RESOURCE · OPERATIONS · RES-003

Agent Incident Runbook

VERSION

v1.0

TRACKED AT

agentmodeai.com/holding/RES-003/

SOURCE

agentmodeai.com/resources/

LICENCE

CC BY 4.0

Agent Mode AI publishes editorial work for senior IT leaders. Every claim including the premise of this template is tracked on a public ledger. Verdict updates and corrections at agentmodeai.com/holding/.

Version 1.0 · Published 4 May 2026 · 60-day review cycle

A four-phase response to agent-mode AI incidents, designed as an overlay on standard SRE incident response.

Source explainer: https://agentmodeai.com/resources/agent-incident-runbook/ Holding-up tracking: https://agentmodeai.com/holding/RES-003/

Phase 1: Detection (target: 4h from action to operator awareness)

Four tripwires that catch agent incidents within the 4-hour window:

Tripwire	Threshold	Pages
Cost-rate	Per-agent spend > 2x rolling-7-day median in any 1h window	Agent owner
Action-rate	Per-agent API call rate > 3x rolling-7-day median, OR call against unused endpoint	Agent owner
Outcome	Customer-visible artifact created with bypassed/sub-threshold human review	Operations queue
Authorisation	Agent invocation of tool/API outside declared inventory	Security team

If the incident surfaces only via downstream consequence (regulator, partner, public), detection clock starts at consequence-surface; post-mortem flags the missing upstream tripwire.

Phase 2: Containment (target: 30s from confirmed harm to all-agents-halted)

Kill-switch must be operationally accessible to at least 3 roles: SRE on-call, security on-call, business owner.

Containment decisions in priority order:

Decision	Use when	Time
All-agents-halt	High blast radius + unclear failure mode	Default
Per-agent-halt	Affected agent isolated, others unaffected	After all-halt verification
Per-tool-halt	Failure is specific tool the agent is calling incorrectly	Agent continues with reduced capability
Per-action-quarantine	Actions queued for human approval rather than executing	Drain + assess each

Containment is the moment the rollback timer starts.

Phase 3: Rollback (procedures by action class)

Class	Rollback procedure	Operator	Time budget	Substitute if impossible
1. Database writes	Transactional rollback OR compensating writes	DBA + agent owner	4h	Customer notification
2. External API calls (payments, bookings, transfers)	Idempotent reversal where supported	Finance + agent owner	24h	Customer comms template + financial reserve
3. Customer communications (emails sent, messages posted)	Sent comms cannot be unsent	Comms lead + legal	4h	Follow-up communication template
4. Document creation/publication	Internal: delete or supersede. Public: unpublish + SEO assess (410-Gone vs 301 to corrections page)	Content lead	1h unpublish, 24h SEO	Public correction notice
5. Code commits/deployments	git revert if not deployed; standard rollback + audit if deployed	Eng on-call	30 min revert, 4h audit	Hotfix forward
6. Identity/access changes	Deprovision; audit what was done with the access	Security + IAM lead	30 min revoke, 24h downstream audit	Incident extension if downstream actions found
7. Knowledge-base/vector-store writes	Delete records; audit which agents retrieved corrupted state	Data eng + agent owner	4h delete, 7d downstream monitoring	Re-index + re-train downstream agents

For each class capture: rollback procedure, authorised operator, time budget, substitute action, verification step.

Phase 4: Post-mortem (MTTD-for-Agents detection chain)

Phase	When	Elapsed
1. Action	Agent took the harmful action	T+0
2. Signal	System that should have detected received the signal	T+?
3. Trigger	Tripwire fired	T+?
4. Page	Human paged	T+?
5. Acknowledgement	Human started working the incident	T+?

MTTD = Phase 5 - Phase 1. Each gap = its own root cause:

- Action → Signal gap = missing instrumentation

- Signal → Trigger gap = tripwire didn't fire
- Trigger → Page gap = alerting hygiene
- Page → Acknowledgement gap = on-call response

Also capture:

Item	Description
Blast radius	Affected: records / dollars / users / regulatory surface
Cost	Financial + time + opportunity impact
Vendor implication	Did vendor kill-switch SLA hold? If not → contract renegotiation material
Tripwire delta	What new tripwire would have caught this 4h earlier? Add + verify against timeline
Runbook delta	What rollback procedure was missing or wrong? Update action-class table
Communication record	Customers, regulators (EU AI Act Art 26 serious-incident reporting), board (if threshold-triggered)
Recurrence prevention	Single change preventing exact recurrence + owner + date

Quarterly fire-drills

Tabletop with agent-system on-call rotation against synthetic scenario from rolling list. Verifies:

- Kill-switch still works
- On-call knows the procedure
- Rollback procedures still match the agent action surface

Roles

Role	Responsibility
Incident commander	Coordinates phases 1-4; declares incident open + closed
Agent owner	Authority over the affected agent's configuration + lifecycle
Security on-call	Kill-switch authority; security implications assessment
SRE on-call	Standard SRE response + agent-system integration
Business owner	Customer + commercial impact decisions; communication approval
Legal counsel	Regulatory reporting decisions (EU AI Act Art 26, GDPR Art 33-34)

Frameworks referenced

- MTTD-for-Agents methodology (https://agentmodeai.com/mttd/)
- EU AI Act Article 26 (deployer obligations) + Article 73 (serious-incident reporting)

- GDPR Article 33 (breach notification to supervisory authority) + Article 34 (data subject notification)
- NIST AI RMF 1.0 Manage function

Licence