

Leiden University

ICT in Business and the Public Sector

Designing an IT Operating Model framework for scaling AI initiatives across organisations

Name:	Imran Meziani
Student ID:	s4290402
Date:	4 August 2026
1st supervisor:	Dr. N. Amat Lefort
2nd supervisor:	Dr. C. J. Stettina
Company supervisor:	N. Nguyen

MASTER'S THESIS

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

Abstract

Background: Organisations are investing heavily in artificial intelligence, yet most struggle to move beyond experimentation, with a large share of AI initiatives abandoned after the proof-of-concept phase. This "pilot-to-scale gap" is often organisational rather than purely technical, which points to the IT Operating Model (ITOM) as a useful structure. However, there are still limited empirically grounded insights into which recurring ITOM patterns help organisations scale AI from pilots toward broader organisational deployment.

Aim: This thesis aims to identify the recurring ITOM patterns that practitioners associate with scaling AI, and to synthesise these patterns into an ITOM framework for AI scaling.

Method: The study follows an exploratory qualitative design, combining a review of the ITOM and AI scaling literature with fifteen semi-structured interviews with experienced practitioners involved in scaling, leading, and operationalising AI. The interviews were analysed through a structured cross-interview comparison across six ITOM domains: governance, people & organisation, processes, technology & infrastructure, data, and culture. The resulting framework was subsequently revised through evaluation with five practitioners.

Results: Practitioners treat AI scaling as an operating model transformation rather than a technology rollout, with recurring patterns emerging in every domain, such as top-down mandate with business ownership, an early Centre of Excellence that later shifts toward federated models, redesign of workflows around AI, shared platforms, data readiness as a hard precondition, and adoption driven by AI literacy and visible leadership. The central finding is that these patterns are interdependent: progress in any one domain is bounded by the maturity of the others. These patterns were consolidated into an ITOM framework that organises the six domains into three layers: Direction, Transformation, and Enabling foundations.

Conclusion: Organisational AI scaling is mainly an IT Operating Model challenge, driven by coordination across domains rather than by the AI tooling alone. The resulting framework offers researchers a domain- and dependency-based view of AI scaling, and gives practitioners a structured way to navigate their ITOM choices for AI scaling.

Acknowledgements

This thesis marks the end of my master's programme in ICT in Business at Leiden University, and it would not have been possible without the help of a number of people.

I am very grateful to my first supervisor, Dr. Natalia Amat Lefort, for the time and energy she put into this project. The regular check-ins kept me motivated and helped me stay on track. I also really enjoyed the business courses she taught, which drew on her academic and consulting experience and gave me a practical perspective that helped during the internship and will remain valuable in the future.

I would also like to thank my second supervisor, Dr. Christoph Johann Stettina, who used his experience in business transformation to give sharp and honest feedback. His input pushed me to raise my standards and work harder, and I am grateful to him for organising consulting projects during the master's programme.

I would also like to thank my company supervisor at PwC Netherlands, Nikita Nguyen. She was always there to guide me, read through my work, and give feedback. Beyond the thesis, she made me feel welcome from day one, taught me a lot in a short time, and trusted me enough to bring me along on one of her IT strategy client projects. I am especially grateful for the energy she put into helping me become a good consultant.

I would also like to thank the rest of the Advisory team for the warm welcome and for making my time at PwC so enjoyable.

Between the courses taught by Dr. Amat Lefort, the consulting projects organised by Dr. Stettina, and the training I got from Nikita, I feel well prepared for my next step: I am happy to say that I will be joining the Advisory team as a consultant starting the 1st of September 2026.

Finally, I want to thank my family and friends, who stayed interested and involved throughout the whole process.

Imran Meziani

Table of Contents

Abstract i

Acknowledgements ii

1 Introduction 1
- 1.1 Problem statement . . . 1
- 1.2 Research questions . . . 2
- 1.3 Overview of the thesis . . . 2

2 Background and related work 3
- 2.1 What is AI? . . . 3
  - 2.1.1 Deep learning . . . 3
  - 2.1.2 Large language models . . . 3
  - 2.1.3 Generative AI . . . 4
  - 2.1.4 Agentic AI . . . 4
  - 2.1.5 Implications for the IT Operating Model . . . 4
- 2.2 IT Operating Model . . . 5
  - 2.2.1 Academic frameworks . . . 5
  - 2.2.2 Practitioner frameworks . . . 6
  - 2.2.3 Domain coverage . . . 7
- 2.3 AI scaling . . . 8
  - 2.3.1 The scaling gap . . . 8
  - 2.3.2 Why organisations struggle to scale AI . . . 8
  - 2.3.3 Existing frameworks for AI scaling . . . 10
- 2.4 Synthesis and research gap . . . 10

3 Method 12
- 3.1 Research approach . . . 12
- 3.2 Research context . . . 12
  - 3.2.1 PwC project context . . . 12
  - 3.2.2 Participant and organisational diversity . . . 13
  - 3.2.3 AI context of participants . . . 13
- 3.3 Conceptual frame . . . 13
- 3.4 Data collection . . . 13
  - 3.4.1 Semi-structured interviews . . . 13
  - 3.4.2 Interview structure . . . 13
  - 3.4.3 Sampling strategy and overview of participants . . . 14
  - 3.4.4 Ethical considerations . . . 15
- 3.5 Data analysis . . . 16
  - 3.5.1 Trustworthiness . . . 16
  - 3.5.2 Framework development procedure . . . 16
- 3.6 Evaluation approach . . . 17

4 Findings 18
- 4.1 Introduction . . . 18
- 4.2 Governance . . . 18

4.3 People & Organisation . . . 20
4.4 Processes . . . 22
4.5 Technology & Infrastructure . . . 24
4.6 Data . . . 26
4.7 Culture . . . 28
4.8 Cross-domain dependencies . . . 30

5 Framework development **32**
5.1 Introduction . . . 32
5.2 Framework structure . . . 32
5.3 Use of the framework . . . 33
5.4 Conclusion . . . 33

6 Evaluation **34**
6.1 Evaluation approach . . . 34
6.2 Evaluation criteria . . . 34
6.3 Evaluation participants . . . 34
6.4 Evaluation procedure . . . 35
6.5 Evaluation results . . . 35
6.6 Revised framework (v2) . . . 36

7 Discussion **39**
7.1 Comparing modern AI scaling with the earlier RPA wave . . . 39
7.1.1 Why compare AI scaling with the RPA wave? . . . 39
7.1.2 What remains similar across the two waves? . . . 39
7.1.3 What is different in modern AI scaling? . . . 40
7.1.4 Lessons from the comparison . . . 40

8 Conclusion **42**
8.1 Answering the research questions . . . 42
8.1.1 SQ1 . . . 42
8.1.2 SQ2 . . . 42
8.1.3 SQ3 . . . 42
8.1.4 SQ4 . . . 43
8.1.5 Main RQ . . . 43
8.2 Contributions . . . 43
8.2.1 Academic contribution . . . 43
8.2.2 Practical contribution . . . 44
8.3 Limitations . . . 44
8.4 Future work . . . 44

Bibliography **46**

Appendices **49**

A Interview Guide **49**

B Interview summaries **52**
B.1 Interview 1 summary . . . 52
B.2 Interview 2 summary . . . 54
B.3 Interview 3 summary . . . 56
B.4 Interview 4 summary . . . 58
B.5 Interview 5 summary . . . 61
B.6 Interview 6 summary . . . 63
B.7 Interview 7 summary . . . 65
B.8 Interview 8 summary . . . 67
B.9 Interview 9 summary . . . 69
B.10 Interview 10 summary . . . 72
B.11 Interview 11 summary . . . 75

B.12 Interview 12 summary . . . 78
B.13 Interview 13 summary . . . 81
B.14 Interview 14 summary . . . 83
B.15 Interview 15 summary . . . 86

C Compact interview summaries **90**
C.1 Interview 1 compact summary . . . 90
C.2 Interview 2 compact summary . . . 91
C.3 Interview 3 compact summary . . . 92
C.4 Interview 4 compact summary . . . 94
C.5 Interview 5 compact summary . . . 95
C.6 Interview 6 compact summary . . . 96
C.7 Interview 7 compact summary . . . 97
C.8 Interview 8 compact summary . . . 99
C.9 Interview 9 compact summary . . . 100
C.10 Interview 10 compact summary . . . 101
C.11 Interview 11 compact summary . . . 102
C.12 Interview 12 compact summary . . . 103
C.13 Interview 13 compact summary . . . 105
C.14 Interview 14 compact summary . . . 106
C.15 Interview 15 compact summary . . . 108

D Full-size framework figures **110**
D.1 Full-size condensed framework figure . . . 111
D.2 Full-size expanded framework figure . . . 112
D.3 Full-size staircase framework figure . . . 113

Chapter 1

Introduction

1.1 Problem statement

In recent years, organisations have significantly increased their investments in artificial intelligence (AI), and it has become a strategic priority at board level. Large-scale surveys show that CEOs see technological change and AI as a primary strategic concern, and that many companies are actively working to embed AI into their operations and business models [29].

In this thesis, AI refers to modern learning-based AI technologies such as deep learning, large language models (LLMs), generative AI (GenAI), and agentic AI systems, rather than traditional rule-based systems. These technologies can transform core processes and enable new ways of creating value, but they also introduce new challenges around the operating model [42].

Organisations typically start with pilots, proofs of concept, and other forms of experimentation to learn the potential value in their specific context. Experimentation is a proven approach for learning under uncertainty and supporting innovation [38].

However, moving from pilots to broader organisational deployment remains difficult; research and surveys indicate that many organisations struggle to move beyond pilots and to realise organisation-wide impact from AI use [12, 24]. Estimates of this pilot-to-scale gap vary but are consistently high: a recent Gartner article states that at least 50% of generative AI projects are abandoned after the proof-of-concept phase [7], and other studies report even higher failure rates (discussed in Section 2.3.1).

Furthermore, organisations must proactively align their structure with developments in IT to remain competitive [15], which is even more pressing in the current wave of AI-driven transformation.

The IT Operating Model (ITOM) describes how an organisation structures and runs its IT to deliver services to the business and reach organisational goals, typically by combining organisational structure, processes, people, technology, and governance [1, 41]. In addition to academic literature, various practitioner frameworks propose ITOM models with similar dimensions. Yet there are still limited empirically supported insights into which recurring ITOM patterns enable organisations to scale AI from pilots to broader organisational deployment, and how practitioners relate these patterns to broader organisational impact.

This research addresses that gap by identifying the recurring ITOM patterns that practitioners associate with AI scaling from pilots to broader organisational deployment across ITOM domains. The expected result is an academic contribution by providing empirically grounded patterns across ITOM domains from real-world experts involved with AI scaling. Furthermore, this research provides a practical contribution in the form of an ITOM framework that supports IT leaders, transformation leads, and consultants in shaping their ITOM for AI impact. IT leaders already make strategic decisions based on the ITOM, which makes this framework recognisable and actionable in their context.

1.2 Research questions

RQ: Which recurring ITOM patterns do practitioners associate with AI scaling initiatives from pilots toward broader organisational deployment across ITOM domains, and how can they be synthesised into an ITOM framework for AI scaling?

- **SQ1:** Which domains are most relevant for an ITOM according to the literature?
- **SQ2:** What challenges, enablers, and existing approaches to scaling AI are identified in the literature?
- **SQ3:** Which recurring ITOM patterns emerge from practitioner interviews on AI scaling?
- **SQ4:** How can the findings from practitioner interviews be synthesised into an ITOM framework for AI scaling?

1.3 Overview of the thesis

This thesis develops an ITOM framework for scaling AI initiatives from pilots toward broader organisational deployment. To make the scope clear, three definitions are used throughout this study:

- **AI** refers to the modern learning-based AI technologies introduced in Section 1.1 (deep learning, LLMs, GenAI, and agentic AI systems), excluding traditional rule-based systems.
- **AI scaling** refers to moving AI initiatives from pilots or proofs of concept toward broader organisational deployment. In this study, an initiative is treated as scaling when it is (1) used beyond a single team or pilot context, (2) used in operational processes rather than experimentation environments, and (3) supported by structural changes rather than temporary or ad-hoc effort.
- **ITOM** refers to how an organisation sets up its governance, people & organisation, processes, technology & infrastructure, data, and culture to reach strategic IT goals.

This thesis is structured as follows. Chapter 2 reviews the literature on AI, the ITOM, and AI scaling, and identifies the research gap. Chapter 3 describes the research design, the interview procedure, analysis approach, and the framework development. Chapter 4 presents the recurring patterns found in the interviews. Chapter 5 translates these patterns into the proposed ITOM framework. Chapter 6 presents the evaluation of the framework and the revised version. Chapter 7 interprets the findings in relation to the earlier Robotic Process Automation (RPA) wave. Finally, Chapter 8 answers the research questions, reflects on the contributions and limitations of this study, and proposes future work.

Chapter 2

Background and related work

2.1 What is AI?

This section introduces the learning-based AI technologies that fall within the scope of this thesis and explains how AI has evolved from early rule-based systems to the foundation models that currently support enterprise AI.

The field has gone through different phases. Early systems were based on handwritten rules and expert knowledge. While effective within narrow domains, they were not suited for use cases without predefined logic, and they were expensive to update in complex environments [14]. Statistical and model-based methods introduced learning from data and created many possibilities, but one step further was the emergence of deep learning. Deep learning uses neural networks trained on large datasets, which makes it possible to process unstructured data such as text and images [36]. Foundational models are large pre-trained 'out of the box' models that can be adapted to specific tasks and use cases. This has made generative technologies more accessible and easier to apply across a broader range of tasks and domains [14].

Although this thesis acknowledges the full range of modern AI, the organisational reality is that most enterprises are currently scaling GenAI and LLM-based applications [7, 24], while agentic AI remains largely in early experimentation [42]. This thesis therefore focuses on the ITOM patterns involved in scaling GenAI-driven initiatives, while also capturing the agentic patterns that practitioners increasingly report as organisations move toward more advanced forms of AI deployment.

2.1.1 Deep learning

Deep learning uses multi-layered neural networks to map unstructured inputs such as text, images, or sensor logs into useful information such as classifications, anomaly scores, or forecasts. In day-to-day operations these capabilities support tasks such as early warning in asset maintenance, image recognition, and fraud flagging [36].

Training deep learning models from scratch requires extensive computational and memory resources and large amounts of labelled data. Transfer learning partly mitigates this by transferring knowledge gained from earlier training to the problem at hand, which saves computing resources and increases training efficiency [36]. According to its implementation partner, C3, Shell's predictive-maintenance program is an example at industrial scale: centrally developed models are adapted and deployed across more than 10,000 assets globally [5].

2.1.2 Large language models

Large language models are a category of foundation model built on the Transformer architecture [2]. They are pre-trained on broad text and can understand and generate natural language. Their importance for enterprises comes from their ability to process large amounts of unstructured language data across domains such as business, customer service, healthcare, and legal [2]. Organisations use LLMs for tasks such as information extraction, summarisation, question answering, and coding assistance.

These applications typically rely on instruction-tuned models that are aligned with human feedback and embedded in existing approval flows [2].

LLMs have many limitations, including (1) hallucination, where models generate inaccurate or fabricated information, (2) bias inherited from training data, (3) and a limited understanding of the physical world [2]. These limitations require guardrails and human oversight.

2.1.3 Generative AI

GenAI refers to AI systems that can create new content, including text, images, video, audio, and synthetic data, by learning patterns and structures from training data [33]. Important model types beyond LLMs include Generative Adversarial Networks (GANs) [11], Variational Autoencoders (VAEs), and Diffusion models, each suited for different types of content generation [33].

Generative AI is used for content generation, design, training and educational materials, and synthetic data generation [33]. Foundation models are prone to hallucination and other errors, which creates risks in practical deployment [14]. Practical deployments require reviews, monitoring, and clear rules for where and how generated outputs may be used. This is especially important in regulated or critical settings.

2.1.4 Agentic AI

Agentic AI systems go further than generating content and move to taking actions. Where a generative AI model produces an output in response to a prompt, an agentic system can decompose a goal into subtasks, use tools and APIs, save context across steps, and execute actions toward that goal with limited human action [32].

An important distinction is between AI agents and agentic AI. An AI agent is a single system built around an LLM that uses external tools, follows a series of reasoning steps, and incorporates real-time information to complete a specific task. Agentic AI extends this to multi-agent architectures where multiple specialised agents coordinate, communicate, and dynamically allocate subtasks within a broader workflow [32]. The core components that define AI agents are perception, reasoning, and action. These are enhanced in agentic AI with persistent memory that preserves context across interactions, advanced reasoning and planning capabilities, and orchestration layers where a 'managing' agent coordinates the subordinate agents and assigns roles [32].

In enterprise settings, this is used for systems that can autonomously manage customer support workflows, or coordinate incident response workflows where agents perform threat classification, log analysis, and compliance assessment in parallel [32].

Agents also change the way operations work in practice. Because they can execute tasks in parallel rather than sequentially, cycle times decrease. Agents can continuously process new data which enables them to adapt during runtime and reroute operations or scale capacity. When combined with customer or user data, agents can also personalise interactions and decisions at a level that is difficult to achieve with fixed human resource models [25]. Beyond operational efficiency, McKinsey notes that agents also open new revenue opportunities, such as subscription models, pay-per-use services, and offering internal expertise as AI-powered services to external clients [25].

Because agentic systems autonomously execute actions, guardrails such as permissions, human-in-the-loop approvals, and logging are crucial for safe agent usage. Challenges include hallucination inherited from the underlying LLMs, coordination bottlenecks between agents, unpredictable behaviour, and governance concerns around accountability [32].

2.1.5 Implications for the IT Operating Model

The AI technologies described in this section have an important commonality: enterprise-level impact depends less on the AI technology itself and more on how organisations go about their deployment. The McKinsey Global Survey on AI, based on 1,491 respondents across 101 countries, finds that more than 78% of organisations now use AI in at least one business function, but many organisations still see limited bottom-line impact from their AI-use cases [25]. The survey finds that workflow redesign is the most important factor in realising impact, followed by CEO oversight of AI governance and tracking well-defined KPIs for AI solutions. Less than one-third of organisations currently follow these practices [25].

Generating value from AI requires organisational change, not just technology deployment. Organisations are beginning to centralise risk, compliance, and data governance. Meanwhile, they are using tech talent and AI solutions while driving AI adoption across business functions [25]. Reskilling has started, but is still limited. Larger organisations are shown to move faster than smaller ones [25].

Several examples from industry show how these organisational decisions play out in practice. Morgan Stanley deployed an AI-powered meeting-debrief assistant that generates notes with client consent, drafts follow-up emails for review, and saves records to Salesforce [27]. GitHub's study of Copilot reports that developers complete tasks 55% faster and experience higher satisfaction when the tool is integrated into development [21]. ING's deployment of GenAI across five strategic business domains within less than a year shows how scope, safety, and speed can be combined through central prioritisation [26].

These cases suggest that the challenges of scaling AI are mainly organisational rather than purely technical. This supports the IT Operating Model focus of this thesis.

2.2 IT Operating Model

An operating model describes how an organisation delivers value through the way work is structured and performed. An IT operating model applies this to the IT function. It concerns how elements such as governance, people and organisation, processes, technology, and data are arranged so that IT can deliver services and support organisational goals.

Despite its importance in practice, the ITOM as a concept has received limited attention in the academic literature. This section reviews both academic frameworks and practitioner models to establish a foundation of ITOM domains that will serve as the foundation for the empirical part of this research.

2.2.1 Academic frameworks

This section reviews academic literature that contributes to understanding what an IT operating model consists of and how its components relate to one another. No academic source provides a complete, integrated ITOM definition, but a combination of studies offer perspectives that together cover key dimensions.

Weill and Ross (2004) define IT governance as the specification of decision rights and accountability to encourage desirable behaviour in the use of IT. They identify five domains related to decision making: **IT principles**, **IT architecture**, **IT infrastructure strategies**, **business application needs**, and **IT investment and prioritisation**. For each domain, decision rights can be allocated through different structures, ranging from centralised to decentralised. Their research shows that organisations with well-designed governance achieve higher profitability than competitors with weaker governance [41].

Ross, Weill, and Robertson (2006) build upon this by introducing four operating model archetypes based on two dimensions: **business process standardisation** and **business process integration**. The four types, Diversification, Coordination, Replication, and Unification, each have different requirements on technology, data, and governance. A Unification model implies centralised governance and standardised processes, while a Diversification model allows more autonomous business units with limited shared infrastructure. The choice of archetype shapes how the IT function is structured and therefore, how AI can be scaled across business units [30].

Alter (2013) provides a structured way for identifying the components of IT-enabled systems through Work System Theory. A work system is defined as a system in which human participants and machines perform work using information, technology, and other resources to produce products or services. The framework identifies nine elements: **processes and activities**, **participants**, **information**, **technologies**, **products/services**, **customers**, **environment**, **infrastructure**, and **strategies**. Alter highlights that these elements must be in alignment for the system to function effectively [1]. Wulf, Mettler, and Brenner (2017) build on this foundation to develop a digital services capability model. They organised the capabilities under eight classes: **consumers**, **services**, **processes and activities**, **organisation**, **information**, **technologies & Infrastructure**, **strategies**, and **environment**. While not an ITOM framework, these classes address dimensions that share similarities to those found in ITOM literature, including people and organisation, data management, technology platforms, and partnership strategy [43].

This thesis does not aim to review each domain individually. Instead, the academic sources above serve to identify an initial set of ITOM-relevant dimensions. Section 2.2.2 examines practitioner models, after which a combined mapping is presented at the end of this section.

2.2.2 Practitioner frameworks

Practitioner frameworks offer more comprehensive and integrated ITOM models than the academic literature. This section reviews widely adopted industry frameworks that define ITOM components.

COBIT 2019, developed by ISACA, is a governance and management framework for enterprise IT. It states seven top-level components of a governance system: **processes**, **organisational structures**, **principles, policies, and procedures**, **information**, **culture, ethics, and behaviour**, **people, skills, and competencies**, **services**, **infrastructure and applications** [20]. In addition to these components, COBIT organises 40 governance and management objectives across five areas which covers the whole IT management cycle: governance oversight, alignment and planning, solution development and implementation, service delivery and support, and performance monitoring and assessment. For the purpose of this analysis, only the seven top-level components are mapped to ITOM domains [20].

Forrester's high-performance IT operating model framework gives an IT-specific perspective on the operating model. The framework is built around six layers: **customer value and needs**, **IT offerings**, **configurable capabilities**, **structure**, **governance**, and **leadership** [8]. The first layers focus outward, starting from stakeholder needs and how IT products and services meet those needs. The middle layers are about how IT capabilities are organised and supported in the organisation. The last layers are about governance and leadership, with the focus on trust, transparency in decision-making, and culture. The framework also states three principles: business-IT alignment, trust, and adaptivity to quickly change and innovate capabilities and resources for taking opportunities [8].

Gartner defines an IT operating model as how the IT organisation contributes to achieving the enterprise's strategic objectives. Their framework identifies nine components organised under three core activities: Engage, Enable, and Deliver, with governance and culture as cross-cutting components that work towards strategy execution [10]. The Engage activity contains **performance**, **decision rights**, and **financials**. Described as "how IT engages with the rest of the enterprise, shaping how priorities are determined, funding is obtained, and success is measured". The Enable activity addresses **organisational structure**, **sourcing and ecosystems**, and **talent**, described as "the resourcing of the IT organisation". The Deliver area goes over the **delivery model**, **tools and platforms**, and **workplaces**, described as "how IT will work and the methods they will use". **Governance** and **culture** cut across all three areas, describing the behaviours needed for strategy execution [10].

ITIL 4, developed by Axelos, is a well-known service-management framework. It defines four top-level dimensions of service management: **organisations and people**, **information and technology**, **partners and suppliers**, and **value streams and processes** [3]. ITIL prescribes these dimensions to be considered at an organisational level for IT service delivery. **External factors** such as political, economic, social, technological, legal, and environmental conditions are noted to influence the four dimensions. Additionally, ITIL's service value system includes **governance** as a distinct component. For the purpose of this analysis, only the four dimensions, external factors, and governance are mapped to ITOM domains, and not the 34 individual management practices or the six service value chain activities. They are too specific for the purpose of this review [3].

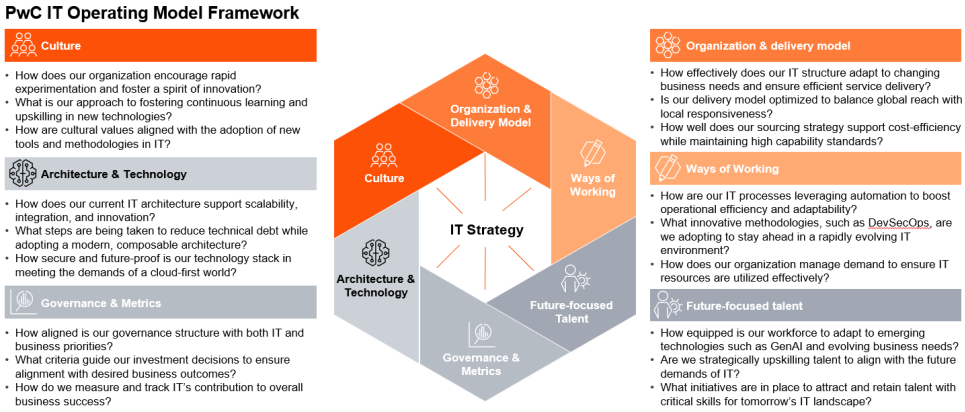


Figure 2.1: PwC Technology Operating Model. Source: PwC (internal).

Lastly, PwC's IT Operating Model Framework positions IT strategy at its centre, surrounded by six domains: **culture**, **architecture & technology**, **governance & metrics**, **organisation & delivery model**, **ways of working**, and **future-focused talent** (see Figure 2.1). Culture covers innovation, continuous learning, and alignment of cultural values with new tools and methodologies. Architecture & technology covers scalability, technical debt reduction, and cloud-first infrastructure. Governance & metrics covers alignment between IT and business priorities, investment decisions, and tracking IT's contribution to business outcomes. Organisation & delivery model covers IT structure, sourcing decisions, and balancing global reach with local responsiveness. Ways of working covers process automation, demand management, and adoption of methodologies such as DevSecOps. Future-focused talent covers workforce readiness for emerging technologies such as GenAI, upskilling, and talent attraction [28].

2.2.3 Domain coverage

Table 2.1: Domain coverage in ITOM frameworks

Domain	Weill & Ross (2004)	Ross et al. (2006)	Alter / Wulf (2013; 2017)	COBIT (2019)	Forrester (2024)	Gartner (2024)	ITIL 4 (2019)	PwC (n.d.)
Technology & Infrastructure	✓	✓	✓	✓	✓	✓	✓	✓
Governance	✓	✓		✓	✓	✓	✓	✓
Processes		✓	✓	✓	✓	✓	✓	✓
People & Organisation			✓	✓	✓	✓	✓	✓
Culture				✓	✓	✓		✓
Data			✓	✓			✓	
Sourcing & Partnerships			✓			✓	✓	
Performance				✓		✓		✓
Environment & Compliance			✓				✓	
Financials	✓					✓		
Strategy			✓					

Table 2.1 maps the ITOM domains identified in the reviewed academic and practitioner sources.

Table 2.2: Domain frequency across ITOM frameworks

Domain	Count
Technology & Infrastructure	8
Governance	7
Processes	7
People & Organisation	6
Culture	4
Data	3
Sourcing & Partnerships	3
Performance	3
Environment & Compliance	2
Financials	2
Strategy	1

Table 2.2 summarises how many times an ITOM-domain was covered in the reviewed sources. Technology & infrastructure, governance, and processes are represented in nearly all reviewed sources (8, 7, and 7 times, respectively). Next, people & organisation follows, appearing in six sources. Together, these four domains form the core of the ITOM across both academic and practitioner literature. Culture is

covered in four sources, appearing mostly in sources that explicitly address behavioural aspects of the IT-organisation, such as COBIT, Gartner, Forrester, and PwC.

Data, sourcing & partnerships, and performance each appear in three sources. Data appears primarily in sources that treat information as a distinct organisational asset, such as COBIT and ITIL. This is noteworthy because data are crucial for AI capabilities. Sourcing & partnerships is most prominent in practitioner frameworks such as Gartner, ITIL, and PwC. Performance appears in COBIT, Gartner, and PwC. Environment and compliance and financials are the least covered, each appearing in only two sources. Strategy appears in only one source, but generally, strategy will be considered as the input which the ITOM works towards.

The practitioner literature appears to address a broader range of ITOM domains than the academic literature. This may be the case because the ITOM concept has largely developed in practical contexts. This mapping is consistent with the findings of Suleman et al., who identified 13 ITOM components through a systematic review of both academic and practitioner literature [35].

The eleven identified domains are used as the basis for selecting the domains used in the empirical part of this research. For this study, the six most represented domains are used: Governance, People & Organisation, Processes, Technology & Infrastructure, Data, and Culture. The first four appear in nearly all reviewed sources and form the core of the ITOM across academic and practitioner literature. Culture is included because it is consistently mentioned in the sources that address behavioural aspects of the IT organisation, and because change and adoption are reported as crucial for AI scaling. Data is included despite its lower frequency because it is repeatedly identified in the literature as a critical precondition for AI scaling (Section 2.3). The remaining domains (Sourcing & Partnerships, Performance, Environment & Compliance, Financials, and Strategy) are treated as part of these six domains rather than as separate dimensions. Section 3 describes how these six domains are used in the research design.

2.3 AI scaling

Section 2.1.5 established that generating value from AI requires organisational change, not just technology deployment. This section examines that challenge in more detail: why organisations struggle to move AI beyond pilots, which challenges the literature identifies, and which frameworks have been proposed for AI scaling. Together with the ITOM domains from Section 2.2, it provides the foundation for the framework design in Chapter 5.

2.3.1 The scaling gap

Despite rising investment and adoption, most organisations struggle to realise returns from AI, and published failure estimates point to the same conclusion even as the exact figures differ. Harvard Business Review notes that some estimates find AI project failure rates as high as 80% [6]. Grebe, Franke, and Heinzl similarly find that only 8% of companies in their global survey report that they can fully leverage AI's potential [12]. A recent Gartner report finds that at least 50% of generative AI projects are abandoned after proof of concept, which is attributed to lack of business value, poor data quality, escalating costs, inadequate risk controls, and poor change management [7].

Although AI-adoption has rapidly increased, fewer than 10% of vertical use cases (within specific business functions) make it past the pilot stage [25]. McKinsey describes this as the "generative AI paradox": horizontal tools such as enterprise copilots scale easily because they require little workflow change, but they deliver gains that are hard to trace to the bottom line. Vertical use cases hold greater potential for measurable impact but require deeper organisational and technical integration [25].

Haefner et al. confirm that competition to leverage AI is fierce, but note that there is a lack of frameworks for implementing and scaling AI successfully [13]. Berndtsson, Ericsson, and Svahn similarly show that scaling data-driven pilot projects requires attention to organisational conditions beyond the pilot phase [4]. The following subsection examines these factors in more detail.

2.3.2 Why organisations struggle to scale AI

The literature finds challenges across several domains. Many technical challenges exist, but multiple studies conclude that organisational factors are the primary bottleneck in AI scaling. Because these

challenges often cut across multiple ITOM domains, they are grouped below into four recurring themes rather than strictly per domain.

Data and infrastructure

Alongside leadership issues, data quality is the most frequently cited barrier to AI scaling. Organisations often lack high-quality data to train effective models [31]. Gartner emphasises that poor data quality affects every department attempting to use generative AI, leading to unreliable outputs and failed retrieval-augmented generation implementations [7]. Beyond quality, organisations face structural data challenges. Khatri and Brown distinguish five data governance decision domains: principles, quality, metadata, access, and lifecycle, which together structure how data is governed as an organisational asset [23].

On the infrastructure side, the inability to deploy and integrate models increases the likelihood of project failure [31]. Based on a case study at a high-tech firm, Tamburri, Tonnarelli, and van Hillegersberg identify infrastructure- and pipeline-related challenges in MLOps adoption [37].

Talent and capabilities

Scaling AI requires a combination of skills that most organisations do not yet have. While many companies have invested in talent, they often lack the engineers needed to independently deploy and maintain models in production [31]. Industry evidence points in the same direction, especially regarding the difficulty of hiring and retaining AI-related roles [24].

Sjödin et al. identify three interdependent capabilities: data pipeline, algorithm development, and AI democratisation. [34]. Their empirical research in manufacturing firms shows that these capabilities must improve together; investing in one while neglecting the others causes scaling initiatives to slow down [34].

Leadership, strategy, and governance

A recurring finding is that AI scaling fails when it lacks strategic direction. Gartner identifies the absence of clear business value as the most fundamental failure point: organisations deploy AI broadly without prioritisation or success metrics [7]. The McKinsey survey finds that CEO oversight of AI governance is one of the factors most associated with bottom-line impact. However, only 28% of respondents report that their CEO is responsible for overseeing AI governance [24]. Similarly, fewer than 30% of companies report that their CEO directly sponsors their AI agenda, contributing to fragmented initiatives with limited enterprise-level coordination [25].

Governance also extends to responsible AI. Gartner warns that treating responsible AI as an afterthought exposes organisations to regulatory risk and project shutdowns [7]. Regulatory pressure such as the EU AI Act reinforces this, making clear ownership and oversight of AI a governance requirement rather than an optional practice.

Process, culture, and change management

As noted in Section 2.1.5, workflow redesign is the organisational attribute most strongly associated with EBIT impact from AI, yet only 21% of organisations report having fundamentally redesigned at least some workflows [24]. This suggests that many organisations insert AI into existing processes rather than adapting those processes to exploit AI capabilities.

Change management is equally important. Gartner identifies it as a key failure point: without it, even technically successful AI tools see low adoption [7]. Grebe, Franke, and Heinzl find that leading companies address this by defining use cases collaboratively with business units and embedding AI into existing value chains rather than treating it as a stand-alone technology initiative [12]. Centres of excellence have accelerated experimentation in many organisations, but they often operate independently from core IT, data, and business functions, which makes their initiatives difficult to scale [25]. Taken together, these challenges show that scaling AI is as much an organisational challenge as a technical one, which is also reflected in the frameworks reviewed in the next subsection.

2.3.3 Existing frameworks for AI scaling

Haefner et al. define a three-phase framework based on an analysis of 22 AI practitioners across industries [13]. The framework identifies both technical components (data pipeline, infrastructure, and AI models) and social components (AI growth vision, AI capabilities, and organisational structure).

The **proving the concept** phase includes assembling data for initial use cases, setting up simple infrastructure, and experimenting with suitable models. Socially, it involves finding a valuable use case, building technical capabilities with external support, and establishing a central AI team. In the **productionizing** phase, proven concepts are scaled up. Data repositories are organised for multiple use cases, a standardised architecture is established, and experimentation is sped up. The central AI team becomes more formalised, vertical teams support specific business lines, agile methods are implemented, and customer value is created. In the **platformizing** phase, data pipelines are optimised, infrastructure is continually improved, and AI capabilities are reused across different contexts. Horizontal teams are developed across business lines, innovation becomes data-driven, and focus on applied research increases.

Holmström and Magnusson offer a framework for AI transformation through three phases: (1) **path framing** to establish the strategic direction, (2) **path narrating** to determine the sequencing of initiatives, and (3) **path stretching** to focus on practical deployment and scaling. Additionally, three core dimensions are defined to navigate AI implementation: (1) automation of routine tasks, (2) augmentation of human capabilities, and (3) data richness [18].

Weber et al. offer a capabilities perspective on AI implementation. Based on 25 exploratory interviews with AI experts, they identify four organisational capabilities essential for successful AI implementation: (1) **AI Project Planning**: identifying and prioritising suitable use cases, (2) **Co-Development of AI Systems**: integrating stakeholders and translating AI to business, (3) **Data Management**: collecting and curating data, and (4) **AI Model Lifecycle Management**: orchestrating development, deployment, and maintenance [40].

Forrester's AI maturity model identifies six dimensions for organisations scaling AI: **principles and strategy**, **governance and processes**, **operating model**, **talent**, **technology and data**, and **roadmap and activation** [9]. Unlike the other sources, this is a maturity model: it describes how far an organisation has progressed in AI maturity rather than which patterns come with AI scaling. It is included here because its dimensions overlap with the ITOM domains used in this study.

These frameworks and the scaling challenges serve as input for the empirical part of this study. They are mapped to their corresponding ITOM domains in the interview structure and help ensure that relevant themes are covered during the interviews.

2.4 Synthesis and research gap

The literature reviewed in this chapter points in two directions that have so far been developed mostly separately.

On one side, the ITOM (Section 2.2) provides a recognisable structure for how organisations configure IT to deliver value. Across academic and practitioner sources, a consistent set of domains comes back: governance, people and organisation, processes, technology and infrastructure, data, and culture. This literature is mature in describing *what* an operating model consists of, but it does not address *how* these domains should be configured for a specific technology, and it does not look at how a new technology moves from experimentation to broader organisational deployment.

The AI scaling literature (Section 2.3) shows that organisations consistently fail to move AI initiatives beyond pilots, and identifies recurring challenges related to data, talent, governance, and organisational change. Existing AI scaling frameworks, such as those proposed by Haefner et al. [13], Holmström and Magnusson [18], and Weber et al. [40], offer useful views on the scaling challenge. However, they focus on capability building, transformation, or value-creation, but do not translate the scaling challenge into the operating model patterns that IT leaders can recognise in practice, and they do not connect these patterns to the ITOM domains through which IT functions are organised in practice.

This thesis addresses this gap. There is currently no empirically grounded framework that:

1. uses the established ITOM domains as the structure,

2. identifies the recurring operating model patterns that practitioners associate with scaling AI, and
3. is expressed in a format that is recognisable and actionable for IT leaders.

Closing this gap is the objective of the empirical part of this study.

Chapter 3

Method

3.1 Research approach

This thesis follows an exploratory qualitative research design aimed at developing an ITOM framework for scaling AI. The study is based on semi-structured interviews with experienced practitioners and focuses on the recurring ITOM patterns that emerge when organisations move AI initiatives beyond pilots and proofs of concept. During the interviews, participants described operating model decisions they had made or observed, as well as the effects they associated with them. These accounts were then analysed across interviews to identify recurring patterns.

A qualitative approach is appropriate because this research aims to understand how practitioners configure the IT operating model when scaling AI, and to identify the patterns that recur across their accounts. This involves interpretation, experience, and organisational complexity, rather than quantitatively measurable variables. The research is exploratory because only partial information is currently available in the literature, and because this is a novel and rapidly evolving field.

The framework developed in this thesis is a synthesis of the empirical findings rather than the result of a formal design science process. This research is not a case study of individual organisations. Instead, it is based on practitioner perspectives across multiple contexts, which provides breadth but less depth within any single organisation.

3.2 Research context

3.2.1 PwC project context

This research was conducted in the context of an internship project within the PwC Netherlands, specifically in the team focused on CIO Advisory. The internship was carried out under the supervision of a Senior Manager with experience in IT strategy, digital transformation, and cost and value management.

PwC serves clients across a wide range of sectors and organisational sizes. The PwC network was used to help identify relevant participants, including both internal consultants and external professionals involved in AI value realisation initiatives. Any PwC- or client-sensitive information was handled confidentially and reported in anonymised form.

Conducting the study in the context of a PwC internship provided access to relevant practitioners, but it may over-represent participants who are already engaged in transformation-oriented or relatively mature AI discussions. It may also create incentives for participants to describe the organisations in a more favourable way. To reduce this risk, the study also includes participants from outside the PwC network. Furthermore, a common interview guide was used across interviews, and the findings were reported in completely anonymised and aggregated form rather than as judgements about specific organisations.

3.2.2 Participant and organisational diversity

Because the study is based on practitioner perspectives rather than on organisational cases, participants may draw on experience from one organisation or from multiple projects and organisational contexts. The sample aims to include variation in organisational size, industry, and participant role, including both larger established organisations and smaller AI-native firms or startups where relevant.

3.2.3 AI context of participants

Participants were selected based on their exposure to AI initiatives that go beyond basic experimentation or the use of simple copilots. The study focuses on participants with experience related to scaling, organising, governing, or operationalising AI in practice. Because participants' level of AI scaling experience may vary across organisations, this context was recorded during the interview and considered in the analysis.

3.3 Conceptual frame

The empirical analysis is structured around the six ITOM domains from Section 2.2 (Governance, People & Organisation, Processes, Technology & Infrastructure, Data, and Culture). Within these domains, participants described operating model decisions and the perceived effects of those decisions in relation to scaling AI.

Across interviews, these domain-based accounts were analysed to identify three types of observations: recurring patterns, frontier observations, and tensions & disagreements. Recurring patterns form the main basis for the framework. Frontier observations and tensions were retained to support analytical transparency, but they are not treated as the main basis for the framework. The findings therefore reflect recurring practitioner perspectives rather than the full reality of any single organisation.

3.4 Data collection

3.4.1 Semi-structured interviews

Semi-structured interviews were conducted with professionals who had direct experience in designing, advising on, or implementing AI initiatives at organisational scale. Participant selection was based on three criteria: (1) hands-on experience with AI technologies including but not limited to large language models, generative AI, and agentic AI systems, (2) a senior or advisory role within or for IT organisations, and (3) exposure to multiple dimensions of how IT organisations operate, not solely technical implementation.

A semi-structured interview format was chosen because it provided enough structure to cover each of the six domains across interviews, while still allowing follow-up questions on issues that proved particularly relevant in a given interview. The semi-structured nature of the interviews also made it possible to use findings from the literature review and points raised during the interviews to ask relevant context-specific follow-up questions.

Participants included management consultants, AI and data leads, product managers, founders, and senior partners. Not all participants held C-level titles, but each participant had practical experience that provided insight into how IT organisations adapt their operating model to scale AI. In total, fifteen interviews of around 60 minutes were conducted. These were recorded with consent and anonymised.

3.4.2 Interview structure

The interviews focused on the six most frequent domains in the cross-framework analysis (Table 2.2), which also align well with the findings in AI scaling (Section 2.3). The remaining five domains were not excluded completely from the analysis. The interviewer asked about them if the participant brought them up or if they turned out to be important in the participant's context.

The full interview guide is provided in Appendix A.

3.4.3 Sampling strategy and overview of participants

Participants were recruited through two channels: (1) the PwC network, including colleagues from different advisory teams and adjacent practices, and (2) the researcher's personal and professional network, including contacts from previous work experience.

This sampling approach is intended to provide access to professionals with relevant AI scaling experience across different roles, organisational sizes, industries, and implementation contexts. Particular attention was given to selecting participants with enough seniority or oversight to speak meaningfully about operating model decisions across domains.

The target sample was fifteen interviews. This sample size was chosen to ensure a sufficient range of perspectives while still being able to do a detailed cross-interview analysis. Because this study relies on practitioner accounts as its main empirical input, it is important to show that the selected participants meet the three sampling criteria defined above. Therefore, Table 3.1 gives an anonymised overview of the participants, their background, and the perspective from which they contribute to the study.

Table 3.1: Anonymised overview of interview participants

ID	Role & context	Background & relevance
I1	Senior manager in AI consulting; enterprise transformation context; 11+ years	• Background in econometrics and business analytics • Experience in IT-oriented advisory and transformation work • Works across data, software engineering, and AI • Relevant for AI strategy, governance, and operating model design
I2	AI product manager and innovation centre lead; 7+ years	• Background in business analytics and product-oriented work • Leads an innovation and experience centre team • Works across design, product, research, and engineering • Relevant for AI adoption, experimentation, and cross-functional ways of working
I3	Co-founder and CEO of an AI-native scaleup; 15+ years	• Former MBB strategy consultant with IT-strategy exposure • COO and CTO experience in multiple technology-driven organisations • Leads an enterprise AI platform company serving large organisations • Relevant for platformisation, governance, and scaling AI in practice
I4	Senior director in AI, data, analytics, and transformation; commercial & industrial sector; 13+ years	• Background in econometrics, operations research, and data analytics • Extensive experience in data, analytics, and AI-related transformation • Works on AI maturity assessments, AI strategy, and use-case prioritisation • Relevant for AI maturity, prioritisation, and the organisational changes needed to move beyond pilots
I5	Former regional head of AI and data in insurance and telecom; 20+ years	• Former regional lead for data and AI strategy execution across a multi-country insurer • Executive-level experience in data, AI, and enterprise IT in telecom • Worked across federated, multi-market environments with differing regulations and maturity levels • Relevant for scaling AI across countries, business units, and regulatory territories
I6	Consulting partner in cloud and digital transformation, with deep financial-services experience; 25+ years	• Long-standing experience in technology and transformation advisory, especially in financial services • Background in operating model change, platform transformation, and end-to-end process improvement • Executive experience in domain operating model design and platform product ownership • Relevant for how cloud, process redesign, and delivery models shape the conditions for AI scaling
I7	Senior AI and cloud consulting manager with AI platform and governance focus; 7+ years	• Background in cloud, data engineering, and AI-related delivery work • Involved in AI platform setup and governance-related initiatives in large organisations • Experience with reusable components, access control, and platform onboarding • Relevant for platform enablement, controls, and AI literacy
I8	Senior consultant in AI and data strategy, adoption, and implementation; 7+ years	• Background in computational science, data analytics, and AI implementation • Works on AI strategy, maturity assessments, implementation planning, and client-facing demos • Experience across consumer and public-sector contexts, including larger organisations and smaller implementation settings • Relevant for early-stage scaling challenges such as value measurement, literacy, and capability gaps

ID	Role & context	Background & relevance
I9	Senior consultant in responsible AI and AI governance; 7+ years	• Background in econometrics, business analytics, and responsible AI • Focuses on AI policy, governance, risk management, and responsible-AI controls in practice • Works on lifecycle design, integration of AI into work processes, and implementation of controls such as security, robustness, and bias mitigation • Relevant for balancing experimentation, centralisation, and responsible use
I10	Consulting director and former Big Tech engineering lead; 20+ years	• Engineering management experience in Big Tech, combined with leadership in data and AI transformation • Background in software engineering, data science, and enterprise AI architecture • Works on scalable, secure, and compliant AI-enabled platforms, with product ownership and integration exposure • Relevant for how governance, data, platforms, and business ownership come together beyond proof-of-concept stage
I11	Chief of Staff in a tech-focused company, with tech strategy consulting experience; tech product / startup context; 8+ years	• Chief of Staff in a tech-focused company with an AI-related platform product • Experience in tech strategy consulting, including work with large organisations • Focus on embedding AI into broader ways of working rather than treating it as a side initiative • Relevant for tensions around data, cost control, organisational structure, and pilot-to-scale thinking
I12	Senior advisor in AI transformation and operating model redesign; 25+ years	• Nearly three decades of experience in IT transformation, digital transformation, big data, and AI • Worked on business-process improvement, IT organisation design, and AI-related transformation across large organisations and regulated sectors • Experience with AI governance, prioritisation, prototyping, roles, MLOps-style testing, data governance, and operating-model evolution • Relevant for disciplined experimentation, cross-functional coordination, and staged operating-model choices
I13	Director in AI transformation and cloud advisory; 15+ years	• Strong technical background in infrastructure, programming, data, analytics, and enterprise transformation • Both early engineering experience and later business and leadership experience across large public- and private-sector organisations • Worked on AI operating model decisions, platforms, governance, benefits management, business-IT, and agentic development • Relevant for how large organisations can operationalise AI at scale
I14	Independent AI advisor and former consulting partner focused on large-scale transformation in energy and industrial sectors; 25+ years	• Extensive experience across energy, industrial operations, manufacturing, and digital transformation • Background includes long-running transformation work, M&A-related experience, and senior roles across consulting, technology, and industry • Explicitly value-driven view of AI, emphasizing process friction, business priorities, and operating-model discipline over technology hype • Relevant for AI scaling in complex industrial and energy environments
I15	Director in AI, innovation, and workforce transformation; 15+ years	• Senior background in workforce transformation, employee experience, and future-of-work topics across leading global technology firms and enterprise environments • Combines experience with predictive models and agent development, data pipelining, and process analysis with organisational psychology and change management • Contributed to work with the World Economic Forum on GenAI and workforce change • Relevant for behavioural adoption, role redesign, and middle-management dynamics

3.4.4 Ethical considerations

All participants were informed that participation is voluntary, the purpose of the study, the use of anonymised quotes, and their right to withdraw at any time without consequence.

To reduce identification risk, quotations are anonymised and potentially identifying organisational details are generalised where necessary. The study does not report commercially sensitive information. Interview recordings are stored securely and are accessible only to the researcher and supervisors. Organisational and participant names are anonymised in all reporting. The data were used only for this research and were not shared with other organisations or participants.

3.5 Data analysis

A structured tabular approach was chosen over formal coding because the interviews followed a consistent domain-based structure. This made a template-based summary more suitable for systematic cross-interview comparison.

The analysis was carried out in three steps.

Step 1: Per-interview structured summary. Each interview was summarised from recordings, using a fixed template that captured, per ITOM domain, the main operating model decisions described by the participant and the perceived impacts the participant associated with them. The purpose of this step was to reduce each interview into a comparable domain-based summary. The full per-interview summaries are included in Appendix B. The full interview summaries are included for transparency and as a readable substitute for raw transcripts, while the compact summaries (Appendix C) contain the more condensed format used in the cross-interview analysis.

Step 2: Cross-interview analysis matrix. The per-interview summaries were consolidated into an analysis matrix in which interview findings were organised by domain. To support systematic comparison, the summaries were condensed into shorter matrix entries while preserving the core findings. This provided the basis for identifying patterns across the interviews.

Step 3: Pattern identification. The matrix was then reviewed across interviews to identify three types of observations:

- **Recurring patterns**: recurring configurations, choices, or mechanisms described by three or more participants, including the effects they were associated with. These form the primary basis for the framework.
- **Frontier observations**: relevant observations that did not recur across enough interviews to qualify as patterns (fewer than three participants). These include findings from single, more advanced contexts and other worthwhile outliers. They are recorded so that relevant insights are not lost, without being treated as recurring patterns.
- **Tensions and disagreements**: cases where participants described the same domain in contradictory ways. These are noted to support analytical transparency and are not carried into the framework.

3.5.1 Trustworthiness

To improve the trustworthiness of the study, the following measures were taken:

- A consistent interview guide was used across interviews.
- Participants were selected for their direct practical experience and according to the criteria defined in Section 3.4.3.
- Patterns were compared across multiple interviews rather than relying on isolated statements.
- The interview guide, summaries, and cross-interview analysis were documented to improve traceability.
- Where possible, participants were asked to verify the summary of their interview to reduce the risk of misinterpretation.

3.5.2 Framework development procedure

The framework was developed in three steps. First, the recurring patterns per ITOM domain were consolidated into the framework. Second, frontier observations were included as examples where relevant; the framework focuses on the recurring patterns rather than on single-context observations. Third, the resulting v1 framework was evaluated with practitioners and revised into v2 based on their feedback (Chapter 6).

3.6 Evaluation approach

The framework was evaluated through structured feedback requests. Evaluation participants included two of the original participants and three additional practitioners not involved in the interviews, reducing the risk of confirmation bias. The framework was evaluated on three criteria: completeness, clarity, and practical usefulness. Feedback was used to revise the framework into v2. The detailed evaluation procedure and results are reported in Chapter 6.

Chapter 4

Findings

4.1 Introduction

This chapter presents the findings from the cross-interview analysis. The findings are organised along the six ITOM domains introduced in Section 2.2. For each domain, recurring patterns are presented first, followed by frontier observations and tensions & disagreements.

Frontier observations and tensions are included for analytical transparency, but they do not form the main basis for the framework.

Following the six domains, Section 4.8 presents the cross-domain dependencies that emerged across the interviews, which serve as input for the layered framework in Chapter 5.

4.2 Governance

The governance findings show that AI scaling is mainly treated as a question of ownership, prioritisation, and organisational control rather than as a purely technical rollout. Across the interviews, the strongest recurring patterns point to top-down mandate, structured use-case selection, and governance arrangements in which business ownership is combined with central coordination and oversight. The findings also show that organisations increasingly need explicit governance around value tracking, cost control, guardrails, and responsible AI, while regulation and the role of compliance and audit make governance more complex.

Recurring patterns

Table 4.1: Cross-interview patterns in Governance

Cross-interview pattern	Supporting interviews
Use-case prioritisation runs through multi-axis funnels (value, feasibility, strategic fit, organisational fit) with standardised templates and staged gates, typically starting with lower-risk cases where data is usable and escalating higher-risk cases where more controls are needed	I1, I4, I5, I6, I9, I10, I12, I14, I15
AI governance is a distinct, multi-disciplinary practice (CIO, CDO, CISO, business, etc.) with joint oversight; traditional IT governance is no longer sufficient	I1, I7, I8, I9, I10, I12, I13
A CoE is the standard early choice to bundle scarce expertise and prioritisation; decentralisation only comes with higher maturity to prevent fragmentation, weak reuse, and governance complexity	I3, I6, I8, I9, I11, I12, I13
Business owns AI demand, value realisation, and risk; IT/data/AI act as enablers to prevent "technology looking for a problem"	I3, I4, I6, I9, I10, I14, I15
AI scaling depends on top-down ownership at board level with mandate and budget, with day-to-day governance delegated, otherwise initiatives stall at middle management and stay stuck in experimentation	I2, I3, I8, I11, I12, I13
Formal, leadership-backed data governance is treated as a prerequisite for AI scaling, requiring explicit ownership rather than informal or purely technical arrangements (the detailed data-ownership mechanisms are covered under Data, Section 4.6)	I2, I3, I6, I8, I11, I12
Governance combines business-value KPIs (ROI, operational figures, adoption, efficiency, NPS, STP) with control/cost metrics against a pre-AI baseline, but value tracking is often still lacking as solutions get built without it	I3, I7, I8, I9, I10, I13
Cost is becoming an explicit governance discipline: FinOps-style monitoring, token tracking, spend visibility to business users, per-user caps, and context-based model routing rather than one model for all tasks	I1, I3, I7, I9, I11, I13
AI must be framed as cross-functional transformation, not isolated pilots; a project/PoC-only mindset leads to failure	I1, I2, I11, I12, I14
AI investment is governed through explicit business-case logic, funding choices, and finance involvement rather than ad hoc local spending; this includes both organisation-wide and divisional funding with stronger focus on expected value	I3, I4, I11, I13
The EU AI Act and other regulation require a dedicated AI governance place and roles, and the pace of tech/AI evolution makes this a particularly big challenge; compliance/audit act as gatekeepers in regulated settings but should not define strategic direction, and involving them too early stalls AI	I7, I12, I13, I14
Prevent stalling at pilot/PoC stage ('PoC hell' / pilot fatigue) from innovations disconnected from operations; requires portfolio ownership with mandate to scale, redirect, or end initiatives	I4, I10, I12, I14
Guardrails and responsible-AI policy are necessary before broad rollout	I1, I4, I9, I12
AI surfaces pre-existing security weaknesses; CISO sometimes acts as initiator rather than a control function; platform/vendor choices (incl. sovereign options) are used to work around security and regulatory constraints	I3, I7, I14

Frontier observations

Table 4.2: Frontier observations in Governance

Frontier observation	**Supporting interviews**
Governance discipline is needed to prevent AI from being used as a superficial patch on deeper organisational, process, or data problems	I14, I15
Treating each AI initiative as a product with joint business/IT ownership and a catalog (ownership, boundaries, intended use)	I10
Three-lines-of-defence adapted to AI, with explicit lifecycle governance, decision moments, and risk-based escalation of high-risk cases	I9
Federated multi-country governance via portfolio funnels prioritising local portfolios to select regional priorities by commonality and critical mass; sometimes non-local priorities are needed for future commonality	I5
Rapid prototyping is used as a governance tool to create evidence, prioritise, and gather executive support; executive attention is the scarce resource, not budget	I12
Employees increasingly act as managers of AI agents, and accountability for AI errors in deliverables, decisions, or client-facing outputs must be assigned, with the biggest issues arising where human sign-off is mature but AI-supported sign-off is not	I15

Tensions & disagreements

Table 4.3: Cross-interview tensions & disagreements in Governance

Tension / disagreement	**Supporting interviews**
Build a new, dedicated AI governance practice vs. extend existing data/product governance to cover AI vs. redesign governance around where responsibility actually sits, since no single model fits all organisations	I1, I7, I8, I12 / I10 / I15
Involve compliance/audit early to set guardrails and act as gatekeepers vs. involve them late, because doing so too soon or letting them steer strategy stalls AI	I2, I3, I9 / I13, I14

4.3 People & Organisation

The findings in People & Organisation show that AI scaling is not just a specialist talent issue, but a broader redesign of capabilities, roles, and responsibility across the organisation. The recurring patterns point to a centralise-first approach to capability building, the emergence of new roles such as AI engineers, product-oriented bridge roles, and platform-related roles, and a combination of broad AI literacy with more focused specialist development. Across the interviews, AI is also more often framed as a way to increase value per employee and reshape work than as an immediate headcount-reduction programme, even though role change, reskilling pressure, and workforce differentiation are clearly part of the picture.

Recurring patterns

Table 4.4: Cross-interview patterns in People & Organisation

Cross-interview pattern	**Supporting interviews**
Scaling AI starts centralised and decentralises once business units reach enough maturity (e.g., CoE → Federated), but there is no universal standard	I3, I5, I6, I7, I9, I10, I11, I12, I13
Common new roles are emerging: AI/agentic engineers, AI architects, AI product owners, and platform engineers. Data steward/data owner roles are being strengthened	I1, I3, I4, I5, I9, I10, I11, I12
Talent strategy combines internal reskilling with external hiring, because neither closes the gap alone; data scientists/ML engineers are the easiest to pivot into agentic/AI engineering roles	I1, I3, I4, I6, I7, I8, I13
AI raises value per FTE faster than headcount grows, creating more value with the same people. Efficiency gains are often redirected to other work instead of cuts, while AI also creates new work in checking, managing, and improving outputs	I1, I6, I10, I11, I13
A recurring role is the AI translator or product-oriented bridge role that connects business needs, process knowledge, and AI delivery	I4, I5, I8, I10, I12, I15
Capability building happens on two levels: broad AI literacy across the organisation and focused hiring or development of scarce specialist roles	I5, I7, I8, I9
Workforce reduction is not an upfront goal, but real "losers" exist who cannot or will not adapt, requiring reskilling or an exit	I1, I4, I5, I12
Build capabilities around management training, business translation, and opportunity-spotting to find AI value	I5, I8, I10
Value realisation still depends on human capability even when AI is technically available; buying in technology does not remove the people challenge, because employees still have to learn to work with it	I13, I14, I15
Non-specialist/citizen builders (no-/low-code personas, "tinkerers") are treated as part of the AI workforce and enabled within guardrails to spread capability while containing operational/financial risk	I1, I3, I14

Frontier observations

Table 4.5: Frontier observations in People & Organisation

Frontier observation	Supporting interviews
Required emerging skills include prompt, context, and harness engineering as a baseline hiring competency; demand is staged per maturity level ranging from prompting to specialised AI-engineering	I1, I7
AI-native hiring filters are being used and self-directed AI experimentation is treated as a core role expectation	I2, I3
Workforce segmented into no-code, low-code, and pro-code personas for strategic workforce planning with a target persona mix	I1
Broad AI scaling is hard to justify on short-term ROI and reskilling engineers is constrained by existing workload, making (expensive) hiring easier	I8
Smaller, flatter, hybrid-role teams with cross-disciplinary engineers as a high-performance organisational model	I2
Expert work is being recorded to capture and transfer expertise into AI models	I3
Heavy AI use by juniors risks losing domain knowledge and critical thinking over time so upskilling must be managed carefully	I9
Cybersecurity roles are becoming increasingly important in AI scaling	I7
Process-knowledgeable employees are positioned as the highest-value spotters of AI opportunities	I11
Pair domain specialists with AI specialists rather than converting domain experts into full AI specialists	I15
Use concrete 18–24 month reskilling pathways (e.g., from a 50% role match) as a structured alternative to talent loss	I15
Redesign reward systems when AI shifts responsibility or task mix; adding accountability without reward fails	I15
Adaptability, learning willingness, and future-orientation are the true scarce resource; unlike technical skill, it cannot be bought in	I15
Middle managers taking on AI responsibility need active support, or middle management becomes the bottleneck rather than the accelerator	I15
Executive-sold demos often don't succeed on the work floor and create resistance in the layers below	I15

Tensions & disagreements

Table 4.6: Cross-interview tensions & disagreements in People & Organisation

Tension / disagreement	Supporting interviews
Inclusive workforce transition (reskilling, ambassadors, managed exit) vs. selective AI-native filtering where non-AI-engaged employees are not hireable	I1, I12, I15 / I2, I3
AI as additive productivity ("more output, same FTE", redirect gains to growth, new oversight work created) vs. AI as a driver of measurable workforce/process reduction (e.g., 100→80 people, 5→4 steps)	I1, I6, I10, I11, I13 / I12
Reskilling existing technical staff as the easier path vs. staff being occupied, making hiring easier but costlier	I6 / I8

4.4 Processes

The process findings show that the strongest value from AI is expected when organisations move beyond isolated task automation and start redesigning work more fundamentally around AI capabilities. Recurring patterns point to AI-first process redesign, an early focus on repetitive and well-documented work,

and a shift in human work toward review, exception handling, and more complex cases. At the same time, the findings show that broader process impact depends on stronger foundations in data, platforms, documentation, and change management, because many organisations otherwise remain stuck at local use cases or pilots rather than achieving end-to-end redesign.

Recurring patterns

Table 4.7: Cross-interview patterns in Processes

Cross-interview pattern	**Supporting interviews**
AI delivers the highest value through AI-first process redesign, decomposing workflows around agentic thinking rather than inserting AI into unchanged steps	I1, I2, I4, I5, I6, I8, I9, I13, I15
Missing data, AI-platform, or process documentation foundations cause pilots to stay too narrow, local, or disconnected, stalling broader organisational scaling	I3, I4, I6, I8, I9, I11, I14
Early priority for text-heavy, repetitive, rule-based, high-volume administrative and well-documented work (reporting, drafting, communication, document extraction, follow-up, finance/procurement)	I4, I6, I7, I9, I11, I12
Claims handling, customer service, and knowledge/legal work are priority high-impact domains	I1, I3, I5, I9, I11, I12
Human work shifts toward review, validation, exception handling, and complex cases, raising the skill profile of remaining staff and facilitating growth without proportional FTE increase	I1, I4, I6, I11, I12, I13
Scaling beyond local experiments requires structured process design with staged delivery gates, lifecycle management, onboarding, change management, value tracking (classical change management and not just tooling)	I7, I8, I9, I10, I12
A/B comparison (work with vs. without AI) is the main evaluation method, often paired with hard process KPIs such as cycle-time, throughput, % volume handled by AI, and reduction in human handovers	I4, I8, I9, I10
Strong efficiency gains in processes, e.g., 60–90% and up to ˜10x time reduction (e.g., claims processing, multi-reviewer text handling collapsed from 5 reviewers to 1, legal intake from ˜60 to ˜6 min)	I1, I2, I6, I12
Software engineering recurs as the most disrupted end-to-end process, with AI-assisted development reshaping technical work	I1, I3, I6, I7
AI value increases with the scope of redesign: from individual productivity tooling, to team and process automation, to cross-silo end-to-end agentic workflows. Real transformation value comes mainly from the top tier, where processes are redesigned from first principles around the final goal and humans are only reintroduced where judgment is needed	I13, I14, I15
Top-down "big bets" combined with bottom-up use cases that emerge from trained business AI-translators or process-knowledge employees	I5, I8, I11

Frontier observations

Table 4.8: Frontier observations in Processes

Frontier observation	Supporting interviews
Operational fit must be validated with actual frontline users, not just technical or design-based preferences	I8, I12
AI shifts where a process starts (e.g., proactive issue detection) rather than only accelerating existing steps	I10
Fully autonomous execution with no human-in-the-loop is possible where guardrails suffice (e.g., autonomous payment arrangements)	I1
An automated test portfolio (unit, integration, system, UAT) as a trust mechanism for AI-generated code without manual review	I1
Reuse mechanisms (templates, shared components, marketplace) and lifecycle management treat AI use cases as living applications requiring monitoring, maintenance, and adaptation rather than one-off deliveries	I7
External pressure forces you to work with AI, e.g., citizens using AI for objections and subsidy applications which increases workload	I2
In-house builds risk falling behind because the market moves faster than ~18-month build cycles; mature use cases often use existing solutions	I3
Emerging advanced frontier use cases include R&D, tailor-made services, and education, alongside legal	I12
Large enhancement-only use cases (e.g., improved asset monitoring) can deliver major gains without redesigning the process	I14
Over-automating customer-facing interactions is a risk where fieldwork or human judgment still matters; excessive automation can create ineffective "chat-on-chat" loops and degrade communication	I14
Process selection by locating friction (waiting time, handoffs, rework, inconsistent data, cross-silo gaps) rather than by task type	I14
Prioritise by economic weight in the cost base / value chain (e.g., supply chain, core operations), not only the highest %-ROI use case, and look beyond familiar starting points like customer contact	I14
Business/market-position dictates how a process should change, not tech: rethink the process when competitors are overtaking you, but optimising remains equally valid where margins or cost still improve	I15

Tensions & disagreements

Table 4.9: Cross-interview tensions & disagreements in Processes

Tension / disagreement	Supporting interviews
Augmentation-first and redesign as AI matures vs. AI-first end-to-end redesign as the primary value source	I10, I14 / I2, I4, I5, I6, I9, I13, I15
Simple, documented, rule-based, repetitive work vs. complex, cognitively demanding work as the higher-value target	I4, I9, I11 / I8, I13
Custom in-house builds vs. existing point solutions given market speed	I5, I6 / I3

4.5 Technology & Infrastructure

The findings in Technology & Infrastructure show that scaling AI depends less on isolated tooling decisions and more on building a shared technical foundation that supports reuse, integration, and control across use cases. Recurring patterns point to shared AI platforms, layered architectures, stronger integration mechanisms, pragmatic build-versus-buy choices, and growing attention to sovereignty, vendor

dependence, and agent identity management. The interviews also show that legacy matters less as a question of system age alone and more as a question of access, integration, and whether existing systems can expose the data and functionality AI needs.

Recurring patterns

Table 4.10: Cross-interview patterns in Technology & Infrastructure

Cross-interview pattern	**Supporting interviews**
Tech and infrastructure are treated as an enabling layer derived from business case, AI ambition, and strategic differentiation rather than technology for its own sake. IT, data, and business decide together at enterprise level, and the hardest challenge is embedding AI into processes and platforms	I2, I4, I5, I6, I8, I10, I11, I14
MCP-style interfaces and broader integration layers are becoming necessary to connect agents safely to enterprise systems and data, including agent-to-agent communication. Strong integration is a prerequisite for higher-level AI value	I2, I3, I7, I10, I11, I12, I13, I14
Sovereignty is increasingly becoming a priority beyond GDPR, treated as strategic and partly geopolitical concern; this constrains deployment in regulated contexts	I2, I3, I7, I11, I12, I13, I15
Organisations generally extend their current strategic vendor stacks rather than replacing them, choosing larger and more stable ecosystems to reduce fragmentation and continuity risk; niche tools are not essential for scaling, and premature re-platforming is avoided	I1, I2, I3, I7, I8, I9, I14
Legacy is no longer an age-based blocker, but an access-and-data problem: whether the system can expose the data and functionality AI needs is the deeper issue, addressed through connectors/APIs and selective modernisation rather than full rebuild; decisions are made case-by-case per system (modernise, keep with agents on top, or use agents to accelerate migration)	I3, I4, I5, I7, I9, I10, I15
Build-vs-buy is case-specific and based on risk-management, strategic differentiation, and market availability; buying certified/supported vendor solutions is preferred for regulated contexts (support, certification, accountability, liability); although AI improves in-house build capabilities, overbuilding is a critical risk	I4, I5, I6, I10, I11, I14, I15
Cloud/hyperscaler is the default deployment model (scalability, cost, convenience, maintainability); on-premise and containerised/local setups are mainly chosen in regulated, public-sector, or sovereignty-sensitive contexts and come with higher maintenance cost, complexity, and skill requirements, often as a hybrid setup	I1, I5, I7, I9, I10, I13
A shared AI platform is the recurring scaling foundation, arranging model/data access, integrations, identity, security, governance, and agent lifecycle management once rather than per use case; without it, organisations stay limited to copilots and support use cases	I2, I6, I7, I9, I10, I11
Scaling depends on a multi-layer architecture (data platform + AI platform + integrations/APIs)	I5, I6, I10, I11, I12
Agent identity and access management is a critical new capability and hard scaling barrier: agents must not inherit human permissions 1:1, requiring separate authentication/authorisation, roles, and boundaries that existing enterprise systems are not designed for	I1, I3, I7, I9, I13
Earlier technology and integration choices constrain AI scaling: fragmented systems, point-to-point integrations, and weak data foundations slow cross-system/end-to-end AI workflows	I4, I5, I8, I11, I15
A layered model strategy is key: match the model to the problem, since frontier or token-heavy models are often suboptimal for simple tasks, while sovereignty, cost, copyright, and security also impact the choice.	I12, I13, I15

Frontier observations

Table 4.11: Frontier observations in Technology & Infrastructure

Frontier observation	**Supporting interviews**
MLOps is extended with semantic testing because static technical tests cannot catch model, data, and context drift; pre- and post-deployment automated testing becomes necessary	I2, I12
Technical scaling depends more on effective use of existing foundation models than on custom training; few use cases justify custom training	I7, I13
Agentic coding tools (e.g., Cursor and Claude Code) reshape IT making pro-code accessible, phasing out low-code, running background/reporting agents, and improving developer output	I1, I2
Agentic systems add new complexity in multi-agent coordination and monitoring of non-deterministic behaviour, not just identity management	I1
Automating compliance checks as a competitiveness lever, since manual governance is costly and slow in the regulated EU context	I12
A reusable "harness" (shared platform/tools/settings) for personalisation on a common foundation	I2
Marketplace/shared-component layer enables cross-organisation reuse of components originally built for individual use cases	I7
Agent memory/context management is ad-hoc and fragmented, but critical	I2
Each application is given its own controlled sub-environment within the shared platform to avoid an uncontrolled shared space	I9
A shift from classic software building toward composing applications from agentic building blocks, expecting an app-less 'conversational' future that deconstructs traditional enterprise applications	I13
AI itself can act as a flexible integration layer, reducing customisation load on heavily customised enterprise platforms	I14
AI oversight tooling for visibility into performance, cost, and usage footprint is introduced as a distinct capability	I14

Tensions & disagreements

Table 4.12: Cross-interview tensions & disagreements in Technology & Infrastructure

Tension / disagreement	**Supporting interviews**
Legacy as a fundamental blocker requiring foundations to be fixed first (e.g. single source of truth, fragmentation) vs. legacy as an access/data issue that can often be worked around or selectively modernised	I4, I11 / I3, I9, I10

4.6 Data

The data findings show that AI scaling still depends heavily on the quality, structure, accessibility, and governance of organisational data. The recurring patterns point to data readiness as a hard condition for scaling, the continued importance of shared definitions, a single source of truth, and a shift toward stronger formal ownership with more responsibility placed closer to business domains. The interviews also suggest that AI changes what counts as usable data, as organisations increasingly need data that works not only for people, but also for models and agents, including unstructured data that was previously governed less strictly.

Recurring patterns

Table 4.13: Cross-interview patterns in Data

Cross-interview pattern	Supporting interviews
Data readiness is a hard prerequisite for AI scaling, not a parallel workstream; weak or fragmented foundations cap AI at limited copilot/support use cases rather than core-process redesign and automation	I1, I4, I5, I6, I7, I8, I11, I15
AI is bounded by reachable data quality and amplifies existing weaknesses in quality, definitions, and fragmentation ("Garbage in, garbage out")	I2, I4, I5, I8, I9, I11, I14
Even where data is broadly available, process- or use-case-specific data selection and access restriction remain essential; agents/use cases should access only the data they need	I1, I3, I7, I9, I11, I13, I14
Single source of truth, shared definitions, and master data management (golden records and no duplications) are treated as critical but difficult, often blocked by fragmented and legacy landscapes and contradictory data which AI then amplifies	I1, I4, I7, I9, I10, I11, I14
Data governance is no longer a background topic, but a strategic discipline with formal roles (CDO and data owners/stewards); it is increasingly a behavioural and accountability change	I1, I4, I6, I8, I9, I10
Data governance increasingly combines central standards and oversight with more local or domain-based ownership and execution, supported by roles such as CDOs, data owners, and data stewards	I5, I6, I8, I11, I12
Ownership is shifting towards domain-based or business-anchored ownership (process/system owners) as organisations mature; purely central ownership is described as ineffective or non-scaling	I1, I5, I6, I9
Layered/medallion architectures are used to map data tiers to use-case criticality (e.g., executive reporting on gold) and to separate shared organisational data from use-case-specific operational data	I1, I7, I13
AI is simultaneously a consumer and a generator of data, transforming speech, email, and field inputs into (structured) data, and using agents to improve quality (e.g., CRM hygiene)	I1, I3, I4
Data quality is defined as fitness for AI or agent use rather than human readability or legacy standards, increasing the importance of agent/LLM-ready and stronger governance of unstructured data	I3, I10, I12

Frontier observations

Table 4.14: Frontier observations in Data

Frontier observation	Supporting interviews
The biggest AI-data challenge is aligning, accessing, and integrating existing data foundations, not creating new data or inventing AI techniques	I13, I14
Counterintuitively, agents tolerate poor data better than humans (patient large-volume reading), reducing but not removing the quality bottleneck, making "fix the data first" sometimes legitimate, but sometimes a delay tactic	I3
"Ambient data" captured passively (meeting transcripts, tool exhaust) replaces manual entry and improves agent context	I2
Knowledge-base use cases (e.g., SharePoint and Docs) scale faster and are less data-constrained than structured-data-heavy ERP/predictive use cases, implying divergent readiness paths by use-case type	I4
Knowledge-graph-like structures explored to improve LLM interpretation of organisational knowledge	I12
Start "where data is clean enough" rather than awaiting enterprise-wide perfection	I12
Weak or incomplete data leads to directionally correct output being perceived as true; error rates should be estimated and output qualities should be distinguished for high-stakes decisions	I15
Generic pretrained models must be grounded in organisational context (internal standards, local language, ethics, communication style) before use in judgment-heavy processes, via prompt engineering, RAG, or heavier methods	I15

Tensions & disagreements

Table 4.15: Cross-interview tensions & disagreements in Data

Tension / disagreement	Supporting interviews
Data is hard precondition that stalls ambitions vs. less of an issue because of agents' tolerance for poor data	I1, I4, I5, I6, I11 / I3
Data readiness as the hard gating constraint vs. the challenge being mainly access of existing data rather than a readiness problem	I1, I4, I5, I6, I11, I15 / I13

4.7 Culture

The culture findings show that AI scaling is widely experienced as a change-management challenge. Recurring patterns point to the importance of leadership behaviour, broad AI literacy, bottom-up experimentation, and narratives that frame AI as support for better and higher-value work rather than only as a cost-reduction tool. The interviews also show that resistance is common and is driven not only by fear of job loss, but also by uncertainty, loss of expertise status, and limited understanding of how to use AI well, which makes psychological safety, learning time, and visible role models important for adoption.

Recurring patterns

Table 4.16: Cross-interview patterns in Culture

Cross-interview pattern	**Supporting interviews**
AI adoption is largely a change-management challenge rather than a pure technology implementation. Without cultural development, AI strategy often fails in execution	I1, I3, I4, I8, I9, I10, I12, I14
Employee resistance comes from job-loss fears, expertise/status loss, and lack of knowledge, which requires direct and open discussions on these concerns	I1, I2, I4, I5, I8, I11, I12, I13
Early adopters, champions, and ambassadors are used to translate adoption into practice, build momentum beyond the central team, and pull others along in their local environment	I3, I4, I6, I7, I8, I11, I13
Effective adoption combines top-down direction with bottom-up experimentation and learning; employees should engage with AI themselves rather than leaving it to a specialist group	I1, I2, I5, I7, I10, I12, I13
AI is an enabler of higher-value, more creative work rather than purely efficiency or FTE reduction, as a narrative choice to lower resistance	I1, I5, I6, I11, I15
Broad AI literacy is a prerequisite for cultural readiness, partly because employees will otherwise use AI without understanding its limits, which leads to organisation-wide risks	I4, I5, I7, I8, I9
Visible and active leadership use of AI is a stronger enabler for adoption than formal top-down programmes or messaging alone; weak and indifferent leadership is a barrier to adoption and scaling. Visible AI use through the hierarchy normalises new ways of working	I2, I6, I8, I13
Time and support must be allocated for learning and experimentation, or adoption slows even when the technology is fully available	I5, I9, I13, I15
Set realistic expectations to prevent early disappointment after a failed first AI attempt	I9, I10, I12
Psychological safety and tolerance of failure are preconditions for AI learning; experimentation enables learning fast enough to scale (while clarifying accountability and acceptable AI use)	I12, I14, I15
There is a growing divide between employees who actively engage with AI and those who avoid it; early adopters gain an advantage, while others risk falling behind if they do not adapt	I1, I12, I13

Frontier observations

Table 4.17: Frontier observations in Culture

Frontier observation	Supporting interviews
"AI-native by default" is used as a hiring filter and individual AI (time) investments are expected	I2, I3
AI treated as an amplifier of existing strengths and weaknesses; the same tool produces different value across users as long as humans stay in the loop	I14, I15
Weak innovation culture and sense of urgency slow the step from interest in AI to actual adoption	I9, I13
A large share of AI success is attributed to individual employees and teams (80%) rather than formal innovation programmes (20%)	I2
A dedicated workstream including reskilling and outplacement for disappearing roles	I1
Middle management whose value rests on information asymmetry or gatekeeping is at risk in AI-native operating models, where individuals are supported by personal agent "swarms" instead of added management layers	I3
AI is used as a sparring partner for senior staff and leaders for cognitive work instead of an output-generation tool	I3
Concrete AI usage signals are being brought into performance evaluations	I2
Adoption depth varies widely even where surface-level training exists; training does not directly equal embedded daily use	I9
Using AI without checking outputs can create more work than it saves, requiring extra guidance during experimentation	I9
"Cassandra" dynamic: parts of the organisation act as if nothing is changing, a denial/inertia barrier distinct from active resistance	I3
Adaptability, learning, and orientation toward the future organisation should be rewarded rather than continuing with today's roles	I15
AI is a response to structural workforce pressure (ageing workforce, retirement-driven knowledge loss), not only cost reduction	I15
Only a small minority of organisations are positioned for true business-model reinvention, for most, the realistic goal is becoming adaptive and value-focused, not revolutionary	I14

Tensions & disagreements

Table 4.18: Cross-interview tensions & disagreements in Culture

Tension / disagreement	Supporting interviews
Actively including less AI-literate employees to avoid alienation vs. building around advocates/change-makers and frame others as individual career risk	I7 / I3, I13

4.8 Cross-domain dependencies

Across the interviews, AI scaling was described not as the expansion of a single technology, but as a broader organisational transformation with dependencies across multiple domains. Successful scaling depended on coordinated change across all of these domains. The findings below summarise the main cross-domain dependencies.

Table 4.19: Cross-domain findings

Cross-domain finding	**Dependent domains**	**Source interviews**
AI scaling is an operating model transformation and not a standalone technology rollout, because progress in any domain depends on coordinated change across governance, people, processes, technology, data, and culture.	Governance; People & Organisation; Processes; Technology & Infrastructure; Data; Culture	I1, I2, I4, I5, I6, I9, I14
Early centralisation helps establish the capabilities, standards, and control mechanisms needed for later decentralisation or federation.	Governance; People & Organisation; Technology & Infrastructure; Data	I3, I6, I9, I12, I13
AI scaling is driven more by business ownership, strategic prioritisation, and expected organisational value than by isolated technical opportunities or ad hoc experimentation.	Governance; Processes; People & Organisation	I1, I3, I4, I5, I14
Process redesign depends on data quality, workflow documentation, integration maturity, and system access, rather than on AI tooling alone.	Processes; Data; Technology & Infrastructure	I2, I4, I5, I6, I8, I9, I10, I11, I15
Governance choices shape the conditions for adoption, because clarity of accountability, leadership commitment, and the way AI is positioned within the organisation determines whether AI scales or stalls.	Governance; People & Organisation; Culture	I3, I11, I12, I15
Process transformation and workforce change are closely interdependent, because process redesign reshapes roles, skill requirements, and the practical conditions for adoption.	Processes; People & Organisation; Culture	I4, I10, I12, I13, I15
Shared AI platforms enable scale by combining reusable infrastructure, controlled access, monitoring, and governance in a common, safe environment.	Technology & Infrastructure; Governance; Data; People & Organisation	I2, I7, I9, I10, I13

Chapter 5

Framework development

5.1 Introduction

This chapter presents the framework developed from the recurring patterns identified in the cross-interview analysis (Chapter 4). The framework synthesises these patterns into a representation of organisational AI scaling across the ITOM domains.

Following the procedure described in Section 3.5.2, the findings were first consolidated into an initial framework version (v1). This first version was then evaluated with practitioners and revised into v2, which is presented in Chapter 6.

This chapter presents the first version (v1), which establishes the underlying structure. Chapter 6 then presents the revised version (v2), which improves and expands it by adding depth and a more actionable view.

5.2 Framework structure

The framework organises the six ITOM domains into three layers around a central objective: organisational AI scaling (Figure 5.1). The layered structure reflects how participants described AI scaling in practice: direction steers, transformation realises value, and enabling foundations make scaling possible.

- **Direction** contains **Governance**, which sets mandate, ownership, prioritisation, and control.
- **Transformation** contains **People & Organisation** and **Processes**, where AI is translated into operational value through redesigned work, roles, and capabilities.
- **Enabling foundations** contains **Culture**, **Technology & Infrastructure**, and **Data**, which provide the conditions needed for scaling.

Direction without foundations stalls, and strong foundations without clear direction also do not lead to transformation. Transformation depends on both: AI creates operational value only where priorities are clear and the necessary organisational, technical, and data conditions are in place.

This layered structure is also consistent with the cross-domain dependencies identified in Chapter 4, where progress in one domain was described as dependent on conditions in others.

Together, the six domains provide a structured representation of the main operating model patterns that practitioners associate with scaling AI. In v1, both the structural relationship between domains and their associated patterns are presented within a single figure. Based on evaluator feedback, this combined view was reorganised in Chapter 6 into a condensed structural view, a more detailed domain-level expansion, and a more actionable view.

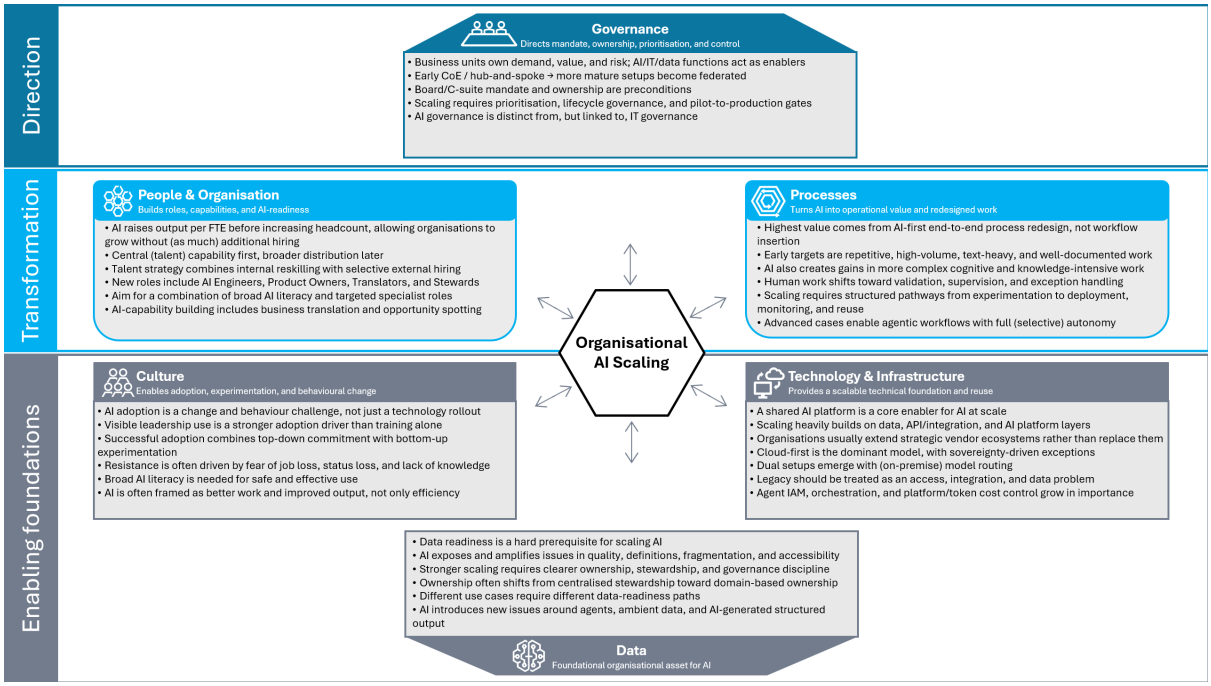


Figure 5.1: IT Operating Model framework for scaling AI (v1).

5.3 Use of the framework

The framework can be used in different ways. As a high-level view, it helps make sense of organisational AI scaling by showing how direction, transformation, and enabling foundations relate to one another. Used diagnostically, it helps identify where scaling stalls, for example when a clear mandate exists but data readiness is weak, or when technical foundations are in place but process redesign and adoption are still behind.

The later revised version extends this structural view through complementary representations. These additions preserve the same underlying framework, while making its domain content more explicit and practical.

5.4 Conclusion

This chapter presented v1 of the ITOM framework for scaling AI. The framework synthesises recurring patterns across the interviews into a three-layer structure of direction, transformation, and enabling foundations around the central objective of organisational AI scaling. As a first version, it primarily represents the structure of the findings. In the next chapter, this version is evaluated with practitioners and revised into an improved v2.

Chapter 6

Evaluation

6.1 Evaluation approach

The v1 framework (Chapter 5) was evaluated with five practitioners to assess its quality and identify improvements for a revised version (v2). Two evaluators had also participated in the interviews (I1 and I12) and three were new evaluators (E1, E2, and E3) who had not been involved in the empirical phase. The evaluators who had participated in the interviews could assess whether the framework reflected the interview input, while the new evaluators assessed it without prior exposure. The evaluation combined an indicative score (1–5) with qualitative feedback on three criteria: completeness, clarity, and practical usefulness.

6.2 Evaluation criteria

The framework was evaluated on three criteria that match the aim of the study: completeness, clarity, and practical usefulness.

- **Completeness:** whether the framework covers the main elements relevant to scaling AI in practice.
- **Clarity:** whether the framework's structure and formulation are understandable.
- **Practical usefulness:** whether the framework is useful in practice for reflecting on or shaping an operating model for AI scaling.

6.3 Evaluation participants

Table 6.1 summarises the evaluation participants and shows that the evaluation included both prior participants and external practitioners with relevant experience.

Table 6.1: Overview of evaluation participants

ID	Role	Sector & experience	Also interviewed
I1	Senior manager in AI consulting	Enterprise transformation / advisory context; 11+ years	Yes
I12	Senior advisor in AI transformation and operating model redesign	Enterprise transformation and digital AI advisory context; 25+ years	Yes
E1	Senior manager in technology consulting with CIO advisory experience	Technology consulting / CIO advisory context; 10+ years	No
E2	Partner in artificial intelligence, with prior C-suite, CEO, and global consulting leadership experience	AI leadership, banking, and global transformation context; 25+ years	No
E3	Director in technology consulting with CIO advisory focus on agile, IT cost, and legacy modernisation	Technology consulting / CIO advisory context; 10+ years	No

6.4 Evaluation procedure

Each evaluator received the v1 framework together with a short explanation and was asked to assess it on the three criteria, providing a score (1–5) and a few points of feedback per criterion. The evaluators reviewed the framework independently at their own pace. The feedback was the primary source for revising the framework into v2.

6.5 Evaluation results

Table 6.2 presents the evaluation results per evaluator, combining the scores on the three criteria with a summary of their feedback. E2 only provided qualitative feedback.

Table 6.2: Evaluation results per evaluator

ID	Compl.	Clar.	Use.	Feedback summary
I1	4	5	2	• Read the framework as an introductory "one-slider" that you could use to discuss with(in) organisations, but it is too abstract to be actionable; they suggested a "roadmap" style element so a change programme can be built around it (example). • Expected governance to reflect the business as the direction-setting party. • Expected regulatory compliance and ethics at governance level, via guardrails. • Expected strategic workforce planning (SWP) under People & Organisation.
I12	3	4	3	• Considered the framework a good starting point. • Identified the main missing information as AI compliance-oriented governance (EU AI Act, ethics, GDPR), particularly relevant in Europe. • Noted the high-level abstraction is strong, but organisations need more detailed elaboration to act on it.
E1	4.5	4	4	• Valued the layered structure and its domains. • Suggested adding an AI architect role under People & Organisation. • Suggested adding security under Technology & Infrastructure or under People & Organisation. • Requested clarification of "central (talent) capability first broader distribution later", "stewards", "AI-first end-to-end process redesign vs. workflow insertion", and "ambient data". • Suggested, (out-of-scope) future work, including a maturity model, a governance growth model, portfolio prioritisation funnel design, and role descriptions. • Requested clarity on what paths are referred to in "... require different data-readiness paths".
E2	–	–	–	• Provided qualitative feedback without numerical scores. • Regarded the framework as almost complete and clearly described. • Suggested cross-checking the "transformation" and "enabling foundations" layers against a practice model that distinguishes foundation and use more explicitly. • Noted that some organisations adapt their ecosystem mainly through AI-based partnerships (external), but still recommends sticking to v1's framework structure and logic. • Practical usability was found to need improvement, because it is hard for leaders or organisations to tell from the framework what to do next. The evaluator acknowledged the risk of scope creep, but suggested a measurement layer per element ("weak", "strong", and "very strong") to improve usefulness.
E3	4	4	2	• Found the model intuitive for a high-level overview. • Suggested adding strategy or vision under Governance. • Noted the importance of a platform with guardrails for cost and risk control, and asked where this sits in the model. • Pointed out that the bullets are a loose mix of benefits, risks, and neutral statements, leaving it unclear what to do with them. • Suggested sub-categories per domain and a subsequent "so what" step (e.g., patterns, actions, or ordering) to improve practical usefulness.

6.6 Revised framework (v2)

This section presents the revised framework (v2) and summarises the changes made in response to the evaluation feedback. Feedback was incorporated when it was raised by multiple evaluators, offered a clear improvement to the framework, and remained within the scope of this research.

The evaluation showed that v1 was appreciated as a clear high-level structure, but was also seen as too abstract to guide action directly. Evaluators asked for greater completeness in some domains, clearer wording, and more practical guidance on how the framework could be used. To address this, v2 was developed as a set of complementary representations of the same underlying framework and findings.

First, the condensed framework view includes the overall structure while making the domain logic visually clearer. Second, a more detailed domain-level expansion makes the main findings per domain more explicit and provides more depth without overloading the core figure. Third, a staircase-style representation adds a more practical progression-style view by showing how organisations can use the framework to think about staged operating model transformation when scaling AI.

Therefore, the revision did not simply add more content to the original figure. Instead, the main visual

was made more abstract, which made it possible to complement it with a separate depth layer and a more actionable usage-oriented view.

Table 6.3: Main changes made in v2 based on the evaluation

Criterion	Feedback	Response in v2
Completeness	Some domains needed more explicit coverage, especially Governance, People & Organisation, and Data.	Added missing content such as compliance and ethics in Governance, workforce planning and roles in People & Organisation, and clearer data-related patterns.
	Evaluators wanted the findings to be represented more fully.	Added a separate domain-level expansion to reflect the main findings in more depth without overloading the main figure.
Clarity	v1 was clear as a high-level structure, but some wording remained too abstract.	Made the main framework more concise and consistent by using shorter and clearer domain statements.
	Some bullets mixed observations, implications, and actions.	Moved detailed content out of the main figure and separated the framework into a condensed view and a more explicit detail view.
Practical usefulness	Evaluators found v1 useful as an overview, but too limited in showing what to do next.	Added a staircase representation to show how the framework can be used in practice.
	Several evaluators asked for more ordering, sequencing, or a clearer "so what".	Translated the framework into a progression-oriented view of the main operating model shifts involved in scaling AI.

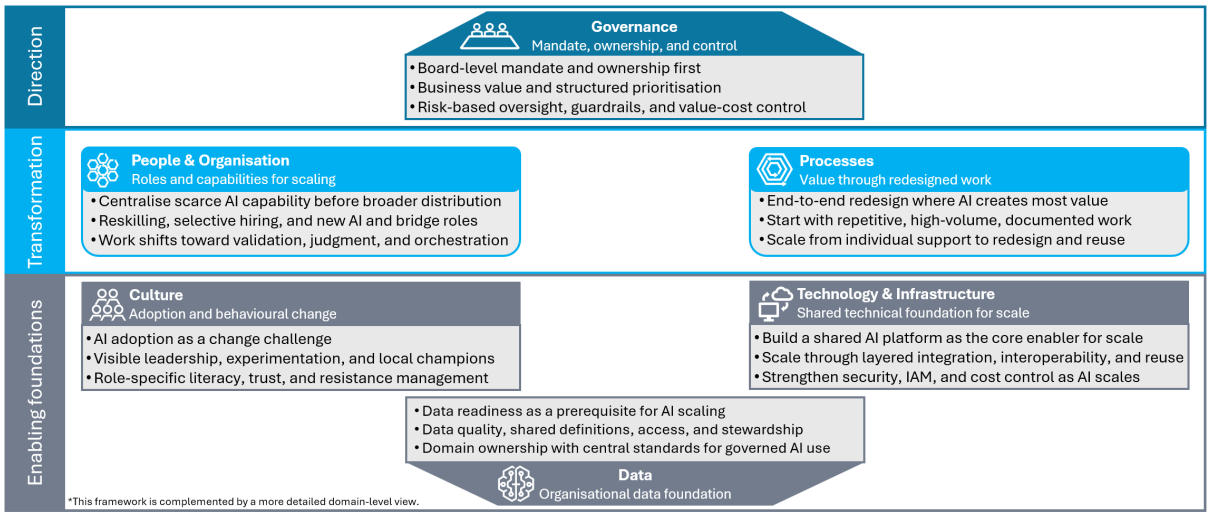


Figure 6.1: Condensed IT Operating Model framework for scaling AI (v2).

Governance Mandate, ownership, and control	People & Organisation Roles and capabilities for scaling	Processes Value through redesigned work
• Start with a board/C-suite mandate and clear ownership, framing AI as cross-functional transformation rather than isolated pilots, before scaling AI • Anchor AI in business strategy, with the business responsible for demand, value, and risk • Run AI governance as a distinct, multi-disciplinary practice (business, CIO, CDO, CISO, risk) with joint oversight, since traditional IT governance alone is insufficient • Use structured portfolio prioritisation based on value, feasibility, strategic fit, and organisational fit • Use governance to steer lifecycle decisions and staged gates from pilot to production • Put AI compliance, ethics, guardrails, and risk-based oversight in place as explicit governance responsibilities • Track value, adoption, and cost, using baseline comparisons, operational KPIs and AI-specific cost controls	• Concentrate scarce AI capability centrally first (e.g., a Centre of Excellence), then distribute it to business units as they mature • Combine reskilling with selective hiring, guided by strategic workforce planning • Introduce roles such as AI engineers, product owners, translators, data stewards (data-quality owners), architects, and platform-related roles • Use bridge roles to connect business needs, process knowledge, and AI implementation • Build both broad AI literacy and targeted specialist capability • Enable non-specialists to build with AI within clear guardrails • Shift work more toward validation, judgment, exception handling, and orchestration • Expect AI to raise value per person rather than cut headcount, redirecting gains to higher-value work and new oversight/quality-control tasks	• Redesign end-to-end processes where AI creates the highest value • Start where friction, repetition, volume, or documentation make AI impact visible • Move from individual support to team productivity, to broader redesign as AI matures • Build pathways from experimentation to deployment, monitoring, and reuse • Recognise that broader redesign depends on documentation, integration, system access, and data readiness • Expand selectively into more complex and knowledge-intensive work • Embed AI across workflows rather than adding it only to isolated tasks • Expect large efficiency gains (60–90% time reduction reported for suitable cases) • Evaluate impact via A/B comparison against pre-AI baselines
Culture Adoption and behavioural change	**Technology & Infrastructure** Shared technical foundation for scale	**Data** Organisational data foundation
• Treat AI adoption as a change challenge, not just a technology rollout • Use visible leadership behaviour as a stronger driver than communication alone • Combine top-down commitment with bottom-up experimentation and early adopters • Treat adoption as an explicit scaling workstream, using champions and ambassadors to translate it into daily practice • Address resistance around job impact, risk, and lack of knowledge directly • Frame AI as enabling higher-value, creative work rather than headcount reduction, to lower resistance • Build broad AI literacy tailored to different roles and needs • Create time, support, and psychological safety for learning and experimentation	• Build a shared AI platform as the core enabler for AI at scale • Build on data, integration, model, and AI platform layers rather than on point solutions • Focus on integration across enterprise systems, tools, and agents • Build for reuse, as fragmented tooling limits scaling beyond copilots and isolated use cases • Extend existing enterprise ecosystems rather than replacing them entirely • Make build-versus-buy decisions based on strategic relevance, regulatory context, and product maturity • Treat legacy mainly as an access, integration, and data-exposure issue rather than just as an age problem • Default to cloud-based scaling, with hybrid or sovereign setups where regulation, security, or sovereignty requires it • Build security and IAM with agent-specific access boundaries, alongside orchestration and cost/token control	• Treat data readiness as a hard prerequisite for AI scaling • Improve data quality, shared definitions, a single source of truth, and master data management • Improve ownership, stewardship, and governance discipline, restricting data access so agents and use cases see only what they need • Move toward more domain-based ownership while maintaining central standards • Recognise that use cases require different data-readiness paths (e.g., knowledge/document-based cases scale faster than structured, data-heavy predictive cases) • Prepare data not only for human use, but also for AI and agent use, including unstructured data and context • Govern both AI-consumed and AI-generated data and output

Figure 6.2: Expanded domain-level representation of the revised framework (v2). Dark blue = direction, light blue = transformation, and grey = enabling foundations.

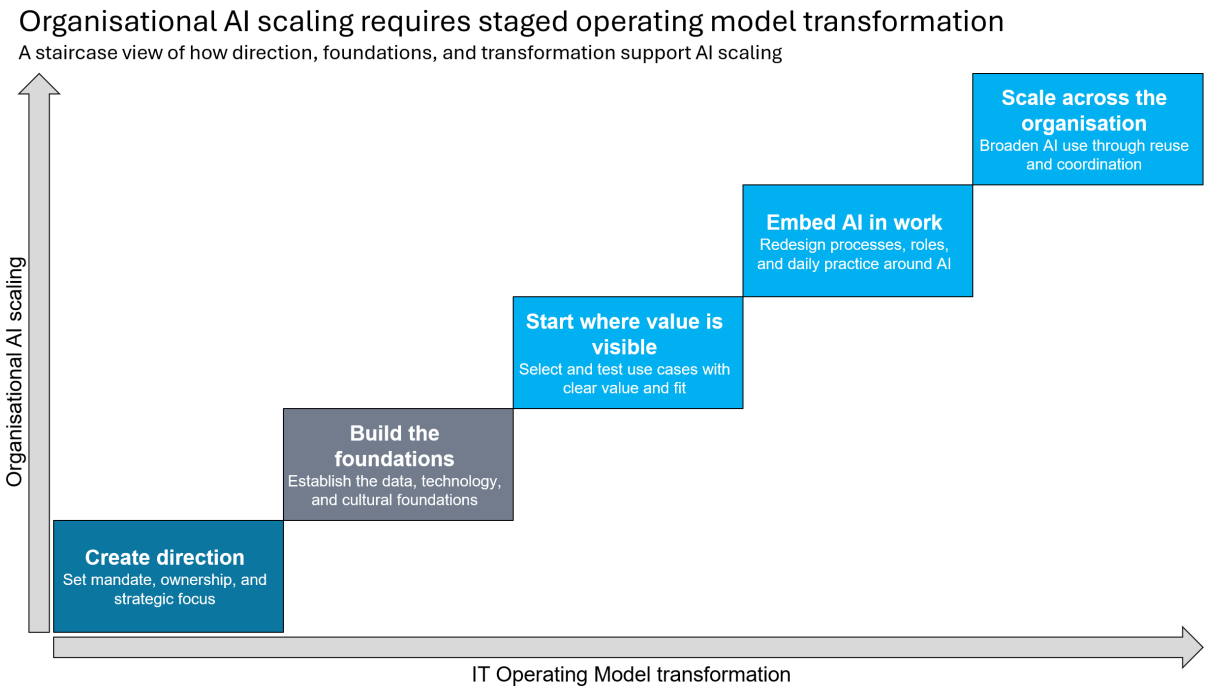


Figure 6.3: Staircase representation of the revised framework (v2). Dark blue = direction, light blue = transformation, and grey = enabling foundations.

For readability, full-size versions of the framework figures are provided in Appendix D.

Chapter 7

Discussion

7.1 Comparing modern AI scaling with the earlier RPA wave

7.1.1 Why compare AI scaling with the RPA wave?

A comparison with the earlier RPA wave helps take a step back and place current AI scaling in a broader automation context. From a strategic perspective, the question of how organisations move from experimentation to broader organisational deployment is not new. Earlier automation waves raised similar questions around where to start, how to organise for scaling, and why local technical success (e.g., a pilot/PoC) often does not translate into broader organisational deployment. The value of the comparison is therefore not in treating AI as a direct continuation of RPA, or in suggesting that the same principles apply again. Instead, it comes from showing which organisational patterns recur, which assumptions become outdated, and what this means for understanding AI scaling today.

The RPA literature provides a useful comparison point because it already documented a recurring gap between local automation success and broader organisational deployment. Early work described RPA as a relatively lightweight form of automation that could interact with existing systems through user interfaces, which made it attractive for repetitive, structured, and rule-based work [39, 17]. At the same time, later studies showed that scaling RPA still depended on more than the tooling itself. Scaling required governance, leadership support, process capabilities, organisational coordination, and continued attention to rollout and adoption [16, 19, 22]. In that sense, the RPA wave already showed that moving from pilots to broader scale is not only a technical issue, but also an organisational one.

This allows the findings of this research to be interpreted against RPA as an earlier automation wave that faced similar pilot-to-scale tensions. The comparison also helps clarify why modern AI scaling should not be reduced to earlier automation approaches. The discussion therefore focuses on what is similar, what is different, and what lessons can be drawn from the comparison.

7.1.2 What remains similar across the two waves?

The findings in this thesis show clear continuity with the earlier RPA wave. The strongest similarity is that scaling again emerges as an organisational challenge rather than as a purely technical one. Across the interviews, participants described governance, ownership, prioritisation, process discipline, and capability concentration as conditions for moving beyond pilots. This is in line with the later RPA literature, which also showed that broader organisational deployment depended on governance, management support, organisational capabilities, and structured implementation rather than on the technology alone [16, 19, 22].

A second similarity is the recurring role of centralisation in the early stages. In the interviews, many participants described central teams, centres of excellence, and shared governance structures as a natural starting point for scaling AI. This resembles the RPA literature, where broader organisational deployment also required concentration of expertise, support, and more coordination once organisations moved beyond isolated initiatives [16, 19]. In both waves, scaling appears to start with a concentration of scarce expertise, standard-setting, and organisational oversight before scaling becomes feasible.

A third similarity is seen in process selection. The RPA literature associated early automation value with repetitive, structured, and well-documented work [39, 17]. The findings in this thesis point to a comparable starting point for the first steps in AI scaling: participants often described the best starting points as areas with visible process friction, high volume, repetition, or strong documentation. In both cases, broader organisational deployment does not begin everywhere at once, but typically starts where value and feasibility can first be demonstrated.

These similarities support the view that the pilot-to-scale problem is not unique to AI. Earlier automation research already showed that local technical success does not automatically translate into scaling. Questions about ownership, prioritisation, process suitability, governance, and support structures were already central in the RPA wave and remain central in AI scaling today.

7.1.3 What is different in modern AI scaling?

Although the findings show important continuity with the RPA wave, they also point to significant differences. The first is the broader scope of what is being scaled. RPA was mainly associated with automating repetitive, structured, and rule-based work, often by mimicking human interaction with existing systems, where the robot essentially takes over manual user actions [39, 17]. In the interviews for this thesis, AI scaling was described more broadly. Participants linked it not only to task automation, but also to knowledge work, business model change, process redesign, shared AI platforms, data readiness, workforce change, and wider operating model alignment.

A second difference is the centrality of data. While data and integration challenges were also present in RPA literature, the findings in this thesis point to data readiness as a harder and more central condition for AI scaling. Participants described shared definitions, a single source of truth, stewardship, access, and the preparation of data for AI and agent use as major challenges. Compared with the earlier RPA wave, the underlying data foundation appears more consequential for scaling.

A third difference concerns the technical and architectural implications. RPA was often seen as a lightweight way to automate work on top of fragmented environments. In contrast, the interviews in this study pointed toward shared AI platforms, integration layers, orchestration, (agent) identity management, and architectural reuse as conditions for scaling. This suggests that modern AI scaling is less easily treated as an add-on layer and more often depends on architectural choices across the enterprise technology landscape.

A fourth difference is the degree of organisational redesign involved. The RPA literature already recognised the importance of process improvement and redesign, especially where poor process quality limited automation value [22, 17]. However, the findings in this thesis suggest that AI scaling pushes this even further. Participants reported that the highest value from AI does not come from inserting it into existing workflows, but from redesigning processes, roles, and decision structures around AI capabilities. In that sense, redesign is not only a supporting condition, but also part of where the value comes from.

Finally, the interviews point to stronger effects on work, capabilities, and adoption. Compared with the more bounded form of automation associated with RPA, modern AI appears to have broader implications for literacy, reskilling, jobs, role design, trust, and organisational change. This does not mean that RPA had no workforce implications, but the findings suggest that AI scaling affects how work is organised and performed across the organisation more deeply.

7.1.4 Lessons from the comparison

The comparison with the earlier RPA wave helps to learn from earlier automation lessons and to see where those assumptions become insufficient. One clear lesson is that the pilot-to-scale problem is not new. RPA already showed that scaling depends on more than promising use cases or local technical success. Governance, ownership, process suitability, support structures, and organisational capability remain necessary conditions.

At the same time, the findings in this thesis suggest that these earlier lessons are no longer enough on their own. Modern AI scaling appears to depend on a broader and more interdependent set of operating model changes. Compared with the RPA wave, the dependencies run more strongly across data, technology and infrastructure, shared platforms, process redesign, workforce change, adoption, and governance at the

same time. This makes it less suitable to approach AI scaling as a simple extension of earlier automation approaches.

This comparison therefore supports the main conclusion of this thesis: organisational AI scaling should be understood as an IT Operating Model challenge, not as a simple automation tool issue. Looking back at the RPA wave offers no template to copy; its value lies in showing which organisational challenges persist and which assumptions no longer hold true for AI. Some organisational challenges clearly return, but the findings of this thesis show that modern AI scaling requires broader and more coordinated transformation across the ITOM. The framework proposed in this research reflects these foundations and provides a structured way of understanding what is required to move beyond the pilot-to-scale gap.

Chapter 8

Conclusion

8.1 Answering the research questions

This thesis aimed to understand how organisations can move AI initiatives beyond pilots and proofs of concept toward broader organisational deployment, using the IT Operating Model (ITOM) as the structure for the analysis. It combined a review of the ITOM and AI scaling literature with fifteen qualitative interviews with practitioners directly involved in AI scaling, and synthesised the results into an ITOM framework that was subsequently revised through practitioner evaluation. The following subsections answer each research question.

8.1.1 SQ1

Which domains are most relevant for an ITOM according to the literature?

The literature describes the ITOM through varying domains, but six recurred consistently enough to structure this study: **Governance**, **People & Organisation**, **Processes**, **Technology & Infrastructure**, **Data**, and **Culture**.

Governance, people and organisation, processes, and technology & infrastructure appeared most consistently across the reviewed ITOM sources and formed the structural core. Culture was included because the transformation and adoption literature repeatedly pointed to behavioural change, leadership behaviour, and organisational acceptance as conditions for success. Data was included for a different reason: although it appears less often as a stand-alone top-level ITOM domain, the AI scaling literature treats it as a critical precondition for deployment. Together, these six domains form the structure for the empirical part of the thesis.

8.1.2 SQ2

What challenges, enablers, and existing approaches to scaling AI are identified in the literature?

The literature confirmed that organisations frequently stall between pilot and scale, and that the obstacles are as much organisational as technical. Recurring challenges included weak data quality and accessibility, difficulty attracting and developing AI talent, unclear ownership and governance, weak direction, limited value measurement, and insufficient process redesign and change support.

The corresponding enablers were clearer prioritisation of use cases, active leadership involvement, stronger governance and data ownership, cross-functional collaboration, more capable technical and data foundations, and explicit redesign of workflows and roles around AI.

Existing AI scaling frameworks explain parts of this picture, but do not translate the scaling problem into recurring operating model patterns across the ITOM domains. That is the gap this thesis addressed.

8.1.3 SQ3

Which recurring IT Operating Model patterns emerge from practitioner interviews on AI scaling?

The interviews showed that practitioners treat AI scaling as an operating model problem rather than a technology deployment, and recurring patterns emerged in every one of the six domains. In governance, scaling depended on a top-down mandate combined with clear business ownership and structured use-case selection. In people and organisation, expertise was typically concentrated first in a central Centre of Excellence and only later distributed toward federated or decentralised models as business units matured. New roles such as AI engineers, product owners, and translators emerged, while also reskilling the rest of the workforce. In processes, the largest gains came not from inserting AI into existing steps but from redesigning workflows around AI from first principles, with practitioners reporting large efficiency gains in specific cases, such as claims handling, legal intake, and payment arrangements, with time reductions in the region of 60–90% for suitable processes. In technology and infrastructure, scaling depended on shared platforms and integration layers rather than isolated tools. In data, readiness was found to be a hard precondition, with participants describing how AI amplifies existing data weaknesses rather than compensating for them. In culture, adoption hinged on AI literacy, visible leadership use, and framing AI as an enabler of higher-value work rather than purely as headcount reduction.

The most important cross-cutting finding was that these patterns are interdependent. A strong mandate stalls without data readiness; strong technical foundations deliver little without process redesign and adoption. Progress in any single domain is bounded by the maturity of the others.

8.1.4 SQ4

How can the findings from practitioner interviews be synthesised into an ITOM framework for AI scaling?

The recurring patterns were consolidated into an ITOM framework that arranges the six domains into three layers around the central objective of organisational AI scaling: **Direction** (Governance), **Transformation** (People & Organisation, Processes), and **Enabling foundations** (Culture, Technology & Infrastructure, Data). This grouping reflects how participants described scaling in practice: governance sets mandate and priorities, people and processes are where value is realised, and culture, technology, and data are the conditions that make deployment possible.

The framework was first built as a structural representation (v1) and then revised into v2 following evaluation by five practitioners. Because a single figure could not carry the full picture, v2 was presented in three views: a condensed structural view, an expanded domain-level view, and a staircase-style progression view. The result is a synthesis of recurring operating model patterns and their dependencies, rather than a maturity model or a fixed sequence of steps.

8.1.5 Main RQ

Which recurring ITOM patterns do practitioners associate with AI scaling initiatives from pilots toward broader organisational deployment, and how can these be synthesised into an ITOM framework for AI scaling?

Combined, the sub-questions answer the main research question. Practitioners associate AI scaling with recurring patterns spread across all six ITOM domains, and the main factor is how well those domains are coordinated: mandate and ownership, workflow and role redesign, capability building, technical foundations, data readiness, and adoption have to improve together rather than in isolation. Scaling therefore depends far less on the AI technology itself than on the surrounding operating model, which is why this thesis frames organisational AI scaling as an ITOM challenge. The framework captures this by organising the recurring patterns across the domains where operating model choices are made and by making their dependencies explicit, offering a structured way to reason about what it takes to close the pilot-to-scale gap.

8.2 Contributions

8.2.1 Academic contribution

This thesis contributes to the literature in three ways. First, it connects two bodies of work that are often kept apart, the IT Operating Model literature and the AI scaling literature, and shows that the challenges and enablers discussed in AI scaling map onto the established ITOM domains.

Second, it offers an empirically grounded framework derived from practitioner interviews. Where existing AI scaling frameworks emphasise capabilities, maturity stages, or implementation paths, this research adds a perspective built on real-world recurring operating model patterns across the six ITOM domains.

Third, the study reframes AI scaling as a coordination problem across the operating model. Its key empirical claim, that progress in one domain is constrained by the maturity of the others, offers a dependency-based view of AI scaling instead of the linear models seen in the current literature.

8.2.2 Practical contribution

The practical value of this thesis lies in the framework and in how it can be applied. It gives IT leaders and practitioners a structured way to see where AI scaling tends to stall and which operating model conditions need attention before further investment pays off. It also offers insights from practitioners with direct experience with AI scaling.

The framework helps practitioners to:

- structure the scaling challenge across the main ITOM domains;
- identify where weakness in one domain is holding back progress in others;
- understand that AI scaling is not a technology issue, but also one of ownership, redesign, capability building, data readiness, and adoption;
- use the staircase view as a starting point for staged operating model change when scaling AI.

8.3 Limitations

This study has several limitations.

- **Exploratory qualitative design.** The research is based on an exploratory qualitative design. This makes it suitable for identifying recurring patterns and developing a framework, but it does not allow causal claims or statistical analysis.
- **Sample size and composition.** The study is based on a limited number of interviews with experienced practitioners. Although this provided strong qualitative insight, the sample does not represent all sectors, organisational sizes, or levels of AI maturity equally.
- **Self-reported impacts.** The findings rely on practitioner accounts of decisions, effects, and organisational patterns. Although valuable, they do not provide direct verifiable evidence of impact in organisations.
- **Consulting environment and access.** The study was conducted in the context of a PwC internship and partly relied on the PwC network for access to participants. This helped provide access to experienced respondents, but it may also have influenced the type of organisations, perspectives, and strategic concerns represented in the sample.
- **European and Dutch context.** Many of the perspectives in this study are based on Dutch and broader European organisational contexts. This matters for topics such as governance, regulation, responsible AI, labour structures, work culture, and data protection. The findings may therefore not transfer directly to non-European contexts.
- **Transferability.** The framework is intended as an empirically grounded and analytically useful representation of recurring patterns, not as a universal truth that applies in the same way to every organisation.

8.4 Future work

This thesis opens several directions for future research.

- **Empirical validation.** The framework could be tested and improved through additional empirical research in other sectors, organisational settings, and levels of AI maturity.

- **Cross-regional comparison.** Future studies could compare AI scaling patterns across different geographical contexts to learn how regulatory environments and traditions shape the operating model.
- **Longitudinal research.** Longitudinal research could help show how organisations move over time from early pilots to mature organisational deployment and how operating model choices evolve during that process.
- **Case-based research.** This study focused on breadth across practitioner perspectives. Future work could instead study individual organisations in more depth.
- **Extension into practical instruments.** Future research could build on the framework by developing complementary practical tools, diagnostics, maturity assessments, or workshop formats for practitioners.

Bibliography

[1] S. Alter. Work system theory: Overview of core concepts, extensions, and challenges for the future. *Journal of the Association for Information Systems*, 14(2):1–70, 2013. doi: 10.17705/1jais.00323.

[2] Y. Annepaka and P. Pakray. Large language models: a survey of their development, capabilities, and applications. *Knowledge and Information Systems*, 67(3):2967–3022, 2025. doi: 10.1007/s10115-024-02310-4.

[3] Axelos. *ITIL Foundation: ITIL 4 Edition*. The Stationery Office, 2019.

[4] M. Berndtsson, A.-M. Ericsson, and T. Svahn. Scaling up data-driven pilot projects. *AI Magazine*, 41(3):94–102, 2020. doi: 10.1609/aimag.v41i3.5307.

[5] A. Bhashyam. How shell scaled ai predictive maintenance to monitor 10,000 pieces of equipment globally, 2022. URL `https://c3.ai/blog/how-shell-scaled-ai-predictive-maintenance-to-monitor-10000-pieces-of-equipment-globally/`. C3 AI Blog.

[6] I. Bojinov. Keep your ai projects on track. *Harvard Business Review*, 101(6):53–59, 2023.

[7] A. Chandrasekaran. Why 50% of genai projects fail — and how to beat the odds, 2026. URL `https://www.gartner.com/en/articles/genai-project-failure`. Gartner.

[8] Forrester Research. The high-performance it operating model, February 2024. URL `https://www.forrester.com/report/high-performance-requires-the-right-it-operating-model/RES165791`.

[9] Forrester Research. Assess your ai maturity: Discover where you stand with your artificial intelligence maturity, September 2025. URL `https://www.forrester.com/report/assess-your-ai-maturity/RES186200?ref_search=4768789_1784463487169`. Maturity Assessment.

[10] Gartner. How cios deliver and show value from the it operating model, 2024. URL `https://www.gartner.com/en/chief-information-officer/insights/it-operating-model`. Gartner Insights.

[11] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, pages 2672–2680, 2014.

[12] M. Grebe, M. R. Franke, and A. Heinzl. Artificial intelligence: How leading companies define use cases, scale-up utilization, and realize value. *Informatik Spektrum*, 46(4):197–209, 2023. doi: 10.1007/s00287-023-01548-6.

[13] N. Haefner, V. Parida, O. Gassmann, and J. Wincent. Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197:122878, 2023. doi: 10.1016/j.techfore.2023.122878.

[14] R. He, J. Cao, and T. Tan. Generative artificial intelligence: a historical perspective. *National Science Review*, 12(5):nwaf050, 2025. doi: 10.1093/nsr/nwaf050.

[15] J. C. Henderson and N. Venkatraman. Strategic alignment: A model for organizational transformation via information technology. In T. J. Allen and M. S. Scott Morton, editors, *Information Technology and the Corporation of the 1990s: Research Studies*, pages 202–220. Oxford University Press, 1994. doi: 10.1093/oso/9780195068061.003.0009.

[16] L. V. Herm, C. Janiesch, A. Helm, et al. A framework for implementing robotic process automation projects. *Information Systems and E-Business Management*, 21:1–35, 2022. doi: 10.1007/s10257-022-00553-8.

[17] P. Hofmann, C. Samp, and N. Urbach. Robotic process automation. *Electronic Markets*, 30:99–106, 2020. doi: 10.1007/s12525-019-00365-8.

[18] J. Holmström and J. Magnusson. Navigating the organizational ai journey: The ai transformation framework. *Business Horizons*, 69(1):89–100, 2026. doi: 10.1016/j.bushor.2025.01.002.

[19] D. M. Q. Huynh and I. Lindgren. Beyond the hype of robotic process automation (rpa): on conditions needed to implement rpa in organizations. *Information Systems and E-Business Management*, 2026. doi: 10.1007/s10257-026-00721-0.

[20] ISACA. *COBIT 2019 Framework: Introduction and Methodology*. ISACA, 2018.

[21] E. Kalliamvakou. Research: Quantifying github copilot's impact on developer productivity and happiness, 2022. URL `https://github.blog/news-insights/research/research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/`. GitHub Blog.

[22] Sommai Khantong and Pankom Sriboonlue. Robotic process automation in business process management: A systematic literature review and an integrated framework. *Technologies*, 14(4):225, 2026. doi: 10.3390/technologies14040225.

[23] V. Khatri and C. V. Brown. Designing data governance. *Communications of the ACM*, 53(1): 148–152, 2010. doi: 10.1145/1629175.1629210.

[24] McKinsey & Company. The state of ai: How organizations are rewiring to capture value, 2025. URL `https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value`. McKinsey Global Report.

[25] McKinsey & Company. Seizing the agentic ai advantage, 2025. URL `https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage#/`. McKinsey Report.

[26] McKinsey & Company. Deploying ai at speed and scale: Talking with ing's marnix van stiphout, 2025. URL `https://www.mckinsey.com/industries/financial-services/our-insights/deploying-ai-at-speed-and-scale-talking-with-ings-marnix-van-stiphout`. McKinsey Interview.

[27] Morgan Stanley. Morgan stanley wealth management announces latest game-changing addition to suite of genai tools, 2024. URL `https://www.morganstanley.com/press-releases/ai-at-morgan-stanley-debrief-launch`. Press Release.

[28] PwC. Information technology operating model, 2025. Internal document, used with permission.

[29] PwC. 29th global ceo survey 2026: Your next move: Reinvent to outperform, 2026. URL `https://www.pwc.com/gx/en/issues/c-suite-insights/ceo-survey.html`. PwC Global CEO Survey.

[30] J. W. Ross, P. Weill, and D. C. Robertson. *Enterprise Architecture as Strategy: Creating a Foundation for Business Execution*. Harvard Business School Press, 2006.

[31] J. Ryseff, B. F. De Bruhl, and S. J. Newberry. *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed*. RAND Corporation, 2024. doi: 10.7249/RRA2680-1.

[32] R. Sapkota, K. I. Roumeliotis, and M. Karkee. Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges. *Information Fusion*, 126:103599, 2026. doi: 10.1016/j.inffus.2025.103599.

[33] S. S. Sengar, A. B. Hasan, S. Kumar, and F. Carroll. Generative artificial intelligence: a systematic review and applications. *Multimedia Tools and Applications*, 84(21):23661–23700, 2025. doi: 10.1007/s11042-024-20016-1.

[34] D. Sjödin, V. Parida, M. Palmié, and J. Wincent. How ai capabilities enable business model innovation: Scaling ai through co-evolutionary processes and feedback loops. *Journal of Business Research*, 134:574–587, 2021. doi: 10.1016/j.jbusres.2021.05.009.

[35] M. Suleman, S. Pekkola, J. Ralyté, and T. Ahola. It operating models: A systematic review of academic and practitioner literature. *Information Systems Management*, 43(1):22–44, 2025. doi: 10.1080/10580530.2025.2513305.

[36] T. Talaei Khoei, H. Ould Slimane, and N. Kaabouch. Deep learning: systematic review, models, challenges, and research directions. *Neural Computing and Applications*, 35(31):23103–23124, 2023. doi: 10.1007/s00521-023-08957-4.

[37] D. Tamburri, M. Tonnarelli, and J. van Hillegersberg. Challenges in adoption and scaling of ai: A case study at a high-tech firm and research roadmap. In M. Schreieck, I. Oshri, J. Kotlarsky, and O. Krancher, editors, *Bridging Digital Sourcing, Platforms, and Ecosystems*, volume 563 of *Lecture Notes in Business Information Processing*, pages 56–77. Springer, 2025. doi: 10.1007/978-3-032-04512-6_4.

[38] S. Thomke. *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*. Harvard Business School Press, 2003.

[39] W. M. P. van der Aalst, M. Bichler, and A. Heinzl. Robotic process automation. *Business & Information Systems Engineering*, 60:269–272, 2018. doi: 10.1007/s12599-018-0542-4.

[40] M. Weber, M. Engert, N. Schaffer, et al. Organizational capabilities for ai implementation—coping with inscrutability and data dependency in ai. *Information Systems Frontiers*, 25:1549–1569, 2023. doi: 10.1007/s10796-022-10297-y.

[41] P. Weill and J. W. Ross. *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Harvard Business School Press, 2004.

[42] World Economic Forum. Ai agents in action: Foundations for evaluation and governance, November 2025. URL `https://www.weforum.org/publications/ai-agents-in-action-foundations-for-evaluation-and-governance/`. White Paper.

[43] J. Wulf, T. Mettler, and W. Brenner. Using a digital services capability model to assess readiness for the digital consumer. *MIS Quarterly Executive*, 16(3):171–195, 2017.

Appendix A

Interview Guide

Introduction

This interview explores the operating model decisions IT leaders and practitioners made when scaling AI initiatives from pilot to broader organisational deployment. The goal is to understand the operating model decisions made by you or your organisation across different domains: governance, people and organisation, processes, technology and infrastructure, data, and culture, and how those decisions affected the scaling process.

The interview will take approximately 60 minutes and is organised into three parts: context, detailed domain decisions and impact, and AI scaling experience and reflection.

Part A: Context (5 minutes)

1. Could you briefly describe your role and your involvement in AI initiatives within your organisation(s)?
2. How would you describe your organisation's overall AI maturity?
3. Which AI initiative or programme are you most closely involved with? What is its current status—pilot, in production, scaling, or something else?

Part B: ITOM-domain decisions and impact (40 minutes)

For each domain below, we explore the operating model decisions made when scaling AI and their impact. Each domain starts with an open question that does not require knowledge of operating model terminology. The interviewer uses the sub-questions to follow up where needed and to make sure all six ITOM domains are covered, even when the participant does not use the same vocabulary.

B1. Governance

Main question: How are AI initiatives steered and overseen in your organisation? For example, who decides which initiatives go ahead, who is accountable, and how are risk and compliance handled? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- Who is responsible for AI investment decisions and prioritisation?
- How do you approach responsible AI governance (ethics, risk, compliance)?
- Has the governance structure changed as AI initiatives scaled?

- How do you measure success of AI initiatives? Do you track specific KPIs?
- *[Environment]* Has regulation (e.g., EU AI Act) affected your governance approach?

B2. People & Organisation

Main question: How are AI initiatives organised in your company in terms of people and teams? For example, what is the organisational structure, do you have a dedicated AI team, are AI activities spread across business units, and what kind of roles or skills are involved? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- Is there a central AI team, distributed teams across business units, or a hybrid model?
- How do AI teams relate to the broader IT organisation?
- Have you created new roles or positions specifically for AI (e.g., ML engineers, AI product managers)?
- How do data scientists, engineers, and business stakeholders collaborate?
- What skills or roles have been hardest to find or develop?
- How has this organisational structure evolved as you scaled?

B3. Processes

Main question: How is AI changing the way actual work gets done in your organisation? For example, are there processes where AI or agents now handle tasks that people used to do, and how has that changed the work, and what changed in the support- and IT-development processes? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- Which core processes have been most affected by AI in your organisation?
- Have you redesigned workflows around AI, or is AI added to existing processes?
- What kind of efficiency or FTE impact have you seen in those processes?
- Are there processes that are handled fully or partly by agents, with limited human involvement?
- How does an AI use case typically move from idea to actual use in a process?
- Do you use specific methodologies for AI delivery (e.g., Agile, MLOps, LLMOps)?

B4. Technology & Infrastructure

Main question: What technology and infrastructure decisions have you made to support AI at scale? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- Do you have a dedicated AI/ML platform?
- Do you use cloud-based AI services? (E.g., AI orchestration, AI workflow tool).
- How is model deployment, monitoring, and retraining handled?
- How do you manage integration with existing enterprise systems?
- Have legacy systems been a constraint? How did you address this?
- *[Sourcing]* How do you make build-versus-buy decisions for AI capabilities? Do you rely on external vendors or cloud AI services?

B5. Data

Main question: How is data handled to support AI initiatives in your organisation? For example, where does the data come from, who manages it, and is it ready to be used? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- How do you ensure data quality for AI use cases?
- Is there a data governance framework? Who owns and manages it?
- How accessible is data across departments and teams for AI initiatives?
- Has data quality or data accessibility been a bottleneck for scaling? How did you address it?
- How has your data architecture evolved as you scaled AI?

B6. Culture

Main question: How do people in your organisation react to AI? For example, is there enthusiasm, concern, or both, and what does leadership do about it? **Follow-up:** what was the impact of these decisions?

Sub-questions:

- Have you taken specific actions to improve AI literacy, trust, or adoption across the organisation?
- Is there broad organisational support for AI, or do you encounter resistance?
- How does leadership signal the importance of AI to the organisation?
- How do employees outside AI teams perceive AI adoption in their work?
- Have there been challenges to adoption? How did you address them?

Part C: AI scaling experience and reflection (15 minutes)

- What were the biggest challenges your organisation faced when moving AI from pilot to broader organisational deployment?
- Were there trade-offs or dependencies between decisions in different domains (e.g., a governance choice that affected how you organised people, or a technology choice that required process changes)?
- If you could redesign your IT operating model for AI from scratch, what would you do differently?
- Looking back at the operating model decisions you made across these domains, which ones were most critical to scaling successfully?
- Are there aspects of scaling AI or operating model decisions that we have not discussed but that you consider important?

Closing

Thank you for your time and insights. I will send you a summary of the interview for review within one week. Do you have any questions?

Appendix B

Interview summaries

The full interview summaries are included for transparency and as a readable substitute for raw transcripts, while the compact summaries reflect the more condensed format used in the cross-interview analysis.

B.1 Interview 1 summary

Table B.1: Structured summary of interview 1

Interview details	Interview 1; 1 May 2026; Senior manager in AI consulting; enterprise transformation context; 11+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Position AI as a default design option for new solutions, but only apply it where the expected benefits clearly outweigh the costs. • Set up AI governance as a separate practice rather than reuse existing IT governance one-to-one. • Invest in risk officer capacity and in their ability to assess AI-specific risks. • Implement tooling and standardised templates to check whether AI applications meet legal and regulatory requirements. • Prioritise AI initiatives along three axes: business impact, feasibility, and how predictive the use case is for broader AI adoption. • Treat AI scaling not as a collection of isolated use cases, but as a broader transformation that requires parallel investment in enabling conditions such as architecture and AI-ready data. • Define readiness conditions per transformation area, such as architecture, AI-ready data, governance, and skills, as a precondition for starting execution.	• AI governance structurally resembles existing IT governance, but the implications are larger and require a different way of thinking. • Risk officers and assessors need new skills to interpret AI-specific risks. • Standardised assessment tooling makes governance scalable across many AI use cases. • Prioritisation based on broader transformation value (not only impact and feasibility) helps select use cases that unlock further AI scaling. • Traditional KPIs remain leading; AI is treated as a means to improve them, not as a goal in itself.

Domain	Key decisions	Reported impact
People & Organisation	• Categorise the entire workforce into three personas: no-code, low-code, and pro-code. • Use these personas as the basis for upskilling, reskilling, training, and strategic workforce planning. • Conduct strategic workforce planning to define the desired persona mix in five years and align hiring and attrition accordingly. • Organise AI delivery in distributed teams linked to business lines for smaller organisations, and move toward a hub-and-spoke model with a centre of excellence in larger organisations. • Introduce new roles such as agentic engineer, agentic platform lead, and additional solution architects to support the agentic development platform. • Treat prompt engineering, context engineering, and (to a lesser extent) harness engineering as required skills for new hires, not as separate job titles. • Hire externally for the first wave of agentic roles and reskill internal staff into these roles in parallel.	• The relationship between value growth and headcount growth changes: value created grows much faster than the number of people hired. • In some organisations, AI investment is partly funded by slowing workforce growth or by redirecting expected personnel expansion into technology-enabled productivity gains. • Some roles become unsustainable and require either reskilling or a managed transition out of the organisation. • Productivity gains in core operational processes can be very large; up to 60–90% time reduction was observed in claim processing at peer organisations. • Distributed AI teams fit smaller organisations; larger organisations need more federated structures, but more complex models also slow decision-making down. • The persona model gives a practical handle on which training and reskilling each group needs, instead of treating the workforce as a single audience. • The stronger focus on AI-ready data increases demand for roles such as data stewards.
Processes	• Build a dedicated internal agentic development platform so that solutions can be developed safely, responsibly, and consistently, without organisational data leaking into public AI tools. • Move away from traditional IT delivery routines toward an agent-oriented way of designing solutions, instead of inserting AI into existing development processes. • Treat guardrails as a first-class part of process design, especially for customer-facing core processes. • Extend DevOps practices into MLOps and LLMOps for AI workloads. • Apply AI to core operational processes such as claim processing and customer interaction, including fully autonomous handling of payment arrangements. • Define a broad testing portfolio (unit, integration, system, user acceptance) as the basis for trusting AI-generated code and agentic behaviour.	• Process design shifts from automating individual steps to orchestrating intelligent workflows in which agents perform tasks. • Agentic systems introduce new layers of complexity, especially around identity and access management, multi-agent coordination, and monitoring of non-deterministic behaviour. • Customer interaction processes can be largely automated when guardrails are strong enough; in some cases the human-in-the-loop is no longer needed. • Strong testing becomes the practical mechanism that allows organisations to trust AI-generated solutions without manually reviewing every line of code. • Process change and infrastructure change happen together; they cannot be separated in practice.
Technology & Infrastructure	• Continue to use existing strategic vendors and extend them with AI capabilities, rather than replace the stack. • Standardise on cloud infrastructure for AI workloads; on-premise is mainly retained in parts of the public sector. • Invest substantially in API management and identity and access management, including agent identities (e.g. Entra agent ID). • Invest in a modern data platform such as Databricks to support AI delivery. • Build and maintain an internal agentic platform, and budget explicitly for ongoing token consumption as a structural cost. • Phase out low-code platforms over time, because agentic coding assistants (Codex, Claude Code, Gemini) make pro-code more accessible.	• AI scaling does not require replacing the existing technology stack, but does require deliberate platform extension. • Cloud is in practice cheaper, more maintainable, and more scalable than on-premise; some organisations are actively decommissioning on-premise infrastructure. • Token consumption becomes a recurring operational cost that must be planned for. • Identity and access management becomes much more complex when agents act on behalf of users and interact with each other. • Leadership must communicate AI ambition, including its financial implications, with cost growth shifting from headcount toward IT and platforms.

Domain	Key decisions	Reported impact
Data	• Treat data readiness as a precondition for AI transformation, not as a parallel workstream. • Invest in monitoring, structuring, and broader availability of data across the organisation. • Buy or develop additional data where existing data is insufficient. • Increase the use of data products to make trusted data more tangible, accessible, and usable for teams working on AI-related initiatives. • Standardise core business concepts through data products with one shared definition per concept. • Strengthen the role of data steward to safeguard data quality at scale. • Adopt a deliberate, layered data architecture, in particular medallion architecture (bronze, silver, gold), so different consumers (executives, finance, agentic systems) can use the layer that fits their purpose. • AI is both a consumer and a generator of data: operational interactions and unstructured inputs can be transformed into structured data, which then becomes part of the enterprise data foundation and requires governance like other business data. • Revisit existing data and probability models, as they often need to be redesigned for AI use.	• Some AI transformations require an explicit preparation phase in which data first has to be brought to the required quality level before use cases can be scaled. • Data emerges as the hardest gating factor for AI scaling: ambitious ideas could not be executed until the data foundation was improved. • Better data design changes how information is represented and consumed across the organisation. • Different stakeholders need access to different data layers; for example agentic systems handling claims need access to raw bronze data, while executive reporting only uses the gold layer. • More structured data work raises requirements on non-specialists, who need to work with cleaner inputs and a more deliberate way of thinking about data. • The stronger focus on AI-ready data increases demand for more people with data responsibilities, including data stewards responsible for quality, availability, and maintainability.
Culture	• Aim for an innovation culture rather than for a separate AI culture. • Combine bottom-up enablement (employees skilling themselves) with explicit top-down commitment from leadership. • Have leadership communicate AI ambition openly and without hedging. • Set up a dedicated workstream that addresses employee concerns about AI: adoption, reskilling, and a clear path for people whose roles disappear, including support to find work elsewhere. • Position AI as a default design option for new solutions: every new initiative must be considered in terms of how AI could be used.	• Some employees feel real uncertainty about their future role, while most accept that AI must be embraced. • Visible top-down commitment, combined with concrete support mechanisms, lowers resistance to AI adoption. • Investing in trainings that some employees do not actually need is a real risk if the persona model is not used. • Leaders are taking the workforce concern seriously, with explicit support for transition, reskilling, and outplacement. • Culture functions in practice as a precondition for scaling rather than as a soft side-topic.

B.2 Interview 2 summary

Table B.2: Structured summary of interview 2

Interview details	Interview 2; 7 May 2026; AI product manager and innovation centre lead; 7+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Move AI ownership out of IT alone and into a central C-suite role, ideally CEO or COO. • Treat AI as a broad transformation across multiple disciplines instead of a tech topic. • Give AI initiatives explicit top-down priority. • Make data privacy, risk, and compliance part of the AI governance setup from the start. • Use mandate and clarity from the top as the main lever to set direction and pace.	• A central C-suite mandate is the main factor for whether AI scaling succeeds or stalls. • Without top-down mandate, AI stays a side topic in IT and the broader transformation does not happen. • Public sector organisations are further behind on this, partly because they lack the modern models and tooling to feel the urgency. • A weak or ambiguous top-down mandate causes AI efforts to stall at middle-management level. • A clear mandate creates direction for middle management and unlocks investment in foundational tooling, talent, and processes.

Domain	Key decisions	Reported impact
People & Organisation	• Move toward smaller, high-performing teams with fewer organisational layers to reach strategic goals faster. • Merge traditional disciplines (design, product, research, engineering) into hybrid roles, where designers do front-end development and engineers act as broader "superpower engineers". • Hire selectively, with a preference for candidates who show high agency, drive, and concrete AI experience over candidates with mainly formal credentials. • Slow down hiring overall as AI raises individual productivity, and reshape the type of profiles being hired. • Treat individual experimentation and self-directed AI use as part of the role, not as optional.	• Traditional role boundaries blur, which leads to smaller and more multidisciplinary teams. • The bar for individual contribution rises: people who refuse to use AI are becoming hard to justify in a high-performing environment. • Hiring shifts toward people who can demonstrate AI fluency in practice, for example through their own agents, rather than toward formal AI qualifications. • AI raises concern about job security; works councils are becoming involved in this discussion.
Processes	• Adopt agentic development environments such as Cursor and Claude Code as the core of the development workflow. • Use background agents that test bugs and front-ends continuously, including outside working hours. • Implement reporting agents to automate recurring reporting tasks. • Redesign core processes with an AI-first mindset rather than inserting AI into existing process steps. • Apply AI to lean processes and high-volume document work, for example rewriting letters or producing memos. • Build a reusable "harness" (platform, tools, settings) so that individuals can personalise their AI setup on top of a shared foundation.	• Process redesign is moving from incremental automation toward fully reimagining processes. • Concrete time reductions are substantial. • Government processes are estimated to allow up to 90% time reduction depending on the process; for example a letter that previously needed five reviewers can be handled by one. • In some contexts, the pressure to automate is increased by the fact that external parties are already using AI to scale interactions, for example in objection-handling or subsidy-application processes by citizens, which can increase workload. • Reaching this level of redesign requires explicit mandate and vision from the top; many organisations are not yet mature enough for this.
Technology & Infrastructure	• Standardise on a small number of foundational tools (such as Cursor, Claude Code, and a chosen harness) and build the rest of the stack around them. • Consolidate vendors, with CIOs moving toward fewer providers to reduce fragmentation. • Use MCP and similar interfaces to connect tools, agents, and enterprise systems. • Invest in a foundational platform with shared infrastructure, settings, and access controls so individuals and teams can build on top of it. • Plan for sovereignty and access to the right models, including local deployment where regulation requires it. • Anticipate that older monitoring and tooling, such as legacy observability stacks, will need to be replaced. • Decide not only which model to make available, but also which harness or working environment is provided. • Treat power capacity as a real constraint for local model deployment in some public sector contexts.	• Central foundational platforms make decentralised building by teams and individuals possible. • Consolidating around a small set of harnesses and providers gives more leverage and reduces complexity. • Agentic coding tools are drastically improving model efficiency and developer output. • Older tooling that does not fit AI workflows is being phased out and replaced. • Local model deployment is constrained by sovereignty requirements and limited power capacity, particularly in the public sector. • Centralised foundational platforms, similar to a Centre of Excellence, are enablers that make decentralised building by teams and individuals possible.

Domain	Key decisions	Reported impact
Data	• Treat "garbage in, garbage out" as the operating principle: AI quality is bounded by the quality of the data it can reach. • Connect AI tools to all relevant systems (email, Slack, CRM, documents) so context, memory, and inputs become available. • Move toward "ambient data": data captured passively from existing tools and meetings (for example transcript folders from note-taking tools) instead of relying on manual data entry. • Use unstructured data more aggressively, since modern AI handles it better than older systems. • Manage agent memory and context explicitly, for example through structured markdown files (claude.md and similar) until better tooling exists.	• Organisations that connect their systems and make data retrievable through standardised interfaces are front-runners in AI scaling. • The urgency of having a data foundation in order is increasing because AI failures in production are often traced back to disconnected or low-quality source systems. • Manual data entry decreases rather than increases; the workforce is filling in less data, while AI agents pull richer context from existing sources. • Memory and context management for agents is currently fragmented and ad-hoc, yet very important for AI scaling.
Culture	• Build a culture in which around 80% of AI success is treated as coming from individual employees and around 20% from formal innovation, in line with the framing in recent industry reporting. • Make permission and active encouragement from senior leaders a deliberate part of the operating model, not an implicit signal. • Create dedicated channels for celebrating AI "cool wins" so good practice spreads. • Bring AI usage into evaluations, for example by reflecting on token usage and concrete examples of AI-assisted work, while being aware that token counts alone are an imperfect metric. • Treat experimentation time as an investment that may not pay off immediately, and budget for that explicitly. • Involve works councils early in the conversation about job impact and AI adoption.	• Permission and encouragement from leadership are one of the main reasons for whether AI adoption succeeds in a team. • Without permission, AI training programmes alone do not change behaviour. • Visible practice from senior leaders pulls others along; absence of it makes adoption stall. • Employees develop real concern about job security, which has to be addressed openly rather than avoided. • Personal investment in AI use becomes a differentiator: those who experiment build a clear advantage in a high-performing environment.

B.3 Interview 3 summary

Table B.3: Structured summary of interview 3

Interview details	Interview 3; 14 May 2026; Co-founder and CEO of an AI-native scaleup; 15+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Treat AI as a board-level topic rather than as an IT-only initiative. • Set up a consortium-style governance structure in which a central data and/or AI platform team is the agenda-driver, surrounded by data privacy officers and security officers as influencers. • Give the CISO an active, sometimes initiating role in AI initiatives, especially around AI security and exposure risks. • Reserve AI funding as part of a broader organisational programme rather than pulling it ad-hoc out of departmental budgets. • Treat the business as the owner of AI demand and IT as the delivery function, so that initiatives start from a real business problem rather than from technology push. • Get cost control structures in place for token consumption, including budget caps per user where consumption runs out of hand.	• AI is no longer a technology side-topic; multiple disciplines (finance, security, IT, business) are now formally involved in steering AI. • A future-proof operating model reduces the cost and time of experimentation, so ideas and hypotheses can be tested, stopped, and revisited quickly as models improve. • Cost is becoming a governance topic in its own right because token consumption can exceed development salaries in some organisations. • KPIs for AI scaling are sometimes still missing. • FTE reduction and topline revenue impact from AI are still limited in many organisations, despite high investment levels. • Change management is a major reason why organisations may invest in AI while still struggling to realise visible enterprise-wide impact. • Without business ownership, AI risks becoming technology looking for a problem instead of solving one.

Domain	Key decisions	Reported impact
People & Organisation	• Set up a Centre of Excellence (CoE) that brings people, process, and technology together centrally, with decentral business-unit teams consuming shared services and localising them. • Use a hub-and-spoke model: the CoE owns best practices, blueprints, shared resources, and core technology; business domains build their own solutions on top of those shared components. • Trends are visible where existing work is being recorded to train AI models (e.g. in radiology). • Use ambassadors for collaboration between engineers and non-engineers. • Plan for the CoE to dissolve over time as AI capabilities mature and move into the domains. • Hire and retain scarce AI-specific roles centrally first (data stewards, ML engineers, agentic engineers), where existing data scientists and ML engineers tend to pivot into agentic roles. • Define new operational roles around agents, for example a "chatbot editor" / agent editor: typically a senior domain expert who curates context, handles edge cases, and continuously improves the agent. • Treat non-technical, communicative people ("tinkerers") as part of the AI workforce, since they often prompt and use LLMs effectively without being software engineers. • Hire AI-native by default in the company itself; candidates who are not AI-native are not hired (start-ups).	• A central CoE is the standard early-stage answer to scarce AI talent and the need for shared standards. • Centralisation is followed in time by decentralisation once domains are mature enough to take over. • Centralise scarce AI talent and knowledge early on, and dissolve the CoE only once capabilities are mature enough to move into the business domains. • The mix of roles around AI is broader than expected: in addition to engineering roles, organisations end up creating editorial / curator roles attached to specific agents. • The bar for individual contribution rises: people who refuse to engage with AI are not suited in a high-performing AI-native environment. • Strict labour protection (especially in the Netherlands) is slowing down workforce reshaping; some employees can stay employed for years even when their work is largely automatable.
Processes	• Treat software engineering as the first process to redesign end-to-end with AI, including how software is specified, written, tested, and deployed. • Apply AI heavily to sales and go-to-market processes: account research, intent analysis, outreach, meeting analysis, follow-ups, sales reporting, often through CRM-native AI features and additional AI tooling. • Use AI for ongoing data quality and CRM hygiene, for example with agents that periodically check meetings, deals, and missing fields and prompt the user. • Invest in claim handling as a primary AI use case in financial services, where existing point solutions can be combined with newer agentic flows. • For mature use cases such as claim handling in insurance, organisations often adopt existing point solutions rather than building everything custom.	• Software engineering is fundamentally disrupted in a matter of months, with a large share of global token usage going to coding agents. • Sales and go-to-market processes are visibly changing through AI tooling, with measurable productivity effects on sales operations. • Legal-tech work is another area where AI tooling is changing knowledge-intensive tasks. • Most organisations are still "staring at the mountain" or "walking to base camp 1"; very few have completed a full AI-driven transformation of a core process, but this is only a matter of time. • Many in-house build attempts in Western Europe fall behind because the market moves faster than 18-month build cycles.

Domain	Key decisions	Reported impact
Technology & Infrastructure	• Buy AI capabilities from existing strategic vendors (Microsoft, Azure, hyperscaler ecosystems) rather than rebuild the stack from scratch. • Add AI orchestration and AI governance tooling on top of the existing stack as deliberate purchases, not as side projects. • Treat sovereignty as a first-class architectural requirement: pick European or local infrastructure for sensitive workloads and rethink reliance on US hyperscalers, especially in public sector. • Decide consciously per legacy system whether to (a) modernise, (b) keep it and let agents work on top of it, or (c) use agents to accelerate migration. • Invest in/build MCP-style interfaces so that agents can reach existing systems and data warehouses (e.g. Azure-based) in a controlled way.	• The technology stack does not need to be replaced for AI scaling, but it does need deliberate extension with orchestration, governance, and agent-related capabilities. • Sovereignty is shifting from a niche concern to a structural design driver, especially in public sector and regulated industries; upcoming European rules may forbid running government workloads on non-European tech. • Existing enterprise systems are often not yet designed to support agent authentication, authorisation, and controlled action execution, creating a significant implementation gap. • Legacy systems are no longer automatically a blocker: agents can interact with old systems through APIs, which can make full modernisation optional rather than mandatory. • Building MCPs and similar integration layers is becoming a new technical capability that organisations need to develop.
Data	• Make data "agent-ready": design for agents as primary consumers of data, in addition to humans and dashboards. • Define explicit access policies for agents, including which data agents may read, which actions they may take, and where human-in-the-loop is required. • Do not let agents simply inherit human user permissions one-to-one; define separate authentication, authorisation, and action boundaries for agents. • Use agents themselves to improve data quality and management, for example by running CRM hygiene agents that flag missing fields, unclear deal status, or stale meetings. • Treat data sovereignty as more than GDPR. • Use existing data warehouses (e.g. in Azure) as the substrate, but expose them to agents via MCPs and similar interfaces.	• Data governance, data access, and data readiness are the hardest challenge for AI scaling, alongside identity and access management. • Authentication and authorisation for agents are effectively being redesigned across the industry, with no mature standard solutions yet. • Some organisations slow down their agent programmes to "fix the data first", which is sometimes legitimate and sometimes used as an excuse to delay. • Counter-intuitively, agents are more tolerant of poor data than humans, since they can read large volumes patiently, which reduces (but does not remove) the data-quality bottleneck.
Culture	• Build the operating model around AI advocates and change-maker rather than around sceptics. • Make "AI-native by default" a hiring filter: candidates who are not actively using AI in their work are not hired into AI-native firms. • Treat refusal to engage with AI as a long-term career risk for individual employees, not as a neutral preference. • Empower individual contributors with their own "swarm" of agents to manage, build, and orchestrate, instead of building more layers of middle management. • Invest in change management explicitly, since AI adoption otherwise stalls even when the technology works. • Use AI as a sparring partner for senior staff (consultants, leaders) to test angles, fill gaps, and accelerate problem-solving, not just as an output generator.	• Cultural openness to AI is the main predictor of which individuals and teams will succeed in AI-scaling organisations. • Middle management layers whose value is mainly information asymmetry and gatekeeping are the first to be at risk in AI-native operating models. • A "Cassandra" dynamic is visible: people who see the disruption coming feel surrounded by colleagues acting as if nothing fundamental is changing.

B.4 Interview 4 summary

Table B.4: Structured summary of interview 4

Interview details	Interview 4; 15 May 2026; Senior director in AI, data, analytics, and transformation; commercial & industrial sector; 13+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Start AI strategy with an AI maturity assessment across dimensions such as data, technology, process, and governance in order to understand the starting point before scaling decisions are made. • Prioritise AI initiatives based on where the most value can be created for the business, using an outside-in perspective on industry trends, challenges, and opportunities. • Translate this into a roadmap that places initiatives along the key operating model dimensions. • Put business ownership on AI use cases, transformation, and value realisation, while IT, data, and AI teams act as enablers and support functions. • Introduce a structured intake and decision process for AI ideas, including a standardised business-case template that weighs expected impact against required effort. • Make an explicit design choice between more central, hybrid/federated, or more autonomous business-led setups, depending on company size, culture, and existing operating model characteristics. • Set central AI policy and guardrails in line with regulation, including the AI Act, and define clearly what is and is not allowed. • Standardise development steps, documentation, and compliance-related ways of working centrally where useful. • Establish an AI Centre of Excellence or similar central body to support knowledge sharing, reuse, and the avoidance of duplicate work across business units. • Align investment logic with the chosen governance structure, for example distinguishing between centrally funded use cases with enterprise-wide impact and divisional use cases with more local value.	• Many organisations are still in a stage of scattered pilots, proofs of value, and copilot rollouts, rather than having a fully embedded governance model for AI. • Stronger business ownership is a necessary step because AI value creation and transformation impact do not sit within IT alone. • Central policy and standardisation help create clarity on compliance, reduce duplication of work across business units, and support knowledge sharing. • The question of how much to centralise versus decentralise AI is context-dependent, with no single model fitting all organisations. • Governance choices affect not only control and compliance, but also team structures, responsibilities, and how organisations prioritise AI investments.
People & Organisation	• Introduce new roles such as AI engineers, AI architects, and AI product owners. • Reskill or upskill existing roles where possible, for example data scientists and software developers who can move into more AI-oriented work. • Upskill the broader organisation so that employees understand how AI changes their role and how to work with AI-supported outputs. • Use workforce-planning logic and skills-gap thinking to assess which capabilities are missing and which existing roles can evolve. • Redesign roles so that employees increasingly validate, supervise, and work with agentic workflows rather than manually producing every intermediate step themselves. • Reconsider how responsibilities are distributed between central teams and divisional teams, especially when AI development capability exists in both places.	• AI is already changing team structures, responsibilities, and required skill profiles. • The impact is twofold: some entirely new roles become more important, while many existing roles change because parts of the work are automated or AI-generated. • Employees increasingly move from making things fully themselves toward checking, validating, and steering AI-supported outputs. • Broader upskilling is necessary across the organisation, not only within specialist AI teams, because AI changes the way many standard functions are performed. • The people impact is not neutral: administrative burden may decrease for some employees, while roles dominated by routine administrative work may come under pressure.

Domain	Key decisions	Reported impact
Processes	• Start process redesign with areas that contain repetitive, manual, and high-volume work, since these are the clearest candidates for automation. • Redesign workflows around agentic logic, where AI can handle more of the operational flow and humans intervene mainly at critical decision or validation points. • Apply AI not only in support processes but also in primary processes where business context and use-case suitability allow it. • Use AI to automate data capture and administrative handling inside processes, for example through speech-to-structured-data flows or automated processing of customer emails. • Use process-level KPI logic to evaluate impact, for example through productivity, throughput, cycle-time reduction, or improved business outcomes in a given function. • Prioritise process transformation based on where AI creates the most concrete business value rather than on technology push alone.	• Most organisations have not yet completed broad process transformation; many are still experimenting with copilots and isolated use cases. • Processes are expected to change substantially, especially where much manual coordination, repetitive work, and administrative handling are still present. • AI can remove administrative burden from employees and allow them to focus more on the core of their role. • At the same time, employees whose jobs are concentrated around repetitive administrative tasks may face stronger disruption. • Human work in processes shifts toward review, validation, and exception handling rather than performing every process step manually. • Support processes are easier early targets across sectors, while primary processes can also change but often require more contextual tailoring.
Technology & Infrastructure	• Start from the current state of the technology landscape and assess whether the organisation already has a data platform or other means of connecting data sources. • If AI is expected to support workflows across systems, introduce an additional platform or integration layer above fragmented point-to-point landscapes. • Decide consciously between buying and building AI solutions, depending on the use case, the firm's strategy, the availability of market solutions, and whether the capability is strategically differentiating. • Involve IT actively in architecture, tooling, and security decisions around AI. • Use existing IT portfolio-management mechanisms where possible, if the organisation already has them. • In a more mature target setup, position IT as the function that supports AI initiatives through security, enterprise architecture, and technology decisions, alongside a dedicated data function and business ownership. • Create an overarching forum or governance body in which business, data, and IT come together around AI-related decisions.	• Many organisations still rely on fragmented system landscapes, point-to-point integrations, exports, and Excel-based workarounds. • This fragmentation limits how far AI can support cross-system or end-to-end workflows unless a more deliberate platform or integration layer is introduced. • Legacy systems remain a significant constraint in some organisations, meaning that some firms first need to improve core foundations before they can scale more sophisticated AI use cases. This is partially an issue due to the lack of a single source of truth. • Security rises higher on the agenda as AI introduces new risks and requires IT to engage differently than with earlier digital tooling. • The rapid expansion of market solutions makes buy-versus-build decisions more pressing and more complex. • AI also changes IT work itself, as developers, IT support functions, and other IT roles start using AI in their own work and become more efficient.
Data	• Treat data readiness, data management, and data governance as core preconditions for AI scaling. • Work toward clearer data definitions and, where possible, a single source of truth. • Recognise that legacy landscapes often prevent a true single source of truth and therefore constrain some AI use cases. • Distinguish between use cases that rely heavily on structured operational data (e.g. ERP-driven cases) and use cases that can already work with less structured knowledge bases such as SharePoint or document repositories. • Focus on monitoring data quality. • Assign data owners, appoint data stewards, and monitor data quality more actively. • Use AI not only as a consumer of data but also as a way to generate more structured data from operational activity, for example by turning speech, emails, or field observations into structured system inputs. • In some contexts, consider using AI to help accelerate migration or reconciliation across legacy systems.	• AI as a catalyst that makes the urgency and value of the data foundation much more visible inside organisations. • Organisations are pushed to invest more seriously in data quality, ownership, stewardship, and monitoring once AI use cases become concrete. • Data quality remains a major bottleneck: "garbage in, garbage out" applies strongly to AI, especially for predictive or heavily data-dependent use cases. • The severity of the bottleneck differs by use case: some knowledge-based applications can move ahead relatively quickly, while use cases dependent on structured enterprise data remain much more constrained. • AI can increase the amount of structured data captured in operations without necessarily requiring employees to perform more manual data entry themselves. • As organisations scale AI, data becomes less of a background issue and more of an explicit management priority.

Domain	Key decisions	Reported impact
Culture	• Invest in AI literacy and practical understanding across the organisation so that employees better understand both the opportunities and the risks. • Treat adoption as an explicit part of AI scaling rather than assuming that deployment alone will lead to usage. • Work with early adopters and ambassadors to create momentum and help other employees engage with AI. • Address concerns about risk, data exposure, and job impact directly rather than leaving them implicit. • Recognise that many organisations have rolled out copilots or early use cases without yet putting the broader adoption effort around them in place.	• Reactions to AI are split between enthusiastic experimenters and more sceptical employees. • Resistance is linked both to fear about job consequences and to concern about risks, data handling, and the unknown. • A large part of resistance is also linked to lack of knowledge: when employees do not understand what AI is or what it does, it is easier for uncertainty to turn into reluctance. • Ambassadors and naturally curious adopters are important for creating traction and helping others start using AI in practice. • Adoption remains a recurring challenge: many organisations are experimenting with AI, but broader embedding in day-to-day work often lags behind the technical rollout.

B.5 Interview 5 summary

Table B.5: Structured summary of interview 5

Interview details	Interview 5; 19 May 2026; Former regional head of AI and data in insurance and telecom; 20+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Position AI as an enablement layer within the broader business strategy rather than as a standalone transformation stream. • Use business-driven governance: steer AI priorities through strategic business workstreams and value streams rather than letting AI determine priorities by itself. • Adopt a federated governance model that balances regional coordination with local business autonomy across countries. • Build a portfolio funnel for AI use cases: collect local backlogs, identify overlaps across territories, assess value potential, and select regional priorities based on critical mass and strategic fit. • Introduce a multi-layer governance structure with regional value-stream leadership and mirrored champions in each country. • Standardise the intake, onboarding, and handling of AI use cases before trying to scale them broadly. • Use staged decision gates to move ideas from concept to proposal, feasibility assessment, business case, and potentially POC/MVP before further scaling. • Evaluate AI initiatives through a mix of financial, strategic, and operational KPIs, including value creation, cost efficiency, NPS-related measures, and straight-through-processing metrics.	• Governance in a multi-country setting is complex because local markets differ in regulation, business models, priorities, and distribution channels. • Business-driven portfolio management helps prevent fragmentation by forcing choices across large numbers of local use cases. • The federated setup makes it possible to combine local relevance with regional scale, but it also requires ongoing negotiation and alignment. • AI governance becomes more practical and scalable once the organisation has aligned a common intake and decision process for use cases. • Not all use cases have equal value across markets, so scaling depends on finding enough commonality to justify regional prioritisation. • In practice, some markets may need to support initiatives that are not local priorities in order to benefit later from regional commonality and shared scaling.

Domain	Key decisions	Reported impact
People & Organisation	• Use a federated hub-and-spoke model with a regional CoE and local country capabilities rather than fully centralising AI execution. • Standardise core AI-related roles across countries, including profiles such as data scientists, AI engineers, and MLOps roles. • Set a minimum baseline of AI capability per country and require countries to staff up to that level through hiring and/or reskilling. • Invest in broad AI literacy and capability building as part of a wider data-driven mindset programme. • Offer layered capability development, including fundamentals for broad audiences, role-specific tracks for specialist profiles, pilot-based learning environments, and management training. • Expand and strengthen the regional CoE as scaling matures and coordination needs increase. • Plan workforce transition only once AI use cases become mature and production-grade rather than starting from headcount reduction as a primary objective.	• AI scaling touches all parts of the people dimension, from literacy and mindset to specialist role design and staffing. • Countries that lacked core AI capabilities had to hire or reskill people into AI-related roles. • The regional CoE grew significantly as AI scaling progressed, reflecting increasing coordination and capability needs. • Workforce reduction is a late-stage consequence of mature automation, not as an upfront goal of the AI strategy. • Scaling AI creates a need for transition planning, because some roles may disappear while new roles and capabilities become necessary. • Cultural and capability-building interventions are necessary to help business users identify opportunities and work effectively with AI.
Processes	• Prioritise AI in business domains with the highest strategic importance, especially retention, acquisition, claims management, underwriting, and selected back-office processes. • Redesign core processes through domain-based transformation rather than simply layering AI onto existing workflows. • Use a greenfield redesign approach: strip out an end-to-end process, decompose it into smaller modular components, redesign these components, and rebuild the process around agents or combinations of agents. • Standardise redesigned processes toward more straight-through-processing logic where possible. • Combine top-down "big bets" based on strategic priorities and best practices with bottom-up use-case generation from trained business staff. • Use literacy programmes such as business translator initiatives to help business people identify AI opportunities in their own day-to-day activities. • Move use cases through a structured pathway from idea to proposal, feasibility assessment, business case, and prototype before embedding them into business operations.	• Claims management and underwriting are examples of the most important processes for AI-enabled redesign. • AI scaling in this context goes beyond local process automation and moves toward domain-based redesign of end-to-end processes. • Bottom-up opportunity identification becomes more effective when business users are trained to recognise where AI can create value. • Process scaling depends on converting many local ideas into a smaller set of governed and validated use cases. • More mature process transformation makes it possible to move from fragmented experimentation toward embedded operational redesign. • In an insurance context, AI is especially relevant in areas where process volume, modularity, and standardisation potential are high.
Technology & Infrastructure	• Re-baseline the technology stack when earlier data or AI choices are no longer fit for scaling. • Prune architectural complexity and avoid extending the stack in ways that make the technology landscape harder to manage. • Build AI on top of a strong data platform foundation rather than trying to create a fully custom stack from scratch. • Prefer best-of-suite choices for core data-platform capabilities and then build selected enabling layers on top of that stack. • Make deliberate build-versus-buy decisions, since internal technology teams tend to overbuild or overengineer solutions. • Separate the question of the data platform from the question of the agentic or AI application layer that sits on top of it. • Modernise legacy incrementally rather than trying to rebuild everything at once, and offset each new addition by stripping out complexity elsewhere where possible. • Use a mix of on-premise and cloud environments in practice, while regional AI teams operate mainly on cloud infrastructure.	• Earlier technology decisions can materially slow down AI scaling, making re-baselining and architectural pruning necessary. • Build-versus-buy decisions are difficult because IT teams often prefer building complex things themselves even when that is not the best path. • Technology complexity can easily push organisations in the wrong direction if stack choices are not disciplined. • AI scaling does not require full modernisation up front, but legacy and infrastructure choices still shape how fast organisations can progress. • The practical reality is hybrid: much of the country-level environment remains on-premise, while regional AI teams work largely on cloud platforms such as Azure. • A clearer separation between foundational data-platform choices and higher-level AI/agentic capabilities supports more manageable scaling.

Domain	Key decisions	Reported impact
Data	• Treat enterprise and data architecture as core AI-scaling decisions rather than secondary technical concerns. • Move away from over-centralised and overly rigid data-model approaches when they no longer fit the cloud and AI context. • Shift toward domain-based data management with clearer ownership of major data domains. • Appoint stronger leadership for data, including a group-level CDO and responsible owners for specific data domains. • Use enabling tooling for data governance and data management where useful, but avoid turning data management into an overly heavy centralised framework. • Improve data quality and data management as foundational conditions for AI scaling. • Ensure that data from different enterprise systems can be connected, read, and retrieved for AI use cases. • Recognise that local data-management practices and global frameworks need to be balanced rather than solved through one universal model.	• Data management is foundational, but difficult to scale because large enterprise-wide frameworks often become too heavy in practice. • Earlier centralised data-model approaches are a sign of immaturity rather than future-proof design. • Data stewardship concepts existed, but did not scale well across business units, which created the need for a renewed domain-based approach. • The organisation had not yet found the sweet spot between full-scale centralised data management and local practicality. • Poor data quality remains a direct constraint on AI performance: "garbage in, garbage out" still applies. • Organisations that make data across systems connected, readable, and retrievable for AI are better positioned to win with AI scaling.
Culture	• Build a data-driven mindset as an explicit cultural component of the AI strategy. • Use literacy and training programmes to help employees understand AI, identify opportunities, and work with AI in practice. • Run broad change-management programmes in parallel with AI scaling because multiple organisational changes overlap. • Position AI as a driver for change and as a way to create more room for creativity and better work, rather than communicating it only through efficiency or headcount narratives. • Encourage employees and leaders to embrace technology as part of the future operating model. • Use bottom-up opportunity generation, but only after employees have enough understanding to recognise meaningful AI use cases.	• AI opportunity generation improves when business users are trained well enough to identify concrete AI opportunities in their own work. • Cultural readiness is more than awareness alone; it also includes capability building, mindset change, and management participation. • Employees do worry about job impact, especially because AI scaling overlaps with broader organisational changes. • Change management therefore becomes necessary not as a side activity, but as a parallel workstream to AI scaling itself. • A stronger AI and data mindset supports the shift from isolated experimentation toward more embedded business adoption.

B.6 Interview 6 summary

Table B.6: Structured summary of interview 6

Interview details	Interview 6; 22 May 2026; Consulting partner in cloud and digital transformation, with deep financial-services experience; 25+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Centralise governance and ownership of AI development in the early stage, until a sufficiently mature AI platform is in place. • Federalise or decentralise AI responsibility only once business units and teams have built enough maturity in using AI and developing agents. • Prioritise AI initiatives using a desirability, viability, feasibility (DVF) logic: user value, business value, and practical doability all need to be strong enough before scaling. • Combine centralised AI enablement with clearer accountability in the business for process and data outcomes.	• Centralisation is useful in the early phase because it helps create coherence, control, and a usable platform foundation. • Decentralisation only becomes workable once local teams are mature enough to use AI responsibly and productively. • AI scaling decisions remain closely tied to practical business value rather than to technology enthusiasm alone. • Premature decentralisation risks fragmentation, weak reuse, and uneven capability development.

Domain	Key decisions	Reported impact
People & Organisation	• Invest in reskilling existing technical staff, especially people with engineering or programming backgrounds, toward AI-related work. • Encourage formal certification and structured upskilling for technical staff moving into AI architecture and engineering roles. • Combine reskilling with selective hiring of younger talent entering the market with more recent AI-related training. • Recognise that scarce AI capabilities cannot be sourced only externally and therefore need to be developed internally as well.	• Scarcity of capable AI talent is a real bottleneck in practice. • People with existing engineering skills are seen as easier to transition into AI-related roles than fully non-technical profiles. • Organisations are using a mix of internal reskilling and external hiring to build AI capability. • AI increases the productivity of strong technical people, especially in engineering-related work. • In the near term, the impact is described more as doing more with the same number of people than as immediate large-scale workforce reduction.
Processes	• Focus AI first on painful, repetitive, and high-effort process areas where the impact is tangible. • Use AI not only for support tasks, but progressively for operational process redesign once the data and infrastructure foundation is strong enough. • Redesign processes AI-first rather than simply layering AI onto existing workflows. • Understand end-to-end processes first, and then reshape them around AI capabilities and automation opportunities. • Apply AI both to core process examples (e.g. intake and process handling in legal insurance, migration/integration work) and to technical delivery work such as software engineering and system conversion.	• A concrete example discussed in the interview is the automation of an intake process in legal insurance, reducing handling time from roughly 60 minutes to 6 minutes. • The greatest early impact is seen in technical and engineering-heavy work, including software production and system integration. • Support processes and lighter AI-enabled work are easier starting points, while deeper operational core-process automation requires a stronger data and platform foundation. • AI is most valuable when organisations move beyond simple add-ons and actually redesign work around AI. • The practical impact is often not direct headcount reduction, but growth or output improvement without proportional FTE growth.
Technology & Infrastructure	• Put in place three enabling layers for scaling AI: a data platform, an API platform, and an AI platform. • Treat the AI platform as the environment in which agents are built, managed, monitored, and governed. • Build on top of existing hyperscaler ecosystems where possible, rather than assuming that scaling requires entirely new specialist vendors. • Make deliberate build-versus-buy choices, while recognizing that large vendors do not always provide everything needed out of the box. • Use AI and platform capabilities to support difficult transformation work such as data migration and system conversion. • Put the enabling platform stack in place before expecting large-scale automation of operational core processes.	• A platform stack is a practical precondition for serious AI scaling. • Existing hyperscaler ecosystems are useful foundations, although they are not always sufficient by themselves. • AI platforms become important not just for model access, but also for monitoring, control, and development of agent-based capabilities. • Stronger infrastructure materially expands the range of processes that can be automated or redesigned with AI. • Without adequate platform and infrastructure foundations, AI remains more limited to lighter use cases and support-oriented applications.
Data	• Treat data infrastructure as a prerequisite for scaling AI in operational processes. • Strengthen data stewardship and data governance as AI use expands. • Move data ownership and stewardship closer to the business, especially to process owners and system owners, rather than keeping it purely centralised. • Use dedicated data leadership where appropriate, for example through CDO-type roles or heads of data/governance. • Recognise proprietary data as a strategic asset and maintain it more actively as a competitive differentiator. • Improve the quality, control, and usability of enterprise data before expecting major AI-driven automation benefits.	• Data stewardship and data governance are becoming more important because proprietary data is one of the few durable sources of differentiation. • Organisations that already invested in a strong data platform have a clear advantage in AI scaling. • Data ownership that sits too centrally can become ineffective; moving it toward the business is more practical. • Major automation gains in operational core processes require previous investment in data infrastructure. • Without that foundation, organisations remain more limited to copilots, support use cases, and lighter forms of AI adoption.

Domain	Key decisions	Reported impact
Culture	• Use leadership behaviour and visible example-setting to stimulate adoption, rather than relying only on formal AI programmes. • Let senior leaders visibly use AI themselves if the goal is to create wider uptake in the organisation. • Encourage a culture in which employees actively experiment, learn, and reposition themselves around AI-enabled ways of working. • Treat AI adoption not only as a technology rollout, but also as a question of whether people embrace more effective ways of working.	• Visible leadership use of AI is a stronger adoption driver than formal top-down messaging. • Employees who embrace AI are becoming more effective, which creates a pull effect on others. • Adoption is therefore shaped both by top-down signalling and by visible local examples of improved work. • AI contributes to a culture in which employees are pushed to adapt, learn, and become more productive with new tools.

B.7 Interview 7 summary

Table B.7: Structured summary of interview 7

Interview details	Interview 7; 29 May 2026; Senior AI and cloud consulting manager with AI platform and governance focus; 7+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Split AI-related governance responsibilities across roles such as CIO, CISO, and CDO, depending on whether the issue concerns infrastructure, security, data governance, or compliance. • Give AI governance an explicit organisational place, especially in larger organisations, because regulation such as the EU AI Act requires clear ownership and oversight. • Introduce governance mechanisms that go beyond traditional IT project governance, because AI creates different infrastructure, compliance, and oversight needs. • Measure AI initiatives not only through technical delivery, but also through business impact analysis, adoption metrics, and operational cost awareness. • Include FinOps-type monitoring as part of AI governance, because usage patterns, costs, and business effects need to be linked more explicitly than in many current AI initiatives.	• AI governance is a combined responsibility rather than something that sits with a single function. • Existing IT governance is insufficient on its own for AI, because AI exposes different risks and requires different organisational controls. • AI initiatives often reach the stage of proof of concept or even production without strong enough measurement of effectiveness and business value. • FinOps and usage-cost monitoring are underdeveloped in many AI initiatives, despite their growing relevance. • AI can expose security weaknesses that were already present in organisations but had not yet been surfaced as clearly through previous IT usage patterns.
People & Organisation	• Raise AI literacy across the business so that employees can use AI responsibly and contribute to broader AI adoption. • Build shared ownership for AI rather than treating it only as a specialist or central team concern. • Use Centres of Excellence, champions, and knowledge-sharing structures to spread capability across the organisation. • Organise AI development and innovation through an inner-source model in which ideas, components, and solutions from one part of the organisation can be shared and reused by others. • Use platform-based enablement so that technical teams can start from boilerplates and reusable components instead of rebuilding everything from scratch. • Adapt skill-building to maturity level: at lower maturity, prompt literacy may be valuable; at higher maturity, more specialised roles become more important. • Continue investing in cybersecurity roles, as it remains one of the most critical and increasingly important areas in AI-scaling organisations.	• AI literacy is a necessary condition for scaling AI more broadly and responsibly across the business. • Champions and inner-source-style knowledge sharing are useful ways to increase reuse, learning, and participation from different business units. • Reusable building blocks and shared platform components reduce duplicate work and make it easier for teams to build on each other's progress. • Existing technical profiles such as data scientists can in some cases transition relatively easily into more AI-engineering-oriented work.

Domain	Key decisions	Reported impact
Processes	• Add explicit lifecycle management for AI use cases after deployment, including continued monitoring, improvement, maintenance, and adaptation when surrounding systems or dependencies change. • Treat AI use cases as living applications rather than deliveries that stop once they reach production. • Include a business process redesign component where useful, especially where organisations want to scale AI beyond local experiments. • Use onboarding, offboarding, and structured use-case handling processes around the AI platform instead of leaving adoption completely informal. • Support reuse by making components, capabilities, and learnings available through a marketplace or shared platform model.	• AI scaling requires more than initial deployment: use cases need continued management once they are live. • Process impact is especially visible in administrative, text-based, and visually structured work, where AI can make tasks significantly faster. • Front-office work, customer interaction, note-taking, reporting, and other high-volume information-handling activities are areas with major short-term potential. • AI-enabled development work is also changing, as coding support from tools such as GitHub Copilot increasingly becomes part of how technical work is done. • Broad process transformation requires not just AI tools, but also supporting processes for onboarding, maintenance, reuse, and iterative improvement.
Technology & Infrastructure	• Build a shared AI platform that supports onboarding of users, boilerplate starting points for teams, integrations, and reuse of building blocks across use cases. • Use a marketplace or shared component layer so that new technical components created for one use case can be reused elsewhere in the organisation. • Leverage the movement of existing enterprise vendors toward AI support, for example through API and MCP-style enablement, rather than assuming that all legacy systems must first be replaced. • Be cautious with highly experimental agentic tooling that acts directly through a user identity, because this creates major security and control risks. • Distinguish between AI use cases that can rely on existing models and those more exceptional cases that still require traditional model training and heavier ML infrastructure. • Use cloud infrastructure for most AI-related compute where possible, while recognizing that some organisations still operate important workloads on-premise. • Treat orchestration and agent management as an identity and access management problem: every agent should have clearly assigned roles and permissions.	• A shared AI platform is a practical way to accelerate delivery, onboarding, and reuse across multiple teams. • Existing vendors and legacy environments are increasingly becoming AI-compatible, which can reduce the need for full replacement before scaling begins. • Experimental agentic setups that act with end-user permissions are creating serious security concerns and are too risky for most organisations. • Most AI workloads are still running in the cloud, although geopolitical and sovereignty concerns may shift some choices over time. • Technical scaling depends less on training custom models for every use case and more on using existing models effectively within a governed infrastructure.

Domain	Key decisions	Reported impact
Data	• Build a data foundation that allows flexible but controlled isolation of access, so that each agent can only reach the specific data needed for its use case. • Apply both role-based and attribute-based access control (RBAC and ABAC) for AI agents and their data interactions. • Strengthen data quality checks within processes rather than treating data quality as an afterthought. • Maintain a shared data platform or, where full centralisation is not possible, assign controlled access across multiple data sources with clear permissions. • Work toward a single source of truth through master data management and golden-record logic. • Use layered data architecture where appropriate, for example through medallion-style or similar structures, while distinguishing between shared organisational data and more operational or use-case-specific data. • Consider stronger data leadership structures in larger organisations, for example through a separate CDO in addition to CIO and CISO roles.	• Controlled access to data is essential for safe and scalable use of agents. • Single source of truth ambitions remain difficult in practice because organisations often still deal with fragmented and older data landscapes. • Master data management is a crucial but difficult foundation for reliable AI use. • Data quality checks remain necessary even when more can be automated, because AI scaling still depends on knowing how data is created, stored, and read. • The data architecture work around AI is partly use-case agnostic: some foundational structures need to be improved regardless of the exact use cases that will later run on top of them.
Culture	• Combine top-down direction with bottom-up participation when trying to change organisational culture around AI. • Actively bring less AI-literate employees along in the transition, rather than assuming they will adapt automatically. • Frame AI adoption as an evolutionary organisational change rather than as a one-time top-down revolution. • Build an innovation culture in which employees are encouraged to explore how AI can improve work, but support this with guidance, literacy, and practical structures.	• A purely top-down push is insufficient for successful AI transformation. • Employees with lower AI literacy may feel alienated if they are not actively included in the change process. • Successful AI transformation depend on whether organisations can combine strategic direction with participation and learning. • AI adoption is an evolutionary journey that depends on starting maturity and the distance between the current and desired state.

B.8 Interview 8 summary

Table B.8: Structured summary of interview 8

Interview details	Interview 8; 1 June 2026; Senior consultant in AI and data strategy, adoption, and implementation; 7+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Establish a central AI governance structure, for example through an AI Centre of Excellence, because AI can no longer be treated as an ad-hoc side responsibility. • Organise AI oversight through a joint governance group around the CIO, CDO, and CISO, rather than assigning full responsibility to one individual function. • Keep final accountability at senior level, but assign dedicated operational roles for day-to-day AI governance and execution because C-level owners cannot manage this in practice themselves. • Measure AI initiatives through a mix of ROI logic, time savings, A/B tests, usage/adoption metrics, quality improvement, and output expansion. • Treat AI governance as distinct from earlier software or office-tool rollouts, because AI is expected to be used broadly across the organisation and carries different risk and compliance implications.	• Many organisations are still immature in both data and AI governance, with no proper central data or AI function in place. • AI introduces broader organisational risk because it is not limited to one department; many employees may use it without understanding the risks. • Larger organisations are cautious with new AI tools and vendors because reliability histories are short and business dependency risk is difficult to assess. • Measuring value remains difficult, but organisations do try to justify AI through time savings, adoption levels, quality gains, and the ability to take on more or more complex work. • Governance maturity is a prerequisite for moving beyond early experimentation.
People & Organisation	• Build an AI Centre of Excellence or similar central capability hub to concentrate scarce knowledge and support the wider organisation. • Train the broader workforce in AI, because AI is expected to become relevant to almost everyone in the organisation rather than to a narrow specialist group only. • Develop or hire scarce AI capabilities, including moving existing data analysts and data scientists further toward ML/AI-oriented work where possible. • Introduce or strengthen intermediary roles that can translate business knowledge into AI opportunities and tools, similar to a product-owner or translator profile. • Work with early adopters and motivated employees first, so that they can help pull the rest of the organisation along. • Accept that capability building is expensive and often difficult to justify through an immediate ROI business case.	• Organisations are facing a two-sided challenge: educating the whole workforce while also acquiring scarce specialist capabilities. • Reskilling existing technical staff is possible, but difficult in practice because those employees are often already fully occupied. • Hiring new AI-capable people is often easier than reskilling, but also expensive and hard to justify financially. • A missing translator role between business needs and AI tooling is a recurring gap in many organisations. • Early adopters are useful for creating momentum, while broader workforce change remains much harder.
Processes	• Apply AI selectively where it adds real value, rather than using it for every task simply because the technology is available. • Evaluate AI use cases partly through A/B testing and observed differences in how work is done with and without AI support. • Use AI especially for more complex or cognitively demanding work where support creates visible productivity gains, rather than only for trivial tasks. • Allow AI to support process redesign across departments where better data and process alignment are needed. • Create enough freedom in the operating model for employees and teams to experiment with AI in their work. • Use templates, work instructions, and process descriptions to make workflows clearer and easier to translate into AI-supported tasks or opportunities.	• Process impact is strongest where AI helps people deal with more complex tasks, not only where it automates the simplest work. • In one procurement example, AI support made it easier to handle more complex work and freed time to target a different market segment. • AI can act as a catalyst for improving collaboration across departments when processes and data need to be aligned. • Many organisations still do not have their own business processes clearly documented, which makes it harder to identify where AI can be implemented effectively. • Organisations that do not invest in AI early risk leaving their workforce behind relative to future ways of working. • Process value is highly dependent on whether AI is linked to genuine business value rather than to novelty.

Domain	Key decisions	Reported impact
Technology & Infrastructure	• Be cautious in adopting very new AI tools and vendors when long-term reliability and continuity are still uncertain. • Avoid replacing recently implemented tools too quickly when organisations have already made major investments, even if newer AI-enabled options emerge. • Treat AI tooling choices as enterprise decisions with implications for many users rather than as isolated team-level software choices. • Use pilots and smaller implementations to learn where AI tools are reliable enough and where organisational caution remains justified.	• Established organisations are hesitant to depend on young AI vendors because the continuity risk is difficult to assess. • Earlier technology investments can slow down the adoption of newer AI tools, because organisations are reluctant to abandon tools they only recently introduced. • AI tool adoption is more sensitive than other software procurement because the tools affect more users and involve higher uncertainty around reliability and risk.
Data	• Treat data governance as a formal responsibility with clear ownership, policy, and execution rather than as something that can simply be added to an existing role. • Strengthen top-down data leadership, for example through a CDO-type role or equivalent governance setup. • Monitor data quality and data discipline actively, including through dashboards and steering mechanisms at departmental level. • Introduce controls and steering to improve data quality, such as mandatory metadata, document structure conventions, and departmental reporting on data quality. • Assign responsibility for checking whether data is entered and maintained correctly, and connect non-compliance to real consequences if needed. • Improve data structure, naming conventions, metadata quality, and versioning discipline so that AI tools can actually retrieve and use information reliably.	• Many organisations are still underestimating the need for formal data governance and structured data responsibilities. • Weak data structure, poor metadata, unclear naming, and weak version control are major barriers to useful AI deployment. • Department-level dashboards and steering discussions can improve data discipline, especially when linked to incentives or accountability. • Data is a cross-cutting prerequisite: without better data discipline, AI cannot create sustainable value at scale. • Data governance is described not only as a technical issue, but also as a matter of discipline, responsibility, and behavioural change.
Culture	• Invest in broad AI literacy because employees across the organisation will use AI, often without fully understanding its risks. • Start adoption efforts with employees who are willing and curious, and use them to pull others along. • Treat change management as a major part of AI scaling rather than as an afterthought. • Build an operating model that gives employees enough space to experiment with AI in practice. • Recognise that leadership hiring and leadership mindset matter: leaders who do not care about AI are unlikely to push AI adoption effectively.	• Employee reactions to AI are highly mixed, ranging from enthusiasm to outright refusal to use AI tools. • Resistance is linked to fears about job loss, loss of expertise status, and discomfort with unfamiliar technology. • Organisations that rely only on formal rollout without cultural support are likely to struggle with adoption. • Change is a major challenge in AI scaling, not just a side issue. • A leadership team that does not take AI seriously is a direct barrier to building an AI-oriented culture.

B.9 Interview 9 summary

Table B.9: Structured summary of interview 9

Interview details	Interview 9; 5 June 2026; Senior consultant in responsible AI and AI governance; 7+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Start AI scaling with a central AI team or Centre of Excellence, so that scarce expertise, prioritisation, and responsible-use guidance are concentrated in one place. • Decentralise AI capability only once the organisation has developed enough maturity and local expertise to work in line with central guidelines. • Define explicit governance around the AI lifecycle, including roles, responsibilities, decision moments, and risk assessment steps. • For higher-risk use cases, use a stronger governance body or escalation mechanism, while lower-risk use cases can remain closer to the business. • Escalate higher-risk AI decisions to a more senior or directive level, rather than leaving them fully to local business or project teams. • Give the business owner explicit decision rights over the use case, while involving technology and control functions such as the CIO/CTO, CISO, CDO, and risk roles in line with the nature and risk level of the AI application. • Apply a three-lines-of-defence logic to AI: the business owns and operates AI use cases in the first line, second-line functions such as risk, security, and data-related control roles define policy and guardrails, and the third line provides independent assurance on whether these controls are working effectively. • Treat responsible AI as a practical precondition for scaling, including policy, risk management, security, robustness, and bias-related controls. • Measure success by comparing to benchmark before AI. • Measure AI initiatives through both value and control logic: business cases, efficiency gains, adoption over time, intervention effects such as training, and cost components including token usage all need attention.	• Centralisation is useful in the early phase because it helps organisations learn faster, prioritise better, and apply responsible-AI guidelines more consistently. • Premature decentralisation creates a risk of "shadow AI", where business units buy or build their own AI solutions outside central standards and governance. • This shadow-AI risk is increased by local purchasing autonomy, because some AI tools can be procured below formal procurement thresholds. • Organisations with weak AI governance are more likely to end up with fragmented tooling, inconsistent controls, and use cases that are not aligned with enterprise policy. • Governance should be connected to business value: risk acceptance for certain use cases is ultimately a business choice, not only a control-function decision. • AI governance relies partly on familiar risk and control structures, but with stronger emphasis on explicit assessments, lifecycle controls, and business accountability for value and risk acceptance.
People & Organisation	• Organise AI capability development as a broad organisational effort rather than as a specialist-only topic, because GenAI will affect a much wider group of employees than traditional AI did. • Invest continuously in awareness, training, and change management, rather than treating AI literacy as a one-off intervention. • Use a central AI team or CoE in the beginning, but evolve toward more distributed capability as maturity grows. • Develop or hire new roles such as AI engineers, ML engineers, and AI product-oriented roles where needed. • Strengthen existing data-related roles, especially data owners and data stewards, because AI scaling increases the importance of data quality and controlled use. • Treat upskilling as important not only for AI use itself, but also to ensure that employees continue to understand their underlying business domain and do not over-rely on AI outputs.	• Organisations are facing a broad people challenge: not only specialist capability building, but also organisation-wide acceptance of AI as part of everyday work. • Resistance is common, especially where employees fear job loss, loss of expertise status, or simply do not know how to work with AI. • Early disappointment is also a barrier: employees may try AI once, expect too much, and stop using it when the first result is underwhelming. • AI scaling therefore requires iteration, learning, and continued support rather than one-time rollout. • If AI does too much of the work for junior employees, organisations may weaken how people build deep domain knowledge and critical thinking over time.

Domain	Key decisions	Reported impact
Processes	• Use proof-of-concept stages first to test whether AI performs at least roughly on the level needed compared with human work, before moving into a pilot. • Move from proof of concept to pilot and then to broader rollout only after required checks and controls have been completed. • In pilot phases, compare AI-supported work against human-only work through shadow-running or A/B-style setups to test effectiveness, quality, and value before scaling further. • Apply AI first in processes that are well-documented, text-heavy, repetitive, or rely on standardised decision rules. • Use AI to extract information from documents, prepare cases for human handlers, and support standardised checks in operational processes. • Where scaling ambition is higher, redesign processes from an end-to-end perspective instead of only automating individual steps. • In more radical redesign scenarios, reconsider not only tasks but also roles, handovers, and the overall process structure from the desired end goal backwards.	• Current AI impact in processes is real but still relatively limited in scale. • AI is already being used to extract relevant information from uploaded documents, support case handling, and in some cases automate straightforward payments or claims decisions when all standard checks are satisfied. • Customer-service-related AI is promising but also risky, because there are visible examples where such implementations go wrong. • Agentic AI is not yet broadly deployed; where it is used, it is mainly in processes where data is available, steps are documented, and work is repetitive or text-heavy. • Pilot-based rollout allows organisations to monitor use cases more closely, evaluate outputs with users, and build a stronger business case before scaling to production. • Scaling decisions after pilot are often business-led: the business sharpens the case based on pilot results and then seeks management approval for broader rollout. • The strongest process redesign is end-to-end reconsideration of how work should be done, rather than simply inserting AI into an unchanged workflow.
Technology & Infrastructure	• Build a shared AI platform instead of deploying AI separately in many isolated environments. • Use the AI platform to pre-arrange core conditions such as access to LLMs, data connections, security layers, and access management. • Give each AI application its own controlled sub-environment within the broader platform rather than allowing all applications to operate in an uncontrolled shared space. • Prefer stable large-vendor ecosystems where possible, especially in larger organisations, and avoid introducing too many new niche tools too quickly. • Use niche orchestration tools mainly for experimentation where needed, but do not assume they are necessary for scaling. • Use cloud as the default for most organisations because of scalability, cost, and maintainability advantages, while reserving on-premise for highly sensitive environments or strict public-sector contexts. • Treat legacy constraints mainly as an access and data-quality issue: the main question is whether systems have connectors/APIs and whether the underlying data is usable.	• A shared AI platform makes scaling easier because core controls and integrations only need to be set up once instead of per use case. • Stable vendor ecosystems are attractive to larger organisations because continuity risk is lower than with many small, fast-moving AI vendors. • Niche agentic or orchestration tools are currently seen mostly as experimental additions rather than essential scaling components. • Legacy systems are not a blocker because of technical age alone; they become a problem when they cannot be connected to or when their data is not in usable shape. • Cloud is superior for most scaling scenarios, with on-premise mainly relevant where data sensitivity or public-sector requirements are unusually strict.

Domain	Key decisions	Reported impact
Data	• Treat data as a strategic asset and not as a "nice to have". • Strengthen traditional data governance rather than bypassing it: AI scaling increases the need for standard checks, data suitability assessments, and clear ownership. • Assign data owners who determine whether and for what purpose data may be used, including in AI applications. • Use data stewards to monitor and improve operational data quality and to ensure that issues are addressed structurally. • Feed data-quality corrections back to the source, rather than adjusting downstream reports or outputs locally. • Work toward a single source of truth and a shared understanding of what data points actually mean across the organisation. • Choose between more centralised and more distributed data environments based on context, since both have advantages and trade-offs around access, ownership, and security.	• AI scaling puts stronger pressure on data governance because poor data quality, conflicting reports, and inconsistent data definitions become even more problematic once AI depends on them. • Data governance roles such as data owners and data stewards are still highly relevant rather than being replaced by AI. • There are recurring problems in organisations where reports contradict each other because underlying data ownership and structure are weak. • Single source of truth ambitions are important, but not always easy to realise in fragmented landscapes. • The core data lesson is that AI magnifies existing weaknesses in data quality, shared definitions, and control over how data is used.
Culture	• Treat AI adoption as a broader change-management process, not just as deployment of a tool. • Use continuous training, webinars, hackathons, and repeated engagement activities to help people learn how to work with AI in practice. • Make space for experimentation, but combine this with guidance so that employees do not become unknowingly incompetent in how they use AI. • Build an innovation-oriented culture in which employees are expected to keep adapting and learning as AI changes the way work is done. • Use leadership communication and signalling to normalise AI as part of everyday work.	• Organisations differ strongly in how deeply they support AI adoption; many do provide webinars and training activities, but the depth and quality vary. • Cultural resistance comes both from fear of job loss and from uncertainty about how to use AI well. • A key risk is that people use AI without understanding its limitations, which creates a form of unconscious incompetence. • Poorly reviewed AI output can create extra work rather than saving time, especially when employees use AI without understanding when its output is reliable and when it needs deeper review. • Organisations with weaker innovation culture and less urgency are slower to translate AI interest into actual adoption. • AI transformation depends on whether people are helped to iterate, learn, and adjust their work habits over time.

B.10 Interview 10 summary

Table B.10: Structured summary of interview 10

Interview details	Interview 10; 12 June 2026; Consulting director and former Big Tech engineering lead; 20+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Treat every AI initiative as a product rather than as a one-off project or experiment. • Assign both a business owner and an IT owner to each AI product, similar to the governance of data products. • Maintain a record or catalogue of AI products, including their ownership, boundaries, and intended use. • Govern AI largely through extensions of existing data and product governance principles rather than by treating AI as entirely separate from the rest of IT and data governance. • Use an AI product owner or technical product owner to define the product vision, coordinate stakeholders, and connect business value, risk, and delivery. • Involve roles such as CIO, CTO, CISO, and other risk/compliance stakeholders in line with the existing governance logic for business-critical digital products. • Start AI initiatives with a business case and explicit value hypothesis, rather than only from technology experimentation.	• Product-style governance as a way to make AI initiatives more scalable, traceable, and better owned. • Without explicit product ownership, AI initiatives risk remaining stuck at proof-of-concept stage instead of being scaled as full products. • Existing executive roles remain relevant; the main challenge is not inventing entirely new governance structures, but organizing impacted roles well around AI rollout. • AI governance is strongest when it combines ownership, traceability, value orientation, and involvement of risk and compliance stakeholders.
People & Organisation	• Move from a CoE to toward broader enablement of business functions, IT, and other domains as AI becomes more developed within the organisation. • Use cross-functional chapters or representative structures where people from business and technology streams co-create AI-related capabilities and initiatives. • Strengthen roles that translate business processes into AI-enabled or AI-led processes, especially AI translators. • Invest in AI engineers and platform engineers to build and operationalise AI capabilities. • Introduce strong AI product-owner roles to hold the end-state vision together and prevent AI initiatives from remaining disconnected experiments. • Build the capability to use AI increasingly for repetitive work, while reserving people more for the complex and exception-heavy work. • Use AI to absorb more repetitive work so that people can focus more on complex tasks, but recognise that concentrating only the hardest work in remaining roles can also create new pressure. • Expect the organisational model to evolve over time as AI moves from augmentation of current work toward more AI-native ways of operating.	• AI translators, AI engineers, platform engineers, and AI product owners are especially important roles for scaling AI. • Product ownership is critical because it helps AI initiatives become scalable products rather than isolated proofs of concept. • Organisations are still largely in an AI-augmented phase, but structures are expected to evolve further if AI-native processes become more common. • Over time, some work is expected to shift away from repetitive tasks toward more complex work, allowing fewer people to handle a larger volume of work.

Domain	Key decisions	Reported impact
Processes	• Start by embedding AI into existing work where it can augment current processes, but move toward deeper redesign as the technology becomes more proven. • Treat AI process transformation as more than adding a tool: process assumptions may need to be removed and workflows reconsidered from a different starting point. • Use A/B-style comparisons between human-only execution and AI-supported execution to assess process impact. • Redesign processes more fundamentally when AI capabilities mature enough to support more autonomous or agentic ways of working. • AI can shift where and how a process begins, for example by detecting and acting on issues earlier rather than waiting for a human to trigger the next step. • Manage adoption, change, and value tracking as part of process change, not as separate afterthoughts. • Build AI initiatives around measurable business cases, for example in customer-service contexts where AI-handled volumes and reduction in transfers to human agents can be tracked.	• The current phase is mainly AI augmentation of existing work rather than full AI-native process redesign. • Process redesign is expected to become more important as AI becomes more proven and trusted in operational settings. • AI can shift where work starts and how it flows, rather than only speeding up existing steps. • In more mature use cases, value can be measured through hard operational numbers, such as the percentage of work handled by AI and reductions in handover to human agents. • AI scaling in processes depends not only on technology, but also on adoption, change management, and the ability to prove value through concrete metrics.
Technology & Infrastructure	• Put in place an AI platform to build, run, and scale AI-enabled products. • Build this AI platform on top of a data platform that supports both structured and unstructured data. • Ensure a clear integration layer, for example through APIs or MCP-style interfaces, so that AI agents and applications can interact with internal and external services. • Consider technology as three connected layers: the data stack, the AI platform stack, and the integration mechanisms connecting them to services. • Make legacy decisions based on the specific process and system landscape involved, rather than treating all legacy systems as a uniform problem. • Use cloud where possible because it is generally the preferred route, while recognizing that some organisations still keep workloads on-premise for greater control. • Balance build-versus-buy choices pragmatically: AI increases what organisations can build themselves, but external tooling remains relevant.	• AI scaling requires more than just a model layer; it also depends on the surrounding data and integration infrastructure. • Integration choices are especially important once agents need to access multiple systems, services, or other agents. • Legacy constraints depend on whether existing systems support the targeted process sufficiently and whether integration support exists. • The technology challenge is broader than AI alone: AI is only one part of a wider platform, data, and integration landscape.

Domain	Key decisions	Reported impact
Data	• Treat data governance as foundational for AI and extend it to AI-specific needs rather than bypassing it. • Ensure that people are clearly responsible and accountable for data. • Invest in data catalogues to improve data discoverability, availability, and lineage. • Govern unstructured data more actively, because AI increasingly depends on it in addition to structured data. • Make common unstructured organisational data available in governed shared forms, for example through common vector stores, instead of having each team repeatedly process the same source material on its own. • Work toward a single source of truth and avoid unnecessary copying of data. • Choose between more centralised and more decentralised data setups pragmatically; both can work if supported by strong data governance, lineage, and ownership. • Use decentralised data products such as function-specific data products where organisational boundaries make this useful, while preserving governance and lineage.	• Data discoverability, availability, and lineage are increasingly important for AI scaling. • Unstructured data becomes a much more serious governance topic with AI than in many earlier data-management setups. • Shared governed access to common data can reduce duplicated effort, for example when teams would otherwise repeatedly create their own vectorised versions of similar content. • Both centralised and decentralised data architectures are workable, but only when data governance is strong enough; without that, both models run into similar problems. • The strongest data principle is to avoid uncontrolled duplication and preserve a reliable source of truth.
Culture	• Invest in training, enabling tools, and sharing of best practices to make AI more accessible across the organisation. • Treat change management as a core workstream in AI scaling, because embedding AI into daily ways of working is harder than the technology itself may suggest. • Combine top-down vision with bottom-up learning and innovation across functions. Both are needed for success. • Use culture-building to help people move from experimenting with AI toward using it effectively in their real work. • Pay attention to first impressions and practical usability, because AI adoption can accelerate or stall depending on early user experience.	• Change management and embedding AI into ways of working is the main bottleneck, more than technology itself. • A top-down push alone is insufficient; broader learning and innovation from different functions are also required. • Training, accessibility, and shared practices are as important for moving from technical availability to real organisational adoption. • Difficulty in proving ROI can also slow adoption, especially when organisations cannot yet connect AI use to hard business outcomes.

B.11 Interview 11 summary

Table B.11: Structured summary of interview 11

Interview details	Interview 11; 19 June 2026; Chief of Staff in a tech focused company, with tech strategy consulting experience; tech product / startup context; 8+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Treat AI as an integral part of the organisation and of the way work is done, rather than as a set of isolated pilot projects. • Put AI rollout under an existing senior technology leadership function such as the CTO or CIO, depending on how tech-focused the organisation is. • Appoint someone with explicit mandate and ownership to drive how AI is rolled out across the organisation instead of leaving it to separated projects. • Use a central approach at least in the early stage in larger organisations, so that security, compliance, cost control, and guardrails are set before broader local experimentation begins. • Centralise initially and decentralise later where appropriate. • Make AI investments conditional on a positive business case, especially if platform or integration investments are big. • Improve visibility into AI consumption and usage costs so that business users and teams understand the financial consequences of model usage.	• Broad AI adoption and value is more likely to fail when it seen as a project instead of a way of working. • Central governance reduces the risk of shadow IT, duplicated work, fragmented experimentation, and uncontrolled costs. • A central rollout model is especially important in large organisations where local units can independently spend budget on AI tools without enough oversight.
People & Organisation	• Add dedicated AI-focused roles where there is enough scale and ambition, such as people focused on process optimisation through AI, AI engineers, or teams that develop internal AI capabilities. • Build different role setups depending on organisational scale: smaller organisations may benefit from one dedicated AI/process role, while larger organisations may require broader internal AI teams. • Use people with process knowledge to identify and improve the work that can be automated or enhanced. • Introduce AI-related roles into existing IT, for example as an additional capability unit within the IT function. • More work can be done without the need for many new people. • Efficiency gains can be redirected toward growth and customer-facing work rather than immediate workforce cuts. • Use change champions from within the business to help drive adoption and communicate change. • Let the people doing the work identify processes for improvement, since they best understand where time is lost and where AI could help.	• Organisations must gather a broader range of AI-focused roles as AI adoption matures. • People with process knowledge are especially valuable because they know where repetitive work sits and where automation could realistically help. • In larger transformations, existing employees who are willing to change are important adoption drivers. • AI-related organisational roles vary by company size and ambition, rather than following one fixed template. • AI will change jobs and people will have to adapt their current work or reskill. • Dedicated ownership increases the likelihood that AI becomes embedded in real organisational improvement work rather than remaining side experimentation.

Domain	Key decisions	Reported impact
Processes	• Focus AI first on text-heavy, repetitive, administratively processes. • Prioritise use cases where AI can free up time for higher-value work rather than simply automate work without a broader goal. • Keep a human in the loop for quality control where needed, especially in reporting, communication, and customer-facing outputs. • Use employees to identify a shortlist of recurring tasks or processes that consume disproportionate time and could be improved with AI. • Apply AI both to support processes such as reporting, email drafting, meeting follow-up, support interactions, and finance-related activities, and to larger core-process examples where the business case is stronger. • For now, in many organisations, use AI mainly for process augmentation and optimisation, but deeper transformation is likely in large-scale operational domains such as banking or customer service.	• There is clear potential in processes that are text-heavy, repetitive, or operationally slow, especially where AI can complete a large share of the work before a human performs the final review. • In some cases, AI enables very substantial efficiency gains, for example by reducing the time spent on recurring reporting or administrative tasks. • Many processes including customer support, finance, and procurement are especially promising areas for AI value. • AI is most useful when it creates time for employees to focus on more important work rather than exclusively being used for FTE-reduction.
Technology & Infrastructure	• If the ambition is to make AI deeply integral to the organisation, embed AI capabilities directly into the IT stack rather than leaving them as separate tools on the side. • Use an internal AI platform or cloud-like organisational environment if AI is to be rolled out broadly and securely across the organisation. • Connect AI to internal data sources through an organised, secure, and restricted access model. • Use APIs and structured access to data so that relevant process-specific information can be selected for particular AI use cases. • Address legacy-system limitations if the aim is to become more AI-native; disconnected legacy environments limit how integration of AI. • Decide on build-versus-buy pragmatically: organisations may buy major platforms, but can also build internal capabilities on top of open models if they have the infrastructure and engineering skills. • Consider on-premise or local alternatives where reducing dependency on large foreign cloud or model providers is strategically important.	• Serious AI scaling is difficult without deep integration into the wider IT stack and access to internal data sources. • A fragmented or poorly connected legacy landscape weakens the ability to become AI-native. • Internal AI platforms are useful when organisations want secure, broad, and organisation-specific access to models and data. • Dependence on a small set of large external AI providers is a strategic risk, especially for organisations that want more control. • Technology choices depend strongly on how integral AI is intended to become; surface-level optimisation requires less infrastructure change than deeper enterprise-wide AI adoption.
Data	• Start by building strong data models and understanding data flows before trying to scale AI internally. • Combine data governance with the organisational ability to store, structure, and access data in the required way. • Use dedicated roles to improve data quality, structure, storage discipline, and clarity about what data exists where. • Work toward a structured and reliable source of truth that can safely feed different AI use cases. • Use centralised data setups in smaller or less complex organisations where that is practical, but allow more decentralised setups in larger organisations where scale and organisational boundaries make this necessary. • Set up data centrally first where possible, and only then decentralise more where organisational complexity requires it. • Use process-specific data selection and access restrictions so that different functions only use the data relevant and permissible for their use case.	• Poor data quality or unclear data organisation is a direct blocker for AI scaling. • Data governance without the actual ability to store and manage data properly is ineffective; organisations must make the desired behaviour possible before enforcing policies. • Structured data management reduces the risk of "garbage in, garbage out" outcomes and is foundational for internal AI use. • Data organisation is context-dependent, but larger and more complex organisations need more deliberate choices around central versus decentralised organisation. • Clear restrictions on who can access which information remain important even when data is made broadly available for AI use.

Domain	Key decisions	Reported impact
Culture	• Communicate clearly that AI should make work more interesting and effective rather than being framed primarily for FTE-reduction. • Position AI as a way to free employees from repetitive and low-value work so they can focus more on judgment and customer value. • Use change champions and internal advocates to help employees adopt AI in practice. • Normalise AI adoption as part of broader technological evolution, while still acknowledging people's fears about role change and job security. • Encourage employees to think proactively about which parts of their work could be improved with AI.	• Fear of job loss is one of the main cultural barriers, especially in larger organisations. • Organisations should aim first at doing more with the same number of people, rather than making headcount reduction the primary narrative. • Cultural adoption improves when employees are helped to see AI as an enabler of better work rather than as a threat.

B.12 Interview 12 summary

Table B.12: Structured summary of interview 12

Interview details	Interview 12; 19 June 2026; Senior advisor in AI transformation and operating model redesign; 25+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Do not rely on innovation departments alone to scale AI, because it leads to a "PoC hell"; a collection of POCs disconnected from operations. • Treat AI as a cross-functional transformation challenge that requires business, IT, compliance, and process ownership to work together. • Active board-level support and steering power for AI transformation, because cross-domain obstacles and painful trade-offs cannot be resolved by bottom-up initiatives alone. • Put clear AI policy and responsible-AI guardrails in place before broad rollout, including clarity on where AI may and may not be used and how higher-risk cases are assessed. • Establish explicit responsible-AI policy, decision routes, and technical boundary conditions up front, so teams know what is allowed, how difficult cases are assessed, and within what (platform) constraints they must work. • Start with strong central coordination in the early stage to keep scarce knowledge, methods, and prioritisation together. • Use a staged governance model: centralise early, and move toward more federated or matrix-like structures as knowledge matures and adoption broadens. • Prioritise use cases explicitly based on value and feasibility rather than trying every idea at once. • Start with lower-risk, more manageable use cases where the data is reasonably in order without too much organisational disruption. • Use rapid prototyping as part of governance to validate whether a use case is worth further investment before committing to full rollout.	• AI governance is difficult because the technology changes faster than large organisations can absorb, especially under strong European regulation. • Innovation-led experimentation without cross-functional governance often leads to proof-of-concept graveyards without broader scaling. • Without strong executive support, organisations fall back into bottom-up experiments that struggle to get priority, funding, and cross-functional cooperation. • Budget is often not the main bottleneck in larger organisations; priority and executive attention are harder to get than money. • A CDO-like role may be useful, but is insufficient on its own; AI scaling needs business, IT, compliance, and executive leadership all actively supporting it. • Rapid prototyping makes prioritisation more evidence-based and allows organisations to secure support and funding more easily by showing realistic outcomes early. • Starting with less controversial and more feasible use cases increases the likelihood of early success and reduces the risk of failing on large moonshots too soon.

Domain	Key decisions	Reported impact
People & Organisation	• Build cross-functional programme teams that combine business, IT, compliance, procurement, and AI expertise. • Organise early AI scaling more as a cross-functional programme team than as a heavy formal hierarchy. • Use a central AI capability group, chapter, or CoE in the beginning to combine scarce knowledge and share methods across the organisation. • Include at minimum an AI engineer, data engineer, business consultant, product owner or project manager, and compliance capability in serious AI delivery teams. • Add infrastructure or cloud specialists where the technical environment is more complex. • Train employees to understand both the possibilities and the limitations of AI, including that AI outputs are not always reliable. • Less people are needed for the same amount of work with proper AI scaling. • Build enough maturity in employees and teams before expecting decentralised AI scaling to work well. • Move more capability toward the business once the central knowledge base is mature enough and the organisation has learned from earlier implementations.	• Work that previously took 5 steps now takes 4, and instead of 100 people you only need 80 if AI is scaled well. • AI engineers are crucial on the application side, supported by data, infrastructure, business, and compliance roles. • Cross-functional teams are the most successful structure because AI almost always cuts across multiple organisational domains. • Centralising knowledge early helps organisations reuse lessons learned and avoid each department independently repeating the same errors. • More mature organisations can shift toward a more matrix-like/decentralised structure, while less mature ones benefit from stronger central coordination. • Employees need both practical AI literacy and realistic expectations about AI imperfection in order to work effectively with it.
Processes	• Use a staged delivery logic of prioritising, prototyping, validating, and then building. • Validate not only technical feasibility, but also whether the use case actually fits the real business process and how employees work in practice. • Start with processes where AI already has strong fit, especially customer service, marketing and communication work, software development, and parts of legal and knowledge-heavy work. • Distinguish between lighter support tools that change little in the organisation and more transformative use cases that alter role structures and process flows significantly. • Treat AI-related process change as classical change management: determine what changes in steps, tasks, skills, and role composition. • Redesign support organisations when AI takes over simpler tasks, so that the remaining human work shifts toward more complex cases. • More advanced use cases such as R&D, tailor-made services, education, and legal are coming up in organisations. • Use rapid prototypes to test not only whether AI can produce an output, but also whether employees actually find the output useful in practice.	• Customer service is the clearest current target area, with transcript generation, summarisation, and progressive automation already delivering visible value. • Marketing and communication work are also as very exposed because AI can increasingly produce and reformat text across channels. • In more transformative cases, the organisation itself changes: if AI handles simpler calls first, the support function needs fewer low-skill handlers and relatively more people able to deal with complex questions. • AI can reduce process steps and labour intensity, for example shifting from five steps to four or reducing staffing needs in certain activities. • Prototyping exposes whether a use case is operationally useful, because process designers may like an idea that frontline employees would never actually use. • Rapid prototypes help not only to validate feasibility and usefulness, but also to mobilise executive support and unlock funding for further development.

Domain	Key decisions	Reported impact
Technology & Infrastructure	• Strengthen automated testing before and after deployment, including not only technical checks but also validation of whether AI outputs remain acceptable. • Extend classic operations and MLOps-style management with additional semantic testing, because model behaviour and the outside world both change over time. • Treat AI integration technically as an API-based architecture problem, but recognise that model choice, sovereignty, and cost make the infrastructure landscape more layered than before. • Use a layered model strategy: deploy larger frontier models only where their capabilities are truly needed, and use smaller or more local models where that is sufficient. • Consider hybrid infrastructure choices based on sovereignty, cost, and task complexity rather than using a single universal setup. • Automate compliance-related checks where possible to reduce the heavy manual governance burden around documentation and approval. • Use infrastructure specialists where secure hosting, scaling, and cloud or on-premise decisions become more complex.	• AI technology is operationally different from normal software because models, data, and surrounding context change way more, making static testing insufficient. • Organisations need additional mechanisms to check whether AI outputs are still semantically correct, not just whether the system is technically running. • Automated testing is increasingly feasible and necessary. • AI infrastructure is becoming more layered, with trade-offs between cloud convenience, sovereignty concerns, model size, and task-specific cost efficiency. • Large frontier models are often unnecessary for more basic routine tasks, bringing opportunities to reduce cost and dependency through smaller or local setups. • Automating parts of compliance and approval work is necessary for competitiveness, especially in regulated European environments where manual governance takes a lot of time.
Data	• Strengthen data governance for both structured and unstructured data, because AI depends on documents, notes, and other less-governed information sources. • Do not wait until all enterprise data is perfect before starting AI, but start in areas where data is already good enough. • Avoid both extremes of negligence and perfectionism in data preparation. • Test if data is suitable for the way LLMs interpret and use information, rather than assuming that human-readable data automatically works well for AI. • Combine central and decentralised data responsibilities: keep platform, a source of truth, and shared standards central, while letting local domains improve and manage their own data quality. • Discipline decentralised data owners and teams to improve the quality and usability of the data they create and maintain. • Use stronger mechanisms to make data entry and data sharing useful for downstream processes, so that data becomes usable across functions. • Explore richer structures such as knowledge-graph-like approaches where they help LLMs interpret organisational knowledge better.	• Enterprise data is often still messy, fragmented, and not governed enough, especially on the unstructured side. • AI scaling adds new pressure because unstructured data that was previously poorly governed now becomes strategically important. • Organisations can still create value by starting in the 'clean corner' of the data landscape, rather than trying to solve everything at once. • LLM suitability is a separate issue from traditional data quality alone: data can be fine for people but still unfit for AI consumption. • Centralised data platforms alone are not enough; local teams must also improve the operational quality of the data they create. • Without some form of discipline and accountability, decentralised units will often fail to maintain data in ways that support downstream AI use.
Culture	• Build a culture in which experimentation, learning, and controlled failure are acceptable parts of improvement. • Create enough psychological and organisational safety for people to try new AI-supported ways of working without being punished immediately when things go wrong. • Balance innovation space with sufficiently clear process targets and business direction. • Treat AI-related change as part of broader organisational learning rather than as a pure technology rollout. • Prepare for resistance and non-adoption, even more so if employees feel threatened or do not want to change how they work.	• AI is similar to earlier digital transformations where people need room to experiment and learn. • Organisations that punish failure too quickly make meaningful AI learning much harder. • Large-scale AI change creates not only winners but also real losers, because some employees will not be able or willing to adapt to the required new skills and ways of working. • Resistance is a major practical issue because employees can quietly undermine the intended new way of working.

B.13 Interview 13 summary

Table B.13: Structured summary of interview 13

Interview details	Interview 13; 30 June 2026; Director in AI transformation and cloud advisory; 15+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Keep the ITSM model, but shift more responsibility and execution toward the business because AI makes self-service development and problem-solving more feasible. • Start with a Centre of Excellence to make early mistakes, form policy, and build shared learning before scaling more broadly. • Move toward a federated model rather than a fully decentralised one, especially in large organisations where regulation, efficiency, and reuse matter. • Use an InnerSource-style model to involve employees from across the organisation in contributing to and spreading AI-related development, instead of keeping it fully centralised in IT. • Prioritise AI initiatives through portfolio management and business-case discipline, with explicit weighting of business value. • Involve finance in AI business cases, for example by linking expected value to headcount reduction, headcount avoidance, and revenue development. • Gain board-level support, primarily from roles such as the CIO and CFO, and increasingly data leadership as well. • Involve compliance in decision-making, but avoid doing so too early in the process to the point that AI stalls. • Improve benefits-management and cost-management disciplines, including model selection, token usage, and value tracking. • Use model routing and make choices based on context rather than applying the best models to all tasks.	• GenAI and agentic AI increase the pressure on the operating model because users expect to build or improve things without waiting on central IT. • A fully decentralised model is ineffective in large organisations because it reduces efficiency, prevents reuse, and increases governance complexity. • Many organisations are weak in benefits management: business cases are approved, solutions are built, but few organisations later track whether the expected value was actually realised. • AI creates broader pressure to build quickly while costs can escalate through poor model choices and unmanaged usage. • Stronger economic steering is necessary because AI is not free or 'automatic' value; value depends on disciplined choices about use cases, models, and scale. • CoE-first pattern is especially relevant for large enterprises, while newer digital-native organisations can structure AI differently from the start.
People & Organisation	• Build the early AI capability group around people with technical expertise, ambition, and clear representation from the business. • Include people with policy and ethical knowledge in the early AI core, especially where regulation and responsible-AI requirements are material. • Combine internal upskilling with external hiring or contracting, rather than assuming one of the two is sufficient on its own. • Recognise that the shift from classical AI to GenAI and agentic AI requires new skills and a different way of building applications. • Focus on people who are willing and motivated to work with AI, because motivation and curiosity are more important than starting skill level alone. • Organise AI delivery in multidisciplinary teams that can build and maintain end-to-end solutions without dependency on separate specialised build teams. • Prioritise technically strong people who can build secure, scalable, enterprise-ready AI platforms, agents, and applications instead of purely experimentation and prototyping.	• A clear split is emerging between people who actively experiment with AI and those who stay away from it. • AI adoption is unavoidable and increasingly central to work. • Less technically skilled but motivated employees are easier to upskill than technically strong employees who resist AI. • Organisations with a strong existing technical capability can scale faster by combining internal talent with external capacity. • The hardest roles to find are people who can translate AI into safe, scalable, and broad organisational value, instead of 'hobbyists'. • Early visible success is crucial for maintaining executive support and momentum during AI transformation.

Domain	Key decisions	Reported impact
Processes	• Apply AI to both support and core processes, without assuming that one category matters more than the other. • Target processes such as HR onboarding/offboarding, finance booking, IT support, software development, claims handling, order processing, and warehouse or pick processes. • Redesign processes from first principles rather than stitching AI onto each existing task step by step. • Start process redesign by asking what the process should look like with minimal human intervention without current assumptions about it. • Use AI to remove unnecessary steps and turn fragmented task flows into more direct end-to-end flows. • Start process redesign from an ambitious end state with minimal human intervention, and only then determine where human involvement is still needed. • Treat AI as relevant not only for deterministic repetitive work, but increasingly also for more agentic and (eventually) highly complex tasks.	• Leading organisations move toward more orchestration-based work. • AI and agentic systems extend current automation (RPA) to a wider set of business activities. • Process redesign that only adds AI onto existing workflows is suboptimal; the biggest gains come from rethinking the entire process. • Core processes such as claims handling, order flow, and warehouse operations are increasingly automatable as AI capabilities combine language, prediction, and agentic coordination. • Organisations that redesign processes radically need fewer human interventions than in their current operating model.
Technology & Infrastructure	• Build AI on top of existing model ecosystems rather than training enterprise models from scratch. • Choose between two broad platform patterns: a hyperscaler-based AI stack or a more containerised agentic stack. • Use hybrid setups where needed, for example combining cloud and more local/containerised deployment to create fallback options and reduce dependency. • Model selection is an architectural choice, including sovereignty, maintainability, copyright, security, and cost implications. • Use model routing so that the system selects the right model for the context. • Integration is a high priority technical concern, including MCP-style patterns, agent-to-agent communication, and the broader integration layer connecting agents to enterprise systems. • Define explicit identity-management rules for agents, including whether they act under human identities or as separate non-human identities with their own permissions. • Prepare for a more app-less and conversational future in which agents handle more of the interaction that is now managed through traditional enterprise applications. • Avoid unnecessary dependence on large external providers where strategic or geopolitical risks make fallback or local alternatives valuable.	• AI creates a shift from classic software building toward a new way of composing applications and agentic building blocks. • Hyperscaler-based platforms are convenient, but more containerised or local setups provide fallback options and greater control with the downside of higher (maintenance) cost, complexity, and skill requirements. • Heavy dependency on one foreign platform/offering can become a strategic business risk due to geopolitical change. • Agentic AI introduces new identity and access-management questions, because agents may need access to data and actions that do not fit existing human-based models. • There will be more deconstruction of traditional applications as work becomes 'mediated' through agents and conversational layers rather than computer screens.

Domain	Key decisions	Reported impact
Data	• Use existing enterprise data sources rather than assuming AI requires entirely new data landscapes. • Treat many AI data challenges primarily as access and integration problems. • Organise data through modern layered architectures such as medallion-style approaches, combined with clear data owners and governance structures. • Follow data minimisation principles, especially where agents access or store context, memory, and personalised information. • Govern temporary memory, operational data stores, and agent memory, including what an agent may store, for how long, and when. • Distinguish between persistent enterprise data and more temporary or interaction-level memory used by agents. • Use integration and access layers to expose the right data safely to the right agent. • Treat data, technology, and infrastructure together as the enabling foundation, rather than assuming that data alone is sufficient.	• Data is not an isolated bottleneck, but part of a broader enabling foundation involving integration and infrastructure. • Modern data architecture patterns such as medallion structures are becoming more common in enterprise settings. • AI comes with new governance challenges around memory, including how much an agent is allowed to remember and when that becomes risky. • The challenge is less about creating new data and more about making existing data reachable, governable, and usable in the right way. • Ownership and lifecycle rules for memory-like data become more important as agents become personalised and deeply embedded in workflows.
Culture	• Build an innovation culture in which AI is actively explored and not treated as something optional. • Create a culture in which employees work with AI themselves rather than leaving experimentation to a small group of specialists. • Ambition and intrinsic motivation are key conditions for successful AI adoption. • Address fear of job loss and broader resistance to change as cultural barriers. • Encourage technical skill development across a much wider group of employees than before. • Use visible adoption from the top and throughout the ranks to normalise the new way of working.	• There is a clear divide between AI adopters and AI avoiders, and this split is likely to grow. • People who engage with AI early gain a structural advantage over those who stay away from it. • Resistance stems from fear, self-doubt, or concern about replacement rather than from lack of capability. • AI scaling requires a broad cultural change to where a larger part of the organisation becomes technically capable and willing to work with AI. • Organisations will see a divide between employees who adapt to AI and those who do not ('winners and losers').

B.14 Interview 14 summary

Table B.14: Structured summary of interview 14

Interview details	Interview 14; 8 July 2026; Independent AI advisor and former consulting partner focused on large-scale transformation in energy and industrial sectors; 25+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Start AI from business problems, friction points, and value opportunities rather than from the question of what the technology can do. • Use AI only where it has clear added business value, rather than pursuing AI for its own sake. • Prioritise AI initiatives through a portfolio funnel with explicit ownership, so that someone has the authority to scale, change direction, or completely stop. • Some boards of companies want 'something with AI', without any value-based reasoning. • Treat AI as a business transformation instead of a technology project. • Use compliance and audit as necessary governance partners when they are needed, but do not let those concerns become the main topic of the AI strategy. • Engage audit and control functions early enough to avoid later blockage, especially in regulated environments. • Use platform and vendor choices, including more sovereign options where needed, to manage security and regulatory constraints pragmatically. • Distinguish explicitly between different levels of AI value creation: task and process value, delivery redesign, and broader value-capture or business-model change. • Run lower-level efficiency initiatives and deeper redesign in parallel instead of assuming value creation moves up a simple staircase.	• A technology-first approach is one of the main reasons organisations fail to realise value from AI. • Lack of ownership leads to pilot fatigue, because initiatives are started but very few are properly scaled or stopped at the right time. • Governance works best when it is tied to business value rather than to AI hype or isolated experimentation. • Compliance is as a gatekeeper that must be managed well, but not as the main source of direction. • Organisations confuse task automation with strategic transformation, which leads them to overestimation of low-level AI initiatives. • The three-level value lens helps distinguish between commodity-level AI use, operating-model redesign, and the much more advanced case of actual business-model change.
People & Organisation	• Build AI capability around people who can solve practical business problems quickly and make value visible in day-to-day work. • Use internal champions who can demonstrate small but meaningful wins and help spread adoption across the organisation. • Combine internal capability building with external expertise where specialised engineering or transformation skill is still scarce. • Be selective in hiring and team composition, especially where strong senior engineering judgment is needed. • Do not remove junior talent too aggressively from the operating model, because junior roles remain important for learning, scaling, and future capability development. • Choose a structure that fits the organisation's broader culture and maturity and do not assume a standard preference for CoEs or federated models. • Broad practical AI learning/engagement is needed, because employees who never learn to work with it risk becoming less relevant over time.	• AI adoption does not automatically reduce headcount; in many cases it creates new work around checking, managing, and improving outputs. • Organisations still need strong people capability even when AI is technically available, because value realisation depends on who can apply it well. • Quick visible wins are important for building organisational momentum and reducing scepticism. • Hollowing out junior jobs weakens long-term learning and scaling capacity. • Structural design is contingent: more centralised or more federated models can both work, but only if they fit the organisation's existing culture and way of operating.

Domain	Key decisions	Reported impact
Processes	• Start process selection by identifying friction, such as waiting time, handoffs, rework, inconsistent data, or other points where value is lost. • Use AI to redesign or remove broken and unnecessary processes rather than automating failing processes as they are. • Avoid over-automating interactions where practical human inspection or intervention remains essential. • Focus on cross-silo process problems rather than isolated local optimisations, because larger value often sits in the gaps between functions. • Use incremental, high-impact wins to build support, but connect them to a broader process-redesign logic. • Prioritise functions with greater share in the cost base or value chain, rather than only selecting use cases with the highest expected percentage ROI. • Look beyond familiar starting points such as customer contact and consider core operational domains where AI can have a bigger economic effect. • Treat Level 1 task improvement as a baseline, but invest in Level 2 delivery redesign where AI makes fundamentally different workflows possible.	• AI pilots fail because they remain too narrow, too local, or disconnected from the broader process context. • Significant gains can come from enhancement-oriented AI use cases that improve existing operations without fundamentally changing the underlying service or business model. • The strongest process value does not come from speeding up existing tasks, but from changing how work is structured. • Cross-silo redesign is much more powerful than siloed experimentation, because organisational friction often sits between teams rather than inside one task. • Functions such as supply chain and core operations may offer more total value than smaller, high-ROI support functions, because their economic weight is larger. • Unfitting chat-based automation lead to "chat talking to chat" loops, while real value still depended on human inspection. • Administrative, rules-heavy work is much easier with AI. • Most organisations remain stuck at task automation, but the bigger operating-model gains come from delivery redesign.
Technology & Infrastructure	• Use platforms and vendor solutions pragmatically rather than a custom-built AI infrastructure for critical use cases. • Prefer certified, supported, and secure solutions where liability, regulation, or criticality are high. • Use vendor- consumption-based access to AI where rather than assuming capacity must be built in-house. • AI requires strong system and data integrations for higher level value. • Use AI as part of the integration layer where it can help bridge across systems and reduce unnecessary customisation in enterprise platforms. • Introduce AI oversight tooling for visibility into how AI initiatives are performing, their cost, and where they are being used. • Allow citizen development within clear boundaries, so that local experimentation does not create disproportionate operational or financial risk. • The technology stack is expected to change significantly over time, but not yet in a way that justifies premature large-scale redesign. • Use sovereign or better-controlled platform options where necessary to manage security, compliance, and data-location concerns.	• Build-versus-buy is a risk-management issue as much as a technical one. • In critical contexts, fully custom internal AI builds are seen as less attractive because support, certification, and accountability become harder. • AI can reduce some integration complexity by acting as a more flexible layer between systems, especially where traditional enterprise platforms have become heavily customised. • Technology on its own is not framed as the main constraint; the harder challenge is embedding it into workable process and platform structures.
Data	• Poor data consistency and weak shared definitions/single source of truth are serious barriers to AI scaling. • Define data products and shared data concepts more clearly so that stakeholders work with the same underlying meaning and structure. • Clarify who and what agents can access which data and under what conditions. • Data governance that balances sovereignty, control, and flexibility rather than aiming for perfect data architecture before acting. • Address the data mess in operational environments directly, for example by standardising naming, structure, and reference logic across sites or business units.	• Data in many organisations is highly inconsistent and fragmented, which limits AI effectiveness. • AI does not solve fundamentally broken data organisation; it often exposes those weaknesses even more. • Shared definitions and access management is essential because even basic disagreement about what data means stalls scaling. • In an organisational example, the hard work is not inventing in new AI techniques, but in aligning existing data foundations well enough to support them.

Domain	Key decisions	Reported impact
Culture	• Build a culture in which experimentation with AI is encouraged and failures do not immediately lead to punishment. • Counter AI hype with realism about where AI can and cannot create meaningful value. • Treat AI as something that amplifies existing strengths and weaknesses instead of as a magic solution. • Encourage employees to engage with AI practically, because adaptation matters more than only having opinions. • Have a pragmatic view of AI-ambition: not every organisation should aim at full business-model reinvention.	• Culture is the main determinant of whether AI strategy can actually be executed. • Organisations that do not allow experimentation safely are less likely to learn fast enough to scale AI well. • AI tends to amplify existing cultural and organisational problems rather than fix them. • AI-transformation narratives are often inflated and only a small minority of companies are positioned for true business-model reinvention through AI. • For most organisations, the more realistic challenge is to become adaptive and value-focused, not revolutionary.

B.15 Interview 15 summary

Table B.15: Structured summary of interview 15

Interview details	Interview 15; 9 July 2026; Director in AI, innovation, and workforce transformation; 15+ years	
Domain	**Key decisions**	**Reported impact**
Governance	• Make AI-specific roles, responsibilities, and accountability explicit because existing governance does not cover it. • Redesign governance around the fact that many employees effectively take on a managerial role when they delegate work to AI agents. • Define clearly who remains accountable when AI contributes to errors in deliverables, decisions, or client-facing outputs. • Build governance that supports scaling by clarifying issues around delegation, AI outputs, escalation, and end responsibility. • There is no one-size-fits-all AI governance model; governance choices depend on where responsibilities sit within the organisation. • Base AI investment choices on business strategy and long-term value rather than on hype, vendor marketing, or short-term efficiency narratives. • Organise AI investment decisions through a funnel in which ideas are filtered based on strategic relevance, expected value, feasibility, and fit with the organisation.	• Unclear accountability for AI-assisted work creates friction and can block scaling because employees do not know who is responsible when something goes wrong. • Governance gaps are especially visible where traditional 'sign-off' structures are clear for human work but not yet for AI-supported work. • Strategic discipline is needed to avoid buying attractive-looking tools that only act as temporary patches on deeper problems. • There is a gap between vendor promises and current product reality, but products do improve over time, so organisations need both critical judgment and patience.

Domain	Key decisions	Reported impact
People & Organisation	• Keep domain specialists and AI specialists working together rather than expecting domain experts to become full AI specialists. • Build reskilling pathways for employees who only partially match future roles, for example through 18- to 24-month learning paths with a starting point of a 50% match. • Redesign reward systems where AI changes responsibility, task composition, or role expectations; don't give people more work and accountability without a reward. • Develop people not only in hard skills, but especially in adaptability, willingness to learn, and the ability to contribute to the organisation of the future. • Support middle managers if they are expected to take on more AI-related responsibility, rather than assuming they will do it automatically. • Recognise that the main bottleneck is often not scarce technical skill, but the human ability to adapt to changing technology. • Recognise that technical skills can often be bought in through vendors or external specialists, but the organisation's own adaptability cannot.	• Workforce change is broader than simply hiring AI engineers or data scientists. • Organisations that do structured transition planning can reduce talent loss while preparing people for future roles. • If responsibility increases without matching clarity, support, or reward, middle management can become a point of stagnation rather than acceleration. • Organisations will struggle more with adaptability and continuous learning than with access to technical products. • Organisations that rely heavily on bought-in technology still face a substantial internal people challenge, because employees must learn to work with it. • Tools embraced by executives, for example because a demo was 'sold' to them, do not always work well on the work floor, which can create resistance and hesitation in the other layers of the organisation.
Processes	• Let the business decide whether AI should optimise the current core process or support the invention of a new one. • Let process priorities follow market position: if the organisation is being overtaken in a core process, it may need to rethink that process rather than only optimise it. • If improving the current core process still strengthens margins or cost position, that can also be a valid strategic play. • Prioritise AI use cases that help the organisation to win in the long term rather than only aiming for short-term headcount efficiencies. • Distinguish between three levels of AI-enabled value: individual productivity, team or process productivity, and organisation-level productivity across multiple connected processes. • Treat individual productivity tools (e.g. Copilot) as a baseline layer that employees expect to have, even when these tools do not always produce strong direct ROI. • Use the middle level to improve team processes and coordinated work rather than only isolated individual tasks. • Use the top level for agent-based or highly automated end-to-end workflows across multiple connected processes around the final goal instead of around the current intermediate steps. • Avoid short-sighted use cases that only patch one step if broader redesign would create stronger strategic value. • Start radical redesign from the desired start point, end point, and human-in-the-loop requirements, rather than from the current assumptions and process design. • Use a greenfield perspective for this redesign and decide explicitly which human skills should still be retained if the system fails.	• There is a clear difference between using AI for individual productivity and for broader organisational transformation; the first can improve work quality, reduce stress, and support work, but does not bring the highest ROI. • Greater value comes when AI connects multiple steps and processes around a broader organisational goal. • Short-term efficiency focused on removing a few FTEs is short-sighted compared to long-term margin improvement or competitive positioning. • End-to-end redesign is riskier, but also a more meaningful way to reach real transformation. • Process priorities differ by organisation: some may need core-process reinvention to remain competitive, while others may benefit more from improving existing core processes. • Organisations should use AI to play to win rather than to play not to lose, meaning that AI choices should support long-term margin improvement, competitiveness, or strategic renewal rather than only short-term efficiency.

Domain	Key decisions	Reported impact
Technology & Infrastructure	• Do not assume that vendor solutions remove the need for internal architectural choices. • Know what you are buying before selecting AI tools or vendors, and first define clearly which problem the organisation is actually trying to solve. • Avoid buying AI as a quick patch for visible pain points without first understanding the underlying problem. • Evaluate vendors and models on a broader set of criteria, including cost, fit with the existing tech stack, long-term pricing behaviour, and EU-related requirements such as sovereignty. • Avoid using expensive models or token-heavy solutions for relatively simple problems where a lighter option would be sufficient. • Compare technical options deliberately rather than defaulting to the most powerful or most visible model or vendor. • Evaluate whether legacy systems can provide sufficient access to data and functionality before layering AI or agents on top of them. • Redesign the tech stack where necessary to make it more future-proof instead of repeatedly adding new AI layers to unsuitable legacy environments. • Treat access to underlying systems and data as a deeper issue than whether agents can technically tolerate slow or awkward interfaces. • Avoid becoming structurally dependent on a vendor that alone understands or controls access to critical legacy systems.	• Vendor offerings are partly product marketing and partly real product improvement. • Products do improve over time, so organisations also need some patience rather than judging too early or taking vendor claims at face value. • Legacy environments are a major constraint when organisations want to move from experimentation to deeper AI-enabled change. • Layering AI repeatedly on top of old systems leads to weak architectures rather than durable transformation. • Vendor dependence becomes more serious when critical processes run on old systems that few internal people still understand. • Poorly governed access to legacy systems is a strategic risk, especially in public-sector or other sensitive core processes. • Expensive models can become a suboptimal choice when organisations fail to weigh costs properly against actual business needs. • Sovereignty and EU-related constraints can act as hard knock-out criteria in vendor and platform decisions.
Data	• Treat data as a foundational precondition for useful AI rather than as a secondary issue. • Ensure that AI use cases can draw from enough relevant data sources before treating outputs as reliable. • Evaluate how often predictions or outputs are likely to be wrong when the available data is incomplete or weak. • Use context-specific data (e.g., through RAG) to shape AI behaviour in domain processes rather than assuming pre-trained models already know local rules, language, or decision logic. • Account for organisational nuance such as language, ethics, and communication culture when applying AI in real work settings. • Avoid treating generic pre-trained models as sufficient for context-heavy decisions such as objections handling or similar judgment-intensive processes.	• AI is fundamentally a prediction mechanism, and its quality depends on the underlying data. • Weak or incomplete data can lead organisations to create false confidence in poor outputs and to mistake flawed AI support for truth. • AI may produce directionally correct outputs without being reliable enough for high-stakes decisions if not used correctly. • Generic models lack organisation-specific context, including local language, ethical nuance, and communication style. • Even strong generic models may not match how work is actually performed if they are not grounded in the organisation's own data and context.

Domain	Key decisions	Reported impact
Culture	• Build a culture that rewards adaptability, continuous learning, and willingness to work toward the future organisation rather than preserving today's roles unchanged. • Create enough time and support for employees to keep learning as AI changes tasks and expectations. • AI is partly a response to structural workforce pressure, such as an ageing workforce, retirement-driven knowledge loss, and the need to preserve productivity despite demographic change. • Improve psychological safety by clarifying accountability and acceptable use when people work with AI. • Avoid treating AI as a magical solution, and instead use it as a strategic means to reach business goals. • Recognise that local communication style, ethics, and work culture shape whether AI outputs actually fit the organisation.	• Adaptability and learning mindset are more difficult and more important than many technical hard skills. • Cultural resistance becomes stronger when employees face more uncertainty, more responsibility, and unclear boundaries around AI use. • Organisations that do not make time for learning risk slowing adoption even when the technology is available. • AI tools can feel misaligned in practice when they reflect external language or norms rather than the organisation's own communication and work culture. • As long as people remain part of the process, AI value cannot be reduced to one fixed number, because different people will still use the same tools differently and create variation in outcomes. • Technology can raise (FTE) productivity, but historically major waves of digital and industrial change have tended to create new work rather than eliminate it.

Appendix C

Compact interview summaries

C.1 Interview 1 compact summary

Table C.1: Compact summary of interview 1

Interview details	Interview 1; Senior manager in AI consulting; enterprise transformation context; 11+ years
Domain	**Compact summary**
Governance	• AI governance set up as separate practice, not 1:1 reuse of IT governance; structurally similar but larger implications and a different mindset. • Risk officer capacity expanded; new skills required for AI-specific risk assessment. • Standardised assessment tooling/templates make legal & regulatory compliance scalable across many use cases. • Guardrails as first-class design element, especially in customer-facing core processes. • Treat AI scaling as a broader transformation instead of a collection of isolated use cases. • Use case prioritisation along three axes: business impact, feasibility, and broader AI adoption. • Readiness conditions per area (architecture, AI-ready data, governance, skills) required before execution starts. • Traditional KPIs remain leading; AI is a means, not a goal.
People & Organisation	• Workforce segmented into three personas (no-code, low-code, pro-code) driving upskilling, reskilling, and workforce planning; without it, real risk of wasted training spend. • Strategic workforce planning sets target persona mix in five years; hiring/attrition aligned. • Delivery model: distributed teams for smaller orgs; hub-and-spoke + CoE for larger orgs: federated models slow decisions. • New roles: agentic engineer, agentic platform lead, extra solution architects, more data stewards. • Data steward role strengthened; demand for data-responsible roles rises. • Prompt, context, and (less so) harness engineering = required skills for new hires, not separate titles. • External hiring for first agentic wave + parallel internal reskilling. • Value grows much faster than headcount; AI partly funded by slowing workforce growth or redirecting planned expansion. • Some roles unsustainable → reskilling or managed exit. • Dedicated workstream for employee concerns: adoption, reskilling, outplacement for disappearing roles.
Processes	• Shift from inserting AI into existing delivery to agent-oriented design orchestrating intelligent workflows where agents perform tasks. • 60–90% time reduction in specific processes like claim processing. • Guardrails as first-class design element, especially in customer-facing core processes. • DevOps extended into MLOps and LLMOps. • AI applied to core ops (claims, customer interaction); e.g., fully autonomous payment arrangements with no human-in-the-loop where guardrails suffice. • Broad test portfolio (unit, integration, system, UAT) = practical trust mechanism for AI-generated code without manual review. • Process and infrastructure change inseparable in practice.

Domain	Compact summary
Technology & Infrastructure	• Existing strategic vendors kept and extended; no stack replacement, but deliberate platform extension. • Internal agentic development platform built for consistent and safe development without data leaking into public AI tools. • Agentic systems add complexity in IAM, multi-agent coordination, and monitoring non-deterministic behaviour. • Cloud standardised for AI; on-prem mainly retained in parts of public sector, actively decommissioned elsewhere (cheaper, more maintainable, more scalable). • Heavy investment in API management and IAM, incl. agent identities (e.g. Entra agent ID) — IAM materially more complex with agents acting on behalf of users and each other. • Modern data platform (e.g. Databricks) is key for AI delivery. • Internal agentic platform maintained; token consumption budgeted as recurring structural cost. • Low-code phased out over time; agentic coding assistants (Codex, Claude Code, Gemini) make pro-code accessible. • Cost base shifts from headcount to IT/platforms; leadership communicates this openly.
Data	• Data readiness = precondition, not a parallel workstream; it is the hardest gating factor; ambitions stall until foundation fixed. • Explicit data prep phase often required before scaling use cases. • Invest in monitoring, structuring, broader availability; buy/build data where insufficient. • Data products expanded for trusted access and to standardise core concepts (one definition each). • Data steward role strengthened; demand for data-responsible roles rises. • AI is both a consumer and a generator of data, turning operational or unstructured inputs into structured data. • Medallion architecture (bronze/silver/gold); agentic claim systems use bronze, exec reporting uses gold. • Existing data and probability models often need redesign for AI. • Raises the bar on non-specialists to work with cleaner inputs and more deliberate data thinking.
Culture	• Aim for innovation culture, not a separate "AI culture". • Bottom-up enablement + explicit, unhedged top-down commitment. • Dedicated workstream for employee concerns: adoption, reskilling, outplacement for disappearing roles. • AI as default design option: every new initiative considered through an AI lens. • Real employee uncertainty exists; visible leadership + concrete support lowers resistance. • Culture is a precondition for scaling, not a soft side-topic.
Cross-domain dependencies	• Data + Technology & Infrastructure + Governance + People & Organisation: AI scaling depends on readiness across data, architecture, governance, and skills before execution can start. • Processes + Governance: Use-case selection is guided by business value and broader transformation potential rather than by isolated AI opportunities.

C.2 Interview 2 compact summary

Table C.2: Compact summary of interview 2

Interview details	Interview 2; AI product manager and innovation centre lead; 7+ years
Domain	**Compact summary**
Governance	• Anchor AI ownership in a C-suite role (CEO/COO), not IT alone; the main factor for whether AI scaling succeeds or stalls. • Frame AI as a multi-disciplinary transformation with explicit top-down priority, not a tech topic. • Embed data privacy, risk, and compliance into AI governance from the start. • Full redesign requires top-down mandate; most organisations are not mature enough yet. • Weak/ambiguous mandate causes AI to stall at middle management; clear mandate unlocks investment in tooling, talent, and processes.
People & Organisation	• Invest in a central foundational platform (CoE) enabling decentralised building by teams/individuals. • Shift to smaller, flatter, high-performing teams. • Hybrid roles and cross-disciplinary engineers become more important. • Hire selectively for high agency, drive, and demonstrated AI experience (e.g., own agents) over formal credentials. • Slow overall hiring as AI raises individual productivity; reshape profiles hired. • Bar for individual contribution rises sharply: employees refusing to use AI become hard to justify. • Job security concerns escalating; works councils becoming involved. • Self-directed AI experimentation treated as part of the role, not optional.

Domain	Compact summary
Processes	• Redesign core processes AI-first rather than inserting AI into existing steps. • Up to 90% time reduction estimated in some text-heavy processes (e.g., a letter that previously required 5 reviewers handled by 1). • External pressure rising: citizens already use AI to scale objections/subsidy applications, increasing inbound workload and forcing automation. • Full redesign requires top-down mandate; most organisations not yet mature enough.
Technology & Infrastructure	• Invest in a central AI/agentic platform for development and reuse. • Standardise on a small set of foundational tools (Cursor, Claude Code, chosen harness); CIOs consolidating vendors to reduce fragmentation. • Build a reusable "harness" (shared platform/tools/settings) for personalisation on a common foundation. • Adopt agentic development environments (Cursor, Claude Code) as core workflow; background agents test bugs/front-ends 24/7; reporting agents automate recurring tasks. • Use MCP and similar interfaces to connect tools, agents, and enterprise systems. • Decide both which model and which harness/working environment is provided. • Phasing out legacy monitoring: traditional tools are unfit for AI workflows. • Sovereignty and local deployment constrained by regulation and power capacity limits, especially in public sector. • Agentic coding tools drastically improve model efficiency and developer output. • Agent memory/context managed ad-hoc via markdown files (claude.md); fragmented but critical; tooling immature.
Data	• "Garbage in, garbage out" as operating principle; AI quality bounded by reachable data quality. • Connect AI to all relevant systems (email, Slack, CRM, docs) for context and memory. • Shift to "ambient data" captured passively (meeting transcripts, tool exhaust) instead of manual entry; manual data entry decreases while AI pulls richer context. • Exploit unstructured data more aggressively than legacy systems allowed. • Front-runners in AI scaling are organisations with connected systems and standardised retrieval interfaces; production AI failures often trace back to disconnected/low-quality source systems.
Culture	• Frame ~80% of AI success as coming from individual employees and teams, ~20% from formal innovation. • Leadership permission and active encouragement is a primary determinant of adoption; without it, training programmes alone do not change behaviour. • Visible senior-leader AI use pulls others along; its absence stalls adoption. • Dedicated channels to celebrate AI wins. • Bring AI use more into evaluations, though metrics like token usage are imperfect. • Budget experimentation time explicitly as investment with delayed payoff. • Address job-security concerns openly. • Personal AI investment becomes a clear individual differentiator. • Self-directed AI experimentation treated as part of the role, not optional.
Cross-domain dependencies	• Processes + Data + Technology & Infrastructure + Culture: AI scaling depends on coordinated change across processes, technology, data, and culture rather than isolated improvements in one area. • Technology & Infrastructure + Governance: A central AI foundation enables safer decentralised development by business teams. • Processes + Data + Technology & Infrastructure: AI-first process redesign only works when tools, agents, and data sources are connected.

C.3 Interview 3 compact summary

Table C.3: Compact summary of interview 3

Interview details	Interview 3; Co-founder and CEO of an AI-native scaleup; 15+ years
Domain	**Compact summary**
Governance	• AI is a board-level topic, not IT-only; without top-down ownership and explicit budget, scaling stalls regardless of technology quality. • Consortium-style governance: central data/AI platform team as agenda-driver, surrounded by DPO and CISO as influencers for security and regulation; CISO sometimes acts as initiator on AI security/exposure. • Business owns AI demand, IT delivers it, this prevents "technology looking for a problem." • AI funded as organisational programme, not ad-hoc from departmental budgets. • Cost control becoming a governance discipline: token consumption can exceed developer salaries; per-user budget caps deployed where usage runs out of hand. • KPIs for AI scaling are needed, yet often still missing. • FTE and topline impact remain limited; change management is more constraining than technology. • Investment in change management crucial; adoption stalls even when tech works. • Multi-disciplinary steering (finance, security, IT, business) formally involved in AI.

Domain	Compact summary
People & Organisation	• Standard early pattern: central CoE + hub-and-spoke; CoE owns blueprints, shared services, scarce talent; domains build on top and localise. • CoE is designed to dissolve over time as domains mature; centralisation to decentralisation is a planned direction to move towards. • Scarce roles (data stewards, ML engineers, agentic engineers) hired centrally first; existing data scientists/ML engineers pivot into agentic roles. • New role emerging: "agent editor"/"chatbot editor"; senior domain expert curating context, edge cases, continuous improvement. • Non-technical "tinkerers" treated as part of the AI workforce; ambassadors bridge engineers and non-engineers. • AI-native hiring filter in startups: non-AI-native candidates not hired; refusal to engage with AI deemed not survivable in high-performing AI-native environments. • Strict Dutch labour protection slows workforce reshaping; employees may stay employed for years even when work is largely automatable. • Expert work is increasingly being recorded to train AI models.
Processes	• Software engineering identified as the first process disrupted end-to-end (spec, code, test, deploy); large share of global token usage going to coding agents. • Sales and Go-To-Market are very AI-augmented: account research, intent, outreach, meeting analysis, follow-ups, reporting; led to measurable productivity effects. • AI agents used for ongoing CRM data quality (flag missing fields, stale deals/meetings). • Claim handling a primary AI use case in financial services; mature use cases often solved via existing point solutions rather than custom builds. • Legal-tech is another visibly disrupted knowledge-work area. • Only a small amount of organisations have completed AI-driven transformation of a core process; in-house builds fall behind because market moves faster than 18-month build cycles.
Technology & Infrastructure	• Consortium-style: central data/AI platform team needed. • Buy from strategic vendors (Microsoft, Azure, hyperscalers); extend stack deliberately with AI orchestration and governance tooling; full stack replacement not needed. • Sovereignty becoming a high priority architectural focus in public sector/regulated; upcoming EU rules may forbid government workloads on non-European tech. • Sovereignty is not just for GDPR, but broader and strategic. • Per legacy system: deliberate choice to 1: modernise, 2: keep it with agents on top, or 3: use agents to accelerate migration; legacy no longer auto-blocker because agents don't 'care' as long as the data is correct. • MCP-style interfaces emerging as a new required capability to expose existing systems/data warehouses to agents safely. • Enterprise systems not yet designed for agent authentication, authorisation, and controlled action execution, but this is crucial.
Data	• Agents must not inherit human permissions 1:1; they require separate authentication/authorisation and action boundaries; human-in-the-loop policies explicit. • Data governance + Identity Access Management is the hardest barrier for AI scaling; agent authentication/authorisation being redesigned industry-wide with no mature standard. • Data must be made "agent-ready"; agents treated as primary consumers alongside humans/dashboards. • Data warehouses (e.g. Azure) used as foundation, exposed via MCPs. • Agents used to improve data quality (CRM hygiene agents). • Counter-intuitive finding: agents more tolerant of poor data than humans (patient large-volume reading); reduces but does not remove the data-quality bottleneck. • "Fix the data first" sometimes legitimate, sometimes an excuse to delay agent programmes.
Culture	• Build the operating model around AI advocates/change-makers, not sceptics. • "AI-native by default" used as hiring filter; refusal to engage with AI framed as long-term individual career risk. • Empower individual contributors with personal "swarm" of agents instead of adding middle-management layers. • Middle management whose value rests on information asymmetry/gatekeeping flagged as first at risk in AI-native operating models. • AI used as sparring partner for senior staff (consultants, leaders), not just for generating output. • Investment in change management crucial; adoption stalls even when tech works. • "Cassandra" dynamic: some people see major disruption coming while others act as if nothing fundamental is changing.
Cross-domain dependencies	• People & Organisation + Governance: Board-level ownership is a prerequisite for scaling workforce redesign. • Technology & Infrastructure + Data + Governance: AI scaling requires explicit sponsorship, budget, and control over access, identity, and data readiness. • People & Organisation + Governance: Early centralisation through a CoE helps work with scarcity, with decentralisation only when business domains/units mature. • Culture + Governance: Leadership commitment and change management shape whether adoption scales beyond experimentation.

C.4 Interview 4 compact summary

Table C.4: Compact summary of interview 4

Interview details	Interview 4; Senior director in AI, data, analytics, and transformation; commercial & industrial sector; 13+ years
Domain	**Compact summary**
Governance	• AI maturity assessment (data + tech + process + governance) as mandatory starting point before scaling decisions. • Prioritise use cases via outside-in value lens; translate into roadmap across operating model dimensions. • Business ownership of AI use cases and value realisation; IT/data/AI act as enablers; this is a necessary shift since value does not sit in IT alone. • Standardised intake and business-case template (impact vs. effort) to filter ideas. • Explicit design choice between central, hybrid/federated, or autonomous business-led setups; context-dependent, no one-size-fits-all. • Central AI policy and guardrails aligned with regulation (incl. AI Act); central standardisation of dev steps, documentation, compliance. • AI CoE for knowledge sharing, reuse, avoiding duplication across business units. • Investments aligned to governance model (central enterprise-wide vs. divisional funding). • Many organisations are stuck in scattered pilots/copilot rollouts without embedded governance. • Governance choices for control, compliance, team structures, responsibilities, and investment prioritisation.
People & Organisation	• Explicit design choice between central, hybrid/federated, or autonomous business-led setups; context-dependent, no one-size-fits-all. • New roles: AI engineers, AI architects, AI product owners. • Reskill data scientists and developers into AI-oriented work; broad upskilling across all employees, not just specialists. • Workforce-planning/skills-gap logic to identify missing capabilities and evolvable roles. • Role redesign: employees move from producing every step to validating, supervising, and steering agentic workflows. • Rebalance responsibilities between central and divisional teams where AI capability exists in both. • Twofold impact: new roles gain importance; many existing roles change as parts of work are automated. • Administrative burden decreases for some employees, while roles with routine administrative work come under pressure.
Processes	• Target repetitive, manual, high-volume work first as clearest automation candidates. • Redesign workflows around agentic logic; AI runs operational flow; humans intervene at critical decision/validation points. • Apply AI in both support and primary processes (primary requires more contextual tailoring). • Automate data capture/admin handling (e.g. speech-to-structured data, automated email processing). • Evaluate via process KPIs: productivity, throughput, cycle-time, business outcomes. • Most organisations have not completed broad process transformation; still copilots and isolated use cases. • Largest short-term gains where there is manual coordination, repetition, and administration. • Human work shifts toward review, validation, exception handling and high impact work.
Technology & Infrastructure	• Assess current landscape and data platform/integration readiness before scaling. • Introduce platform/integration layer above point-to-point systems when AI has to work cross-system. • Conscious buy-vs-build decisions based on use case, strategy, market availability, and strategic differentiation; this is becoming more and more complex because of rapid market expansion. • IT is actively involved in architecture, tooling, security; leverage existing IT portfolio management where available. • Target setup: IT enables security/EA/tech decisions alongside dedicated data function and business ownership; cross-functional effort (business + data + IT) for AI decisions. • Bottleneck: fragmented systems, point-to-point integrations, Excel workarounds limit cross-system/end-to-end AI workflows. • Legacy systems are a significant constraint; some organisations must fix foundations (incl. single source of truth gaps) before scaling AI. • AI can accelerate migration of legacy systems. • Security becomes a higher priority because of new AI-specific risks. • AI also reshapes IT itself; developers and IT support gain efficiency through AI use.

Domain	Compact summary
Data	• Data readiness, management, and governance treated as core preconditions for AI scaling. • Push toward clearer data definitions and single source of truth (often blocked by legacy). • Distinguish structured-data-heavy use cases (ERP-driven) from knowledge-base cases (SharePoint/docs). Knowledge-based use cases scale faster. • Assign data owners and data stewards; actively monitor data quality. • AI acts as both consumer and generator of structured data (e.g. speech/email/field input transformed into structured data). • "Garbage in, garbage out" is a big bottleneck, particularly for predictive/data-heavy use cases; knowledge-based use cases are less constrained. • AI acts as catalyst for making data foundation urgency visible.
Culture	• Invest in AI literacy across the organisation (opportunities and risks). • Treat adoption as scaling workstream; deployment alone does not lead to usage. • Leverage early adopters and ambassadors. • Address risk, data exposure, and job-impact concerns directly. • Many copilot rollouts lack accompanying adoption effort. • Workforce split between enthusiastic experimenters and sceptics; resistance is driven by job fears, risk concerns, and mostly a lack of knowledge. • Adoption lag is a recurring challenge; broad embedding in daily work comes after the technical rollout.
Cross-domain dependencies	• Processes + People & Organisation + Governance: AI scaling is an operating model challenge in which business ownership and structured prioritisation direct the transformation. • Processes + Data + Technology & Infrastructure: Process redesign depends on data quality and integration rather than on AI tooling alone. • People & Organisation + Processes + Culture: Role redesign toward validation and exception handling depends on both process redesign and employee adoption.

C.5 Interview 5 compact summary

Table C.5: Compact summary of interview 5

Interview details	Interview 5; Former regional head of AI and data in insurance and telecom; 20+ years
Domain	**Compact summary**
Governance	• AI positioned as enablement layer inside business strategy, not standalone transformation; priorities steered by business value streams, not AI itself. • Federated hub-and-spoke: regional CoE + local country capabilities; regional CoE grew significantly as scaling matured. • Federated governance balances regional coordination with local autonomy, enabling local relevance and regional scale but requiring continuous cross-market alignment. • Portfolio funnel aggregates local backlogs, identifies overlaps, and selects regional priorities based on critical mass, strategic fit, and broader impact. • Standardised staged pathway with gates from idea/concept to business case, prototype/POC, and embedding before scaling. • KPI mix: value creation, cost efficiency, NPS, straight-through-processing. • Multi-country complexity (regulation, business models, channels) makes governance inherently hard; some markets must support non-local priorities to unlock later regional commonality. • Scaling depends on finding enough commonality across markets to justify regional prioritisation.
People & Organisation	• Federated hub-and-spoke: regional CoE + local country capabilities; regional CoE grew significantly as scaling matured. • Standardised core roles (data scientists, AI engineers, MLOps) and minimum AI capability baseline per country, achieved via hiring and reskilling. • Capability building is not only technical: management training and business-translator roles are needed to generate and absorb AI opportunities effectively. • Layered capability building: broad fundamentals, role-specific tracks, pilot-based learning, management training; business translator initiatives for opportunity spotting. • Workforce reduction framed as late-stage consequence of mature automation, not upfront goal; transition planning needed as some roles disappear and new ones emerge. • Countries lacking AI capabilities forced into hire/reskill push; cultural + capability interventions needed for business users to engage effectively.

Domain	Compact summary
Processes	• Priority domains: retention, acquisition, claims management, underwriting, selected back-office; highest strategic AI impact. • Greenfield end-to-end redesign: decompose processes into modular components, rebuild around agents/agent combinations, push toward straight-through-processing. • Hybrid demand generation: top-down "big bets" + bottom-up use cases from trained business staff (translators). • AI scaling moves beyond local automation toward domain-based redesign; effectiveness was highest where volume, modularity, and standardisation potential are high.
Technology & Infrastructure	• Re-baseline technology stack when previous choices block scaling; prune architectural complexity; offset each new addition by stripping complexity elsewhere. • Build AI on strong data platform foundation; prefer best-of-suite for core data platform, build selective enabling layers on top. • Explicit separation between data-platform layer and agentic/AI application layer for manageable scaling. • Deliberate build-vs-buy discipline; IT teams tend to overbuild/overengineer; flagged as a real risk. • Incremental legacy modernisation, not full rebuild; in practice the setup remains hybrid, with country-level environments often on-premise and regional AI teams mainly on cloud platforms such as Azure. • Earlier tech decisions can slow AI scaling; bad stack choices push organisations in the wrong direction.
Data	• Group-level CDO and domain owners provide stronger leadership without requiring a heavy centralised framework. • AI must be built on a strong data platform as a required foundation. • Enterprise/data architecture treated as core AI-scaling decision, not secondary concern. • Shift from over-centralised rigid data models to domain-based data management with clear domain ownership. • Earlier centralised models reflected immaturity; previous data stewardship did not scale across business units. • Sweet spot between centralised and local data management not yet found. • Poor data quality is a direct constraint ("garbage in, garbage out"); cross-system connectivity, readability, retrievability are competitive factors for AI scaling.
Culture	• Data-driven mindset embedded as explicit cultural pillar of AI strategy. • AI needs to be seen as driver for creativity and better work, not purely efficiency/headcount; deliberate narrative choice. • Change management run as parallel workstream, not side activity, because AI overlaps with multiple parallel organisational changes. • Employees actively worried about job impact given overlapping changes. • Bottom-up opportunity generation only works after sufficient literacy; cultural readiness consists of: awareness + capability + mindset + management participation.
Cross-domain dependencies	• Processes + People & Organisation + Governance: Business-led portfolio management links prioritisation to process transformation and capability deployment. • Processes + Data + Technology & Infrastructure: Process redesign depends on data quality, connectivity, and appropriate technology stack choices. • People & Organisation + Data + Processes + Governance: Literacy, role standardisation, local maturity, and change management need to develop together.

C.6 Interview 6 compact summary

Table C.6: Compact summary of interview 6

Interview details	Interview 6; Consulting partner in cloud and digital transformation, with deep financial-services experience; 25+ years
Domain	**Compact summary**
Governance	• Centralise AI ownership early to build platform maturity, coherence, and reuse; decentralise/federalise only once business units reach sufficient AI maturity. • Premature decentralisation risks fragmentation, weak reuse, and uneven capability development. • Prioritise use cases via desirability, viability, and feasibility (DVF): user value + business value + executability must all hold before scaling. • Pair central AI enablement with clearer business-side accountability for process and data outcomes. • AI scaling decisions tied to practical business value, not technology enthusiasm.

Domain	Compact summary
People & Organisation	• Build AI capability centrally first, then shift more responsibility to business units once local teams are mature enough to use and develop AI effectively. • AI talent scarcity is a real bottleneck; cannot be solved by external hiring alone. • Reskill existing technical/engineering staff (formal certification, structured upskilling); easier transition than non-technical profiles. • Combine internal reskilling with selective hiring of younger talent with recent AI training. • AI most visibly raises the productivity of strong technical staff. • Short-term effect is "more output with same FTE" rather than large-scale headcount reduction.
Processes	• Start AI on painful, repetitive, high-effort processes for tangible impact; expand to operational core redesign once data/platform foundations are ready. • Redesign AI-first end-to-end, don't just add it on; this is where AI delivers most value. • Quantified case: legal insurance intake automation cut handling time from ~60 min to ~6 min (~10x). • Largest short term gains seen in technical/engineering-heavy work (software production, system conversion). • Support processes are easy entry points; deep core-process automation requires stronger data + platform foundations. • Business impact is output growth without proportional FTE growth, rather than direct headcount cuts.
Technology & Infrastructure	• Three enabling layers required for scaling: data platform, API platform, AI platform. • AI platform is an environment to build, manage, monitor, and govern agents (not just model access). • Build on hyperscaler ecosystems where possible, but recognise vendors don't cover everything; focused build-vs-buy needed. • Platform stack is a prerequisite for AI scaling, without it, organisations stay limited to copilots and support use cases. • Stronger infrastructure expands the range of processes that can be impacted within AI scaling; AI/platform capabilities also support transformation work like migrations.
Data	• Use dedicated data leadership (CDO / head of data & governance) where appropriate. • Data infrastructure is a hard prerequisite for AI automation in operational core processes; focus on data quality, control, usability first. • Strengthen data stewardship and governance as AI scales; move ownership closer to business (process/system owners); only central ownership becomes ineffective. • Data is treated as strategic asset and a competitive differentiator. • Organisations with mature data platforms have a clear AI scaling advantage; without foundations, AI stays limited to copilots and support use cases.
Culture	• Visible leadership use of AI is a stronger adoption driver than formal top-down messaging. • Encourage active employee experimentation and repositioning around AI-enabled ways of working. • Employees embracing AI become noticeably more effective, creating pull effect on peers. • Treat AI adoption as a behavioural/working-method shift, not just a technology rollout.
Cross-domain dependencies	• Technology & Infrastructure + People & Organisation + Governance: AI scaling follows a staged model from early centralisation and platform building toward later decentralisation. • Processes + Data + Technology & Infrastructure: Ambitious process redesign depends on data, API, and AI-platform foundations. • People & Organisation + Technology & Infrastructure: Talent scarcity stalls scaling, while reskilling depends on access to usable platforms for practical learning. • Culture + People & Organisation + Technology & Infrastructure: Visible leadership use of AI helps turn AI investments into actual adoption and productivity gains.

C.7 Interview 7 compact summary

Table C.7: Compact summary of interview 7

Interview details	Interview 7; Senior AI and cloud consulting manager with AI platform and governance focus; 7+ years
Domain	**Compact summary**
Governance	• Split AI governance across roles like CIO, CISO, CDO based on infra/security/data/compliance scope; AI governance treated as combined responsibility, not single-function. • Consider dedicated CDO alongside CIO/CISO in larger organisations. • Give AI governance an explicit organisational place, especially in larger organisations driven by EU AI Act requirements. • Traditional IT governance is insufficient for AI: different infrastructure, compliance, and oversight needs require new mechanisms. • Many initiatives reach PoC/production without robust effectiveness measurement, creating scaling risk despite technical progress. • Measure AI initiatives via business impact, adoption metrics, and operational cost, not just technical delivery. • Embed FinOps-style usage/cost monitoring into AI governance. • AI surfaces pre-existing security weaknesses that previous IT usage had not exposed yet.
People & Organisation	• Raise AI literacy across business as a necessary condition for responsible scaling; build shared ownership rather than central-team-only model. • Use CoEs, champions, and inner-source knowledge sharing to spread capability, increase reuse, and enable cross-BU participation. • Adapt skills to maturity: prompt literacy at low maturity; specialised roles (incl. AI engineering) at higher maturity; data scientists can transition relatively easily into AI engineering work. • Continued investment in cybersecurity roles is one of the most critical and increasing needs in scaling organisations.
Processes	• Platform-based enablement with boilerplates/reusable components reduces duplicate work and accelerates teams. • Add explicit lifecycle management: monitoring, maintenance, improvement, adaptation to dependency changes; AI use cases as living applications, not one-off deliveries. • Pair AI rollout with business process redesign to scale beyond local experiments. • Formalise onboarding/offboarding and structured use-case handling around the AI platform rather than informal adoption. • Enable reuse via marketplace/shared component model. • High near-term process impact is expected in administrative, text-based, visually structured, front-office, customer-interaction, note-taking, and reporting work. • AI-assisted development (e.g., GitHub Copilot) reshaping technical work. • Broad transformation requires onboarding, maintenance, reuse, and iterative improvement processes, not just tools.
Technology & Infrastructure	• Build shared AI platform with onboarding, boilerplates, integrations, and reusable building blocks; key for delivery and reuse. • Marketplace/shared component layer enables cross-organisation reuse of components built for individual use cases. • Leverage existing enterprise vendors moving toward AI (API/MCP-style enablement) to avoid full legacy replacement before scaling. • Be careful with experimental agentic tooling acting under user identity rights, this a serious security risk. • Distinguish use cases relying on existing models vs. exceptional cases needing custom training and heavier ML infra; technical scaling depends more on effective use of existing models rather than custom training. • Most AI workloads run in the cloud, while some on-prem persists due to geopolitical and sovereignty concerns. • Treat agent orchestration as an identity access management problem: every agent gets specific roles and permissions.
Data	• Build data foundation enabling flexible but controlled isolation; agents access only data needed for their use case. • Apply both RBAC and ABAC to AI agents and their data interactions, this is essential for safe scalable agent use. • Strengthen in-process data quality checks rather than treating quality as afterthought; quality remains critical even with more automation. • Shared data platform preferred; where centralisation not possible, controlled access across multiple sources with clear permissions. • Pursue single source of truth via master data management/golden records, this is crucial but difficult; fragmented/legacy data landscapes are a big challenge. • Use layered/medallion-style architecture, distinguishing shared organisational data from operational/use-case-specific data. • Foundational data architecture work is partly use-case agnostic: some data structures must improve regardless of future use cases.

Domain	Compact summary
Culture	• Combine top-down direction with bottom-up participation; a purely top-down push is insufficient. • Raise AI literacy across the business with shared ownership. • Actively include less AI-literate employees to reduce risk of alienation. • Frame AI adoption as evolutionary organisational change, not one-time revolution; success depends on combining strategic direction with participation and learning. • Build innovation culture supported by guidance, literacy, and practical structures; journey depends on starting maturity and gap to desired state.
Cross-domain dependencies	• Technology & Infrastructure + Governance + People & Organisation: AI scaling combines a shared platform, governance, and broad business contribution. • People & Organisation + Processes + Technology & Infrastructure: Inner sourcing enables teams to build, share, and reuse AI components across functions and departments. • Technology & Infrastructure + Data + Governance: AI requires stronger guardrails than traditional IT, especially around compliance, security, cost control, and identity and access management. • Data + Technology & Infrastructure: Shared AI infrastructure has to support reuse and AI-ready access control.

C.8 Interview 8 compact summary

Table C.8: Compact summary of interview 8

Interview details	Interview 8; Senior consultant in AI and data strategy, adoption, and implementation; 7+ years
Domain	**Compact summary**
Governance	• Establish central AI governance with joint CIO/CDO/CISO oversight; ad-hoc ownership is no longer good enough since AI spans the whole organisation. • Central AI/data governance functions and mindset are prerequisite; without them, organisations remain stuck in experimentation. • Keep (financial) C-level accountability but assign dedicated operational AI governance roles, since executives cannot manage day-to-day in practice. • Measure AI value via mix of ROI, time savings, A/B tests, adoption metrics, quality gains, and ability to handle more/complex work. • Treat AI governance as different from previous software rollouts due to broader user base, risk, and compliance exposure.
People & Organisation	• Build AI CoE to concentrate scarce expertise; combine with broad workforce AI training since AI will affect nearly everyone. • Two-sided talent challenge: educate entire workforce while hiring scarce specialists. • Broad capability building is necessary but hard to justify with short-term ROI logic, creating a scaling dilemma which is hard to accept. • Reskilling existing data analysts/scientists toward ML/AI is possible but hard since they are already fully occupied. • Hiring AI-talent often easier than reskilling but expensive and hard to justify via immediate ROI. • A translator/product-owner role between business and AI tooling is a recurring organisational need. • Use early adopters to create momentum; broader workforce change remains much harder.
Processes	• Apply AI selectively where real business value exists and evaluate via A/B tests comparing work with and without AI (agents). • Strongest process impact in complex/cognitively demanding work, not only trivial automation (e.g., procurement case: AI freed capacity to target new market segment). • Use AI to redesign processes, not just automate existing steps. • AI acts as catalyst for cross-departmental process and data alignment. • Many organisations lack documented business processes which is a major barrier to identifying viable AI use cases. • Operating model must allow space for team-level experimentation. • Templates, work instructions, and process descriptions help translate workflows into AI opportunities. • Organisations that do not invest early risk leaving their workforce behind relative to future ways of working.
Technology & Infrastructure	• Larger organisations are cautious with new AI vendors due to short reliability histories and unclear business dependency risk, so most stay with current vendors. • Recent previous tech investments slow AI tool adoption but organisations are reluctant to abandon recently deployed tools. • Treat AI tooling as enterprise-level decisions, not team-level software picks, due to broad user impact and higher uncertainty. • Use pilots to validate reliability before scaling.

Domain	Compact summary
Data	• Treat data governance as formal role with ownership, policy, and execution; not an add-on responsibility. • Strengthen top-down data leadership (CDO or equivalent). • Weak data structure, poor metadata, unclear naming, and weak versioning are major practical barriers to AI working reliably. • Use department-level dashboards, steering discussions, and accountability (with consequences) to enforce data discipline. • Mandatory metadata, document structure conventions, and departmental quality reporting recommended. • Data is a cross-cutting prerequisite; without discipline, AI cannot create sustainable value at scale. • Data governance framed as discipline and behavioural change, not just technical fix.
Culture	• Broad employee use without risk awareness creates organisation-wide exposure. • Broad AI literacy is essential since employees will use AI without understanding risks. • Employee reactions highly mixed: enthusiasm to outright refusal; resistance tied to job loss fears, expertise status loss, and discomfort with new tech. • Change management is a major scaling challenge, not a side issue. • Leadership mindset critical; leaders indifferent to AI are a direct barrier to AI culture. • Formal rollout without cultural support likely to fail. • Start with willing/curious employees to pull others along.
Cross-domain dependencies	• Data + Governance: AI scaling stalls where data governance is immature or where there is no clear (central) data ownership. • Data + Technology & Infrastructure: Poor metadata and information management reduce the reliability of AI retrieval. • People & Organisation + Processes: Missing translator roles and poorly documented processes limit the identification and redesign of AI use cases.

C.9 Interview 9 compact summary

Table C.9: Compact summary of interview 9

Interview details	Interview 9; Senior consultant in responsible AI and AI governance; 7+ years
Domain	**Compact summary**
Governance	• Start scaling with central AI team/CoE to concentrate scarce expertise, prioritisation, and responsible-use guidance; decentralise only after maturity is built. • Premature decentralisation leads to "shadow AI" risk, amplified by local procurement autonomy below thresholds; weak governance leads to fragmented tooling and policy misalignment. • Define explicit AI-lifecycle governance: roles, decision moments, risk assessments; escalate higher risk use cases to senior/directive level. • Business owner holds decision rights; CIO/CTO, CISO, CDO, risk involved per risk level; risk acceptance is ultimately a business choice, not only control-function. • Apply three-lines-of-defence to AI (1: business operates, 2: risk/security/data leadership sets guardrails, 3: assures through audits). • Treat responsible AI (policy, risk, security, robustness, bias) as a practical precondition for scaling, not an ethics nice-to-have. • Measure via both value (business case, efficiency, adoption, training effects) and controls (incl. token cost); benchmark against pre-AI baseline. • Use a staged governance logic for AI initiatives, moving from proof of concept to pilot and broader rollout only after required checks, controls, and value proof are done.
People & Organisation	• GenAI affects far broader employee base than traditional AI; capability-building must be organisation-wide, not specialist-only. • Central CoE early and evolve to distributed capability with maturity. • Hire AI/ML engineers and AI product roles; strengthen data owners/stewards as they are highly relevant. • Continuous awareness, training, change management; not one-off literacy interventions. • Resistance driven by fear of job loss, status loss, or not knowing how to use AI; early disappointment after weak first try with AI is a key adoption barrier. • Risk: heavy AI use by juniors may erode deep domain knowledge and critical thinking over time; upskilling must protect business-domain understanding.

Domain	Compact summary
Processes	• Staged rollout: PoC (test parity vs. human) → pilot (shadow-run/A/B vs. human-only) → broader rollout after controls completed; pilots help the business case and management approval. • Best early fit: well-documented, text-heavy, repetitive, rule-based processes (document extraction, case prep, standardised checks, simple payments/claims). • Customer service AI seen as promising but sometimes risky. • Agentic AI is limited to data-rich, documented, repetitive contexts. • For higher ambition: end-to-end redesign of tasks, roles, handovers, and workflow; not just AI inserted into unchanged process. (transformation vs. incremental improvements) • Current AI process impact is already real, but still limited in scale.
Technology & Infrastructure	• Shared AI platform as core scaling enabler: pre-arranges LLM access, data connections, security, identity access management; set up once vs. per use case. • Each application gets its own controlled sub-environment within platform to avoid uncontrolled shared space. • Prefer stable large-vendor ecosystems (lower continuity risk); avoid proliferation of niche tools; niche orchestration/agentic tools are experimental but not essential for scaling. • Cloud as default (scalability, cost, maintainability); on-prem only for highly sensitive or strict public-sector contexts. • Legacy is not blocked by age but by lack of connectors/APIs and unusable data quality which reframes legacy as access and data issue.
Data	• Data as strategic asset; AI scaling increases existing weaknesses in quality, definitions, and control. • Strengthen (not bypass) traditional data governance: data owners decide use/purpose incl. AI; stewards drive operational quality; corrections fed back to source, not patched downstream. • Push toward single source of truth and shared data definitions; fragmented landscapes cause contradictory reports that AI then amplifies. • Centralised vs. distributed data environments chosen by context (access, ownership, security trade-offs). • Legacy is not blocked by age but by lack of connectors/APIs and unusable data quality which reframes legacy as access and data issue.
Culture	• AI adoption is primarily a change-management challenge rather than just tool deployment; use ongoing training, webinars, hackathons, and leadership signalling to normalise AI. • Combine experimentation space with guidance to prevent unsafe or uninformed AI use. • Poorly reviewed AI output can create more work than it saves when reliability judgment is weak. • Weaker innovation culture / lower urgency leads to slower translation of AI interest into adoption. • Adoption depth varies widely across organisations even when surface-level training exists.
Cross-domain dependencies	• Technology & Infrastructure + People & Organisation + Governance: AI scaling works best when central capability and governance are established first before decentralisation. • Governance + Data + Technology & Infrastructure + Processes: Responsible AI is a scaling condition rather than a separate concern. • Technology & Infrastructure + Data + Governance: A shared AI platform centralises controls, secure access to models and data, and reusable infrastructure. • Processes + Data: Process value depends on documented workflows and data readiness.

C.10 Interview 10 compact summary

Table C.10: Compact summary of interview 10

Interview details	Interview 10; Consulting director and former Big Tech engineering lead; 20+ years
Domain	**Compact summary**
Governance	• Treat every AI initiative as a product (not project/PoC) with both business and IT ownership; maintain catalogue of AI products with ownership, boundaries and intended use. • Extend existing data and product governance to AI instead of inventing parallel structures; involve CIO, CTO, CISO, other C-suite roles and risk/compliance accordingly. • Start initiatives from explicit business case and value hypothesis, not tech experimentation. • Measure AI initiatives through explicit business cases and operational KPIs, such as percentage of work handled by AI and reductions in handovers to human agents. • Without product-style ownership, AI initiatives stall at the PoC stage. • AI governance is strongest when combining ownership, value orientation, traceability, and compliance.

Domain	Compact summary
People & Organisation	• Start with CoE and move toward broader enablement of business, IT, and other domains as AI matures; use cross-functional 'chapters' for co-creating with business and tech. • Critical roles for scaling: AI translators, AI engineers, platform engineers, and AI product owners. • Mandate an AI product owner to define vision, coordinate stakeholders, and link value, risk, and delivery. • AI absorbs repetitive work; people shift toward complex/exception-heavy tasks: fewer people can do more work, but concentration of only hard work creates new pressure. • The current focus is on 'AI-augmentation', but this is expected to evolve toward an 'AI-native' operating model as AI matures.
Processes	• Start by embedding AI into existing processes (augmentation); move to deeper redesign (including removing process assumptions) as AI matures toward agentic operation. • AI can shift where processes start (e.g., proactive issue detection) rather than only speeding existing steps. • Use A/B comparisons (human-only vs. AI-supported) to assess process impact; track hard metrics like % volume handled by AI and reduction in handovers to human agents (e.g., customer service). • Adoption, change management, and value tracking must be integral to process redesign, not afterthoughts.
Technology & Infrastructure	• Put in place an AI platform built on a data platform supporting structured and unstructured data, plus a clear integration layer (APIs, MCP-style interfaces) for agents to reach internal/external services. • Three-layer model: data stack + AI platform stack + integration mechanisms; AI scaling depends on all three, not the model layer alone. • Cloud preferred where possible; some organisations have on-prem for control. • Legacy decisions made case-by-case based on process fit and integration support, not as uniform problem. • Build-vs-buy: AI expands what can be built in-house, but external tooling remains very relevant. • Integration becomes critical once agents must access multiple systems, services, or other agents.
Data	• Data governance is foundational for AI and must be extended; requires clear data accountability, catalogs, lineage, and discoverability. • Unstructured data governance becomes a much more serious topic with AI than in previous data setups. • Make shared unstructured data available in governed forms (e.g., common vector stores) to avoid each team re-vectorizing the same source material. • Both centralised and decentralised (e.g., function-specific data products, mesh-like) architectures can work, but only with strong governance, lineage, and ownership. Without it, both fail. • Strongest principle: avoid uncontrolled duplication, preserve single source of truth.
Culture	• Change management and embedding AI into daily work are described as the main bottleneck, more than the technology itself. • Combine top-down vision with bottom-up learning/innovation across functions; top-down alone is not enough. • Invest in training, accessibility, shared best practices; first impressions and usability strongly affect adoption. • Difficulty proving ROI can stall adoption when AI use cannot yet be tied to hard business value.
Cross-domain dependencies	• People & Organisation + Governance: AI product owners and translators connect the business, process redesign, risk, and delivery. • Technology & Infrastructure + Data: The value of the AI platform depends on underlying data-platform quality, integration, and governed access to structured and unstructured data. • Data + Governance: Data architecture choices matter less than the strength of data governance in making AI scalable. • Processes + People & Organisation + Culture: Scaling typically moves from augmentation toward redesign, which requires adoption and role change.

C.11 Interview 11 compact summary

Table C.11: Compact summary of interview 11

Interview details	Interview 11; Chief of Staff in a tech-focused company, with tech strategy consulting experience; tech product / startup context; 8+ years
Domain	**Compact summary**
Governance	• Treat AI as integral way of working, not isolated pilots; a 'project mindset' leads to failure. • Place AI ownership under existing senior tech leadership (CTO/CIO) with explicit mandate to drive enterprise scaling. • Centralise first (security, compliance, cost, guardrails), decentralise later as maturity grows, especially critical in large organisations where local units can independently spend on AI tools. • Require positive business case for major platform/integration investments. • Build cost visibility into AI consumption so business users see financial consequences of model usage. • Central governance reduces shadow IT, duplicated work, fragmented experimentation, and uncontrolled spend.
People & Organisation	• Start with centralised talent teams, and move to decentralise later. • Add dedicated AI roles (process-optimisation leads, AI engineers, internal AI capability teams); role mix scales with organisation size and ambition. • Embed AI roles inside existing IT as an additional capability unit. • Process-knowledgeable employees are highest-value identifiers of automation opportunities; they know where time is actually lost. • Efficiency gains should be redirected to growth and customer-facing work rather than headcount cuts (more work and value with the same number of people) • AI will change jobs; reskilling and adaptation are unavoidable.
Processes	• Prioritise text-heavy, repetitive, administrative processes first (e.g., reporting, communication drafting, meeting follow-up, support, finance, and procurement). • Strongest near-term value: AI completes a large share of the work, human does final review. • Substantial efficiency gains are seen in recurring reporting/admin tasks. • Deeper transformative potential in large-scale operational domains (banking, customer service); most other organisations are currently at augmentation stage. • Let employees shortlist disproportionately time-consuming recurring tasks. • Gather input from employees with strong process knowledge, because they know where time is lost and where AI can help.
Technology & Infrastructure	• AI ambition impacts infrastructure: surface optimisation needs little change; enterprise-integral AI requires embedding into the core IT stack. • Internal AI platform / cloud-like environment needed for broad, secure organisational rollout. • Connect AI to internal data via APIs with structured, restricted, process-specific access. • Legacy fragmentation is a direct blocker to becoming AI-native because of disconnected sources. • Build-vs-buy is pragmatic: buy major platforms, but you can run open source models internally if the infra/engineering skills are there. • Strategic risk in dependence on a few large foreign cloud/model providers; on-prem/local alternatives worth considering for sovereignty.
Data	• Poor data quality and unclear data organisation is a direct barrier to AI scaling. • Build strong data models and map data flows before scaling AI internally. • Data governance without the actual capability to store, structure, and access data is ineffective; organisations must enable the desired behaviour next to enforcing policy. • Work toward a structured and reliable single source of truth feeding multiple use cases. • Centralise data first where feasible; decentralise only where organisational complexity requires it. • Process-specific data selection and access restrictions remain essential even when data is broadly available for AI. • "Garbage in, garbage out" risk is a foundational concern.
Culture	• Fear of job loss is the primary cultural barrier, especially in large organisations. • Frame AI as freeing time for higher-value work, not purely FTE reduction. • Use change champions and internal advocates to translate adoption into practice. • Encourage employees to proactively find their own AI-improvable tasks. • Adoption improves when AI is framed as enabler, while acknowledging legitimate role-change fears.
Cross-domain dependencies	• Culture + Governance: Framing AI as FTE reduction stalls employee adoption. • Data + Technology & Infrastructure: AI scaling depends on strong data models, data flows, storage, and integrated access. • Technology & Infrastructure + Governance: AI ambition determines whether AI is used as an add-on or embedded in the core IT stack. • People & Organisation + Processes: Employees with strong process knowledge are crucial for translating AI potential into real implementation.

C.12 Interview 12 compact summary

Table C.12: Compact summary of interview 12

Interview details	Interview 12; Senior advisor in AI transformation and operating model redesign; 25+ years
Domain	**Compact summary**
Governance	• Avoid an innovation-department-only model; leads to "PoC hell" disconnected from operations. • Treat AI as cross-functional transformation requiring joint business/IT/compliance/process ownership. • Active board-level sponsorship essential; bottom-up cannot resolve difficult cross-domain trade-offs. • Establish responsible-AI policy, decision routes, and technical boundary conditions before broad rollout, so teams know what is allowed and how difficult cases are handled. • Staged governance: centralise early (scarce knowledge, methods, prioritisation), federate/matrix later as maturity grows. • Prioritise use cases explicitly on value + feasibility and start with lower-risk cases where data is usable. • Rapid prototyping as governance tool: evidence-based prioritisation, unlocks executive support and funding. • Budget is rarely the main bottleneck in large organisations; executive priority and attention are scarcer than money. • CDO role alone is insufficient; needs business + IT + compliance + executive co-ownership. • EU regulation + tech/AI change speed make AI governance structurally hard.
People & Organisation	• Cross-functional programme teams over heavy hierarchy; central AI capability/CoE early to pool scarce skills. • Core delivery team: AI engineer, data engineer, business consultant, PO/PM, compliance; add infra/cloud where needed. • Shift toward decentralised/matrix structures only after central maturity is built; premature decentralisation fails. • AI scaling can reduce workforce needs and process steps (e.g. from 100 to 80 people; from 5 to 4 steps). • Real losers exist: some employees unable/unwilling to adapt to new skills and roles. • Employees need both practical AI literacy and realistic expectations of AI imperfection/unreliability.
Processes	• Staged delivery: prioritise → prototype → validate → build. • Validate operational fit with frontline employees, not just technical feasibility (designers may like ideas users reject). • Highest-fit areas now: customer service, marketing/communications, software development, legal/knowledge work. • Customer service is clearest current target (transcripts, summarisation, progressive automation delivering visible value). • Distinguish lightweight support tools from transformative use cases that restructure roles. • When AI absorbs simple cases (e.g., support), remaining human work shifts to complex cases → fewer low-skill, more high-skill staff. • Treat AI rollout as classical change management on steps, tasks, skills, and role composition. • Emerging advanced use cases: R&D, tailor-made services, education, and legal.
Technology & Infrastructure	• Extend MLOps with semantic testing/monitoring: static technical checks aren't enough since models, data, and context drift. • Automated testing is becoming feasible and necessary across pre- and post-deployment. • API-based integration architecture, but model choice, sovereignty, and cost make infra more layered than before. • Layered model strategy: frontier models only where needed; smaller/local models for routine tasks to cut cost and dependency. • Hybrid infra decisions driven by sovereignty, cost, task complexity; no universal setup. • Automate compliance/approval checks; manual governance is a major competitiveness drag in regulated EU context.
Data	• Use a hybrid data model: keep platform, a source of truth, and shared standards central, while letting local domains manage and improve their own data quality. • Strengthen governance for unstructured data, this is newly strategic for AI, previously neglected. • Avoid both negligence and perfectionism; start where data is clean enough rather than waiting for enterprise-wide perfection. • LLM-suitability is not the same as traditional data quality: human-readable data may still be unfit for AI consumption. • Hybrid model: central platform/single source of truth/standards + decentralised domain quality ownership. • Decentralised units require discipline/accountability or data quality erodes. • Explore knowledge-graph-like structures to improve LLM interpretation of organisational knowledge.

Domain	Compact summary
Culture	• Build experimentation and controlled-failure culture; psychological safety is precondition for AI learning. • AI scaling requires large-scale change management rather than a pure technology rollout. • Punishing failure too quickly destroys meaningful AI learning. • Balance innovation space with clear process targets and business direction. • Expect resistance and quiet non-adoption: employees can undermine intended new ways of working. • AI change produces real winners and losers, not just upskilling opportunities.
Cross-domain dependencies	• People & Organisation + Data + Governance: Early centralisation creates the control and the capability basis needed for later federation. • Processes + Culture + Governance: Rapid prototyping leads to executive commitment by demonstrating value and creating momentum for further learning and scaling. • Processes + People & Organisation + Culture: Workforce impact from process redesign can trigger resistance when automation removes steps or roles. • Technology & Infrastructure + Data + Governance: Semantic and context drift and changing data conditions require monitoring as part of the (compliant) AI operation.

C.13 Interview 13 compact summary

Table C.13: Compact summary of interview 13

Interview details	Interview 13; Director in AI transformation and cloud advisory; 15+ years
Domain	**Compact summary**
Governance	• Federated model preferred over full decentralisation due to challenges in efficiency, reuse, and governance complexity. • Board sponsorship and investment required (CIO, CFO, and increasingly data leadership); involve compliance but not too early or AI stalls. • Retain ITSM but shift execution toward business as AI enables self-service development and problem solving; AI increases pressure because users expect to build without waiting on central IT. • Portfolio management uses explicit business-value weighting; finance is involved through headcount reduction/avoidance and revenue impact. • Use an InnerSource-style model to involve employees across the organisation in AI development rather than keeping it centralised in IT. • AI benefits management is key, but often still lacking because of solutions getting built without value tracking. • Stronger economic steering is needed with benefits management, model selection, token usage monitoring, value tracking, and model routing (context-based instead of "best model for all tasks"). • AI value is not automatic, but requires disciplined choices on use cases, models, and scale.
People & Organisation	• Early AI core including: technical experts + ambitious talent + clear business representation + policy/ethics/responsible-AI expertise (especially where regulation matters). • Start with a CoE to make early mistakes, form policy, and build shared learning, while recognizing that digital-native organisations may structure AI differently from the start. • Combine internal upskilling with external hiring/contracting; both are needed. • Broad technical skill development is required across a much wider employee base than before. • The shift from classical AI/IT to modern AI requires new skills and a different way of building applications. • Motivation and curiosity are more important than skill level; less skilled, but motivated employees are easier to upskill than skilled AI resisters. • Have multidisciplinary teams building and maintaining initiatives end-to-end to reduce dependency on specialised teams. • Prioritise engineers who can deliver secure, scalable, enterprise-ready solutions, not just experimentation/prototyping. • Hardest roles to hire: translators who convert AI into safe, scalable, broad organisational value, not "hobbyists". • Organisations with strong existing technical capability scale faster by combining internal talent with external capacity. • Clear emerging split between AI experimenters and avoiders; AI adoption is unavoidable and increasingly central to work.

Domain	Compact summary
Processes	• Apply AI to both support and core processes, neither should be assumed more important. • Example target processes: HR on/offboarding, finance booking, IT support, software development, claims handling, order processing, and warehouse picking. • Redesign from first principles, not by adding AI onto existing steps; the biggest gains come from rethinking the entire process. • Start redesign by asking what the process should look like with minimal human steps, then reintroduce humans only where needed. • Use AI to remove unnecessary steps and turn fragmented flows into direct end-to-end flows. • AI is not only relevant for repetitive work but (increasingly) for agentic and more complex tasks too. • There is a shift toward orchestration-based work with AI/agentic extending RPA automation to a wider set of business activities. • AI makes core processes increasingly automatable through a combination of language, prediction, and agentic coordination. • Radically redesigned processes need far fewer human interventions.
Technology & Infrastructure	• Build on existing models and do not train models from scratch. • Two main platform types: hyperscaler AI stacks vs. containerised agentic stacks; hybrid setups combine both for fallback and reduced dependency. • Hyperscalers are more convenient and containerised/local setups come with higher maintenance cost, complexity, and skill requirements. • Heavy dependency on a foreign platform can be strategic business risk under geopolitical change. • Model selection is key: sovereignty, maintainability, copyright, security, and cost issues are at play. • Model routing so system selects right model based on the context. • Integration is a top technical concern: MCP-style patterns, agent-to-agent communication, integration layer connecting agents to enterprise systems. • Explicit identity-management rules for agents: acting under human identities vs. separate non-human identities with own permissions. • Prepare for an app-less, conversational future where agents do work traditionally done in enterprise applications; this will lead to deconstruction of apps. • Shift from classic software building toward composing applications from agentic building blocks. • Avoid dependence on large external providers where geopolitical risks make local alternatives valuable.
Data	• Use existing enterprise data sources; AI does not require entirely new data landscapes. • Many AI data challenges are primarily access and integration problems, not data creation. • Modern layered architectures (e.g. medallion-style), combined with clear data owners and governance are becoming more common with AI. • Follow data minimisation principles, especially for agent context, memory, and personalisation. • Govern (agent) memory and operational data stores: what may be stored, how long, and when. • Distinguish persistent enterprise data vs. temporary/interaction-level agent memory. • Integration/access layers expose the right data safely to the right agent. • Data, technology, and infrastructure treated together as one enabling foundation. • Ownership and lifecycle rules for memory-like data grow in importance.
Culture	• Build innovation culture where AI is actively explored, not treated as optional. • Early visible success is critical for sustaining executive support and broader organisational momentum. • Employees should work with AI themselves rather than leaving experimentation to a small specialist group. • Ambition and intrinsic motivation are key conditions for successful AI adoption. • Address fear of job loss and resistance to change as cultural barriers. • Visible AI-adoption through the hierarchy normalises new ways of working. • Growing divide between AI adopters and avoiders (“winners and losers”); early engagers gain structural advantage. • Resistance stems from fear, self-doubt, or concern about replacement, not from lack of capability. • AI scaling requires broad cultural shift to where a larger part of the organisations becomes technically capable and willing to work with AI.
Cross-domain dependencies	• People & Organisation + Culture + Governance: A CoE-first, federated, InnerSource-style model supports self-service while maintaining reuse and control. • Technology & Infrastructure + Governance: Economic steering through model routing, usage monitoring, and business-case discipline is needed to keep AI scalable. • Processes + People & Organisation + Culture: Radical process redesign changes roles substantially, creates resistance, and requires people to learn new ways of working. • Data + Technology & Infrastructure + Governance: Agent identity, memory, and integration access require governance rules that go beyond traditional controls.

C.14 Interview 14 compact summary

Table C.14: Compact summary of interview 14

Interview details	Interview 14; Independent AI advisor and former consulting partner focused on large-scale transformation in energy and industrial sectors; 25+ years
Domain	**Compact summary**
Governance	• Anchor AI in business problems and friction, not tech; a technology-first approach is a big cause of AI value failure. • Use a portfolio funnel with explicit ownership that has the power to scale, change direction, or end initiatives; otherwise, pilot fatigue comes in. • Treat AI as business transformation, not a technology project. • Use AI only when it has a clear added business value, some boards want AI for its own sake and 'for no reason'. • Do not let compliance/audit define (strategic) direction, but use them as gatekeepers in regulated settings. • Use platform/vendor choices (including sovereign options) to work around security and regulatory constraints where needed. • Apply an explicit three-level lens: task/process value, delivery redesign, business-model change; this prevents overestimating low-level wins. • Run (low level) efficiency and deeper redesign in parallel; value creation does not follow a linear staircase. • Organisations confuse task automation with strategic transformation.
People & Organisation	• Build capability around people who translate AI into visible day-to-day business value; internal champions drive adoption with small but impactful wins. • Combine internal capability building with external expertise where senior engineering/transformation talent is scarce and be selective in hiring. • Do not hollow out junior roles because this weakens long-term learning, scaling capacity, and future capability pipeline. • AI adoption does not automatically reduce headcount; it also creates new work in checking, managing, and improving outputs. • Structural choice (CoE vs. federated) is based on culture/maturity and there is no universal standard. • Value realisation still hinges on strong human capability even when AI is technically available. • Broad practical AI engagement across employees is essential; those who never learn to work with AI risk becoming less relevant. • Use visible small wins to build momentum across teams and reduce scepticism.
Processes	• Select processes by locating friction: waiting time, handoffs, rework, inconsistent data, and cross-silo gaps. • Redesign or remove broken processes rather than automating them as is. • Prioritise cross-silo problems over local optimisations; the biggest value is found between functions, not in single tasks. • Prioritise by economic weight in cost base/value chain (e.g., supply chain, core operations), not just highest %-ROI use case and look beyond familiar starting points like customer contact. • Treat Level 1 task automation as baseline; the real operating-model gains come from Level 2 delivery redesign. • Use incremental high-impact wins, but connect them to a broader process-redesign logic. • Most pilots fail because they stay too narrow, local, or disconnected from the broader (organisational) context. • The biggest value comes from changing how work is structured, not from speeding up existing tasks. • Large enhancement-only AI use cases (e.g., improved asset monitoring) can deliver major gains without being new. • Do not over-automate customer-facing interactions where practical fieldwork or human judgment still matters, because this can create ineffective "chat-on-chat" loops and poor communication.
Technology & Infrastructure	• Prefer platforms/vendor solutions over custom-built AI for critical use cases; build-vs-buy is a risk-management issue. • In critical/regulated contexts, custom internal AI is risky due to support, certification, and accountability issues. • Expect the tech stack to change a lot, but not yet; avoid premature large-scale re-platforming. • Use certified, supported, secure, (sovereign) solutions to manage liability, compliance, and data location. • Strong system/data integrations are prerequisite for higher-level AI value. • AI itself can act as a flexible integration layer, reducing customisation load on heavily customised enterprise platforms. • Introduce AI oversight tooling for visibility into performance, cost, and usage footprint. • Allow citizen development within guardrails to contain operational/financial risk. • Embedding AI into workable process and platforms is the main challenge, not technology. • Access AI capacity through vendor token/consumption models rather than building foundational capacity in-house.

Domain	Compact summary
Data	• Inconsistent, fragmented data and missing shared definitions/single source of truth is a big barrier to AI scaling. • AI does not fix broken data foundations; it exposes and amplifies them. • Define data products and shared concepts so stakeholders work with the same meaning. • Clarify who and which agents can access which data and under what conditions. • Balance sovereignty, control, and flexibility in data governance rather than waiting for a perfect architecture before acting. • Tackle operational "data mess" directly by standardising across business units. • The hard work is aligning existing data foundations, not inventing new AI techniques.
Culture	• Build a culture where AI experimentation is encouraged and failures are not immediately punished; safe experimentation is needed for learning fast enough to scale. • Counter AI hype with realism about where AI can and cannot create meaningful value. • Treat AI as an amplifier of existing strengths and weaknesses, not a 'magic fix'. • Encourage employees to engage practically with AI; adaptation matters more than just having opinions. • Adopt AI pragmatically; not every organisation should aim for business-model reinvention. • Culture is the main determinant of whether AI strategy can actually be executed. • Transformation narratives are often inflated; only a small minority of companies are positioned for true business-model reinvention. For most, the realistic goal is becoming adaptive and value-focused, not revolutionary.
Cross-domain dependencies	• Processes + Culture + Governance: Starting from business friction helps avoid pilot fatigue and supports AI ambition in operational value. • Processes + Technology & Infrastructure + Governance: Prioritisation, redesign ambition, and platform choices are dependent on each other rather than separate decisions. • Processes + People & Organisation + Governance: Cross-silo redesign requires ownership and coordination beyond individual functions. • People & Organisation + Culture + Processes: Preserving junior roles matters for maintaining learning capacity during AI-driven process redesign.

C.15 Interview 15 compact summary

Table C.15: Compact summary of interview 15

Interview details	Interview 15; Director in AI, innovation, and workforce transformation; 15+ years
Domain	**Compact summary**
Governance	• Anchor AI investment in business strategy and long-term value, not vendor hype, demo-driven enthusiasm, or short-term efficiency cuts. • Run AI ideas through a funnel filtered on strategic relevance, expected value, feasibility, and organisational fit; discipline here prevents "attractive patches" on deeper problems. • Treat AI governance as a redesign problem, not an extension of existing sign-off: employees now delegate work to AI agents and effectively act as managers of those agents. • Define explicitly who is accountable when AI contributes to errors in deliverables, decisions, or client-facing outputs; without this, scaling stalls at the friction point. • Expect no universal AI governance model; design follows where responsibilities actually sit in the organisation. • Governance gaps are sharpest where human sign-off is mature but AI-supported sign-off is undefined. • Balance critical judgment with patience: AI vendor claims are partly marketing, but products do genuinely improve over time.
People & Organisation	• Pair domain specialists with AI specialists; do not expect domain experts to convert into full AI specialists. • Build 18–24 month reskilling pathways, e.g., starting from a ~50% role match; this is a concrete alternative to talent loss. • Redesign reward systems when AI shifts responsibility or task mix; adding accountability without reward doesn't work. • Treat adaptability, learning willingness, and future-orientation as the scarce resource, not technical skill. • Buy in technical skill via vendors and externals if needed; internal adaptability cannot be bought. • Support middle managers taking on AI responsibility; without support, middle management becomes the bottleneck rather than the accelerator. • Do not assume that bought-in tech reduces the people challenge, employees still have to learn to work with it. • Beware executive-sold demos: tools embraced at the top often misfire on the work floor and breed resistance in the layers in between.

Domain	Compact summary
Processes	• Let the business, not IT, decide whether AI optimises the current core process or supports inventing a new one; the choice follows market position. • If competitors are overtaking you in a core process, rethink it; if margins/cost still improve, optimizing it is just as valid. • Apply an explicit three-level value lens: (1) individual productivity, (2) team/process productivity, (3) organisation-wide end-to-end across connected processes. • Treat individual productivity tools (e.g. Copilot) as a baseline layer employees now expect, even if direct ROI is weaker. • Locate the real value at Level 3: agent-based end-to-end workflows designed around the final goal, not around current intermediate steps and assumptions. • Avoid single-step patches when broader redesign would create stronger strategic value. • Redesign greenfield from desired start point, end point, and human-in-the-loop requirements; explicitly decide which human skills to retain. • Play to win, not to play not to lose: short-term FTE-cutting is short-sighted versus long-term margin, competitiveness, and strategic renewal. • End-to-end redesign comes with more risk but is the only way to reach real transformation.
Technology & Infrastructure	• Do not assume vendor solutions remove internal architectural choices; define the problem before selecting the tool. • Evaluate vendors on cost, stack fit, long-term pricing behaviour, and EU sovereignty; sovereignty can act as a hard knock-out criterion for some organisations. • Match model to problem: expensive/token-heavy models are a suboptimal choice for simple problems. • Test whether legacy can expose the data and functionality AI needs before layering agents on top. • Repeatedly stacking AI on unsuitable legacy produces suboptimal architectures, not durable transformation; redesign the stack where needed. • Treat access to underlying systems and data as the deeper issue, not whether agents tolerate awkward interfaces. • Vendor lock-in becomes a strategic risk when only externals still understand critical legacy, especially in public-sector or sensitive processes. • Combine patience with critical judgment on vendor roadmaps: products improve, but don't believe every claim.
Data	• Treat data as a foundational precondition; AI is fundamentally a prediction mechanism and its quality is dependent on its data. • Verify AI use cases can work with enough relevant data before treating outputs as reliable. • Estimate expected error rates when data is incomplete; weak data creates false confidence and turns wrong outputs into perceived truth. • Distinguish directionally correct outputs from reliable outputs; the first is not sufficient for high-stakes decisions. • Ground models in organisational context, including internal standards, local language, ethics, and communication style, because generic pretrained models miss it. • Avoid relying on generic models for judgment-heavy processes (e.g. objections handling); this requires grounding first (e.g., prompt engineering, RAG or more intensive methods).
Culture	• Reward adaptability, continuous learning, and orientation toward the future organisation, not preservation of today's roles. • Frame AI partly as a response to structural workforce pressure, such as an ageing workforce and retirement-driven knowledge loss, not only as a cost-reduction initiative. • Allocate time and support for learning, or adoption slows even when the technology is fully available. • Improve psychological safety by clarifying accountability and acceptable AI use. • Treat AI as a strategic means to business goals, not a magical solution. • Don't go with a simple FTE-substitution: as long as humans stay in the loop, the same tool produces different value across users; historically, technology reshaped work more than eliminate it.
Cross-domain dependencies	• People & Organisation + Culture + Governance: Unclear accountability for AI-supported work reduces psychological safety and slows adoption. • Processes + Governance: End-to-end AI redesign requires new accountability structures and sign-off mechanisms beyond traditional governance. • Data + Technology & Infrastructure + Processes: Weak data, limited legacy access, and vendor lock-in stall process redesign. • People & Organisation + Processes + Culture: Process gains remain dependent on human capability and adoption, which makes fixed ROI assumptions unreliable.

Appendix D

Full-size framework figures

D.1 Full-size condensed framework figure

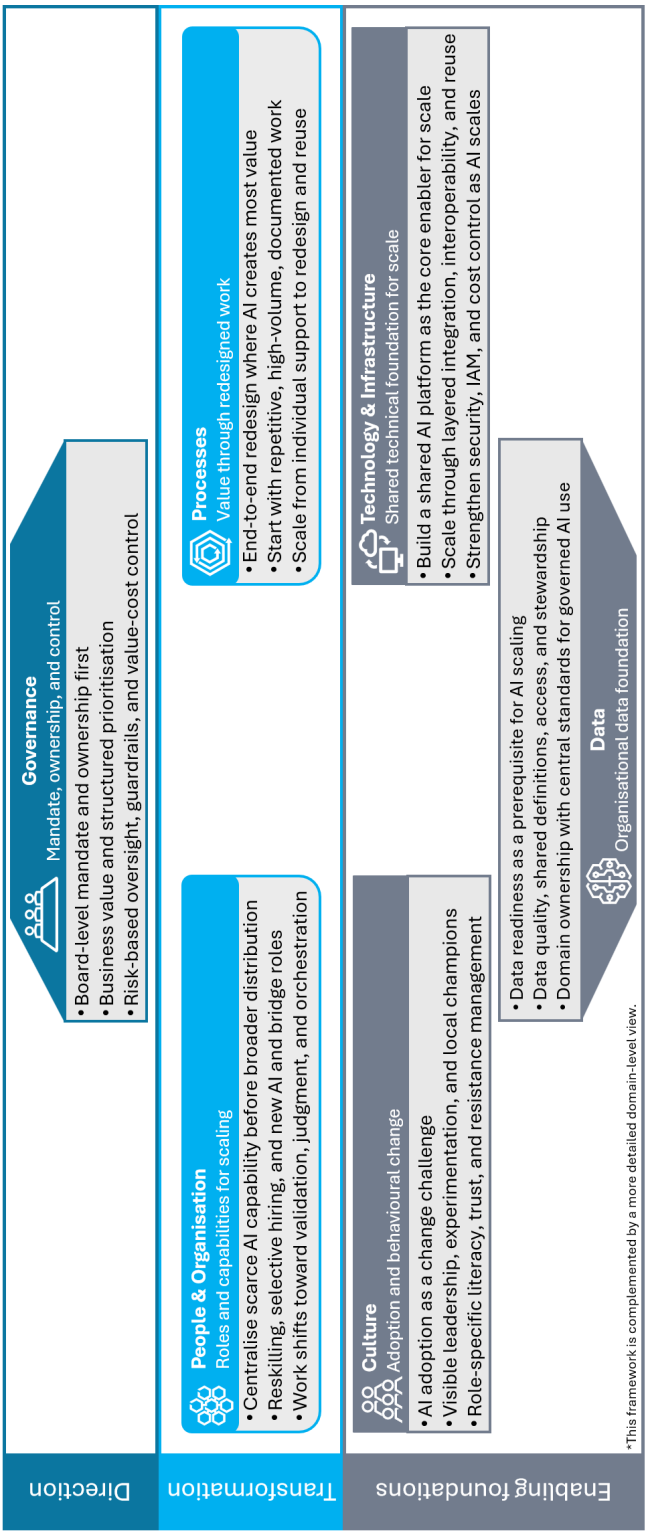


Figure D.1: Full-size condensed IT Operating Model framework for scaling AI (v2).

D.2 Full-size expanded framework figure



Figure D.2: Full-size expanded domain-level representation of the revised framework (v2).

D.3 Full-size staircase framework figure

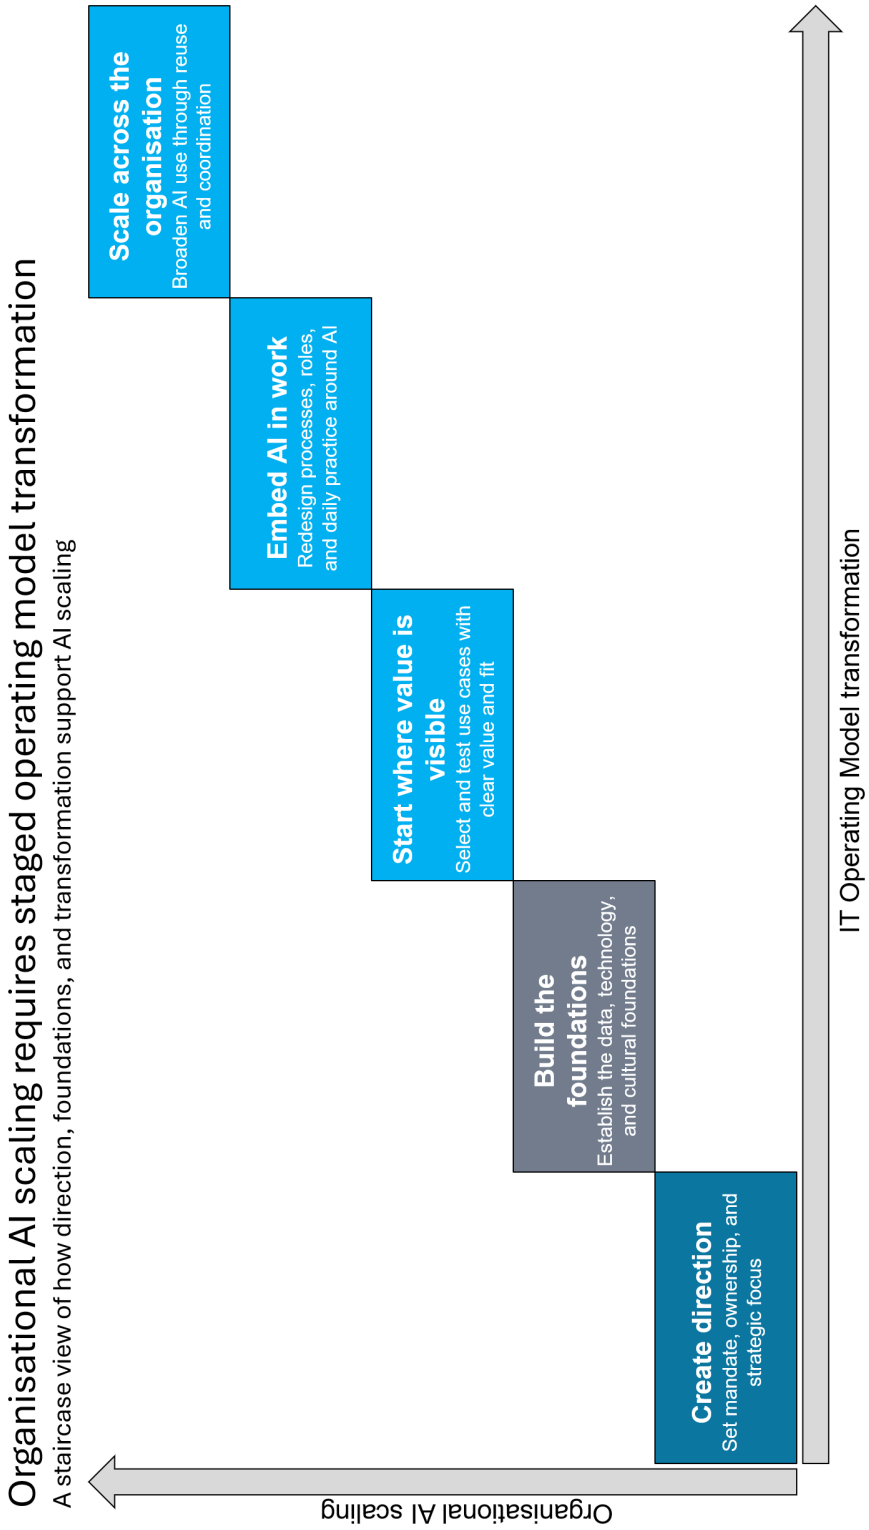


Figure D.3: Full-size staircase representation of the revised framework (v2).