


The reliability stack is changing.

Autonomy changes the operating model—not the accountability for outcomes.



Autonomous IT is outgrowing human-only correlation

More changes. More signals. More dependencies. Less time to assemble the story by hand.

3.7×

average telemetry growth
in the last year

9.5×

expected agentic workload
growth
over the next two years

The constraint is no longer seeing a signal.

It is building enough context to decide—before the
next change arrives.

Reliability teams own outcomes— inside an operating model built for handoffs

The protagonist is still human: SRE, platform and operations leaders remain accountable when automation moves faster than the queue.

- 01  dashboard see the symptom

- 02  alert queue collect fragments

- 03  manual triage rebuild context

- 04  handoff find authority

The job is no longer just to explain what happened. It is to
govern what happens next.

Reliability must move from observing systems to governing action

The shift is not a smarter dashboard. It is a closed loop with evidence, context and explicit authority.



That is where reliability becomes a control plane for autonomous IT.

The market is funding the shift

\$2.25B


disclosed capital

57

financings

45

companies



+28%

2025 vs 2024


94%

of 2025 already in 2026 YTD

Capital signals conviction, not efficacy. Public-announcement floor; 2026 through Sep 10.



Capital is funding the chain from telemetry to safe action



The foundation still absorbs most capital. The emerging bet is to close the loop around it.

Analytical classification by primary problem stated in each announcement; categories can overlap in practice.

AI on top of observability is not the same as a new operating model

Adding AI to first-generation observability can improve discovery and explanation. It may still preserve workflows, incentives and authority boundaries designed before agentic systems.

Agentic systems raise a different design question

- 01 — usable telemetry
- 02 — causal context
- 03 — verifiable reasoning
- 04 — governed action
- 05 — operational learning

The fair counterpoint

Established platforms bring scale, integrations, operational maturity and customer trust. New architectures still need to prove efficacy and broad autonomous remediation.

When reliability can act, security becomes the boundary

Moving from recommendation to change turns authority into part of the reliability system.

OBSERVE & RECOMMEND

read signals
investigate
explain
recommend

ACT IN THE ENVIRONMENT

restart
roll back
change configuration
open a pull request

BOUNDARY

— identity

— least privilege

— containment

— authorization

— auditability

— rollback

Operational autonomy needs explicit limits, reversible actions and human accountability.

New problems

require us to evaluate new architectures—not only new features.

Reliability remains the promise: understand the system, protect the outcome, govern the change.

What evidence would you require before an agent moves from recommendation to action?