

# DISASTER RECOVERY POLICY

## PURPOSE

This document defines the procedures for restoring the FDD platform to full working order in the event of a disaster.

---

## DETAILS

### Platform Availability

The FDD platform is deployed across two AWS Availability Zones in the Sydney region. All storage is replicated across at least two Availability Zones.

If one Availability Zone becomes unavailable, the platform is designed to continue operating at full functionality with no noticeable impact on end users.

---

## LOSING A REGION

AWS provides a standard uptime guarantee of **99.99%**.

If the AWS Sydney region becomes unavailable, the FDD platform will be commissioned in another AWS region using the same deployment tools and procedures used for:

- Releasing new software versions
- Managing infrastructure configuration

Data will be restored from the latest available backups (as documented in the Database Backup and Restore procedures).

The **target Recovery Time Objective (RTO)** for full platform restoration is **72 hours**.

---

## RECOVERY POINT OBJECTIVES (RPO)

The Recovery Point Objective depends on the type of data.

- **Full data backups** are performed every **24 hours**.
- The most important data—such as **mechanical equipment data**—is cached on the data collection device.

- In the event of a disaster (e.g., region loss), data will be restored first from the database backup.
- Any data collected after the backup snapshot will be retrieved from the data collection device.

Additional data sources will be rebuilt as follows:

- **Meter data** will be rebuilt from import files.
- **Weather data** will be rebuilt from external weather data APIs.

Secondary data, such as **user and group records**, will have a maximum RPO of **24 hours**.