AIOps WITHOUT GOVERNANCE IS JUST AUTOMATION YOU CAN'T EXPLAIN

The same explainability gap the industry is fighting in generative AI already runs your infrastructure, and no one is watching it.

exactmarket℠

ELEPHANT IN THE ROOM EBOOK SERIES

TABLE OF CONTENTS

WELCOME TO THE ELEPHANT IN THE ROOM!

You have in your hands a guide designed to call out the elephant in the room. A topic that's too important to ignore but maybe isn't getting the attention it deserves.

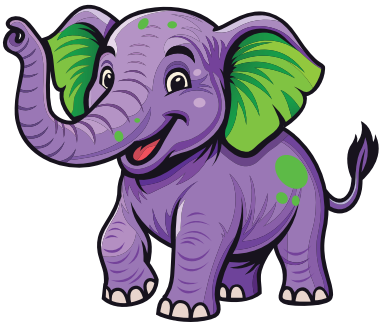

INTRODUCING THE ELEPHANT

Operational AI now makes remediation and suppression decisions your organization cannot explain, audit, or defend.

The governance frameworks you've built for generative AI stop at the edge of the infrastructure layer, and AIOps platforms have been running autonomously inside that gap for years. Generative AI is something that you or your customers could interact with, but AIOps stayed invisible because it sits quietly automating decisions.

AIOps was a monitoring feature when most AI governance frameworks got written. Today it makes thousands of decisions per hour, from alert suppression to issue remediation. The gap between AIOps operational authority and its governance oversight has become a regulatory concern and a legitimate internal risk for serious damage.

AI governance councils have sprung up in companies in the last few years, attempting to properly quantify, classify and regulate this new wave of technology. AI policies and questionnaires have been developed primarily with data security and generative AI as focus, leaving AIOps tools' autonomy as an unrealized risk.

This guide highlights four accountability gaps hiding in the operations stack, reframes AIOps as the production AI system it already is, and lays out a four-part framework, CAMP, for bringing autonomous IT decisions under the same governance discipline the enterprise already applies to generative AI.

INTRODUCTION

The AI You're Governing Isn't the Only AI Making Decisions

Every board now asks about generative AI governance, and what follows is practically a template where policies get written to appease audits and vendor questionnaires.

Meanwhile, AIOps platforms have been making autonomous decisions in production for years. They correlate alerts, suppress noise, trigger runbooks, and scale infrastructure on their own, and they sit in almost none of those governance reviews.[1]

Most operations leaders already sense this gap, but few have named it out loud.

The Visibility Asymmetry

If a chatbot hallucinates to a customer, this makes headline news, but if an AIOps platform suppresses a real alert at 2 a.m., it's overlooked as another log entry.

The reason is infrastructure-layer invisibility. Generative AI lives where users see it, and AIOps lives in the operations stack, where the work happens before anyone notices it happened.

SAME AI FAILURE MODE. DIFFERENT AMOUNT OF ATTENTION.

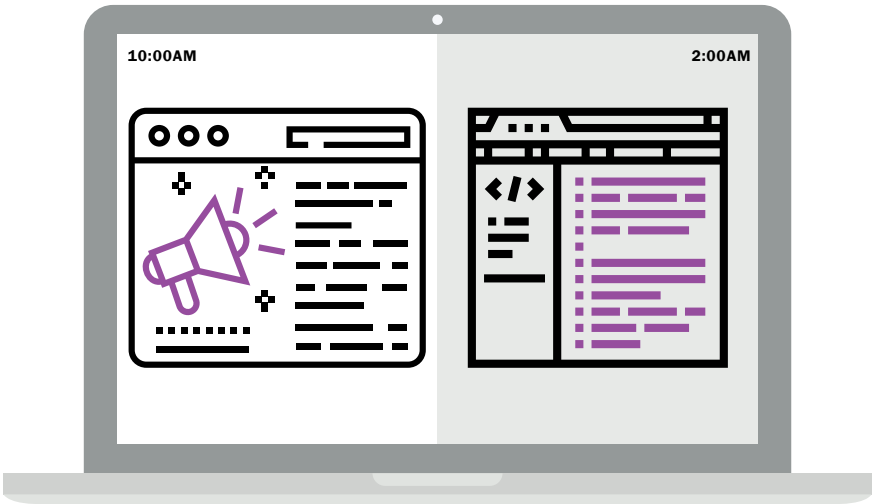

How Much Has Already Been Delegated

The scale of autonomous activity inside an enterprise's operations stack is staggering. AIOps can cut alert volume by up to 99% through correlation and suppression, which sounds like a benefit until you reframe it: 99% of what the platform sees never reaches a human, and the platform is the one deciding which 1% does.[2] With enterprise environments generating millions of events per day, the number of autonomous decisions happening per hour is a figure that humans can't handle. But their AI also can't explain everything.

What's Actually Being Decided

Modern AIOps platforms now span from detection to automated remediation, delivering faster response times and freeing up human teams for other work. The tradeoff is hard to see because an automated decision you can't explain is a decision you can't defend.

An automated decision you cannot explain is a decision you cannot defend.

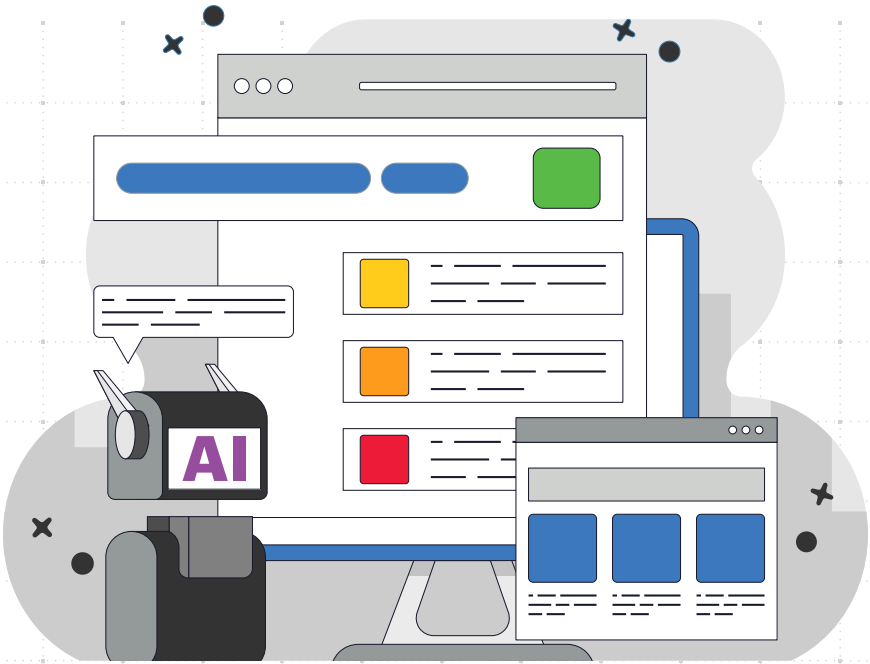

CHALLENGE

Four Accountability Gaps Hiding in Your Operations Stack

These four gaps share a pattern. Each one is a place that an AI decision gets made without a documented reason, exposing both auditing and material risk. The first three are structural problems around how AIOps got deployed and governed. The fourth is a security problem that attackers are already exploiting.

01 DECISIONS WITHOUT DECISION-MAKERS

Governance reviews miss the autonomous-action scope

02 THE AI THAT ARRIVED THROUGH THE UPDATE CHANNEL

AIOps features ship via product updates, bypassing AI councils

03 THE EXPLAINABILITY DEFICIT

Vendors treat model logic as proprietary IP

04 SUPPRESSION IS A SECURITY EVENT

Alert-fatigue rules are where attackers hide

Decisions Without Decision-Makers

When AI governance reviews an AI tool, the questions have good security intention. Where and how is the data stored? Will the data be used for training new models? Can it hallucinate into a customer conversation? The triad of privacy, cybersecurity, and legal has built a playbook for these questions. But the boilerplate questionnaires are missing questions specific to AIOps. Nobody is asking if the actions taken are reversible, or what fail-safes protect critical systems and data.

A financial services firm deploys a new monitoring tool that includes AIOps capabilities. The procurement team runs the vendor through the AI governance questionnaire. Data residency checks out. The vendor confirms training data is scoped appropriately. The legal team approves the terms. Nobody asks what autonomous actions the platform is authorized to take once deployed. Six months later, when an internal audit tests readiness against the U.S. Department of the Treasury Financial Services AI Risk Management Framework, the controls on decision-path auditability and explainability thresholds cannot be satisfied for the AIOps platform because the review that preceded deployment never required them.[3]

Regulated industries now face this directly. In sectors where the EU AI Act applies, “the model learned it” is not defensible in an audit.[4] An AIOps platform that suppresses an alert in a high-risk environment is making a decision the operator is required to be able to explain.

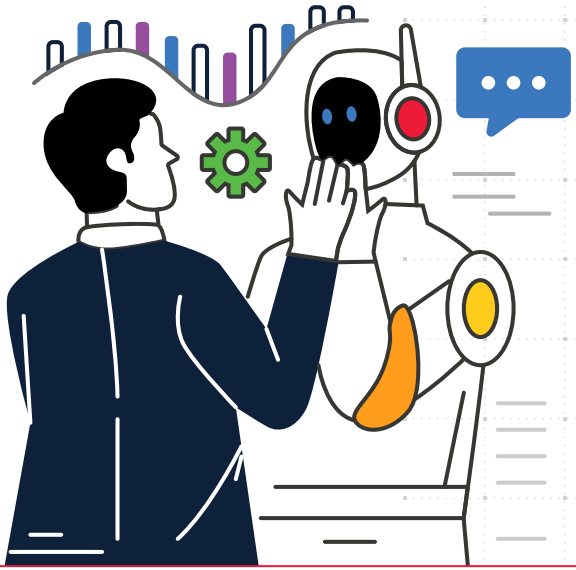

“The model learned it” is not defensible in an audit.

The AI That Arrived Through the Update Channel

As companies adopt generative AI offerings such as ChatGPT, Claude, or Gemini, the tools are passing through newly established AI councils or specific AI compliance and governance policies. Meanwhile AIOps features have been added quietly to tools already running in production, most without any of those gates. Gartner projects that 40% of enterprise applications will include task-specific AI agents by the end of 2026, up from under 5% in 2025.[5] ServiceNow acquired Moveworks for $2.85 billion in March 2025 specifically to embed agentic AI into a platform already installed across thousands of enterprise operations environments.[6]

40%

of enterprise apps will include task-specific AI agents by end of 2026, up from under 5% in 2025[5]

$2.85B

ServiceNow's March 2025 acquisition of Moveworks to embed agentic AI into platforms already running in production[6]

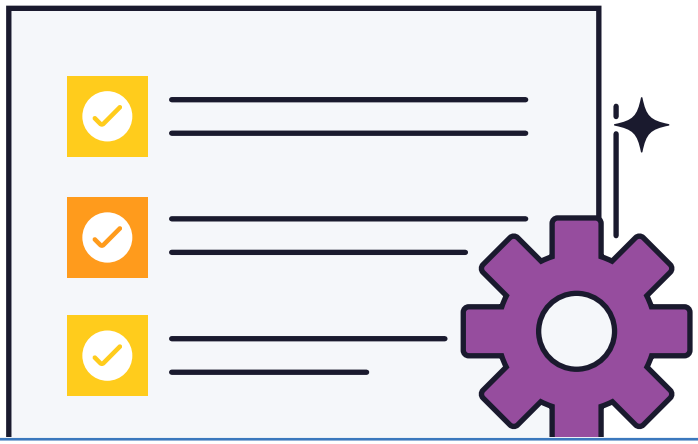

The Explainability Deficit

For years, most AIOps platforms shipped without meaningful explainability, and at best a customer might get some overall metadata that was useless for governance review. While explainable AI (XAI) has been around since as early as 2015, the recent mass adoption of gen AI has increased demand for XAI, and now regulatory pressure is pulling AIOps in the same direction.

SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) have been standard XAI techniques for years, and both have been applied to anomaly detection and incident response systems in production. Walmart's real-time AIOps pipeline uses SHAP-based explainers as surrogate models on top of its predictive engines, so every anomaly call can be traced to the specific features that drove it.[7] The techniques exist, but using them adds computational costs and can slow down incident response which is why most vendors choose not to implement them by default.

ALERT #48291 – SUPPRESSED

Decision	Suppressed
Timestamp	2026-04-15 02:47:13
Matched pattern	Recurring deploy-window noise (pattern_id 7c4a)
Confidence	0.87
Similar suppressions	Last 30 days: 412
Escalation threshold	0.95

What a defensible suppression record looks like.

More vendors are starting to implement XAI techniques like reasoning chains and advertising explainability as a selling feature, but some still treat any information on the model as proprietary IP and refuse to expose the features driving a decision.

If the vendor for your AIOps tooling does not have explainability today and it is not on their roadmap, you may be answering to an auditor before they catch up.

Suppression Is a Security Event

Alert suppression keeps SOC teams functional, but it is also where the highest-consequence decisions get made. Suppression exists because of alert fatigue, and enterprise SOCs are overwhelmed by alert volumes, with many teams unable to keep pace. Suppression rules are written aggressively because analysts cannot function otherwise, and any system that decides which signals humans do not see is a system attackers want to influence.

IBM has documented that attackers already weaponize alert fatigue through a tactic called “alert storming,” where they launch high volumes of low-priority events to hide malicious activity in the noise they create.[8] Researchers have also demonstrated adversarial manipulation of ML-based detection pipelines, including evasion attacks that tweak inputs so malicious traffic appears normal to an intrusion detection system.[9]

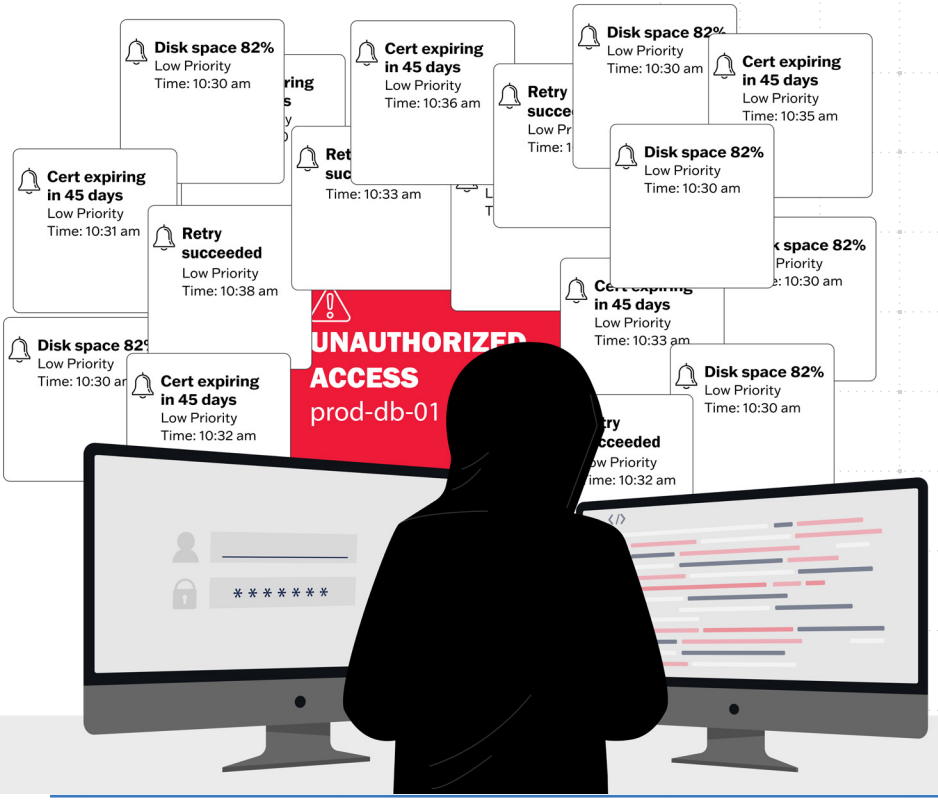

TRANSFORMATION

From Automation You Trust to Automation You Can Defend

Treat AIOps as a Governed AI System, Not a Tool

AIOps is often a feature of existing software, or an add-on functionality, rather than a separate application. Organizations must identify which applications or platforms have these capabilities and manage them as a governed AI system.

The vendor paperwork most procurement teams already have does not cover this. A SOC 2 report audits data handling and operational controls, but the framework was not designed to capture the risk of probabilistic, autonomous systems. A fully compliant SOC 2 program can coexist with material AI risk that no audit procedure will detect.[10] IBM's Cost of a Data Breach Report 2025 found that 97% of organizations experienced an AI-related security incident and lacked proper AI access controls, and 63% of organizations lacked AI governance policies.[11]

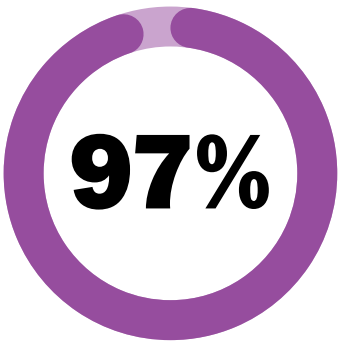

of organizations with an AI-related security incident lacked proper AI access controls[10]

of breached organizations overall lacked AI governance policies[11]

Build the Inventory Before You Need It

The IT Asset Management (ITAM) market is over $2B USD and growing because it's difficult for companies to keep a record of all of their hardware, software, and SaaS.[12] This problem gets compounded when you not only need to know what you have, but also if it has AI, and then if that AI can make autonomous decisions.

That means a configuration review of every monitoring, ITSM, and endpoint platform, focused on what automation rules, agentic capabilities, and autonomous remediation options are enabled. It means an audit of the platform's action logs for the past quarter to see what the system did on its own, not just what it is theoretically allowed to do. And it means interviewing the on-call engineers about the runbook steps they did not have to execute this month because something else executed them first.

For regulated financial industries, this is not optional. The U.S. Department of the Treasury Financial Services AI Risk Management Framework includes an AI inventory as a named control objective (GV-1.6), spanning six sub-objectives from shadow IT to portfolio-level risk analysis, and most institutions cannot even complete the framework's adoption stage questionnaire without one. The inventory is the foundation everything else in this framework depends on.

Shift the Question from "Did It Work?" to "Can We Explain It?"

Most operations teams measure mean time to resolution (MTTR), incidents per analyst, and alert fatigue scores as KPIs. Those are the right metrics for operational performance but not for governance. The governance question is how many autonomous actions can be traced back to a documented reason. A firewall that makes autonomous block decisions is expected to produce logs that prove which rule enforced it; AIOps should meet the same bar.

Operational reports, produced on a set schedule, should contain the inputs and logic that produced any autonomous AIOps decisions. These same items should be used when selecting a new AIOps vendor or feature to your existing software. What on-demand output does the platform produce when an auditor asks why an action was taken?

SOLUTION

The AIOps Accountability Framework

Four controls are essential to bringing autonomous IT decisions under enterprise AI governance—collectively referred to as **CAMP**.

Control	What it does
C — Catalog	Inventory every autonomous decision path
A — Authorize	Define blast radius and approval tiers
M — Monitor	Watch decisions over time for drift and manipulation
P — Prove	Produce evidence for every autonomous action

CATALOG:

Inventory Every Autonomous Decision Path

Start with a complete list of what the platform is allowed to do on its own. Any automated action that is AI controlled or influenced gets written down, and most organizations discover the list is longer than they thought once they force it onto paper.

Each entry gets three pieces of information attached to it:

- An owner, so accountability has a name
- A business justification, so the action's existence is defensible
- A blast radius, so the impact scope is known before something goes wrong

AUTHORIZE:

Define Blast Radius and Approval Tiers

Not every autonomous action deserves the same trust level. The framework ties each action to a tier, and the tier determines what approval the action needs before it runs.

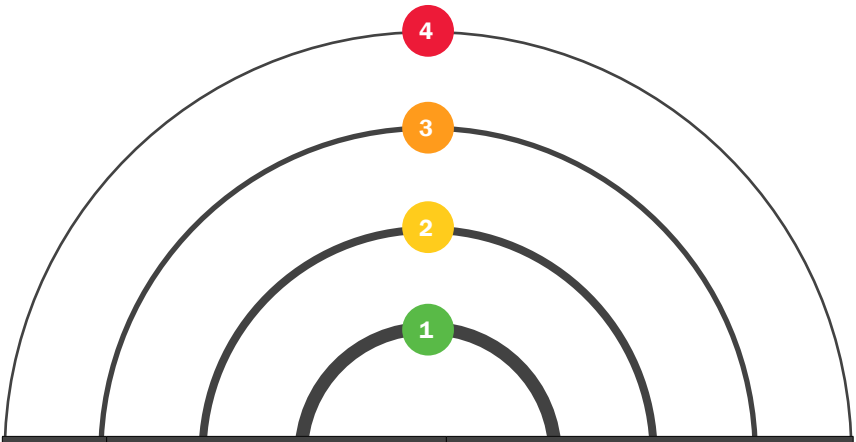

TIER	EXAMPLE ACTION	APPROVAL
1	Restart a stateless container	Automatic, logged
2	Scale infrastructure capacity	Automatic, logged, owner notified
3	Fail production workloads over to a secondary region	Human approval required
4	Delete production data or backups	Human approval, no exceptions

The tiering definitions will be individual to each company's regulations and risk profile. The output is a written authorization matrix that can be read by the board and enforced.

MONITOR:

Watch Decisions Over Time for Drift and Manipulation

The platform's behavior changes over time, and most of that change happens inside the vendor's model where the customer cannot see it.

Two signals get watched continuously:

Activity Volume
A sudden change in how often the platform acts on its own

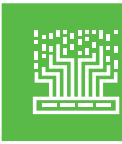

Activity Pattern
A shift in what kinds of actions the platform is taking

PROVE:

Produce Evidence for Every Autonomous Action

For every autonomous action the platform takes, the customer needs to be able to answer three questions on demand: what the platform did, when it did it, and what information it had at the time.

Three pieces of evidence get captured for each action:

THE ACTION
What the platform did, on what target, at what time

THE INPUT STATE
The data the platform was working from when it decided to act

EXPLAINABILITY
The reason the decision was made, with the relevant details

Governance Has to Catch Up to Autonomy

The gap between what AIOps is allowed to do and what the enterprise can explain is widening. AIOps is the place the gap is hardest to see because the work happens below the surface of what governance usually reviews.

The improved response times AIOps provides are real—and pulling the humans back into every decision seemingly defeats the purpose of using AI automation in the first place. The CAMP framework is a way to bring existing AIOps autonomy under real governance.

New SaaS features arrive through update channels, models retrain quietly in the background, and the tier-one action of today becomes the tier-three action of tomorrow as the business learns what the platform is actually doing in production. Until an audit, or worse a catastrophic incident occurs, organizations without governance of their AIOps won't be able to explain why.

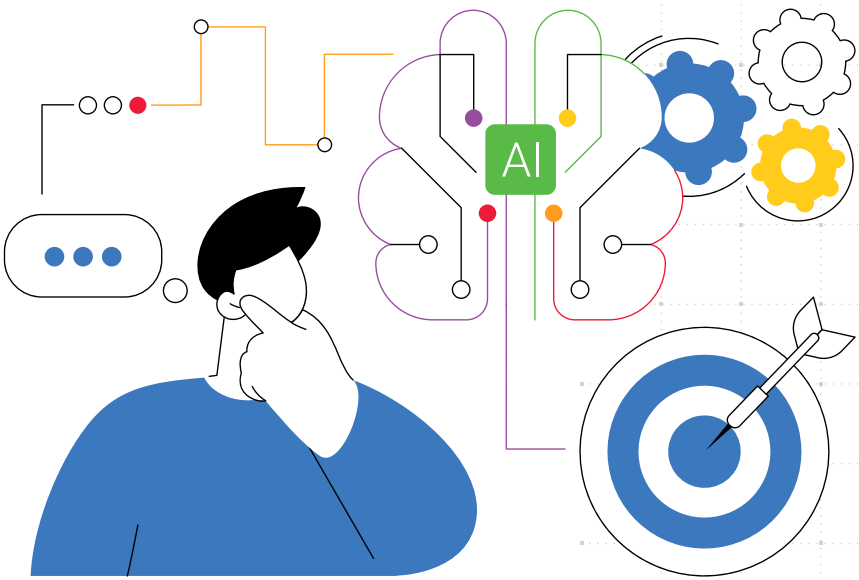

GLOSSARY

TERM	DEFINITION
agentic AI	AI systems capable of taking autonomous actions toward a goal without human input at each step. In AIOps, agentic capabilities let the platform decide and execute remediation, not just detect and recommend.
AIOps	Artificial intelligence for IT operations. A capability layer that applies machine learning to monitoring, event correlation, and remediation, typically embedded in existing monitoring, ITSM, or endpoint tools rather than deployed as a separate product.
alert storming	An attacker tactic that floods an environment with low-priority alerts to hide malicious activity in the noise. Documented as a technique that weaponizes alert fatigue against the SOC.
authorization tier	The level of pre-approval granted to an autonomous action. Tiers range from fully automatic for low-risk actions to human-required for high-blast-radius actions such as data deletion.
autonomous action	A decision or operation the AIOps platform takes without human review. Includes alert suppression, correlation, remediation, scaling, and ticket closure
blast radius	The scope of systems, data, or services impacted by an action. Used to determine what approval an autonomous action requires before it runs.
EU AI Act	European Union regulation governing AI systems, with specific obligations for high-risk applications including audit trails, explainability, and human oversight requirements.
explainability	The ability to produce a reviewable record of why an AI system reached a given decision. Required by governance and increasingly by regulation for autonomous AI actions.
Financial Services AI Risk Management Framework (FS AI RMF)	A U.S. Treasury and Financial Services Sector Coordinating Council framework released February 2026 that defines control objectives for AI governance in financial institutions.

TERM	DEFINITION
Local Interpretable Model-Agnostic Explanations (LIME)	An XAI technique that explains individual AI decisions by approximating the model locally with a simpler, interpretable model. Used to show why a specific anomaly was flagged or a specific action was taken.
SHapley Additive exPlanations (SHAP)	An XAI technique that calculates the contribution of each input feature to an AI decision using game theory. SHAP produces both local explanations for individual decisions and global explanations for the model as a whole, and is used in production AIOps deployments to make autonomous actions auditable.
SOC 2	An AICPA audit framework that evaluates a service organization's controls around data security, availability, processing integrity, confidentiality, and privacy. Does not evaluate whether the platform makes autonomous decisions or whether those decisions are governed.
Explainable AI (XAI)	A field of AI research and tooling focused on producing human-readable explanations for the decisions and outputs of machine learning models. XAI techniques include SHAP, LIME, reasoning chains, feature attribution, and confidence scoring. Required by governance and increasingly by regulation for autonomous AI systems.

CITATIONS

SOURCES

1. Thoughtworks, "AIOps: What we learned in 2025," January 2026
2. ServiceNow, "ServiceNow's AIOps and Gartner's Event Intelligence Solutions: A Perfect Match," March 2025
3. U.S. Department of the Treasury, "Treasury and Financial Services Sector Coordinating Council Release AI Risk Management Framework," February 2026
4. European Union, "Regulation (EU) 2024/1689 (Artificial Intelligence Act)," June 2024
5. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," August 2025
6. ServiceNow, "ServiceNow to Extend Leading Agentic AI to Every Employee for Every Corner of the Business with Acquisition of Moveworks," March 2025
7. Walmart Labs, "Anomaly Detection for Incident Response at Scale," April 2024
8. IBM, "What Is Alert Fatigue?," November 2025
9. ISACA, "Combating the Threat of Adversarial Machine Learning to AI-Driven Cybersecurity," August 2025
10. Michel Hjazeen, "SOC 2 for AI Systems: The Missing Controls Framework," February 2026
11. IBM, "Cost of a Data Breach Report 2025," July 2025
12. Mordor Intelligence, "IT Asset Management Market Size & Share Analysis - Growth Trends and Forecast (2026 - 2031)," March 2026

ABOUT

exactmarket™

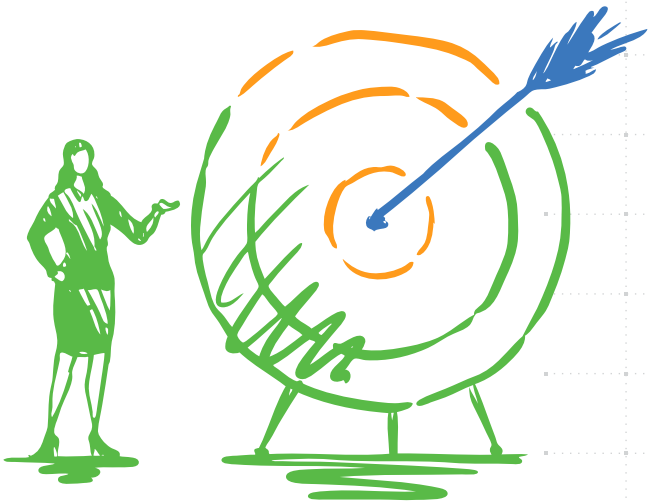

Founded in 2007, Exact Market is a woman-owned, WBENC-certified business focused on unifying marketing and technology around a shared vision to help enterprises innovate with confidence.

We bring together strategy, creativity, and data to help our clients stand out and stand for something. Our team of strategists, writers, designers, and technologists understands that the future of storytelling and software shares the same foundation: clarity, authenticity, and human connection.

We don't just help you talk about innovation.

WE HELP YOU LIVE IT.

Find out more: www.exactmarket.com

exactmarket℠

ELEPHANT IN THE ROOM EBOOK SERIES