


The background is decorated with several 3D wireframe cubes of varying sizes and orientations, some floating in the light blue space and others resting on a diagonal plane that splits into orange and teal sections at the bottom. Small orange squares are also scattered throughout the design.


# 2026 trends and predictions report

Architecture for the agentic era: How AI will reshape data, security, and observability



## CONTENTS

- 03 Introduction: The coming shockwave of AI-generated data
- 04 Prediction: AI-assisted data saturation breaks observability budgets
- 06 Prediction: Agentic AI pipelines become attack surfaces
- 07 Prediction: Production AI deployments fail because of fragile data platforms
- 08 Prediction: AI-driven vendor consolidation upends procurement
- 09 Prediction: Private credit turmoil will curb AI data center growth and force leaner AI innovation
- 10 Conclusion: Architect an AI-ready future



## INTRODUCTION

# The coming shockwave of AI-generated data



AI has quickly evolved from experimental copilots to **fully agentic systems**: autonomous entities that generate, consume, and analyze telemetry at unprecedented scale.

As these AI systems proliferate across IT and security environments, they're creating a data shockwave few infrastructures are ready to handle.

The world's telemetry pipelines, built for human-speed operations, are now buckling under machine-scale workloads. Data is exploding faster than budgets can grow, security attack surfaces are multiplying, and legacy observability tools are reaching their breaking point.

**The result?** Too much data is too costly.

In 2026, IT, security, and observability leaders face a pivotal moment: Modernize architectures to handle AI-scale data efficiently, or face spiraling costs, blind spots, and systemic fragility as agents become a core part of enterprise operations.

This report explores five predictions that reveal how AI will reshape the economics, security, and architecture of data in the coming years — and what leaders must do now to prepare.




## AI-assisted data saturation breaks observability budgets

**For CIOs, the implication is clear:** The old “collect and keep everything” mindset is financially unsustainable.

By 2027, 35% of enterprises will see observability costs consume more than 15% of their overall IT operations budget, forcing CIOs to adopt data tiering strategies to cope.

Observability has become indispensable in modern IT, but the costs are spiraling out of control. Analyst firms such as Gartner have tracked the steep rise in telemetry data volumes fueled by cloud-native architectures, microservices, and distributed applications. The result: The median spend on observability platforms now exceeds \$800,000 annually for a single vendor, with year-over-year increases topping 20%.

Today, most enterprises dedicate between 3–7% of their IT operations budget to monitoring and observability. By 2027, however, over a third of enterprises will see these costs surpass 15% of their overall IT operations spend. This shift represents more than ballooning line items. It signals a structural imbalance between data growth and IT budgets.



The arrival of AI intensifies this imbalance. Large language models, machine learning pipelines, and the high-performance infrastructure needed to train and run them generate massive amounts of telemetry data. Organizations face not only need to monitor these new workloads but also must retain data for security investigations, model performance tracking, and regulatory compliance. Telemetry growth, already acute, will reach crisis levels.

For CIOs, the implication is clear: The old “collect and keep everything” mindset is financially unsustainable. Enterprises can no longer afford to funnel every log, metric, and trace into premium platforms where costs scale linearly with data volume. Instead, forward-looking IT leaders are turning to data tiering strategies.



## AI-assisted data saturation breaks observability budgets *(continued)*


**Data tiering** allows enterprises to separate high-value, frequently accessed telemetry from lower-value or longtail data.

Critical, real-time streams can flow into high-performance observability platforms for rapid detection and response. Less urgent or historical data can be routed into lower-cost storage tiers — such as object storage or security data lakes — where it remains accessible for compliance audits, investigations, or to replay into analytic tools when needed.

The payoff is twofold: Organizations regain control of their IT budgets and still meet the operational, security, and compliance requirements of a data-intensive era. As observability costs rise, CIOs who fail to implement data tiering risk not only budget overruns but also diminished agility in managing AI-driven environments.

### DATA TIERING

## Regain control of IT budgets





# Agentic AI pipelines become attack surfaces

The adoption of AI in cybersecurity is projected to drive **~\$93B in additional spend by 2030.**

In 2026, 20% of Fortune 2000 companies will suffer material security incidents that originate from compromised or manipulated AI-driven observability or control pipelines.

The adoption of AI in cybersecurity is projected to drive ~\$93B in additional spend by 2030. Already today, teams are experimenting with early AI SOC and AI SRE prototypes for everything from incident detection to automated runtime error resolution. If these prototypes live up to expectations, their advantages are clear, ranging from faster remediation and reduced downtime, to lowered stress on overworked IT and security teams.

However, a key element of effective AI adoption in IT Ops and the SOC is high quality, trustworthy telemetry data used to both train and execute models. Because of the importance of data in training and operating AI models, attackers are leaning in and treating telemetry data as the next-generation avenue of supply chain attack.

This is not theoretical. MITRE ATLAS, a knowledge base of adversarial attack patterns against AI systems, has already documented data

poisoning and model evasion methods that map directly to telemetry workflows. These attack methods alter the cybersecurity surface area from common targets, like endpoints and user credentials, to the telemetry workflows feeding AI-based systems.

Black-box AI-driven telemetry data pipelines will facilitate this data poisoning attack vector. By removing essential human insight, understanding, and experience from telemetry data management, black-box approaches invite manipulation with little to no oversight over the kind of data ingested into downstream systems, where it can impact model training and use. Worse, these compromises can escalate into corrupted control planes, leading to downtime, stealthy data exfiltration, or even manipulated incident response.

Avoiding this attack vector requires keeping the human in the loop, supported by instrumentation, metrics, and monitoring, when applying AI to telemetry data pipelines.



# Production AI deployments fail because of fragile data platforms

By 2027, **90% of production AI deployments will fail** due to unsolvable scaling challenges at the data tier.

# 90%


Obstacles around data quality, organization, and governance have hampered early AI efforts in half of companies. It isn't a surprise, as these same difficulties were common during the Big Data era, when heavily siloed enterprises first tried leapfrogging from rudimentary dashboards into predictive analytics without modernizing their data strategy. Admittedly, it is disheartening to realize that enterprises have made little progress at managing their data over the intervening years, but here we are. Again.

While AI developers grapple with the blocking and tackling of data management for their early prototypes, an insurmountable problem awaits them on the other side of the AI deployment lifecycle: production-class workloads. Today, AI developers are mainly deploying small-scale prototypes, in which a few dozen agents query underlying datastores, gather large amounts of data, and share it with users and fellow agents.

Tomorrow, these agents will number in the thousands, or tens of thousands, with each of them requesting data dozens or hundreds of times per second as they execute their tasks.

This sudden spike in demand will far outstrip the capacity of the underlying datastores to accommodate it. Datastores ranging from log analytics platforms and object stores to databases, data warehouses and data lakes were architected mainly for batch processing and human interaction speeds. These legacy tools are ill-equipped for an agentic future that requires both massive levels of concurrency and consistent data throughput. As enterprises move from the prototype phase into production, they run a significant risk of systemic failures — not because of the LLMs, but because of the inadequate data infrastructure used to train and run them.






## AI-driven vendor consolidation upends procurement

# 15%

By 2027, 15% of enterprises will have switched their primary observability or security platforms — not because of features, but because AI agent ecosystems **forced vendor consolidation**.



For years, the market has balanced between vendor platformization and enterprise preference for best-of-breed freedom. The promise of agentic AI is breaking that balance. Vendors are embedding autonomous agents for detection, response, and analytics that depend on proprietary data models and closed APIs. To ensure performance and protect revenue, these vendors are building AI assistants that work only inside their own ecosystems — effectively locking out third-party agents and integrations.

Enterprises that once connected multiple tools now find their data fragmented and workflows brittle. Agents trained on a single platform will be unable to reason across diverse datasets without custom connectors that degrade reliability. What began as innovation becomes a constraint, pushing teams toward consolidation under vendors that offer “fully integrated” AI-driven stacks.

This shift mirrors earlier platformization waves in observability, but the stakes are now higher. Agentic AI thrives on broad, contextual

data access — something closed ecosystems inherently limit. Enterprises that lose data control also lose the ability to experiment with or adopt rapidly evolving AI models. While vendor lock-in may simplify short-term adoption, it risks strategic stagnation as AI capabilities advance beyond any single provider’s boundaries.

The winners in this transition will prioritize data independence — decoupling data control from platform choice. Open telemetry pipelines, schema-on-need architectures, and neutral storage layers will become essential safeguards. These capabilities let organizations adopt or replace agents as needed without overhauling their data foundations.

In the near term, however, agent ecosystems will dictate consolidation. The gravitational pull of AI-native platforms will outweigh interoperability ideals, driving one in six enterprises to abandon long-standing vendors simply to stay compatible in an agent-dominated future.



# Private credit turmoil will curb AI data center growth and force leaner AI innovation

Instead of relying on ever-larger data centers, AI firms will focus on **doing more with less**, developing more intelligent algorithms, optimizing data usage, and creating systems that learn efficiently.

AI data center expansion is fueled by private credit, which comes from largely unregulated investment funds rather than traditional sources of corporate funding, such as banks and bonds. Because private credit operates outside traditional banking oversight, it carries higher risk from limited transparency, concentrated exposure, and the potential for defaults or liquidity shortages when markets tighten.

Private lenders have been eager to fund massive AI projects that need billions of dollars to build power-hungry, high-performance data centers. A recent example is a joint venture between Meta and Blue Owl Capital to develop the Hyperion campus in Richland Parish, Louisiana. This source of easy money may not last.

Many private credit funds are now facing pressure. Some have overextended, others are struggling to raise new capital, and a few are close to bankruptcy. If these problems spread, the amount of available funding will shrink. When that happens, data center projects that depend on private credit will be delayed or cancelled. The entire pace of AI infrastructure expansion could slow down as funding dries up. All it takes is a severe recession or a black-swan event to trigger a potential credit meltdown.

In the short term, this will mean fewer new facilities and tighter access to GPUs and compute power. Companies that counted on endless capacity will be forced to adapt. In the long run, this could change how AI is built and deployed. Instead of relying on ever-larger data centers, AI firms will focus on doing more with less, developing more intelligent algorithms, optimizing data usage, and creating systems that learn efficiently.

A slowdown in private credit will hurt AI's rapid expansion, but it may also make the industry stronger. When easy money dries up, creativity and efficiency take its place. The next generation of AI may not come from a giant data center; it may come from a smarter idea.



## CONCLUSION

# Architect an AI-ready future



**AI is no longer a distant promise** — it's here, embedded in every layer of infrastructure, from log ingestion to incident response. But most organizations are still managing data architectures designed for a pre-AI world.

To thrive in the agentic era, enterprises need more than smarter models — they need smarter telemetry pipelines, and architectures that can:



**Control costs** by intelligently routing, reducing, and storing only what matters.



**Enhance security** by maintaining visibility and integrity across AI-driven data flows.



**Increase agility** by embracing open, interoperable platforms rather than proprietary AI silos.

The next generation of observability and security leaders won't just monitor systems — they'll engineer the data foundations that make intelligent automation possible, trustworthy, and affordable.

Those who succeed will turn data chaos into clarity, and AI complexity into competitive advantage.

## REFERENCES

Gartner, June 25, 2025, *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027*

Mitre Atlas, 2021-2025, *Atlas Matrix*

Grand View Research, 2025, *AI In Cybersecurity Market (2025 - 2030)*

Avepoint, 2024, *AI and Information Management Report 2024*

Federal Reserve Bank of Boston, José L. Fillat, Mattia Landoni, John D. Levin, and J. Christina Wang, May 21, 2025, *Could the Growth of Private Credit Pose a Risk to Financial System Stability?*

abfjournal, Rita Garwood, May 20, 2025, *AI Boom Reportedly Presents \$1.8T Opportunity for Private Credit*

McKinsey & Company, April 28, 2025, *The cost of compute: A \$7 trillion race to scale data centers*

Reuters, Matt Tracy, May 7, 2025, *Moody's warns of risk posed by rising retail exposure to private credit*



Cribl, the Data Engine for IT and Security, empowers organizations to transform their data strategy. Powered by a data processing engine purpose-built for IT and Security, Cribl's product suite is a vendor-agnostic data management solution capable of collecting data from any source, processing billions of events per second, automatically routing data for optimized storage, and analyzing any data, at any time, in any location. With Cribl, IT and Security teams have the choice, control, and flexibility required to adapt to their ever-changing data needs. Cribl's offerings — [Stream](#), [Edge](#), [Search](#), and [Lake](#) — are available either as discrete products or as a holistic solution.

Learn more: [cribl.io](https://cribl.io) | Try now: [Cribl sandboxes](#)

Join us: [Slack community](#) Follow us: [LinkedIn](#) and [X \(Twitter\)](#)

©2025 Cribl, Inc. All Rights Reserved. 'Cribl' and the Cribl Flow Mark are trademarks of Cribl, Inc. in the United States and/or other countries. All third-party trademarks are the property of their respective owners.