selected and adopted, and ex post evaluation in the phase where policy effects are assessed after implementation [17]. In that sense, ex ante evaluation helps structure selection and prioritisation upfront, whereas ex post evaluation provides feedback on what actually happened after implementation.

Ex ante assessment is constrained by the information available before implementation. In practice, this often means that anticipation is translated into explicit assessment categories. Jukić, Vintar, and Benčina show that many ex ante and ex post evaluation approaches are organised around recurring parameters such as costs, benefits, and risks [17].

A complementary insight comes from research on e-government implementation in developing countries, that evaluation frameworks can be used not only to "judge" or "audit" outcomes, but also to prioritise what should receive attention under resource constraints. Choi et al. propose an assessment method that explicitly aims to support both ex ante and ex post evaluation and argue that prioritisation is especially relevant where resources are limited. Their approach consists of a so called multi-dimensional framework (strategy, technology, organisation, people, environment) with a prioritisation technique called analytic hierarchy process (AHP) to give weight to factors, thereby supporting decisions about where corrective action or prevention measure are most needed [7].

However, tools and procedures for ex ante assessment do not, by themselves, guarantee meaningful anticipation. In the study of Wanckel of a digital tool for sustainability impact assessments in the German Federal Government, tool use varies with intra-ministerial coordination practices and portfolio-specific routines. The study also shows that, in that specific setting, prevailing norms and routines make in-depth consideration of policy alternatives and consequences unlikely, making the tool more inclined to facilitate formal compliance than to encourage genuine substantive reflection [36].

Meyerhoff Nielsen describes the Danish case as one in which strategy cycles are linked to both ex ante and ex post evaluation. In that context, e-government strategies are supported by mandatory business cases and recurring reporting on progress, effects, and goals, while high-risk initiatives may be redesigned, delayed, or stopped at an early stage [23]. At the same time, this case also points to the reporting burden such arrangements can create and to remaining limitations in relation to change management and longer-term post-implementation learning [23].

A similar cyclical logic appears in Canedo et al.'s ICT governance kits for Brazilian state-owned enterprises. Their approach combines maturity levels, associated practices, and artefacts with a continuous-improvement cycle in which organisations monitor whether improvement actions produced the expected results and use detected deviations as input for corrective action in a new cycle [5]. This frames accountability not only as reporting, but also as iterative verification and adjustment [5].

2.5 Fragmentation of governance across reports, time, and decision arenas

A recurring challenge in governing recommendations is that follow-up is rarely organised as a single, integrated process. Rather, recommendations move across different report types and oversight bodies, unfold over different time horizons, and depend on decisions taken in different organisational arenas. As a result, ownership, monitoring, and steering can become fragmented: recommendations may be registered in one place, discussed in another, and only much later implemented, reinterpreted, or de-prioritised in yet another setting [3, 4, 27, 19].

Fragmentation is visible at the level of basic monitoring infrastructure. The ANAO report on the Australian Department of Health shows that even a well-established audit follow-up routine can remain fragmented at the level of monitoring infrastructure [3]. In its review of recommendation follow-up in the Australian Department of Health, the ANAO found that the monitoring process relied on iterative spreadsheets and emailed documents. That arrangement created risks of data loss through manual transfer and made it difficult to trace the implementation course of a recommendation without consulting multiple spreadsheets, written documents, and hard-copy files. The ANAO also found that several recommendations from audited reports had not been captured in the spreadsheet at all [3]. This fragmentation is not only an efficiency problem; it also affects the ability to prioritise consistently, to compare progress across units and time, and to maintain a shared picture of what is still "open" and why.

The difficulty becomes sharper when recommendations travel across decision arenas and depend on actors beyond the organisational unit it currently concerns. In the UNFPA review of Joint Inspection Unit recommendations, some recommendations remained outstanding for five years or more, partly because implementation depended on inter-agency harmonisation or political processes rather than on internal organisational decisions alone [4]. Although this is a UN inter-organisational setting, it is analytically relevant here because it shows how recommendations can become long-lived coordination objects when implementation depends on actors beyond the focal organisation. The review also points to a structural limitation of the web based tracking system (WBTS): when recommendations are accepted in principle, but cannot yet be implemented (for example, due to lack of opportunity or funding constraints), the system's binary status categories ("implemented" vs. "not implemented") can leave such recommendations marked "in progress" indefinitely, without capturing the underlying reason for non-implementation [4]. This shows clearly how a specific mechanism with recommendations with roots in different arenas, can translate into long-lived ambiguous follow-ups and delayed decisions.

At the organisational level, fragmentation can also be understood as a coordination and institutionalisation problem. The OECD toolkit discusses this in the adjacent field of public policy evaluation: without clear guidance on who does what and when, evaluation-related activities can take place in isolation and drift away from decision-making. While the toolkit is not about recommendation follow-up specifically, the mechanism is transferable: where responsibilities and coordination routines are unclear, follow-up activities are more likely to become dispersed and weakly steerable over time [27]. The same toolkit explicitly characterises low maturity situations as decentralised and fragmented, with strong differences in approaches and standards and no central visibility over what has been evaluated and for what purpose [27]. Whilst this is discussed in the context of evaluation, the described failure mode closely matches the governance challenge of recommendations that originate from multiple sources, if responsibilities and coordination mechanisms are unclear, then follow-up becomes dispersed, inconsistent, and hard to steer over time.

Lappi, Aaltonen, and Kujala reinforce this point in the context of government digitalisation: in the Finnish central-government setting they studied, authority still largely followed existing organisational layers, centralised steering and reporting remained limited, and multiple regular and irregular reports and statements burdened projects, making integrated governance more difficult [19].

Fragmentation also appears in the organisation of follow-up itself. A governance arrangement can spread the follow-up on purpose over a network of actors. Umor, Zakaria, and Sulaiman study this in the Malaysian performance-audit context. There, the follow-up on audit issues was organised through coordination, cooperation, and collaboration between auditors, auditees, regulators, and other actors. A dedicated follow-up division had a central role and coordinated both inside and outside the audit institution [34]. This is a context-specific case, and it gives most attention to issues with punitive elements such as crime, corruption, and disciplinary matters. It is still relevant here, because it shows that follow-up can become a distributed coordination process across actors and forums, and not only an administrative tracking exercise [34]. This need for coordination is also where things can go wrong. The information about a recommendation must be connected to the moments where decisions are made. The information side is what is recommended, by whom, and with which dependencies. The decision side is who decides the priority, the scope, the resourcing, and the closure. When these two sides are not connected by a clear structure, the follow-up becomes a set of separate activities and not a process that can be steered.

Reading these sources together, fragmentation does not come from one cause. Tracking is often spread across different artefacts and handovers, which makes the longitudinal picture incomplete [3]. Some recommendations depend on processes that sit between organisations or in the political arena, and a simple status system cannot represent this well [4]. Responsibility is often not clearly assigned, and the coordination between institutions is weak [27]. On top of that, in multi-level governance structures, authority, reporting, and steering stay distributed across the layers without enough central oversight [19]. For ICT-related recommendations, where dependencies and context change over time, these mechanisms together make it hard to monitor consistently, to prioritise correctly, and to steer the work as a connected governance cycle. The follow-up then becomes a set of separate tasks instead.

2.6 Temporal decay, obsolescence, and learning failure

Beyond fragmentation across reports and decision arenas, the literature also suggests that recommendation follow-up can weaken over time, as implementation trajectories extend across changing organisational, political, and technological contexts. Recommendations are typically issued at a specific moment, based on a specific risk picture, organisational context, and set of dependencies. Yet, implementation of recommendations can often extend over multiple years, and as times passes, the meaning of an "open" recommendation can become ambiguous: it may still be relevant and blocked by constraints, it may have become outdated by changing conditions, or may have been informally superseded without being explicitly reconsidered.

A first issue related to time is that recommendations remain formally open for years without an actual representation of why. In the UNFPA review context, Callejas shows that 20 JIU recommendations had remained outstanding for five years or more and that some recommendations stayed "in progress" because implementation depended on factors such as inter-agency harmonisation, political processes, lack of opportunity, or funding constraints. He further notes that the tracking system's binary implementation choice could leave such recommendations marked as "in progress" without adequately representing the underlying reason for continued non-implementation [4].

A second issue related to time is that recommendations may lose fit with the context in which they were originally formulated. This can be illustrated indirectly in domains such as information security governance. Magnusson et al. do not study recommendation follow-up itself, but their review shows that public-sector information security governance depends on continuous measurement, risk assessment, and adjustment in an increasingly complex threat landscape [22]. This suggests that governance mechanisms should support periodic reassessment, rather than assuming that a recommendation retains the same meaning, priority, or feasibility over time.

A third issue concerns learning failure. This happens when the follow-up focuses on administrative closure and not on adaptation and feedback. In their theory-building study of public-sector digitalisation, Di Giulio and Vecchi argue that ICTs act as ordering forces. They reshape power, authority, and the interdependencies between actors [14]. The authors also stress that policy actors must learn how to govern these technologies, and that they activate learning processes in response to failures [14]. The same point applies to recommendation governance. The governance needs mechanisms to revisit the assumptions, to reinterpret the actions taken, and to adjust the ownership when implementation experience gives new information. Without such mechanisms, two problems can occur. The first is premature closure, where a recommendation is closed before the consequences are understood. The second is indefinite openness, where progress cannot be brought to a clear decision.

Relatedly, the OECD toolkit links impact to governance design: evaluation recommendations are unlikely to lead to tangible change unless there are institutional mechanisms that ensure ownership, accountability, and integration into decision-making processes. The same toolkit also argues for a learning-oriented evaluation culture, in which unexpected outcomes are treated as opportunities for learning rather than as failures to be punished. Taken together, these points show that follow-up over time requires both decision-making integration and room for adaptation and learning [27]. This connects directly to temporal decay: without structured follow-up mechanism that keep recommendations connected to decision moments over time, recommendations can drift endlessly into administrative routines that neither close them meaningfully nor generate learning about what should change.

A practical illustration of time-structured steering appears in Denmark's governance arrangement for large ICT initiatives. Meyerhoff Nielsen describes how systematic risk reviews and recurring progress reporting can trigger the redesign, delay, or termination of high-risk ICT initiatives, while also linking project-level follow-up to benefit realisation and broader strategic adjustment cycles [23]. At the same time, this case also highlights a tension as periodic reporting and review can be perceived as an extra burden, indicating that temporal governance needs to balance learning and steering benefits with organisational capacity constraints [23].

What comes out of these sources is that temporal decay is not just slow implementation. It is a structural risk in governance. Recommendations can stay open for years without their status really saying why. They can lose fit with the context in which they were written. And they can fail to produce learning when the follow-up never revisits the original assumptions. This is a reason to treat recommendations as longitudinal objects in the follow-up. They need periodic reconsideration, interpretation against the

current context, and a review focus on learning, instead of being handled as a task that is either done or not done.

Table 2.1 provides a compact overview of the main challenges identified in the literature regarding the implementation and monitoring of recommendations. It is intended as a brief problem-oriented summary of the literature discussed above.

Table 2.1: Main literature-based challenges in the implementation and monitoring of recommendations

Challenge	**What the literature points to**	**Indicative references**
Translation into action	Recommendations do not automatically become implementation actions. Even where formal follow-up mechanisms exist, actual execution can remain partial, delayed, or dependent on management support, resources, and workable organisational embedding.	[15], [6], [19]
Uncertainty of effects and feasibility	ICT-related recommendations are often issued as if their effects and feasibility are reasonably knowable upfront, while in practice outcomes, constraints, and dependencies may shift during implementation.	[8], [14], [25]
Ex ante / ex post disconnect	Anticipation and accountability are often weakly connected. Ex ante assessment may structure choices beforehand, while ex post evaluation and feedback remain underdeveloped or insufficiently connected to later steering and correction.	[17],[7], [23],[6]
Fragmented follow-up and monitoring	Follow-up is often not organised as one integrated process. Recommendations may move across reports, actors, and decision arenas, while tracking information is dispersed and responsibilities remain unclear.	[4], [27], [34], [19]
Temporal drift over time	Recommendations may remain formally open for years, lose fit with changing contexts, or fail to generate learning. Status categories can mask delayed feasibility, while review cycles may be too weak to support reconsideration and adaptation.	[4]; [22], [14], [27], [23]

2.7 The effect chain of advisory research: a lens from Dutch Rekenkamer practice

The sources discussed so far describe what we could call follow-up, implementation, and impact mainly as separate parts of a broader governance problem. A recent Dutch contribution by Lemmens integrates these parts into one analytical framework [20]. Lemmens writes from the context of decentral 'Rekenkamers', which are the audit bodies of Dutch municipalities and provinces. That setting is clearly different from the setting of a central-government shared service provider. The actors, the political arena, and the legal basis of the recommendations are not the same. The mechanics that Lemmens develops, however, describe how a recommendation moves from a published report toward actual change, and these mechanics are not tied to the municipal context alone. We therefore use the framework here as a conceptual lens, while keeping the contextual limits in mind. A producer-side interview with Lemmens is conducted later in this thesis as part of the empirical work in section 4.3: this section only uses his published work.

Separating follow-up, implementation, and impact

Lemmens argues that the terms follow-up ("opvolging"), implementation ("doorwerking"), and impact are often used interchangeably in practice, but should be kept apart analytically [20]. He proposes an effect chain ("effectketen") in which the three terms are three consecutive stages, and he defines each stage separately. Figure 2.1 summarises the chain, together with the characteristics, time horizon, and measurability that Lemmens assigns to each stage.

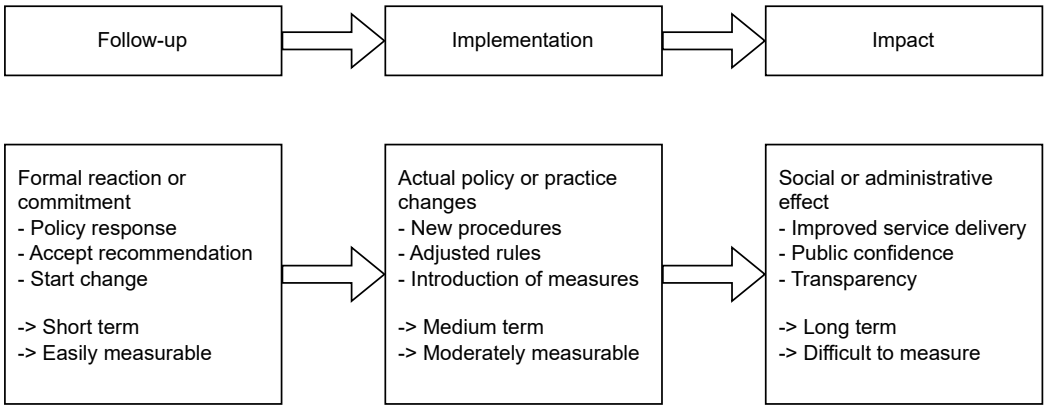


Figure 2.1: The effect chain of advisory research: follow-up, implementation, and impact. Adapted from Lemmens [20].

Follow-up, in the definition of Lemmens, is the first stage of the chain and covers the direct, formal reaction to the recommendations. In his municipal context this means agenda setting in the council, a formal decision to adopt, reject, or amend the recommendations, and any commitments, motions, or stated intentions that follow from that decision [20]. He describes this stage as relatively easy to measure, because it consists mainly of documentable formal acts, and as visible within a short time frame of weeks to a year. Lemmens is explicit that follow-up is only meaningful when commitments are "clear, specific, and measurable [...], with realistic timelines and adequate resources, rather than non-committal intentions" ("helder, specifiek en meetbaar [...], voorzien van realistische tijdslijnen en adequate middelen, in plaats van vrijblijvende intenties") [20]. Equally important, according to Lemmens, is the quality of the system in which follow-up is recorded.

Implementation is the second stage of the chain. Lemmens reserves this term for the actual translation of formal commitments into structural organisational change: new work protocols, training of staff, improved control systems, redesigned responsibilities, and adjusted financial procedures [20]. He stresses that this stage goes beyond "paper promises" ("papieren beloftes") and requires real anchoring in "the daily practice" ("de dagelijkse praktijk") of the organisation. The time horizon is longer, typically one to three years, and measurability is only moderate.

Impact is the third and final stage. Lemmens describes impact as the eventual societal or governmental effects of the implemented changes: better public services, higher effectiveness and efficiency of policy, more transparency, or improved public trust in government [20]. He acknowledges that this stage is difficult to measure, that it requires long time horizons of two to five years or more, and that it is sensitive to many mitigating factors that lie outside the producer's reach. He also notes that impact is not always positive: recommendations may fail to produce the intended effect, cause unwanted side effects, or be overtaken by time [20].

For the present thesis, the main contribution of this separation is that it makes visible what several earlier sources were implicitly treating as the same thing. Reports of "open" recommendations (Section 2.6) or of "implementation status" (Section 2.5) can in fact refer to different stages of the chain. A governance model for recommendation follow-up therefore needs to be explicit about which stage it is tracking.

Stake, influence, and interest along the process chain

Alongside the effect chain, Lemmens applies the "logic model" of input, throughput, output, outcome, and impact to the Rekenkamer process itself [20]. His key analytical move is to map each stage of that process onto what the Rekenkamer can actually affect. The input stage (topic selection, research design, resources) is fully under the Rekenkamer's own control. The throughput (research execution, interviews, right of reply) and the output (the report, presentation, communication) are also largely under its control. From the outcome phase and further, the situation is different. How the report is used in the political arena, and how it is later carried out and put into practice, is no longer under direct control, the Rekenkamer can only influence them. The final impact lies outside even that influence, and the Rekenkamer only has an interest in it. Lemmens summarises this distinction as a distinction between control, influence, and interest along the chain, as visualised in Figure 2.2 [20]. He explicitly notes that "the link between a Rekenkamer recommendation and an improvement in practice as perceived by the citizen is generally not directly demonstrable" ("de link tussen de aanbeveling van de Rekenkamer en een door de burger waargenomen verbetering van de praktijk is over het algemeen niet direct te leggen") [20]. Causality becomes less strong when you go further along the chain.

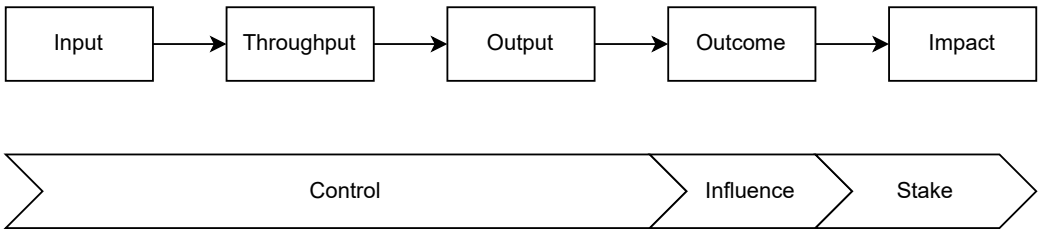


Figure 2.2: Control, influence, and interest along the stages of the rekenkamer process, mapped onto the "logic model" of input, throughput, output, outcome, and impact. Adapted from Lemmens [20].

This mapping is analytically useful for the present thesis, even though SSC-ICT is not a producer of advice but a receiver of it. The same logic applies, but mirrored. Where a Rekenkamer loses control after publication, a receiving organisation only *gains* control once a recommendation enters its own sphere. A governance model for recommendation follow-up on the receiver side therefore has to answer a complementary question: which stages of the incoming recommendation are under the organisation's own control, which can only be influenced (for example, through cross-organisational dependencies), and which lie outside both.

Contextual limits and transferability

Two limitations of the Lemmens framework should be made clear before we use its elements further. First, the decentral Rekenkamer context that Lemmens describes is placed in a political arena where one report goes to one council, which then decides on it with a formal vote, and the result of this vote is open for the public. SSC-ICT receiving advice through multiple audit and advisory channels sits in a different kind of political arena. External visibility is not absent: SSC-ICT falls under the Ministry of the Interior and Kingdom Relations, and some ICT-related reports do reach the Second Chamber, senior civil service governance, or the broader CIO-system. This visibility, however, goes in an indirect way through the ministerial line and different control bodies, and is not brought together in one clear decision place like a municipal council. The point of the analysis is not that external political pressure is smaller, but that it comes into the organisation in a more diffuse and less predictable way. Because of this, the role of the external anchor in the chain cannot be copied one-to-one. Second, the effect chain as Lemmens describes it, assumes there is one producer (the Rekenkamer) that gives recommendations to one political arena (the council). In a receiver-side situation with several producers who give recommendations into a multi-layered governance system, the chain maybe must be set up for each source separately, and not only once for the whole organisation.

2.8 The Dutch context

In the Netherlands, for many years we see problems with managing government ICT not as loose events, but as structural issues. In "Parlementair onderzoek naar ICT-projecten bij de overheid" the parliamentary inquiry presents ICT not as a narrowly technical issue, but as a matter that cuts across government organisation, control, and implementation, and links weaknesses in ICT governance to broader problems in responsibility, steering, and oversight [33]. The diagnosis already confirms two core patterns highlighted earlier in this review: (1) Governance fragmentation across decision arenas and (2) the difficulty of producing sustained follow-up that survives changes in time horizons and institutional attention.

A key theme in "Parlementair onderzoek naar ICT-projecten bij de overheid" is fragmentation in responsibilities and steering information. The report discusses fragmented and unclear responsibilities, limited portfolio steering, and shortcomings in steering information that can create an overly reassuring impression of project control [33]. The report also points to fragmented information provision and to the dispersal of ICT issues across departments and parliamentary committees, which means that oversight and decision-making in the Dutch central-government context remain spread across multiple arenas rather than being consolidated in one place [33]. This spread across institutions reflects a common point in the literature: recommendations may be documented in one setting, formally approved in another, and later put into practice under different interpretations.

At the same time, the Dutch case also illustrates that formal monitoring infrastructures can exist without actually resolving the underlying governance problem. The parliamentary report describes annual reporting obligations for large and high-risk ICT projects to the CIO Rijk, letting the reports be processed in public reports (the Jaarrapportage Bedrijfsvoering Rijk) and the Rijks ICT-dashboard [33]. At the same time, that visibility remains partial: only projects above specific thresholds are included, and the reported information is often limited or uneven, for example, with respect to expected management and maintenance costs and the specification of benefits [33].

More recent Dutch discussions frame these issues within a broader execution perspective. Although the "Staat van de Uitvoering 2024" is not limited to ICT governance, it does identify outdated IT, limited data exchange, and strong organisational autonomy as recurring system-level barriers in public service delivery [32]. The report argues that many bottlenecks are well known while cross-organisational solutions remain limited, and therefore calls for stronger direction ("regie") and less non-commitment ("vrijblijvendheid") in inter-organisational digital coordination [32]. This aligns with the earlier emphasis that ICT-related recommendations often depend on interdependencies and shifting context and therefore require governance mechanisms that connect information, authority, and decision moments across organisational boundaries.

The temporal dimension also appears in the more specific Dutch case of the basisregistraties (base registers). In "Grip op gegevens" the Algemene Rekenkamer shows that earlier recommendations from 2014 had only been followed up to a limited extent, as also reflected in the 2017 follow-up monitor, which formed part of the reason for the 2019 follow-up investigation [1]. Although this report concerns the system of base registers ("basisregistraties stelsel") rather than Dutch ICT governance as a whole, it is relevant here because it shows how governance issues can persist across multiple review cycles when coherence, central steering, and concrete follow-up remain insufficient [1].

Another specific mechanism relevant to recommendation follow-up appears in cabinet responses to external advice. The report "Kabinetsreacties, nader bekeken", a study of cabinet responses to reports on the functioning of government more broadly, and not only ICT-related reports, finds that responses often acknowledge the problem and refer to actions already underway, while concrete follow-up to specific recommendations remains limited [18]. At the same time, the study also notes that limited concreteness in the immediate cabinet response does not necessarily mean that the advice has no effect, since advice may still have indirect, agenda-setting, or delayed influence over a longer period [18]. This nuance reinforces the need for governance models that track not only "implementation status", but also reinterpretation, absorption into other programmes, and explicit closure rationales over time.

A concrete Dutch ICT case also shows how recommendation follow-up can remain underdeveloped despite substantial review activity. In its evaluation of the follow-up to BIT and Gateway advice on the DSO, the Auditdienst Rijk concludes that progress had been made, but that monitoring and steering of recommendation follow-up were not yet adequately organised and that measures were often too weakly specified to determine when advice had actually been implemented [2]. This is a single case in the Omgevingswet context, but it illustrates how recommendations can remain difficult to govern when ownership, intended effects, and monitoring arrangements are insufficiently structured.

The more recent Dutch digitalisation agenda also shows attempts to organise follow-up in a more programmatic way. The "Verzamelbrief digitalisering december 2025" describes how recommendations from "Ons Digitaal Fundament" (Our digital fundament) are incorporated into the execution of the "Nederlandse Digitaliseringsstrategie" (Dutch digitalisation strategy) through explicit prioritisation, scope choices, and attention to implementation consequences [35]. The same letter provides examples where advice triggers concrete governance and monitoring actions, such as strengthening steering arrangements and financial accountability structures following advice on Digipoort [35]. At the same time, the letter records delays and longer timelines in other trajectories due to complexity, multiple stakeholders, and unresolved dependencies, and it refers to planned ex ante and ex post measurement in the development of policy instruments [35]. Because this is a government policy letter rather than an independent evaluation, it is best used here as evidence of current follow-up practice rather than as proof of effectiveness.

The Dutch sources together show a context in which responsibilities are fragmented, ICT-project visibility depends on thresholds and is only partial, the same problems return in several diagnostic cycles, and cabinet reactions do not always translate into commitments that can be monitored [33, 18, 1]. At the same time, the Dutch case also shows recent attempts to improve the steering: calls for binding agreements and central direction, programme-based prioritisation, and more explicit monitoring arrangements. In this kind of environment, a structured follow-up and review model is both relevant and well-timed [32, 35, 2].

Chapter 3

Methodology

This chapter explains how the research was designed and conducted to answer the research questions and develop the artefact. It outlines the overall research strategy, the case context, the execution of the design science process, the data sources, and the evaluation approach.

3.1 Research Design

This thesis follows a design science research approach because the central aim is not to only analyse a problem, but to design and evaluate and artefact that supports the actual practice. This artefact aims to be a useful structured follow-up and governance method for ICT-related recommendations over time. Design science research builds and evaluates an artefact through a structured, repeatable process, which separates it from consultancy or an unsystematic case description. For this, the Design Science Research Methodology (DSRM) process model is followed [29].

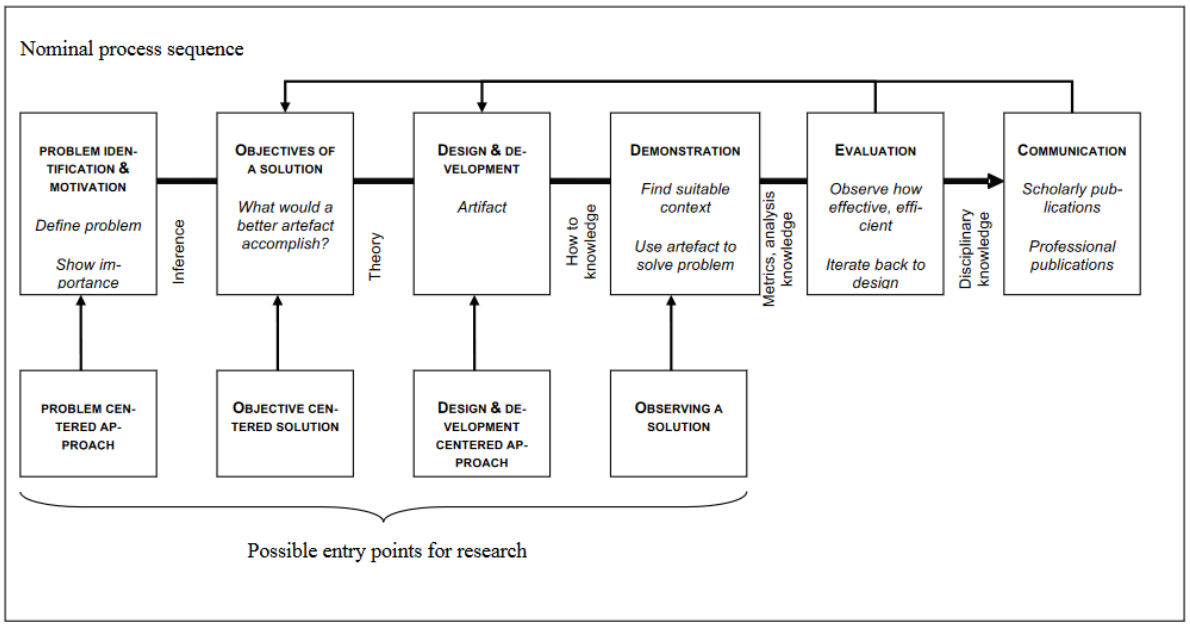


Figure 3.1: Design science research process (DSRM) model [29]

The DSRM process model gives the overall structure of the research, but it does not specify how stakeholder goals, requirements and design choices are linked together inside that process. For this link, we follow Wieringa, who decomposes the design task into problem investigation, treatment design, and treatment validation, and he introduces the contribution argument as the structured reasoning that ties an artefact in its context to the goals of stakeholders.

Design science is combined with qualitative case study. The case study component is to ground the artefact in an empirical organisational setting, to capture how follow-up currently works in practice, and to evaluate the method's usability and perceived value with stakeholders who would realistically use it. The design science component is required to put the empirical insights and literature synthesis into context to be able to create a coherent testable governance method and to evaluate it against its intended purpose [29].

3.2 Case context and scope

The setting is SSC-ICT, a shared service provider in the Dutch central government landscape. The analysis performed here is the organisational follow-up and governance of ICT-related recommendations originating from internal and/or external audits, evaluations, and advisory reports, including how these recommendations are monitored, prioritised, revisited over time and translated (or not) into steerable implementable actions.

The Dutch context is really relevant, because Dutch public-sector ICT governance is characterised by issues on unclear responsibilities and recurring challenges of translating acknowledgments of issues into monitorable commitments. These issues and dynamics are visible in public governance and audit sources that discuss long ongoing ICT concerns, follow-up patterns and attempts to improve the current governance situation [33, 1, 18, 32, 35].

3.3 Execution of DSRM

The research activities are structured according the six DSRM activities. This does not mean that it is a strictly linear execution. The artefact is refined iteratively as insights from analysis and evaluations feed back into the design [29].

1. Problem identification and motivation: The practical and governance problems are established by analysing how ICT-related recommendations are currently tracked and governed, and where fragmentation, loss of visibility or temporal obsoletion occur. This information is gathered by document analysis of relevant Dutch public-sector reports and follow-up documentation, and qualitative interviews within the organisation. The used sources in the literature review sketch the landscape of why follow-up and steering of ICT advise is difficult and why structured monitoring and review is relevant.

2. Objectives of a solution: The problem is analysed into explicit design objectives and requirements. According to Wieringa, this is the requirements specification step within treatment design, in which the requirements are documented together with a contribution argument that links each requirement to a stakeholder goal. These objectives are derived from three inputs, the literature review (Chapter 2), the empirical analysis at SSC-ICT as well as analysis of contributions from producers of recommendation reports. The intention is to define what the method must enable in practice, without describing a tool or other arrangements [29].

3. Design and development: The primary artefact is a follow-up and governance method. This phase consists of iteratively specifying the information structure needed to make recommendations governable, the governance cycle by which recommendations are prioritised, reviewed and escalated and the decision moments and roles that connect recommendations to implementation planning and steering. The design is informed by the literature review and the empirical observations at SSC-ICT.

4. Demonstration: The artefact is demonstrated by applying it to selected historical recommendation sets (cases) drawn from the documents relevant to the SSC-ICT context. Demonstration is used to show the method can actually structure and represent real recommendation streams and to make the artefact concrete enough for meaningful stakeholder evaluation [29]. Evaluation: The artefact is

evaluated through a stakeholder validation workshop, after it has been demonstrated on the cases. Participants work with the demonstrated method and then assess it with the Method Evaluation Model of Moody, which measures perceived ease of use, perceived usefulness, and intention to use [26]. This is read together with the requirement coverage from the demonstration and with the open feedback given during the workshop. The evaluation thus checks whether the method meets the design requirements in Chapter 4 and whether stakeholders find it usable and useful in practice [37, 29].

5. Communication: The communication of results is done through this thesis document and through the artefact that is provided to SSC-ICT.

3.4 Data sources and collection

3.4.1 Document analysis

Policy-related and advisory documents relevant to ICT governance and recommendation follow-up in the Dutch government are analysed. Public Dutch sources are used to both contextualise the governance environment and provide realistic recommendation cases for demonstration. Examples of such documents include performance and oversight (audit) reports and follow-up analyses addressing ICT governance and implementation from internal governmental as well as external third party advisory sources.

3.4.2 Semi-structured interviews

Semi-structured interviews are conducted with two distinct categories of stakeholders, reflecting two different analytic purposes.

The first category consists of interviews with stakeholders inside SSC-ICT who are involved in ICT governance, implementation, oversight, and recommendation follow-up. These interviews are used to reconstruct and validate how follow-up currently works in practice, including pain points and information gaps, and to surface recurring problems across roles. They primarily inform RQ1 and RQ2 and later serve as the basis for stakeholder evaluation of the demonstrated artefact. The interview outline can be found in Appendix A

The second category consists of producer-side interviews with stakeholders at organisations that generate ICT-related recommendations directed at the Dutch central government. These interviews serve three main purposes: (i) to cross-check the receiver-side results with an external viewpoint on the same governance context, (ii) to uncover producer-side perspectives on recommendation quality, feasibility, and follow-up that are not accessible from within SSC-ICT, and (iii) learn from a well-established follow-up practice.**The interview outline can be found in Appendix B**

The two categories differ not only in position but in sample size and role heterogeneity. The receiver-side sample consists of interviews across partially overlapping roles inside one organisation, which is sufficient to identify recurring cross-role patterns. The producer-side sample consists of interviews across four different institutional positions, where each respondent represents a distinct position in the Dutch advisory and oversight landscape rather than a replication within a single role. This difference in data character motivates two different analytical approaches, described in Section 3.4.3.

3.4.3 Data analysis and synthesis approach

The first interview material is analysed using a qualitative coding approach informed by Miles, Huberman, and Saldaña. In line with their view of qualitative analysis as an iterative process of data condensation, data display, and conclusion drawing and verification, the analysis did not proceed as a strictly linear sequence, but moved back and forth between coding, clustering, comparison, and interpretation [24].

In a first coding cycle, interview fragments were coded inductively with short descriptive codes close to the language of the respondents. Following Miles, Huberman, and Saldaña, these codes functioned as labels that captured the primary meaning of relevant data segments and supported the retrieval and comparison of fragments across interviews. In a second cycle, related first-cycle codes were grouped into higher-order categories and themes through pattern coding. Miles, Huberman, and Saldaña describe pattern coding as a way of condensing larger amounts of qualitative material into a smaller number of more meaningful analytic units, such as themes, categories, or constructs [24].

The analytical approach just described fits the receiver-side interviews, where the analytic aim is to identify recurring patterns across partially comparable roles inside one organisation. It does not fit the producer-side interviews in the same way. Each of the producer-side respondents occupies a distinct institutional position, so the purpose of the analysis is not to identify patterns through repetition across comparable cases, but to compare the perspective each respondent brings from their specific position in the advisory and oversight landscape. For a small heterogeneous sample of this kind, inductive pattern coding produces misleading results: the apparent frequency of a code carries no analytic weight when each interview represents a different institution.

For the producer-side interviews we therefore use the role-ordered matrix approach described by Miles, Huberman, and Saldaña as a data-display technique for comparing role perceptions across issues of interest to a study [24]. In this method, the rows stand for the people in the roles (here, the producer-side respondents). The columns stand for the analytic categories that come from the research questions and from the gaps we found in the receiver-side analysis. In each cell there is a short summary of what the respondent said about that category, with some direct quotes when they make the meaning more clear. The matrix is already an analytic activity, and not only some step before the real analysis. Miles, Huberman, and Saldaña say very clear that choosing the rows and columns of a matrix, and choosing which data to put in the cells, together is already part of data condensation and analysis [24]. The matrix allows two forms of reading: down a column to compare how respondents in different institutional positions see the same issue, and across a row to reconstruct one respondent's coherent position across all categories. The category set used for the columns of the matrix follows the structure of the interview protocol.

One part of the protocol needs to be explained separately, because it decides what is in one of the columns of the matrix. In the confirmatory section, each respondent gets a short list of obstacles that came up in the receiver-side analysis at SSC-ICT.

During the analysis, analytic memos were used to document reflections on recurring issues, emerging relationships between codes, and possible interpretations of the data. This supported the transition from coded fragments to broader thematic patterns and helped make the analytical decisions more explicit [24]. The resulting themes were then compared across interviews and linked back to the research questions. In this way, the analysis served both to structure the empirical findings and to synthesise design requirements for the artefact.

Given the design science orientation of this thesis, the goal of the analysis is not to produce statistical generalisations or to claim saturation in a grounded-theory sense, but to reconstruct how recommendation follow-up is currently understood and organised within SSC-ICT and to surface recurring problems across roles. The thematic structure resulting from the two coding cycles therefore feeds directly into the requirements synthesis (Section 4.4), where each requirement is grounded in both the empirical themes and the literature reviewed in Chapter 2.

3.4.4 Evaluation criteria

The artefact produced is a governance method, therefore evaluation focusses on the practical and organisational value, rather than the technical performance. In the terms of [37], the evaluation in this thesis is a validation, not an implementation evaluation, as the artefact is applied to a historical case and assessed by stakeholders before any real-world transfer to operational use, so the goal is to predict how the artefact would interact with the SSC-ICT context rather than to study a real-world implementation [37]. The validation is done against the synthesised requirements, derived from the empirical results and from the base of Dutch government documents.

The artefact is a method: it describes how recommendations are registered, assessed, owned, reviewed and closed over time. A method cannot be judged as true or false, only as effective or not, and this can only be shown through use in practice [26]. For the evaluation an instrument is therefore needed that shows whether the method helps the work and whether stakeholders would want to use it.

For this we use the Method Evaluation Model (MEM) of Moody [26]. The MEM was made to validate information systems design methods. It combines two dimensions of method "success": the actual efficacy, whether the method improves the task, and the adoption in practice, whether people want to use it [26]. A method that improves the work but is not used has no effect, and a method that is used but lowers performance has a negative effect; both dimensions are therefore needed [26]. This fits the artefact of this thesis: the method only helps SSC-ICT when it both supports the work and is accepted by the people who must use it. The analyses in Section 4.1 point in the same direction, where several respondents said that an instrument felt as administrative burden will not be used.

The MEM builds on the Technology Acceptance Model (TAM) [12] and translates its three constructs to the domain of methods [26]. These are Perceived Ease of Use, how much effort a person thinks the method costs. Perceived Usefulness, how effective a person thinks the method is in reaching its goals. And Intention to Use, how much a person intends to use the method [26]. In the MEM these perceptions, and not the measured efficacy, determine whether a method is adopted [26]. The validation in this thesis takes place before real use, so actual efficiency and effectiveness cannot yet be measured. Therefore, the perception-based constructs form the correct level for measurement.

The instrument is the post-task survey of Moody, adapted to this artefact. It has sixteen items: six for Perceived Ease of Use, eight for Perceived Usefulness and two for Intention to Use, the same division as Moody [26]. The Perceived Usefulness items are written in terms of the goals of the method, namely keeping recommendations visible over time, making ownership explicit and supporting steering, because usefulness can only be judged against the goals of the method [26]. The wording is changed from the use of a system to the use of the register. Each item uses a five-point Likert scale. Following Moody, half of the items are negatively formulated and the items are placed in random order, so that respondents keep attention for each statement [26]. The full form is in Appendix C

The survey is filled in once, after the participants have applied the method to the case, including the update from the intake to the later review moment. This follows the procedure of Moody [26] and makes sure that the three constructs refer to the same experience with the method. The questionnaire is filled in individually and before the plenary discussion, so that each answer is not influenced by the opinions of others. The analysis of the survey results, together with the group exercise, the observation and the validation, is described in Chapter 7.

3.4.5 Ethics and data handling

All participation in this study was voluntary. Participants were informed of the purpose of the research and the recording before each session and verbal consent was obtained before the recording started. Participants had the opportunity to ask for the recording to be paused or stopped at any time. The receiver-side interviews, the producer-side interviews, and the validation workshop were recorded for analysis. The recordings are stored on the Leiden University OneDrive and are not shared outside of research. They will be deleted after finalisation of the thesis trajectory.

Receiver-side participants are reported anonymously, by job title only, and quotations are not attributed to individual participants. The producer-side respondents speak for their institutional position: they are identified at the level of the institution. Lemmens is named, because the interview builds directly on his published work, and he agreed to this.

The internship at SSC-ICT is covered by a non-disclosure agreement. For this reason, internal documents are not reproduced in this thesis, and the demonstration in Section 6.1 is built on public reports only. Where internal context was needed to fill the register, it was provided by the company supervisors and is described at the level of the register fields, without reproducing the underlying documents.

3.5 Scope and intended generalisability

The artefact is designed and evaluated in the SSC-ICT case, but the method is specified at a level intended to be transferable to similar government organisations that deal with recurring ICT-related recommendations across multiple oversight and advisory channels. The thesis therefore distinguishes between elements that address general follow-up governance needs and case-specific elements tied to the SSC-ICT environment and Dutch central-government practice.

Chapter 4

Requirements synthesis

This chapter translates the empirical findings and literature review into design requirements for the artefact. It first summarises the key insights from the interview analysis and then synthesises these insights into a coherent set of requirements for the follow-up and governance method.

4.1 Interview analysis

For this thesis, eight participants have participated in an exploratory interview which form one of the basis for the synthesis of requirements.

Participant	**Job Title**	**Date**
Participant 1	Department Manager Finance and Control	17-02-26
Participant 2	Head Internal Audit and Advisory	18-02-26
Participant 3	Executive Secretary to the COO	26-02-26
Participant 4	Coach Transition and Transformation	26-02-26
Participant 5	Concern Controller Planning, Strategy and Management Information	25-03-26
Participant 6	Project-controller	30-03-26
Participant 7	Board secretary and Advisor Business unit Infrastructure Services	30-03-26
Participant 8	Concern Controller	09-04-26

Table 4.1: Overview of interview participants

To analyse how ICT-related recommendations are currently followed up, these interviews were held in different parts of the organisation. In line with the methodological approach, the interviews are not used to measure frequencies or produce statistical generalisations, but to reconstruct how recommendation follow-up is understood and organised in practice, where this differs across contexts and which problems recur across roles. The process schemes should therefore be read as rough analytical reconstructions of described practise rather than as official process documents.

4.1.1 Theme coding overview

Table 4.2: Thematic structure derived from inductive coding of interview data

Theme	RQ	Codes	Focus
Analytical themes			
T1 — Current follow-up practices	RQ1	41	Existing processes, instruments and routines for tracking audit findings and recommendations
T2 — Governance structure and roles	RQ1	29	Formal governance bodies, escalation paths, and authority structures
T3 — Responsibility and ownership gaps	RQ2	28	Unclear, dispersed, or avoided ownership of recommendations
T4 — Monitoring and visibility gaps	RQ2	33	Absence of consolidated status views, logging and proactive tracking
T5 — Prioritisation and decision-making	RQ2	23	How recommendations are (not) prioritised and what drives escalation
T6 — Recommendation characteristics	RQ2	43	Quality, abstraction, scope and origin of recommendations
T7 — Organisational capacity and culture	RQ2	63	Capacity constraints, silo behaviour, and cultural resistance to process
T8 — Tools and governance instruments	RQ3	35	Existing and desired tooling for recommendation management
Contextual themes			
T9 — Agile transformation context	—	27	Ongoing Agile/OKR transition shaping the governance environment
T10 — External oversight and pressure	—	21	CIO-Rijk, Rekenkamer, and consultancy as external accountability mechanisms

Note: Codes may appear in more than one theme where they address overlapping dimensions.

Following Miles, Huberman, and Saldaña, the themes summarised in Table 4.2 are the result of a second, pattern-coding cycle that condensed the 300 first-cycle descriptive codes generated from the interview material into a smaller set of analytically meaningful categories. To make this condensation traceable, the full code-to-theme mapping is provided in Appendix D (Table D.1). Each row shows the first-cycle codes that are put together under the same pattern-coded theme. The codes keep the short description labels that were used in the first analysis cycle. The display serves three purposes: (i) it shows the empirical basis for each theme, (ii) it makes clear the internal differences inside one theme (for example, sub-patterns that can justify more detailed separation in the narrative), and (iii) it helps the traceability matrix from themes to design requirements that we introduce in Section 4.4.6.

Cross-theme codes Codes that appeared in more than one theme (for example, *Dispersed monitoring* in T3 and T4; *Escalation happens too little* in T2 and T5; *Hiring external consultancy firm* in T6 and T10) were on purpose kept under each related theme during pattern coding. This is because the original interview parts had analytic meaning on more than one dimension. In line with Miles, Huberman, and Saldaña, we see these cross-theme codes as analytic links and not as mistakes in classification. They show the moments where themes meet in the data. Several of these codes come back in Section 4.1.2, where we discuss how fragmentation, dispersed ownership, and capacity limits strengthen each other. A consolidated traceability matrix linking these themes onwards to the design requirements is provided in Chapter 4.

4.1.2 Cross-cutting observations from the interviews

This section shows the main findings from the eight interviews. They are grouped under the ten pattern-based themes (T1–T10) in Table 4.2 and explained more in detail in Table D.1. The goal here is not to describe the follow-up process itself; that is done with the four process models in the next section. Instead, the goal is to show the patterns that come back in different roles and settings, and that together form the basis for the design requirements in Chapter 4.

In line with the analytical approach in Chapter 3, we report our findings on the level of repeating patterns, not on the level of single participants. When we use a quotation, it is to make the main point more clear. If an observation comes from only one interview, we say this clearly, and we use this observation only as an example, not as proof of a general pattern.

T1 - Current follow-up practices

Across the interviews, people said that follow-up practices were not the same for everything. For results from the financial and IT audit cycle (*accountant findings*, *audit log for IT*, *auditlog and auditmonitor*), several participants separately described almost the same procedure: the finding is registered in a central log, classified, assigned to a line manager, worked on by a task owner, and only closed after independent verification (*closure requirement*, *recommendation closure conditions*, *response cycle*).The procedure was explained as being in a phase where it sustains itself with only limited management involvement. In one account, it was said: "the audit side, via the audit log and audit monitor, is relatively well in place" ("de auditkant via auditlog en auditmonitor is relatief goed ingeregeld").

For the broader category of advisory and policy-driven recommendations, there was no similar built-in procedure explained. One participant summarised the contrast as follows: "there is no tight process at the back-end through which we have secured this" ("[er is] niet een heel strak proces aan de achterkant waarmee we dat geborgd hebben"). A related observation was that "*follow-up*" was not understood as a synonym for implementation: an explicitly documented decision not to act on a recommendation was treated as a valid form of follow-up, provided the rationale was recorded (*administrative response*, *recommendation has been followed*). This already signals a design choice that the artefact will need to support: closure rationales other than "done".

T2 - Governance structure and formal roles

Governance was described as having many levels and being spread over quite a large number of different bodies (*many governance bodies*, *gremia*, *board structure includes DMT*, *customers' council*). One account listed a sequence of meetings ("overleggen"); bestuurlijk overleg, afnemersberaad, financiële commissie dienstverlening, technisch overleg, and questioned whether coordination between these forums was strong enough so that the recommendations arrive at the right level. Final decision-making on priority and ownership was consistently described as a directieteam (DT) matter (*management-team has final say in priority*, *CEO has final say*); escalation otherwise followed the hierarchical line (*strategic-tactical-operational*, *hierarchical structure governs escalation*).

Two repeating conditions were added to this picture. First, formal escalation was described as something that "in our organisation, this actually happens too rarely" ("gebeurt bij ons dan eigenlijk soms te weinig in de organisatie"), suggesting that the formal route is available but underused (*escalation happens too little*). Second, several accounts noted that the distance between the level at which a recommendation is received and the level at which it is executed is substantial: typically several organisational layers (*too much distance to top governance*, *managers are people, process and technology*). This observation returns under T3 and T7 as one of the mechanisms through which ownership and capacity problems come into existence.

These two observations combine, together with one already noted under T2, into a single mechanism. The translation distance between the C-level, at which a recommendation is received, and the team that ultimately performs the work was described as several organisational layers (T2). At the execution layer, the link back to the originating recommendation tends to have disappeared by the time the work has been packaged into a plan or project: the work is performed against the plan, not against the report. Taken together, this means that follow-up can stall without anyone noticing, the people closest to the work do not know they are following up on a specific recommendation, and the people who do know are too far removed to monitor progress directly. For artefact design, this means that only having a central register is not enough; the register must also be connected downwards to the workflow tools and plans through which the work is actually done.

T3 - Responsibility and ownership gaps

Ownership was the topic that came back the most often in the interviews. Two related but still different patterns appeared. The first concerned the assignment of ownership: when ownership was not written down clearly with "name and jersey number" ("naam en rugnummers"), discussions about who was responsible kept coming back, and the work did not move forward (*need for clear responsibility assignment*, *need for responsibility assignment is the basis*, *unclear responsibility*). One account described the downstream effect of unclear assignment in concrete terms: "and when that employee then realises... I haven't done anything with this, and starts saying, well, that's not mine, that's Justin's responsibility" ("en als die medewerker dan op een gegeven moment erachter komt van ... ik heb hier niks mee gedaan, en begint te blaten van, ja daar ben ik niet van, daar is Justin van").

The second pattern was about ownership that goes over the borders of one organisation. Some of the recommendations were seen as not fully under SSC-ICT's direct sphere of influence (*cross-organisation responsibility*, *own sphere of influence*, *recommendation area shift*). These included, for example, procurement that must follow government-wide framework contracts, service-chain dependencies with the parent ministry, or agreements between two parties that SSC-ICT cannot decide on alone. As one account put it: "*[het] zit niet in onze cirkel van invloed, het zit bij een opdrachtgever of bij een kadersteller ... daar houd mijn verantwoordelijkheid wel op*" (it is not within our sphere of influence, it sits with a client or a framework-setter, that is where my responsibility ends). For these recommendations the question is not only *who* owns the action, but also *how far* that ownership reaches.

In one account, a condition was added that is relevant even when a tracking system exists: action lists made with SMART goals, deadlines, and reminders still did not lead to regular follow-up. The organisation described action tracking as something they generally struggle with, not only in certain specific areas. This points to ownership as more than an administrative attribute: explicit assignment is a precondition for follow-up, but on its own it does not guarantee that the action is carried through.

T4 - Monitoring, visibility and transparency gaps

Outside the audit cycle, no consolidated overview of outstanding recommendations was described. Tracking was reported to be spread out across documents, lists, and individual notes (*no central logging of outstanding actions*, *no standard log*, *dispersed monitoring*, *fragmented follow-up*). One account characterised the overall pattern as follows: "each time some form of monitoring is improvised... but it is not uniform and not brought into mutual coherence" ("ik denk dat iedere keer wel weer een vorm van monitoring wordt geïmproviseerd ... maar het is niet eenduidig en niet in onderling verband met elkaar gebracht").

Two consequences of this fragmentation were described. First, because there is no shared overview, it is hard to set priorities between the different recommendation streams, since each list is shown on its own (*shared overview*, *unclear overview*, *need for transparency*). One participant summarised this as a need for "all commitments in one system... in one overview system" ("alle toezeggingen in één systeem ... in één overzichtsysteem"). Second, the link between an originating recommendation and the eventual operational work was described as weak: by the time a recommendation has been translated into a plan or project, awareness of the original report tends to be lost (*operations does not have recommendations in mind*, *operations only look at the plan handed to them*, *recommendation traceability*). This makes it difficult to verify, after the fact, whether the recommendation was actually addressed by the work that followed.

T5 - Prioritisation and decision-making

Prioritisation across recommendations was described as underdeveloped. For audit findings, severity classifications and deadlines provided a partial ordering (*weighted findings*, *deadline setting*, *technical priorities are well established*). For the broader recommendation portfolio, however, no shared mechanism for weighing recommendations against one another was reported (*no structured value assessment of recommendations*, *difficult priority assessment*, *need for prioritisation*). One account framed the underlying difficulty as an "everything is important" issue: because there are no clear weights, every list seems like it has the highest priority, and choices between the lists are then made in a more hidden, indirect way.

A second point was about how long we need to keep the prioritisation going. People explained that recommendations are more like living artefacts and not just one-time tasks:"sometimes a trajectory of several years to actually resolve it" ("soms een traject van meerdere jaren ... om het echt opgelost te krijgen"), during which urgency, feasibility, and scope can shift (*years-long trajectory*, *recommendation stays open over time*, *recommendations stalling due to time and energy constraints*). This suggests that prioritisation cannot be a one-time act at intake, but needs to be revisited at logical moments, such as around governance meetings or reporting cycles (*need for logical assessment moments*).

T6 - Recommendation characteristics and quality

A first observation, distinct from the quality of any individual recommendation, concerned the diversity of *sources* from which recommendations enter the organisation (*recommendation source*, *recommendation initiative*, *reasons for outsourcing advice*). Across the interviews, inputs were described as arriving through the yearly ADR audit cycle, the Rekenkamer, ACICT reports, BIT and Gateway reviews, statutory *doorlichtingen* commissioned by another directorate, ad-hoc external studies (the accounts mentioned KPMG, Deloitte, Sopra Steria, and PwC in different contexts), and internally-requested studies through IAA. Each of these channels was explained as having its own rhythm, its own way of connecting things, and its own way of receiving input. One account framed the absence of a connecting intake point directly: "in government, it always comes from different directions. There is not one central point where you can line up all the studies; it does not work that way" ("[het] komt bij de overheid altijd uit verschillende hoeken. Er is niet ergens één centraal punt waar je alle onderzoeken kan naastgaan, zo werkt het niet"). This wide range of origins is the practical basis for the observation, mentioned again under T4, that there is no single main overview to prioritise against.

Beyond source diversity, three sub-patterns surfaced about the recommendations themselves. The first concerned abstraction: recommendations were described as sometimes too high-level to act on without further interpretation (*recommendations are too abstract*, *recommendation abstraction leaves room*, *need for clarification of recommendations*). One account illustrated this with the example "improve the security of BZ" ("verbeter de security van BZ"), a recommendation that, as formulated, can be decomposed in many different ways before implementation can begin.

The second sub-pattern concerned recommendations that restate what the organisation already knows, sometimes referred to as open goals ("inkoppers"). One account described the limited value of such recommendations as follows: "those are rather obvious points, and the rest, I cannot do anything with that" ("dus dat zijn nogal van die inkoppers en de rest, daar kan ik helemaal niks mee"). The same account contrasted these recommendations with the advice from outside experts, who had more specific knowledge. Their advice was described as more directly useful.

The third sub-pattern concerned timing and contextual fit: where the underlying object of a recommendation is itself still developing quickly, for example a cost-price model under active revision, formalising procedures around it was described as inefficient, because the resulting procedure can be overtaken within months (*recommendations difficult to execute*, *recommendations often hang around*). Together, these sub-patterns suggest that the artefact will need to make the *quality and fit* of an incoming recommendation explicit, rather than treating every recommendation as equally implementable by default.

T7 - Organisational capacity and cultural barrier

This was the largest theme by code count, and it surfaced in every interview. Three sub-patterns are worth distinguishing.

Capacity concentration
Follow-up work tended to be delegated to people who were already heavily loaded with operational responsibilities (*already busy people are the people being asked*, *colleagues being over-asked*, *lacking time*). One account observed: "for this kind of work, the same people are usually asked those with the relevant knowledge, and those are typically the people who are already very busy" ("vaak in dit soort dingen [worden] dezelfde mensen ... gevraagd, die wel die kennis hebben van de zaken. Dat zijn veelal de mensen die vaak al heel druk zijn"). Another account noted that the team was currently being "somewhat over-asked" ("een beetje overvraagd"), and that prioritisation by line and C-level management needed to act as a stronger filter on incoming work.

Silo behaviour
Cooperation within organisational units was described as functioning reasonably well, but cooperation across units less so (*within-silo cooperation works well*, *outside-of-silo cooperation is not well organised*, *working outside of own island is difficult*). One account captured the pattern as "it is all fairly organised in little silos, in islands... working across islands doesn't go all that well" ("het [is] allemaal redelijk georganiseerd ... in silootjes, in eilandjes ... over de eilandjes heen gaat samenwerking niet zo heel goed"). For recommendations that span multiple units, which is typical for cross-cutting advisory recommendations, this logically raises the coordination cost. A concrete instance was described in one account in which multiple security programmes were independently asking the same people for the same deliverables, without coordinating with each other; the duplication only got noticed when one of the teams flagged it. For recommendations that span multiple units, this raises the coordination cost and creates the risk of parallel work that the organisation only notices late.

Reactive default Several accounts described an organisational default of acting on explicit instruction rather than on initiative (*organisation is reactive instead of proactive*, *organisation culture requires concrete instructions*, *organisation driven on assumptions*).One person also mentioned that it is often not correct to assume that the person who gets a delegated task already understands the whole context: "the assumption is made that Pietje Puk knows exactly what to do; that is not the case" ("de aanname wordt gedaan dat Pietje Puk precies weet wat hij moet doen. Dat is niet zo"). A related observation was that structured, procedural work was not seen as attractive (*formalised processes are not sexy*, *little attention for processes*). For artefact design, the implication is that a method that relies on implicit ownership or self-organising follow-up is unlikely to function in this context: instructions and ownership need to be explicit.

T8 - Tools and governance instruments

A range of tools was described as being in active use across domains: the audit log and audit monitor (Excel-based) for audit findings (*audit log and monitor is an Excel-file*, *auditlog and auditmonitor*); TopDesk for change and ticket workflows (*TopDesk platform monitoring*); Confluence and Jira on the agile side; FCC for portfolio change; a Qlik Sense dashboard for OKR and capacity overviews; and various working documents for taskforce-style follow-up (*transformatiebord tool*, *task management in lists*). One account also recalled an earlier dedicated registration system, Plato, with sub-modules for different commitment streams (parliament, ADR, Rekenkamer) in an Oracle environment.

Two observations cut across this tool landscape. First, no single tool was described as covering recommendation follow-up across sources; tooling tended to be domain-specific (*unclear tool use*, *need for clear governance tool*, *need for central governance of recommendations*). Second, several accounts emphasised that any new instrument should plug into existing workflows rather than form a parallel track: "we should not invent a new way of working, but connect to the existing one" ("we moeten geen nieuwe werkwijze bedenken, maar aansluiten [bij] de bestaande werkwijze"), and "the overview must be linked to work streams, so that it actually surfaces in a team's work and is treated as work" ("het overzicht moet wel gekoppeld zijn aan werkstromen, dus dat het ergens opkomt bij je team, dat het daadwerkelijk als werk wordt gezien"). This means for the artefact that just adding a register is not enough; the register must also be connected to the workflow tools where the actions are already followed.

T9 - Agile transformation context

The interviews were conducted while a transition to a SAFe-style agile way of working was being introduced top-down across the organisation (*agile working is top-down mandated*, *from structure transforming governance*, *translation strategic to tactical OKR*, *workflows are being organised in lean portfolios*). Strategic OKRs are being translated into tactical OKRs and from there into portfolio steering; work is being decomposed into epics on Jira; value streams are being defined. Several accounts framed this transition as both an opportunity and a constraint for recommendation governance.

As an opportunity, the transition reorganises responsibilities in ways that may create more room for governance work: managers currently combining responsibility for people, process and technology are expected to be split into separate roles such as product owner, scrum master, and people lead (*managers are people, process and technology*). As a constraint, the transition itself absorbs attention and capacity, and the organisation was described as still in the early phase of the change (*agile working being implemented*, *agile working requires mature organisation*). For the artefact, this means that its design should fit with the new tools that are coming (Jira, Lean Portfolio Management, Qlik Sense) and not stay linked to the old situation before the transition.

T10 - External oversight and pressure

External control systems were seen as an important reason why people keep to the follow-up. The audit cycle has a fixed external rhythm: the ADR comes back every year, puts the results in categories, and looks again at points that are still not solved (*external audit CIO-Rijk*, *ADR*, *Rekenkamer*). For one-off advisory reports, this external rhythm is largely absent (*external pressure does not function*, *no mandatory follow-up*). One account observed that even where external pressure exists, follow-up is not always organised accordingly: "even with external pressure, it is not properly allocated" ("zelfs met externe druk wordt het niet goed belegd").

A second observation concerned the role of external advisory parties as both producers and amplifiers of recommendations. One account described how external reports can give weight to points that the organisation had already raised internally (*external pressure is good*, *hiring external consultancy firm*), while also noting the limits of external independence (*independence consultancy firm*), meaning they are paid to do the job according to the initiators' will. For the artefact, this implies that both the origin of a recommendation and any associated external accountability cycle should be clearly included as metadata, since they determine how frequently follow-up must be reported.

A third observation, raised in only one interview but concrete enough to record as a design constraint, concerned the confidentiality of recommendation content and the need for controlled cross-organisational visibility (*confidentiality discussion*). Reports and findings were described as carrying differing sensitivity classifications, and visibility across organisational boundaries. For example a control sitting at one ministry but relevant to another, was described as a question the organisation has to settle case by case: "what is confidential and may everyone see, and what is not confidential and may not everyone see" ("wat is geheim en mag iedereen zien en wat is niet geheim en mag niet iedereen zien"). Because this point came up only in one interview, it is written here more as a limitation to keep in mind for artefact design. Access control and visibility across boundaries should be clearly shown in the register, but we do not see this yet as a confirmed pattern for the whole organisation.

Synthesis of observations

Read together, the ten themes describe an organisation in which follow-up is well embedded for one specific category of input, the recurring financial and IT audit cycle, and largely informal, person-dependent, and dispersed for the broader landscape of advisory and strategic recommendations. The core mechanisms behind this asymmetry are visible across the themes and reinforce one another: the absence of a consolidated overview (T4), the under-specified ownership of recommendations that fall outside the audit channel (T3) and outside the organisation's direct sphere of influence (T2, T3), the lack of a shared way of weighing recommendations against one another (T5), the uneven quality and fit of incoming recommendations themselves (T6), the concentration of follow-up work on already-loaded staff (T7), the domain-specific nature of available tooling (T8), and the absence of an external rhythm to mark review moments outside the audit cycle (T10). The ongoing agile transition (T9) reshapes the environment in which any artefact will have to function. These observations form the empirical input for the design requirements developed in Section 4.4; the four models presented in the next section provide the procedural input.

4.2 Processmodeling

What emerged from the interviews is not a single shared process, but four distinguishable patterns for how recommendations enter SSC-ICT and how they are handled afterwards. The four patterns are not choices for the same input. Each one belongs to another type of recommendation source and to another level of how much it is built into the process. Interpretation 1 describes the handling of audit findings from the Auditdienst Rijk (ADR), which is the only category with a formalised end-to-end process. Interpretation 2 describes the handling of external advisory reports commissioned by the directieteam[1] (DT), for which a dedicated programme or steering group ("stuurgroep") is typically created. Interpretation 3 explains how to deal with the other strategic and policy-related recommendations, for which there is no similar structural integration yet. Interpretation 4 is not its own process, but more a feature that can be used with recommendations from any of the first three. Some recommendations cannot be handled only inside SSC-ICT's own sphere of influence, so they need coordination between different organisations.

Only Interpretation 1 is supported by documented process descriptions and by consistent accounts across roles. The remaining three are reconstructions of how participants described actual practice in their own domains; they should be read as analytical reconstructions rather than as official process descriptions. Therefore, the diagrams below are on purpose quite high-level. Each one shows the main structural feature that makes one pattern different from another, instead of showing all detailed operational steps. In this way, the asymmetry in follow-up coverage becomes visible, and this works as the empirical basis for the design requirements that we develop in Section 4.4.

4.2.1 Formalised audit follow-up

The most structured follow-up pattern concerns audit findings from the ADR, covering both financial and IT audits. Several participants with direct responsibility for this process described it in largely consistent terms, which provides reasonable confidence that this pattern reflects embedded organisational practice rather than an individual interpretation. This pattern aligns primarily with theme T1 (current follow-up practices) and T2 (governance structure and formal roles) in Section 4.1.1.

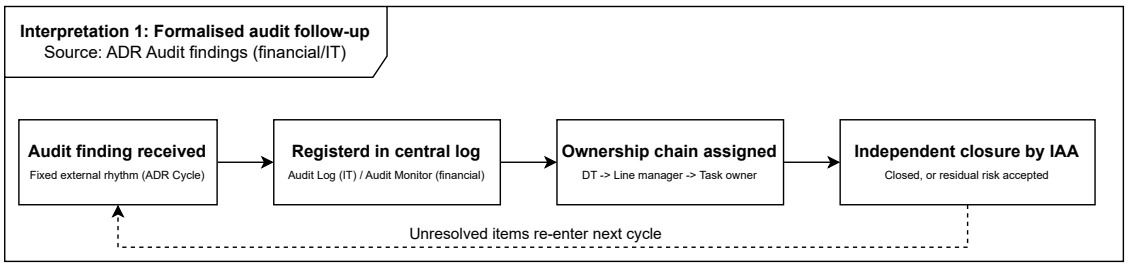


Figure 4.1: Interpretation 1: Formalised audit follow-up.

Figure 4.1 summarises the pattern. Findings arrive through the ADR's fixed external auditing schedule, are registered in one of two central instruments (the Audit Log for IT findings, the Audit Monitor for financial and operational findings), and are then assigned along an ownership chain that begins with the responsible DT member, passes through a line manager, and ends at a task owner. The key feature of this pattern is that closing the task is not decided by the task owner or the line manager: evidence of resolution must be verified independently by Internal Audit and Advisory (IAA) before the finding is formally closed in the log. One account described the current cycle as largely self-sustaining, requiring only annual escalation under normal conditions, while the same process had earlier required bi-weekly reporting to higher governance during a period of elevated backlog in the past before the use of the auditlog and -monitor.

[1] At SSC-ICT, three related governance bodies are referenced throughout this chapter. The *directieteam* (DT) consists of the CEO, COO, CTO, CFO, and the directiesecretaris. The *managementteam* (MT) is a distinct body, consisting of the COO together with the directors of the three business units (BIS, BST, BAS). Together, DT and MT form the *directiemanagementteam* (DMT). Departmental management teams under individual DT members (e.g. the management team under the CFO, referred to as "MT-CFO" in interview material) are distinct from the organisation-wide MT and are specifically defined at each point where they are mentioned.

Two features of this pattern are worth making explicit. First, closure does not require full implementation: where residual risk is explicitly accepted and documented, a finding can be closed with rationale. One participant illustrated this with an example in which, once a particular compliance rate had reached a qualitatively acceptable level, the remaining residual risk was formally accepted rather than pursued further to be improved. Second, the formal process explicitly distinguishes between *acknowledging* a recommendation, *acting on it*, and *deciding not to act on it* as three valid forms of follow-up, provided the decision and its reasoning are recorded. Both features assume the existence of a record where such reasons can be stored, compared, and checked again later.

Multiple participants were explicit that this pattern applies specifically to ADR findings and not to recommendations from other sources. The same people who explained the audit process with confidence also made it clear that strategic and advisory recommendations, even when they come from related sources, do not go through this process. One account summarised the contrast directly: "there is no tight process at the back-end through which we have secured this" ("[er is] niet een heel strak proces aan de achterkant waarmee we dat geborgd hebben").

4.2.2 Programmatic follow-up via taskforce

The second pattern is about external advisory reports that the DT ordered, like the Deloitte report about the transformation programme and the Sopra Steria report about strategic security advice. This pattern was mainly explained by participants who worked in these specific programmes. It was also confirmed independently by people in other roles, including one participant who had seen the earlier KPMG-led process that resulted in the current business-unit structure. It aligns with theme T1 in part, and more strongly with T2 and T4 (monitoring and visibility gaps).

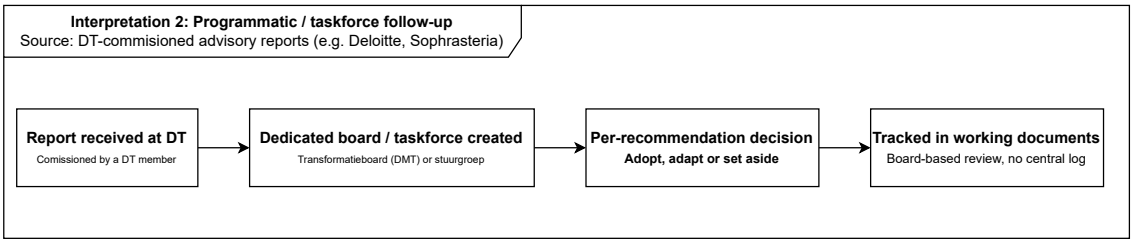


Figure 4.2: Interpretation 2: Programmatic follow-up via taskforce.

Figure 4.2 captures the pattern. After a report has been received at DT level, a dedicated governance body is created or designated to steer follow-up. The scope of this body varies with the scope of the report: the Deloitte trajectory was steered through a transformation board ("transformatieboard") with full DMT participation, while the Sopra Steria trajectory ran through a CTO-chaired security steering group ("stuurgroep") at programme-team level. Within these bodies, recommendations are discussed individually and a decision is taken per recommendation (adopt, adapt, or set aside) with the rationale recorded in the working documents of the programme. For the Deloitte trajectory, an external re-audit by CIO Rijk took place approximately eighteen months after the original report, and its results were fed back into the same transformation board. The Sopra Steria trajectory had no comparable external cycle.

The main weakness of this pattern, as many participants explained, is that the tracking only happens in the working documents of that specific programme and is not recorded or organised together with recommendations from other sources. One account described this directly: "there is not one uniform way in which we handle this" ("er is niet één manier hoe we hiermee omgaan"). The practical consequence is that continuity of follow-up depends on the continued involvement of whoever owns the programme. The same account noted that cultural and organisational issues identified in the earlier KPMG report had resurfaced years later without appearing in any ongoing tracking instrument, adding that there is "not a reminder somewhere indicating this comes from that report and still needs to be followed up" ("niet ergens een tikkertje [...] dit komt uit dat rapport en dit moet je nog opvolgen"). A participant from another function confirmed the same pattern from the business unit side. They explained that the organisational change from that KPMG report was only partly carried out, and that the other recommendations were later more or less changed in meaning or replaced after several leadership changes later.

4.2.3 Strategic follow-up without structural embedding

The third pattern describes recommendations that fit neither of the first two patterns. This includes strategic and policy-driven recommendations from sources such as ACICT, the Algemene Rekenkamer, statutory screenings, and ad-hoc studies commissioned from external consultancy firms, where no dedicated programme is created and where no formalised register applies. Across the interviews, this was the pattern most frequently cited as problematic, and it is the pattern through which the majority of non-audit recommendations enter the organisation. It aligns with themes T3 (responsibility and ownership gaps), T4 (monitoring and visibility gaps), T5 (prioritisation and decision-making), and T7 (organisational capacity and cultural barrier).

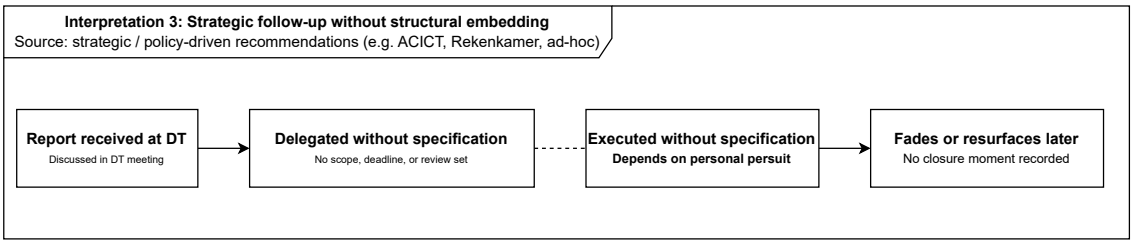


Figure 4.3: Interpretation 3: Strategic follow-up without structural embedding. The dashed segment marks the point at which the link to the originating recommendation tends to be lost.

Figure 4.3 summarises the pattern. A report arrives at the DT, is discussed, and actions are delegated downward. Beyond that point, the pattern becomes difficult to reconstruct, because the interviews describe no consistent follow-up mechanism. Several accounts independently described delegation as an "over the fence" ("over de schutting") handover in which the scope of the recommendation, the expected result, and the review moment are not fully specified at the point of delegation. One account stated the underlying assumption explicitly: "the assumption is made that Pietje Puk knows exactly what to do; that is not the case" ("de aanname wordt gedaan dat Pietje Puk precies weet wat hij moet doen. Dat is niet zo"). The same account provided a concrete illustration of what can happen in the absence of structured follow-up: a governance framework of more than ninety pages that was written several years earlier, formally adopted, and effectively never implemented, characterised as having "vanished into a drawer" ("in de la").

The dashed segment in Figure 4.3 marks the point at which the link between the downstream work and the originating recommendation tends to be lost. This observation is supported by accounts from the execution side of the organisation. One participant whose work begins once plans are already underway noted that they are involved only "at the moment we start working with those plans" ("op het moment dat we met die plannen aan de slag gaan"), with no structural guarantee that the plan still carries an identifiable link back to the report it originated from. One account explained the same kind of distance from the governance side. They said that action-tracking in this situation is only done around the meeting cycles and does not go further to follow up on recommendations across different reports. The result is that recommendations can be only partly followed, completely changed, or just dropped without a clear ending point and without a rationale that people can look up later.

A similar point in this pattern is that progress, when it happens, usually comes from personal effort and not from formal structures. One account stated the mechanism directly: "if you delegate this and then lean back, it simply does not happen" ("als je dat eruit zet en achterover gaat lenen, dan gebeurt het gewoon niet"). The same account characterised the organisation as reactive: staff act on explicit instructions rather than on initiative, means that tasks you give to others, without clear context at the moment you delegate, almost never move forward by themselves. This shows why the pattern is not always unproductive. Some items are solved, but this happens more because of the specific person, and not because of the general structure. In practice, this means the pattern is weak when staff change, when the workload changes, or when people change their own priorities. If the person who was responsible for an item leaves, has less time, or puts focus somewhere else, there is no system in place that keeps this item clearly visible.

A further observation from multiple accounts is that this pattern disproportionately affects the same individuals: follow-up work tends to concentrate on colleagues who already carry the relevant domain knowledge, who are also typically the colleagues already heavily loaded with operational work. One account summarised this as "for this kind of work, the same people are usually asked, and they are typically already very busy" ("vaak in dit soort dingen [worden] dezelfde mensen gevraagd [...] die vaak al heel druk zijn"). The result is that, even when someone is clearly responsible, the work can still stop for reasons that do not relate to the real content of the recommendation, which makes the weakness explained above even bigger.

4.2.4 Chain-dependence

The fourth interpretation differs in kind from the first three. Rather than describing a separate follow-up route, it captures a property that can apply to recommendations arriving through any of the other patterns: some recommendations cannot be resolved within SSC-ICT's direct sphere of influence, because their resolution depends on decisions, resources, or actions held by other organisations in the government IT chain. It aligns primarily with theme T2 and with the cross-organisation dimension of T3.

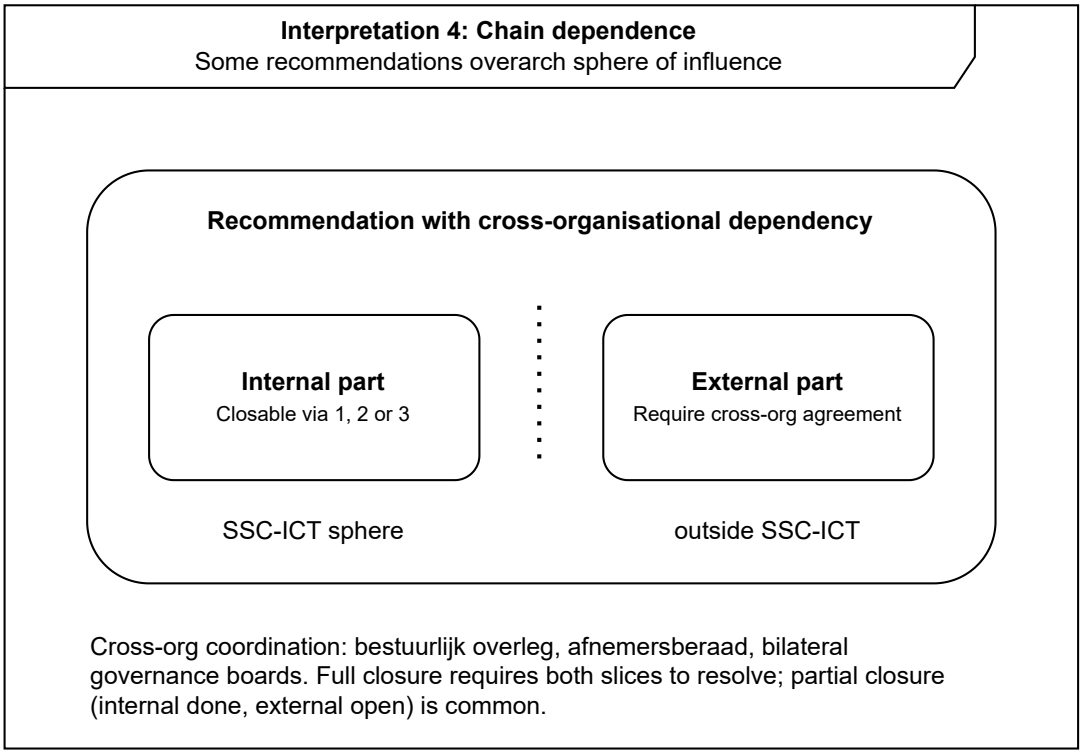


Figure 4.4: Interpretation 4: Chain-dependence as a modifier on Interpretations 1-3. The recommendation is represented as a single object that straddles the SSC-ICT boundary, with each slice carrying its own resolution pathway.

Figure 4.4 represents this property as a structural overlay rather than a flow. A recommendation with a cross-organisational dependency is shown as a single object that spans the SSC-ICT sphere of influence: the internal slice sits within the sphere and follows whichever of Interpretations 1–3 applies, while the external slice sits outside and cannot be closed by SSC-ICT on its own. The split is a property of the recommendation itself, not a step that is introduced after intake, the two slices exist from the moment the recommendation arrives and are resolved through two different pathways.

The external slice is resolved through cross-organisational governance bodies: the bestuurlijk overleg, the afnemersberaad, the financiële commissie dienstverlening, technical overleggen, and bilateral governance boards with partner ministries (gremia). One account summarised the position directly: "it is not within our sphere of influence, it sits with a client or a framework-setter, that is where my responsibility ends" ("het zit niet in onze cirkel van invloed, het zit bij een opdrachtgever of bij een kadersteller [...] daar houdt mijn verantwoordelijkheid wel op"). A second account, from a role with control responsibility, said that there are many of these bodies and that the coordination between them is not always strong enough to make sure that a specific recommendation always ends up in the right group. Concrete examples that came back in several interviews are recommendations about government-wide framework agreements, the two-sided signing of DFAs, and the arrangements for returning devices by clients. For each of these topics SSC-ICT has influence, but cannot make the final decision alone.

The practical consequence of this overlay is that closure on the SSC-ICT side does not imply closure of the recommendation as a whole. A recommendation can be fully addressed within the internal slice while the external slice remains unresolved, sometimes over multi-year timespans. One account described a concrete example in which a finding that had originated several years earlier was only resolved after a former director returned in an interim capacity and could confirm at a governance meeting that the issue had finally been resolved. This illustrates both that such recommendations can persist across leadership cycles and that their eventual closure can depend on specific individuals rather than on structural follow-up mechanisms.

4.3 External perspectives on recommendation follow-up

The process models in Section 4.2 describe the follow-up as it is organised inside SSC-ICT. On purpose, they do not say anything about how the organisations that make ICT-related recommendations see this follow-up from their side of the handover, and they also do not describe how recommendation follow-up is arranged in similar public-sector contexts. This section introduces that outside view. The goal has three parts: first, to compare the results from the receiver side with an independent view on the same governance situation; second, to learn from how these organisations themselves handle follow-up, since at least some of their practices may be relevant as design input, and third, to bring forward producer-side points about recommendation quality, feasibility, and follow-up that are not visible from inside SSC-ICT alone.

Four interviews were conducted for this purpose: with the Algemene Rekenkamer (Rekenkamer), the Adviescollege ICT-toetsing (ACICT), CIO Rijk, and a representative speaking from the Vereniging van Rekenkamers (VvR) and the decentral rekenkamer practice, who is also the author of *De effectketen van rekenkameronderzoek* [20].The four respondents have different roles in institutions for advice, oversight, and policy direction. We see them here as four comparable sources, even if their roles are not completely the same. Each contributes a different angle: the Rekenkamer as an independent "high council of state" ("hoog college van staat") with its own follow-up system; ACICT as a statutory advisory college that focusses on the ex ante assessment of IT projects; CIO Rijk as the central policy directorate of the CIO-system, which itself does not produce advice but sets frameworks ("kaders") through which much of the advisory output is expected to be absorbed; and Lemmens from the VvR as a combined practitioner–researcher on the effect chain concept. A overview of the intervies held are found in Table 4.3.

Institution	**Date**
De Algemene Rekenkamer	14-04-2026
Adviescollege ICT-Toetsing	17-04-2026
CIO-Rijk	22-04-2026
Lemmens (VvR)	23-04-2026

Table 4.3: Overview of producer-side participants

All four were interviewed using the same producer-side semi-structured protocol, as described in Section 3.4.2 (Appendix B), adapted for Lemmens into a more concept-oriented variant that uses his published effect chain framework as the leading structure. Following Miles, Huberman, and Saldaña, the data are condensed into a role-ordered matrix, presented in Table 4.4. The matrix is organised around five categories, which together reflect the structure of the interview protocol and the analytical purpose of this section. The first category, *institutional position and mandate*, describes what each organisation is,

on what legal basis it operates, and how its mandate relates to that of the other three. The second, *own follow-up practice*, describes what each organisation itself does, or does not do, after a recommendation has been issued. The third, *observed patterns in receiving organisations*, captures what each respondent sees happening at the receiving side; this category is used to triangulate the patterns identified in the receiver-side analysis in Section 4.1. The fourth, *design prescriptions for a follow-up model*, records what each respondent considers necessary for a follow-up model to work. The fifth, *warnings and caveats*, records what each respondent considers dangerous or non-transferable.

Table 4.4: Role-ordered matrix of producer-side perspectives on recommendation follow-up. Each cell condenses the respondent's position on the category indicated by the preceding section header; direct quotes are retained in the original Dutch with an English translation in parentheses.

Respondent	Condensed perspective
Category 1: Institutional position and mandate	
Rekenkamer	Independent high council of state, anchored in the Constitution and the Government Accounts Act (*Comptabiliteitswet*); reports to the Tweede Kamer and the minister. Conducts research on own initiative, determined by its own three-member college. Three research types are distinguished: the annual accountability research (*verantwoordingsonderzoek*), policy research (*beleidsonderzoek*) published throughout the year, and focus research without recommendations for current events. Findings from accountability research that touch the statutory task, termed imperfections ("onvolkomenheden"), are binding, and the minister is obliged to address them; recommendations from policy research and performance audits are non-binding. The respondent was explicit about terminology: "they are not pieces of advice. I would really not use that word. Recommendations" ("het zijn geen adviezen. Dat woord zou ik echt niet gebruiken. Aanbevelingen").
ACICT	Independent advisory college established by law (*Instellingswet Adviescollege ICT-toetsing*) in 2019, succeeding the earlier Bureau ICT-toetsing (BIT). Its creation was itself one of the 34 recommendations in the parliamentary Elias report. Mandate covers three types of work: IT projects with an IT component above 5M EUR (mandatory notification), maintenance of important systems and portfolios, and translation of new policy or legislation into systems. Advice is given before the project starts, on the rationale that adjustment is easiest when no irreversible steps have yet been taken. Produces 15–20 reviews per year. Positioned as forward-looking and strategic, in explicit contrast to the Auditdienst Rijk (backward-looking, compliance) and the Rekenkamer (system of checks and balances). The college itself uses a principle-based assessment framework rather than a checklist.
CIO Rijk	Part of the Ministry of the Interior and Kingdom Relations (BZK), within the DG Digitalisation and Government Organisation, directorate CIO Rijk and Digitalisation Policy. The respondent framed the directorate explicitly as a policy directorate ("beleidsdirectie"), not as an advisory body: "we are not an advisory body" ("we zijn geen adviescirkel"), a distinction used to separate CIO Rijk from bodies such as ACICT or the Rekenkamer. The formal authorities are delegated from the minister or staatssecretaris through two legal instruments, the Coordination Decree on Organisation, Operations and Information Systems of the National Government ("Coördinatiebesluit organisatie, bedrijfsvoering en informatiesysteem rijksdienst") and the CIO-System Decree ("Besluit CIO-stelsel"). Produces binding policy frameworks ("kaders"), but the respondent noted that the choice of instrument depends on the goal: lighter instruments such as a guidance ("handreiking"), a guideline ("richtlijn"), or a instruction ("werkinstructie") are sometimes used when binding rules are not needed. The CIO-system was itself revised in 2025 and took effect on 1 January 2026, adding roles for data, privacy, and technology alongside the existing CISO role. The model is explicitly layered ("gelaagd"): CIO Rijk communicates through the departmental CIOs, who in turn work toward their own execution organisations; CIO Rijk itself does not verify individual execution organisations directly.
Lemmens (VvR)	Speaks from a combination of positions rather than a single institution: chair and member of three decentral rekenkamers in the south of the Netherlands, board secretary ("bestuurssecretaris") of the Vereniging van Rekenkamers (VvR), and independent researcher with a socio-economic history background. He has been working with rekenkamers for approximately twenty years, since the introduction of mandatory decentral rekenkamers in 2005–2006 under the municipal dualisation reform. The VvR is not itself a producer of recommendations but a supporting organisation for roughly 250 decentral rekenkamers, which together cover around 430 municipalities, and is funded through a subsidy from BZK. Its main functions are quality support for small and medium-sized rekenkamers, knowledge pooling through meta-dossiers and meta-analyses, and interest representation ("belangenvertegenwoordiging") toward the Rijk. Decentral rekenkamers are positioned differently from the Algemene Rekenkamer: they do not cover financial compliance, which is left to the accountant, but investigate effectiveness, efficiency, and compliance of the "policy conducted" ("gevoerd beleid"). Lemmens is also the author of the effect chain article used as the basis for this interview.

Continued on next page

Table 4.4 continued from previous page

Respondent	Condensed perspective
Category 2: Own follow-up practice	
Rekenkamer	Maintains two distinct follow-up regimes in parallel, matched to the legal status of the recommendations. Binding findings from accountability research are monitored until resolved: the ministry must produce an improvement plan, and the Rekenkamer returns annually. Non-binding recommendations from policy research are monitored by a dedicated team using an Excel-based register extended with a custom query tool, recording per recommendation the source, the minister's commitment, the gradation of current progress, and the full timeline of what has happened over time. The arrangement is structurally reinforced by fixed contact persons (*AR-contactpersonen*) inside each ministry, whose explicit task the respondent described as "it is their job to make sure people get to work on it" ("het is hun baan om te zorgen dat mensen daarmee aan de gang gaan"). Large implementation organisations such as the Belastingdienst and UWV maintain their own counterpart contact person. No prioritisation inside the register yet, but the respondent considered introducing it.
ACICT	No formal role in checking whether advice is followed up. After an advisory report is issued and the responsible minister has produced an administrative reaction, the college's formal involvement ends. What does take place: evaluation conversations with the commissioning organisation, framed as process learning for ACICT itself; informal aftercare, including further clarification conversations when requested; and periodic conversations between ACICT's internal account managers and the ministerial CIO to maintain awareness of what is happening at each ministry. Follow-up itself is expected to run through the CIO system (*CIO-stelsel*), where the ministerial CIO is formally responsible. One procedural safeguard exists in the college's own law: if a minister deviates from an ACICT advice, the deviation must first be discussed in the Council of Ministers (*ministerraad*) before the administrative reaction is sent to the Second Chamber. In reality, however, the answers from the administration are often written so general that they hide the non-compliance and do not start this procedure.
CIO Rijk	No fully implemented central follow-up system for own kaders yet. The respondent acknowledged this directly: "we should be doing that, but we have not set it up 100% yet. So it has happened that we have issued kaders without monitoring on them" ("dat zouden we moeten doen, maar we hebben dat nog niet 100% ingericht. Dus het is ook wel gebeurd dat wij kaders hebben gesteld zonder dat daar monitoring op zit"). The current direction is that every new kader should be accompanied by a monitoring and evaluation provision. Performance data is published annually on the Rijks ICT-dashboard on accountability day ("gehaktdag"). It does not include indicators such as the percentage of cloud contracts with a risk analysis or the number of published algorithms, and the respondent was careful not to describe these as real KPIs. A maturity model is in development: building on existing models for information security and information management, with the intention to expand coverage so that organisations can be scored over time on the firmness ("stevigheid") of their governance arrangements. Decision-making on new kaders runs through inter-departmental bodies (CISO-raad, CIO-beraad), where departmental CIOs and CISOs contribute before adoption. Multi-year effect monitoring for each advice is not considered necessary; the respondent preferred to verify whether a specific implemented measure achieved its intended result, but did not see value in long-running monitoring of all recommendations structurally.
Lemmens (VvR)	The VvR developed a format that colleges (local executives) can use to report on recommendation follow-up under the 2023 Gemeentewet article 185a obligation. Lemmens explicitly noted that the article covers only follow-up, not implementation or impact, and that the VvR format deliberately broadens the reporting scope: "they [the legislator] have said nothing about effect" ("ze hebben dus niks gezegd over doorwerking"). Lemmens also co-authored the VvR checklist for formulating a good recommendation ("die heb ik mede opgesteld"). The VvR itself does not monitor individual recommendations; follow-up is left to each decentral rekenkamer and, where relevant, to the griffie (municipal council clerk) of each municipality. When griffies already track motions, amendments, and commitments, they could in principle extend this to rekenkamer-recommendations, but practice varies. Lemmens himself advocates annual follow-up conversations between the rekenkamer and the council clerk or relevant civil servants, with the aim of clarifying intent, contextualising recommendations after publication, and discovering misinterpretations early.

Continued on next page

Table 4.4 continued from previous page

Respondent	Condensed perspective
Category 3: Observed patterns in receiving organisations	
Rekenkamer	Ownership named as the central bottleneck: "if it lands somewhere in nothing, you can forget about it" ("als het een beetje ergens in een niets landt, dan kan je het wel vergeten"). High personnel turnover extends the problem: in the respondent's estimate, a single dossier can pass through four owners in two years, breaking institutional memory. Cabinet changes introduce further discontinuities, since an incoming minister may hold different views. Political sensitivity rises closer to the minister; implementation organisations respond more pragmatically than central departments. A specific structural pattern surfaced for ICT: the Ministry of BZK holds a coordinating role for IT and operational management across ministries, but this coordinating role is often filled insufficiently to produce change in other ministries' territory, since "ministries are in charge of their own matters" ("ministeries gaan over een eigen ding") and work in silos. Short-term capacity pressure tends to crowd out long-term concerns. A positive observation: substantive impact often occurs during the research itself rather than after the report, with departments adjusting in response to the line of questioning before the final report appears.
ACICT	Confirmed, in the Confirmatory section of the interview, every receiver-side obstacle presented to them. The distinction between systematic follow-up of audit findings and largely absent follow-up of strategic advice was described as a government-wide pattern, not an SSC-ICT-specific characteristic: "where you see in the private sector that audit or advisory reports have a central register, checked periodically at executive-board level, you see here that this follow-up sometimes sits at too low a level" ("waar je in het bedrijfsleven ziet dat audit rapporten of advies rapporten een centraal register hebben, periodiek op raad van bestuursniveau gecheckt, zie je hier eigenlijk dat die opvolging soms op een te laag niveau hangt"). Fragmented registration and unclear ownership were recognised with the observation that "government is also good at fragmenting responsibilities" ("de overheid is daar ook goed in om soms verantwoordelijkheden te versnipperen"). External pressure was confirmed as the main driving factor, with the note that recurrence by itself is not enough "you can make the same findings for years, but if there are no consequences attached [...]" ("je kunt jarenlang dezelfde bevindingen doen, maar als er geen consequenties aan verbonden zijn [...]"). Compliance-style advice tends to enter the organisation at a lower level and is easier to follow up; strategic advice requires more thinking time and sometimes political or policy-level decisions, which makes it inherently harder to follow up. The capacity limits at the top become worse because there is not enough IT knowledge at the executive level.
CIO Rijk	Confirmed most receiver-side obstacles, often framing them through the lens of organisational maturity rather than individual behaviour. The distinction between systematic follow-up of audit findings and under-developed follow-up of strategic advice was recognised as a maturity gap: the financial audit chain has a controlsystematiek (control system) that has existed for decades, while the IT governance chain is "still a young domain" ("nog een jong terrein"). The "over the fence" ("over de schutting") pattern from the receiver-side interviews was recognised, with the explanation that Tweede Kamer moties are often not preceded by an executability check ("uitvoeringstoets"), and that even when the uitvoeringstoets is formally part of the standard procedure it is government-wide often not done or done too little: "sometimes there is simply far too much haste" ("er is soms gewoon veel te veel haast"). When asked which single factor most obstructs follow-up, the respondent named capacity first, and ownership second, and at a deeper level named culture as the underlying cause: "if you really want to know why advice is seldom followed up, I think it is really a culture thing" ("als je echt zou willen weten van waarom worden adviezen weinig opgevolgd, dan denk ik dat dat echt een cultuur ding is"). External pressure was confirmed as an effective driver: ACICT reports are "more feared" ("meer gevreesd") than internal CIO-oordelen because the former go public and to the Chamber, while the latter remain internal. One observation specifically on SSC-ICT: the respondent was surprised that an IT delivery organisation does not have its own dedicated CIO, and noted that the CIO-stelsel explicitly allows a departement to appoint one for organisations with heavy IT dependence.
Lemmens (VvR)	Recognised the audit-versus-strategic-advice distinction from municipal practice: accountants also perform a light IT audit at municipalities, whose results tend to be picked up consistently, while strategic advice is "somewhat more fluid" ("wat vloeibaarder") and more likely to remain open over time. Confirmed that "public visibility" ("publieke zichtbaarheid") actively supports follow-up; without it, reports can end "at the bottom of the drawer" ("onder in de la"). Independently reported a pattern also observed by the Rekenkamer: during a rekenkamer investigation, departments often already pick up the concrete findings, so that the recommendations which survive into the final report tend to be the more abstract ones, over which the receiver can more easily "roll over" ("overheen walsen"). A second pattern relates to the 39% figure in his article: even where recommendations are formally rejected by the council for political reasons, public servants may still implement them because they find them useful. For the SSC-ICT context specifically, the absence of a Second Chamber equivalent to trigger external political pressure was named as the core ownership question: without an external political anchor, ownership must be taken explicitly inside the organisation, at the directieteam (DT) level. Defensive routines in organisations ("defensive routines") suppress unwelcome information; this dynamic is the same mechanism that Lemmens will discuss in a forthcoming "Policy research" ("Beleidsonderzoek") Online article, with a parallel to the Dutch "Child benefit scandal" ("toeslagenaffaire").

Continued on next page

Table 4.4 continued from previous page

Respondent	Condensed perspective
Category 4: Design prescriptions for a follow-up model	
Rekenkamer	A central register is the minimum, with the following fields drawn from the respondent's own operational experience: the source of the recommendation, the recommendation itself, the owner, the lead time, internal monitoring, and a link to the plan under which the recommendation is being carried out. A dashboard at administrative level is recommended, presented at audit-committee meetings so follow-up becomes a recurring agenda item. Ownership must be explicitly assigned and actively maintained; administrative dialogue between the oversight body and the receiving organisation at senior level functions as an activator ("collegeleden in gesprek met de minister"). Success factors for follow-up, drawn from a best-practices report the Rekenkamer contributed to: quality of the recommendation, the relationship with the party being audited, media strategy and political sensitivity, the mere existence of a follow-up mechanism, and public transparency; "shame is a negative word, but making it public also helps" ("shame is een negatief woord, maar dat ook openbaar maken helpt"). Transparency is important because it is the mechanism through which ownership can be held accountable inside the organisation.
ACICT	A three-layer structure ("stappengelaagdheid") is necessary: (1) a central register of all recommendations, (2) monitoring of actual follow-up, and (3) measurement of the effect of the follow-up. The respondent was emphatic that the third layer is where most public-sector follow-up in fact stops, and that this is the layer that matters, since "following up on recommendations is not a goal in itself [. . .] what you actually want is to achieve a particular effect" ("het opvolgen van aanbevelingen is niet een doel op zich [. . .] je wil juist een bepaald effect bereiken"). Priority must come explicitly from the top: "as long as an executive board does not prioritise this, so does not show from the top that it is important [. . .] then nothing comes of it" ("zolang een directie hier geen prioriteit aan heeft, dus te tonen dat het vanuit de top belangrijk is [. . .] dan kun je er ook niks aan hebben"). Recommendations should be linked to "de why", the intended effect and the underlying goal, not only to SMART-formulated instrumental action items, so that people executing the work understand why they are doing what they are doing.
CIO Rijk	Core prescription was to design minimally and to embed the follow-up mechanism in existing structures, rather than building a parallel system. Two principles recurred in the interview: the organisation should design its own system; "just design a system yourself that works for your organisation" ("gewoon zelf een systeem voor bedenken wat werkt voor jouw organisatie"), and this system should be linked to an existing process: "the best option is to connect it to planning and control" ("het beste is om het te laten aansluiten bij planning en controle"). On the register itself, the respondent was explicit that list-keeping quickly turns into administrative burden ("lijstjes bijhouden is alleen maar heel veel administratie"), so the register must stay "lean en mean" and contain only the information that the people who actually steer need, decided together with those people. The respondent further pointed to the three lines model (three-lines-of-defence / governance) as a useful frame for an organisation like SSC-ICT, and noted that a maturity-model approach, rather than tracking each individual recommendation, could be a meaningful instrument to score organisational "firmness" on follow-up discipline over time. The value of a dedicated CIO function at SSC-ICT was acknowledged, but qualified: effectiveness depends on how the role is filled, on leadership, and on organisational position, not on the title alone.
Lemmens (VvR)	The minimum set of fields earlier proposed by the Algemene Rekenkamer (source, recommendation, owner, lead time, internal monitoring, link to implementation plan) was confirmed as a sufficient base: "this really is the minimum set, you can dress it up further but I think this is the basis" ("dat is echt wel de minimumset, je kunt het allemaal aankleden maar ik denk dat dit wel de basis is"). Three design principles were emphasised. First, any register must be accompanied by interne transparantie (internal transparency), so that open items are visible across the organisation, and by an interne controller function that actively signals which items have not yet been followed up, effectualised, or produced impact. Second, recommendations must themselves be "concrete" ("concreet"), "executable" ("uitvoerbaar"), and addressed ("geadresseerd") at a named role or function: "if you do not address it, then nobody can feel ownership, and it stays hanging somewhere" ("als je het niet adresseert, dan is er niemand die zich daar eigenaar van kan voelen en dan blijft het ergens hangen"). The receiving organisation itself, rather than the producer, should translate each recommendation into measurable goals; this is where effectualisation and impact become measurable in practice. Third, aftercare requires both producer engagement (annual conversations between rekenkamer and council secretary, or equivalent) and receiver-side processes (the counsil clerk, or in the SSC-ICT context the secretariat). For SSC-ICT specifically, the absence of a Second Chamber equivalent (in the decentral rekenkamer contextual level) means that ownership has to be embraced internally by the DT: "I think ownership should sit there, and someone should then also embrace it" ("ik denk dat daar het eigenaarschap zou moeten zitten en dat iemand dat dan ook omarmt"). The register itself must also allow organisations to add "explanation" ("toelichting") per item, including for recommendations that have been overtaken by time, so that the register stays "alive" ("levend") rather than becoming a tick-list.

Continued on next page

Table 4.4 continued from previous page

Respondent	Condensed perspective
Category 5: Warnings and caveats	
Rekenkamer	Formulate recommendations at a high enough level of abstraction that they remain valid over time, which is the respondent's own guard against temporal decay. Do not formulate so prescriptively that the organisation's own judgement is replaced, characterised as "sitting in the minister's chair" ("op de stoel van de minister gaan zitten"); this also preserves the Rekenkamer's ability to evaluate the result objectively later. Do not expect short timelines: "you also need the courage and understanding, as a Rekenkamer, to accept that it is not sorted out in three months" ("je moet ook de ballen en het begrip als rekenkamer hebben dat dat dan niet in 3 maanden voor de bakker is"). A register or follow-up process is work, not a one-time design: "it does not happen automatically [. . .] it is not that the gnomes came by at night and solved it for you" ("het gaat niet vanzelf [. . .] het is niet dat de kabouters 's nachts zijn langsgekomen en die hebben dat dan voor je opgelost"). Design the process together with people in the organisation, not alone in isolation; otherwise, they will not follow it. Finally, the respondent was clear that the Rekenkamer's arrangement is partly non-transferable: "we really are pampered creatures" ("wij zijn wel echt luxebeestjes"), because the Rekenkamer is taken seriously in ways that smaller advisory bodies are not.
ACICT	The register must not become "a bookkeeping exercise" ("een boekhoudkundige exercitie"), where registration becomes a goal in itself and the reason for maintaining it is forgotten. The respondent drew an explicit parallel with risk management, where registers are often maintained out of habit without a clear purpose, turning into an administrative process. Recommendations should be formulated to be SMART, but must leave the receiving organisation room to act: "there are multiple roads to Rome" ("er zijn meerdere wegen naar Rome"). If an organisation proposes a different measure that addresses the same risk, that is acceptable and even desirable; the advisory body should not take over the organisation's responsibility to think for itself. Organisations that stop thinking and just execute what an advisory body prescribes are a failure mode, not a success.
CIO Rijk	The central warning was that the register must not turn into "a chore in the reporting cycle" ("een corvée klusje met de rapportagecirkel"), where the act of filling it becomes a bureaucratic task detached from steering value. The respondent was explicit about minimising reporting load ("as few reporting burdens as possible, as few administrative burdens as possible"; "zo min mogelijk rapportage lasten, zo min mogelijk administratie lasten") and warned that if the register becomes too large, people are made "driven up the wall" ("gillend gek") by it.A second, more role-oriented warning: the respondent warned not to depend only on formal hard powers ("harde bevoegdheden") when you set up or make stronger a CIO function, because the real effectiveness depends at least as much on leadership, the place in the organisation, and on cooperation. The long-term monitoring of effects for each advice was also approached carefully: checking at specific moments if one measure reached its planned outcome was seen as useful, but ongoing and systematic monitoring of all recommendations was considered too much.
Lemmens (VvR)	Receiving organisations often have a structural incentive toward vague recommendations, because "then you cannot be held accountable if you fail to achieve it" ("dan kun je er niet op afgerekend worden als je niet haalt"). A follow-up design that does not require concreteness therefore will not fix the underlying behaviour. Defensive routines inside organisations suppress unwelcome information and preserve the status quo, a dynamic that Lemmens will discuss in aforementioned fore-coming research. Registers with only binary status markers quickly collapse into "tick-lists" ("vinklijstjes"); the design must force organisations to supply an explaination per item, including for recommendations that have become time-obsolete, so that the register continues to carry meaningful information. The absence of public visibility inside SSC-ICT removes a natural external check, which means that the internal controller function has to do more work than in a municipal setting. Finally, a learning culture ("een lerende cultuur") is a required condition for any of this to function; without the willingness to hear "contradicting views" ("tegenspraak"), the register and the controller together will still not produce learning.

4.3.1 Synthesis of producer-side observations

Before moving to the design requirements, we first look at the four columns of Table 4.4 together with the receiver-side patterns from Section 4.1 and the four process interpretations from Section 4.2. Because each of the four producer-side respondents has a different institutional position, the goal is not to count how often something is repeated, but to see where the external view confirms the internal picture, where it shows something that is not visible from inside SSC-ICT, and where it gives a warning for patterns that cannot be transferred.

Points where the producer-side view confirms the receiver-side picture

The asymmetry between audit follow-up and strategic advisory follow-up was recognised by all four respondents, and they framed it explicitly as a government-wide pattern, not as something that is specific to SSC-ICT. ACICT used almost the same words as several receiver-side accounts: a central register tends to exist for audit findings, but not for strategic advice. Lemmens recognised the same split at the municipal level, where the light IT-audits of the accountant are picked up relatively consistently, while strategic advice is "somewhat more fluid" ("wat vloeibaarder"). CIO Rijk framed the same gap as a difference in maturity: the financial-audit chain has had decades to build up its control systematics, while the IT-governance chain is still a young domain ("nog een jong terrein"). The dispersion of ownership, the "over the fence" ("over de schutting") pattern, and the concentration of follow-up work on colleagues who are already loaded, were all confirmed. External pressure was confirmed by all four respondents as a mechanism that supports follow-up, and the Rekenkamer, ACICT, and Lemmens all point to public visibility as a part of that mechanism.

Points that the producer-side view adds

Three observations only become visible when the external positions are included. First, the three-layer structure that ACICT described, namely (1) registration, (2) monitoring of the actual follow-up, and (3) measurement of the effect, is a cleaner way of ordering what the receiver-side interviews described in more scattered terms. Lemmens's effect chain covers the same distinction with the words "follow-up" ("opvolging"), "implementation" ("doorwerking"), and "impact", and he warns explicitly that registers that stop at layer 1 fall back into "tick-lists" ("vinklijstjes"). This is not present in the receiver-side material, but it fits the empirical patterns that were described there, which makes it usable as a shared framing for the rest of the thesis (see Section 2.7 in Chapter 2). Second, the Rekenkamer and Lemmens independently reached consensus on a minimum set of register fields: source, recommendation, owner, lead time, internal monitoring, and a link to the plan under which the recommendation is carried out. The consensus is useful, because the two positions are independent: the Rekenkamer speaks from the operational practice of its own register, and Lemmens co-authored the VvR checklist for good recommendations. Third, the requirement that recommendations themselves should be "concrete, executable, and addressed" ("concreet, uitvoerbaar en geadresseerd") was raised by Lemmens as a condition on the producer side, not only on the receiver side. This broadens the follow-up problem upstream of SSC-ICT: when a recommendation arrives without a named role or function, no internal register can repair that afterwards.

Points that warn against direct transfer

Not everything on the producer side can be transferred to the SSC-ICT context. The follow-up arrangement of the Rekenkamer depends on the fixed AR-contactpersoon inside each ministry, whose specific task is to keep recommendations moving, and on the administrative dialogue between college members and the minister. The respondent self said that the Rekenkamer sits in a relatively privileged position; "we really are pampered creatures" ("wij zijn wel echt luxebeestjes"), and Lemmens made the same point from the other side: for SSC-ICT there is no Tweede Kamer (Second Chamber) equivalent that produces external political pressure in its direct sphere of influence (see the notes on transferability in section 2.7), and ACICT does not come back after the advice is issued. This has a concrete consequence for the artefact design: the monitoring function that public visibility and a returning producer provide in the Rekenkamer case, has to be built internally at SSC-ICT, through an internal controller function (Lemmens) and through explicit ownership at DT level (Lemmens, ACICT). CIO-Rijk adds a warning in the other direction: systematic long-term effect monitoring of every recommendation is seen as too much, and any register must stay "lean en mean" and must connect to existing planning-and-control cycles, instead of becoming a parallel track.

Implications for the requirements

When the four columns are read against each other, a coherent external picture comes out. The receiver-side analysis showed what is missing in the current practice; the producer-side analysis shows what a follow-up arrangement needs at minimum to function. Four ingredients come back across the four respondents: a central register with a minimal set of fields, explicit ownership at a level that is high enough to carry accountability, an internal monitoring function that compensates for the absence of external return, and a framing of follow-up that distinguishes between follow-up, effect, impact, instead of treating these three as one category. These ingredients, together with the receiver-side themes from Section 4.1 and the process patterns from Section 4.2, form the input for the design requirements synthesized in Section 4.4.3 to 4.4.5.

4.4 Synthesis: design requirements for a follow-up and governance method

The literature review (Chapter 2), the receiver-side themes (Section 4.1), the four process interpretations (Section 4.2), and the producer-side matrix (Section 4.3) together indicate that ICT-related recommendations should not be treated as single tasks that can be governed only through administrative tracking. Recommendations are more longitudinal governance objects that cross time, decision area and organisational boundaries,. The core challenge is not only implementing the actual recommendations, but also if the organisation can keep them visible and interpretable over time, prioritise and steer them under shifting constraints, and close, revise or retire them with an explicit rationale when the contexts change that initially provided the recommendations.

This section translates the empirical and literature findings into a set of design requirements that the artefact must satisfy. The requirements make the design choices traceable as every requirement is connected back to one or more more observations or literature sources and they create the basis for the evaluation in Chapter 7, where the artefact is validated against these requirements together with perceived usability and acceptance.

4.4.1 Design problem and contribution argument

The requirements below sit inside a design problem that follows the template of Wieringa [37, p. 16]:

> Improve the follow-up of recommendations from external advisory and audit reports at SSC-ICT, by designing a structured follow-up and governance method that satisfies the requirements PC1-PC5, FR1-FR13, and NFR1-NFR7, in order to help the management, the audit committee, and the executing teams achieve more complete, traceable, and learning-oriented follow-up.

Each requirement is justified by a contribution argument. In the requirements below, the *Rationale* field carries the contribution argument; it states which stakeholder goal the requirement supports and under which assumptions about the SSC-ICT the connection is valid. The *Sources* field records the empirical and literature evidence that supports the contribution argument and the *Fit criterion* field states the observable condition under which the requirement is seen as satisfied in the validation phase.

4.4.2 Approach to requirements synthesis

The synthesis follows the Volere requirements process described by Robertson and Robertson. Volere is chosen for two reasons. First, it provides an explicit *atomic requirement* structure (the Volere "snow card") that captures, per requirement, not only a description but also a rationale, a source, a fit criterion, and a priority. This makes the link between empirical observation and design choice traceable, which is required for a design science research output. Second, Volere distinguishes between *project drivers*, *project constraints*, *functional requirements*, and *non-functional requirements*, which fits the situation here: several of the empirical findings are not really things that the method must *do* (functional), but conditions under which it must *operate* (constraints) or how well it must do its work (non-functional)

The Volere shell is followed in adapted form, because the artefact in this thesis is a governance method and not a software product. Three adaptations are made:

1. Robertson and Robertson uses "originator" field to identify the stakeholder who raised the requirement [30]. In this case, the requirement is grounded from evidence; interview themes, producer-side respondents and the literature review. Therefore the source field records evidence instead of a single named originator.
2. For a software requirement, the fit criterion is a quantified test that the system can be measured against [30]. In a governance setting, for an artefact testing is not always possible. So the fit criterion is therefore stated as an observable condition, that the artefact, when applied to a recommendation case, must satisfy.
3. The Quality Gateway is a checkpoint that filters out requirements that are out of scope, not testable, or otherwise unfit [30]. In this case, the Gateway is applied during the synthesis'; a requirement is only kept if it is supported by one of the three evidence sources, has a description that is independent of the others (atomicity), and has a fit criterion that can be evaluated when the artefact is applied on a case.

The synthesis results in a set of project constraints (PC), functional requirements (FR), and non-functional requirements (NFR), each expressed in a compact Volere-style table.

Each requirement is presented in the same template throughout this section, given by Table 4.5.

Table 4.5: Adapted Volere requirement template used throughout this section. The structure follows Robertson and Robertson, with the originator field replaced by a source field that records the three evidence streams used in this thesis.

Attribute	**Content**
Identifier	`PCx`, `FRx`, or `NFRx`
Type	Project constraint, functional, or non-functional
Description	What the artefact must satisfy.
Rationale	Why this requirement exists.
Sources	Interview themes (T1–T10) / producer-side respondents / literature references.
Fit criterion	Observable condition under which the requirement is satisfied during evaluation.
Priority	MoSCoW: *Must*, *Should*, *Could*, *Won't*.

The MoSCoW priority is used as follows: "Must" marks requirements without which the artefact cannot fulfil its core purpose; "Should" marks requirements that are strongly supported but where partial satisfaction is acceptable; "Could" marks useful extensions that are not essential. No requirement is marked "Won't", since this thesis covers a single artefact iteration.

4.4.3 Project drivers and constraints

The project drivers (the purpose of the artefact, the stakeholders, and the goals) are described in Chapter 1 and not repeated here. The project constraints, however, are not always made explicit in earlier chapters and are therefore listed below. They restrict the design space inside which the functional and non-functional requirements operate.

PC1 - Dutch central-government environment
The artefact is designed for SSC-ICT, which sits in the Dutch central-government context. The artefact must be compatible with this context: with the CIO-system as revised in 2025, with the role of CIO Rijk as a policy directorate that issues frameworks, and with the broader political arena in which external visibility is mediated through the ministerial line rather than through one direct political body [33, 35]. The artefact must therefore not introduce roles, decision moments, or escalation paths that conflict with this environment.

Identifier	`PC1`
Type	Project constraint
Description	The artefact must operate within the Dutch central-government CIO-system, including the revised CIO-system decree in force from 1 January 2026.
Rationale	SSC-ICT is part of this system; recommendations flow through it; the artefact cannot bypass it.
Sources	Lit [33, 35]; producer-side (CIO Rijk); T2, T10.
Fit criterion	Roles and decision moments referenced in the artefact match those defined in the CIO-system; no role is introduced that conflicts with an existing one.
Priority	Must

PC2 - No parallel tracking system
CIO Rijk was explicit in the producer-side interview that an effective follow-up arrangement should connect to existing planning-and-control cycles rather than form a parallel track, and that adding a new system quickly turns into “een corvée klusje” (a chore). Several receiver-side accounts (T8) repeated this point in different words. The artefact must therefore not create a separate tracking track that runs next to the audit log, the auditmonitor and other used tooling.

Identifier	`PC2`
Type	Project constraint
Description	The artefact must connect to existing planning-and-control cycles and tools, not replace or duplicate them.
Rationale	A parallel track adds administrative burden without adding steering value; existing tools already carry the operational work.
Sources	T8; producer-side (CIO Rijk).
Fit criterion	For each tool currently in use (audit log, auditmonitor, TopDesk, Jira, FCC, Qlik Sense), the artefact specifies how it links to that tool rather than how it replaces it.
Priority	Must

PC3 - Compatible with organisational change
At the time of the interviews, SSC-ICT is in a transition to a SAFe-style way of working (T9). The organisational environment is therefore shifting during the period in which the artefact is intended to function. The artefact must not be tied to the situation as it was before the Agile transition, and should therefore also not on routines that are being phased out.

Identifier	PC3
Type	Project constraint
Description	The artefact must remain functional during and after the SAFe/agile transition at SSC-ICT, without being dependent on either the pre-transition or the post-transition operating model.
Rationale	A design tied to the old structure would lose relevance during the transition; a design tied too tightly to the new structure would over-extend assumptions that are not yet stable in the organisation.
Sources	T9.
Fit criterion	The artefact does not break, and does not require redesign, when the way of working changes around it.
Priority	Should

PC4 - Confidentiality classification
One receiver-side account discussed that reports and findings carry different sensitivity classifications, and that visibility across organisational boundaries is a question that the organisation has to settle case by case (T10).

Identifier	PC4
Type	Project constraint
Description	The artefact must support access control aligned with the confidentiality classification of the underlying report.
Rationale	Reports differ in sensitivity; not every reader should see every recommendation.
Sources	T10 (one interview).
Fit criterion	A recommendation can be marked as restricted to a defined role or unit; visibility test passes for that restriction when applied on a case.
Priority	Should

PC5 - Effect chain as conceptual frame
The artefact uses the effect chain of follow-up, implementation, and impact as its conceptual frame for the stages a recommendation can travel through after intake [20]. The frame is adopted because it cleanly separates three things that the empirical material treated as one: the formal commitment, the operational change, and the eventual effect. The frame is applied with the contextual limits in Section 2.7 in mind.

Identifier	PC5
Type	Project constraint
Description	The artefact must use the three-stage effect chain (follow-up, implementation, impact) as its conceptual frame for tracking the state of a recommendation over time.
Rationale	A single "status" field cannot represent these three stages at once; without an explicit frame, the artefact would reproduce the conflation that the empirical material already shows.
Sources	Lit [20]; producer-side (Lemmens).
Fit criterion	The artefact's status model and reporting explicitly reflect the three stages.
Priority	Must

4.4.4 Functional requirements

The functional requirements explain what the artefact must actually *do*. We base them on the empirical patterns in Section 4.1, on the four process interpretations in Section 4.2, and on the producer-side perspectives in Section 4.3. They are also connected to the literature reviewed in Chapter 2.

FR1 - Central recommendation register with history
Multiple sources show that organisations cannot reconstruct the implementation trajectories when follow-up information is fragmented across files, spreadsheets and email handovers. The ANAO example shows that when we depend on changing spreadsheets and sending documents by email, there is a risk that some points are missed and details get lost. This makes it difficult to keep a clear long-term overview [3]. T4 confirms this pattern at SSC-ICT, where outside the audit cycle no consolidated overview of outstanding recommendations exists, and tracking is improvised each time. Both the Rekenkamer (own register with full timeline) and Lemmens (VvR practice) confirmed on the producer side that a single persistent register is the minimum requirement for follow-up to function.

Identifier	`FR1`
Type	Functional
Description	The artefact must provide a single, persistent register that holds all ICT-related recommendations addressed to the organisation, together with the full history of decisions, ownership changes, and status changes per recommendation.
Rationale	Without a persistent log, longitudinal visibility cannot be guaranteed and recommendations risk being lost between handovers, leadership changes, or programme cycles.
Sources	Lit [3, 4]; T4; producer-side (Rekenkamer, Lemmens).
Fit criterion	For any recommendation in the register, the full timeline of changes (status, owner, decisions) can be retrieved in one query without consulting other systems.
Priority	Must

FR2 - Minimum data per recommendation
From the producer side, the interviews with the Rekenkamer and with Lemmens both came, independent from each other, to the same small set of register fields: the source, the recommendation itself, the owner, the lead time, the internal monitoring, and a link to the plan under which the recommendation is executed. This consensus is informative because the two positions are independent. T6 showed that the diversity of recommendation sources, ranging from ADR audits to ACICT, the Rekenkamer, statutory screenings, and ad-hoc consultancy reports, makes such data necessary if comparison across sources is to be possible.

Identifier	`FR2`
Type	Functional
Description	Each recommendation in the register must record, at minimum: the source, the recommendation text, the owner (named role or function, not person), the lead time, the internal monitoring arrangement, and a link to the plan under which the work is carried out.
Rationale	These six fields are the minimum set on which independent producer-side respondents agreed; below this set, recommendations cannot be compared, prioritised, or steered consistently.
Sources	Producer-side (Rekenkamer, Lemmens); T3, T4, T6.
Fit criterion	A sample of at least 5 recommendations from different sources can be entered with all six fields filled; any missing field is flagged at intake.
Priority	Must

FR3 - Decision-relevant status model
Callejas shows that simple binary status categoriesfail when recommendations stay open over long periods because of external dependencies, funding constraints, or political processes; recommendations can be marked "in progress" indefinitely without representing the underlying reason for non-progress. T1 confirmed at SSC-ICT that follow-up is not the same as implementation: an explicitly documented decision not to act on a recommendation can be a valid form of follow-up, provided the rationale is recorded. Interpretation 4 showed further that closure on the SSC-ICT side does not always imply closure of the recommendation as a whole.

Identifier	`FR3`
Type	Functional
Description	The artefact must support a status model that captures both the progress of a recommendation and the reason for non-progress. At minimum, the model must be able to distinguish the following situations: open, in progress, blocked by external dependency, deprioritised, superseded, under reassessment, closed (implemented), and closed (residual risk accepted with rationale).
Rationale	A binary status hides the underlying conditions that determine whether a recommendation can move forward, and therefore prevents meaningful steering.
Sources	Lit [4]; T1, T5; Interpretation 4.
Fit criterion	Each of the eight situations can be expressed in the register through the status fields and their value sets. Closure of any kind requires a rationale field to be filled.
Priority	Must

FR4 - Internal/external dependence Interpretation 4 (Section 4.2.4) showed that some recommendations cannot be resolved within SSC-ICT's direct sphere of influence, because their resolution depends on decisions, resources, or actions held by other organisations. The producer-side interview with the Rekenkamer added that ICT recommendations involving the coordinating role of BZK across ministries can stay open for years for exactly this reason. T2 and T3 both confirmed from the receiver side that these kinds of recommendations need another way of handling than the fully internal recommendations.

Identifier	`FR4`
Type	Functional
Description	For recommendations with cross-organisational dependencies, the artefact must separately track the internal slice (closable by SSC-ICT) and the external slice (requiring cross-organisational agreement), and must prevent closure of the recommendation as a whole until both slices are resolved or explicitly retired.
Rationale	Treating chain-dependent recommendations as a single object hides the part of the work that SSC-ICT cannot finish on its own and creates the risk of premature closure on paper.
Sources	Interpretation 4; T2, T3; producer-side (Rekenkamer); Lit [4].
Fit criterion	For a chain-dependent recommendation case, the artefact records the two slices separately; closure of the internal slice does not automatically close the recommendation as a whole.
Priority	Must

FR5 - Three-layer follow-up
ACICT was explicit in the producer-side interview that an effective arrangement needs a three-layer structure ("stappengelaagdheid"): (1) a central register of all recommendations, (2) monitoring of actual follow-up, and (3) measurement of the effect. Lemmens makes the same distinction with the terms follow-up ("opvolging"), implementation ("doorwerking"), and impact [20]. He also warns that this is usually the point where most follow-up in the public sector actually stops, so the effect level stays not well developed.

Identifier	`FR5`
Type	Functional
Description	The artefact must distinguish, per recommendation, three layers: follow-up (formal commitment), implementation (actual operational change), and effect (intended outcome). Each layer must have its own status field and its own evidence requirement.
Rationale	Conflating the three layers makes "open" and "closed" ambiguous and prevents the organisation from knowing whether an item has only been acknowledged, has actually been implemented, or has produced its intended effect.
Sources	Lit [20]; producer-side (ACICT, Lemmens).
Fit criterion	For any recommendation, the register holds three distinct status fields; reporting can show counts per layer (e.g. "followed up but not yet implemented", "implemented but effect not yet measurable").
Priority	Must

FR6 - Intake assessment of recommendation quality
T6 showed that incoming recommendations are not equally implementable. Some are too abstract to act on without further interpretation "improve the security of BZ" ("verbeter de security van BZ"), some restate what the organisation already knows "open goals" ("inkoppers"), and some are formulated in moments where the underlying object is itself still under change. Lemmens added on the producer side that recommendations should be "concrete, executable, and addressed to a named role" ("concreet, uitvoerbaar en geadresseerd"). ACICT warned in parallel that recommendations should leave the receiving organisation room to choose the path "multiple roads to Rome" ("meerdere wegen naar Rome"), so this requirement is not about prescribing a single solution.

Identifier	`FR6`
Type	Functional
Description	At intake, each recommendation must be assessed on three quality attributes: concreteness, addressability (named owner role), and executability. Recommendations that fail one or more of these attributes are flagged as "decomposition needed" before they can be assigned an owner.
Rationale	Abstract or unaddressed recommendations cannot be owned in any meaningful sense; the register would only record their existence, not move them forward.
Sources	T6; producer-side (Lemmens, ACICT).
Fit criterion	For a sample of 5 recommendations from different sources, the intake check produces a decomposition action for those that fail the three-attribute test; entries flagged "too abstract" cannot be assigned an owner before decomposition.
Priority	Should

FR7 - Transparent prioritisation logic
Because ICT recommendations target multiple values (risk reduction, service continuity, public-value dimensions), prioritisation cannot rely on ad-hoc judgement alone. Choi et al. propose multi-dimensional assessment combined with an explicit prioritisation technique, motivated by resource constraints and the need to direct attention where corrective measures matter most. T5 confirmed at SSC-ICT that prioritisation across recommendations is currently underdeveloped and that the result is the "everything is important" problem. The Verzamelbrief Digitalisering points in the same direction at the level of central-government ICT policy [35].

Identifier	`FR7`
Type	Functional
Description	The artefact must support an explicit prioritisation logic that uses transparent criteria and weightings, that can be revised when context changes, and that produces an ordering across recommendations from different sources.
Rationale	Without explicit weightings, every list looks equally urgent and trade-offs become implicit; explicit prioritisation makes trade-offs discussable and traceable.
Sources	Lit [7, 32, 35]; T5.
Fit criterion	For any prioritisation decision, the register records the criteria used, the weighting applied, and the rationale; revision moments can be triggered without re-entry of the recommendation.
Priority	Must

FR8 - Mapping to P&C cycle
T8 was explicit that any new instrument should connect to existing workflows rather than form a parallel track. CIO Rijk repeated the same point on the producer side: "the best option is to connect it to planning and control" ("het beste is om het te laten aansluiten bij planning en controle"). Lappi, Aaltonen, and Kujala support this in the wider literature: ICT change is executed through projects, and governance should connect operational initiatives to the higher strategic level through vertically aligned portfolio practices.

Identifier	`FR8`
Type	Functional
Description	Each recommendation in the register must be linkable to the location where the work is actually carried out (programme, project, or whichever planning-and-control instrument is in use), and that link must surface in the workflow tool the team uses.
Rationale	Without this link, follow-up at register level and execution at team level remain disconnected; T4 already showed that the link to the originating recommendation tends to disappear once a plan has been written.
Sources	Lit [19, 35, 33]; T8, T9; producer-side (CIO Rijk).
Fit criterion	For a sample recommendation, the link from the register to the corresponding work item can be opened directly; missing links are visible in the register.
Priority	Must

FR9 - Review moments scaled to risk

Meyerhoff Nielsen describe how mandatory business cases and recurring progress reporting in the Danish digital governance arrangement create review moments that can trigger redesign, delay, or termination of high-risk initiatives. The same case also notes the perceived burden of periodic reporting, which means that review intensity needs to be balanced against capacity. T5 underlined the same balance from the receiver side: recommendations are sometimes "a trajectory of several years" ("een traject van meerdere jaren"), and their urgency, feasibility, and scope can shift along the way.

Identifier	`FR9`
Type	Functional
Description	The artefact must schedule periodic review moments per recommendation, with review intensity scaled to risk and priority. At each review, continuation, revision, or closure must be decided deliberately.
Rationale	Without scheduled reviews, recommendations drift; with uniform scheduling, the review process itself becomes a burden. Scaling review intensity to risk allows the organisation to spend effort where it matters.
Sources	Lit [23, 22]; T5.
Fit criterion	For each recommendation, a next-review date is set at intake; overdue reviews are flagged; review intensity differs by risk band.
Priority	Must

FR10 - Learning loop

Di Giulio and Vecchi argue that ICTs act as ordering forces that reshape interdependencies, and that policy actors must learn how to govern them by activating learning processes in response to failure. The OECD toolkit [27] reaches asimilar conclusion: a learning-oriented evaluation culture treats unexpected outcomes as opportunities for learning rather than as failures. Lemmens also warned that registers that do not have a space to explain each item become only simple tick-lists and supports active learning culture.

Identifier	`FR10`
Type	Functional
Description	The artefact must capture, per recommendation, the assumptions made at intake, the deviations observed during implementation, and the lessons drawn from these deviations. Reviews must update these fields rather than overwrite them.
Rationale	Implementation experience often differs from the situation assumed at intake; without an explicit learning loop, this knowledge is lost between cycles.
Sources	Lit [14, 27]; producer-side (Lemmens).
Fit criterion	For any recommendation, the register holds at least one "assumptions / deviations / lessons" field; closure rationale is mandatory and is not allowed to be empty.
Priority	Should

FR11 - Ex ante and ex post

Jukić, Vintar, and Benčina emphasise that ex ante and ex post evaluation should function as one integrated learning cycle: ex ante anticipates consequences and compares alternatives, while ex post identifies actual effects and feeds back into future ex ante assessments. Choi et al. make the same point with an assessment method that explicitly supports both. In practice, this means that the artefact must record what is expected at intake (ex ante) and compare it later to what actually happened (ex post).

Identifier	`FR11`
Type	Functional
Description	At intake, the artefact must record expected outcome, expected risks, and key assumptions per recommendation. At later review, it must record observed results and the deviation against the ex-ante prediction.
Rationale	Without ex-ante / ex-post linking, the organisation cannot tell whether a closed recommendation produced its intended effect, and cannot use that information to improve future intake.
Sources	Lit [17, 7, 23, 5]; producer-side (ACICT, CIO Rijk).
Fit criterion	Per recommendation, ex-ante prediction and ex-post observation are recorded in distinct fields; deviation between the two is computable.
Priority	Should

FR12 - Ownership

T3 was the largest theme in the receiver-side analysis and identified ownership as the central bottleneck. The Rekenkamer reinforced the same point on the producer side, framing assignment as the earlier mentioned fitting quote, a question of "name and jersey number" ("naam en rugnummers"). Lemmens added that recommendations must be addressed to a specific role or function ("geadresseerd"), or else nobody can feel ownership. Umor, Zakaria, and Sulaiman show in the Malaysian audit context that follow-up is best understood as a governance network practice in which a dedicated coordination role plays a central part.

Identifier	`FR12`
Type	Functional
Description	For each recommendation, the artefact must record an explicit owner expressed as a named role or function (not a person), a defined escalation path up to a decision body, and, for chain-dependent items, a cross-unit dependency handler.
Rationale	Person-bound ownership breaks under staff turnover; role-bound ownership survives. Without a defined escalation path, blocked items have nowhere to go.
Sources	T2, T3; producer-side (Rekenkamer, Lemmens); Lit [34, 27].
Fit criterion	Each register entry identifies a named role, not a named person; escalation path is set at intake and is revisable; chain-dependent items show their cross-unit dependency handler.
Priority	Must

FR13 - Closure with explicit rationale
Closure of a recommendation must not be silent. T1 already showed that in the audit cycle a finding can be closed because residual risk has been formally accepted, and that this counts as a valid form of closure provided the rationale is recorded. Lemmens made a related point on the producer side: register entries that have been overtaken by time should be retired with an explanation rather than dropped. Without this, the register loses its meaning and becomes a record of items that simply "vanished".

Identifier	`FR13`
Type	Functional
Description	Closure of any recommendation must require selection of a closure type (implemented; residual risk accepted; obsolete due to context; superseded by [identifier]) together with a rationale text, and must be authorised by the owner role.
Rationale	A register without rationale-bound closure cannot distinguish recommendations that were resolved from recommendations that were forgotten.
Sources	T1; producer-side (Lemmens, ACICT); Lit [4].
Fit criterion	No recommendation can be moved to "closed" without a closure type and a non-empty rationale; closure records are auditable later.
Priority	Must

4.4.5 Non-functional requirements

The non-functional requirements describe how well the artefact must do its work. They concern usability, administrative burden, internal transparency, and the cultural and organisational conditions under which the artefact is meant to function.

NFR1 - Embedding in existing moments
Wanckel shows that even a well-designed digital tool can become symbolic if it does not fit existing coordination practices and portfolio routines, and that dominant norms can lead to assessments being done late in the process mainly as a formality. CIO Rijk gave the practical version of the same warning: the register must not turn into "een corvée klusje in de rapportagecirkel" (a chore in the reporting cycle), just like Lemmens, the Rekenkamer and ACICT all mentioned in a similar way. The artefact must therefore surface in pre-existing meetings and reporting cycles, not in a separate tracking ritual.

Identifier	`NFR1`
Type	Non-functional
Description	Follow-up actions in the artefact must surface in pre-existing decision moments and reporting cycles, rather than create new meetings or reports for their own sake.
Rationale	Embedding in real decision moments protects against symbolic compliance and against the "box-ticking" failure mode.
Sources	Lit [36, 33]; T1, T8; producer-side (CIO Rijk, Rekenkamer, ACICT, Lemmens).
Fit criterion	For each recommendation, at least one review moment is mapped to a pre-existing meeting (DT meeting, audit committee, board, OKR review); no new recurring meeting is introduced specifically for the artefact.
Priority	Must

NFR2 - Low administrative burden
CIO Rijk was emphatic on this point: "as few reporting burdens as possible, as few administrative burdens as possible" ("zo min mogelijk rapportage lasten, zo min mogelijk administratie lasten") and all other producers had statements that aligned with this one. The Danish case is consistent: even mandatory business cases and recurring reporting are perceived as a burden when they grow beyond what is needed for steering [23]. T7 showed that capacity is already concentrated on colleagues who are heavily loaded; any addition will land on the same people.

Identifier	`NFR2`
Type	Non-functional
Description	The artefact must keep mandatory fields and reporting actions to the minimum needed for steering. Optional fields are allowed; mandatory fields must be justified by explicit steering use.
Rationale	High administrative burden directly reduces the probability that the artefact will be used.
Sources	Producer-side (CIO Rijk, ACICT, Lemmens, Rekenkamer); Lit [23]; T7.
Fit criterion	In stakeholder evaluation, the construct means of the Method Evaluation Model are above the neutral midpoint of the scale, and the open answers do not name the administrative burden as a reason not to use the artefact.
Priority	Must

NFR3 - Internal transparency
The Rekenkamer linked follow-up effectiveness to public visibility: "shame is a negative word, but making it public also helps" ("shame is een negatief woord, maar dat ook openbaar maken helpt"). For SSC-ICT this external visibility is partly absent (Section 2.7); Lemmens therefore argued that the corresponding mechanism must be built internally, through an internal-controller function and through visibility across the organisation. T4 confirmed that the absence of a shared overview is one of the main sources of fragmentation today.

Identifier	`NFR3`
Type	Non-functional
Description	All open recommendations within the organisation's sphere of influence must be visible across organisational boundaries inside SSC-ICT, subject to the confidentiality constraints in PC4.
Rationale	Without internal transparency, there is no organisational equivalent of the public-visibility mechanism that supports follow-up at the Rekenkamer.
Sources	Producer-side (Lemmens, Rekenkamer); T4; Lit [3].
Fit criterion	Any DT member can retrieve, in one query, an overview of all open recommendations within their scope; access control is aligned with classification.
Priority	Must

NFR4 - Resilience against staff turnover
The Rekenkamer estimated that a single dossier can pass through four owners in two years, breaking institutional memory. T3 said this on the receiver side, where confidence in follow-up was tied to specific people rather than to the process. Interpretation 3 showed the same movement in the strategic-follow-up pattern: progress usually depends on one person who pushes it, and it stops when this person leaves or changes focus.

Identifier	`NFR4`
Type	Non-functional
Description	The continuity of follow-up must not depend on any specific person. When the owner role changes incumbent, the new incumbent must be able to resume from the current register state without reconstructing context from outside.
Rationale	Person-bound continuity is the failure mode behind the strategic-follow-up pattern in Interpretation 3.
Sources	Producer-side (Rekenkamer); T3; Interpretation 3.
Fit criterion	For a sample recommendation, simulated owner transition leaves the new owner with full history, current status, current dependencies, and next-review date directly from the register.
Priority	Must

NFR5 - Capacity awareness
T7 showed that follow-up activities focus on colleagues who already have the needed domain knowledge and who are usually already very busy. The Rekenkamer added that short-term pressure on capacity often pushes away attention for long-term issues. The artefact cannot solve the capacity problem itself, but it can make the concentration of capacity visible at the level where we can redistribute it.

Identifier	`NFR5`
Type	Non-functional
Description	The artefact must signal when a single owner role carries an unusual concentration of open recommendations, so that redistribution or escalation can be triggered.
Rationale	Without this signal, the same colleagues end up with all the work, and follow-up stalls for capacity reasons rather than for content reasons.
Sources	T7; producer-side (Rekenkamer).
Fit criterion	The register can show, per owner role, the current count of open recommendations weighted by priority; a configurable threshold flags concentration.
Priority	Should

NFR6 - Learning culture compatibility
Lemmens warned that a follow-up arrangement only works inside "a learning culture" ("een lerende cultuur") where "contradicting views" ("tegenspraak") is welcome. Without that, the register and the controller together still will not produce learning. The OECD toolkit makes the same point at the level of evaluation maturity: low-maturity contexts are characterised by fragmentation and absence of central visibility, and follow-up becomes effective only when ownership, accountability, and decision-making integration are strengthened [27].

Identifier	`NFR6`
Type	Non-functional
Description	The artefact must support, not replace, a learning-culture environment. Closure rationales must be reviewable by others; deviations between ex-ante and ex-post records must be discussable without blame attached.
Rationale	The artefact cannot create a learning culture by itself, but it must not work against it; in particular, it must not turn ex-post deviation into a punishment mechanism.
Sources	Producer-side (Lemmens); Lit [27, 14].
Fit criterion	In stakeholder evaluation (MEM constructs), respondents do not perceive the artefact as a control or compliance tool; perceived usefulness reflects steering and learning value rather than reporting compliance.
Priority	Should

NFR7 - Transferability beyond SSC-ICT
In Section 3.5 we state that the artefact is designed and evaluated at SSC-ICT, but specified at a level intended to be transferable to similar government organisations that deal with recurring ICT-related recommendations across multiple oversight and advisory channels. ACICT confirmed in the producer-side interview that the underlying patterns at SSC-ICT are not SSC-ICT-specific, but government-wide. So in the description of the artefact, we must clearly separate elements that are about general follow-up governance needs from elements that are specific for the SSC-ICT situation.

Identifier	`NFR7`
Type	Non-functional
Description	The artefact must be specified in such a way that general elements (register structure, status model, three layers, role-bound ownership, scheduled reviews) can be reused by similar shared-service organisations without redesign, while context-specific elements (CIO-system roles, agile portfolio attachment, specific Dutch governance bodies) are explicitly marked as such.
Rationale	The intended contribution of this thesis is both practical (for SSC-ICT) and academic (for similar contexts); the second contribution requires an explicit separation between general and case-specific elements.
Sources	Methodology (Section 3.5); producer-side (ACICT, CIO Rijk); Lit [33, 32].
Fit criterion	The artefact specification distinguishes, per element, whether it is general or case-specific; an external reader can identify which parts apply outside SSC-ICT.
Priority	Should

4.4.6 Traceability matrix

Table 4.6 maps each requirement back to the three evidence streams that motivate it: receiver-side interview themes (T1–T10, see Section 4.1) together with the four process interpretations (Section 4.2), the producer-side interviews (Rekenkamer, ACICT, CIO Rijk, Lemmens, see Section 4.3), and the literature references (Chapter 2). The matrix has two analytical purposes. First, it shows on how many evidence streams each requirement rests. Most requirements rest on two or three streams. Two constraints rest on the receiver-side material alone, and the difference between them matters: PC3 is supported by one stream, but the theme behind it (T9) is built on codes from several interviews, while PC4 rests on a single interview and is marked as such. Second, it makes it possible for the reader to check, for each requirement, whether the empirical and theoretical sources really fit together or point in different directions. When a requirement is supported by all three streams, this counts as relatively strong evidence in the synthesis: when the support is not equal, this is visible in the priority (*Should* instead of *Must*).

The producer-side interview column lists the four respondents with an abbreviation: Rek = Rekenkamer, AC = ACICT, CIO = CIO Rijk, Lem = Lemmens. Lemmens appears both as a producer-side interview respondent (column "Producer-side interviews", abbreviated as Lem) and as the author of the effect chain article [20] (column "Literature"). Both are mentioned in cases where both are relevant, because the interview and the article have a different strength as evidence.

Table 4.6: Traceability matrix linking each requirement to its evidence sources.

Req.	**Receiver-side themes / interpretations**	**Producer-side interviews**	**Literature**
PC1	T2, T10	CIO	[33, 35]
PC2	T8	CIO	—
PC3	T9	—	—
PC4	T10 (one interview)	—	—
PC5	—	Lem	[20]
FR1	T4	Rek, Lem	[3, 4]
FR2	T3, T4, T6	Rek, Lem	—
FR3	T1, T5, Interp. 4	—	[4]
FR4	T2, T3, Interp. 4	Rek	[4]
FR5	—	AC, Lem	[20]
FR6	T6	AC, Lem	—
FR7	T5	—	[7, 32, 35]
FR8	T8, T9	CIO	[19, 35, 33]
FR9	T5	—	[23, 22]
FR10	—	Lem	[14, 27]
FR11	—	AC, CIO	[17, 7, 23, 5]
FR12	T2, T3	Rek, Lem	[34, 27]
FR13	T1	AC, Lem	[4]
NFR1	T1, T8	CIO, Rek, AC, Lem	[36, 33]
NFR2	T7	CIO, Rek, AC, Lem	[23]
NFR3	T4	Rek, Lem	[3]
NFR4	T3, Interp. 3	Rek	—
NFR5	T7	Rek	—
NFR6	—	Lem	[27, 14]
NFR7	—	AC, CIO	[33, 32]

Chapter 5

Artefact design

This chapter translates the requirements synthesised in Chapter 4 into the artefact that this thesis produces. The artefact is a governance method. Therefore, it is shown as one clear set of parts that together tries to meet the requirements PC1–PC5, FR1–FR13, and NFR1–NFR7. Each component is described with reference to the requirements it covers, so that the link between requirement and design choice stays visible and the artefact can be evaluated against the fit criteria from Chapter 4 during the demonstration in Chapter 6.

5.1 Recommendation register

The register sits in the middle of the artefact, so we describe it first. We specify it on the level of fields, not on the level of a tool. The register is a logical structure that can be implemented in the tools that are already in use at SSC-ICT (audit log, auditmonitor, TopDesk, Jira, FCC, Qlik Sense), in line with PC2. A public open source template is made available on j-en-t.nl.

5.1.1 Schema

Per recommendation, the register holds the fields listed in Table 5.1. The fields are grouped in four logical groups: identification, ownership, state, and learning. We use *state* for everything in the entry that changes over time. This includes the three *status* fields, but also the dates, the priority, and the closure record. We keep the word *status* for the three effect chain layers described in Section 5.2. The mandatory minimum within these groups follows directly from FR2 and gives the six-field set on which the Rekenkamer and Lemmens independently agreed: source, recommendation text, owner, lead time, internal monitoring, and a link to the plan under which the work is carried out. The other fields are added because the other functional requirements (FR3, FR5, FR7, FR9, FR10, FR11, FR12, FR13) need them.

Table 5.1: Register schema. Per field, the table lists the type or value set, whether the field is mandatory at intake (Yes), conditionally required (Cond.), or optional (No), and the requirement that motivates the field.

Field	**Type / value set**	**Mand.**	**Req.**
Identification			
Identifier	text (auto-generated, unique)	Yes	FR1
Source	enum: ADR, Rekenkamer, ACICT, internal advice, other	Yes	FR2
Source report	text (report title) + date of issue	Yes	FR2
Recommendation text	text (as it appears in the source report)	Yes	FR2
Date received	date	Yes	FR1
Confidentiality	enum (per organisational classification scheme)	Yes	PC4, NFR3
Intake quality			
Intake quality flags	three booleans: concrete? addressed? executable?	Yes	FR6
Ownership			
Owner role	text (named role or function, not a named person)	Yes	FR2, FR12
Escalation path	text (named decision body, e.g. DT, audit committee)	Yes	FR12
Cross-unit handler	text (named role responsible for the external slice)	Cond. (a)	FR4, FR12
State (append-only: every change writes a new row)			
Follow-up status	enum (5 values)	Yes	FR5
Implementation status	enum (5 values)	Yes	FR5
Impact status	enum (4 values)	Yes	FR5
Status reason	enum (3 values)	Cond. (b)	FR3
Lead time	date (target completion)	Yes	FR2
Next-review date	date	Yes	FR9
Review intensity	enum (low / medium / high), derived from priority	Yes	FR9
Internal monitoring	text (named arrangement: who monitors, with which cadence)	Yes	FR2
Link to plan / work item	reference to the implementation plan or work item	Yes (c)	FR2, FR8
Priority	enum (high / medium / low)	Yes	FR7
Priority record	text (criteria used, weighting applied, rationale)	Cond. (d)	FR7
Closure record	enum (4 closure types) + text (rationale)	Cond. (e)	FR13
Learning			
Ex-ante record	text (expected outcome, expected risks, key assumptions)	Yes	FR11
Ex-post record	text (observed results, deviation against ex-ante)	Cond. (f)	FR11
Lessons log	text, append-only (assumptions, deviations, lessons)	No	FR10

The mandatory fields in Table 5.1 are not chosen by default. Each field group has its own steering purpose, and a field is only marked mandatory when no steering decision can be made without it. The identification group is mandatory because without source, text, and date, recommendations cannot be compared across reports (**FR1**, **FR2**). The ownership group is mandatory because, as T3 shows, follow-up stops when ownership is unclear (**FR12**). In the state group, the three status fields, the lead time, the next-review date, and the link to the plan are mandatory because they together carry the steering signal of the entry (**FR5**, **FR8**, **FR9**); the priority follows the same logic (**FR7**). The intake-quality flags form their own group, directly after identification, because they are the gate between registration and ownership. An entry that fails one of the flags cannot get an owner (**FR6**). The ex-ante record is mandatory in the learning group because it is written once at intake and cannot be added later without losing the original meaning (**FR11)**. All other fields are conditional or optional, so the register does not ask for information that is not needed at that moment.

The conditional cases in Table 5.1 apply as follows.

(a) The cross-unit handler is required only when the recommendation has a cross-organisational dependency, for fully internal items the field stays empty.

(b) The status reason is required when one of the three status fields takes a value that is not the default (open / not started / not yet measurable) and the value alone does not yet say why; it stays empty when the default applies.

(c) The link to plan or work item is mandatory in the schema, but missing links are allowed during the short window between intake and plan creation; missing links are flagged in the register rather than blocking entry.

(d) The priority record is required at every prioritisation decision; between decisions the field is not updated.

(e) The closure record is required at closure and the rationale must be non-empty; before closure the field is empty.

(f) The ex-post record is required from the first review onwards; at intake only the ex-ante record exists.

An *enum* field can only take one value from a fixed list. The value sets for the enum fields in Table 5.1 are the following.

- **Follow-up status:** not yet handled; acknowledged; decided to act; decided not to act; superseded.
- **Implementation status:** not started; in progress; paused; blocked by external dependency; delivered.
- **Impact status:** not yet measurable; under measurement; effect observed; effect not observed.
- **Status reason:** deprioritised; under reassessment; on hold (other, specified in the lessons log).
- **Closure types:** implemented; residual risk accepted; obsolete due to context; superseded by [identifier].

5.2 Status model

The register holds three status fields per recommendation. They reflect the three stages of the effect chain: follow-up, implementation, and impact. Using one combined status field would hide the difference between a recommendation that is only acknowledged, a recommendation that is implemented in the operation, and a recommendation that already produced its intended effect. The interview data treated these three aspects as one, and this is the confusion that the method tries to remove. This perspective deals with **PC5**, **FR3**, and **FR5**.

5.2.1 The receiver-side reading of the effect chain

The effect chain in Section 2.7 is described from the producer side: the Rekenkamer produces a recommendation, and it loses control after publication. SSC-ICT sits on the other side of the same chain. The same three stages are used, but the place where things happen and how control is organised are different. Figure 5.1 shows the two views next to each other, so the difference is visible.

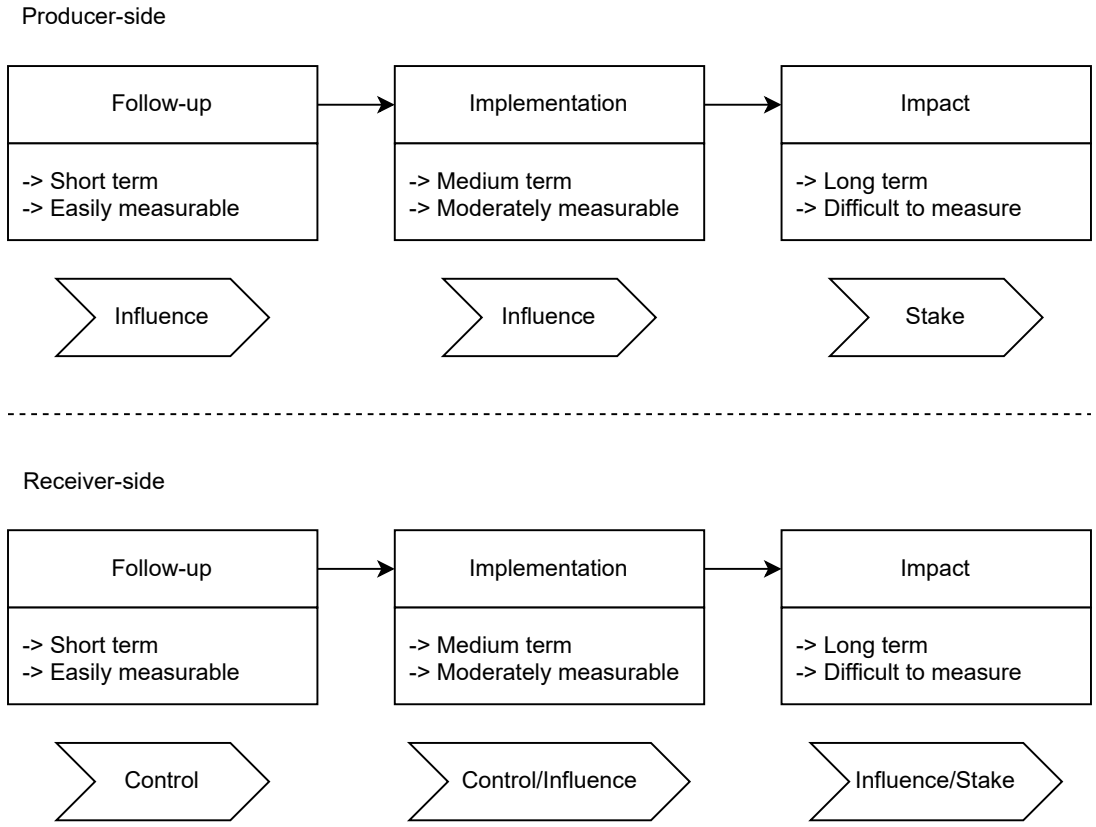


Figure 5.1: The same effect chain seen from two sides. Top: producer side according to [20]: the Rekenkamer has only influence on follow-up and implementation, and a stake in impact. Bottom: receiver side: SSC-ICT has control over follow-up and implementation (with influence where cross-organisational dependencies apply), and influence or stake on impact. Time horizons and measurability match across the two sides, but the control profile differs. In [20], the producer-side bands belong to the logic model (Figure 2.2), in which the outcome stage covers follow-up and implementation together. They are shown on the effect chain here for visual parallel. Adapted from [20].

For SSC-ICT as receiver, the three stages mean the following:

- *Follow-up.* The internal formal commitment: the register entry is created, the owner role is assigned, and the decision (act, not act, superseded) is recorded. This is the moment SSC-ICT *gains* control over the recommendation.
- *Implementation.* The operational change inside SSC-ICT: a linked work item is delivered, a procedure is adjusted, a system is updated. SSC-ICT keeps full control here, except where a cross-organisational dependency reduces it to influence.
- *Impact.* The operational or service-level effect: improved service continuity, reduced residual risk, user-side improvements. Here SSC-ICT only has influence (and partly stake), because the actual effect depends on users, partner organisations, and the technical environment.

The two sides of the effect chain meet at the handover point: where the producer loses control, the receiver gains it. Two more points come out of this comparison. First, measurability is differently distributed: the producer can measure follow-up easily through documentable formal acts at the council, while the receiver measures follow-up through register entries and tooling. Implementation is harder for the producer to measure, and easier for the receiver, because the receiver owns the work items. Impact is hardest for both. Second, for the receiver side, the chain has the same contextual limit that Section 2.7 already mentioned for the multi-source case: SSC-ICT gets advice from different producers in a governance system with several layers, so the chain maybe needs to be organised per source and not one time for the whole organisation.

5.2.2 Status fields and value sets

Each layer has its own small value set. The value sets are short on purpose. This keeps the method simple and stops the register from growing with extra status types over time. Table 5.2 lists the value set per layer.

Table 5.2: Status values per layer of the receiver-side effect chain.

Layer	**Value set**	**Default at intake**
Follow-up	`not yet handled`, `acknowledged`, `decided to act`, `decided not to act`, `superseded`	`not yet handled`
Implementation	`not started`, `in progress`, `blocked by external dependency`, `paused`, `delivered`	`not started`
Impact	`not yet measurable`, `under measurement`, `effect observed`, `effect not observed`	`not yet measurable`

Two rules apply across the three layers. First, every status change writes a new row; the previous row stays. This gives the full timeline that **FR1** asks for, and it does not depend on a person remembering the history. Second, when a status takes a non-default value, the `status reason` field must be filled. For the defaults, no reason is needed.

A recommendation can sit in different combinations across the three layers at the same time. The combination `decided to act / blocked by external dependency / not yet measurable` describes a recommendation that is formally accepted but cannot move forward because of cross-organisational factors. A single status field cannot express this combination. The three-layer model does, which satisfies **FR4**.

The three layers also work like gates. Layer 2 only becomes active when layer 1 reaches `decided to act`, layer 3 only when layer 2 reaches `delivered`. This implements **FR5**, that closure of one layer does not imply closure of another.

The value sets above cover the eight situations from **FR3**, but not as one flat list. The situations are expressed through combinations of the fields. An entry is *open* when the follow-up status is still *not yet handled* or *acknowledged*. *In progress* and *blocked by external dependency* are values of the implementation layer. *Deprioritised* and *under reassessment* are status reasons, because they say why an entry does not move, not where it stands. *Superseded* is a value of the follow-up layer and also a closure type. The two closed situations are closure types. During the design one value was added that FR3 does not name: *paused*, for work that is stopped inside the organisation without an external dependency, for example because of capacity (T7). In a flat list, the progress and the reason for non-progress would sit in the same field. The layered model keeps them apart, and that is what **FR3** asks for.

5.3 Intake

The register holds entries and intake is the process step that produces those entries. At intake, a recommendation arrives, is checked for usability, gets an owner, gets a starting priority, and is linked to the plan or work item where the work happens. Without an explicit intake step, the register fills with whatever arrives. The problems from the receiver-side process patterns (abstract recommendations, missing owners, no link to operations) then return inside the artefact. Intake therefore performs three checks, in order.

The first check tests the recommendation itself on three attributes; is it concrete (specific enough to act on), is it addressable (a named role or function can tak responsibility for it) and is it executable (the action is realistic, given the known dependencies and constraints). These three attributes follow **FR6** and are stored in the register as three booleans. A recommendation that fails one of these checks is not refused. It is flagged as "decomposition needed" and cannot get an owner before it is split into sub-recommendations that do pass the test. This is how Lemmens's rule that recommendations should be "concrete, executable, and addressed" ("concreet, uitvoerbaar en geadresseerd") is applied at intake. It also stops the abstract items and open goals ("inkoppers") that T6 describes.

The second check fills the owner role, the escalation body, and, when applicable, the cross-unit handler. The owner role is a named role or function and not a named person, so that personell change does not break any continuity (**FR12, NFR4**). The escalation body is the named decision body to which the owner can escalate when the recommendation is blocked at the role's level. The cross-unit handler is mandatory when the recommendation has a cross-organisational slice. The trigger for the handler field is the intake question "can the organisation close this on its own"; a "no" makes the handler field mandatory and marks the entry as depended on the broader organisational chain.

The third check sets the initial priority. The priority is not a single number, but a triple, scored at the intake on three dimensions:

- Risk and compliance: Residual risks, audit obligation, security exposure
- Strategic alignment fit: Fit with the organisation's goals and with the planning and control instrument/cycle
- Feasibility and capacity; required capacity against available capacity, dependency situation, technical readiness

Each dimension is scored High, Medium or Low. The three dimensions are derived from Choi et al., who argue that prioritising of ICT-related changes in the public sector must be multidimensional, transparent, and usable both ex ante and ex post. Their original STOPE x AHP method is a bit too heavy for this usecase, considering **FR2**, so the three dimensions used here are an abstraction of their five: Strategy and Organisation are collapsed into Strategic Alignment, Environment and the security-side of Technology are collapsed into Risk and compliance, and the competency-side of technology together with people is collapsed into Feasibility and capacity. The relative weight between the three dimensions is set once per prioritisation cycle by the body that already does prioritisation work in the organisation, and not per recommendation. Per recommendation, only the three scores and a one-line rationale are recorded. The priority field in the register holds these. The rule that collapses the three scores into the single priority value is not fixed by the artefact. The same body that sets the weights also sets this rule, so that it fits its own organisational context, in line with the position of CIO-Rijk that an organisation should design a system that works for its own context (Table 4.4).The rule in use is part of the priority record, so every priority stays traceable to the rule that produced it.

The output of intake is one of two states. Either the recommendation is registered with all mandatory fields filled, and is ready to move into the implementation layer, or it is registered with the intake-quality flags showing one or more failures and the "decomposition needed" flag set, and the entry stays at the follow-up layer until the decomposition is done. Figure 5.2 shows the flow of intake decisions.

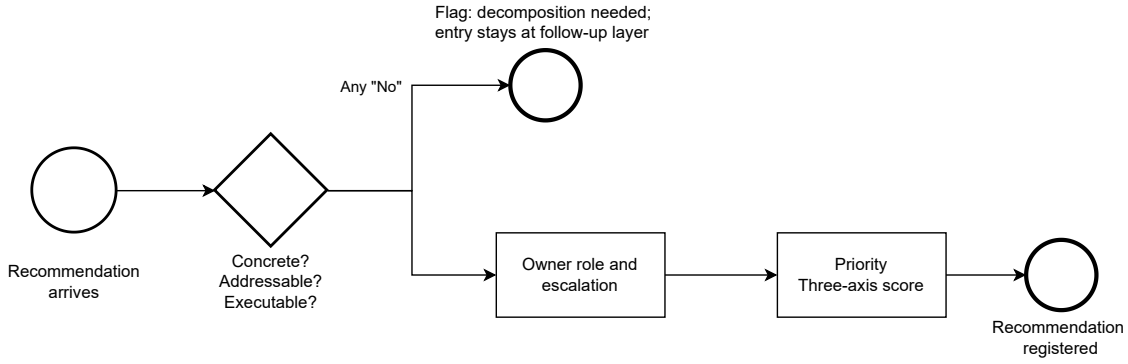


Figure 5.2: The intake decision flow of a new recommendation

5.4 Review and closure

A recommendation that is registered does not move forward by itself. Two moments determine whether the register stays current: The periodic review while the entry is open, and the closure when it is no longer open. Without scheduled reviews, recommendations drift, exactly the pattern that the receiver-side process descriptions for non-audit recommendations already showed in Section 4.1. Without explicit closure, recommendations vanish "into the drawer" instead of being retired with reasoning, which is the failure mode that Lemmens described from the producer side. The artefact handles the two moments together. At each review, the entry either continues, is revised, or is closed. Each option needs a recorded action. Closure needs more than the other two.

The review cycle is not uniform. It is scaled to risk and priority, as **FR9** requires. The review-intensity field in the register has three values; low, medium, high, and is derived from the priority field, which itself is single low / medium / high, derived at intake from the three-dimension score in the priority record. The mapping from priority to review intensity is aligned one-to-one. The next review date is set at intake and updated at every review. Overdue reviews are flagged in the register. The flag is not a missed deadline for the owner of the recommendation, but it is a signal for the escalation body to get a heads-up that that entry needs attention.

The frequency does not determine when in the time the review actually takes place. Creating a separate review meeting per recommendation is exactly what **NFR1** prevents. The artefact therefore does not introduce its own review moments, but pins each review to a pre-existing decision moment in the organisation. Which moment depends on the source of the entry. For audit-cycle entries, the existing ADR cycle already provides the review moment, with the audit log, the audit monitor, and IAA verification. The artefact follows that rhythm. For strategic and policy-driven entries, the DT meeting and (where applicable) the audit committee carry the review. For programmatic entries that run under a transformation board or steering group, the review takes place in that body. For entries linked to an OKR-style instrument, the OKR review cycle is the anchor. The escalation-path field in the register stores the named decision body per entry, and that same body carries the formal review of the entry. In this way, every recommendation has at least one review moment in an already existing decision body, and no new recurring meeting is added for the artefact itself.

At each review moment, three outcomes are possible: continue, revise, or close. Continue updates the status fields and the next-review date, and leaves the entry open. Revise updates the priority record (and through it the priority and the review intensity), the linked plan, or the cross-unit handler, and writes the change reason in the lessons log. Close requires the closure procedure described next. Closing a recommendation always needs a rationale. As **FR13** requires, no entry can move to "closed" without a closure type and a non-empty rationale, and the move must be authorised by the owner role. The four closure types are the ones already listed in the status-model value sets in Section 5.1.1.

- Implemented: The recommendation is delivered in the operation and, where measurable, the intended effect is observed. The rationale states what was delivered and on which evidence.
- Residual risk accepted: The recommendation is not fully implemented, but the remaining risk is formally accepted at a named level. This is the closure type that the audit cycle already uses, as T1 explicitly described; it is valid only when the rationale records who accepted the risk and on which basis.
- Obsolete due to context: The recommendation has been overtaken by changes in the underlying situation (organisational, technical, or regulatory). The rationale records what changed and why the recommendation no longer applies.
- Superseded by [identifier]. The recommendation is replaced by another recommendation that addresses the same concern. The rationale records the identifier of the superseding entry, so that the chain stays traceable.

For chain-dependent entries with an internal and an external slice, **FR4** adds one extra rule on top of **FR13**: closure of the internal slice alone is not enough; the external slice must also be resolved or explicitly retired; otherwise, the entry as a whole stays open. The four closure types apply per slice.

Figure 5.3 shows the closure decision flow.

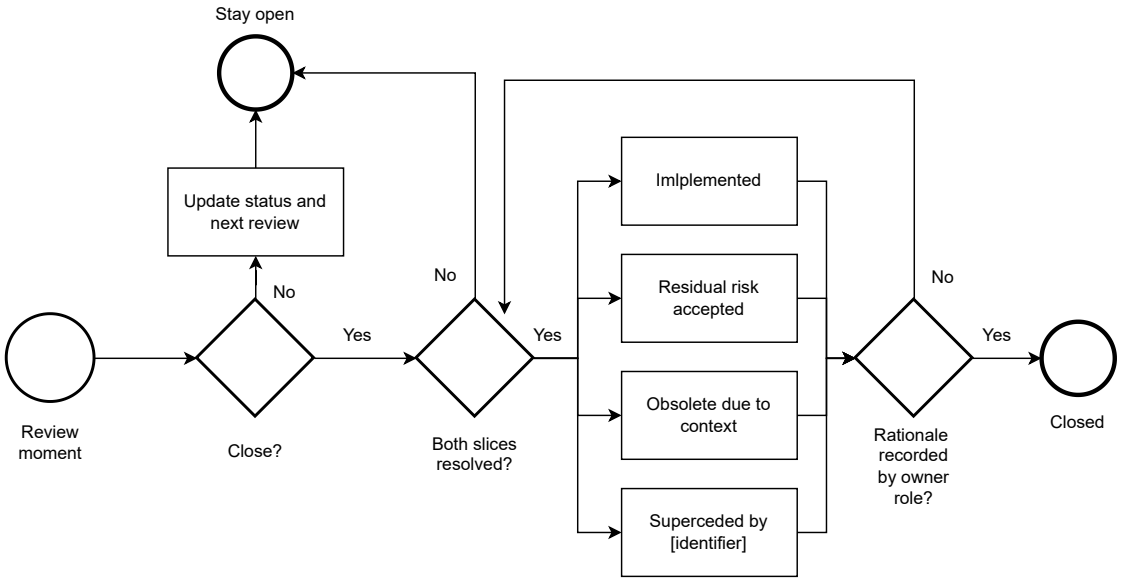


Figure 5.3: The closure decision flow

5.5 Learning loop

The artefact has so far recorded what arrives, what is decided, what is owned and what is closed. None of these records by themselves tell whether the organisation has learned anything across cycles. A register that does not learn over time fills with closed entries that nobody remembers the rationale of. Then a similar recommendation comes back a few years later. The learning loop is part of the artefact that prevents this. It does not need a separate process. It uses three register fields that were already listed (ex-ante record, ex-post record, lessons log), and a small rule for how those fields are updated.

The first part of the loop is the ex-ante record. At intake, the artefact records three things per recommendation: the expected outcome, the expected risks, and the key assumptions that were made when the entry was registered. The expected outcome is a one-line statement of what the organisation thinks will be different once the recommendation is implemented. The expected risks lists the conditions under which this outcome might not be reached. The assumptions name what is being taken for granted at intake, for example, that a depending unit will deliver on time, that a regulation will not change or that a team has the capacity to absorb the work. Together, the three form the ex-ante record. The owner role writes them once, at intake and they are not overwritten later.

The second part of the loop is the ex-post record. From the first review onwards, the artefact records the observed results against the ex-ante prediction. Deviation is the gap between what was expected at the start and what is actually seen now. Deviation is recorded as information. It is not a judgment. It can be positive (the benefit is bigger than expected), negative (the benefit is smaller than expected, or there is a side effect), or zero (the expectation was correct). All three are valid outcomes of a review. The ex-post record is updated at every review, but the ex-ante record stays as it was at intake, so a reader can compare the two side by side and trace each difference back to a specific assumption. This is the combination that **FR11** requires: ex-ante and ex-post in separate fields, where the deviation can be calculated.

The third part is the lessons log. It is the field where the assumptions, the deviations, and the lessons drawn from them are written in append-only form. Append-only means that a new entry does not replace an older one; both stay visible. This satisfies the "reviews must update rather than overwrite" rule that FR10 asks for. An entry in the log is not bound to a specific review moment; it can be added at any time and carries the date and the owner role. Closure does not need a separate lessons-log entry on its own; the closure rationale already captures the lessons of the closing item. Across the register, the collected lessons-log entries become a kind of informal knowledge base: a new intake of a similar recommendation can use the results of earlier cycles, without needing to rebuild the context from outside the artefact.

Two design choices try to keep this loop light. First, the three fields are not scored or graded; any deviation is only noted as information and not as a judgment. **NFR6** asks for exactly this: ex-post deviation must be discussable without blame attached. If we would turn deviation into a score, the loop changes into a control mechanism. The rationale fields then get filled in a defensive way, and the loop works against the learning culture instead of supporting it. Lemmens called this culture "a learning culture" and warned that the register on its own cannot create it; the register can only avoid working against it. Second, the lessons log is marked as optional in the schema, so the owner role decides whether a particular insight is worth recording. The result is that the learning loop adds no extra meeting and no extra required field on top of intake and review.

5.6 Tool mapping

The register schema described in Section 5.1.1 is a logical specification. It says which fields a recommendation has, not where those fields are stored. The choice of which tool implements which part of the register is the responsibility of the organisation using the artefact. PC2 only requires that the register does not run as a parallel track next to the tools the organisation already uses for adjacent purposes; **PC3** requires that the mapping can change over time without breaking the register schema.

The register itself is the main record for the fields that only exist for recommendation governance: identification, ownership, the follow-up status, priority and priority record, closure record, and the learning fields. None of these have a natural home in adjacent tooling, and putting them elsewhere would re-create the fragmentation that the empirical material in Section 4.1 already showed.

Two field groups do not store their own content, but instead refer to external tools. The implementation status lives in the work-item tool that the responsible team already uses for the actual work; the register's link-to-plan field stores the reference, so the user can open the work item directly. The register does not prescribe which work-item tool is used; it only needs the reference. **FR8** asks for exactly this. The impact status, when it can be measured, is taken from the metric system that the organisation already has in place (for example, an OKR cycle or service metrics that are combined in a dashboard layer). The register links the result to an already named metric; it does not create a second copy of the metric definition.

For audit-source items, the organisation's existing audit registration tools already contain a second record of the same recommendation. In the mapping, these tools are seen as the main source for the follow-up status of audit-source items; the register only reflects the closure decision and does not make its own decision. At SSC-ICT, these tools are the audit log (for IT findings) and the audit monitor (for financial and operational findings).

NFR3 and **NFR5** do not need new tools. Both are overviews built from the data the register already holds. The organisation can render these overviews in whichever reporting layer is already in use; nothing in the register design depends on a specific reporting tool.

NFR3 covers internal transparency. Within SSC-ICT, any DT member should be able to see, in one place, the open recommendations that fall within their area of responsibility. Without this cross-boundary visibility, recommendations get lost between organisational units, which **T4** identified as a main source of fragmentation. The overview is a list of open register entries, filtered by the confidentiality field so a reader only sees items their classification level allows. **PC4** governs the filter logic.

NFR5 covers capacity awareness. When one owner role carries many open recommendations at the same time, the work often stalls because of capacity, not because of content (**T7**). The overview is a per-role count of open recommendations, with high-priority items counted heavier than low-priority ones. When the total for a role crosses a configurable threshold, the role is flagged. The threshold itself can be adjusted by the organisation as the workload changes.

5.7 Governance cycle

The previous four sections describe individual parts of the artefact. This section puts them together. Per recommendation, the artefact runs through a small cycle with three moments: intake, review and closure, and a learning loop that run across these moments and feeds into future cycles. The cycle is not a calendar schedule for the organisation as a whole. It runs per recommendation. At any moment, recommendations can be on different points in the cycle.

Each of the three moments is pinned to a pre-existing decision body, so that no new meeting is created specifically for the artefact (NFR1). Intake is pinned to the body that already receives the source. For audit-source items, this is the audit committee. For all other sources (DT-commissioned reports, strategic advice, policy-driven items), this is the DT meeting. The body chosen at intake is then recorded in the escalation-path field of the register entry. Review is pinned to the same body. The escalation-path field, set at intake, also names the body for later reviews. The recommendation stays in the same governance line that produced the entry, so no handover between bodies is needed. Closure is pinned to the same body again, with one addition. For audit-source items, an independent verification step by IAA is added before the entry can be moved to "closed". This follows the existing audit closure practice and is not a new step for the artefact. The learning loop does not need a separate decision body. The ex-ante record, the ex-post record, and the lessons-log entries are filled by the owner role between meetings, and are visible to anyone who reads the entry, including future intakes of similar recommendations. Figure 5.4 shows the cycle.

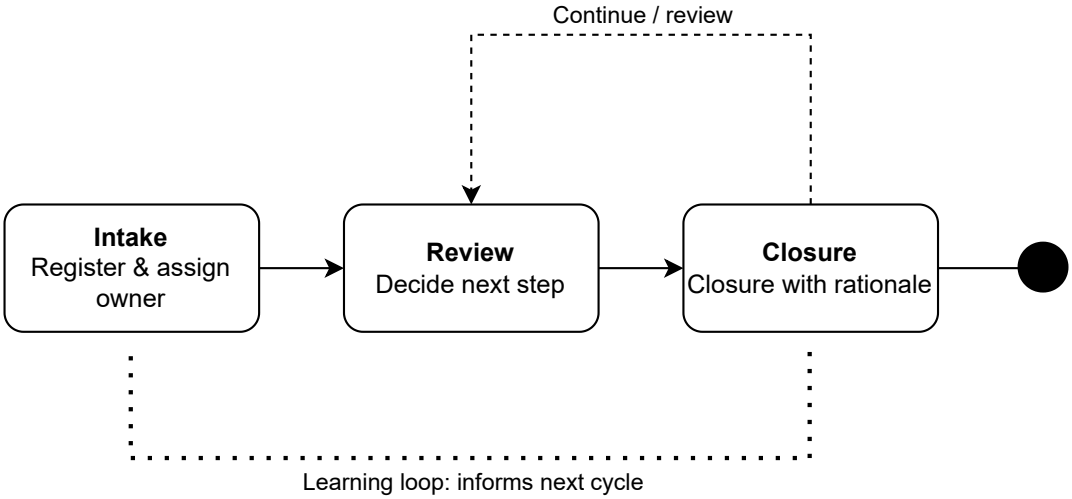


Figure 5.4: The per-recommendation governance cycle

Not every recommendation completes the three moments in the same way and it therefore does not try to imply that. Recommendations can be closed at "obsolete due to context" without reaching implementation; in that case the intake assumptions are still recorded and can be compared against the rationale for obsolescence in the closure record. Recommendations can also re-enter the cycle: when a closed recommendation is superseded by a new one, the new one starts a new cycle, and the link to the superseded entry is recorded in the closure record (**FR13**). So, each cycle has a clear limit per entry, but the total register becomes bigger over time when we start new cycles.

For the governance bodies, the cycle has a small visible effect. They do not get new meetings; they get a few extra agenda items from the register, in meetings they already hold. The intake-quality flags, the priority record, and the closure rationale are the elements that the bodies actually discuss; the rest of the register is recorded by the owner role between meetings. This split keeps **NFR1** (embedding in existing moments) and **NFR2** (low administrative burden) consistent with each other, even though in theory the two requirements tend to work against each other a bit.

5.8 Generalisability

The artefact is designed and evaluated at SSC-ICT, but it is specified at a level intended to be reusable by similar government organisations that receive recurring ICT-related recommendations through multiple oversight and advisory channels. **NFR7** requires that the description distinguishes the general elements of the artefact from the case-specific ones, so that an external reader can identify which parts apply outside SSC-ICT. The general elements depend only on the receiver-side governance problem itself: the recommendation register and its field schema, the three-layer status model based on the effect chain, the intake gate with its quality check and the three-dimension priority, the rationale-bound closure with the four closure types, the learning loop with ex-ante and ex-post pairing, the principle that field groups map onto pre-existing tools, and the per-recommendation governance cycle. These elements describe what receiver-side follow-up needs as a function. They do not describe which body or tool does the function locally. Similar organisations can therefore reuse them without redesign.

The case-specific elements depend on the Dutch central-government context (PC1) and on SSC-ICT's situation at the time of writing. These include the specific decision bodies named in the artefact (DT, audit committee, MT, DMT, IAA, the cross-organisational "bestuurlijk overleg" and "afnemersberaad"), the attachment to the CIO-system as revised in 2025, the SAFe-style agile transition and the related OKR cycle, and the specific tools mentioned by name (audit log, auditmonitor, and the dashboard layer used at SSC-ICT). A similar organisation that uses the artefact would replace each of these with its own version: other named decision bodies, another planning-and-control cycle, other tools. The general elements stay the same. This separation matches the position that ACICT took in the producer-side interview, where the patterns at SSC-ICT were described as not SSC-ICT-specific but government-wide.

Chapter 6

Demonstration

This section applies the artefact from Chapter 5 to a set of real recommendations that were addressed to SSC-ICT. The reason for this follows the validation logic in Section 3.4.4. Wieringa connects the DSRM step "demonstration" to the task of treatment validation. The demonstration looks at the effect, the trade-offs and the sensitivity of the artefact in its context. The later "evaluation" step then asks if these effects meet the requirements [37]. There is no operational implementation yet, so the treatment validation is done by making a *validation model* of the artefact in its context and by studying that model [37]. In this thesis the set of historical cases is that validation model. It shows how the artefact would work in the SSC-ICT context, and by running the cases through it, the mechanism becomes visible: how the artefact registers, structures and steers recommendations. This happens before the expert-opinion validation with stakeholders in Chapter 7. The demonstration has two goals. The first goal is to show that the method can hold the recommendations that SSC-ICT really receives, and that the differences between sources, levels and themes show up in the register fields as the design predicts. The second goal is to make the artefact concrete, so that the workshop participants can read and discuss a working example and not only an abstract description.

6.1 Demonstration definition

The demonstration is *not* an implementation evaluation. Wieringa puts the border between validation and evaluation as follows: in an evaluation the stakeholders use the artefact in the real world, separate from the researcher, while in a validation the researcher studies a model of how the stakeholders would use it [37]. Here cases are entered into the register, based on the content of public reports, and not the organisation in daily use, considering the confidentiality of the execution. The extra context needed to fill in the register is in this case based on the context the company supervisors have given. This places the demonstration on the validation side of that border. When a report does not give the information that a field asks for, the field is marked as "to be filled at intake" and stays an open question for the workshop. The samples below show the information that was available up to the interim audit of October 2025.

6.1.1 Inventory and case selection

The demonstration uses an inventory of 126 ICT-related recommendations from nine public reports. These reports are addressed to SSC-ICT or to the wider BZ–SSC-ICT transition. The inventory contains the three Auditdienst Rijk (ADR) audit reports of 2023, 2024 and the interim 2025 cycle; three ADR *onderzoeksrapporten* (research reports) on *licentiebeheer* (license management), the *kostprijsmodel* (cost-price model) and *LCM-kosten* (LCM costs); one ACICT advice on the BZ ICT transition; the matching *Kamerbrief* from the Minister and State Secretary; and one BIT-advies on DWR Next. Table 6.1 gives an overview. The full per-recommendation inventory is given in Appendix E. These reports were provided by SSC-ICT themselves at the start of the interview based on their perceived relevance to the thesis.

Table 6.1: Overview of the recommendation inventory used for the demonstration.

Code	Report	Type	Date	Rec.
AR2023	Auditrapport 2023 SSC-ICT	Audit report	15-03-2024	26
AR2024	Auditrapport 2024 SSC-ICT	Audit report	14-03-2025	20
INT2025	Interim-auditrapport 2025 SSC-ICT	Interim audit	27-10-2025	24
LIC2020	Licentiebeheer SSC-ICT	Research report	20-10-2020	14
KPM2020	Onderzoek Kostprijsmodel SSC-ICT	Research report	21-12-2020	3
LCM2024	Robuustheid LCM-kosten in kostprijsmodel	Research report	10-10-2024	7
ACICT2025	Advies AcICT: Transitie ICT-dienstverlening (BZ)	Advisory	03-02-2025	15
TK1417	Kamerbrief 26 643, nr. 1417	Cabinet letter	01-10-2025	4
BIT2017	BIT-advies programma DWR Next	Advisory	21-03-2017	13
			Total	**126**

The demonstration does not cover all 126 recommendations. The goal is to show the method on the patterns that it is designed for, and not to repeat the whole inventory. The cases are chosen on three criteria, already named in the introduction of this chapter. First, the set must contain at least one source for each process interpretation in Section 4.2, so the artefact is shown on the patterns it must cover. Second, the set must contain at least one chain-dependent recommendation, because **FR4** cannot be shown without one. Third, the set must contain at least one recommendation that does not pass the intake quality check of **FR6**, so the decomposition flag is used.

The public inventory does not contain the consultancy report that the DT commissioned (for example the Deloitte and Sopra Steria reports) that informed Interpretation 2 in Section 4.2, because those reports are internal. For the programmatic pattern the demonstration therefore uses the DWR Next trajectory. The DWR Next recommendation came from a BIT review of an IT programme rather than from a commissioned advisory report, but it was still followed up through a dedicated programme with go/no-go steering, so it shows the same programmatic pattern. Five cases are chosen. Table 6.2 links each case to the interpretation it stands for and to the requirements that it mainly exercises.

Table 6.2: Selected demonstration cases and the elements each one exercises.

Case	ID	Interp.	Primarily exercises
A	AR2023-DFA-04	1	Baseline audit path; clean intake; audit-log + IAA closure (FR1, FR13, NFR2)
B	AR2024-IT-02	1	Abstract recommendation: decomposition (FR6) and chain-dependence (FR4); recurrence/history (FR1)
C	ACICT2025-01	3	Strategic advisory intake; chain-dependence (FR4); confidentiality (PC4); external reporting cadence (FR9)
D	BIT2017-DWRN-02	2	Programmatic follow-up; link to plan (FR8); long-horizon closure (FR13) and learning loop (FR10, FR11)
E	KPM2020-01	3	Temporal decay; overdue-review flag (FR9); re-prioritisation at review (FR7); superseded-by closure (FR13); ex-ante/ex-post gap and learning loop (FR10, FR11)

Each case is registered with the schema in Section 5.1. To keep the tables readable, each register sample shows only the fields that carry the relevant signal for that case; conditional and optional fields that are empty at this moment are left out. The three-layer status is written as *follow-up / implementation / impact*.

6.1.2 Case A - baseline audit finding (Interpretation 1)

Case A is a concrete operational finding from the financial audit cycle: AR2023-DFA-04, "record the control for input verification of correct tariffs (four-eyes principle) in a visible way" ("Leg de beheersmaatregel voor invoercontrole op juiste tarieven (vier-ogen-principe) zichtbaar vast"). This is the kind of finding that Interpretation 1 already handles well. It is added to show that the method stays light for simple findings and does not add extra burden (NFR2). Table 6.3 shows the register sample.

Table 6.3: Register excerpt, Case A (AR2023-DFA-04). State shown at closure.

Field	**Value**
Identification	
Identifier	AANB-2024-014
Source	ADR
Source report	Auditrapport 2023 SSC-ICT (15-03-2024)
Recommendation	"Record the control for input verification of correct tariffs (four-eyes principle) in a visible way." ("Leg de beheersmaatregel voor invoercontrole op juiste tarieven (vier-ogen-principe) zichtbaar vast.")
Confidentiality	Internal (DT and direct staff)
Intake quality	
Concrete / addressed / executable	yes / yes / yes (flag "decomposition needed": off)
Ownership	
Owner role	Manager Finance & Control
Escalation path	DT (CFO)
State	
Priority (3D)	Risk & compliance: M; Strategic fit: L; Feasibility: H
Review intensity	Low
Follow-up / impl. / impact	decided to act / delivered / effect observed
Lead time	31-12-2024 (closed within target)
Internal monitoring	Audit log; annual ADR re-check; IAA verification
Link to plan	Finance control change item
Closure record	*implemented* — control documented and verified by IAA

The case passes the three intake checks without problems, so no decomposition is needed and an owner role is given at the start. The closure does not depend on the task owner: as in the formalised audit pattern, the evidence is verified by Internal Audit and Advisory (IAA) before the finding is closed in the log. The entry is closed as *implemented*, with a rationale. The same closure mechanism also allows the *residual risk accepted* type, which the audit cycle uses when a control reaches a level that is qualitatively acceptable and the rest of the risk is formally accepted (Section 4.2.1); Case A does not need this type. For such an item the register adds little above what the audit log already does. That is the point, the method does not add extra steering work where the existing pattern is already enough.

6.1.3 Case B - abstract recommendation requiring decomposition (Interpretation 1)

Case B is AR2024-IT-02, "establish a uniform set of standards, including clear minimum security requirements, and test periodically whether the standards are met" ("Stel een uniform normenkader op, inclusief duidelijke minimale beveiligingseisen, en toets periodiek of aan die normen wordt voldaan"). It enters through the same ADR cycle as Case A, but it is strategic and abstract, and the solution goes beyond SSC-ICT. So this case exercises the intake quality check (FR6) and chain-dependence (FR4) at the same time. It is also the case that is used in the validation workshop (Section 7).

At intake the recommendation does not pass two of the three quality attributes. The uniform set of standards ("uniform normenkader") is not concrete enough to act on without more interpretation, and it is not addressed to a named role. So the entry gets the flag "decomposition needed" and cannot receive an owner. The decomposition splits the recommendation into an internal slice (an SSC-ICT set of standards with minimum security requirements, plus a testing cycle) and an external slice (alignment with CIO Rijk on what becomes government-wide ("rijksbreed") in relation to the BIO renewal and NIS2). After the decomposition the intake checks are run again and pass, the flag is removed, and the CTO (who is also the CIO at SSC-ICT) is confirmed as owner. Table 6.4 shows the register excerpt after decomposition.

Table 6.4: Register excerpt, Case B (AR2024-IT-02), as of the interim audit (October 2025). The State fields show the append-only history: each change is a new dated line and the older lines stay. Dates are illustrative.

Field	**Value**
Identification	
Identifier	AANB-2025-014
Source	ADR
Source report	Auditrapport 2024 SSC-ICT (14-03-2025); re-issued as INT2025-IT-02 (27-10-2025)
Recommendation	"Establish a uniform set of standards, including clear minimum security requirements, and test periodically whether the standards are met." ("Stel een uniform normenkader op, inclusief duidelijke minimale beveiligingseisen, en toets periodiek of aan die normen wordt voldaan.")
Confidentiality	Internal (DT and direct staff)
Intake quality	
At intake	concrete: no; addressed: no → flag "decomposition needed": on
After decomposition	concrete: yes; addressed: yes; executable: yes → flag off
Ownership	
Owner role	CTO (also CIO) — assigned 12-06-2025, after decomposition
Escalation path	DT
Cross-unit handler	CISO / CTO toward CIO Rijk (external slice)
State (append-only: each change is a new dated line)	
Follow-up status	[14-03-2025] not yet handled — intake; no owner yet [12-06-2025] decided to act — after decomposition; CTO assigned
Implementation status	[12-06-2025] not started [18-09-2025] in progress — internal concept drafted [27-10-2025] **blocked by external dependency** — waits on the government-wide framework (Q1 2026)
Impact status	[14-03-2025] not yet measurable
Status reason	[27-10-2025] testing cycle waits on the government-wide framework (expected Q1 2026)
Lead time	[12-06-2025] 31-12-2025 [27-10-2025] revised to 31-03-2026 — follows the external framework
Priority (3D)	Risk & compliance: H; Strategic fit: H; Feasibility: M
Review intensity	High
Internal monitoring	CISO, monthly to CTO; annual ADR re-check
Link to plan	Security programme / *stuurgroep* work item
Learning	
Ex-ante record	Expected outcome: one set of standards with minimum requirements and a recurring test. Assumption: the government-wide framework arrives in time.

The State fields in Table 6.4 show the append-only rule of the register (Section 5.1.1): each status change is a new dated line in the same cell, and the older lines stay. So the full timeline can be read back from

one entry (FR1). The impact layer stays *not yet measurable* over the whole period, so it has only one line; a line is added only when a field changes.

The case shows three things. First, the intake quality check (FR6) stops an abstract recommendation from getting an owner before it is workable. Without this check the register would only note that the recommendation exists, but would not move it forward. Second, chain-dependence (FR4) is shown inside one entry: the internal slice moves on while the external slice waits at CIO Rijk, so the implementation layer has the value *blocked by external dependency* with a status reason, and the cross-unit handler field names who holds the external slice. One combined status field could not say "internally progressing but blocked externally"; the three-layer model can. Third, the same recommendation comes back in the interim audit (INT2025-IT-02). The persistent register (FR1) keeps both occurrences under one timeline and shows the "unresolved items re-enter next cycle" loop of the audit pattern, instead of starting a new and disconnected entry. The dated lines in Table 6.4 make this timeline explicit.

6.1.4 Case C - strategic advisory recommendation (Interpretation 3)

Case C is ACICT2025-01, "jointly take measures to reduce urgent security risks in the short term and report on progress every three months to the Minister of Foreign Affairs and the State Secretary" ("Tref samen maatregelen om urgente beveiligingsrisico's op korte termijn terug te dringen en rapporteer over voortgang iedere drie maanden aan de minister van Buitenlandse Zaken en de staatssecretaris"). This is a strategic advisory recommendation. In current practice such advice enters without a formalised register, so it has the highest risk of the "over the fence" ("over de schutting") pattern of Interpretation 3. The register is what gives it the structure that it otherwise misses. The case has two features that the earlier cases do not have: the measures are in a confidential annex (PC4), and the producer sets a fixed external reporting rhythm of three months, which shapes the review rhythm (FR9). Table 6.5 shows the register sample.

Table 6.5: Register excerpt, Case C (ACICT2025-01). State as of the interim audit (October 2025).

Field	**Value**
Identification	
Identifier	AANB-2025-031
Source	ACICT
Source report	Advies AcICT, Transitie ICT-dienstverlening BZ (03-02-2025)
Recommendation	"Jointly take measures [. . .] to reduce urgent security risks and report every three months [. . .]." ("Treff samen maatregelen om urgente beveiligingsrisico's [. . .] terug te dringen en rapporteer iedere drie maanden [. . .].")
Confidentiality	Restricted — measures detailed in a confidential annex; entry visible to DT and security roles only
Intake quality	
Concrete / addressed / executable	yes / yes / yes (joint with BZ) flag off
Ownership	
Owner role	CISO (SSC-ICT side)
Escalation path	DT
Cross-unit handler	BZ security lead, via programme SHIELD (external slice)
State	
Priority (3D)	Risk & compliance: H; Strategic fit: H; Feasibility: M
Review intensity	High
Follow-up / impl. / impact	decided to act / in progress / not yet measurable
Lead time	31-12-2025 ("short term" in the advice)
Next review	Bound to the three-monthly progress report to the Minister and State Secretary
Internal monitoring	CISO, monthly; external producer cadence every three months
Link to plan	Programme SHIELD work item

The case shows four points. First, a strategic advisory recommendation that has no audit-style register of its own gets the same structure as an audit finding. The asymmetry between Interpretation 1 and Interpretation 3 that Section 4.2 described is removed at the registration. Second, confidentiality (PC4) is held as a field as the entry can be marked as restricted, so the measures are only visible to DT and

security roles, while the recommendation can still be steered. Third, chain-dependence (FR4) returns here, with the external slice held through programme SHIELD and a BZ counterpart. SSC-ICT owns the internal slice but cannot close the recommendation alone. Fourth, the review moment is constrained to follow the rhythm determined by the producer (the three-monthly report). This shows how the review schedule (FR9) attaches to an existing rhythm and does not create a new one.

6.1.5 Case D - programmatic follow-up over a long horizon (Interpretation 2)

Case D is BIT2017-DWRN-02, "use, for the completion and go-live of DWR Next, a firm project-based approach with clear go/no-go criteria" ("Hanteer voor de voltooiing/livegang van DWR Next een stevige projectmatige aanpak met [...] heldere go/no-go-criteria"). It comes from a BIT-advies, the statutory predecessor of ACICT, and it was followed up through a dedicated programme and not through the audit log, which is the programmatic handling of Interpretation 2. The case is from 2017, so it also shows the long-horizon behaviour of the method: link to plan, closure with rationale, and the learning loop. Table 6.6 shows the register excerpt.

Table 6.6: Register excerpt, Case D (BIT2017-DWRN-02). State shown at closure.

Field	**Value**
Identification	
Identifier	AANB-2017-002
Source	BIT (statutory predecessor of ACICT)
Source report	Definitief BIT-advies programma DWR Next (21-03-2017)
Recommendation	"Use, for the completion and go-live of DWR Next, a firm project-based approach [...] clear go/no-go criteria." ("Hanteer voor de voltooiing/livegang van DWR Next een stevige projectmatige aanpak [...] heldere go/no-go-criteria.")
Confidentiality	Internal
Intake quality	
Concrete / addressed / executable	yes / yes / yes flag off
Ownership	
Owner role	Programme manager DWR Next
Escalation path	*Stuurgroep* / DT
Cross-unit handler	Departments (*afnemers*), for their own preparation
State	
Priority (3D)	Risk & compliance: M; Strategic fit: H; Feasibility: M
Review intensity	Medium
Follow-up / impl. / impact	decided to act / delivered / effect observed
Lead time	31-12-2017 (planned go-live DWR Next)
Link to plan	DWR Next programme plan (releases, go/no-go gates)
Learning	
Ex-ante record	Expected outcome (2017): predictable delivery of DWR Next with go/no-go control. Assumptions: Windows 10/Citrix migration and the Rijnstraat 8 move as the surrounding context.
Ex-post record	Programme delivered with go/no-go control. The object of the recommendation, DWR Next, was later replaced by a next workplace generation; the recommendation itself was not re-issued.
Closure record	*implemented* — approach delivered and applied at the go/no-go moments; the later replacement of DWR Next is recorded in the ex-post record.

The case shows three points. First, the link to plan (FR8) connects the register entry to the programme plan where the work was really steered. This is exactly the connection that was missing in the programmatic pattern, where the tracking lived only in the programme working documents and was lost when the programme ended. Second, closure with rationale (FR13) records not only that the work was done, but also what happened to the object of the recommendation afterwards. A reader who opens this entry years later sees that the approach was delivered, and also that DWR Next itself was later replaced. Without the rationale and the ex-post record, only "closed" would remain. Third, the learning loop (FR10, FR11) keeps the 2017 ex-ante assumptions next to the ex-post observation, so the closure records

why the recommendation was overtaken and not only that it was overtaken. Without the register none of this history would stay, because in practice this trajectory was tracked only inside the programme.

6.1.6 Case E - research-report recommendation that decays over time (Interpretation 3)

Case E is KPM2020-01, "perform a post-calculation (nacalculatie, i.e. checking realised costs against the budget) over 2021 and check whether bundling can lead to simplification without losing the desired management information" ("Voer [...] over 2021 een nacalculatie uit en ga [...] na of bundeling [...] tot vereenvoudiging kan leiden, zonder dat daardoor gewenste managementinformatie verloren gaat"). It comes from an ADR research report ("onderzoeksrapport"), not from the audit cycle, and no programme followed it up. So it is handled in the same way as Interpretation 3 (Section 4.2): there is no fixed follow-up process, and the recommendation is easy to lose. The content here is operational and financial, not strategic. The case therefore shows that the same problem also hits a financial recommendation, and not only strategic advice. The case is added for two reasons. First, it is the only case where the recommendation decays over time, so it is the one place where the overdue-review flag (FR9) fires and the temporal decay of Section 2.6 becomes concrete. Second, the same concern returns four years later in LCM2024, in a sharper and BZ-specific form, so the case also shows re-prioritisation at review (FR7), the *superseded by* closure type (FR13), and a real ex-ante/ex-post gap (FR10, FR11). Table 6.7 shows the register excerpt at closure.

Table 6.7: Register excerpt, Case E (KPM2020-01). State shown at closure (October 2024). The State fields show the append-only history: each change is a new dated line and the older lines stay. Dates between the two report dates are illustrative.

Field	**Value**
Identification	
Identifier	AANB-2020-007
Source	ADR
Source report	Onderzoek Kostprijsmodel SSC-ICT (21-12-2020)
Recommendation	"Perform [...] a post-calculation over 2021 and check [...] whether bundling [...] can lead to simplification without losing the desired management information." ("Voer [...] over 2021 een nacalculatie uit en ga [...] na of bundeling [...] tot vereenvoudiging kan leiden, zonder dat daardoor gewenste managementinformatie verloren gaat.")
Confidentiality	Internal (DT and direct staff)
Intake quality	
Concrete / addressed / executable	yes / yes / yes (flag "decomposition needed": off)
Ownership	
Owner role	Concern Controller (Planning & Control)
Escalation path	DT (CFO)
State (append-only: each change is a new dated line)	
Follow-up status	[21-12-2020] decided to act [10-10-2024] superseded — see closure record
Implementation status	[21-12-2020] not started — not progressed; no actual costing carried out (confirmed in LCM2024)
Impact status	[21-12-2020] not yet measurable
Status reason	
Lead time	[21-12-2020] 31-12-2021
Next-review date	[21-12-2020] 31-12-2021 overdue flag raised 2022–2024; entry not reviewed
Review intensity	Medium
Internal monitoring	Concern Controller; annual P&C cycle
Link to plan	— (no implementation plan created; missing link flagged)
Priority (3D)	[21-12-2020] Risk & compliance: L; Strategic fit: M; Feasibility: M [10-10-2024] reassessed — Strategic fit: H
Priority record	[10-10-2024] weight on strategic fit raised after LCM2024 put *tariefopbouw* (tariff build-up) back on the agenda
Learning	
Ex-ante record	Expected outcome (2020): one actual costing shows whether the cost-price model can be simplified. Assumption: the model is stable enough to formalise a procedure around it.
Ex-post record	[10-10-2024] the actual costing need returned in LCM2024, sharper and BZ-specific; the 2020 assumption of a one-off analysis did not hold; no actual costing over 2023 was carried out.
Closure record	[10-10-2024] *superseded by* AANB-2024-0xx (LCM2024-05); the rationale records the successor entry and that the same concern returned in a broader, BZ-specific form.

The case shows four points. First, it is the only case where the overdue-review flag (FR9) fires on a real recommendation: the next-review date was set for 2021, the review did not happen, and the entry stayed unreviewed for years. This is the temporal decay of Section 2.6, now on a concrete recommendation. Second, the priority is reassessed at the 2024 moment (FR7): in the other cases the priority is scored once at intake, while here the priority record holds a changed weighting after the topic returned. Third, the entry closes with the *superseded by* type (FR13) and a real identifier chain to the LCM2024 successor entry, it is the only case in the set that uses a closure type other than *implemented*. Fourth, the learning loop (FR10, FR11) holds a real ex-ante/ex-post gap: the 2020 expectation that one post-calculation would settle the simplification question did not hold, and the need returned four years later in a sharper form. The append-only State fields keep this four-year history readable from one entry (FR1).

6.1.7 What the demonstration shows

Read together, the five cases show that the artefact can hold recommendations from each process interpretation, and that the distinguishing features of the design appear in the register fields as intended. Table 6.8 maps every requirement to where it is addressed: in the demonstration cases of this section, in the stakeholder workshop of Section 7, or not yet assessed in this thesis.

Two limits of the demonstration should be named. First, it is a study of a validation model: the researcher enters the cases from the content of public reports, so the register values are a plausible reconstruction and not operational data, and several fields are marked "to be filled at intake". In the terms of Wieringa, the demonstration addresses the effect and sensitivity questions of validation, that is, whether the artefact and the context together can represent recommendations that differ in source, level and theme. It does not yet address the requirements-satisfaction question, that is, whether the method is seen as usable and would be used [37]. That question needs the people who would use it, and it is the goal of the expert-opinion validation in the next section. Second, and for the same reason, the demonstration shows representational capability (the method can hold the cases) and not realised effect (the method improves the follow-up in practice). This is in line with the perception-based validation chosen in Section 3.4.4. The next section therefore takes Case B into a stakeholder workshop, where the participants apply the register to the same recommendation themselves and judge its usability, usefulness and intention to use.

Table 6.8 shows which requirements each case covers. Together the five cases cover all thirteen functional requirements, plus confidentiality (PC4), the effect chain frame (PC5), and low administrative burden (NFR2). The other requirements are not covered by a case: PC1 is context, PC2, PC3, NFR1, NFR3, NFR5 and NFR7 are design-level, NFR4 follows from role-based ownership (FR12), and NFR6 is a perception requirement that is taken to the workshop.

Table 6.8: Requirements covered by each demonstration case. Cases: A = AR2023-DFA-04, B = AR2024-IT-02, C = ACICT2025-01, D = BIT2017-DWRN-02, E = KPM2020-01.

Requirement	**A**	**B**	**C**	**D**	**E**
Project constraints					
PC4 Confidentiality			✓		
PC5 Effect chain frame	✓	✓	✓	✓	✓
Functional requirements					
FR1 Register with history	✓	✓			✓
FR2 Minimum data	✓	✓	✓	✓	✓
FR3 Status model		✓			
FR4 Internal/external dep.		✓	✓		
FR5 Three-layer follow-up		✓		✓	✓
FR6 Intake quality check		✓			
FR7 Prioritisation logic					✓
FR8 Mapping to P&C				✓	✓
FR9 Review scaled to risk			✓		✓
FR10 Learning loop				✓	✓
FR11 Ex ante / ex post				✓	✓
FR12 Ownership	✓	✓	✓	✓	✓
FR13 Closure with rationale	✓			✓	✓
Non-functional requirements					
NFR2 Low admin. burden	✓				
Not covered by any case (see text)					
PC1 Dutch central-gov context			—		
PC2 No parallel system			—		
PC3 Compatible with org. change			—		

Continued on next page

Table 6.8 continued

Requirement	A	B	C	D	E
NFR1 Embedding in existing moments			—		
NFR3 Internal transparency			—		
NFR4 Resilience to turnover			—		
NFR5 Capacity awareness			—		
NFR6 Learning-culture fit			—		
NFR7 Transferability			—		

Chapter 7

Validation

The demonstration in Section 6.1 shows that the register can hold the recommendations that SSC-ICT receives, and that it still works for cases that differ in source, level, and theme. In the terms of Wieringa, this answers two of the treatment-validation questions: what the method does with a case (the effect question), and whether it keeps working when the case changes (the sensitivity question) [37]. The demonstration does not show whether the result meets the requirements as the people who would use the method see it. Wieringa calls this the requirements-satisfaction question. This section answers it with a stakeholder workshop, of which the workshop outline can be found in Appendix F.

The workshop is an expert-opinion validation. Wieringa describes expert opinion as the simplest way to validate an artefact. The design is given to people who know the context. They imagine how the method would work in that context and say what they think the effect would be [37]. So the judgement comes from how the experts imagine the method would work, not from real use. This has two consequences for how we read the results. First, the goal is not to collect all opinions, but to see how realistic users think the method would work. Second, Wieringa warns that positive opinions can be polite answers, or can mean that a person cannot imagine why the method would fail, while negative opinions are more useful, because they show problems in practice that the designer did not think of [37]. For this reason we give more weight to the critical points than to the positive scores, and in the plenary part we ask why participants judged the method as they did.

7.1 Why the MEM and not a software-usability test

The artefact is a method and not a software product, so the validation cannot use a normal usability test for software. Following Moody, a method cannot be true or false, only effective or not, and this can only be shown through use [26]. The Method Evaluation Model (MEM) joins two parts of method success: the actual efficacy, whether the method improves the task, and the adoption in practice, whether people want to use it [26]. The actual efficacy cannot be measured before real use. The MEM therefore measures the perceptions that, in the model, lead to adoption: perceived ease of use, perceived usefulness, and intention to use [26]. The validation in this thesis is done before real use, so these perceptions are the right level to measure, as argued in Section 3.4.4. The MEM is built on the Technology Acceptance Model from Davis [12].

7.2 Workshop design

The workshop has one goal: to validate if the method works in practice. This means two things at the same time. The first is usability: can people who would use the method in their work fill in the register? This is tested by letting participants fill in the register themselves on a real case. The second is acceptance: would they want to use the method? This is tested with the MEM questionnaire after the exercise and the open discussion.

The workshop uses one case, AR2024-IT-02, the same abstract and chain-dependent recommendation that the demonstration uses as Case B (Section 6.1.3). For the workshop the case is extended with a short fictional storyline that runs from a moment of intake (May 2025) to the half-yearly review

moment (December 2025). The participant-facing materials use plain Dutch and do not use the T0/T1 labels from the design chapters. In the first step the groups register the recommendation from zero: identification, intake-quality check, owner, priority, and the ex-ante expectation. In the second step the groups get a short description of what happened between May and December 2025. In this storyline the recommendation is decomposed into an internal slice (an SSC-ICT set of standards plus a testing cycle) and an external slice (alignment with CIO Rijk on what becomes government-wide), the CTO is confirmed as owner, the internal work goes on, and the external slice is blocked while it waits on the government-wide framework. The groups then update the register to the December review, with the append-only rule.

Participants worked in groups of two or three on one laptop per group, and filled in the register together from their own expertise. The session lasted 60 minutes: a short introduction with a presentation of the method, the two-step exercise, the questionnaire, and a plenary discussion. During the session no more information was given than the case description. When a group could not go on without help, a short neutral clarification was provided. One observer, a fellow peer from Leiden University, followed the groups and wrote down which fields were filled in fast, which fields gave discussion, which questions were asked, and short literal quotes. Each participant filled in the questionnaire alone, before the plenary discussion, so that the answers were not influenced by the opinions of others. This follows the post-task design of Moody and matches the one-group post-test design of his field experiment with practitioners [26]. The session was recorded for qualitative analysis. The recording and the observation notes are the main evidence and the questionnaire is supporting evidence.

A test run was held on 31 May 2026 with two peers, to check the materials and the timing. The main change after this was to put all participant-facing terminology in Dutch, instead of leaving design terms in English, so that the words match the language the participants use in their daily work. The case and the register schema were not changed.

The questionnaire is the MEM post-task survey of Moody, adapted to the register. It has sixteen items: six for perceived ease of use, eight for perceived usefulness, and two for intention to use, the same division as Moody [26]. Each item uses a five-point Likert scale. Half of the items are negatively formulated and the items are placed in random order, so that respondents keep attention for each statement. The wording is changed from the use of a system to the use of the register, and the perceived-usefulness items are written in terms of the goals of the method (keeping recommendations visible, making ownership explicit, supporting steering). Three open questions follow the scale: what the register would change in the participant's work, what their biggest doubt is, and what is still missing for the method to be usable in their work. The full form is in Appendix C.

The session was offered to 37 people across Finance and Control, Internal Audit and Advisory, and Organisation Strategy and Development. Seven people came, among them the Head of Finance and Control and the Head of Internal Audit and Advisory. The seven formed three groups of two, two, and three people. The questionnaire was filled in alone, so there are seven questionnaires. The register was filled in per group, so there are three register entries. The people who came are a self-chosen expert group, not a sample for statistical generalisation. This fits the expert-opinion validation above. The two department heads are relevant, because the method puts ownership and steering at a senior level, and these two roles stand for that level and for the audit follow-up.

7.3 Analysis approach

The questionnaire is analysed descriptively. The negatively formulated items are first recoded, so that a high score always means a positive judgement, by replacing the score s with $6 - s$. Per construct, the recoded item scores are averaged per respondent and over all respondents. Following Moody, the direction of a construct is read against the neutral point of the scale (3): a mean above 3 is a positive perception, a mean below 3 a negative one [26]. Moody tests the construct means against this midpoint with one-sample t-tests. With seven respondents the sample is too small for these tests, and too small for the factor analysis or reliability analysis that Moody runs on his larger samples [26] [1]. The construct

[1] Moody removed his item Q13, a perceived-usefulness item, from the analysis in both of his experiments, because it failed the validity test. He describes this as a systematic flaw in that item [26]. With seven respondents we cannot run such a test, so we keep all sixteen items. Our adapted item 13 shows no deviating pattern: its score (4.14) is in line with the other perceived-usefulness items.

means therefore only show the direction, not a final score. The open questions and the observation are used to explain the pattern.

7.4 Results

7.4.1 Questionnaire

Table 7.1 reports the recoded construct means per respondent and over all respondents. Table 7.2 reports the recoded mean per item.

Table 7.1: MEM construct scores per respondent, recoded so that high = positive judgement. Neutral = 3. SD is the sample standard deviation over the seven respondents.

Construct	R1	R2	R3	R4	R5	R6	R7	Mean	SD
Perceived ease of use	2.83	3.33	2.50	4.00	4.33	3.50	3.00	3.36	0.65
Perceived usefulness	3.38	4.62	3.62	4.25	4.25	3.62	3.88	3.95	0.44
Intention to use	3.50	4.50	3.50	4.50	4.50	3.50	3.50	3.93	0.53
All items	3.19	4.12	3.19	4.19	4.31	3.56	3.50	3.72	0.48

The pattern is the same for most respondents. Perceived usefulness and intention to use are above the neutral midpoint for all seven respondents, with means of 3.95 and 3.93. No respondent had an intention to use score lower than 3.50. The average perceived ease of use is above the midpoint (3.36), but it is the lowest and also the most unstable dimension: two of the seven respondents (R1 and R3) scored it below neutral. In Moody's words, the participants think the register is useful and they want to use it, but the part about ease of use still needs to get better.

Table 7.2: MEM item scores, recoded so that high = positive judgement. Direction shows if the item was positive (+) or negative (−) before recoding. Min–Max and SD (sample) are over the seven respondents.

#	Item (short)	Construct	Mean	Min–Max	SD
5	Easier to establish who is responsible for what	PU	4.43	4–5	0.53
7	Overall, found the register useful	PU	4.29	4–5	0.49
10	Would not use the register (−)	ITU	4.29	4–5	0.49
3	Harder to keep overview (−)	PU	4.14	4–5	0.38
13	Easier not to lose sight of recommendations	PU	4.14	4–5	0.38
8	Harder to follow over longer time (−)	PU	4.00	2–5	1.00
1	Procedure complex to follow (−)	PEOU	3.57	2–4	0.79
11	Rules of the register clear and understandable	PEOU	3.57	2–4	0.79
12	No effective solution (−)	PU	3.57	2–5	0.98
15	Improvement over current follow-up	PU	3.57	3–5	0.79
16	Intend to use it in the future	ITU	3.57	3–5	0.79
2	Would reduce the effort of follow-up	PU	3.43	3–4	0.53
4	Overall hard to use (−)	PEOU	3.43	2–5	1.27
6	Easy to learn	PEOU	3.29	1–5	1.38
14	Not sure I could apply it on my own (−)	PEOU	3.29	2–5	1.11
9	Hard to apply to this case (−)	PEOU	3.00	2–4	0.82

At item level, the highest scores are about ownership and overview: making clear who is responsible (item 5, 4.43), the overall usefulness (item 7, 4.29), the low rejection of future use (item 10, 4.29), and keeping recommendations in view (items 3 and 13, both 4.14). The lowest scores are the ease-of-use items about applying the method to this case (item 9, 3.00), learning it (item 6, 3.29), and applying it alone (item 14, 3.29). One respondent (R6) scored "easy to learn" at the lowest possible level, which pulls the metric down. This matches the open answer of the same respondent, who asked for more explanation of the classification and of when something is urgent.

7.4.2 Observation

All three groups read the case carefully and finished both steps in time. The identification fields and the recommendation text were filled in fast. The field that gave the most discussion was "status reason" ("status-reden"): Group 2 asked what it meant in step 1, and Group 3 asked the same in step 2. A short neutral clarification was given each time. Group 1 also asked what "executability" ("uitvoerbaarheid") and the "cross-organisational owner" ("cross-organisatorische eigenaar") field meant, but went on without problems after a short explanation. These fields are the fields that the method adds on top of a plain register. The questions came on these added fields and not on the basic ones.

Group 2 recognised the recommendation and knew the process behind it, so they filled in the register from their own experience, not as a hypothetical case. Their entries follow how the recommendation was handled in the organisation. This makes their input close to real practice. It also means the groups did not all work from the same perspective.

The intake-quality check gave a real discussion in step 1. In the source report the recommendation is abstract, but the case description says for whom it is meant. One group therefore judged it as concrete enough and addressed, and did not flag it for decomposition, while the design intent is that it fails the check (Case B, Section 6.1.3). So the check did its work, because it put the question on the table, but the result depends on how the reader reads the recommendation (Figure 7.1).

When the three layers (follow-up, implementation, impact) were explained, a participant linked them to "design, existence, operation" ("opzet, bestaan en werking"), the terms the audit people already use and the same terms came back in the plenary discussion. So the three-layer model was recognised in terms the organisation already knows. At the same time, it again becomes clear that the participants choose to formulate things in their own words.

One point came back in the exercise and in the plenary discussion. Group 3 said that the register reminds them of "Plato", an older registration system at SSC-ICT: the set-up was good and gave a full overview, but it was not kept up to date. The group said the same in the plenary: the system is good, but people will have to keep it up to date. This is the main acceptance risk the participants named. The method is seen as good, but it only works if people keep updating it.

Groep 1 — stap 1 (intake)

Veld	Type	Waarde
Intake-kwaliteit		
Concreet? *Specifiek genoeg om op te handelen? (ja / nee)*	V	Ja
Geadresseerd? *Is er een rol of functie aan te wijzen? (ja / nee)*	V	Ja
Uitvoerbaar? *Realistisch gegeven afhankelijkheden? (ja / nee)*	V	Ja
Vlag 'decompositie nodig' *Zet aan als één of meer attributen 'nee' zijn*	V	

Groep 2 — stap 1 (intake)
afwijkend: niet concreet → decompositie-vlag aangezet

Veld	Type	Waarde
Intake-kwaliteit		
Concreet? *Specifiek genoeg om op te handelen? (ja / nee)*	V	Nee
Geadresseerd? *Is er een rol of functie aan te wijzen? (ja / nee)*	V	Ja
Uitvoerbaar? *Realistisch gegeven afhankelijkheden? (ja / nee)*	V	Ja
Vlag 'decompositie nodig' *Zet aan als één of meer attributen 'nee' zijn*	V	Ja

Groep 3 — stap 1 (intake)

Veld	Type	Waarde
Intake-kwaliteit		
Concreet? *Specifiek genoeg om op te handelen? (ja / nee)*	V	ja
Geadresseerd? *Is er een rol of functie aan te wijzen? (ja / nee)*	V	ja
Uitvoerbaar? *Realistisch gegeven afhankelijkheden? (ja / nee)*	V	ja
Vlag 'decompositie nodig' *Zet aan als één of meer attributen 'nee' zijn*	V	

Figure 7.1: Intake quality check in step 1, as the three groups filled in the register. Groups 1 and 3 judged the recommendation concrete and left the decomposition flag empty. Group 2 judged it not concrete and set the flag (in red), which is the design intent for this case.

7.4.3 Open questions

The open answers fall into four points. First, the participants in control, management, and audit roles named a clear value: a shared overview. R1 called it "a common source of truth for monitoring and follow-up" ("een gemeenschappelijke bron van waarheid voor monitoren & follow-up"), R3 said it makes it "clearer for control and management what the follow-up status is" ("overzichtelijker voor control + management wat de status van opvolging is"), and R6 said "easier to monitor" ("beter te monitoren"). Second, some participants who do not work with recommendations in their daily work said the register would change little for them (R2: "not much" ["niet veel"]; R4: "personally not" ["persoonlijk niet"]; R5: "I do not deal with recommendations"

[“ik houd mezelf niet bezig met aanbevelingen”]; R7:
“no” [“nee”]). This says something about the group, not
about the method. The usefulness scores of these participants carry less weight.

Third, the biggest doubt was the administrative burden, together with the question if owners would keep the register up to date. R1 and R4 named the administrative burden directly; R1 added that the method must become more user-friendly. R2 named complexity. R3 and R6 named the risk in behaviour: R3 wrote that “keeping the status updated by action holders/owners. Experience teaches that this does not happen” (“het bijhouden van de status door actiehouders/eigenaren . . . ervaring leert dat dat niet gebeurt”), and R6 doubted “that people feel responsibility and do something with it” (“dat mensen verantwoording voelen en er iets mee doen”). This is the same risk as the “Plato” comparison. Fourth, for what is still missing, R1 asked to connect to the terminology used inside the organisation, R6 asked for an explanation of the classification and of when something is urgent, and R5 wrote about where the register is stored, the danger of overwriting old data, and re-prioritising with hindsight. The append-only rule is made exactly to prevent the worry of R5, so this rule was not fully understood and must be explained more clearly. It was also explained that this methodology must be seen as a way of working that is tool independent, meaning the organisation may pick a tool that is most fitting for this way of working and storing the data.

7.5 Interpretation against the requirements

Read against the requirements from Chapter 4, the validation gives a mixed result.

On acceptance, usefulness and intention are positive. Perceived usefulness and intention to use are above neutral for all seven respondents. In the open answers, the value is the overview and the steering for control and management. No participant saw the register as a control or punishment tool. R6 saw it as a way to monitor and to give feedback. Within this small sample, this supports the fit criterion of NFR6 (must not be seen as a control or compliance tool). It also supports the main argument that the method fills the ownership and overview gaps from Chapter 4 (FR12, NFR3).

The weak points are ease of use and the administrative burden. Perceived ease of use is the lowest value, and the administrative burden is the most common doubt in the open answers. This concerns requirement NFR2 (low administrative burden). The validation does not let us say that NFR2 is met. It shows that the method is seen as useful, but that the felt effort must go down before people would use it. The fields the method adds for steering, mostly status reason and the cross-organisational owner, are the fields that gave the questions in the exercise.

The doubt about keeping the register up to date is the most useful finding considering Wieringa, because it is a negative prediction that comes from the participants’ own experience. It corresponds to NFR4 (the follow-up must not depend on one specific person) and on the strategic follow-up pattern in interpretation 3. The participants see that the register takes away the dependence on the memory of one person. But they say, correctly, it does not take away the dependence on someone who actually updates it. The parts of the design that are meant to show when an entry is not updated, namely internal transparency (NFR3), the capacity flag (NFR5) and the overdue-review flag (FR9), makes this risk smaller, but cannot remove it. We read this as a real limit of the method, not as a design mistake, as updating the method stays a condition for it to work.

The participants did not only name the risk, they also named a possible remedy. When it was asked how to make sure something is actually done with the register, a participant from group 1 said: “commitment from the top, so that updating the register is expected and not optional”. They gave a recent example from their own work, where a periodic report stalls because some contributors do not deliver their part in time. this points back to the design choice to put ownership and the review moments at a senior decision body (FR12, NFR1). The register needs a body that asks for the overview, or the updating will not happen.

The exercise also showed that the harder parts of the method can be used in practise. The groups did the decomposition step (FR6), split the internal and external slice of a chain-dependent recommendation (FR4) and used the confidentiality field (PC4) without big problems. The append-only rule (FR1, status model) was the part that one participant did not fully understand (Figure 7.2). This is a point of explanation, not a gap in the design.

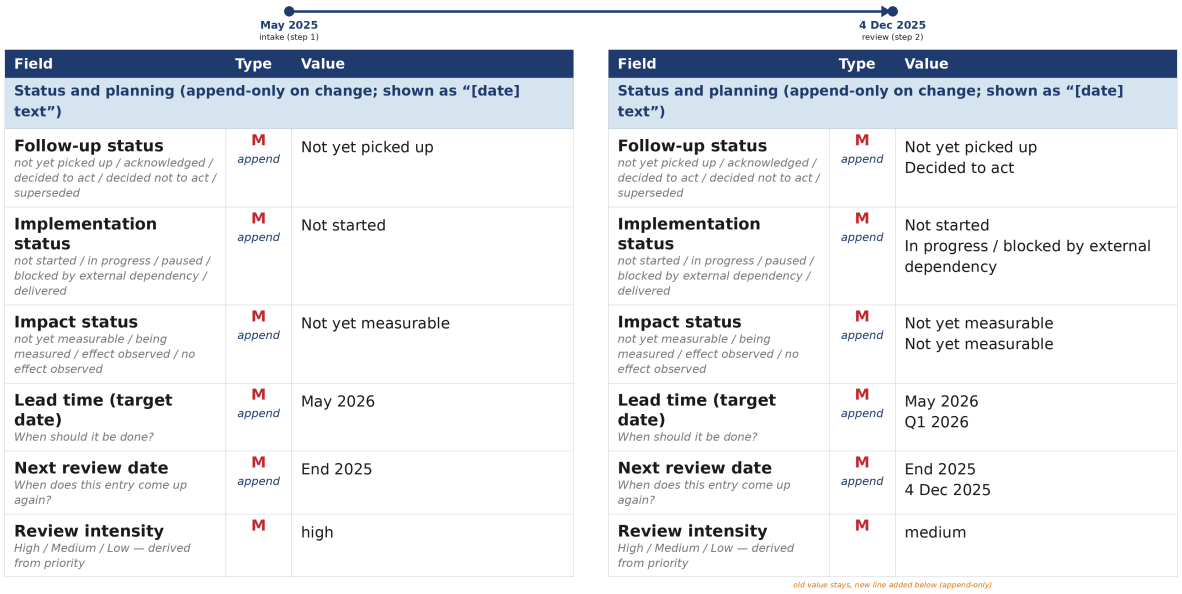


Figure 7.2: The same register entry of group 3 at intake (step 1, May 2025) and at the review (step 2, December 2025). On a change the old value stays and the new value is added on a new line (append-only); the impact status stays the same.

Last, the terminology point is concrete and easy to fix. The rest run put the materials in Dutch, but two paricipants still asked to match the words SSC-ICT uses inside the organisation. During the exercise, this pair was at times translating the terms for each other, and they named decomposistion as a word that is not used n the organisation. So this can be changed easily.

One group also suggested to link the method to a tool that the organisation has recently acquired, which already has the ability to accommodate the fields and flows that are required for the register. This fits PC2, stating the artefact is a method, not a tool, and is meant to be put on top of tooling that is already in use. The suggestion shows that this design choice is understood and seen as useful by the participants.

Table 7.3 shows which requirements the workshop tested and the result. The participants judged ownership (FR12), internal transparency (NFR3) and learning-culture fit (NFR6) as positive, the administrative burden (NFR2) as not met, and several added fields (FR1, FR3, FR6, FR7) and the effect chain frame (PC5, FR5) as only partly clear. Some requirements were not tested in this session: closure (FR13), link to plan (FR8), the overdue-review flag (FR9) and the capacity flag (NFR5) did not come up in the exercise, and PC1, PC3 and NFR7 are out of reach of a single workshop.

Table 7.3: Requirements tested in the validation workshop (case AR2024-IT-02, $n = 7$, three groups), with the result.

Req.	Result	Rationale
Positive		
PC2	Positive	A group proposed putting the register on a tool the organisation already has.
PC4	Positive	Confidentiality field used without problems.
FR2	Positive	Identification and base fields filled in fast.
FR4	Positive	Internal/external slice split without problems.
FR10/11	Positive	Ex-ante recorded at intake, ex-post at the review step.
FR12	Positive	Top-scored value; clearest who is responsible (item 5, 4.43).
NFR1	Positive	A participant pointed to "commitment from the top" as what makes updating happen.
NFR3	Positive	Participants valued a shared overview most (R1, R3, R6).
NFR6	Positive	No participant saw the register as a control or punishment tool.
Partial		
PC5, FR5	Partial	Participants matched the three layers to the audit terms "design, existence, operation" ("opzet, bestaan en werking"). The idea is familiar, but we did not test whether they applied the layers correctly, and the two do not map one-to-one.
FR1	Partial	Append-only used, but one participant did not fully grasp it.
FR3	Partial	Status reason was the field that gave the most discussion.
FR6	Partial	The check raised the question, but one group did not flag what the design intends; the result depends on the reader.
FR7	Partial	Priority set at intake; one participant worried about re-prioritising with hindsight.
NFR4	Partial	Removes dependence on one person's memory, not on someone updating it.
Negative		
NFR2	Negative	Lowest construct (PEOU 3.36); administrative burden the most common doubt; not shown to be met.
Not tested in the validation		
FR8	—	Link to plan not exercised in the session.
FR9	—	We use the overdue-review flag for the not-updated risk, but the workshop did not exercise it.
FR13	—	Closure was not part of the exercise.
NFR5	—	We use the capacity flag for overloaded owners, but the workshop did not exercise it.
PC1	—	Context constraint.
PC3	—	Not relevant to a single session.
NFR7	—	Single organisation; not tested.

7.6 Limits of the validation

Three limits must be kept in mind when reading the results. First, expert opinion is the weakest of the validation methods in Wieringa as positive opinions can be polite answers, or can mean that the participants could not imagine why the method would fail [37]. We make this risk smaller by reading the scores together with the open questions and the plenary discussion and by giving more weight to critical points mentioned before than to the positive scores.

Second, the sample is small and self-chosen, as well as the case. Seven of the 37 invited people came, and the people who came may be the people who care most about this topic. The questionnaire is filled in per person ($n = 7$), but the register is filled in per group ($n = 3$), so the register evidence is at group level. Some of the people who came do not work with recommendations in their daily work, which explains the "not much change for me" answers and makes their usefulness scores count for less. At the same time, two department heads came, which makes the panel stronger for assessing the ownership and audit follow-up questions and the open discussion. The numbers in the Table 7.1 only show the direction. With seven respondents we cannot do the factor analysis, or the one-sample t-test that Moody does, and we do not claim adoption, only a positive perception and intention on the constructed case. One participant also found the negatively-worded items in the questionnaire difficult, because of the double negation. Moody uses negated items on purpose, to keep the respondent's attention, but with the reverse-coded items, this can add noise to the answers, which is a further reason to read the scores only as a direction.

Third, the limits of the validation model are the same as in Section 6.1. The case is built from public reports, with an added storyline, and it is used once, before any real use. So, the workshop answers the requirements satisfaction questions only at the level of the expected use, not real use and effect. A test of real effect, for example through action research in the sense of Wieringa, where the method is used on a real recommendation inside the organisation, is future work.

Chapter 8

Discussion

This chapter brings the results of the thesis together. The chapters before this one each describe one step of the design science process: the current follow-up at SSC-ICT and the requirements (Chapter 4), the artefact (Chapter 5), the demonstration (Chapter 6), and the validation with stakeholders (Chapter 7). Here we discuss these results next to each other. We first give the key findings. After that we take the main risk that came out of the validation, namely the discipline of keeping the register up to date, and we place it in the context of the organisation with a deepening interview. Then we discuss what the results mean for practice and what they add to the theory. The chapter ends with the limitations.

8.1 Key findings

The current follow-up at SSC-ICT works in two ways. Findings from the financial and IT audit cycle are followed up relatively well. The finding is put in a log, classified, given an owner, and only closed after an independent check by IAA. This part keeps running with little attention from management. For the other recommendations, the advisory and policy ones, there is no such fixed process at the back-end. These recommendations are handed over "over the fence" ("over de schutting"), without a clear scope, a clear result, or a review moment, and can end up "in the drawer" ("in de la"). The problem that comes back the most is ownership. When it is not clear who is the owner, the discussion about who is responsible keeps coming back and the work does not move. A clear owner is needed, but it is not enough by itself. We saw that even action lists with deadlines and reminders did not lead to follow-up. One more point: follow-up is not the same as implementation. A written decision to not act on a recommendation, with the reason, is also a good form of follow-up.

So the problem is not a missing tool. What is missing is a way of working that keeps a recommendation visible and steerable over time. The artefact gives this as a method: a register for recommendations with a field schema. The design rests on three sources: the interviews inside SSC-ICT, the interviews with the producers of recommendations, and the literature. The heart of the design is the status model with three layers from the effect chain (follow-up, implementation, impact). In this way "open" and "closed" do not hide anymore in which layer a recommendation is. Next to this status, the register has an intake check that looks at the quality of a recommendation before it gets an owner, an owner and a review moment at a senior body that already meets, a closure with a clear reason and four closure types, and a history that you can only add to, with an ex-ante and an ex-post note. The producer-side interviews gave two limits for this design. The register must stay small, and it must hang on the planning-and-control cycle that is already there, instead of running as a parallel track.

The demonstration shows that the register can hold the recommendations that SSC-ICT really gets. We made an inventory of 126 recommendations from nine public reports and put five cases through the register, one for each process interpretation. The fields that make the design different come back in the register as the design says: the decomposition flag, the split between an internal and an external part, the overdue-review flag, the closure by "superseded by", and the gap between the ex-ante and the ex-post note. This only shows that the method and the context together can hold recommendations that differ in source, level and theme. It does not yet show that the method works in real use.

The validation gives a mixed but clear picture. The stakeholders find the method useful and say they would use it. Perceived usefulness and intention to use are above the neutral midpoint for all seven respondents (3.95 and 3.93). The value they name is a shared overview and a clear owner, for control and for management. No one saw the register as a control or punishment tool. The weak point is the ease of use and the administrative burden. Perceived ease of use is the lowest of the three (3.36), and for two respondents it is below neutral. So we cannot say that the low-burden requirement (NFR2) is met. The biggest risk the participants name is the discipline of keeping the register up to date. They point to the older "Plato" register: well set up, but not kept up to date. As a solution they name commitment from the top. So the method is strong on the parts it is built for, the ownership and the steering at a senior level, and it is weak there where the organisation is already weak. This risk of keeping the register up to date is the point that the rest of this chapter works out further.

8.2 Adoption in context: a deepening interview

The validation in Chapter 7 left one question open. The participants judged the method as useful and stated the intention to use it, but named the update discipline as the main risk, with the earlier Plato register as the precedent of a well-designed instrument that was not kept up to date. This is a question about the organisation, not about the method. The open question is, whether the organisation in the middle of an Agile transformation has the maturity and adoption capacity to keep an instrument like this alive. To place the workshop finding in that context, one deepening interview was held after the workshop, with an external consultant involved in the Agile transformation at SSC-ICT (12-06-2026). The interview was semi-structured and recorded with consent.

This interview is interpretive, not confirmatory. It does not change the validation result. It describes the situation in which the adoption risk from the workshop is valid. Two things shape how we interpret it. First, with one respondent, there is no pattern to establish, so the account is used to deepen and contextualise, not to generalise. Second, the respondent is an external hire whose work is the success of the transformation. It gives an outside view on the organisation, but also an interest in it, which we keep in mind.

The ownership gap is recognised, and explained as culture. The respondent recognised the workshop finding directly. His team is running a process analysis of the Agile delivery process at SSC-ICT, and ownership came out as one of two core bottlenecks across the organisation. The mechanism he described matches the receiver-side material in Section 4.1. Work is handed over and then treated as no longer one's own, and no one owns the whole trajectory from start to end. He explained this as culture, attitude, and behaviour, reinforced by the structure of three business units that each run their own way of working and pursue their own goals. He called these units silos, with further silos inside them, which is the same pattern as theme T7. The "distance between wanting and doing" that motivated this interview fits here too. People see that change is needed, but they locate the problem, and the solution, somewhere other than with themselves.

Maturity is judged low, and tied to the same behaviour. Asked how mature the organisation is in working with fixed methods and instruments, the respondent judged this maturity as low. The reason connects back to ownership. Systems such as TopDesk are used to push work away instead of to coordinate, and email is used to shift responsibility ("I have sent it, so it is no longer mine"). He contrasted this with the SAFe principle of people over process, and noted that at SSC-ICT the process tends to lead instead. This matches the producer-side warnings against a register that becomes a "bookkeeping exercise" or a "chore". In a low-maturity setting, an instrument risks being filled to satisfy the process rather than to steer. He extended the producer-side point on defensive behaviour as well. Individuals and teams were described as engaged and full of good ideas, but once an issue rises into the management layer, people turn defensive to avoid owning the problem. This is the same mechanism that Lemmens named as "defensive routines" on the producer side (Section 4.3), now reported from inside the receiving organisation.

The method fits the SAFe structure, as an enabler in the prioritisation cycle. On fit with the transformation, the respondent was positive and concrete. He framed an audit or advisory finding that requires action as an "enabler" in SAFe terms. It does not deliver value itself, but it is needed to deliver value well, like architecture or security work. The artefact around it he described as the way of working that makes such findings visible and steerable. He saw no conflict with available capacity. In SAFe, team capacity is fixed, so a new finding does not start a capacity discussion but a prioritisation discussion,

taken each quarter at team, train, and portfolio level. A finding would enter that existing quarterly prioritisation instead of a separate track. This supports PC2 (no parallel track), NFR1 (embedding in existing decision moments), and FR7 (explicit prioritisation). It also puts NFR5 in a different light. The work concentration that NFR5 is meant to surface does not appear as a separate capacity signal in SAFe, but as a prioritisation question inside the train. On tooling, the respondent confirmed the position of PC2 and FR8. Concrete work that follows from a finding belongs in Jira, where all work is planned, while the register itself may live in whatever instrument the owner keeps, as long as the work has a named owner and stays traceable. This matches the link-to-plan field in the schema (Section 5.1.1).

Adoption depends on ownership at the right level, and on felt urgency. Asked what determines whether a new instrument lands or fades, the respondent named two things. People must believe it solves a real problem, and they must see it work in small steps. Culture he described as a slow bottleneck that cannot be switched on at a fixed moment. Asked how the method would fail, he answered: through ownership, and through people not understanding why something is blocked or urgent. Asked the reverse, how to make it land, he gave the same answer the other way around. There must be an explicit owner who treats the item as theirs and as urgent, positioned inside the portfolio near the framework-setting disciplines, for example the CISO for security findings. The reasoning was that the person who feels the pain when a finding is ignored should also be the person who carries it. This points back to the design choice to place ownership and review at a senior decision body (FR12, NFR1), and it gives that choice an organisational reason. He also raised one point the protocol had not covered. There is a structural quality problem at SSC-ICT, where what counts as "ready" or "good enough" is judged differently across the process. He linked this to the SAFe notions of a "definition of ready" and a "definition of done". The intake gate in FR6 is related but not the same. FR6 checks the incoming recommendation before it gets an owner, on concreteness, addressability, and executability. A definition of ready checks a unit of work before a team starts it, and a definition of done checks the result after delivery. FR6 therefore guards the input, but the method does not set an explicit standard for when the resulting work is complete. The closure types in FR13 record *that* an item is closed and why, but not whether the delivered work met a quality bar. This is a gap the respondent's point exposes, and we return to it in the limitations (Section 8.4).

What this means for the artefact. The interview is in line with the producer-side and receiver-side findings. It confirms the design choices that carry most of the adoption risk. Follow-up is anchored in the existing quarterly prioritisation rather than a new meeting (NFR1, PC2) and ownership sits at a level where the consequence of inaction is felt (FR12). It also confirms the limit of the artefact, already stated in NFR6. The register can make the absence of follow-up visible, but the willingness to take ownership and act is a cultural condition that the method cannot create. In a low-maturity, silo-structured organisation, that condition is exactly where the risk sits. This is the same conclusion the workshop participants reached from their side.

8.3 Practical implications

A risk comes back across the producer-side warnings, the workshop, and the deepening interview: the register can become just an administrative exersice. ACICT warned against a "bookkeeping exercise" ("boekhoudkundige exercitie") CIO-Rijk against a "chore" ("corvée klusje"), and the participants pointed to Plato, a register that was not kept up to date. The design answers this on one side, by keeping the mandatory fields to the steering minimum **NFR2** and linking updates to meetings that already take place **NFR1**. There is also another side, which this thesis does not test but notes for practice: make the update itself cost less effort.

Reducing administrative burden is one of the main reasons why governments adopt AI [38], and recent large language models can turn free text into structured records [10]. The meetings to which reviews are linked (**NFR1**) already produce written notes. An AI assistant could read those notes and propose a draft update for the entries involved, for example a status change with its reason, or a closure rationale. The owner role then confirms or corrects the draft, instead of typing it themselves. The ownership decision and the append-only history stay the same. The register then becomes something that is updated as a by-product of a meeting that happens anyway, rather than a separate task someone has to remember.

This is not a solution and it has its own risks. The notes have to be reliable, and a wrong machine-proposed update is a new error that the confirmation step has to catch. Whether this lowers the felt burden enough to change the behaviour the participants doubted, is a question for future work.

8.4 Limitations

Some limits are already named where they come up, in the demonstration (Chapter 6) and in the validation (Chapter 7). This section names them shortly and adds the limitations that were not yet discussed.

The method is studied before real use, in line with Wieringa on validation, whereas evaluation would look at the artefact during or after real use [37]. The demonstration puts historical cases in the register, and the workshop lets stakeholders fill the register for one constructed case. It is an expert-opinion validation, a recognised validation method where the experts judge how the method would work. In both, no real recommendation was followed up with the method inside the organisation. So the thesis shows that the register can hold the recommendations and that stakeholders expect it to be useful. It does not show that the follow-up gets better in practice. A test of real effect, for example through action research [37], is future work.

The internal information at SSC-ICT is confidential, so internal documents are not used in the thesis and the demonstration is built on public reports only (Section 3.4.5). The context that was needed to fill the register came from the company supervisors and is given at the level of the register fields. The reader can therefore not check the register values against the source documents. The values show how the method would work, but they are a reconstruction and not data from real use.

The questionnaire is read as a direction, not as a statistical proof. With seven respondents the sample is too small for the t-tests and the factor analysis that Moody runs on his larger samples [26]. This still fits the MEM, which measures perceptions before real use, and it fits Moody himself, who replaced a factor analysis by a simpler inter-item check when his own sample was too small [26]. We read the construct means only as a direction against the neutral point. The signal is still clear: usefulness and intention to use are positive for all seven respondents, read together with the open questions and the observation (Section 7). So the result carries no statistical weight, but it supports the perception that the method is useful and would be used.

The method checks the input, but not the result. The intake gate (FR6) checks the incoming recommendation before it gets an owner. The closure types (FR13) record that an item is closed and why. Neither one checks if the delivered work is good enough. The method follows the trajectory of a recommendation and leaves the quality of the work to the delivery process that is already there. So the method can show that the follow-up happened, but not that it was done well. A next version could add a quality check, a "definition of done", to the closure step.

The main limit is one the method cannot solve by itself. The register makes it visible when a recommendation is not followed up, but it cannot make people take ownership and keep it up to date. This risk came back through the whole study. In the interviews people said that action tracking is something the organisation struggles with, and that lists with deadlines and reminders still did not lead to follow-up (Section 4.1). The producer-side respondents warned that a register turns into a "bookkeeping exercise" or a "chore" when filling it becomes the goal. In the workshop the participants pointed to "Plato", an older register at SSC-ICT that was set up well but was not kept up to date. The design tries to lower this risk, for example by keeping the fields to a minimum (NFR2) and by making a missing update visible (NFR3, FR9), but it stays a condition for the method to work. This is the case for every register and every governance instrument, not only for this design. Whether SSC-ICT keeps the register alive is a question of culture and management, and this thesis could not test that over time.

The generalisability is only discussed, it is not checked in practice. The artefact separates the general parts from the case-specific parts (Section 3.5), and the AC-ICT respondent said the patterns at SSC-ICT are government-wide and not SSC-ICT-specific. Still, the method was designed and evaluated in one organisation only. No second organisation was used to test the transfer. The context is also not fixed: SSC-ICT is in the middle of a SAFe-style agile transition. The method was judged against the organisation as it is now. When the transition changes the decision bodies or the planning-and-control cycle, the case-specific mapping has to be done again, while the general parts stay the same. That poses no issue, as the rest of the method holds.

Chapter 9

Conclusions

This chapter closes the thesis. The goal was to design and evaluate a follow-up and governance method for ICT-related recommendations at SSC-ICT. We first answer the research questions, one by one. After that we give the contributions of the work, for practice and for theory. The chapter ends with the future work.

9.1 Answers to the research questions

RQ1 asked how ICT-related recommendations are now followed up and governed at SSC-ICT. The follow-up runs in two ways. Recommendations from the audit cycle are followed up relatively well, through an audit log and an audit monitor, with an owner and a closure only after an independent check by IAA. For the advisory and policy recommendations there is no such fixed process. They are handed to a unit and then followed up in a loose way, and there is no central overview of all recommendations together.

RQ2 asked which problems the stakeholders feel in this current way of working. The problem that comes back the most is ownership. It is often not clear who is the owner of a recommendation, so the work stays lying around and the discussion about who is responsible keeps coming back. Recommendations are handed over "over the fence", without a clear scope, result, or review moment, and can end up "in the drawer". There is also no shared overview of which recommendation is in which stage. And a clear owner alone is not enough, because the organisation has trouble to keep action lists up to date.

RQ3 asked which requirements for an effective method follow from the literature and from practice. We derived 25 requirements from three sources: the interviews inside SSC-ICT, the interviews with the producers of recommendations, and the literature. The requirements are written with the Volere template and split into project constraints (PC), functional requirements (FR), and non-functional requirements (NFR). The most important ones ask for a clear owner, a status that shows the stage of a recommendation, an intake check on the quality, a low administrative burden, and a register that hangs on the structures that are already there instead of a parallel system.

RQ4 asked how these requirements are translated into a method. The answer is the recommendation register: a method with a field schema, and not a software tool. The heart of it is a status model with three layers from the effect chain (follow-up, implementation, impact), so the register shows in which layer a recommendation is. Around this the register sets a clear owner at a senior body, a review moment in meetings that already take place, an intake check before an owner is set, a closure with a reason and four closure types, and a history that you can only add to. For the parts that another tool already covers, like the implementation status in Jira, the register links to that tool and does not copy it.

RQ5 asked to what extent the method supports the stakeholders, measured as perceived ease of use, perceived usefulness, and intention to use after they used it on a case. The method scores positive on usefulness (3.95) and on intention to use (3.93). The stakeholders find it useful and would use it. Ease of use scores lower (3.36) and is the weak point, together with the worry about the administrative burden. So the method is seen as useful and wanted, but the ease of use and the discipline to keep it up to date still need attention.

Together these answers give an answer to the main question: how a structured follow-up and governance method can be designed and evaluated for a government organisation like SSC-ICT. The method does this by making each recommendation a separate object that is followed over time, with a clear owner, a status in three layers, and a review and closure at a senior body. The demonstration shows that the method can hold the recommendations that SSC-ICT really gets. The validation shows that the stakeholders see it as useful and want to use it. The method does not yet show its effect in real use, and it stands or falls with the discipline to keep the register up to date.

9.2 Contribution to theory

This thesis makes three contributions to the literature reviewed in Chapter 2.

First, it gives a receiver-side view of the effect chain. Lemmens describes the chain from the producer side, where a Rekenkamer loses control after it publishes a recommendation [20]. We turned the chain around for the receiving organisation, where control is gained at the handover instead of lost (Section 5.2). This matters for a part of the literature that treated "open" and "implementation status" as the same thing. Callejas shows that a binary status leaves recommendations marked "in progress" without saying why [4], and the fragmentation sources describe follow-up that is spread across reports and arenas without a shared status [27, 19]. The three-layer status model (follow-up, implementation, impact) answers that problem. It separates the formal commitment, the operational change, and the effect, so that "open" and "closed" no longer hide which stage a recommendation is in. We also found a limit of the original chain. It assumes one producer and one political arena. When several producers feed a multi-layer receiver, the chain has to be set up per source, not once for the whole organisation (Section 2.7).

Second, it treats recommendation follow-up as a long-term governance object and not as a task that is either done or not. The literature already pointed this way. Di Giulio and Vecchi argue that ICT changes interdependencies and that actors have to keep learning [14], Dawes shows that feasibility and desirability change as the policy context changes [13], and Magnusson et al. show that security governance needs continuous reassessment [22]. This thesis puts that view into a design. The register keeps a recommendation under periodic review, with ex-ante and ex-post records and an append-only history, so that a recommendation can be reinterpreted, reprioritised, or retired with a reason instead of staying open without one (Section 5.1).

Third, it uses the Method Evaluation Model outside its usual setting. Moody built the MEM, on top of the Technology Acceptance Model [12], to validate information systems design methods [26]. Here it is used to validate a public-sector governance method. This shows that the MEM constructs (perceived ease of use, perceived usefulness, intention to use) also work for a method whose object is organisational follow-up and not software or a modelling technique.

9.3 Future work

The most important next step is to test the method in real use. This thesis stops at validation, so the effect on the real follow-up is not shown yet. For a real-use test the field-level register first has to be built on the tools that SSC-ICT already uses (PC2). Then the method can be used on real recommendations inside the organisation and followed over a longer time, as an evaluation after real use in the sens of Wieringa, for example through action research [37]. Such a study would also test the parts that one workshop could not reach, found in Table 7.3. And it would show whether the administrative burden stays low when the register is used over time, and not on one case.

A second direction is to lower the administrative burden, because it is the weak point of the method and the main risk for keeping the register up to date. One way is an AI assistant that reads the notes of the meetings to which the reviews are linked (NFR1) and proposes a draft update, a status change with

its reason or a closure rationale, which the owner then confirms or corrects [38, 10]. The register would then be updated as a by-product of a meeting that happens anyway. Whether this really lowers the felt burden, and whether it changes the behaviour, has to be built and tested.

A third direction is the quality of the delivered work. The method checks the incoming recommendation (FR6) and records why an item is closed (FR13), but is does not check whether the work behind a closure is good enough. A next version could add a "definition of done" to the closure step, so that the register also says something about the quality of the follow-up and not only that it happened.

A fourth direction is the transfer to other organisations. The method was designed and evaluated only at SSC-ICT, and the split between the general parts and the case-specific parts (NFR7) was argued but not tested. A next study could apply the method in another government organisation and see which parts stay the same and which parts have to be set up again.

Bibliography

[1] Algemene Rekenkamer. *Grip op gegevens: het stelsel van basisregistraties voor burgers en bedrijven*. Den Haag: Algemene Rekenkamer, Dec. 6, 2019.

[2] Auditdienst Rijk (ADR). *Opvolging BIT en Gateway adviezen DSO*. Onderzoeksrapport 2021-0000148449. Den Haag: Auditdienst Rijk (ADR), July 26, 2021. URL: https://www.eerstekamer.nl/overig/20210930/onderzoeksrapport_opvolging_bit_en/document3/f=/vlmooer486zx_opgemaakt.pdf.

[3] Australian National Audit Office. *Implementation of audit recommendations: Department of Health*. Performance Audit 8 2014–15. OCLC: 902612820. Canberra, 2014.

[4] Jorge T Flores Callejas. *REVIEW OF THE ACCEPTANCE AND IMPLEMENTATION OF JIU RECOMMENDATIONS BY THE UNITED NATIONS POPULATION FUND (UNFPA)*. JIU/ML/2016/5. Geneva: United Nations, 2016.

[5] Edna Dias Canedo et al. "Best Practices Kits for the ICT Governance Process within the Secretariat of State-Owned Companies of Brazil and Regarding these Public Companies". In: *Information* 9.6 (June 2018), p. 141. ISSN: 2078-2489. DOI: 10.3390/info9060141. URL: https://www.mdpi.com/2078-2489/9/6/141.

[6] Edna Dias Canedo et al. "Proposal of an Implementation Methodology of ICT Processes". In: *Information* 10.11 (Oct. 23, 2019), p. 327. ISSN: 2078-2489. DOI: 10.3390/info10110327. URL: https://www.mdpi.com/2078-2489/10/11/327.

[7] Hyeri Choi et al. "Rethinking the assessment of e-government implementation in developing countries from the perspective of the design–reality gap: Applications in the Indonesian e-procurement system". In: *Telecommunications Policy*. The promise and reality: assessing the gap between theory and practice in ICT4D 40.7 (July 1, 2016), pp. 644–660. ISSN: 0308-5961. DOI: 10.1016/j.telpol.2016.03.002. URL: https://www.sciencedirect.com/science/article/pii/S0308596116300015.

[8] Antonio Cordella and Carla M. Bonina. "A public value perspective for ICT enabled public sector reforms: A theoretical reflection". In: *Government Information Quarterly* 29.4 (Oct. 2012), pp. 512–520. ISSN: 0740624X. DOI: 10.1016/j.giq.2012.03.004. URL: https://linkinghub.elsevier.com/retrieve/pii/S0740624X12001001.

[9] Antonio Cordella and Andrea Paletti. "ICTs and value creation in public sector: Manufacturing logic vs service logic". In: *Information Polity* 23.2 (June 29, 2018). Ed. by Soon Ae Chun, Nabil R. Adam, and Beth Noveck, pp. 125–141. ISSN: 15701255, 18758754. DOI: 10.3233/IP-170061. URL: https://journals.sagepub.com/doi/full/10.3233/IP-170061.

[10] John Dagdelen et al. "Structured information extraction from scientific text with large language models". In: *Nature Communications* 15.1 (Feb. 15, 2024), p. 1418. ISSN: 2041-1723. DOI: 10.1038/s41467-024-45563-x. URL: https://www.nature.com/articles/s41467-024-45563-x.

[11] Nooranita Dain and Mohd Mohid Rahmat. "Factors Influencing Public Sector Auditees on Implementing Audit Recommendations." In: *Jurnal Pengurusan* 51 (Dec. 15, 2017), p. 1. ISSN: 0127-2713. DOI: 10.17576/pengurusan-2018-51-17. URL: https://openurl.ebsco.com/contentitem/doi:10.17576%2Fpengurusan-2018-51-17?sid=ebsco:plink:crawler&id=ebsco:doi:10.17576%2Fpengurusan-2018-51-17.

[12] Fred Davis. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology". In: *MIS Quarterly* 13 (Sept. 1, 1989), p. 319. DOI: 10.2307/249008.

[13] Sharon S. Dawes. "The Evolution and Continuing Challenges of E-Governance". In: *Public Administration Review* 68 (s1 2008), S86–S102. ISSN: 1540-6210. DOI: 10.1111/j.1540-6210.2008.00981.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6210.2008.00981.x.

[14] Marco Di Giulio and Giancarlo Vecchi. "Implementing digitalization in the public sector. Technologies, agency, and governance". In: *Public Policy and Administration* 38.2 (Apr. 2023), pp. 133–158. ISSN: 0952-0767, 1749-4192. DOI: 10.1177/09520767211023283. URL: https://journals.sagepub.com/doi/10.1177/09520767211023283.

[15] Loveness Innocent and Edwin Kwesigabo. "The Impact of Management Support on the Timely Execution of ICT Audit Recommendations: A Case Study of Tanzania Posts Corporation". In: *EJSMT* 1.2 (June 30, 2025), pp. 49–60. ISSN: 3083-6565. DOI: 10.59324/ejsmt.2025.1(2).04. URL: https://ejsmt.com/index.php/journal/article/view/12.

[16] Zahir Irani et al. "The impact of legacy systems on digital transformation in European public administration: Lesson learned from a multi case analysis". In: *Government Information Quarterly* 40.1 (Jan. 2023), p. 101784. ISSN: 0740624X. DOI: 10.1016/j.giq.2022.101784. URL: https://linkinghub.elsevier.com/retrieve/pii/S0740624X22001204.

[17] Tina Jukić, Mirko Vintar, and Jože Benčina. "Ex-ante evaluation: Towards an assessment model of its impact on the success of e-government projects". In: *Information Polity* 18.4 (Nov. 1, 2013), pp. 343–361. ISSN: 1570-1255. DOI: 10.3233/IP-130320. URL: https://journals.sagepub.com/action/showAbstract.

[18] Auke Klein Entink. *Kabinetsreacties, nader bekeken*. Den Haag: ICTU, Aug. 2024. URL: https://staatvandeuitvoering.nl/onderzoek/kabinetsreacties-nader-bekeken/.

[19] Teemu Mikael Lappi, Kirsi Aaltonen, and Jaakko Kujala. "Project governance and portfolio management in government digitalization". In: *Transforming Government: People, Process and Policy* 13.2 (May 20, 2019), pp. 159–196. ISSN: 1750-6166, 1750-6166. DOI: 10.1108/TG-11-2018-0068. URL: http://www.emerald.com/tg/article/13/2/159-196/374967.

[20] Etienne Lemmens. "De effectketen van rekenkameronderzoek". In: *Beleidsonderzoek Online* (Dec. 2025). ISSN: 2213-3550. DOI: 10.5553/BO/221335502025017. URL: https://www.boomportaal.nl/doi/10.5553/BO/221335502025017.

[21] Shuhua Monica Liu and Qianli Yuan. "The Evolution of Information and Communication Technology in Public Administration". In: *Public Administration and Development* 35.2 (2015). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/pad.1717, pp. 140–151. ISSN: 1099-162X. DOI: 10.1002/pad.1717. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/pad.1717.

[22] Lars Magnusson et al. "Information security governance in the public sector: investigations, approaches, measures, and trends". In: *International Journal of Information Security* 24.4 (July 16, 2025), p. 177. ISSN: 1615-5270. DOI: 10.1007/s10207-025-01097-x. URL: https://doi.org/10.1007/s10207-025-01097-x.

[23] Morten Meyerhoff Nielsen. "Governance lessons from Denmark's digital transformation". In: *Proceedings of the 20th Annual International Conference on Digital Government Research*. dg.o '19. New York, NY, USA: Association for Computing Machinery, June 18, 2019, pp. 456–461. ISBN: 978-1-4503-7204-6. DOI: 10.1145/3325112.3329881. URL: https://dl.acm.org/doi/10.1145/3325112.3329881.

[24] Matthew B. Miles, A. M. Huberman, and Johnny Saldaña. *Qualitative data analysis: a methods sourcebook*. Third edition. Thousand Oaks, Califorinia: SAGE Publications, Inc, 2014. 381 pp. ISBN: 978-1-4522-5787-7.

[25] Gianluca Misuraca and Gianluigi Viscusi. "Shaping public sector innovation theory: an interpretative framework for ICT-enabled governance innovation". In: *Electronic Commerce Research* 15.3 (Sept. 2015), pp. 303–322. ISSN: 1389-5753, 1572-9362. DOI: 10.1007/s10660-015-9184-5. URL: http://link.springer.com/10.1007/s10660-015-9184-5.

[26] Daniel Moody. "The Method Evaluation Model: A Theoretical Model for Validating Information Systems Design Methods". In: *ECIS 2003 Proceedings* (Jan. 1, 2003). URL: https://aisel.aisnet.org/ecis2003/79.

[27] OECD. *Implementation Toolkit for the OECD Recommendation on Public Policy Evaluation*. OECD Public Governance Reviews. OECD Publishing, Feb. 3, 2025. DOI: 10.1787/77faa4fe-en. URL: https://www.oecd.org/en/publications/implementation-toolkit-for-the-oecd-recommendation-on-public-policy-evaluation_77faa4fe-en.html.

[28] Wanda J. Orlikowski and C. Suzanne Iacono. "Research Commentary: Desperately Seeking the "IT" in IT Research—A Call to Theorizing the IT Artifact". In: *Information Systems Research* 12.2 (June 2001), pp. 121–134. ISSN: 1047-7047, 1526-5536. DOI: 10.1287/isre.12.2.121.9700. URL: https://pubsonline.informs.org/doi/10.1287/isre.12.2.121.9700.

[29] Ken Peffers et al. "A design science research methodology for information systems research". In: *Journal of Management Information Systems* 24 (Jan. 1, 2007), pp. 45–77.

[30] Suzanne Robertson and James Robertson. *Mastering the Requirements Process: Getting Requirements Right*. Google-Books-ID: yE91LgrpaHsC. Addison-Wesley, Aug. 6, 2012. 579 pp. ISBN: 978-0-13-294284-3.

[31] Nur Imroatun Sholihat. "The Implementation of COBIT 4.1 and COBIT 5-Based IT Governance Audits in the Ministry of Finance of Indonesia". In: *The Indonesian Journal of Accounting Research* 21.2 (2018). ISSN: 2086-6887. DOI: 10.33312/ijar.341. URL: https://ijar-iaikapd.or.id/index.php/ijar/article/view/341.

[32] Stuurgroep Staat van de Uitvoering. *Staat van de Uitvoering 2024 — Doorbreek de status quo*. Den Haag: Stuurgroep Staat van de Uitvoering, 2024. URL: https://open.overheid.nl/documenten/d1256815-d92e-414c-81b8-d627d1e0cb4a/file.

[33] Tijdelijke commissie ICT-projecten. *Parlementair onderzoek naar ICT-projecten bij de overheid: Eindrapport*. 33326, nr. 5. Den Haag: Tweede Kamer der Staten-Generaal, Sept. 16, 2014, p. 5. URL: https://www.tweedekamer.nl/sites/default/files/field_uploads/33326-5-Eindrapport_tcm181-239826.pdf.

[34] Sarimah Umor, Zarina Zakaria, and Noor Adwa Sulaiman. "Effectiveness of Follow-Up on Performance Auditing Issues in Practice: A Governance Network Perspective". In: *Asian Journal of Accounting Perspectives* 13.1 (Feb. 28, 2020), pp. 1–25. ISSN: 0128-0384. DOI: 10.22452/AJAP.vol13no1.1. URL: https://ajap.um.edu.my/index.php/AJAP/article/view/22452.

[35] Eddie Van Marum. *Verzamelbrief digitalisering december 2025*. Letter. Last Modified: 2026-01-09T14:00. Dec. 18, 2025. URL: https://www.rijksoverheid.nl/documenten/kamerstukken/2025/12/18/verzamelbrief-digitalisering-december-2025.

[36] Camilla Wanckel. "Introducing a digital tool for sustainability impact assessments within the German Federal Government: A neo-institutional perspective". In: *International Review of Administrative Sciences* 89.2 (June 1, 2023), pp. 433–449. ISSN: 0020-8523. DOI: 10.1177/00208523211047093. URL: https://doi.org/10.1177/00208523211047093.

[37] Roel J. Wieringa. *Design Science Methodology for Information Systems and Software Engineering*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014. ISBN: 978-3-662-43838-1 978-3-662-43839-8. DOI: 10.1007/978-3-662-43839-8. URL: https://link.springer.com/10.1007/978-3-662-43839-8.

[38] Anneke Zuiderwijk, Yu-Che Chen, and Fadi Salem. "Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda". In: *Government Information Quarterly* 38.3 (July 1, 2021), p. 101577. ISSN: 0740-624X. DOI: 10.1016/j.giq.2021.101577. URL: https://www.sciencedirect.com/science/article/pii/S0740624X21000137.

Appendix A

Exploratory interview protocol (receiver side)

The interview protocol was used in Dutch. It has been translated to English for this appendix.

This appendix contains the protocol of the exploratory receiver-side interviews at SSC-ICT (Section 3.4.2). The protocol is semi-structured: the questions are indicative, the order and wording were adapted to the role of the respondent, and the *probes* were only used where relevant.

Introduction

- Thanks + goal: I want to understand how follow-up currently works and where it breaks down, so that I can design a workable model.
- Takes at most about an hour.
- Descriptive, not an assessment.
- May I record? The recording will not be used for publication, at most for an anonymous quote, and serves as a reference for me.

Context & role

- What is your role in relation to recommendations/advice? (receiving, translating, executing, monitoring, reporting, deciding)
  - *Probe:* at which meeting tables/committees/consultations are you present where this comes up?
- Which types of recommendations do you come across? (audit findings vs. strategic advice vs. security/operations)
- From which sources do these recommendations mainly come? (internal/external, audit/evaluation/advisory)

Practice

- Walk me through it: from the moment of publication/receipt → to "completed/closed". What are the steps?
  - *Probe:* who picks it up first? Who translates it into actions? Who monitors?
- Where is it decided: "do we do this or not / now or later"?
  - *Probe:* is that an explicit decision, or does it "slip away"?

- How does a recommendation eventually reach execution down the line (teams, backlog, projects, line)?
  - *Probe:* does this happen via agile teams, projects, line management, or ad hoc?
- What is a realistic definition of 'following up' at your organisation?
  - *Probe:* fully implementing, partly, including it in an ongoing programme, or "addressed with a rationale"?

Registration and labeling

- Where is the overview of open recommendations kept? (Excel, Trello/Jira, SharePoint, personal lists, audit tooling)
  - *Probe:* is there one single source of truth or several?
- Which statuses do you use (open/in progress/closed/etc.) and what does each status mean in practice?
  - *Probe:* how do you show "evidence" of progress/closure?
- How are recommendations referred to? (IDs, report paragraph, theme/category)
  - *Probe:* is there a labeling practice that helps or rather causes confusion?
- In what way are recommendations made SMART (or not)?
  - *Probe:* what is typically too abstract ("the what") vs. too detailed (unworkable)?
- Which information fields are really needed to be able to steer? (minimal set)
  - *Probe:* owner, deadline, priority, dependencies, risk/impact, evidence, decision moment, rationale, "obsolete yes/no".

Governance: prioritisation, decision-making, review cycle

- How is priority determined when there are many recommendations?
  - *Probe:* criteria? (risk, compliance, customer impact, feasibility, dependencies, time-sensitivity)
- Which review moments/cadence exist? (monthly/quarterly/yearly; per domain; per board)
- What happens to recommendations that do not clearly fall under Finance (CFO) or Tech (CTO)?
  - *Probe:* "on whose plate" does this land? How is ownership enforced?
- How does escalation work when something stays stuck? (who decides, what is the trigger?)

Bottlenecks & causes

- Does follow-up ever go wrong? If so, where?
  - in the advice itself (too abstract / no "how" / too much / not fitting), or
  - in the follow-up process/governance (ownership / visibility / prioritisation / review), or
  - in execution (capacity / dependencies / legacy)?
  - *Probe:* give one concrete example per chosen category.
- Does it sometimes become unworkable (too many recommendations, too much administration)? Where is that limit?
- To what extent does (a lack of) (external) pressure play a role in follow-up?
  - *Probe:* the difference between "annually recurring" audits/monitoring versus one-off advice.

Improvement needs & design input

- Suppose you could set it up again. What are the 3 most important changes?
- Has an improvement register / central monitor been attempted before?
  - *Probe:* what worked / did not work?
- What could a workable minimal solution (MVP) be?
- When is follow-up "good enough" for you? (measurable/observable)
  - *Probe:* what would you then need to be able to show to MT/boards/auditors?

Closing

- Word of thanks.
- Who are 2–3 others I absolutely should speak to (another angle of the process)?

Appendix B

Producer-side interview protocol (advisory and oversight bodies)

The interview protocol was used in Dutch. It has been translated to English for this appendix.

This appendix contains the semi-structured protocol used for the producer-side interviews with advisory and oversight bodies (Section 3.4.2); the four interviews are analysed in the producer-side matrix in Section 4.3. The protocol is semi-structured: the questions are indicative, the order and wording were adapted to the role of the respondent, and not every question was asked in every interview. The interview with Lemmens used a more concept-oriented variant of this protocol, built around his effect chain framework; it is not reproduced separately here.

Introduction

- General introduction: who I am, what I do, and why this research matters.

Role

- Can you tell me what your organisation does and what role you yourself have in it?
- How long have you been working in this role?
- Which types of assignments do you or your team carry out?
- What is the task of your organisation, what do you try to achieve, and what do you deliver?
- Which types of issues do you focus on? (compliance, strategy, risk, governance)
- Do you mainly signal, advise, assess, or actually steer?
- Who are your typical clients or recipients?
- Is there a specific focus on ICT or digitalisation themes?

Mandate and form

- On what basis does your organisation give advice; that is, do you have a formal mandate or a legal basis?
- Is advice from your organisation binding, or is it non-committal for the receiving party?
- How does that differ (if at all) per type of report or assignment?
- How does your mandate relate to that of other oversight or advisory bodies (e.g. ADR, Rekenkamer, ACICT)?

- Is your advice more focused on compliance (testing against rules and standards), or do you also give broader, strategic advice on how risks can be mitigated or systems improved?
- What are the advantages and risks of this approach from your perspective?
- Do you notice a difference in how receiving parties deal with compliance findings versus broader strategic advice?
- What does the communication around an advisory report look like? Which steps do you go through from research to delivery?
- Is there a right of reply (*hoor en wederhoor*) before publication?
- Who receives the report and who is responsible for a response?
- What is the form of the delivery (presentation, written report, debriefing)?

Production of the advice

- How is the end product usually produced? That is, which sources, conversations, analyses, or assessment frameworks do you use?
- How do you determine whether a recommendation is sufficiently concrete, usable, and executable?

Follow-up

- Looking back at the advice or reports you have issued: how do you experience the follow-up of these by the receiving organisations in general?
- Can you give a concrete example of advice where follow-up went well?
- And an example where it went less well?
- What made the difference?
- Does your organisation pay explicit attention to the follow-up of advice issued earlier? Do you monitor whether recommendations have been adopted?
- Is there a formal follow-up process on your side?
- Do you track multi-year KPIs or indicators based on your advice?
- How does that relate to whether or not your advice is followed up?
- Which obstacles or barriers do you see at receiving organisations that hinder the follow-up of advice?
- Do factors such as capacity, prioritisation, ownership, cross-organisational dependency, or political will play a role?
- What is your impression of the difference between advice with strong external pressure (e.g. an annual audit) versus one-off advice?
- Do recommendations sometimes lose relevance because technology, policy, or context changes?

More ACICT-related:

- Specifically about ICT advice: do you see differences in how ICT-related recommendations are picked up versus other types of advice?
- What makes ICT advice more complex or, on the contrary, easier to follow up?
- Does the technical complexity or the dependency on other parties play a role?
- Are there striking examples that have stayed with you?

Confirmatory section

From my earlier research at SSC-ICT, a number of obstacles emerged in the follow-up of advice. I would like to hear whether you recognise these from your perspective as an advising party, and whether you would like to add to or nuance them.

Obstacle	**Description / verification question**
No uniform follow-up process for non-audit advice	Formal audit findings are usually followed up more systematically than strategic advice. Do you recognise that distinction? What is your explanation for it?
Fragmented registration and no central overview	Recommendations are kept in separate systems, in working documents, or not at all. Do you have a picture of this, and do you experience it as a problem?
Unclear ownership and responsibility	It is not always clear who owns the follow-up of a recommendation. Do you recognise this? What would be a solution for it?
External pressure as the main driver	Follow-up seems to depend partly on recurring external pressure (e.g. an annual audit). Do you share that observation? How do you view it?
No formal follow-up obligation for advisory reports	Unlike accountants' or audit findings, follow-up of advisory reports is often non-committal. In your experience, does this lead to less discipline in follow-up?
No multi-year monitoring of the effect of advice	Multi-year KPIs or indicators to measure whether advice has structurally changed anything are hardly kept. Do you recognise that? Is that something you would want to change?

Improvement

- What would help organisations, or what should they improve, to monitor, prioritise, and govern recommendations better over time?
- Which information should, in your view, be in a good follow-up model or register as a minimum?
- Which pitfalls should I avoid when designing a governance and/or follow-up model?

Closing

- Is there anything you would still like to add that is relevant for this research, or that has not yet been covered?

Appendix C

Validation questionnaire (Method Evaluation Model)

The questionnaire was administered in Dutch. It has been translated to English for this appendix.

This appendix contains the post-task survey used in the validation workshop (Section 3.4.4 and Section 7.2). The instrument is adapted from the Method Evaluation Model of Moody [26]. It has sixteen statements: six for perceived ease of use, eight for perceived usefulness, and two for intention to use. Each statement uses a five-point Likert scale, half of the statements are negatively worded, and the statements are presented in random order. Three open questions follow the scale. The wording is changed from the use of a system to the use of the register, and the perceived-usefulness statements are written in terms of the goals of the model (keeping recommendations visible, making ownership explicit, and supporting steering).

Questionnaire as presented to participants

This questionnaire measures your experience with the register in terms of ease of use, usefulness, and intention to use. For each statement, indicate to what extent you agree, on a scale from strongly disagree to strongly agree. Note that some statements are deliberately negatively worded. There are no right or wrong answers. Three open questions follow at the end.

Scale: 1 = strongly disagree, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree.

Statement	1	2	3	4	5
1. I found the procedure for filling in the register complex and hard to follow.	☐	☐	☐	☐	☐
2. I think the register would reduce the effort needed to follow up recommendations.	☐	☐	☐	☐	☐
3. With this register it would become harder to keep an overview of recommendations.	☐	☐	☐	☐	☐
4. Overall, I found the register hard to use.	☐	☐	☐	☐	☐
5. The register would make it easier to establish who is responsible for what.	☐	☐	☐	☐	☐
6. I found the register easy to learn.	☐	☐	☐	☐	☐
7. Overall, I found the register useful.	☐	☐	☐	☐	☐
8. Using this register would make it harder to keep following recommendations over a longer period.	☐	☐	☐	☐	☐
9. I found it hard to apply the register to this case.	☐	☐	☐	☐	☐

Statement	1	2	3	4	5
10. I would certainly not use this register to follow up recommendations.	☐	☐	☐	☐	☐
11. I found the rules of the register clear and understandable.	☐	☐	☐	☐	☐
12. Overall, I think this register does not offer an effective solution for following up recommendations.	☐	☐	☐	☐	☐
13. Using this register would make it easier not to lose sight of recommendations.	☐	☐	☐	☐	☐
14. I am not sure that I could apply the register on my own in practice.	☐	☐	☐	☐	☐
15. Overall, I think this register is an improvement over how we currently follow up recommendations.	☐	☐	☐	☐	☐
16. If I work with recommendations in the future, I intend to use this register.	☐	☐	☐	☐	☐

Open questions

A. What would the register change for you in your work?

B. What is your biggest doubt or objection to this model?

C. What is still missing in the model for it to be usable in your work?

Questionnaire design: construct mapping

Table C.2 shows, per statement, the construct it measures, whether it is positively or negatively worded, and the corresponding item in the original survey of Moody. The statements are presented to participants in random order, the mapping is used only for analysis. Before analysis, the negatively worded statements are reverse-coded so that a high score always means a positive judgement.

Table C.2: Mapping of each questionnaire statement to its construct, direction, and source item.

Statement	Construct	Direction	Source (Moody)
1	Perceived ease of use	Negative	Q1
2	Perceived usefulness	Positive	Q2
3	Perceived usefulness	Negative	Q3
4	Perceived ease of use	Negative	Q4
5	Perceived usefulness	Positive	Q5
6	Perceived ease of use	Positive	Q6
7	Perceived usefulness	Positive	Q7
8	Perceived usefulness	Negative	Q8
9	Perceived ease of use	Negative	Q9
10	Intention to use	Negative	Q10
11	Perceived ease of use	Positive	Q11
12	Perceived usefulness	Negative	Q12
13	Perceived usefulness	Positive	Q13
14	Perceived ease of use	Negative	Q14
15	Perceived usefulness	Positive	Q15
16	Intention to use	Positive	Q16

In total: six statements for perceived ease of use (1, 4, 6, 9, 11, 14), eight for perceived usefulness (2, 3, 5, 7, 8, 12, 13, 15), and two for intention to use (10, 16). Eight statements are negatively worded and eight positively worded.

Appendix D

Code-to-theme inventory

This appendix contains the full code-to-theme inventory. Table 4.2 in the body gives the condensed overview of the ten themes; the inventory below shows, per pattern-coded theme, the first-cycle descriptive codes that were grouped under it during the second (pattern-coding) cycle, following Miles, Huberman, and Saldaña [24]. Codes that appeared in more than one theme were deliberately kept under each related theme; these cross-theme codes are discussed in Section 4.1.2.

Table D.1: Code-to-theme inventory: first-cycle codes grouped under each pattern-coded theme.

Pattern-coded theme (n codes)	**First-cycle codes**
Analytical themes	
T1 — Current follow-up practices (RQ1; 41 codes)	Accountant findings; Administrative response; Application for external review; Audit format discussion; Audit log and monitor is an Excel-file; Audit log for IT; Audit owner; Audit phases; Audit scope; Auditlog and auditmonitor; Bi-weekly meeting; Broader use possible of audit monitor; Closure requirement; Confident in auditlog; Control cycle; Control cycle right of reply; Finding; Finding closure; Follow-up improved; High-level check; Implicit risk assessment in report; Initial risk-analysis; Internal audit; Management response; Monitoring status; Norm assessment; Recommendation check; Recommendation closure conditions; Recommendation follow-up; Recommendation has been followed; Recommendation update; Resolve interim findings; Response cycle; Secured findings; Weighted findings.
T2 — Governance structure and formal roles (RQ1; 29 codes)	ACICT; ADR; Board structure; Board structure includes DMT; CEO has final say; CIO-BZK; CTO is also the CIO; Customers' council; Dispersed governance responsibilities; Escalation happens too little; Explicit decision-making; External audit CIO-Rijk; From structure transforming governance; Governance structure; Governance transparency; Gremia; Hierarchical structure governs escalation; Low need for exceptional escalation; Management-team has final say in priority; Management-team responsibility; Managers are people, process and technology; Many governance bodies; Programme; Rekenkamer; Reporting structure; Steering group responsibility; Steering groups are project based; Strategic-tactical-operational; Too much distance to top governance.
T3 — Responsibility and ownership gaps (RQ2; 28 codes)	Claiming no responsibility; Cross-organisation discussion; Cross-organisation responsibility; Dispersed monitoring; Management needs to address ownership; More need for delegation of responsibility; Need for clear responsibility assignment; Need for responsibility assignment is the basis; Need to take ownership of own responsibilities; Operations does not have recommendations in mind; Operations only look at the plan handed to them; Own sphere of influence; Own work is most important; Ownership assignment and delegation; Problem is that responsibility is delegated low in the organisation; Recommendation area shift; Responsibility governance; Responsibility; Responsibility discussion; Shared responsibility; Shifting responsibility to another; Smooth ownership assignment; Unclear responsibility; Working responsibility assignment.

Continued on next page

Table D.1 continued from previous page

Pattern-coded theme (n codes)	First-cycle codes
T4 — Monitoring, visibility and transparency gaps (RQ2; 33 codes)	Advice written away; Dispersed monitoring; Fragmented follow-up; Governance transparency; Implemented external progress monitor; Impression that recommendations are seen as side-issue; Knowing what is an issue in the organisation; Lack of active coordination and finding each other; Lack of active follow-up reporting; Lacking follow-up; Mandatory recording of follow-up steps; Need for logical assessment moments; Need for overview to be connected to workflows; Need for proactive monitoring; Need for process formulation; Need for transparency; No central logging of outstanding actions; No clear status of follow-up; No formal process description; No implementation tracking; No mandatory follow-up; No standard log; Operational level box-ticking; Programme manager; Recommendation traceability; Report thrown over the fence; Report vanished in drawer; Shared overview; TopDesk platform monitoring; TopDesk traceable workflows; Unclear overview.
T5 — Prioritisation and decision-making (RQ2; 23 codes)	"Everything is important" problem; Accepting residual risk; Deadline setting; Difficult priority assessment; Escalation happens too little; Explicit decision-making; Feels panic for the workload of changes needed; Hierarchical structure governs escalation; Huge recommendation backlog; Impact of report gives urgency; Implementation of changes is lacking; Low need for exceptional escalation; Management-team has final say in priority; Need for logical assessment moments; Need for prioritisation; Need for realistic planning; No structured value assessment of recommendations; Plans based on business case; Political influence driven by madness of the day; Technical priorities are well established; Years-long trajectory.
T6 — Recommendation characteristics and quality (RQ2; 43 codes)	Disalignment of recommendations and IT technicians; Every report may induce political questions; Getting recommendations to right people is difficult; Hiring external consultancy firm; Hope for unambiguous recommendations; Huge recommendation backlog; Impact of report gives urgency; Implicit risk assessment in report; Independence consultancy firm; Lacking service from consultancy firm; Need for clarification of recommendations; No impact quantification; No structured value assessment of recommendations; Qualitative advice is more difficult than quantitative; Re-inventing the wheel; Reasons for outsourcing advice; Recommendation abstraction leaves room; Recommendation area shift; Recommendation effects; Recommendation governance proposed solution; Recommendation initiative; Recommendation outsourcing; Recommendation scope; Recommendation source; Recommendation stating to work in BUs; Recommendation stays open over time; Recommendations are too abstract; Recommendations difficult to execute; Recommendations often hang around; Recommendations stalling due to time and energy constraints; Reformulation of old advice; Reports stating the obvious; Soprasteria report; Useless recommendations.
T7 — Organisational capacity and cultural barrier (RQ2; 63 codes)	Already-busy people are the people being asked; Assumption that others understand what you mean; Capacity constraints; Colleagues being over-asked; Confidence tied to person, not process; Cooperation is difficult because of own silos; Coordination needs improvement; Culture change needed in way of working; Culture change to underpromise and overdeliver; Differing worlds of experience; Don't make things unnecessarily complicated; Everyone is too busy to do governance; Formalised processes are not sexy; From policy-side no attention for processes; Good preparation is half the job; If we weren't government, we would be bankrupt; Importance of organisational side; Knowledge is there, framework is not; Lack of active coordination and finding each other; Lacking policy-to-execution translation; Lacking time; Little attention for processes; Multiple groups working on the same thing without each other knowing; Need for better time management of colleagues; Need for coordination; Need for holding each other accountable; No interest in processes; Not everyone is good at communicating; Not everyone is good at governing; Organisation culture requires concrete instructions; Organisation driven on assumptions; Organisation is not mature; Organisation is reactive instead of proactive; Organisation needs to professionalise; Organisational change is difficult; Organisational maturity is important; Outside-of-silo cooperation is not well organised; People communicate from their own box; Policy has financial impact; Policy thrown over the fence; Regular work is already a lot of work; Resistance to change is just part of the process; Setting a good example; Technology is more important than processes; Tight structure lessens need for written processes; Way of working ingrained; Within-silo cooperation works well; Working outside of own island is difficult.
T8 — Tools and governance instruments (RQ3; 35 codes)	Audit log and monitor is an Excel-file; Auditlog and auditmonitor; Broader use possible of audit monitor; Change capacity is registered in Qlik Sense dashboard; Confluence tool; FCC tool; Jira in relation to Agile; Jira tool; Jira tool possible solution; Mandatory recording of follow-up steps; MS Project for waterfall; Need for better information management cross-organisation; Need for central governance of recommendations; Need for clear governance tool; Need for clear workflow insights; Need for overview to be connected to workflows; No standard log; PLATO tool; Process aligned with existing workflows; Qlik Sense dashboard as inspiration; Task management in lists; TopDesk platform monitoring; TopDesk traceable workflows; Transformatiebord tool; Unclear tool use; Visual representation of transition on talking board; Workflows are being organised in Lean portfolios.
Contextual themes	

Continued on next page

Table D.1 continued from previous page

Pattern-coded theme (n codes)	First-cycle codes
T9 — Agile transformation context (context; 27 codes)	Agile change is window of opportunity for broader changes; Agile working being implemented; Agile working is top-down mandated; Agile working requires mature organisation; Agile working was used unofficially; Delphi sessies; From structure transforming governance; Jira in relation to Agile; MS Project for waterfall; Objective Key Results; Pen testing methodology; Portfolio management of OKR; Product delivery is not flexible; Programme; Project waterfall; Red teaming methodology; Required periodic agency screening; Scrum epics; SSC-ICT is not customer focussed, unpredictable and has long delivery times; Taking initiative without communicating; Translation strategic to tactical OKR; Unclear enterprise architecture; Value streams; Workflows are being organised in Lean portfolios; Working back-to-front is an issue.
T10 — External oversight and pressure (context; 21 codes)	Application for external review; CIO-BZK; CIO-Rijk; Confidentiality discussion; External audit CIO-Rijk; External pressure; External pressure does not function; External pressure is good; Hiring external consultancy firm; Independence consultancy firm; Lacking service from consultancy firm; Reasons for outsourcing advice; Recommendation outsourcing; Rekenkamer; Required periodic agency screening; Requiring useless external validation; Soprasteria report.

Appendix E

Recommendation inventory

This appendix lists the inventory of recommendations that we collected from public reports about SSC-ICT and the BZ transition. The inventory holds 126 recommendations from nine source documents. We use this inventory as the empirical base for the demonstration cases in Chapter **??**. The source data are in Dutch; we translated them to English for this appendix. Dutch terms are quoted with an English translation in parentheses on first use. The two source columns that are not relevant for this thesis, level ("Niveau") and theme ("Thema"), are left out here.

For each recommendation we give six fields. **ID** is the stable identifier used across the thesis. **Finding / paragraph** points to the section of the source report where the recommendation was placed. **Recommendation** is the translated text. **Object** is the short subject tag from the source. **Owner** is the target party that the source addresses. **Page** is the page in the source document.

The organisational abbreviations are: SSC-ICT (Shared Service Centre ICT), BZ (Ministry of Foreign Affairs), BZK (Ministry of the Interior and Kingdom Relations), CIO-Rijk (Chief Information Officer of the central government), ADR (Audit Service of the Central Government, "Auditdienst Rijk"), AcICT (Advisory Council on ICT Assessment, "Adviescollege ICT-toetsing") and BIT (Bureau for ICT Assessment, "Bureau ICT-toetsing").

E.1 Source documents

Table E.1 lists the nine reports. The report code in the first column is the prefix of the recommendation IDs in Section E.2.

Code	Title	Type	Issuer	Date	Reference	Pp.
AR2023	Audit Report 2023 – Shared Service Centre ICT (SSC-ICT)	Audit report	ADR	15-03-2024	2024-0000174102	24
AR2024	Audit Report 2024 – Shared Service Centre ICT (SSC-ICT)	Audit report	ADR	14-03-2025	2025-0000063257	33
INT2025	Interim Audit Report 2025 – Shared Service Centre ICT (SSC-ICT)	Interim audit report	ADR	27-10-2025	2025-0000480708	30
LIC2020	License Management SSC-ICT	Research report	ADR	20-10-2020	2020-0000203716	21
KPM2020	Study of the Cost-Price Model SSC-ICT	Research report	ADR	21-12-2020	–	21
LCM2024	Robustness of LCM cost inclusion in the cost-price model SSC-ICT	Research report	ADR	10-10-2024	2024-0000475839	19
ACICT2025	AcICT Advice: Transition of ICT Service Provision (BZ)	Advice	AcICT	03-02-2025	2025-0000122205	8
TK1417	Parliamentary letter 26 643, no. 1417 (follow-up to AcICT advice)	Parliamentary letter	MinBZ + State Secretary BZK	01-10-2025	26 643, no. 1417	8
BIT2017	Final BIT Advice, project DWR Next	BIT advice	BIT	21-03-2017	2017-0000153638	6

Table E.1: Source documents of the recommendation inventory.

E.2 Recommendations

Table E.2 lists all 126 recommendations, grouped per source document. Where the source text was cut off in the original document, we mark this with "[text in source cut off]".

Table E.2: Recommendation inventory (translated from Dutch).

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
AR2023	*Audit Report 2023 – Shared Service Centre ICT (SSC-ICT)*				
AR2023-EI-01	3.4 External hiring – testing of knock-out criteria	Ensure accurate recording of the requirements and wishes in the digital system, so that discrepancies with the request form are prevented; record information only once so that contradictions cannot occur.	Process/control	SSC-ICT	14
AR2023-EI-02	3.4 External hiring – testing of knock-out criteria	Record transparently and unambiguously that the digital system is leading for requests, and (re)confirm this line towards suppliers.	Process/agreements	SSC-ICT + suppliers	14
AR2023-EI-03	3.4 External hiring – testing of knock-out criteria	Hold suppliers accountable on the basis of the sanction model when they offer candidates who do not meet the hard requirements.	Governance/ contract	SSC-ICT + suppliers	14
AR2023-EI-04	3.4 External hiring – process control	Steer towards a lasting change of behaviour in which maintaining the quality of procurement management is self-evident.	People/culture	SSC-ICT	14
AR2023-EI-05	3.4 External hiring – process control	Obtain sufficient insight into the working of the systems and the possible risks within procurement management, in order to prevent errors.	System + risk control	SSC-ICT	14
AR2023-DEV-01	3.5 Non-returned devices	Take account of additional capacity to process returned devices at the return desks adequately (in parallel with regular service and exchange requests).	Capacity/operation	SSC-ICT	15
AR2023-DEV-02	3.5 Non-returned devices	Ensure a frequent monitoring process around the number of non-returned devices, so that timely adjustment is possible.	Process/monitoring	SSC-ICT	15
AR2023-DEV-03	3.5 Non-returned devices	Consider the measure of issuing new devices only after the old devices have been handed in.	Policy/process	SSC-ICT + customers	15
AR2023-DFA-01	3.6 Lead time for signing DFAs	Hold customers within the BZK group accountable for signing the DFA (Dossier of Financial Agreements, “Dossier Financiële Afspraken”) on time.	Governance/steering	BZK (owner) + SSC-ICT + customers	16
AR2023-DFA-02	3.6 Lead time for signing DFAs	Communicate actively with customers that timely signing of the DFA is separate from substantive disputes.	Communication/ process	SSC-ICT + customers	16
AR2023-DFA-03	3.7 Revenue process DFA	Set up a control measure that ensures that settled corrections are adjusted in the source administration on time.	Control measure	SSC-ICT	17
AR2023-DFA-04	3.7 Revenue process DFA	Record visibly the control measure for input control on correct rates (four-eyes principle).	Control measure	SSC-ICT	17
AR2023-DFA-05	3.7 / Non-allocable service provision	Give priority to cleaning up accounts and other aspects that lead to non-allocable service provision, and adjust the source administrations accordingly.	Data/process	SSC-ICT	17
AR2023-DFA-06	3.7 / Non-allocable service provision	Strengthen the change management of the DFA generator by recording the procedures, having them approved, and updating them periodically.	System/change management	SSC-ICT	17
AR2023-ASSET-01	3.8 Assets outside the regular process	Draw up a process description for the goods receipt and the registration procedure of the relevant goods flow.	Process	SSC-ICT	18

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
AR2023-ASSET-02	3.8 Assets outside the regular process	Provide a work instruction with essential control measures, and take account of shielding secret information in the financial administration.	Process/work instruction	SSC-ICT	18
AR2023-ASSET-03	3.8 Assets outside the regular process	Make formal agreements with all owners of devices about the responsibility for management, so that an integral asset-management process arises; emphasise the integral responsibility of the owner.	Governance/ agreements	SSC-ICT + asset owners	18
AR2023-RISK-01	4.1 Insight into risks and irregularities	Set up a risk-management system for financial management processes and IT management processes (GITCs), with recorded roles/ responsibilities, (fraud) risk analyses, AO/IB (administrative organisation/internal control) and a closed PDCA cycle with periodic testing and reporting.	Governance/model	SSC-ICT (as part of BZK)	20
AR2023-TRANS-01	4.2 Transformation continued without earmarked fund	Investigate, in coordination with (FEZ) BZK and the Ministry of Finance, how the financing of the transformation can be safeguarded until the end of the programme.	Governance/ financing	SSC-ICT + owner (BZK) + Ministry of Finance	20
AR2023-FRAUDE-01	4.3 Fraud risks still to be anchored in policy	Adopt an anti-fraud policy as soon as possible and keep it up to date, by paying continuous attention to the control of fraud risks from a risk/ control framework.	Policy/risk control	SSC-ICT	21
AR2023-TOP-01	4.4 Migration of the Topdesk asset module	Investigate the possibilities for a strong connection between the financial administration (Exact) and Topdesk, in order to optimise asset management.	System integration	SSC-ICT	21
AR2023-TOP-02	4.4 Migration of the Topdesk asset module	Ensure that objects that stay behind in the Topdesk asset module are written off / adjusted in the financial administration and are correctly disposed of.	Process/ administration	SSC-ICT	21
AR2023-TOP-03	4.4 Migration of the Topdesk asset module	Ensure that the functionalities made available by Topdesk match the needs of the key users.	System/ requirements	SSC-ICT	21
AR2023-TOP-04	4.4 Migration of the Topdesk asset module	Safeguard that the transition to the line organisation is formally recorded, through handover and acceptance/approval by the key users.	Governance/ handover	SSC-ICT	21
AR2023-CHAIN-01	4.5 Chain dependency of the (procurement) administration	Strengthen the cooperation in the chain to optimise performance and user-friendliness and to serve customers better; attention and steering from BZK remain required here.	Governance/ cooperation	SSC-ICT + chain partners (OBF, RIS, O&P Rijk) + BZK	22
AR2023-LIC-01	4.7 License policy SSC-ICT	Adopt the license policy formally (as long as this has not yet happened), so that the finding can be regarded as resolved.	Policy	SSC-ICT	23
AR2024	*Audit Report 2024 – Shared Service Centre ICT (SSC-ICT)*				
AR2024-EI-01	3.3 External hiring	Keep carrying out the control measures, so that the lawfulness of external hiring is structurally safeguarded.	Process/control	SSC-ICT	14
AR2024-DEV-01	3.4 Return of devices	Keep ensuring a frequent monitoring process around the number of non-returned devices, so that timely adjustment is possible.	Process/monitoring	SSC-ICT	14
AR2024-DEV-02	3.4 Return of devices	Keep applying the firm measures (such as resetting passwords), in order to keep the return rate on level.	Policy/process	SSC-ICT + customers	14

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
AR2024-DFA-01	3.5 Lead time for signing DFAs	Keep communicating to customers the importance of the new DFA working method and the agreements made in the Administrative Consultation ("Bestuurlijk Overleg", BO), among others through the Customer Council ("Afnemersberaad").	Communication/ governance	SSC-ICT + customers	17
AR2024-DFA-02	3.5 Lead time for signing DFAs	Hold customers within the Ministry of BZK accountable for signing the DFA on time and for providing the purchase-order number on time for invoicing.	Governance/ agreements	SSC-ICT + BZK customers	17
AR2024-ALLOC-01	3.6 Non-allocable service provision	Carry out a root-cause analysis of the fluctuations (second half of 2024) at the business services with the largest share that leads to non-allocable service provision.	Analysis	SSC-ICT	20
AR2024-ALLOC-02	3.6 Non-allocable service provision	Keep giving priority to cleaning up accounts that lead to non-allocable service provision, and adjust the source administrations accordingly.	Data/process	SSC-ICT	20
AR2024-ALLOC-03	3.6 Non-allocable service provision	Strengthen the change management of the DFA generator by recording the procedures, having them approved, and updating them at fixed moments.	System/change management	SSC-ICT	20
AR2024-ASSET-01	3.7 Assets outside the regular process	Adopt the process description for the goods receipt and the registration procedure of this goods flow.	Process	SSC-ICT	23
AR2024-ASSET-02	3.7 Assets outside the regular process	Provide a work instruction with essential control measures, and take account of shielding secret information in the financial administration.	Work instruction/ control	SSC-ICT	23
AR2024-ASSET-03	3.7 Assets outside the regular process	Make formal agreements with all owners of the assets about the responsibility for management, so that an integral asset-management process arises; emphasise the responsibility of the owner.	Governance/ agreements	SSC-ICT + asset owners	23
AR2024-ASSET-04	3.7 Assets outside the regular process	Ensure that all assets are registered (on time) in the CMDB, in order to improve the view of building-bound assets and network equipment.	Data/system	SSC-ICT	13
AR2024-LIC-01	4.3.7 License policy SSC-ICT adopted	Communicate the license policy well with clients, because it affects licenses that are part of SSC-ICT service provision to clients.	Communication/ policy	SSC-ICT + clients	28
AR2024-ARCH-01	4.3.8 Central archiving of performance evidence	Communicate unambiguous guidelines on which information and evidence should preferably be recorded in Cifas to substantiate invoices (among others packing slips).	Process/guidelines	SSC-ICT	28
AR2024-ARCH-02	4.3.8 Central archiving of performance evidence	Monitor the quality of the performance evidence by taking samples and adjusting where necessary.	Control/monitoring	SSC-ICT	28
AR2024-ARCH-03	4.3.8 Central archiving of performance evidence	Make clear internally, per cost category, which evidence (quality requirements) must be recorded (preferably in Cifas), so that procedures are followed consistently.	Process/guidelines	SSC-ICT	28
AR2024-IT-01	4.3.9 Recurring themes in IT audits	Apply an integral and risk-based approach for the IT management processes of systems.	Governance/ICF	SSC-ICT	29
AR2024-IT-02	4.3.9 Recurring themes in IT audits	Draw up a uniform standards framework, including clear minimum security requirements, and test periodically whether these standards are met.	Standards framework/ compliance	SSC-ICT	29

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
AR2024-IT-03	4.3.9 Recurring themes in IT audits	Process the lessons learned from earlier audits actively into policy, and set up a continuous improvement cycle in which vulnerabilities are systematically evaluated and addressed.	PDCA/ improvement cycle	SSC-ICT	30
AR2024-IT-04	4.3.9 Recurring themes in IT audits	Define central ownership and take security by design as a starting point, in order to strengthen the IT management processes structurally and to minimise future findings.	Governance	SSC-ICT	30
INT2025	*Interim Audit Report 2025 – Shared Service Centre ICT (SSC-ICT)*				
INT2025-DFA-01	1.1 Signing of DFAs	Keep communicating the importance of the new DFA working method and the agreements made in the Administrative Consultation (BO) with customers (among others through the Customer Council).	Communication/ governance	SSC-ICT + customers	10
INT2025-DFA-02	1.1 Signing of DFAs	Hold the Ministry of Foreign Affairs and customers within BZK accountable for signing the DFA on time and for providing the purchase-order number on time, for timely invoicing.	Governance/ agreements	SSC-ICT + BZ + BZK customers	10
INT2025-DFA-03	1.2 Revenue process DFA / DFA validation process	Make the compliance and demonstrability of the followed process steps visible, so that they are internally accountable and verifiable; set up internal control and file-building accordingly.	Process/controls	SSC-ICT	12
INT2025-DFA-04	1.2 Revenue process DFA / relation with customers	Keep ensuring close involvement, coordination and communication with customers, and respond on time to signals and requests.	Communication/ process	SSC-ICT + customers	12
INT2025-ALLOC-01	Non-allocable service provision	Actively fulfil the monitoring role to limit fluctuations and promote transparency; signal variations and address them adequately.	Monitoring/ governance	SSC-ICT	13
INT2025-ALLOC-02	Non-allocable service provision	Keep giving priority to cleaning up accounts that lead to non-allocable service provision, and adjust the source administrations accordingly.	Data/process	SSC-ICT	13
INT2025-DEV-01	1.3 Non-returned devices	Take additional measures for the remaining percentage of non-returned devices.	Policy/process	SSC-ICT + customers	15
INT2025-DEV-02	1.3 Non-returned devices	Set, in consultation with the directing parties, a framework for the maximum permissible percentage of non-returned devices.	Governance/ frameworks	SSC-ICT + directing parties	15
INT2025-DEV-03	1.3 Non-returned devices	Organise explicit (administrative) decision-making on the process and the follow-up actions for the remaining part, so that the lifecycle-management process stays transparent and manageable.	Governance/ decision-making	SSC-ICT + directing parties	15
INT2025-ASSET-01	1.4 Assets outside the regular process	Make formal agreements with all owners of the assets about responsibility for management, so that an integral asset-management process arises; emphasise the responsibility of the owner.	Governance/ agreements	SSC-ICT + asset owners	17
INT2025-CHAIN-01	2.3 Chain cooperation / service quality	Keep cooperating well with chain partners and set priorities in consultation for the optimisation of systems and process steps (among others invoice processing).	Governance/ cooperation	SSC-ICT + chain partners (OBF, O&P Rijk) + BZK	22
INT2025-ARCH-01	2.4 Central archiving of performance evidence	Ensure that administrative procedures around performance evidence are followed consistently (recording evidence such as packing slips on time and in good quality), in order to minimise the risks of incorrect or unlawful payments.	Process/controls	SSC-ICT	23

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
INT2025-ARCH-02	2.4 Central archiving of performance evidence	Make clear internally, per cost category, which evidence (preferably in Cifas) must be recorded; more explicit recording of approval and evidence becomes the norm.	Guidelines	SSC-ICT	23
INT2025-TRANS-01	2.6 Topdesk migration / safeguarding of the transformation	Plan formal review moments within the line, in order to evaluate the effects of safeguarded transformation activities.	Governance/ monitoring	SSC-ICT	26
INT2025-TRANS-02	2.6 Topdesk migration / safeguarding of the transformation	Include a risk analysis at handover, in order to increase awareness of risks and control.	Risk management	SSC-ICT	26
INT2025-TRANS-03	2.6 Topdesk migration / safeguarding of the transformation	Ensure central storage of discharge forms, so that they are available for possible audits.	Documentation/ archiving	SSC-ICT	26
INT2025-TRANS-04	2.6 Topdesk migration / safeguarding of the transformation	Safeguard structural financing of the safeguarded activities: are the activities included in regular budgets?	Financing/ governance	SSC-ICT + owner	26
INT2025-TRANS-05	2.6 Topdesk migration / safeguarding of the transformation	Make a final account over 2025 with achieved goals, lessons learned, safeguarding and recommendations.	Accountability	SSC-ICT	26
INT2025-IT-01	2.8 Recurring themes in IT audits	Apply an integral and risk-based approach for the IT management processes of systems.	Governance/ICF	SSC-ICT	27
INT2025-IT-02	2.8 Recurring themes in IT audits	Draw up a uniform standards framework with minimum security requirements, and test periodically whether these standards are met.	Standards framework/ compliance	SSC-ICT	27
INT2025-DBA-01	2.1/2.2 Compliance with the DBA Act (outlook)	Anticipate stricter enforcement from 2026 by updating, at the latest by the end of Q1 2026, the action plan aimed at full compliance with the DBA Act ("Wet DBA").	Action plan	SSC-ICT	21
INT2025-TOP-01	2.6 Topdesk migration	Pay extra attention to a complete and correct data transfer during the migration to the new Topdesk asset module.	System migration	SSC-ICT	25
INT2025-TOP-02	2.6 Topdesk migration	Minimise the period in which the system is unavailable during the migration (process continuity).	Continuity	SSC-ICT	25
INT2025-TOP-03	2.6 Topdesk migration	Ensure timely training of the end users for a smooth transition.	People/training	SSC-ICT	25
LIC2020	*License Management SSC-ICT*				
LIC2020-6.2	Recommendations 6.2	6.2 Draw up a license policy. To arrive at a license policy, the following steps can be taken: appoint one (or more) organisational unit that is responsible for drawing up the license policy; [text in source cut off]	Policy/process	SSC-ICT	15
LIC2020-6.3	Recommendations 6.3	6.3 Define license management and assign the responsibility. Define what SSC-ICT understands by license management and which tasks and responsibilities the various organisational units fulfil within it. Then determine which organisational unit is best equipped and [text in source cut off]	Policy/process	SSC-ICT	15

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
LIC2020-6.4	Recommendations 6.4	6.4 Introduce periodic consultation. To align the processes better, introducing periodic consultation between all parties involved in license management could help. During this consultation it is important that all involved organisational units think and [text in source cut off]	Policy/process	SSC-ICT	15
LIC2020-6.5	Recommendations 6.5	6.5 Consider introducing a workflow management system. The introduction of a workflow management system can possibly also contribute to an improvement in the alignment of the processes. A workflow management system gives insight into the status and progress of [text in source cut off]	System/control	SSC-ICT	15
LIC2020-6.6	Recommendations 6.6	6.6 Formalise the advisory role of the license desk. The License Desk has the insight into the numbers and types of available licenses and has the license-technical knowledge to give advice, for example on architecture questions. It would therefore be logical to [text in source cut off]	Policy/process	SSC-ICT	15
LIC2020-6.7	Recommendations 6.7	6.7 Carry out structural analyses that can lead to cost savings. Given the insights and possible cost savings that the analyses of the License Desk have yielded, it is recommended to carry out such analyses structurally. It is quite possible that more analyses can [text in source cut off]	Policy/process	SSC-ICT	15
LIC2020-6.8	Recommendations 6.8	6.8 Have the license rates in the PDC (Product and Services Catalogue, "Producten- en Dienstencatalogus") checked by the License Desk. Submitting the license rates to the License Desk for checking before publication in the PDC is a possibility that can be realised fairly easily, by including this step in the authorisation process.	System/control	SSC-ICT	15
LIC2020-6.9	Recommendations 6.9	6.9 Ensure correct and complete documentation of the building blocks of the products and services. A first requirement for detailed budgets is that the building blocks used to build a product and/or service are correctly documented and [text in source cut off]	Policy/process	Central government (frameworks) + SSC-ICT	15
LIC2020-6.10	Recommendations 6.10	6.10 Ensure that the costs and revenue of licenses are visible from the Business Services and the Technical Services. A translation key, based for example on the distribution key for allocating license costs as used in the integral cost-price model, would [text in source cut off]	System/control	SSC-ICT	15
LIC2020-6.11	Recommendations 6.11	6.11 Use the project code to gain insight into the costs and revenue of licenses of custom products. Drawing up budgets for the license costs and the related revenue for the custom category should be fairly simple, for example on the basis of [text in source cut off]	System/control	SSC-ICT	15
LIC2020-6.12	Recommendations 6.12	6.12 Formalise the obtaining of access to systems and networks. To deploy the tooling for measuring and monitoring the use of licenses, access to the client's systems and networks is necessary. What can help is that this is arranged formally when concluding the license agreement [text in source cut off]	Policy/process	SSC-ICT	15

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
LIC2020-6.13	Recommendations 6.13	6.13 Process the license costs and the revenue on the same basis. To gain insight from the revenue into the financial coverage of the costs, a choice must be made in the basis for booking the revenue and the costs, so that they are mutually comparable. To [text in source cut off]	System/control	SSC-ICT	15
LIC2020-6.14	Recommendations 6.14	6.14 Use a standard layout for the description of the booking line when purchasing licenses in DigiInkoop. A possible solution for the lack of (correct) information in the booking line when purchasing licenses would be a standard layout of the [text in source cut off]	System/control	SSC-ICT	15
LIC2020-6.15	Recommendations 6.15	6.15 Safeguard the use of the correct Technical Services codes. The reason that the check on the use of the correct Technical Services codes does not take place is that Finance & Control does not have the necessary information to assess this. A possible solution could therefore be [text in source cut off]	System/control	SSC-ICT	15
KPM2020	*Study of the Cost-Price Model SSC-ICT*				
KPM2020-01	Follow-up steps – post-calculation/ simplification	Carry out a post-calculation in 2021 (in any case over 2021) to check whether simplification of the cost-price model is desirable and which simplifications are useful (without loss of the desired management information).	Analysis/PDCA	SSC-ICT	20
KPM2020-02	Follow-up steps – management & procedures	Record procedures for the use, management and changes of the cost-price model (including control measures, roles and responsibilities) and include more direct references to source documents / source data in the cost-price model, in order to make the model more robust.	Process/governance	SSC-ICT	20
KPM2020-03	Follow-up steps – safeguarding input	Investigate whether manual input can be replaced by automated interfaces with source applications, in order to reduce the risk of errors in the input information.	System integration	SSC-ICT	20
LCM2024	*Robustness of LCM cost inclusion in the cost-price model SSC-ICT*				
LCM2024-01	1.4 Courses of action	Make visible how the annual quotation and its reflection in the Dossier of Financial Agreements (DFA) follow from the Service Provision Agreements (DVA, "Dienstverleningsafspraken").	Transparency/ finance	SSC-ICT + BZ	5
LCM2024-02	1.4 Courses of action	Indicate, per generic but especially also per client-specific Business Service, how the quoted amount is built up and which (technical) services fall under it (including definition of the TS, rate, quantities).	Transparency/ finance	SSC-ICT + BZ	5
LCM2024-03	1.4 Courses of action	Give (if desired for the insight that BZ needs to assess the quotation) a further explanation for the Technical Services that are allocated from the stack.	Transparency/ finance	SSC-ICT + BZ	5
LCM2024-04	1.4 Courses of action	Give insight into the IT landscape (and its age) that is specific to BZ and for which SSC-ICT provides the LCM. SSC-ICT has started with a tool (the DFA Generator) to unlock the CMDB for BZ; ensure that this tool becomes available soon.	System/ transparency	SSC-ICT + BZ	5

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
LCM2024-05	1.4 Courses of action	Improve the robustness of quotations by carrying out post-calculations and using these to further improve the plausibility of the team budgets.	PDCA/finance	SSC-ICT	5
LCM2024-06	1.4 Courses of action	Make the post-calculation specific to BZ and share it with BZ.	Transparency/ finance	SSC-ICT + BZ	5
LCM2024-07	1.4 Courses of action	Check, in the elaboration of the Vision on Life Cycle Management of SSC-ICT, whether the LCM by now covers the entire IT landscape of SSC-ICT.	Governance/ monitoring	SSC-ICT	5
ACICT2025	*AcICT Advice: Transition of ICT Service Provision (BZ)*				
ACICT2025-00	Main advice	Do not push through the exit decision (BZ leaving SSC-ICT) now; first give top priority to the urgent improvement of the current ICT service provision.	Decision-making	BZ + SSC-ICT + BZK	1
ACICT2025-01	Advice 1 – urgent security risks	Jointly take measures to reduce the urgent security risks in the short term, and report on the progress every three months to the Minister of Foreign Affairs and the State Secretary for Digitalisation and Kingdom Relations. *Note: Details of the measures are in a confidential annex (not in the public advice).*	Steering/monitoring	BZ + SSC-ICT	5
ACICT2025-02a	Advice 2 – governance of BZ / risk management	Set up a closed risk-management process: map the critical processes / information flows and the threat level, and on that basis define the specific security need.	Risk management	BZ	5
ACICT2025-02b	Advice 2 – governance of BZ / requirements for service providers	Improve the process for specifying requirements for service providers: draw up clear requirements based on the security need, in which the international context and the threat picture are weighed in.	Process/ requirements	BZ	5
ACICT2025-02c	Advice 2 – direction / contract management	Improve the contract management of BZ, so that agreements are made, monitored and adjusted in a structured way; ensure central insight into which ICT services are obtained from which supplier, and into compliance with agreements by both the supplier and BZ.	Process/governance	BZ	5
ACICT2025-02d	Advice 2 – role consistency / task demarcation	Act with role consistency: address SSC-ICT and other service providers as a predictable partner, avoid ad-hoc steering, and realise clear task demarcation, including the division of responsibilities for central coordination and prioritisation.	Roles & responsibilities	BZ	5
ACICT2025-02e	Advice 2 – internal security processes	Ensure effective compliance with the requirements from the BZ security baseline and a structured information-security process within BZ; improve incident handling, with root-cause investigation in all cases.	Process/operations	BZ	6
ACICT2025-02f	Advice 2 – accreditation policy	Monitor compliance with the rules from the accreditation policy; when assessing systems that process classified information, take at least the relevant (technical) information-security frameworks into account.	Compliance/process	BZ	6
ACICT2025-03a	Advice 3 – government-wide approach (top priority)	Realise, with high priority and from BZK, a government-wide approach for security against state actors (BBN3), and assign the commissioning role to the Secretaries-General Consultation (SGO).	Governance	BZK + CIO-Rijk + departments	6

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
ACICT2025-03b	Advice 3 – government-wide framework ($\leq$ 6 months)	Realise, under the direction of CIO-Rijk and within six months, an elaboration of a government-wide framework for security against state actors; mobilise specialist knowledge (AIVD, MIVD, NCSC, BZ, other ministries).	Framework development	BZK + CIO-Rijk + departments	6
ACICT2025-03c	Advice 3 – ISO/BIO & risk-analysis guidelines	Take ISO27001 as a starting point and make the government-wide framework additional to it (in line with the renewal of the BIO following NIS2); include guidelines for the risk-analysis process and the selection of security measures.	Standards framework	BZK + CIO-Rijk	6
ACICT2025-03d	Advice 3 – measures & detection	Elaborate a selection of applicable security measures (among others crypto solutions, data diodes, virtualisation) and keep these up to date; include detection rules and apply the Assume Breach principle.	Measures/detection	BZK + CIO-Rijk + implementing parties	6
ACICT2025-03e	Advice 3 – government-wide BBN3 service provision	Realise BBN3 ICT service provision government-wide, step by step: determine which government party develops this and facilitate government-wide availability of at least one BBN3 workplace; cooperation between ICT service providers and data centres is crucial.	Service provision/ governance	BZK + CIO-Rijk + central-government service providers	6
ACICT2025-04a	Advice 4 – evaluation	Evaluate after two years (or earlier if opportune) the follow-up of the advice and the effectiveness of the measures, at both BZ and SSC-ICT; select an independent party for this.	Evaluation	BZ + SSC-ICT + BZK	6
ACICT2025-04b	Advice 4 – follow-up steps / exit decision	Determine the follow-up steps on the basis of the evaluation, including how to deal with the postponed exit decision; only elaborate a transition scenario then, so that the gains from advice 2 and 3 can be taken into account.	Decision-making	BZ + SSC-ICT + BZK	6
TK1417	*Parliamentary letter 26 643, no. 1417 (follow-up to AcICT advice)*				
TK1417-ACICT-1	Repetition of AcICT recommendations (letter)	The letter repeats AcICT recommendation 1 (see ACICT2025-01 to -04): ensure that SSC-ICT and BZ reduce the urgent security risks. *Note: The letter also mentions programme SHIELD as elaboration of the follow-up.*	Communication/ policy	BZ + SSC-ICT + BZK	2
TK1417-ACICT-2	Repetition of AcICT recommendations (letter)	The letter repeats AcICT recommendation 2 (see ACICT2025-01 to -04): ensure that the governance at BZ fundamentally improves. *Note: The letter also mentions programme SHIELD as elaboration of the follow-up.*	Communication/ policy	BZ + SSC-ICT + BZK	2
TK1417-ACICT-3	Repetition of AcICT recommendations (letter)	The letter repeats AcICT recommendation 3 (see ACICT2025-01 to -04): realise, from BZK, a government-wide approach for security against state actors. *Note: The letter also mentions programme SHIELD as elaboration of the follow-up.*	Communication/ policy	BZ + SSC-ICT + BZK	2
TK1417-ACICT-4	Repetition of AcICT recommendations (letter)	The letter repeats AcICT recommendation 4 (see ACICT2025-01 to -04): evaluate the follow-up of the advice jointly after 2 years. *Note: The letter also mentions programme SHIELD as elaboration of the follow-up.*	Communication/ policy	BZ + SSC-ICT + BZK	2

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
BIT2017	*Final BIT Advice, project DWR Next*				
BIT2017-DWRN-01a	Advice 1	Elaborate with urgency a fallback scenario for the relocation of I&M to Rijnstraat 8, so that the relocation does not depend on DWR Next (taking DWR 1.3 as starting point). Determine the final moment at which this scenario must be chosen. *Note: Specifically aimed at I&M as launching customer.*	Relocation of I&M to Rijnstraat 8 + workplace continuity (DWR 1.3 vs DWR Next)	SSC-ICT + customers	5
BIT2017-DWRN-01b	Advice 1	If the fallback scenario is chosen: organise this in a separate project and include all ICT-related work that is needed at SSC-ICT and I&M for the relocation to go well.	Fallback project organisation around the I&M relocation	SSC-ICT + customers	5
BIT2017-DWRN-02	Advice 2	Apply, for the completion / go-live of DWR Next, a firm project-based approach with predictable timelines, predictable service provision, clear roles/responsibilities and clear go/no-go criteria (SSC-ICT and the departments).	DWR Next (digital workplace) – development and implementation	SSC-ICT + customers	5
BIT2017-DWRN-02a	Advice 2a	Make an overarching planning on the basis of a concrete and complete functional and technical product description (divided into releases). Include all relevant dependencies (ODC/infrastructure, network adjustments at locations, application migration, orders, taking into management and acceptance). Update it regularly and assign maintenance to a firm programme manager.	Planning, release and dependency management DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-02b	Advice 2b	Name all stakeholders explicitly and completely (not only customers/ users, but also infrastructure managers and functional managers of applications). Record acceptance criteria per stakeholder, including departmental exceptions on privacy impact and security, and set up an acceptance process.	Stakeholder and acceptance management DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-02c	Advice 2c	Determine, with the manager of building-bound facilities (often SSC-ICT), whether the networks in the buildings meet the connection conditions, and have these adjusted if necessary.	Building-bound network facilities for DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-02d	Advice 2d	Make clear agreements about the testing of DWR Next with all involved parties (among others functional managers, ODC, managers, suppliers, customers and users), and use the product description as a starting point.	Test and acceptance approach DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-02e	Advice 2e	Use the product description, the acceptance criteria and the building/ network conditions as a basis for reporting, steering and go/no-go decisions, both at SSC-ICT and at the departments.	Steering & decision-making (go/no-go) DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-03	Advice 3	Actively monitor the preparations by the departments; successful implementation of DWR Next depends on timely fulfilment of preconditions by the customers.	Customer preparation for DWR Next	SSC-ICT + customers	5
BIT2017-DWRN-03a	Advice 3a	Have the departments draw up an issuance policy and choose – within the device list offered by SSC-ICT – a device mix (cloudbooks/tablets/ laptops/fixed) that fits their budget and needs.	Device portfolio/ issuance policy for DWR Next	SSC-ICT + customers	5

continued on next page

Table E.2 – continued from previous page

ID	Finding / paragraph	Recommendation	Object	Owner	Pg.
BIT2017-DWRN-03b	Advice 3b	Have the departments draw up a complete list of applications to be adjusted; direct the suppliers to make applications suitable for Windows 10 and virtualisation; and ensure that functional managers are available for testing.	Application migration to Windows 10/Citrix	SSC-ICT + customers	6
BIT2017-DWRN-03c	Advice 3c	Support the departments actively in carrying out the above work (issuance policy, application adjustment and testing).	Support/enablement of customers for DWR Next	SSC-ICT + customers	6
BIT2017-DWRN-03d	Advice 3d	Ask the departments to draw up a planning for their own work and adjust the planning of SSC-ICT accordingly; on that basis, make agreements about service provision and planning.	Alignment of planning SSC-ICT ↔ departments (DWR Next)	SSC-ICT + customers	6

Appendix F

Workshop materials

The workshop materials were used in Dutch. They have been translated to English for this appendix.

This appendix contains the materials used in the validation workshop (Chapter 7): the workshop outline, the case used in the exercise, the description of what happened between intake and the review moment, and the blank register template that the participants filled in. The questionnaire that the participants completed afterwards is given separately in Appendix C.

Workshop outline

Goal

The workshop has one goal: to validate whether the designed follow-up and governance model works in practice. Validation here means two things at the same time. First, usability: can the people who would use the model in their work actually work with it? This is tested by letting the participants fill in the register themselves on a real case. Second, acceptance: would they want to use the model in their work? This is tested with a questionnaire based on the Method Evaluation Model [26].

Setup and participants

The session is suitable for two to about ten participants. The participants work in groups of two or three people, depending on the number present. The register is filled in together on one laptop per group. Each participant contributes from his or her own expertise. The session lasts 60 minutes. The facilitator facilitates; one observer follows the groups and keeps observation notes in the Excel register.

Time	Block	Content
0–5	Arrival	Participants arrive. The materials are sent by e-mail at this moment so that they can open the file on their own laptop. Short consent for recording (spoken, captured in the recording itself).
5–15	Introduction	Short explanation: the goal of the session, what the model is, how it works (presentation slide), and how the two steps of the exercise run. At most 10 minutes.
15–40	Exercise	The groups fill in the register on the case in two steps. Step 1 (10 min): intake from zero (May 2025). Step 2 (10 min): they continue with what they made themselves and apply the append-only rules with new case information.
40–45	Questionnaire	Each participant fills in the questionnaire individually, before the plenary discussion, so that opinions do not influence each other.
45–60	Discussion & close	Plenary discussion of the exercise and the questionnaire.

Role of the facilitator

- Give no more information than is in the case description.
- Whether a group asks for clarification, and which questions they ask, is itself observation material.
- If a question is really needed in order to continue (for example, an unclear field in the template), give a short neutral clarification. The facilitator does not note this; the observer records which group asked what and which clarification was given.
- Do not try to defend the model; stay open to all remarks and suggestions.

Role of the observer

The observer follows the groups during the exercise and the discussion. The main things to watch are:

- Which fields are filled in quickly, and which fields raise discussion?
- Which fields are skipped, and why?
- Literal statements where possible.

Questions for the plenary discussion

The questionnaire has already been filled in at this point. The discussion is open and is about the experience with the model, not about individual questionnaire answers. Two questions are used:

- What was your experience with the model while filling it in?
- To what extent does this model solve the current problems around follow-up, as you know them in your work?

Case description

Context

The Auditdienst Rijk (ADR) audits SSC-ICT every year. In May 2025 the Audit Report 2024 appears, which, under the heading "recurring themes in IT audits" ("rode draden IT-audits"), makes a number of recommendations about the way SSC-ICT organises its IT control. One of these recommendations concerns drawing up a uniform set of standards for information security. This workshop follows that one recommendation from receipt in May 2025 to the half-yearly review moment in December 2025.

The following information is relevant for filling in the register and does not follow from the recommendation itself:

- The audit report is classified internal (visible to the DT and direct staff).
- The internal monitoring of audit findings normally runs through the CISO, who reports monthly to the CTO.
- The recommendation comes from the yearly ADR audit cycle; the ADR returns every year and re-assesses open items.

Where the case gives no information, fill in what you consider most logical from your own expertise. There is no "right" answer; it is precisely your reasoning that is of interest for this research.

The recommendation as received

Recommendation AR2024-IT-02 (from the Audit Report 2024 of the Auditdienst Rijk, May 2025):
"Establish a uniform set of standards, including clear minimum security requirements, and test periodically whether the standards are met."
(Original Dutch: *"Stel een uniform normenkader op, inclusief duidelijke minimale beveiligingseisen, en toets periodiek of aan die normen wordt voldaan."*)
Level: strategic **Theme:** IT control & security **Target party:** SSC-ICT

Properties of this recommendation

Three remarks are relevant for this recommendation:

- The recommendation is formulated at a high level of abstraction: a "uniform set of standards" can be worked out in several ways.
- The recommendation partly touches SSC-ICT's own sphere of influence and partly government-wide frameworks that CIO Rijk is responsible for (the renewal of the BIO, the Cybersecurity Act / NIS2).
- Activity around information security is already running that partly overlaps (among others a security programme with its own steering group).

What you will do in this workshop

You will process this recommendation in the register in two steps, at different stages.

- In Step 1 your group fills in the register at the moment of receipt (May 2025): you register the recommendation, decide the owner and the priority, and record what you expect to happen.
- In Step 2 your group receives a short description of what happened between May and December 2025. You then update the register to the review moment in December 2025.

Help with filling in

This recommendation has a part that SSC-ICT can pick up itself (the internal part) and a part that depends on another organisation, CIO Rijk (the external part). You keep this as one entry in the register. You record the external part in the field "Cross-organisational owner". If the internal work cannot continue because it is waiting on that external part, you can show this through the implementation status "blocked by external dependency" with a short status reason. The part that SSC-ICT does itself you follow through the normal status fields.

You derive the review intensity (high / medium / low) from the priority: the heavier the three priority scores weigh together, the higher the review intensity.

After filling in the first round, send the document as it stands to the researcher by e-mail before continuing with Step 2.

Step 2 — What happened between intake and the review moment?

May 2025 – December 2025. Read this carefully before you begin Step 2.

- At the end of May 2025 the CTO first had a decomposition carried out by the directiesecretaris and the CISO. The recommendation was split into:
  - an internal part: drawing up an SSC-ICT set of standards with minimum security requirements, plus a testing cycle;
  - an external part: alignment with CIO Rijk on what becomes "government-wide" and what SSC-ICT may fill in independently.
- After the decomposition the intake checks were run again: Concrete yes, Addressed yes, Executable yes. The flag "decomposition needed" was removed. The CTO was confirmed as the definitive owner on 12 June 2025. (Update the intake-quality fields and the flag to this new situation, and note briefly in the lessons log that this changed after decomposition, so that the original assessment from Step 1 stays visible.)
- The CTO instructed the CISO to draw up a draft set of standards based on the BIO and the NIS2 renewal. A first draft exists and was discussed internally in the CISO council in August 2025.
- Alignment with CIO Rijk on what should be "government-wide" and what SSC-ICT may fill in itself started in September 2025, but runs more slowly than expected. CIO Rijk indicates that the government-wide framework will only be available in draft in Q1 2026.

- Internally, the first test of the draft has not yet started, because there is uncertainty about which standard becomes leading: the organisation's own SSC-ICT version or the government-wide framework that is coming.
- A receiving ministry asked in the meantime for clarity about which standards will apply to their systems. This has increased the visibility of the topic, but also the pressure.
- It is now 4 December 2025. The half-yearly DT review is scheduled. You update the register to this moment.

Register template

The participants filled in the register template below, in Word on their laptops. It is the participant-facing form of the schema specified in Section 5.1.1 (Table 5.1).

V = mandatory field; C = conditional (fill in only if applicable); O = optional. Append = append-only field: on a change, the old value stays and a new line is added below it with [date].

Field	**Req.**	**Value**
Identification		
Identifier *(auto-ID, e.g. AANB-2025-XXX)*	V	
Source *(ADR / Rekenkamer / ACICT / internal advice / other)*	V	
Source report *(title + date of issue)*	V	
Recommendation text *(verbatim as in the source report)*	V	
Date received	V	
Level / theme *(strategic / tactical / operational + theme)*	V	
Confidentiality *(per classification scheme)*	V	
Intake quality		
Concrete? *(specific enough to act on? yes / no)*	V	
Addressed? *(can a role or function be named? yes / no)*	V	
Executable? *(realistic given the dependencies? yes / no)*	V	
Flag "decomposition needed" *(set if one or more attributes are "no")*	V	
Ownership		
Owner (role) *(a role or function, not a person's name)*	V	
Escalation path *(decision body, e.g. DT, audit committee)*	V	
Cross-organisational owner *(only if there is an external dependency)*	C	
Status and planning (append-only on change, shown as "[date] text")		
Follow-up status *(not yet handled / acknowledged / decided to act / decided not to act / superseded)*	V *append*	
Implementation status *(not started / in progress / paused / blocked by external dependency / delivered)*	V *append*	
Impact status *(not yet measurable / under measurement / effect observed / no effect observed)*	V *append*	
Status reason *(fill in only if the status is non-default)*	C	
Lead time (target date) *(when should it be done?)*	V *append*	
Next review date *(when does this entry come up again?)*	V *append*	
Review intensity *(high / medium / low — derived from priority)*	V	
Internal monitoring *(who monitors, with what cadence)*	V	
Link to plan / work item *(reference; may still be open at intake)*	V	
Priority (3D) *(risk & compliance / strategic fit / feasibility — each H/M/L)*	V *append*	
Priority rationale *(one line: why this score)*	C *append*	

Field	Req.	Value
Learning		
Ex-ante expectation *(expected outcome, expected risks, key assumptions. NOT overwritten later.)*	V	
Ex-post observation *(fill in only at review: what deviated from ex-ante?)*	C *append*	
Lessons log *(assumptions, deviations, lessons. One line per entry with [date, owner].)*	O *append*	
Closure (fill in only when closing the entry)		
Closure type *(implemented / residual risk accepted / obsolete due to context / superseded by [ID])*	C	
Closure rationale *(why closed this way, authorised by which owner)*	C	

Appendix G

Glossary

This appendix lists the institutional and SSC-ICT-specific Dutch terms that are used throughout the thesis, with their English translation, for easy lookup. The entries below are short reference descriptions, not definitions. Section G.1 first describes the management structure of SSC-ICT, which is referred to in several chapters.

G.1 Management structure at SSC-ICT

SSC-ICT is an executive agency ("agentschap") of the Dutch central government. The executive board ("directieteam", DT) consists of the CEO, the COO, the CTO, the CFO, and the board secretary ("directiesecretaris"). The management team ("managementteam", MT) consists of the COO and the directors of the three business units: IT Infrastructure Services (BIS), IT Standard Services (BST), and IT Application Services (BAS). The DT and the MT together form the combined board ("directiemanagementteam", DMT). Internal Audit and Advice (IAA) carries out the independent verification step in the audit follow-up cycle.

G.2 Institutional terms

Dutch	English	Note
Algemene Rekenkamer	Netherlands Court of Audit	Independent high council of state that audits the central government.
Rekenkamer	Netherlands Court of Audit (short form)	In the thesis, the short form of the Algemene Rekenkamer.
decentrale rekenkamers	decentral courts of audit	Independent audit bodies of municipalities, provinces, and water boards. Since a 2023 law each of these must have its own (or a joint) independent rekenkamer. They check whether local policy is effective, efficient, and lawful. The thesis uses the short plural "rekenkamers" for these.
Tweede Kamer	House of Representatives	The elected chamber of the Dutch parliament.
Adviescollege ICT-toetsing (ACICT)	Advisory Council on ICT Assessment	Assesses central-government IT projects before they start.
Bureau ICT-toetsing (BIT)	Bureau for ICT Assessment	Statutory predecessor of ACICT.
Auditdienst Rijk (ADR)	Audit Service of the Central Government	Runs the yearly financial and IT audit cycle.

Dutch	English	Note
CIO Rijk	CIO of the central government	Central policy directorate within BZK that sets binding frameworks.
Vereniging van Rekenkamers (VvR)	Association of Courts of Audit	Supporting body for the decentral rekenkamers.
BZK	Ministry of the Interior and Kingdom Relations	Hosts CIO Rijk.
BZ	Ministry of Foreign Affairs	Customer ministry in the case material.
griffie	municipal council clerk	Can track the follow-up of rekenkamer recommendations.
ministerraad	Council of Ministers	Discusses a deviation from an ACICT advice before the response goes to the Tweede Kamer.
Comptabiliteitswet	Government Accounts Act	Anchors the task of the Algemene Rekenkamer.
Gemeentewet	Municipalities Act	Article 185a obliges reporting on rekenkamer follow-up.
Instellingswet Adviescollege ICT-toetsing	ACICT Establishment Act	Established ACICT in 2019.
Besluit CIO-stelsel	CIO-System Decree	Delegates the authorities of CIO Rijk. Revised in 2025.
Wet DBA	DBA Act	Law on the assessment of employment relationships.

G.3 SSC-ICT and governance terms

Dutch	English	Note
agentschap	executive agency	The legal status of SSC-ICT within the central government.
directieteam (DT)	executive board	The CEO, COO, CTO, CFO, and the board secretary.
managementteam (MT)	management team	The COO and the directors of the three business units.
directiemanagementteam (DMT)	combined board	The DT and the MT together.
directiesecretaris	board secretary	Supports a member of the executive board.
BIS	Business Unit IT Infrastructure Services	One of the three business units.
BST	Business Unit IT Standard Services	One of the three business units.
BAS	Business Unit IT Application Services	One of the three business units.
Interne Audit en Advies (IAA)	Internal Audit and Advice	Performs the independent verification step in the audit cycle.
stuurgroep	steering group	Set up to steer the follow-up of a programme.
transformatieboard	transformation board	The governance body for the transformation programme.

Dutch	English	Note
gremia	governance forums	The various coordination bodies through which steering runs.
bestuurlijk overleg	administrative consultation	A cross-organisational governance body.
afnemersberaad	customers' council	A cross-organisational governance body of the customer ministries.
financiële commissie dienstverlening	financial service-provision committee	A cross-organisational governance body.
dienstverleningsafspraken (DVA)	service-provision agreements	Agreements between SSC-ICT and its customers.
Dossier Financiële Afspraken (DFA)	Dossier of Financial Agreements	The financial agreements recorded per customer.
effectketen	effect chain	Lemmens's frame of follow-up, implementation, and impact.
opvolging	follow-up	The first stage of the effect chain.
doorwerking	implementation	The second stage of the effect chain.
kaders	frameworks	Binding policy frameworks set by CIO Rijk.