



EIXO 4 – INTELIGÊNCIA ARTIFICIAL E TRANSFORMAÇÃO DIGITAL NA AUDITORIA PÚBLICA

TOWARD AI-READY AUDITING: A MODEL CONTEXT PROTOCOL (MCP) NETWORK FOR OVERSIGHT INSTITUTIONS

Adriano Marabuco de Albuquerque Lima

Ph.D. (in progress) and an M.Sc. in Computer Science from the Federal University of Pernambuco (UFPE). He also holds degrees in Law (USP) and Business Administration (FGV/EAESP), with specializations in Public Law and Machine Learning. His research focuses on artificial intelligence, time series forecasting, and auditing applications, with publications in *Procedia Computer Science* and leading Brazilian computing conferences (SBSI, ENIAC). Professionally, he serves as an External Control Analyst at the Court of Accounts of Pernambuco (TCE-PE) and has prior experience as an Internal Control Auditor in São Luís. His interdisciplinary background combines computing, management, and law, with a particular interest in applying AI and data science to enhance public sector accountability.

Lattes: <http://lattes.cnpq.br/5477365718517046>

E-mail: adrianomarabuco@tcepe.tc.br

ABSTRACT

Public audit institutions face growing complexity in analyzing procurement, payroll, and budget data, while current AI applications remain fragmented and raise concerns about transparency and trust. This paper proposes a **federated network of Model Context Protocol (MCP) servers**, operated by oversight bodies to expose validated datasets and analytic workflows through standardized, policy-aware interfaces. AI agents interact with MCP servers under controlled authentication and logging, producing reproducible and auditable outputs. We outline a reference architecture, an explainable agent workflow, governance safeguards, and an illustrative use case in payroll auditing, arguing that MCP



infrastructures can enable an **AI-ready, citizen-centric audit ecosystem** that enhances both efficiency and accountability.

Keywords

Model Context Protocol (MCP); Artificial Intelligence Agents; Public Sector Auditing; Digital Government; Accountability; Interoperability

1 INTRODUCTION

Public audit institutions must analyze ever-expanding data sources, from procurement contracts to payroll systems. Manual processes and fragmented IT solutions hinder efficiency, while AI-based initiatives remain isolated pilots. Without interoperability and governance, these tools risk opacity, bias, and erosion of public trust.

This paper addresses that gap by proposing a **federated MCP network** for oversight institutions. By combining AI agents with protocol-level safeguards, we move beyond isolated applications toward a transparent and secure digital infrastructure for auditing.

2. BACKGROUND AND RELATED WORK

AI Agents. AI agents are autonomous systems capable of perceiving inputs, reasoning, and taking actions to achieve defined goals (Karpas et al., 2022). Recent *tool-augmented agents* integrate large language models with external tools and APIs, enabling dynamic, multi-step tasks such as anomaly detection or report drafting.

Model Context Protocol (MCP). Introduced in 2024, MCP is an open standard that governs how agents securely access external datasets and workflows. It specifies uniform interfaces for authentication, provenance, and logging, ensuring



reproducibility and accountability. MCP allows institutions to publish validated services without losing sovereignty over sensitive data.

Relevance to Auditing. For oversight institutions, these developments are crucial. Current AI applications in auditing focus on narrow tasks such as fraud scoring (Zhang & Yang, 2021) or predictive analytics (Janssen & Kuk, 2016), often in isolation. Policy studies stress the need for interoperability and governance in digital government (Dawes, 2010; OECD, 2020). However, no prior work links agent-based AI architectures with protocol-level safeguards tailored to auditing. This paper fills that gap.

3 Proposed Approach: MCP Network for Oversight Institutions

Vision and Architecture. We propose a **federated network of MCP servers**, each managed by an oversight institution. These servers expose three service layers: (i) **Data services** (e.g., procurement, payroll, budget datasets), (ii) **Workflow services** (standardized auditing routines), and (iii) **Governance services** (authentication, logging, provenance). Servers interconnect using shared standards, allowing cross-institutional queries while preserving data sovereignty.

AI Agent Interaction Model. Agents follow a secure workflow: (1) discover available services via MCP registries; (2) invoke them with authenticated tokens; (3) execute workflows within institutional servers, ensuring sensitive data remains local; (4) return outputs with provenance metadata for reproducibility.

Governance and Security. The network is governed jointly by participating institutions. Safeguards include federated identity management, immutable logging, periodic certification of workflows, and access tiers aligned with privacy regulations.

Use Cases: Payroll and Procurement Auditing. Payroll fraud, such as ghost employees, continues to undermine efficiency and trust in the public sector. Traditional approaches rely on isolated institutional databases, manual checks, and



fragmented scripts, which are costly and prone to error. In a federated MCP network, each oversight body publishes validated payroll datasets through standardized schemas, ensuring interoperability without requiring centralization of sensitive data. This foundation allows agents to seamlessly discover and access relevant records across multiple institutions.

Once payroll datasets are available, AI agents invoke standardized auditing workflows exposed through MCP servers. Examples include **cross-database identity matching**, which compares employee identifiers across institutions to detect duplication or impersonation; **HR-payroll reconciliation**, which links attendance records and HR logs to salary payments; and **anomaly detection routines**, which flag unusual patterns in overtime or allowances. These workflows run within institutional servers, ensuring that confidential data never leaves its secure environment.


Procurement auditing presents a complementary challenge. Oversight bodies must review vast amounts of contract data, supplier registries, and bidding processes, often with limited analytical tools. The MCP network allows institutions to expose procurement datasets in harmonized formats and provide reusable workflows such as **bid pattern analysis**, which detects collusive bidding behavior; **price benchmarking**, which compares unit costs across contracts to identify overpricing; and **timeline analysis**, which flags tenders with unusually short submission periods. Agents can securely invoke these workflows across multiple MCP servers, producing insights that would be difficult to obtain within isolated silos.

The results of both payroll and procurement analyses are returned with structured rationales and detailed provenance metadata. This guarantees that auditors can reproduce every finding, compare methods across institutions, and communicate results transparently to legislatures and citizens. By supporting payroll fraud detection and procurement anomaly monitoring under a common protocol, the MCP network illustrates its capacity to scale across different audit domains. This



not only improves efficiency and accuracy but also establishes a collaborative and citizen-centric model for digital auditing.

Figure 1 illustrates the federated architecture: institutions operating MCP servers, agents accessing workflows, and auditors/citizens consuming accountable outputs.



4 CONCLUSION

We proposed a federated MCP network that allows oversight institutions to securely share datasets and auditing workflows for AI agents. By embedding governance, provenance, and interoperability at the protocol level, the approach advances from fragmented pilots to a reproducible, citizen-centric audit ecosystem. An illustrative use case in payroll demonstrates its potential. Future work includes prototyping and piloting with selected institutions to evaluate performance, security, and adoption.

REFERENCES



DAWES, S. Interagency information sharing: Expected benefits, manageable risks. *Journal of Policy Analysis and Management*, v. 19, n. 2, p. 253–277, 2010.

JANSSEN, M.; KUK, G. Big and open linked data (BOLD) in government: A challenge to transparency and accountability. *Government Information Quarterly*, v. 33, n. 4, p. 563–567, 2016.

KARPAS, E. et al. What can AI agents do? A survey. *Proceedings of AAI*, 2022.

OECD. Digital Government Review. Paris: OECD Publishing, 2020.

POWER, M. *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press, 2021.

ZHANG, J.; YANG, L. Detecting anomalies in public sector data with machine learning. *ACM Transactions on Management Information Systems*, v. 12, n. 3, p. 1–18, 2021.