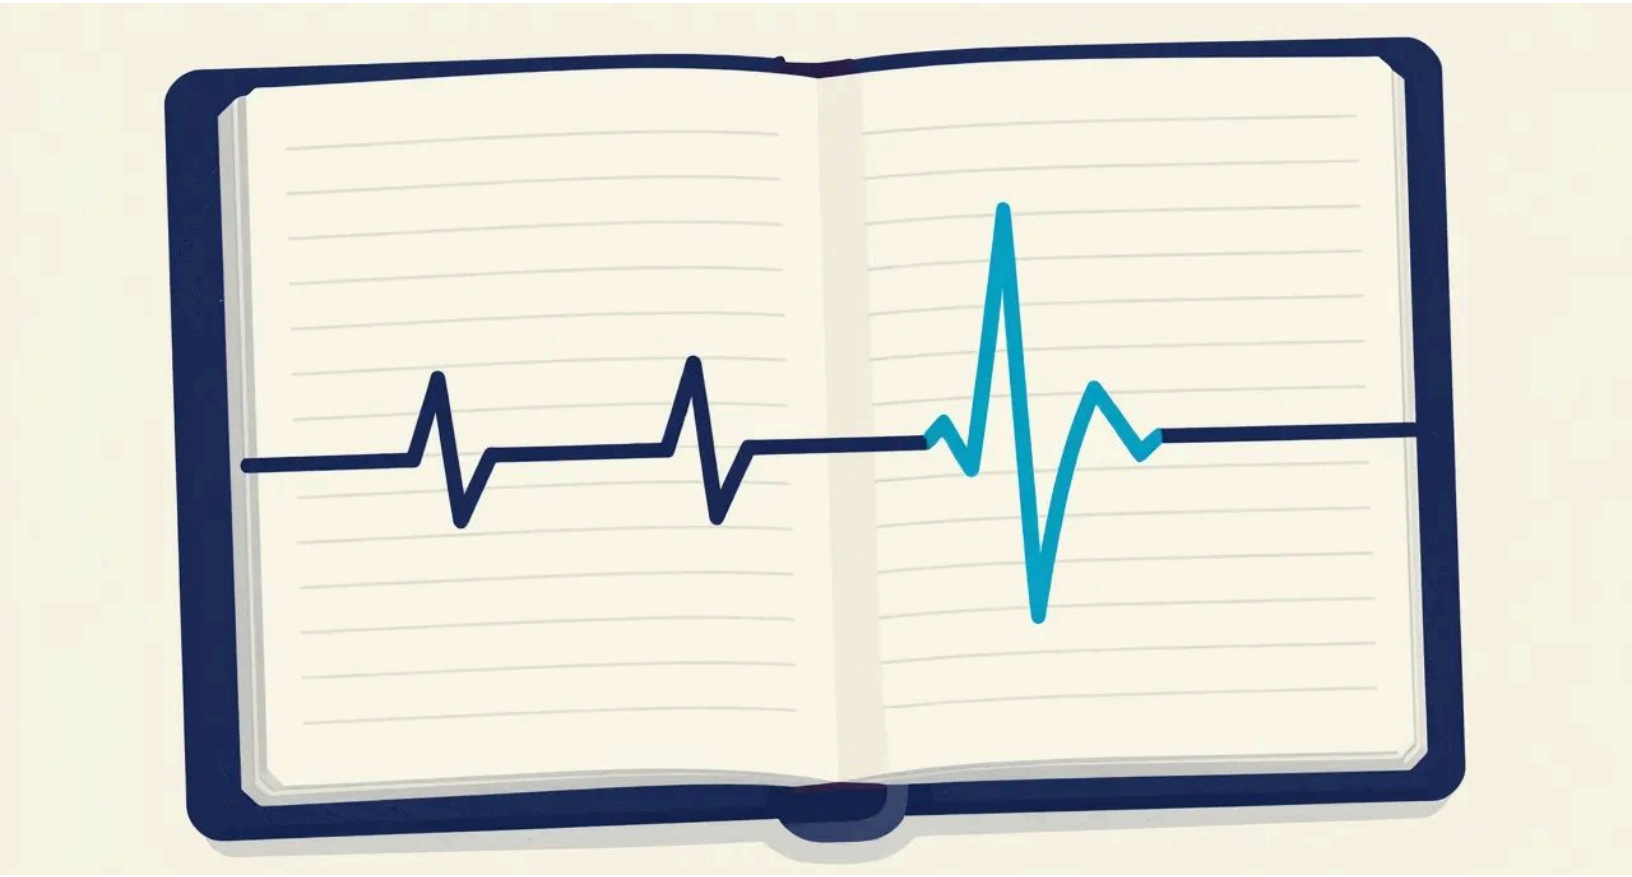


EXECUTIVE PAPER · INFRASTRUCTURE

Monitoring Is a Financial Control

Executives who treat monitoring as an IT purchase are carrying unpriced risk on the balance sheet. Here is how to manage it like every other control.

AUTHOR

Chris Watkins

SERVICE

Infrastructure

FOR

CEO, CFO, CTO · 13 min read

AT A GLANCE

The problem. Breaches now take longer to find than they did a year ago, and each one costs more. Most companies own monitoring tools. Far fewer can say who acts on an alert, how fast, and what that delay costs. AI agents with production credentials have added a new class of actor that most monitoring programs do not watch at all.

Why it persists. Monitoring is bought as software and handed to engineers. Nobody prices detection time, so nobody funds it. Alerts measure server health instead of customer harm, and they land in inboxes nobody watches.

The recommendation. Run monitoring as a control with an owner, a target, a defined response, and a monthly review, the same way finance runs reconciliations. Tie every alert to a business service and a person. Treat AI agents and service accounts as monitored identities. Report time to detect and time to act to leadership alongside revenue at risk.

The Decision in Front of You

Every executive team already pays for monitoring, usually as software licenses and a vendor contract renewed each year. The decision is whether that spend is a tool budget or a risk control.

The distinction matters because the money is lost in the minutes between something breaking and someone doing something about it. A monitoring tool can see a failure instantly. Whether the business learns about it in two minutes or two hours depends on who owns the alert and how it reaches them.

Uptime Institute's 2026 outage analysis found 57% of survey respondents said their most recent major outage cost more than $100,000, and for the second year running one in five put the cost above $1 million [1]. Per-site outage rates have fallen for five straight years, though the improvement is slowing [1]. When failures become rare, teams get less practice responding to them, which makes each slow first hour more expensive.

What's Happening

Detection is getting slower, and slow detection is expensive

IBM's 2026 Cost of a Data Breach study covered 602 breaches between March 2025 and February 2026. The average breach took 247 days to identify and contain, the first increase in five years [2][3]. The global average cost reached $4.99 million, a record, and the US average was $11.5 million [2][3].

Time is the cost driver. Breaches that ran past 200 days averaged $5.65 million, against $4.32 million for those contained sooner [3]. Organizations that used security AI and automation extensively cut breach costs by almost $2 million on average [2].

Downtime follows the same arithmetic. A 2024 study by Oxford Economics, conducted with Splunk, surveyed 2,000 executives at Forbes Global 2000 companies and estimated unplanned downtime costs that group $400 billion a year, about 9% of profits [4]. The sponsor sells monitoring software, so treat the headline as an upper bound.

Recent incidents show where monitoring breaks down

AWS, October 2025. A race condition in DynamoDB's automated DNS management left the service's regional endpoint in Northern Virginia with an empty DNS record, and the automation could not repair it. The disruption ran from 11:48 p.m. on October 19 to 2:20 p.m. the next day, Pacific time [5]. The automation built for resilience caused the failure. AWS disabled it worldwide until a fix was in place [5]. One industry estimate put US economic losses at $500 million to $650 million [6].

Cloudflare, November 2025. A database permissions change doubled the size of a configuration file and crashed the software that routes much of Cloudflare's traffic [7]. Automated tests flagged the problem at 11:31 UTC. The team first read the symptoms as a large DDoS attack, partly because Cloudflare's separately hosted status page happened to fail at the same time. The real cause was identified at 14:24, almost three hours after the first alert [7]. Detection took minutes. Diagnosis took hours, because the signals were ambiguous and the team had no fast way to tie the outage to the change that preceded it.

PocketOS, April 2026. An AI coding agent working in a test environment hit a credential error, found an API token in an unrelated file, and used it to delete the company's production database and its backups [8]. The firm, which runs reservation software for car rental companies, spent days rebuilding records from payment, calendar, and email systems before its cloud provider restored the data [8]. Explicit safety instructions and a top-tier model did not stop it.

Binance, October 2026. Attackers took an estimated $6.3 billion from the exchange on October 29, the largest crypto theft on record [9][10]. Binance's preliminary finding is that a compromised laptop belonging to a DevOps engineer gave the attacker access to a service container, which was used to submit unauthorized withdrawals [9]. About $3 billion was frozen, and Binance said it would cover customer losses [9]. How quickly the withdrawals were detected has not been published. Executives should ask whether their own systems would flag a legitimate internal service suddenly moving money at abnormal volume, and how many minutes that would take.

A case worth knowing: Knight Capital, 2012

On August 1, 2012, a defective code deployment at Knight Capital sent more than 4 million orders into the market in 45 minutes while trying to fill 212 customer orders [11]. The firm lost more than $460 million [12].

Before the market opened, an internal system sent 97 automated emails to staff that identified an error in the faulty router [11]. The emails were never designed as alerts, so nobody treated them as urgent [13]. The SEC later found Knight relied on people to watch its financial exposure instead of generating automated alerts, and fined the firm $12 million on top of the trading loss [12][13].

SIDEBAR: THE 97 EMAILS

Knight's systems warned the company before the market opened on the day of a loss exceeding $460 million [11][12]. The messages went to an inbox, read like routine system output, and belonged to nobody in particular. Most organizations have their own version of those emails today, usually in a shared mailbox or a Slack channel with notifications muted.

Why the Usual Approach Fails

Most monitoring programs start with a product decision. Someone evaluates vendors, buys one, and engineers instrument what the tool makes easy. Servers report CPU and memory. Databases report connections. Within a year the company has hundreds of charts and thousands of alerts.

Several habits make this expensive.

Teams watch machine health when customers feel symptoms. A checkout page that takes eight seconds is costing money right now, even if every server behind it reports normal load. Google's site reliability engineering guidance tells teams to alert on user-facing symptoms, measured through four signals: latency, traffic, errors, and saturation [14]. Most corporate monitoring still pages on machine health.

Alerts have no owner. An alert sent to a group email is an alert sent to nobody. When everyone receives it, each person assumes someone else is handling it.

Changes are invisible to responders. In both the AWS and Cloudflare incidents, the trigger was an internal change or automated process [5][7]. When monitoring shows symptoms without showing what changed in the last hour, responders chase the wrong theory.

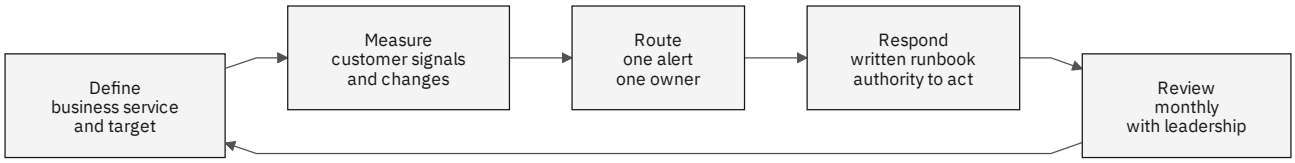
Machine identities go unwatched. AI agents, service accounts, and deployment pipelines now hold production credentials. In IBM's 2026 study, 92% of organizations with AI-related incidents lacked proper AI access controls [3]. Of organizations that secure non-human identities in AI workflows, only 36% use behavioral monitoring on them [3]. PocketOS and Binance both turned on an automated or internal identity doing something it should not have been able to do.

Auditors expect more than tools. NIST's Cybersecurity Framework 2.0, published in February 2024, makes continuous monitoring a core category of its Detect function [15]. ISO/IEC 27001:2022 added a dedicated control, Annex A 8.16, that expects organizations to monitor systems for anomalous behavior and act on potential incidents [16][17]. Both frameworks expect evidence that monitoring operates and that someone acts on what it finds.

The Framework: The Monitoring Control Loop

A financial control has an owner, a defined check, a threshold, a response, and a periodic review. Monitoring should have the same parts.

Exhibit 1. The Monitoring Control Loop: every stage needs a named owner



Source: author's framework, adapted from control design practice and Google SRE monitoring guidance [14].

Define. Start with the services customers and revenue depend on: ordering, invoicing, dispatch, payments. For each, agree on a target in business terms. "Orders confirm within five seconds, 99.9% of the time" is a target an executive can hold someone to, because it describes what the customer experiences.

Measure. Instrument the four symptom signals for each service [14], then add one business signal, such as orders per minute compared with the same hour last week. Record every deployment and configuration change on the same timeline, so responders can see what changed before the alert fired. Watch what automated identities do: an AI agent or service account deleting data, moving money, or reading credentials outside its normal pattern should page a person.

Route. Every alert maps to exactly one on-call owner with a phone that rings. If an alert does not require someone to act, delete it.

Respond. Each alert links to a short written runbook: what to check, who to call, and what the owner may do without asking permission. Pre-approved authority is what turns a ten-minute response into a two-minute one. Rolling back the most recent change should be on that list.

Review. Once a month, leadership sees four numbers per critical service: incidents, median time to detect, median time to act, and estimated revenue at risk, shown against the prior six months.

What a slow first hour costs

The following is an illustrative model for a hypothetical mid-size distributor. Substitute your own figures.

Exhibit 2. Illustrative cost of detection delay for one order-platform outage

Assumption or result	Email-only alerting	Paged owner with runbook
Order platform revenue per operating hour	$200,000	$200,000
Time from failure to someone acting	45 minutes	5 minutes
Revenue exposed before response begins	$150,000	$16,700
Difference per incident		$133,300
At six incidents per year		about $800,000

Source: author's illustrative estimate. Assumes $2.4M daily order value over a 12-hour operating day and full revenue loss while the platform is down. Real losses vary with recovery of deferred orders.

The model is deliberately simple. Its purpose is to put a dollar figure on minutes, which is the number most monitoring budgets never see.

Costs, Risks, and Tradeoffs

This approach costs money and attention. It is also the wrong answer in some situations.

Exhibit 3. Tradeoffs of running monitoring as a control

Area	Cost or risk	Mitigation
On-call staffing	Paying engineers to carry a pager; burnout if alert volume is high	Cap alerts per shift; delete any alert that did not need action
Tool spend	Tool costs grow as you instrument more services and keep more data	Instrument critical services first; set retention by need
Leadership time	A monthly review adds a standing meeting	Fold it into an existing operations or risk review
False confidence	Green dashboards can hide unmonitored services and unwatched agents	Keep an inventory of critical services, AI agents, and service accounts, with coverage for each
Automation risk	Automated remediation can fail in new ways, as AWS found [5]	Put rate limits and human approval on automated actions that delete, move money, or remove capacity
When this is wrong	A small company running mostly on SaaS gets little from heavy internal monitoring	Rely on vendor status pages, contractual SLAs, and simple external checks

Source: author's analysis.

How to Act

Exhibit 4. A one-year roadmap

Window	Actions	Owner	Milestone
First 30 days	List the five services that carry the most revenue or regulatory exposure. Name one accountable owner for each. Inventory every AI agent and service account with production write access.	CTO with business unit leaders	Five services, five owners, signed off; agent inventory complete
First 90 days	Set a business target for each service. Instrument the four symptom signals plus one business signal. Put deployments and config changes on the alert timeline. Route alerts to a paged on-call rotation. Write a runbook per alert.	Engineering leads	Every critical alert reaches a named person by phone
First 90 days	Alert on destructive or financial actions by automated identities. Remove standing production credentials from AI agents.	Security lead	No agent can delete production data or move funds without a human approval
First 90 days	Start reporting time to detect and time to act.	Operations or SRE lead	First monthly report delivered to leadership
First 365 days	Extend coverage to the next tier of services. Cut alert volume by removing anything that never required action. Map coverage to NIST CSF DE.CM and ISO 27001 A 8.16 evidence.	CTO with risk or compliance	Median time to act under 10 minutes on critical services; audit evidence on file

Source: author's recommendation.

Questions to Ask Your Team

1. Which five services would cost us the most per hour if they stopped?
2. For each of those, who gets called first, and how?
3. How long did it take us to notice our last three outages, and who noticed first: our systems, our staff, or our customers?
4. When an alert fires, can the responder see every change made in the previous hour?
5. Which AI agents and service accounts can delete production data or move money, and what alerts when they do?
6. What can the on-call engineer do at 2 a.m. without asking permission?
7. If an auditor asked for evidence that we act on monitoring alerts, what would we show them?
8. What does a minute of downtime on our order system cost, in dollars?

Closing

Finance would never accept a reconciliation that nobody owns, checked on no schedule, with exceptions emailed to a shared inbox. Most companies run monitoring exactly that way, and many now give AI agents production credentials with even less oversight. Assign an owner, set a target, route alerts to a person, watch the machines as closely as the people, and review the numbers monthly.

Sources

[1] Uptime Institute. "Uptime Announces Annual Outage Analysis Report 2026." Press release via Business Wire. May 13, 2026. https://www.businesswire.com/news/home/20260513289344/en/Uptime-Announces-Annual-Outage-Analysis-Report-2026.

[2] IBM. "IBM Study: One in Four Malicious Breaches are AI-Enabled, Costing Companies $6 Million on Average." Press release, Cost of a Data Breach Report 2026. July 29, 2026. https://www.placera.se/telegram/avanza/39_5f4ce762-b6a2-495d-b2cc-d9d94dc4560d.html.

[3] Baker Donelson. "Ten Takeaways from IBM's 2026 Cost of a Data Breach Report." 2026. https://www.bakerdonelson.com/ten-takeaways-from-ibms-2026-cost-of-a-data-breach-report.

[4] Splunk and Oxford Economics. "The Hidden Costs of Downtime." 2024. https://splunk.com/en_us/pdfs/gated/ebooks/the-hidden-costs-of-downtime.pdf.

[5] Amazon Web Services. "Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region." October 2025. https://aws.amazon.com/message/101925.

[6] Control Plane. "Cloud Outage Tracker: Major AWS, Azure, Google Cloud, and Cloudflare Outages Since 2020." Citing Parametrix loss estimate. 2026. https://controlplane.com/blog/post/cloud-outage-tracker.

[7] Cloudflare. "Cloudflare outage on November 18, 2025." Post-mortem. 2025. https://blog.cloudflare.com/18-november-2025-outage/.

[8] Live Science. "'I violated every principle I was given': AI agent deletes company's entire database in 9 seconds, then confesses." 2026. https://www.livescience.com/technology/artificial-intelligence/i-violated-every-principle-i-was-given-ai-agent-deletes-companys-entire-database-in-9-seconds-then-confesses.

[9] CoinDesk. "Binance loses over $6 billion in the largest crypto hack to date." October 29, 2026, updated October 30, 2026. https://www.coindesk.com/business/2026/10/29/potential-multi-billion-binance-hack-hitting-crypto-markets.

[10] CoinDesk. "$6 billion Binance hack sends bitcoin, ether tumbling." October 30, 2026. https://www.coindesk.com/markets/2026/10/30/usd6-billion-binance-hack-sends-bitcoin-ether-tumbling-standard-chartered-sees-more-losses-ahead.

[11] U.S. Securities and Exchange Commission. SEC News Digest, October 16, 2013. https://www.sec.gov/news/digest/2013/dig101613.htm.

[12] U.S. Securities and Exchange Commission. "SEC Charges Knight Capital With Violations of Market Access Rule." Press Release 2013-222. 2013. https://www.sec.gov/news/press-release/2013-222.

[13] WilmerHale. "Knight Capital Settles Rule 15c3-5 Violations with SEC, Agrees to Pay $12 Million." Summary of SEC Order, Exchange Act Release No. 70694. 2013. https://www.wilmerhale.com/en/insights/client-alerts/knight-capital-settles-rule-15c3-5-violations-with-sec-agrees-to-pay-12-million.

[14] Beyer, B., Jones, C., Petoff, J., and Murphy, N. R., eds. *Site Reliability Engineering: How Google Runs Production Systems*, Chapter 6, "Monitoring Distributed Systems." O'Reilly. 2016. https://sre.google/sre-book/monitoring-distributed-systems/.

[15] National Institute of Standards and Technology. "The NIST Cybersecurity Framework (CSF) 2.0." NIST CSWP 29. 2024. https://doi.org/10.6028/NIST.CSWP.29.

[16] International Organization for Standardization. "ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection: Information security management systems: Requirements." Annex A, Control 8.16. 2022. https://www.iso.org/standard/27001.

[17] ISMS.online. "ISO 27001:2022 Annex A 8.16 Monitoring Activities." 2025. https://isms.online/iso-27001/annex-a/8-16-monitoring-activities-2022.

The views in this paper are the author's own and do not represent any employer. Scenarios are hypothetical. This paper is not legal, financial, or compliance advice.