Point of view

Operationalize Your Enterprise AI

5 Requirements for Moving From Experimentation to Trusted, Scalable Enterprise Operations



teradata.

Table of Contents

3 Introduction

4 The operationalization gap

5 A 5-part foundation for operational enterprise AI

6 1. Understand: Give AI the context of the business

7 2. Deploy: Different models. One enterprise AI strategy.

8 3. Scale: Match compute to changing demand

9 4. Control: Extend sovereignty beyond data residency

10 5. Protect: Secure the intelligence behind AI

11 Why Teradata for operational enterprise AI

12 From fragmented AI to an enterprise operating model

13 A practical path forward

14 Sources

Introduction

AI has moved beyond experimentation. Enterprise operations are the next challenge.

Enterprises have demonstrated that AI can generate insights, automate tasks, and improve productivity. But adoption is not the same as operationalization. McKinsey's State of AI 2025 found that 88% of respondents reported AI use in at least one business function, while approximately one-third said their organizations had begun scaling AI programs.[1] The strategic question is no longer whether AI works in a pilot, but whether enterprises can generate repeatable business value with predictable costs, acceptable risk, and measurable performance.

The challenge intensifies as AI evolves from isolated applications and human-directed interactions toward autonomous agents participating in critical business processes. AI must do more than generate accurate responses. It must understand enterprise business context, access trusted information, operate across production environments, scale with unpredictable demand, and maintain the security, governance, and resilience that enterprises require. Gartner forecasts that more than 40% of agentic AI projects will be canceled by the end of 2027 due to escalating costs, unclear business value, or inadequate risk controls.[2]

This creates a new challenge: operationalizing AI across the enterprise.

Traditional enterprise data and infrastructure architectures were built to support established analytical and operational workloads. AI introduces new demands for persistent business context, variable compute, integrated deployment, operational control, and protection. Meeting these requirements calls for a more comprehensive approach to enterprise AI architecture.

The organizations that succeed will be those that establish the foundation to operate AI reliably, economically, and securely—not simply those that deploy the most models or agents.

> The question is no longer simply whether we can build AI. It's whether we can operate AI across the enterprise.

5 requirements emerge as AI becomes operational

1. AI needs governed business context so it can understand the meaning, relationships, policies, and rules behind enterprise data
2. Organizations need integrated production environments that reduce the work required to assemble data, analytics, AI software, compute, storage, networking, and operations
3. AI and agentic workloads need compute that can respond to variable demand without turning every temporary peak into permanent capacity
4. Enterprises need control over where data and AI operate, who can administer the environment, and which technologies and external dependencies critical workloads rely on
5. The intelligence behind AI must remain secure, private, governed, available, and recoverable as AI becomes embedded in mission-critical operations

The operationalization gap

The distance between a promising demonstration and a production-grade AI capability is an operating-model gap. IBM's 2025 CEO study reported that only 25% of AI initiatives achieved expected ROI and only 16% scaled enterprise-wide.[3] That gap has a business cost: delayed time to value, duplicated integration effort, unpredictable operating expense, and exposure when critical workflows fail. Models alone can't close it. Business context, production architecture, workload economics, sovereignty, and protection must work together. Consider an illustrative service-operations agent that can summarize incidents in a sandbox but can't safely operate across customer records, operational systems, and sensitive policies in production.

AI operates in a business context vacuum

Enterprise knowledge is distributed across metadata, business definitions, relationships, lineage, policies, workflows, applications, and documents. An agent may retrieve an accurate document yet misunderstand which customer, asset, policy, or service level applies. Gartner found that 63% of surveyed data-management leaders either lacked or were uncertain whether they had the right data-management practices for AI.[4]

That creates a context tax

When each application must rediscover business meaning, reconstruct relationships, and compensate for missing context through repeated prompts and retrieval, latency, cost, and inconsistency can compound. In the illustrative incident scenario, an agent shouldn't infer the definition of a critical outage afresh on every request. Governed business relationships, provenance, and reusable industry knowledge can help establish consistent grounding. The opportunity isn't simply to retrieve more data, but to reuse what the enterprise already knows.

Scale creates a different infrastructure problem

Analytics, data engineering, model inference, and agentic workloads have different concurrency patterns, service levels, and demand curves. Persistent capacity may suit predictable production services, while temporary bursts need more granular and isolated resources. Flexera's 2026 State of the Cloud reports that 85% of respondents identify managing cloud spend as a top challenge and estimate 29% of IaaS/PaaS spending is wasted across cloud workloads generally.[5] Workload-aware economics matter as demand becomes more variable.

And the protection boundary expands

As AI participates in consequential business operations, protection must cover more than stored data. Identity-aware access, sensitive information, business rules, context, availability, and recovery all matter. In Uptime Institute's 2026 outage analysis, 57% of respondents reported that their most recent major IT outage cost over $100,000.[6] Critical AI-dependent workflows need to be designed for failure and recovery.

A 5-part foundation for operational enterprise AI

Operationalizing AI requires more than deploying models and applications. Enterprises need a foundation that brings together trusted business context, integrated infrastructure, flexible compute, operational control, and enterprise-grade protection.

Five interconnected requirements provide a practical framework for moving AI from experimentation to sustained enterprise operations: understand, deploy, scale, control, and protect.

These requirements aren't sequential steps or separate technology decisions. They represent interconnected capabilities that help organizations address the barriers preventing AI from delivering repeatable business value. Different workloads will have different priorities, making it essential to identify and address the most significant constraint first.

Requirement	AI Objective	Business Outcome	Teradata Solution
Understand	Give AI trusted business context to act intelligently	Improve AI accuracy, explainability, and actionability	Tera Context Engine
Deploy	Accelerate AI from development to production	Reduce deployment complexity and accelerate time to value	Teradata Factory
Scale	Scale AI efficiently as demand changes	Maintain performance while improving infrastructure efficiency and cost control	Teradata Cloud
Control	Maintain control over where and how AI operates	Support sovereignty, operational independence, and deployment choice	Teradata Sovereign
Protect	Protect the intelligence powering enterprise AI	Strengthen security, privacy, governance, and operational resilience	Teradata privacy, security, and resiliency

1. Understand:
Give AI the context of the business

Traditional AI approaches often rely on retrieving relevant information and asking models to infer its meaning. But enterprise intelligence depends on more than retrieving the right documents or data.

AI must understand how customers, products, transactions, policies, and business processes relate to one another—and apply that understanding consistently across applications and agents.

Tera Context Engine addresses this challenge by transforming fragmented enterprise knowledge into governed, reusable business context. Through native context graphs, Industry Knowledge Models, and deterministic retrieval, organizations can give AI a more consistent foundation for understanding and acting on enterprise information.

Rather than repeatedly reconstructing business meaning for every AI interaction, organizations can build, govern, and reuse trusted context across their AI ecosystem.

Why this matters

By making business context reusable across applications and agents, Tera Context Engine helps organizations move beyond isolated AI implementations toward a more consistent foundation for enterprise intelligence. This enables AI to interpret information within the context of the business, improving the reliability, explainability, and relevance of AI-driven insights and actions.

Build enterprise knowledge once.
Govern it consistently.
Reuse it wherever AI operates.



Native context graph
Connect enterprise entities, relationships, definitions, lineage, provenance, and business meaning in a governed graph rather than relying on disconnected metadata records.



Deterministic retrieval
Follow governed relationships and predictable retrieval paths to limit repeated inference and unnecessary retrieval work.



Industry Knowledge Models
Combine industry expertise, ontologies (formal maps of business concepts and relationships), and governed relationships to help AI interpret entities, business rules, and how information connects.



Governed, open context reuse
Connect business context across existing data platforms, applications, and agents, with clear human-review and access-control boundaries.

2. Deploy: Different models. One enterprise AI strategy.

Moving AI into production requires more than deploying models. Organizations must integrate enterprise data, analytics, AI development tools, compute, networking, security, and operations. Fragmented architectures increase complexity, slow deployment, and make it harder to operate AI reliably at scale.

Teradata enables organizations to operationalize analytics and AI across cloud and customer-controlled infrastructure, matching deployment approaches to business, technical, and regulatory requirements. Rather than prescribing a single deployment model, Teradata provides complementary architectures designed to support enterprise performance, flexibility, and operational control. **Teradata Factory** combines Teradata analytics and **AI Studio** with **Dell** infrastructure and **NVIDIA** acceleration in a configured, tested, and validated on-premises stack. Its integrated architecture brings enterprise data, AI capabilities, and workload-specific CPU and GPU resources together, reducing the effort required to assemble and integrate the underlying infrastructure.

For cloud deployments, **Teradata Cloud** provides enterprise analytics, data access, workload management, and elastic compute in a managed cloud environment, enabling organizations to operationalize analytical and AI workloads without managing the underlying infrastructure.

For organizations with heightened sovereignty requirements, **Teradata Sovereign** extends deployment choice through customer-controlled on-premises and sovereign cloud architectures designed to address data residency, operational control, and infrastructure isolation.



3. Scale: Match compute to changing demand

AI workloads introduce new patterns of demand that traditional capacity planning alone may struggle to accommodate. Enterprises need infrastructure that can respond to changing workloads without compromising critical production performance or requiring permanent investment in peak capacity.

AI and analytics do not consume compute uniformly. A steady production workload, an ad hoc data-engineering job, and an agentic workflow with unpredictable bursts should not automatically share the same scaling model. The operating objective is to match resources to workload characteristics and service-level needs rather than paying for idle capacity or accepting avoidable contention.

Teradata Cloud combines persistent Active Compute for sustained demand with on-demand **Elastic Compute** for variable workloads. Elastic Compute can add isolated capacity in smaller, node-based increments, helping organizations align resources with individual workload requirements. Workload-aware management helps balance service priorities and demand across ongoing analytics and bursts of AI and agentic activity.

- Align capacity with workload demand rather than provisioning for every peak.
- Isolate variable workloads to help protect the performance of critical services.
- Extend analytics and AI to compatible open table formats without creating unnecessary data copies.
- Use scheduling, prioritization, auto-suspend, and consumption visibility to improve control over operating costs.



4. Control: Extend sovereignty beyond data residency

For governments and regulated enterprises, sovereignty extends beyond the geographic location of stored data. It concerns who operates the environment, whether workloads can run without public-internet connectivity, which external dependencies are required, and how organizations preserve operational choice. Thales's 2025 Data Threat Report found that 55% of respondents pursuing digital sovereignty cited customer, regional, or global privacy mandates as drivers.[7]

- **Data sovereignty:** Control where data and workloads reside, how information is accessed, and how governance and policy are applied.
- **Operational sovereignty:** Maintain greater control over authorized operators, infrastructure, continuity, recovery, and day-to-day operations.
- **Digital sovereignty:** Preserve choice across infrastructure, deployment environments, AI models, inference, tools, and platforms.

Teradata supports enterprise sovereignty through complementary deployment architectures designed to address different operational and regulatory requirements.

Teradata Factory enables integrated, customer-controlled on-premises deployments, while Teradata Sovereign Cloud with **Google Distributed Cloud** supports physically isolated, air-gapped environments with a cloud operating experience.

Together, these approaches provide organizations with greater choice over where AI operates, how infrastructure is managed, and which external dependencies critical workloads require.



5. Protect: Secure the intelligence behind AI

As agents move from generating answers to taking actions, protecting enterprise intelligence requires more than data encryption. Organizations must manage what agents can access and do, protect sensitive information, and preserve the ability to recover critical services.

Resiliency
Combine immutable backup protection with database-level high availability and tested recovery procedures to help restore critical analytical and AI-dependent services after disruption.

Security
Apply identity-aware access and policy-based permissions to enterprise data and applications. As agents invoke data and tools, their identities, authorizations, and permitted actions must be governed by enforceable controls.

Privacy
Protect sensitive identities while preserving the relationships and analytical utility required for authorized AI and analytics. PolyAnonymization and PII Vault, powered by **Anonomatic**, help organizations reduce exposure of sensitive information without unnecessarily sacrificing the business intelligence derived from it.

Governance
Maintain business definitions, lineage, and auditability; distinguish the context agents use to interpret information from the runtime authorization required when they access data or invoke tools.



Why Teradata for operational enterprise AI

Operationalizing AI requires more than connecting models to enterprise data. It demands a foundation that combines trusted business intelligence, scalable infrastructure, and enterprise control.

Teradata brings together three critical capabilities:

- **Trusted enterprise intelligence:** Connect enterprise data with governed, reusable business context to help AI deliver more accurate, explainable, and actionable outcomes.
- **Performance and efficiency at scale:** Support demanding analytics and AI workloads with flexible compute, workload isolation, and resource management designed to balance performance and economics.
- **Enterprise choice and control:** Operationalize AI across cloud, on-premises, and sovereign environments while addressing enterprise requirements for security, privacy, governance, and resiliency.

Together, these capabilities help organizations move beyond isolated AI initiatives toward a more integrated approach to operational enterprise intelligence.



From fragmented AI to an enterprise operating model

The five requirements for operational enterprise AI are interconnected. Trusted business context enables more reliable AI outcomes, integrated deployment accelerates production readiness, flexible compute supports changing workloads, sovereignty provides operational choice and control, and enterprise-grade protection safeguards critical intelligence and services.

Together, these capabilities establish a foundation for operating AI reliably, efficiently, and securely across the enterprise.

Understand

Deploy

Scale

Control

Protect

Consider the illustrative service-operations agent: governed relationships help identify the right customer and incident; integrated deployment connects data and AI; isolated compute accommodates investigation spikes; an appropriate deployment model addresses operational constraints; and data access policies and recovery capabilities help protect continuity. The same foundations support conventional analytics and agent-assisted data work. Each dimension addresses a distinct obstacle to reliable operations.

What this enables

- Reduce repeated context-building work across analytical and agentic use cases.
- Shorten integration effort and clarify the path from pilot to operational use.
- Match compute consumption to workload demand to improve cost visibility and control.
- Choose deployment architectures that meet performance, operational, and sovereignty requirements.
- Reduce operational exposure with privacy, security, availability, and recovery practices.

Where's your constraint today: context, deployment, scale, control, or protection?

The next question

The five dimensions provide a way to identify the constraint that matters most: missing business context, deployment complexity, inefficient scaling, insufficient operational control, or gaps in protection. The priority is to resolve the limiting factor—not to undertake a broad platform transformation.

A practical path forward

Begin with a priority workload, its business owner, expected outcome, production dependencies, and acceptable risk.

Evaluate it against five questions:

- **Understand:** Does the AI have governed access to the business definitions, relationships, policies, lineage, and industry knowledge it needs?
- **Deploy:** Can the workload move into production without creating a complex, bespoke infrastructure integration project?
- **Scale:** Can compute expand and isolate around changing demand without permanently scaling the entire environment?
- **Control:** Do you have the required control over where the workload runs, who operates it, and which external technologies it depends on?
- **Protect:** Can the enterprise intelligence behind the workload remain private, secure, governed, available, and recoverable?

Operationalizing enterprise AI isn't a single product decision. It's the discipline of making context, infrastructure, deployment, control, and protection work together.

Start with one workload.
Leave with one decision to operationalize it.

Bring your priority AI initiative to a working session with Teradata experts.

Together, we'll evaluate its requirements for business context, deployment, scale, control, and protection—identifying the most significant barrier to production and defining a practical next step.

Whether that means a context review, deployment architecture discussion, workload-sizing exercise, or security and recovery assessment, the goal is simple: leave with a clear action, an accountable owner, and a measurable success criterion.

Schedule your enterprise AI workload assessment.



Sources

1. “The State of AI in 2025: Agents, Innovation, and Transformation,” McKinsey, 2025, https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-state-of-ai-2025.
2. “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027,” Gartner, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027.
3. “IBM Study: CEOs Double Down on AI While Navigating Enterprise Hurdles,” IBM, 2025, https://newsroom.ibm.com/2025-05-06-ibm-study-ceos-double-down-on-ai-while-navigating-enterprise-hurdles.
4. “Lack of AI-Ready Data Puts AI Projects at Risk,” Gartner, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk.
5. “Flexera Finds Cloud Value is Rising While AI Waste Grows,” Flexera, 2026, https://www.flexera.com/about-us/press-center/flexera-finds-cloud-value-is-rising-while-ai-waste-grows.
6. “Uptime Announces Annual Outage Analysis Report 2026, ” Uptime Institute, 2026, https://professionalservices.uptimeinstitute.com/about-ui/press-releases/uptime-announces-annual-outage-analysis-report-2026.
7. “2025 Thales Data Threat Report,” Thales, 2025, https://cpl.thalesgroup.com/2025/data-threat-report.

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

The Teradata logo is a trademark, and Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

© 2026 Teradata Corporation All Rights Reserved. Produced in U.S.A. 11.26

teradata.