NOVA IMS
INFORMATION MANAGEMENT SCHOOL

Master's Degree in
Data Science and Advanced Analytics

Specialization in
Data Science

Agentic AI in Central Public Administration

Exploring Opportunities and Challenges for Workflow Integration and Task Automation

José António Rodrigues Cavaco

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

Agentic AI in Central Public Administration

Exploring Opportunities and Challenges for Workflow Integration and Task Automation

by

José António Rodrigues Cavaco

Master Thesis presented as partial requirement for obtaining the Master's degree in Data Science and Advanced Analytics, with a specialization in Data Science

Supervised by

Vítor Santos, NOVA Information Management School

April, 2026

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism, any form of undue use of information or falsification of results along the process leading to its elaboration. I further declare that I have fully acknowledged the Rules of Conduct and Code of Honor from the NOVA Information Management School.

[Lisbon, 2026]

José Cavaco

ACKNOWLEDGEMENTS

I would like to express gratitude to all those who helped me with the completion of this thesis.

First I would like to thank my supervisor, Professor Vítor Santos for the guidance and support provided throughout this work.

I am also grateful to the professionals who participated in the evaluation phase of this study. Thank you for your availability and practical insights.

Finally I would like to thank my mom, André, Ana Rita and Sofia for all the support, patience, and encouragement throughout this journey

ABSTRACT

This thesis examines the application of Agentic AI in Central Public Administration and proposes practical guidelines for its adoption. The study is motivated by persistent inefficiencies in public administration, including fragmented workflows, coordination challenges, and repetitive administrative tasks. While AI adoption has increased, most implementations remain limited to assistive or predictive systems, with limited exploration of autonomous, agent-based approaches. To address this gap, a Design Science Research methodology was applied. A literature review was conducted to assess the current state of Agentic AI in public administration, identifying key challenges, technologies, and use cases. Based on these findings, a set of guidelines was developed to identify suitable processes for Agentic AI application, along with a structured implementation framework covering use case selection, readiness assessment, implementation, deployment, and evaluation. The framework was illustrated through a simplified use case and evaluated through semi-structured interviews with three experts. The results indicate that the proposed approach is useful and applicable, particularly due to its structured and gradual implementation. Some improvements were identified, namely in process selection and monitoring. Overall, this research contributes a practical framework to support the integration of Agentic AI into public administration workflows.

KEYWORDS

Agentic AI; Digital Transformation; Process Redesign; Public Sector Innovation.

Sustainable Development Goals (SDG):

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

AI	Artificial Intelligence
Agentic AI	Agentic Artificial Intelligence
CPA	Central Public Administration
DSR	Design Science Research
EU	European Union
Gen AI	Generative Artificial Intelligence
IT	Information Technology
LLM	Large Language Model
OECD	Organization for Economic Cooperation and Development
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
ROI	Return on Investment
SLR	Systematic Literature Review
SLRQ	Systematic Literature Review Question

1 INTRODUCTION

1.1 CONTEXT

In recent years, Artificial Intelligence (AI) has become a central theme across multiple sectors, driven by its potential to transform processes, enhance decision-making, and increase efficiency, with its growth clearly visible within corporate and business environments (Singla et al., 2025).

As AI has become more pervasive, extending beyond the field of information systems into everyday organizational and governmental contexts (Costa et al., 2024), one of its most visible advancements has been the development of Generative AI, a class of models capable of producing content such as text, images, and code (Schneider, 2025). These systems are particularly valuable because they can transform large volumes of raw data into actionable insights, supporting employees in decision-making and task execution more effectively. Today, Generative AI stands out as the most recognizable form of artificial intelligence, having captured both public attention and corporate investment due to its diverse use cases.

However, the frontier of AI adoption is shifting toward a more ambitious paradigm: Agentic AI. Unlike traditional or generative AI, which primarily provide outputs in response to user prompts, agentic systems are designed to act autonomously on behalf of users. They go beyond content generation to include the ability to plan, reason, and execute sequences of actions in pursuit of defined objectives (Schneider, 2025). In practice, this means that such systems could take over repetitive, highly standardized tasks or even coordinate across different organizational units with minimal human intervention. By automating such tasks, organizations stand to gain considerable efficiency, saving both time and resources while improving the management of complex workflows (Corvalán, 2025).

The public sector represents a particularly interesting context for exploring the applications of Agentic AI. Governments, like other large organizations, face ongoing pressure to modernize, deliver services more efficiently, and respond to citizens’ needs in a timely manner. Yet, the introduction of new technologies in the public sector tends to be slow, (Alvarenga et al., 2020), often hindered by rigid procedures, legal constraints, and bureaucratic structures. In Portugal, recent initiatives have shown some progress in digital transformation and AI adoption, but overall integration remains at an early stage (Martinho, 2024). The challenge, therefore, lies not only in the technological capabilities themselves but also in the institutional and cultural readiness of public administrations to absorb and scale such innovations, making the study of Agentic AI in this field both timely and relevant.

While Agentic AI is already gaining traction in the private sector, with emerging tools demonstrating its ability to streamline business operations and customer service, there remains limited research and practical experimentation within public administration

(Blanchet, 2025). Yet, the potential benefits are considerable. Public workflows often involve multiple agencies, with overlapping competencies and dependencies that can lead to inefficiencies, delays, and redundancies. Integrating agentic systems into these workflows could help improve coordination, reduce repetitive administrative tasks, and ultimately enhance service delivery to citizens.

Given this gap, it becomes important to investigate the role of Agentic AI in the redesign and optimization of public sector processes.

Therefore, the following Research Question can be defined:

RQ: How can Agentic AI be applied in Central Public Administration through a structured implementation framework and guidelines?

1.2 OBJECTIVES

Thus, the goal of the research is to produce a set of guidelines for the use of Agentic AI technologies in Central Public Administration.

In order to achieve this goal, the following intermediate objectives were defined:

1) Comprehensive literature review on Agentic AI

2) Analyze the current state of AI use in public administration

3) Identify cases of poor coordination and inefficiency and mapping current limitations.

4) Create set of guidelines for the use of Agentic AI technologies in Central Public Administration.

5) Simulate or model possible use cases;

6) Evaluate the proposed model.

1.3 EXPECTED RESULTS AND CONTRIBUTIONS

This project is expected to generate valuable contributions for the modernization of public administration by examining how Agentic AI can be integrated into existing workflows to improve efficiency, reduce redundancy, and enhance the quality of services delivered to citizens (Salah et al., 2025). By analyzing concrete use cases, it will highlight opportunities where agentic systems can take over repetitive and bureaucratic procedures, facilitate coordination between agencies, and help mitigate the delays and inefficiencies that often arise in highly regulated environments. At the same time, the research will extend current knowledge on AI in government, addressing an area where studies remain scarce. While many administrations have experimented with predictive analytics, chatbots, or digital assistants, the step toward systems that can plan, and act autonomously represents a more ambitious transformation that is still largely unexplored.

The project will also produce guidelines and recommendations for policymakers and managers, emphasizing not only where Agentic AI can create value but also the safeguards needed to ensure its responsible adoption. Questions of accountability, transparency, and ethical compliance are especially important in the public sector, where decisions carry significant social impact and must respect democratic principles and legal standards. By integrating these dimensions into the analysis, the research will contribute to a more comprehensive understanding of both the opportunities and the risks of agentic systems.

Ultimately, this work aims to position public administration as a meaningful testing ground for Agentic AI. The public sector provides a challenging but valuable environment for this technology, given its reliance on complex, rule-based systems and high-stakes decision-making (Kurhayadi, 2025). Lessons drawn from this context can contribute not only to improving governance and public service efficiency, but also to advancing the broader scientific and practical understanding of agentic systems.

2 CENTRAL PUBLIC ADMINISTRATION (RELEVANCE CYCLE)

2.1 Overview

The Central Public Administration is the national government's executive branch, responsible for implementing laws and public policies. It consists of the ministries, secretaries of state, directorates-general, and central services in charge of executing government decisions and ensuring the functioning of national public programs on a continuous basis. It includes the bodies that design regulations, coordinate sectoral strategies, and manage national services. Unlike regional or local administrations, which focus on territorial specific responsibilities, the Central Public Administration operates at the national level, overseeing policies and services whose coordination and impact extend across the entire country.

Its governance structure follows a hierarchical model, with decision-making distributed across this administrative structure. In order to ensure coherence, effective coordination and communication between these units is essential both horizontally, across sectors, and vertically, throughout the administrative hierarchy. Traditionally, the central administration has been guided by principles of compliance, stability, and accountability, reflecting its mandate to maintain continuity and uphold the rule of law.

Over the past decades, the Portuguese Central Public Administration has undergone several waves of modernization, shaped both by international administrative reform trends and by national initiatives aimed at improving efficiency and service quality such as the creation of transversal coordination structures like the Centre for Planning, Policy and Foresight (PlanAPP), designed to strengthen strategic planning, policy coherence, and evidence-based decision-making across government (*Decreto-Lei n.º 21/2021 | DR*, 2021) and, at an European level, the Coordinated Plan on Artificial Intelligence (European Comission, 2018). These reforms sought to increase efficiency, transparency, and service quality while reducing bureaucratic burden. The transition from paper-based procedures to digital workflows, that is happening to this day, marked a significant shift, progressively integrating data-driven decision-making and expanding the use of information and communication technologies. Interoperability platforms, shared services, and standardized digital infrastructures have become central components of administrative modernization (OECD, 2025).

Given its scale, complexity, and strategic importance, the Central Public Administration plays a pivotal role in national digital transformation efforts (Xiao et al., 2023). It concentrates the institutional capacity required for large reforms, and its inter-ministerial reach enables the coordination of digital policies across domains such as health, education, finance, and social protection. Its ability to modernize internal processes, standardize digital infrastructures, and integrate advanced technologies has implications for the entire public sector.

In the following sections, I will detail the structure of central government, explore key policy areas, examine the evolution of e-government, and discuss the main challenges and opportunities that lie ahead.

2.2 CENTRAL GOVERNMENT

The Central Government serves as the core coordinating layer of the state, where national priorities are defined and translated into policies that guide the work of ministries and public services. Compared to Public Administration in general, its role is therefore less operational and more strategic: ensuring coherence across sectors, allocating resources, and ensuring that public action remains aligned with national objectives. In this coordinating function, bodies with cross-government mandates, such as ARTE (Agência para a Reforma e Transformação do Estado), play an important role in operationalizing national strategies and supporting the adoption of shared digital infrastructures across the administration.

From a modernization perspective, the Central Government acts as the catalyst for reform. National strategies on digital transformation, data governance, public service quality, and administrative simplification typically originate at this level, before being implemented by ministries and agencies (OECD, 2025).

This positioning makes the Central Government essential for understanding the evolution of Portugal's public sector. It provides the link between high-level policy goals and the concrete administrative domains explored in the next section. At the same time, it shapes the conditions under which digital public services and emerging technologies, such as AI, can be adopted at scale.

2.3 AREAS

The Central Government is structured around several sectors that correspond to the main functions of the state, such as health, education, social security and justice. These areas reflect the complexity of public intervention and illustrate how different parts of the administration respond to social needs, deliver services, and manage public resources.

These domains differ significantly in the types of processes they manage. For example, the health sector is characterized by real-time, data-intensive operations involving patient records and service delivery, whereas education relies on periodic administrative cycles such as enrollment, evaluation, and certification. Social security processes are typically highly standardized and eligibility-based, involving large-scale application handling, while justice involves more complex, case-specific procedures with strong legal constraints.

Despite these differences, they share common challenges: large volumes of administrative processes, strong interdependence across organizations, and the need to coordinate policies that affect millions of citizens. These sectors are where the consequences of limited digital integration become most visible. When systems are fragmented or operate on incompatible platforms, citizens experience the inefficiencies directly: repeated document requests, delays, inconsistent records, or the need to navigate several separate portals for related services. For public institutions, this translates into duplicated effort, manual checks, and operational rigidity that restricts their ability to adapt or innovate (Xiao et al., 2023).

As a result, these domains rely heavily on complex workflows, cross-institutional information exchanges, and the availability of accurate, accessible, and interoperable data. Understanding these policy areas is therefore essential for interpreting how national priorities translate into everyday administrative practice.

2.4 E- GOVERNMENT

The digital transformation of the state has progressively reshaped how public institutions operate, exchange information, and deliver services. In Portugal, e-government initiatives have aimed to reduce administrative burden, promote interoperability, and provide citizens and businesses with simpler, more transparent access to public services (Alvarenga et al., 2020). Rather than focusing only on digitizing existing procedures, these efforts increasingly involve rethinking how information flows between institutions and how services can be designed around the needs of users.

A central principle of e-government is the transition from institution-centered processes to integrated service pathways. This requires reliable digital identification mechanisms, shared technological infrastructures, and common data standards that allow entities to communicate efficiently. It also requires re-engineering processes to remove unnecessary steps, automate routine tasks, and ensure that information is collected once and reused across the administration whenever legally possible (Xiao et al., 2023).

Digital identity systems have become a foundational component of this transformation, enabling secure authentication, digital signatures, and access to online services. At the same time, the expansion of digital public services has shifted many activities that once required in-person interactions into online environments, streamlining citizen experiences and reducing operational pressure on public institutions. This shift is the source of the evolution in modernization goals.

2.5 Challenges and Opportunities

Despite progress in digitalization, the Central Public Administration continues to face persistent structural challenges. Many systems remain fragmented, with limited interoperability and uneven adoption of shared digital infrastructures. This results in duplicated data collection, manual verification tasks, and workflows that depend heavily on human coordination and manual tasks. High volumes of repetitive, rule-based procedures continue to consume significant administrative capacity, creating delays and reducing service efficiency. At the same time, strict regulatory requirements and cautious innovation environments slow the experimentation and adoption of emerging technologies (Banfi & Pedroni, 2026).

These constraints, however, also create opportunities. Advances in automation and AI, particularly Agentic AI, offer new ways to streamline workflows, reduce manual effort, and improve coordination across institutions. Leveraging interoperable platforms, standardized data structures, and intelligent agents could help integrate services more effectively and support faster, more reliable administrative processes. By redesigning workflows around automation and data exchange, the public sector can improve efficiency and free employees for higher-value activities, while strengthening consistency, transparency, and service quality (Salah et al., 2025).

3 LITERATURE REVIEW

3.1 AGENTIC AI

Alan Turing's 1950 paper framed computational intelligence as the ability of a computer to perform actions that are indistinguishable from those of a human, proposing that, in order to evaluate the Intelligence of a system we should take into account the actions it can perform (Turing, 1950). Turing's contribution and proposal is key to the definition and the rise of research in the field of Artificial Intelligence today. This field has emerged greatly due to the advancement in computing capacity of systems and large availability of training data.

Having established that AI refers to technologies that mimic human behavioral and deducting skills for reasoning and decision making, it follows that, with computing power and memory capabilities that far exceed those of humans, AI naturally has the capacity to improve complex, information-heavy processes, particularly in domains characterized by formal bureaucratic rules, large volumes of data and repetitive decision making. From an organizational perspective this potential is highly appealing, however, this resource is still not being used to its full extent, despite recent technological advances in computing and the development of increasingly capable models (Costa et al., 2024).

The concept of intelligent agents is has its roots in earlier research within Artificial Intelligence, particularly from the 1980s onwards. In this context, an agent is typically defined as an entity that perceives its environment through sensors and acts upon it in order to achieve specific goals. Early work in agent theory explored autonomous behavior, decision-making, and interaction between multiple agents, leading to the development of multi-agent systems capable of coordinating actions in distributed environments. However, these systems were often limited by rigid rule-based logic and lacked the flexibility required to operate effectively in complex, real-world scenarios. Recent advances in machine learning, particularly with large language models and improved computational capabilities have significantly expanded the practical applicability of these concepts. This evolution has given rise to what is now referred to as Agentic AI, where agents are no longer constrained to predefined rules but can plan, reason, and adapt their actions dynamically within complex workflows (Ren et al., 2025).

As mentioned previously, today the most common form of AI generally used is Generative AI, but for these types of tasks and actions, the necessity for Agentic AI emerges. Now that these systems have solidified their capability in pattern recognition, content generation and action planning, the next step is to allow them to interact with existing information systems. AI is increasingly capable of operating within structured workflows rather than merely supporting isolated decisions. This marks the transition from AI as a passive assistant (Gen AI) to AI as an active component of organizational processes (Agentic AI) (Hosseini & Seilani, 2025).

These Agentic AI systems are particularly relevant for contexts such as public administration, where workflows are governed by formal rules, span multiple steps and often require coordination across organizational units or institutions. By embedding decision-making and action execution within the same system, Agentic AI offers an approach better aligned with the complexity and rigidity of administrative processes than traditional automation or assistive AI solutions (Schneider, 2025).

3.1.1 CONCEPTS

Firstly, it is important to distinguish between the different types of AI systems most commonly used today, as Artificial Intelligence is not a monolithic technology but rather a broad set of computational approaches designed to perform different functions. In practice, organizations deploy a variety of AI systems with distinct operational roles. Predictive AI systems are primarily aimed at prediction and classification tasks, analyzing historical and real-time data to estimate outcomes or support decision-making under uncertainty. In parallel, generative AI systems focus on the production of new content, such as text, summaries, recommendations, or code, based on learned data representations, enabling more flexible interaction with information and cognitive support in organizational contexts. Additionally, many information systems rely on deterministic or rule-based automation, where predefined rules trigger specific actions without learning or adaptation. These approaches differ substantially in their capabilities, levels of autonomy, and roles within organizational processes (Feuerriegel et al., 2023; Schneider, 2025).

Despite recent advances, most widely adopted AI systems remain limited to supporting specific tasks or isolated decisions. Predictive models estimate outcomes, generative models produce content or recommendations, and rule-based systems automate predefined procedures. In all cases, the execution of decisions and the coordination of multi-step processes typically remain external to the AI system itself, relying on human intervention or separate workflow mechanisms.

Agentic AI emerges at this boundary between decision support and action execution. Rather than merely producing outputs, agentic systems are designed to pursue objectives through the planning and execution of actions within structured environments. This introduces the concept of agency into AI systems, understood as the capacity to act autonomously under constraints, monitor outcomes, and adapt behavior accordingly. In this project, the focus is placed on Agentic AI as a distinct paradigm that extends beyond assistive or predictive functions toward direct participation in organizational workflows (Hosseini & Seilani, 2025).

3.1.2 ARCHITECTURES, PROTOCOLS AND FRAMEWORKS

The practical realization of Agentic AI systems relies on three interrelated layers that together enable autonomous action: architectures, protocols, and frameworks. Architectures define the high-level structural organization of an agentic system; protocols specify how agents communicate and coordinate actions; and frameworks provide the software infrastructure and tooling through which these architectural and interaction principles are implemented in practice (Derouiche et al., 2025)

The architecture of an agentic AI system refers to its high-level structural design, determining how components are organized and how they interact to perceive information, make decisions, and execute actions. Rather than focusing on individual models, architectural design defines how reasoning modules, memory components, decision logic, and interfaces to external systems are coordinated within a single operational framework. Architectures can be distinguished along several key dimensions that shape how autonomous behavior is organized. One such dimension is the distinction between single-agent and multi-agent architectures. In single-agent systems, a single autonomous entity is responsible for planning and executing tasks end-to-end, while multi-agent architectures distribute responsibilities across multiple specialized agents that coordinate to achieve shared objectives. Another important dimension concerns control structures, which may be centralized, where a single coordinating agent or controller directs actions, or decentralized, where agents operate with greater autonomy and coordination emerges through interaction rather than top-down control.

Architectures can also differ in how action execution is organized. Workflow-oriented orchestrators guide agents through predefined procedural steps, reflecting established process logic and fixed sequences of actions, while tool-oriented orchestrators dynamically select and invoke available tools based on contextual reasoning and situational needs. These architectural choices influence scalability, transparency, and controllability, making them particularly relevant in public administration contexts characterized by formal procedures, inter-organizational coordination, and strict accountability requirements.

Beyond architectural structure, protocols enable coordination and communication among agents by defining how tasks are delegated, how information is exchanged, and how actions are authorized. In multi-agent systems, protocols traditionally govern agent-to-agent interactions, specifying how responsibilities are distributed, multi-agent research introduced formal mechanisms such as the Contract Net Protocol for structured delegation and coordination.

In addition to these coordination mechanisms, recent developments in agentic systems highlight the importance of agent-to-environment protocols, which define how agents access external tools, data sources, and services. One emerging standard in this space is the Model Context Protocol (MCP), which provides a structured interface for exposing tools and contextual data to AI agents. Through MCP, external systems are described in a standardized format, allowing agents to dynamically discover available resources, invoke tools, and retrieve information in a consistent and reusable manner (Adimulam et al., 2026).

This distinction between agent-to-agent and agent-to-environment interaction is particularly relevant in the context of public administration, where processes depend on multiple interconnected information systems and require strict control over data access and execution. By making interactions explicit and standardized, protocols such as MCP contribute to ensure auditability, traceability, and clear responsibility attribution across complex workflows (Derouiche et al., 2025).

Finally, frameworks constitute the practical layer through which agentic architectures and protocols are implemented. They provide reusable software components for building, orchestrating, and deploying autonomous agents, including support for memory, tool integration, coordination, and execution control. Existing frameworks differ in their design philosophies, ranging from modular and flexible orchestration to more structured multi-agent collaboration, and many remain oriented toward experimentation rather than mature operational use. As a result, current solutions often lack built-in support for transparency, enforceable constraints, and systematic oversight, highlighting the gap between experimental agentic systems and production-ready deployments (Shakudo, 2025, p. 9).

3.1.3 Areas of Application

Agentic AI applications are emerging most visibly in domains where processes are complex, data-rich, and governed by formalized rules or compliance requirements, making them highly relevant for public administration. Recent work highlights use cases such as proactive eligibility assessment for social benefits, automated case handling, and cross-agency orchestration of life-event services (Salah et al., 2025). In these settings, agents not only support decision-making but also coordinate multi-step workflows, interact with legacy information systems, and trigger follow-up actions, thereby reducing processing times and administrative burden while preserving auditability.

In citizen-facing services, Agentic AI can operate as goal-oriented virtual case workers that combine conversational interfaces with backend process execution. This includes, for example, guiding citizens through complex procedures, pre-filling forms from existing records, and automatically routing requests to the appropriate organizational units, while maintaining a transparent record of decisions and system actions. In parallel, back-office applications focus on workflow optimization and resource allocation, where multi-agent systems monitor workloads, reassign tasks across units, and adjust process parameters in response to real-time conditions, an approach that is increasingly explored in enterprise contexts and can be adapted to public-sector workflows (Salah et al., 2025).

Beyond administrative procedures, Agentic AI is being piloted in policy implementation and regulatory compliance, where agents continuously monitor regulatory changes, interpret obligations, and propose or execute updates to procedures and internal controls. In critical domains such as healthcare, urban mobility, and cybersecurity, agentic systems are leveraged to integrate heterogeneous data sources, detect anomalies, and coordinate incident response across organizational boundaries, while keeping human oversight in decision points with high ethical or legal impact. These application areas illustrate how Agentic AI extends the reach of automation from isolated tasks to end-to-end public services, aligning technical capabilities with the multi-actor, rule-intensive nature of public administration.

3.1.4 CHALLENGES AND OPPORTUNITIES

The integration of Agentic AI into organizational workflows raises a set of technical, organizational, and governance challenges that go beyond those associated with conventional AI systems. On the technical side, orchestrating multiple agents, tools, and workflows in a controllable and verifiable manner remains difficult: multi-agent architectures must balance autonomy with reliability, avoid cascading failures across long execution chains, and provide clear monitoring of intermediate states. Many existing frameworks are still geared toward experimentation and lack robust mechanisms for enforcing constraints, guaranteeing predictable behavior, or providing systematic audit trails, capabilities that are essential in public administration. In addition, the heavy reliance on large language models introduces risks related to hallucinations, distributional shift, and uncertain performance in edge cases, which can be problematic when agents are entrusted with action execution rather than recommendation (Derouiche et al., 2025) .

From an organizational and legal perspective, Agentic AI amplifies longstanding concerns around transparency, accountability, fairness, and data protection in the public sector. When agents plan and execute multi-step actions, assigning responsibility for outcomes becomes more complex, especially if decisions emerge from interactions among multiple agents, systems, and human actors. Public administrations must therefore design explicit interaction protocols, logging mechanisms, and human-in-the-loop safeguards that ensure traceability, enable ex post review, and clarify the responsibility for each decision and action. Furthermore, the deployment of agentic systems interacts with existing administrative law, procurement rules, and public-sector values, requiring careful alignment with principles such as legality, proportionality, and equal treatment of citizens (Salah et al., 2025).

At the same time, the opportunities associated with Agentic AI are substantial. By shifting from static, process-driven automation toward outcome-oriented, adaptive workflows, agentic systems can reduce administrative burden, shorten processing times, and free civil servants to focus on complex judgement, relational work, and oversight. In areas such as benefit administration, licensing, and regulatory supervision, agents can proactively identify eligible cases, detect anomalies or non-compliance, and coordinate remedial actions across agencies, contributing to more responsive and data-driven governance. More broadly, Agentic AI offers a technological basis for rethinking the “operating system” of public administration, enabling continuous improvement of workflows and more integrated service delivery, provided that appropriate governance frameworks, technical safeguards, and institutional capacities are put in place (Robinson, 2025).

3.2 RELATED WORK

3.2.1 PRISMA PROTOCOL

This literature review followed the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) protocol to ensure a transparent and replicable process for identifying, screening, and selecting relevant studies. PRISMA was used as a methodological framework to document how the initial pool of records was progressively refined into a focused set of core references aligned with the objectives of this thesis on Agentic AI in Public Administration, with particular emphasis on processes, workflows, and service delivery.

In the identification stage, records were retrieved from two major academic databases to ensure broad coverage. A total of 950 publications were collected from Scopus and 500 from ScienceDirect, resulting in 1 450 records. After merging the two datasets, 126 duplicate entries were removed, yielding a consolidated corpus for screening.

The screening stage was conducted at the title and abstract level using three relevance dimensions. First, papers were required to explicitly address artificial intelligence, with particular attention to agentic or autonomous AI systems. Second, studies needed to be situated in a public-sector context, including public administration, government, or public services. Third, the focus had to extend beyond isolated technical solutions to encompass processes, workflows, services, or automation. Only publications simultaneously meeting all three criteria were retained. This systematic filtering reduced the dataset to 403 papers.

Finally, a manual analysis was performed to assess conceptual fit and substantive relevance to the research questions. This step excluded papers that were only tangentially related or that used similar terminology in unrelated contexts. The resulting shortlist of 64 papers was then examined through full-text reading, leading to the identification of 13 core references. These final studies constitute the analytical foundation of the literature review and are summarized in Table 6, representing the most relevant contributions for conceptualizing and contextualizing agentic AI in public administration.

3.2.2 PRISMA EXECUTION

Table 1 - Literature Reviews Research Questions

SLRQ1	What is the current status of research in Agentic AI use in public administration?
SLRQ2	What are the major issues of AI in public administration?
SLRQ3	What kind of Agentic AI technologies are currently useful in this area?
SLRQ4	What are the most relevant use cases of the use of Agentic AI in public Administration?

These research questions define the analytical scope of the systematic literature review. They are designed to capture both the current state of research and its practical implications.

Table 2 - Systematic Literature Review's Keywords

	Public Administration	**Agentic AI**
Keywords	Data	Artificial Intelligence
	Process	Software Agents
	Governance	Architectures
	Services	Protocols

The search string used was: **"Public Administration" AND ("Data" OR "Process" OR "Governance" OR "Services" OR "Compliance" OR "Use Case") AND ("Agentic AI" OR "Artificial Intelligence" OR "Software Agents" OR "Architectures" OR "Protocols"**).

The search was conducted in January 2026 on the following scientific information resource databases:

Table 3 - Systematic Review's Resource Databases

Resource Database	**Resource URL**
Scopus	https://www.scopus.com/home.uri

Science Direct	https://www.sciencedirect.com/

These databases were used to apply the defined search string and retrieve the initial set of records considered in the identification stage of the review. The results obtained were then subjected to a structured filtering process based on predefined inclusion and exclusion criteria.

Table 4 - Systematic Review's inclusion and exclusion criteria

Inclusion Criteria	**Exclusion Criteria**
Any scientific article showing evidence of Agentic AI in public administration	Papers focusing on Agentic AI but without potential interest for public administration
Paper must be a peered reviewed conference or journal paper written in English	Articles not in English and duplicate papers
Paper is published between 2021 and 2026	Articles published before 2021
Paper is published between 2021 and 2026	Non-academic or non-scientific papers (e.g., websites, magazines reports, newspapers, consulting articles, books, citations)
	Papers with titles outside the scope of this work

These criteria were applied systematically during the screening and eligibility stages to ensure consistency in the selection process. By restricting the scope to recent, peer-reviewed, and contextually relevant studies, the review maintains a strong alignment with the research objectives. This approach also helps ensure that the final set of papers reflects both the current state of the field and its most relevant developments.

- PRISMA Protocol Selection Phases

The overall selection process is summarized in Figure 1, which follows the standard PRISMA structure. This diagram provides a visual representation of how the initial set of records was progressively filtered through the different stages, from identification to final inclusion.

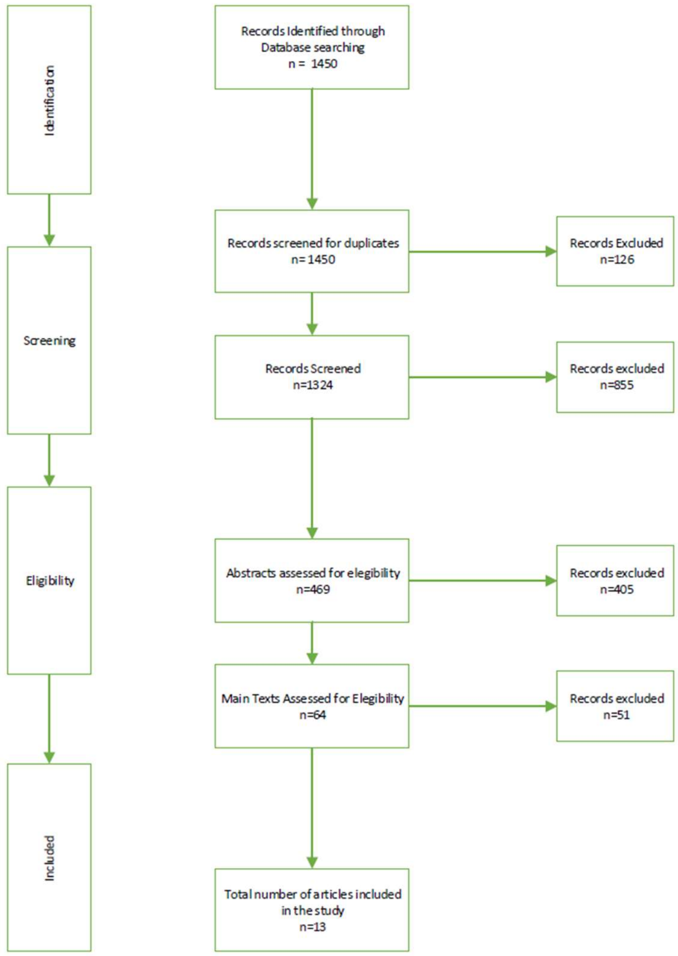

Figure 1 - PRISMA Protocol Selection Phases

As illustrated in the figure, the review process involved multiple stages of filtering, combining automated and manual screening. This structured approach ensures transparency in how the final set of studies was obtained and reinforces the replicability of the methodology adopted in this research.

Table 5 - Relevant Papers

#	Authors	Article	Contribution	Publication Type
#1	(Aljuneidi et al., 2023)	Did that AI just Charge me a Fine? Citizens' Perceptions of AI-based Discretion in Public Administration	Citizen's perceptions and acceptance of autonomous AI systems exercising	Conference Article

#	Authors	Article	Contribution	Publication Type
			authority in decision making in public services	
#2	(Gong et al., 2025)	Recovery from AI government service failures: Is disclosing the identity of the AI agent an effective strategy?	Improving user tolerance and trust in AI-based government services after failures through transparent disclosure that the service agent is AI	Article in Scientific - journal
#3	(Nikiforova et al., 2026)	Proactive Public Services in the Age of Artificial Intelligence: Towards Post-Bureaucratic Governance	Process-level concept of proactive public services, showing how AI enables automated services across workflows. The role of agentic AI as an operator in public administration rather than a passive decision-support tool.	Article in Scientific journal
#4	(Kurhayadi, 2025)	Evaluating the readiness of public institutions for AI-Driven decision making: A framework for adaptive governance	Public Sector AI readiness framework that clarifies the organizational and governance prerequisites for deploying agentic AI-driven decision-making in public administration.	Article in Scientific journal
#5	(Xiao et al., 2023)	Digital Government Transformation Through Artificial Intelligence: The Mediating Role of Stakeholder Trust and Participation	AI-enabled automation and decision support in public administration requires stakeholder trust and participation to achieve institutional transformation, informing governance conditions for agentic AI in public workflows.	Article in Scientific journal
#6	(Alhefaity et al., 2026)	Building trust through Technology: AI, public service perception, and citizen satisfaction in Abu Dhabi policing	Citizen trust and satisfaction as enabling conditions for AI-driven policing, showing how perceived service quality and institutional trust shape the acceptance of AI-enabled public services in law enforcement.	Article in Scientific journal
#7	(Banfi & Pedroni, 2026)	Towards a conversational public administration? Public services, chatbots, and new organizational challenges for local administrations	Analysis of AI as a relational and communicative agent in local public administrations, highlighting the importance of data governance, and organizational capacity	Article in Scientific journal
#8	(Al-Ansi et al., 2024)	Elevating e-government: Unleashing the power of AI and IoT for enhanced public services	Systematic synthesis of AI+IoT in e-government that helps frame where agentic AI fits within broader digital governance, especially around automation, data governance, and ethical constraints	Research paper

#	Authors	Article	Contribution	Publication Type
#9	(Robles & Mallinson, 2023)	Catching up with AI: Pushing toward a cohesive governance framework	Comparative analysis of EU, US, and private-sector AI governance frameworks that informs conditions under which agentic AI can be responsibly deployed in public administration	Research paper
#10	(Phan et al., 2026)	AI capacity in the public sector: Pathways from the environmental context to an organizational impact	Evidence that public sector AI adoption depends on regulatory support, government incentives, and citizen pressure, and that performance gains require AI capability plus strong management and an innovation culture.	Research paper
#11	(Wang et al., 2025)	Citizens' intention to follow recommendations from a government-supported AI-enabled system	Experimental evidence that algorithmic transparency, privacy concerns, and communication strategy shape citizen trust and compliance with government AI recommendations, supporting legitimacy requirements for agentic public services.	Article in Scientific journal
#12	(Nzobonimpa et al., 2025)	Automating public policy: a comparative study of conversational artificial intelligence models and human expertise in crafting briefing notes	limits of generative AI as an autonomous policy-writing agent, showing that while it supports efficiency, human judgment remains essential for substantive, accountable decision-making	Article in Scientific journal
#13	(Bondarenko et al., 2025)	Modernization of Administrative control over the legality of decisions of local self government bodies of Ukraine	Conceptual and comparative framework for agentic administrative control, showing how AI-enabled risk assessment and automated oversight can be embedded in multi-level public governance workflows while balancing legality, security, and local autonomy.	Article in Scientific journal

3.2.3 Results Analysis

Following the PRISMA-based selection process described in the previous section, this stage synthesizes the findings of the 13 included studies to identify their main contributions and to answer the systematic literature review questions. The analysis focuses on recurring themes, dominant challenges, enabling technologies, and representative applications of Agentic AI in public administration.

SLRQ1 - What is the current status of research in Agentic AI use in public administration?

The Literature indicates the research in Agentic AI in Public Administration is **recent and still evolving,** reflecting the early stage of Agentic AI adoption within public administration. Most studies do not analyze fully autonomous, end-to-end agentic systems integrated into public workflows. Instead, research primarily focuses on **AI-enabled public services** and governance arrangements that can be seen as precursors to more agentic configurations. In particular, a substantial share of the literature examines **citizen-facing applications**, investigating how citizens perceive, trust, and respond to AI-supported government services, including policing contexts and AI-based recommendation systems. These works emphasize behavioral outcomes such as trust, satisfaction, intention to comply, and willingness to follow AI-supported decisions, highlighting the importance of public acceptance as a precondition for adoption (Alhefaity et al., 2026; Wang et al., 2025) .

A smaller set of studies focuses on AI governance and institutional readiness. These works show that public administrations often lack the organizational capacity, skills, and governance maturity required for more advanced and autonomous AI systems. As a result, current implementations tend to remain limited in scope, and research on Agentic AI in public administration is still largely exploratory rather than empirical (Kurhayadi, 2025; Phan et al., 2026; Robles & Mallinson, 2023)

An important pattern is that most of the included studies focus on citizen-facing services and measure outcomes such as trust, satisfaction and perceptions. This does not necessarily mean that Agentic AI is mainly being implemented in citizen-facing settings. A likely reason is that these services are more visible and easier to study using surveys and experiments, while internal workflow automation and cross-agency orchestration are harder to observe and evaluate due to access restrictions and data sensitivity. This research may over-represent citizen-facing discussions compared to back-office and coordination-focused applications. This imbalance reinforces the relevance of this thesis, which focuses on workflow integration and task automation in Central Public Administration.

SLRQ2 - What are the major issues of Agentic AI in public administration?

Across the included studies, the most prominent issues relate to **trust, legitimacy, and accountability**, particularly in contexts where AI systems influence important decisions, sanctions, or access to public services. Several studies show that citizens are highly sensitive to the perceived role of AI in decision-making and that concerns over fairness, discretion, and human oversight directly affect acceptance and compliance. This is especially evident in cases where AI systems are perceived as replacing or constraining human judgement, rather than supporting it (Aljuneidi et al., 2023; Bokhari et al., 2025; Wang et al., 2025).

The literature also highlights the importance of **transparency and failure management**, with specific attention to how governments respond when AI-supported services fail or produce poor outcomes (Gong et al., 2025).

Finally, multiple studies underline **organizational readiness and capacity** as persistent challenges. Limitations in skills, governance maturity, and institutional learning mechanisms are identified as barriers to responsible and scalable deployment. Rather than technological constraints alone, these works point to deeper organizational and governance issues that hinder the effective supervision and integration of more autonomous AI systems in public administration (Kurhayadi, 2025; Phan et al., 2026; Robles & Mallinson, 2023).

SLRQ3 - What kind of Agentic AI technologies are currently useful in this area?

The most visible technologies in the reviewed literature are **conversational AI systems and chatbots**, which are said to be entry points for AI adoption in public services. These systems are primarily used to support citizen interaction, information report, and service navigation, while also raising new organizational challenges related to coordination, responsibility allocation, and service quality management. (Banfi & Pedroni, 2026)

Beyond conversational interfaces, some studies analyze **AI-enabled decision-support and service delivery systems**, including recommendation systems endorsed by government agencies and broader e-government infrastructures. These technologies are typically designed to better administrative processes rather than operate autonomously, supporting tasks such as prioritization, recommendation, or service optimization (Al-Ansi et al., 2024; Wang et al., 2025).

There is also emerging evidence of AI being applied to **policy support activities**, such as the drafting of briefing notes, where AI models are compared directly with human experts. While these systems are not fully agentic, they illustrate how AI can increasingly participate in knowledge-intensive public-sector tasks, potentially paving the way for more advanced agentic configurations in the future (Nzobonimpa et al., 2025).

SLRQ4 -What are the most relevant use cases of the use of Agentic AI in public Administration?

The most recurrent use cases identified in the literature relate to **citizen-facing government services**, where AI systems influence service satisfaction, compliance behavior, and uptake. These include applications in policing, AI-supported recommendations, and interactive service channels, where trust and perceived fairness play a central role in determining outcomes (Alhefaity et al., 2026; Bokhari et al., 2025; Wang et al., 2025).

Another set of use cases concerns **personalized public services**, reflecting a shift from just reactive, request-driven bureaucratic models. In this context, AI is utilized as enabling governments to anticipate citizen needs, automate eligibility checks, and coordinate service delivery across administrative boundaries, although empirical evidence remains limited (Nikiforova et al., 2026).

Finally, the literature includes applications related to **administrative control, legality oversight, and governance frameworks**, where AI supports monitoring, compliance, and control functions within public administration. These use cases emphasize the role of AI in strengthening institutional capacity and accountability, rather than replacing core administrative authority (Bondarenko et al., 2025; Robles & Mallinson, 2023).

4 METHODOLOGY

Given that the objective of this thesis is to design and propose a set of guidelines for the adoption of Agentic AI in Central Public Administration, a Design Science Research (DSR) methodology was selected. DSR is suitable for this study because its goal is to create and evaluate artifacts that address identified organizational problems. The proposed guidelines will constitute an artefact that translates evidence from the literature and the analysis of administrative workflows into recommendations for implementation.

4.1 DESIGN SCIENCE RESEARCH

Design Science Research is a research method in the field of Information Systems that focuses on the creation and evaluation of artifacts that solve identified organizational problems. This design is based on, initially, a literature review to assess the current state of research on the topic and later the creation of an artifact, this can be a model, a construct, a framework, etc. which is designed to solve a problem brought up by the existing literature. DSR emphasizes purposeful design and problem-solving relevance.

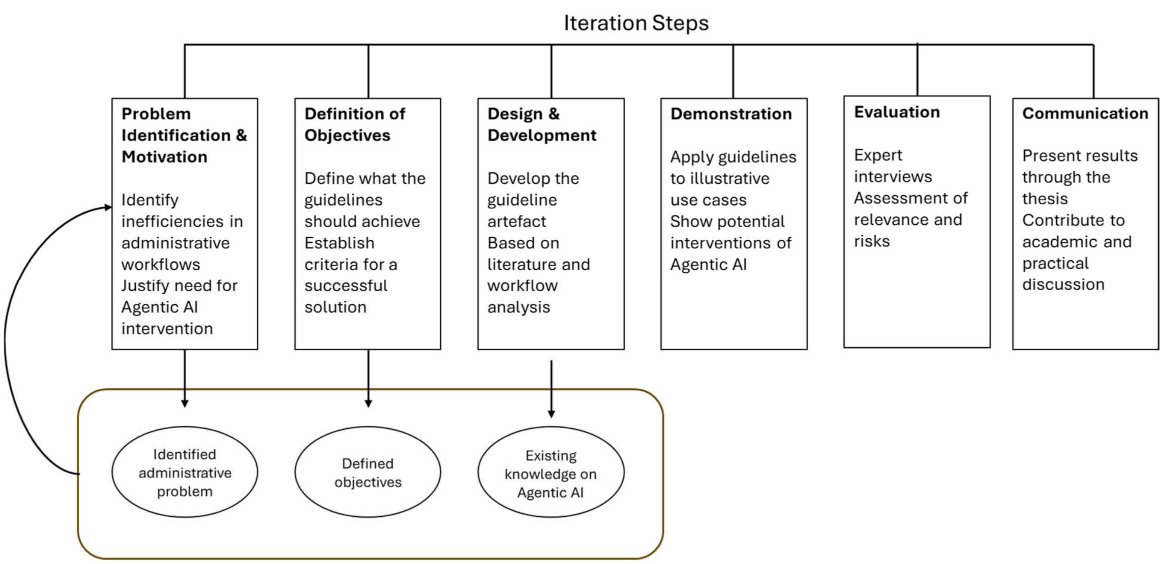

Figure 2 - DSR Process Model (Peffers) Adapted

According to Hevner et al (2004) Design Science Research is based on two complementary cycles:

The relevance cycle links the research to a concrete organizational problem, ensuring that the artefact addresses a real and practical need
The rigor cycle connects the design of the artefact to the existing scientific knowledge base.

This dual focus makes DSR particularly appropriate for the public administration context, where practical relevance, institutional constraints, and accountability requirements must be balanced with methodological rigor (Hevner et al., 2004).

In this thesis, the proposed guidelines for the use of Agentic AI in Central Public Administration are treated as an artefact, because they are designed based on the theoretical insights from the literature review.

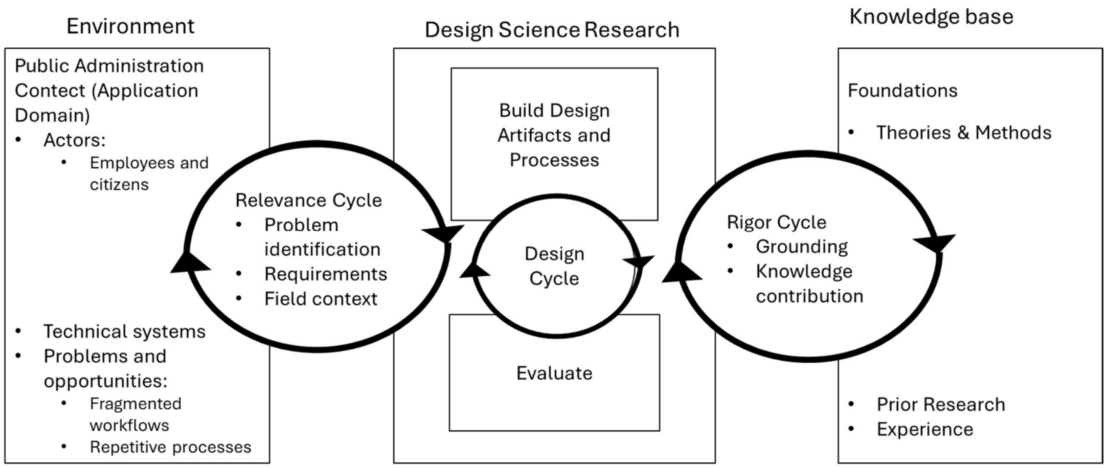

Figure 3 - Design Science Research Cycles in Public Administration Context

To operationalize DSR, this study follows the Design Science Research Process Model proposed by Peffers et al. (2007). This model defines a sequence of six iterative activities that guide the design, development, and evaluation of artefacts (Peffers et al., 2007).

The six activities of the Peffers et al. (2007) model are:

1. Problem identification and motivation
2. Definition of objectives for a solution
3. Design and development
4. Demonstration
5. Evaluation
6. Communication

These activities do not necessarily follow a strictly linear order and may be revisited iteratively as new insights emerge during the research process.

The first step, Problem identification and motivation focuses on clearly defining the problem to be addressed and justifying why a solution is needed. In DSR, problems are typically rooted in organizational practice and reflect inefficiencies, limitations, or unmet needs.

Based on the identified problem, the objectives of the solution are then defined. This step specifies what a successful solution should achieve and establishes the criteria that will guide the design of the artefact and its expected contribution.
In the Design and Development phase, the artefact is created. This may involve conceptual modelling, framework design, or the development of methods or guidelines informed by theory and empirical evidence.

Once developed, the artefact is demonstrated in a relevant context to illustrate how it can address the identified problem. Demonstration may involve simulations, hypothetical scenarios, or illustrative use cases that show the artefact's applicability in practice.

The Evaluation assesses how well the artefact meets the defined objectives. This may involve analytical reasoning, expert validation, or comparison with existing approaches.

Finally, the results of the research, including the artefact and its evaluation, are communicated to relevant audiences, such as researchers, practitioners, or policymakers.

4.2 Design Science Research implementation

The research begins with the identification of a practical problem in Central Public Administration, namely the existence of fragmented workflows, coordination difficulties between organizational units, and a high administrative burden in processes that are largely rule-based and repetitive. Although digital transformation initiatives are ongoing, many procedures still rely heavily on manual intervention and sequential task execution, which limits efficiency and responsiveness. These observations, supported by findings in the literature, motivate the exploration of Agentic AI as a potential means to improve workflow integration and task automation.

Based on this problem, the objectives of the research are defined around the development of a set of practical guidelines for the adoption of Agentic AI in Central Public Administration. These guidelines aim to support in identifying suitable use cases, understanding risks and constraints, and designing agentic solutions aligned with public-sector values.

The design and development phase is grounded in a comprehensive literature review on Agentic AI and its existing or proposed applications, with particular attention to the public sector. This review follows a PRISMA-based systematic approach and is complemented by an analysis of current administrative workflows to identify inefficiencies and points where agentic systems could realistically provide added value. Drawing on this combined evidence, a structured set of guidelines is developed, translating theoretical insights and empirical observations into concrete recommendations for implementation.

To demonstrate the applicability of the proposed guidelines, simplified models and illustrative use cases are developed. These examples show how Agentic AI could intervene in selected administrative contexts, such as coordinating tasks across organizational units or automating standardized procedural steps. The proposed guidelines are then evaluated through analytical reasoning and consistency checks against the literature, ensuring alignment with identified challenges, governance requirements, and best practices for AI adoption in public administration as well as through expert interviews.

Finally, the results of the research, including the guidelines and their underlying rationale, are communicated through this thesis, contributing both to academic discussion and to practical reflection on the future role of Agentic AI in public administration.

5 GUIDELINES FOR THE USE OF AGENTIC AI IN CENTRAL PUBLIC ADMINISTRATION

This chapter presents a set of guidelines for the use of Agentic AI technologies in Central Public Administration and its implementation process.

5.1 ASSUMPTIONS

The guidelines proposed in this research are based on the main findings of the literature review on Agentic AI and its potential applications in public administration. Several observations in the literature help identify the types of situations where these systems provide the most value.

The studies show that Agentic AI is more useful in environments with structured rules, repetitive procedures, and large volumes of administrative data. Many public administration processes follow formal workflows, which makes them compatible with systems capable of executing predefined steps.

Agentic AI differs from traditional AI tools because it is able to coordinate multiple steps of processes instead of simply providing predictions or recommendations. This makes it relevant for administrative procedures with several stages and many documents.

Another common finding is that many inefficiencies in public administration arise from fragmented workflows and coordination challenges between departments or institutions. Agentic systems can help address this issue by transferring information between systems and triggering the next procedural step automatically.

The literature also indicates that applications generally fall into two broad areas. On one side, agents can support citizen-facing services, helping individuals navigate procedures, submit applications and obtain information. On the other, they can assist internal administrative work, such as case handling or task allocation across units/entities.

Another point emphasized in the literature is that the use of AI in public administration must respect requirements of transparency, accountability, and human oversight. Administrative decisions often have legal implications, which means automated actions must remain traceable and understandable.

For this reason, most studies suggest introducing Agentic AI gradually within existing workflows, supporting specific steps of administrative procedures rather than replacing entire processes.

Based on these observations, the guidelines developed in this research focus on identifying the types of administrative processes where Agentic AI applications are most appropriate.

5.2 GUIDELINES

Based on the assumptions identified in the previous section, this research proposes a set of guidelines for identifying suitable applications of Agentic AI in public administration. The guidelines are structured as a mapping between common public administration processes and the types of Agentic AI capabilities that can support them

The table is separated between external and internal to distinguish the two different roles of agentic ai in Public Administration: Internal processes refer to administrative and operational activities within public institutions, external processes relate to interactions between public administration and citizens or organizations.

Table 6 - Opportunities of intervention of Agentic AI in PA processes

			Agentic AI Applications			
			Automated case handling	Cross-agency coordination	Workflow Assistance and Process Guidance	Monitoring And adaptive Management
Public Administration Processes	Internal	Interministerial Processes		1)	1)	
		Decision Process	2)			
		Application Processing	3)		3)	
		Regulatory monitoring				4)
		Resource allocation/workload management				5)
	External	Citizen Communication			6)	
		Decision Process			7)	
		Citizen application Submission	8)			8)

The table presents a cross-mapping between administrative processes and typical Agentic AI applications discussed in the literature. Each entry describes how agent-based systems could support a specific type of process by automating tasks, assisting employees, or coordinating workflows across organizational units.

<u>Guidelines</u>

1) Interministerial processes – These processes involve coordination between multiple ministries or public entities, often requiring the exchange of information, sequential approvals, and shared procedures. They can be fragmented across different information systems, which creates delays and increases the complexity of information exchange that could be simplified.

Agentic AI systems can support these processes primarily through **cross-agency coordination capabilities**. Agents can automatically exchange data between information systems, ensure that required information is available to each entity, and trigger the next step in the workflow once predefined conditions are met. This reduces the need for manual communication and minimizes delays caused by transferring cases between organizations.

In addition, **workflow assistance** agents can help guide employees involved in these processes by identifying required steps, validating inputs, and ensuring that procedures are followed consistently across institutions. This is particularly relevant in complex cases where responsibilities are distributed across units.

By overseeing the entire workflow, an agent can contribute to improving the process by helping detect bottlenecks, delays or redundancies that are not easily visible in fragmented system, which improves overall efficiency.

2) Internal Decision Processes - These processes refer to administrative decisions taken within institutions, often based on predefined rules, internal guidelines, and the analysis of administrative data. They typically involve the evaluation of cases such as approvals or classifications, and may require consistency, traceability, and compliance with legal frameworks.

Agentic AI systems can support these processes primarily through **automated case handling.** Agents can analyze administrative data from similar past cases, apply predefined rules or decision criteria, and generate decision recommendations. This reduces the time required for routine evaluations and supports more consistent decision-making across similar cases.

The role of the agent in this context is to handle high-volume and standardized cases, particularly those with clear and well-defined criteria. This allows human workers to focus on more complex cases that require contextual judgement, interpretation, or exception handling.

3) Application Processing – This process refers to the handling of administrative applications submitted within public institutions, including tasks such as data verification or document validation. These processes are typically repetitive, rule-based, and involve the interaction with multiple internal records and information systems.

Agentic AI systems can support these processes primarily through **workflow assistance and process guidance**. Agents can guide the processing of administrative applications by validating submitted information, retrieving relevant records from internal systems, and assisting employees in completing routine procedural steps. This helps ensure that applications are processed consistently and in accordance with established procedures, while reducing the likelihood of errors or missing information.

In addition, **automated case handling** can be applied to simpler or highly standardized applications. In these cases, agents process applications end-to-end by applying predefined rules, verifying eligibility criteria, and generating preliminary decisions or recommendations. This reduces processing time and allows human workers to focus on more complex or exceptional cases.

4) Regulatory Monitoring – This refers to tracking legal and regulatory changes and ensuring that internal procedures, rules, and administrative practices remain aligned with the applicable legal framework.

Agentic AI systems can support in these situations primarily through **monitoring and adaptive management** by monitoring legal and regulatory updates and identifying changes that may affect existing procedures or internal rules, by analyzing these updates in relation to current workflows, agents would be able to detect areas that require revision and ensure a faster and more efficient response to regulation changes. The agent could also suggest the necessary adjustments. This ensures that regulatory changes are incorporated in the administrative processes.

5) Resource allocation/workload management - Involves distributing tasks and managing workloads across organizational units to ensure efficient use of resources. It requires continuous visibility over task volumes, processing times, and available capacity.

Agentic AI systems can support this function through **monitoring and adaptive management**. Agents can track workloads across units in real time, analyse their capacity, and dynamically reassign cases to balance workload, contributing to reduce bottlenecks while preventing delays and improving efficiency.

6) Citizen Communication - Involves interactions between public institutions and citizens, such as providing information and supporting users in navigating administrative services or procedures.

Agentic AI systems can support this function through **workflow assistance and process guidance**. They can interact with citizens through conversational interfaces, guiding them through administrative procedures and helping them understand requirements and next steps. They can also provide personalized guidance based on the citizen's situation, improving accessibility and reducing the need for direct human intervention in routine interactions.

7) External Decision Process – Refers to the evaluation stage of administrative processes that directly affect citizens, such as approving requests or determining eligibility. These processes require a high level of accuracy and consistency, as decisions must follow established rules and directly impact individuals.

Agentic AI systems can support these processes through **workflow assistance and process guidance**. Agents can assist employees during the evaluation stage by validating information, highlighting relevant criteria, and ensuring that all required steps are followed. This supports more consistent and reliable decisions, while keeping human oversight in the final outcome.

8) Citizen Application Submission – In any process where citizens submit applications for public services or benefits, in which they provide information or complete forms, there are often errors or incomplete submissions, which can delay processing.

Agentic AI systems can support this function through **workflow assistance and process guidance**. These systems could guide citizens through the submission process, pre-fill forms using existing government records, and validate inputs in real time. This reduces errors and improves the quality of submitted applications.

In addition, **monitoring and adaptive management** can be used to track submission patterns and identify recurring issues in the input process. This allows institutions to improve how applications are received and handled over time.

5.3 Implementation Process

With some of the suitable applications for Agentic AI in Public Administration identified, it is also necessary to define how to introduce these technologies in practice. The workflow presented below is intended to guide the implementation of this technology across the types of processes discussed in the previous section, even though specific requirements may vary depending on the nature of the process. This structured approach ensures that adoption is carried out in a controlled and consistent manner.

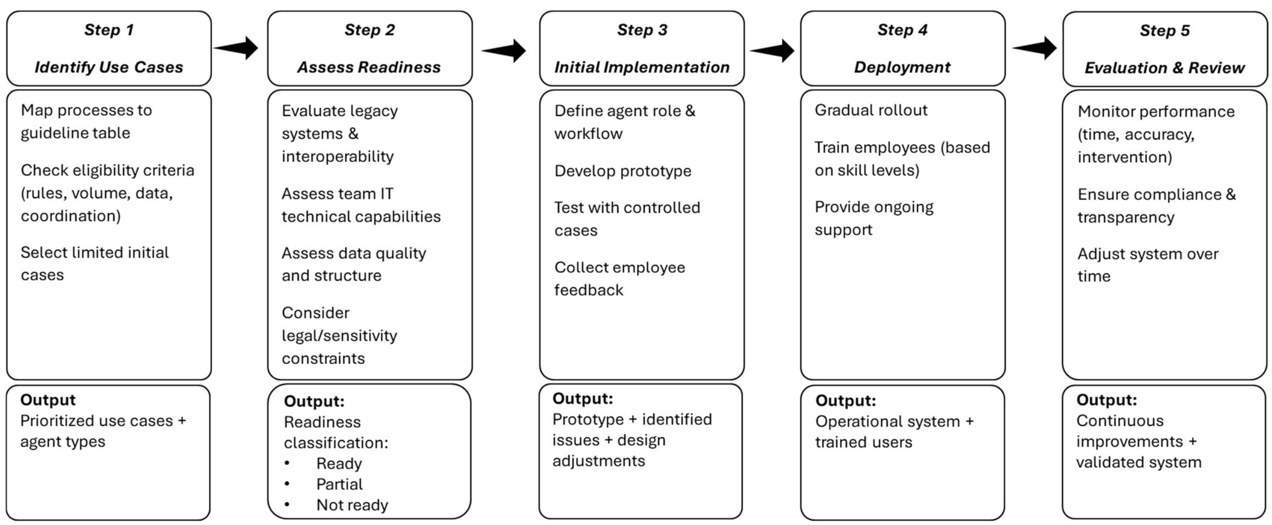

Figure 4 - Implementation Process Workflow

Step1: Identifying Potential Use Cases

The first step consists of identifying administrative processes that are suitable for the application of Agentic AI. This involves reviewing the types of processes referred to in the guideline table (Table 6 - Opportunities of intervention of Agentic AI in PA processes) and matching them to the organization.

These processes should meet a set of criteria, derived from the literature, particularly they should:

- Follow clearly defined rules and procedures
- Involve repetitive and high volume tasks
- Require coordination across multiple steps or units
- Rely on structured and accessible data

Based on this mapping, a small number of use cases should be selected for further evaluation, keeping the initial implementation limited. After completing this step, a prioritized list of candidate processes is obtained, including the types of agents required for each process, as defined in the guideline table.

Step 2: Assess Organizational Readiness

This step consists of evaluating if the organization has the technical and structural conditions to support the implementation of Agentic AI, considering dimensions such as the state of internal information systems and the technical abilities of the teams. This one of the most critical stages of the process, as the feasibility, cost, and duration of implementation depend largely on these factors.

The assessment should be structured across four main dimensions.

The state of the legacy systems, including flexibility and compatibility with modern technologies, as well as the organization of data structures can limit the ability to introduce new solutions or require significant adaptation efforts. Although these changes may be complex and resource-intensive, they can also generate broader organizational benefits, as agent-based systems rely on well-structured data and efficient information systems, which are beneficial to the organization more broadly.

It is also important to evaluate the technical capabilities of the teams and the organization's data management practices. The ability to develop, integrate, and maintain these systems depends not only on infrastructure, but also on the availability of technical expertise. Where these capabilities are limited, implementation may require strengthening the IT function, which represents an additional investment but can contribute to long-term improvements in organizational capacity.

Finally, the nature of the process should be considered, including the sensitivity of the data involved, legal constraints, and requirements for transparency and accountability, as these factors may impose additional limitations on how Agentic AI can be applied.

The outcome of this step is a readiness assessment.

To operationalize the readiness assessment described above, a structured evaluation framework is proposed. The first table defines the key dimensions of readiness and provides qualitative criteria to assess each dimension across three levels. This allows for a standardized evaluation of the organization's conditions.

The individual assessments are then aggregated into an overall readiness score, which is used to classify the organization into one of three readiness levels. This classification provides a clear indication of whether the implementation of Agentic AI can proceed directly or whether preparatory actions are required.

Finally, the framework links identified weaknesses to targeted improvement actions.

Each dimension is evaluated on a three-level scale (low, medium, high), allowing for a simplified but consistent assessment across different organizational contexts.

Table 7 - Readiness Assessment Criteria Scoring

Dimension	Criteria	Low (1)	Medium (2)	High (3)
Information Systems	Interoperability and flexibility	fragmented, low compatibility	Partial integration, some constraints	interoperable, adaptable
Data Management	Structure, quality, accessibility	Unstructured and inconsistent	Partially structured, some gaps	Structured, consistent, accessible
Technical Capabilities	Skills, IT capacity, AI readiness	Limited expertise, no AI experience	Basic technical capacity	Strong IT team, AI/data expertise
Process Constraints	Legal, sensitivity, transparency	Highly restrictive, sensitive	Moderate constraints	Clearly defined, manageable constraints

The total score is calculated as the sum of the values assigned to each dimension, resulting in a minimum score of 4 and a maximum score of 12.

Table 8 - Readiness assessment Score Interpretation

Score Range	Readiness Level
10–12	Fully prepared
7–9	Partially prepared
4–6	Not prepared

While the overall score provides a general indication of readiness, the assessment is not fully compensatory. A low score in any individual dimension may represent a critical limitation that cannot be offset by strong performance in other areas.

Therefore, the main interpretation of this score should be the expected complexity of implementation, particularly in terms of required time, cost, and organizational effort.

In particular, if any dimension is assessed at the lowest level, the organization should not be considered fully prepared, regardless of the total score. In such cases, targeted improvements in the identified dimension are required before proceeding.

In cases where the organization is only partially prepared, targeted improvements may be required depending on the dimension where limitations are identified. These may include upgrading or adapting legacy systems to improve interoperability, improving data quality and accessibility, or strengthening technical capabilities through training or recruitment. Addressing these gaps is necessary to ensure that the implementation can proceed effectively.

This classification informs whether the process can move forward to the next stage or whether preparatory actions are required.

Table 9 - Identified Issues and possible actions

Weak Dimension	Typical Issue	Recommended Action
Information Systems	Legacy systems, low interoperability	System modernization or redesign
Data Management	Poor quality, fragmentation	Data cleaning, standardization, integration
Technical Capabilities	Lack of expertise	Training programs, targeted hiring
Process Constraints	Legal/sensitivity barriers	No direct solution; Higher implementation complexity, requiring stricter controls and limited flexibility in system design

Step 3: Initial Implementation – In the third step a controlled, small-scale version of the proposed solution is developed in order to test its behavior before the full deployment. The objective is to evaluate how the system performs in a controlled setting and how it interacts with existing workflows and data.

The first activity consists of defining the system design. This includes specifying the role of the agent within the process, the sequence of actions it will perform, and how it interacts with existing systems and employees. At this stage, it is important to clearly define the scope of automation, identifying which tasks are to be fully automated and which require human validation or intervention, it is also important to assess how the agent will access the data and services required to perform its tasks. This includes identifying whether existing systems can be exposed through standardized interfaces, such as the Model Context Protocol (MCP), or whether additional custom integration mechanisms need to be developed, which would delay the process.

The system should be tested using simulated or controlled cases that reflect realistic conditions. This allows the identification of potential issues, including errors in decision logic, limitations in data availability, or difficulties in integration with existing systems.

At this stage, the involvement of employees with practical experience in the process is very important. Their knowledge of procedural details, exceptions, and informal practices can help ensure that the system reflects how processes operate in reality rather than only how they are formally defined. Their input is therefore essential for identifying limitations and improving the system before deployment.

Outputs of this step:

- A functional prototype of the proposed solution.
- Identification of practical limitations
- Adjustments to the process design based on testing and employee feedback

Step 4: Deployment

This step is the introduction of the system into the real workflows in a gradual and controlled manner. Deployment should expand progressively as the system demonstrates reliability.

A key component of this stage is the training and support provided to employees who will interact with the system, which should take in account the varying technical capabilities of the employees. Assessing these is essential so that training is appropriately designed and answers the needs of users. This should enable the understanding of how to use the system, how to interpret its outputs, and how to act in situations where intervention is required. Training should be aligned with the level of interaction each role has with the system.

In addition to training, continuous support should be provided during the initial stages of deployment to address difficulties, clarify usage, and ensure that the system is properly integrated into daily workflows.

Step 5: Continuous Evaluation and Review

This step consists of monitoring the system after deployment to ensure that it continues to operate as intended.

Performance should be assessed through key indicators such as processing time, accuracy of decisions, and frequency of required human intervention. In parallel, the system should be periodically reviewed to ensure continued compliance with administrative rules, as well as requirements for transparency and accountability.

Based on these observations, adjustments may be made to improve performance, correct unintended behavior, or adapt the system to changes in regulations or processes.

6 USE CASE

This subchapter shows how the application of the framework proposed in the previous section in practice. To demonstrate this, a fictional use case is used as an example.

Due to its illustrative nature, the example does not aim to represent a complete or fully realistic implementation. The elements described are simplified and intended only to demonstrate how the different steps of the framework can be applied in a structured way. As a result, some aspects of the process are not explored at the level of detail that would be required in a real-world implementation.

This example should be seen as an initial demonstration of how the framework can be applied. Further work, including real-world testing and validation with experts, would be needed to properly assess its effectiveness. The framework is further evaluated in the next chapter through expert interviews.

Use Case Description

In this simplified example of a public administration process, a citizen applies for a public financial support program. The eligibility for this support depends on a set of predefined criteria, such as income level, assets, and household composition. Based on this evaluation, the application may result in full support, partial support, or rejection.

This type of process is common in public administration and typically involves multiple steps, including data collection, verification of information, and decision-making. In many cases, these steps are performed manually and require coordination between different units or information systems, which can lead to delays and inefficiencies.

Step 1: Identifying Potential Use Cases

In the context of this example, the process has already been defined as the eligibility assessment for a public financial support program. Therefore, this step focuses on evaluating its suitability for the application of Agentic AI.

Following the criteria defined in the framework, this process can be considered a strong candidate, as it:

- follows clearly defined rules and procedures,
- involves repetitive and high-volume tasks,
- requires coordination across multiple steps, including data collection and verification,
- relies on structured and accessible data from different sources.

Based on this assessment, the process is selected for implementation.

According to the classification proposed in the guideline table, this use case primarily involves two types of agentic applications:

- **Automated case handling**, responsible for processing individual applications and producing eligibility decisions,
- **Monitoring and adaptive management**, due to the need to coordinate and validate information from multiple sources and to handle variations in case complexity, such as differences in income, assets, household composition, and other relevant factors.

This assessment and selections should take about 1 to 2 weeks.

Step 2 - Assess Organizational Readiness

In this step, an assessment is conducted to evaluate whether the organization has the necessary conditions to support the implementation of the agent.

The evaluation follows the four dimensions defined in the framework: Information Systems, Data Management, Technical Capabilities, and Process Constraints.

Firstly, the state of internal information systems is evaluated. In this process, the eligibility decision depends on information such as income, assets, and household composition, which are typically stored in different systems and in some cases, managed by different public entities. If these systems are not interoperable, or if data exchange mechanisms are limited, the implementation of an automated solution may be constrained or require significant adaptation efforts. Given this distribution, lets consider these systems are fairly interoperable so for this dimension the attributed note is assessed as **Medium (2).**

The data quality is then evaluated in terms of structure and accessibility. In this case, the data is fragmented across different databases, with inconsistencies in formats and missing values in some records. These issues limit its direct use in automated decision-making and require significant preprocessing before implementation. As a result, this dimension is assessed as **Low (1)**.

The technical capabilities of the organization are also considered. The IT team demonstrates general technical competence and is capable of supporting system integration and maintenance. However, it has limited experience with AI-based systems, particularly in the context of automated decision processes. This dimension is therefore assessed as **Medium (2)**.

Finally the nature of the process introduces additional constraints. This use case involves sensitive personal and financial data, which requires strict compliance with legal and regulatory requirements. In addition, decisions must be transparent and justifiable, meaning that the system should be able to explain how a given outcome was reached. This may limit full automation and require human oversight in certain cases, particularly when data is incomplete or when exceptions occur. This dimension is therefore assessed as **Medium (2)**.

Based on this evaluation, the readiness scores are as follows:

- Information Systems: Medium (2)
- Data Management: Low (1)
- Technical Capabilities: Medium (2)
- Process Constraints: Medium (2)

The total score is 7, corresponding to a classification of **Partially Prepared** according to the framework.

To address these limitations, two main preparatory actions are defined. First, a reorganization and standardization of data storage is required to ensure that data is consistent, complete, and accessible across systems. This includes data cleaning, format standardization, and the consolidation of fragmented data sources. Second, targeted training programs are implemented to strengthen the IT team's capabilities in AI-related technologies, particularly in areas such as model integration and system maintenance.

These preparatory actions are assumed to take approximately **2 to 4 months**, after which the organization reaches a sufficient level of readiness to proceed to the initial implementation phase.

Step 3: Initial Implementation

The goal of step 3 is to develop a small scale version of the system.

The system is designed to follow the main steps of the process. It retrieves the relevant information available, applies the predefined eligibility rules, and produces a preliminary outcome/decision. In situations where data is incomplete, or inconsistent should be used to evaluate the ability of the agent to solve those problems, these cases require a more thorough experienced review so it is possible to establish the scope of what can or can't be automated. This ensures that the system supports the process without fully replacing existing decision mechanisms at this stage.

The system is tested using a set of simulated or past applications, allowing the evaluation of its performance under realistic conditions. This testing helps identify potential issues, such as incorrect application of rules, missing or inconsistent data, or difficulties in retrieving information from different sources.

Testing is carried out using representative cases, and the results are compared with expected decisions. Based on these tests, adjustments are made to improve performance and ensure alignment with the process.

This phase is conducted over approximately **1 to 2 months**.

Step 4: Deployment

In this stage, the system is introduced into the real workflow of the eligibility assessment process.

Its use becomes part of the standard procedure, supporting the processing of applications by generating preliminary decisions. Employees remain responsible for reviewing cases that require intervention, particularly those involving incomplete or uncertain information.

At the same time, employees receive training adapted to their role in the process. For example, caseworkers are trained on how to submit and review applications through the system or handle cases flagged for review. Supervisors may receive additional training focused on monitoring outputs and managing exceptions.

In addition to training, support mechanisms are made available during the initial phase of operation. This may include user guidelines or documentation, a support contact for reporting issues or clarifying system behavior, and internal follow-up sessions to address recurring difficulties and ensure consistent use of the system.

As the system becomes integrated into daily operations, its role within the process becomes more clearly defined, and the interaction between the system and employees stabilizes.

Step 5: Continuous Evaluation and Review

After deployment, the system is monitored to ensure that it continues to operate effectively within the eligibility assessment process.

Performance indicators such as processing time, consistency of decisions, and the frequency of cases requiring human intervention are evaluated. A reduction in processing time and manual workload is observed, indicating improved efficiency in the handling of applications. At the same time, periodic reviews are conducted to ensure that decisions remain compliant with eligibility rules and that outcomes are transparent and justifiable. Some inconsistencies are identified in cases involving incomplete or conflicting data, requiring further refinement of the decision logic.

Based on these observations, adjustments are made to improve system performance and address identified issues. The system is also updated when necessary to reflect changes in eligibility criteria or administrative procedures. Over time, this continuous evaluation process supports the stabilization and improvement of the system within the workflow.

7 EVALUATION AND DISCUSSION

7.1 EVALUATION

To evaluate the proposed guidelines and implementation framework, semi-structured interviews were conducted with three experts from different professional backgrounds. These include a data scientist in the banking sector, a public administration employee, and a professional working in digital transformation within the public sector. The objective of selecting participants with distinct profiles was to obtain diverse perspectives on the applicability and relevance of the proposed work.

Table 10 - Interviewees Profiles

Expert	Background	Line of work
E1	Data Scientist in the Banking Sector	Banking
E2	Administrative Public Worker	Public Sector
E3	Professional in Digital Transformation in the public Sector	Public Sector

All participants were presented with the guidelines and the implementation process prior to the interviews. The same four questions were used to guide the discussion, focusing on usefulness, criticism, implementation potential, and suggestions for improvement. This structure ensured consistency across interviews while allowing flexibility for additional clarification when necessary.

Q1: Do you consider the proposed framework and guidelines as useful and why? If not, why do you believe it is not?

Q2: Do you have any criticism towards the proposed framework and guidelines? Please explain.

Q3: Would you consider implementing the proposed framework and guidelines? Please clarify why/ why not.

Q4: Do you have any recommendation or suggestions for further improvements?

The full interview scripts and corresponding transcripts are presented in Appendix A.

It is important to note that, although both the guidelines and the framework were presented, the feedback provided by the experts focused almost exclusively on the implementation process. The guidelines themselves were generally perceived as clear, comprehensive, and representative of real public administration processes, and therefore did not generate significant criticism.

Overall Assessment of the Framework

All three experts considered the framework useful and applicable. A common observation was that the structured, step-by-step approach helps translate a conceptual understanding of Agentic AI into a practical implementation process. The division into stages was seen as particularly valuable in guiding organizations through adoption in a controlled and gradual manner.

The inclusion of a readiness assessment phase was consistently highlighted as one of the strongest aspects of the framework. Experts emphasized that evaluating organizational conditions before implementation is critical, especially in public administration contexts where legacy systems, data limitations, and institutional constraints can significantly affect feasibility. This aligns with the initial motivation of the research, which identifies structural inefficiencies and fragmentation in public sector workflows.

Additionally, the phased approach, particularly the inclusion of a pilot or initial implementation stage before full deployment, was recognized as good practice. Experts noted that this reduces risk and allows for adjustments based on real-world testing.

7.2 Discussion

The positive assessment of the step-by-step structure indicates that the framework offers a coherent approach. In particular, the relevance attributed to the readiness assessment phase reinforces the importance of considering organizational conditions and the validation of the phased implementation approach suggests that gradual deployment and initial testing are appropriate strategies for reducing implementation risk in complex administrative environments.

Despite the positive assessment, several limitations were identified, with strong convergence across interviews

Identified Limitations and Areas for Improvement

Step 1 – Identification of Use Cases

Two of the experts highlighted that the process of selecting and prioritizing use cases is not sufficiently detailed. While the framework provides general criteria (e.g., rule-based processes, high volume), it does not clearly define how to choose between multiple eligible processes in practice.

This was identified as a gap, as the initial selection of use cases directly influences the success of the implementation. Suggestions included introducing a more structured prioritization method, such as scoring mechanisms or weighted criteria, and explicitly incorporating input from employees who work directly with the processes. These employees are seen as having essential knowledge of operational details and potential inefficiencies.

Step 5 – Monitoring and Continuous Evaluation

Another recurring point of criticism concerns the level of detail in the monitoring phase. Although the framework includes performance evaluation and compliance checks, experts considered this step too generic.

Specifically, they suggested:

- Defining clearer and more concrete performance indicators (e.g., processing time or error rates)
- Including mechanisms for ongoing feedback collection from system users

The inclusion of user feedback was particularly emphasized, as it provides a practical perspective on how the system performs in real workflows and complements technical metrics.

Organizational and Human Factors

One of the most relevant insights from the interviews relates to the importance of organizational culture and human factors in the implementation process.

Experts highlighted that successful adoption of Agentic AI does not depend solely on technical readiness, but also on how the system is perceived and accepted by employees. Resistance to change, lack of familiarity with technology, and concerns about automation can significantly affect implementation outcomes.

This reinforces the importance of:

- Involving employees early in the process (particularly in Step 1 and Step 3)
- Adapting training to different levels of technical capability
- Monitoring not only technical performance but also user experience and acceptance

The need to account for these factors suggests that implementation should be understood not only as a technical process, but also as an organizational change process.

The evaluation conducted through expert interviews not only validates the usefulness of the proposed framework, but also provides a set of practical insights regarding its limitations and applicability.

7.3 Revised Model

These findings serve as the basis for refining the initial design. Given the feedback from the evaluation interviews, some of the steps of the implementation process were rewritten:

Table 11 - Summary of Changes in the Revised Model

Framework Steps	Changes
Step 1	Improved case selection
Step 3	Introduce metrics to the prototype
Step 5	More detailed Monitoring

Updated Framework

Step 1: Identifying Potential Use Cases

The first step consists of identifying administrative processes that are suitable for the application of Agentic AI. This involves reviewing the types of processes described in the guideline table (Table 6 - Opportunities of intervention of Agentic AI in PA processes) and matching them to the organization's activities.

To support a more structured selection, this step should combine general criteria with input from employees who are directly involved in the processes. These individuals have practical knowledge of day-to-day operations and can help identify tasks that are particularly repetitive, time-consuming, or prone to inefficiencies. This input could be obtained via scheduled meetings with representatives from the operational areas of the institution.

The identification of candidate processes should be guided by a set of criteria derived from the literature. In particular, suitable processes should:

- Follow clearly defined rules and procedures
- Involve repetitive and high-volume tasks
- Require coordination across multiple steps or units
- Rely on structured and accessible data

Based on this assessment, a small number of processes should be selected for further evaluation, keeping the initial implementation limited in scope. The output of this step is a prioritized list of candidate processes, including the types of agents required for each case, as defined in the guideline table.

Step 2: Assess Organizational Readiness

This step remains unchanged from the initial framework. The proposed readiness assessment model continues to provide a structured evaluation of the organization's conditions for implementation.

Step 3: Initial Implementation – In the third step a controlled, small-scale version of the proposed solution is developed in order to test its behavior before the full deployment. The objective is to evaluate how the system performs in a controlled setting and how it interacts with existing workflows and data.

The first activity consists of defining the system design. This includes specifying the role of the agent within the process, the sequence of actions it will perform, and how it interacts with existing systems and employees. At this stage, it is important to clearly define the scope of automation, identifying which tasks are to be fully automated and which require human validation or intervention.

The system should be tested using simulated or controlled cases that reflect realistic conditions. This allows the identification of potential issues, including errors in decision logic, limitations in data availability, or difficulties in integration with existing systems.

At this stage, the involvement of employees with practical experience in the process is very important. Their knowledge of procedural details, exceptions, and informal practices can help ensure that the system reflects how processes operate in reality rather than only how they are formally defined. Their input is therefore essential for identifying limitations and improving the system before deployment.

The development of the initial prototype also allows for the collection of preliminary performance data regarding the redesigned process. Indicators such as processing time and estimated human effort required per case can be used to assess the efficiency gains associated with the proposed solution.

In addition, it is possible to collect initial feedback from employees. This will help in the next stage as it helps identify potential gaps in user understanding and varying levels of technical readiness among employees. As a result, it supports the design of more targeted training and support measures, aligned with the needs and capabilities of different user groups.

Outputs of this step:

- A functional prototype of the proposed solution.
- Identification of practical limitations

- Adjustments to the process design based on testing and employee feedback

Step 4: Deployment

This step remains unchanged from the initial framework. The proposed approach to gradual deployment continues to provide a structured transition from prototype to real-world implementation. No significant limitations were identified in this stage during the evaluation phase.

Step 5: Continuous Evaluation and Review

This step consists of monitoring the system after deployment to ensure that it continues to operate as intended.

To support this, structured monitoring routines should be established based on a set of clearly defined performance indicators. These indicators should capture both technical performance and operational impact. In particular, the evaluation may include:

- Accuracy of the system's outputs or decision recommendations
- Processing time and overall improvements in process speed
- Estimated efficiency gains, including reductions in human effort
- Return on investment (ROI), to ensure the investment is delivering financial benefits
- Employee feedback and user experience
- Compliance with legal and regulatory requirements

Monitoring should be conducted on a regular basis, with the frequency depending on the criticality and complexity of the process. This ensures that potential issues are identified early and that the system remains aligned with organizational needs and constraints.

Based on the results of this evaluation, adjustments may be made to improve performance, correct unintended behavior, or adapt the system to changes in regulations, data, or administrative procedures.

8 CONCLUSIONS

8.1 SYNTHESIS OF THE DEVELOPED WORK

This thesis set out to explore how Agentic AI can be applied in Central Public Administration and to propose practical guidelines for its adoption. The research was motivated by the growing relevance of AI in organizational contexts and by the challenges that continue to affect public administration, namely fragmented workflows, coordination difficulties, repetitive administrative tasks, and limited process efficiency.

The literature review showed that research on Agentic AI in public administration is still recent and relatively limited. Most studies do not focus on fully autonomous systems operating in real workflows, but rather on adjacent applications such as conversational AI, AI-supported decision systems, and governance conditions. Even so, it was possible to identify consistent patterns. Agentic AI is particularly relevant in contexts with clear rules, repetitive procedures, structured data, and multi-step workflows involving several actors. At the same time, its use must remain aligned with transparency, accountability, legal compliance, and human oversight.

Based on these findings, this thesis developed a set of guidelines mapping public administration processes to Agentic AI applications. An implementation process was also proposed to support gradual and structured adoption, including use case identification, readiness assessment, initial implementation, deployment, and continuous evaluation.

To illustrate the framework, a fictional use case was presented. This example demonstrated how the implementation steps can be applied and how factors such as interoperability, data quality, technical capabilities, and process constraints affect feasibility. The evaluation phase, based on semi-structured interviews with three experts, supported the usefulness of the framework. The interviews confirmed that the structure is coherent and practical, particularly in its emphasis on readiness assessment and phased implementation. They also highlighted areas for improvement, namely a more structured approach to use case selection and a more detailed monitoring stage, which were incorporated into the revised model.

Overall, the thesis contributes to the discussion on Agentic AI in public administration by translating a recent and fragmented body of literature into a practical implementation framework. Rather than proposing full automation, it supports a gradual and controlled approach, where Agentic AI is introduced in selected processes to reduce administrative burden, improve coordination, and increase efficiency while maintaining institutional safeguards.

8.2 LIMITATIONS

This research has some limitations that should be considered.

First, the framework was not applied in a real public administration context. The use case presented is fictional and was only intended to illustrate how the framework can be used, which does not fully capture real-world complexity.

Second, the evaluation was based on a small number of interviews. Although the participants had relevant backgrounds, the sample is limited. A larger number of interviews, especially across different areas of public administration, would provide more robust validation.

Third, the topic of Agentic AI in public administration is still recent. As a result, the available literature is limited and often focuses on related technologies rather than fully agentic systems, meaning some parts of the framework rely on broader interpretations.

Finally, the framework was designed to be general across central public administration. While this increases applicability, it does not account for the specific characteristics of individual sectors.

8.3 Future Work

A key next step would be to apply the framework in a real public administration setting. Testing it in practice would allow a better understanding of its feasibility, impact, and limitations, and would help refine the proposed steps.

Future research could also expand the evaluation phase by conducting more interviews with a wider range of professionals. This would provide stronger validation and help identify differences across institutions and roles.

Another possible improvement would be to develop more detailed tools to support specific steps of the framework. For example, a more structured method for selecting and prioritizing use cases could be created, as well as clearer performance indicators for the monitoring stage.

Finally, future work could focus on specific areas of public administration, such as social services or taxation, to adapt the framework to more concrete contexts and make it more operational.

BIBLIOGRAPHICAL REFERENCES

Adimulam, A., Gupta, R., & Kumar, S. (2026). *The Orchestration of Multi-Agent Systems: Architectures, Protocols, and Enterprise Adoption* (arXiv:2601.13671). arXiv. https://doi.org/10.48550/arXiv.2601.13671

Al-Ansi, A. M., Garad, A., Jaboob, M., & Al-Ansi, A. (2024). Elevating e-government: Unleashing the power of AI and IoT for enhanced public services. *Heliyon*, *10*(23), e40591. https://doi.org/10.1016/j.heliyon.2024.e40591

Alhefaity, S. R. S. A., Mohamad, E., Jamli, M. R., Ito, T., Larasati, A., & Mohamad, N. A. (2026). Building trust through Technology: AI, public service perception, and citizen satisfaction in Abu Dhabi policing. *Multidisciplinary Science Journal*, *8*(4), 2026277–2026277. https://doi.org/10.31893/multiscience.2026277

Aljuneidi, S., Heuten, W., Tepe, M., & Boll, S. (2023, September 6). Did that AI just Charge me a Fine? Citizens' Perceptions of AI-based Discretion in Public Administration. *Proceedings of the 2023 ACM Conference on Information Technology for Social Good, GoodIT '23*. https://doi.org/10.1145/3582515.3609518

Alvarenga, A., Matos, F., Godina, R., & C. O. Matias, J. (2020). Digital Transformation and Knowledge Management in the Public Sector. *Sustainability*, *12*(14), 5824. https://doi.org/10.3390/su12145824

Banfi, G., & Pedroni, M. L. (2026). Towards a conversational public administration? Public services, chatbots, and new organisational challenges for local authorities. *RDBCI: Revista Digital de Biblioteconomia e Ciência Da Informação*, *24*, e026005–e026005. (Italy; Contemporary). https://doi.org/10.20396/rdbci.v24i00.8680284

Blanchet, A. (2025, September 23). Agentic AI and the future of work: Navigating technological promise and the risk of increased automation. *Equal Times*. https://www.equaltimes.org/agentic-ai-and-the-future-of-work

Bokhari, S. A. A., Park, S. Y., & Manzoor, S. (2025). Digital Government Transformation Through Artificial Intelligence: The Mediating Role of Stakeholder Trust and Participation. *Digital*, *5*(3), 43. https://doi.org/10.3390/digital5030043

Bondarenko, S., Niziaieva, V., Kravchenko, M., Kaliuha, Y., Kolisnichenko, R., & Tsumariev, M. (2025). Modernization of Administrative Control over the Legality of Decisions of Local Self-Government Bodies of Ukraine. *European Political and Law Discourse*, *12*(1), 31–44. https://doi.org/10.46340/eppd.2025.12.1.4

Corvalán, J. G. (2025, May 1). *Artificial intelligence agents and agentic workflows: The new frontier of automation*. ResearchGate. https://www.researchgate.net/publication/389738694_Artificial_intelligence_agents_and_agentic_workflows_the_new_frontier_of_automation

Costa, C. J., Aparicio, M., Aparicio, S., & Aparicio, J. T. (2024). The Democratization of Artificial Intelligence: Theoretical Framework. *Applied Sciences*, *14*(18), 8236. https://doi.org/10.3390/app14188236

Decreto-Lei n.º 21/2021 | DR. (2021). https://diariodarepublica.pt/dr/detalhe/decreto-lei/21-2021-159432384

Derouiche, H., Brahmi, Z., & Mazeni, H. (2025). *Agentic AI Frameworks: Architectures, Protocols, and Design Challenges*. https://arxiv.org/html/2508.10146v1

European Comission. (2018). *EU Coordinated Plan on Artificial Intelligence*. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52018DC0795

Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2023). *Genrative AI*.

Gong, Z., Han, X., & Zheng, Y. (2025). Recovery from AI government service failures: Is disclosing the identity of the AI agent an effective strategy? *Government Information Quarterly*, *42*(4), 102087. https://doi.org/10.1016/j.giq.2025.102087

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). *Design Science in Information Systems Research*.

Hosseini, S., & Seilani, H. (2025). The role of agentic AI in shaping a smart future: A systematic review. *Array*, *26*, 100399. https://doi.org/10.1016/j.array.2025.100399

Kurhayadi, K. (2025). Evaluating the readiness of public institutions for AI-Driven decision making: A framework for adaptive governance. *Edelweiss Applied Science and Technology*, *9*(4), 1569–1580. https://doi.org/10.55214/25768484.v9i4.6335

Martinho, D. R. P. (2024). *Building Trust in Artificial Intelligence* [NOVA Inforation Management School]. http://hdl.handle.net/10362/174866

Nikiforova, A., Rodriguez Müller, A. P., Tangi, L., & Martin-Bosch, J. (2026). Proactive Public Services in the Age of Artificial Intelligence: Towards Post-Bureaucratic Governance. In I. Lindgren, M. P. Rodríguez Bolívar, M. Janssen, E. Loukis, F. Mureddu, P. Panagiotopoulos, G. Viale Pereira, & E. Tambouris (Eds.), *Electronic Government* (pp. 395–413). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-01589-1_25

Nzobonimpa, S., Savard, J.-F., Caron, I., & Lawarée, J. (2025). Automating public policy: A comparative study of conversational artificial intelligence models and human expertise in crafting briefing notes. *AI & SOCIETY*, *40*(5), 3627–3639. https://doi.org/10.1007/s00146-024-02103-x

OECD. (2025). *Goverment at a Glance: Portugal*. https://doi.org/https://doi.org/10.1787/0efd0bcd-en

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*, 45–77.

Phan, T. T., Van Nguyen, P., Corvello, V., & Pham, T. T. P. (2026). AI capacity in the public sector: Pathways from the environmental context to an organizational impact. *Journal of Innovation & Knowledge*, *13*, 100931. https://doi.org/10.1016/j.jik.2025.100931

Ren, Y., Liu, Y., Ji, T., & Xu, X. (2025). *AI Agents and Agentic AI-Navigating a Plethora of Concepts for Future Manufacturing*. https://doi.org/10.48550/arXiv.2507.01376

Robinson, R. (2025, October 14). The Agentic State: A Global Framework for Secure and Accountable AI-Powered Government. *EDRM*. https://edrm.net/2025/10/the-agentic-state-a-global-framework-for-secure-and-accountable-ai-powered-government/

Robles, P., & Mallinson, D. J. (2023). Catching up with AI: Pushing toward a cohesive governance framework. *Politics & Policy*, *51*(3), 355–372. https://doi.org/10.1111/polp.12529

Salah, M., Alnoor, A., Abdelfattah, F., & Al Halbusi, H. (2025). *Agentic Artificial Intelligence in Public Administration: Foundations, Applications, and Governance*.

Schneider, J. (2025). *Generative to Agentic AI: Survey, Conceptualization, and Challenges* (arXiv:2504.18875). arXiv. https://doi.org/10.48550/arXiv.2504.18875

Shakudo. (2025, December 1). *Top 9 AI Agent Frameworks as of December 2025 | Shakudo*. https://www.shakudo.io/blog/top-9-ai-agent-frameworks

Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (2025, November 5). *The State of AI: Global survey | McKinsey*. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Turing, A. M. (1950). *Computing machinery and intelligence. Mind, LIX(236), 433–460.* https://doi.org/https://doi.org/10.1093/mind/LIX.236.433

Wang, Y.-F., Chen, Y.-C., & Chien, S. (2025). *Citizens' intention to follow recommendations from a government-supported AI-enabled system*. https://doi.org/10.1177/09520767231176126

Xiao, J., Zhang, H., & Han, L. (2023). How Digital Transformation Improve Government Performance: The Mediating Role of Partnering Agility. *IEEE Access*, *11*, 59274–59285. https://doi.org/10.1109/ACCESS.2023.3284793

APPENDIX A

Entreview 1 – Data Scientist (Banking Sector)

Interveiwer: Agora que viste a apresentação, vou-te fazer 4 perguntas sobre o framework proposto. Consideras o framework proposto útil? Porquê?

Expert 1: Sim, considero o framework útil, apresenta uma abordagem estruturada para a implementação de agentes de inteligência artificial. A divisão em etapas ajuda a perceber como passar de uma fase mais conceptual para uma aplicação prática. Também é relevante o foco na identificação de processos candidatos e na avaliação do nível de preparação da organização, porque isso é muitas vezes um ponto crítico na adoção de novas tecnologias. No geral, parece-me um modelo aplicável e com potencial para melhorar eficiência em processos administrativos.

Interveiwer: Tens alguma crítica ao framework?

Expert 1: Acho que há alguns pontos que podiam ser mais desenvolvidos. Na fase 1, a parte de seleção dos processos podia estar mais detalhada. Ou seja, apesar de existirem critérios, não é totalmente claro como priorizar ou escolher entre vários processos elegíveis na prática. Poderia haver instruções mais especificas de seleção. Outro ponto é a fase 5. A parte de monitoring e avaliação contínua está presente, mas está um pouco genérica. Seria importante explicar melhor como esse acompanhamento é feito na prática, que tipo de métricas são usadas de forma mais específica, e com que frequência se faz essa avaliação.

Interveiwer: Considerarias implementar este framework? Porquê / por que não?

Expert 1: Sim, consideraria implementar, principalmente porque o framework está alinhado com práticas já utilizadas em projetos de automação. A abordagem faseada, com piloto antes de deployment completo, faz muito sentido e reduz risco. Além disso, o foco na avaliação da organização antes da implementação é muito importante e muitas vezes negligenciado. Com alguns ajustes nos pontos que referi, acho que poderia ser aplicado em contexto real.

Interveiwer: Tens recomendações ou sugestões de melhoria?

Expert 1: Diria que seria importante aprofundar dois pontos principais: Como já referi, tornar mais claro o processo de seleção de use cases no início, talvez com um método mais estruturado de priorização (por exemplo, critérios ponderados ou scoring). Também desenvolver melhor a componente de monitoring no Step 5, incluindo exemplos de indicadores e mecanismos de controlo. De resto, o framework está bem estruturado e cobre os aspetos mais importantes da implementação.

Entrevista 2 – Public worker

Interveiwer: Agora que viu a apresentação, vou fazer quatro perguntas sobre o proposto. Considera as guidelines propostas úteis? Porquê?

Expert 2: Sim, considero que o framework é útil, principalmente porque organiza bem as várias fases de implementação e ajuda a estruturar um processo que normalmente não é muito claro dentro das instituições públicas. A divisão em etapas ajuda a perceber o que tem de ser feito em cada fase e torna mais fácil pensar numa possível aplicação.

Interveiwer: Tem alguma crítica ao framework e guidelines?

Expert 2: A parte das guidelines acho que está bem explicada e engloba muitos tipos de processos diferentes, todos presentes no dia a dia de instituições publicas, em relação as guidelines tenho algumas observações. Na fase 1, acho que devia incluir de forma mais explícita a recolha de opinião dos funcionários que trabalham diretamente nos processos. São essas pessoas que conhecem melhor os detalhes e conseguem identificar onde é que a aplicação de agentes pode realmente trazer benefícios. Na fase 2, fiquei com uma dúvida: essa avaliação e score é feito para cada processo ou é só um para toda a instituição?

Interviewer: Nesse caso, dos quatro critérios que avalio — Information Systems, Data Management, Technical Capabilities e Process Constraints — apenas o último varia claramente de processo para processo, e em alguns casos também a parte dos dados. Os outros dois são mais transversais à instituição. Além disso, o output da fase 1 é uma lista de processos, e cada um desses processos já terá um score atribuído na fase 2.

Expert 2: Ah ok, não tinha percebido isso.

Interviewer: Sim, de outra entrevista também já percebi que tenho de explicar melhor o processo de seleção na fase 1, mas enfim, próxima questão. Consideraria implementar estas guidelines? Porquê / por que não?

Expert 2: Sim, consideraria implementar, porque acho que o framework tem uma lógica estruturada que faz sentido no contexto da administração pública. A abordagem gradual, com fases bem definidas, ajuda a reduzir o risco e a adaptar a implementação à realidade da instituição. Também é importante o facto de considerar limitações como sistemas existentes e legais, que são muito relevantes neste contexto.

Interviewer: Tens recomendações ou sugestões de melhoria?

Expert 2: Sim, além do que já referi sobre envolver mais os funcionários na fase inicial, acho que no Step 5 podia incluir de forma mais clara a recolha de feedback dos próprios utilizadores do sistema ao longo do tempo. Por exemplo, entrevistar ou recolher opiniões dos funcionários depois da implementação para perceber como o agente está a funcionar na prática e identificar pontos de melhoria. Isso ajudaria a complementar os indicadores mais técnicos com uma perspetiva mais operacional do uso do sistema.

Entrevista 3 –Digital Transformation Professional (Public Sector)

Interviewer: Agora que viste a apresentação, vou-te fazer 4 perguntas sobre o as guidelies e framework propostos. Considera-los úteis? Porquê?

Expert 3: Sim, considero. Está bem estruturado e cobre as principais fases necessárias para introduzir agentes de inteligência artificial numa organização pública. Acho particularmente relevante o facto de incluir uma fase de implementação inicial com testes antes do deployment completo, porque isso permite perceber melhor como a solução funciona na prática e reduzir riscos. No geral, parece-me um modelo aplicável e alinhado com iniciativas de transformação digital.

Interviewer: Tens alguma crítica?

Expert 3: Sim, tenho algumas sugestões. Na fase 3, acho que podias incluir uma avaliação mais clara de como a instituição se está a adaptar a novas tecnologias. Ou seja, perceber não só a capacidade técnica, mas também como estas mudanças são vistas internamente, qual é a aceitação por parte dos colaboradores e o impacto na forma de trabalhar. Isso é importante porque a adoção destas soluções não depende apenas da tecnologia, mas também da forma como as pessoas as recebem.

Interviewer: Considerarias implementar este framework? Porquê / por que não?

Expert 3: Sim, consideraria implementar. O framework segue uma lógica consistente e inclui etapas importantes como testes iniciais, avaliação da organização e implementação gradual. Isso está alinhado com boas práticas em projetos de transformação digital. Além disso, permite ajustar a solução ao longo do tempo, o que é essencial neste tipo de iniciativas.

Interviewer: Tens recomendações ou sugestões de melhoria?

Expert 3: Sim, recomendaria incluir uma componente mais quantitativa na avaliação dos resultados. Por exemplo, medir indicadores como o tempo de processamento antes e depois da introdução dos agentes. Esse tipo de KPI ajuda não só a avaliar o impacto real da solução, mas também a demonstrar valor à organização e aos colaboradores. Este tipo de análise poderia ser feito após o desenvolvimento do protótipo e antes do rollout final, permitindo tomar decisões mais informadas sobre a expansão da solução.

ANNEX A (ETHICS COMMITTEE REPORT)

This is to certify that
Project No.: **DSCI2026-4-197419**
Project Title: **Agentic AI in Central Public Administration**
Principal Researcher: **José Cavaco**

according to the regulations of the Ethics Committee of NOVA IMS and MagIC Research Center this project was considered to meet the requirements of the NOVA IMS Internal Review Board, being considered **APPROVED** on 4/19/2026.

It is the Principal Researcher's responsibility to ensure that all researchers and stakeholders associated with this project are aware of the conditions of approval and which documents have been approved.

The Principal Researcher is required to notify the Ethics Committee, via amendment or progress report, of
- Any significant change to the project and the reason for that change;
- Any unforeseen events or unexpected developments that merit notification;
- The inability of the Principal Researcher to continue in that role or any other change in research personnel involved in the project.

Lisbon, 4/19/2026

NOVA IMS Ethics Committee
ethicscommittee@novaims.unl.pt

Data with Purpose.

Campus de Campolide, 1070-312 Lisboa, Portugal
novaims.unl.pt