



1. Overview

This Business Continuity and Disaster Recovery document describes the mechanisms, policies and procedures used by Renato Software Ltd, trading as Senso.cloud, to continue operating and recover services in the event of disruption to the platform or its supporting infrastructure.

Senso.cloud is hosted on Microsoft Azure. In most cloud infrastructure disruption scenarios, Microsoft is responsible for restoring the underlying cloud platform and maintaining the resilience of Azure services. Senso.cloud is designed to use the resilience, backup, deployment and recovery capabilities available within that cloud environment.

2. Scope and objectives

The purpose of this document is to describe the business continuity and disaster recovery arrangements applicable to the Senso.cloud platform. The objective is to reduce the likelihood of a dependent service failure affecting the whole platform, protect customer data from accidental corruption or deletion, and provide practical recovery options for service disruption scenarios.

RSL considers the following primary disruption scenarios when planning business continuity and disaster recovery activities:

1. Application failure
2. Data corruption and restoration
3. Network outage
4. Dependent service failure
5. Region-wide service disruption

3. Platform design and service separation

Senso.cloud is designed using a framework of modular components that separates responsibilities across key service levels, including Client, Modules and Portal. This separation reduces the likelihood that a disruption in one dependent service will bring down the entire application.

The platform uses different update and recovery methods depending on the part of the service affected:

Platform area	Continuity approach
Client	Client source code and related components are backed up daily with check-in procedures. Backups are stored in secure, geo-located locations to support recovery in the event of catastrophic failure.
Modules	Module source code and related components are backed up daily with check-in procedures. The modular design helps isolate module-level issues from the wider platform where possible.
Portal	Portal updates use deployment slots, allowing changes to be validated in a staging slot before being swapped into production. If a production swap does not perform as expected, the previous production slot can be swapped back to restore the last known good site.

4. Application failure

Application failures may be recoverable or non-recoverable. Recoverable failures can often be mitigated through fault handling, alternate actions or rollback procedures. Non-recoverable failures may render a component or, in more severe circumstances, the application unavailable until recovery action is completed.

The following application failure responses are used where applicable:

- Some failures are addressed transparently by handling faults automatically and taking alternate action. For example, Azure services include automatic failover capabilities within the service itself.
- Where appropriate, the application may continue to handle user requests with reduced functionality, allowing the service to remain impaired but operational.
- More severe service disruptions may render the application unavailable until the issue is remediated, rolled back, or recovered using the appropriate continuity process.

5. Portal deployment and rollback

Deployment slots are used when updating the Senso.cloud portal. This approach allows RSL to validate changes before moving them into production and gives the team a rapid rollback mechanism if an update does not behave as expected.

When deployment slots are used:

- Application changes are validated in a staging deployment slot before being swapped with the production slot.
- After a swap, the slot containing the previously staged application holds the previous production version.
- If the changes swapped into production are not as expected, the same swap can be performed again to return to the last known good site.

6. Data corruption and restoration

Data corruption or deletion is typically caused by an application defect, upgrade or maintenance issue, or human error. Such failures are application-specific and cannot always be detected or mitigated automatically by the infrastructure.

Examples of data-related disruption scenarios include:

- A datacentre outage caused by a natural disaster, requiring geo-redundancy and recovery using an alternate datacentre or paired region.
- Upgrade or maintenance errors during planned upgrades or maintenance to an application or database, which may require rapid rollback to a prior database state.

Database backups are an essential part of the Senso.cloud business continuity and disaster recovery strategy because they protect customer data from accidental corruption or deletion.

7. Backup strategy and point-in-time restore

RSL creates full, differential and transaction backups for the purpose of point-in-time restore (PITR). Backup routines vary by database and depend on compute size, database activity and the relevant area of the platform.

Backup component	Frequency / approach
Transaction log backups	Generally occur every 5-10 minutes.
Other transaction routines	Generally run every 2 hours.
Database-specific backup routines	Vary from 5 minutes to 2 hours depending on the database and area of the platform.
Differential backups	Generally occur every 12 hours, with frequency based on compute size and database activity.
Backup location and replication	Backups are stored in containers replicated to a paired datacentre to protect against a datacentre outage in the same geolocation as the hosted data.

These backups are used to support the following restore scenarios:

- Restore a database to a point in time within the applicable retention period.
- Restore a deleted database to the time it was deleted or to another available point in time within the retention period.
- Restore from a paired datacentre in the event of a geographic disaster.

8. Network outage

If parts of the Azure network are inaccessible, customers may be unable to access the Senso.cloud application or their data. In those circumstances, reducing functionality may not be a viable option. The available responses are application downtime while Microsoft restores the network service, or failover to an alternate region where available and appropriate.

Scenario	Recovery approach
Application downtime	Service resumes after Microsoft restores the affected network service or underlying Azure dependency.
Failover to alternate region	Where regional failover capability is available and appropriate, service recovery may be performed using an alternate region or paired datacentre.

9. Region-wide service disruption

Many failures are manageable within the same Azure region. In the unlikely event of a region-wide service disruption, locally redundant copies of data may not be available. Senso.cloud uses geo-replication so that additional copies of data are available in another region.

If Microsoft declares a region lost, Azure may remap DNS entries to the secondary region as part of the regional recovery process. Recovery timing will depend on the nature of the Microsoft Azure incident, the affected services, customer tenancy location, and the availability of the relevant paired region or recovery target.

For Australian tenancies, the Senso.cloud Australia privacy addendum states that Australian customer Service Data is provisioned and stored in Microsoft Azure Australia regions and that backups and disaster recovery replicas for Australian tenancies are retained within Australia.

10. Dependent service failure

The Senso.cloud platform relies on multiple components, including client software, modules, the web portal, databases and supporting Azure services. The modular service design is intended to reduce the likelihood that a dependent service failure will affect the full platform.

Where a dependent service is impaired, RSL will assess the affected platform area, determine whether service can continue in a reduced or impaired state, and apply the appropriate recovery process, which may include rollback, restoration from backup, service restart, vendor escalation, or failover where available.

11. Complete catastrophic failure

In the event of a complete catastrophic failure affecting every datacentre used by RSL, RSL can reassemble the Senso.cloud framework once a new datacentre is available. The source recovery position states that this reassembly can be completed within approximately 24-48 hours once a suitable datacentre is available.

In such a catastrophic scenario, historic data may not be available immediately or at all, depending on the nature of the failure and the availability of replicated backups. Once the framework is operational, new data would begin logging from the point the platform resumes operation.

12. Responsibilities and dependencies

Business continuity and disaster recovery depend on cooperation between RSL, Microsoft Azure and, where relevant, the customer. Microsoft is responsible for the resilience and recovery of the underlying Azure cloud infrastructure. RSL is responsible for Senso.cloud application recovery, backup and restore processes, deployment rollback and customer communications where service-impacting incidents occur. Customers are responsible for their own local networks, devices, identity systems, internet connectivity and any customer-managed configuration required to use the service.

13. Communication and support

Customer support and incident handling are provided through the Senso.cloud support process and Support Service SLA. In the event of a service-impacting incident, RSL will triage the incident, assess the affected platform area, apply the appropriate recovery action and provide updates in accordance with the applicable support process.