Operational Resilience Statement

1. Purpose

Many of our customers are regulated firms that must keep their important business services within impact tolerances, including under the FCA's and PRA's operational resilience rules. This statement sets out how the Subcontext platform is built and run to stay available, and what customers need to know to map Subcontext into their own resilience frameworks.

It is written for the people in a customer firm who map important business services, set impact tolerances and run scenario testing under the FCA's operational resilience rules (PS21/3, SYSC 15A) and the PRA's equivalent requirements, and who manage outsourcing and third-party risk under SYSC 8 and SS2/21.

2. Definitions

Term	Meaning
Important business service	A service a customer provides to its own clients which, if disrupted, could cause intolerable harm to those clients or to market integrity. The customer decides which of its services are important
Impact tolerance	The maximum disruption a customer can tolerate for an important business service, set by the customer
Recovery Time Objective (RTO)	The target time to restore the Subcontext platform after a disruption
Recovery Point Objective (RPO)	The maximum period of data that could be lost in a disruption
Severe but plausible scenario	A disruption that is serious but realistic, used to test whether a service can stay within tolerance

3. Responsibilities

Operational resilience is shared between Subcontext and its customers.

Subcontext	Customer
Designs, runs, monitors and recovers the Subcontext platform	Identifies its important business services and sets their impact tolerances
Sets and tests the platform's recovery objectives	Decides how Subcontext fits into those services, and plans for its unavailability
Tells customers about incidents that affect them	Keeps its own team able to take over conversations, and its own fallback processes
Manages its own suppliers and sub-processors	Maintains its own exit plan, using our Exit Plan & Data Return Policy

Subcontext	Customer
Provides information for the customer's mapping and testing	Includes Subcontext in its own scenario testing where relevant

Within Subcontext, the directors are accountable for operational resilience. The Security Officer owns the incident response and recovery procedures, and the engineering team runs the platform day to day.

4. Services Customers Rely On

Subcontext typically supports these customer business services:

- Customer conversations on web chat, SMS, WhatsApp, email and voice
- Insurance quoting, applications and policy servicing
- Collection of payments and Direct Debit mandates
- Lead management, follow-up and outreach

Mapping the dependencies

Customer service	Main Subcontext components	Main external dependencies
Web chat conversations	Application servers, database, AI agents	AWS, AWS Bedrock
SMS, WhatsApp and voice conversations	Application servers, database, AI agents	AWS, AWS Bedrock, Twilio, Meta (WhatsApp), OpenAI and ElevenLabs (voice)
Email conversations	Application servers, database, inbound email processing	AWS
Quoting, applications and policy servicing	Application servers, database, integrations	AWS, Iress (market quoting, where enabled), Data8 (address lookup, where used)
Payments and Direct Debit mandates	Application servers, database	Stripe
Lead management and outreach	Application servers, background workers, database	AWS, Twilio

The full list of sub-processors, with their locations and purposes, is in the Sub-Processor List.

5. How the Platform Is Built for Resilience

- Hosted on Amazon Web Services in the EU (Ireland) region, with application servers spread across two availability zones
- Application and background-processing servers scale automatically with demand, and unhealthy servers are replaced automatically
- A managed PostgreSQL database with automated backups and point-in-time recovery for the last 7 days, encrypted at rest
- A web application firewall and per-IP rate limits protect the platform from malicious and excessive traffic
- Every release is an immutable, versioned image, so a previous release can be restored quickly

- Releases are approved before deployment, and database migrations run before the new release goes live
- AI inference runs on models hosted in AWS Bedrock in the UK and EU, and agents can be moved to an alternative model

In more detail:

Area	Arrangement
Application	Web and background workers run as containers in private network subnets across two availability zones, scaling between 1 and 6 tasks each. Health checks replace unhealthy tasks
Database	Managed PostgreSQL with deletion protection, required TLS and encryption at rest. The production database runs in a single availability zone, and is recovered from automated backups or point-in-time recovery if that zone is lost
Cache	Managed Redis, encrypted at rest and in transit, running as a single node
File storage	Amazon S3 with server-side encryption
Infrastructure	Defined as code and version-controlled, so the environment can be rebuilt consistently
Releases	A failed database migration stops the release before it goes live. Previous images are kept so a release can be rolled back
Protection	AWS WAF managed rule sets and a per-IP rate limit at the edge, and application rate limits per IP, per sign-in and per feature

6. Recovery Objectives

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

Customers should compare these targets with the impact tolerances of the services that depend on Subcontext. Where a tolerance is shorter than our RTO, the customer's plan should include a fallback, for example handing conversations to its own team, as described in section 7.

7. Severe but Plausible Scenarios

We plan for, and periodically review our response to, these scenarios:

Scenario	Response
Loss of an AWS availability zone	Application servers continue in the remaining zone; the database is recovered by RDS or restored from backup
AI model provider outage	Agents are moved to an alternative model; conversations can be handed to the customer's own team
Messaging or telephony provider outage	Conversations continue on unaffected channels; the customer's team is told which channels are affected

Scenario	Response
Payment provider outage	Payment steps are paused and resumed when the provider recovers; no card or bank details are lost
Cyber incident	Handled under the Incident Response Plan, including containment, recovery from clean backups and customer notification
Loss of key personnel	Infrastructure is defined as code and the platform is documented, so others can operate and restore it
Faulty release	The release is rolled back to the previous image; a failed database migration stops the release before it goes live
Loss of the AWS EU (Ireland) region	The platform is rebuilt from infrastructure code and restored from backups; recovery would take longer than the RTO, as the platform does not run in a second region

The last scenario is included so customers can plan for it. We do not currently run the platform in a second AWS region, and we will review whether multi-region recovery is needed as the platform grows.

What a customer can do while Subcontext is unavailable

- Hand conversations to its own team, using its usual phone and email channels
- Continue to use its own systems of record, which remain the source of truth for policies and payments
- Use its own Stripe account and records for payment questions
- Rely on its most recent self-service exports of leads and its insurance book

8. Incident Response and Recovery Procedure

1. **Detect:** monitoring, alarms or a customer report identify a disruption.
2. **Assess:** the incident lead decides the severity under the Incident Response Plan and which customer services are affected.
3. **Communicate:** affected customers are told by email, with updates until the incident is resolved.
4. **Contain and recover:** the platform is restored by the quickest safe route, such as replacing failed tasks, rolling back a release, failing over to an alternative AI model, or restoring the database.
5. **Protect the platform:** where needed, the platform is put into maintenance mode with a holding page, or an emergency lockdown is applied at the load balancer.
6. **Review:** after a Critical incident, we review what happened, what it meant for customers, and what we will change, and share a summary with affected customers.

9. Monitoring and Communication

- Application errors and performance are monitored continuously, with alerts to the engineering team
- Infrastructure alarms cover the database, cache, load balancer and application servers
- Customers affected by an incident are told by email, with updates until it is resolved
- Planned maintenance is announced in advance where practicable, and a holding page is shown while it runs

Error alerts are raised the first time a new error occurs, when a resolved error comes back, and when error counts rise. A daily heartbeat message to the team confirms that scheduled processing is running.

10. Third-Party Dependencies

Our suppliers are selected, contracted and monitored under the Supplier Risk Management Policy. Critical suppliers are reviewed at least annually, including their resilience. Where a supplier provides a single channel, such as SMS or voice, its outage affects that channel only, and we keep the other channels running.

11. Testing

Backup restoration and the scenarios above are tested at least annually, and after significant changes to the platform.

Each test records the scenario, what was done, how long recovery took against the RTO and RPO, any problems found, and the actions taken to fix them. Actions are tracked to completion by the Security Officer.

12. Supporting Your Resilience Mapping

On request, we provide the information customers need to map Subcontext into their own resilience frameworks, including dependencies, recovery objectives and the results of our testing.

We will also, on reasonable request:

- Answer due diligence and resilience questionnaires
- Take part in a customer's scenario testing where Subcontext is part of the scenario
- Tell customers about material changes to how the platform is hosted or recovered

13. Related Documents

- Business Continuity & Disaster Recovery Policy
- Incident Response Plan
- Service Levels & Support Policy
- Supplier Risk Management Policy
- Sub-Processor List
- Exit Plan & Data Return Policy
- Change Management Policy

14. Document Control

Item	Detail
Owner	Security Officer
Approved by	The directors of Subcontext Ltd
Review	At least annually, and after significant changes to the platform

Item	Detail
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.