QualityAI



AGENTIC AI IN DATA GOVERNANCE

Architecting the autonomous future of trusted data

GOVERNANCE THAT FINALLY DOES SOMETHING

Data governance has always been a discipline built on human judgment, manual process, and controls that kick in after something's already gone wrong. Agentic AI changes that.

These are systems that can plan, reason across multiple steps, and act on their own – which means, for the first time, governance teams have technology that doesn't just surface problems but does something about them. In real time. Without someone having to pick up the ticket.

This POV argues for a different model of data governance: one where AI agents are active participants – enforcing policies, catching anomalies, enriching metadata, fixing quality issues, and keeping the bridge open between human stewards and automated pipelines. The goal isn't to take humans out of the loop. It's to give governance teams leverage over a data estate that no human team, however talented, could watch alone.

Organizations that deploy agentic AI in their governance programs report up to 60% reductions in manual data quality remediation effort and a 3–5x increase in policy coverage across their data estate (Gartner, 2025).



 | Copyright © 2026 QualityAI. All rights reserved.

THE CRACKS IN HOW WE GOVERN DATA

A scale problem nobody designed for

The volume of data modern organizations generate looks nothing like what traditional governance frameworks were built for. A midsized financial institution might operate across tens of thousands of data assets - cloud, on-premise, SaaS, everything in between. A global retailer may process billions of customer events a day. Against that reality, governance models built on spreadsheet inventories, quarterly quality reports, and manual stewardship reviews are structurally overmatched. They were never designed for this.

Governance that only shows up after the damage

Most governance is reactive: problems surface after they've already caused harm; a regulatory breach, a bad business decision, a failed audit. The tools most governance teams rely on - catalogues, policy documents, quality dashboards are mostly observation tools. They tell you what happened. They don't stop it happening, and they don't fix it. The gap between knowing and acting is where governance value disappears.

Stewards drowning in the wrong work

Data stewards are skilled people with finite time. In most organizations, the ratio of governed assets to dedicated stewards is bad, really bad. Stewards end up spending most of their time on low-value, repetitive work: answering metadata queries, chasing data owners for updates, manually profiling datasets. That leaves almost nothing for the work that actually matters - designing governance frameworks, resolving cross-domain problems, advising on data strategy.

Regulation just made automation non-negotiable

Regulatory pressure on data keeps increasing. GDPR and its equivalents require accountability and traceability of records, consent management, and breach notification. BCBS 239 demands accurate risk data aggregation. The EU AI Act adds governance requirements for AI training data. Keeping pace with all of this through manual processes is no longer a choice. Automation has moved from nice-to-have to regulatory necessity.

Copyright © 2026 QualityAI. All rights reserved.

WHAT WE MEAN WHEN WE SAY "AGENTIC"

Beyond the chatbot

Agentic AI is not just a better chatbot. These systems go beyond single-turn question answering to exhibit genuinely goal-directed behavior. An agentic system can take a high-level objective, break it into sub-tasks, pick and use the right tools, assess intermediate results, adjust its approach, and keep going until the job is done – without needing a human to approve every step.

Three things set agentic AI apart from conventional AI tools:

Tool use: The ability to use external tools – databases, APIs, code executors, communication platforms – to take action in the world, not just generate text about it.

Multi-step reasoning: The ability to plan sequences of actions and adapt as results come in, rather than treating each input as a fresh start.

Memory and context: The ability to hold state across a workflow – remembering what it found in step two when it gets to step seven.



 | Copyright © 2026 QualityAI. All rights reserved.

INSIDE AN AGENTIC GOVERNANCE SYSTEM

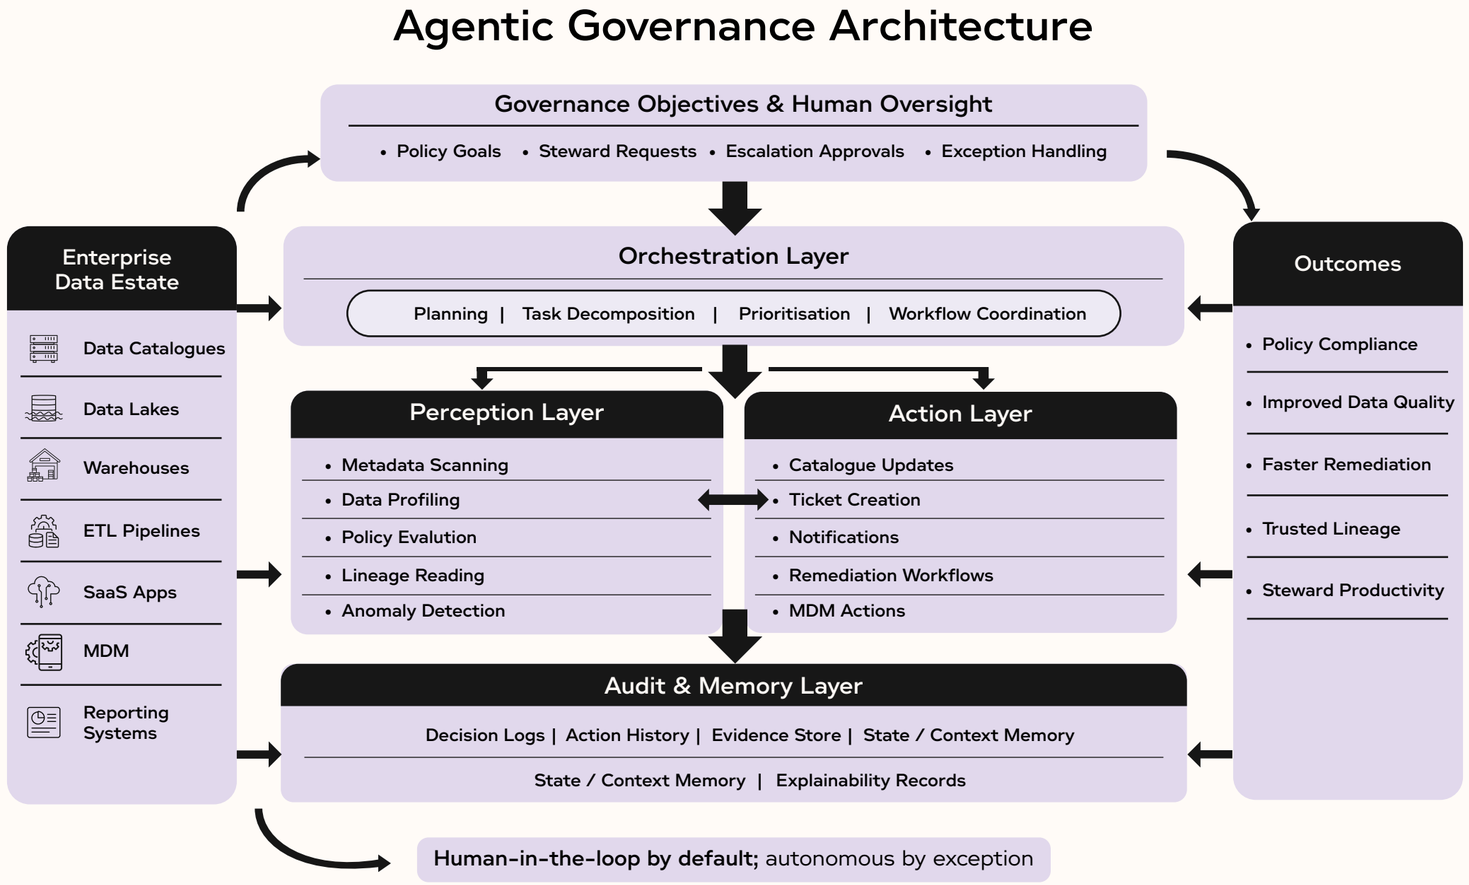


Figure 1: Agentic governance architecture showing how orchestration, perception, action, and audit layers work together under human oversight.

Copyright © 2026 QualityAI. All rights reserved.

A governance-focused agentic system typically has four layers:

Orchestration layer: Takes governance objectives from humans or orchestration systems and breaks them into executable plans.

Perception layer: Profiles datasets, assesses quality, scans for policy violations, and reads metadata from catalogues and lineage systems.

Action layer: Connects to data catalogues, MDM platforms, ticketing systems, notification services, and remediation pipelines to do things.

Audit and memory layer: Logs every action, decision, and outcome. Non-negotiable in regulated industries.

The critical design principle here is "human in the loop by default, autonomous by exception." Agents act independently on low-risk, well-defined tasks. They escalate to humans when confidence is low, stakes are high, or policy is ambiguous.



 | Copyright © 2026 QualityAI. All rights reserved.

WHAT THIS LOOKS LIKE IN PRACTICE

Agentic AI Data Governance Use Case Map

DATA QUALITY & REMEDIATION	METADATA & LINEAGE	COMPLIANCE, PRIVACY & AI GOVERNACE	ENTERPRISE ORCHESTARTION & STEWARDSHIP
3.1 Autonomus Data Quality Monitoring & Remediation Monitor → Quarantine → Notify → Remediate	3.2 Intelligent Metadata Enrichment	3.3 Policy Compliance Monitoring	3.9 Steward Copilot and knowledge Management
3.5 Master Data Conflict Resolution	3.4 Automated Data Lineage Discovery & Maintenance	3.6 Regulatory reporting Preparation	3.10 Cross Domain Governance Orchestartion
		3.7 Privacy and Consent Management	3.11 Governing Unstructured Data
		3.8 AI Model Governance	

Figure 2: Use case map grouping eleven agentic governance applications by data quality, metadata, compliance, and enterprise co-ordination objectives.

 | Copyright © 2026 QualityAI. All rights reserved.

1. Autonomous data quality monitoring and remediation

An agentic system can watch data pipelines continuously against defined quality rules, completeness thresholds, referential integrity, statistical distribution bounds and move immediately when something breaks. Rather than logging an alert that sits in a queue for days, the agent quarantines affected records, notifies the data owner, triggers a root cause investigation, proposes a fix, and where it's confident enough, executes the remediation and logs the action. A quality issue lifecycle that used to take days can compress to minutes.

2. Intelligent metadata enrichment

Metadata completeness is the chronic weak point of enterprise data catalogues. Agents can tackle this by automatically profiling new datasets as they arrive, inferring data types, detecting PII through pattern recognition, suggesting business glossary mappings, identifying probable data owners from lineage and access patterns, and drafting descriptions for human review. Work that used to take weeks of manual stewardship can happen in hours.

3. Policy compliance monitoring

Governance policies are only as good as their enforcement. Agentic systems can scan the data estate continuously for violations: datasets held past their retention date, sensitive data accessed outside approved roles, unclassified assets sitting in production environments. Findings get automatically escalated, ticketed, and tracked to resolution, with compliance rates visible in real time.

4. Automated data lineage discovery and maintenance

Lineage – knowing where data comes from, how it's been transformed, and where it ends up, is foundational to governance and notoriously hard to maintain by hand. Agentic systems can parse SQL scripts, ETL job logs, API call records, and notebook execution histories to build and update lineage graphs automatically. Impact analysis that previously took days of manual tracing becomes instant.



 | Copyright © 2026 QualityAI. All rights reserved.

5. Master data conflict resolution

In organizations running distributed systems, master data entities, customers, products, counterparties, frequently exist in conflicting forms across domains. Agents can monitor continuously for conflicts, apply configurable matching and survivorship rules to propose golden record resolutions, surface cases needing human judgment to stewards in a prioritised queue, and update downstream systems once a decision is made. MDM maintenance no longer requires an army.

6. Regulatory reporting preparation

Regulatory submissions require accurate, lineage, backed data assembled from multiple systems on tight deadlines. Agentic systems can automate the assembly, validation, and reconciliation steps, checking data against prior-period submissions for anomalies, flagging missing or stale source data, and generating audit trails that demonstrate provenance to regulators. Human reviewers focus on exceptions and signoff, not assembly.

7. Privacy and consent management

GDPR and equivalent regulations require organizations to honour subject access requests, right-to-erasure requests, and consent withdrawal within fixed timeframes. Agents can automate discovery of personal data across the estate in response to subject requests, co-ordinate erasure or restriction workflows across multiple systems, maintain consent audit logs, and generate statutory response documentation.

8. AI model governance

As AI and machine learning deployments grow, governing the training data underneath them becomes a genuine risk. Agentic governance systems can monitor training datasets for representational bias, data drift, and provenance issues, flagging when a model's training data no longer reflects production, or when a dataset contains records that should have been deleted under retention policy. The EU AI Act is making this mandatory.



 | Copyright © 2026 QualityAI. All rights reserved.

9. Steward copilot and knowledge management

Not all the value here comes from full automation. An agent working as a steward copilot accessible in plain English can make an individual steward significantly more productive. Ask it: 'which datasets contain customer date of birth and haven't been reviewed in six months?' or 'draft a data quality issue report for the orders table.' The agent retrieves, composes, and drafts. The steward reviews and decides.

10. Cross domain governance orchestration

Large organizations govern data across multiple domains such as finance, risk, customer, product, each with its own stewards, tools, and policies. Agentic orchestration can co-ordinate across all of them, propagating a policy change to every affected domain, running a multi domain impact assessment, managing a cross-domain quality remediation programme. Enterprise-wide governance without a proportional increase in central headcount.

11. Governing unstructured data

Unstructured data accounts for roughly 80% of enterprise data by volume yet it remains the most under-governed category. Traditional approaches which rely on periodic sampling, keyword search, and folder-level access controls doesn’t produce the continuous, auditable coverage that GDPR or the EU AI Act requires. Agentic AI changes this as agents can apply NLP-based classification, entity recognition, and sensitivity scoring across the entire unstructured corpus continuously, not on a sample. Organizations who are holding significant volumes of sensitive unstructured content, this shift from periodic sampling to continuous agentic coverage is one of the most consequential governance improvements available today.



Copyright © 2026 QualityAI. All rights reserved.

GETTING THERE: A FOUR-PHASE PATH

Implementation Roadmap for Agentic Data Governance

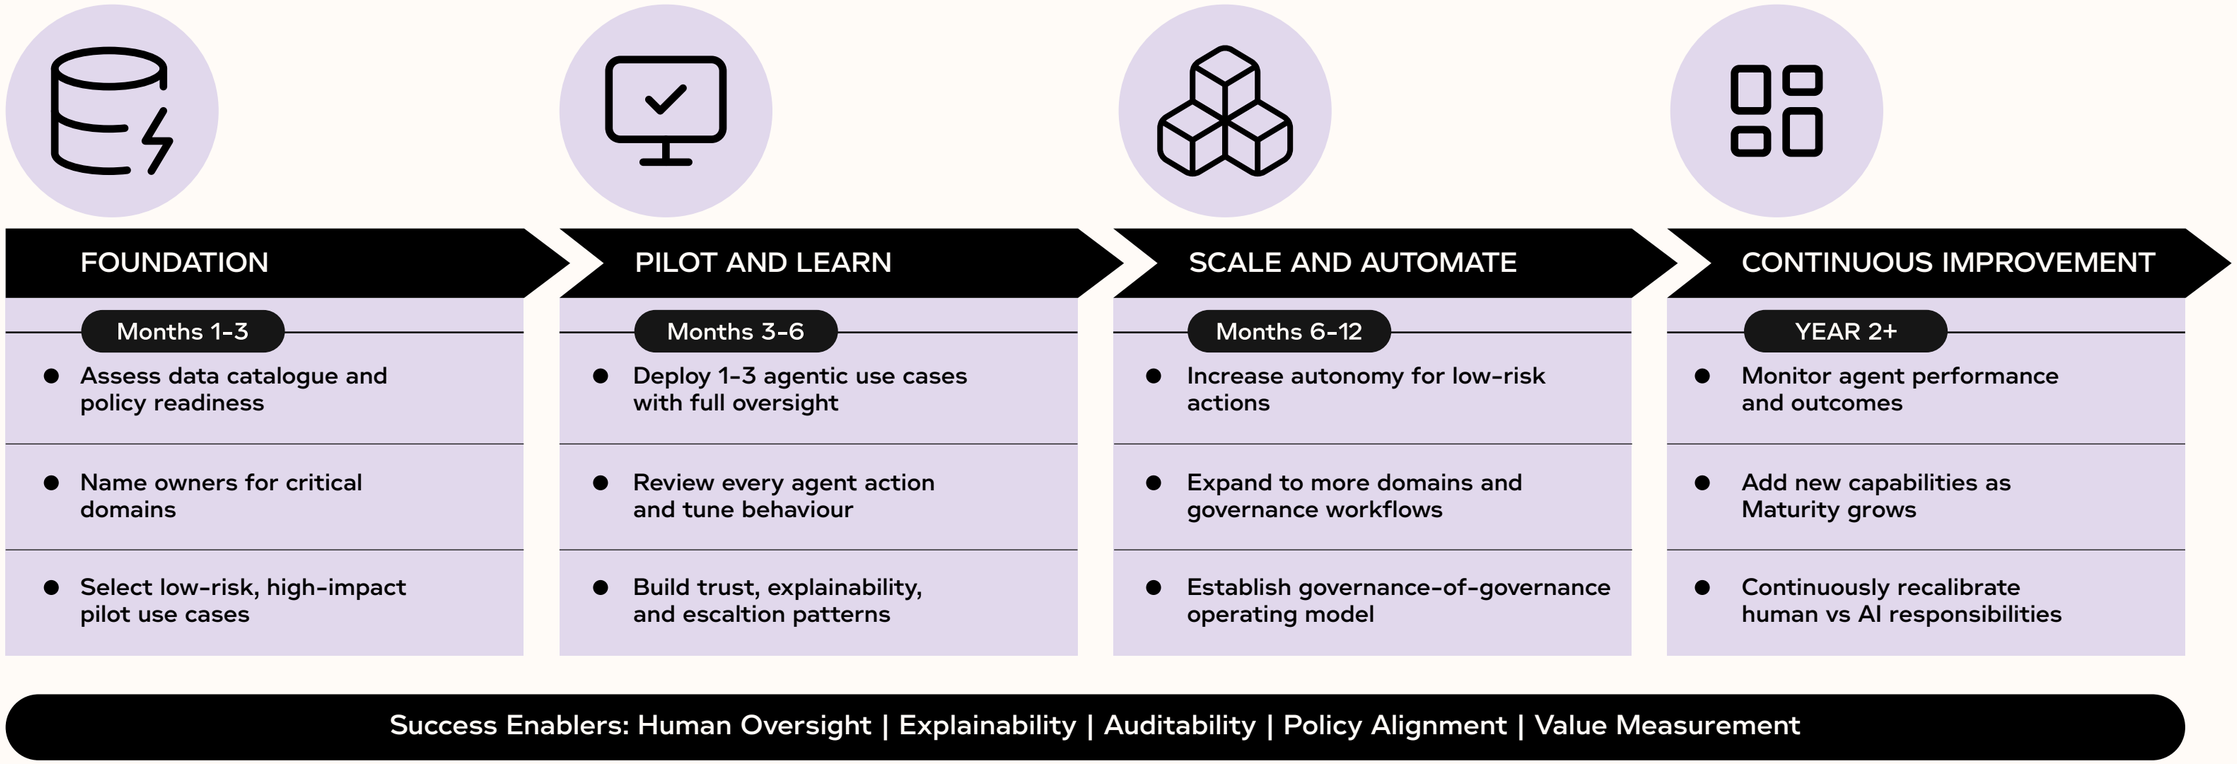


Figure 3: Implementation roadmap outlining a phased path from foundation and pilot to scaled automation and continuous improvement.

Copyright © 2026 QualityAI. All rights reserved.

GETTING THERE: A FOUR-PHASE PATH

Phase 1 - Foundation (months 1-3)

No agentic system can govern data it can't see. Before deploying agents, organizations need a sufficiently mature foundation: a functional data catalogue with reasonable coverage, a defined policy framework, and named data owners for critical domains. Assess readiness honestly and fix critical gaps first. For pilot use cases, picking something high-impact and low-risk - automated quality monitoring on a well-understood domain is usually the right starting point.

Phase 2 - Pilot and learn (months 3-6)

Deploy one to three use cases in a controlled environment with full human oversight. Every agent action should be reviewed before execution in this phase - the goal is calibration, not speed. Identify the edge cases the agent handles badly. Build trust with stewards and business stakeholders. Invest in explainability: agents should be able to say, in plain language, why they did what they did.

Phase 3 - Scale and automate (months 6-12)

As confidence grows, progressively increase autonomy for well-understood, low-risk actions. Expand to additional domains and use cases. Establish a governance-of-governance model: who oversees the agents, how behaviors get updated as policies change, how errors get escalated and fixed. Start measuring and reporting the value the program is delivering.

Phase 4 - Continuous improvement (year 2+)

Mature programs operate in a continuous improvement loop: agent performance is monitored, new capabilities are added as the technology develops, and the human-automation boundary is regularly recalibrated. Governance leaders become orchestrators of agents and human stewards, rather than managers of purely human teams.



 | Copyright © 2026 QualityAI. All rights reserved.

THE GUARDRAILS THAT MAKE THIS SAFE TO TRUST

Agentic AI sounds like the obvious way forward - but it's worth stepping back to understand the guardrails that need to be in place before it is.

The risk of automated error at scale

The same capability that makes agentic AI valuable makes it capable of propagating errors at scale. A misconfigured quality rule enforced by an agent across millions of records can cause more damage than the same rule applied manually to a sample. Robust testing, staged rollout, and mandatory human review for high-impact actions aren't optional safeguards.

Transparency and explainability

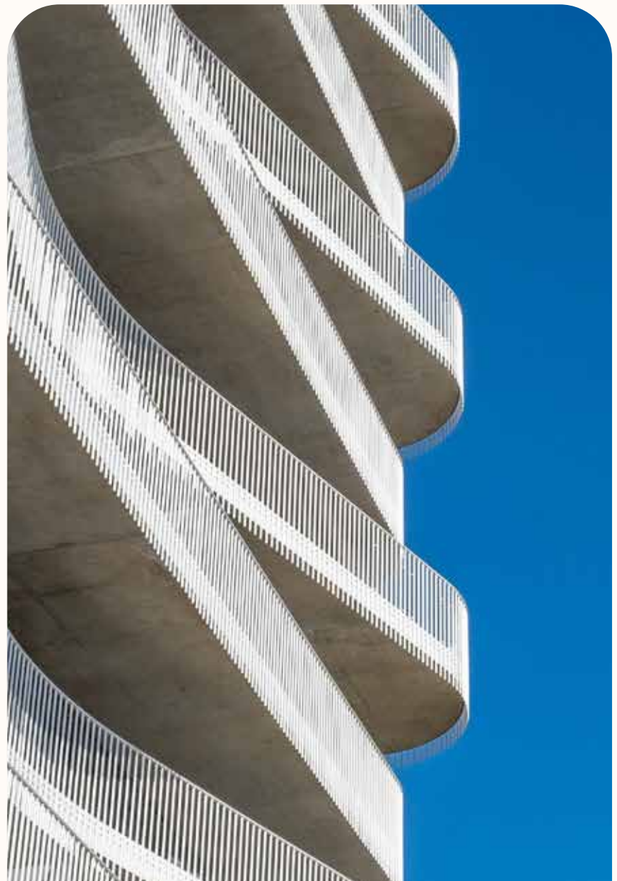
Regulated organizations must be able to explain their governance decisions to auditors and regulators. Every agent action needs an auditable, human-readable explanation: what the agent observed, what rule it applied, what it did, and what happened as a result. Black-box agents don't clear that bar.

Bias in governance decisions

Agentic systems configured on historical governance patterns may reproduce existing biases - applying stricter quality monitoring to some domains than others based on where past attention has focused. Governance leaders should periodically audit agent behavior for consistency and fairness across the estate.

Human oversight as a non-negotiable

The most important governance risk is the quiet erosion of human accountability. As agent autonomy increases, the temptation to remove human oversight entirely becomes real. Resist it. The governance operating model should define clear human accountability for every class of agent decision, with escalation paths that are tested and maintained. Agents are accountable to humans. They are not a substitute for them.



 | Copyright © 2026 QualityAI. All rights reserved.

Auditing decisions that unfold over many steps

Auditing a single-step AI decision is manageable: one input, one output, one rationale. Multi-step agentic workflows are a different problem. An agent pursuing a governance objective might execute ten or fifteen sequential actions, each defensible on its own, whose combined effect nobody necessarily sanctioned. That shifts where accountability sits – away from individual actions and onto the whole decision chain.

Every multi-step workflow therefore needs a structured, human-readable execution log: the initial trigger, the reasoning behind each step, every tool invoked and what it returned, any confidence thresholds or escalation rules that fired, and the outcome of human review. Crucially, that log must be repeatable – an auditor needs to reconstruct why the agent did what it did at any point in the chain, not just confirm what the final action was. Systems that record what happened without recording why will not hold up under regulatory scrutiny.

In the pilot phase, every checkpoint gets a sign-off. As confidence grows, routine checkpoints shift to exception review. But they stay visible and searchable. That part doesn't change.

The guiding ethical principle for agentic governance: automate the routine, amplify the human. Agents handle volume and speed. Humans retain judgment, accountability, and strategic direction.

MEASURING WHAT MATTERS

The following table presents a recommended set of KPIs for an agentic data governance programs:

KPI	Description	Target (year1)
Data quality issue MTTR	Mean time to remediate quality issues	< 4 hours
Policy coverage rate	% of data assets covered by active governance policy	85% +
Metadata completeness	% of catalogue assets with complete core metadata	>80%
Steward productivity	Governed assets per steward FTE	+50% vs baseline
Automated resolution rate	% of quality issues resolved without human action	40-60%
Regulatory breach incidents	Number of governance related regulatory incidents	YoY reduction
Agent accuracy rate	% of agent decisions confirmed correct on review	>95%
Escalation rate	% of agent tasks escalated to human review	<15%

Copyright © 2026 QualityAI. All rights reserved.

THE ROAD AHEAD

Early adopters are already deploying agents for quality monitoring, metadata enrichment, and compliance scanning. The results are real: material reductions in manual effort, step-change improvements in g overnance coverage. The organizations moving now are building capabilities that will become a genuine competitive advantage - faster, more reliable data, lower regulatory risk, and a data estate that business leaders can trust.

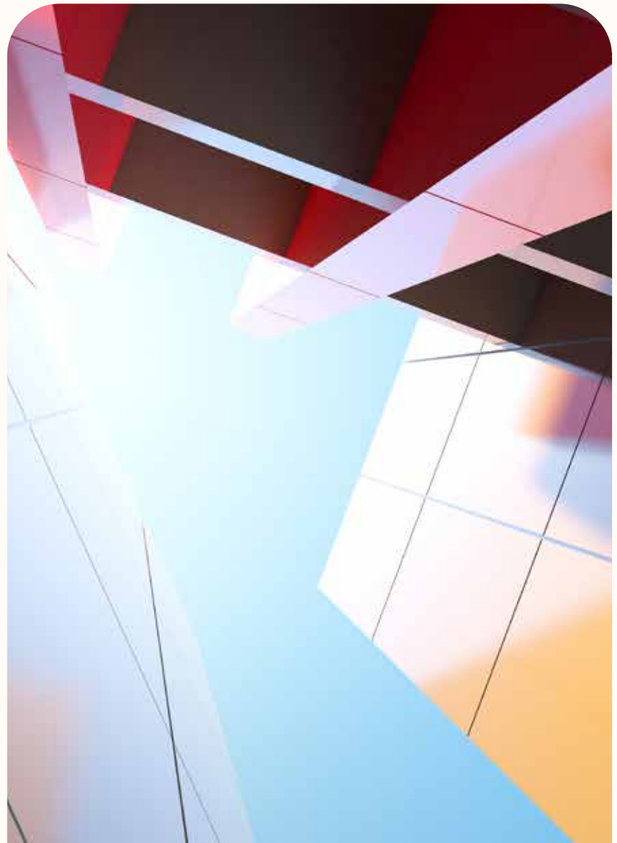
The governance professional's role isn't diminished by this. It's elevated. The Senior Data Governance Manager of 2027 is an orchestrator of human and AI capability - setting strategy, designing agent behavior, holding the accountability line, and continuously improving a program that runs at machine speed and human scale. That's a more powerful role than manual governance ever offered.

The question is no longer whether agentic AI will reshape data governance. It's whether your organization leads the change or reacts to it.

WHY QUALITYAI

QualityAI brings a proven knowledge base and experience delivering complex solutions - expertise other consultancies don't carry into the room. That background matters when the conversation shifts from "what should agentic governance look like" to "how do we run it across our data estate."

QualityAI's AI-driven engineering practice, paired with deep regulatory domain knowledge, maps agentic workflows to real governance work: policy validation, metadata enrichment, lineage tracking, data quality triage. What QualityAI offers clients isn't just a framework and a roadmap deck. It's the engineering to make governed agentic data pipelines real.



Copyright © 2026 QualityAI. All rights reserved.

WHAT'S NEXT

This is just the beginning. The next generation of agentic systems will shift governance from a detect-and-fix regime to predict-and-prevent.

Anticipatory governance models trained on historical failure patterns, lineage graphs, and upstream pipeline behavior that can signal degradation before it registers as an incident.

Context-aware classification sensitivity and PII detection that reads intent, not just pattern.

Governance-as-code policies defined in version-controlled, machine-executable form.

Continuous compliance state real-time regulatory compliance scoring rather than point-in-time audit snapshots. These will be the norm.

The governance discipline that emerges from this looks less like a control function and more like an autonomous immune system - always running, continuously calibrating, learning from every decision it makes. Organizations building agentic foundations today aren't just solving the current scale problem. They're positioning themselves to absorb the next generation of capability as it arrives, rather than scrambling to retrofit it later.

THE INFLECTION POINT

Data governance is at an inflection point. The traditional model - manual stewardship, reactive controls, spreadsheet inventories - cannot keep up with the scale, speed, and complexity of a modern data estate. Agentic AI offers a practical path to something better: governance that is proactive, continuous, measurable, and genuinely enterprise wide.

The organizations that get this right treat agentic governance as a strategic transformation, not an IT project. They invest in the foundations. They deploy thoughtfully. They hold rigorous human oversight. They measure what they're getting and keep improving. The prize - a trusted, governed data estate that enables confident decision-making and demonstrates regulatory compliance - is worth the work.

Copyright © 2026 QualityAI. All rights reserved.

AUTHOR(S)



Krishna Jyoti Das
Senior Architect - Delivery

Krishna Jyoti Das is a Senior Architect at QualityAI and a senior member of the Data Centre of Excellence (CoE). With more than 23 years of experience in data governance, data management, and in designing of large-scale data pipelines, he specializes in transforming complex data into meaningful, actionable insights.



Anish Bishen
Vice President, Head of Data - EMEA

Anish Bishen is the Head of Data for the EMEA region, bringing over 26 years of experience in data migration, modernization and democratization. Anish has led large-scale Data and AI transformation programs, helping organizations modernize complex data estates and bring a data centric approach to AI.

With his expertise spanning strategy and architecture, delivery and operating models, governance and compliance, Anish brings together business priorities and technology execution. Anish's experience across Data Engineering, Test Data Management (TDM), Analytics and AI enables clients to move from strategy through to scalable, governed implementation.

This Point of View was brought to life through the collaborative efforts of the following contributors:

- **Chloe Hibbert**
- **Aadeesh Balodi**
- **Raghu Kumar**

Copyright © 2026 QualityAI. All rights reserved.



www.quality-ai.com

Copyright© 2026 QualityAI. All rights reserved.