TrendCandy

State of Context Management Report 2026

SPONSORED BY DataHub

Table of Contents

High confidence, weak foundations: The state of context management in 2026

This year marks a turning point for enterprise AI, but not in the way most organizations think.

The data reveals striking contradictions:

- 88% claim to have fully operational context platforms, yet 61% frequently delay AI initiatives due to a lack of trusted data.
- 90% of organizations say they're "AI-ready," yet 87% cite data readiness as their biggest impediment to putting AI into production.

This report maps the gap between what organizations believe they have and what they actually need to scale agentic AI. The findings show that context engineering is emerging as a core discipline and context platforms are transitioning from claimed capability to recognized strategic priority.

For leaders navigating this transition, the message is clear: understanding the difference between having some context management and having mature context management will determine who captures AI value and who gets caught in the chasm between aspiration and reality.

Context management is the organization-wide capability to reliably deliver the most relevant data to AI context windows, enabling the governed and enterprise scale deployment of agents.

Shirshanka Das
Co-founder & CTO,
DataHub

Who should read this report

Role	What you care about	What you'll learn
Chief Data Officers (CDOs)	Strategic AI infrastructure decisions and enterprise data governance	Why 88% claim operational context platforms, yet 89% plan major infrastructure investments in the next 12 months—and what this reveals about true maturity vs perceived maturity
AI/ML Platform Leaders	Production-scale AI deployment and model reliability	What the gap between success metrics and operational challenges reveals about infrastructure readiness for production-scale AI
Data Engineering Leaders	Operational efficiency, data readiness, and team productivity	Why 86% of data teams spend considerable time searching for the right data today, and how context engineering is emerging as a core discipline
Enterprise Architects	System integration, governance frameworks, and long-term scalability	Why 93% of organizations plan to treat context as shared infrastructure, and what the obstacles to scaling agentic AI reveal about the need for a unified context platform

Survey methodology

DataHub commissioned an independent research firm, TrendCandy, to survey 250 IT and data team leaders on the topic of data context management for AI agents.

Responses were collected online using national B2B panels using generally accepted survey methodology practices.

The sample was randomly selected, and all respondents were compensated for participation. Double-screening and data verification was employed to yield high-quality data.

The margin of error for this study is +/-4% at the 95% confidence level.

Respondent breakout by role, company size, and industry

Role
(level)

- 32% DIRECTOR
- 32% MANAGER
- 30% EXECUTIVE
- 6% VP

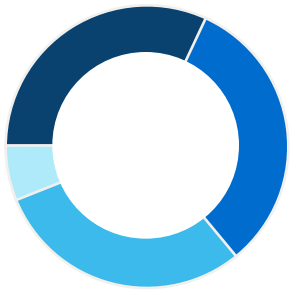

Size
(employees)

- 40% 100 TO 1,000
- 36% 1,001 TO 5,000
- 14% 5,001 TO 10,000
- 10% MORE THAN 10,000

Industry

- 43% TECH AI-NATIVE
- 28% TECH NON-AI-NATIVE
- 14% FINANCIAL SERVICES
- 10% RETAIL, E-COMMERCE
- 4% TELECOM, MEDIA, ENT
- 1% INSURANCE

CHAPTER 01

The aspiration-reality gap

Organizations self-assess high context management maturity while admitting fundamental infrastructure gaps.

High self-assessment masks operational struggles

Organizations report strong confidence in their context management capabilities:

- **39% self-report at Stage 4 maturity:** “We have accurate context for AI agents to use and manage data and AI assets”
- **88% are confident or very confident they have a fully operational context platform** that provides context to both humans and AI agents
- **89% say they’re effective or very effective at using AI** to organize and contextualize internal data

Which stage best describes your organization’s current situation when it comes to understanding the context of your data?

STAGE 01	STAGE 02	STAGE 03	STAGE 04
16%	21%	24%	39%
We primarily use spreadsheets, Slack or Teams, and tribal knowledge	We have harnessed metadata so humans can use and manage data	We have a single pane of glass for humans and machines to use and manage data and AI assets	We have accurate context for AI agents to use and manage data and AI assets

Operational capabilities tell a different story

But when asked about specific operational capabilities, a different picture emerges.

Risk tolerance for agents is lenient

82%

would somewhat or completely trust AI agents with high-stakes tasks without reliable context, lineage, observability, and governance.

AI outputs are inaccurate

66%

report AI models in their organizations generating biased or misleading insights due to low maturity of data infrastructure in providing sufficient context.

Authoritative sources remain elusive

57%

find it challenging or very challenging to identify authoritative sources of truth for data.

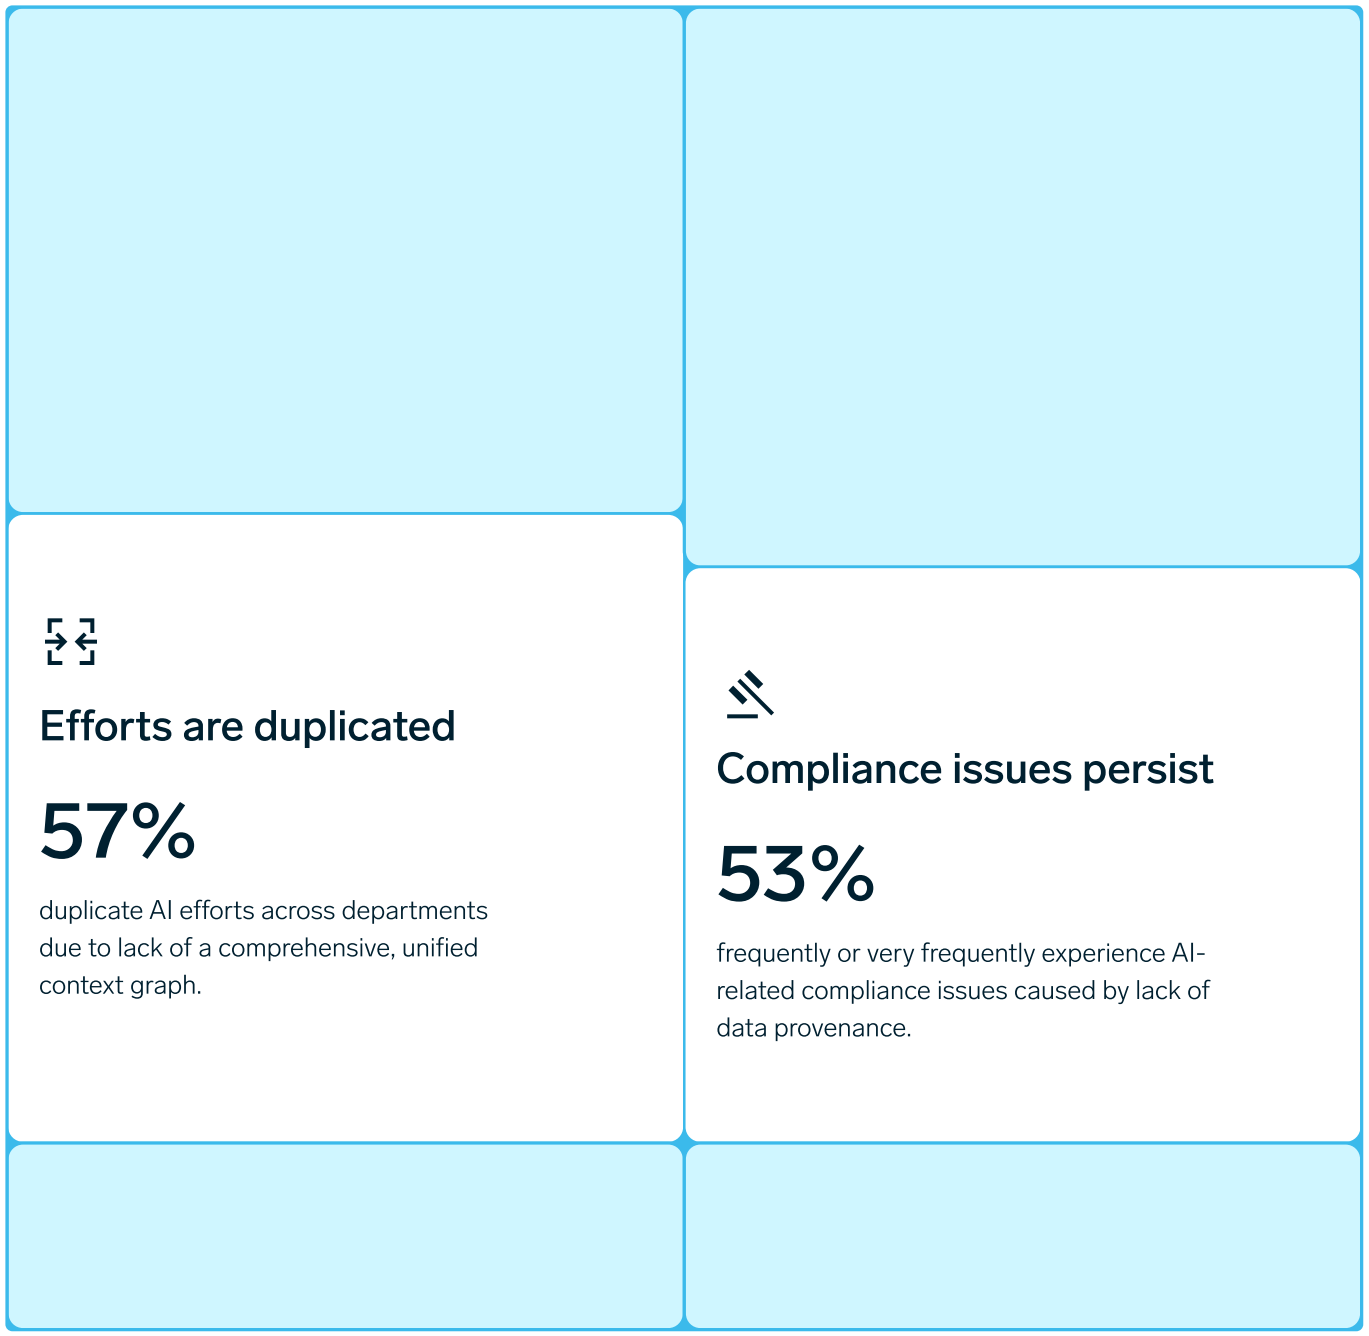

Organizations claim to have fully operational context platforms, yet struggle with the fundamentals those platforms should provide: authoritative sources, reliable outputs, governed data, and unified context.

In an emerging category without established benchmarks, these gaps may be understandable, but could also lead to a loss of trust and confidence as AI is more widely deployed.

Organizations may be assessing their maturity against peer perception rather than against what production-scale AI actually demands.

ROI from AI can be high

Organizations report strong returns from AI and data initiatives while simultaneously reporting operational issues that, once addressed, could further yield significant value.

Reported returns from AI and data initiatives are strong

251%

Average estimated ROI of their organization's AI initiatives over the past year.

95%

Would rate the ROI from their organization's data initiatives as good or excellent.

92%

Say AI initiatives at their organization are likely or very likely to be completed on time and within budget.

Metrics organizations track for AI impact

Which metrics does your team actively use to measure the impact of AI on data organization and context?

(SELECT ALL THAT APPLY)

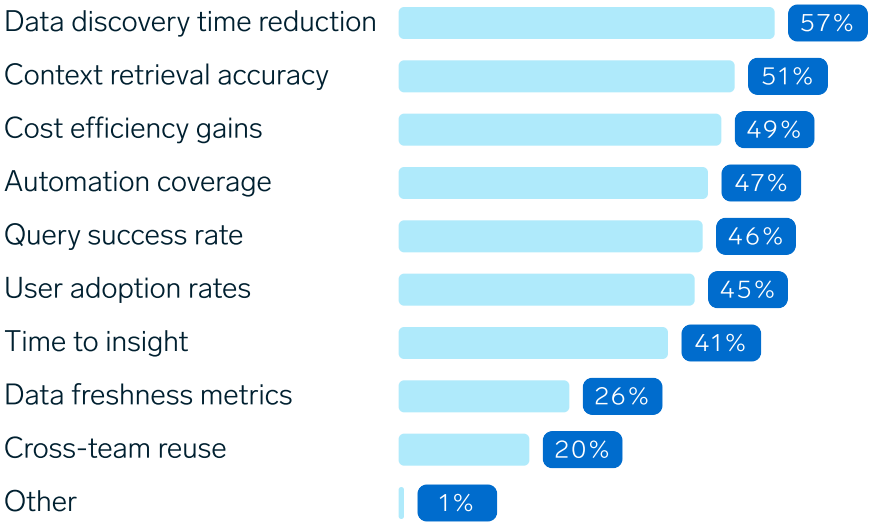

Outcomes organizations prioritize for AI agents

What outcomes are most important to your team when using AI agents to organize and contextualize internal data?

(SELECT ALL THAT APPLY)

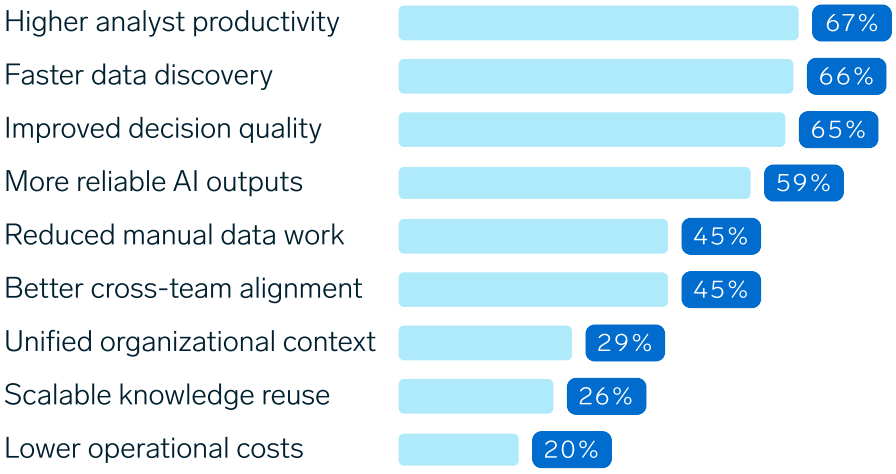

Plenty of value is yet to be unlocked

These self-reported returns should be noted alongside lingering operational challenges documented on the following page, suggesting there is **plenty of value yet to be derived from AI implementations.**

86%

of teams spend some time or a great deal of time searching for the right data today

Only 2% report no time lost

81%

of analysts lose some time or a great deal of time per month due to reliance on spreadsheets and tribal knowledge for data discovery

Only 2% report no time lost

Top agentic AI priorities vs. reported experiences

Priority for AI Agents	% Prioritizing	Reported Operational Experience	% Reporting
Higher analyst productivity	67%	Analysts lose time to spreadsheets/tribal knowledge	81%
Faster data discovery	66%	Teams spend considerable time searching for data	86%
Improved decision quality	65%	AI models frequently generate biased/ misleading insights	66%
More reliable AI outputs	59%	AI models frequently generate biased/ misleading insights	66%

CHAPTER 02

From pilot to production: where 'good enough' breaks

As AI initiatives move from pilot to production, infrastructure gaps become bottlenecks and 'good enough' solutions reveal their limitations.

AI-ready confidence meets data readiness reality

90% of organizations say they're "AI-ready," yet 87% state data readiness is their biggest impediment to putting AI into production at their organization.

In practice, this contradiction suggests that most organizations have cleared the bar for AI experimentation while the foundational data challenges that determine production success remain unresolved.

How accurately does the term "AI-ready" describe your organization's data today?

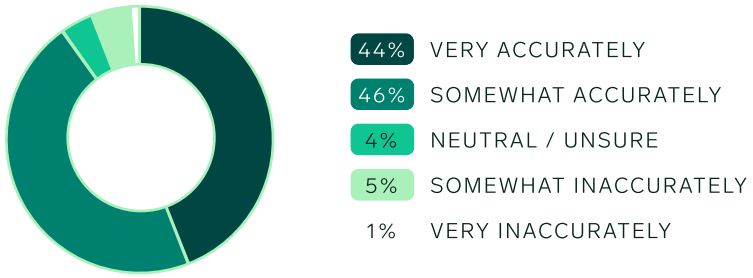

To what extent is data readiness your biggest impediment to putting AI into production at your organization?

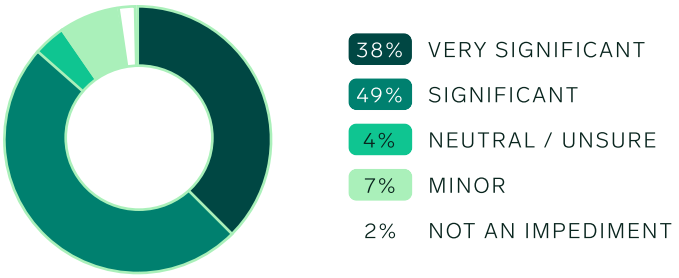

Expectations outpace infrastructure

92% of organizations expect on-time AI delivery, yet 61% admit delays from untrusted data.

The most likely explanation: delay incidents are being treated as one-off exceptions rather than systemic signals. When 61% report recurring delays from untrusted data, the problem isn't individual project management, it's that timeline expectations haven't been recalibrated to account for infrastructure gaps.

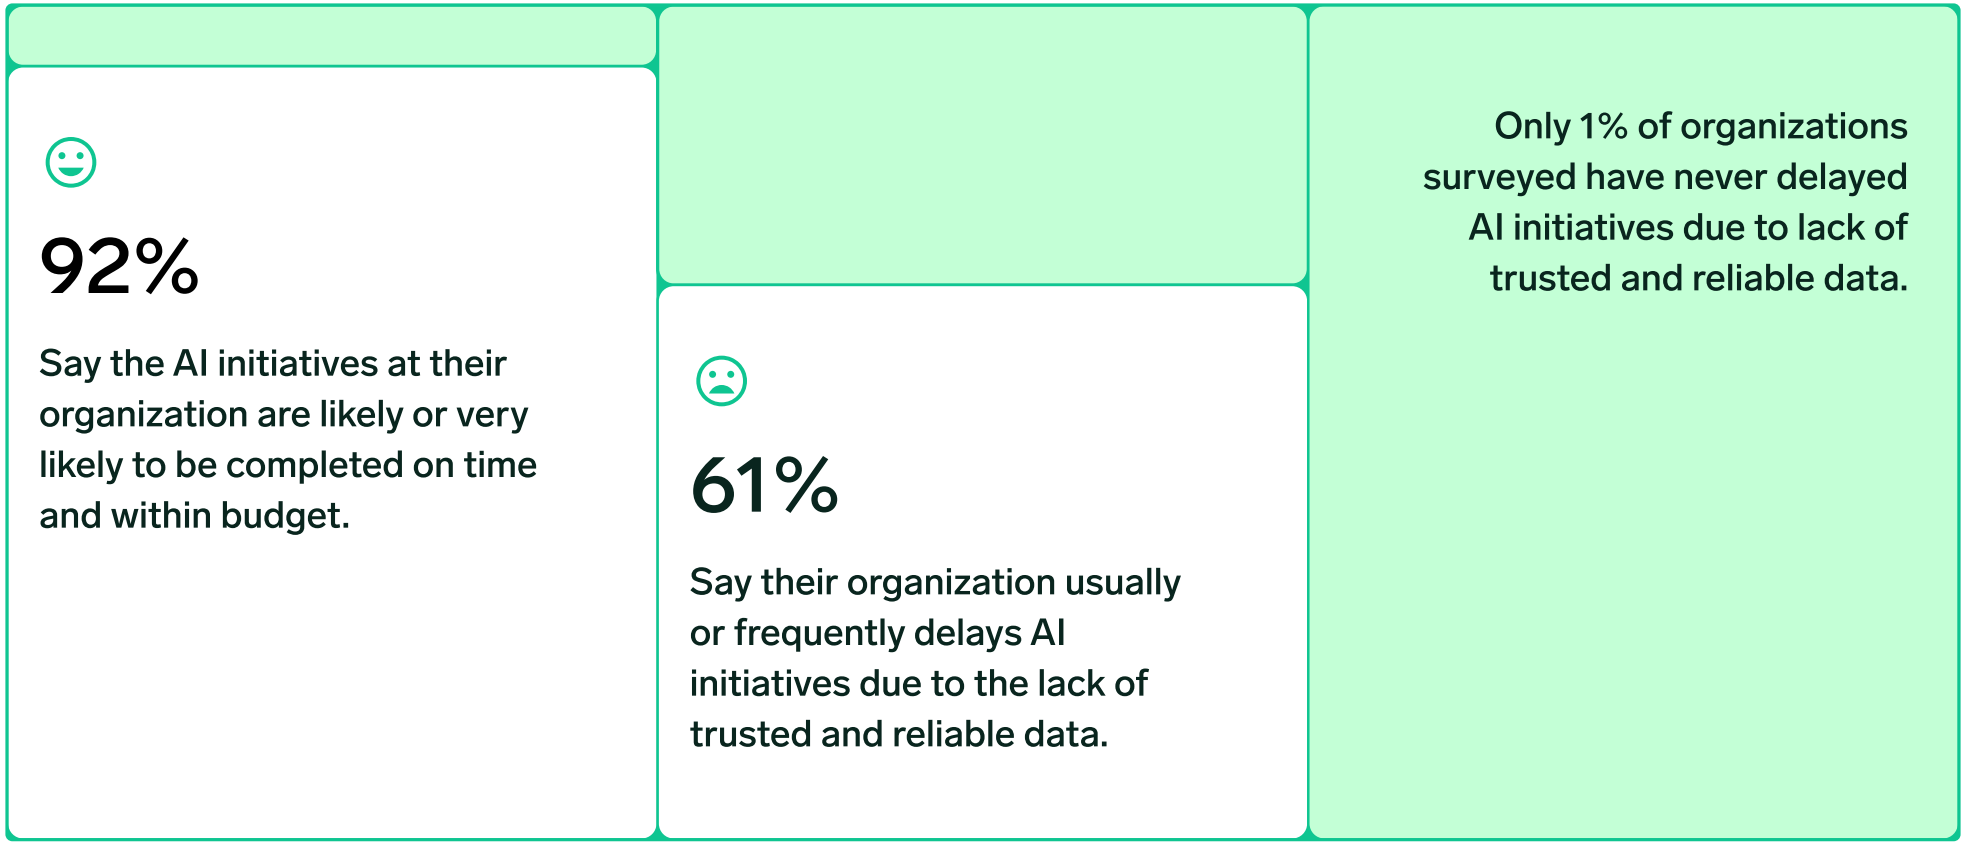

Obstacles to scaling AI agents

When AI agents move from pilot to production, systemic challenges emerge. These obstacles reveal where **context infrastructure gaps become bottlenecks** to enterprise-wide deployment.

Biggest obstacles to scaling AI agents

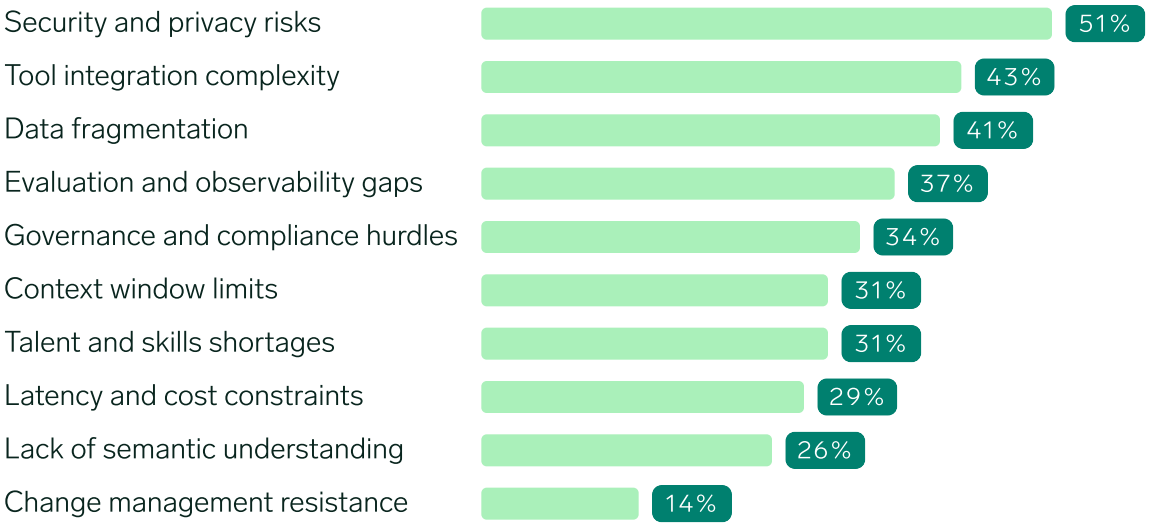

What are the biggest obstacles your team faces when scaling AI agents in production?

(SELECT ALL THAT APPLY)

Security leads risks

More than half cite security concerns as an obstacle to scaling AI agents. Security frameworks are a foundational component of operational context platforms.

Integration pain

Organizations report difficulty integrating tools, suggesting multiple point solutions rather than unified context platform infrastructure.

Fragmented data

Fragmented data across silos prevents AI agents from accessing comprehensive context across the organization.

Blind spots persist

More than one-third lack the capability to effectively monitor AI agents and the data they use.

Organizations self-reporting Stage 4 maturity (39%) and fully operational context platforms (88%) would not typically encounter these foundational challenges at these levels. The obstacles suggest infrastructure maturity is lower than self-assessments indicate.

CHAPTER 03

The infrastructure correction

Chapters 1 and 2 documented the gap between what organizations claim and what they experience. This chapter reveals what they're doing about it.

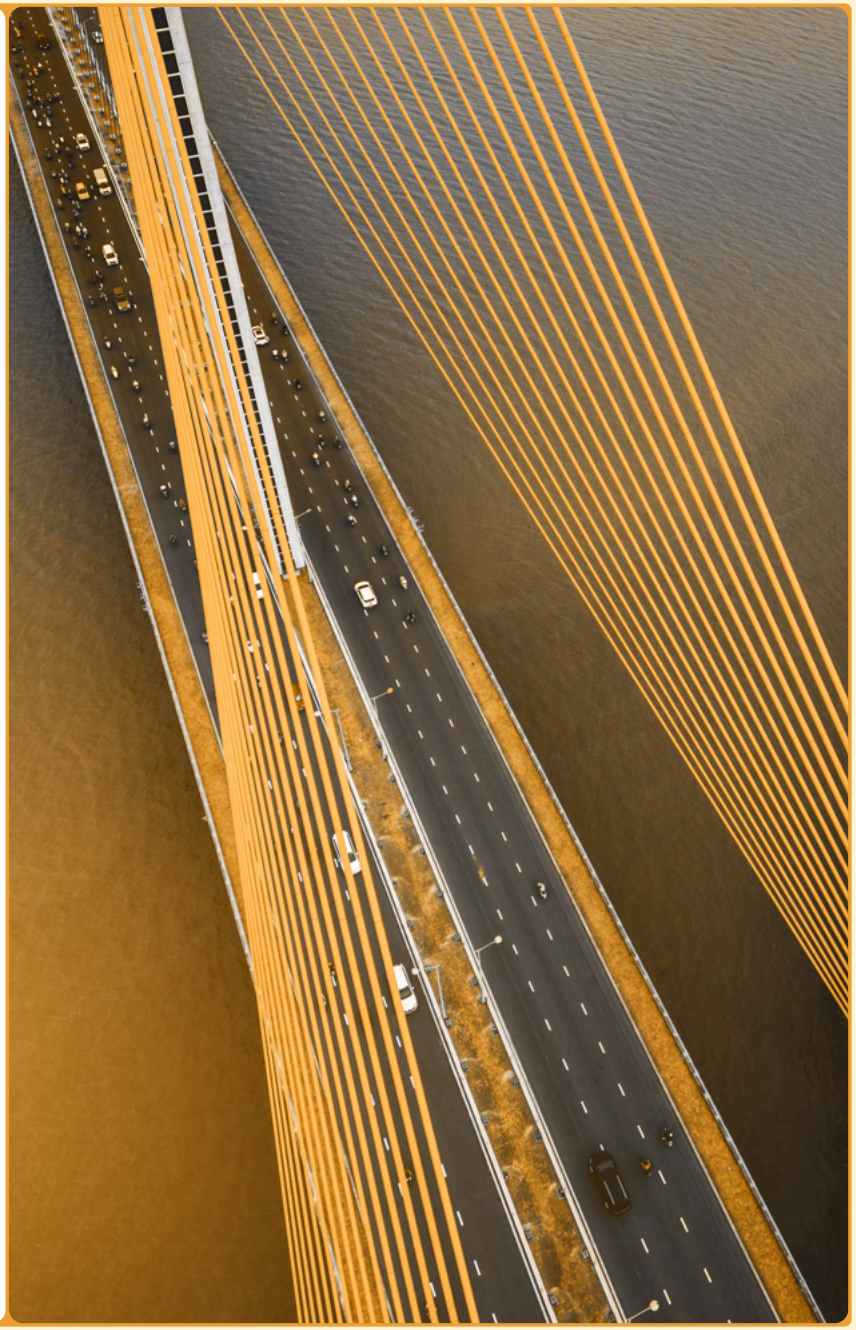

What IT and data leaders now accept about production-ready AI

IT and data leaders agree that context management, not prompt engineering or RAG alone, is essential for agentic AI success.

Agents need context

83%

agree that agentic AI cannot reach production value without a context platform

Prompts are not enough

82%

agree that prompt engineering alone is no longer sufficient to power AI at scale

RAG is not enough

77%

agree that retrieval-augmented generation (RAG) alone is insufficient for accurate and reliable AI deployments in production

Context platform is required

- 26% STRONGLY AGREE
- 57% AGREE
- 11% NEUTRAL/UNSURE
- 6% DISAGREE
- 0% STRONGLY DISAGREE

Prompt engineering alone is insufficient

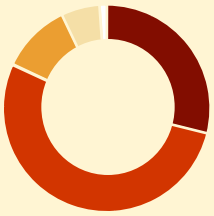

- 29% STRONGLY AGREE
- 53% AGREE
- 11% NEUTRAL/UNSURE
- 6% DISAGREE
- 1% STRONGLY DISAGREE

RAG alone is insufficient

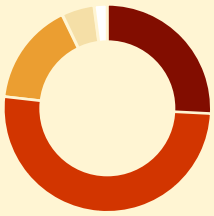

- 26% STRONGLY AGREE
- 51% AGREE
- 16% NEUTRAL/UNSURE
- 5% DISAGREE
- 2% STRONGLY DISAGREE

Context management reaches the C-suite

Context management is emerging from a technical concern to an **enterprise-wide strategic priority**, with leadership alignment and formal strategy both reaching critical mass.

Strategy established

88%

have formally included context management architecture in their organization's AI strategy today

Executive priority locked

91%

see context management as an important priority at the executive or C-level over the next 1–3 years

Shared infrastructure model takes hold

93%

of organizations are likely to treat context as shared infrastructure rather than team-specific tooling

How organizations are investing in context management

Organizations are backing their recognition of context's importance with concrete action:

investing in context platforms and building relevant skills.

Building the infrastructure: Context platforms become top priority

89%

of teams are likely to invest in context management infrastructure within the next 12 months

91%

of organizations are likely to build or buy tools to help create a context platform within the next 12 months

92%

expect investment in context platforms to increase year over year

Building the capabilities: Context engineering training surge imminent

95%

agree that context engineering is important to power AI agents at scale

95%

of data teams are likely to invest in training for context engineering skills during 2026

2026 context management priorities reveal the foundational work ahead

The priority list confirms what the investment surge revealed: organizations are focused on building foundational context management infrastructure.

Top priorities for 2026

What goals are most important for your data team's context management strategy in 2026? (SELECT ALL THAT APPLY)

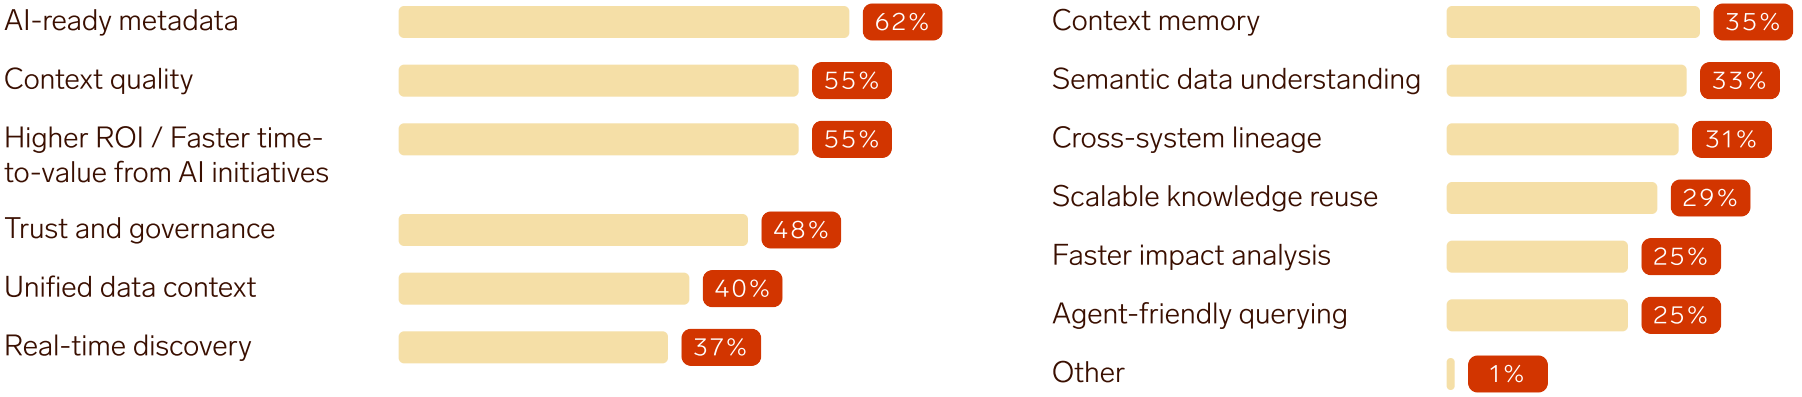

The 2026 context management agenda isn't about optimization: it's about building the basics.

CONCLUSION

The path forward: Closing the gap between context management confidence and capability

Organizations face a choice: continue calibrating maturity against peer perception, or benchmark against what production-scale AI actually requires.

From self-assessment to self-correction

This research captures a market inflection point: the moment context management shifts from adopted capability to **recognized infrastructure discipline.**

The confidence

What organizations claimed

- **39% self-assess** at Stage 4 maturity
- **88% have** fully operational context platforms
- **90% describe** their data as AI-ready
- **89% effective** at using AI to organize and contextualize data
- **251% average** ROI from AI initiatives
- **92% expect** on-time AI delivery

The reality

What organizations experience

- **87% cite** data readiness as significant impediment to production
- **61% frequently delay** AI initiatives due to lack of trusted data
- **66% frequently get** biased or misleading AI insights
- **53% frequently face** compliance issues from lack of data provenance
- **57% duplicate** AI efforts across departments
- **57% struggle** to identify authoritative data sources
- **86% of data teams spend** considerable time searching for data

The correction

What organizations are doing

- **89% investing** in context infrastructure in next 12 months
- **91% building or buying** context platform tools
- **92% increasing** investment year-over-year
- **95% investing** in context engineering training
- **95% agree** context engineering is important to power AI agents at scale

The opportunity ahead

The aspiration-reality gap documented in this report is not a failure. It is what a category looks like in the moment before it matures. Organizations have **identified context management as the right destination**; the survey data shows they are now redirecting investment, formalizing strategy, and building skills to match.

Context management is the new competitive moat

As AI models continue to evolve, the differentiator is no longer which model an organization uses; it is the quality, reliability, and governance of the context that model receives. **Organizations that invest in mature context infrastructure will compound advantages across every AI initiative**: more accurate outputs, faster time to production, fewer compliance failures, and less duplicated effort. Those that don't will continue to scale the very problems that stall their AI ambitions today.

The window is now

The data tells a clear story: the industry is at an inflection point. With 89% of teams planning to invest in context infrastructure within the next 12 months and 95% prioritizing context engineering training, **2026 is shaping up to be the year context management moves from aspiration to operational reality**. But investment alone won't close the gap. Organizations that treat context management as a foundational discipline, rather than a feature to bolt on after deployment, will separate from the pack.

Three imperatives for data and IT leaders in 2026

01

Audit honestly, not optimistically.

The data shows that self-assessed maturity dramatically overstates actual capability. Benchmark your context infrastructure against what production AI actually demands: authoritative sources, reliable lineage, unified context, and governed access.

02

Invest in context as shared infrastructure.

With 57% of organizations duplicating AI efforts across departments, the cost of fragmented, team-specific tooling is already measurable. A unified context platform eliminates redundancy, accelerates every downstream AI initiative, and creates the single source of truth that 93% of organizations say they want.

03

Build the context engineering muscle.

Context engineering is emerging as the discipline that bridges the gap between having data and having AI-ready data. Organizations that develop this capability will unlock the production-scale AI value that remains out of reach for most today.

From confidence to capability

The state of context management in 2026 is defined by a paradox: **organizations have never been more confident in their AI readiness, and they have never been more willing to invest in closing the gaps they're beginning to see.**

That combination is exactly what a maturing market looks like. The question is no longer whether context management matters. It is how quickly your organization can close the distance between the context infrastructure it has and the context infrastructure production-scale AI requires.

About DataHub

DataHub transforms enterprise data into trusted context, enabling intelligent decision making by humans and AI agents. As the only open-source, AI-native data catalog, DataHub Cloud unifies AI-powered discovery, governance, and observability —helping enterprises scale data operations efficiently while ensuring data quality, compliance, and AI-readiness across their entire data estate.

Backed by a thriving open-source community of 14,000+ members and proven by over 3,000 organizations worldwide, DataHub offers enterprise-grade security, flexible deployment options, and battle-tested APIs that power world-class data platforms. Organizations rely on DataHub to accelerate time-to-value from their data investments, ensure AI system reliability, and implement unified governance—enabling AI & data to work together and bring order to data chaos.

DataHub

See DataHub Cloud in action

BOOK A DEMO

Join the conversation

- Follow DataHub on LinkedIn
- Join the DataHub open-source community on Slack
- Visit us at DataHub.com