MESG

MESTRADO EM ENGENHARIA
DE SERVIÇOS E GESTÃO

Evaluating BPMN-Based Municipal Workflow Digitalisation:

The CityFlow Case

Gabriel Martins Rosa

Master Thesis

Supervisor at FEUP: Prof. José António Rodrigues Pereira de Faria

Supervisor at Associação Porto Digital: Mafalda Machado

U.PORTO
FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

2026-07-31

Abstract

Municipal governments manage diverse operational services, including animal welfare, green-space maintenance, market administration, and occurrence reporting, that have traditionally relied on paper forms, spreadsheets, and informal coordination. Digitalising these workflows through Business Process Management (BPM) and a unified platform based on Business Process Model and Notation (BPMN) promises traceability, accountability, and data-informed management, but requires standardised process models to accommodate variable municipal field work.

This dissertation evaluates CityFlow through an embedded mixed-methods case study conducted within Associação Porto Digital. Six analytical dimensions structure the evaluation: operational fit and functional coverage; adoption and usability; perceived value; data and monitoring; governance; and technological evolution. Evidence from 14 interview sessions with 26 participants across three operational service areas is triangulated with technical platform data and project documents.

RQ1: How effectively does CityFlow support municipal workflow digitalisation across different operational contexts and user profiles? The evidence shows differentiated effectiveness. Fit was strongest in the Centro de Recolha Oficial de Animais (CROA), where regulated work and repeated process refinement were closely aligned, and partial in Espaços Verdes and Mercados, where field variability or functional gaps required workarounds. Frequent use did not by itself establish good operational fit. Effectiveness therefore depends more on the fit between process models and operational context than on technical capability alone.

RQ2: Which improvements to existing capabilities, missing developments, and governance arrangements should be prioritised? Duplicate work was the broadest recurring friction, followed by missing functionality and monitoring and governance gaps. Existing workflows and authoritative-record arrangements should be improved before platform scope expands.

RQ3: Which technological evolution path is most promising for CityFlow, including monitoring, interoperability, and a bounded role for AI-enabled assistance? Accessible operational monitoring and interoperability form the strongest near-term path; process mining remains conditional on better event logs, while a human-approved voice-to-record artificial intelligence (AI) scenario remains exploratory.

The reusable contribution is a six-dimensional evaluation protocol comprising evidence registers, coding and triangulation rules, diagnostic maturity logic, and an application procedure. The CityFlow capability profile and roadmap demonstrate the protocol in one case; the AI scenario remains outside its reusable core.

Keywords: BPMN, municipal workflow digitalisation, mixed-methods case study, maturity assessment, public-sector BPM, CityFlow

Resumo

Os governos municipais gerem serviços operacionais diversos, incluindo bem-estar animal, manutenção de espaços verdes, administração de mercados e reporte de ocorrências, que tradicionalmente dependem de formulários em papel, folhas de cálculo e coordenação informal. A digitalização destes fluxos de trabalho através da gestão de processos de negócio (BPM, *Business Process Management*) e de uma plataforma baseada na notação *Business Process Model and Notation* (BPMN) promove rastreabilidade, responsabilização e gestão informada por dados, mas exige que modelos de processo normalizados acomodem a variabilidade do trabalho de campo municipal.

Esta dissertação avalia o CityFlow através de um estudo de caso integrado de métodos mistos realizado na Associação Porto Digital. A avaliação é estruturada por seis dimensões: adequação operacional e cobertura funcional; adoção e usabilidade; valor percecionado; dados e monitorização; governação; e evolução tecnológica. A evidência de 14 sessões de entrevista com 26 participantes em três áreas de serviço operacional é triangulada com dados técnicos da plataforma e documentos do projeto.

RQ1: Com que eficácia o CityFlow apoia a digitalização dos fluxos de trabalho municipais em diferentes contextos operacionais e perfis de utilizador? A evidência mostra uma eficácia diferenciada. A adequação foi mais forte no Centro de Recolha Oficial de Animais (CROA), onde o trabalho regulado e o refinamento repetido dos processos estavam mais alinhados, e parcial nos Espaços Verdes e Mercados, onde a variabilidade do trabalho ou lacunas funcionais exigiam soluções alternativas. A utilização frequente não demonstrou, por si só, uma boa adequação operacional. A eficácia depende, assim, mais da adequação entre os modelos de processo e o contexto operacional do que apenas da capacidade técnica.

RQ2: Que melhorias às capacidades existentes, desenvolvimentos em falta e mecanismos de governação devem ser priorizados? O trabalho duplicado foi a fricção recorrente mais abrangente, seguido por funcionalidades em falta e lacunas de monitorização e governação. Recomenda-se melhorar primeiro os fluxos existentes e definir os registos oficiais antes de expandir a plataforma.

RQ3: Qual é o caminho de evolução tecnológica mais promissor para o CityFlow, incluindo monitorização, interoperabilidade e um papel limitado para assistência baseada em IA? A monitorização operacional acessível e a interoperabilidade constituem o caminho mais sólido a curto prazo; a mineração de processos depende de melhores registos de eventos e um cenário de inteligência artificial (IA), com aprovação humana, para converter voz em registos permanece exploratório.

A contribuição reutilizável é um protocolo de avaliação em seis dimensões, composto por registos de evidência, regras de codificação e triangulação, lógica diagnóstica de maturidade e um procedimento de aplicação. O perfil de capacidades e o roteiro do CityFlow demonstram o protocolo num caso; o cenário de IA permanece fora do seu núcleo reutilizável.

Palavras-chave: BPMN, digitalização de fluxos de trabalho municipais, estudo de caso de métodos mistos, avaliação de maturidade, BPM no setor público, CityFlow

Acknowledgments

My thanks go to Associação Porto Digital for the internship opportunity and for providing access to the CityFlow platform, project documentation, and operational stakeholders. I am especially appreciative of Mafalda Machado and Elisabete Monteiro for their openness, practical guidance, operational context, and day-to-day support during the work. I also thank the product management and service design teams for sharing institutional knowledge.

I am also grateful to DMSI, and specifically Octávio Filgueiras and Eduardo Fernandes, for helping coordinate contacts with the municipal organic units and supporting access to the operational context needed for the dissertation.

My sincere thanks also go to Professor José António Faria for his supervision and constructive feedback throughout the development of this dissertation.

I thank the operational teams of CROA, Espaços Verdes, Mercados, and the internal stakeholders who participated in the interviews for their time, candour, and operational insight. Without their willingness to share their experience, the work carried out in this dissertation would not have been possible.

I also wish to acknowledge INESC TEC, specifically the Innovation, Technology and Entrepreneurship Centre (CITE) team, for providing me with the opportunity to share my experience throughout the dissertation and for the very valuable feedback that had a significant impact on the work carried out.

Finally, I am deeply thankful to my family for their unconditional support throughout my life, and to my friends for their support and the laughter when it was needed most. Special thanks go to Filipe Marques and Carolina Almeida.

Contents

1 Introduction **1**

1.1 Problem Context 1

1.1.1 Research Context: Embedded Internship at APD 2

1.2 Research Objectives and Questions 3

1.3 Scope and Exclusions 3

1.4 Contributions 4

1.5 Dissertation Structure 4

2 Literature Review and Analytical Framework **5**

2.1 BPM and Workflow Digitalisation in Public Services 5

2.2 Municipal Digital Transformation 6

2.3 BPM Governance, Maturity, and Capability Assessment 6

2.4 Evaluation, Indicators, and Public Value 7

2.5 Adoption, Usability, and Service Fit 7

2.6 Design Science Research Positioning 8

2.7 Synthesis: Analytical Dimensions 8

3 Methodology **10**

3.1 Methodological Positioning 10

3.2 Research Design 10

3.2.1 Artifact Construction and Design Evolution 10

3.3 Case Selection and Units of Analysis 11

3.4 Evidence Sources 13

3.5 Data Preparation, Coding, and Descriptive Analysis 13

3.6 Technical and Documentary Analysis 14

3.7 Triangulation, Maturity Synthesis, and AI Scenario 14

3.7.1 Formative Review and Evaluation Boundary 15

3.8 Research Quality, Ethics, and Limitations 16

4 Case Study Context: CityFlow **17**

4.1 Organizational Setting: Associação Porto Digital 17

4.2 CityFlow: A BPMN-Based Municipal Workflow Platform 17
4.3 Operational Service Areas (“Verticals”) 20
4.3.1 CROA: Centro de Recolha Oficial de Animais 20
4.3.2 Espaços Verdes: Green Spaces and Environment 21
4.3.3 Mercados: Markets and Fairs 21
4.3.4 Ocorrências/Reporta, Occurrence and Incident Reporting 22
4.3.5 Other Verticals and Modules 23
4.3.6 Vertical Summary 23
4.4 Technical Architecture 24
4.5 Stakeholder Landscape 25
4.6 Project Timeline and Current Status 26

5 Results and Discussion **28**
5.1 Current Value and Operational Fit (RQ1) 28
5.1.1 Usage Frequency 28
5.1.2 Operational Fit 28
5.1.3 Value Recognition 30
5.1.4 Stratification by Stakeholder Category 31
5.1.5 Usage Motivation 31
5.2 Improvement Priorities (RQ2) 31
5.2.1 Duplicate Work 31
5.2.2 Missing Functionality 32
5.2.3 Improvement and Development Priorities 33
5.2.4 Barriers 33
5.3 Future Evolution (RQ3) 34
5.3.1 Operational Monitoring: The Primary Answer to RQ3 34
5.3.2 Process Mining Feasibility 35
5.3.3 Bounded AI Scenario 35
5.4 Discussion: Triangulation and Maturity Profile 36
5.4.1 Triangulation Synthesis 36
5.4.2 Case-Specific Maturity Profile 38
5.5 Recommendations Roadmap 40
5.5.1 Evidence-to-Requirement Chain 40

5.5.2 Framework Application Protocol 40
5.5.3 CityFlow Improvement Register 41

6 Conclusions and Future Work **44**
6.1 Research Questions Revisited 44
6.2 Summary of Contributions 45
6.3 Methodological Reflection 46
6.4 Limitations 46
6.5 Future Work 47
6.6 Final Remarks 47

References **48**

7 Sustainable Development Goals Analysis **51**

Appendices **52**

A Qualitative Codebook and Indicator Definitions **53**
A.1 Code Hierarchy 53
A.1.1 Dimension 1: Operational Fit and Functional Coverage 53
A.1.2 Dimension 2: Adoption, Usability, and Digital Literacy 53
A.1.3 Dimension 3: Value and Perceived Impact 53
A.1.4 Dimension 4: Data, Indicators, and Monitoring 53
A.1.5 Dimension 5: Governance, Ownership, and Sustainability 54
A.1.6 Dimension 6: Technology, Integrations, and Future Evolution 54
A.1.7 Supplementary: Recommendation Type 54
A.2 Indicator Definitions 54
A.3 Coding Procedure 55

B Interview Session Register **56**

C Participant Register **57**

D Session-Level Indicator Grid **58**

E Triangulation Matrix **59**
E.1 Service-Area Entries 59

E.1.1 T01: CROA Operational Fit and Value . . . 59
E.1.2 T02: CROA Data, Indicators, and Monitoring . . . 59
E.1.3 T03: Espaços Verdes Operational Fit and Adoption . . . 59
E.1.4 T04: Espaços Verdes Data Quality and Monitoring . . . 60
E.1.5 T05: Mercados Operational Capability . . . 60
E.2 Cross-Case and Platform Entries . . . 60
E.2.1 T06: Governance and Sustainability . . . 60
E.2.2 T07: Technology and Future Evolution . . . 60
E.2.3 T08: Technical Implementation Versus Adoption . . . 60
E.2.4 T09: Ocorrências/Reporta Scope Limitation . . . 61
E.2.5 T10: Platform Portfolio Governance . . . 61
E.2.6 T11: Platform Evaluation and Prioritisation . . . 61

F PSPP Descriptive Statistics 62
F.1 Session Composition (N=14) . . . 62
F.2 Usage Frequency (N=10 operational) . . . 62
F.3 Value Recognition (N=14) . . . 62
F.4 Operational Fit (N=14) . . . 62

G Claim-Evidence Traceability 63

H Quantitative Analysis Report 65
H.1 Elasticsearch Volume and Coverage . . . 65
H.1.1 Elasticsearch Volume and Capture Coverage . . . 65
H.1.2 BPMN/Camunda Inventory . . . 65
H.1.3 Espaços Verdes Process-Definition Inventory . . . 65
H.1.4 Additional Technical Annex Service Areas . . . 66
H.1.5 Deployments by Year . . . 67
H.2 BPMN/Camunda Technical Footprint and Limits . . . 67
H.3 CROA Module Coverage . . . 67

I Qualitative Coding Frequencies 68
I.1 Code Frequencies by Category . . . 68
I.2 Category Totals . . . 68

J Generative AI Use Note **69**

List of Tables

2.1 Analytical dimensions mapped to literature and indicators 9

3.1 Documented design evolution and disposition of alternatives. 11
3.2 Stakeholder coverage, exclusions, and claim consequences. 12
3.3 Evaluation framework architecture and artifact boundary. 15
3.4 Formative review history and limits. 15

4.1 Active Camunda process definitions checked in the 15 June 2026 targeted technical 17
4.2 Simplified CityFlow technical architecture. 19
4.3 Summary of production-relevant verticals examined 23
4.4 CityFlow stakeholder governance structure. 26
4.6 Timeline milestones and their analytical relevance. 27

5.1 CityFlow usage frequency by context. 28
5.2 Operational fit by context. 29
5.3 Value recognition by vertical (N = 14 sessions) 30
5.4 Duplicate work in operational sessions 32
5.5 Missing functionality by vertical. 32
5.6 Barrier evidence across 14 sessions (26 participants) 33
5.7 Case-specific capability profile by service area and dimension. 39
5.8 Evidence-to-requirement-to-action chain. 40
5.9 Systematic improvement register with ownership, constraints, and closure criteria. 42
5.10 Reusable artifact component responsibilities and completion conditions. 43

6.1 Platform-level and service-area-specific conclusions and boundaries. 45

D.1 Full session-level indicator grid. 58

E.1 Triangulation entry index. 59

H.3 Espaços Verdes process-definition metrics from the 29 April 2026 inventory. . . 66

List of Figures

3.1 Evidence-to-artifact pathway. 11

4.1 CROA occurrence BPMN process, based on geralOcorrencia version 12 captured on 15 June 18
4.2 CityFlow information system architecture and primary data flows. 19
4.3 CROA module home screen. 21
4.4 Mercados module home screen. 22
4.5 Reporta occurrence list with status and action controls. 23
4.6 BPMN inventory characteristics by service area. 25
4.7 CityFlow project and study timeline. 26

5.1 Bounded voice-to-structured-record AI scenario. 36
5.2 Normalised coded-theme coverage by context and analytical dimension. 37
5.3 Operational fit and process-definition iteration 38

Abbreviations

APD Associação Porto Digital

BPM Business Process Management

BPMN Business Process Model and Notation

BPMM Business Process Maturity Model

CMP Câmara Municipal do Porto

CROA Centro de Recolha Oficial de Animais

DMGA Departamento Municipal de Gestão de Ativos

DMSI Direção Municipal de Sistemas de Informação

DSR Design Science Research

ES Elasticsearch

FAGOP Ferramenta de Apoio à Gestão Operacional (CityFlow predecessor project)

FEUP Faculdade de Engenharia da Universidade do Porto

GenAI Generative Artificial Intelligence

KPI Key Performance Indicator

LDAP Lightweight Directory Access Protocol

MESG Mestrado em Engenharia de Serviços e Gestão

MinIO Object storage system

NiFi Apache NiFi data pipeline management

NodeJS JavaScript runtime environment

OAuth Open standard for access delegation

PDS Ponto de Situação (project status meeting)

PostgreSQL Relational database management system

Power BI Microsoft Power BI (business intelligence platform)

PSPP GNU PSPP (open-source SPSS-compatible statistical software)

ReactJS JavaScript library for building user interfaces

ROPI Rotas, Itinerários e Percursos (event integration for itineraries)

RQ Research Question

SDG Sustainable Development Goal

SPSS Statistical Package for the Social Sciences

1 Introduction

1.1 Problem Context

Municipal governments manage a diverse portfolio of operational services: animal welfare and veterinary care, green space and park maintenance, market and fair administration, occurrence and incident reporting, mobility asset management, and environmental compliance. Each of these services generates operational records, task assignments, inspection reports, intervention logs, planning schedules, status updates, that have traditionally been managed through paper forms, spreadsheets, and informal communication channels.

Digitalising these workflows presents both an opportunity and a challenge (Dumas et al. 2018; Pittaway and Montazemi 2020). The opportunity is that structured digital workflows can improve traceability, accountability, data quality, and management visibility. When every animal intervention, park maintenance task, and market inspection is recorded in a single platform with consistent data structures, operational managers can understand what work is being done, where, and by whom. Municipal leadership can move from anecdotal oversight to evidence-informed management. The challenge is that municipal operational work is inherently variable. Field workers operate in diverse physical environments with different connectivity, mobility, and task characteristics. An animal handler cleaning cages and a gardener maintaining a public park face fundamentally different work contexts, yet both need the platform to support their tasks without creating additional friction. Standardizing processes across verticals as different as animal shelter management and park maintenance requires balancing the consistency of a workflow platform against the flexibility that operational work demands.

The Associação Porto Digital (APD) is a private non-profit association that coordinates the Porto Municipality's urban digitalisation strategy. In partnership with the municipality, APD has developed CityFlow: a web platform for municipal workflow digitalisation. CityFlow models operational processes in the Business Process Model and Notation (BPMN) standard (Dumas et al. 2018), executes them through Camunda, the workflow-management software, and stores operational records in Elasticsearch, the searchable document database. The platform is deployed across multiple municipal operational service areas. APD project vocabulary calls these areas "verticals"; the term denotes a service domain, such as animal welfare or green spaces, rather than a formal level in the municipal hierarchy. By mid-2026, the platform supported operational workflows in the Centro de Recolha Oficial de Animais (CROA), green spaces and environment (Espaços Verdes), markets and fairs (Mercados), and occurrence reporting (Reporta/Ocorrências), with additional verticals, including mobility asset management, event integration for itineraries, and asset management, at earlier stages of requirement elicitation or prototype development.

CityFlow's development model embeds service design and product management alongside engineering, meaning that each vertical is approached through user research, process modelling, iterative prototyping, and production deployment. This model represents a significant investment by the municipality in moving from fragmented, paper-based, or siloed digital practices toward an integrated workflow platform. Understanding whether this investment is delivering value, where it is generating friction, and how it should evolve is the problem that motivates this

dissertation.

This dissertation evaluates CityFlow's support for municipal workflow digitalisation through an embedded mixed-methods case study (Yin 2018; Creswell and Plano Clark 2018) conducted within APD's product management team. The evaluation combines interview evidence from 14 distinct interview sessions with 26 registered participants across operational and stakeholder groups, technical evidence from platform extracts, and documentary evidence from project materials, requirements records, governance communications, and supporting analysis outputs. The reusable artifact developed through the evaluation is the CityFlow Evaluation and Maturity Assessment Framework: a six-dimensional evaluation protocol comprising evidence registers, coding and triangulation rules, diagnostic maturity logic, and an application procedure. Its application to CityFlow produces a case-specific capability profile and recommendations roadmap. These are demonstration outputs rather than reusable components of the framework, while the bounded artificial intelligence (AI) scenario is a separate exploratory output.

The framework examines six dimensions introduced here before the results: operational fit and functional coverage; adoption, usability, and digital literacy; value and perceived impact; data, indicators, and monitoring; governance, ownership, and sustainability; and technology, integrations, and future evolution. Chapter 2 grounds these dimensions in the literature, Chapter 3 operationalises them through the interview guide, codebook, and indicator grid, and Chapter 5 reports the evidence for each research question.

1.1.1 Research Context: Embedded Internship at APD

The research was conducted during a five-month internship (February to June 2026) within APD's product management team. This embedded position provided direct access to the CityFlow platform, project documentation, stakeholder communications, weekly product and project governance meetings, and the operational teams using CityFlow across multiple municipal verticals. The researcher participated in weekly PDS (Ponto de Situação) meetings with APD, the Direção Municipal de Sistemas de Informação (DMSI) of Câmara Municipal do Porto (CMP), and supplier representatives, progress meetings with the product management lead, and intermediate presentations to APD and municipal stakeholders.

The embedded position shaped the evidence base in three ways. First, it enabled the collection of interview evidence from 26 participants across operational users (CROA, Espaços Verdes, Mercados), internal stakeholders (APD product management, service design, and one front-end tester/quality-assurance stakeholder), a DMSI/CMP governance representative, and one supplier-side stakeholder whose hybrid project/development role included direct programming experience. Second, it provided access to platform evidence, Camunda process definitions, BPMN models, Elasticsearch aggregate data, and front-end interfaces, that would not have been available to an external researcher. Third, it enabled direct observation of governance dynamics, prioritization processes, and the practical constraints that shape municipal platform development. The internship was conducted under APD and Porto Municipality confidentiality conditions, and all platform evidence was captured read-only, without creating, editing, or exporting operational records.

1.2 Research Objectives and Questions

To address the problem outlined above, this dissertation pursues three research objectives:

1. To assess how effectively CityFlow supports municipal workflow digitalisation across different operational contexts and user profiles.
2. To identify and prioritise improvements to existing capabilities, missing developments, and governance arrangements.
3. To evaluate the most promising technological evolution path for CityFlow, including monitoring, interoperability, and a bounded role for AI-enabled assistance.

These objectives are addressed through three research questions:

RQ1: How effectively does CityFlow support municipal workflow digitalisation across different operational contexts and user profiles?

RQ2: Which improvements to existing capabilities, missing developments, and governance arrangements should be prioritised?

RQ3: Which technological evolution path is most promising for CityFlow, including monitoring, interoperability, and a bounded role for AI-enabled assistance?

1.3 Scope and Exclusions

Interview findings are reported at session level for the assessed case. Chapter 3 defines the sampling, counting, and interpretation rules used throughout the analysis.

The dissertation evaluates CityFlow as a municipal workflow platform deployed in a specific organizational context. It is not an evaluation of APD as an organization, of the Porto Municipality's overall digital strategy, or of Camunda, Elasticsearch, or BPMN as technologies independent of their use in CityFlow.

The evaluation covers four production-relevant verticals. CROA, Espaços Verdes, and Mercados are assessed through interview and technical evidence. Ocorrências and Reporta are included in the technical scan but lack dedicated operational interview coverage. Additional verticals at earlier stages of development are referenced for platform-evolution context but are not primary units of analysis.

Process mining uses event logs to discover, monitor, and improve actual process behaviour, including process discovery and conformance checking (W. Van Der Aalst 2016). It was assessed here for feasibility but is treated as exploratory because the available logs do not meet the quality required for conclusive discovery or conformance analysis. The bounded AI scenario draws on exploratory research into large-language-model assistance across the BPM lifecycle, for which task context, input quality, factual errors, and expectations remain material concerns (Vidgof, Bachhofner, and Mendling 2023). In this dissertation, *bounded* means that the assistant only converts dictated field notes into a proposed structured record: a person must review and approve it, and the assistant neither makes operational decisions nor writes an authoritative record autonomously. It is a design exploration, not a production implementation or near-term development recommendation.

1.4 Contributions

The dissertation makes two types of contribution.

The **practical contribution** is a structured, evidence-based evaluation of CityFlow that identifies where the platform creates value, where it generates friction, and which improvements should be prioritised. Its case-specific outputs are a diagnostic capability profile and a sequenced recommendation roadmap grounded in triangulated interview, technical, and documentary evidence.

The **academic contribution** is the reusable CityFlow Evaluation and Maturity Assessment Framework: an adaptable six-dimensional protocol that may be applied and tested in comparable contexts. It combines evidence registers, session-level qualitative coding, descriptive and technical indicators, documentary evidence, triangulation rules, diagnostic maturity logic, and a repeatable application procedure. It does not include the CityFlow-specific capability assignments, roadmap decisions, or exploratory AI scenario as universally reusable results.

1.5 Dissertation Structure

The dissertation is organised in six chapters.

Chapter 2 reviews the literature on BPM and workflow digitalisation in public services, municipal digital transformation, BPM governance and maturity, evaluation and public value, adoption and usability, and the DSR positioning of the evaluation artifact. It synthesizes the literature into six analytical dimensions that structure the empirical evaluation.

Chapter 3 presents the methodology: an embedded mixed-methods case study design, preceded by a comparison of alternative methodological approaches. The chapter describes the case and units of analysis, the five evidence streams, data preparation and coding procedures, QualCoder qualitative coding, GNU PSPP and IBM SPSS Statistics-compatible descriptive analysis, technical and documentary analysis, triangulation logic, and research quality, ethics, and limitations.

Chapter 4 describes the case context: APD as an organization, CityFlow as a BPMN-based municipal workflow platform, the four production-relevant operational verticals with a vertical summary table, the technical architecture, the stakeholder landscape, and the project timeline.

Chapter 5 presents the results and discussion, organised around the three research questions: current value and operational fit (RQ1), improvement priorities (RQ2), and future evolution (RQ3). The chapter reports the empirical results before developing the triangulation synthesis, case-specific maturity profile, integrated recommendations roadmap, and limitations of the evaluation.

Chapter 6 concludes the dissertation with a summary of contributions, methodological reflection, limitations, directions for future work, and final remarks.

The appendices contain the qualitative codebook and indicator definitions, session and participant registers, indicator grid, triangulation matrix, PSPP descriptive statistics, claim-evidence traceability table, and technical data inventory.

2 Literature Review and Analytical Framework

This chapter uses a targeted, structured narrative review to frame CityFlow's empirical evaluation. Iterative backward and forward snowballing identified the concepts needed to assess process fit, adoption, value, monitoring, governance, and future technological evolution. Because reproducible database queries and screening counts were not retained, this is not presented as a systematic review.

Search lanes were defined around the dissertation's analytical needs: BPM/BPMN workflow digitalisation, municipal digital transformation, BPM governance and maturity, evaluation and public value, adoption and usability, design science positioning, case-study methods, and mixed-method triangulation. Candidate sources were selected when they were peer-reviewed articles, academic books, standards, or official institutional material that directly supported the evaluation framework or methodology. Generic smart-city, vendor, or sector-distant material was excluded unless it provided transferable concepts for workflow governance, data quality, or technical evolution. The literature base was re-reviewed against the final research questions and empirical findings.

2.1 BPM and Workflow Digitalisation in Public Services

Business Process Management (BPM) provides methods for discovering, analysing, redesigning, executing, and monitoring organizational processes (Dumas et al. 2018). BPMN extends this logic by making process models both human-readable and executable. When BPMN models are deployed to workflow engines such as Camunda, tasks, gateways, events, and states become operational traces rather than only diagrams (Dumas et al. 2018). This makes BPMN valuable for traceability and monitoring, but also exposes misalignment between the modelled process and the work actually performed.

The quality of process models matters. Mendling et al.'s seven process modelling guidelines (Mendling, Reijers, and Van Der Aalst 2010) emphasise clarity, minimalism, and structural correctness. In CityFlow, model quality is therefore part of technical capability: process-definition versions, task granularity, and gateway structures influence whether operational work can be executed and understood through the platform.

Public-sector BPM differs from private-sector process improvement. Local government operates under political oversight, legal mandates, budget cycles, procurement constraints, transparency obligations, and diverse citizen-facing service contexts (Niehaves, Plattfaut, and Becker 2013). Standardisation can improve reliability and accountability, but it can also conflict with professional discretion and variable field conditions. This tension is central to the CityFlow case, where verticals such as CROA, Espaços Verdes, Mercados, and occurrence reporting share a workflow platform but differ in operational predictability.

Process mining offers a possible extension to BPMN-based workflow management by using event logs to discover, monitor, or compare process behaviour (W. Van Der Aalst 2016). Its value depends on clean logs with reliable case identifiers, activity labels, timestamps, and resource information (IEEE 2016; Van Der Aalst et al. 2012). Government applications frequently face

fragmented or noisy data (Rawiro et al. 2022). Van Voorst tot Voorst (Voorst tot Voorst 2025) shows that municipal process-mining projects can require substantial preprocessing before conformance or discovery claims are defensible. CityFlow event-log evidence is therefore assessed against these same feasibility criteria.

2.2 Municipal Digital Transformation

Municipal digital transformation involves more than automating existing tasks. It changes how services are coordinated, recorded, monitored, and governed. Pittaway and Montazemi (Pittaway and Montazemi 2020) highlight constraints that shape public-sector digitalisation: legacy systems, limited specialist capacity, procurement rules, public accountability, and users with varied digital literacy. These constraints make incremental value, institutional fit, and operational continuity more important than rapid disruption.

BPM culture is also a condition for sustained process improvement. Kregel, Distel, and Coners (Kregel, Distel, and Coners 2022) show that shared process-oriented values and behaviours help local governments sustain BPM initiatives beyond initial technical deployment. Szedmák (Szedmák, Varga, and Szabó 2025) similarly emphasises the tension between centralized platform governance and operational autonomy. CityFlow's vertical-by-vertical deployment, service-design involvement, and user research reflect this public-sector logic: adoption depends on demonstrating value in each operational setting.

The smart-city and digital-government literature frames workflow platforms as part of wider data and governance ecosystems. Gil-Garcia, Pardo, and Nam (2015) identify information, technology, organizational, and institutional dimensions of smart-city initiatives. Cross-sectoral process modelling work likewise positions BPM as a way to connect city systems and support integration, flexibility, and real-time information access (Dessai and Javidroozi 2021). CityFlow sits at this intersection: it combines workflow technology, operational data, cross-departmental coordination, and municipal accountability.

2.3 BPM Governance, Maturity, and Capability Assessment

BPM governance concerns the roles, routines, and decision mechanisms that guide process management (Santana et al. 2011). In a multi-vertical platform such as CityFlow, governance determines how priorities are selected, how user feedback reaches product teams, how releases are communicated, and how new verticals are sequenced.

Maturity models provide a structured way to assess capabilities and identify improvement paths. Established models such as BPMM (Object Management Group (OMG) 2008), Hammer's process audit (Hammer 2007), and Rosemann and de Bruin's BPM maturity model (Rosemann and de Bruin 2005) emphasise dimensions such as process design, metrics, ownership, governance, methods, people, and information technology. Becker, Knackstedt, and Pöppelbuß (Becker, Knackstedt, and Pöppelbuß 2009) add that maturity assessment must be transparent and evidence-based, while de Bruin and Rosemann (de Bruin and Rosemann 2007) show how capability areas can be grounded in expert judgement. Van Looy, De Backer, and Poels (Looy, Backer, and Poels 2014) further support a multi-dimensional capability profile rather than reducing organisational maturity to one technical indicator or aggregate score.

The caution in this literature is important. Fisher (Fisher 2004) and Smajli, Feldman, and Cox

(Smajli, Feldman, and Cox 2024) warn that maturity models can become box-ticking exercises when scores are detached from operational evidence or when technical deployment is confused with maturity. This dissertation therefore uses maturity diagnostically: each vertical is assessed across dimensions, unknown values are retained where evidence is insufficient, and no single maturity score is presented as a validated universal measure.

2.4 Evaluation, Indicators, and Public Value

Evaluating CityFlow requires both measurement and judgement. Performance measurement supplies descriptive indicators; evaluation interprets what those indicators mean for operational work and public value. Van Der Kolk (Van Der Kolk 2022) shows that public-sector performance measurement must handle multiple stakeholders, goal ambiguity, attribution challenges, and the risk that measured indicators displace unmeasured values. For CityFlow, no single metric can capture effectiveness, so the evaluation combines operational fit, value recognition, usability, data quality, monitoring, governance, and technical evidence.

Indicators also shape attention. Lnenicka et al. (Lnenicka et al. 2024) show that public data indicators make infrastructure visible and governable. This supports the use of a structured indicator grid rather than a purely narrative account of interviews. The grid enables comparison across verticals while keeping evidence rules explicit in Appendix A.

Public value is not limited to efficiency. Moore (Moore 1995) and Benington and Moore (Benington and Moore 2011) frame public value as collectively valued outcomes authorised through public governance. For a municipal workflow platform, value includes traceability, accountability, institutional learning, and management visibility. Tejedo-Romero (Tejedo-Romero et al. 2022) and Harrison, Pardo, and Cook (Harrison, Pardo, and Cook 2012) reinforce that digital transparency depends not only on data existence but also on accessible monitoring and governance structures able to act on the evidence.

Service design literature adds the user and service-system perspective. Sangiorgi (Sangiorgi and Prendiville 2014), Teixeira et al. (Teixeira et al. 2012), and Patrício et al. (Patrício et al. 2011) show that public services are co-produced across encounters, operational experiences, and broader service systems. This is directly relevant to CityFlow because the same platform must serve field workers, operational managers, APD product teams, and municipal leadership.

2.5 Adoption, Usability, and Service Fit

Technical capability does not guarantee adoption or appropriation. In mandated workplace systems, users may be required to use a platform, but acceptance and embedded use depend on whether it fits the work. Simonsen and Kensing's MUST method (Simonsen and Kensing 1997) argues that fit-to-work requires direct engagement with operational practice. This supports the dissertation's focus on how CityFlow fits field conditions, not only whether it is available.

Service-design research also links user research to service concept and evaluation. Foglieni et al. (Foglieni, Maffei, and Villari 2014) connect contextual inquiry, journey mapping, service blueprinting, and prototyping to service design decisions. CityFlow's development process included APD service design involvement, which helps explain why users often recognise value even where operational fit remains partial.

The methodological literature supports the embedded mixed-methods case-study design used here. Yin (Yin 2018) provides the case-study logic; Creswell and Plano Clark (Creswell and Plano Clark 2018) support integration of qualitative and descriptive quantitative evidence; Jick (Jick 1979) frames triangulation; Braun and Clarke (Braun and Clarke 2006) ground the thematic coding approach.

2.6 Design Science Research Positioning

The dissertation also uses Design Science Research (DSR) language to position the CityFlow Evaluation and Maturity Assessment Framework as a design-oriented artifact. DSR treats artifacts as carriers of design knowledge and connects problem environments to the knowledge base through relevance and rigour cycles (Hevner et al. 2004). Peffers et al. (Peffers et al. 2007) describe a sequence of problem identification, objectives, design, demonstration, evaluation, and communication. The dissertation follows a similar structure, but the primary methodological label remains an embedded case study: CityFlow was designed, deployed, and maintained by APD before and throughout the study period.

Gregor and Hevner (Gregor and Hevner 2013) classify design knowledge contributions by solution and domain maturity. The CityFlow framework is best understood as an improvement contribution: it adapts BPM evaluation and maturity concepts to municipal workflow digitalisation. Jones and Gregor (Jones and Gregor 2007) help describe its components: purpose, constructs, principles, justificatory knowledge, and mutability. The reusable artifact contributed here is the six-dimensional evaluation protocol: evidence registers, coding and triangulation rules, diagnostic maturity logic, and the application procedure. The capability profile and recommendations roadmap are outputs of demonstrating that protocol on CityFlow. The bounded AI scenario is a separate exploratory case output, and the pre-existing CityFlow platform is not an artifact created by this study.

The application of the framework to CityFlow is a naturalistic demonstration: it is exercised in a real organisational setting with genuine evidence constraints. Dated academic, organisational, and supervisor reviews provide formative feedback on its framing and reporting. They do not constitute a summative stakeholder acceptance test of the final artifact. Formal evaluation with intended users and across multiple municipal platforms remains future work. This boundary limits generalisability but does not undermine the case-based contribution.

2.7 Synthesis: Analytical Dimensions

The literature converges on six analytical dimensions that structure the empirical evaluation and are operationalized through the indicator grid, codebook, and evidence matrix.

1. **Operational fit and functional coverage.** Alignment between BPMN processes, interface affordances, and operational work (Dumas et al. 2018; Mendling, Reijers, and Van Der Aalst 2010).
2. **Adoption, usability, and digital literacy.** Factors that enable or constrain daily use (Simonsen and Kensing 1997; Pittaway and Montazemi 2020).
3. **Value and perceived impact.** Operational, organizational, and public value attributed to CityFlow (Van Der Kolk 2022; Tejedo-Romero et al. 2022).

4. **Data, indicators, and monitoring.** Availability and accessibility of operational data for management (Lnenicka et al. 2024; Harrison, Pardo, and Cook 2012).
5. **Governance, ownership, and sustainability.** Structures that guide prioritization, release communication, and feedback (Santana et al. 2011; Rosemann and de Bruin 2005).
6. **Technology, integrations, and future evolution.** Architecture, interoperability, process-mining feasibility, and bounded AI assistance (W. Van Der Aalst 2016; Vidgof, Bachhofner, and Mendling 2023).

Table 2.1 maps each dimension to the primary literature sources, empirical indicators, research questions, and artifact components.

Table 2.1: Analytical dimensions mapped to literature, indicators, research questions, and artifact components

Dimension	Primary literature	Empirical indicators	Primary RQ	Artifact component
Operational fit and functional coverage	Dumas et al. 2018; Mendling et al. 2010	Usage frequency, operational fit, missing functionality	RQ1	Indicator grid, codebook
Adoption, usability, and digital literacy	Simonsen & Kensing 1997; Pittaway & Montazemi 2020	Usability barriers, digital literacy barriers, training needs, resistance	RQ1	Codebook
Value and perceived impact	Van Der Kolk 2022; Tejedo-Romero 2022; Harrison et al. 2012	Value recognition	RQ1, RQ2	Indicator grid, triangulation matrix
Data, indicators, and monitoring	Lnenicka et al. 2024; Harrison et al. 2012	Monitoring gaps, Elasticsearch data coverage, status-field availability	RQ2, RQ3	Technical data inventory, PSPP outputs
Governance, ownership, and sustainability	Santana 2011; Rosemann & de Bruin 2005; de Bruin & Rosemann 2007	Governance gaps, improvement priorities	RQ2	Recommendation roadmap
Technology, integrations, and future evolution	Van Der Aalst 2016; Vidgof et al. 2023	Process-mining feasibility, AI scenario, interoperability	RQ3	AI scenario, process-mining assessment

The framework is not a hypothesis-testing model. It is an organizing structure that connects the literature to the empirical evidence, gives each finding a theoretical home, and makes the scope and limits of the evaluation transparent.

Existing BPM, public-sector digital transformation, maturity-assessment, process-mining, and smart-city process-modelling literature provides concepts for analysing workflow platforms, process governance, monitoring, and operational improvement. However, the literature offers more limited directly applicable guidance for evaluating an already deployed BPMN-based municipal workflow platform across heterogeneous operational verticals by combining interview evidence, technical platform evidence, and documentary or governance evidence. This dissertation addresses that gap by developing and applying a case-specific evaluation and maturity-assessment framework to CityFlow. The framework is not proposed as a universally validated model, but as a documented and adaptable structure for analysing comparable municipal workflow-digitalisation contexts.

3 Methodology

3.1 Methodological Positioning

This dissertation uses an embedded mixed-methods case-study design to evaluate CityFlow in its municipal setting (Yin 2018; Creswell and Plano Clark 2018). The choice follows from the research problem: CityFlow is not only a technical system but a workflow platform used inside operational routines, professional cultures, and governance constraints.

Alternative designs were less suitable. A survey would not explain the work practices shaping use, while ethnography would not provide the required structured technical and cross-vertical comparison. Design Science Research informs the evaluation framework (Peffers et al. 2007; Gregor and Hevner 2013), but CityFlow itself preceded the study. Multiple-case, action-research, and longitudinal designs were infeasible within the internship. The selected design treats CityFlow as the main case and its verticals, stakeholders, and evidence streams as embedded units, allowing interview, technical, and documentary triangulation (Jick 1979).

3.2 Research Design

The three research questions address effectiveness (RQ1), improvement priorities (RQ2), and technological evolution (RQ3). A structured indicator grid, defensible descriptive comparisons, and aggregate CityFlow/Camunda evidence connect them to the analysis. Interviews explain operational experience; technical data describes platform structure and activity, without purporting to validate each statement against individual workflow instances.

The dissertation's design-oriented contribution is the CityFlow Evaluation and Maturity Assessment Framework. Its reusable core is a six-dimensional evaluation protocol comprising evidence registers, coding and triangulation rules, diagnostic maturity logic, and a repeatable application procedure. The case-specific capability profile and recommendations roadmap are outputs of demonstrating that protocol on CityFlow. The bounded AI scenario is a separate exploratory output outside the reusable core. The diagnostic logic is not a validated universal maturity model.

3.2.1 Artifact Construction and Design Evolution

The artifact emerged through documented scope and method decisions rather than from a single design event. The initial January 2026 proposal emphasised key performance indicator (KPI) monitoring. A 3 March academic review exposed the risk of equating available indicators with adoption, value, or maturity. APD and DMSI scope reviews on 4-5 March redirected the empirical focus from prospective setup cases to active production service areas. Evidence collection then showed that neither interviews nor technical counts could explain CityFlow alone, leading to the mixed-method protocol, explicit triangulation rules, and a case-specific capability profile. The full decision trail is retained in the artifact-design decision register.

Table 3.1: Documented design evolution and disposition of alternatives.

Stage	Alternative considered	Disposition	Evidence-based rationale	Effect on the artifact
Initial framing	KPI-only monitoring	Superseded	Monitoring cannot establish fit, adoption, value, or governance	Six analytical dimensions replace a metric-only view
Case selection	Parking or prospective ENTI setup cases	Replaced	Active service areas offered observable workflows and contrasting conditions	CROA, Espaços Verdes, and Mercados became embedded units
Synthesis	Universal or composite maturity score	Narrowed	One embedded case cannot validate general weights or thresholds	Dimension-specific diagnostic levels with evidence boundaries
Evidence strategy	Qualitative or quantitative evidence alone	Rejected	Work practice and implementation state require different evidence streams	Registers, coding, descriptive indicators, and technical evidence are triangulated
Technical emphasis	Process-mining-led evaluation	Deferred	Available logs lack sufficient lifecycle, assignee, and multi-event case quality	A log-quality gate replaces unsupported discovery or conformance claims
Product direction	Broad feature expansion	Deferred	Duplicate work and incomplete authoritative records recur across service areas	Existing workflows, ownership, and monitoring precede expansion
Advanced technology	Autonomous AI or migration-led implementation	Narrowed or deferred	Evidence supports one input friction but no autonomous decisions, production readiness, or migration case	A human-approved non-production scenario remains outside the core artifact

Figure 3.1 makes the construction chain explicit. Evidence is registered and assessed through the six dimensions, coded and triangulated under declared rules, and converted into bounded case decisions. Reapplication begins by scoping a process and assembling new evidence; it does not reuse CityFlow's capability assignments as if they were universal results.

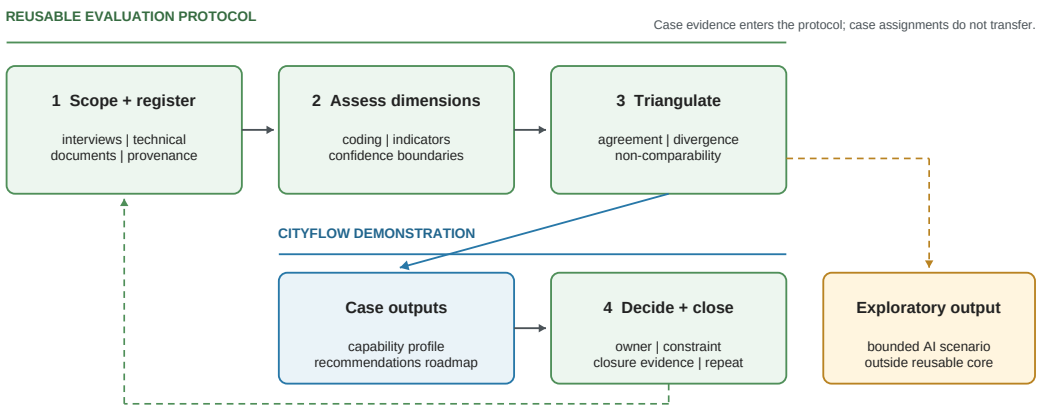


Figure 3.1: Evidence-to-artifact pathway. The evaluation protocol is reusable; the CityFlow capability profile and roadmap are demonstration outputs, and the AI scenario remains exploratory.

3.3 Case Selection and Units of Analysis

CityFlow was selected as an operational municipal workflow platform with production-relevant verticals under varied conditions, enabling analysis of municipal digitalisation beyond technical deployment alone.

The corpus comprises 14 sessions and 26 registered participants. Several sessions were joint interviews, making the session the primary unit of analysis; participant-level interpretation is used only when speaker or role attribution is reliable.

Sampling was purposive and access-constrained, seeking contrasting services and stakeholders. Participants did not assess non-participating colleagues, and group sessions prevent role comparisons where attribution is uncertain.

The operational comparison focuses on CROA, Espaços Verdes, and Mercados, with additional internal stakeholder interviews from APD, DMSI/CMP, and the supplier side. The technical scan also includes Ocorrências/Reporta, but that vertical is not used for operational interview comparison because no operational-user interview was collected.

The sampling sequence was staged but adaptive. CROA was the initial pilot because it was an established production service area with mature documentation and sufficient access to test the interview and evidence procedures. Espaços Verdes followed as a contrasting variable-field context in which connectivity, task variability, and handoffs could challenge a standardised workflow. Mercados provided a third operational context with daily tablet use and a different planning and reporting pattern. Internal-stakeholder interviews then overlapped with the later operational sessions in May as governance, quality, supplier, and monitoring questions became more specific. The order was therefore purposeful and iterative, not a strict experimental sequence.

Table 3.2 separates covered stakeholder perspectives from deliberate or access-constrained omissions and states the consequence for interpretation.

Table 3.2: Stakeholder coverage, exclusions, and claim consequences.

Stakeholder group	Coverage	Sampling role or reason for absence	Claim consequence
Operational users	S02, S05-S06, S08-S11	Daily and field-work perspectives across three service areas	Supports session-level fit and friction findings across the sampled service areas
Supervisors and technical-superior users	S01, S03, S07-S11	Operational oversight, coordination, and record-use perspectives	Supports management-value findings within sampled service areas
APD service design/former product ownership	S12	Added design history and service perspective after initial operational work	Supports context and governance interpretation, not final-artifact acceptance
APD tester / quality assurance	S13	Direct account of front-end manual and automated testing, acceptance criteria, release quality, and feedback practices	One tester; no backend/API-testing coverage or representative QA sample
Supplier project/development perspective	S14	Requirements elicitation, service/process design, supplier-team coordination, and direct programming of CityFlow functionality	One hybrid role; no team-wide developer view or dedicated BPMN-modeller evidence
DMSI governance and indicator perspective	S15	Added municipal governance, monitoring, and prioritisation context	Supports governance interpretation within one stakeholder account
Additional developers, dedicated BPMN process modellers, backend/API testers, and integration owners	Not separately interviewed	Access and internship-time constraint; S13 and S14 cover bounded facets rather than role cohorts	No team-wide consensus or specialised backend, modelling, or integration-effort claim
Ocorrências/Reporta operational users	Not interviewed	No dedicated operational session was obtained	Technical inventory only; no operational fit, value, or adoption conclusion

Stakeholder group	Coverage	Sampling role or reason for absence	Claim consequence
Current product leadership	Embedded/documentary context only	No dedicated research interview was recorded with the current role	No claim of formal product-owner endorsement or acceptance of the final artifact
Citizens	Not interviewed	Citizens are not direct CityFlow operators and citizen-outcome evaluation was outside scope	No direct citizen-experience or public-impact conclusion

3.4 Evidence Sources

The study combines five evidence streams. First, interview evidence captures daily use, operational fit, perceived value, barriers, and improvement priorities. Interviews included operational users, supervisory or technical-superior profiles, APD stakeholders including one front-end tester/quality-assurance participant, one DMSI/CMP stakeholder, and one supplier-side participant with a hybrid project/development role and direct programming experience, following semi-structured interview practice (Brinkmann and Kvale 2018; McGrath, Palmgren, and Liljedahl 2019).

Second, documentary evidence reconstructs project scope, governance, priorities, requirements, and analytical scope decisions. This includes APD communications, CityFlow requirement/context material, project records, course and formatting requirements, and supporting analysis outputs.

Third, BPMN and Camunda evidence provides technical indicators of implemented workflow structure: process definitions, versions, task and gateway counts, service-task indicators, deployment counts, and historic-instance counts. These indicators describe technical implementation and activity; they do not by themselves prove operational maturity.

Fourth, Elasticsearch and related CityFlow extracts provide evidence of local operational data volume, index coverage, module-level records, and available metadata. These extracts are documented as local, pre-production, or pipeline-derived evidence and are interpreted as one input to triangulation rather than as standalone measurements of production impact. Additional validation work separated raw or private platform material from dissertation-facing derived outputs. Its role was claim safety: to test whether front-end, Camunda, and Elasticsearch evidence could support claims about product structure, monitoring feasibility, status semantics, and process-mining readiness. Because the accessible live platform endpoint was pre-production unless independently confirmed otherwise, validation evidence strengthens technical feasibility claims but does not convert them into production usage or impact measurements.

Fifth, an event-log audit assesses process-mining feasibility. The available event log contains many one-event cases, many cancelled events, limited assignee information, and mixed process coverage. Process mining is scoped as an exploratory technical assessment rather than conclusive process discovery, bottleneck analysis, cycle-time analysis, or conformance checking, consistent with established quality criteria for event logs (W. Van Der Aalst 2016).

3.5 Data Preparation, Coding, and Descriptive Analysis

Before analysis, each anonymized session was linked to its recording, participants, context, and analytical unit to preserve source traceability (Patton 2015). Direct quotations were checked

against the session recordings. Group-session evidence was attributed at session level unless the speaker or role could be identified reliably. Portuguese quotations are followed by the author's English translation in square brackets.

Interview evidence was coded using the adopted structured codebook organised around the six analytical dimensions defined in Chapter 2, with a seventh category for recommendation type. QualCoder, an open-source computer-assisted qualitative data analysis application, was used to attach codes to source passages and retrieve them by case and category. The maintained project contains 14 sources and 14 cases. The final synthesis contains 96 validated participant-evidence segments organised into 7 categories and 22 codes; non-participant speech identified during source review was excluded. The codebook and resulting frequencies are reproduced in Appendices A and I. QualCoder supported organisation and retrieval; analytical interpretation remained with the researcher (Braun and Clarke 2006).

The coded indicators were transferred to a session-level indicator grid covering fifteen dimensions: usage frequency, usage motivation, value recognition, operational fit, usability barriers, digital-literacy barriers, training needs, resistance, duplicate work, missing functionality, and governance, monitoring, improvement, new-development, and evidence-confidence indicators. The full grid is reproduced in Appendix D. Unknown values were retained where evidence was insufficient for classification, to avoid false absence.

Descriptive statistics were computed from the indicator grid using frequencies, valid percentages, and crosstabs by vertical where cell sizes permitted (Pallant 2016). They describe the 14 assessed sessions and are used for case comparison rather than statistical generalisation.

3.6 Technical and Documentary Analysis

Technical analysis compares production-relevant verticals using BPMN, Camunda, Elasticsearch, and event-log indicators. The core indicators include process counts, task counts, service-task counts, gateway counts, version indicators, deployment counts, historic-instance counts, Elasticsearch document counts, index coverage, and event-log quality metrics.

The analysis deliberately avoids equating technical activity with operational success. A vertical can be technically active but misaligned with field work, or operationally valued while still requiring governance, training, integration, or data-quality improvements. Technical evidence is therefore interpreted together with interview and documentary evidence.

Documentary analysis supports the interpretation of governance, priorities, and institutional constraints, drawing on official communications, requirements records, and project documentation.

3.7 Triangulation, Maturity Synthesis, and AI Scenario

Table 3.3 summarises how the analytical dimensions, evidence streams, and research questions are integrated through the evaluation framework.

Table 3.3: Evaluation framework architecture and artifact boundary.

Framework layer	Inputs	Main outputs	Research-question link
Literature-derived dimensions	Operational fit, adoption, value, data/monitoring, governance, technology	Analytical categories	Structures RQ1, RQ2, and RQ3
Reusable evidence protocol	Session and participant registers, coding rules, indicator grid, technical and documentary registers	Traceable evidence base	Supplies RQ1, RQ2, and RQ3
Reusable synthesis protocol	Evidence matrix, agreement/divergence rules, diagnostic levels, claim boundaries	Triangulated and confidence-bounded findings	Integrates RQ1, RQ2, and RQ3
Reusable application procedure	Process scope, evidence assembly, assessment, prioritisation, ownership, closure evidence	Repeatable decision cycle	Converts findings into bounded action
CityFlow demonstration outputs	CityFlow evidence processed through the protocol	Capability profile and recommendations roadmap	Answers RQ1, RQ2, and RQ3 for this case
Exploratory case output	Field-entry friction, technical constraints, and governance boundary	Human-approved voice-to-record scenario	Bounded RQ3 exploration outside the core artifact

Triangulation integrates qualitative evidence, PSPP descriptive outputs, technical indicators, and documents. Agreement across sources increases confidence; disagreement is retained as an analytical finding because gaps between technical deployment, recorded data, and lived operational use are central to evaluating CityFlow.

The maturity synthesis uses four levels: fragmented; deployed but constrained; operationally embedded; and managed and evolving. A level is assigned only when the evidence supports the relevant dimension. Technical volume alone cannot establish operational maturity, and unknown dimensions remain unscored.

The bounded AI scenario belongs to the future-evolution dimension. It explores whether a human-approved assistant could transform dictated or free-text field accounts into proposed structured records, a narrow case-specific application of the broader LLM-supported BPM tasks discussed by Vidgof et al. (Vidgof, Bachhofner, and Mendling 2023). The scenario is scoped as a non-production design exploration, separate from maturity scoring; it does not involve legal, clinical, or enforcement decisions.

3.7.1 Formative Review and Evaluation Boundary

The evolving research design and artifact were exposed to four documented formative review points. These reviews changed framing, scope, and reporting, but none administered a summative acceptance protocol to intended users of the final framework.

Table 3.4: Formative review history and limits.

Date	Reviewer role/context	Version reviewed	Resulting change	Evaluation limit
2026-03-03	Academic review in the CITE context	KPI-oriented proposal	Separated monitoring, adoption, value, and maturity; realigned the questions and literature	Concept review, not artifact use
2026-03-04 to 2026-03-05	APD and DMSI scope discussions	Candidate cases and data plan	Prioritised active production service areas and contrasting operational contexts	Scope review, not framework acceptance

Date	Reviewer role/context	Version reviewed	Resulting change	Evaluation limit
2026-06-15	APD interim-presentation audience	Preliminary evidence synthesis	Clarified case selection, chronology, participant profiles, and practical recommendations	Interim communication review before the final artifact
2026-07-16	Academic supervisor	Full dissertation draft	Tightened evidence boundaries, capability logic, figures, roadmap actions, and closure tests	Manuscript review; no final stakeholder acceptance test

The interviews themselves were diagnostic: they elicited CityFlow use, operational fit, value, and improvement needs before the final framework was assembled. Participants were not shown an alpha version of the framework, and the study does not retrospectively claim interview-based artifact iteration. Evaluation therefore consists of formative review of successive research versions and a naturalistic demonstration on CityFlow. Summative stakeholder validation of the final reusable protocol remains unperformed.

3.8 Research Quality, Ethics, and Limitations

Research quality is supported through source traceability, explicit counting rules, documented interpretation criteria for technical data, and separation between source evidence and instrumental tool use. Major findings are linked to the coded interviews, descriptive tables, technical outputs, documentary evidence, or triangulation matrix.

Participants are identified by anonymized labels. No personally identifiable or confidential operational data is disclosed in the dissertation.

The intended framework users–APD product management, municipal service owners, DMSI, and supplier representatives–did not complete a final acceptance test of the assembled artifact. The formative reviews support design evolution and the CityFlow application demonstrates practical use, but neither is reported as summative validation. This distinction is carried into the limitations and future-work sections.

The internal-stakeholder sample included direct but bounded technical perspectives. S14 described personal programming experience, requirements work, and coordination of the supplier development team; S13 described front-end manual and automated testing, acceptance criteria, and release-quality work, while explicitly excluding backend/API testing from that role. These accounts support first-hand claims about those participants' practices, but they are not distinct or representative developer and QA samples. The study therefore does not claim team-wide developer or QA consensus, comprehensive implementation difficulty, dedicated BPMN-modeller experience, backend/API-test coverage, or an integration-owner perspective. A follow-up design should broaden each technical role across multiple participants and add task-based observation of operational users.

Chapter 4 presents the case context against which this methodology was applied.

4 Case Study Context: CityFlow

4.1 Organizational Setting: Associação Porto Digital

The Associação Porto Digital (APD), founded in 2004, coordinates the Porto Municipality's urban digitalisation strategy through product management, service design, engineering, data, infrastructure, and project competences. CityFlow sits within its Innovation, Experimentation and Digital Services unit; product management coordinates priorities and stakeholders, supported by engineering and service design.

This organizational positioning matters for the case study. CityFlow is not an external consultancy project or an academic prototype. It is an internal municipal platform developed iteratively across verticals, shaped by operational needs, resource constraints, and institutional priorities. The researcher's internship was embedded in the product management team, which provided direct access to project documentation, stakeholder communications, platform evidence, and operational interviews under APD and Porto Municipality confidentiality conditions.

4.2 CityFlow: A BPMN-Based Municipal Workflow Platform

CityFlow is a web platform for municipal workflow digitalisation built on a business process model and notation (BPMN) engine (Dumas et al. 2018). It aims to integrate municipal operations, standardise workflows, improve coordination, and record execution with traceable responsibility.

Deployment proceeded by service area. Phase 1 introduced CROA; Phase 2 extended the platform to Espaços Verdes, Mercados, Reporta/Ocorrências, and supporting administrative modules.

Table 4.1 summarizes selected production-relevant process definitions checked in the 15 June 2026 targeted technical scan. All were active (not suspended) at the time of capture. The complete 29 April process-count inventory is reported in Appendix H.

Table 4.1: Active Camunda process definitions checked in the 15 June 2026 targeted technical scan.

Vertical	Main process definitions	Version / status
CROA	modulosDoCROA; geralOcorrencia	111; 12, active
Espaços Verdes	modulosDoAmbiente; geralPlanMensalAmbiente	42; 12, active
Mercados	geralPlanSemMercados	9, active
Ocorr./Reporta	geralOcorrenciasReportaPorto	9, active

Detailed task, gateway, service-task, and historic-instance counts are reported in Appendix H.

Figure 4.1 presents the captured CROA occurrence process (geralOcorrencia, version 12) in a compact BPMN layout. It retains direct archival, continued field analysis, and conditional reopening. The initial gateway evaluates the request status; subsequent gateways evaluate a task decision. The Archived step is a user task with onward transitions, rather than a terminal event.

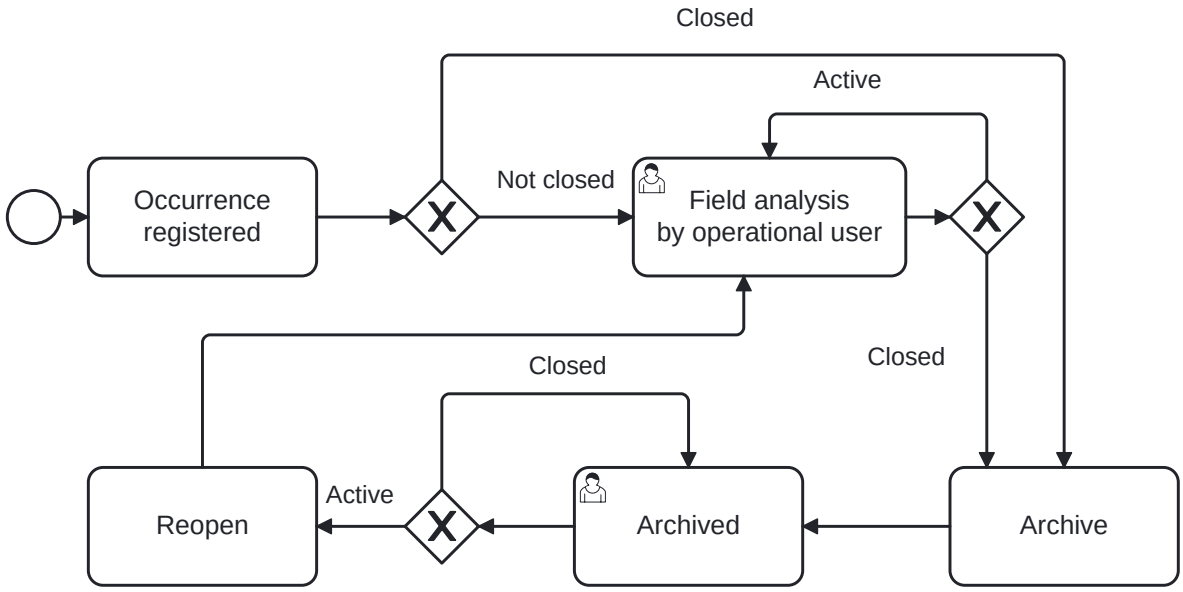


Figure 4.1: CROA occurrence BPMN process, based on geralOcorrencia version 12 captured on 15 June 2026.

Version numbers record deployment iteration, not maturity. CROA's main module reached version 111, Espaços Verdes version 42, and the Mercados and Ocorrências/Reporta processes version 9.

CityFlow layers Camunda Platform 7.15.0 over Elasticsearch 7.15.2, exposed through a ReactJS interface and NodeJS backend. PostgreSQL stores supplementary state; NiFi and MinIO support pipelines and objects; OAuth/LDAP integrates municipal identity; and Power BI forms the documented reporting layer.

This architecture reflects a deliberate choice to assemble the platform from established open-source components rather than building proprietary equivalents. Figure 4.2 shows the layered architecture and primary data flows.

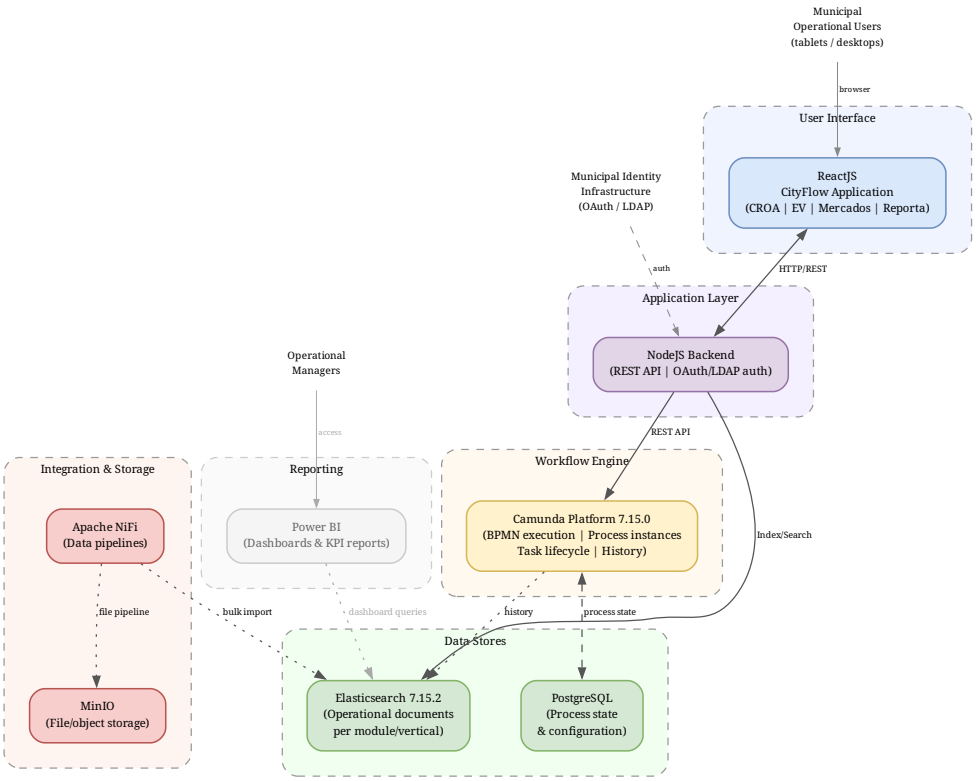


Figure 4.2: CityFlow information system architecture and primary data flows.

Table 4.2 summarises the component stack and main data flows.

Table 4.2: Simplified CityFlow technical architecture. Component versions as documented at the 15 June 2026 targeted capture.

Layer	Component	Role in CityFlow
User interface	ReactJS CityFlow application	Module-specific browser interface for CROA, Espaços Verdes, Mercados, Ocorrências/Reporta, and related verticals
Application layer	NodeJS backend	Mediates REST API calls, authentication, and access to platform services
Workflow engine	Camunda Platform 7.15.0	Executes BPMN process definitions, process instances, task lifecycles, and history queries
Operational data store	Elasticsearch 7.15.2	Stores module and vertical documents for operational recording and aggregate analysis
Supporting state	PostgreSQL	Stores supplementary process state and configuration data
Integration and storage	Apache NiFi and MinIO	Supports data pipelines, file/object storage, and integration with legacy or external data sources
Reporting layer	Power BI	Business intelligence dashboard layer for operational monitoring and KPI reporting

The study directly examined selected interface views, Camunda definitions and metadata, and Elasticsearch aggregates. Other components provide system context but were not performance-tested. Interface capture confirms functional presence, not task-based usability. The modular

architecture accommodates varied municipal records but requires cross-stack monitoring and integration.

The platform's strategic rationale reflects the Porto Municipality's broader digital transformation objectives: move from paper-based or siloed digital practices toward integrated, traceable, and auditable workflow management. In practice, this means each municipal vertical that adopts CityFlow replaces or augments its existing operational recording methods with structured BPMN processes and centrally indexed operational data.

4.3 Operational Service Areas ("Verticals")

In APD project vocabulary, a "vertical" is an operational service domain or module family, not a level of organisational hierarchy. The dissertation examines four such service areas that represent distinct municipal operational contexts, following the embedded case design logic in which sub-units are selected to illuminate different facets of the main case (Yin 2018). Three verticals are assessed through both qualitative interview evidence and quantitative technical indicators. The fourth vertical, Ocorrências/Reporta, is included in the technical scan but does not have dedicated operational interview coverage.

4.3.1 CROA: Centro de Recolha Oficial de Animais

CROA is the Porto Municipality's official animal collection centre responsible for the capture, shelter, veterinary care, and adoption management of stray and abandoned animals. It operates under regulated procedures with public health, animal welfare, and legal compliance obligations.

CROA was the first vertical to go into production on CityFlow (Phase 1, Version 1.0). Its operational processes include preventive interventions, cage and cadaver management, occurrence reporting, and scheduled inspections. The Camunda process definitions include modulosDoCROA (CROA:GERAL, version 111) and geralOcorrencia (CROA:OCORRENCIA, version 12), confirming an extensive and repeatedly deployed BPMN footprint. At the 15 June capture, the corresponding Elasticsearch indices contained approximately 52,000 cadaver documents and 5,500 preventive-intervention documents. These figures describe the indexed footprint rather than unique completed cases. Figure 4.3 shows the CROA module home screen, confirming the presence of task, notification, and calendar affordances in the CityFlow application shell.

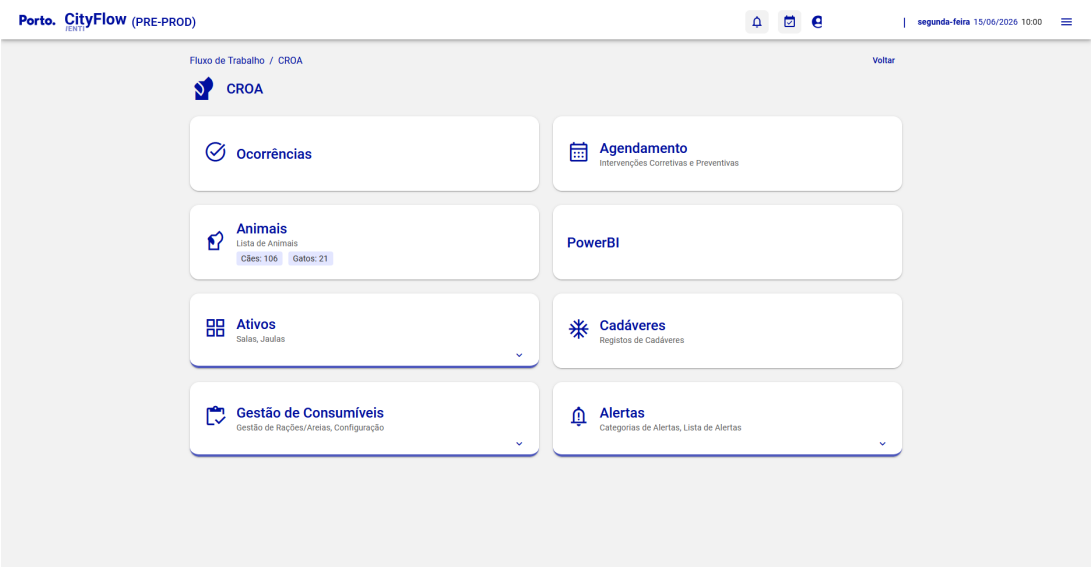


Figure 4.3: CROA module home screen.

Source: Phase 3 technical capture, 15 June 2026. Detail areas removed for publication.

CROA also has the deepest qualitative coverage in this study: six interview sessions across three interview rounds, covering operators, veterinary staff, supervisors, and the operational supervisor.

4.3.2 Espaços Verdes: Green Spaces and Environment

The Espaços Verdes vertical covers the municipal management of green spaces, parks, gardens, and related environmental maintenance. This vertical spans multiple operational teams distributed across the city, including the Palácio, Carvalhido, and Campinas locations.

The Espaços Verdes module was developed in Phase 2 and includes modulosDoAmbiente (Modulos Do Ambiente, version 42) and geralPlanMensalAmbiente (AMB:PLAN_MENSAL, version 12). The monthly planning process definition supports structured scheduling of maintenance activities. The Elasticsearch index for the monthly ambient plan contained approximately 9,400 records at the time of capture.

Three interview sessions were conducted with Espaços Verdes operational teams at the Palácio, Carvalhido, and Campinas locations. Each session included two participants: S08 at Palácio, S09 at Carvalhido, and S10 at Campinas. A further interview (S15) was conducted with the DMSI/CMP stakeholder responsible for environmental management coordination. Because the Espaços Verdes interviews were conducted as group sessions, their evidence is reported at session level.

4.3.3 Mercados: Markets and Fairs

The Mercados vertical covers the municipal management of markets and fairs, including stall allocation, fee management, inspections, and operational coordination. CityFlow supports the Mercados module through geralPlanSemMercados (MERCADOS:PLAN_SEM, version 9), covering weekly market planning.

At the time of technical capture, the Mercados weekly-plan index contained approximately 155,000 planning documents, making it the largest examined index by document count. It also contained distinct execution states (not executed, executed, archived, created, and cancelled). Quality-status values were present in a smaller subset and require service-owner validation before use as quality measures. Figure 4.4 shows the Mercados list view with the main module areas for market records, verification lists, planning, intervention execution, and alerts.

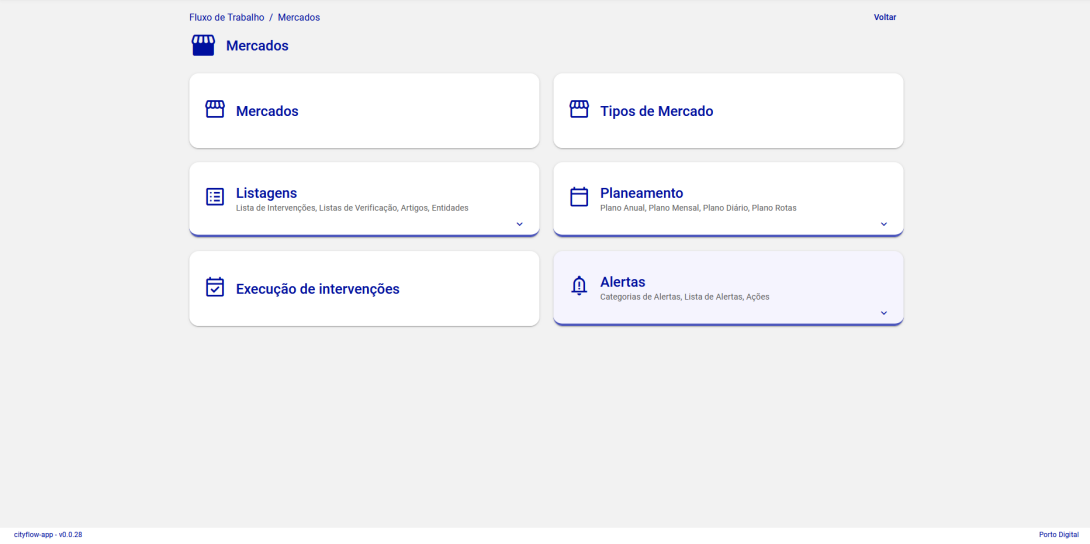


Figure 4.4: Mercados module home screen.

Source: Phase 2 technical capture, 15 June 2026. Browser and user context cropped.

One group interview session (S11) was conducted with Mercados operational staff, covering three street and tablet-based market operators and two technical-superior users who work primarily with the PC-based interface. One attendee was excluded from the participant count, resulting in five registered participants for this session. The single session limits within-vertical interpretation, but the group format captured interactions between operational and technical-superior perspectives that individual sessions would not have revealed.

4.3.4 Ocorrências/Reporta, Occurrence and Incident Reporting

Ocorrências/Reporta is a cross-cutting occurrence and incident reporting function that spans multiple verticals. It is implemented as geralOcorrenciasReportaPorto (rep:OCORRENCIA, version 9). A separate legacy `geralocorrencia` index contained approximately 5,800 documents and no records dated after December 2023; it is therefore used as inventory context rather than as a current Reporta activity measure. Figure 4.5 shows the Reporta occurrence list, including search, filter, status, and action affordances.

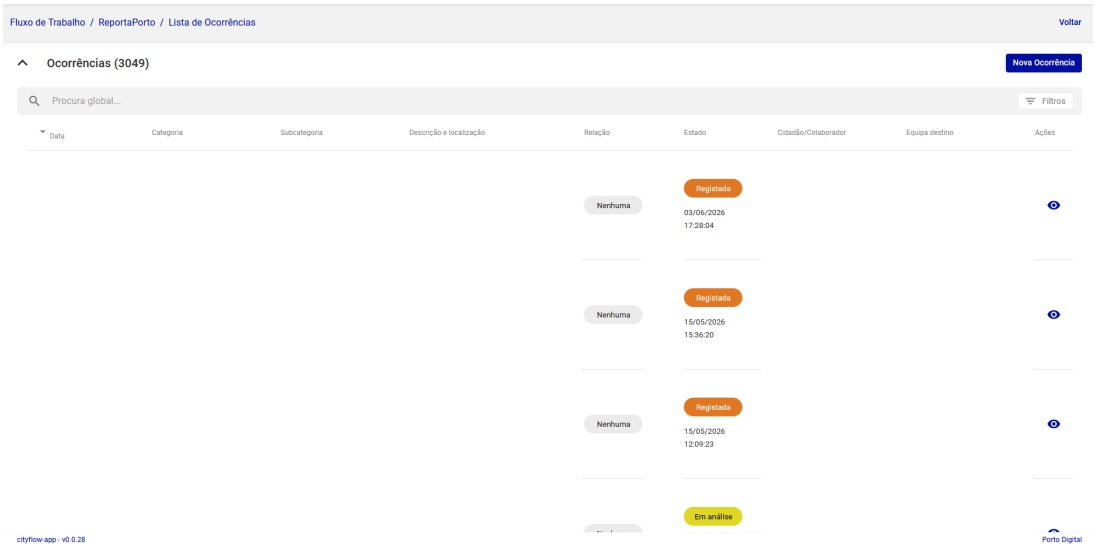


Figure 4.5: Reporta occurrence list with status and action controls.

Source: Phase 2 technical capture, 15 June 2026. Description and location fields removed for publication; browser and user context cropped.

Reporta appears in the technical scan because of its cross-vertical relevance. For this vertical, the evidence comes from technical indicators; the embedded case design concentrated interview coverage on CROA, Espaços Verdes, and Mercados.

4.3.5 Other Verticals and Modules

Additional verticals exist in the CityFlow project scope, including Mobilidade (mobility asset management), ROPI (event integration for itineraries and neighbourhoods), and DMGA (asset management). These verticals were at varying stages of requirement elicitation, prototype validation, or early development during the internship period. They are referenced where relevant to the platform's evolution path and recommendation roadmap, but they are not primary units of analysis in this dissertation.

4.3.6 Vertical Summary

Table 4.3 summarizes the key characteristics of the four production-relevant verticals examined in this study.

Table 4.3: Summary of production-relevant verticals examined

Characteristic	CROA	Espaços Verdes	Mercados	Ocorrências/Reporta
Primary process family	CROA modules	Ambiente modules	Mercados weekly plan	Reporta occurrences
BPMN version (main)	111	42	9	9
Captured ES documents (29 April)	60,000	13,000	229,000	5,800
Interview sessions	6	3 (+1 DMSI)	1	0
Registered participants	11	6 (+1 DMSI)	5	0
Production since	Phase 1 (2023-24)	Phase 2 (2024-25)	Phase 2 (2024-25)	Phase 2 (2024-25)
Interview coverage	Deepest	Moderate	Single group session	None

The asymmetry in interview coverage reflects the study's embedded case design. CROA received the deepest coverage because it combined the longest production history, the richest technical evidence, and the most accessible operational participants. Espaços Verdes and Mercados provide comparison at progressively less intensive levels, and Ocorrências/Reporta is included through technical evidence only.

4.4 Technical Architecture

CityFlow's technical architecture reflects a modern workflow-platform pattern assembled from open-source and commercial components.

The **Camunda BPMN engine** manages the process layer through deployed BPMN XML process definitions (Dumas et al. 2018). The engine creates process instances, manages task lifecycles, and records execution state. The six targeted process definitions checked in the 15 June scan are a representative subset of the broader process inventory reported in Appendix H. All six were active (not suspended) at the time of capture. Their version numbers describe deployment iteration; they do not, without operational and governance evidence, establish maturity.

The **Elasticsearch data layer** provides the operational data store. Each vertical or sub-module is represented by one or more indices with status-like and date fields suitable for aggregate monitoring. Several indices contain large bulk loads: 146,822 of 155,135 Mercados weekly-plan documents were dated April 2026, while 52,266 of 52,389 CROA cadaver documents were dated June 2025. The counts therefore describe stored footprint and possible migration or generation patterns, not completed operational activity. Analysis used aggregate counts and bucket distributions rather than record-level exports.

The **CityFlow front-end** is organised by vertical and exposes task, notification, calendar, search, filter, detail, and action controls. Phase 3 screenshots confirm functional presence across the examined modules, but not usability, execution correctness, or navigation completeness; those require task-based testing. Published figures omit operational detail.

Power BI dashboards and KPI reporting are part of CityFlow's documented architecture. The layer was not available for direct assessment, so findings on dashboard accessibility rely on interview evidence (Section 5.3.1).

Figure 4.6 compares BPMN inventory characteristics across the four service areas. The Espaços Verdes value of 23 is the number of latest process definitions in the 29 April 2026 inventory. Appendix H lists their keys and names and reports historic-instance probes separately. Mean version summarises deployment iteration.

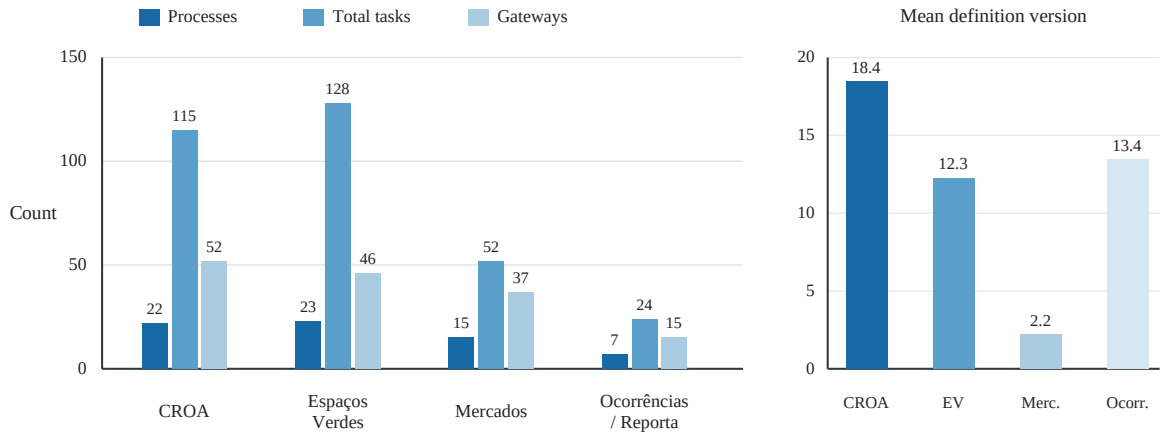


Figure 4.6: BPMN inventory characteristics by service area. Left panel shows latest process-definition, task, and gateway counts; right panel shows mean definition version as a descriptive iteration indicator. Inventory dated 29 April 2026; selected main definitions were refreshed on 15 June 2026 (Appendix H).

4.5 Stakeholder Landscape

CityFlow involves several stakeholder groups with different responsibilities: municipal operational users, APD product and service teams, DMSI/CMP governance stakeholders, and the supplier/development team. These roles shape how the platform is developed, governed, and used.

Municipal operational users include animal-care, green-space, market, and occurrence teams working through field tablets and office PCs. Differences in mobility, connectivity, task variety, regulation, and digital literacy make their comparison analytically useful.

APD product and service teams translate operational needs into priorities, process models, and interfaces.

DMSI/CMP coordinates municipal information systems and bridges APD's platform work with operational divisions.

The supplier team implements and deploys the platform stack and BPMN models.

Within the interview sample, S13 provided a direct APD front-end testing perspective covering manual and automated interface tests, acceptance criteria, and release quality. S14 combined requirements and service/process-design work, operational coordination of the supplier development team, and direct programming experience. These are first-hand but single-participant technical accounts; they do not represent either team as a whole.

CityFlow's governance operated through several parallel structures during the internship period. Weekly PDS (Ponto de Situação) meetings brought together APD product management, DMSI/CMP representatives, and the supplier's development team to review progress, escalate issues, and plan upcoming work. These meetings covered vertical status, server infrastructure planning, pipeline migration status, and resource allocation. APD product management and

service design conducted separate progress meetings to track against the product roadmap and coordinate user research activities. The researcher's internship was embedded within the product management function, providing direct observation of these governance structures.

Table 4.4: CityFlow stakeholder governance structure.

Governance layer	Participants	Primary function	Meeting cadence
Strategic coordination (PDS)	APD product management, DMSI/CMP, supplier development team	Progress review, issue escalation, resource allocation, vertical status	Weekly
Product oversight (Acompanhamento)	APD product management lead, service design lead	Roadmap tracking, user research coordination, requirement prioritisation	Weekly
Operational use (Vertical teams)	Municipal operational users, team supervisors	Daily platform use, feedback, friction reporting	Continuous

The governance structure reveals a potential gap between strategic coordination and operational use: PDS and progress meetings address platform development and roadmap progress, but the reviewed evidence did not identify a formal governance layer dedicated to operational monitoring, KPI review, or systematic feedback aggregation from vertical teams. This structural gap is consistent with the monitoring gaps identified in Chapter 5 and informs the governance recommendations in Section 5.5.

This dissertation focuses on the operational user perspective and the triangulation of interview evidence with technical and documentary evidence. Supplier, APD, and DMSI/CMP perspectives are included as internal stakeholder evidence that contextualizes governance, technical delivery, testing, priorities, and constraints.

4.6 Project Timeline and Current Status

Deployment age and evidence date provide context for interpreting differences between service areas. Figure 4.7 separates the requirements, implementation, extension, internship, and technical-capture periods.

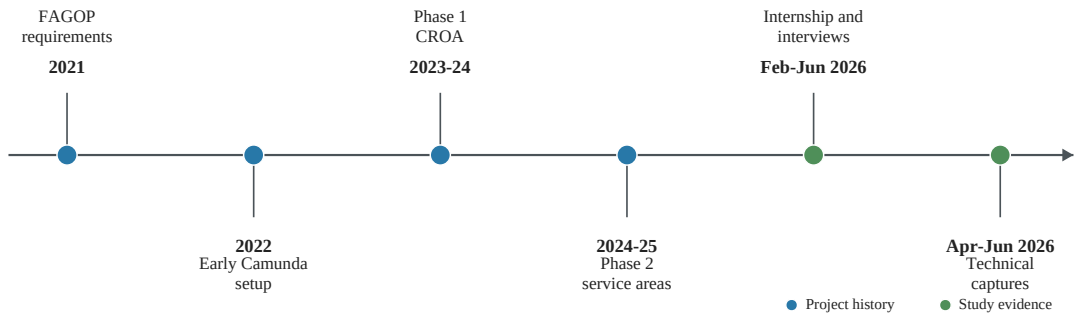


Figure 4.7: CityFlow project and study timeline. Dates are drawn from the project records and technical captures used in this dissertation.

Table 4.6: Timeline milestones and their analytical relevance.

Period	Project or study milestone	Relevance to the evaluation
2021	FAGOP service-design and requirements work	Established the initial workflow and requirements baseline
2022	Early Camunda setup represented in the retained deployment export	Provides technical history; operational meaning was not independently reconstructed
2023-2024	CityFlow Phase 1 and CROA production introduction	CROA has the longest observed production and refinement history
2024-2025	Phase 2 extension to Espaços Verdes, Mercados, Reporta, and administrative modules	Establishes the shorter comparison history of the Phase 2 service areas
February-June 2026	Embedded internship, interviews, documentary review, and platform analysis	Defines the period of primary evidence collection
29 April and 15 June 2026	Full inventory and targeted technical refresh	Defines the dates attached to the reported BPMN, Camunda, Elasticsearch, and interface evidence

The retained Camunda export contains 500 deployments: 85 dated 2022, 201 in 2023, 199 in 2024, 14 in 2025, and 1 in 2026. These counts describe deployment activity in that export; they do not reveal the purpose or operational impact of each deployment. During the internship, several service areas were in production, others were in quality assurance or prototype validation, and requirements work continued for ROPI and DMGA. A NiFi pipeline migration and a Camunda version update were also under discussion. These facts provide governance and technology context rather than outcome measures.

The 15 June 2026 targeted capture provides the closing technical snapshot of the selected definitions, indices, and modules described in this chapter.

Having described the case context, Chapter 5 presents the empirical results from the evaluation, organised around the three research questions, and discusses their implications for CityFlow's current value, improvement priorities, and future evolution.

5 Results and Discussion

This chapter reports value and fit (RQ1), improvement priorities (RQ2), future evolution (RQ3), triangulation and capability synthesis, and the resulting roadmap. Chapter 6 synthesises the limitations.

The results point to a differentiated pattern: CityFlow is consistently used in the sampled operational contexts, but its effectiveness depends on the fit between standardised BPMN workflows and local work practices. The analytical unit is the interview session. Counts and percentages refer to the 14 assessed sessions unless stated otherwise; participant counts describe the composition represented by those sessions. Ocorrências/Reporta is included in the technical scan but not in the operational results because no dedicated operational interview was collected for that service area.

5.1 Current Value and Operational Fit (RQ1)

5.1.1 Usage Frequency

Usage frequency was coded from the period and routine described in each interview, not from platform telemetry. *Daily* means use on each normal working day as part of the participant group's routine; *regular* means recurring use, usually several times per week or whenever a periodic workflow is due, but not every working day. These labels therefore describe cadence rather than an expected intensity: a periodic process can be valuable without daily execution.

Of the 10 sessions where operational use could be assessed, 7 reported daily use (representing 16 of 22 operational participants) and 3 reported regular use (representing 6 of 22). No operational session reported rare or non-existent use. The four internal-stakeholder sessions were classified as not applicable because those participants interact with CityFlow through governance, requirements, and oversight rather than routine operation. The sampled operational contexts therefore show consistent engagement. Whether that use translates into effective support is assessed separately through fit and value.

Table 5.1 presents the distribution by vertical.

Table 5.1: CityFlow usage frequency by context (N = 14 interview sessions).

Context	Interview sessions	Daily	Regular	Not applicable
CROA	6	4	2	0
Espaços Verdes	3	2	1	0
Mercados	1	1	0	0
Internal stakeholder	4	0	0	4
Total	**14**	**7**	**3**	**4**

The distribution supports routine use in the assessed workflows, not municipality-wide adoption. Sustained adoption also depends on operational value rather than mandate alone (Dumas et al. 2018; Kregel, Distel, and Coners 2022).

5.1.2 Operational Fit

Operational fit captures how well CityFlow aligns with the actual work practices, task sequences, and information needs of each operational context. Assessment used a three-level ordinal indicator. *Good fit* required explicit evidence that the main workflow was supported without a material workaround. *Partial fit* required evidence that core tasks were supported but that one

or more consequential gaps, parallel practices, or field-condition mismatches remained. *Poor fit* was reserved for evidence that the core workflow could not be completed meaningfully in CityFlow or was conducted mainly outside it. Insufficient evidence was coded unknown rather than inferred as poor. The complete rules are reproduced in Appendix A.

Of the 10 operational sessions, 7 were assessed as partial fit (17 of 22 operational participants) and 3 as good fit (5 of 22); no operational session was assessed as poor fit. All three good operational assessments came from CROA (S01, S03, and S07). The four internal-stakeholder sessions supplied contextual assessments rather than operator-use evidence: three were partial and S14, the supplier participant with a hybrid project/development role, was good. The combined distribution in Table 5.2 is retained for corpus transparency, but the internal row does not substitute for operational fit evidence.

Table 5.2: Operational fit by context (N = 14 interview sessions).

Vertical	Sessions	Partial fit	Good fit
CROA	6	3	3
Espaços Verdes	3	3	0
Mercados	1	1	0
Internal stakeholder	4	3	1
Total	**14**	**10**	**4**

The qualitative evidence explains the gap between partial and good fit. In CROA, participants described CityFlow as well aligned with the core animal registration, intervention recording, and occurrence management workflows. As one participant summarized: "O CityFlow em si é funcional e é prático" [CityFlow itself is functional and practical] (S05, CROA operational supervisor). The three good-fit CROA sessions came from participants who worked primarily with these core workflows and had longer experience with the platform. This is consistent with the finding in the BPM maturity literature that process familiarity and iterative refinement improve user-perceived alignment (Rosemann and de Bruin 2005).

In Espaços Verdes, all three sessions assessed fit as partial. Participants described situations where the platform captured planned maintenance activities but did not fully accommodate the variability of field conditions, ad-hoc repairs, and multi-team coordination. One participant noted that the monthly planning process worked well for scheduled tasks but that the mobile interface did not always reflect the granularity of field-level decisions. Another noted a specific limitation of the platform's structured data entry: predefined task categories did not capture the full range of field work reported by the Palácio team (S08). The same constraint was echoed in the Campinas session (S10), where participants also described recurring disruption of planned weekly tasks by urgent field events. This finding echoes the tension between process standardisation and operational flexibility identified in the municipal BPM literature (Niehaves, Plattfaut, and Becker 2013; Pittaway and Montazemi 2020).

In Mercados, the single session assessed fit as partial. Participants recognised that CityFlow introduced structure to market planning and stall management, but they identified gaps in how fee calculations, inspection records, and historical market data were handled.

The pattern concerns workflow alignment rather than non-use: every operational session reported regular or daily use, and none had poor fit. Good fit was more common in the assessed CROA procedures; variable field conditions left Espaços Verdes partial despite frequent use. This interpretation remains bounded to the purposive sample.

5.1.3 Value Recognition

Value recognition captures the degree to which participants perceive CityFlow as delivering operational or organizational value. Assessment used a three-level indicator: limited, clear, or strong value recognition.

Of the 14 sessions, 2 reported limited value (4 of 26 participants), 6 reported clear value (10 of 26 participants), and 6 reported strong value (12 of 26 participants). Value recognition varied markedly by vertical.

Table 5.3: Value recognition by vertical (N = 14 sessions)

Vertical	Sessions	Limited	Clear	Strong
CROA	6	0	3	3
Espaços Verdes	3	2	1	0
Mercados	1	0	0	1
Internal stakeholder	4	0	2	2
Total	**14**	**2**	**6**	**6**

CROA participants uniformly recognised value at clear or strong levels. Three mechanisms materialise that value. **Traceability:** digital records replace or consolidate parts of the paper trail and make animal histories retrievable. **Operational coordination:** structured occurrence and intervention records make current work visible across roles. **Management visibility:** recorded histories support faster responses to operational and managerial queries. One CROA participant described the transition from paper-based cage records to digital registration as the single most important improvement, noting that the platform had become indispensable to daily operational oversight: “Já no dia é constante porque a gente como encarregados temos que ir vendo sempre se há serviços” [On a daily basis it is constant because we as supervisors have to keep checking whether there are services] (S07, CROA supervisor). Another CROA participant described the platform’s value for management visibility outside normal hours, explaining that CityFlow records made it possible to answer a decision-maker’s query almost immediately rather than waiting until the next working day (S01). This pattern is consistent with the public-value literature, which identifies traceability and accountability as dimensions of value that matter independently of efficiency gains (Harrison, Pardo, and Cook 2012).

Espaços Verdes participants reported weaker value recognition. Both operational sessions at the Palácio and Carvalhido locations assessed value as limited. Participants acknowledged that the platform brought structure to monthly planning, but they did not perceive it as transformative for their daily field work. Paper-based or informal coordination practices persisted alongside CityFlow for certain task types. The Campinas session assessed value as clear, reflecting a more positive experience.

The Mercados group session assessed value as strong. Participants emphasized the benefit of having a unified market planning system replacing multiple spreadsheets and paper records. As one participant explained, the platform enables real-time operational awareness that was previously impossible: “recebemos em tempo real e conseguimos agir logo” [we receive information in real time and can act immediately] (S11, Mercados group session). The same session also described the office-side benefit: field photos and updates became visible to back-office staff quickly enough to support coordination between field and PC-based users. A CROA participant described the application simply as “É um aplicativo simples, fácil de trabalhar” [It is a simple application, easy to work with] (S07, CROA supervisor). Internal stakeholders assessed value

at clear or strong levels, valuing the platform's role in standardizing processes across verticals and enabling data-driven management conversations. One stakeholder distinguished between operational and management value: "o maior ganho está ... nesta parte mais macro" [the greatest gain is ... at the macro level, enabling managers to have a big-picture view of their departments rather than only improving day-to-day task completion] (S12, APD service design stakeholder). The same stakeholder proposed a concrete before-and-after KPI: "métricas de tempo de realização de uma tarefa antes do CityFlow e depois do CityFlow" [metrics of task completion time before CityFlow and after CityFlow, measuring whether the platform actually reduced task completion time compared to the paper-based baseline] (S12).
The divergence between CROA and Espaços Verdes in value recognition, despite similar usage frequency, is an important finding. It suggests that value recognition is not simply a function of use. It depends on whether the platform addresses a concrete operational pain point, paper records being replaced by digital traceability in CROA, rather than adding a digital layer on top of work practices that were already functioning adequately.

5.1.4 Stratification by Stakeholder Category

The indicator grid permits only a coarse profile-level comparison: operational end-user sessions (S01-S03 and S05-S11) and individual internal-stakeholder sessions (S12-S15). Finer operator-versus-manager stratification is not defensible because several operational interviews were group sessions and speaker-to-role attribution is not consistently available. The detailed grid is in Appendix D. Two patterns matter for the analysis. First, value recognition is strong or clear in 8 of 10 operational sessions (18 of 22 operational participants) and in all 4 internal stakeholder sessions. The two limited-value sessions are both in Espaços Verdes, yet both still report daily or regular use, showing that lower value recognition does not equal non-use.
Second, fit is assessed using the ordinal scale defined in Section 5.1.2, applied consistently across both groups. Only 3 of 10 operational sessions (representing 5 of 22 participants) were assessed as good fit, while 7 of 10 (representing 17 of 22) were partial. Internal stakeholders similarly split one good and three partial. This similar session-level distribution corroborates the existence of a fit gap in the collected evidence, but it does not prove that the gap is role-independent or shared by unobserved peers; neither fine-grained role comparison nor peer-agreement measurement was part of the study.

5.1.5 Usage Motivation

Usage motivation further explains the value-recognition pattern. In CROA, use is predominantly mixed (5 of 6 sessions), combining organisational expectation with practical recognition that the platform replaced inadequate paper records. The one CROA session coded as obligation was the supervisory/administrative perspective (S07). In Espaços Verdes, all three operational sessions were coded as obligation, while Mercados was coded as mixed. This helps explain why Espaços Verdes can show daily or regular use but lower value recognition and uniformly partial fit: the issue is not resistance, but a mandated tool layered over field workflows that were not fully broken before digitalisation.

5.2 Improvement Priorities (RQ2)

5.2.1 Duplicate Work

Duplicate work captures whether participants reported performing the same or similar task in both CityFlow and another system or format. Of the 10 sessions where duplicate work could be

assessed, 9 reported duplicate work as present (21 of 22 operational participants) and 1 reported it as absent (1 of 22 operational participants). One CROA session (S03, the encarregada) was the only operational session assessed as having no duplicate work.

Table 5.4: Duplicate work in operational sessions (N = 10 sessions, 22 operational participants; internal stakeholders classified as not applicable)

Service area	Operational interview sessions	No	Yes
CROA	6	1	5
Espaços Verdes	3	0	3
Mercados	1	0	1
Operational total	**10**	**1**	**9**

Duplicate work is the most consistent negative finding across verticals. It appears as paper-digital duplication, Excel or legacy-system coexistence, and internal platform duplication caused by irreversible states. Espaços Verdes participants described planning or recording on paper before entering the result in CityFlow (S08), while CROA participants described clinical checklist duplication and having to recreate cancelled occurrences rather than reopen them. This is both an efficiency problem and a data-quality risk: parallel records can diverge, and users may continue to trust the older record system (Tejedo-Romero et al. 2022).
Figure 4.1 includes conditional reopening in the captured model. Whether that route was available for the reported cancelled records requires a targeted workflow and interface test.
The recommendation roadmap therefore separates two pathways. External parallel work requires integration, workflow redesign, or policy decisions about the authoritative record. Internal duplication caused by CityFlow's own state-transition behaviour requires configuration or product changes such as controlled state reversal or record merging.

5.2.2 Missing Functionality

Missing functionality captures whether participants identified gaps in what CityFlow currently provides for their operational needs. It was reported as present in 10 of 14 sessions, representing 21 of 26 participants.

Table 5.5: Missing functionality by vertical (N = 14 sessions, 26 participants).

Context	Interview sessions	Yes	Partial	Unknown
CROA	6	5	1	0
Espaços Verdes	3	2	0	1
Mercados	1	1	0	0
Internal stakeholder	4	2	0	2
Total	**14**	**10**	**1**	**3**

The qualitative evidence provides concrete examples. In CROA, participants identified missing functionality in clinical observation recording, stock management integration, and adoption workflow tracking. In Espaços Verdes, participants wanted better support for ad-hoc task creation, photo capture for field reports, and integration with external asset databases. The Palácio session specifically indicated that the predefined task list did not cover all actual field activities (S08). In Mercados, the main gaps concerned fee calculation, inspection scheduling, and historical market data querying. Internal stakeholders identified gaps in reporting, dashboard configurability, and cross-vertical data aggregation.

Missing functionality is expected in an actively developed platform, but here it also explains duplicate work: users often maintain parallel records because CityFlow does not yet cover a task they must document. Addressing the highest impact gaps would therefore improve fit and reduce duplication at the same time.

5.2.3 Improvement and Development Priorities

Two indicators captured the urgency of improvement: priority to improve existing capabilities and priority for new development.

For improving existing capabilities, 13 out of 14 sessions (25 of 26 participants) assessed the priority as high. The single medium-priority session was S03 (CROA, the encarregada who assessed operational fit as good and duplicate work as absent), representing 1 of 26 participants. This near-unanimity supports the interpretation that operational users and stakeholders see CityFlow as a platform worth improving, not as a failed experiment. It is consistent with Fisher's observation that process maturity assessment is most useful when it identifies improvement opportunities rather than just certifying current capability (Fisher 2004).

For new development, the picture was more varied: 1 session assessed it as low (1 of 26 participants), 8 as medium (13 of 26 participants), and 5 as high (12 of 26 participants). High new-development priority was concentrated in CROA (S06, S07) and Mercados (S11), while Espaços Verdes and internal stakeholders were predominantly medium. This pattern suggests that the verticals with longer platform experience have more developed ideas about what new capabilities would add value, a finding that supports the iterative, vertical-by-vertical development model.

5.2.4 Barriers

Table 5.6 summarizes four barrier types across the interview corpus. The counts combine the session-level indicator grid with coded interview evidence. Unknown indicates that the session did not provide enough evidence for classification.

Table 5.6: Barrier evidence across 14 sessions (26 participants)

Barrier type	Yes	Partial	Unknown	Not applicable
Usability barriers	7	3	0	4
Digital literacy barriers	2	4	5	3
Training needs	4	2	7	1
Resistance or inertia	2	5	5	2

Usability barriers were present or partial in all 10 operational sessions. The qualitative evidence contextualizes these: they concerned primarily the mobile interface on tablets used in the field, small touch targets, slow loading times in areas with limited connectivity, and navigation patterns that required more steps than the task warranted. These are consistent with the usability challenges identified in the service design literature for field-based municipal work (Foglieni, Maffei, and Villari 2014).

Digital literacy barriers were present or partial in 6 sessions, with 5 remaining unknown. A CROA session framed this partly as a generational issue, with older workers finding the platform more challenging to use (S07). Training needs were present or partial in 6 sessions and unknown in 7. The identified needs ranged from basic orientation to support after new modules and version changes.

Resistance or inertia was present in 2 sessions and partial in 5, while 5 remained unknown. Combined with daily or regular use in every operational session, the evidence points more strongly to workflow and usability friction than to outright rejection (Pittaway and Montazemi 2020; Niehaves, Plattfaut, and Becker 2013).

Additional friction patterns cut across these barrier types. CROA participants reported connectivity or platform reliability issues, including cases where media appeared to save but was later missing. Internal stakeholders connected training and release support to governance. The direct APD tester/QA account noted the absence of a formally designated product manager ("não temos gestora de produto" [we do not have a product manager]), quality concerns before production release, limited testing capacity, and informal user feedback after releases (S13). These observations match the governance gaps reported as present or partial in 6 of 14 sessions. They describe front-end manual and automated testing, not backend/API quality assurance or a team-wide QA view.

5.3 Future Evolution (RQ3)

The answer that emerges from the evidence has a clear hierarchy. Operational monitoring, dashboards, KPI reports, and cross-vertical views built on the platform's existing data foundation, is the nearest-term, highest-confidence priority. It is the path for which the evidence is strongest, the technical prerequisites are already partially met, and the operational demand is most consistent across verticals. Process mining, by contrast, is an exploratory assessment: the event-log quality deficits documented in Section 5.3.2 mean that process discovery and conformance checking cannot currently support actionable conclusions. The bounded AI scenario is furthest from production readiness, it addresses a real operational friction but is conditional on stronger data quality, monitoring, integration, and governance foundations. The remainder of this section presents these three paths in that order of maturity, grounding each in the evidence rather than in technological ambition.

5.3.1 Operational Monitoring: The Primary Answer to RQ3

The technical scan completed on 15 June 2026 confirmed that selected indices contain status-like and date fields suitable for aggregate monitoring. The legacy `geralocorrencia` index, whose dated records ended in December 2023, contained 3,517 documents labelled open and 2,219 labelled closed. The `geralplansemmercados` index contained execution-state buckets (not executed: 107,779; archived: 34,705; created: 7,458; executed: 5,057; cancelled: 136). Quality status was populated for 7,985 of its 155,135 documents (ok: 3,496; ne: 2,372; nok: 1,926; cond: 191).

These are cross-sectional bucket counts, not reconstructed state transitions. They establish that monitorable fields exist, but they do not establish case throughput, transition rates, process duration, or the operational meaning of every source label. A monitoring implementation would first need a validated data dictionary, timestamp semantics, denominator rules, and reconciliation against source-system totals.

However, monitoring gaps were reported in 4 of 10 operational sessions, with a further 4 classified as unknown. The qualitative evidence suggests that operational users and supervisors do not currently have reliable access to dashboards, KPI reports, or cross-vertical monitoring views. The technical scan identified Power BI in the documented architecture but did not independently test operational-manager access or dashboard usability; those claims therefore rest on the reviewed interview evidence, not on the researcher's access conditions. A DMSI stakeholder summarized

the automation gap: “Neste momento não temos nada, nada automático” [At this moment we have nothing, nothing automatic] (S15, DMSI stakeholder). The supplier participant, whose hybrid role included direct programming and development-team coordination, confirmed that technical usage and availability measures could be collected, while distinguishing these from meaningful impact measurement (S14). Registered-user counts alone do not distinguish active from inactive users, daily from occasional users, or usage by module and vertical. A Mercados participant described the data extraction limitation directly: “Ainda está limitado um bocado na tal extração de dados” [It is still somewhat limited in that data extraction] (S11, Mercados group session). The same group session articulated the integration vision: “O futuro tem que caminhar para todas as aplicações conseguirem comunicar” [The future has to move toward all applications being able to communicate] (S11). CROA participants acknowledged that stable macro indicators existed in Power BI but described them as insufficient for operational decision-making. This gap between data availability and data accessibility is consistent with the public-sector performance measurement literature, which identifies the distance between measurement infrastructure and measurement use as a persistent challenge (Van Der Kolk 2022; Lnenicka et al. 2024).

5.3.2 Process Mining Feasibility

The event-log audit confirmed that process mining cannot currently support conclusive analysis. The available event log contained many one-event cases, many cancelled events, limited assignee information, and mixed process coverage. These characteristics are consistent with logs extracted from a general-purpose workflow engine without dedicated log-quality governance (W. Van Der Aalst 2016). Recent work on process mining in municipal contexts (Voorst tot Voorst 2025) similarly found that municipal information systems record detailed execution data but that this data is fragmented and noisy, requiring substantial preprocessing before process mining can yield actionable insights, a finding that directly corroborates the exploratory treatment of process mining in this study. Process mining is retained as an exploratory technical assessment and is not presented as a current analytical capability, consistent with the established requirements for event-log quality in process mining methodology (W. Van Der Aalst 2016; W. M. P. Van Der Aalst 2011).

Case volume is an additional suitability condition. Automated discovery is not useful for definitions with only a handful of comparable completed cases; those processes are better examined through BPMN review, task walkthroughs, and targeted conformance questions. A future process-mining study should therefore select a process with sufficient repeated cases *after* cancellation-heavy bulk records are excluded, rather than aggregate low-volume and heterogeneous definitions merely to increase the event count.

5.3.3 Bounded AI Scenario

The bounded AI scenario explored a human-approved assistant that transforms dictated or free-text operational accounts into proposed structured CityFlow records. It uses non-identifiable test data and leaves correction and approval with the worker.

The scenario connects to a concrete operational need identified in the qualitative evidence: field workers in Espaços Verdes and CROA described situations where dictating a short operational note would be faster and more practical than navigating the mobile interface, particularly when wearing gloves, working in wet conditions, or handling animals. A voice-to-structured-record pipeline could reduce this specific friction while preserving human control over what is recorded. Recent work on large language models in business process management (Vidgof, Bachhofner,

and Mendling 2023) provides a technical foundation for such scenarios, but the gap between research prototypes and production-ready municipal systems remains substantial.
Figure 5.1 translates the scenario into a conceptual workflow: the assistant creates a proposed record and surfaces uncertainty before worker review and approval.



Figure 5.1: Bounded voice-to-structured-record AI scenario.

The architecture also constrains the scenario. CityFlow currently runs on Camunda Platform 7.15.0 with a NodeJS backend and Elasticsearch data layer. Camunda 8 would offer connectors, language-agnostic job workers, and conversation-oriented task patterns that are more directly aligned with an AI-assisted workflow, but migration from Camunda 7 would require process-model translation, API adaptation, backend changes, and validation across verticals. The migration effort keeps this scenario separate from maturity scoring and in the medium-term research horizon rather than the near-term roadmap.

5.4 Discussion: Triangulation and Maturity Profile

5.4.1 Triangulation Synthesis

The triangulation matrix (Appendix E, T01-T11) integrates interview indicators, technical platform data, and documentary evidence. It follows one explicit interpretive rule: technical breadth is not a proxy for use, fit, value, or maturity. Process-definition versions indicate deployment iteration; active flags indicate configuration state at capture; document counts indicate stored volume. None of these measures, alone, demonstrates how frequently people use a workflow or how well it supports their work.

Cross-stream corroboration. T01 and T02 show that CROA combines a broad technical footprint with interview evidence of traceability, rapid historical retrieval, and management visibility. The same interviews also identify paper or signature duplication and clinical gaps. The streams therefore corroborate that an implemented and valued platform can remain operationally incomplete; they do not establish that process-version iteration caused the reported value. At platform level, documentary portfolio gaps and internal-stakeholder accounts similarly point to incomplete ownership, prioritisation, and outcome-review routines (T06, T10, and T11).

Requirement and meeting records also document parallel practices described in the duplicate-work interviews.

Cross-stream divergence. Espaços Verdes combines 23 inventoried process definitions and high historic-instance probe counts with uniformly partial fit in the three sampled sessions (T03 and T04). Mercados combines substantial stored volume with partial fit in one group session, and part of that volume comes from a bulk import rather than sustained transactional use (T05). The mixed event log is large but analytically weak for process mining (T07). These cases demonstrate why technical activity cannot validate interview outcomes. Likewise, the fact that six targeted definitions were active at capture cannot explain, or independently confirm, the absence of poor-fit assessments.

Ocorrências/Reporta. This service area remains in the technical inventory but outside the operational comparison because no dedicated operational interview was collected (T09). This separation prevents a technical footprint from being reported as observed operational value.

Corroboration, divergence, and non-comparability together prevent any one stream from carrying the interpretation (Jick 1979).

Figure 5.2 shows the distribution of coded interview evidence across contexts and analytical dimensions, normalised by the number of eligible sessions in each context.

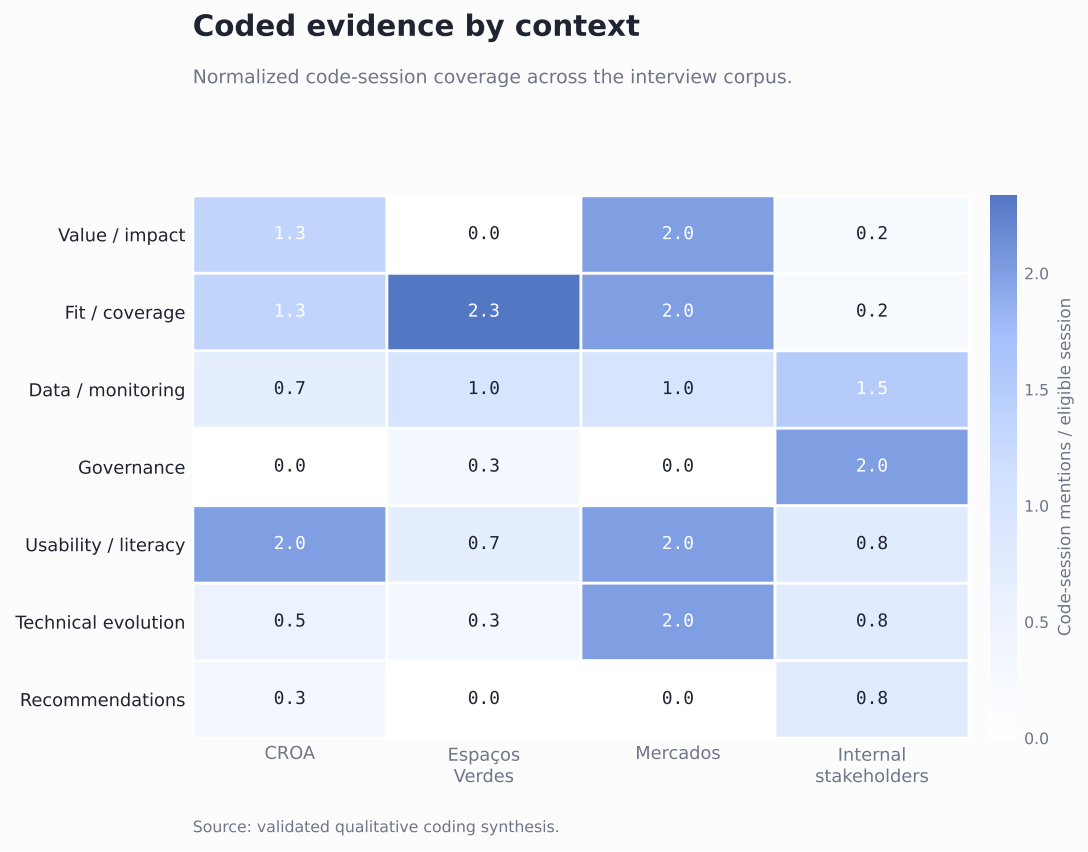


Figure 5.2: Normalised coded-theme coverage by context and analytical dimension.

The indicator grid supports differentiated cross-context profiles. CROA shows the strongest sampled pattern across value and fit and also has the broadest technical implementation history; Espaços Verdes shows regular use but uniformly partial fit and weaker value recognition;

Mercados combines daily use and strong value recognition with partial fit based on one group session. Two findings cut across the operational verticals: duplicate work is present in 9 of 10 operational sessions, and no operational session assessed fit as poor. These recurring session-level findings have broader evidence coverage than a single-service observation, but remain bounded by the purposive sample. The detailed indicator grid is retained in Appendix D.
Figure 5.3 juxtaposes session-level operational-fit assessments with process definition iteration for the three service areas with interview coverage. It visualises divergence between evidence streams without treating version as maturity; marker size shows gateway count.

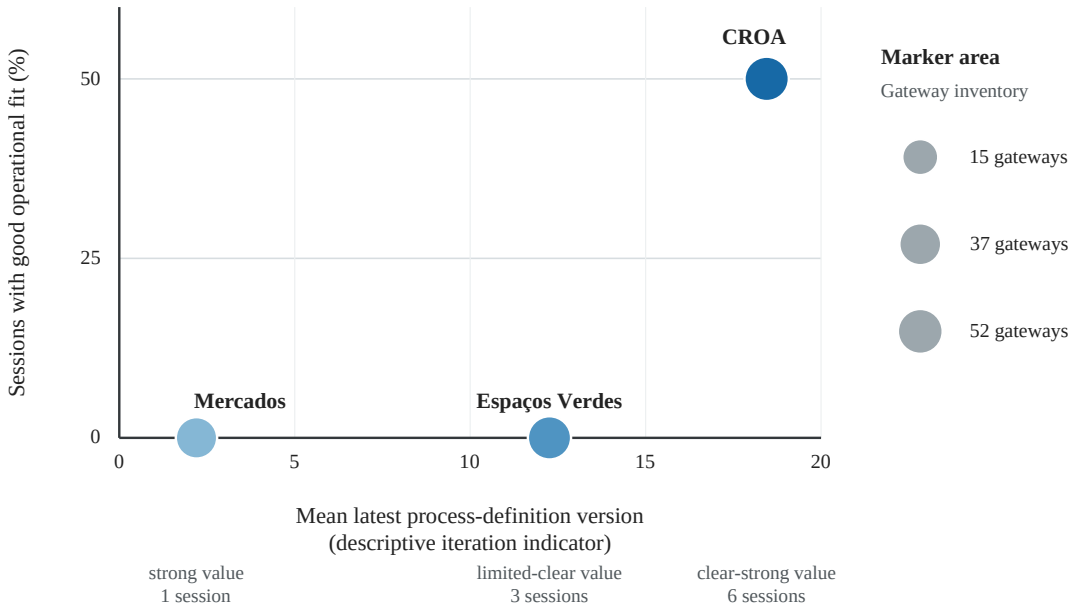


Figure 5.3: Operational fit and process-definition iteration by service area. Mercados rests on one group session.

5.4.2 Case-Specific Maturity Profile

The maturity profile is a diagnostic, case-specific synthesis grounded in the capability areas identified in Section 2.3: process design and implementation, metrics, ownership and governance, people, and information technology (Hammer 2007; Rosemann and de Bruin 2005). In line with evidence-based maturity-model design (Becker, Knackstedt, and Pöppelbuß 2009), each assignment is linked to the triangulation matrix (Appendix E), indicator grid (Appendix D), and claim register (Appendix G). Process-version numbers are inventory evidence only and never determine a level by themselves.
Four diagnostic levels are used:

- **Fragmented:** evidence confirms that the capability is absent, inaccessible, or handled mainly through disconnected practices.

- **Deployed but constrained:** the capability exists and is used, but material gaps, workarounds, reliability issues, or incomplete coverage remain.
- **Operationally embedded:** routine use and realised benefit are supported across the available interview and technical evidence, without a material contradiction for that dimension.
- **Managed and evolving:** operational embedding is supplemented by evidence of repeated improvement, ownership, measurement, or a feedback routine.

An evidence gap produces *unknown*, not *fragmented*. The levels are ordinal within a dimension but are not combined into a numeric score or a universal ranking. Mercados assignments carry medium confidence because operational evidence comes from one group session; CROA and Espaços Verdes assignments carry high confidence at the service-area level represented by the sample.

Table 5.7: Case-specific capability profile by service area and dimension.

Dimension	CROA	Espaços Verdes	Mercados	Key evidence
Process iteration	Operationally embedded	Deployed but constrained	Deployed but constrained	BPMN inventory plus T01, T03, T05, T08; versions are descriptive only
Operational data availability	Operationally embedded	Deployed but constrained	Deployed but constrained	Index coverage and T01-T05; Mercados bulk-import boundary retained
User-perceived operational fit	Operationally embedded	Deployed but constrained	Deployed but constrained	Appendix D; C04-C06
Value recognition	Operationally embedded	Deployed but constrained	Operationally embedded	Appendix D; C04-C06
Duplicate-work reduction	Deployed but constrained	Deployed but constrained	Deployed but constrained	9/10 operational sessions; C03
Monitoring and governance	Deployed but constrained	Fragmented	Fragmented	T02, T04-T06, T10-T11; C07-C08
Training and digital literacy	Deployed but constrained	Deployed but constrained	Deployed but constrained	Appendix D and reviewed session evidence

CROA has the strongest combined evidence for operational fit, value, and iteration, consistent with its longer implementation history. Its duplicate work and monitoring constraints show why technical iteration is not equivalent to whole-platform maturity.

CROA is *operationally embedded*, not *managed and evolving*: routine use and realised value accompany multiple deployments, but the reviewed evidence does not establish a complete repeated-improvement, ownership, measurement, or feedback routine. Version inventory documents history, not level.

Espaços Verdes is routinely used but remains constrained by variable field conditions, incomplete task categories, parallel handoffs, and limited monitoring. Mercados combines strong value recognition with partial fit and fragmented monitoring. Its large document count partly reflects an April 2026 bulk import and is therefore classified as deployed but constrained for data availability rather than as independently embedded.

The profile identifies where technical implementation, operating experience, and governance evidence align or diverge. It is an improvement diagnostic, not a ranking of service-area quality or a validated maturity model.

5.5 Recommendations Roadmap

The framework connects evidence to action through an explicit chain rather than treating recommendations as a separate brainstorming exercise. Table 5.8 shows how recurring findings become requirements, reusable decision rules, and CityFlow-specific actions.

5.5.1 Evidence-to-Requirement Chain

Table 5.8: Evidence-to-requirement-to-action chain.

Evidence and research question	Derived requirement	Reusable artifact response	CityFlow demonstration decision
C03; RQ2: duplicate work in 9/10 operational sessions	Identify the record, handoff, cause, and authoritative owner before proposing change	Authoritative-record and no-new-parallel-record rule	Audit external duplication and correct internal state behaviour before expansion
C04-C06; RQ1-RQ2: fit and gaps vary by service area	Test process suitability and exceptions in the operating context	Six-dimensional assessment plus workflow-suitability screen	Specify CROA clinical, Espaços Verdes field-task, and Mercados fee-calculation gaps separately
T02, T04-T05, C08; RQ2-RQ3: data exist but monitoring is constrained	Define semantics, denominators, lineage, and reconciliation before dashboard use	Data-evidence register and triangulation rules	Prototype CROA monitoring first, then validate each later service area
T06, T10-T11, C07; RQ2-RQ3: ownership and feedback are incomplete	Assign a decision owner and verifiable closure condition	Prioritisation, ownership, feedback, and closure protocol	Operate feedback, decision, release, and interoperability registers
T07, C09; RQ3: event-log volume exceeds analytical quality	Pass a declared quality gate before advanced analysis	Evidence-confidence rule and process-mining gate	Defer discovery and conformance analysis
T07, C12; RQ3: one field-input friction is evidenced but production readiness is not	Bound automation by human approval, privacy, and non-production testing	Exploratory-output boundary	Retain the voice-to-record scenario outside the core artifact

5.5.2 Framework Application Protocol

The intended users of the protocol are APD product management and quality assurance, municipal service owners, DMSI, and supplier technical representatives. Their responsibilities are complementary: product management coordinates the cycle and preserves the evidence trail; service owners validate operating consequences and accept workflow changes; DMSI validates data, integration, and municipal governance constraints; supplier technical representatives assess implementation evidence; and APD quality assurance defines and reviews test evidence within its agreed testing scope. Applying the framework consists of seven steps:

1. **Scope the process and decision.** Name the service area, workflow, operating boundary, decision owner, and question to be answered.
2. **Assemble the evidence.** Register interview or observation evidence, technical extracts, documents, dates, provenance, and known gaps.
3. **Assess the six dimensions.** Apply the codebook and indicator definitions without converting unknown evidence into a negative score.
4. **Triangulate.** Record corroboration, divergence, and non-comparability; technical volume never overrides evidence of operating fit.
5. **Prioritise a bounded response.** Apply evidence breadth, operational consequence, dependency, feasibility, and the workflow-suitability screen.
6. **Assign ownership and closure evidence.** Record the service owner, technical owner, constraint, target decision, and test that would close the item.
7. **Review and repeat.** Reassess after the agreed development cycle or a material process, data, or ownership change, retaining the previous evidence state for comparison.

The minimum outputs are an evidence register, dimension assessment, triangulation entry, bounded decision, named ownership, and closure result. A quarterly review is a practical default for stable service areas, while active changes should be reviewed at the end of the relevant development cycle. The cadence is an implementation proposal, not an observed CityFlow governance practice.

Worked CROA example (prospective application). C03 and the CROA evidence in T01-T02 identify duplicate handoffs and constrained monitoring, including the cancelled-occurrence state problem described in Section 5.2.1. The derived requirement is to designate the authoritative record and preserve an auditable state history without creating a replacement spreadsheet, paper register, or second CityFlow record. APD product management and the CROA service owner would scope the affected workflow and choose between an audited reopen, controlled reversal, or record-linking rule; the supplier would prototype the selected behaviour, APD quality assurance would test the interface behaviour against agreed acceptance criteria, and DMSI would review any integration or record-governance effect. Closure would require a manager-reviewed before/after handoff test, retained audit history, reconciled counts, and confirmation that no new parallel record was introduced. This is a worked use of the protocol, not evidence that the change was implemented or accepted.

5.5.3 CityFlow Improvement Register

The roadmap uses four explicit prioritisation criteria, following the process improvement logic in the literature (Hammer 2007; Fisher 2004):

1. **Evidence breadth:** repeated evidence across sessions and service areas is prioritised over a single request.
2. **Operational consequence:** duplicate entry, incomplete authoritative records, service interruption, and decision visibility take precedence over convenience features.
3. **Dependency:** data semantics, reliability, ownership, and monitoring precede process mining, AI, or cross-service aggregation.
4. **Feasibility:** changes to known workflows and existing data are sequenced before work that requires platform migration or unresolved integrations.

The outcomes below are expected consequences to be tested, not impact already measured by this study.

Before applying the roadmap to another municipal process, a workflow-suitability screen is required. A repeatable, stateful process with a clear authoritative record and manageable exceptions is a candidate for improving or extending an existing CityFlow workflow. A process with a stable objective but frequent exceptions should first be redesigned with configurable task categories, explicit exception paths, and field-compatible capture. Highly discretionary or poorly bounded work should not be forced into a rigid BPMN sequence; flexible record capture, integration with an existing authoritative system, or further process discovery may be more appropriate.

Table 5.9 operationalises three horizons. The next development cycle addresses external and internal duplication, the three highest-consequence functional gaps, and a CROA monitoring prototype. The following two cycles extend validated monitoring and establish feedback, prioritisation, release, and interoperability controls. Process mining, AI exploration, and cross-service aggregation remain medium-term because they depend on the earlier data-quality, governance, and user-control foundations.

Table 5.9 is the systematic improvement register. It maps every action to a research question, evidence, proposed decision owner, implementation constraint, and the evidence needed to close the item. The named roles are recommendations for accountability, not evidence of an organisational assignment already made.

Table 5.9: Systematic improvement register with ownership, constraints, and closure criteria.

# / horizon	Scoped action	RQ and evidence	Proposed decision owner	Principal constraint	Closure evidence
1a / Immediate	Resolve external paper, Excel, and legacy duplication	RQ2; C03	Service owner with APD product management; DMSI where integrated	A legal or external-system obligation may require a retained parallel record	Approved authoritative-record decision and before/after handoff test
1b / Immediate	Correct internal state-behaviour duplication	RQ2; C03-C04	APD product management and CROA service owner; APD QA validates	Auditability and legal state history	Tested reversal, reopen, or linkage with audit history and no duplicate record
2 / Immediate	Implement the three named functional-gap specifications	RQ1-RQ2; C04-C06	Respective service owner with APD product management; APD QA validates	Different verticals require distinct exception and record rules	Service-owner acceptance test and no-new-parallel-record check
3 / Immediate	Validate and expose a CROA monitoring prototype	RQ2-RQ3; T02, C08	CROA service owner with APD and DMSI data roles	Field semantics, reconciliation, and access control	Data dictionary, reconciled totals, role access, and manager review
4 / Near-term	Extend validated monitoring to Espaços Verdes and Mercados	RQ2-RQ3; T04-T05, C08	Respective service owners with APD/DMSI	Service-specific status meanings and the Mercados bulk-import boundary	Semantic validation and a reconciled service-area dashboard
5 / Near-term	Operate a structured user-feedback register	RQ2; T06, C07	APD product management	Named capacity and recurring review cadence	Owner, review record, decision, and closed-loop user validation
6 / Near-term	Operate evidence-based prioritisation and release communication	RQ2-RQ3; T10-T11	APD product leadership with service owners	Decision authority and transparent prioritisation criteria	Decision log and release note linked to evidence and validation
7 / Near-term	Complete the interoperability register	RQ2-RQ3; C03, C06, T11	DMSI and each authoritative-system owner	Data ownership, identifiers, privacy, and failure handling	Owner-approved mappings, error handling, and reconciliation control
8 / Medium-term	Reassess process mining after a log-quality gate	RQ3; T07, C09	APD technical/data owner with process owner	Stable cases, ordered activities, lifecycle semantics, and resource fields	Repeatable extraction passing the declared quality criteria
9 / Medium-term	Run a bounded AI design study	RQ3; T07, C12	APD product/QA and service owner with governance review	Approved non-identifiable data, privacy, error handling, and human confirmation	Usability, error, privacy, and approval results on non-production data
10 / Medium-term	Define validated cross-service measures	RQ3; T02, T04-T05, T11	DMSI data governance with service owners	Shared semantics, denominators, and source lineage	Metric definitions, lineage, reconciled values, and management review

The reusable artifact also requires operational custody. Table 5.10 maps each component to its research purpose, evidence, proposed custodian, principal constraint, and a completion condition. Custodianship means responsibility for maintaining the component during an application; it does not imply that these roles have already accepted the framework.

Table 5.10: Reusable artifact component responsibilities and completion conditions.

Reusable component	RQ and evidence	Proposed custodian	Principal constraint	Completion condition
Scope and evidence registers	RQ1-RQ3; C01, C13-C14	APD product management or designated evaluation lead	Provenance, confidentiality, capture dates, and access	Sources are dated, authorised, traceable, and explicit about gaps
Codebook and indicator grid	RQ1-RQ2; C02-C08	Evaluation lead with service-owner review	Session-level evidence and single-coder interpretation	Definitions are applied consistently, unknowns retained, and disputed classifications reviewed
Technical and documentary inventory	RQ1-RQ3; T01-T10, C13-C14	APD/DMSI technical or data owner	Cross-sectional and potentially pre-production evidence	Lineage and capture scope are recorded and reported totals reconciled
Triangulation and claim-boundary rules	RQ1-RQ3; T01-T11	Evaluation lead with the relevant decision owner	Evidence streams may be non-comparable or contradictory	Agreement, divergence, non-comparability, confidence, and boundary are recorded
Diagnostic maturity logic	RQ1-RQ2; C11	APD product management with service owners	Case-specific levels; no composite or universal score	Every assignment is evidence-linked or explicitly unknown
Prioritisation and roadmap logic	RQ2-RQ3; C03, C07-C10, T10-T11	APD product leadership with municipal decision owners	Authority, capacity, dependency, and service-specific semantics	Every action records its RQ, evidence, owner, constraint, and closure test
Application and review procedure	RQ1-RQ3; C15-C16	APD product management coordinating all intended-user roles	Summative usability and acceptance remain untested	One bounded cycle preserves inputs, decision, ownership, closure result, and review date

The results and their discussion lead to the concluding chapter, which synthesises the contributions, reflects on the methodology, acknowledges the limitations, and outlines directions for future work.

6 Conclusions and Future Work

6.1 Research Questions Revisited

The dissertation was structured around three research questions. This section summarizes the answer to each.

RQ1: How effectively does CityFlow support municipal workflow digitalisation across different operational contexts and user profiles? The evidence supports a differentiated answer, consistent with the BPM literature's finding that process-aware information systems deliver value in proportion to the alignment between process models and operational task structures (Dumas et al. 2018; Kregel, Distel, and Coners 2022). CityFlow supports workflow digitalisation effectively where process models align closely with operational task structures and where sustained investment in iterative refinement has occurred, as demonstrated in the CROA vertical. In newer or more operationally variable contexts such as Espaços Verdes, the platform is regularly used and valued but has not yet achieved the same depth of operational alignment. The cross-vertical pattern is not about resistance or non-use, every operational session reported regular or daily use, and no operational session assessed operational fit as poor. The pattern is about the distance between standardised BPMN workflows and the variability of municipal field work. CityFlow should therefore not be assumed to fit every municipal process. Repeatable, stateful work with a clear authoritative record is the strongest candidate. Work with frequent but knowable exceptions requires workflow redesign and flexible capture before automation; highly discretionary or poorly bounded work may be better served by integration, flexible record capture, or further process discovery.

RQ2: Which improvements to existing capabilities, missing developments, and governance arrangements should be prioritised? Three improvement clusters emerged from the evidence, ordered by strength and consistency. First, duplicate work, reported in 9 of 10 operational sessions, is the most pervasive friction and should be addressed through authoritative-record decisions, integration where appropriate, and correction of the platform's own state-transition behaviour. Second, missing functionality, reported in 10 of 14 sessions, directly causes some of the duplicate work and should be filled where it blocks complete operational workflows within a single system. Third, governance arrangements, prioritization criteria, release communication, and user feedback mechanisms, should be strengthened to support reliable improvement before extension to additional service areas.

RQ3: Which technological evolution path is most promising for CityFlow? Operational monitoring is the nearest-term priority, supported by an existing data foundation of status-like fields, date fields, and record volumes that are not yet translated into consistently accessible, semantically validated operational views. Process mining should remain exploratory until cleaner, process-specific event logs pass a documented quality gate. The bounded AI scenario addresses a real operational friction, hands-free input in field conditions, but should be treated as a human-approved design exploration rather than a near-term development priority.

Table 6.1 separates platform-level conclusions from service-area evidence. This prevents a finding from one context from being presented as a platform-wide property.

Table 6.1: Platform-level and service-area-specific conclusions and boundaries.

Scope	Supported conclusion	Boundary
Platform level	Routine use is present in all sampled operational sessions; duplicate work and improvement demand recur across contexts	Purposive 14-session sample; no population estimate
CROA	Strongest combined evidence of fit, value, traceability, and iterative implementation	Material duplicate work, clinical gaps, and monitoring constraints remain
Espaços Verdes	Routine use and planning support coexist with partial fit and weaker value recognition	Three sampled sessions; variable field work and incomplete task categories
Mercados	One group session reports daily use, strong value, and real-time coordination benefits	Partial fit; single-session evidence; document volume includes a bulk import
Ocorrências/Reporta	Technical implementation is present and suitable for scoped inventory analysis	No dedicated operational interview; no operational-value conclusion

6.2 Summary of Contributions

This dissertation set out to evaluate CityFlow, a BPMN-based municipal workflow digitalisation platform, through an embedded case study. The dissertation also developed the CityFlow Evaluation and Maturity Assessment Framework as a design-oriented artifact. Its reusable core comprises the six analytical dimensions, evidence registers, coding and triangulation rules, diagnostic maturity logic, and application procedure. The case-specific capability profile and recommendations roadmap are CityFlow demonstration outputs, while the bounded AI scenario is an exploratory output outside the core artifact.
The main empirical contributions are:

1. **A differentiated account of CityFlow's effectiveness across verticals.** The study found that CityFlow supports municipal workflow digitalisation effectively where process models align closely with operational task structures and where sustained investment in iterative refinement has occurred, as in CROA. In newer or more variable operational contexts, such as Espaços Verdes, the platform is used and valued but has not yet achieved the same depth of operational alignment.

2. **A structured, evidence-based improvement roadmap.** The study identified duplicate work as the most consistent negative finding across verticals, missing functionality as the most frequently cited gap, and governance arrangements as a maturing but still partial foundation. The recommendations are sequenced to address frictions before expanding ambition.

3. **A bounded feasibility assessment of advanced technical directions.** Process mining was assessed as exploratory due to event-log limitations. Operational monitoring was identified as the nearest-term priority with an existing data foundation. A bounded AI scenario was developed as a future-evolution design exploration tied to a specific operational friction.

4. **The CityFlow Evaluation and Maturity Assessment Framework.** The artifact provides an adaptable evaluation structure for comparable BPMN-based municipal workflow platforms by combining structured evidence registers, qualitative coding, descriptive indicators, technical evidence, triangulation, diagnostic maturity logic, and a repeatable procedure from process scoping to ownership and closure evidence. The CityFlow maturity profile and roadmap are results of applying the artifact, not reusable assignments. The profile is explicitly case-specific and diagnostic, not a validated universal model. In DSR terms, the

application of the framework to CityFlow constitutes a naturalistic demonstration of its utility (Peffers et al. 2007); formal stakeholder and multi-site evaluation is identified as a priority for future work.

6.3 Methodological Reflection

The embedded mixed-methods case-study design proved appropriate for the research problem. The combination of interview evidence, documentary evidence, and technical platform evidence enabled triangulation that no single evidence stream could have supported alone. The indicator grid provided a structured bridge between qualitative interview material and descriptive quantitative analysis. The six dimensions derived from the literature gave the interview guide, codebook, indicator grid, and final synthesis a common structure. QualCoder was used for the maintained coding project. The final synthesis contains 96 validated participant-evidence segments across 14 sources and 14 cases, organised into 7 categories and 22 codes. GNU PSPP, using SPSS-compatible syntax and data structures, enabled reproducible descriptive analysis.

The main methodological lesson is that evidence diversity is both a strength and a constraint. Corroboration can increase confidence, but technical volume, configuration state, and interview judgement measure different phenomena and must not be presented as mechanical validation of one another. Retaining divergence, unknown values, source dates, and claim boundaries was therefore as important as identifying agreement.

The design history also separates three forms of evaluation. Academic, organisational, and supervisor reviews between March and July 2026 were formative: they changed the problem framing, case selection, evidence boundaries, and reporting. Applying the completed analytical procedure to CityFlow was a naturalistic demonstration because it produced a traceable capability profile and roadmap from real case evidence. No interview participant was shown a final framework alpha, and intended users did not administer a summative acceptance test of the assembled artifact. Formative review and demonstration therefore support relevance and practical use without establishing stakeholder acceptance or external validity.

A replication should combine individual interviews, direct task observation, workflow walkthroughs, and longitudinal measures. Stratified coverage should retain operators, supervisors, and service owners while broadening the technical sample to multiple developers and QA staff, dedicated BPMN process modellers, backend/API testers, and data or integration owners. External, preferably multi-site validation remains necessary.

6.4 Limitations

The interview corpus is a purposive, access-based sample sized for an embedded case study. It is not representative and does not support statistical generalisation (Yin 2018). Participants were not asked whether unobserved colleagues would agree with their assessments. The Mercados service area is represented by a single group session, while Ocorrências/Reporta has no dedicated operational interview coverage. Group sessions are interpreted at session level, consistent with the dyadic joint-interview methodology described in Chapter 3. All direct quotations used in the dissertation were checked against session recordings.

The internal-stakeholder interviews included one supplier participant with direct experience of requirements work, development-team coordination, and programming CityFlow functionality, together with one APD tester responsible for front-end manual and automated testing. These interviews provide bounded developer-experienced and quality-assurance perspectives, but not distinct or representative samples of software developers or QA staff, and no dedicated BPMN

process modeller, backend/API tester, or integration owner was interviewed. Implementation and quality claims are therefore limited to these hybrid roles and triangulated with technical and documentary evidence; the study does not claim team-wide developer or QA consensus.
The technical platform evidence is based on local, pre-production, or pipeline-derived extracts captured in the 29 April inventory and targeted 15 June 2026 refresh. These extracts characterise the platform's technical footprint at those points in time, not current production performance. Status buckets are snapshots rather than reconstructed transitions, and process versions are implementation-iteration indicators rather than maturity scores. Process mining is scoped as exploratory because the event log has gaps that prevent reliable discovery and conformance analysis (see Section 5.3.2).
As noted in Chapter 5, the maturity profile is a diagnostic synthesis specific to the CityFlow case and its verticals.
Because the framework was assembled after the diagnostic interviews, the study contains no retrospective alpha test, final stakeholder acceptance test, or comparison of decisions made with and without it. The CityFlow demonstration, owners, closure tests, and formative reviews support practicality, but summative usability, decision usefulness, and adoption remain unmeasured.

6.5 Future Work

Within CityFlow: immediate priorities are the CROA monitoring prototype and a production-wide duplicate-work audit. Process mining should proceed only when a process-specific log passes the quality gate defined in Section 5.3.2.
Methodologically: APD product management and QA, service owners, DMSI, and supplier technical representatives should first apply the reusable protocol to one bounded decision, assessing traceability, usability, assignment consistency, decision usefulness, ownership, test evidence, and closure feasibility. A later multi-site evaluation could test transferability without importing CityFlow-specific assignments.
For the bounded AI scenario: conduct user research and a governed, non-production prototype before considering platform integration.
For the research community: comparative municipal studies should test whether high use with partial fit, and the gap between standardised models and variable field work, recur beyond CityFlow. Dutch municipal process-mining work also reports fragmented, noisy data requiring substantial preprocessing (Voorst tot Voorst 2025); this supports further comparison but does not establish transferability.

6.6 Final Remarks

CityFlow is an ambitious platform operating in a demanding institutional environment. Municipal workers use it daily in operational contexts that range from regulated animal shelter procedures to variable field maintenance and market administration. It has achieved real operational embedding in CROA and meaningful operational presence in Espaços Verdes and Mercados. It has not yet resolved the structural frictions of duplicate work, missing functionality, and partial operational fit that accompany that embedding.
The contribution is not a maturity certification, but a reusable protocol for identifying value, friction, priorities, and closure evidence. The CityFlow profile and roadmap demonstrate it in one setting, providing case-specific decision support and a basis for further evaluation of municipal BPMN-based platforms.

References

Becker, J., R. Knackstedt, and J. Pöppelbuß. 2009. "Developing Maturity Models for IT Management: A Procedure Model and Its Application." *Business & Information Systems Engineering* 1 (3): 213–222.

Benington, J., and M. H. Moore. 2011. "Public Value in Complex and Changing Times." In *Public Value: Theory and Practice,* edited by J. Benington and M. H. Moore, 1–30. Basingstoke: Palgrave Macmillan.

Braun, V., and V. Clarke. 2006. "Using Thematic Analysis in Psychology." *Qualitative Research in Psychology* 3 (2): 77–101.

Brinkmann, S., and S. Kvale. 2018. *Doing Interviews.* SAGE Publications Ltd.

Creswell, J. W., and V. L. Plano Clark. 2018. *Designing and Conducting Mixed Methods Research.* 3rd ed. Los Angeles: SAGE Publications.

de Bruin, T., and M. Rosemann. 2007. "Using the Delphi Technique to Identify BPM Capability Areas." *ACIS 2007 Proceedings,* 47.

Dessai, A., and V. Javidroozi. 2021. "Cross-sectoral Process Modelling for Smart City Development." *Business Process Management Journal* 27 (7): 2051–2074.

Dumas, M., M. La Rosa, J. Mendling, and H. A. Reijers. 2018. *Fundamentals of Business Process Management.* Berlin, Heidelberg: Springer Berlin Heidelberg.

Fisher, D. M. 2004. "The Business Process Maturity Model A Practical Approach for Identifying Opportunities for Optimization."

Foglieni, F., S. Maffei, and B. Villari. 2014. "A Research Framework for Service Evaluation." In *ServDes 2014: Service Future, Proceedings of the Fourth Service Design and Service Innovation Conference,* 99:423–428. Linköping Electronic Conference Proceedings 44. Lancaster, United Kingdom: Linköping University Electronic Press.

Gil-Garcia, J. R., T. A. Pardo, and T. Nam. 2015. "What Makes a City Smart? Identifying Core Components and Proposing an Integrative and Comprehensive Conceptualization." *Information Polity* 20 (1): 61–87.

Gregor, S., and A. R. Hevner. 2013. "Positioning and Presenting Design Science Research for Maximum Impact." *MIS Quarterly* 37 (2): 337–355.

Hammer, M. 2007. "The Process Audit." *Harvard Business Review* 85 (4): 111–123.

Harrison, T. M., T. A. Pardo, and M. Cook. 2012. "Creating Open Government Ecosystems: A Research and Development Agenda." *Future Internet* 4 (4): 900–928.

Hevner, A. R., S. T. March, J. Park, and S. Ram. 2004. "Design Science in Information Systems Research." *MIS Quarterly* 28 (1): 75–105.

IEEE. 2016. *IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams.* 1849-2016. IEEE.

Jick, T. D. 1979. “Mixing Qualitative and Quantitative Methods: Triangulation in Action.” *Administrative Science Quarterly* 24 (4): 602–611.

Jones, D., and S. Gregor. 2007. “The Anatomy of a Design Theory.” *Journal of the Association for Information Systems* 8 (5): 312–335.

Kregel, I., B. Distel, and A. Coners. 2022. “Business Process Management Culture in Public Administration and Its Determinants.” *Business & Information Systems Engineering* 64 (2): 201–221.

Lnenicka, M., A. Nikiforova, M. Luterek, P. Milic, et al. 2024. “Understanding the Development of Public Data Ecosystems: From a Conceptual Model to a Six-Generation Model of the Evolution of Public Data Ecosystems.” *Telematics and Informatics* 94:102190.

Looy, A. V., M. D. Backer, and G. Poels. 2014. “A Conceptual Framework and Classification of Capability Areas for Business Process Maturity.” *Enterprise Information Systems* 8 (2): 188–224.

McGrath, C., P. J. Palmgren, and M. Liljedahl. 2019. “Twelve Tips for Conducting Qualitative Research Interviews.” *Medical Teacher* 41 (9): 1002–1006.

Mendling, J., H. Reijers, and W. Van Der Aalst. 2010. “Seven Process Modeling Guidelines (7PMG).” *Information and Software Technology* 52 (2): 127–136.

Moore, M. H. 1995. *Creating Public Value: Strategic Management in Government.* Cambridge, MA: Harvard University Press.

Niehaves, B., R. Plattfaut, and J. Becker. 2013. “Business Process Management Capabilities in Local Governments: A Multi-Method Study.” *Government Information Quarterly* 30 (3): 217–225.

Object Management Group (OMG). 2008. *Business Process Maturity Model (BPMM) Version 1.0.* Technical report formal/2008-06-01. Milford, MA, USA: Object Management Group (OMG).

Pallant, J. 2016. *SPSS Survival Manual: A Step by Step Guide to Data Analysis Using IBM SPSS.* 6th ed. Maidenhead: Open University Press.

Patrício, L., R. P. Fisk, J. Falcão e Cunha, and L. Constantine. 2011. “Multilevel Service Design: From Customer Value Constellation to Service Experience Blueprinting.” *Journal of Service Research* 14 (2): 180–200.

Patton, M. Q. 2015. *Qualitative Research & Evaluation Methods: Integrating Theory and Practice.* Fourth edition. Thousand Oaks, California: SAGE Publications, Inc.

Peffers, K., T. Tuunanen, M. A. Rothenberger, and S. Chatterjee. 2007. “A Design Science Research Methodology for Information Systems Research.” *Journal of Management Information Systems* 24 (3): 45–77.

Pittaway, J. J., and A. R. Montazemi. 2020. “Know-How to Lead Digital Transformation: The Case of Local Governments.” *Government Information Quarterly* 37 (4): 101474.

Rawiro, D., F. L. Gaol, S. Supangkat, and B. Ranti. 2022. “Process Mining Applications in Government Sector: A Systematic Literature Review.” In *Proceedings of the International Conference on Industrial Engineering and Operations Management,* 3110–3118. Rome, Europe: IEOM Society International.

Rosemann, M., and T. de Bruin. 2005. “Towards a Business Process Management Maturity Model.” *ECIS 2005 Proceedings of the Thirteenth European Conference on Information Systems,* 1–12.

Sangiorgi, D., and A. Prendiville. 2014. “A Theoretical Framework for Studying Service Design Practices: First Steps to a Mature Field.” In *Proceedings of the 19th DMI International Design Management Research Conference,* 2042–2063. London: Design Management Institute.

Santana, A. F. L., C. F. Alves, H. R. M. Santos, and A. De Lima Cavalcanti Felix. 2011. “BPM Governance: An Exploratory Study in Public Organizations.” In *Enterprise, Business-Process and Information Systems Modeling,* edited by T. Halpin, S. Nurcan, J. Krogstie, P. Soffer, et al., 81:46–60. Berlin, Heidelberg: Springer Berlin Heidelberg.

Simonsen, J., and F. Kensing. 1997. “Using Ethnography in Contextual Design.” *Communications of the ACM* 40 (7): 82–88.

Smajli, E., G. Feldman, and S. Cox. 2024. “Exploring the Limitations of Business Process Maturity Models: A Systematic Literature Review.” *Information Systems Management* 42 (1): 2–21.

Szedmák, B., L. Varga, and R. Z. Szabó. 2025. “Digital Transformation of Public Services: The Case of the Document Management Application.” *International Journal of Public Administration,* 1–18.

Teixeira, J., L. Patrício, N. J. Nunes, L. Nóbrega, et al. 2012. “Customer Experience Modeling: From Customer Experience to Service Design.” *Journal of Service Management* 23 (3): 362–376.

Tejedo-Romero, F., J. F. F. E. Araujo, Á. Tejada, and Y. Ramírez. 2022. “E-Government Mechanisms to Enhance the Participation of Citizens and Society: Exploratory Analysis through the Dimension of Municipalities.” *Technology in Society* 70:101978.

Van Der Aalst, W. 2016. *Process Mining.* Berlin, Heidelberg: Springer Berlin Heidelberg.

Van Der Aalst, W., A. Adriansyah, A. K. A. De Medeiros, F. Arcieri, et al. 2012. “Process Mining Manifesto.” In *Business Process Management Workshops,* edited by F. Daniel, K. Barkaoui, and S. Dustdar, 99:169–194. Springer Berlin Heidelberg.

Van Der Aalst, W. M. P. 2011. *Process Mining: Discovery, Conformance and Enhancement of Business Processes.* Berlin, Heidelberg: Springer Berlin Heidelberg.

Van Der Kolk, B. 2022. “Performance Measurement in the Public Sector: Mapping 20 Years of Survey Research.” *Financial Accountability & Management* 38 (4): 703–729.

Vidgof, M., S. Bachhofner, and J. Mendling. 2023. *Large Language Models for Business Process Management: Opportunities and Challenges.* arXiv: 2304.04309 [cs].

Voorst tot Voorst, F. van. 2025. “Analyzing Municipal Process Data from the Social Domain: A Process Mining-Based Approach.” Master Thesis, Utrecht University.

Yin, R. K. 2018. *Case Study Research and Applications: Design and Methods.* 6th ed. SAGE Publications.

7 Sustainable Development Goals Analysis

This chapter analyses the alignment between this dissertation and the United Nations Sustainable Development Goals (SDGs). The analysis follows MESG dissertation guidance to identify relevant goals, targets, contributions, and performance indicators, and uses the United Nations framing of the 2030 Agenda and its 17 SDGs. The strongest alignment is with SDG 9 and SDG 16 because the work evaluates digital infrastructure quality, operational traceability, and institutional accountability in a municipal workflow platform.

SDG 9: Industry, Innovation and Infrastructure. Target 9.1 concerns quality, reliable, sustainable, and resilient infrastructure. CityFlow represents a digital infrastructure investment by the Porto Municipality. The dissertation evaluates technical implementation, data coverage, and operational embedding through BPMN process-definition versions, Elasticsearch index coverage, operational data volumes, and process-mining feasibility.

SDG 16: Peace, Justice and Strong Institutions. Target 16.6 concerns effective, accountable, and transparent institutions. By digitising operational workflows, CityFlow creates auditable records of who did what and when. The evaluation assesses transparency, accountability, and evidence-based municipal management through status/state-field availability, duplicate-work patterns, and stakeholder assessments of governance gaps.

SDG 8: Decent Work and Economic Growth. Target 8.2 concerns productivity through technological upgrading and innovation. The dissertation evaluates a platform intended to improve municipal operational work by replacing paper-based and informal coordination with structured digital processes. Relevant indicators include the number of operational verticals in production use, user-reported usage frequency, and recognition of practical value by operational staff.

SDG 11: Sustainable Cities and Communities. Target 11.3 concerns integrated and sustainable urban management. CityFlow supports municipal service domains such as animal welfare, green spaces, markets, and occurrence reporting. The evaluation assesses cross-vertical data aggregation, managerial monitoring access, and stakeholder-reported governance and coordination improvements.

The dissertation's primary SDG contributions are to Goals 9 and 16, through the evaluation of digital infrastructure quality and institutional transparency. Goals 8 and 11 are addressed through the productivity and integrated urban management dimensions of the evaluation.

Appendices

A Qualitative Codebook and Indicator Definitions

A.1 Code Hierarchy

The qualitative coding framework was organized around six analytical dimensions derived from the literature review (Chapter 2, Section 2.7) and one supplementary code family for recommendation types.

The coding cues in the final column are illustrative and are not interview quotations.

A.1.1 Dimension 1: Operational Fit and Functional Coverage

Code	Description	Inclusion criteria	Illustrative coding cue
FIT-01: Good alignment	Platform aligns well with operational task structure	Participant describes CityFlow as fitting their work without significant gaps or workarounds	"The system does exactly what we need for cage registration"
FIT-02: Partial alignment	Platform supports some tasks but not others; workarounds required	Participant identifies specific tasks where CityFlow fits and specific tasks where it does not	"The monthly plan works well, but for ad-hoc repairs we still use paper"
FIT-03: Missing functionality	Specific functional gap identified	Participant names a concrete function or feature that is absent	"We cannot attach photos to field reports"
FIT-04: Duplicate work	Same task performed in CityFlow and another system or format	Participant describes entering the same data in two places	"We fill the paper form in the field and then type it into CityFlow at the office"

A.1.2 Dimension 2: Adoption, Usability, and Digital Literacy

Code	Description	Inclusion criteria	Illustrative coding cue
ADP-01: Daily use	CityFlow used as part of daily operational routine	Participant reports using CityFlow every working day	"I open CityFlow first thing every morning"
ADP-02: Regular use	CityFlow used regularly but not daily	Participant reports frequent but not daily use	"I use it a few times a week, when I need to check plans"
ADP-03: Usability barrier	Interface or interaction design creates friction	Participant describes difficulty with navigation, input, or screen layout	"The buttons are too small on the tablet"
ADP-04: Digital literacy barrier	Limited digital skills constrain effective use	Participant or colleague described as struggling with basic digital interaction	"Some of the older workers find it hard to learn"
ADP-05: Training need	Participant identifies need for training	Explicit mention of training gap or request	"We only had one training session when it launched"
ADP-06: Resistance or inertia	Active or passive resistance to platform use	Participant describes reluctance, avoidance, or preference for old methods	"Some colleagues just refuse to use it"

A.1.3 Dimension 3: Value and Perceived Impact

Code	Description	Inclusion criteria	Illustrative coding cue
VAL-01: Strong value	Platform perceived as delivering significant operational or organizational value	Participant expresses clear positive judgement about platform's contribution	"It changed completely how we track animals"
VAL-02: Clear value	Platform perceived as useful but not transformative	Participant acknowledges benefits but without strong positive affect	"It helps with organization, yes"
VAL-03: Limited value	Platform perceived as having marginal or unclear value	Participant struggles to identify concrete benefits	"I don't see much difference from before"

A.1.4 Dimension 4: Data, Indicators, and Monitoring

Code	Description	Inclusion criteria	Illustrative coding cue
DAT-01: Data quality concern	Issues with accuracy, completeness, or timeliness of data in CityFlow	Participant identifies specific data that is wrong, missing, or outdated	"The counts don't match what we actually have"
DAT-02: Monitoring gap	Lack of accessible dashboards, reports, or monitoring views	Participant describes inability to see operational summaries or KPIs	"I can't see how many occurrences are open across my team"
DAT-03: Data value	Recognition that platform data enables new insights or management conversations	Participant describes using CityFlow data for decisions or reporting	"Now I can show my manager exactly what we did this month"

A.1.5 Dimension 5: Governance, Ownership, and Sustainability

Code	Description	Inclusion criteria	Illustrative coding cue
GOV-01: Prioritization gap	Unclear or inconsistent criteria for which features or verticals receive development resources	Participant expresses uncertainty about why some things are developed and others are not	"We asked for this feature months ago and nothing happened"
GOV-02: Communication gap	Insufficient communication about releases, changes, or platform roadmap	Participant describes being surprised by changes or unaware of new features	"Nobody told us about the update"
GOV-03: Feedback gap	Lack of structured mechanism for users to provide input	Participant describes not knowing how to report problems or suggest improvements	"I told the person who came for the interview but nothing changed"

A.1.6 Dimension 6: Technology, Integrations, and Future Evolution

Code	Description	Inclusion criteria	Illustrative coding cue
TEC-01: Integration need	Platform needs to connect with external system	Participant describes data that lives in another system and should be in CityFlow	"The asset database is separate, we have to check both"
TEC-02: Mobile/field limitation	Platform's mobile or field-use capability is insufficient	Participant describes issues specific to tablet or field use	"Sometimes there is no signal and I cannot save"
TEC-03: Future direction	Participant expresses view on platform's future evolution	Participant discusses where CityFlow should go next	"It would be great if we could just speak and it would record"

A.1.7 Supplementary: Recommendation Type

Code	Description
REC-01: Improve existing	Recommendation to enhance current functionality
REC-02: Develop new	Recommendation to build new capability
REC-03: Strengthen governance	Recommendation to improve prioritization, communication, or feedback structures

A.2 Indicator Definitions

The following indicators were assessed for each interview session.

Indicator	Scale	Definition	Scoring rule
Usage frequency	rare / periodic / regular / daily / not applicable / unknown	How often the participant or participant group uses CityFlow in the workflow discussed	Daily = each normal working day; regular = recurring but not every working day; periodic = used when a scheduled process is due; rare = exceptional use. Cadence is context-dependent and is not an expected-intensity score
Usage motivation	voluntary / obligation / mixed / not applicable / unknown	Whether CityFlow use is voluntary or mandated	Assessed from interview evidence

Indicator	Scale	Definition	Scoring rule
Value recognition	limited / clear / strong / not applicable / unknown	Degree to which participant perceives CityFlow as delivering value	Limited = no concrete or only marginal benefit; clear = at least one concrete recurring benefit; strong = material change to traceability, coordination, decision visibility, or completion of work, supported by explicit evidence
Operational fit	poor / partial / good / not applicable / unknown	Degree to which CityFlow aligns with operational task structures	Poor = core workflow cannot be completed meaningfully or occurs mainly outside CityFlow; partial = core tasks are supported but material gaps, parallel practices, or exception mismatches remain; good = main workflow is supported without a material workaround
Usability barrier	no / partial / yes / not applicable / unknown	Interface or interaction friction present	Assessed from descriptions of difficulty with the platform interface
Digital literacy barrier	no / partial / yes / not applicable / unknown	Digital skills gap constraining effective use	Assessed from descriptions of digital skill limitations
Training need	no / partial / yes / not applicable / unknown	Training gap identified	Assessed from explicit training mentions or skill-gap descriptions
Resistance or inertia	no / partial / yes / not applicable / unknown	Active or passive resistance to platform use	Assessed from descriptions of reluctance, avoidance, or preference for old methods
Duplicate work	no / partial / yes / not applicable / unknown	Same task performed in CityFlow and another system or format	Assessed from descriptions of parallel record-keeping
Missing functionality	no / partial / yes / not applicable / unknown	Specific functional gap identified	Assessed from concrete feature or function requests
Monitoring gap	no / partial / yes / not applicable / unknown	Lack of accessible monitoring, dashboards, or reports	Assessed from descriptions of monitoring or reporting limitations
Governance gap	no / partial / yes / not applicable / unknown	Prioritization, communication, or feedback mechanism gap	Assessed from descriptions of governance-related concerns
Improve existing priority	low / medium / high / not applicable / unknown	Urgency of improving current capabilities	Assessed from priority statements and improvement emphasis
New development priority	low / medium / high / not applicable / unknown	Urgency of developing new capabilities	Assessed from new-feature requests and development emphasis
Evidence confidence	low / medium / high / not applicable / unknown	Confidence in the indicator assessment based on source quality	Assessed from evidence completeness, human review status, and attribution grounding

A.3 Coding Procedure

Coding followed a hybrid deductive-inductive approach. The six analytical dimensions were derived deductively from the literature (Chapter 2). Individual codes within each dimension were developed inductively from the interview evidence, following the thematic analysis procedures described by Braun and Clarke (2006). The codebook was iteratively refined as coding proceeded. Each coded passage was classified by session, dimension, code, polarity, claim summary, attribution unit, and confidence. The final synthesis contains 96 validated participant-evidence segments across 14 sources and 14 cases, organised into 7 categories and 22 codes.

B Interview Session Register

Session ID	Date	Context	Participants	Profile summary
S01	2026-03-20	CROA	2	veterinarian doctor and engineer/chief of division
S02	2026-04-17	CROA	2	operator/tratadora and veterinary nurse
S03	2026-04-17	CROA	1	encarregada
S05	2026-05-14	CROA	2	operator/tratador and encarregado operacional
S06	2026-05-14	CROA	2	two tratadores
S07	2026-05-14	CROA	2	supervisory and administrative perspectives
S08	2026-04-23	Espaços Verdes	2	operational and supervisory field perspectives
S09	2026-04-29	Espaços Verdes	2	operational and supervisory field perspectives
S10	2026-05-04	Espaços Verdes	2	operational and supervisory field perspectives
S11	2026-05-15	Mercados	5	three street/tablet operational users and two technical-superior users
S12	2026-05-13	Internal stakeholder	1	former product owner and service designer
S13	2026-05-21	Internal stakeholder	1	APD front-end tester and quality-assurance stakeholder
S14	2026-05-21	Internal stakeholder	1	supplier project/development stakeholder with programming experience
S15	2026-05-22	Internal stakeholder	1	municipal governance and KPI perspective

Source: Interview session register. N = 14 sessions.

C Participant Register

Alias	Session	Context	Stakeholder category	Role group
CROA-P01	S01	CROA	End user	veterinarian doctor
CROA-P02	S01	CROA	End user	chief of division / engineer
CROA-P03	S02	CROA	End user	operator / tratadora
CROA-P04	S02	CROA	End user	veterinary nurse
CROA-P05	S03	CROA	End user	encarregada
CROA-P08	S05	CROA	End user	operator / tratador
CROA-P09	S05	CROA	End user	operational supervisor
CROA-P10	S06	CROA	End user	tratador
CROA-P11	S06	CROA	End user	tratador
CROA-P12	S07	CROA	End user	supervisor
CROA-P13	S07	CROA	End user	supervisor
EV-P01	S08	Espaços Verdes	End user	chief supervisor
EV-P02	S08	Espaços Verdes	End user	operational gardener
EV-P03	S09	Espaços Verdes	End user	operational supervisor
EV-P04	S09	Espaços Verdes	End user	operational gardener
EV-P05	S10	Espaços Verdes	End user	operational staff
EV-P06	S10	Espaços Verdes	End user	supervisor
MER-P01	S11	Mercados	End user	tablet operational user
MER-P02	S11	Mercados	End user	tablet operational user
MER-P03	S11	Mercados	End user	tablet operational user
MER-P04	S11	Mercados	End user	technical PC user
MER-P05	S11	Mercados	End user	technical PC user
STK-P01	S12	Internal stakeholder	APD	service design / former product owner
STK-P02	S13	Internal stakeholder	APD	front-end tester / quality assurance
STK-P03	S14	Internal stakeholder	Supplier	hybrid project/development role
STK-P04	S15	Internal stakeholder	DMSI	governance and analytics

Source: Participant register. N = 26 registered participants.

D Session-Level Indicator Grid

Indicator definitions and scoring rules are provided in Appendix A. The full grid with per-session indicator values is reproduced below. N = 14 sessions, 26 registered participants.

Table D.1: Full session-level indicator grid.

Sess	Usage	Motv	Value	Fit	Usability	DigLit	Train	Resist	Dup Work	Missing	Improve	New Dev	Conf
S01	daily	mixed	strong	good	yes	partial	partial	yes	yes	yes	high	medium	high
S02	regular	mixed	clear	partial	yes	unknown	unknown	unknown	yes	yes	high	medium	high
S03	daily	mixed	strong	good	partial	partial	partial	partial	no	partial	medium	low	medium
S05	daily	mixed	clear	partial	yes	unknown	unknown	unknown	yes	yes	high	medium	medium
S06	regular	mixed	clear	partial	yes	partial	unknown	partial	yes	yes	high	high	medium
S07	daily	oblig.	strong	good	partial	yes	yes	partial	yes	yes	high	high	medium
S08	daily	oblig.	limited	partial	partial	unknown	unknown	unknown	yes	yes	high	high	high
S09	regular	oblig.	limited	partial	yes	unknown	unknown	unknown	yes	unknown	high	medium	medium
S10	daily	oblig.	clear	partial	yes	partial	unknown	partial	yes	yes	high	medium	medium
S11	daily	mixed	strong	partial	yes	unknown	unknown	unknown	yes	yes	high	high	medium
S12	n/a	n/a	strong	partial	n/a	yes	yes	partial	n/a	yes	high	high	high
S13	n/a	n/a	clear	partial	n/a	n/a	yes	n/a	n/a	unknown	high	medium	medium
S14	n/a	n/a	strong	good	n/a	n/a	n/a	n/a	n/a	unknown	high	medium	high
S15	n/a	n/a	clear	partial	n/a	n/a	yes	yes	n/a	yes	high	medium	medium

Column abbreviations: Usage = usage frequency (daily/regular/not_applicable); Motv = usage motivation (mixed/oblig. = obligation/not_applicable); Value = value recognition (strong/clear/limited); Fit = operational fit (good/partial); Usability = usability barrier (yes/partial/not_applicable); DigLit = digital literacy barrier; Train = training need; Resist = resistance or inertia; Dup Work = duplicate work; Missing = missing functionality; Improve = improve existing priority; New Dev = new development priority; Conf = evidence confidence. Unknown indicates insufficient evidence for classification; `n/a` indicates that the indicator does not apply to that session.

E Triangulation Matrix

This appendix summarises the evidence combined in each triangulation entry.

Table E.1: Triangulation entry index.

ID	Scope	Dimension	Confidence	Research questions
T01	CROA	Operational fit and value	High	RQ1; RQ2
T02	CROA	Data, indicators, and monitoring	High	RQ1; RQ2; RQ3
T03	Espaços Verdes	Operational fit and adoption	High	RQ1; RQ2
T04	Espaços Verdes	Data quality and monitoring	High	RQ1; RQ2; RQ3
T05	Mercados	Operational capability	Medium	RQ1; RQ2; RQ3
T06	All production service areas	Governance and sustainability	High	RQ2; RQ3
T07	All production service areas	Technology and future evolution	High	RQ2; RQ3
T08	Cross-case	Technical implementation versus adoption	High	RQ1; RQ2
T09	Ocorrências/Reporta	Scope limitation	High	RQ1
T10	Platform level	Portfolio governance	High	RQ2; RQ3
T11	Platform level	Evaluation and prioritisation	High	RQ2; RQ3

E.1 Service-Area Entries

E.1.1 T01: CROA Operational Fit and Value

- **Quantitative evidence:** 22 BPMN processes, 115 tasks, 60,646 probed historic instances, and 60,164 Elasticsearch documents.
- **Qualitative evidence:** users describe rapid historical retrieval, management response, traceability, and legal evidence; they also report paper/signature duplication and clinical workflow gaps.
- **Interpretation:** CROA is the strongest operational-value case, but a broad technical implementation and realised value coexist with incomplete coverage and parallel work.

E.1.2 T02: CROA Data, Indicators, and Monitoring

- **Quantitative evidence:** 55.0% Elasticsearch index coverage and stable operational records across multiple CROA modules.
- **Qualitative evidence:** management reports stable macro indicators, but filters and exports make ad-hoc monitoring difficult.
- **Interpretation:** monitoring value exists, but flexible self-service analysis and data access remain improvement priorities. External reporting layers were not evaluated directly.

E.1.3 T03: Espaços Verdes Operational Fit and Adoption

- **Quantitative evidence:** 23 BPMN processes, 128 tasks, and 1,721,560 probed historic instances.
- **Qualitative evidence:** all three sessions show routine or regular use but partial fit: paper or WhatsApp handoffs persist, task catalogues are incomplete, connectivity constrains field entry, and changing priorities displace planned work.
- **Interpretation:** high technical activity and routine use do not by themselves demonstrate strong operational fit; CityFlow remains primarily a constrained registration and execution layer in the assessed contexts.

E.1.4 T04: Espaços Verdes Data Quality and Monitoring

- **Quantitative evidence:** 12,986 Elasticsearch documents and 52.2% index coverage; the dominant planning event log is largely single-event and cancelled.
- **Qualitative evidence:** across the three sessions, missing or generic task categories, free-text limitations, paper or WhatsApp handoffs, and unpredictable work reduce record completeness and monitoring value.
- **Interpretation:** data volume and workflow activity require qualitative interpretation; taxonomy, handoffs, connectivity, and event-log structure constrain reliable monitoring.

E.1.5 T05: Mercados Operational Capability

- **Quantitative evidence:** 15 BPMN processes, 228,309 probed historic instances, 228,989 Elasticsearch documents, and 86.7% index coverage.
- **Qualitative evidence:** the group describes daily execution, real-time operational visibility, substantial improvement, and decision-support potential; planning and report analysis remain partly in Excel, while reliability, extraction, alerting, and integration gaps persist.
- **Interpretation:** Mercados is operationally used and valued but remains deployed with constraints rather than fully integrated or managed. The operational interpretation rests on one group session.

E.2 Cross-Case and Platform Entries

E.2.1 T06: Governance and Sustainability

- **Quantitative evidence:** 500 deployments are recorded, concentrated in 2022-2024, with 14 in 2025 and 1 in 2026.
- **Qualitative evidence:** APD, supplier, and partial DMSI evidence identify ownership, financing, staffing, training, QA capacity, and impact measurement as sustainability conditions. The technical-role evidence includes one APD front-end tester and one supplier participant with a hybrid project/development role and direct programming experience.
- **Interpretation:** platform sustainability depends on governance and delivery capacity, not only deployed process volume. The technical-role and DMSI accounts are individual perspectives rather than team consensus.

E.2.2 T07: Technology and Future Evolution

- **Quantitative evidence:** the local event log has 816,000 events; 89.73% of cases have one event, 83.79% of events are cancelled, and only 4 events contain an assignee.
- **Qualitative evidence:** stakeholders prioritise integration, robust existing service areas, monitoring, reliable impact measurement, and data-access improvements.
- **Interpretation:** the strongest near-term technology path is monitoring and interoperability built on better data quality; the bounded AI scenario remains exploratory.

E.2.3 T08: Technical Implementation Versus Adoption

- **Quantitative evidence:** production service areas differ substantially in process versions, coverage, document counts, and historic-instance probes.
- **Qualitative evidence:** CROA and Mercados report clear realised value but different constraints; all three Espaços Verdes sessions show routine use with partial fit, local workarounds, and weaker recognised value.
- **Interpretation:** capability assessment must distinguish technical deployment, execution activity, operational fit, adoption, value, data use, and governance.

E.2.4 T09: Ocorrências/Reporta Scope Limitation

- **Quantitative evidence:** 7 BPMN processes, 49,260 probed historic instances, 5,835 Elasticsearch documents, and 14.3% index coverage.
- **Qualitative evidence:** no operational interviews were collected for this service area.
- **Interpretation:** Ocorrências/Reporta can be included in the quantitative technical scan, but not in the operational cross-case comparison.

E.2.5 T10: Platform Portfolio Governance

- **Quantitative evidence:** the digital-solutions portfolio records owners for 81% of products, suppliers for 19%, project completion dates for 42%, and contract references for 0%.
- **Qualitative evidence:** internal stakeholders describe unclear product ownership, fragmented responsibilities, staffing dependencies, and weak decision continuity.
- **Interpretation:** interview governance concerns are consistent with wider portfolio metadata and decision-traceability gaps. Portfolio evidence provides organisational context rather than a CityFlow implementation measure.

E.2.6 T11: Platform Evaluation and Prioritisation

- **Quantitative evidence:** the DMSI evaluation matrix scores citizen impact, strategy alignment, interoperability, effort, migration, infrastructure, cost, return, and non-implementation risk.
- **Qualitative evidence:** interview evidence calls for stronger indicators, documented feedback, prioritisation, support, training, and improvement of existing service areas.
- **Interpretation:** a CityFlow roadmap should connect operational evidence and platform indicators to an explicit recurring evaluation and prioritisation process. Routine use of the referenced prioritisation matrix was not observed.

Source: maintained triangulation matrix, 11 entries.

F PSPP Descriptive Statistics

Generated from PSPP output using GNU PSPP 2.0.1.

F.1 Session Composition (N=14)

Vertical	Count	%
CROA	6	42.9
Espacos Verdes	3	21.4
Mercados	1	7.1
Internal Stakeholder	4	28.6

F.2 Usage Frequency (N=10 operational)

Level	Count	Valid %
Regular	3	30.0
Daily	7	70.0

F.3 Value Recognition (N=14)

Level	Count	%
Limited	2	14.3
Clear	6	42.9
Strong	6	42.9

F.4 Operational Fit (N=14)

Level	Count	%
Partial	10	71.4
Good	4	28.6

G Claim-Evidence Traceability

This appendix links the principal claims to the evidence summaries used in the analysis. Evidence identifiers refer to the qualitative register (E), triangulation matrix (T), and artifact-design register (AD).

ID	Principal claim	Evidence	Scope
C01	The corpus contains 14 interview sessions and 26 registered participants.	Appendices B-C	Group interviews are analysed at session level unless speaker or role attribution is reliable.
C02	All ten assessed operational sessions describe daily or regular CityFlow use.	Appendix D; S01-S03, S05-S11	Assessed service areas only.
C03	Nine of ten operational sessions report duplicate work or parallel handoffs.	E004, E009, E014, E018, E041, E046, E073, E081, E086; Appendix D	Session-level finding.
C04	CROA provides the strongest evidence of realised operational and management value while retaining material constraints.	E001-E012, E031-E057; T01-T02	Applies to the sampled CROA workflows.
C05	The three Espaços Verdes sessions show routine use with partial fit, workarounds, and incomplete task coverage.	E013-E020, E083-E086, E088-E090; T03-T04	Session-level interpretation across three locations.
C06	Mercados is used daily and provides operational value, with remaining planning, extraction, reliability, and integration constraints.	E070-E082; T05	One group session; stored volume includes bulk-loaded planning records.
C07	Internal stakeholders identify governance and ownership constraints.	E021, E027, E028, E058, E064, E069, E100; T06, T10-T11	Four individual perspectives, including one front-end QA and one supplier project/development participant.
C08	Monitoring access and measurement remain incomplete despite monitorable platform data.	E029-E030, E059-E060, E067, E095-E096; T02, T07, T11	Dashboard usability was not assessed directly.
C09	The available event log is unsuitable for conclusive process discovery or conformance analysis.	T07; Appendix H	Process mining remains conditional on a process-specific quality gate.
C10	Existing workflows, data semantics, and governance should be strengthened before broad platform expansion.	E025, E048, E057, E063, E069; Table 5.9	A prioritisation conclusion, not a record of an approved organisational decision.
C11	CROA is operationally embedded; Espaços Verdes and Mercados are deployed but constrained.	Appendix D; T01-T11; Table 5.7	Case-specific diagnostic profile; Mercados has medium confidence.
C12	A human-approved voice-to-record assistant is a defensible exploratory scenario.	E024, E046, E048; T07; Figure 5.1	No model trial or production implementation was conducted.
C13	Technical captures support claims about platform structure, data fields, and analysis feasibility.	29 April inventory; 15 June targeted capture; Appendix H	They do not independently establish operational outcomes.
C14	Process versions describe deployment iteration and active flags describe suspension state at capture.	BPMN inventory; Table 4.1; Appendix H	Neither is a maturity or usage measure.
C15	The reusable artifact is the six-dimensional evaluation protocol; the capability profile and roadmap are CityFlow demonstration outputs.	AD01-AD10; Tables 3.1, 3.3, 5.8-5.10	The AI scenario remains outside the reusable core.

ID	Principal claim	Evidence	Scope
C16	Formative reviews changed the artifact, and its CityFlow application demonstrates practical use.	Table 3.4; Tables 5.8-5.10	Summative stakeholder acceptance and multi-site validation remain future work.

H Quantitative Analysis Report

The full Camunda, BPMN, and Elasticsearch inventory was captured on 29 April 2026. Selected definitions and interface evidence were refreshed on 15 June 2026; results below retain their respective capture dates.

H.1 Elasticsearch Volume and Coverage

H.1.1 Elasticsearch Volume and Capture Coverage

Vertical	Indices	With data	Empty	Coverage %	Captured documents	Analytical role
MERCADOS	15	13	2	86.70	228989	operational comparison
CROA	20	11	9	55.00	60164	deep-dive case study
ESPAÇOS VERDES / AMBIENTE	23	12	11	52.20	12986	operational comparison
REPORTA	7	1	6	14.30	5835	scoped technical context

H.1.2 BPMN/Camunda Inventory

Task and gateway counts refer to **BPMN XML elements**. Version fields identify the latest deployment iteration visible in the 29 April inventory.

Vertical	Processes	Avg. version	Max. version	Tasks	Service tasks	Gateways	Hist. instances	Role
ESPAÇOS VERDES / AMBIENTE	23	12.26	42	128	12	46	1721560	operational comparison
MERCADOS	15	2.20	9	52	12	37	228309	operational comparison
CROA	22	18.45	111	115	25	52	60646	deep-dive case study
REPORTA	7	13.43	45	24	6	15	49260	scoped technical context

H.1.3 Espaços Verdes Process-Definition Inventory

The 23 reported Espaços Verdes definitions are listed below so that the process count in Chapter 4 is inspectable. Historic-instance values are non-uniform probes from the 29 April inventory; zero means no instance was returned by the probe, not proof that the process was never used. Selected main definitions were checked again in the 15 June targeted refresh.
The inventory identifiers and source names are:

- **P01:** `ambientegeralplaneamento`; Ambiente Geral Planeamento; version 8.
- **P02:** `geralAlertasAmbiente`; AMB:ALERTAS; version 1.
- **P03:** `geralAlertasCatAmbiente`; AMB:ALERTAS_CAT; version 1.
- **P04:** `geralAlertasHistAmbiente`; AMB:ALERTAS_HIST; version 1.
- **P05:** `geralCasasPontoAmbiente`; AMB:CASA_PONTO; version 7.
- **P06:** `geralElementosAmbiente`; AMB:ELEM; version 12.
- **P07:** `geralEntradasAmbiente`; Geral Entradas Ambiente; version 11.

- **P08:** `geralEquipamentosAmbiente`; AMB:EQUIP; version 11.
- **P09:** `geralExecucaoSemAmbiente`; AMB:EXEC; version 10.
- **P10:** `geralIntCorrAmbiente`; AMB:INT_CORRECTIVA; version 21.
- **P11:** `geralIntPrevAmbiente`; Geral Intervenção Preventiva Ambiente; version 13.
- **P12:** `geralJardinsAmbiente`; AMB:JARDIM; version 23.
- **P13:** `geralListaVerificacaoAmb`; AMB:LISTA_VERIFICACAO; version 1.
- **P14:** `geralPlanAnualAmbiente`; AMB:PLAN_ANUAL; version 31.
- **P15:** `geralPlanMensalAmbiente`; AMB:PLAN_MENSAL; version 10.
- **P16:** `geralPlanSemAmbiente`; AMB:PLAN_SEM; version 20.
- **P17:** `geralSaidaAmbiente`; Geral Saida Ambiente; version 11.
- **P18:** `geralTipoJardimAmbiente`; AMB:TIPO_JARDIM; version 6.
- **P19:** `geralTipoOperacaoAmbiente`; AMB:TIPO_OPERACAO; version 6.
- **P20:** `geralZonasAmbiente`; AMB:ZONA; version 15.
- **P21:** `mapasActivosAmbiente`; Mapa Ativos Ambiente; version 8.
- **P22:** `modulosAmbientePlaneamento`; Modulos Planeamento do Ambiente; version 13.
- **P23:** `modulosDoAmbiente`; Modulos Do Ambiente; version 42.

Table H.3: Espaços Verdes process-definition metrics from the 29 April 2026 inventory. IDs map to the complete key and source-name list above.

ID	User tasks	Service tasks	Gateways	Historic-instance probe
P01	3	0	0	0
P02	2	2	2	0
P03	2	0	2	0
P04	2	0	2	0
P05	2	0	2	0
P06	2	0	2	0
P07	2	0	2	0
P08	2	0	2	0
P09	4	1	4	3,251
P10	2	0	2	28
P11	2	0	2	0
P12	2	3	2	1,181
P13	2	0	2	0
P14	5	4	6	84
P15	2	0	1	1,709,475
P16	3	0	2	6,975
P17	2	0	2	0
P18	2	0	2	0
P19	2	0	2	85
P20	2	2	2	0
P21	5	0	1	0
P22	5	0	1	0
P23	54	0	1	481

H.1.4 Additional Technical Annex Service Areas

The remaining service areas continue to be recorded in the complete CSVs but are not used to support the main operational comparison.

H.1.5 Deployments by Year

Year	Deployments
2022	85
2023	201
2024	199
2025	14
2026	1

H.2 BPMN/Camunda Technical Footprint and Limits

- **Exported BPMN files:** 134 files. Basis for process-definition inventory and XML element counts, not a standalone maturity assessment.
- **Camunda deployments:** 500. Basis for temporal evolution and iteration.
- **Camunda historical instances:** 2,505,623. Captured API total across the retained history.
- **Camunda historical activities:** 12,433,632. Captured API total across the retained history.
- **Open user tasks:** 311,832. Captured task-state total, not a validated active-user backlog.
- **Local event log:** 816,000 events. Useful for feasibility assessment, but mixed.
- **Local Elasticsearch:** 322,879 documents across all inventoried index groups; the four production-scope rows in Table H.1 account for 307,974. Counts are grouped by index naming conventions, not repeated per service area.

H.3 CROA Module Coverage

- **CROA deep-dive** as the primary case study, being the vertical with the most complete technical, operational, and qualitative evidence.
- **Comparison with Espaços Verdes/Ambiente and Mercados**, both as operational production cases.
- **Ocorrências/Reporta as a scoped case**, due to its cross-cutting, migration/integration nature and absence of dedicated operational interview coverage.
- **Exploratory process mining**, only after filtering the event log by `processDefinitionKey` and separating cancelled events.
- **Interviews as qualitative analysis**, used to explain adoption, friction, and perceived value.

Source: full quantitative inventory captured on 29 April 2026; targeted technical refresh completed on 15 June 2026.

I Qualitative Coding Frequencies

The final synthesis contains 96 validated participant-evidence segments across 14 sources and 14 cases, organised into 7 categories and 22 codes. The coding identifies recurring patterns in perceived value, operational fit, workflow friction, monitoring needs, and future improvement opportunities.
The codebook (Appendix A), the code frequencies below, and the claim-evidence traceability matrix (Appendix G) reproduce the relevant coding summaries.

I.1 Code Frequencies by Category

Code	Category	Count
Duplicate work	Operational fit and functional coverage	9
Good or partial alignment	Operational fit and functional coverage	6
Missing functionality	Operational fit and functional coverage	6
Daily or regular use	Adoption, usability, and digital literacy	8
Digital literacy barrier	Adoption, usability, and digital literacy	2
Resistance or inertia	Adoption, usability, and digital literacy	2
Training need	Adoption, usability, and digital literacy	1
Usability barrier	Adoption, usability, and digital literacy	7
Strong value recognition	Value and perceived impact	7
Clear value recognition	Value and perceived impact	6
Data quality concern	Data, indicators, and monitoring	5
Monitoring available	Data, indicators, and monitoring	2
Monitoring gap	Data, indicators, and monitoring	9
Governance clarity	Governance, ownership, and sustainability	1
Ownership or governance gap	Governance, ownership, and sustainability	5
Sustainability risk	Governance, ownership, and sustainability	4
Integration need	Technology, integrations, and future evolution	3
Reliability concern	Technology, integrations, and future evolution	5
Technical quality	Technology, integrations, and future evolution	3
Improve existing	Recommendation type	2
Develop new	Recommendation type	1
Strengthen governance	Recommendation type	2

I.2 Category Totals

Category	Total codes
Operational fit and functional coverage	21
Adoption, usability, and digital literacy	20
Data, indicators, and monitoring	16
Value and perceived impact	13
Governance, ownership, and sustainability	10
Technology, integrations, and future evolution	11
Recommendation type	5

Source: validated qualitative coding outputs prepared in QualCoder.

J Generative AI Use Note

In accordance with the FEUP and University of Porto guidance on responsible use of generative artificial intelligence, this dissertation discloses limited instrumental use of GenAI-supported tools during the research and writing process.

Speech-to-text tools assisted preparation of working interview transcripts. Session recordings remained the source for direct quotations, and group evidence was reported at session level where individual attribution was not reliable.

GenAI tools were not used to generate original research data, select evidence autonomously, interpret findings, make methodological decisions, or formulate the dissertation conclusions. All factual claims, quotations, statistics, citations, and final prose were reviewed by the author against primary sources. Authorship and academic responsibility rest exclusively with the author.