



THE **AI** TRUST GAP REPORT

For the Public Sector



GLOBAL EXECUTIVE
INSIGHTS

What public sector leaders reveal







Table of Contents

Executive Summary & Key Findings	3
Introduction	4
Live Data for Operational Awareness	5
The Right Data: Relevance, Context, and Semantic Consistency	7
Guardrails: Policy, Security, Compliance, and SLA-Bound Autonomy	9
The Role of Lakehouses	11
Hybrid, Multi-Source Reality Is the Norm for AI Initiatives	12
Cost and Performance Optimization Is a Critical Requirement for AI ROI	15
Ownership, Funding, and Why It Changes the Data Conversation	15
What This Means for Public Sector Leaders	17
What “Good” Looks Like and the 90-Day Blueprint	18
Methodology	19



Executive Summary

The Outcome Problem, Not the AI Problem

AI in the public sector is no longer experimental. Governments and agencies are investing heavily in real-time safeguarding and vulnerability detection, fraud prevention across benefits and taxation systems, faster and more responsive citizen services, and cross-agency decision-making and coordination.

However, not all organizations have been able to turn agentic AI projects into successful mission outcomes. Successful agentic AI initiatives require access to live data, comprehensive business context, and compliance with governance and regulatory guardrails, all of which can be challenging for many public sector organizations to provide.

This paper summarizes findings conducted for us by Arlington Research to shed light on these and other requirements, that warrant a fundamental rethink in enterprise data management strategies and architecture in order to ensure agentic AI success.

Live Data Gap

73%

73% of public sector organizations say AI data must be real-time or within-the-minute to be trustworthy — yet traditional data architectures, designed for historical analytics, cannot support this sub-minute latency, creating a critical bottleneck.

This is notably higher than the global average of 66%, reflecting the real-time demands of safeguarding, fraud detection, and citizen service delivery.

Relevance Gap

61%

61% of public sector organizations struggle to identify trustworthy data or prepare and integrate it for AI initiatives.

Even when data exists, determining which data is authoritative across fragmented agency systems remains a major challenge.

Complexity at Scale

479
DATA SOURCES

Enterprise AI initiatives in the public sector draw on an average of 479 data sources, with over one in five organizations accessing more than 1,000.

This distributed landscape increases the difficulty of ensuring consistency, governance, and performance across agencies.

Guardrail Gap

72%

72% of respondents in the public sector struggle with AI data security and access controls — the highest of any sector surveyed and 5 points above the global average of 67%.

As AI systems interact with sensitive citizen data, consistent policy enforcement becomes critical.

Performance Gap

68%

Even among organizations using data catalogs and lakehouses, 68% of public sector leaders report difficulty optimizing performance for AI workloads — significantly above the global average of 60%.

Performance and cost efficiency are critical barriers to scaling AI sustainably in resource-constrained environments.

KEY FINDINGS

Survey responses from **117 enterprise leaders in the public sector** reveal several gaps between what trustworthy AI requires and what most organizations can currently deliver.

Where AI Is Expected to Deliver Value in the Public Sector

AI is increasingly embedded into **mission-critical operational workflows**, not just analytics. Priority outcomes include safeguarding vulnerable individuals and communities in real time, detecting fraud across benefits and taxation systems, delivering faster and more responsive citizen services, and enabling cross-agency decision-making and coordination. **These are decision-making and action-oriented use cases in which AI must operate in the moment, within live public service environments — often under time-sensitive, high-stakes conditions.**

Introduction

Agentic AI is moving enterprise AI from answering to doing. Agents perceive conditions, decide what to do next, and then execute actions, often by calling tools and autonomously interacting with systems of record. This shift raises the bar for data management, as the quality and relevance of data used by agents directly impacts the trustworthiness and effectiveness of those agents. More specifically, agents break trust when they (a) lack **live situational awareness**, (b) act on incomplete or wrong data due to missing business context and inconsistent semantics, or (c) operate outside **compliance and business-SLA guardrails**. [1]

Independent research reinforces the same pattern: many organizations are struggling to translate “AI enthusiasm” into durable production outcomes. Gartner forecasts that **over 40% of agentic AI projects will be canceled by the end of 2027**, citing escalating costs, unclear business value, or inadequate risk — three constraints that are frequently downstream of data readiness, governance, and integration complexity. [2]

McKinsey, likewise reports broad AI adoption but limited scaling; in its 2025 survey results, only 7% of respondents said AI had been fully scaled across their organizations. [3]


In the Public Sector, this translates to an **AI Reality Check**: safeguarding initiatives are constrained by data silos across agencies, fraud detection programs lack real-time cross-departmental visibility, and compliance automation is blocked by fragmented governance frameworks.

In a survey commissioned by Denodo, Arlington Research conducted primary research into AI adoption patterns, which points to why agentic AI programs often stall. The operational bar is high (freshness expectations skew towards real-time), the “right data” problem is not solved by traditional data management solutions (relevance and trust are contextual), and implementing guardrails across distributed sources remains difficult, especially for organizations still planning or piloting AI initiatives.

This report structures the evidence around the three requirements most predictive of agentic AI success:

- **Live data:** agents can only act safely when they have live situation awareness, informed by data that mirrors real-time reality as closely as possible. [1]
- **The right data:** agents need source-of-truth clarity based on consistent semantics that identifies the data most relevant to the situation at hand. [1]
- **Guardrails:** agents must operate under security, privacy, and other regulatory compliance constraints, as well as internal business SLA constraints, consistently across all tools and sources. [1]

To overcome these barriers, this report also explores how to optimize AI cost and performance across distributed, hybrid data environments, and examines how distributed AI ownership impacts public sector data strategy. Finally, it outlines exactly what this means for public sector leaders, providing a clear 90-day blueprint to close the AI trust gap and turn stalled pilots into measurable mission outcomes.



Live Data for Operational Awareness



SURVEY INSIGHT

73% of public sector organizations say AI data must be real-time or within-the-minute to be trustworthy.

This demand is notably higher in the public sector than globally — 73% vs. 66% — reflecting the zero-latency requirements of safeguarding decisions, fraud detection, and emergency citizen services.

Public sector AI operates in uniquely high-stakes environments. Safeguarding systems must identify vulnerable individuals across multiple disconnected agency records in real time.

How Up to Date Does Data Need to Be for Trustworthy AI Results? (Public Sector Data)




Figure #1

Fraud detection must flag suspicious claims before payments are made. Citizen services must respond to live eligibility conditions, not last week's batch extract. In all these cases, stale data is not just inconvenient — it can have serious consequences for individuals and communities.

Most public sector data architectures were designed for analytics and reporting, where data is aggregated and prepared in advance. Cross-agency integration typically relies on batch processing and scheduled pipelines, introducing latency precisely where AI systems need live visibility most.

As AI systems move into operational workflows, access to live data becomes a prerequisite for safe decision-making. Organizations already running AI in production are significantly more likely to require real-time or near-real-time data access than those still planning initiatives.

For public sector AI embedded in safeguarding, fraud, and service delivery workflows, real-time awareness is not optional — it is essential for responsible and effective decision-making. Data architectures built primarily for reporting cannot support AI that must sense and act on live citizen conditions.

Data Freshness Requirements Differ by AI Maturity

Organizations already running AI in production are significantly more likely to require real-time or near-real-time data than those still planning initiatives. This suggests that the need for live operational data becomes clearer once AI systems move from experimentation into real-world workflows.

This pattern is particularly significant in the public sector, where only 56% of organizations currently have AI in production — compared to the global average of 67% — indicating that a large share of the sector has yet to encounter the full force of live data demands. As more public sector organizations move to production, real-time data requirements will become an even more acute constraint.

Many enterprises have historically separated operational systems from analytics platforms for good reason. Traditional data warehouses are widely considered to be designed for reporting and analysis rather than operational control loops, emphasizing curated historical insight rather than real-time decision-making. [4]

Agentic AI blurs this boundary. When AI systems must sense and act on changing business conditions, the most current and authoritative data often remains in operational applications and SaaS platforms rather than in analytics stores. Once data is copied from its original source, it immediately begins to lose its “liveness.”

Even modern analytics architectures can reintroduce staleness through pipeline structures. Coordinating lakehouse and warehouse environments often involves multiple processing stages, and additional extract, transform, and load (ETL) layers can introduce latency and data quality risks.

As a result, organizations deploying AI in operational environments often find that providing **live access to distributed operational data** becomes a core architectural requirement.

Providing live data for AI systems therefore involves more than reducing ingestion latency. In practice it often requires:

- On-demand access to operational data sources without always copying the data into a central store
- Performance optimization mechanisms such as caching or pushdown to support high-volume queries
- Governance controls to ensure real-time access does not bypass security or policy enforcement

At the same time, the broader AI ecosystem is evolving toward architectures where models retrieve data and invoke tools dynamically at runtime. For example, Anthropic’s Model Context Protocol (MCP) standardizes how AI systems connect to external tools and data sources, acting as a common interface between models and enterprise systems. [6]

This shift further increases the importance of governed, real-time access to enterprise data across distributed environments.

The Right Data: Relevance, Context, and Semantic Consistency

Survey responses indicate that identifying and preparing the “right data” is one of the most significant barriers to trustworthy AI in production. In particular, respondents frequently cite the challenge of identifying relevant and trustworthy data and preparing or integrating it for AI systems.

What Do Your AI Teams Consider as the Top Data-Related Challenges for Trustworthy AI in Production? (Respondents were able to choose their top two challenges.)




Figure #2

In the public sector, cost and performance optimization ranks as the top challenge (40%) — significantly above the global average of 33%. This reflects the resource constraints facing public sector organizations, where AI investments must demonstrate clear efficiency gains and tight budget accountability. Requesting access to data is also a notably higher concern (21% vs. 14% globally), reflecting the siloed, permission-heavy data environments common across government agencies.

While access to live data is essential, it does not guarantee that AI systems are acting on the correct data. Even accurate data can become misleading when business definitions differ across systems or when the data lacks the appropriate business context for a given decision.

In the public sector, this manifests as inconsistent definitions of citizen, case, risk, or eligibility across central government, local authorities, healthcare, and social care systems. When definitions vary across agencies, automated decisions can quickly become unreliable — or unfair.

In these situations, AI systems may rely on incomplete or inconsistent information, increasing the risk of incorrect actions or recommendations.

Differences in how terms such as “customer,” “account,” or “risk” are defined across systems can lead to conflicting interpretations of the same underlying data.

Traditional data catalogs and governance tools typically capture static metadata — business terms, ownership, and policy definitions defined in advance by data stewards. However, operational AI systems often encounter dynamic situations where context must be interpreted at runtime, making static definitions alone insufficient.

As a result, organizations must increasingly address not only **data availability**, but also **data relevance and semantic consistency** across systems.

THE “RIGHT DATA” CHALLENGE

The survey results reinforce that identifying and preparing the “right data” remains a major challenge for organizations deploying AI systems. In particular, respondents frequently cite difficulties identifying relevant and trustworthy data sources and preparing or integrating that data for AI use cases.

The “right data” challenge is also reflected in how organizations say they currently ensure relevance and trustworthiness. Many rely on a combination of BI and analytics tools, governance platforms, and lakehouse environments. However, most respondents selected multiple approaches.

This suggests that ensuring trustworthy data for AI remains a **multi-layered problem** for many organizations, often requiring multiple tools and processes to identify, validate, and prepare the most relevant data for AI systems.

How Do You Ensure AI Uses the Most Relevant and Trustworthy Data?

(Respondents were able to choose more than one.)




Figure #3

Why Semantic Consistency Becomes a Gating Factor for Agents

Agentic AI systems rely on consistent business meaning across data sources. When definitions vary across systems, automated decisions can quickly become unreliable.

Semantic ambiguity becomes particularly problematic when agents execute actions rather than simply generate responses:

- If terms such as “customer,” “account,” “active,” “revenue,” or “risk” are defined differently across systems, an agent may select incorrect records, invoke the wrong tools, or trigger unintended workflow actions.
- If an agent cannot determine which source is authoritative for a given task, it cannot reliably justify or automate decisions.

These challenges explain why many organizations are investing in semantic-layer approaches that centralize business definitions and ensure consistent interpretation of data across distributed systems.

Platforms such as Denodo position a **universal semantic layer** as a way to harmonize business meaning across complex data environments while preserving access to live operational sources. [7] This aligns with a broader industry shift: many data platforms now emphasize semantic layers as a foundation for governed self-service analytics and AI systems that rely on consistent context at runtime. [8]

“Right Data” as a Retrieval and Provenance Problem

Modern AI architectures increasingly address the “right data” challenge through retrieval-based approaches. In foundational research on retrieval-augmented generation (RAG), the authors note that parametric models alone struggle to maintain up-to-date knowledge and reliable provenance. Instead, they propose retrieving information from external sources at runtime to ground AI outputs in authoritative data. [9]

This principle becomes even more important in enterprise environments. AI agents must not only retrieve live data from original sources, but also determine which sources are the most authoritative for the given task. The most relevant or trusted source, out of potentially hundreds or thousands in a typical enterprise, often may only be determined at runtime.

These challenges reinforce why many organizations are exploring metadata-driven approaches that help AI systems identify the most relevant and authoritative data sources dynamically.

Denodo describes its approach to this problem as **Query RAG**, a metadata-driven framework that uses semantic context to determine which live data sources should be accessed at runtime. [10]

Guardrails: Policy, Security, Compliance, and SLA-Bound Autonomy



SURVEY INSIGHT

72% of respondents in the public sector struggle with AI data security and access controls — the highest guardrail gap of any sector surveyed and 5 percentage points above the global average of 67%.


This reinforces a consistent pattern in early-stage AI programs: before production processes and governance stabilize, organizations experience guardrail friction most acutely — particularly around consistent access controls, metadata standardization, and governance across distributed systems.

AI in the public sector must operate within strict legal, ethical, and societal constraints: data protection legislation, transparency and explainability requirements, fairness in automated decision-making, and accountability for outcomes that affect citizens’ rights and entitlements. Guardrails are the defining difference between a “helpful assistant” and a “trusted agent.”

Agents break trust when they operate outside compliance and business-SLA guardrails, and fragmented governance and security controls across sources are a common cause of failure. [1]

Survey results indicate that governance and guardrail challenges remain among the most difficult obstacles when deploying AI systems into production environments. This risk is increasingly codified in security and risk frameworks for AI systems:

- The Open Security Foundation's OWASP Top 10 for LLM Applications explicitly lists "Excessive Agency" and "Insecure Plugin Design," reflecting the reality that tool-connected agents can take unintended actions or expand the attack surface if permissions and control paths are weak. [11]
- Industry leaders such as Amazon Web Services (AWS) similarly identify excessive agency as a key threat typically introduced through agentic architectures, and AWS recommends least-privilege and permission boundaries for agentic workflows. [12]
- The National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF 1.0) is designed to help organizations manage AI risks and promote trustworthy AI through lifecycle practices, and NIST has also published a generative AI profile aligned to AI RMF practices. [13]



The survey data indicates that organizations consider guardrail-adjacent problems among the **most difficult** to address — especially security/access control consistency and metadata standardization across systems:

How Difficult or Easy Is It to Address the Following Data-Related Issues? (Public Sector Data)




Figure #4

The public sector stands out on security and access control complexity: 16% describe it as “very difficult” — nearly double the global average of 9% — reflecting the sensitivity of citizen data and the multi-agency governance challenge.

The “guardrail gap” is particularly visible early, before production systems and governance processes stabilize.

When comparing organizations planning AI with those already in production, respondents at the planning stage report greater difficulty across several guardrail-related capabilities, particularly around access control consistency, metadata standardization, and data governance.

What Guardrails Require in Practice

For agentic AI, guardrails extend beyond model alignment or safety prompts. In enterprise environments, they typically require several operational capabilities:

- **Consistent access controls and policy enforcement across sources**, so governance cannot be bypassed in downstream systems.
- **Auditability and traceability** of what data agents accessed and what actions they executed.
- Least-privilege tool access, ensuring agents can only invoke permitted actions under defined conditions. [12]
- **Semantic clarity**, so policies map consistently to business meaning (for example what constitutes “sensitive,” “regulated,” “PII,” or “citizen” data). [11]

These requirements explain why many organizations are exploring architectures that centralize governance and policy enforcement across distributed data environments. Platforms such as Denodo position logical data management as one approach to enabling these guardrails consistently across heterogeneous systems. [1]

The Role of Lakehouses, and the Criticality of Logical Data Management



SURVEY INSIGHT

Even among organizations using data catalogs and lakehouse environments, 68% of public sector leaders report difficulty optimizing performance for AI workloads.

This signals that modernized central platforms are necessary infrastructure, but not sufficient on their own for operational, agentic AI — where cost, performance, and real-time access patterns increasingly determine whether pilots can scale sustainably.

Modern lakehouses and cloud platforms provide scalable compute and storage foundations for analytics and AI workloads. Vendors such as Databricks and Snowflake define lakehouse architectures as unifying aspects of data lakes and data warehouses to reduce silos and support analytics/AI workloads.

However, emerging evidence suggests that traditional data modernization alone does not fully address the operational requirements of agentic AI.

Gartner notes that many organizations still lack AI-ready data management practices and predicts that a substantial share of AI projects will be abandoned when those foundations are missing, despite significant investments in data modernization. [14]

Similarly, even leading data platform vendors acknowledge that pipeline complexity can introduce new operational constraints. For example, IBM notes that long processing times and additional ETL layers can contribute to data staleness and data quality risks — an architectural tension when systems must support both accurate analytics and real-time operational decisions. [5]

These signals suggest that while lakehouses remain essential infrastructure for analytics and AI development, additional architectural capabilities are often required when AI systems must operate across distributed enterprise environments.

Agentic AI systems typically require:

- Access to live operational data
- Consistent business semantics across distributed systems
- Centralized governance and policy enforcement
- Performance optimization and cost controls for AI workloads
- Access across hybrid, multi-cloud, SaaS, on-premises, and third-party data environments

Hybrid, Multi-Source Reality Is the Norm for AI Initiatives

In the global survey of enterprise AI leaders commissioned by Denodo, public sector organizations report that AI initiatives often draw from large numbers of original data sources, not a small number of consolidated systems.

The same survey suggests that organizations use multiple data access patterns simultaneously, including embedded data product teams, data marketplaces, central request processes, and direct-to-source access.

How Many Original Data Sources Feed AI Initiatives?




Figure #5

This diversity of access paths, combined with the sheer volume of original data sources, significantly increases the difficulty of applying consistent guardrails and semantics across enterprise data environments.

How Do Your AI Builders Access the Data They Need? (Respondents were able to choose more than one.)




Figure #6

Logical Data Management as “Supplementation,” Not Replacement

Lakehouses continue to serve a critical role in most enterprises for analytics, reporting, and many AI development workloads, while logical data management complements these environments when AI systems require live access to distributed operational data.

A useful way to interpret this architectural implication is as follows:

- **Lakehouses and warehouses** provide essential storage/compute and are well suited for many analytics and AI tasks, but they are often not designed to provide governed, real-time, business-ready data to every operational consumer across every system including most AI agents embedded in business workflows. [15]
- **Logical data management** is widely defined as providing a virtual layer that enables access to and integration of data across disparate sources “without requiring replication or movement,” creating unified representations and accelerating integration for analytics and AI. [16]
- Platforms such as Denodo position logical data management as a way to centralize data practices — including discovery, access management, security, integration, monitoring, and sharing — through logical representations rather than duplicating governance and integration logic across individual systems. [16]

This “supplement, don’t replace” positioning is also consistent with external analyst framing that modernization often involves layering capabilities rather than replacing existing platforms. [18]

Cost and Performance Optimization Is a Critical Requirement for AI ROI

Agentic AI can be cost-sensitive because agents may call multiple tools, perform repeated retrieval, iterate, and require logging/auditability for each task. This raises the bar for data performance and cost controls, as AI-intensive workloads can exceed the data consumption patterns that preceded it. [1] Gartner similarly cites escalating costs as a key driver for agentic AI project cancellations. [2]

Industry analysis increasingly highlights cost and performance optimization as critical factors in determining whether AI initiatives can scale economically. In the public sector, this pressure is especially acute: budget constraints and public scrutiny mean AI initiatives must demonstrate clear mission value to justify investment. The data reflects this — cost and performance optimization is the top data challenge for public sector AI teams (40% vs. 33% globally).

For example, an independent study commissioned by Denodo and conducted by analyst firm Vektor8 compared lakehouse-only architectures with environments supplemented by a logical data management layer, reporting improvements in ROI, cost avoidance, and time-to-insight when organizations introduced a unified logical access layer. [17]

Such findings illustrate how organizations evaluate architectural approaches not only in terms of technical capability, but also in terms of economic sustainability for AI workloads.

Ownership, Funding, and Why It Changes the Data Conversation



SURVEY INSIGHT

Survey results indicate that responsibility for developing and maintaining AI initiatives is often distributed across multiple functions, rather than owned by a single central team.

This distributed ownership changes the data conversation: the data foundation must support not only traditional analytics and governance workflows, but also fast-moving AI delivery teams and business sponsors who need timely, governed access to operational data within real workflows.

A defining operational reality of enterprise AI is that ownership is shared beyond traditional data teams. While this sometimes occurs, AI is more frequently driven by central AI teams within technology leadership or by distributed business initiatives aligned to specific operational outcomes (safeguarding, fraud reduction, service transformation, productivity). These AI teams often operate under different mandates than traditional data programs, specifically to bring innovation to public services faster and deliver measurable mission outcomes.

This creates a practical requirement: the data foundation must work for fast-moving AI teams and business sponsors, not only for traditional analytics and data governance programs often prioritized by data teams. The survey evidence below shows how organizations describe responsibility for developing and maintaining AI initiatives.

The survey indicates that responsibility for AI initiatives is often distributed across multiple functions rather than centralized in a single team.

While only 56% of public sector organizations currently have AI in production (vs. 67% globally), a comparable share report sponsorship distributed across functions — suggesting governance frameworks are still being built at a critical moment.

Who Is Responsible for Developing and Maintaining AI Initiatives? (Respondents were able to choose more than one.)




Figure #7

These ownership patterns help explain why agentic AI can expose limitations in existing data access and governance models. When AI investment is owned closer to mission outcomes, there is less tolerance for slow access cycles, prolonged data onboarding, or unclear governance pathways.

This aligns with Gartner’s warning that many projects will be canceled due to cost, unclear value, and inadequate risk controls, each of which can be amplified when organizations struggle to deliver timely, governed access to the data required for AI systems. [2]

Who Is the Executive Sponsor for AI Initiatives?

(Respondents were able to choose more than one.)




Figure #8

What This Means for Public Sector Leaders

Agentic AI success criteria tend to converge on a small set of architectural outcomes:

- **Business & Operations Leaders (Policy, Services, Safeguarding):** Cannot rely on AI for critical safeguarding or citizen service decisions due to data latency and fragmentation. Struggle to demonstrate measurable mission outcomes. Face increasing scrutiny around AI fairness, transparency, and accountability.
- **AI & Innovation Teams:** Lack access to live, cross-agency data. Spend significant time reconciling inconsistent definitions of citizen, case, eligibility, and risk. Struggle to operationalize AI within real-world service delivery and safeguarding workflows.
- **Data & Technology Leaders:** Face the challenge of managing an average of 479 distributed data sources across agencies. Must address the highest governance complexity of any sector surveyed: 72% cite security and access controls as difficult, with 16% calling it “very difficult.” Performance and cost pressures are growing: 68% report difficulty optimizing AI workload performance.

What “Good” Looks Like and the 90-Day Blueprint


Closing the AI trust gap requires a shift from fragmented data silos and batch aggregation to real-time, governed data collaboration across agencies. A modern AI-ready public sector data foundation must provide live cross-agency data access, shared context with semantic consistency, and governed access aligned to legal and ethical requirements.

Leading public sector organizations are adopting a 90-day blueprint:

- **Days 1 – 30 (Define Outcomes):** Align on mission KPIs (safeguarding improvement, fraud reduction rate, service speed) and map required data sources and governance gaps across agencies.
- **Days 31 – 60 (Build Data Products):** Create governed, reusable data products aligned to use cases to enable real-time cross-agency access and semantic consistency — defining shared entities like citizen, case, eligibility, and risk.
- **Days 61 – 90 (Operationalize):** Integrate AI into live service and safeguarding workflows, enable real-time decisioning, and monitor compliance, fairness, and mission outcomes.

Closing the AI Trust Gap in the Public Sector

The future of AI in government is not about more models or more pilots. It is about trusted data, real-time cross-agency access, and governed decision-making that is explainable, fair, and accountable. Organizations that close this gap will improve safeguarding outcomes and public safety, deliver faster and more effective citizen services, and build public trust through transparency and compliance. Those that don't will remain constrained by silos, unable to scale AI, and unable to realize its full public value.




Methodology

Arlington Research conducted this online survey in June 2025. Respondents consisted of 850 executives and business decision-makers in organizations with **1000+ employees**. The data presented in this specific report is drawn exclusively from a targeted subset of **117 enterprise leaders within the Public Sector**.


The survey was conducted across the Americas, Europe / Middle East and Asia-Pacific regions, in countries with significant IT presence:

- 75 each in the US, the UK, France, Germany, Japan, and Australia
- 50 each in Canada, Brazil, Italy, Spain, China, Singapore, South Korea, and the UAE



Citations

- [1]. [Delivering Trusted Agentic AI](#)
- [2]. [Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027](#)
- [3]. [AI at work but not at scale](#)
- [4]. [What is a Data Warehouse? | Microsoft Azure](#)
- [5]. [What is a data warehouse? | IBM](#)
- [6]. [Model Context Protocol \(MCP\) - Anthropic](#)
- [7]. [Semantic Layer: Your Key to AI-Ready Data Insights | Denodo](#)
- [8]. [Universal Semantic Layer | Denodo](#)
- [9]. [Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks | UCL NLP](#)
- [10]. [Query RAG | Denodo](#)
- [11]. [OWASP Top 10 for Large Language Model Applications | OWASP Foundation](#)
- [12]. [Excessive agency - Generative AI Lens](#)
- [13]. [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\) | NIST](#)
- [14]. [Lack of AI-Ready Data Puts AI Projects at Risk](#)
- [15]. [Supercharge Your Lakehouse with AI-Ready Data](#)
- [16]. [What is Logical Data Management? | Denodo](#)
- [17]. [Denodo Users Gain 345% ROI, 3–4x Faster Time-to-Insight, and Other Advantages, When Deployed alongside a Data Lakehouse | Denodo](#)
- [18]. [Gartner's Data Leaders Checklist for the Agentic AI Era | Provided by Denodo](#)



Visit denodo.com | Email info@denodo.com | Discover community.denodo.com