StackGen

CONTEXT ATLAS

CONTEXT ATLAS V0.4 13 CONTEXT FAMILIES • 3 TIERS • 9 SUB-TYPE QUALIFIERS

Did your team have the right information when it mattered?

Map the diagnostic context used in every incident — and close the gaps before the next one hits.

DERIVED FROM THE STATE OF RELIABILITY 2026 ANALYSIS OF:

27K+	178K+	392	2018-26
classified incidents	status-page entries	companies	date range

A practical reference for SRE, platform and reliability teams.

Context Atlas v0.4

SSOR 2026 CLASSIFIER REFERENCE • 13 FAMILIES + AIPACK • 3 TIERS • 9 V0.4 SUB-TYPE QUALIFIERS • OPERATIONAL (PROMOTED FROM EXPERIMENTAL 2026-05-31)

HOW TO USE

1. Tag Context_Used:
F-codes the operator clearly referenced during detection → diagnosis → fix (pipe-separated: F4|F7|F11a)

2. Tag New_Context_Added:
F-codes from preventative commitments in the postmortem (forward-looking; use qualifiers: F4(synthetic-checks))

3. Tier Patterns:
Tier 1 dominates Context_Used · Tier 2 dominates New_Context_Added · Tier 3 under-represented in both · AIPack = AI-specific only

4. Only tag what's explicit or strongly implied:
Leave empty if body discloses nothing specific.

TIER 1 — Direct Signals — "what's happening right now"

F3, F4, F5, F7

F3 External Dependency Posture

Upstream health, inventory, track record, blast-radius map, dependency contracts. Crucial when a named upstream is involved.

upstream health live · SLA/quota contracts · blast-radius map · dependency criticality

F4 Telemetry & Observability

Golden signals (traffic/latency/errors/saturation), log streams, distributed traces, dashboards, alerts, synthetic checks, security telemetry.

Datadog/CloudWatch dashboards · log streams · traces · alerts · synthetic checks

F5 Service Behaviour

Behavioural baselines ("what's normal"), norm-deviation signals, behavioural profile (request mix, dep-call shape). One level above raw telemetry.

request-mix baseline · error-rate norm · dep-call profile

F7 Change & Runtime Config History

Deploy events, release provenance, config changes, feature flag flips, IaC drift, credential/cert changes. The "what changed" context — most-cited diagnostic type in the SSOR corpus.

deploy event · config change · feature flag flip · IaC drift · cert rotation

TIER 2 — Structured State & Knowledge — "what's the system / who owns it / what have we learned"

F1, F2, F6, F9, F10, F11a, F11b

F1 Topology & Ownership

Environment, cloud accounts, service inventory + depends_on graph, compute/datastore inventory, source code, identity/access config.

- service graph
- cloud account
- source code
- access config

F2 Declared State (Architecture & IaC)

IaC manifests (Terraform/Pulumi/Helm), architecture docs (RFCs/ADRs/design docs), identity/access policy (declared rules). The "what we intended" view.

- Terraform manifests
- Helm charts
- ADRs
- RFC docs

F6 SLO & Service Contract

SLO definitions & targets, SLI track record + error budget state, SLA/contractual obligations. The operator's commitment.

- error budget state
- SLO targets
- SLA obligations

F9 Incident Memory

Incident inventory, recurring pattern recognition, resolution lineage ("what worked last time"), near-miss inventory, cross-service cascade map.

- prior similar incidents
- what worked last time
- near-miss log

F10 Ownership, Escalation & Org Context

Org/team hierarchy, on-call rotation (PagerDuty/Opsgenie), escalation topology, subject-matter expert map. WHO knows what / who's on call.

- on-call rotation
- escalation path
- SME map

F11a Procedural Knowledge ("what to do")

Runbooks & playbooks, troubleshooting guides, remediation step docs, known-fix repos. Stepwise instructions.

- runbook
- playbook
- troubleshooting guide
- known-fix repo

F11b Narrative Knowledge ("what was said")

Postmortem narratives, Slack/Teams/email history, call transcripts, wiki/KB articles. Different trust profile from Procedural.

- postmortem doc
- Slack history
- call transcript
- wiki article

TIER 3 — Impact, Decision & Action Intelligence — "what matters / what can we safely do"

F8, F12 + AIPACK (DOMAIN PACK)

F8 Action Readiness, Authorization & Constraints

Rollback capability + path validity, failover + rehearsal outcomes, feature-flag inventory + test results, DR posture (RTO/RPO) + drill outcomes. Pairs CAN-we with HAVE-we-tested-it.

- rollback-rehearsal
- failover test
- DR drill
- auth constraints

F12 Customer & Business Context

Service-to-customer impact model, account-tier / business-criticality tags, live customer-impact estimation.

- customer impact model
- account tier tags
- live impact estimate

AI PACK AI Operator Context & Guardrails

Model inventory, production model performance, training data lineage, inference endpoint map, fallback model tests, AI action authorization, human-in-the-loop triggers. Use for AI-specific context in AI-incident RCAs.

- model inventory
- fallback model tests
- HITL triggers
- AI auth gates

Context Atlas is one of three dimensions: FM/RC (what failed+why) · Fix/RM (what was done) · Context Atlas (what info was used/committed)

NEW V0.4

v0.4 Sub-type Qualifiers

9 QUALIFIERS · USE IN PARENS ON F-CODES

F4(observability-fate-sharing)	Observability system has fate-share with the thing being observed (Roblox Consul anti-pattern)
F8(destructive-action-confirmation)	Confirmation before destructive operations (Atlassian site-delete, GitLab rm-rf wrong host)
F4(context-source-deletion)	The catalog state itself was the deleted context (TravisCI GCE image deletion)
F10(cross-team-credential-coordination)	Credential rotation in one team's calendar, downstream config didn't track it (F7–F10 boundary)
F4(stack-trace-capture)	Raw crash-artefact capture (distinct trust from running telemetry)
F4(stale-monitoring-data)	Data exists but is wrong (Indian grid 2012 frozen readings, Northeast 2003 silent EMS alarm processor)
F3(regulated-upstream)	Operator can't switch providers (Zerodha NSE-mandated leased-line vendor)
F11b(operator-disclosure-discipline)	Meta-context: when/how operators disclose (Sun 1999-2000 cosmic-ray bug withheld)
F8(safety-context-redundancy)	Silent safety context channel masking races (Therac-25 hardware interlocks)
F4(predictive-degradation-model)	Multiple operator commitments (Anthropic, Inngest). Pairs with RM-46 once any operator publishes implementation. `tracked — not new v0.4`

Turn information gaps into observability investments

Tag what your team actually used during the incident

At post-incident review, assign F-codes to Context_Used — the information your team referenced while detecting, diagnosing, and fixing the incident.

EXAMPLES:

- **F4 (Telemetry & Observability)** – you used Datadog dashboards and error-rate alerts
- **F7 (Change & Runtime Config History)** – you checked the deploy log to confirm timing
- **F11a (Runbooks)** – you followed an existing playbook to execute the fix

These are the context types that are working for your team right now.

Tag what your team committed to adding

Assign F-codes to New_Context_Added — the information gaps your postmortem identified and committed to closing.

Common patterns:

- **F4(synthetic-checks)** – you didn't have data validation checks in place
- **F9 (Incident Memory)** – you reinvented a fix your team had already solved before
- **F11a (Runbook)** – there was no playbook; someone had to figure it out from scratch

The gap between Context_Used and New_Context_Added is your observability investment roadmap.

Check your tier coverage

Tier 1 (F3, F4, F5, F7) –

direct signals, "what's happening right now." Should dominate Context_Used.

Tier 2 (F1, F2, F6, F9, F10, F11a, F11b) –

structured knowledge, "what's the system / what have we learned." Should dominate New_Context_Added.

Tier 3 (F8, F12, AIPack) –

decision and impact intelligence. Consistently under-represented — most teams don't know blast radius or rollback status until well into the incident.

If Tier 1 is absent from Context_Used, your team is diagnosing blind.

PRACTICAL PROMPT

"In your last major incident, which F-code was most notably missing from Context_Used? That's your highest-priority observability gap."

From the State of Reliability 2026

The Context Atlas is one part of a connected incident-learning framework developed through the analysis of public reliability data.

Failure Mode	What failed?
Root Cause	Why did it fail?
Remediation	What was done about it?
Context Atlas	What information was required to diagnose and act?

Use the framework to create a more consistent language for postmortems, incident trends, and reliability investment.

DOWNLOAD THE FULL FRAMEWORK

FM Cheat Sheet
Failure Mode Taxonomy

RC Cheat Sheet
Root Cause Taxonomy

RM Cheat Sheet
Remediation Taxonomy

Download The State of Reliability 2026 full report.

Follow StackGen for the report release and additional reliability research.

stackgen.com/state-of-reliability-2026