Incident Response Plan

1. Purpose

This plan establishes procedures for detecting, responding to, and recovering from security incidents affecting Subcontext systems, data, or services. Its aims are to limit harm to customers and the people whose data they entrust to us, to restore normal service quickly, to meet our legal and contractual notification duties, and to learn from every incident so it is less likely to happen again.

2. Scope

This plan covers all security incidents including but not limited to: unauthorised access, data breaches, malware, denial of service, insider threats, and vulnerabilities actively being exploited.

It applies to:

- The Subcontext platform, including the web application, REST API, MCP server, embeddable widgets and the iOS, Android and Mac apps
- The AWS infrastructure the platform runs on, and the source code and build pipeline used to produce it
- Data held for customers (where Subcontext acts as processor) and data about our own users and contacts (where Subcontext acts as controller)
- Incidents at a sub-processor that affect Subcontext data or services
- All Subcontext personnel, contractors and anyone else with access to Subcontext systems

Service outages that have no security cause are handled under the Business Continuity & Disaster Recovery Policy, using the same roles and communication steps set out here where they are useful.

3. Definitions

Term	Meaning
Security event	An observable occurrence that may have security significance, such as a failed sign-in, a firewall block or an alarm. Most events are not incidents.
Security incident	An event, or series of events, that has compromised or is reasonably likely to compromise the confidentiality, integrity or availability of Subcontext systems or data.
Personal data breach	A security incident leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data, as defined in the UK GDPR.
Incident Lead	The person who owns the response to a particular incident from identification to closure.
Awareness	The point at which Subcontext has a reasonable degree of certainty that a personal data breach has occurred. The 72-hour ICO clock runs from this point.

4. Incident Classification

Severity	Description	Response Time
Critical	Active data breach, system compromise, or widespread service outage	Immediate (within 1 hour)
High	Confirmed vulnerability being exploited, significant unauthorised access	Within 4 hours
Medium	Suspected breach, anomalous activity, contained security event	Within 24 hours
Low	Minor policy violation, unsuccessful attack attempt, security advisory	Within 72 hours

The response time is the time within which the Incident Lead is assigned and containment work begins. Severity is set during identification and is reviewed as facts emerge: an incident is raised to a higher severity as soon as new information justifies it, and is lowered only with the Incident Lead's agreement. Where there is doubt between two levels, the higher one is used.

4.1 Worked examples

Situation	Typical severity
Evidence that one customer organisation can read another organisation's data	Critical
A production secret (for example a database credential) is found in a public place	Critical
A known vulnerability in a dependency is being exploited against the platform	High
A single user account shows sign-ins from an unexpected location after a lockout	Medium
A spike of requests blocked by the web application firewall, with no sign of success	Low
A team member emails a document containing customer data to the wrong internal colleague	Low, raised if the data left Subcontext

5. Detection

- Automated monitoring of application logs, security events, and infrastructure metrics
- Audit trail review for anomalous user activity
- Dependency vulnerability scanning
- Reports from employees, customers, or external parties

Sources in practice include infrastructure alarms on the database, cache, load balancer and application servers (which email hello@subcontext.com), error monitoring alerts on new, reopened and rising errors, the web application firewall and its logs, account lockouts and rate limiting, the audit log, container image scanning, and reports received through the Vulnerability Disclosure Policy or the support channels set out in the Service Levels & Support Policy.

5.1 How to report an incident

Anyone who suspects an incident must report it straight away to the Security Officer or another member of the founding team, and must not attempt to investigate further on their own. Customers and members of the public can report to hello@subcontext.com or support@subcontext.com, or through an in-app support

ticket. Reports are acknowledged and assessed even when they turn out to be false positives, and nobody is criticised for reporting in good faith.

6. Response Procedure

6.1 Identification

- Confirm the incident is genuine (not a false positive)
- Determine the type, scope, and severity
- Identify affected systems, data, and users
- Assign an incident lead

The Incident Lead opens an incident record at once and keeps a timestamped log of facts, decisions and actions throughout the response. The record notes whether personal data is or may be involved, which customer organisations are affected, and whether Subcontext is acting as processor or controller for that data.

6.2 Containment

- Isolate affected systems to prevent further damage
- Preserve evidence for investigation
- Implement temporary controls (block IPs, revoke credentials, disable accounts)
- Assess whether the incident is ongoing

Containment options available to the team include: rate limiting or blocking traffic at the web application firewall; locking user accounts and ending their sessions; revoking API tokens and OAuth access; rotating secrets; replacing affected application tasks; placing the platform in maintenance mode for a planned holding period; and, where the application itself cannot be trusted or reached, an emergency lockdown at the load balancer. Evidence (relevant logs, audit entries, container images and database snapshots) is preserved before it can expire or be overwritten, and is kept until the post-incident review is closed.

6.3 Eradication

- Identify and remove the root cause
- Patch vulnerabilities
- Reset compromised credentials
- Verify that the threat has been eliminated

Fixes made during an incident follow the emergency change route in the Change Management Policy.

6.4 Recovery

- Restore affected systems from clean backups where necessary
- Verify system integrity before returning to service
- Monitor for recurrence

Recovery decisions, including restoring the database to an earlier point in time, are taken by the Incident Lead with a director, because they may affect data written since that point. Heightened monitoring continues for an agreed period after the incident is closed.

6.5 Notification

- **Customers:** notify affected customers without undue delay in accordance with the Data Processing Agreement
- **ICO:** notify the Information Commissioner's Office within 72 hours of becoming aware of a personal data breach that is likely to result in a risk to individuals' rights and freedoms
- **Data subjects:** notify affected individuals without undue delay where the breach is likely to result in a high risk to their rights and freedoms

Where Subcontext is the processor, the customer is the controller and decides whether to notify the ICO and the individuals concerned. Our role is to tell the customer without undue delay and to give them the information they need: the nature of the breach, the categories and approximate number of people and records affected, the likely consequences, the measures taken or proposed, and a contact point. Information is provided in stages if it is not all available at once. Notifications to the ICO and data subjects described above are made by Subcontext where it is the controller.

Customer notifications are sent by email to the organisation's contacts, with updates until the incident is resolved. Customers who are regulated firms may have their own reporting duties to the FCA or PRA, and we will support them with timely, accurate information. The decision whether a breach is notifiable, and the reasons, are recorded even where no notification is made.

7. Roles and Responsibilities

- **Incident Lead:** coordinates the response, makes containment decisions, manages communications
- **Engineering:** investigates technical aspects, implements fixes, restores services
- **Founding Team:** approves external communications, engages legal/regulatory advisers as needed

The Security Officer acts as Incident Lead unless they assign the role to someone else, and owns this plan. Because the team is small, one person may hold more than one role, but every Critical or High incident involves at least two people so that decisions are checked. All personnel must report suspected incidents promptly and follow instructions from the Incident Lead.

8. Communication

- Internal discussion of an incident is kept to those who need to know
- No statement about an incident is made to customers, the press or the public except as approved by the founding team
- Communications are factual and do not speculate about cause or impact before it is established
- A record is kept of every external notification: who was told, when, and what was said

9. Post-Incident Review

- Conduct a post-incident review within 5 business days of resolution
- Document: timeline, root cause, impact, response actions, and lessons learned
- Update policies, procedures, and controls as needed
- Share findings (appropriately redacted) with the team

Reviews are blameless: they look at how systems and processes allowed the incident, not at assigning fault. Each improvement action has an owner and a target date, and is tracked to completion. Relevant risks are added to or updated in the risk register. Low severity incidents may be reviewed together in a periodic summary rather than individually.

10. Testing

This plan is reviewed and tested at least annually through tabletop exercises or simulated incidents.

Exercises use realistic scenarios, such as a leaked credential, a cross-tenant data exposure or a sub-processor breach, and test notification timings as well as technical steps. Findings are recorded and used to improve this plan.

11. Compliance

The Security Officer checks that incident records are complete, that notification decisions are documented and that review actions are closed. Failure to report a known or suspected incident is a breach of policy and may result in disciplinary action.

12. Related Documents

- Information Security Policy
- Business Continuity & Disaster Recovery Policy
- Vulnerability Management Policy
- Vulnerability Disclosure Policy
- Logging & Monitoring Policy
- Change Management Policy
- Supplier Risk Management Policy
- Operational Resilience Statement
- Service Levels & Support Policy

13. Document Control

Item	Detail
Owner	Security Officer
Approved by	The directors of Subcontext Ltd
Review	At least annually, and after any Critical or High incident
Effective	October 2026

Subcontext Ltd, registered in England and Wales, Company No. 17019645. Registered office: 124 City Road, London, EC1V 2NX. The current version of this document is published at subcontext.com.