Business Continuity & Disaster Recovery Policy

1. Purpose

This policy establishes Subcontext's approach to maintaining service availability and recovering from disruptive events. It sets out how the platform is built to withstand failure, how data is backed up, the recovery targets we work to, and who does what when a serious disruption happens.

Many of our customers are regulated firms that rely on Subcontext for important business services. This policy is the internal counterpart to the Operational Resilience Statement, which summarises the same arrangements for customers.

2. Scope

This policy covers:

- The production Subcontext platform: web application, background processing, REST API, MCP server, embeddable widgets and the services the mobile apps depend on
- Production data stores: the database, cache and file storage
- The source code, infrastructure code and build pipeline needed to rebuild the platform
- Third-party services the platform depends on
- The people and knowledge needed to operate and restore the service

Security incidents are handled under the Incident Response Plan. Where an incident causes a disruption, both documents apply.

3. Definitions

Term	Meaning
Disruption	Any event that stops or materially degrades the service, whatever the cause
Disaster	A disruption that cannot be resolved by the platform's automatic recovery and needs restoration of infrastructure or data
Recovery Time Objective (RTO)	The target maximum time from declaring a disaster to restoring service
Recovery Point Objective (RPO)	The target maximum period of data that could be lost, measured back from the moment of failure

4. Roles and Responsibilities

Role	Responsibility
Security Officer	Owns this policy, keeps recovery procedures current and arranges testing
Recovery Lead	A member of the founding team who takes charge during a disaster: declares it, directs recovery and decides when service is restored. Normally the Security Officer.
Engineering	Carries out restoration and redeployment, and verifies the service afterwards
Founding team	Approves customer communications and decisions with commercial or contractual impact

5. Infrastructure

- The platform is hosted on Amazon Web Services (AWS) in the EU (Ireland) region
- AWS provides redundant infrastructure with multiple availability zones
- The platform uses managed database services (RDS) with automated backups
- Static assets and file storage use AWS S3 with cross-region durability

In addition:

- Application and background-processing servers run in two availability zones and scale automatically between 1 and 6 tasks with demand
- Health checks detect unhealthy servers and replace them automatically, without manual action
- Infrastructure alarms on the database, cache, load balancer and application servers alert the team by email
- The database has deletion protection enabled, so it cannot be deleted by accident

6. Backup Strategy

- **Database:** automated daily backups with point-in-time recovery capability, retained for 7 days
- **File storage:** S3 with 99.999999999% durability (11 nines)
- **Configuration:** infrastructure managed as code, version-controlled in Git
- **Secrets:** encrypted and stored separately from application code

Database backups are encrypted at rest. Application releases are stored as immutable, versioned container images, so any recent release can be redeployed exactly as it was. Backups are not copied into development or test environments.

7. Recovery Objectives

Metric	Target
Recovery Time Objective (RTO)	4 hours
Recovery Point Objective (RPO)	24 hours (database), near-zero (file storage)

These targets apply to the core platform. Functions that depend on a third party (for example sending SMS or taking card payments) are restored as soon as that third party is available again.

8. Disaster Recovery Procedures

- Database restoration from automated backups or point-in-time recovery
- Application redeployment from version-controlled infrastructure code
- DNS failover to alternative infrastructure if primary region is unavailable
- Communication to affected customers via email and status page

8.1 Steps in a disaster

1. **Assess.** Whoever notices the disruption alerts the founding team. The Recovery Lead confirms the cause and scale, and decides whether automatic recovery is enough or a disaster should be declared.
2. **Declare and communicate.** The Recovery Lead declares the disaster, starts a timestamped log, and tells affected customers by email what is affected and when the next update will be.
3. **Protect.** Where needed, maintenance mode shows a holding page while work is done, or an emergency lockdown at the load balancer stops traffic reaching an application that cannot be trusted.
4. **Restore.** Data is restored using the most suitable route (point-in-time recovery or a backup), and the application is redeployed from a known-good image and the infrastructure code.
5. **Verify.** Engineering confirms the service is working, data is consistent, and alarms and monitoring are healthy before traffic is returned.
6. **Resume and update.** Service is restored, customers are told, and monitoring is heightened for an agreed period.
7. **Review.** A review is held within 5 business days, following the format in the Incident Response Plan, and actions are tracked to completion.

8.2 Worked example

If the database became corrupted at 14:00 by a faulty data change, the Recovery Lead would declare a disaster, place the platform in maintenance mode, and restore the database to a point shortly before 14:00 using point-in-time recovery. The application would then be pointed at the restored database, checked, and returned to service. Customers would be told which period of activity, if any, could not be recovered, and asked to re-enter anything essential.

9. Service Dependencies

The platform depends on third-party services including AWS, LLM providers, Twilio, Stripe, and SES. Outages in these services may affect platform functionality. Subcontext monitors third-party service status and communicates known issues to customers.

Dependency	Effect of an outage	Response
AI model provider	Agents cannot respond	Agents can be moved to an alternative model; conversations can be handed to the customer's own team
Messaging or telephony provider	Affected channels stop sending or receiving	Other channels continue; customers are told which channels are affected

Dependency	Effect of an outage	Response
Payment provider	Payment steps cannot complete	Payment steps are paused and resumed when the provider recovers; no card or bank details are lost
Email delivery	Outbound and inbound email are delayed	Other channels continue; customers are told email is affected

Suppliers are assessed for resilience under the Supplier Risk Management Policy, and the current list is published in the Sub-Processor List.

10. People and Knowledge

Infrastructure is defined as code and the platform's operation is documented, so that more than one person can operate and restore it. Access needed for recovery is held by more than one member of the founding team, and recovery procedures are kept where they can be reached if the platform itself is unavailable.

11. Testing

- Backup restoration is tested at least annually
- Recovery procedures are reviewed and updated as infrastructure changes
- Third-party dependency failover scenarios are assessed periodically

Restore tests record the time taken and are compared with the RTO and RPO. Any shortfall is treated as an improvement action with an owner and a target date. Test results are available to customers on request to support their own resilience mapping.

12. Exceptions

Any planned departure from this policy, such as temporarily reducing redundancy during a migration, must be approved in advance by a director, recorded with its reason and duration, and reversed by the agreed date.

13. Related Documents

- Operational Resilience Statement
- Incident Response Plan
- Information Security Policy
- Change Management Policy
- Supplier Risk Management Policy
- Sub-Processor List
- Service Levels & Support Policy
- Exit Plan & Data Return Policy

14. Document Control

Item	Detail
Owner	Security Officer
Approved by	The directors of Subcontext Ltd
Review	At least annually, and after any disaster or significant infrastructure change
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.