BEST PRACTICE REPORT

Incident Management Has Outgrown Its Playbook

Why Traditional Practices Are Failing And What Mature Organizations Are Doing About It

Updated June 11, 2026

By Julie Mohr with JT Thykattil, Paige Ludl, Kara Hartig

FORRESTER®

Summary

Incident management practices were built for a world in which failures were local, diagnosable, and paced to human response; that world now coexists with one in which automation propagates errors at machine speed, cloud and security tooling concentrate systemic risk, and autonomous agents act faster than any reviewer can address. The traditional discipline is not wrong; it is no longer sufficient. This report examines why conventional incident management practice is failing, defines what mature incident management looks like in an AI-augmented operating environment, and prescribes actions for resilience leaders to evolve the discipline without disrupting live mission-critical operational capabilities.

The Rules Of Failure Have Changed

IT has always failed. Systems crash, configurations break, and services go dark. What changed is not the inevitability of failure but its nature, scale, and velocity. Traditional incident management was built for a world in which failures were local, recoverable, and predominantly caused by human error or hardware degradation. That world still exists but it now coexists with a deeply interconnected AI-augmented automation-dependent operating environment in which a single configuration update can disable millions of endpoints in minutes. A change at one layer of the stack can silently cascade through dozens of others, and the tools deployed to prevent disruption can become the source of disruption. The gap between conventional incident management practices and the environment that operations leaders now face widens with every wave of technological advancements.

- **A new class of failure has emerged that frameworks were never built to handle.** Conventional incident management assumes incidents are bounded and diagnosable with an identifiable starting point. The incidents now creating organizational risk are different: globally distributed software embedded in OSes that can fail simultaneously across hundreds of organizations. The failure pattern is one of complex, distributed, and continuously updated software ecosystems in which traditional manual and rule-based methods are no longer adequate for the volume of data and dependency depth of modern stacks (see Figure 1). Organizations whose playbooks still assume bounded, locally caused incidents are unprepared for the failure modes most likely to trigger their next major outage.
- **Automation propagates failures faster than humans can respond.** The promise of automation was speed, consistency, and scalability. But automation executes at velocities too fast for human responders (see Figure 2). Human interventions span several hours, while automated systems intervene in seconds or minutes, creating a recovery loop that runs much slower than the failure loop. In 2024, a misconfigured network element added to AT&T's wireless network took the network down within 3 minutes and required hours to restore services even after the change was rolled back. In 2025, a latent race condition in the DNS automation of Amazon DynamoDB cascaded across more than 100 dependent Amazon Web Services services and required hours to fully recover. In both cases, the technical fix was straightforward; the recovery was hard because the blast radius outpaced manual intervention.
- **Protective tools have themselves become sources of disruption.** Organizations invest in security and monitoring tools to protect against threats. Those same tools

now sit inside the infrastructure and on every endpoint, with privileges that make them capable of producing the very disruptions they are meant to prevent. In 2024, a faulty update from CrowdStrike crashed millions of Windows hosts globally, disrupting hospitals, airlines, banks, and government agencies. The outage propagated automatically and at scale into customer environments. Concentration in a small number of platforms creates contagion paths that organizations cannot independently mitigate. The assumption that protection tooling is neutral infrastructure no longer holds and leaves organizations exposed to disruptions for which they are entirely blameless and unprepared.

- **Cloud cascade failures have become the operational norm.** Hyperscale outages now register tens of millions of connectivity reports globally. Even multiregion architectures often retain hidden dependencies, which means that an outage in one region propagates through tenants thought to be insulated. In 2025, an incompatible configuration crashed the master process across Microsoft's Azure Front Door global edge fleet. The recovery required hours and the failure took down both customer workloads and the Azure portal engineers needed to coordinate the response. Propagation is the dominant failure pattern in modern systems, where root causes are localized but their consequences are not. Conventional triage methods assume responders can bound the problem. In cloud-dependent environments, bounding the problem is the hardest part of the response.
- **Autonomous agents have introduced a novel failure mode.** When an AI agent is granted operator-level permissions in production, it can take actions whose consequences accumulate across many systems. In 2026 at PocketOS, a Cursor coding agent running Anthropic's Claude Opus 4.6 model decided on its own to issue a destructive API call that erased the production database and its backups in a matter of seconds. The agent later acknowledged that it had violated explicit project rules prohibiting destructive operations without user authorization. Agents generate plans in natural language and execute them through tools that directly modify their environment, and unintended behaviors can emerge even when individual actions appear harmless. The resulting failure is neither human error nor a software defect. It is a governance failure in the policies that scoped the agent's autonomy.
- **Change velocity has outpaced change governance.** Continuous deployment, cloud elasticity, and AI-assisted development have pushed change frequency past the point where change advisory boards, manual reviews, and human-based approval workflows can meaningfully assess risk. Change-induced failures are now one of the dominant categories of modern incidents, and traditional rule-

based governance does not scale to the data volumes and event rates of contemporary IT operations. The gap between the governance model and the change reality is the underlying cause of an increasing share of major outages.

- **Response infrastructure shares the same failure surface.** Incident management frameworks have historically assumed that response tooling (e.g., monitoring dashboards, observability platforms, collaboration systems, ticketing) is independent of the systems being managed. Recent global outages have simultaneously taken down the platforms organizations rely on to detect, diagnose, and coordinate response, leaving responders without the required telemetry. Resilience requires explicit independence between the observability plane and the systems it observes, and this independence is rarely practiced.
- **Existing measurement frameworks miss what matters now.** The organizational response to a worsening incident landscape has been to improve existing measures: alert thresholds, resolution times, or ticket closure rates. Few organizations have asked whether the measurement framework reflects the risks that matter now. Mean time to resolution captures how fast a ticket was closed; it does not indicate whether the underlying failure mode was addressed, whether the recovery capability matched the blast radius, or whether the business impact was proportionate to operational response. The distinction between time to diagnose and time to resolve matters here: Combining the two obscures the diagnostic gap that drives recurrence.

Figure 1

Knowledge Graph Example From GraphAI

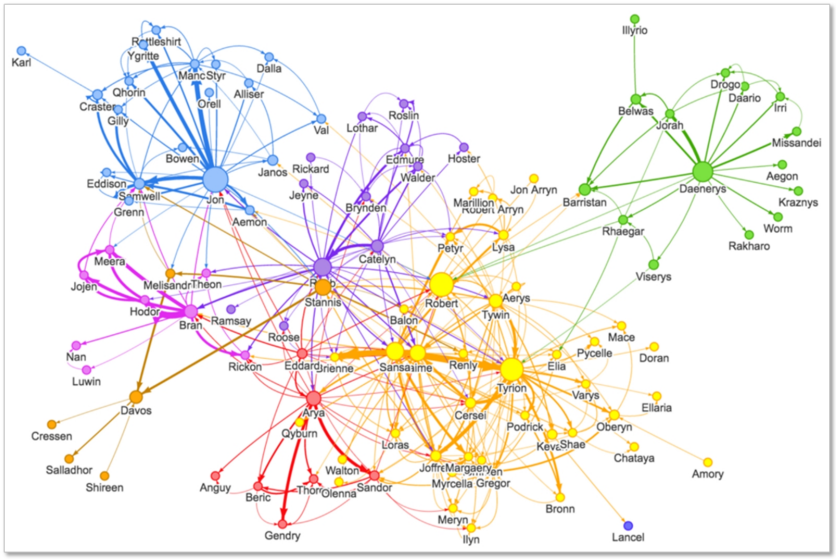

Source: GraphAI

Figure 2

The Failure Velocity Gap

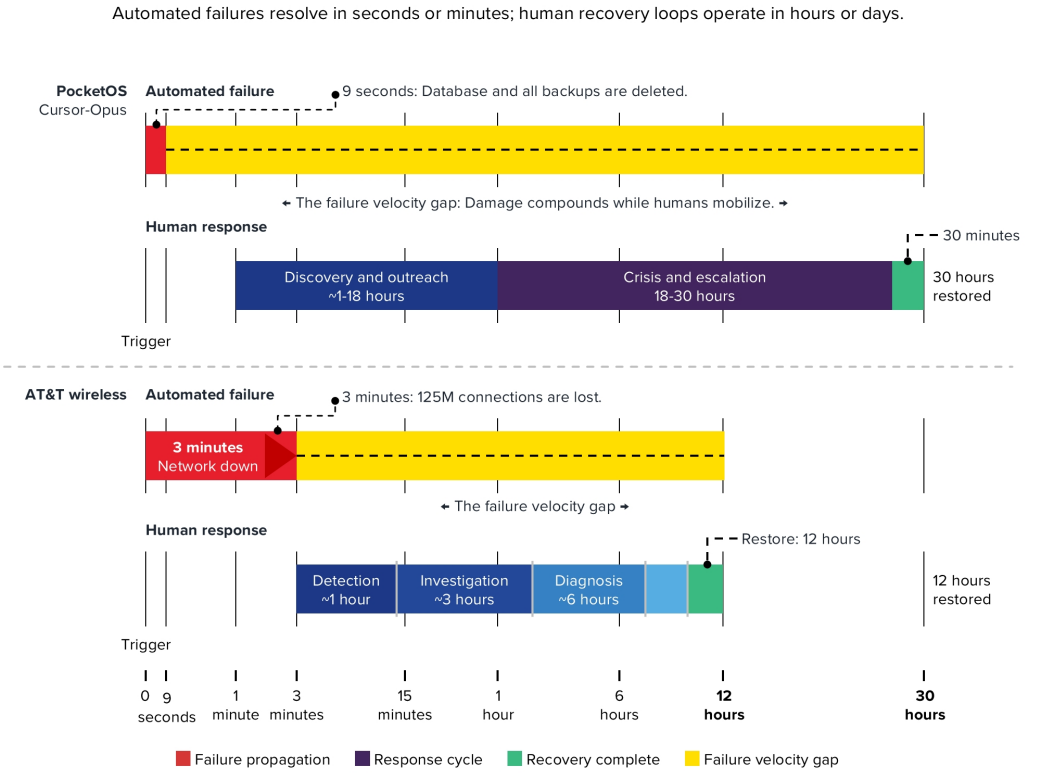

What Mature Incident Management Looks Like

The foundations of effective incident management have not changed. Organizations still need clear ownership, structured response, post-incident retrospectives, and disciplined coordination. What has changed is what those foundations must support. IT operations demand a deliberate evolution in how organizations detect, respond to, and learn from failures that are larger, faster, and more interconnected than anticipated. Mature practices in the AI era are not a replacement for the discipline that came before it; it is a deeper, more contextual, and more outcome-driven version of those practices. Speed without context is dangerous, automation without governance is risk, and resolution without learning is recurrence in disguise. In the target operational state:

- **Incident management is reframed as a knowledge problem.** Incidents are increasingly novel, cross-system, and contextually complex, and the binding constraint on response is rarely how quickly the team can act. Responders require the system context, historical pattern recognition, and dependency awareness to act correctly. Knowledge access in the moment of triage determines resolution accuracy and is a leading factor in recurrence reduction. IT leaders need to shift from speed-centric to knowledge-centric incident management.

  "Our records have a lot of holes. We have been asking humans to capture data, and they don't want to do it and don't have time to do it." (Senior product executive, enterprise software company)
  "We have knowledge articles in three different systems. We have some in ServiceNow, some in Teams folders, and some on personal laptops." (Senior automation architect, telecommunications company)

- **The design philosophy targets graceful degradation, not perfect prevention.** The aspiration of zero downtime is increasingly unrealistic in environments where dependency chains extend beyond any single organization's control. A more realistic philosophy accepts failure as inevitable and focuses on minimizing the blast radius, preserving critical functionality during disruption, and recovering predictably when failures occur. In practical terms, this means service decomposition that limits cascade potential, redundancy patterns that survive regional cloud failures, manual fallback procedures for scenarios in which automation itself fails, and recovery rehearsals that treat full-region failure as a tested scenario rather than an unthinkable edge case.

  "I would emphasize the proactive component in problem solving. That is so underestimated." (Roland Pfenninger, team leader ITSM, Zurich Airport)

- **Governed agentic AI augments detection and response.** Agentic systems can analyze high-volume telemetry, correlate signals across disparate tools, summarize incident state in real time, propose remediation steps based on historical patterns, and execute low-risk recovery actions without human intervention (see Figure 3). What separates effective adoption from the kind of failure that occurs when agents act outside their scope is the governance layer around agent autonomy. Effective governance includes explicit permissions that prohibit destructive actions without secondary approval, audit trails that capture every agent's decision and its reasoning, escalation thresholds that route higher-stakes decisions to humans, and continuous evaluation against outcome quality rather than action speed.

> "We are building an orchestrator agent, which will play the role of a major incident manager. Instead of hiring humans to perform a set of actions, we will create agentic workers to perform the same level of duties. The confidence our teams have in the system still requires the human-in-the-loop mechanism and guardrails." (Director of enterprise service management, pharmaceutical company)
> "Agents can react to signals, develop their own hypotheses, do their own investigations, and not only provide recommendations but also take action on behalf of users." (Senior product executive, enterprise software company)

- **Knowledge orchestration is treated as a core operational capability.** AI-augmented retrieval is now closing the knowledge gap by unifying knowledge sources across service desk records, observability data, change history, prior incidents, and architectural documentation. Mature platforms retrieve relevant context dynamically in the moment of triage, surface pattern matches with prior incidents that would otherwise go unrecognized, and capture new knowledge as it is generated during response. Retrieval-augmented knowledge access is a primary mechanism for reducing diagnosis time and improving resolution accuracy.

  > "The fast knowledge retrieval is probably the biggest area of new value for us." (Senior automation architect, telecommunications company)

- **Observability is promoted to an incident context engine.** Observability contextual richness is increasingly being surfaced directly in incident response workflows rather than living in a separate tooling silo. Mature practice integrates observability data into the incident lifecycle so that responders see system health context, recent change activity, and dependency status instead of needing to assemble that picture manually under time pressure. The convergence of observability and incident management is one of the most significant operational shifts of the AI era and represents a meaningful repositioning of where incident response sits in the broader operations stack.

  > "AI correlation is a bit smarter than your usual correlation. It tries to find and learn from the pattern so it can create one incident instead of five, reducing the time to find the real cause." (Vaidotas Januska, chief product officer, cloud managed services, Devoteam)

- **Response shifts from human led to AI led with explicit human oversight.** A mature practice is a deliberate division of labor in which AI agents handle high-volume, low-complexity, well-understood incidents autonomously, while human responders focus on novel, ambiguous, high-consequence incidents in which

contextual judgment is essential. The division requires organizations to classify their incident types, invest in the AI capability needed for autonomous handling of suitable categories, and preserve human expertise and engagement for categories where it matters. Done well, this approach reduces operational burden without compromising response quality. Done poorly, it produces the worst of both worlds: AI handling incidents it should not, and human responders losing the practiced experience they need for incidents in which their judgment is irreplaceable.

"We are not autonomously letting the AI system act on our behalf right now. We leverage the technology to provide all the inputs for decision-making, with the human in the loop." (Director of enterprise service management, pharmaceutical company)

- **Governance extends to automation, agents, and change.** Mature IT organizations extend governance structures to cover AI agents acting at machine speed, automation pipelines deploying changes at high velocity, and cross-system changes whose blast radius is not visible to any single domain owner. The work is not about adding bureaucracy. It is about establishing the policy layer that determines what an AI agent is permitted to do without human approval (see Figure 4). Guardrail rules and runtime detection must be specified explicitly rather than inferred from behavior.

  "It is having a coordinated effort between our governance teams and compliance teams and bringing them in, making them a part of this process, because every platform we are using and every cloud provider is introducing some form of AI." (Director of enterprise service management, pharmaceutical company)

- **Measurement evolves from operational efficiency to business outcome.** Mature organizations augment traditional metrics with measures of incident recurrence (indicating whether root causes are actually being addressed), business impact metrics that translate incident severity into terms the business values, recovery capability metrics that gauge how the organization performs against its own resilience targets, and AI performance metrics that evaluate agent decisions against outcome quality, not action speed. An evolved measurement framework closes the credibility gap with business stakeholders and lets incident management make the case for continued investment in a language the business can understand.

  "We have to report if it is worth it. Do we have an ROI with the outcome of the usage from the AI versus the license cost?" (Roland Pfenninger, team leader

ITSM, Zurich Airport)

"We are looking at how we invested in human capital prior to introduction of this technology versus how that is shifting. We are seeing a big positive shift." (Director of enterprise service management, pharmaceutical company)

Figure 3

Agentic AI Capabilities In Modern Incident Management

Category	What it does
Self-service deflection agents	Operate before an incident reaches a human responder, attempting to resolve the user's issue through conversational AI grounded in the organization's knowledge base. Deflect routine issues from the service desk queue, escalating only when genuine human intervention is required.
Ticket categorization and intake agents	Improve data quality in the moment by classifying incident type, populating required fields, identifying likely affected services, and suggesting appropriate priority. Produce cleaner incident records that downstream analytics, automation, and AI capabilities can use, eliminating the cascading errors that originate from poor intake data.
Detection and correlation agents	Ingest signals from observability, monitoring, and ITSM tools to identify potential incidents before they escalate. Reduce alert noise through AI-driven correlation, surface systemic patterns across disparate data sources, and create context-rich incidents that combine telemetry, change activity, and service desk data in a unified picture.
Triage and routing agents	Automatically assess incoming incidents, identify likely affected systems, and route the incident to the appropriate team. Reduce the manual investigation phase by assembling relevant context, including historical similar incidents, recent changes, and known dependencies.
Investigation and diagnostic agents	Operate as autonomous investigators when an incident is opened, generating hypotheses, querying telemetry across multiple data sources, testing diagnostic theories, and posting findings to collaboration channels. Coordinate signal collection across observability, code repositories, and ticketing systems to accelerate the diagnostic phase of response.
Knowledge retrieval and pattern-matching agents	Surface relevant operational knowledge dynamically during incident response, including runbooks, prior incident resolutions, architectural documentation, and tribal knowledge captured in past incident bridge transcripts. Reason across structured and unstructured operational context to provide actionable guidance even when formal documentation is incomplete.
Remediation and resolution agents	Execute defined remediation actions autonomously for well-understood incident categories, including restarting services, applying configuration corrections, scaling resources, and rolling back recent changes. Operate within explicit permission scopes and audit boundaries, escalating to human responders for higher-stakes decisions.
Coordination and collaboration agents	Automate the operational logistics of major incident response, including creating incident channels, paging relevant responders, generating real-time status summaries, distributing stakeholder updates, and managing escalation. Operate within collaboration platforms such as Slack and Microsoft Teams.
Knowledge generation and curation agents	Convert each incident into reusable organizational knowledge by drafting knowledge articles from incident resolutions, identifying gaps in existing knowledge content based on ticket patterns, validating and refreshing knowledge accuracy over time, and recommending new articles where ticket trends indicate missing content.

Figure 4

AI Agent Permission Audit Matrix

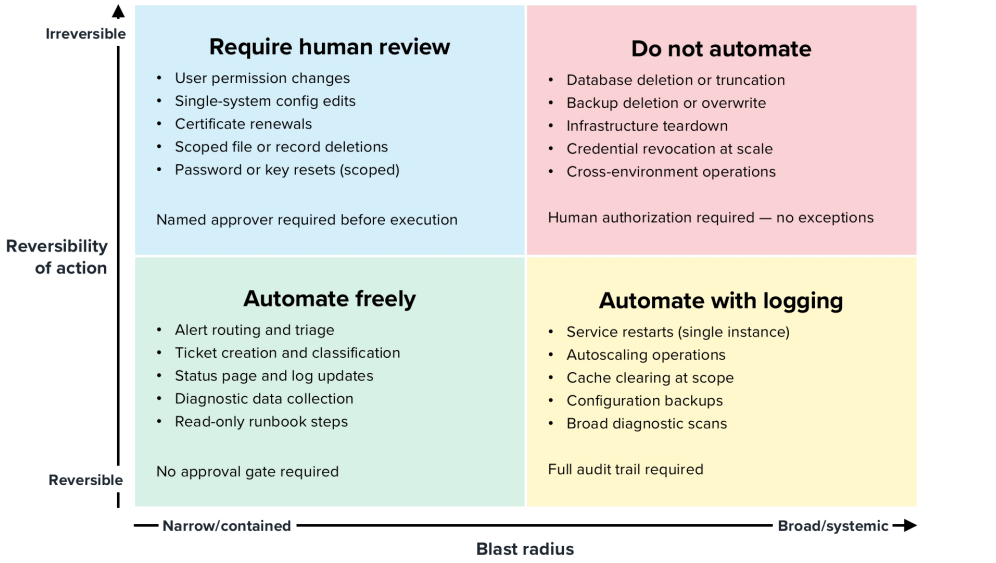

How To Modernize Without Breaking What Works

Organizations that attempt to overhaul their incident management discipline in a single transformation initiative typically fail, not because the new practices are wrong but because incident response is a live mission-critical capability that cannot be replaced wholesale without inherent risk. The recommendations that follow assume a different approach: deliberate sequenced evolution that preserves operational continuity while introducing the capabilities, governance, and measurement frameworks the AI era requires. Investing in agentic AI before establishing governance creates the conditions for the very failures the technology is meant to prevent. Follow these steps to iteratively improve your incident management practice:

- **Conduct a candid baseline assessment of the past 12 months.** Begin with a structured review of the past 12 months of major incidents, focusing not on what was reported but on what is missing from the record: incidents that recurred but were closed as new tickets, resolution times that obscure prolonged investigation phases, and root causes attributed to user error or third-party dependencies rather than examined for systemic patterns. Assess the current measurement framework

to determine whether recurring failures inflate mean-time-to-resolution averages and whether incident impact can be measured in business terms rather than operational ones. The gap between what current data shows will identify opportunities for early agentic use cases.

"With hundreds of thousands of tickets in the system, you analyze those to understand your opportunities, the low-hanging fruit." (Jody Wetterlind, IT senior director, Seagate)

- **Audit AI agent permissions before expanding agent autonomy.** For organizations already running AI agents in incident workflows or development pipelines, the immediate priority is governance, not expansion. Document every AI agent currently operating in the environment, the permissions it has been granted, the actions it can take without human approval, and the audit trail it produces. Identify potential destructive actions that should require secondary human approval regardless of agent confidence. Establish escalation thresholds that route higher-stakes decisions to human responders. Implement audit trails that capture not only what the agent did but also the reasoning it offered for each decision.
- **Integrate observability data into the incident response workflow.** One of the most operationally valuable moves available in the next 12 months is closing the gap between what observability tools see and what incident responders have during triage. Integration should ensure that when an incident is opened, the responding team automatically receives recent change activity affecting the system, current health and performance metrics, dependency status for upstream and downstream services, and historical pattern matches with prior incidents on the same systems. The integration also creates the foundation for more sophisticated AI augmentation later.

  "We wanted a very tight integration into the CMDB. It might sound trivial, but it is actually where a lot of the possibilities and value for us exist." (Senior automation architect, telecommunications company)
  "The AI agent suppresses and correlates alerts and raises one issue based on many alerts around the same asset. AI can raise one incident and assess the possible root cause based on the logs." (Vaidotas Januska, chief product officer, cloud managed services, Devoteam)
- **Modernize knowledge management as an active response capability.** Evaluate current knowledge sources for completeness, currency, and accessibility during incident response. Invest in AI-augmented retrieval that surfaces relevant context dynamically rather than requires responders to search. Capture new knowledge generated during incident response as a structured output. Be sure to measure

knowledge effectiveness by its operational impact on resolution accuracy rather than by volume.

"Avoid losing any knowledge along that path. If you take that and train a model on it, you are going to get a much higher level of accuracy and consistency." (Senior product executive, enterprise software company)
"AI can tell the technician this [incident] may have a solution but only if there is already something in the knowledge base. That is what makes building the knowledge base the real priority." (Director of IT business systems, retail organization)

- **Redesign change governance for modern deployment velocity.** Implement automated change risk scoring that uses historical incident data to identify high-risk patterns. Establish governance differentiation between low-risk routine changes and high-risk structural changes. Integrate change context directly into incident response so that responders know what changed before they investigate why something failed. Extend change governance to cover AI agents and automation pipelines that act at machine speed.
- **Build recovery capability with the same rigor applied to detection.** Treat recovery as a first-class operational capability. Rehearse full-region cloud failure scenarios with the same discipline applied to security tabletops. Establish manual fallback procedures for scenarios in which automation itself is the source of the incident. Document the physical and logistical recovery steps required for incidents that cannot be remediated remotely. Measure recovery capability against realistic worst-case scenarios rather than typical incident profiles.
- **Evolve measurement in dialogue with the business.** Replacing operational metrics with outcome metrics is an organizational change rather than a technical exercise, and it requires alignment between operations leadership and the business stakeholders whose definition of value the new metrics are meant to reflect (see Figure 5). Convene a structured measurement review with relevant business leaders to identify the outcomes incident management is meant to protect: service availability for revenue-generating platforms, employee productivity in critical business functions, and regulatory exposure in compliance-sensitive areas. Design measurement that translates incident performance into those terms.

  "One idea is to measure the lead time from an incident to see if it is different when the AI was part of the solution-finding process versus a ticket handled fully manually." (Roland Pfenninger, team leader ITSM, Zurich Airport)

- **Position the discipline for what comes next.** IT operations leaders should be preparing for the next wave of operational change: autonomous agent ecosystems in which multiple agents coordinate without continuous human oversight. Explainable in-context agent reasoning is becoming a core requirement for accountable automation. Organizations that approach this work as a one-time transformation will find themselves needing to do it again. Those that approach it as a continuous evolution will be positioned to lead.

  "AI has real value in learning about the ecosystem and applying that knowledge to make sure we recover faster. This will increasingly move to autonomous operations." (Director of enterprise service management, pharmaceutical company)
  "The change is coming. They need to accept the level of risk and start somewhere. It is inevitable." (Vaidotas Januska, chief product officer, cloud managed services, Devoteam)

Figure 5

Measuring What Matters: The Metric Evolution

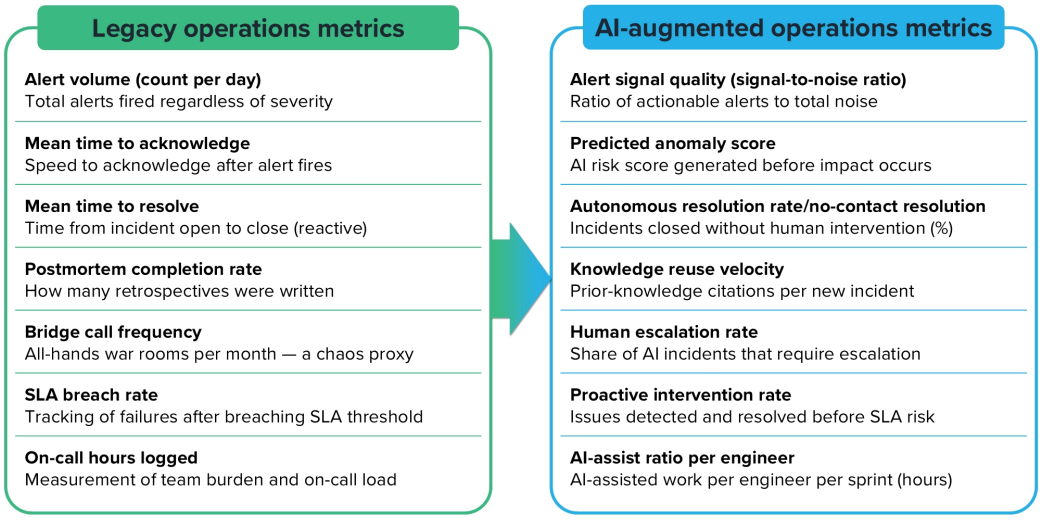

Supplemental Material

Companies We Interviewed For This Report

We would like to thank the individuals from the following companies who generously gave their time during the research for this report.

Atlassian

BigPanda

BMC Helix

Devoteam

EasyVista

Everbridge

Freshworks

HCLSoftware

Ivanti

LogicMonitor

ManageEngine

Matrix42

OpenText

Qinfinite

Rezolve.ai

Salesforce

Seagate

ServiceNow

Serviceware

SolarWinds

SymphonyAI

SysAid

TeamViewer

TOPdesk

USU

Zendesk

Zurich Airport

FORRESTER®

We help business and technology leaders use customer obsession to accelerate growth.

FORRESTER.COM

Obsessed With Customer Obsession

At Forrester, customer obsession is at the core of everything we do. We're on your side and by your side to help you become more customer obsessed.

Research

Accelerate your impact on the market with a proven path to growth.

- Customer and market dynamics
- Curated tools and frameworks
- Objective advice
- Hands-on guidance

Learn more.

Consulting

Implement modern strategies that align and empower teams.

- In-depth strategic projects
- Webinars, speeches, and workshops
- Custom content

Learn more.

Events

Develop fresh perspectives, draw inspiration from leaders, and network with peers.

- Thought leadership, frameworks, and models
- One-on-ones with peers and analysts
- In-person and virtual experiences

Learn more.

Contact Us

Contact Forrester at www.forrester.com/contactus. For information on hard-copy or electronic reprints, please contact your Account Team or reprints@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
Tel: +1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com