



Grid Dynamics white paper

Smarter data and stronger governance for real-world AI agents in finance

www.griddynamics.com



Grid Dynamics

Contents

Executive summary	3
Data-centric foundation for agentic AI readiness	4
Active metadata management and observability	5
Unify unstructured data	6
Access controls for quality data	7
Security and governance for agentic AI productionization	8
Agent lifecycle management	8
Enterprise integration	10
Agentic runtime	12
Observability	13
Governance & security	14
Conclusion	15
About Grid Dynamics	16

Executive summary

Research shows that 76%¹ of financial services companies plan to implement agentic AI within the next 12 months. However, not all of these initiatives will succeed in an industry where sensitive data, strict regulation, and deep third-party dependencies leave no margin for error.

The real differentiator now is how effectively leaders move agentic AI from pilot territory into true enterprise-scale production.

Financial services' leaders who've taken agentic AI into production have stories about bad data, agent sprawl, authentication hurdles, black-box reasoning, and the fixes that followed. Leaders who haven't usually circle back to a polished pilot and call it progress.

Pilots create ideal conditions using perfect datasets, controlled access, and temporary governance, making it seem like every AI agent interprets information accurately and executes flawlessly. Productionization exposes the truth when agents autonomously make decisions and take action in the real world. Today, neither data quality nor AI governance is mature enough at most financial institutions to support safe, scalable agentic deployment.

 ***It's no surprise that 57%² of financial services organizations are still building the foundations required to operationalize agentic AI.***

Most companies have data spread across mainframes, clouds, legacy applications, and M&A inheritances. Structured and unstructured information rarely flow together, and without metadata, lineage, and context, autonomous agents can't interpret information correctly. Manual access rules and limited data observability only increase drift and compliance risk.

At the same time, most firms don't yet have the governance and security framework needed to support autonomous AI agents in production. They lack clear controls for managing agent lifecycles, securing integrations, setting role-appropriate access, ensuring durable execution, explaining decisions, and monitoring every action an agent takes.

In this white paper, we'll explore how financial institutions can make their data, security, and governance frameworks agentic AI-ready to safely operationalize autonomous AI agents across the front, middle, and back office, turning pilots into real-world solutions.

¹ *The Rise of Agentic AI in Financial Services | SS&C Blue Prism*

² *Agentic AI for Capital Markets: Building a Foundation for Success | AWS for Industries*


Data-centric foundation for agentic AI readiness

In recent months, many companies discovered the hard way that pilots stalled not because the models were wrong, but because their data wasn't ready. Even the best models struggle without clean, reliable, and business-specific data behind them. That's why the focus is shifting toward making data—structured and unstructured—trustworthy and representative of how the business actually operates and its regulatory commitments. This data-centric AI development approach improves outcomes by refining and optimizing the data itself, making it a key driver of model success.

For example, a model may catch obvious fraud patterns on its own, but when transaction data is enriched with merchant category codes, device fingerprints, geolocation, spending history, and behavioral metadata, the agent becomes far more precise, reducing false positives while still blocking genuine threats.

To connect, enrich, and govern every data source, many firms have stacked tools over time: data warehouses, data lakes, integration platforms, custom scripts, dashboards, and one-off fixes. Eventually, these layers become fragile and overloaded. Technical debt builds silently, pulling IT teams into constant firefighting instead of innovation. Progress slows, and AI initiatives stall while maintenance demands keep growing.

Businesses instead need to centrally manage every data source across the organization. Such data modernization for AI allows AI agents to gain the context they need to make reliable decisions, accelerate workflows, reduce operational risk, and deliver value across functions, from Anti Money Laundering (AML) to customer service.




Grid Dynamics' Data Observability Starter Kit monitors data quality in real time using rule-based and ML-driven checks, tracking ingestion, schema changes, and pipeline health across structured and unstructured data.

Active metadata management and observability

As AI adoption becomes more pervasive across the enterprise, the ability to trust and trace your data becomes non-negotiable. Active metadata management and data observability form the backbone of trustworthy AI. They ensure that data pipelines are transparent, quality issues are caught early, and every dataset is used responsibly.

Think of metadata as the nervous system of your data ecosystem, delivering the context, visibility, and control that make your data easily consumable by AI agents. But manual classification doesn't scale. Automated tagging does, because it applies structure at the speed of ingestion, capturing lineage, sensitivity, and business meaning, with human review only where needed.



Active metadata management techniques involve continuous capture, integration, analysis, and consumption of metadata, all governed by policies that ensure security, compliance, and visibility throughout the data lifecycle.

Data observability complements this by enabling organizations to monitor the health and quality of data in real-time. It detects anomalies, tracks schema changes, and verifies that data is complete, consistent, timely, unique, and valid, so issues can be identified and resolved before they affect downstream AI outcomes, such as retrieval and compliance.


In short, metadata turns raw assets into governed, contextual knowledge before anyone even asks for it. But data intelligence is what unlocks its value by exposing its source, how it changed, who touched it, and where it flows next. Cataloging, lineage, quality scoring, and productization make data transparent, trusted, and defensible. Compliance and cybersecurity depend on this. And agent reliability depends on this.

In wealth and asset management, an AI agent can determine investment suitability using dozens of inputs: risk tolerance questionnaires, liquidity needs, income stability, portfolio objectives, KYC information, and historical product performance. If metadata isn't accurate or current, an AI agent could recommend a product that no longer matches a client's profile. Active metadata and observability ensure every suitability input is fully tracked, updated, and auditable, so agents recommend products that align with regulations, client needs, and internal policy.

Active metadata also enables financial institutions to automatically log who accessed what data and when. This supports AI model transparency for regulators, ensures compliance with privacy laws like GDPR, and strengthens internal governance by reducing unauthorized access and data drift.

Unify unstructured data

Large financial institutions have data spread across multiple systems, departments, and formats. This includes unstructured data, the largest underused asset that holds enormous value: invoices, emails, PDFs, voice transcripts, logs, images, and documents. Automated metadata management also gives structure and meaning to this data through clear lineage and real-time observability so AI agents, analysts, and engineers can use it confidently.



Access controls for quality data

Now imagine you have quality data available with active metadata management and data observability. Does that guarantee access?

Financial institutions still save access controls in spreadsheets and PDFs that don't travel with the data, making valuable data a liability. Whether a record moves from storage to analytics to fine-tuning, policies should follow automatically based on identity, role, and purpose. That keeps the right people in the loop, prevents accidental exposure, and reduces compliance fire drills.

String data governance is essential, but not all architectures are built for it to scale. Financial institutions increasingly require an open and hybrid design because most already span multiple clouds and on-prem environments. This requires:

- Separating compute and storage to avoid costly migrations; and
- Using open file formats such as Apache Iceberg so applications can read/write data in place, without vendor lock-in.

Finally, the unified data foundation isn't just meant for one-time metadata, observability, and governance setup. Financial institutions must ensure data integration is continuous, not episodic. Data pipelines should adapt on every run, adjust for drift, and optimize for quality, performance, and cost. And teams need visibility into that movement to ensure AI success in production.



Security and governance for agentic AI productionization

With the right data in place, agentic AI can make informed autonomous decisions and take actions that align with an organization's goals and compliance requirements. However, high-quality data alone doesn't guarantee safe execution.

Financial services environments were built for human operators, not AI collaborators. As a result, most firms lack clear policies for agent rollouts and retirements, duplicate agents across departments, undefined integration paths with internal and external systems, and inconsistent access permissions. They also have limited visibility into how agents work, why they make certain decisions, or what triggers specific actions.

To run agents safely in production, businesses must put security and governance systems in place that are purpose-built for agentic AI.

Agent lifecycle management


Most financial institutions often begin with isolated agent experiments: fraud workflow assistants, underwriting helpers, onboarding copilots, AML triage bots, and call-center assistants.

But without lifecycle discipline, these prototypes drift into:

- Agent sprawl across units, desks, and geographies
- Multiple unsynchronized versions of “the same agent.”
- No traceability over changes or ownership
- Difficulty sunseting faulty or outdated agents

This can result in reputational damage, regulatory exposure, inconsistent customer communication, and operational failures. To avoid this, leading institutions implement three core controls that bring order to agent growth and ensure safe scaling:

- **Central agent registry:** A single authoritative catalog that tracks purpose, owner-approved, version, current status, permission scope, and tooling access. Administrators can instantly disable or restrict any agent that goes rogue or falls out of compliance. Such a structure enforces discovery, reuse, accountability, and security at scale.



A central agent registry replaces fragmented agent sprawl with a single source of truth for every deployed agent.

- **Controlled rainbow deployment:** Old and new agent versions run side-by-side, with traffic gradually shifting to the new one. Teams can test real-world behavior, compare performance across versions, and roll back seamlessly if needed, maintaining reliability in production.
- **Simulation testing:** Instead of unit tests, firms can simulate synthetic scenarios such as adversarial cases, rare financial events, edge regulatory scenarios, multi-agent escalations, and abnormal market states. These simulations integrate directly into CI/CD pipelines and trigger automatically with each commit. Many teams also include human-in-the-loop reviews to capture qualitative feedback and reinforce governance with every release.

Enterprise integration


Agent demos often look impressive in isolation. But when they're connected to production systems in financial services, cracks appear fast. Banks, insurers, and asset managers sit on one of the most complex technology landscapes in the world that includes core banking, loan origination, CRMs, ERPs, KYC and AML platforms, trading systems, compliance engines, document archives, and dozens of knowledge stores.

Without a unified strategy, integration complexity grows exponentially:

- Every agent requires custom integrations with every tool ($M \times N$)
- Internal systems evolve faster than agents can adapt
- Manual authentication workflows block automation
- Maintenance demand skyrockets across teams


This $M \times N$ connector explosion drains engineering capacity and slows every AI initiative to a crawl. To overcome this, leading financial institutions are adopting three foundational integration standards that reduce complexity and improve agent communication and access.

- **MCP (Model Context Protocol):** A unified interface layer that lets agents connect to enterprise tools through a shared standard, instead of bespoke point-to-point integrations. It collapses complexity from $M \times N$ down to $M + N$. One connector update benefits every agent immediately, accelerating rollout speed while reducing integration debt across the board.



The Model Context Protocol (MCP) standardizes connections between AI models and enterprise tools, replacing complex point-to-point integrations with a unified interface.

- **A2A communications:** Standardized agent-to-agent messaging enables coordinated workflows across multiple business functions. Finance sees the largest gains here in AML escalations, fraud investigations, dispute resolution workflows, underwriting pipelines, and call-center follow-through.



MCP connects agents to tools and data through a shared server, while A2A enables direct communication between agents, creating a strong foundation for real-world agentic AI applications.

- **OAuth-style delegated access:** Treats agents as first-class identities rather than anonymous scripts. Each agent receives scoped permissions, short-lived credentials, enforced traceability, and immediate revocation options. This ensures agents access only what they need and no more, while creating a full audit trail for every integration.

Agentic runtime


Agentic runtime is where financial AI agents actively think, plan, and execute tasks. And unlike traditional applications, these agents don't just respond to requests—they orchestrate long-running workflows, interact with multiple systems, and often trigger irreversible actions. In financial services, runtime failures carry material consequences: broken workflows, lost state, duplicate trades, unclosed alerts, inconsistent customer handling, regulatory exposure, and costly remediation.

Without the right architecture, the biggest risks quickly surface:

- Long-running processes lose state mid-flow and fail silently
- Distributed agents drift out of sync
- AI agent-generated code executes without proper containment
- Agentic failures cascade through multi-system workflows

These issues can rapidly erode trust across risk, compliance, and operations teams. To prevent this, leading institutions are reinforcing their runtime environment with two core architectural safeguards designed specifically for autonomous systems:

- **Durable execution:** A resilient runtime model that checkpoints progress at defined intervals, persists state externally, and automatically retries transient failures. If an agent crashes mid-task, the workflow resumes from the last safe point instead of restarting. High-risk actions can also trigger human approval gates, ensuring critical activities, such as reconciliation, payments, or record adjustments, remain controlled and auditable.



Durable execution ensures agents can recover from failures and quickly resume workflows by maintaining state, managing task queues, and coordinating timed actions.

- **Controlled sandbox execution:** Secure, ephemeral execution environments that isolate generated code from core systems. Containers operate with constrained CPU, memory, network access, and time limits. Once execution completes, they self-destruct, preventing persistence or lateral movement. This enables agents to code, test, and adjust logic safely, without jeopardizing production systems.

Observability

Financial regulators, such as FINRA and the SEC in the US, demand clarity, auditors require traceability, executives demand certainty, and risk teams demand control. But agent behavior is often opaque, especially once systems move beyond prototypes and into real customer and transactional flows. Traditional logging falls short because it can't capture reasoning paths, contextual actions, or regulatory-relevant decision chains.

When observability is weak, financial teams are left guessing:

- Why did the agent take that action?
- What data did it rely on?
- Which tools did it call and why?
- Was the decision compliant and fair?
- How much did this workflow cost to run?

To eliminate this blind spot, leading financial institutions are adopting three observability practices that turn black-box agent behavior into transparent, governed workflows:

- **Semantic tracing:** Financial institutions are capturing prompts, retrieved data, tool calls, reasoning paths, model responses, and metadata across the entire workflow. These traces create a complete narrative of how decisions are formed, essential for audits, regulatory reporting, and post-event reviews. It turns agent behavior into a visual, trackable process rather than a hidden one.
- **Outcome-based evaluation:** Instead of validating whether an agent followed a predetermined script, teams focus on whether it delivered a safe, compliant, and accurate result. These evaluation frameworks test outcomes across edge cases and integrate directly into CI/CD pipelines, helping teams detect drift early and reduce risk.



Outcome-based evaluation focuses on whether the agent achieved the right result efficiently and safely, measuring task completion, output quality, tool usage, and safety.

- **Real-time cost controls:** Dashboards track token usage, API calls, and latency per agent. When costs spike or performance drops, teams can intervene immediately. Finance teams gain predictability, and operations teams gain control.

Governance & security

Agentic AI introduces a fundamentally different risk profile for financial institutions. These systems can act autonomously, trigger irreversible actions, chain tools together, transfer value, access regulated data, and produce compliant-relevant outputs. When governance is weak, a single misstep can create regulatory breaches, customer harm, or material financial loss.

Unlike traditional applications, agents require continuous oversight, not periodic review. Without strict controls, firms face:

- over-privileged access
- untracked high-impact actions
- regulatory exposure
- undetected policy violations
- opaque decision trails

To strengthen safety and accountability at scale, leading institutions are shifting from ad-hoc security controls to a layered, automated governance model built around three core pillars:

- **Policy-as-code guardrails:** Instead of relying on documented procedures or human judgment, runtime guardrails enforce authorization policies programmatically. When an agent attempts to perform an action, the system evaluates whether it is allowed, denied, or requires approval.
- **Non-human identities (NHIs):** Each agent is treated as a first-class identity with its own role, permissions, authentication method, and revocation rights. This ensures least-privilege access, clear ownership, and complete traceability across all systems and workflows, rather than shared credentials or static keys.
- **Immutable decision logs:** Comprehensive, tamper-proof records capture every action, decision rationale, and access request. These logs support compliance audits, regulatory reviews, fraud investigations, and incident response.

An immutable database such as XTDB can retain full snapshots of all actions behind every decision, allowing financial institutions to reproduce responses to regulatory enquiries perfectly. In contrast, traditional databases make storing granular history and reproducing historical queries complex and expensive.



Conclusion

Financial institutions that win with agentic AI won't be the ones with the flashiest demos or the most experimental pilots. They will be the ones that build the strongest data foundation, the clearest oversight, and the most disciplined governance to support autonomous agents operating in the real world.

In finance, the stakes are higher than anywhere else. Every action leaves a regulatory footprint. Every decision has customer impact. Every failure exposes institutional risk. But you have a choice: sit back and watch competitors make progress—or ask, “How fast can we make this safe, compliant, observable, and scalable?”

Firms that succeed will unlock entirely new operating models:

- Autonomous regulatory workflows
- Self-correcting fraud and AML processes
- Continuous onboarding and suitability assessments
- Proactive compliance monitoring
- High-velocity customer servicing
- Automated operations across the front, middle, and back office

However, an increased focus on reliability often drives up costs and slows innovation. That's why Grid Dynamics built an enterprise-grade Agentic AI Platform to eliminate that trade-off. Built on Temporal, it provides a robust foundation with guardrails, observability, model management, and a ready-to-use development environment, so you can orchestrate, monitor, and scale thousands of AI agents while keeping costs in check.

If you want support moving from pilots to production safely, we're here to help. Contact us.

About Grid Dynamics


Grid Dynamics empowers enterprises to drive growth, boost efficiency, and transform their digital capabilities through expert technology consulting, platform engineering, AI, and advanced analytics.


Fusing technical vision with business acumen, Grid Dynamics (NASDAQ: GDYN) fuels a relentless drive toward sustainable digital transformation. As a forefront provider of technology consulting, platform and product engineering services, and bespoke software development, we draw from over 8 years of leadership in Enterprise AI, coupled with profound expertise in cloud, data, and advanced analytics.

Our commitment to engineering excellence, R&D leadership, a co-innovation ethos, globally efficient follow-the-sun delivery model, and an unwavering "whatever it takes" dedication to client success empower us to solve even the most complex enterprise challenges, ensuring profitable business outcomes and future-proof growth.

Founded in 2006, Grid Dynamics is headquartered in Silicon Valley and has a global talent pool of intellectually curious problem solvers in offices across the Americas, Europe, and India.

LET'S TALK AGENTIC AI DEPLOYMENT →





Grid Dynamics

trusted engineering partner for digital transformation

Grid Dynamics Holdings, Inc.

6101 Bollinger Canyon Road, Suite 465,

San Ramon, CA 94583

www.griddynamics.com