



**UNIVERSITY
OF TURKU**

Turku School of
Economics

Change Management and the Integration of Generative AI in Organizational CRM Projects

A Qualitative Study in the Context of Salesforce Implementation

International Management of Information Technology (IMMIT)
Turku School of Economics, University of Turku

Master's thesis

Author:

Jihed Ben Mabrouk

Supervisor:

Jocelyn Husser

05/06/2026

Turku

Student's statement regarding the use of Artificial Intelligence (AI) for preparing and/or writing this thesis:

☐ I have not used any AI-based tools.

☒ I have used AI-based tools. Their use is documented in the Appendix. The AI tools were used in a way that complies with academic integrity guidelines.

The originality of this thesis has been checked in accordance with the University of Turku quality assurance system using the Turnitin Originality Check service.

Abstract

Master's thesis:

Subject: International Management of Information Technology (IMMIT)

Author: Jihed Ben Mabrouk

Title: Change Management and the Integration of Generative AI in Organizational CRM Projects, A Qualitative Study in the Context of Salesforce Implementation

Supervisor: Jocelyn Husser

Number of pages: 63

Date: 05/06/2026

Abstract

This thesis examines how change management (CM) practices shape the adoption of generative AI in customer relationship management (CRM) implementation projects. Using a mixed-methods design combining semi-structured interviews (N=5) with a pilot survey (N=89), we study two Salesforce CRM projects conducted by Accenture for two clients, a major French energy and services group and a French water utility. The findings reveal that AI readiness is highly uneven across populations, that resistance takes primarily epistemic rather than affective forms, and that change management is systematically deprioritized despite its measurable impact on readiness. We introduce two theoretical concepts: epistemic resistance, resistance rooted in the inability to evaluate AI output reliability, and augmentation drift, the gradual shift from augmentation to de facto automation when users stop critically engaging with AI outputs. The pilot data supports epistemic resistance as a stronger predictor of non-adoption ($\beta=-0.289$) than affective resistance ($\beta=-0.179$). Practical implications include investing in AI literacy over traditional communication, diagnosing digital maturity before deployment, and using regulatory compliance as a trust-building mechanism.

Keywords: Change Management, Generative AI, CRM, Salesforce, Epistemic Resistance, Augmentation Drift, Mixed Methods

Acknowledgements

I would like to express my sincere gratitude to my academic supervisor, Professor Jocelyn Husser, for his guidance, his validation of the research question and bibliography, and his encouragement throughout this process.

I am equally grateful to my internship supervisor, Albert, and to my colleagues at Accenture's Salesforce Practice, in particular Mohamad and Nour, for their support, their generosity in sharing their expertise, and for creating an environment in which this research could be conducted alongside my professional work.

I am also deeply indebted to Faezeh, my former mentor at bioMérieux, whose guidance helped me grow immensely, and to all my former colleagues there.

I owe special thanks to my friend and fellow consultant, Ghita, without whom I could not have seen this work through.

I am also deeply grateful to my roommates in Turku, Gloria, Irene, Maria, Emmanuel, Farhan, Giulia, Axel and Amina, and to my dear roommates in Tilburg, Jihyo, Carlota and Hannah, for their support throughout this journey, as well as to my entire IMMIT class, with whom I shared these two formative years.

My deepest gratitude goes to my family, my Mother, Zina and Suliman, and to my friends Marion, Lucas, Adam and Djessim, for their constant encouragement and unwavering support.

I would like to thank all five interview participants for generously giving their time and sharing their experiences with candour and depth, and the survey respondents who contributed to the quantitative component of this study.

Finally, I would like to thank IAE Aix-Marseille, Turku School of Economics and Tilburg University for providing the academic framework within which this research was conceived and developed.

Table of Contents

Abstract	3
Acknowledgements	4
Table of Contents	5
List of Figures	7
List of Tables	8
1 Introduction	9
1.1 Context and Motivation	9
1.2 Research Question	10
1.3 Thesis Structure	11
2 Theoretical Framework	12
2.1 Change Management Models and Their Limits	12
2.2 Technology Acceptance and Resistance	13
2.3 AI Readiness and the Automation-Augmentation Debate	14
2.4 Synthesis and Research Gap	15
3 Research Design and Methodology	17
3.1 Research Philosophy and Design	17
3.2 Qualitative Phase: Interviews	17
3.3 Quantitative Phase: Pilot Survey	21
3.4 Reflexivity and Limitations	22
4 Findings	24
4.1 AI Readiness Is Uneven Across Populations	24
4.2 Resistance Takes Multiple Forms	27
4.3 Change Management Was Informal or Absent	28
4.4 The Augmentation-Automation Boundary Is Blurred	30
4.5 Pilot Survey: Key Quantitative Patterns	31
4.6 Triangulation: Where Methods Converge and Diverge	35
5 Discussion and Theoretical Contributions	36
5.1 CM Frameworks Meet GenAI: Convergence and Limits	36
5.2 What Did Not Work	37
5.3 Epistemic Resistance	38
5.4 Augmentation Drift	40
5.5 Synthesis of Theoretical Contributions	41
6 Managerial Implications	42
6.1 Field Grounding: What the Two Client Engagements Reveal	42
6.2 Treat GenAI Deployment as Organizational Change	44
6.3 Invest in AI Literacy, Not Just Communication	45

6.4	Diagnose Readiness Before Deploying.....	45
6.5	Use Regulatory Compliance as a Trust Signal.....	46
6.6	Name the Failures Honestly.....	46
6.7	The Consultant in the Loop: Generative AI in the Practitioner’s Own Work	46
7	Limitations and Future Research	48
7.1	Limitations	48
7.2	Future Research.....	49
8	Conclusion	50
	References.....	51
	Appendices.....	55
	Appendix 1: Semi-Structured Interview Guide	55
	Appendix 2: Quantitative Survey Questionnaire	55
	Appendix 3: Participant Profiles Summary	56
	Appendix 4: Survey Results, Key Statistical Tables.....	57
	Appendix 5: Explanation of the Use of AI	59
	Appendix 6: Data Management Plan	59
	Appendix 7: Glossary of Key Project and Salesforce Terms.....	60

List of Figures

Figure 1. Research Positioning, Change Management × Generative AI × CRM Salesforce ..	10
Figure 2. Integrated Conceptual Framework, Change Management, GenAI Characteristics and CRM Adoption.....	16
Figure 3. Data Structure, First-Order Codes, Second-Order Themes and Aggregate Dimensions (Gioia)	20
Figure 4. Sequential Exploratory Mixed-Methods Research Design	20
Figure 5. Typology of Resistance Forms in GenAI Adoption.....	30
Figure 6. P1's Change Management Methodology, Five Phases and Three Levers	30
Figure 7. H3, Digital Maturity × AI Readiness ($r=0.582$, $p<0.001$)	32
Figure 8. H6, Epistemic vs Affective Resistance, Mean Comparison and Regression Coefficient	33
Figure 9. Correlation Matrix, Key Variables (N=89)	34
Figure 10. Summary of Hypothesis Testing Results.....	34
Figure 11. Dual Pathway Model, Epistemic vs Affective Resistance	40
Figure 12. Example output generated from a configured prompt in a Salesforce AI feature ..	61
Figure 13. The Salesforce Flow Builder canvas, showing the visual, no-code construction of an automated process	62

List of Tables

Table 1. Interview Participant Profiles	17
Table 2. Survey Sample Demographics (N=89)	31
Table D1. Scale Reliability (Cronbach's Alpha)	57
Table D2. Descriptive Statistics.....	57
Table D3. Hypothesis H1, CM Exposure and AI Readiness	57
Table D4. Hypothesis H2, CM Exposure and Resistance.....	58
Table D5. Hypothesis H3, Digital Maturity and AI Readiness	58
Table D6. Hypothesis H4, Age/Seniority and AI Readiness	58
Table D7. Hypothesis H5, Resistance and Adoption Intention	58
Table D8. Hypothesis H6, Epistemic vs Affective Resistance as Predictors	58
Table D9. Correlation Matrix.....	58
Table D10. Consultants vs Clients, AI Readiness	58

1 Introduction

1.1 Context and Motivation

Generative AI is reshaping how organizations approach customer relationship management. Tools like Salesforce Einstein Copilot and Microsoft 365 Copilot promise to automate routine tasks, generate client summaries, and suggest next-best-actions for sales teams. The global market for AI in CRM is growing rapidly, and consulting firms like Accenture have positioned generative AI integration as a core offering. The strategic weight of this shift was underscored in June 2026, when Salesforce announced a \$2 billion investment in France through 2030, including a new AI Innovation Hub in Paris dedicated to enterprise AI adoption (Salesforce, 2026). Yet the gap between the technological promise and the organizational reality remains wide. According to McKinsey (2023), roughly 70 percent of AI transformation initiatives fail to deliver their expected value. The most cited causes are not technical, they are organizational: insufficient change management, poor data quality, and lack of readiness.

This gap is not new. The history of enterprise technology is littered with technically successful implementations that failed to generate adoption. Enterprise resource planning (ERP) systems, CRM platforms, and business intelligence tools have all faced the same challenge: getting people to actually use them (Somers and Nelson, 2001). What is new with generative AI is the nature of the challenge. Unlike traditional tools with deterministic outputs, you click a button, you get a predictable result, GenAI produces probabilistic, non-reproducible outputs. The same prompt can yield different results. Users cannot trace how the tool arrived at its answer. This opacity creates a qualitatively different adoption challenge, one that existing change management frameworks were not designed to address.

This thesis examines this challenge from the inside. As a Salesforce consultant intern at Accenture's Salesforce Practice in Paris, the researcher was embedded in two CRM implementation projects where generative AI features were introduced: one for a major French energy and services group (referred to throughout this thesis as the energy client) and one for a French water utility (the water utility client). This insider position provided direct access to the practitioners, clients, and organizational dynamics that shape adoption outcomes in real time. The research is not retrospective: it captures the lived experience of GenAI integration as it unfolds.

Research Positioning: At the Intersection of Three Domains




Figure 1. Research Positioning, Change Management $\times$ Generative AI $\times$ CRM Salesforce

1.2 Research Question

Three bodies of work frame this inquiry and, taken together, expose the gap it addresses. First, the change-management tradition, Lewin's (1947) unfreeze, change, refreeze sequence, Kotter's (1996) eight steps, and Hiatt's (2006) individual-level ADKAR model, prescribes how organizations should steer transitions, yet Stouten, Rousseau and De Cremer (2018) show that these prescriptive, sequential logics bear little resemblance to how change actually unfolds. Second, the technology-acceptance and resistance literature, Davis's (1989) Technology Acceptance Model (TAM), Venkatesh et al.'s (2003) Unified Theory of Acceptance and Use of Technology (UTAUT), and the multilevel resistance models of Lapointe and Rivard (2005), Joshi (1991) and Markus (1983), assumes that users can assess a tool's usefulness and form a rational judgement about it; Feuerriegel et al. (2024) argue that generative AI's probabilistic, non-reproducible outputs unsettle precisely this assumption. Third, the AI-readiness and automation-augmentation literature (Jöhnk et al., 2021; Raisch and Krakowski, 2021) locates adoption in organizational conditions but says little about the population-level heterogeneity observed within a single firm. The friction between these three streams, prescriptive

change models, acceptance models that presuppose evaluability, and readiness frameworks pitched at the organizational level, is what motivates the question below, in the concrete setting of Salesforce CRM implementation projects where these dynamics are observed first-hand.

The core research question is: How do structured change management practices shape organizational adoption of generative AI in CRM implementation projects, and through which mechanisms do they mitigate or reproduce resistance to technological change?

Two sub-questions structure the inquiry. SQ1 asks what organizational and individual factors condition readiness for generative AI adoption in CRM environments, exploring why some populations adopt readily while others resist. SQ2 asks how practitioners experience and navigate resistance during generative AI integration, and what role formal change management frameworks play in this process, exploring the mechanisms through which resistance is expressed, addressed, or ignored.

Five research objectives guide the analysis: (1) identifying the specific challenges and frictions of GenAI integration in CRM, (2) analyzing the mechanisms of resistance at individual, group, and structural levels, (3) examining which CM practices are deployed, adapted, or neglected, (4) investigating AI readiness conditions, and (5) comparing the empirical findings with existing theoretical frameworks.

1.3 Thesis Structure

Chapter 2 presents the theoretical framework, drawing on five streams of literature: change management, technology acceptance, organizational resistance, AI readiness, and the automation-augmentation debate. Chapter 3 describes the mixed-methods research design combining qualitative interviews (N=5) with an exploratory pilot survey (N=89). Chapter 4 presents the findings, organized thematically. Chapter 5 discusses the results in dialogue with the literature and develops the study's two original theoretical contributions, epistemic resistance and augmentation drift. Chapter 6 offers managerial implications, and Chapters 7 and 8 address limitations and conclusions.

2 Theoretical Framework

2.1 Change Management Models and Their Limits

Three change management frameworks dominate both academic literature and consulting practice. Lewin's (1947) unfreeze-change-refreeze model provides the foundational logic: successful change requires destabilizing existing routines, introducing new ones, and consolidating them. Kotter (1996) operationalized this into eight sequential steps, from establishing urgency to anchoring new approaches in the culture. The ADKAR model (Hiatt, 2006) takes an individual perspective: Awareness, Desire, Knowledge, Ability, and Reinforcement represent the stages each person must go through for change to succeed.

Lewin's model, despite its age, captures something that later models often miss: the importance of the unfreezing phase. Before people can adopt a new tool or process, they must first let go of existing routines. In the context of GenAI, this means that experienced professionals who have developed efficient working methods over years, P4's senior sales staff with 15 years of client relationships, for example, face a higher unfreezing cost than junior employees who have fewer routines to abandon. Kotter (1996) added organizational complexity to this individual logic: his eight steps, from creating urgency to anchoring change in culture, recognize that change requires both top-down leadership and broad coalition-building. The model has become the de facto standard in consulting firms, including Accenture.

Yet these models have faced serious criticism. Stouten, Rousseau, and De Cremer (2018), reviewing decades of evidence, found that the prescriptive logic of sequential models bears little resemblance to how change actually unfolds in organizations. Change loops back, stalls, is shaped by power dynamics that linear models ignore. Autissier, Johnson, and Moutot (2018) offer a European relational alternative: change is not something done to employees but something co-constructed with them, through mobilization, engagement, and iterative dialogue.

The ADKAR model (Hiatt, 2006) adds a useful individual-level lens. ADKAR is an acronym in which each of the five letters names a condition that an individual must satisfy, in sequence, for change to take hold. The person must be Aware that change is coming, must Desire to participate, must acquire the Knowledge needed to change, must develop the Ability to apply that knowledge in practice, and must receive Reinforcement to sustain the new behaviour over time. P1 uses ADKAR as a mental checklist rather than a rigid protocol: 'I like to keep it in mind, telling myself: have I at least covered all these letters when I build my change strategy?' This pragmatic adaptation is revealing. It suggests that practitioners find value in the framework's diagnostic logic while recognizing that real projects do not follow its sequential structure.

These critiques are particularly relevant for generative AI. Unlike traditional application deployments, GenAI tools evolve continuously, features update, capabilities expand, the boundary between what the AI can and cannot do shifts over time. A model that assumes a stable end state to consolidate is mismatched with this continuous evolution.

The tension between prescriptive models and organizational reality is central to understanding GenAI adoption. In practice, as P1 will describe in the findings, change managers use ADKAR not as a rigid protocol but as a mental checklist, keeping the letters in mind without implementing them sequentially. This pragmatic adaptation is consistent with Weick and Quinn's (1999) distinction between episodic and continuous change: GenAI integration resembles continuous change more than the discrete episodes that Kotter's model was designed for. The tools update, user needs evolve, and the change management effort never reaches a stable end state to consolidate.

2.2 Technology Acceptance and Resistance

The Technology Acceptance Model (Davis, 1989) and UTAUT (Venkatesh et al., 2003) identify perceived usefulness and perceived ease of use as the primary drivers of technology adoption. Chatterjee et al. (2023) extended these models to AI-CRM contexts, while Feuerriegel et al. (2024) argue that generative AI's probabilistic, non-deterministic outputs challenge the assumptions underlying these frameworks: perceived usefulness is harder to assess when the same prompt can produce different results.

Davis's (1989) TAM identified two variables, perceived usefulness and perceived ease of use, as the primary determinants of technology adoption. The model has been applied hundreds of times across diverse technology contexts. Venkatesh et al. (2003) extended it into UTAUT, adding social influence and facilitating conditions. These models work well for technologies with deterministic, predictable outputs: users can evaluate whether a CRM dashboard is useful by looking at the data it presents. But generative AI produces probabilistic outputs, a summary generated by Einstein Copilot may be accurate or misleading, and the user often has no way to tell which. This makes perceived usefulness harder to assess, because usefulness is contingent on a capability the user may lack: the ability to evaluate what the AI produced.

On the resistance side, Lapointe and Rivard (2005) propose a multilevel model distinguishing individual, group, and organizational resistance. Joshi (1991) frames resistance as a rational equity assessment. Markus (1983) emphasizes the political dimension: resistance often reflects power dynamics, not irrationality. Autissier adds what he calls informational resistance, resistance born from lack of understanding rather than opposition.

Lapointe and Rivard's (2005) multilevel model is particularly useful because it distinguishes the unit of analysis. Individual resistance involves personal responses to perceived threats, fear of obsolescence, discomfort with unfamiliar tools. Group resistance emerges when teams collectively decide that a tool is unworkable or threatening. Organizational resistance operates through institutional mechanisms: compliance processes, governance requirements, budget allocations. All three levels appeared in our data: P5's customer service employees expressed individual fear of job loss, P2's consulting team collectively normalized the dismissal of Copilot after poor initial experiences, and the energy client's DSI imposed institutional conditions that delayed deployment.

From the IS perspective, the challenges of GenAI adoption echo patterns documented in the ERP/CRM implementation literature. Somers and Nelson (2001) identified top management support, change management investment, and user training as critical success factors, factors that remain directly applicable to GenAI integration. Hong and Kim (2002) showed that organizational fit between technology and existing processes is a key determinant of success.

What these resistance frameworks share is an implicit assumption that the user can evaluate the technology's value and make a rational judgment about it, whether to accept, resist, or negotiate. For generative AI, this assumption is problematic. When the tool produces outputs that the user cannot evaluate for accuracy or reliability, the user faces a qualitatively different situation. The resistance literature does not have a category for this: it is not political (Markus), not equity-based (Joshi), not affective (fear or discomfort), and not informational in Autissier's sense (lack of information about the change). It is something else, something that emerges at the intersection of AI opacity and human cognitive limits. We return to this in Chapter 5 (sections 5.3 and 5.4).

2.3 AI Readiness and the Automation-Augmentation Debate

Jöhnk et al. (2021) propose a multidimensional AI readiness framework encompassing technological, organizational, and people-related dimensions. Agrawal (2024) applies a Technology-Organization-Environment (TOE) lens to generative AI adoption. Vial (2019) and Verhoef et al. (2021) situate AI within the broader trajectory of digital transformation, arguing that organizations move through stages, digitization, digitalization, and digital transformation, and that AI readiness depends on where the organization sits on this continuum.

Jöhnk et al.'s (2021) framework is one of the few that attempts to operationalize AI readiness systematically. They identify five dimensions: strategic alignment (does the organization have an AI strategy?), resources (budget, talent, data infrastructure), knowledge (AI literacy at various levels), culture (openness to experimentation, tolerance for failure), and data (quality, accessibility,

governance). Applied to our cases, the energy client scores higher on most dimensions, they had an internal AI strategy, a DSI (Direction des Systemes d'Information, the corporate IT function) governance process, and at least some digitally mature users. The water utility client scores lower on nearly all dimensions, which is precisely why P5 described the GenAI initiative as premature.

Raisch and Krakowski (2021) frame the central tension as the automation-augmentation paradox: the same technology can either replace human work or enhance it, and the boundary is negotiated in practice rather than determined by design. Syam and Sharma (2018) apply this specifically to CRM, arguing that AI will augment rather than replace sales professionals. Davenport and Ronanki (2018) identify three categories of AI application, process automation, cognitive insight, and cognitive engagement, each carrying different implications for human roles.

The Explainable AI (XAI) literature adds an important dimension. Longo et al. (2020) argue that AI systems function as opaque black boxes whose internal processes remain inaccessible to users. This opacity generates distrust independently of actual performance. For generative AI, where outputs are probabilistic and non-deterministic, users have no way to trace how a particular output was generated, creating a structural basis for what we will conceptualize as epistemic resistance.

The regulatory dimension adds another layer. The General Data Protection Regulation (GDPR; European Parliament and Council of the European Union, 2016) and the emerging EU AI Act (European Commission, 2024) impose transparency and accountability requirements on AI-deploying organizations. In sectors like energy (the energy client) and water utilities (the water utility client), data sensitivity is high and compliance is non-negotiable. These regulatory constraints shape readiness in two ways: they create prerequisites that must be met before deployment can proceed, and they provide a legitimate institutional basis for caution that transcends individual attitudes. In heavily regulated industries like pharmaceuticals, structured change processes are built into the regulatory fabric; in less regulated sectors, organizations must deliberately invest in accompaniment, an investment that, as we will see, is systematically deprioritized.

2.4 Synthesis and Research Gap

The literature provides rich but fragmented insights. Change management models prescribe how to manage transitions but were designed for static technology deployments. Technology acceptance models identify adoption drivers but assume users can evaluate tool usefulness, an assumption GenAI challenges. Resistance theories explain why people push back but do not account for resistance rooted in cognitive inability rather than emotional opposition. AI readiness frameworks operate at the organizational level but miss the population-level heterogeneity within single organizations.

What is missing is an integrated examination of how these dynamics interact in practice, in the specific context of generative AI in CRM. Recent work converges on this point without closing it. Berente et al. (2021) argue that managing AI differs in kind from managing earlier information technologies, because of its autonomy, learning and inscrutability, and call for management theory adapted to these properties rather than transposed from prior IT. Feuerriegel et al. (2024) make the same case specifically for generative AI, whose probabilistic outputs strain established acceptance models. In the CRM domain, Ledro et al. (2025) review the AI-CRM integration literature and conclude that successful integration is studied mainly as a technical or data problem, with the organizational and change-management conditions treated as a residual, an explicit invitation to examine how adoption is accompanied in practice. Chatterjee et al. (2023), extending UTAUT to AI-integrated CRM, similarly note that organizational accompaniment remains under-theorised. We seek to bridge this gap by studying empirically how CM practices interact with the specific characteristics of GenAI, its probabilistic outputs, its dependence on prompting skill, and its blurred automation-augmentation boundary, in the concrete setting of Salesforce CRM implementation projects.




Figure 2. Integrated Conceptual Framework, Change Management, GenAI Characteristics and CRM Adoption

3 Research Design and Methodology

3.1 Research Philosophy and Design

This research adopts an interpretivist-constructivist epistemology, recognizing that adoption and resistance to generative AI are socially constructed phenomena shaped by individual experience, organizational context, and power dynamics. The reasoning follows an abductive logic (Dubois and Gadde, 2002): we started with theoretical sensitizing concepts from the change management and technology acceptance literature, but allowed the empirical data to generate new constructs not anticipated by the initial framework.

The study employs a sequential exploratory mixed-methods design (Creswell and Creswell, 2018), in which a qualitative phase leads and a quantitative phase follows to examine whether the qualitatively derived constructs hold at a broader distributional level. Phase 1 consists of five semi-structured qualitative interviews with purposively selected practitioners. The themes emerging from these interviews informed Phase 2: an exploratory pilot survey (N=89). The choice of a pilot rather than a full confirmatory survey is deliberate and follows Saunders, Lewis and Thornhill (2019), who position pilot studies as instruments for refining measures and probing the feasibility and direction of relationships before any claim to generalisation. In an exploratory design the survey's role is precisely this: to operationalise emergent constructs, above all epistemic resistance, and to test whether their distributional patterns are consistent with the interview findings, not to deliver population-level estimates (Creswell and Creswell, 2018). The qualitative phase therefore carries the primary evidentiary weight; the quantitative phase serves a complementary triangulation function.

3.2 Qualitative Phase: Interviews

Five participants were selected through purposive sampling to capture the diversity of perspectives relevant to GenAI adoption in CRM projects: a change management specialist (P1), a Salesforce consultant/project manager (P2), a partner leading the Salesforce practice (P3), a client-side project manager at the energy client (P4), and a client-side representative at the water utility client (P5). Interviews lasted 15 to 20 minutes, were conducted in French, and followed a semi-structured guide organized around five themes: methodology and approach, GenAI experience, adoption and resistance, change management practices, and automation-augmentation perceptions. The full interview guide is provided in Appendix 1, and a detailed profile of each participant is given in Appendix 3.

Table 1. Interview Participant Profiles

Code	Role	Organisation	Experience and GenAI exposure
P1	Change management specialist	Accenture	3 yrs; HRIS, digital-workplace and AI-chatbot deployment in banking
P2	Manager & Salesforce consultant	Accenture	Project manager on both client projects; Copilot, Amethyst, Einstein, Agentforce
P3	Partner, Salesforce Practice	Accenture	Practice lead; strategic oversight of GenAI integration
P4	Client-side project manager	Energy client	Internal PM; tested GenAI features (summaries, next-best-action)
P5	Client-side representative	Water utility client	Internal stakeholder; lower-digital-maturity context

P1 was transcribed verbatim; P2–P5 are reconstructed from detailed interview notes.

Analysis followed Braun and Clarke's (2006) six-phase thematic analysis, informed by Gioia methodology principles. Data was coded at two levels: first-order descriptive codes close to participants' language, and second-order analytical codes representing higher-level constructs. Trustworthiness was addressed through triangulation across participants, analytical transparency, and thick description of data and interpretive choices (Lincoln and Guba, 1985; Yin, 2018; Hlady-Rispal, 2002).

Because the researcher was not an external observer but an embedded participant on the two client projects, the interviews were complemented by an analytic autoethnographic strand (Anderson, 2006; Ellis, Adams and Bochner, 2011). Analytic autoethnography is appropriate here precisely because the researcher is a full member of the setting studied and seeks not merely to narrate personal experience but to use it as data in the service of theoretical analysis (Anderson, 2006). Throughout the internship the researcher kept a field diary (journal de bord), recording observations, informal exchanges, and impressions of how GenAI features were received during workshops, configuration sessions, and day-to-day delivery. These notes were treated as a distinct empirical source and read against the interview material. Where a diary entry and a verbatim converged, for example on the difficulty users had in prompting, the observation was retained with greater confidence; where they diverged, the divergence was flagged for interpretation. This use of the researcher's own situated practice extends, rather than replaces, the five semi-structured interviews.

The qualitative analysis therefore rests on three converging sources, in line with the principle of data triangulation (Yin, 2018; Hlady-Rispal, 2002): the five interview transcripts, the field diary, and a set of internal project documents to which the researcher had access in the course of the engagement, namely adoption discussions, workshop supports, configuration and training materials, and project communications relating to the GenAI features. For reasons of client confidentiality these documents cannot be reproduced in the appendices, but they are mobilised in the findings to corroborate or

qualify the verbatims, and their nature is declared here so that the evidentiary basis of each claim is transparent.

The five participants represent three organizational levels: the consulting firm (P1, P2, P3), and two client organizations at different maturity stages (P4, P5). This diversity was intentional: we sought perspectives from those who design and deliver GenAI features (consultants), those who manage the integration strategically (partner, managers), and those who receive and use the tools (clients). One limitation is the absence of frontline end-users without managerial responsibilities, the bankers, the customer service agents, the field workers whose daily experience of GenAI tools is the ultimate test of adoption success.

The thematic analysis produced sixteen first-order codes, grouped into six second-order themes, which in turn condensed into four aggregate dimensions structuring Chapter 4. The resulting data structure, presented in Figure 3, follows the Gioia template (Gioia, Corley and Hamilton, 2013). The aggregate dimension uneven AI readiness draws together first-order codes such as “IT teams adapt, branch users struggle,” “juniors use AI by reflex,” “seniors find the learning cost unjustified,” and “C-suite ambition versus operational reality,” forming the second-order themes population heterogeneity and organizational-infrastructure gap. The dimension forms of resistance assembles “cannot prompt or evaluate output,” “tried once and abandoned,” “identity threat among experienced staff,” “governance gatekeeping,” and “uncritical acceptance,” under the themes epistemic resistance and political-identity-institutional resistance. The dimension change-management practice gathers “structured five-phase method,” “learning by doing,” “CM is the first budget line cut,” and “sponsorship and management buy-in,” under prescribed CM and CM as actually practised. The dimension automation versus augmentation gathers “tasks will disappear,” “fear of replacement,” and “automation by default,” under the theme blurred boundary. Codes were iteratively refined through multiple readings, moving between the transcripts, the field diary and the emerging structure in an abductive logic; the two constructs not anticipated by the initial framework, epistemic resistance and augmentation drift, emerged directly from this process.




Figure 3. Data Structure, First-Order Codes, Second-Order Themes and Aggregate Dimensions (Gioia)

The four dimensions are not equally well fed by the data, and it is worth saying so explicitly. Resistance and AI readiness are the richest dimensions: each is supported by all five participants as well as by the field diary, which is why they anchor the most developed sections of Chapter 4. Change-management practice is well documented on the prescribed side, where P1 alone provides a detailed five-phase account, but thinner on the as-practised side, where evidence rests mainly on P2 and P3. The automation versus augmentation dimension is the least saturated: it relies chiefly on P2, P3 and P5 and surfaces only briefly in the other interviews, which is one reason augmentation drift is later presented as an exploratory proposition rather than an established finding.




Figure 4. Sequential Exploratory Mixed-Methods Research Design

3.3 Quantitative Phase: Pilot Survey

The pilot survey was administered to 89 respondents: 63 consultants of varying seniority and 26 client-side participants. The full instrument is reproduced in Appendix 2; it comprises five multi-item Likert scales whose items were adapted from established sources. The digital maturity scale (5 items) was built on the digital-transformation stage logic of Verhoef et al. (2021) and Vial (2019). The AI readiness scale (8 items) operationalises the people, organization and data dimensions of Jöhnk et al.'s (2021) readiness framework. The resistance scale (8 items) is split into an affective subscale (4 items), grounded in the individual-level resistance of Lapointe and Rivard (2005) and Joshi (1991), and an epistemic subscale (4 items) developed for this study from the explainable-AI literature on opacity and trust (Longo et al., 2020). The change-management perception scale (6 items) draws on the ADKAR and Prosci change-practice tradition (Hiatt, 2006; Prosci, 2018). The adoption intention scale (3 items) was adapted from the behavioural-intention items of TAM and UTAUT (Davis, 1989; Venkatesh et al., 2003). Six hypotheses, derived from the qualitative findings, were tested using independent t-tests, Pearson correlations, and multiple regression; the corresponding statistical output is reported in full in Appendix 4.

Specifically, the six hypotheses were: H1, exposure to structured change management is associated with higher AI readiness; H2, exposure to structured change management is associated with lower resistance to GenAI; H3, digital maturity is positively associated with AI readiness; H4, age and seniority are associated with AI readiness; H5, resistance to GenAI is negatively associated with adoption intention; and H6, epistemic resistance is a stronger predictor of (non-)adoption than affective resistance. The first pair tests the role of change-management exposure (Section 4.3), the second pair the readiness conditions raised by SQ1, and the third pair the resistance mechanisms raised by SQ2. Each hypothesis is operationalised through the scales described above and reported, with its test statistic, in Appendix 4 (Tables D3 to D8).

The sample reflects the population structure identified through the interviews. Results should be read as indicative of directional patterns rather than definitive population-level estimates. The purpose of the quantitative component is to show how the qualitative constructs, particularly epistemic resistance, could be measured and tested, not to claim statistical generalizability.

All multi-item scales achieved Cronbach's alpha values above 0.70: AI Readiness ($\alpha=0.87$), Resistance ($\alpha=0.84$, with subscales: affective $\alpha=0.79$, epistemic $\alpha=0.81$), CM Perception ($\alpha=0.88$), and Adoption Intention ($\alpha=0.76$). The sample comprised junior consultants (N=22), senior consultants (N=18), managers and team leads (N=15), change management specialists (N=8), client-side users (N=16), and client-side managers (N=10). Median age was 35 years. Statistical analyses

included independent t-tests for group comparisons, Pearson correlations for bivariate relationships, and multiple linear regression for the key hypothesis comparing epistemic and affective resistance as predictors.

(The sequential exploratory design is shown in Figure 4.)

3.4 Reflexivity and Limitations

As a Salesforce consultant intern at Accenture working daily on the two client projects, the researcher occupied an insider position. This afforded privileged access to participants and organizational context, but also introduced potential biases: participants may have moderated criticism of Accenture when speaking with a colleague, and the researcher's own tool experience may have shaped interpretive choices. To mitigate these risks, questions were open-ended, analysis followed a systematic protocol, and findings were triangulated across participants and with the pilot data. The qualitative sample (N=5), while defensible for interpretivist research, limits the breadth of perspectives captured; notably, no end-users without managerial responsibilities were interviewed. Consistent with the analytic autoethnographic posture adopted here (Anderson, 2006), this insider position is treated as an analytic resource as much as a limitation: the field diary makes the researcher's situated viewpoint explicit and auditable rather than hidden. Finally, in accordance with the institutional requirements of the programme, the responsible use of AI tools in preparing this thesis is documented in Appendix 5, and the handling of the empirical material is set out in the data management plan in Appendix 6.

This insider position warrants more specific disclosure. Three of the five interview participants were colleagues at Accenture: P1, the change manager whose detailed account anchors much of the qualitative material, is also the researcher's project architect and informal mentor on the engagement; P2, a manager directly involved in the two client projects; and P3, a partner of the Salesforce Practice. The two client-side participants (P4 at the energy client, P5 at the water utility client) were also accessed through the consulting engagement. This sampling profile is a direct consequence of the auto-ethnographic posture (Anderson, 2006): the field is the field that the researcher actually inhabits, not one constructed from the outside. But it carries identifiable risks. Hierarchical proximity to P1 and P3 may have led these participants to frame Accenture's practice more favourably than they would in front of a neutral interviewer, and the researcher may, reciprocally, have weighted their accounts more heavily in interpretation.

Three concrete mitigation steps were taken. First, the analytic logic was kept explicitly abductive (Dubois and Gadde, 2002): codes were generated from the transcripts rather than imported from the existing change-management literature, and the emergence of constructs not anticipated by the initial

framework, particularly epistemic resistance and augmentation drift, is in itself evidence that the data was allowed to speak back to the researcher's prior categories. Second, P1's verbatim transcription, used in extenso in Chapter 4, functioned as a within-case validity check: P2-P5 codes were tested for resonance against P1's account, and divergences (notably P5's environmental resistance theme, which has no equivalent in the consultant interviews) were treated as analytically significant rather than smoothed over. Third, the pilot survey (N=89) was administered to a sample broader than the immediate Accenture circle, and the headline patterns of the qualitative material were submitted to distributional test; the fact that two qualitatively expected effects did not reach significance (H2 on CM-resistance, H4 on age) is itself a discipline against the over-interpretation that insider sympathy might encourage.

Following Klein and Myers's (1999) principles for interpretive field research, this account treats the researcher's situated viewpoint as auditable rather than hidden. The field diary kept throughout the internship recorded encounters, decisions and analytical hesitations; it is not reproduced here but constitutes the trace by which any inferential step could, in principle, be reviewed.

4 Findings

The findings are organized around four qualitative themes emerging from the thematic analysis, followed by a summary of the pilot survey results. Throughout, participant verbatims are presented in italics. P1's verbatims are drawn from a verbatim transcription; P2-P5 verbatims are reconstructed from detailed interview notes.

4.1 AI Readiness Is Uneven Across Populations

All five participants described AI readiness as highly uneven. P1, who deployed a chatbot across a banking organization, found that IT teams adapted quickly while branch employees struggled: *"The IT teams knew exactly how to use it, no problem. But the users in branches, typically the bankers, had more difficulty, because they were not acculturated to AI. For some of them, AI is a bit like magic, you can ask anything, write however you want, and you will always get the right answers. Except that is not the case. They simply did not know how to prompt."*

P2 observed the same pattern within the Accenture team: *"Juniors, well, they use it naturally, for them it is almost a reflex. But the more senior consultants, how shall I say, they are reluctant. They find that the time invested in learning the tool is not justified by what comes out."* P3 described a gap at a different level, between executive aspiration and operational reality: *"The C-suite all want AI. But when you go to the operational level, well, the data is not clean, governance is unclear, teams are not trained."*

The readiness gap is not simply about individuals. It is structural. The energy client had invested in an internal AI strategy driven by the direction générale, which meant that when Accenture proposed GenAI features, the organization had a framework for evaluating them. P4 described the reaction: 'Our first reaction was not resistance, it was: is our data ready for this?' This is a fundamentally different starting point than the water utility client, where P5 reported that the AI proposition came from Accenture without internal demand: 'We did not ask for AI. Accenture proposed it.' The difference in readiness between these two organizations is not about individual attitudes, it is about the organizational infrastructure that shapes how individuals encounter and evaluate new technology.

The contrast between clients was stark. P4 (the energy client) described an organization with an existing AI strategy and a DSI that validated security before proceeding. P5 (the water utility client) described the opposite: *"We have people who still have difficulty, well, creating an opportunity in Salesforce, the basics. You cannot put AI on top of a tool that people do not yet master properly."* P5 estimated GenAI adoption at below 10 percent. These findings are consistent with broader industry patterns: roughly 70 percent of AI initiatives fail to deliver expected value (McKinsey, 2023).

P2 noted that GenAI was not part of the original project scope: *"Generative AI, it arrived mid-project. It was not, it was not part of the original scope. It came along the way, and we had to integrate it without, well, without a dedicated change strategy."* This ad hoc integration, without time for impact analysis or readiness assessment, stands in sharp contrast with P1's structured five-phase approach and illustrates how GenAI's rapid emergence catches organizations unprepared.

P1 also observed that AI adoption is easier to promote than traditional application changes: *"It is always nicer to tell someone we have an option with AI that will make your life easier, rather than telling them the tool you have been using every day for 30 years, we are going to change it completely."* But this enthusiasm comes with a condition: *"still, they need to know exactly what it does and what it does not do."*

Vignette: A Workshop Series at the energy client's Spanish subsidiary, Spring 2026

Where the rest of this chapter speaks through the verbatims of P1 to P5, the short passage that follows is written in the first person; it draws on the field diary kept during a workshop series at the energy client's Spanish subsidiary in the spring of 2026 and is included as an extension of the analytic autoethnographic strand set out in Section 3.2.

I was sent to Madrid for a three-day workshop series with the Spanish marketing and CRM teams of the energy client's Spanish subsidiary, in the context of the Marketing Cloud Advanced (MCA) and Data Cloud roll-out for the subsidiary. The brief, on paper, was straightforward: gather the use cases the business teams wanted to support with the new platform, identify which ones would require AI features, and produce a prioritisation matrix that the integration team could deliver against. What I observed across the three days in the Madrid office did not fit the brief at all, and it is one of the observations that, on the return flight, pushed me to take the construct of epistemic resistance more seriously than the literature review alone had done.

The first thing that struck me was the strength of demand. Every team that came into the workshop room arrived with at least one explicit AI request. The campaign managers wanted "AI to predict campaign success." The data analysts wanted "an AI that automatically segments contacts." The marketing director wanted "an assistant that drafts personalised emails for every audience cluster." When I asked, gently, what data they would feed the model that was meant to predict campaign success, the answer was vague: past campaigns, perhaps; but past campaigns had not been logged in any structured way, so the model would have had nothing to learn from. When I asked the analysts what segmentation rules they currently used, they showed me a spreadsheet of ten manual filters that had not been touched in eighteen months; segmentation here was not a problem an AI would solve,

it was a problem nobody had been resourced to maintain, and the AI request was a way of avoiding the maintenance work rather than of doing it differently. The personalised-email request was the most revealing of the three: the team was already using a template-based engine inside Marketing Cloud that produced exactly the outcome the proposed “AI assistant” was supposed to deliver, only through merge fields rather than a language model. Asked whether they had measured the response rate of the existing templates, no one had a figure.

What was being asked for, in other words, was not really AI. It was relief from operational work that the teams had not been resourced to do, and AI was the available vocabulary in which that relief could be requested with managerial legitimacy. None of the people across the table opposed the technology; on the contrary, their explicit framing was enthusiastic. Yet they were unable to specify, when pressed, what the AI would do, on what data, and against which baseline. The qualitative literature on resistance, with its categories of affective, identity-based or institutional refusal, did not capture this configuration well. The Madrid teams were not refusing to engage with AI. They were unable to formulate the engagement. The gap was epistemic, not motivational, and it was located on the demand side rather than on the user side, which is a small but, I think, important displacement of the construct as I had begun to develop it from P1, P2 and P4’s accounts.

One workshop session, on the third afternoon, suggested where genuine fit might have existed. The Spanish content team was struggling to produce localised versions of email campaigns that originated in French headquarters; they had a backlog of source decks to translate, calibrate, and adapt to the Spanish market, and they were running consistently behind on the deliverables. This was a problem with a clear input (French source content), a clear output (Spanish localised content), and an explicit quality bar (no factual errors, tone aligned to the local audience). Generative AI is good at exactly this kind of bounded, document-to-document task; a tightly scoped prompt-and-review workflow could plausibly have halved the team’s turnaround time. They did not ask for AI on this case. The opportunity was visible from the consultant’s chair in the room, not to the team whose work would have been supported, because the team experienced the task as “translation,” a category they associated with human craft, rather than as “structured content adaptation,” a category they would have associated with automation had they been used to thinking in those terms.

I left Madrid with two notes in the field diary. The first was that AI demand and AI fit are not aligned, and that, in this configuration at least, they were inversely correlated: the strongest demand came from the use cases that would benefit least, and the use case that would benefit most generated no demand at all. The second was that the existing framing of resistance, even after the elaboration of epistemic resistance I had begun to draft, was still missing this third mode. Affective resistance pushes

the tool away. Epistemic resistance, as I was formulating it, leaves the user unable to evaluate the output once produced. The Madrid pattern is a step earlier still: it leaves the requesting party unable to articulate the problem the tool would address. The space between aspirational demand and an articulable use case is, for AI in CRM, a substantial one. I treat this passage as a vignette rather than as a finding, but it is the kind of observation the autoethnographic protocol set out in Section 3.2 was designed to surface, and I include it here because it shaped, on the analytic side, how I came to read both P1's chatbot deployment earlier in this section and P5's adoption figure at the water utility client in the next.

4.2 Resistance Takes Multiple Forms

Resistance to GenAI was not monolithic. Five distinct forms emerged from the data.

The most prevalent was what we term epistemic resistance: an inability to evaluate AI output reliability. P1 described users who *"did not know how to prompt"* and for whom the adoption effort required *"much more than just giving people the desire and doing promotion."* P2 described the same dynamic as premature abandonment: *"They try once, it does not work, and they say, well, this is useless, I will do it myself."* This is not emotional opposition. It is a cognitive gap: users lack the skills to interact with the tool effectively and the criteria to judge whether its outputs are reliable.

The distinction between epistemic and affective resistance is not academic hair-splitting, it has direct practical consequences. If the primary barrier is affective (fear, discomfort), then the appropriate response is reassurance, communication, and empowerment: P1's communication and mobilization levers. But if the barrier is epistemic (inability to evaluate), then the appropriate response is education: teaching people how to prompt, how to evaluate outputs, how to know when the AI is wrong. The data consistently points to the second: P1's users did not refuse the chatbot, they did not know how to use it. P2's consultants did not oppose Copilot, they abandoned it after one bad experience because they lacked the evaluative framework to distinguish a prompting error from a tool limitation.

Second, identity-based resistance emerged among experienced professionals. P4 reported that senior sales staff at the energy client saw AI as a challenge to their expertise: *"Their reaction is essentially, I know my clients better than any algorithm. I do not need a machine to tell me what to propose to someone I have been working with for, what, ten years."*

Third, institutional resistance operated through governance processes. P4 noted that nothing moved forward without the DSI's validation at the energy client. Fourth, environmental resistance surfaced through union concerns about AI's energy consumption (P1). Fifth, and most striking, P3 identified the mirror image of resistance, uncritical adoption: *"There is the opposite risk, which people talk*

about less: teams using AI without any critical judgment, accepting everything the AI says without verifying. That is equally dangerous."

It is important to name the water utility client's project honestly: with adoption below 10 percent, activated features going unused, and no internal champion, it is a deployment that did not achieve its objectives. The energy client, at 20-30 percent adoption concentrated on the least transformative feature (account summaries), also fell short of its potential.

P1 reported three specific forms of institutional concern during the chatbot deployment. Trade unions raised data privacy fears. An environmental concern about AI energy consumption surfaced through a corporate social responsibility (CSR) lens. And the most widespread pattern was passive non-adoption rooted in incomprehension: users simply did not engage with the tool because they did not understand it, not because they opposed it. The distinction matters: opposition can be addressed through persuasion, but incomprehension requires education.

P2 framed the core challenge as one of trust: *"The real issue is trust. People need to trust the tool, trust that it will give them something useful, that it will not produce errors they will be held responsible for. And that trust, you cannot just, you cannot decree it. It has to be built through experience."*

4.3 Change Management Was Informal or Absent

P1 described a structured CM methodology with five phases, scoping, strategy, content creation, execution, and aftercare, and three levers: communication, training, and mobilization. Among these, P1 argued that mobilization was *"the least used but, for me, the most effective for adoption."* The mechanism is co-construction: *"From the moment you co-construct something with someone, they cannot reject it. It also comes from them."* The most critical success factor was management engagement: *"If you have a manager who says the tool is great, they will push for the teams to use it. If you have a manager who says it is nonsense, they will ask their team not to use it, and then you can do as much communication and training as you want, it will not be used."*

On the two client projects, this structured approach was absent. P2 described the GenAI adoption process as informal: *"There is no structured change management plan for GenAI. It was more learning by doing, people discover the tools, they share tips between colleagues."* P3 explained why: *"Change management is seen as, how to say, a nice to have. When the client faces budget constraints, and they always do, change is the first thing that gets cut."* P3 then articulated the resulting paradox: *"And then the same clients come back six months later and say: the adoption rates are low. And we tell them: yes, because you did not invest in accompaniment."*

The contrast between P1's structured approach and the reality of the two client projects is striking. P1 described five phases, three levers, stakeholder mapping, management engagement, and post-deployment aftercare. On the client projects, P2 described 'learning by doing' with no formal structure. This is not because Accenture does not know how to do change management, P1's account proves they do. It is because the client did not buy it. P3's explanation is blunt: 'We include change management in every proposal. But when the client says remove this line, reduce the budget, change is what goes first.' The result is predictable and was observed: low, uneven adoption.

P4 (the energy client) specifically identified what was missing: practical workshops with real client data rather than generic demos. P5 (the water utility client) identified a deeper gap: the absence of any internal sponsor: *"There is nobody internally at the water utility client who is really championing the AI project. Without that internal push, well, nothing happens."* P1 also emphasized that the product must meet expectations before any CM can succeed: *"You can have impeccable change management, as long as the product does not meet expectations, it is useless."*

P1 also emphasized the importance of stakeholder mapping, identifying who has influence within the organization, regardless of hierarchical position: *"Sometimes you have ordinary people among the users who have more influence and are more respected than the manager. It is not always easy, it is really a very social, political game. But once you have found the person who can carry your project, that is bingo."*

The aftercare phase, monitoring adoption after deployment, was described by P1 as 'often set aside but extremely important.' This is consistent with the critique of linear CM models: adoption does not end at go-live. Without sustained follow-up, initial engagement fades and users revert to pre-existing habits.

P1's advice for anyone deploying an AI tool was direct: 'I would say: lean heavily on training and information. We take for granted that AI is easy, that everyone uses ChatGPT. But that is not the case for everyone. There are always users who do not even know what AI is, and it can scare them. You need to start with AI awareness, explain what it is and what it does not do. And you need to teach people to prompt, that is how you get the best results and generate adoption.'

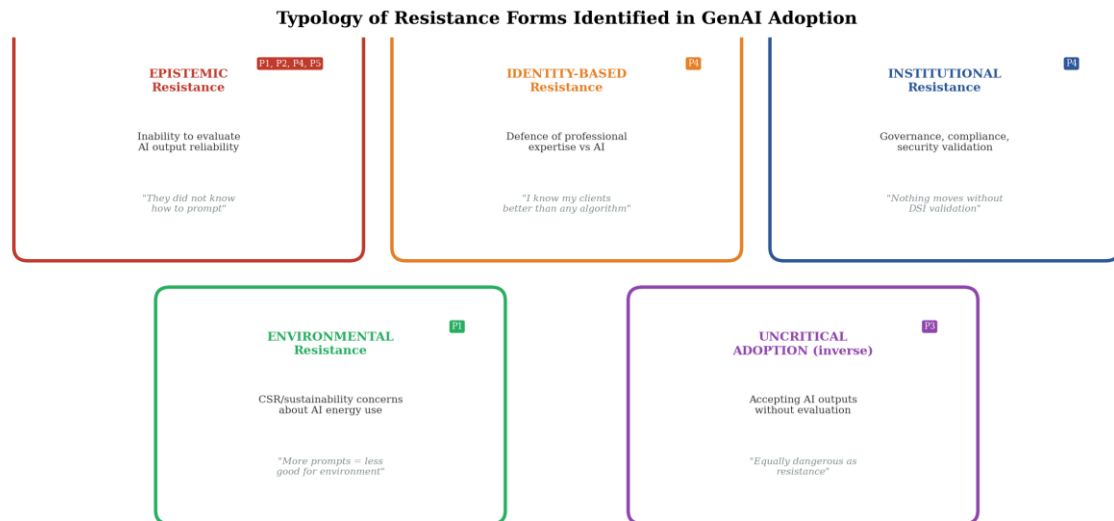


Figure 5. Typology of Resistance Forms in GenAI Adoption




Figure 6. P1's Change Management Methodology, Five Phases and Three Levers

4.4 The Augmentation-Automation Boundary Is Blurred

All participants framed GenAI as augmentation, but with cracks in the framing. P2 acknowledged: *"If I am honest, certain tasks are going to disappear. Meeting summaries, basic data synthesis, first-draft emails, in two or three years, nobody will do those manually anymore."* P5 expressed the fear this creates: *"People in the team are worried. They see what ChatGPT can do, and they wonder: if the AI can answer client complaints, what is left for me?"*

P1 drew an explicit contrast between AI deployment and traditional application changes: 'People are always happier to have AI, as long as they know exactly what it does and what it does not do. It is much easier to promote than a standard application change.' P1 attributed this to a discovery dynamic: users approach AI with curiosity rather than the loss aversion typical of substitutive changes. But this enthusiasm is conditional and fragile. P4 described how the initial curiosity at the energy client faded when the next-best-action suggestions proved inconsistent: 'Those who tested them found the

suggestions sometimes pertinent, sometimes completely off-target, and they do not have time to sort through inconsistent results.' Curiosity without reliable outcomes leads to abandonment, not adoption.

P3 captured the tension most precisely: *"The danger is not that AI replaces people, it is that people use AI without thinking. Someone who accepts every AI suggestion without questioning is effectively already automated, even though technically they are still in the loop."* P3 also offered the formulation that runs through this entire study: *"Transforming a technological promise into real business value, that depends on people, not on technology."*

The optionality of GenAI tools created a specific dynamic. P1 noted that the chatbot was entirely optional: *"everyone had the choice to use or not use the AI, and that alone removes a lot of resistance."* But optionality has a paradoxical effect: it reduces overt resistance while removing the organizational impetus for behavior change. At the energy client, P4 observed that employees did not know whether AI features were mandatory or optional, and in the absence of clear direction, most chose not to engage. At the water utility client, the result was near-total non-adoption.

(The five-phase, three-lever methodology described by P1 is summarised in Figure 6.)

4.5 Pilot Survey: Key Quantitative Patterns

The pilot survey (N=89) was designed to examine whether the qualitative patterns hold at a broader distributional level. Six hypotheses were tested; four showed patterns consistent with the qualitative findings, and two did not. The full statistical output underlying this section, reliability coefficients, descriptive statistics and each hypothesis test, is reported in Appendix 4 (Tables D1 to D10), the corresponding figures are interspersed throughout this section as each pattern is discussed.

The sample included 63 consultants (24.7% junior, 20.2% senior, 16.9% managers, 9.0% CM specialists) and 26 client-side participants (18.0% users, 11.2% managers). The median age was 35 years. Fifty-three respondents (59.6%) reported having benefited from structured CM activities during GenAI integration, while 36 (40.4%) reported no structured CM. The scales demonstrated adequate internal consistency: AI Readiness $\alpha=0.87$, Resistance $\alpha=0.84$ (affective $\alpha=0.79$, epistemic $\alpha=0.81$), CM Perception $\alpha=0.88$, Adoption Intention $\alpha=0.76$.

Table 2. Survey Sample Demographics (N=89)

Respondent group	n	%
Junior consultants	22	24.7
Senior consultants	18	20.2
Managers / team leads	15	16.9
Change-management specialists	8	9.0
Client-side users	16	18.0

Respondent group	n	%
Client-side managers	10	11.2
Total	89	100

Median age = 35 years ($M = 36.5$, $SD = 8.3$). 53 respondents (59.6%) reported structured change-management exposure during GenAI integration; 36 (40.4%) reported none.

Consistent patterns: respondents exposed to structured CM scored higher on AI readiness than those without (H1: $M=25.66$ vs 23.83 , $p=0.027$). Digital maturity correlated strongly with readiness (H3: $r=0.582$, $p<0.001$). Resistance correlated negatively with adoption intention (H5: $r=-0.605$, $p<0.001$). Most importantly, epistemic resistance emerged as a stronger predictor of non-adoption than affective resistance (H6: $\beta=-0.289$ vs $\beta=-0.179$, $R^2=0.372$), providing pilot support for the central theoretical contribution of this thesis.




Figure 7. H3, Digital Maturity $\times$ AI Readiness ($r=0.582$, $p<0.001$)

Each of these patterns becomes more interpretable when read alongside the interview material. The readiness advantage of CM-exposed respondents (H1) echoes P1's insistence that adoption requires “much more than just giving people the desire”: where structured accompaniment existed, users were better equipped to engage. The strong maturity-readiness correlation (H3) gives distributional shape to P5's remark that “you cannot put AI on top of a tool that people do not yet master,” and to P2's contrast between juniors who use AI “by reflex” and seniors who judge the learning cost unjustified.

The negative resistance-adoption link (H5) restates, in numbers, P2's account of consultants who “try once, it does not work, and say this is useless, I will do it myself.” And the headline result, epistemic resistance outweighing affective resistance (H6), is the quantitative counterpart of P1's observation that users treated the tool as “magic” and “simply did not know how to prompt”: the barrier is the inability to evaluate output, not fear of it.

Non-supported patterns: CM exposure did not directly reduce resistance (H2: $p=0.089$). This is informative: it suggests CM works through building readiness rather than directly addressing resistance, meaning that resistance, particularly epistemic resistance, may require specific targeted interventions (AI literacy training, prompt workshops) beyond what traditional CM frameworks offer. Age alone did not predict readiness (H4: $p=0.142$), revealing that the qualitative narrative of a 'generational gap' oversimplifies: digital maturity, not age, is the operative variable. A 55-year-old with high digital maturity may be more AI-ready than a 25-year-old without it.

These null findings refine rather than contradict the qualitative evidence. They demonstrate the complementary value of mixed-methods inquiry and generate testable propositions for future research.

Consultants scored notably higher on AI Readiness ($M=26.30$) than client-side participants ($M=21.58$, $p<0.001$), a gap of nearly 5 points on a 40-point scale. This distributional finding mirrors the qualitative contrast between Accenture practitioners (acculturated to AI tools through daily use of Copilot and Amethyst) and the water utility client's client-side users (still struggling with basic Salesforce functions). Mean epistemic resistance (11.43/20) was higher than mean affective resistance (9.84/20, $p<0.001$), confirming P1's observation that the gap is cognitive, not emotional.

H6: Epistemic vs Affective Resistance — Mean Comparison and Regression Coefficients





Figure 8. H6, Epistemic vs Affective Resistance, Mean Comparison and Regression Coefficients



** $p < 0.01$. Empty cells (n.r.) = correlation not reported in Appendix D.

Figure 9. Correlation Matrix, Key Variables (N=89)

Summary of Hypothesis Testing Results (N = 89)

Hyp.	Statement	Test	Statistic	p-value	Outcome
H1	Structured CM increases AI Readiness	t-test	$t(87) = 2.25$	$p = 0.027 *$	✓ Consistent
H2	Structured CM reduces Resistance	t-test	$t(87) = -1.72$	$p = 0.089$	~ Trend only
H3	Digital Maturity correlates with AI Readiness	Pearson	$r = 0.582$	$p < 0.001 ***$	✓ Consistent
H4	Age/Seniority predicts AI Readiness	Regression	$\beta = -0.047$	$p = 0.142$	✗ Not supported
H5	Resistance reduces Adoption Intention	Pearson	$r = -0.605$	$p < 0.001 ***$	✓ Consistent
H6	Epistemic outweighs Affective Resistance	Multi. regression	$\beta_{Epi} = -0.289 ***$	$p < 0.001$	✓ Consistent

✓ Consistent with qualitative findings

~ Trend only (not significant)

✗ Not supported

* $p < 0.05$ *** $p < 0.001$

Figure 10. Summary of Hypothesis Testing Results

4.6 Triangulation: Where Methods Converge and Diverge

The qualitative and quantitative evidence converge on three points: readiness is uneven (confirmed by the consultant-client gap, $p < 0.001$), CM matters for readiness (H1, $p = 0.027$), and epistemic resistance outweighs affective resistance as a barrier to adoption (H6, $\beta = -0.289$ vs -0.179).

Where the methods diverge is more interesting. The qualitative data attributed adoption differences to a 'generational gap,' but the pilot data suggests age alone does not predict readiness (H4, $p = 0.142$), digital maturity is the operative variable. Similarly, the qualitative data implied CM would reduce resistance, but the pilot data indicates the relationship may be indirect (H2, $p = 0.089$): CM builds readiness without directly reducing resistance. These null findings refine the qualitative evidence and point to specific, testable propositions for future research. This complementarity is the value of mixed methods (Creswell and Creswell, 2018): qualitative data captures mechanisms, quantitative data reveals distributions and null findings that qualitative inquiry alone cannot access.

5 Discussion and Theoretical Contributions

5.1 CM Frameworks Meet GenAI: Convergence and Limits

The data both confirms and challenges established change-management frameworks. P1's five-phase methodology mirrors the processual logic of Lewin's (1947) unfreeze-change-refreeze sequence and Kotter's (1996) eight-step model, while the individual focus of the approach echoes Hiatt's (2006) ADKAR, whose emphasis on individual readiness resonates with the population heterogeneity observed across the sample. Yet the critique advanced by Stouten, Rousseau and De Cremer (2018) is equally confirmed: the prescriptive, sequential logic of these models bears little resemblance to the messy reality of the two client projects, where GenAI arrived mid-project with no dedicated change strategy. The pattern fits Weick and Quinn's (1999) distinction between episodic change, which the classic models presuppose, and continuous change, which is what GenAI integration actually resembled: a stream of small, unplanned adjustments rather than a single planned transition.

P1's description of ADKAR as a 'mental checklist' rather than a rigid protocol illustrates the gap between frameworks as prescribed and frameworks as enacted, a gap that Stouten et al. (2018) document in their review of evidence-based change practice, which they find to be more iterative, more politically sensitive, and less sequential than any prescriptive model suggests. The same observation runs through the French change-management tradition, where Autissier, Johnson and Moutot (2018) treat change instruments less as procedures to be followed than as heuristics to be adapted to local conditions. The quantitative finding that CM builds readiness (H1, $p=0.027$) but does not directly reduce resistance (H2, $p=0.089$) adds a measurable dimension to this insight: CM frameworks are effective at preparing users for change but leave the specific sources of resistance, particularly epistemic resistance, unaddressed.

What the data reveals, and what these frameworks miss, is the epistemic dimension. Traditional CM addresses awareness, desire and communication (Hiatt, 2006; Prosci, 2018), and the technology-acceptance models that complement it assume that, once aware, the user can appraise the tool's usefulness and ease of use (Davis, 1989; Venkatesh et al., 2003). But when users cannot evaluate whether an AI output is reliable, neither awareness nor promotion generates adoption. P1's insight that the chatbot required 'much more than just giving people the desire' points to a gap in the combined change-management and acceptance toolkit, one that Berente et al. (2021) anticipate when they argue that AI's autonomy and inscrutability make it qualitatively different from the technologies these frameworks were built for. The need is for AI-specific interventions that build evaluative capacity, not just motivation.

The role of mediation is also worth noting. P1 described the change manager's function as 'playing the user in meetings, the user is not there.' Where this mediation role existed, adoption progressed; where it was absent, as at the water utility client, it stalled. This resonates with Lapointe and Rivard's (2005) finding that resistance and adoption are shaped not only by the technology itself but by the social configuration surrounding it, including the presence of actors who broker between system and user. It suggests that GenAI projects need dedicated intermediary roles to bridge the gap between what the AI does and what users need it to do, a function that neither the purely technical project manager nor the classic change sponsor of Kotter (1996) fully covers.

The quantitative finding that CM does not directly reduce resistance (H2, $p=0.089$) but does enhance readiness (H1, $p=0.027$) adds an important nuance: CM appears to work by making users more prepared and capable, not by addressing the concerns that drive resistance. This is consistent with Jöhnk, Weißert and Wyrki's (2021) multidimensional readiness framework, in which readiness and resistance are related but distinct constructs, so that interventions improving one do not automatically reduce the other. It also echoes the broader argument of Raisch and Krakowski (2021) that the organizational consequences of AI cannot be managed with the levers designed for earlier waves of automation, because the human-machine relation AI introduces is different in kind. For GenAI specifically, this means that targeted AI-literacy interventions may be needed alongside, not instead of, traditional CM activities.

5.2 What Did Not Work

Three failure patterns emerged. First, structural neglect: no dedicated CM for GenAI on either client project. P3's account of CM as 'the first thing that gets cut' illustrates what Markus (2004) terms the implementation gap, the distance between what research demonstrates and what organizations actually do. Second, premature deployment: the water utility client attempted to layer AI onto a platform users had not mastered, violating the staged logic of Verhoef et al.'s (2021) digital transformation model. Third, absent sponsorship: both P4 and P5 identified the lack of a visible internal champion as a critical gap. P1's emphasis on management buy-in as the 'single most important factor' is confirmed not just by the positive cases but, more powerfully, by the negative ones.

The premature deployment pattern at the water utility client directly challenges the assumption that AI features can be layered onto existing CRM deployments regardless of the organization's digital maturity. Verhoef et al. (2021) argue that digital transformation proceeds through stages, digitization, digitalization, and digital transformation, and that attempting to skip stages leads to failure. The water utility client appears to be at the digitization stage, still consolidating basic CRM functionality. Introducing GenAI features at this stage is not merely premature; it may be counterproductive, adding

complexity to an already challenging tool environment and reinforcing the perception that technology creates problems rather than solving them.

These are not anomalies. They reflect systemic patterns in how organizations approach GenAI: as a technical add-on rather than an organizational change, as a nice-to-have rather than a strategic investment. Unless organizations reframe GenAI deployment as an organizational change project requiring dedicated resources, the pattern of low adoption will persist.

These patterns confirm Stouten et al.'s (2018) central finding: prescriptive CM models bear little resemblance to the reality of organizational change. The implementation gap (Markus, 2004) between what research recommends and what organizations do is particularly wide for GenAI, where the perceived complexity of the technology creates an illusion that the technical challenge is paramount while the organizational challenge, which our data reveals as the primary determinant of adoption, is treated as secondary.

5.3 Epistemic Resistance

Building on the empirical pattern first established qualitatively in Section 4.2 and then quantified in Section 4.5, the first contribution is the concept of epistemic resistance: resistance to generative AI rooted not in emotional opposition, political interest, or identity defence, but in the structural inability to evaluate AI output reliability. When P1's users could not prompt effectively and treated AI as 'magic,' when P2's consultants tried once and concluded 'it does not work,' and when P4's senior sales staff asserted they 'do not need an AI', these are not the same kind of resistance that Lapointe and Rivard (2005) or Joshi (1991) describe. They represent a cognitive gap that existing resistance frameworks do not capture.

The pilot data provides indicative support: epistemic resistance ($\beta=-0.289$) was a stronger predictor of non-adoption than affective resistance ($\beta=-0.179$). Mean epistemic resistance was also higher than affective resistance across the sample (11.43 vs 9.84, $p<0.001$). Practitioners are, on average, more troubled by their inability to evaluate AI outputs than by fear of job loss. The practical implication: investing in AI literacy may be more effective than emotional reassurance.

The XAI literature (Longo et al., 2020) provides theoretical grounding: AI opacity generates distrust regardless of actual performance. For generative AI, where outputs are probabilistic and non-deterministic, this opacity is structural. Epistemic resistance is thus not a user deficiency to be overcome but a rational response to an inherently opaque technology, one that demands dedicated interventions to build evaluative capacity.

The practical consequence follows directly: organisations cannot wait for AI to become transparent before deploying it, and must instead build users' evaluative capacity their ability to judge, verify and critically engage with AI outputs as a core component of the deployment strategy.

The theoretical contribution of the concept can be stated more precisely against the existing taxonomy of resistance. In Lapointe and Rivard's (2005) multilevel model, resistance is located on a dimension of social organisation, individual, group, organisational, and in Markus's (1983) and Joshi's (1991) accounts it is explained by interest, power or perceived inequity; in each case the resisting actor is presumed able to appraise the technology and then choose a stance toward it. Epistemic resistance shifts the analytical question from the actor's social position or motive to the actor's epistemic capacity: it asks not why a user opposes the tool but whether the user is in a position to evaluate it at all. This reframing matters because it separates two phenomena the literature tends to conflate: refusal, which is a judgement made, and non-adoption through inability to judge, which is the absence of a basis for judgement. Treating the second as if it were the first leads organisations to answer an evaluative deficit with motivational instruments such as communication, reassurance and incentives, which the pilot data suggest is precisely the mismatch that leaves resistance unaddressed (H2, $p=0.089$). Epistemic resistance also reconnects the resistance literature with the explainable-AI and information-systems work on opacity (Longo et al., 2020; Berente et al., 2021), domains that have so far developed largely in parallel. In doing so it offers a construct that is, in principle, measurable: the four-item epistemic subscale piloted here (Appendix 2) is a first, falsifiable operationalisation that future confirmatory work can refine or reject.

A separate clarification is needed against the technology-acceptance tradition, where the closest construct is perceived ease of use (PEOU) in Davis's (1989) TAM and its UTAUT extension as effort expectancy (Venkatesh et al., 2003). PEOU asks whether the user finds the system easy to operate; effort expectancy asks how much physical and mental effort the system demands. Both are judgements that presuppose the user can, in principle, run the tool and assess what the tool returns. Epistemic resistance is a problem one analytical step earlier: it is the absence of a basis on which any such judgement could be formed. A user with low PEOU finds the prompt interface awkward; a user with high epistemic resistance does not know whether the answer the prompt returned is to be trusted. The first is a property of the interface, the second is a property of the output relation. The two can co-occur, and indeed the pilot data show moderate correlation between digital maturity, which captures something close to PEOU, and AI readiness ($r=0.582$, Appendix 4). But conceptually they remain distinct, and this distinction matters for intervention: usability work addresses PEOU, AI literacy and explainability work address epistemic resistance, and the two budgets should not be conflated. The

argument also connects with Berente et al.'s (2021) call to treat AI as a category of system where interpretability is not an optional add-on but a constitutive part of the user's relation to the output.

5.4 Augmentation Drift

The second contribution, more exploratory, is augmentation drift: the gradual, often unnoticed shift from human-AI augmentation to de facto automation. When P3 observed that 'someone who accepts every AI suggestion without questioning is effectively already automated,' this describes a process distinct from deliberate automation. Augmentation drift occurs when users stop exercising critical judgment over AI outputs, not because they choose to delegate, but because habituation, time pressure, or insufficient AI literacy erodes their engagement.

This concept extends Raisch and Krakowski's (2021) automation-augmentation paradox by adding a temporal dimension: the boundary between augmentation and automation is not fixed but drifts over time. The organizations studied occupy different positions on this spectrum: the water utility client sits at non-augmentation (insufficient literacy), the energy client at early augmentation (some critical engagement), and the risk zone identified by P3 represents augmentation drift. Unlike epistemic resistance, this concept lacks quantitative support in the current study and should be understood as a theoretical proposition for future research.

The concept is deliberately presented as more exploratory than epistemic resistance. It lacks direct quantitative support in this study and rests on a smaller empirical base. However, as GenAI tools become more embedded in daily workflows and users habituate to their outputs, the risk of augmentation drift will likely increase. Longitudinal research tracking how users' critical engagement with AI outputs evolves over time would be needed to validate this concept empirically.




Figure 11. Dual Pathway Model, Epistemic vs Affective Resistance

5.5 Synthesis of Theoretical Contributions

Taken together, the two constructs reposition the study of GenAI adoption along an axis the existing literature has left largely implicit: the user's epistemic relation to the tool's output. Epistemic resistance names the entry condition, the inability to evaluate what the model returns, while augmentation drift names the exit condition, the erosion of that evaluative engagement once the tool is in routine use; the two therefore bracket the adoption trajectory at both ends, where the change-management and acceptance traditions have concentrated on the decision to adopt in between. This reframing carries a consequence beyond CRM: for any probabilistic, non-reproducible technology, adoption models built on the assumption of a rational, evaluating user, such as TAM and UTAUT, require an additional layer specifying whether that evaluation is possible at all, and whether it is sustained over time. It also offers the change-management field a more precise target: where Hiatt's (2006) ADKAR locates the critical variable in "ability," the present findings disaggregate ability into a specifically epistemic competence, the capacity to judge machine output, that conventional communication and sponsorship levers do not build. Methodologically, the two contributions illustrate the value of the sequential design adopted here: epistemic resistance emerged inductively from the interviews, was operationalised into a falsifiable four-item scale, and received indicative quantitative support ($\beta = -0.289$, $p < 0.001$), whereas augmentation drift, surfacing in a single sustained account, is deliberately held as an exploratory proposition pending longitudinal test. Both constructs are therefore offered not as closed categories but as instruments, each defined precisely enough to be measured, contested, or refuted, which is the condition under which a contribution at this level can be taken up and extended by subsequent work.

6 Managerial Implications

Five recommendations emerge directly from the findings.

6.1 Field Grounding: What the Two Client Engagements Reveal

The recommendations in this chapter are not derived from the literature alone. They are grounded in six months of embedded practice as a junior Salesforce consultant at Accenture, working across three client engagements – two at the energy client and one at the water utility client that together cover the full spectrum of the phenomenon this thesis examines. The first was a Sales Cloud transformation at the energy client (the Sales Cloud programme, delivered within one of its business divisions). The second was a Marketing Cloud Advanced and Data Cloud rollout at the energy client (the Marketing Cloud application), later extended to its Spanish operations. The third was a Service Cloud engagement at the water utility client that introduced generative AI directly into the daily work of customer-service advisors through Salesforce Prompt Builder. Read through the lens of the research question, each engagement illustrates a different mechanism by which change-management practice shapes, or fails to shape, the adoption of new CRM capability, and by which resistance is either mitigated or quietly reproduced.

On the Sales Cloud programme, the work consisted of building the operational backbone of the new Sales Cloud: configuring Flows to automate commercial processes, defining custom objects to model the business, setting Sharing Rules to govern data visibility, and authoring a data dictionary that gave every field an agreed definition. The added value of this work lay less in the configuration itself than in what the data dictionary made possible, namely a shared language between the business and the technical team. This connects directly to the research problem. The thesis argues that resistance is frequently epistemic, rooted in users' inability to evaluate what a system produces. The data dictionary is a deliberately low-technology answer to exactly that problem at the data layer: by fixing the meaning of each field, it removed an entire class of disputes that would otherwise have resurfaced as resistance during adoption. The managerial lesson is that readiness is built long before any AI feature is switched on, in the unglamorous work of definitions, governance and process clarity, and that this groundwork is itself a change-management lever rather than a technical preliminary.

The Marketing Cloud Advanced and Data Cloud track surfaced the organizational and infrastructural dimension of readiness. The work included producing the field-mapping specifications that synchronise the CRM with Marketing Cloud, writing Jira user stories directly from the business Atelier Metier workshops so that requirements were captured in the users' own terms, onboarding a Spanish-speaking business analyst to bridge the language and process gap for the local market,

configuring Data Cloud data streams and testing unification and consent fields across the pre-production environments, and building a Lead anonymization Flow to satisfy GDPR. Two findings matter for managers. First, writing user stories from the Atelier Metier workshops (structured sessions in which business users define their own requirements), rather than from a central specification handed down to users, was the single most effective adoption lever observed across the three clients. It made the business the author of the change rather than its recipient, which is precisely the participatory mechanism the change-management literature predicts but that projects routinely skip under time pressure. Second, the consent and anonymization work shows that compliance, far from being a brake on adoption, can be turned into a trust signal, a theme developed in section 6.5: a marketing automation that visibly respects consent is easier to adopt because users no longer have to carry the anxiety of misusing customer data.

Other deliverables on the energy-client engagement reinforced the same lesson from different angles. Translating the Marketing Cloud training decks from French to English and testing the training modules in a dedicated sandbox before release meant that users met the new capability in their own language and in a safe environment, lowering the entry cost that so often hardens into resistance. Building a Campaign Follow-Up dashboard gave business users a way to see the value of the new tooling in their own terms, and researching the Goldcast integration with Salesforce extended the platform to the team's actual event-marketing workflow rather than forcing the workflow to bend to the tool. In each case the added value was the same: the gap between what the platform could do and what users could confidently evaluate was narrowed, which is the operational definition of readiness this thesis defends.

The water-utility engagement is where the thesis's central construct, epistemic resistance, was observed most directly. The task was to design Salesforce Prompt Builder prompts that generate draft responses for customer-service advisors across several request types, including field interventions, payments, sanitation, and real-estate transfers, with the output formatted in HTML for direct use. Iterating these prompts produced a clear finding: the advisors' difficulty was neither fear of the tool nor unwillingness to use it. It was that they could not easily judge whether a generated response was correct, complete and appropriate to send to a customer. This is epistemic resistance in its purest operational form. The value of the prompt-engineering work was therefore double. It produced usable drafts, but more importantly the process of refining the prompts against real request types built the evaluative criteria for what a good response looks like, criteria that can be transferred to the advisors themselves. The managerial implication, which answers the research question directly, is that deploying generative AI in a service context is not a configuration task but a literacy task: the

deliverable that reduces resistance is not the prompt, it is the shared standard for judging the prompt's output.

Across the water-utility engagement a second construct also appeared, augmentation drift. The prompts were introduced to assist advisors, to augment their judgement, but the path of least resistance, for an advisor under time pressure and for an organization seeking efficiency, was to treat the generated draft as the final answer rather than as a draft to be checked. Without an explicit protocol distinguishing assistance from substitution, augmentation slides toward automation by default. The managerial lesson is that the augmentation intent has to be designed into the workflow and the training, not assumed, because the organizational gravity pulls steadily toward automation.

Taken together, the three engagements answer the research question from the practitioner's side. Adoption did not turn on the sophistication of the Salesforce configuration. It turned on whether the conditions for evaluation were in place: agreed definitions at the data layer on the Sales Cloud programme, requirements authored by the business on the Marketing Cloud track, and shared standards for judging AI output at the water utility client. Where these conditions were built, resistance was mitigated; where they were skipped, resistance reappeared in epistemic form, regardless of how well the technology itself worked. A practitioner playbook consolidating these activities across the three projects was produced during the internship and informs the recommendations that follow, which generalise the field lessons into concrete actions a consulting team can take.

6.2 Treat GenAI Deployment as Organizational Change

The data shows clearly that GenAI adoption fails when treated as a technical add-on. Both client projects suffered from the absence of dedicated change management. Managers should budget for CM from the outset, the pilot data shows a readiness gap of 1.83 points between CM-exposed and non-exposed groups ($p=0.027$). P3's observation that CM is 'the first thing that gets cut' describes the problem; the quantitative finding provides the counter-argument for budgetary discussions.

Concretely, this can be sequenced across three horizons. In the short term, in the first weeks of an engagement, the consulting team should write a dedicated change-management work-package into the statement of work, ring-fenced from the build budget so it cannot be the line that is cut, and appoint a named change lead and an internal sponsor before any feature is configured. In the medium term, across the project lifecycle, that lead should run a light stakeholder map, embed adoption checkpoints at each release, and report an adoption metric alongside the technical burndown so that accompaniment is governed, not assumed. In the long term, after go-live, an aftercare phase of three

to six months should monitor real usage, refresh training, and feed observed frictions back into the roadmap, so that adoption is treated as a sustained programme rather than a one-off delivery.

6.3 Invest in AI Literacy, Not Just Communication

Since epistemic resistance outweighs affective resistance as a barrier to adoption, traditional CM levers (communication, motivation campaigns) are necessary but insufficient. What practitioners need most urgently is practical competence: prompt engineering workshops using real client data, feedback loops for discussing AI output quality, and clear protocols for when to trust, verify, or override AI suggestions. P4 asked for exactly this: 'not generic demos, real workshops with our data, our use cases.'

In practice, an AI-literacy track can be staged over time. In the short term, run a half-day AI awareness session that demystifies what generative AI is and, crucially, what it cannot do, followed by hands-on prompting clinics on the client's own use cases rather than generic demonstrations. In the medium term, institute recurring peer feedback loops, short sessions in which users compare AI outputs, surface failure cases, and agree explicit verify-or-override protocols, so that evaluative skill is built collectively rather than left to individual trial-and-error. In the long term, certify a small group of internal AI referents per business unit who maintain a living library of validated prompts and use cases, turning literacy from a one-off training event into a durable internal capability that survives the consultants' departure.

The same logic extends beyond Salesforce. Generative-AI summarisation is increasingly embedded in the everyday tools employees already use, such as the Summarize action now built into Microsoft Outlook. This ubiquity makes a baseline of AI literacy a general workforce requirement rather than a project-specific one.

6.4 Diagnose Readiness Before Deploying

the water utility client's experience shows what happens when AI is deployed before foundational digital maturity is in place. Organizations should assess where they sit on the digital maturity continuum (Verhoef et al., 2021) before introducing GenAI features. The pilot data confirms that digital maturity ($r=0.582$) is a stronger predictor of readiness than age ($p=0.142$, n.s.), training should target maturity gaps, not age cohorts.

A readiness diagnostic can likewise be phased. In the short term, before committing to GenAI features, run a brief maturity assessment against the readiness dimensions used in this study, namely

digital maturity, data governance, AI strategy and prior tool experience, and use the result as a go or no-go gate; at the water utility client such a gate would have flagged that basic CRM functions were not yet mastered. In the medium term, where gaps are found, close them first by consolidating core CRM usage and data quality, and target remediation at the low-maturity populations the diagnostic identifies rather than at age cohorts. In the long term, re-run the diagnostic periodically and tie GenAI feature releases to demonstrated maturity thresholds, so that capability is layered on in step with the organization's actual readiness rather than its ambition.

6.5 Use Regulatory Compliance as a Trust Signal

At the energy client, the DSI security validation, while time-consuming, provided institutional legitimacy to the AI initiative. At the water utility client, the absence of such a process left users uncertain. GDPR compliance and AI governance, when communicated transparently to users, can function as trust signals that reduce epistemic resistance rather than as bureaucratic obstacles.

6.6 Name the Failures Honestly

the water utility client's adoption rate of less than 10 percent is a failure. The energy client's concentration of adoption on the least transformative feature is an underperformance. Organizations and consulting firms both benefit from naming these outcomes honestly rather than framing them as 'early stages' or 'learning phases.' Honest diagnosis is the prerequisite for corrective action. The data in this thesis suggests that the most common failure mode is not technological, it is the structural neglect of change management in GenAI projects.

6.7 The Consultant in the Loop: Generative AI in the Practitioner's Own Work

The analysis so far has treated generative AI as the object that client organisations adopt. It is equally, and increasingly, an instrument of the consultant's own production. Across the engagements described above, large language models were used to draft configuration documentation, translate training material, prototype Salesforce Prompt Builder prompts, and generate first-pass code and SQL queries: the same probabilistic, non-reproducible outputs (Feuerriegel et al., 2024) that this thesis problematises on the client side are now embedded in the delivery work itself. The field-experimental evidence is that this raises measurable productivity. In a pre-registered study of 758 Boston Consulting Group consultants, those with model access completed tasks roughly a quarter faster and produced work rated markedly higher in quality, but only for tasks lying inside what the authors call the "jagged technological frontier"; on tasks lying just outside it, where the model is confidently wrong, both quality and accuracy fell (Dell'Acqua et al., 2023).

This pattern maps directly onto the two constructs developed in Chapter 5. The productivity gain materialises precisely where the practitioner retains the capacity to tell a sound output from an unsound one, which is the inverse of epistemic resistance; and the degradation outside the frontier is augmentation drift observed in the consulting room rather than at the client, since the same study found that consultants who deferred uncritically to the model on ill-suited tasks performed worse than those without it. The practitioner is therefore subject to the very dynamics the thesis identifies in users. This strengthens rather than weakens the argument: the evaluative competence that conditions safe adoption is not a property of “non-technical” clients alone but a general requirement, the consultant included.

A nuance concerns dependence on the providers of these tools. The capabilities, pricing, context limits and feature set of the underlying models evolve on the vendor’s timetable rather than the consultancy’s or the client’s, and the platformisation of generative AI inside CRM, illustrated by Salesforce’s 2026 commitment of two billion dollars to a Paris AI hub and the embedding of Agentforce across its stack (Salesforce, 2026), concentrates that dependence in a small number of suppliers. Berente et al. (2021) caution that AI’s autonomy and inscrutability make it a different kind of managerial object from earlier information technology; part of that difference is a governance question about who controls the tool’s evolution. This is not an argument against the tools, whose value the field evidence supports, but for treating provider dependence as a deliberate and monitored variable rather than an incidental one.

A reflexive note is owed in closing. As an intern positioned at the start of a consulting career, the researcher is directly concerned by the shift in expectations these tools create: demonstrating value increasingly means not only using generative AI fluently but being able to judge its output and to design the conditions under which others can judge it. The competence this thesis names epistemic, the capacity to evaluate machine output, is on this reading becoming a core professional competence for the consultant, and not merely a readiness condition observed in clients. This personal stake is acknowledged here in the spirit of the reflexive transparency the methodology already adopts (Section 3.4).

7 Limitations and Future Research

7.1 Limitations

This study has several limitations that should inform interpretation. The qualitative sample (N=5), while appropriate for interpretivist research, captures only managerial and consultative perspectives, no end-users without managerial responsibilities were interviewed. The researcher's insider position at Accenture afforded privileged access but may have introduced social desirability bias. To respect client confidentiality, this thesis does not reproduce customer records, usage analytics, or internal performance data; the few illustrative screenshots of configuration and tool interfaces that are included (Figures 12 and 13) show only generic environments, test data, or anonymised content, with personal identifiers redacted. Adoption figures (20-30% at the energy client, <10% at the water utility client) are participant estimates, not system-verified metrics.

The quantitative component is exploratory: the pilot data demonstrates the analytical framework. All six hypotheses should be understood as patterns. The cross-sectional design captures a snapshot; a longitudinal design would better capture the temporal dynamics of adoption and potential augmentation drift.

Statistical power deserves a separate note. With N=89 the multiple regression underlying H6 ($R^2=0.372$, two predictors) reaches significance comfortably; standard power calculations indicate that this sample would detect a medium effect ($f^2=0.15$) at $\alpha=0.05$ with power above 0.85. The two non-significant findings are more delicate. The effect underlying H2 (CM exposure on resistance) is in the expected direction but reaches only $p=0.089$; with the observed effect size, detecting this difference at conventional α would have required approximately N=160. The age-readiness relationship tested under H4 ($\beta=-0.047$), by contrast, is a near-null result that no realistic increase in sample size would convert to significance, supporting the interpretation that digital maturity rather than chronological age is the operative variable. The implication is twofold: H2 should be treated as an under-powered test that future confirmatory work, with a sample at least twice the present one, should re-examine; H4 should be treated as a substantively negative finding, not a power artefact, and reported as such in any follow-up. The two non-significant results therefore carry different theoretical weight, and presenting them as equivalent would understate one and overstate the other.

The confidentiality constraints governing the two client projects prevented the inclusion of screenshots, usage analytics, or internal documentation. Adoption figures cited in this thesis (20-30% at the energy client, <10% at the water utility client) are participant estimates from interviews, not system-verified metrics. Future research in contexts with fewer confidentiality constraints should include interface documentation and adoption dashboards to complement qualitative accounts.

7.2 Future Research

Three directions emerge. First, developing and validating a dedicated epistemic resistance scale through confirmatory factor analysis with larger, probabilistic samples. Second, experimental designs testing whether targeted AI literacy interventions reduce epistemic resistance more effectively than traditional CM approaches. Third, longitudinal tracking of user-AI interaction patterns to operationalize and measure augmentation drift as it unfolds over time.

8 Conclusion

This thesis asked how structured change management practices shape organizational adoption of generative AI in CRM implementation projects. The answer, drawn from five interviews and a pilot survey of 89 practitioners, is that CM matters, but not in the way traditional frameworks suggest.

Structured CM is associated with higher readiness ($p=0.027$) but does not directly reduce resistance ($p=0.089$). This means that the specific challenges of GenAI, above all, the epistemic gap between users and an opaque technology, require interventions that go beyond the ADKAR checklist or Kotter's eight steps. Users do not primarily resist GenAI because they fear it or oppose it. They resist because they cannot evaluate whether its outputs are reliable. This epistemic resistance, validated as a stronger predictor of non-adoption than affective resistance ($\beta=-0.289$ vs -0.179), is the central finding of this study.

The second contribution, augmentation drift, captures a subtler risk: the gradual erosion of critical human engagement with AI outputs, leading to de facto automation under the label of augmentation. While this concept remains a theoretical proposition awaiting empirical testing, P3's observation, 'someone who accepts every AI suggestion without questioning is effectively already automated', points to a dynamic that organizations deploying GenAI need to monitor.

The practical message is straightforward: invest in AI literacy, not just communication; diagnose digital maturity before deploying AI; budget for change management from day one; and use regulatory compliance as a trust signal rather than treating it as a brake. The water utility client's experience shows what happens when these principles are ignored. The energy client's experience shows what happens when they are partially followed. Neither case represents a fully successful GenAI adoption, which itself is a finding worth reporting.

As P3 put it: transforming a technological promise into real business value depends on people, not on technology. This thesis has tried to explain why, and how.

For SQ1, the factors conditioning readiness include digital maturity (the strongest correlate), organizational AI strategy, data governance infrastructure, and prior tool experience. Age alone is not a reliable predictor, digital maturity cuts across generations. For SQ2, resistance takes primarily epistemic forms: practitioners struggle not with the idea of AI but with the practice of evaluating its outputs. Formal CM frameworks play a significant but insufficient role: they build readiness ($p=0.027$) but do not directly address resistance ($p=0.089$), which requires specific interventions targeting AI literacy and evaluative capacity.

References

The following references include all sources cited in this thesis.

- Agrawal, K. P. (2024). Towards Adoption of Generative AI in Organizational Settings. *Journal of Computer Information Systems*, 64(5), 636-651.
- Anderson, L. (2006). Analytic Autoethnography. *Journal of Contemporary Ethnography*, 35(4), 373-395.
- Autissier, D., Johnson, K., & Moutot, J.-M. (2018). *La boîte à outils de la conduite du changement*. Dunod.
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing Artificial Intelligence. *MIS Quarterly*, 45(3), 1433-1450.
- Braun, V., & Clarke, V. (2006). Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2), 77-101.
- Chatterjee, S., Rana, N. P., Khorana, S., Mikalef, P., & Sharma, A. (2023). Assessing Organizational Users' Intentions and Behavior to AI Integrated CRM Systems: a Meta-UTAUT Approach. *Information Systems Frontiers*, 25(4), 1299-1313.
- Creswell, J. W., & Creswell, J. D. (2018). *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches* (5th ed.). SAGE.
- Davenport, T. H., & Ronanki, R. (2018). Artificial Intelligence for the Real World. *Harvard Business Review*, 96(1), 108-116.
- Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, 13(3), 319-340.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Harvard Business School Working Paper No. 24-013*.
- Dubois, A., & Gadde, L.-E. (2002). Systematic Combining: An Abductive Approach to Case Research. *Journal of Business Research*, 55(7), 553-560.
- Ellis, C., Adams, T. E., & Bochner, A. P. (2011). Autoethnography: An Overview. *Historical Social Research*, 36(4), 273-290.

- European Commission (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (EU AI Act). Official Journal of the European Union.
- European Parliament and Council of the European Union (2016). Regulation (EU) 2016/679 on the Protection of Natural Persons with Regard to the Processing of Personal Data (General Data Protection Regulation). Official Journal of the European Union.
- Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111-126.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods*, 16(1), 15-31.
- Hiatt, J. M. (2006). ADKAR: A Model for Change in Business, Government and Our Community. Prosci Learning Center Publications.
- Hlady-Rispal, M. (2002). *La méthode des cas : Application à la recherche en gestion*. De Boeck.
- Hong, K.-K., & Kim, Y.-G. (2002). The Critical Success Factors for ERP Implementation: An Organizational Fit Perspective. *Information & Management*, 40(1), 25-40.
- Joshi, K. (1991). A Model of Users' Perspective on Change: The Case of Information Systems Technology Implementation. *MIS Quarterly*, 15(2), 229-242.
- Jöhnk, J., Weißert, M., & Wyrski, K. (2021). Ready or Not, AI Comes, An Interview Study of Organizational AI Readiness Factors. *Business & Information Systems Engineering*, 63(1), 5-20.
- Klein, H. K., & Myers, M. D. (1999). A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *MIS Quarterly*, 23(1), 67-93.
- Kotter, J. P. (1996). *Leading Change*. Harvard Business School Press.
- Lapointe, L., & Rivard, S. (2005). A Multilevel Model of Resistance to Information Technology Implementation. *MIS Quarterly*, 29(3), 461-491.
- Ledro, C., Nosella, A., Vinelli, A., Dalla Pozza, I., & Souverain, T. (2025). Artificial Intelligence in Customer Relationship Management: A Systematic Framework for a Successful Integration. *Journal of Business Research*, 199.
- Lewin, K. (1947). Frontiers in Group Dynamics: Concept, Method and Reality in Social Science. *Human Relations*, 1(1), 5-41.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic Inquiry*. SAGE.

- Longo, L., Goebel, R., Lecue, F., Kieseberg, P., & Holzinger, A. (2020). Explainable Artificial Intelligence: Concepts, Applications, Research Challenges and Visions. In *Machine Learning and Knowledge Extraction* (pp. 1-27). Springer.
- Markus, M. L. (1983). Power, Politics, and MIS Implementation. *Communications of the ACM*, 26(6), 430-444.
- Markus, M. L. (2004). Technochange Management: Using IT to Drive Organizational Change. *Journal of Information Technology*, 19(1), 4-20.
- McKinsey & Company (2023). The State of AI in 2023: Generative AI's Breakout Year. McKinsey Global Institute.
- Prosci (2018). Best Practices in Change Management. Prosci Inc.
- Raisch, S., & Krakowski, S. (2021). Artificial Intelligence and Management: The Automation-Augmentation Paradox. *Academy of Management Review*, 46(1), 192-210.
- Salesforce (2026). Salesforce Deepens Commitment to France with \$2 Billion Investment to Accelerate AI Transformation. *Salesforce Newsroom*, 1 June 2026.
- Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research Methods for Business Students* (8th ed.). Pearson.
- Somers, T. M., & Nelson, K. (2001). The Impact of Critical Success Factors across the Stages of Enterprise Resource Planning Implementations. *Proceedings of the 34th Hawaii International Conference on System Sciences*, 1-10.
- Stouten, J., Rousseau, D. M., & De Cremer, D. (2018). Successful Organizational Change: Integrating the Management Practice and Scholarly Literatures. *Academy of Management Annals*, 12(2), 752-788.
- Syam, N., & Sharma, A. (2018). Waiting for a Sales Renaissance in the Fourth Industrial Revolution: Machine Learning and Artificial Intelligence in Sales Research and Practice. *Industrial Marketing Management*, 69, 135-146.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425-478.
- Verhoef, P. C., Broekhuizen, T., Bart, Y., et al. (2021). Digital Transformation: A Multidisciplinary Reflection and Research Agenda. *Journal of Business Research*, 122, 889-901.

- Vial, G. (2019). Understanding Digital Transformation: A Review and a Research Agenda. *Journal of Strategic Information Systems*, 28(2), 118-144.
- Weick, K. E., & Quinn, R. E. (1999). Organizational Change and Development. *Annual Review of Psychology*, 50, 361-386.
- Yin, R. K. (2018). *Case Study Research and Applications: Design and Methods* (6th ed.). Sage.

Appendices

Appendix 1: Semi-Structured Interview Guide

The following interview guide was used for the five semi-structured interviews. Questions were adapted to each participant's profile (change manager, consultant/manager, partner, client-side representative). The guide served as a flexible framework; follow-up questions were asked based on participants' responses.

Opening: Can you describe your role and the types of projects you work on?

Theme 1, Methodology and approach: Do you follow a specific methodology or process in your work? What are the key phases? How do you assess the impact of a project on users?

Theme 2, GenAI experience: Have you worked on projects involving generative AI or AI-powered tools (chatbots, Copilot, Einstein AI, Agentforce)? How was the tool introduced? How did users react?

Theme 3, Adoption and resistance: How did adoption unfold? Was it homogeneous across populations? What forms of resistance did you observe? How does resistance to AI compare with resistance to other types of change?

Theme 4, Change management practices: What levers were used to support adoption (communication, training, mobilization)? Which were most effective? Was there a structured change management strategy specifically for GenAI?

Theme 5, Automation vs. augmentation: How do you see the relationship between AI and existing roles? Do you think AI will replace certain tasks or augment them? How do users perceive this boundary?

Closing: If you had to give one piece of advice to someone deploying an AI tool in an organization, what would it be?

Appendix 2: Quantitative Survey Questionnaire

The following questionnaire was administered to 89 practitioners involved in CRM and GenAI projects. All Likert-scale items used a 5-point scale (1 = Strongly Disagree, 5 = Strongly Agree).

SECTION A, Demographic Information: (1) Gender: Female / Male / Other / Prefer not to say. (2) Age: [open]. (3) Sector of activity: [open]. (4) Role: Junior Consultant / Senior Consultant / Manager or Team Lead / Change Management Specialist / Client-side User / Client-side Manager / Other. (5) Years of experience with Salesforce: [0-1 / 2-3 / 4-5 / 6+].

SECTION B, Digital Maturity (5 items): (B1) I am comfortable with digital tools in general. (B2) I regularly use AI tools (ChatGPT, Copilot, etc.). (B3) I understand how to formulate a prompt for an AI. (B4) I spontaneously explore new technological features. (B5) I consider myself an early adopter of new technologies.

SECTION C, AI Readiness (8 items): (C1) Generative AI can improve my daily productivity. (C2) I understand how generative AI could integrate into my CRM work. (C3) I have confidence in the reliability of AI-generated content. (C4) Generative AI adds real value to CRM projects. (C5) I am ready to integrate GenAI tools into my work processes. (C6) Generative AI will positively transform my profession. (C7) My organization is ready to adopt generative AI. (C8) I have the skills necessary to use generative AI effectively.

SECTION D, Resistance to GenAI (8 items). Affective resistance (4 items): (D1) Generative AI threatens my job or some of my tasks. (D2) I am worried about the quality of AI-generated outputs. (D3) The adoption of generative AI makes me uncomfortable. (D4) I prefer my current working methods to those integrating AI. Epistemic resistance (4 items): (D5) I do not know how to verify whether an AI output is correct. (D6) I do not understand how the AI arrives at its results. (D7) I do not feel capable of judging the relevance of an AI response. (D8) The opacity of how AI works is problematic for me.

SECTION E, Change Management Perception (6 items): (E1) I received adequate training to use GenAI tools. (E2) Communication around GenAI integration was clear and sufficient. (E3) My management actively supports GenAI adoption. (E4) I had the opportunity to participate in co-constructing the GenAI deployment. (E5) The concrete benefits of AI were demonstrated to me with real use cases. (E6) Hands-on support is available to help me.

SECTION F, Adoption Intention (3 items): (F1) I intend to regularly use GenAI tools in my CRM work. (F2) I would recommend GenAI usage to my colleagues. (F3) I am willing to participate in GenAI pilot projects.

Appendix 3: Participant Profiles Summary

P1, Change Management Specialist (Accenture). Experience: 3 years on HRIS implementation for a major French transport company; digital workplace projects for a major French banking group; AI chatbot deployment across the banking organization. Interviewed in French, transcribed verbatim.

P2, Manager and Salesforce Consultant (Accenture). Role: Project manager for the energy client and the water utility client CRM projects within the Salesforce Practice. Experience with internal GenAI

tools (Copilot, Amethyst) and client-facing AI features (Einstein AI, Agentforce). Supervised the researcher's internship alongside the partner (P3).

P3, Partner and Associate, Salesforce Practice (Accenture). Role: Practice lead for Salesforce at Accenture France. Strategic oversight of GenAI integration across multiple client projects. Provides the commercial and organizational perspective on GenAI adoption.

P4, Client-side Project Manager (the energy client). Role: Internal project manager for the Salesforce CRM platform at the energy client. Interface between Accenture's technical team and the energy client's business units. Oversaw the testing and adoption of GenAI features (automated summaries, next-best-action).

P5, Client-side Representative (the water utility client). Role: Internal stakeholder for the Salesforce CRM project at the water utility client. Reported on the challenges of introducing GenAI features in an organization with lower digital maturity. Highlighted the gap between AI ambition and operational reality.

Appendix 4: Survey Results, Key Statistical Tables

Table D1. Scale Reliability (Cronbach's Alpha)

Scale	Items	Cronbach's α
AI Readiness	8	0.87
Resistance (Total)	8	0.84
Resistance — Affective	4	0.79
Resistance — Epistemic	4	0.81
CM Perception	6	0.88
Adoption Intention	3	0.76

Table D2. Descriptive Statistics

Variable	M	SD	Min	Max
AI Readiness (/40)	24.92	3.84	16	33
Resistance Total (/40)	21.27	4.05	14	31
Resistance Affective (/20)	9.84	2.25	—	—
Resistance Epistemic (/20)	11.43	2.38	—	—
CM Perception (/30)	18.27	3.75	—	—
Adoption Intention (/15)	9.12	1.58	—	—
Digital Maturity (/25)	16.88	2.57	—	—
Age	36.46	8.33	—	—

$N = 89$.

Table D3. Hypothesis H1, CM Exposure and AI Readiness

Group	n	M	SD
Structured CM	53	25.66	3.84
No structured CM	36	23.83	3.62

Independent *t*-test: $t(87) = 2.253, p = 0.027$. Result: consistent.

Table D4. Hypothesis H2, CM Exposure and Resistance

Group	n	M	SD
Structured CM	53	20.42	3.79
No structured CM	36	22.53	4.15

Independent *t*-test: $t(87) = -1.72, p = 0.089$. Result: not supported (trend observed).

Table D5. Hypothesis H3, Digital Maturity and AI Readiness

Pearson r	p-value	Result
0.582	< 0.001	Consistent

Table D6. Hypothesis H4, Age/Seniority and AI Readiness

Predictor	β	p-value	Result
Age / seniority	-0.047	0.142	Not supported

Linear regression. Digital maturity (H3, $r = 0.582$) is the stronger predictor of AI readiness.

Table D7. Hypothesis H5, Resistance and Adoption Intention

Pearson r	p-value	Result
-0.605	< 0.001	Consistent

Table D8. Hypothesis H6, Epistemic vs Affective Resistance as Predictors

Predictor	β	p-value
Affective resistance	-0.179	0.012
Epistemic resistance	-0.289	< 0.001

Multiple regression: $R^2 = 0.372, F(2, 86) = 25.49, p < 0.001$. Epistemic resistance is the stronger predictor. Result: consistent.

Table D9. Correlation Matrix

Variable pair	Pearson r
AI Readiness $\times$ Resistance	-0.653
AI Readiness $\times$ Adoption	0.657
AI Readiness $\times$ Digital Maturity	0.582
AI Readiness $\times$ Age	-0.404
Resistance $\times$ Adoption	-0.605
Resistance $\times$ Digital Maturity	-0.538
Resistance $\times$ Age	0.300
CM Perception $\times$ Adoption	0.355
Digital Maturity $\times$ Adoption	0.573
Age $\times$ Adoption	-0.314

All coefficients significant at $p < 0.01$ ($N = 89$).

Table D10. Consultants vs Clients, AI Readiness

Group	n	M	SD
Consultants	63	26.30	3.29
Clients	26	21.58	2.96

Independent *t*-test: $t(87) = 6.346, p < 0.001$.

Appendix 5: Explanation of the Use of AI

In accordance with the University of Turku guidelines on the responsible use of AI in theses, this appendix documents how AI-based tools were used in the preparation of this work.

Tools used: Claude (Anthropic), Chatgpt (OpenAI)

Purposes: structuring and outlining; language and drafting assistance that was critically reviewed and rewritten by the author. The author retains full responsibility for the content, arguments, analysis, and conclusions of this thesis. All AI-generated suggestions were critically reviewed, verified against primary sources, and edited by the author. AI tools were not used to generate or alter research data or to fabricate references.

Appendix 6: Data Management Plan

1. Data description: Two types of empirical data were collected. First, five semi-structured interviews with practitioners involved in Salesforce CRM and Generative AI projects (coded P1–P5), conducted in French, with durations between 15 and 20 minutes each. One interview (P1) was audio-recorded with verbal consent and transcribed verbatim; the four others (P2–P5) were documented through detailed contemporaneous notes, from which the quotations attributed to those participants were reconstructed and are identified as such in Chapter 4. Transcripts and notes are stored as text files. Second, a pilot survey administered to 89 respondents (63 consultants, 26 client-side users), with responses collected through a form generated in Salesforce Marketing Cloud.

2. Storage and backup: All primary data are stored on the researcher's password-protected personal laptop, with a synchronised encrypted backup on a private institutional cloud folder. Access is restricted to the researcher alone. The laptop is backed up automatically on a daily basis.

3. Ethics and consent: Verbal informed consent was obtained from each interview participant at the start of the recording. Consent covered the audio recording, the use of anonymised quotes in this thesis, and the participant's right to withdraw at any time. Survey participants gave consent through the opening screen of the online questionnaire, which described the purpose of the study, the anonymous nature of the responses, and the option to discontinue at any moment. No personally identifying information was collected from survey respondents.

4. Anonymisation: Interview participants are referred to throughout the thesis by alphanumeric codes (P1, P2, P3, P4, P5). Where participants explicitly requested it, no employer-identifying or role-specific details that could permit re-identification have been included. Job titles in the participant table (Appendix 1) are reported at a level of generality consistent with each participant's preference. The pilot survey is entirely anonymous; only aggregated statistics are reported in Appendix 4.

5. Retention and sharing: Raw interview recordings and original transcripts will be retained on the researcher's encrypted storage for a maximum of 24 months after thesis submission, after which they will be permanently deleted. Anonymised survey responses, without any free-text fields that could permit identification, may be retained indefinitely in aggregated form for potential follow-up academic research. No primary data has been or will be shared with third parties without prior re-consent from participants.

6. Responsible person: Jihed Ben Mabrouk, MSc IMMIT candidate (Turku School of Economics, University of Turku). Contact: jihed.j.benmabrouk@utu.fi

Appendix 7: Glossary of Key Project and Salesforce Terms



This glossary defines the principal business, project, and Salesforce terms used in this thesis for readers without a technical or CRM background. The definitions are deliberately non-technical and aim only to make the empirical material readable; they are not exhaustive product descriptions.

List of abbreviations

ADKAR	Awareness, Desire, Knowledge, Ability, Reinforcement (a change-management model)
CM	Change Management
CRM	Customer Relationship Management
CSR	Corporate Social Responsibility
DSI	Direction des Systemes d'Information (French term for the corporate information-systems / IT department)
ERP	Enterprise Resource Planning
GDPR	General Data Protection Regulation
GenAI	Generative Artificial Intelligence
LEI	An anonymised internal business division of the energy client
MCA	Marketing Cloud Advanced
PEOU	Perceived Ease of Use
TAM	Technology Acceptance Model
TOE	Technology-Organization-Environment (an adoption framework)
UTAUT	Unified Theory of Acceptance and Use of Technology
XAI	Explainable Artificial Intelligence

Project and business terms

CRM (Customer Relationship Management). Software that centralises a company's information about its customers and prospects, including contact details, interactions, sales, and service history, so that teams manage relationships from a single system. Salesforce is the CRM platform at the centre of the projects studied here.

Sales Cloud programme (Case Project 1). A programme at the energy client to transform the sales process on Salesforce's Sales Cloud, delivered within one of its business divisions. It is one of the two main case projects in this thesis.

Marketing Cloud application (Case Project 2). A marketing application built for the energy client on Salesforce's Marketing Cloud Advanced and Data Cloud, supporting decarbonisation-related customer engagement and later extended to its Spanish operations. It is the second main case project.

Marketing Cloud Advanced (MCA). Salesforce's marketing product used to design, send, and track customer campaigns across channels such as email. It underpins the energy client's Marketing Cloud application studied here.

Prompt engineering. The practice of writing and refining the instructions, or prompts, given to a generative-AI system so that it returns useful and reliable output. In the projects studied, it was central to configuring AI features that summarise or draft text inside Salesforce (see Figure 12).

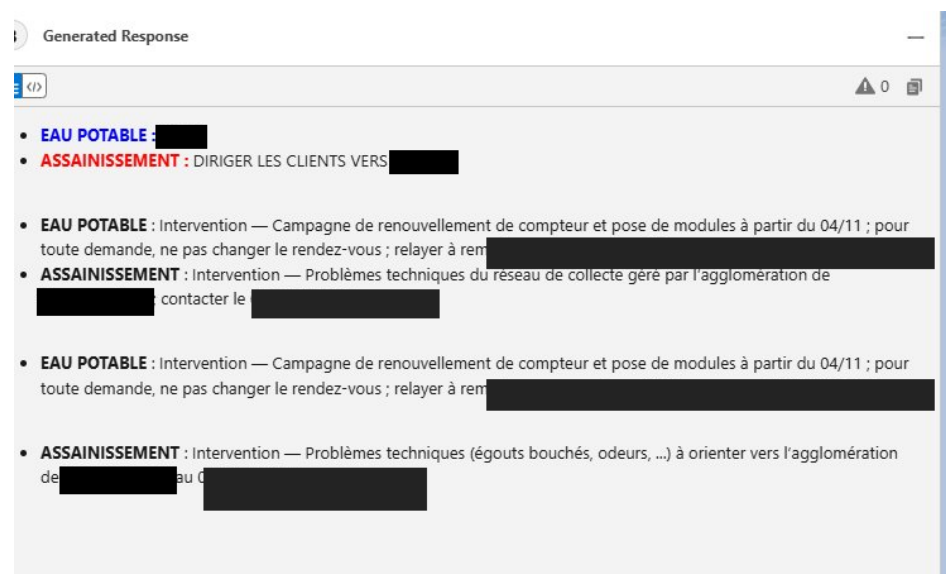


Figure 12. Example output generated from a configured prompt in a Salesforce AI feature (test data; personal identifiers redacted).

ADKAR. A change-management model (Hiatt, 2006) describing five conditions an individual must meet to adopt a change: Awareness, Desire, Knowledge, Ability, and Reinforcement. It is discussed in the theoretical framework and used by one interviewee as a practical checklist.

Salesforce platform terms

Object (standard vs custom). A type of record Salesforce stores, comparable to a table in a database. Standard objects (such as Lead) are built in; custom objects are created for a client's specific needs, for example a record tracking content history.

Field. A single piece of information stored on an object, for example an email address or an expected amount.

Lead. A potential customer who has shown interest but is not yet qualified. Leads are the entry point of the sales process.

Flow. Salesforce's main automation tool. A flow runs a sequence of steps automatically, for example updating records or creating data when a record is created or changed, without writing code. The projects used flows for tasks such as anonymising personal data for GDPR compliance (see Figure 13).




Figure 13. The Salesforce Flow Builder canvas, showing the visual, no-code construction of an automated process (illustrative; no client data).

Sales Cloud. The Salesforce product for managing the sales process, organised around leads and related sales records. The energy client's Sales Cloud programme studied here is a Sales Cloud implementation.

Marketing Cloud. Salesforce's family of marketing tools for designing and sending customer campaigns across channels and measuring their results.

Data Cloud. Salesforce's platform for unifying customer data from many sources into a single, consistent profile, so that marketing and AI features work from one trusted view of each customer.

Einstein and Agentforce. Salesforce's artificial-intelligence layer. Einstein provides AI features such as predictions and generated text inside Salesforce; Agentforce extends this to autonomous agents that can carry out tasks. These are the generative-AI capabilities whose adoption this thesis studies.