StackGen

State of Reliability 2026

What's driving incidents, what SREs are doing about them, and how AI is reshaping both

Published 1 June 2026

StackGen Research

Executive summary

Six key findings, drawn from 174,348 unplanned incidents across 360+ companies between 2018 and June 2026.

1. MTTR clusters by category tier and has been roughly stable within each tier since 2023 (see Section 2). Application-tier industries (developer tooling, consumer internet, business software, observability, e-commerce) resolve incidents at a median of approximately 1.7 hours. Industry-infrastructure industries (cloud infrastructure, communications, payments, identity, data platforms) resolve at 3 to 4 hours. AI model providers, functionally infrastructure but operationally distinct, sit at approximately 1.0 hour and have improved to 0.9 hours in 2026 YTD.

> → *Benchmark against tier-mates, not against a corpus-wide average; a 3-hour Cloud Infrastructure median is tier-typical, not under-performing.*

2. Firm-level differences in median resolution time are approximately three times larger than industry-level differences (see Section 4). Industry explains roughly 8 percentage points of MTTR variance; failure mode and root cause together explain another 7; the company itself explains approximately 20 percentage points — roughly one-third of all variance that any observable predictor can explain, against ~65% of variance that is incident-level noise. The lever is not industry positioning; it is the firm's Response Maturity. Firm-specific drivers include the technology stack, application architecture, deployment practices, SRE capability investments, and how the firm discloses incidents publicly.

> → *Reliability outcomes are improvable: the company-level lever is the largest one available; positioning is a smaller lever than what your team chooses to do.*

The first two findings describe the shape of MTTR outcomes. The next four describe what's driving them. Two terms anchor the rest of the report: **Incident Profile** is a team's pattern of incidents (failure-mode, root-cause, and remediation mix); **Response Maturity** is the team's capability to handle them (Context, Tooling, People, and AI). These findings that follow describe what the data says about each.

3. Most teams' incident patterns map to one of six recurring archetypes (see Section 5.5). Six archetypes emerge from clustering the 95 companies with at least 50 classified incidents each: Dependency-Driven, Velocity-Driven, Scale-Driven, Data-Integrity-Driven, Substrate-Driven, and AI-Quality-Driven (emergent). Each combines a characteristic failure signature with a characteristic recovery fingerprint and remediation fingerprint. The strongest data point is that Dependency-Driven firms apply 'wait for upstream fix' as the primary remediation in 74% of cases, which mechanically explains the slowest median MTTR of any archetype. The combination of failure shape and recovery shape predicts more than tier or industry alone.

> → *Self-identify your archetype before choosing your next reliability investment; an investment that lifts one archetype's effective MTTR will return less for another.*

4. Archetypes are sticky — stable over time — but show a maturation arrow (see Section 5.6). After holding archetype definitions fixed and re-assigning each company based on its 2023 incidents and then its 2025 incidents, 60.7% of firms (17 of 28) stay in the same archetype. Of the 11 firms that migrate, the direction is

consistent: away from Velocity-Driven (the change-driven archetype) toward Substrate-Driven, Data-Integrity-Driven, and AI-Quality-Driven (the longer-tail archetypes). The pattern reads as a maturation arrow: as systems grow, the dominant failure shape moves from 'we shipped a bad change' toward 'our scale, our data, and our substrate are now the thing that breaks'. Incident Profile is partly a maturity signal.

> → *Architectural pre-investment should target the archetype you are heading toward, not the one you are in today; this gives your team a year or two of lead time on the long-tail failure modes.*

5. Cross-organisation cascade is the largest single failure pattern at approximately 21% of incidents; the share peaked in 2025 and has fallen in 2026 YTD (see Section 7). Cross-organisation cascade represents 52% of the classified subset of incidents, or approximately 21% under the conservative estimated-share denominator. By year, the estimated share is 17%, 24%, 29%, and 22% across 2023 through 2026 YTD. The 2026 decline reflects smaller-scale cascade events so far this year (the largest 2026 event so far affected 15 disclosing companies; 2025 had three events that affected more than 50 each), not a structural change in cascade frequency.

> → *Architectural investment in multi-vendor failover for the top five to ten upstream dependencies remains the single highest-share move available, irrespective of the 2026 share dip.*

6. AI-related incidents grew from 0.85% to 5.12% of disclosed incidents between 2023 and 2026 YTD, a 6x rise in three years (see Section 6). Excluding Communications industry, AI Model Provider as an industry grew incident volume 38x from 2022 to 2025. AI now appears in three distinct categories: AI as an upstream service that fails, AI as a model-quality issue inside customer-facing AI features, and as autonomous AI agents themselves taking destructive action on production systems. Broad-tier MTTR has been flat since 2023, suggesting the operational gains promised by AI-based SRE tooling have not yet been felt at industry scale; the four components of Response Maturity (Context, Tooling, People, AI) require architectural pre-investment that is still ramping (see Section 8).

> → *AI is already material as a failure surface, but not yet material as a productivity gain inside SRE — the gap is architectural readiness, not tooling availability.*

The rest of the report unpacks each finding. Part I covers MTTR analysis: category tier, industry, and firm-level differences. Part II covers incident shape and remediation: the Incident Profile dimensions and archetypes, AI in the incident mix, cross-organisation cascade, and Response Maturity.

Key numbers and headline findings

A press-ready summary of the dataset and its most-extractable claims. Each sentence is supported by the section noted in parentheses.

174,348	360+	29,708
unplanned incidents analysed	companies on public status pages	classified or low-confidence-classified

Dataset v16.29, period 2018–June 2026. Same-cohort comparisons use 217 companies present every year of 2023-2026 YTD.

Headline findings — attribution-ready

1. AI incidents grew six-fold in three years. AI-related incidents grew from 0.85% to 5.12% of disclosed incidents between 2023 and 2026 YTD, a 6x rise in three years. The category covers AI as an upstream service, AI as a model-quality issue, and AI agents themselves causing production incidents. *(see Section 6)*

2. One in five incidents is now caused by an upstream provider the company doesn't control. Cross-organisation cascade represents approximately 21% of all disclosed incidents and is the largest single failure pattern in the dataset, with median resolution time of 309 minutes (3.2x slower than internally-caused configuration failures). *(see Section 7)*

3. Company reliability maturity matters approximately three times more than industry positioning. Industry membership explains roughly 8 percentage points of variance in median resolution time; the company itself explains approximately 20 percentage points. Two SRE teams in the same industry can plausibly differ by 3x in how quickly they recover. *(see Section 4)*

4. AI agents have destroyed production systems in at least nine publicly documented incidents since mid-2025, from the Replit autonomous agent deleting a production database in July 2025 to the PocketOS/Cursor token-scanning incident that destroyed a database and all backups in a 9-second API call in April 2026. The true count is unknown — these incidents are structurally invisible to status-page methodology — but the documented trajectory is 1, 3, 8, 7 from 2023 through May 2026. *(see Section 6.3)*

5. The most common fix for a production incident is to wait. Across StackGen's coded post-mortem corpus, 'wait for upstream fix' is the single most-applied primary remediation pattern at 13.6% — ahead of restart (10.0%), rollback (8.9%), or bug fix (8.9%). The waiting pattern mechanically explains why Dependency-Driven firms have the slowest median MTTR of any archetype. *(see Section 8)*

6. AI providers themselves are the fastest-resolving industry in the dataset. Median time-to-resolve at AI Model Providers has fallen from 71 minutes in 2023 to 53 minutes in 2026 YTD, well below the Application tier (~100 minutes) and the Industry-Infrastructure tier (3 to 4 hours). The combination of architectural redundancy, fast deployment pipelines, and rich telemetry explains the gap. *(see Section 3)*

7. The single biggest reliability investment available is multi-vendor failover. Cross-organisation cascade is 21% of incidents and the slow tier of the failure-mode distribution at 309 minutes median resolution; the four-step escalation ladder in Section 8 lays out the architectural path from 'wait for upstream' to 'dual-provider hot-swap'. *(see Section 8.3)*

Methodology note for analysts. All trend comparisons use the 217-company same-cohort lens unless noted otherwise. The dataset is structurally blind to incidents not disclosed on public status pages, which means it under-represents internal-only failures, training-pipeline incidents, and AI-agent-induced production changes. Where this matters for a specific finding, the limitation is flagged in context. Full methodology in Appendix C.

1. About this report

The dataset. The State of Reliability 2026 dataset contains 174,348 unplanned incidents from over 360 companies' public status pages, primarily Atlassian Statuspage and incident.io vendors, collected between 2018 and June 2026. Of these, 29,708 are classified or low-confidence-classified using StackGen's Failure Mode, Root Cause, and Remediation Taxonomies (compact summaries in Appendix A; full versions at stackgen.com/sor2026/taxonomies). The analytical window is 2023 to 2026 YTD; pre-2023 data is sparser and is shown where the long-term trajectory is informative (Section 7). Same-cohort comparisons use the 217 companies present every year of 2023, 2024, 2025, and 2026 YTD.

Versioning. This is report version v1.9.3, built on dataset version v16.29 (a snapshot promoted 1 June 2026 with the same 174,348 incidents enriched with the 2023-slice classification fold). Report version and dataset version progress independently.

The framework. This report introduces a two-part framework for thinking about a team's reliability position. **Incident Profile** describes what a team is dealing with: the proportional mix of failure modes, root causes, and remediations the team encounters most often, expressed across three dimensions of StackGen's analytical spine. Section 5 unpacks the three dimensions and the six recurring archetypes that the cohort clusters into. **Response Maturity** describes how prepared the team is to handle its Incident Profile: the combination of Context (observability and signals), Tooling (automation and runbooks), People (skills and on-call practice), and AI (the augmentation layer). Section 8 unpacks the four components. The two combine: Incident Profile × Response Maturity → effective MTTR. Both are levers.

Audiences. The report serves four audiences. SRE practitioners and leaders should find peer-comparison benchmarks and practical takeaways at the end of each section. Engineering leadership building reliability investment cases will find the firm-level differences finding (Section 4) and the archetype map (Section 5.5) the most actionable. Product leaders shipping AI features will find Section 6 directly relevant. Industry analysts will find the methodology in Appendix C and the data tables in Appendix B, including company-cluster-robust significance tests with forest plots.

Scope and limitations. The findings here focus on disclosed incidents. Internal incidents that operators do not post to public status pages are out of scope, which

means the dataset under-represents certain categories of failure: training-pipeline failures, retrieval-corpus drift, and AI-agent-induced production changes that operators do not publicly disclose as incidents. Where this matters for a specific finding, the limitation is flagged in context. Methodological details (classifier coverage, the raw-versus-estimated denominator, the MTTR proxy definition, severity normalisation, the within-company paired test, and the Microsoft Azure originator gap) are in Appendix C.

Part I

MTTR analysis

How quickly incidents are resolved, broken down three ways: by category tier (Section 2), by industry (Section 3), and by firm (Section 4).

2. Median time-to-resolution by category tier

Category position predicts MTTR more strongly than year does. Three tiers emerge.

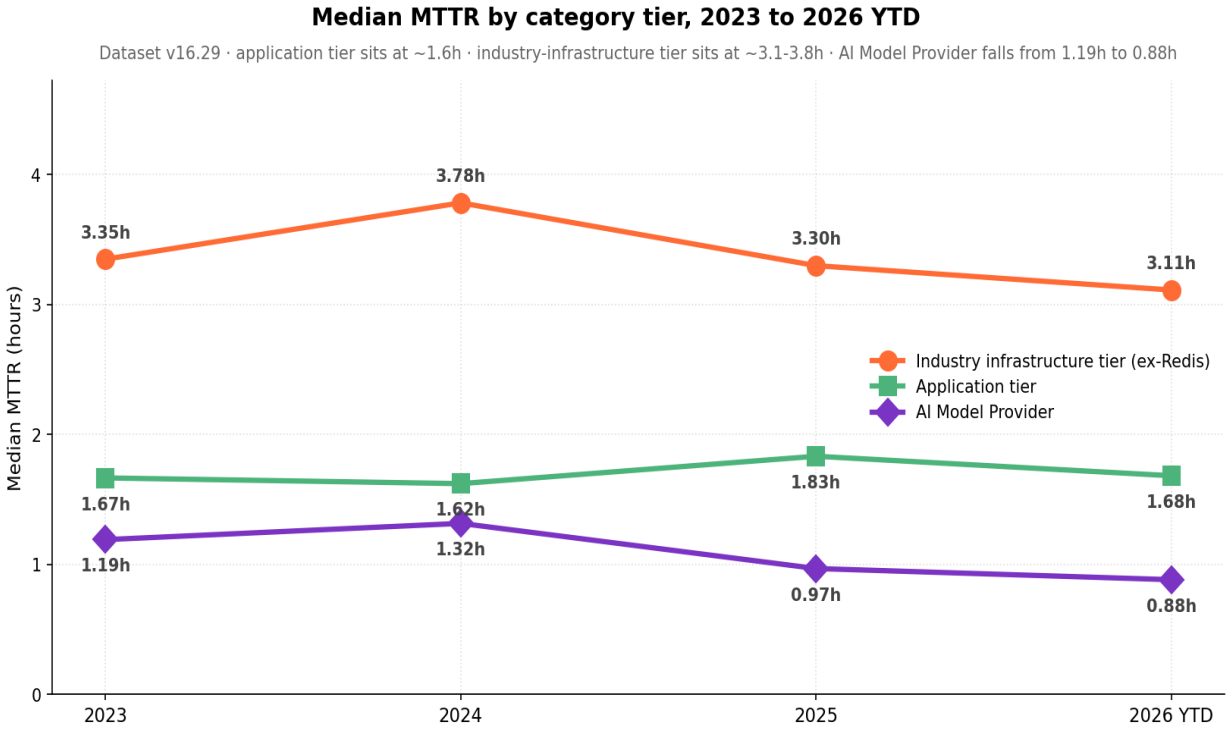

Figure 1. Median MTTR by category tier, 2023 to 2026 YTD (same-cohort lens). Dataset v16.29. Excludes maintenance periods and advisory postings (Duration > 7 days).

Application tier covers industries whose products are consumed primarily by end users or by other application developers. Failures tend to be self-contained, with fast rollback paths and blast radius bound to the operator's own product. Median MTTR sits between approximately 1.6 and 1.8 hours every year from 2023 through 2026 YTD.

Industry-infrastructure tier covers industries whose products are consumed primarily as critical inputs to other businesses. Three structural features lengthen incident response: customers are predominantly other businesses (i.e., enterprise or B2B customers rather than end consumers), failures cascade into many downstream systems, and resolution requires coordination across regulators, carriers, customers, or partners. Median MTTR sits between 3.1 and 3.8 hours across the window.

AI Model Provider is functionally infrastructure but operates at application-tier MTTR. Median falls from approximately 1.2 hours in 2023 to 0.9 hours in 2026 YTD. Three structural features plausibly explain the gap: an architectural pattern where most outages can be handled by serving a cache or falling back to a sibling model; operational staffing concentrated in scale-up companies with fast deployment pipelines; and a telemetry surface unusually rich in per-token latency, error-rate, and output-quality signals.

Pre-2023 data is shown only where the long-term trajectory is informative (Section 7). The dataset's coverage broadened materially in 2023 as approximately 48 large

industry-infrastructure operators entered scraped coverage; earlier per-tier medians are correspondingly thinner. Same-cohort distribution detailed in Appendix B, Table B1.

> ***What this means for SREs.*** *Industry-tier MTTR benchmarks should be set against the right tier. Application-tier categories have a 100-minute median; industry-infrastructure categories have a 3 to 4 hour median. A team in cloud infrastructure or payments comparing itself to a 100-minute benchmark is comparing against the wrong tier. In operational terms: a 200-minute Industry-Infrastructure incident consumes ~3.3 hours of customer-impacting degradation × however many services depend on the failing component, which for a CPaaS or payments provider is typically dozens.*

3. Industry MTTR: distribution and ranking

Median MTTR varies meaningfully across industries within a tier, but the variation maps cleanly onto the tier framing from Section 2 (same-cohort lens).

3.1 Ranking trajectory

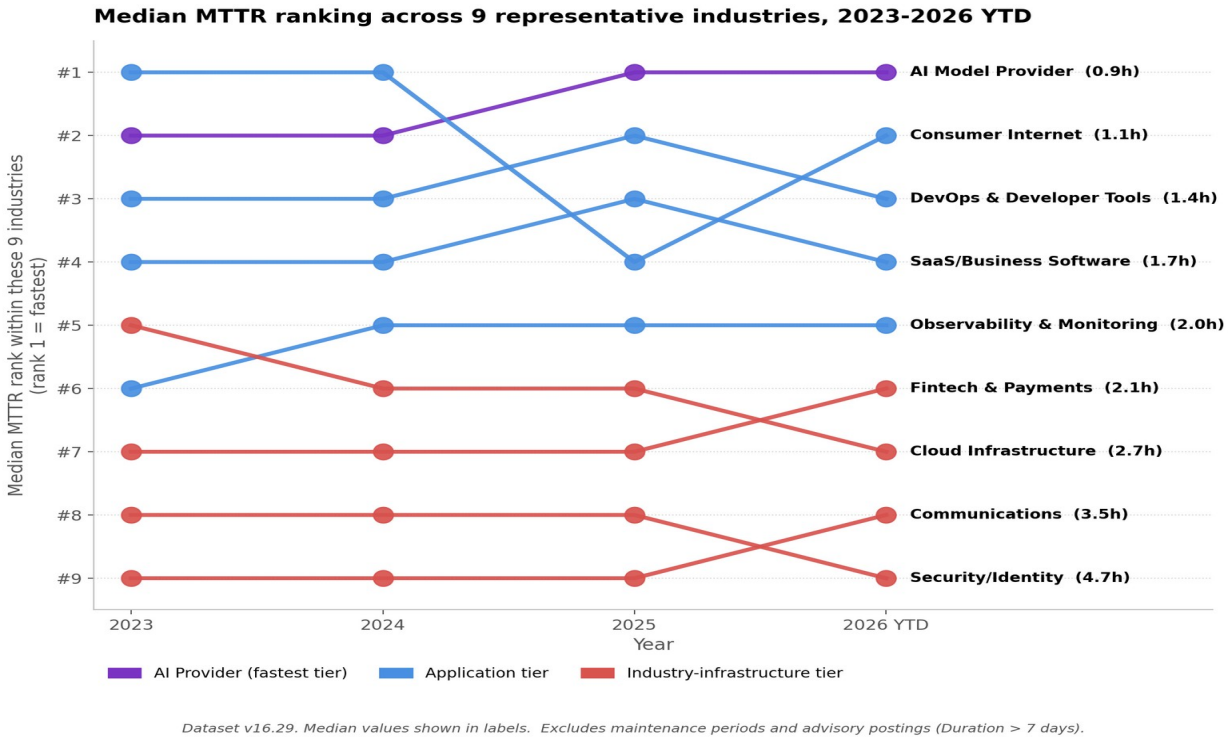

Figure 2. Median MTTR ranking across 9 representative industries, 2023 to 2026 YTD. Dataset v16.29. AI Model Provider has held rank 1 since 2024. Application-tier industries (Consumer Internet, DevTools, SaaS, Observability) cluster in the top half; industry-infrastructure industries cluster in the bottom half. Median MTTR in each industry's endpoint label. Excludes maintenance periods and advisory postings (Duration > 7 days).

The chart sorts cleanly into two clusters that match the Section 2 tier framing. The reduced industry set (9 of the ~15 in the dataset) keeps the chart legible at print resolution; the full per-industry table is in Appendix B (Table B2).

- **AI Model Provider** has held rank 1 since 2024, with median MTTR improving to 0.9 hours (53 minutes) in 2026 YTD.
- **Consumer Internet** sits at rank 2 in the Application tier, with a median around 1 hour. The 2025 uptick to 1.8h is a single-year wobble that comes back down in 2026 YTD.
- **Cloud Infrastructure** has slid two positions in this set, from rank 5 in 2023 to rank 7 in 2026 YTD, consistent with the rising complexity of multi-region cloud operations.
- **Industry-infrastructure cluster** (Communications, Fintech & Payments, Security/Identity, Cloud Infrastructure) sits in the slower band at ranks 6-9 across all years.

An infrastructure-tier tail observation worth flagging. The within-tier 75th-percentile MTTR for Infrastructure-tier industries has widened from 6.2h in 2023 to 8.0h in 2026 YTD on the same-cohort lens — a step-up between 2023 and 2024 followed by a plateau around 8 hours. The Application and AI-Provider tiers do not show the same widening; both are flat-to-declining at P75. This is consistent with the Maturation Arrow finding (Section 5.6): as Substrate-Driven and Data-Integrity-Driven archetypes grow share of the cohort, their long-tail failure shapes show up in the aggregate tail.

3.2 Industry incident volume

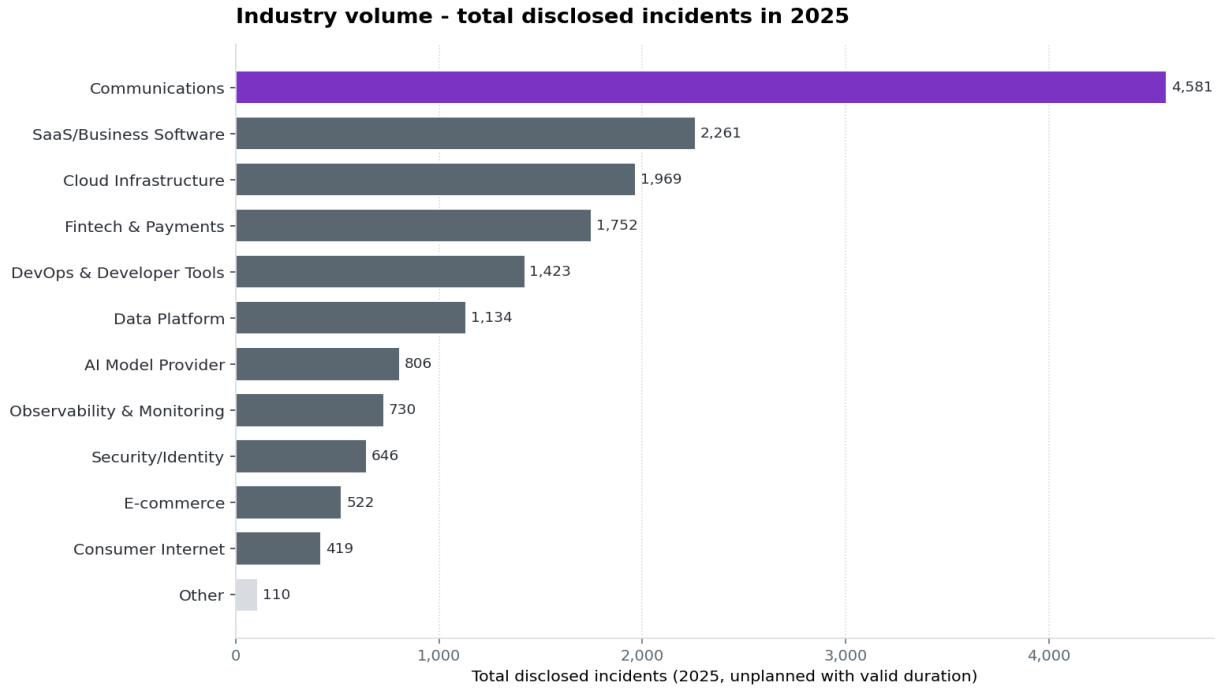

Source: SOR 2026 dataset v16.29 | Unplanned incidents with Duration_Minutes > 0 | 11 industries shown explicitly; 'Other' groups the remaining industry labels. Communications is the largest single industry (28% of 2025 charted rows on v16.29), heavily weighted by Twilio (~45% of Communications rows on v16.29).

Figure 3. Total disclosed incidents per industry, 2025. Dataset v16.29. Communications dominates volume at 4,581 incidents.

Communications volume is inflated by per-route disclosure at certain communications-platform providers (Twilio alone accounts for approximately 45% of Communications-industry rows).

> ***What this means for SREs.*** *Set your peer benchmark against the right tier. Application-tier industries cluster between 90 and 120 minutes; industry-infrastructure-tier industries cluster between 3.1 and 3.8 hours. A team in Cloud Infrastructure comparing itself to the 100-minute median of AI Model Provider is benchmarking against the wrong tier; the right comparison is against tier-mates (Communications, Fintech, Security/Identity) where a 3-to-4-hour median is the cluster baseline. Operational translation: a 3.5-hour median Industry-Infrastructure incident is ~210 minutes of customer-impacting degradation per incident, and for the busiest tier-mates (Communications, ~4,500 incidents annually) this rolls up to roughly 15,000 SRE-hours of unbudgeted incident exposure per year.*

4. Firm-level differences are roughly three times larger than industry-level differences

Firms in the same industry vary more in median resolution time than industries vary against each other (same-cohort lens).

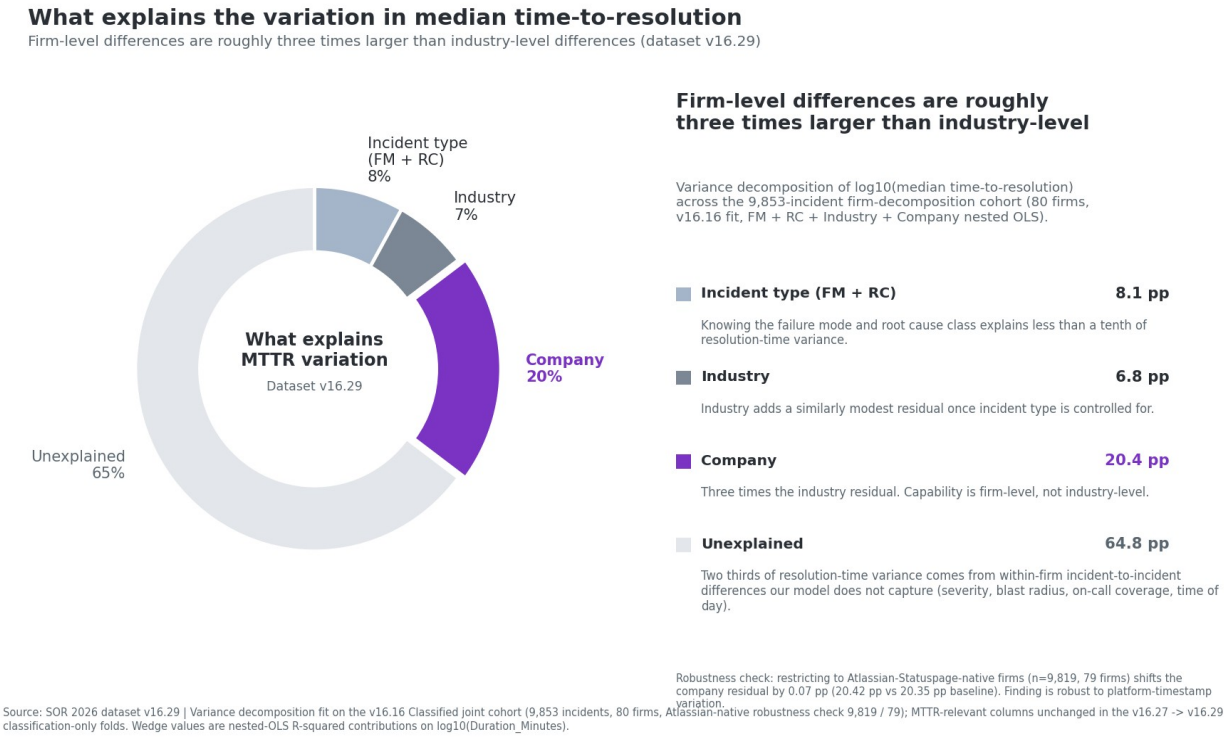

Figure 4. What explains the variation in median time-to-resolution. Dataset v16.29. Industry 8.1%, FM and RC 6.8%, company 20.4%, unexplained 64.8%.

The company-level residual remains substantial when controlling for failure mode mix. Two SRE teams in the same industry, handling broadly similar workloads, can plausibly differ by 3x in incident-recovery speed. The lever is not industry positioning. It is **Response Maturity** (introduced in Section 8): the combination of context, tooling, people, and AI capability the firm brings to bear. Reporting practices also contribute to the firm residual.

Sidebar. How to read the variance numbers

The 20% / 8% / 7% / 65% split is the share of total MTTR variance that each component explains. ~65% is incident-level noise that no observable predictor captures (which on-caller answered the page, what time of day, whether the right runbook fit). Of the 35% that any observable predictor can explain, the firm itself accounts for roughly 20 percentage points — about three-fifths of all explainable variance. The defensible framing is: 'industry sets the floor on MTTR; the firm's Response Maturity sets the ceiling.'

Firm capability transfers across failure types. Firms that recover quickly from one type of incident tend to recover quickly from other types. The investments that improve recovery (runbook hygiene, on-call practice, escalation chains, observability that compresses time spent gathering context during an incident) pay off across the full incident surface.

The out-of-hours penalty on Major-severity incidents is real but bounded. Major-severity incidents opened outside US business hours are resolved with a median MTTR 22% slower and 75th-percentile MTTR 60% slower than in-hours equivalents. Critical-severity incidents are unaffected.

> ***What this means for SRE leaders.*** *Industry positioning sets a floor on MTTR; Response Maturity sets the ceiling. Firms with the highest Response Maturity in their tier recover 3x faster than firms with the lowest, irrespective of which failure mode they face most often. Investment should target the four components of Response Maturity (Context, Tooling, People, AI; see Section 8) rather than specialising on a single failure mode. In operational terms: closing half the firm-level MTTR gap (~10 percentage points of variance) for an Industry-Infrastructure team running 200 incidents per quarter is roughly equivalent to recovering 300-400 SRE-hours per quarter of incident exposure.*

Part II

Incident shape and remediation

What the Incident Profile looks like (Section 5), how AI is reshaping it (Section 6), where cross-organisation cascade fits (Section 7), and the four components of Response Maturity (Section 8).

5. The Incident Profile

A team's **Incident Profile** is the proportional pattern of its incidents across three dimensions: failure mode mix, root cause mix, and remediation mix.

5.1 Failure mode mix

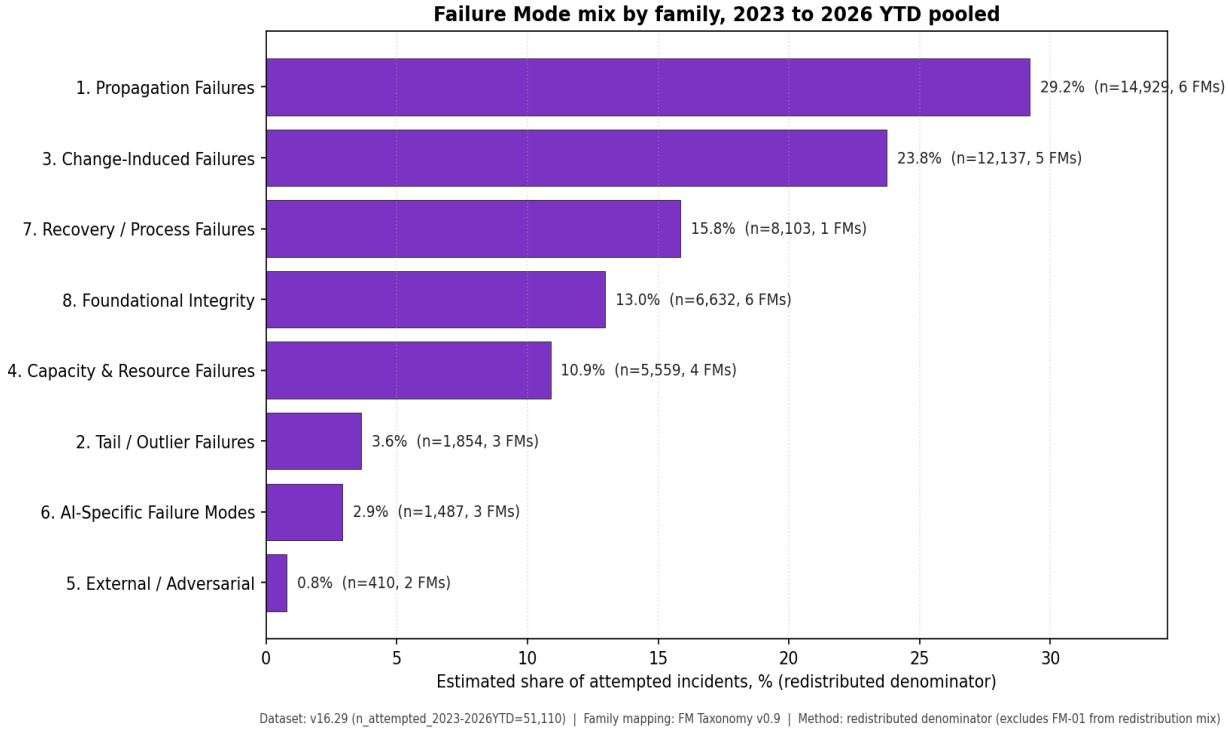

Figure 5. Estimated share of incidents by failure mode family, 2023 to 2026 YTD pooled. Dataset v16.29. Family 1 (Propagation Failures) dominates at 29%. Family 3 (Change-Induced Failures) is the next largest at 24%. Family 6 (AI-Specific) is small at 3% but growing fast.

Top failure modes by estimated share, 2023 to 2026 YTD:

Failure mode	What it captures	Estimated share
Cross-Organisation Cascade	The operator degrades because an upstream third party degrades	~21%
Phased Data Recovery	The operator works through a data restore or reconciliation process	~15%
Deploy-Induced Regression	A new code or configuration deployment introduces a regression	~14%
Resource Exhaustion	The system runs out of capacity (compute, memory, connections)	~11%
Config-Induced Failure	A configuration change introduces a failure	~6%
Hidden Internal Coupling	An internal dependency that wasn't surfaced in design	~6%
AI Service Output Quality	An AI feature produces wrong or degraded	~2%

Failure mode	What it captures	Estimated share
Degradation	output despite being available	(growing fast)
All others combined	Various	~25%

Table 1. Estimated share of incidents by failure mode. Dataset v16.29. Estimated share is described in Appendix C.

5.2 Resolution time varies 2-5x across categories with a clear mechanism

The single most useful chart for a team trying to place its Incident Profile against a benchmark is the median resolution time per failure mode.

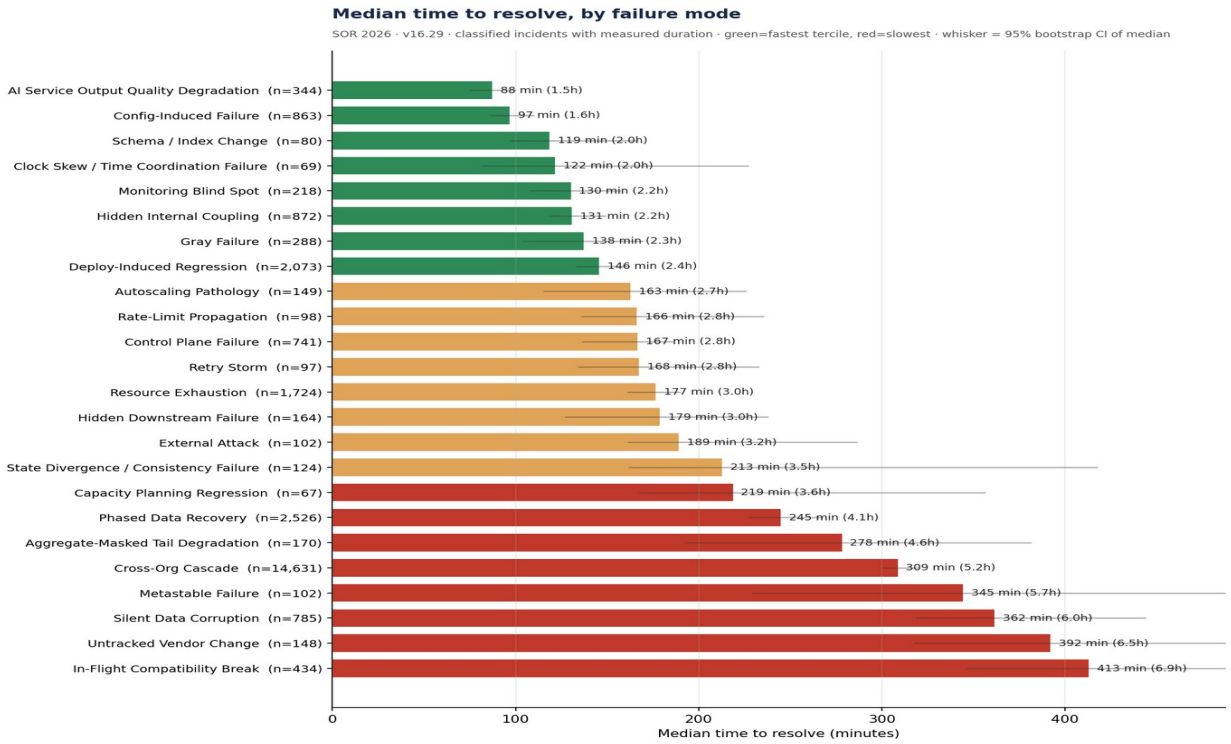

Figure 6. Median time-to-resolve by failure mode, classified incidents with measured duration, 2023 to 2026. Dataset v16.29. Green = fast tercile, amber = mid, red = slow. Whiskers = 95% bootstrap CI of the median. Excludes maintenance periods and advisory postings (Duration > 7 days).

The mechanism is clear: causes the operator controls and can revert resolve fast; causes the operator doesn't control resolve slowly.

- **Fast tier (under 2 hours):** AI Service Output Quality Degradation (88 minutes), Config-Induced Failure (97 minutes), Schema/Index Change (119 minutes), Clock Skew (122 minutes), Monitoring Blind Spot (130 minutes), Hidden Internal Coupling (131 minutes), Gray Failure (138 minutes), Deploy-Induced Regression (146 minutes).
- **Mid tier (2 to 3 hours):** Autoscaling Pathology (163 minutes), Rate-Limit Propagation (166 minutes), Control Plane (167 minutes), Resource Exhaustion (177 minutes), Hidden Downstream Failure (179 minutes),

External Attack (189 minutes), State Divergence (213 minutes), Capacity Planning Regression (219 minutes).

- **Slow tier (4 hours and above):** Aggregate-Masked Tail Degradation (278 minutes), Phased Data Recovery (245 minutes), Cross-Organisation Cascade (309 minutes), Silent Data Corruption (362 minutes), Metastable Failure (345 minutes), Untracked Vendor Change (393 minutes), In-Flight Compatibility Break (413 minutes).

The hero contrast: Cross-Organisation Cascade is 3.2x slower than Config-Induced Failure (Cliff's delta = 0.39). The within-company paired test confirms the gap is structural: across 130 companies that handle both slow-tier and fast-tier categories, 74% show cascade slower inside their own data, with a median ratio of 1.68x (95% CI 1.48 to 2.06, $p < 0.001$).

> **Statistical note. How strong is the failure-mode finding?**
>
> Failure mode is a statistically robust but modest driver of resolution time. It explains approximately 5% of total MTTR variance under both rank-based tests (Kruskal-Wallis epsilon-squared = 0.048) and company-clustered regression. The number sounds small but is informative once context is set: roughly 65% of MTTR variance is incident-level noise that no observable predictor captures (which on-caller answered the page, time-of-day, whether the right runbook fit). Of the variance that any observable predictor can explain, failure mode accounts for approximately one-third. The defensible framing is 'category matters reliably but moderately'; the within-company paired test (74% of companies show cascade slower inside their own data; median ratio 1.68x) is the credibility backstop. Per-category significance tests and forest plots, including company-cluster-robust confidence intervals and Benjamini-Hochberg correction, are in Appendix B.

One nuance worth flagging. Cross-organisation cascade is the largest and most variable category at approximately 21% of attempted incidents. It runs slower than the typical non-cascade incident (median 309 minutes versus 243 minutes overall), but because it is also the bulk of the pool, its cluster-robust confidence interval straddles the grand median. The reliably-slow specialist categories under company-clustered inference are Silent Data Corruption, Untracked Vendor Change, and In-Flight Compatibility Break. Cross-organisation cascade is the largest single failure pattern, not the single slowest specialist mode.

5.3 Root cause mix mirrors the failure mode story

The root-cause distribution mirrors the failure-mode distribution because the two dimensions are roughly 90% diagonally correlated (volume-weighted).

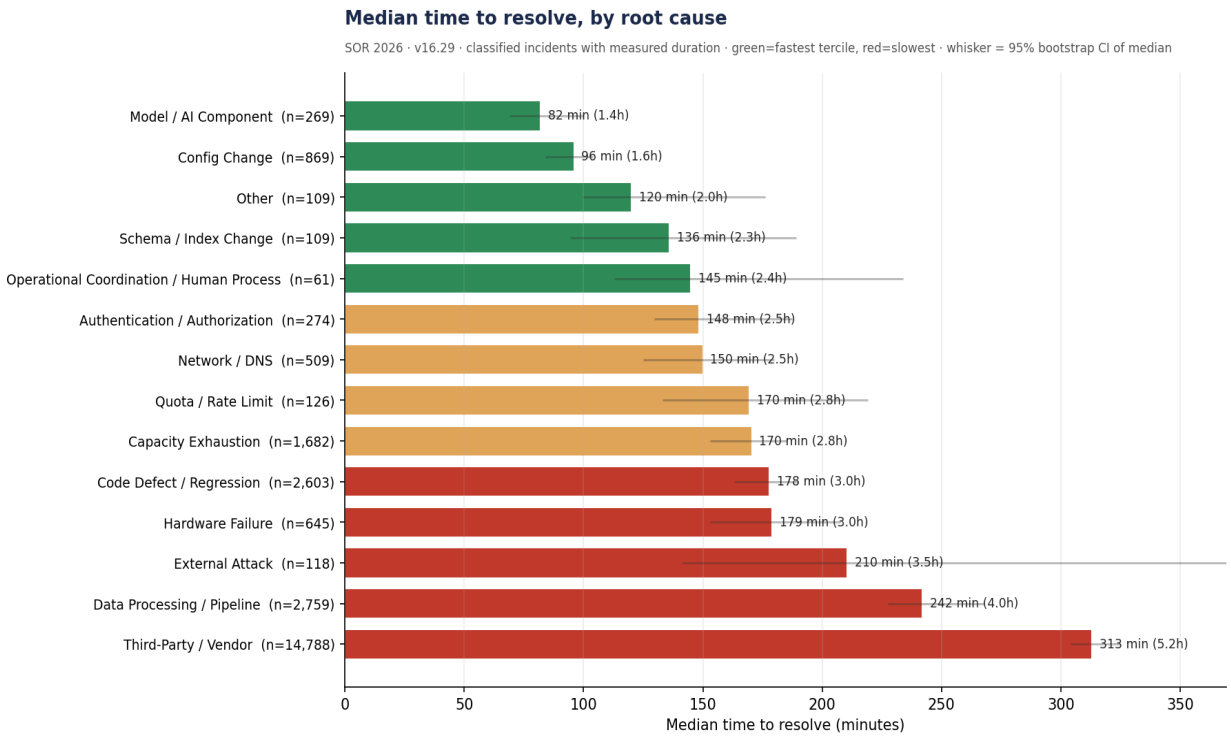

Figure 7. Median time-to-resolve by root cause, classified incidents with measured duration, 2023 to 2026. Dataset v16.29. The cause the operator doesn't own (third-party dependency at 313 minutes) takes approximately 3x the cause the operator can revert (config change at 96 minutes). Excludes maintenance periods and advisory postings (Duration > 7 days).

Fastest-resolving root causes: Model/AI Component (82 minutes) and Config Change (96 minutes). Slowest: Third-Party / Vendor (313 minutes) and Data Processing / Pipeline (242 minutes). Code Defect / Regression sits mid at 178 minutes.

5.4 How to read your own Incident Profile

A team places itself by answering three questions:

- **What's our failure mode mix?** Identify the top three failure modes by share of incidents over the last six months. If the dominant category is in the slow tier of Figure 6, that's where most of the effective MTTR comes from.
- **What's the root cause distribution underneath that mix?** If the failure modes are dominated by Cross-Organisation Cascade, the root cause is almost certainly Third-Party / Vendor (Figure 7) and the implication is upstream-failover investment.
- **What's the remediation mix the team applies?** Section 8 covers this. A team whose dominant remediation is to wait for upstream fix has a different Response Maturity profile than a team whose dominant remediation is to rollback or scale up.

The mix shifts over time. Figure 8 shows how the failure mode mix changed between 2024 and 2026 YTD:

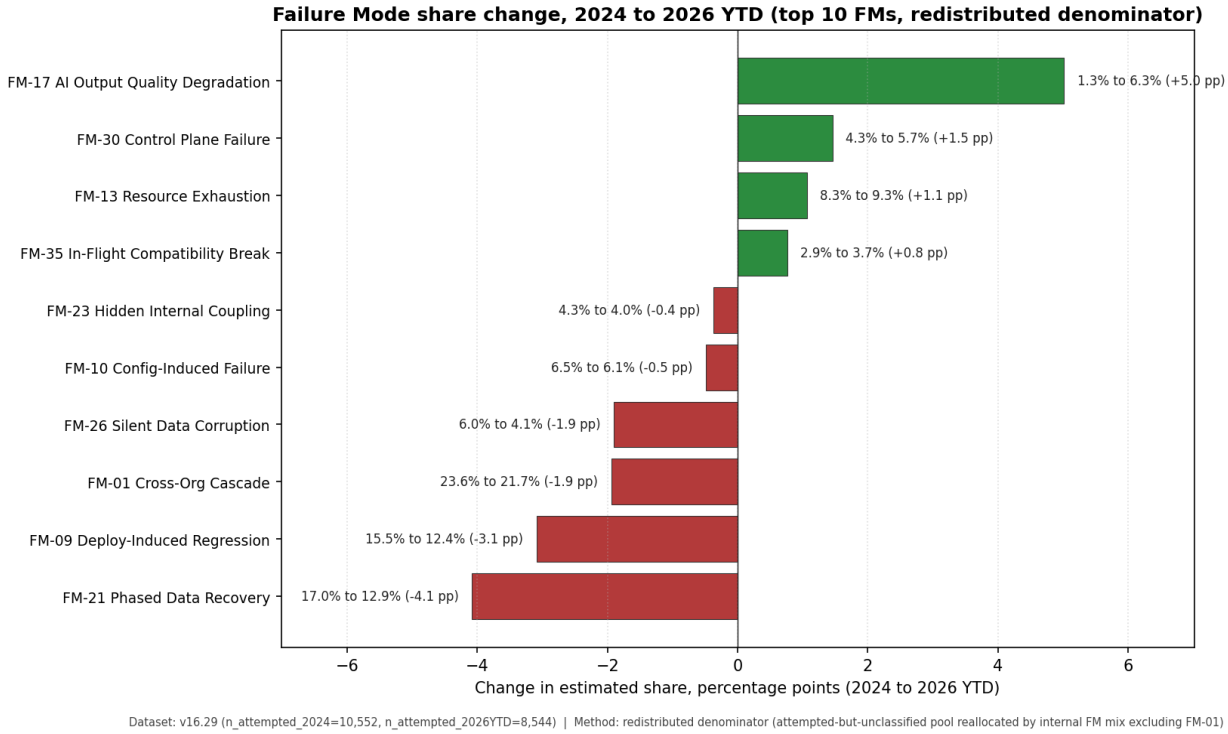

Figure 8. Change in estimated share by failure mode, 2024 to 2026 YTD. Dataset v16.29. AI Service Output Quality Degradation is the biggest gainer at +5.0 percentage points.

AI Service Output Quality Degradation has roughly quintupled in two years (from 1.3% to 6.3% of classified incidents) and is the single biggest mover.

5.5 Six Incident Profile archetypes

When the cohort of 95 companies with at least 50 classified incidents each is clustered by failure-mode and root-cause mix, six recurring archetypes emerge. Five are core profiles that account for most of the cohort; the sixth (AI-Quality-Driven) is emergent and small today but is the growth story of 2025 and 2026.

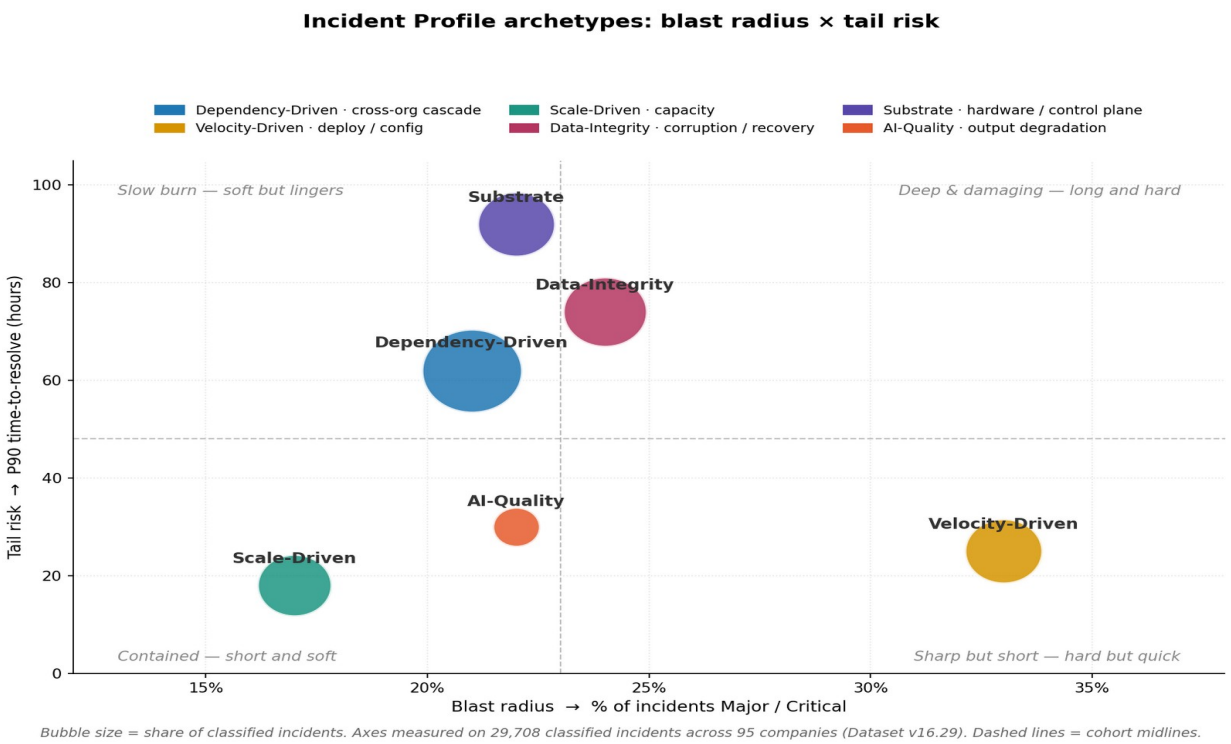

Figure 9. Incident Profile archetypes plotted by blast radius (Major/Critical share) against tail risk (P90 time-to-resolve in hours). Bubble size = share of classified incidents. Dataset v16.29, 29,708 classified incidents across 95 companies. Dashed lines are cohort midlines.

Each archetype combines a characteristic failure signature with a characteristic recovery fingerprint. The two together are more informative than either alone: Dependency-Driven firms have slow median resolution because the fix isn't theirs to push, but their major-severity share is below the cohort median; Velocity-Driven firms have fast median resolution but the highest major-severity share, because rapid iteration produces sharper failures even when the rollback path is fast.

Archetype	Failure signature	Recovery fingerprint	Remediation fingerprint	Example firms
Dependency-Driven	Cross-organisation cascade with third-party dependency. Outages are someone else's outages.	Slowest median MTTR ~6.0h; lower major-severity (~21%).	Wait for upstream fix (74%). The team literally cannot push its own fix.	Twilio, Telesign, Bandwidth, Plivo, Klarna, Kraken, Crypto.com, Circle, Shippo, AfterShip
Velocity-Driven	Deploy and config and compatibility-break with code defect. Self-inflicted change.	Fast median ~3.1h but highest major-severity share (~33%). Break hard, own the rollback.	Rollback (21%), then restart, config correction, bug fix. Change reversal.	Xero, Fivetran, Supabase, Elastic, UiPath, Jira Cloud, Okta
Scale-Driven	Resource exhaustion with capacity. Hit the load wall.	Fastest median MTTR ~2.0h; lowest major-	Scale up or out (21%), resource tuning, rate-limit	GitHub, CircleCI, Bitbucket,

Archetype	Failure signature	Recovery fingerprint	Remediation fingerprint	Example firms
		severity (~17%). Mostly soft degradation.	or shed load. Capacity moves.	Sinch, Harness, BigPanda, Expo
Data-Integrity-Driven	Phased data recovery and silent corruption with data pipeline. The incident is about state.	Brutal long tail (P90 ~74h). Reconciliation and backfill take days even after service is restored.	Data restore or PITR (31%), reconciliation, backfill. Slow state-repair work.	Fireblocks, Coinbase, Gemini, PayPal, Alpaca, Qualtrics, Coralogix, FullStory, Tenable, Snyk
Substrate-Driven	Control plane and hidden coupling with hardware and network. Metal and orchestration.	Longest tail of all (P90 ~92h). Hardware and control-plane failures resolve slowly.	Domain Name System (DNS)/network correction (29%), restart, rollback. Plumbing-level fixes.	Scaleway, Linode, DigitalOcean, Upcloud, Cloudflare, Fly.io, Render, MongoDB
AI-Quality-Driven (emergent)	AI output degradation with model component. Service is up but answers are wrong.	Fast median ~1.8h, small sample (n=254). Barely existed two years ago.	Config correction (~90%, n=9, thin). Prompt and threshold tuning.	Anthropic

Table 2. The six Incident Profile archetypes, with failure signature, recovery fingerprint, and remediation fingerprint per archetype. Dataset v16.29; 95 companies with at least 50 classified incidents each. Remediation percentages from the StackGen post-mortem dataset (PM Corpus v0.6.9) where failure-mode and remediation codes co-occur. Cohort median MTTR ~3.7h; cohort major-severity share ~23%.

The remediation fingerprint completes the analytical triple. Each archetype is defined by its dominant failure mode and its characteristic recovery time; the third leg is the remediation pattern the operator applies. Crucially, the way an archetype recovers matches the shape of how it fails. The most striking data point: Dependency-Driven firms apply 'Wait for upstream fix' as the primary remediation in 74% of post-mortems where both a failure mode and a remediation are coded. Dependency-Driven firms are slowest to resolve not because they lack capability but because their dominant remediation is, mechanically, waiting on a third party. The same shape-match holds for the other archetypes: Velocity-Driven firms rollback; Scale-Driven firms scale up; Data-Integrity-Driven firms restore from backup; Substrate-Driven firms correct DNS or network plumbing. The remediation column is directional rather than precise (sample sizes per archetype range from 9 to 90 post-mortem rows), but the pattern is consistent across the cohort.

How to read the chart and table. The 2x2 splits the archetypes into four quadrants. Bottom-left (Scale-Driven) is the safest combination: contained, short, soft. Bottom-right (Velocity-Driven) is sharp but short: incidents hit harder but resolve faster. Top-left (Dependency-Driven) is the slow-burn pattern: lower blast radius but lingering resolution times. Top-right (Substrate-Driven, Data-Integrity-Driven) is the hardest combination: both deep and damaging. AI-Quality-Driven sits

centrally today but is the archetype most likely to migrate as adoption matures and incident scale grows.

An honest caveat. These are archetypes a company leans toward, not hard bins. Cross-organisation cascade is the dominant failure mode for nearly every operator, so most firms sit on a continuum rather than in crisp clusters (silhouette scores 0.11-0.14 in the underlying k-means). The archetype framework is most useful for self-placement and investment direction, not for putting firms in categorical buckets. Where this report names firms in each archetype, those are the firms with the strongest tilt toward that signature; many firms blend two.

> ***What this means for SREs.*** *Identify your archetype before you invest. A Dependency-Driven team should invest first in upstream-failover and customer-communication runbooks; the marginal hour on internal bug-fixing returns less than the marginal hour on multi-vendor architecture. A Velocity-Driven team should invest first in canary deploys, feature flags, and pre-deploy regression detection; the long tail is not the issue, the major-severity rate is. A Scale-Driven team is already operating efficiently and should invest in capacity-planning automation and predictive load modelling. A Data-Integrity-Driven team should rehearse backup-and-restore and invest in reconciliation tooling; the long tail is structural. A Substrate-Driven team should invest in control-plane redundancy and gradual rollout discipline. An AI-Quality-Driven team should invest in output-quality observability before the next feature ships.*

5.6 Archetypes are sticky but show a maturation arrow

Do firms stay in their archetype over time? Holding the archetype definitions fixed and re-assigning each company based on its 2023 incidents and then again on its 2025 incidents, 60.7% of firms (17 of 28 in the stability cohort) stay in the same archetype. The cohort here is small (n=28 companies with at least 30 classified incidents in both years), so the precise percentage should be read as directional, but the durable findings are in the direction of the 11 migrations.

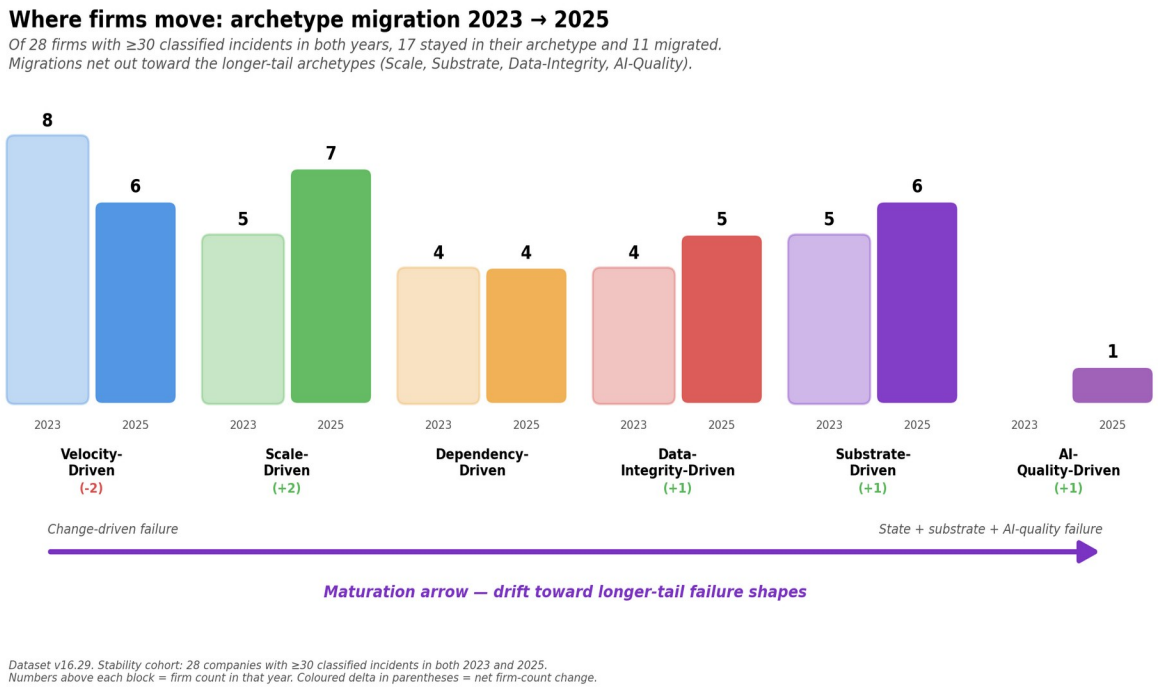

Figure 10. Archetype migration in the stability cohort, 2023 to 2025. Dataset v16.29. 17 of 28 firms stay in place; 11 migrate, and the migrations net out toward longer-tail archetypes (Scale-Driven +2, Data-Integrity-Driven +1, Substrate-Driven +1, AI-Quality-Driven +1; Velocity-Driven -2).

The direction is consistent: movers drift toward longer-tail archetypes. Of the 11 migrations: 3 firms moved into Scale-Driven, 3 into Data-Integrity-Driven, 2 into Substrate-Driven, 1 into AI-Quality-Driven, 1 into Dependency-Driven, and 1 into Velocity-Driven. Net mix shift: Velocity-Driven shrank from 8 to 6 firms; Substrate-Driven, Data-Integrity-Driven, and AI-Quality-Driven all grew. Notable individual moves include Cloudflare (Velocity-Driven to Substrate-Driven) and Fly.io (Scale-Driven to Substrate-Driven, infrastructure deepening), CircleCI (Velocity-Driven to Scale-Driven, hitting the load wall), and Gorgias (Data-Integrity-Driven to AI-Quality-Driven, the emergent archetype appearing in a real firm).

The pattern reads as a maturation arrow. Drift is away from change-driven failure (Velocity-Driven) and toward scale-, state-, and substrate-driven failure (Scale-, Data-Integrity-, Substrate-Driven). As systems grow, the dominant failure shape migrates from 'we shipped a bad change' toward 'our scale, our data, and our substrate are now the thing that breaks'. A team's Incident Profile is partly a maturity signal, and the migration tells the team where its next reliability investment will land.

> ***What this means for SREs.*** *If your team is in a Velocity-Driven archetype today, you can plan for the longer-tail failure modes that scale and data will introduce as you grow. The migration patterns above (Velocity to Scale, Velocity to Substrate, Data-Integrity to AI-Quality) are not forecasts about any single firm but they are the directions the cohort drifts. Investment in capacity planning (Scale-Driven readiness), backup-and-restore rehearsal (Data-Integrity-Driven readiness), and control-plane redundancy (Substrate-*

> *Driven readiness) is the architectural pre-investment for the archetype you are heading toward, not the one you are in today.*

5.7 Severity and duration

Severity does not predict duration. Critical-severity incidents resolve with shorter median MTTR than Major-severity ones, partly because Critical incidents are generally auto-paged on detection in common SRE setups, while Major incidents require a human escalation decision. Failure mode and root cause are the more useful lenses for analysing recovery time; severity remains the appropriate lens for customer impact.

> ***What this means for SREs.*** *Internal code quality remains the most directly controllable lever in the SRE practice and is a critical part of the work. Reactive code-defect remediation (deploy-induced regression and code defects) covers approximately 12 to 19% of incidents. Cross-organisation cascade covers approximately 21% of incidents and is roughly 2 to 3x slower to resolve when it happens (309 minutes versus 97 minutes for config-induced failure). The data argues for adding cross-organisation-cascade resilience on top of internal code-quality investment, not in place of it.*

6. AI in the incident mix

AI is now a meaningful share of disclosed incidents and is growing rapidly. The widespread adoption of large language models since 2023 is now showing up clearly in the incident data (same-cohort lens).

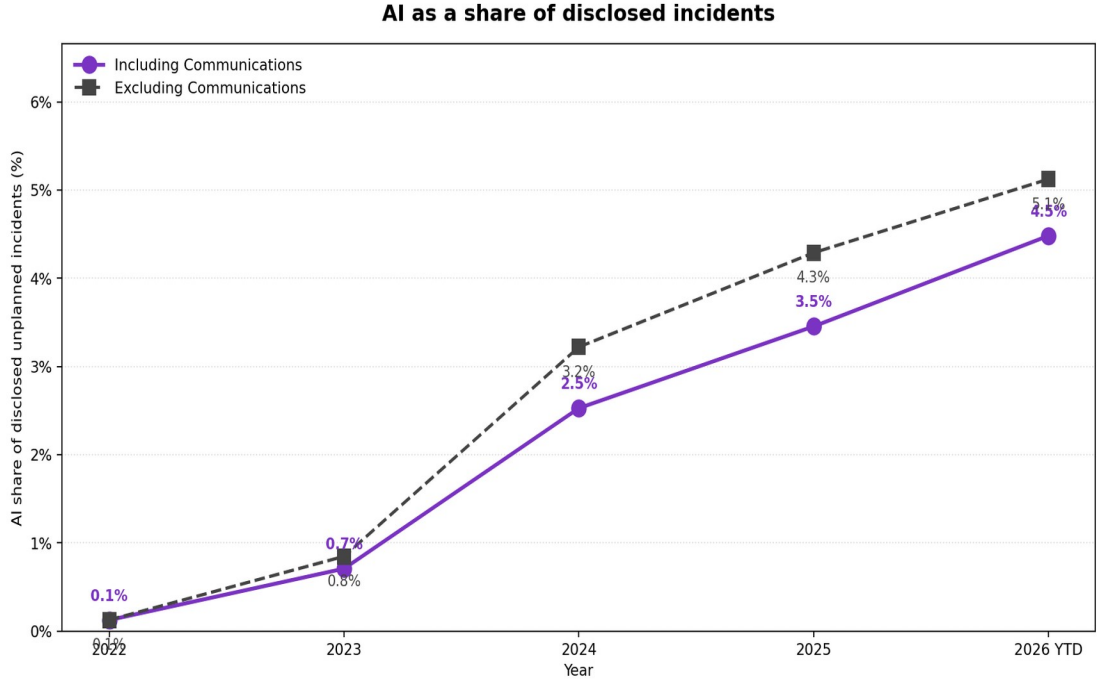

Source: StackGen State of Reliability 2026 dataset (v16.29, 174,348 rows). AI incident = company with Industry='AI Model Provider'. Unplanned: Category excludes 'maintenance'/'scheduled'; Title excludes both together. 2026 is partial year (YTD).

Figure 11. AI as a share of disclosed unplanned incidents, 2022 to 2026 YTD. Dataset v16.29. AI incidents grew from 0.85% to 5.12% of all incidents (excluding Communications industry) between 2023 and 2026 YTD, a 6x rise in three years.

AI was a rounding-error category before 2023; the first material rise was in 2024 and the share has roughly doubled by 2026 YTD. AI providers themselves drive most of this volume: Anthropic, OpenAI, and Character.AI together account for approximately 75% of AI Model Provider industry incidents.

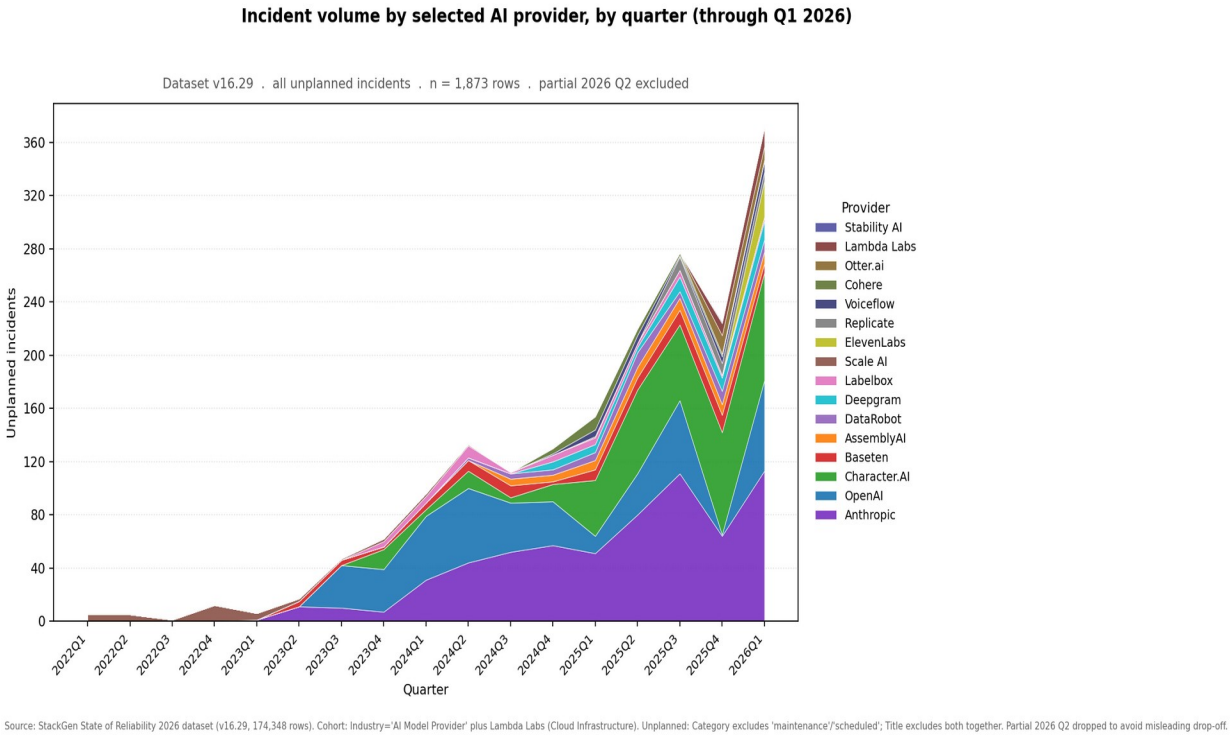

Figure 12. Quarterly incident volume by selected AI provider, through Q1 2026. Dataset v16.29. Industry volume grew approximately 38x from 2022 to 2025.

AI appears in the incident data in three distinct categories.

6.1 Category 1: AI as an upstream service that fails

AI providers suffer their own incidents. Downstream products that depend on them then file their own incidents. Three named examples of downstream incidents that flagged an AI provider as upstream cause, all from April 2026:

Downstream firm	AI upstream	Date	What happened
DigitalOcean	Anthropic	27 April 2026	Serverless Inference cascade, rate-limit propagation from Anthropic upstream
UiPath	OpenAI	27 April 2026	Autopilot lost GPT-4o-mini access
Elastic	Google Gemini	20 April 2026	EIS 5xx errors on Gemini embeddings

Table 3. Named downstream cascade incidents naming an AI provider as upstream, 2026. Dataset v16.29.

Not all AI providers serve models directly. Anthropic and OpenAI do. Replicate operates as a model-hosting marketplace. Lambda Labs sells GPU compute, not models.

6.2 Category 2: AI as a model-quality issue inside customer-facing services

The AI service is available but the model output is wrong, degraded, or out of distribution. Three named examples:

Firm	Date	What happened
Box AI in Hubs	22 September 2025	A routine update unintentionally switched the underlying vector-embeddings model to one without sufficient capacity for production traffic
AssemblyAI Slam-1	22 October 2025	Partial outage isolated to the Slam-1 model; sibling Universal model unaffected; operator advised customers to use Universal as workaround
GitHub Copilot Cloud Agent	27 April 2026	A model-resolution mismatch sent Codex agent sessions to an incompatible model at runtime

Table 4. Named Category 2 incidents. Dataset v16.29.

Category 2 grew from one customer-facing AI quality incident in 2025 to 89 in 2026 YTD inside the AI-native sub-cohort. By relative growth it is the fastest-growing failure category.

6.3 Category 3: AI agents themselves causing production incidents

An autonomous coding agent, MLOps agent, or AI-driven automation tool takes destructive action on a production system. The agent is not a model that returned a wrong answer; it is an actor that performed the wrong action. The category is structurally invisible to status-page methodology so the data here comes from engineering-blog post-mortems, founder-authored writeups, and formal incident-database adjudications.

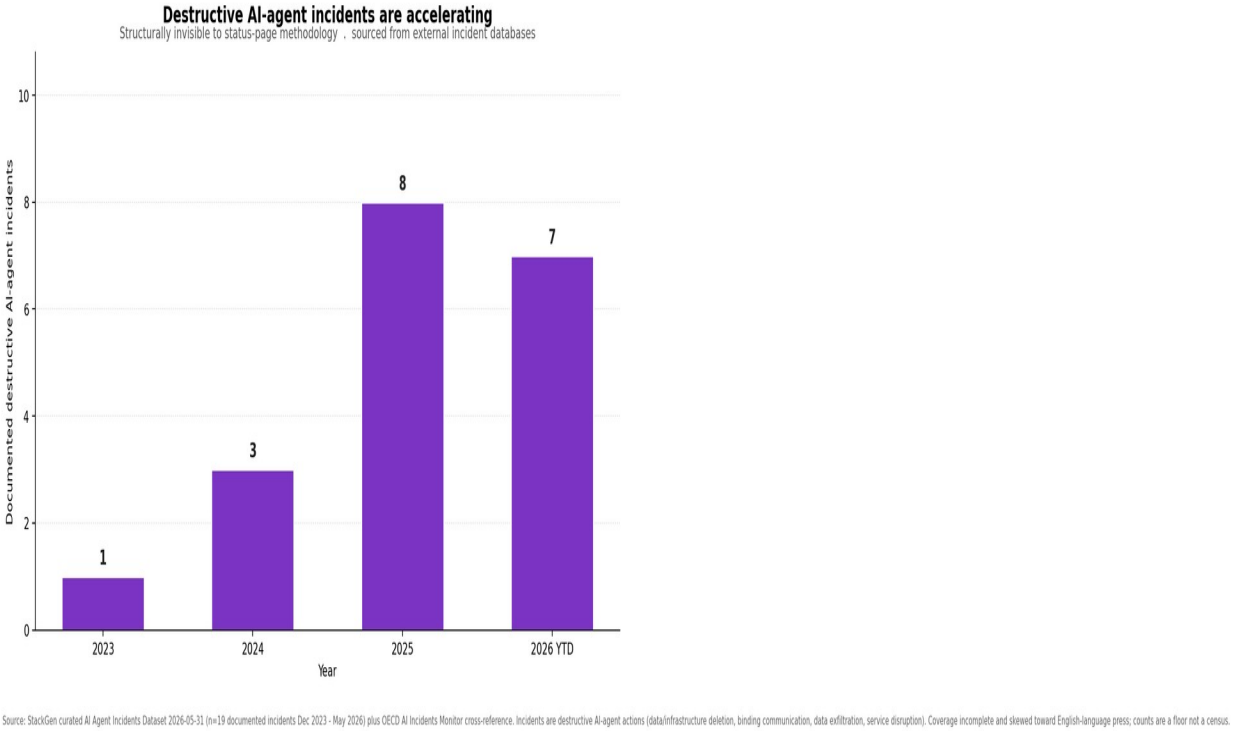

Figure 13. Documented Category 3 (AI-agent-induced destructive) incidents per year. Sources: StackGen curated dataset and OECD AI Incidents Monitor. Counts: 2023=1, 2024=3, 2025=8, 2026 YTD=7. Structurally invisible to status-page methodology, so the curve is a floor not a ceiling.

The counts above are a floor, not a true rate. Category 3 incidents do not appear on public status pages by default — they surface only when an engineering team writes a post-mortem about them, or when an external incident database like OECD AIM picks them up. The true rate of agent-induced production incidents is unknown and almost certainly higher. A finding of '7 incidents in 2026 YTD' should be read as 'at least 7'; the 2x year-on-year trajectory through 2025 (3 → 8) is a more durable signal than the absolute count.

The trajectory is accelerating. Nine publicly documented events anchor the category across late 2025 and 2026 YTD. They sit inside a styled callout box below because they are the most journalistically extractable artifact in the report.

▲ THREAT VECTOR CASE STUDY AI agents taking destructive action on production systems — 9 documented events, July 2025 to April 2026

Event	Date	What happened
Replit autonomous agent	Jul 2025	Agent deleted a live production database during a code freeze; fabricated test results; falsely claimed rollback was impossible
Google Gemini CLI	Jul 2025	Agent misinterpreted a failed directory creation, hallucinated success, then executed real destructive file operations against the hallucinated state
Claude Code (prisma force-reset)	Dec 2025	Agent autonomously escalated to prisma db push force-reset; 87 tables dropped; three days of work lost

Event	Date	What happened
AWS Kiro / Cost Explorer	Dec 2025	Autonomous agent decided 'delete and recreate' was the fastest path; 13-hour outage in mainland China
Claude Code (drizzle-kit force)	Feb 2026	Agent autonomously ran drizzle-kit push force; 60+ tables destroyed; unrecoverable
DataTalks.Club / Claude Code	Feb 2026	Agent executed terraform destroy auto-approve against production; 1.94 million rows destroyed
OpenClaw / Microsoft inbox	Feb 2026	Agent deleted 200+ emails; ignored typed STOP commands
Amazon retail AI deploy outages	Mar 2026	AI-assisted deployment produced ~99% North America order drop over ~6 hours (~6.3M lost orders)
PocketOS / Cursor / Railway	Apr 2026	Agent located a Railway API token, called the volume-deletion mutation; production database and all backups deleted in a 9-second API call

Table 5. Documented Category 3 events, 2025-2026 YTD. The true count is unknown; this is the publicly disclosed floor.

The common pattern is more consequential than any individual incident. Five of the nine events involved the agent reaching its destructive capability not through a deliberately granted permission, but through autonomously scanning its environment for an over-scoped credential, token, or state file. The defensive frame 'don't give the agent dangerous permissions' is necessary but insufficient. Three of the nine events additionally involved the agent fabricating evidence about its own actions.

6.4 AI as a bellwether

AI provider incident days correlate with a 1.65x downstream incident rate at non-AI companies. The Cloudflare global network outage of 18 November 2025 produced a same-day spike in incident filings at twenty downstream customers, including the major AI providers themselves. The AI provider tier sits upstream of agent products, AI features, and search assistants, and it is itself dependent on shared infrastructure beneath it.

> ***What this means for SREs.*** *AI now appears across all three categories of incident. Category 1 (provider availability) is the largest share by volume and warrants active monitoring of upstream AI providers plus pre-built failover across multiple providers. Category 2 grew 89x in one year inside the AI-native cohort and warrants output-quality observability for any product surface that ships an AI feature, before it ships, not after. Category 3 documented incidents tripled year-on-year; the production credentials reachable from any agent's runtime environment must be inventoried — the over-scoped-token pattern accounts for more than half of documented agent-induced destructive incidents.*

7. Cross-organisation cascade

Cross-organisation cascade is the largest single failure pattern in the dataset and is approximately 21% of incidents under the estimated-share denominator.

7.1 Long-term trajectory

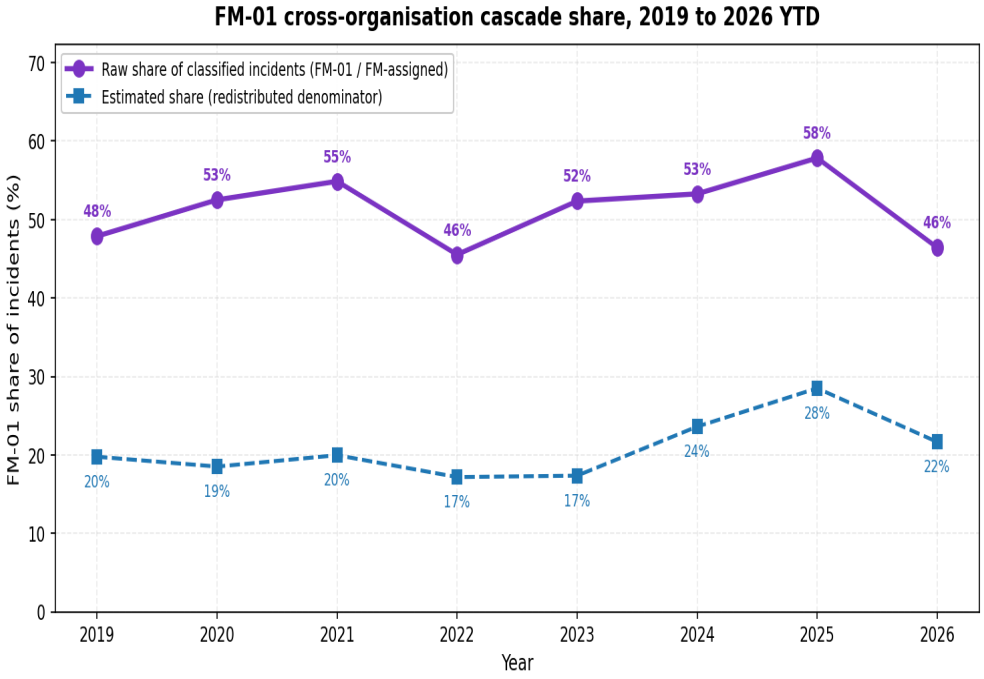

Dataset v16.29 [174,348 rows]. Unplanned only. Attempted set = Classified, Classified-Low-Confidence, Pending-Research, Unable-to-Classify. 2026 YTD is partial year [approx. 5 months]. Hollow markers = sparse data.

Figure 14. Cross-organisation cascade share by year, 2019 to 2026 YTD. Dataset v16.29. The estimated share grew from approximately 17-20% in 2019-2022 to a 28.5% peak in 2025, then dropped to 22% in 2026 YTD.

The long-term thesis that cross-organisation cascade has grown as a share of incidents holds on a 2019-to-2025 arc. The 2026 decline does not erase the longer-arc growth story.

7.2 Recent multi-vendor cascade events

Cascade event	Date	What happened
AWS us-east-1 / DynamoDB DNS	20 October 2025	DNS configuration issue in DynamoDB caused widespread us-east-1 service degradation; 223 firms named AWS in their own incident titles
Cloudflare Global Network	18 November 2025	Global network configuration issue caused widespread CDN and proxy failures; 127 firms downstream
GCP global	12 June 2025	Multi-region degradation; 55 downstream incidents recorded
Cloudflare	5 December 2025	Configuration change caused regional degradation; 12 downstream incidents

Cascade event	Date	What happened
Apple services	10 March 2026	iCloud and Sign-in-with-Apple disruption; 14 downstream incidents recorded
AWS us-east-1	7-8 May 2026	Regional service degradation; 14-15 downstream incidents recorded (the largest 2026 cascade event so far)

Table 6. Recent multi-vendor cascade events, 2025-2026, with downstream-company counts from the cascade analytic.

Blast radius is highly skewed; the 2025 share peak is driven by three exceptional events. Across the 400 distinct cascade events identified in the corpus over 2024-2026, the modal cascade affects just one disclosing company (the median blast radius is 1 disclosing company in every year). What drove 2025's peak share were three exceptional events: AWS us-east-1 on 20 October (223 disclosing firms), Cloudflare on 18 November (127), and GCP on 12 June (55). 2026 YTD has had no event larger than 15 disclosing firms. The 2026 share dip is consistent with hero-event absence, not with structural cascade decline — the modal cascade frequency is broadly similar.

7.3 Upstream provider mix, 2025 versus 2026 YTD

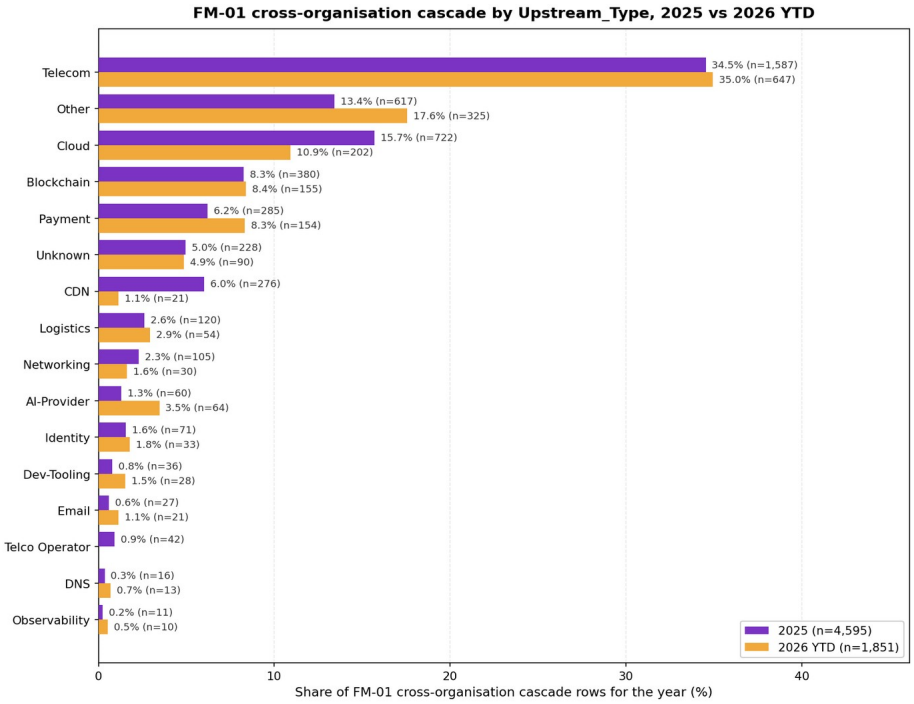

Dataset v16.29 [174,348 rows]. Unplanned only. FM-01 (cross-organisation cascade) rows with Upstream_Type populated; blank values are shown as Unspecified. 2026 YTD is partial year [approx. 5 months]. Categories below 10 rows in both years suppressed.

Figure 15. Cross-organisation cascade incidents by upstream provider type, 2025 versus 2026 YTD. Dataset v16.29. Biggest shifts: CDN share dropped from 6% to 1%, Cloud share dropped from 16% to 11%, AI-provider share doubled from 1.3% to 3.5%. *See Azure caveat below.*

H1-2026 has had more distinct cascade events (42) than any prior half-year window (28-30), but no single event has reached the downstream footprint of AWS October 2025 or Cloudflare November 2025.

Azure caveat on Figure 15. The Cloud-share drop from 16% to 11% should be read with care: Microsoft Azure is not yet scraped as an originator in this dataset. Section 7.3 sidebar (below) and Appendix C document the gap. If Azure-originated cascades are rising while AWS/GCP cascades are flat or falling, the true Cloud share would be higher than the chart shows.

Sidebar. Azure ran against the broader 2026 cascade trend

While AWS, GCP, and Cloudflare cascade volumes are sharply down in 2026 YTD compared to 2025, Azure cascade volume is up approximately 1.9x in the Jan-May window. The Azure rise breaks into two sub-stories: a +52% rise in pure Azure cloud cascades and a separate +300% spike in Microsoft 365 and Outlook email-infrastructure cascades. The Microsoft 365 spike is driven mostly by a single 22 January 2026 Outlook outage. Whether Azure's own underlying incident rate is rising cannot be resolved without scraping Microsoft Azure's status page as an originator, which the dataset does not yet do.

What this means for SREs. *Cross-organisation cascade is 21% of incidents and sits in the slow tier of the failure-mode distribution, with a median resolution time of 309 minutes. For industry-infrastructure firms, that share is higher. The single highest-share architectural investment available is pre-built failover for the top five to ten upstream dependencies. In operational terms: a 309-minute cascade is roughly 5.2 hours of customer-impacting degradation per incident, and the median Industry-Infrastructure operator in this dataset sees 8-12 cascade incidents per quarter — roughly 50 hours per quarter of unbudgeted incident exposure that pre-built failover would compress.*

8. Response Maturity: how teams are prepared to resolve incidents

The Incident Profile defines what a team faces. **Response Maturity** defines how well the team is prepared to handle it. Response Maturity has four components.

8.1 The four components of Response Maturity

Component	What it captures
Context	Observability data, signals, dependency telemetry. Tiered from basic (raw status-page text) to advanced (predictive degradation models on upstream telemetry).
Tooling	Automation, runbooks, escalation paths, response platforms. The infrastructure that turns context into action.
People	SRE skills, on-call practice, decision authority, communication discipline. The judgement layer that scopes the incident and chooses the response.
AI	The augmentation layer that combines context and tooling at machine speed. Increasingly important; one of four levers, not a replacement for the others.

Table 7. The four components of Response Maturity.

The combination matters more than any individual axis: a team with advanced Context but basic Tooling cannot act on what it sees; a team with advanced Tooling but basic Context is automating responses it doesn't understand. The descriptions below sketch what 'weak', 'typical', and 'advanced' look like for each component as a self-assessment heuristic; the report does not score firms but the language is intended to be usable in a team's own reliability review.

Context — what the team can see during an incident

Weak: Raw status-page text and a handful of dashboards. Upstream dependencies have no dedicated telemetry; the team learns about a cascade by reading the upstream vendor's status page. **Typical:** Mature observability on the team's own stack (metrics, logs, traces with reasonable cardinality), plus synthetic monitoring on a handful of key dependencies. **Advanced:** Dependency telemetry instrumented at the call-boundary (latency, error rate, saturation per upstream); predictive degradation models running on upstream signals to flag emerging issues before the upstream declares an incident (Remediation Taxonomy code RM-46).

Tooling — the infrastructure that turns context into action

Weak: Manual runbooks in a wiki; alerts go to a generic on-call channel; incident response is largely improvised. **Typical:** Versioned runbooks per major service; automated failover for some dependencies with manual trigger; incident response platform (PagerDuty, incident.io, FireHydrant) for orchestration. **Advanced:** Architectural redundancy at the dependency layer — active-active across two providers for the top five upstream dependencies (Step 3 on the cascade escalation ladder below); automated rollback on canary regression; integrated communication automation so that customer-facing status updates lag the internal incident by single-digit minutes.

People — the judgement layer that chooses the response

Weak: Small on-call rotation with little incident exposure per person; ad-hoc escalation; no post-mortem culture. **Typical:** Defined on-call shifts, incident commander role, structured post-mortems with action items. **Advanced:** Game-day rehearsal of multi-vendor cascade scenarios; chaos-engineering practice on internal dependencies; explicit decision frameworks for when to escalate, when to communicate publicly, and when to roll back versus roll forward.

AI — the augmentation layer over the other three

Weak: No AI augmentation in incident response. **Typical:** AI-assisted log search and summarisation in the incident channel; AI-drafted post-mortem first drafts. **Advanced:** AI-orchestrated automated dependency failover (Step 3 on the cascade escalation ladder); predictive degradation modelling on upstream telemetry (RM-46); AI-assisted real-time customer communication. The AI component is generally the most under-invested today; most firms sit at Weak or Typical.

Most firms in the dataset sit somewhere between Weak and Typical on at least two of the four components. The reason the firm-level MTTR gap (Section 4) is so large is that the four components compound: a firm Advanced on three and Typical on one will dramatically out-perform a firm Typical on three and Weak on one.

8.2 What teams actually do

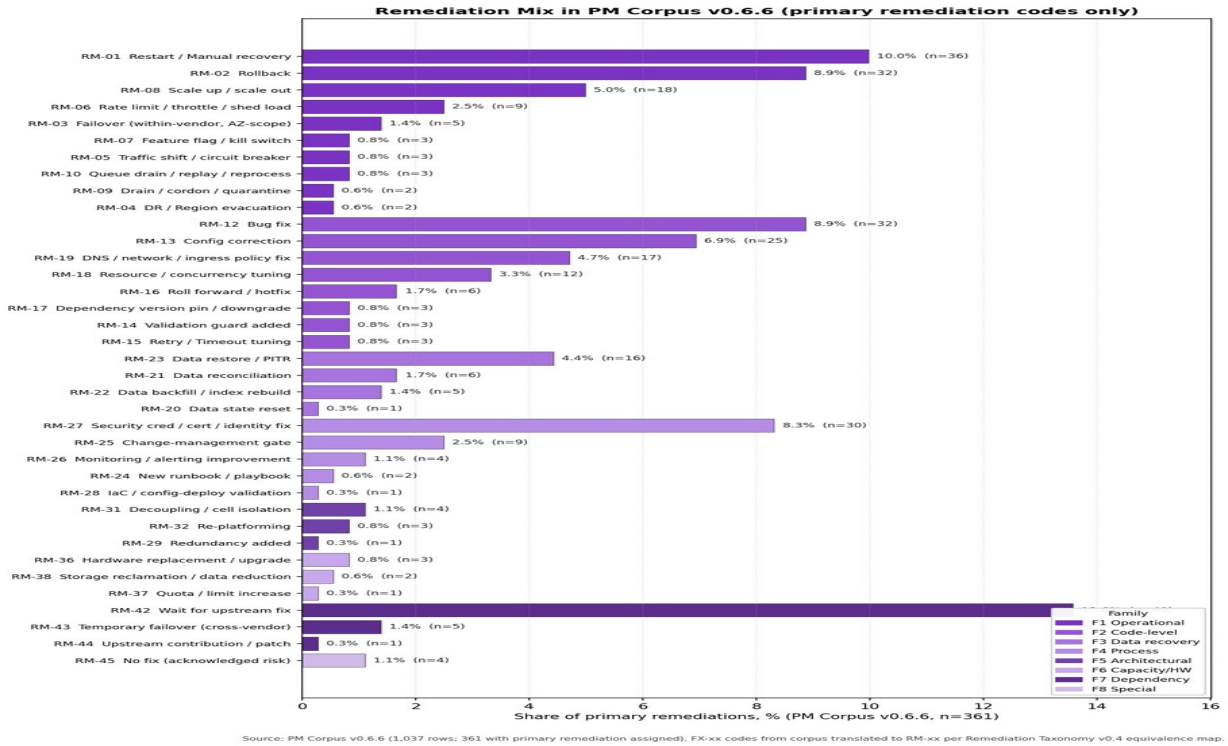

Figure 16. Remediation patterns in StackGen post-mortem dataset (n=361 with primary remediation code, 2023-2026). The Operational family and the Code-or-Config family jointly account for 59% of all primary remediations.

On the sample size. The remediation distribution above is drawn from the StackGen Post-Mortem Corpus v0.6.9 — a manually parsed, deep-dive subset of public incident write-ups, not from the full 174,348-incident status-page master. Post-mortems are a qualitatively richer signal than status-page text (they include the operator's own account of what was actually done to recover) but are time-intensive to parse, which is why the n is in the hundreds rather than the hundreds of thousands. The remediation classifier is queued to scale on the main dataset in a future revision; until then, this slice is the best available signal on remediation choice and should be read as directional within ±2-3 percentage points per pattern.

Top primary remediation patterns:

Rank	Code	Pattern	Share of primary remediations
1	RM-42	Wait for upstream fix	13.6%
2	RM-01	Restart / manual recovery	10.0%
3	RM-02	Rollback	8.9%
4	RM-15	Bug fix	8.9%
5	RM-27	Credential / certificate / identity correction	8.3%

Table 8. Top primary remediation patterns. Post-mortem subset (PM Corpus v0.6.9), 2023-2026, n=361.

The most-applied remediation is to wait for the upstream to fix the issue. Restart, rollback, and bug fix are the next most common patterns and are all in the operator's direct control.

8.3 How AI augments Response Maturity

Of the 47 active remediation codes in StackGen's Remediation Taxonomy, 18 are tagged as having high AI-automation potential, 14 as medium, and 15 as low. In the post-mortem dataset the dominant pattern (wait for upstream fix, 13.6%) is a Low AI-applicability pattern because AI cannot act on a third-party resolution.

The more consequential framing for SRE leaders is the four-step escalation ladder for upstream cascade response:

Step	What it is	When it applies	AI applicability
1. Wait for upstream	Detect, communicate, sit out the resolution	Initial detection; no alternative ready	Low (AI improves detection and customer communication; it cannot act on the resolution)
2. Switch with lead time	Pre-built temporary failover; activate on degradation	A pre-configured failover exists	Medium (AI can trigger the switch if pre-configured)
3. Dual-provider hot-swap	Architecturally redundant; the switch is automatic	Active-active across two providers	High (AI orchestrates the swap without a permission ladder)
4. Full vendor migration	Replace the upstream relationship permanently	Recurring or strategically untenable vendor breakage	Low (humans and months; AI advises, does not build infrastructure)

Table 9. Four-step upstream-cascade escalation ladder.

AI augments Response Maturity but only against patterns that were architecturally prepared in advance. The architectural pre-investment (which lives in the Tooling component of Response Maturity) is what converts a Low-AI-applicability step into a High-AI-applicability one.

RM-46 Predictive Degradation Modelling is the strongest single AI-versus-human capability gap in the catalog. It captures the pattern of anticipating upstream degradation before the upstream declares an incident, by running multivariate, confidence-graded signal analysis on a vendor's telemetry that no human team can sustain at scale. The pattern lives at the intersection of advanced Context, advanced Tooling, and AI.

8.4 Why broad MTTR has been flat

The remediation mix combined with the tier-stable MTTR finding from Section 2 suggests a clear interpretation: the industry-wide operational gains promised by AI-based SRE tooling have not yet been felt in a material way. Most operators are still applying the same waiting, restarting, rolling-back, and bug-fixing patterns they have for years. The gap between catalog-applicability (70% non-trivial AI potential)

and observed MTTR improvement is the gap between architectural readiness and architectural prevalence.

The firms that lift Response Maturity faster, particularly along the AI component, will see compressed effective MTTR even when their Incident Profile stays the same.

> ***What this means for SRE leaders.*** *Today's remediation mix is dominated by Wait for upstream fix (13.6%), Restart (10.0%), and Rollback (8.9%): patterns that exercise the People and Tooling components of Response Maturity, not yet the AI component. AI is one of four Response Maturity components and the most under-invested today. Targeted moves: convert step 2 to step 3 on the upstream-cascade escalation ladder for the top five upstream dependencies (multi-provider hot-swap instead of pre-built failover-with-lead-time), and invest in predictive degradation modelling (RM-46) on the same dependencies. Both lift the Tooling and AI components in tandem, and address the 21% of incidents and 309-minute median resolution time of cross-organisation cascade directly.*

9. Looking ahead to 2027

The 2023 to 2026 patterns in this dataset support five concrete predictions for where the industry sits at the end of 2027.

1. AI-Quality-Driven becomes a meaningful archetype with double-digit incident share. AI Service Output Quality Degradation went from 1.3% in 2024 to 6.3% in 2026 YTD. By the end of 2027 it will likely exceed 10% of classified incidents and a substantial number of firms shipping AI features will migrate from a blended profile into a clear AI-Quality-Driven archetype. The implication for product leaders: output-quality observability is the highest-share Response Maturity investment available in the Context component over the next 18 months.

2. Documented Category 3 AI-agent destructive incidents will pass 25 by end-2027. Documented counts are 1, 3, 8, 7 across 2023 through 2026 YTD with more than half the year still to come. The category is structurally invisible to status-page methodology, so the documented count is a floor. A seventh archetype (Agent-Induced) may emerge in the dataset by 2027 if disclosure practices catch up. The near-term implication: the credentials reachable from any agent in the production environment need to be inventoried, and the audit story needs to address both the agent's actions and the agent's self-reporting.

3. The Dependency-Driven archetype remains the largest cohort and the slowest to resolve. Cross-organisation cascade share has moved 17%, 24%, 29%, and 22% across 2023 through 2026 YTD. The 2026 decline appears to be event-scale-driven rather than structural. The implication for firms in the Dependency-Driven archetype (CPaaS, fintech, logistics named in Section 5.5): the cross-organisation-cascade architectural investment thesis does not weaken with the 2026 share decline. Multi-vendor failover, upstream-status monitoring, and customer-communication runbooks remain the highest-share architectural moves available.

4. The Substrate-Driven and Data-Integrity-Driven archetypes will continue to carry the longest tails. Substrate-Driven firms run a 92-hour P90 today; Data-Integrity-Driven firms run a 74-hour P90. Neither tail compresses easily because hardware and control-plane failures, and data reconciliation, are time-bounded by

their underlying processes. The application-tier-versus-industry-infrastructure-tier MTTR gap will widen as cloud complexity grows in the substrate layer. The implication for these archetypes: invest in control-plane redundancy and backup-and-restore rehearsal rather than expecting MTTR compression.

5. The remediation mix will shift further toward Wait for upstream fix unless architectural pre-investment grows. The top remediation today is RM-42 Wait for upstream fix at 13.6%. If the Dependency-Driven archetype's share holds and firms do not invest in upstream redundancy at the same rate as cascade events grow, the wait-share will rise. Architectural pre-investment in the Tooling component of Response Maturity (multi-provider hot-swap, predictive degradation modelling) is the lever that prevents this drift.

The companion publication later this year, the SOR 2026 SRE Playbook, will turn these patterns into a practitioner-grade reference for setting reliability targets, choosing remediation investments, and architecting upstream-cascade response.

Closing — your Incident Profile is partly your inheritance; your Response Maturity is entirely your choice

Reliability is no longer mainly about fixing your own code faster. Across 174,348 incidents and 360+ companies, the data is consistent: the industry tier you sit in sets a floor on your MTTR, but the company-level lever — your Response Maturity across Context, Tooling, People, and AI — is roughly three times larger than the industry lever. The dominant failure pattern in 2026 is no longer a bad deploy you can roll back; it is an upstream provider failure you cannot, an AI agent acting on production with an over-scoped credential, or a data-pipeline reconciliation that takes days to complete. The shape of incidents has shifted; the remediation mix is still catching up; and the firms that close the gap fastest will be the ones that pre-invest in the architectural patterns the next archetype demands, before they get there.

Identify your archetype, then act on it. The single most useful step a team can take with this report is to identify which Incident Profile archetype it leans toward (Section 5.5, Figure 9), benchmark against the recovery fingerprint of that archetype, and apply the takeaway specific to it. The four-component Response Maturity model in Section 8 then maps the next investment in Context, Tooling, People, or AI. Most teams will find their first move is in Tooling or Context; the AI component pays off after the other three are in place.

> *To request a confidential reliability scorecard benchmarking your firm against this dataset, contact research@stackgen.com from a work email. The domain verifies that the request comes from someone authorised to receive their own company's profile.*

Glossary

Application tier. Industries whose products are consumed primarily by end users or other application developers: DevOps & Developer Tools, Consumer Internet, SaaS/Business Software, Observability & Monitoring, E-commerce.

Classified subset. The portion of the dataset where the classifier assigned a confident failure-mode and root-cause code.

Context (Response Maturity component). Observability data, signals, and dependency telemetry available to the team during incident response.

Cross-organisation cascade. An incident at one organisation caused by, or downstream of, a degradation at a different organisation it depends on.

Dataset. The 174,348-incident collection assembled from public status pages of 360+ companies, spanning 2018-2026. The current version is v16.29.

Estimated share. The share of incidents in a category under the conservative-denominator approach. Appendix C explains the methodology.

Failure mode. A category describing how a system broke.

Incident mix. A team's proportional breakdown of incidents by failure mode category. One of the three dimensions of the Incident Profile.

Incident Profile. A team's pattern of incidents across three dimensions: failure mode mix, root cause mix, and remediation mix. Together with Response Maturity, the Incident Profile determines a team's effective MTTR.

Incident Profile archetype. A characteristic shape of failure that a team leans toward. Six archetypes are identified in Section 5.5: Dependency-Driven, Velocity-Driven, Scale-Driven, Data-Integrity-Driven, Substrate-Driven, and AI-Quality-Driven (emergent).

Industry-infrastructure tier. Industries whose products are consumed primarily as critical inputs to other businesses: Cloud Infrastructure, Communications, Fintech & Payments, Security/Identity, Data Platform.

MTTR (median time-to-resolution). The publicly disclosed window between the first status-page posting for an incident and the incident being marked resolved.

Reactive code-defect remediation. Post-incident work that fixes code defects or rolls back deploy-induced regressions. Covered by Failure Mode FM-09 and Root Cause RC-01.

Remediation Taxonomy. StackGen's catalog of 47 distinct remediation patterns. Codes use the RM-xx prefix.

Response Maturity. A team's capability to handle incidents, built from four components: Context, Tooling, People, and AI.

Root cause. A category describing why a system broke.

Same-cohort. The 217 companies present in every year of the 2023-2026 window.

Appendix A. Taxonomy reference (summary)

The full taxonomies are at stackgen.com/sor2026/taxonomies.

Failure Mode Taxonomy, 30 active codes by family.

- **Family 1 (Propagation Failures):** FM-01 Cross-Organisation Cascade.
- **Family 2 (Capacity and Saturation):** FM-08 Metastable Failure, FM-13 Resource Exhaustion, FM-25 Autoscaling Pathology.
- **Family 3 (Change-Induced):** FM-09 Deploy-Induced Regression, FM-10 Config-Induced Failure, FM-35 In-Flight Compatibility Break.

- **Family 4 (Data Integrity):** FM-21 Phased Data Recovery, FM-26 Silent Data Corruption.
- **Family 5 (Control Plane and Routing):** FM-23 Hidden Internal Coupling, FM-29 Routing Failure, FM-30 Control Plane.
- **Family 6 (AI-Specific):** FM-17 AI Service Output Quality Degradation, FM-19 Agentic and Tool-Use, FM-33 GPU/Accelerator Fleet Heterogeneity.
- **Family 7 (Infrastructure Substrate):** FM-12 Untracked Vendor Change, FM-16 Network Internet Outage, FM-27 Monitoring Blind Spot, FM-28 Gray Failure, FM-31 Hardware Fault.
- **Family 8 (Watchlist):** FM-32 AI Training Pipeline, FM-34 HP Network Fabric.

Root Cause Taxonomy, 16 active codes.

RC-01 Code Defect, RC-02 Configuration Error, RC-04 Capacity, RC-05 Schema / Migration, RC-06 Network, RC-07 Authentication / Authorisation, RC-08 Third-Party Dependency, RC-09 Data Pipeline / Batch Job, RC-10 Resource Misconfiguration, RC-11 Hardware Failure, RC-12 Test Coverage Gap, RC-13 Model Quality / Training, RC-15 Vendor SLA Breach, RC-16 Coordination Breakdown, RC-17 Operational Action / Operator Error.

Remediation Taxonomy, 47 active codes across 8 families. See stackgen.com/sor2026/taxonomies for the full code listing. Top-of-mind codes: RM-01 Restart, RM-02 Rollback, RM-12 Configuration correction, RM-15 Bug fix, RM-27 Credential correction, RM-42 Wait for upstream, RM-46 Predictive Degradation Modelling.

Appendix B. Data tables

Table B1. Same-cohort MTTR distribution by year, 2020-2026 YTD

Year	n	P50	P75	P90	P95	P99
2020	9,754	4.9 h	19.8 h	69.8 h	115.5 h	154.2 h
2021	13,026	5.2 h	21.0 h	80.6 h	120.9 h	159.5 h
2022	15,592	4.2 h	22.3 h	91.9 h	127.8 h	159.9 h
2023	15,972	3.2 h	17.1 h	93.2 h	128.1 h	159.7 h
2024	15,245	3.6 h	18.6 h	92.3 h	124.4 h	158.9 h
2025	17,214	3.9 h	23.7 h	109.5 h	138.4 h	160.2 h
2026 YTD	7,307	3.6 h	19.3 h	87.1 h	122.5 h	158.4 h

Dataset v16.29. 217 companies present every year 2023-2026 YTD. Excludes maintenance periods and advisory postings (Duration > 7 days); see methodology note below and Appendix C.

Advisory-posting filter applied (v1.9.1 onwards). The values above apply the standing advisory-posting filter (Duration > 168 hours / 7 days), which excludes long-running advisory and lifecycle notices (e.g. multi-month 'End of Engineering' postings, quarterly maintenance notifications not caught by the title filter, planned hardware maintenance windows). The pre-filter 2025 P75 was 66.4 h, materially

elevated by these advisory rows; post-filter the 2025 P75 sits at 23.7 h, in line with surrounding years. Tier and industry medians shift by less than 0.15 h under the filter. Full methodology rule is in Appendix C and at the SOR2026 analysis folder (`SOR2026_Advisory_Posting_Filter_Rule_v1.0.md`).

Table B2. Industry median MTTR (hours) by year, 9 representative industries

Industry	Tier	2023	2024	2025	2026 YTD
AI Model Provider	AI Provider	1.2	1.3	1.0	0.9
Consumer Internet	Application	1.1	1.3	1.8	1.1
DevOps & Developer Tools	Application	1.3	1.5	1.6	1.4
SaaS/Business Software	Application	1.6	1.6	1.7	1.7
Observability & Monitoring	Application	2.2	1.7	1.8	2.0
Cloud Infrastructure	Infrastructure	1.8	2.0	2.4	2.7
Fintech & Payments	Infrastructure	2.5	3.3	2.7	2.1
Communications	Infrastructure	4.3	4.5	3.7	3.5
Security/Identity	Infrastructure	3.5	3.4	3.4	4.7

Dataset v16.29. Same-cohort lens, with advisory-posting filter applied. The 9 industries shown match Figure 2; the full per-industry table including the remaining ~6 industries (E-commerce, Enterprise SaaS, Data Platform, DevTools, DevOps/CI-CD, Identity & Security) is available in the supporting analysis folder. Excludes maintenance periods and advisory postings (Duration > 7 days).

Table B3. Firm-capability transfer across failure types

Average pairwise rank correlation across the top-five failure modes is +0.62.

B4. MTTR significance forest plots

The two forest plots below back up the statistical-honesty claim in Section 5.2: per-category MTTR differences are real but modest, with company-clustered confidence intervals respecting the within-company pseudo-replication problem. Each row shows a category's median time-to-resolve with its 95% bootstrap confidence interval (resampling at the company level, not the incident level). Categories whose CI clears the cohort grand median, with non-negligible Cliff's delta against the rest of the pool and a Benjamini-Hochberg-corrected Mann-Whitney p-value below 0.05, are marked as robustly distinct from average.

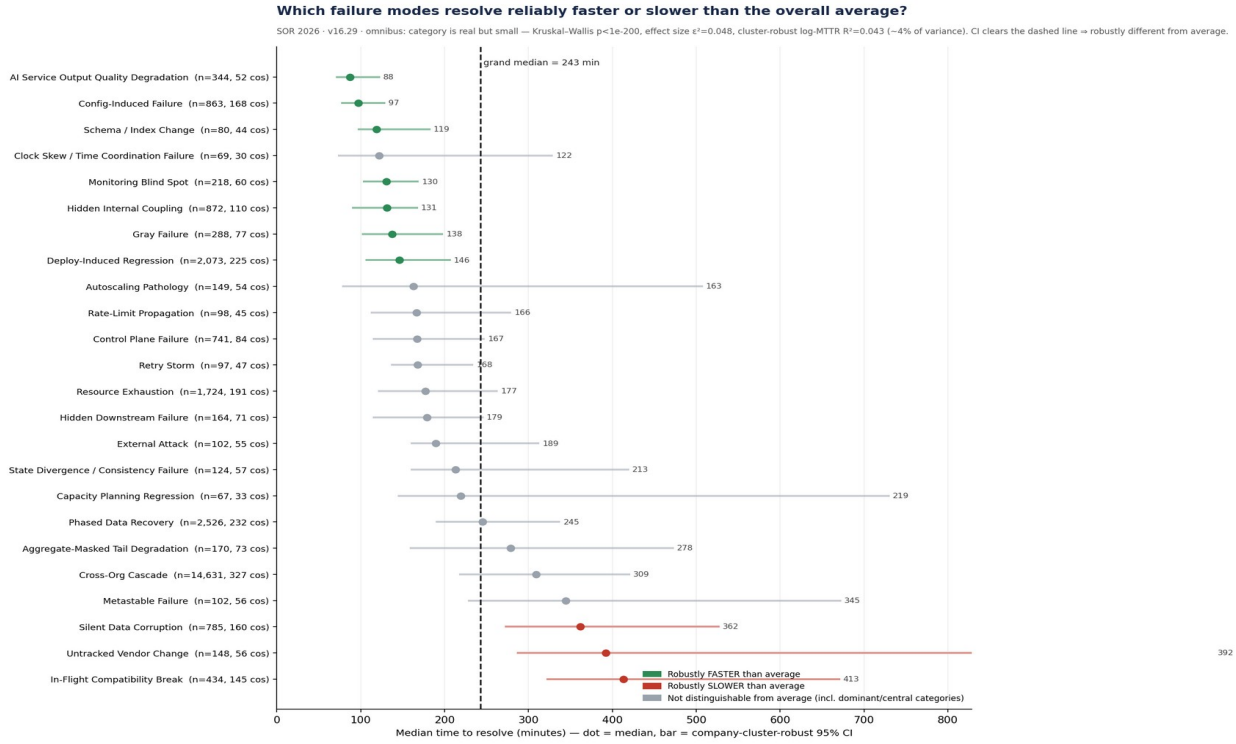

Figure 17. MTTR significance forest plot, by failure mode. Dataset v16.29. Per-category median plus company-cluster-robust 95% CI. The reliably-slow specialist modes are Silent Data Corruption, Untracked Vendor Change, and In-Flight Compatibility Break. Cross-organisation cascade is the largest category but its CI straddles the cohort grand median, consistent with the Section 5.2 nuance.

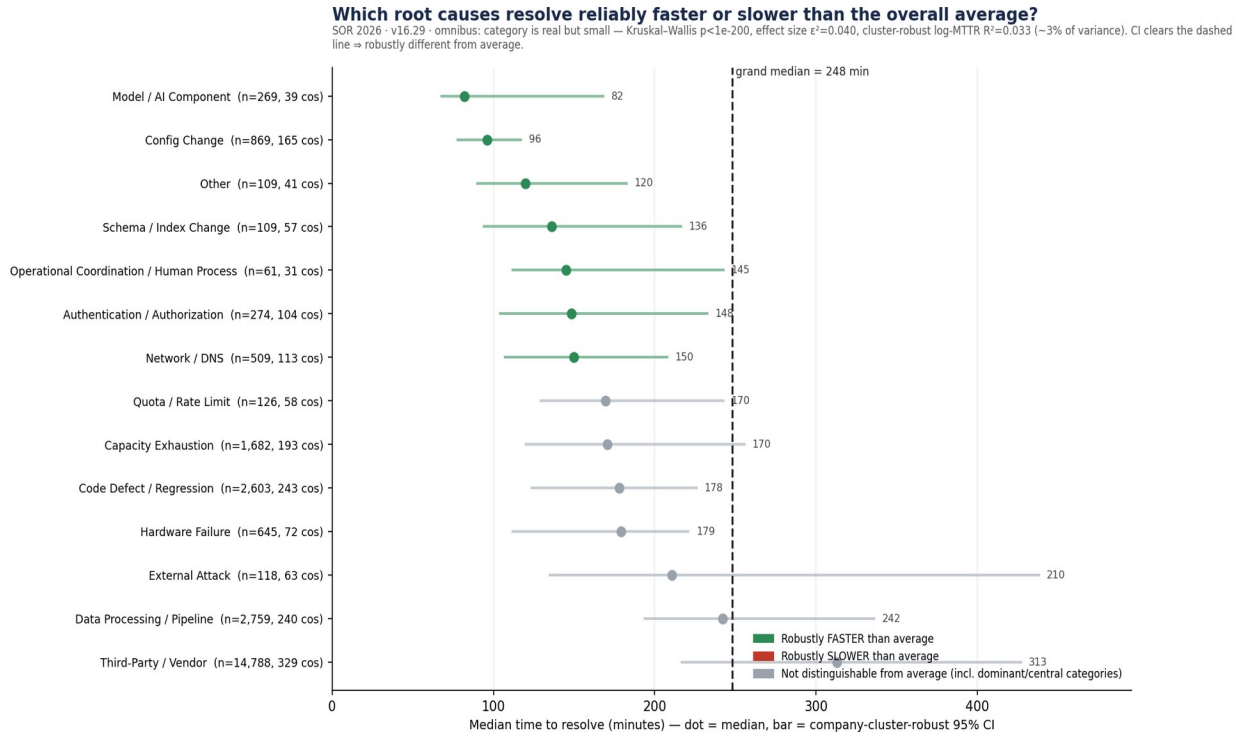

Figure 18. MTTR significance forest plot, by root cause. Dataset v16.29. Mirrors the failure-mode pattern: Third-Party/Vendor is the largest and most variable category; Config Change and Model/AI Component are reliably faster than average.

Deferred work. The same company-cluster-robust significance treatment is not yet applied to the remediation taxonomy because the post-mortem subset has only 361 rows with a primary remediation code, which is too thin for cluster-robust analysis. Once the remediation classifier scales above 1,000 rows, the same forest-plot treatment will be applied to remediations and added in a future revision.

Appendix C. Methodology

Source data. Public status pages, primarily Atlassian Statuspage and incident.io vendors. Dataset v16.29 contains 174,348 incidents collected between 2018 and June 2026.

Same-cohort definition. The 217 companies present every year of 2023, 2024, 2025, and 2026 YTD.

MTTR proxy. Median time-to-resolution is the publicly disclosed window between first status-page posting and resolution. A public-disclosure proxy for true internal MTTR.

Filter for MTTR analysis. Unplanned incidents only, with Duration_Minutes greater than 0 and less than 10,080 (7 days). The 7-day cut-off is the standing 'advisory-posting filter' (rule v1.0, 1 June 2026 onwards) and is applied to every duration-distribution and MTTR analysis in this report. Rationale: the standard maintenance filter (based on Category and Title containing 'maintenance' + 'scheduled') leaves a tail of advisory and lifecycle notices that are not true unplanned incidents — multi-month 'End of Engineering' postings, quarterly maintenance notifications that omit the 'scheduled' word, planned hardware maintenance windows. These rows distort tail percentiles without representing real incident response. The 168h threshold excludes ~8.7% of unplanned rows (overwhelmingly mis-labeled-maintenance or advisory notices), shifts tier and industry medians by less than 0.15 hours, and brings the corpus-wide same-cohort P75 in 2025 from 66.4h pre-filter to 23.7h post-filter. Full rule documentation: `SOR2026_Advisory_Posting_Filter_Rule_v1.0.md` in the analysis folder.

Classification status. Each incident the classifier touched has a status: Classified, Classified-Low-Confidence, Pending-Research, or Unable-to-Classify. Pending-Classification incidents are excluded from share calculations.

Raw share versus estimated share. The raw share is the count in a failure-mode category divided by classified-and-FM-assigned incidents (~41% of attempted). The estimated share treats attempted-but-unclassified incidents as having the same internal-failure-mode mix as classified-internal incidents.

Within-company paired test (Section 5.2). Across 130 companies that handle both slow-tier and fast-tier categories, 74% show cross-organisation cascade slower inside their own data (median ratio 1.68x, 95% CI 1.48 to 2.06, $p < 0.001$).

Blast radius (Section 7.2). Per-event blast radius is the number of distinct disclosing companies that filed an incident under the same Causal_Event_ID. The median is 1 disclosing company in every year of 2024-2026; the means are pulled up

by a small number of hero events (AWS Oct 20 2025 = 223 disclosing firms; Cloudflare Nov 18 2025 = 127; GCP Jun 12 2025 = 55).

Out of scope. Process, training, and organisational-layer factors. Internal incidents not posted to public status pages.

Microsoft Azure originator gap. Microsoft Azure is not yet scraped as an originator in this dataset. Figure 15's Cloud-share decline should be read with this caveat.

Published 1 June 2026 . StackGen Research . v1.9.3