Atlan's Context Conference

Atlan Brings Software Discipline to Enterprise Context

Abstract

At its Context Conference on October 28th, Atlan made three announcements aimed at managing AI context with a software engineering approach. Trident is a shared workspace where people and agents create, version, govern, and distribute the context artifacts AI depends on. AI Registry is an open specification, released under the Apache 2.0 license, for publishing, versioning, and syncing those artifacts across agent harnesses. The Forward Deployed Playbook, now open source, documents Atlan's method for building talk-to-data agents, and an accompanying accelerator program offers hands-on implementation help.

Trident organizes skills, metrics, ontologies, and agents into workspaces aligned to business teams and adds the lifecycle management, governance, and operating visibility an enterprise needs. AI Registry provides the open distribution layer: it defines how organizations register, version, and share those artifacts on infrastructure they own, so any compliant client can install and update them. The playbook documents how teams create, test, and improve context for analytics agents through a three-step loop of bootstrapping context, managing it under version control, and improving it against evals.

Context

Enterprise AI now runs across multiple harnesses, models, and teams. Each platform and toolset comes with its own instructions, skills, and definitions, so the same business knowledge needs to be rebuilt or transferred, and drifts when one copy changes. The consequences show up in production. Sixty-one percent of data and analytics leaders report that missing, incomplete, or inconsistent context has caused AI to misinterpret data or produce an inappropriate result at least occasionally. Where those failures happen very frequently, 90% also report a material wrong action in the past six months.[1]

Much of this is a problem of management versus model. Business rules change constantly, and when one does, 81% of organizations must update it in two or more systems. Every copy is another place for an agent to pick up an outdated policy or metric. Software teams solved the same problem with versioning, ownership, and release management. Most context programs lack that discipline.

Verizon's conference session described the result: third-party agent studios that each manage context locally, with no reuse, no shared standards, and no way to carry curated context forward when a tool changes. Enterprises need to manage context like software, keep it portable, and build it in forms agents can use. Atlan's announcements target all three.

Key Ramifications

The following are the key ramifications of Atlan's Context Conference announcements.

Context artifacts get a release lifecycle with shared ownership. Trident manages skills, metrics, ontologies, and agents the way engineering teams manage code. Workspaces mirror business structure, so the care team owns its playbook and finance owns its metrics. Owners can pin a critical agent to a tested version or let it take new releases automatically, see which consumers still run an older version, and trace dependencies (in Atlan's demo, a call-prep

[1] EMA, Data & Analytics Research, September 2026

skill relied on five others). A built-in agent finds, compares, and merges duplicate skills, and usage data drives suggested improvements that owners review and release. For analytics, Trident manages metric definitions with formulas, grain, dimensions, and SQL examples, and native apps for Snowflake and Databricks generate starting context from existing data environments. Trident scans each skill on arrival, and policy can block skills with high or critical vulnerabilities from being served through APIs, MCP, or SDKs. Fleet observability reports activity, cost per run, latency, and token spend by harness. In Atlan's own environment, 53 of 130 production agents and roughly 300 of 1,100 skills saw use in the prior 30 days.

Context becomes portable across the tools that consume it. AI Registry sits above the formats that express business context and defines how artifacts are registered, versioned, and distributed. It borrows two patterns from data infrastructure. For distribution, it works like a package manager: an owner publishes an artifact once, every consumer installs the same version, and a bad release can be pulled. For interoperability, it follows the open table format model, in which any compliant client reads any compliant registry. A registry consists of files on storage the customer controls, such as a git repository or object storage bucket. The underlying artifacts, such as skill markdown files, stay unchanged, and the registry adds the metadata raw formats lack: owner, version, digest, and status. AI Registry complements rather than replaces formats such as Google's Open Knowledge Format for unstructured context and Apache OSI for semantic models, which Google stated it intends to join. The initial release covers skills and agents, with MCP servers and OSI-compatible metrics planned.

Analytics accuracy becomes a repeatable build discipline. The Forward Deployed Playbook documents the method Atlan uses to build talk-to-data agents. Teams first bootstrap context from existing metadata, documentation, and BI assets, then manage it through version control, ownership, testing, and release, and then improve the agent against evals that expose specific misses. The method rests on three kinds of artifacts: metrics define what the agent should calculate, an ontology captures business concepts and how they relate, and skills guide how the agent works through a task or investigation. A database schema supplies none of these. Atlan Labs reports that across 18 months of testing, frontier models working from bare schema fell short of production accuracy, while context mined from documentation and BI layers lifted accuracy to about 83% without human intervention. In the same tests, tables described with DDL and known traps scored roughly 20 points higher than plain-language descriptions. The accelerator gives teams hands-on help applying the playbook to their own data.

EMA Perspective

Organizations that combine context assembly methods (versioned context prepared ahead of time alongside context assembled at request time) report frequent context failure at roughly a third the rate of those relying on a single method (11% versus 31%), and material wrong actions at less than half the rate (20% versus 48%).[1] These findings reinforce the value of combining prepared, governed context with information assembled when an agent runs a task. Atlan's three announcements go straight at that management problem, and together they move the company from arguing that context matters to defining how an enterprise should operate it.

Trident supplies the operating model: decentralized ownership with centralized control. The people who understand business rules, metrics, and operating procedures usually sit outside the central data organization. Trident's workspaces make those experts responsible for maintaining their context while the enterprise keeps versioning, dependency, and security controls in one place. Software organizations scaled the same way, with teams shipping independently against shared repositories and shared release rules. Context programs that depend on a central data team to create every artifact will not keep pace with the volume of skills and agents the business now produces.

AI Registry addresses the portability issue. Enterprises are accumulating real business knowledge in context artifacts, and they need that investment to survive changes in models and tools. Practitioners at the conference made the point repeatedly: they will not lock that knowledge into a single platform. By publishing the registry as an open spec on customer-owned storage, Atlan separates the upkeep of context from the environments that consume it, including its own. The spec exists today, and Atlan demonstrated its implementation across harnesses. A standard requires independent implementations, not only an open specification, so adoption by harness vendors and registry builders will determine how far it travels.

The Forward Deployed Playbook lowers the barrier to getting started, and the findings behind it support a practical shift in where data leaders invest. Before replacing a model to fix wrong answers, improve metric definitions, table documentation, and evals, and measure the change. Atlan's finding that an

open-weight model with strong context performed comparably to a frontier model warrants testing against each organization's own workloads. If the result holds across enterprise workloads, better context could widen model choice and help control inference costs.

Buyers should test three processes separately. For change ingestion, ask Atlan to show how a revised rule in a source document becomes an updated context artifact, and where people must step in. For lifecycle management, have Trident identify the skills and agents that depend on that artifact, test the update, and show which pinned and unpinned agents consume which release across two harnesses. For interoperability, retrieve and use the same artifact through AI Registry from a client that relies on no Trident-specific capability. Together, those walkthroughs show whether the three announcements form a complete operating model for context.

Atlan's approach gives enterprises a credible operating model for managing AI context independently of the agents and models that consume it, starting with skills and agents. Trident runs that context, AI Registry moves it between tools, and the playbook shows teams how to build it. Together they define the layer where enterprise context gets versioned, governed, and shared, and Atlan has opened that layer to the rest of the market.

Additional Resource: **Atlan Platform Documentation** provides technical information about Trident and Atlan's enterprise AI context management capabilities.



About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com or follow EMA on X or LinkedIn.

4813.103026