

The Agentic Shift: Redefining Observability for the AI Era

How AI is reshaping systems — and
what observability must become

splunk>
a CISCO company



The rules didn't change. They stopped applying.

Remember the first time you watched *The Fifth Element*? For some of us, it was the first time a movie managed to truly capture the chaos of a city in flight.

We take it all in alongside an understandably overwhelmed Leeloo. In every direction, lanes of flying cars intersect one another. Buses scramble up and down buildings, and wait ... was that a flying frigate?

It was an effective moment. From that very first glimpse of a 23rd century NYC, we know for certain that the rules of the road — at least as we know them — do not apply here.

That's the shift we're living through right now. AI didn't just make our systems faster. It made them different.



Observability's next evolution

Before AI, it was a lot simpler to answer a critical question:
If something goes wrong, what's the impact on the business?

Infrastructure served as a reliable proxy for business health. Criticality was understood from the bottom up, complexity lived in silos, and people connected the dots. In most cases, you could look at infrastructure to determine what was broken, what wasn't, and how it affected the bottom line. That's the environment traditional monitoring was built for — one where systems behaved predictably and signals were clear.

As systems grew more distributed and complex, monitoring evolved into observability, giving teams the ability to understand environments they could no longer fully predict.

Now, that complexity has compounded again.

AI introduces a new class of systems that are non-deterministic, adaptive, and increasingly autonomous. As a result, system health and business impact are no longer tightly aligned.

A system can appear healthy while quietly degrading or produce an answer that looks correct while being fundamentally wrong (anyone who's asked a chatbot for math homework help knows this to be true).

The challenge isn't just complexity. It's that the signals we trust no longer tell the full story.

To observe AI effectively, observability must evolve. That evolution is **agentic observability** — a term we're starting to hear more often but also one that isn't always clearly defined.

What we mean by 'business impact'

Business impact connects system behavior to what actually matters to the organization.

This can include things like:

- Checkout completion
- Revenue at risk
- Customer experience

In other words, it's not just what's happening in the system, but how it affects the business.



What is agentic observability?

Depending on who you ask, agentic observability can mean different things. For some, it means observing AI. For others, it's about using AI to improve observability — helping detect, investigate, and resolve issues.

Both point to the same underlying shift.

Agentic observability extends the principles of observability into the agentic AI era. It does this not just by adding automation, but by changing how systems are understood and operated.

At its core, agentic observability is built on three capabilities.

It starts with **aligning performance problems to business impact**. Without that foundation, teams or AI agents can't effectively prioritize what matters, and even the most advanced AI-driven capabilities risk accelerating the wrong work.

From there, **AI should help teams act** — improving how issues are detected, investigated, and resolved. This shifts teams away from reactive, fire-drill operations toward a more proactive and intentional model where many issues are resolved before customers ever feel the impact.

And as AI becomes more central to how systems operate, **teams must also be able to observe and trust those systems**. Without that trust, the risk of using AI to innovate can outweigh the potential benefit. The result? AI adoption is slowed or stopped altogether.

Together, these capabilities help teams move from seeing signals to understanding them in minutes— and from understanding to action.

In practice, this means using AI agents and a unified data foundation to find, fix, and help prevent issues, while also observing and governing AI systems. The result is a more focused approach to operations that helps teams spend less time chasing signals and more time innovating.

We'll explore these capabilities in more detail throughout this guide. But before we get there, it's worth grounding this shift in two pressures: scale and trust.



Pressure 1: Issues of scale

As AI becomes more deeply embedded across systems and workflows, the first challenge is scale.


What once felt manageable starts to converge into something harder to contain, driven by sustained workloads and rapidly increasing data volumes.

Infrastructure is pushed harder with AI workloads that are more demanding than before. They no longer arrive in predictable bursts, but run continuously.

At the same time, data continues to grow faster than teams can realistically process or understand. Signals (and costs) multiply, but context doesn't always keep up.

As a result, even simple questions become harder to answer. What's happening? Where is it happening? And what actually matters?

A new level of inferencing demand with Agentic AI



Pop quiz: What does “non-deterministic” mean?

You run the same action twice ... and get two different outcomes.

Why?

- a. Buttercup spilled coffee on the server again.
- b. Your *We go to 11* t-shirt is influencing system performance.
- c. The system doesn't always produce the same result from the same input.
- d. Someone definitely changed something ... maybe ...
frantically messages Team Lead Bob, drinking a daiquiri in Bora Bora.

Answer: c

Non-deterministic systems, like many AI models, don't behave the same way every time. The same input can lead to different outputs depending on context, timing, or internal state.

That's what makes them powerful ... and what makes them harder to predict, test, and trust.



Pressure 2: Issues of trust

Today's companies aren't just using AI to move faster. They're using it to innovate and create entirely new kinds of applications that power exceptional customer experiences.

As organizations move from using AI to depending on it, the expectations placed on these systems change as well. They're no longer just expected to run — they're expected to be right and be safe to rely on. Availability does not guarantee outcomes, and performance does not guarantee quality.

Since many AI models are non-deterministic, the signals teams rely on become harder to interpret and harder to trust. An output may look correct and still be fundamentally wrong. A system may appear healthy while quietly drifting away from expected behavior. Issues don't always announce themselves; they may surface as subtle changes that are easy to miss and difficult to diagnose. In some cases, the impact isn't just performance. It's exposure.

Instead of outages, teams are increasingly dealing with silent inaccuracies:

- A system may remain fully operational while producing incorrect or incomplete results.
- An AI agent might generate a confident, well-structured response that is factually wrong.
- A recommendation engine might begin to influence customer behavior in ways that reduce conversion or erode trust — without triggering any traditional alerts.

From a system perspective, everything appears nominal. From a business perspective, it is not.

Without a clear way to understand how systems are behaving, why decisions are made, and how those decisions impact outcomes, teams are left to operate with uncertainty, even when everything appears to be working as expected. For some organizations — especially in highly regulated industries like finance and healthcare — that uncertainty becomes a barrier. The risk is too great. Without trust, they may hesitate to adopt AI at all, limiting their ability to take advantage of its benefits.

In social media buying, a recommendation system can keep impressions flowing and clicks coming while quietly drifting toward lower-quality audiences — driving up acquisition costs and eroding return on ad spend without triggering a single alert.



The emergence of agentic observability

The pressures facing modern systems don't point to a single missing tool or feature. They point to a shift in how systems need to be understood and operated.

Observability isn't being replaced — it's being pushed to evolve. And not all approaches will keep pace.

Just as observability emerged to make sense of distributed systems, the rise of AI is pushing it to evolve again. The challenge is no longer just to collect and correlate signals. It's to connect them to outcomes, interpret them in context, and act on them in time to make a difference.

This is where agentic observability begins to take shape.

There's still some ambiguity around what that means. For some, it's about adding AI to existing workflows — automating queries or accelerating investigation. Those steps matter, but they don't fully address the shift underway.

Agentic observability represents something more fundamental.

It extends observability beyond visibility into understanding, action, and risk reduction — helping teams de-risk AI-powered systems so they can confidently use them to move the business forward.

It starts with the ability to prioritize based on business impact, creating a foundation for AI-driven operational assistance and deeper visibility into AI-powered systems — ensuring teams can focus on what matters most, act on it effectively, and trust the systems they rely on.

The goal isn't to replace human judgment, but to augment it — helping teams bridge the gap between visibility, understanding, and action in increasingly complex environments.



Three essential capabilities of agentic observability

Align to business impact

Observability only becomes effective when teams can prioritize based on what matters to the business. Aligning system behavior to outcomes like revenue and customer experiences gives teams the context to focus on the right issues, not just the loudest ones.

Without that foundation, even the most advanced AI-driven capabilities will fall short — accelerating activity without improving results.

1

Fix and prevent with AI

Once teams can prioritize what matters, AI agents can help them act on it.

AI agents help teams move beyond reactive workflows by assisting across the incident lifecycle — from detection and investigation to remediation and prevention. Instead of chasing alerts, teams can resolve the issues that matter most, faster and more consistently.

Over time, this not only accelerates response, but reduces the likelihood of issues recurring.

2

Observe AI systems

Observing AI systems provides visibility into how models and agents behave in terms of quality, cost, and impact — helping teams identify issues early, reduce risk, and use AI with greater confidence.

3

CAPABILITY 1:

Align to business impact

Agentic AI helps teams connect and synthesize telemetry across their entire environment — helping them prioritize issues based on business impact and focus on what matters most.

This requires more than just isolated signals. It depends on having access to connected data from across infrastructure, applications, and networks — all in one proverbial room — so teams can understand how system behavior translates into real outcomes.

Without that foundation, even the most advanced AI-driven capabilities fall short — accelerating activity without improving results.

What this means for teams

As systems grow more complex and data volumes increase, it becomes harder to determine what actually matters. Aligning observability to business impact helps organizations focus on the signals that affect customer experience, performance, and revenue — rather than reacting to everything.

For teams, this means shifting from responding to alerts to understanding which issues truly require attention, and why.

Where teams can get stuck

In many environments, system data is fragmented across tools and domains. Teams can see what's happening within individual systems, but lack the connected view needed to understand how those signals impact the business, making it difficult to prioritize effectively.


Where to start

- Define a small set of key outcomes to anchor observability efforts
- Map signals across infrastructure, applications, and networks to a critical user journey or workflow
- Use that context to surface and prioritize the issues with the greatest business impact

Example: From signal to impact

A spike in API latency may not trigger a critical alert on its own. But when connected to application performance and user behavior, it may reveal a checkout slowdown that is directly impacting conversion and revenue.

Instead of responding to dozens of alerts, teams can focus on the few issues that matter most — and act on them first.



CAPABILITY 2: Fix and prevent with AI

Agentic AI enables teams to move beyond reactive workflows by helping to interpret signals, guide investigations, and support action across the incident lifecycle — from detection through remediation and prevention.

What this means for teams

Modern environments were already complex. With AI, that complexity has expanded — more services, more dependencies, more signals, and more ways for issues to emerge and propagate across systems.

What used to be a two-dimensional problem is now multi-dimensional (think a present-day freeway versus the flying frigates in 23rd-century New York). An issue in application performance might originate in the network, while a spike in infrastructure usage might trace back to an AI-driven workflow. The number of possible connections — and failure points — grows fast.

At that scale, manual investigation simply doesn't hold up. Even large teams can struggle to identify the issues that matter, resolve them quickly, and do so cost effectively.

Agentic AI-driven assistance helps teams operate within that complexity. It can summarize incidents, connect related events across domains, and guide investigations toward likely root causes and next steps — helping teams move faster from detection to resolution, and reduce the likelihood of issues recurring.

Where teams get stuck

Execution can break down at scale.

As environments grow more complex, signals become harder to correlate across systems and domains. Investigations stall as teams move between tools, manually piecing together context to understand what's happening.

Without the ability to connect signals quickly and consistently, even experienced teams struggle to resolve issues efficiently.


Where to start

- Introduce AI into investigation workflows to summarize alerts and correlate signals
- Use AI to surface likely root causes and suggest next steps
- Begin with guided remediation, where AI supports decisions rather than replacing them

Example: From signal to resolution

A spike in infrastructure usage appears alongside intermittent performance issues across several services. On its own, each signal seems manageable. But together, they point to a misconfigured AI-driven workflow repeatedly triggering resource-intensive processes.

Instead of manually tracing the issue across systems, AI can correlate the signals, identify the likely root cause, and guide next steps — helping teams resolve the issue quickly and prevent it from recurring.



CAPABILITY 3: **Observe AI systems**

Agentic AI enables organizations to observe not just whether systems are running, but how AI-powered systems behave — helping teams build trust and reduce risk as they rely more heavily on AI.

What this means for teams

As AI becomes embedded in how systems operate, it introduces new uncertainty. These non-deterministic systems can produce outputs that appear correct while being fundamentally wrong or shift behavior over time without obvious signals.

This makes them harder to monitor using traditional approaches — and harder to trust.

Observing AI systems helps close that gap. It requires expanding observability beyond applications and infrastructure to include the full AI stack — from GPUs and model endpoints to vector stores, orchestration layers, and the agent frameworks coordinating work.

It also means capturing new kinds of telemetry. In addition to traditional signals, teams need visibility into prompts and responses, tool calls, context windows, and how behavior varies across users, workflows, or environments.

With that broader view, teams can begin to monitor not just performance, but quality and safety. This includes tracking issues like model drift, hallucination risk, policy violations, data leakage, and prompt injection attempts — all the while tracking cost and efficiency signals like token usage, model performance, and the value generated by AI-driven interactions.

This level of visibility isn't optional. As AI systems take on a more active role in customer experience and decision making, they need to be observed just as strictly as application and infrastructure layers, with added focus on quality, trust, and cost over time.

With that visibility, organizations can use AI more confidently — unlocking its potential to drive innovation while reducing the risk of unintended outcomes.

Where teams get stuck

Understanding AI behavior is not straightforward.

Traditional observability signals don't capture whether outputs are correct, whether models are drifting, or whether decisions align with expectations. Without that visibility, teams are left to infer behavior from incomplete signals — making it difficult to diagnose issues or validate performance.

Where to start

- Instrument AI systems alongside existing infrastructure
- Track outputs and monitor performance trends over time
- Establish checks for quality, consistency, and cost.

Example: From signal to insight

An AI-powered support assistant continues to respond quickly and accurately based on traditional metrics. But over time, its responses begin to drift, introducing subtle inaccuracies that impact customer trust.

By observing output quality and behavior over time, teams can detect the issue early, adjust the system, and prevent broader impact — maintaining both performance and confidence.

What comes next

The shift to agentic observability isn't something organizations will decide to adopt someday. It's already happening in the systems being built, the expectations being set, and the pace at which teams are expected to respond.

The question isn't whether observability will evolve. It's whether teams will evolve with it.

For many organizations, that starts with small, practical steps: Connecting system behavior to real business outcomes, introducing AI into investigation workflows, and gaining visibility into AI systems. Over time, those steps build toward something more powerful: a model of observability that helps teams not just see what's happening, but understand it and act with confidence.

But not every approach — or every provider — will keep pace.

As these demands grow, observability providers will need to ensure their observability strategy supports all three capabilities. Without that foundation, teams risk falling behind, even as their systems continue to evolve.

You don't have to solve this all at once.

But getting started matters.

See how leading organizations are approaching observability in the AI era.

Get the analyst report

Talk through your approach to agentic observability.

Contact a Splunk expert



Splunk, Splunk>, Data-to-Everything, and Turn Data Into Doing are trademarks or registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners. © 2026 Splunk LLC. All rights reserved.

26_CMP_ebook_The-agentic-shift-redefining-observability-for-the-AI-era_v8

