AI Incident Response Playbook

Structured Response Procedures for AI System Incidents | TeraSystemsAI

EMERGENCY INCIDENT HOTLINE

[YOUR 24/7 HOTLINE NUMBER]

For Critical/P1 incidents - immediate response required

1. Severity Classifications

Classify incidents immediately upon detection to determine response urgency and escalation path.

CRITICAL (P1)

Response: 15 min

System down, safety breach, data exposure, active harm to users

HIGH (P2)

Response: 1 hour

Major functionality impaired, bias incident affecting users, compliance violation

MEDIUM (P3)

Response: 4 hours

Degraded performance, non-critical feature failure, potential compliance issue

LOW (P4)

Response: 24 hours

Minor issues, cosmetic problems, feature requests misclassified as bugs

2. Incident Response Workflow

Phase 1: Detection & Triage (0-15 minutes)

1 Acknowledge & Log Incident

Create incident ticket, assign severity level, timestamp all actions

2 **Assemble Response Team**

Page on-call engineer, notify incident commander, engage relevant SMEs

3 **Initial Assessment**

Determine blast radius, affected users/systems, potential root causes

Phase 2: Containment (15-60 minutes)

4 **Implement Immediate Mitigation**

Activate kill switch if needed, rollback to last known good state, enable fallback systems

5 **Communicate Status**

Notify stakeholders, update status page, prepare customer communications

Phase 3: Resolution (1-24 hours)

6 **Root Cause Analysis**

Investigate logs, traces, and metrics to identify root cause

7 **Implement Fix**

Deploy permanent fix, validate resolution, restore full functionality

Phase 4: Recovery & Review (24-72 hours)

8 **Post-Incident Review**

Conduct blameless retrospective, document lessons learned, update runbooks

3. Incident Type Playbooks

3.1 AI Model Failure / Performance Degradation

- ☐ Check model serving infrastructure health
- ☐ Review recent deployments or config changes
- ☐ Analyze input data distribution for drift
- ☐ Compare current metrics to baseline
- ☐ Consider rollback to previous model version
- ☐ Enable fallback/rule-based system if available

3.2 Bias Incident / Fairness Violation

- ☐ Document affected users and decisions
- ☐ Preserve evidence for investigation
- ☐ Disable or limit affected functionality
- ☐ Notify legal/compliance team
- ☐ Prepare remediation plan for affected users
- ☐ Conduct fairness audit on model

3.3 Security Breach / Adversarial Attack

- ☐ Isolate affected systems immediately
- ☐ Preserve logs and evidence
- ☐ Notify security team and CISO
- ☐ Assess data exposure scope
- ☐ Engage incident response retainer if needed
- ☐ Prepare regulatory notifications (GDPR, etc.)

3.4 Data Pipeline Failure

- ☐ Identify pipeline failure point
- ☐ Check upstream data source availability
- ☐ Assess impact on model predictions
- ☐ Enable cached/stale data mode if safe
- ☐ Coordinate with data engineering team
- ☐ Plan data backfill if needed

4. Communication Templates

4.1 Internal Status Update

```
INCIDENT STATUS UPDATE
━━━━━━━━━━━━━━━━━━━━━

Incident ID: [INC-XXXX]
Severity: [P1/P2/P3/P4]
Status: [Investigating/Identified/Mitigating/Resolved]
Commander: [Name]

SUMMARY:
[Brief description of the incident]

IMPACT:
- Users affected: [Number/Percentage]
- Systems affected: [List]
- Business impact: [Description]

CURRENT ACTIONS:
- [Action 1 - Owner - ETA]
- [Action 2 - Owner - ETA]

NEXT UPDATE: [Time]
━━━━━━━━━━━━━━━━━━━━━
```

4.2 Customer Communication

```
Subject: [Service Name] - Incident Update

We are aware of an issue affecting [brief description].
Our team is actively working to resolve this.

Current Status: [Status]
Estimated Resolution: [Time if known, or "We are working to resolve this as quickly as possible"]

Impact: [What customers may experience]

We will provide updates every [30 minutes/1 hour/as available].

We apologize for any inconvenience this may cause.

- The [Company] Team
```

4.3 Post-Incident Summary

```
POST-INCIDENT REPORT
═══════════════════

Incident: [INC-XXXX]
Date: [Date]
Duration: [X hours Y minutes]
Severity: [Level]

EXECUTIVE SUMMARY:
[2-3 sentence summary]

TIMELINE:
[HH:MM] - Incident detected
[HH:MM] - Response team assembled
[HH:MM] - Root cause identified
[HH:MM] - Fix deployed
[HH:MM] - Incident resolved

ROOT CAUSE:
[Detailed explanation]

IMPACT:
- [Metric 1]
```

```
- [Metric 2]

REMEDIATION ACTIONS:
☐ [Action 1] - Owner - Due Date
☐ [Action 2] - Owner - Due Date

LESSONS LEARNED:
- [Lesson 1]
- [Lesson 2]
━━━━━━━━━━━━━━━━━━━━
```

5. Escalation Contacts

Role	Primary Contact	Backup Contact	Escalation Criteria
On-Call Engineer	[Name/Contact]	[Name/Contact]	First responder for all incidents
Incident Commander	[Name/Contact]	[Name/Contact]	P1/P2 incidents, coordinates response
ML Engineering Lead	[Name/Contact]	[Name/Contact]	Model-specific issues
Security Officer	[Name/Contact]	[Name/Contact]	Security breaches, data exposure
Legal/Compliance	[Name/Contact]	[Name/Contact]	Regulatory issues, bias incidents
Executive Sponsor	[Name/Contact]	[Name/Contact]	P1 incidents, external comms needed

---