Market Guide for AI Site Reliability Engineering Tooling

26 January 2026 - ID G00836089 - 16 min read

By: Daniel Betts, Chris Saunderson, Hassan Ennaciri

Initiatives: Technologies and Markets

Organizations struggle to justify costs to adopt site reliability engineering practices to deliver on their reliability and resilience goals. Heads of I&O must strategically evaluate and invest in AI SRE tooling to lower cost of SRE adoption, meet operational demands and deliver effective reliability, efficient products, platforms and services.

Overview

Key Findings

- Many organizations consider site reliability engineering (SRE) approaches but struggle to justify the investment into hiring SREs, training internal teams and purchasing tools.
- Traditional SRE teams and operations teams cannot keep up with the technology and operational demands required of them to deliver effective reliability and efficiency outcomes.
- Organizations that select AI SRE tooling focused on operations only will become better at reactively fixing incidents but not improving system reliability.

Recommendations

- Kickstart SRE introduction by evaluating AI SRE tooling as an initial investment into SRE, limiting cost exposure and determining real-life improvements or gaps.
- Augment existing SRE and operations teams by investing in AI SRE tooling to enable them to focus on proactive reliability improvement activities.
- Build an AI SRE strategy beyond just operations by utilizing insights surfaced through tooling capabilities (such as telemetry, event correlations and root cause analysis) to proactively prevent incidents through reliability-focused design and delivery workflows.

Strategic Planning Assumptions

- By 2029, 85% of enterprises will use AI SRE tooling to optimize operations in order to meet organizational and customer reliability demands. This will increase from less than 5% in 2025.
- By 2029, 75% of organizations will integrate AI-distilled SRE lessons into product design and delivery, up from 10% in 2025.
- By 2030, 60% of new infrastructure designs will be validated by AI agents that use historical data to identify potential points of failure before development begins.
- By 2029, 90% of organizations will experience an AI-caused outage such as erroneous recommendations, yet continue AI SRE for its speed and scalability gains.

Market Definition

This document was revised on 29 January 2026. The document you are viewing is the corrected version. For more information, see the Corrections *page on gartner.com.*

The AI site reliability engineering (SRE) tooling market enables the adoption of SRE practices by utilizing AI approaches and techniques to focus on improving reliability, resilience, and the customer experience of products and platforms. These tools help organizations scale the value from SRE approaches by setting and managing reliability goals, and surfacing monitoring and observability insights and performance demands. The tools are delivered as stand-alone task-specific tools or as part of platforms with broader capabilities. AI SRE tools are essential for ensuring the reliability, performance, and overall health of software systems, providing valuable insights and automation capabilities that help scale teams managing complex systems effectively.

AI SRE tools focus on delivering capabilities in the following areas:

- Analysis of contextual data and dynamically adjusting thresholds, enabling intelligent alerting
- Rapid information retrieval during critical incidents, and insights into reliability metrics and routine operational tasks
- Delivery of rapid analysis of telemetry change and incident data, identifying root causes to accelerate remediation
- Forecasting of load, capacity requirements, and service-level objective (SLO) breaches to enable predictive analytics
- Automatic detection of issues and remediation without human intervention
- Synthesis of performance data and optimal thresholds for service-level indicators (SLIs)/SLOs
- Providing comprehensive service-health monitoring and insights to stakeholders
- Integrating into developer and engineering practices to embed reliability

Mandatory Features

Mandatory features include:

- Autonomous, semiautonomous AI SRE decision making and actions with humans in the loop to enable SRE at scale and be proactive
- Support for hybrid infrastructure operational environments across on-premises, private and public cloud, edge, and colocation

Common Features

Common features include:

- Enhanced detection and correlation techniques to aid with incident detection and remediation
- Insights on incident cause and remediation provided to end users
- SLO and SLI definition, measurement, management, and insight generation

- Management and support for error budget policies
- End-to-end dependency mapping
- Predictive analytics for reliability breaches
- Generated post-mortem analysis
- Agentic AI and AI agents
- Assessment of testing reliability objectives
- Self-service reporting for stakeholders
- Integration APIs to embed into workflows
- Deployment of tooling as on-premises, cloud or SaaS
- Self-service portal or API access for end-user interaction
- Accelerated root cause analysis
- Run book creation and generation
- Capacity forecasting and autoscaling

Market Description

AI SRE fundamentally transforms how organizations manage reliability and resilience within their distributed complex products and platforms in hybrid, on-premises and cloud-native environments by introducing predictive capabilities and intelligent automation. Tools in this market also offer a path to evaluation, introduction and maturation of SRE practices.

AI SRE tooling is currently focused on assisted, semiautonomous and autonomous AI to revolutionize critical SRE functions by automating and augmenting SRE processes; these technologies focus on accelerating remediation; improving incident detection, anomaly detection and response; powering root cause analysis, postmortem activities, predictive maintenance; and enhancing overall system resilience. Feature sets need to move beyond this to focus on proactive, preventative reliability impacts on customer experience, based on reliability metrics such as SLIs and SLOs (see Figure 1).

Figure 1: AI SRE Tooling Capability Overview

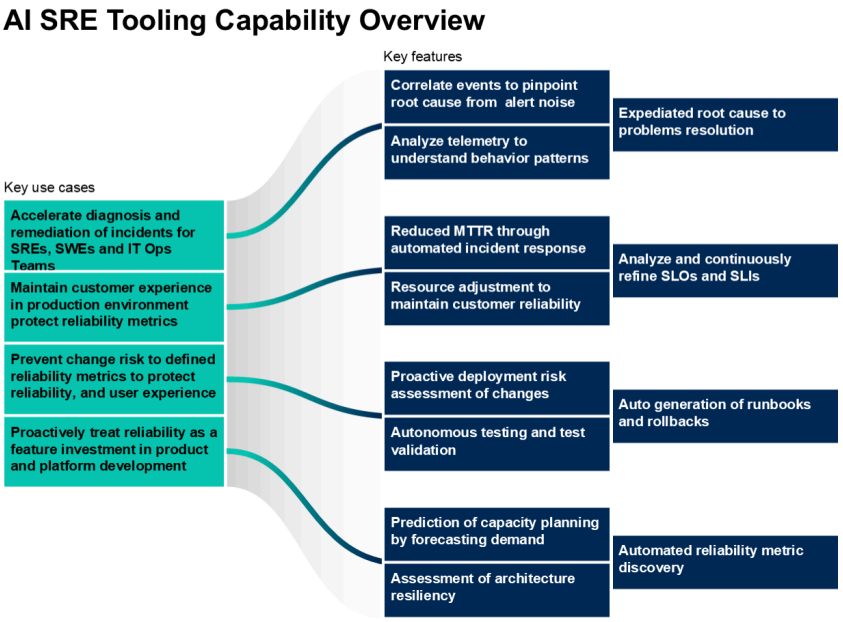

Source: Gartner
836089

Gartner

Market Direction

Adoption

AI SRE tools make introducing and scaling SRE initiatives viable for more organizations as they reduce the investment requirements to hire and train staff to implement SRE practices. Modern distributed systems have increased reliability and performance requirements but are also increasingly complex, producing vast amounts of telemetry data and alerts, which can overwhelm teams, resulting in missed signals, alert fatigue, and slower incident resolution. The addition of AI-driven innovations and capabilities will magnify these problems. AI SRE solutions are emerging to tackle these challenges by automating incident response and directly addressing alert fatigue and delayed resolutions.

Embedding AI at the heart of operations provides the potential to transform the function from reactive to proactive, driven by structured approaches rather than firefighting. This also helps to build an internal narrative about the adoption of AI in the I&O domain, demonstrating adaptation and mindful introduction of valuable AI capabilities.

AI SRE tools have initially concentrated on operational tasks for three main reasons:

- **Clearly defined challenges:** Problems such as alert fatigue and manual root cause analysis are specific, data-heavy issues that AI can effectively address. The scope and control of what AI is applied to is also well-defined.
- **Return on investment and reduced costs:** These solutions promise benefits like reduced downtime, fewer engineering hours spent on incidents, and prevention of major outages. However, experimentation is required to carefully measure ROI to justify adoption and proliferation.
- **Availability of data:** Operations continuously generate logs, metrics, and traces, providing data to train machine learning models, care needs to be taken with data quality.

To enable the successful adoption of AI SRE tooling, the tools must be accompanied by an organizational culture change. Reliability must be treated as a critical feature and not seen as someone else's problem. This shift should focus on reliability improvements across the delivery life cycle, AI enablement, embracing risk, continuous improvement, and collaboration between I&O and product teams to ensure that signals from the tools are used to drive change. Investing in tools alone will not deliver success.

Product Landscape

AI SRE tools are currently available as:

- Stand-alone solutions, often startups that integrate into existing toolchain and operational investments
- A capability set in a broader platform, such as observability, DevOps or IT service management (ITSM) tools

This market remains formative, and we expect to see additional products launched, as well as an increasing number of new startups, acquisitions of established tools and retirements.

Capabilities

The focus of vendors who are building and releasing AI SRE tools is on ensuring reliability goals can be defined, met, measured, communicated and improved upon based on business goals and customer expectations.

AI SRE tools have evolved to offer these key features (see Table 1).

Table 1: SRE Tooling Features and Descriptions

(Enlarged table in Appendix)

Feature	Description
Correlated events in production environments	Gather and examine substantial amounts of raw event data such as logs, metrics, and status updates from various IT systems. Filter out irrelevant information and uncover significant patterns and connections, enabling teams to identify security risks or operational issues before they develop into serious problems.
Proactive telemetry analysis	Proactively examining raw data such as metrics, logs, and traces from IT systems to gain a comprehensive understanding of their behavior. This practice is fundamental to observability, enabling teams to spot performance trends, troubleshoot complex and unforeseen issues, and enhance system reliability before problems generate alerts or affect users.
Accelerated root cause analysis	Speed up problem resolution by leveraging techniques such as AI, machine learning, and focused data correlation to quickly pinpoint the underlying causes of incidents. Eliminating the need for lengthy manual investigations, this method enables faster remediation and minimizes impact on the business and customer. Improves the detection of a change that was implemented that caused an incident.
Generated run books	Develop and maintain run books by automatically generating content from multiple sources, standardizing procedures, and maintaining accuracy through continuous updates. This approach lessens the manual workload for SRE and operations teams, reduces the risk of outdated information, and boosts operational efficiency by ensuring that accurate, up-to-date procedures are always available for handling various incidents.
Automated response to incidents	Leverage intelligent data correlation, incident prediction, and dynamic generation of context-aware remediation steps. Significantly reduce MTTR, minimize manual effort by automating run book execution and alert triage, and enable proactive prevention. Streamlines SME onboarding and produces audience-specific incident communications, optimizing incident response and enhancing system resilience and efficiency. Provide recommendations based on findings to SREs, SWEs and operations teams.
Optimized workload resource and configuration	Lower cloud costs by automatically detecting and applying efficient resource configurations for workloads. Continuously adjusting resources to match real-time needs and performance goals. Ensure optimal utilization, prevent overprovisioning, and deliver cost savings while enhancing resource management for SRE and operations teams.
Managed reliability metrics	Utilize to define, monitor, and proactively manage service reliability. Automate analysis and prediction to lessen manual effort and enhance the precision of maintaining SLIs and SLOs. More dependable services, decreased downtime, and improved alignment with user expectations through data-driven insights and proactive measures.
Assessed deployment risk	Evaluate deployment data and predictive analytics to detect possible failures, SLO risks, or increased operational effort associated with new releases. Empowers SREs to implement proactive mitigation measures, adopt safer deployment strategies, and make informed go/no-go decisions.
Validation and testing	Test case generation, execution, and system validation carried out with minimal human involvement. Intelligent creation of unit and integration tests, dynamic maintenance of test suites, and autonomous validation of deployments to ensure they meet expected behavior.
Autonomous debugging	Automated error detection, root cause analysis, and recommended application of fixes, greatly reducing the need for manual debugging. Improve the efficiency and accuracy of debugging, allowing developers to concentrate on more complex challenges while AI manages routine error resolution.
Generated rollback plans	Accurate rollback procedures focused on problematic components, including detailed data migration steps. Reduced MTTR minimizes risk of human error during high-pressure situations, consistent, safer recovery processes supporting SLO compliance.
Predictive capacity planning	Enhanced predictive capacity planning through AI/ML data analysis to accurately forecast demand, automate resource adjustments, and deliver clear insights. Enables proactive resource optimization and greater cost-efficiency to ensure resource availability.
Architecture resilience assesment	Assesses a system's resilience to failures and its capacity to maintain SLOs. Proactively uncover weaknesses in design, fault tolerance, and recovery mechanisms, allowing for targeted enhancements that improve overall reliability and reduce the impact of service disruptions.
Agent integration	Facilitate seamless integration of large language model (LLM)-based AI applications with external data sources, tools, and services through model context protocol (MCP) and Agent2Agent (A2A) protocol.
Audit trail of actions and reasoning	Provide auditable output of the AI agent's activities, decisions, and reasoning.

Source: Gartner (January 2026)

AI SRE tooling capabilities continue to be developed, and many vendors are expanding feature sets that their tools can deliver. Key roadmap capabilities planned for delivery over the next one to three years are highlighted in Table 2 below.

Table 2: AI SRE Tooling Capabilities Roadmap

(Enlarged table in Appendix)

Roadmap item	Capability description
Proactive incident avoidance	Foresee potential incidents by analyzing trends and patterns. This predictive capability is critical as companies increasingly demand resilience over mere reactive measures.
Protection of SLOs and composite SLOs	Ensure that agreed customer SLOs and composite SLOs are protected to maintain customer experience.
Increased automation with human-AI partnerships	Paradigm shift toward platforms that support a collaborative model. Instead of completely replacing human intervention, these systems work as digital teammates automating routine tasks while leaving complex decision making to experienced engineers.
Integration across multicloud and hybrid environments	Manage complex, multicloud environments. As organizations continue to scale their digital operations, the integration capabilities offered by these platforms become essential for maintaining stability and cost-efficiency.
Enhanced contextual awareness through knowledge graphs	Use of knowledge graphs to provide deeper insights into system behavior. Reduces the "click operations" that have long plagued traditional incident response workflows.
Focus on reducing operational toil	Predictive MTTR reductions and improved automation; enterprises can allocate engineering resources to strategic projects, ultimately fostering innovation alongside stability.
Controlled failure injection	Automated chaos experiment generation streamlining the design of complex experiments, minimizing manual effort. Learning from system data to quickly uncover deeper and previously unknown risks, support proactive resilience strategies and improved SLO compliance building system confidence and reduced downtime.
Multiagent architectures and cross-functional collaboration	AI SRE platforms are moving toward multiagents. Coordination of several specialized AI routines that manage detection, diagnosis, remediation, and governance, ensuring a more holistic approach to operational stability.
Agent reliability testing	Testing AI SRE agents and their reliability within agentic architecture at completing tasks.

Source: Gartner (January 2026)

Risks

AI agent behavior can be unpredictable, and even well-designed systems may introduce new points of failure if not properly managed. Strengthened governance with robust oversight, proactive monitoring, and clearly defined mechanisms for human intervention minimize reliability issues.

Market Analysis

The promise of AI SRE is appealing to organizations, as they are driven by the need to deliver highly reliable yet scalable customer experiences, especially in their most critical features. However, the adoption of AI SRE is still in its infancy for many organizations, as they grapple with the foundational concepts of establishing and measuring reliability focused on the customer experience, and developing trust. The adoption of AI SRE tooling represents an investment in protecting user experience and enabling scalability for organizations seeking to increase the reliability they deliver through their products, platforms and services.

SRE tooling vendors have developed feature sets and solutions to provide:

- Visibility
- Rapid insight
- Assessment
- Reliability enforcement
- Management
- Optimization
- Remediation
- Governance

Driving efficient and effective reporting, protection and enablement of reliability focused on customer experience remains challenging, particularly across all phases of the delivery process and in the running production environment. I&O leaders must either invest in dedicated AI SRE tooling solutions or leverage AI SRE tooling capabilities from platform vendors offering reliability feature sets. They must integrate these tools with existing tools and processes to gain full reliability insight for all development and delivery activities.

An I&O organization establishing the value of foundational AI SRE practices can gain value by adopting features from existing platforms they have access to. However, they may not deliver the specific target solution to a problem that a dedicated AI SRE tooling solution may provide. Experimentation with tools is critical to find the most cost-effective solution to meet your organization's reliability goals.

Market Dynamics

The current formative state of the AI SRE tooling market includes a mix of vendors. Stand-alone solutions offer integrations with the most commonly used platforms and focus on the most distinguishing features that reliability engineers and their stakeholders require. Conversely, platforms with either features or modules that implement AI SRE capabilities also offer a subset of these features, albeit more tightly integrated into the capabilities of the broader platform. Many of the representative vendors in this market are startups, and investment should be done with caution. Most vendor solutions are currently targeting the reactive operations environment with their capabilities rather than proactive prevention in delivery and design phases of the SDLC.

Pricing models are still in flux and developing. They are often token-based, task-based (such as investigations like RCA or questions asked), data-volume-based, or compute-utilization-based. Many providers offer tiered, bundled and package-type pricing solutions.

There has been limited acquisition activity in the AI SRE tooling market, as it remains early in its cycle.

Market Sizing

Given the early stage of this market and the relatively short time that tooling has been available, it is challenging to identify revenue driven specifically by AI SRE investments. However, we will continue to review and evaluate the market size as announcements and changes occur in this market.

Representative Vendors

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

Vendor Selection

The AI SRE tooling market is in an early state, and the vendors selected met at least two capabilities as outlined in this guide (see Table 3).

Table 3: Representative Vendors in Site Reliability Engineering Tooling

(Enlarged table in Appendix)

Vendor	Product name
Amazon Web Services (AWS)	AWS DevOps Agent
Anyshift	Annie
AutonomousOps AI	HealR
Bacca AI	Bacca
BigPanda	BigPanda Platform
Cast AI	Cast AI
Causely	Causely
Ciroos AI	Ciroos
Cleric.ai	Cleric
Coralogix	Olly
Dash0	Agent0
Datadog	Bits AI
DeductiveAI	DeductiveAI
DrDroid	DrDroid
Dynatrace	Dynatrace
Fabrix.ai	Agentic AI Platform
Firefly.ai	Agentic Cloud Automation Platform
GrafanaLabs	Grafana Assistant
Harness	Harness AI SRE
Honeycomb.io	Observability Platform
Ilert	Ilert AI SRE
Incident.io	AI SRE
Komodor	Klaudia AI
Last9	Last9
Lightrun	Lightrun
Microsoft	Azure SRE Agent

Market Recommendations

Heads of I&O who want to implement AI SRE into reliability initiatives must:

- Augment existing SRE and operational resources by investing in AI SRE tooling to improve operational reliability and incident response.
- Determine their goals and ground their AI SRE intuitive around them.
- Assess and evaluate their existing observability, automation, DevOps or ITSM tools, as they may already possess some of the AI SRE capabilities described in this market guide.

- Prioritize tool investments that can elevate impacts on customer-focused reliability as part of product and platform workflows to prevent poor customer experience in production environments.
- Balance speed, innovation and risk by implementing AI SRE tools into IT operations with humans in the loop to trace change, incident and release impacts on reliability metrics.
- Evaluate AI SRE tooling investments and suitability by regularly reviewing vendor roadmaps for required features and regularly updating their capabilities based on customer reliability needs.
- Validate AI SRE-tool-based insights with direct feedback about customer sentiments from product owners.
- Enhance product and platform design and delivery by applying insights gained from AI-driven SRE to incorporate reliability improvements throughout all stages of the product life cycle.

Evidence

This analysis of SRE inquiries is based on a review of Gartner client inquiries discussing reliability from end-user organizations for 2024 and 2025. The total number of inquiries in which SRE and AI SRE was mentioned in 2025 has grown by 85% compared to 2024. Gartner has fielded more than 600 client interactions (inquiries and one-on-one conference meetings), discussing questions such as:

- “How do I prevent failure?”
- “How do I succeed with an SRE initiative?”
- “What tooling enables SRE at scale?”
- “How do I measure reliability in SRE?”
- “Where should I start with AI SRE?”

Document Revision History

Market Guide for Site Reliability Engineering Tooling - 17 December 2024

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

AI, Autonomy, and Architects: The Future of Site Reliability Engineering

Innovation Insight: AI-Augmented SRE

Justifying Your SRE Investments

7 Steps to Start and Evolve Your SRE Practice

Top AI and GenAI Use Cases to Augment SRE

Improve Customer Experience and Reliability With SLIs and SLOs

---

Table 1: SRE Tooling Features and Descriptions

Feature	Description
Correlated events in production environments	Gather and examine substantial amounts of raw event data such as logs, metrics, and status updates from various IT systems. Filter out irrelevant information and uncover significant patterns and connections, enabling teams to identify security risks or operational issues before they develop into serious problems.
Proactive telemetry analysis	Proactively examining raw data such as metrics, logs, and traces from IT systems to gain a comprehensive understanding of their behavior. This practice is fundamental to observability, enabling teams to spot performance trends, troubleshoot complex and unforeseen issues, and enhance system reliability before problems generate alerts or affect users.
Accelerated root cause analysis	Speed up problem resolution by leveraging techniques such as AI, machine learning, and focused data correlation to quickly pinpoint the underlying causes of incidents. Eliminating the need for lengthy manual investigations, this method enables faster remediation and minimizes impact on the business and customer. Improves the detection of a change that was implemented that caused an incident.
Generated run books	Develop and maintain run books by automatically generating content from multiple sources, standardizing procedures, and maintaining accuracy through continuous updates. This approach lessens the manual workload for SRE and operations teams, reduces the risk of outdated information, and boosts operational efficiency by ensuring that accurate, up-to-date procedures are always available for handling various incidents.

Automated response to incidents	Leverage intelligent data correlation, incident prediction, and dynamic generation of context-aware remediation steps. Significantly reduce MTTR, minimize manual effort by automating run book execution and alert triage, and enable proactive prevention. Streamlines SME onboarding and produces audience-specific incident communications, optimizing incident response and enhancing system resilience and efficiency. Provide recommendations based on findings to SREs, SWEs and operations teams.
Optimized workload resource and configuration	Lower cloud costs by automatically detecting and applying efficient resource configurations for workloads. Continuously adjusting resources to match real-time needs and performance goals. Ensure optimal utilization, prevent overprovisioning, and deliver cost savings while enhancing resource management for SRE and operations teams.
Managed reliability metrics	Utilize to define, monitor, and proactively manage service reliability. Automate analysis and prediction to lessen manual effort and enhance the precision of maintaining SLIs and SLOs. More dependable services, decreased downtime, and improved alignment with user expectations through data-driven insights and proactive measures.
Assessed deployment risk	Evaluate deployment data and predictive analytics to detect possible failures, SLO risks, or increased operational effort associated with new releases. Empowers SREs to implement proactive mitigation measures, adopt safer deployment strategies, and make informed go/no-go decisions.
Validation and testing	Test case generation, execution, and system validation carried out with minimal human involvement. Intelligent creation of unit and integration tests, dynamic maintenance of test suites, and autonomous validation of deployments to ensure they meet expected behavior.

Autonomous debugging	Automated error detection, root cause analysis, and recommended application of fixes, greatly reducing the need for manual debugging. Improve the efficiency and accuracy of debugging, allowing developers to concentrate on more complex challenges while AI manages routine error resolution.
Generated rollback plans	Accurate rollback procedures focused on problematic components, including detailed data migration steps. Reduced MTTR minimizes risk of human error during high-pressure situations, consistent, safer recovery processes supporting SLO compliance.
Predictive capacity planning	Enhanced predictive capacity planning through AI/ML data analysis to accurately forecast demand, automate resource adjustments, and deliver clear insights. Enables proactive resource optimization and greater cost-efficiency to ensure resource availability.
Architecture resilience assesment	Assesses a system's resilience to failures and its capacity to maintain SLOs. Proactively uncover weaknesses in design, fault tolerance, and recovery mechanisms, allowing for targeted enhancements that improve overall reliability and reduce the impact of service disruptions.
Agent integration	Facilitate seamless integration of large language model (LLM)-based AI applications with external data sources, tools, and services through model context protocol (MCP) and Agent2Agent (A2A) protocol.
Audit trail of actions and reasoning	Provide auditable output of the AI agent's activities, decisions, and reasoning.

Source: Gartner (January 2026)

Table 2: AI SRE Tooling Capabilities Roadmap

Roadmap item	Capability description
Proactive incident avoidance	Foresee potential incidents by analyzing trends and patterns. This predictive capability is critical as companies increasingly demand resilience over mere reactive measures.
Protection of SLOs and composite SLOs	Ensure that agreed customer SLOs and composite SLOs are protected to maintain customer experience.
Increased automation with human-AI partnerships	Paradigm shift toward platforms that support a collaborative model. Instead of completely replacing human intervention, these systems work as digital teammates automating routine tasks while leaving complex decision making to experienced engineers.
Integration across multicloud and hybrid environments	Manage complex, multicloud environments. As organizations continue to scale their digital operations, the integration capabilities offered by these platforms become essential for maintaining stability and cost-efficiency.
Enhanced contextual awareness through knowledge graphs	Use of knowledge graphs to provide deeper insights into system behavior. Reduces the “click operations” that have long plagued traditional incident response workflows.
Focus on reducing operational toil	Predictive MTTR reductions and improved automation; enterprises can allocate engineering resources to strategic projects, ultimately fostering innovation alongside stability.
Controlled failure injection	Automated chaos experiment generation streamlining the design of complex experiments, minimizing manual effort. Learning from system data to quickly

	uncover deeper and previously unknown risks, support proactive resilience strategies and improved SLO compliance building system confidence and reduced downtime.
Multiagent architectures and cross-functional collaboration	AI SRE platforms are moving toward multiagents. Coordination of several specialized AI routines that manage detection, diagnosis, remediation, and governance, ensuring a more holistic approach to operational stability.
Agent reliability testing	Testing AI SRE agents and their reliability within agentic architecture at completing tasks.

Source: Gartner (January 2026)

Table 3: Representative Vendors in Site Reliability Engineering Tooling

Vendor	Product name
Amazon Web Services (AWS)	AWS DevOps Agent
Anyshift	Annie
AutonomousOps AI	HealR
Bacca AI	Bacca
BigPanda	BigPanda Platform
Cast AI	Cast AI
Causely	Causely
Ciroos AI	Ciroos
Cleric.ai	Cleric
Coralogix	Olly
Dash0	Agent0
Datadog	Bits AI
DeductiveAI	DeductiveAI
DrDroid	DrDroid
Dynatrace	Dynatrace

Fabrix.ai	Agentic AI Platform
Firefly.ai	Agentic Cloud Automation Platform
GrafanaLabs	Grafana Assistant
Harness	Harness AI SRE
Honeycomb.io	Observability Platform
Ilert	Ilert AI SRE
Incident.io	AI SRE
Komodor	Klaudia AI
Last9	Last9
Lightrun	Lightrun
Microsoft	Azure SRE Agent
NeuBird	Hawkeye
New Relic	New Relic AI
NoFire.ai	NoFire.ai
Nudgebee	Troubleshooting Assistant
Observe	Observe
PagerDuty	PagerDuty AI
Phoebe	Phoebe

Plural	Plural
Resolve.ai	Resolve AI
Rootly ai	Rootly AI SRE
RunWhen	RunWhen
Sedai	Sedai Autonomous Cloud Platform
SRE.AI	SRE.AI
TierZero AI	TierZero
Traversal	Traversal

Source: Gartner (January 2026)