DOWNLOADABLE ENGINEERING REFERENCE / PDF EDITION

Incident Readiness and Telemetry Checklist

This checklist is about operational readiness. If you cannot reconstruct a failure safely, you cannot contain it quickly.

The goal is high-signal telemetry that supports incident response without leaking secrets or sensitive data.

> [!OPERATIONAL] Scope and privacy This checklist is written for teams shipping real systems. It favors correlation, reason codes, and auditability over invasive tracking or raw payload logging.

1) Logging Policy (Before You Log Anything)

- Define a denylist of sensitive data:
  - tokens, credentials, secrets
  - raw personal data not needed for debugging
  - full request/response bodies by default
- Define what you *do* log:
  - correlation IDs
  - principal identifiers (stable, non-secret)
  - reason codes and policy identifiers
  - boundary transitions (auth events, privilege changes)

2) Tracing and Correlation

- Ensure you can follow a request across:
  - client request ID
  - gateway request ID
  - service trace/span IDs
  - queue/job IDs
- Verify correlation survives retries and partial failures.

3) Security-Relevant Signals

- auth failures grouped by reason (not raw payload)
- token refresh failures and rotation anomalies
- suspicious rate patterns (enumeration, brute force, automation)
- permission denials that indicate policy drift
- unexpected admin operations or privilege changes

4) Mobile-Specific Signals (Privacy-Safe)

- build version, OS version, device class (coarse)
- network type and connectivity errors
- crash rates per version
- critical flow success rates (login, onboarding, purchase, admin)

Avoid device fingerprinting or invasive tracking.

5) Availability and Degradation

- define SLOs for critical endpoints/flows
- monitor:
  - latency percentiles
  - error rate
  - saturation (CPU, memory, queue depth)
- have a clear “degraded mode” plan for non-critical features

6) Incident Playbook Inputs

During an incident, ensure you can answer:

- what changed (deploys, config, dependencies)
- which users are affected
- where failures originate (boundary)
- what mitigation is safest (rollback, flag, capacity)

7) Retention and Access

- define retention for logs and traces based on operational need
- restrict access to sensitive operational tooling
- ensure audit logging exists for access to logs and incident systems

8) “Good telemetry” in one table

Goal	Prefer	Avoid
Reconstruct incidents	Correlation IDs, reason codes, structured events	Raw secrets, full bodies by default
Understand auth failures	Token kind + deny reason + boundary	Logging tokens or credentials
Detect abuse	Rate patterns, policy identifiers, anomaly counters	Fingerprinting or invasive user tracking
Debug mobile issues	Coarse device class, build version, flow success rates	Persistent identifiers not needed for ops

Public engineering reference, not a completed assessment, legal advice, audit, certification, or customer deliverable.

Editable Markdown source: demonicbinary.com/downloads/incident-readiness-telemetry-checklist.md