

# Open Contracting Partnership Backup Testing

21 Oct 2024

This document describes the backup testing process that Dogsboddy Technology does for Open Contracting Partnership.

## Schedule

All backups will be tested every six months.

We typically do this around 10th March and September each year.

## Overview

Backup testing typically falls into checking that they are there and look good...

- Is there the correct number of backups? How far do they go back?
- Do the backups store the files we think they do?
- Do (MySQL) database backups finish ok? Does the final line say "dump complete"?
- Are there any other files or databases on the server that need backing up?
- Does the system have enough disk space to uncompress a backup?

## General Instructions for all servers

- Are all of the backups defined in the maintenance plan present

Checking new sites haven't been added and missed from backups. If they're intentionally excluded that's OK but has this been agreed?

- Are all sites set-up on the webserver being backed up?
- Are all databases in the database server being backed up?

Making sure we can restore without the server

- Are the credentials and location details needed to access the backups recorded in the wiki?
- Are the backups appropriately named that you'd know which ones you need without access to the server?

Making sure the customer isn't over paying for backups - falls more into the resource review if we say we are maintaining their S3 buckets

- Is a retention period set for backups?

# Specific Instructions for some servers

## Redmine

The Redmine server backups are stored in AWS.

Follow normal steps [General Instructions for all servers](#).

Backups are in AWS S3.

- Get the login credentials from the server
- AWS API access is quite locked down and I believe only James has full access and so cli commands need to be very specific e.g. regular `aws s3 ls` won't work due to insufficient permissions and so in the below example, the exact bucket is stated

```
cat /home/sysadmin-tools/aws-settings.local
AWS_ACCESS_KEY_ID="redacted"
  AWS_SECRET_ACCESS_KEY="redacted" # REMEMBER TO ADD SPACES
BEFORE RUNNING THIS, TO PREVENT IT BEING WRITTEN TO BASH HISTORY
AWS_DEFAULT_REGION="eu-west-2"
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION

mkdir backup-testing
cd backup-testing

aws s3 ls s3://ocp-redmine-backup/
* Download a site backup and a database backup and test as normal
aws s3 ls s3://ocp-redmine-backup/database/
aws s3 cp s3://ocp-redmine-backup/database/$BACKUP .
zless $BACKUP.sql.gz | tail
# Look for Dump completed on
aws s3 ls s3://ocp-redmine-backup/site/
aws s3 cp s3://ocp-redmine-backup/site/$BACKUP .
tar -xzvf $BACKUP
# Compare against site files on server /home/redmine/public_html

cd
rm -rf backup-testing
REMOVE the backups you have downloaded!
```

# PostgreSQL

On [ocp23.open-contracting.org](https://ocp23.open-contracting.org) we are taking PostgreSQL backups using **pgBackRest**. The backups are managed by pgbackrest and stored in S3.

We are taking weekly "full" backups and daily incremental. Backups are kept for 4 weeks.

This is defined in `/etc/pgbackrest/pgbackrest.conf` and `/etc/cron.d/postgres_backups`

You can confirm backups exist with:

```
su - postgres
pgbackrest info
```

We can't test Kingfisher DBs because they are ~2TB!

Other PostgreSQL servers:

- **ocp13** - Registry. No current backups. Backup script in place and in cron - check s3 for files and possibly follow ocp23 process above (untested).
- **ocp22** - Credere. Dev server no DB backups.
- **ocp27** - Supposed to be the same as ocp13 but still under development. When live the backup testing process will need fleshing out.

# MySQL

ocp21 hosts MySQL used by their WordPress site.

Backups are in AWS S3.

- Get the login credentials from the server
- AWS API access is quite locked down and I believe only James has full access and so cli commands need to be very specific e.g. regular `aws s3 ls` won't work due to insufficient permissions and so in the below example, the exact bucket is stated

```
cat /home/sysadmin-tools/aws-settings.local
AWS_ACCESS_KEY_ID="redacted"
  AWS_SECRET_ACCESS_KEY="redacted" # REMEMBER TO ADD SPACES
BEFORE RUNNING THIS, TO PREVENT IT BEING WRITTEN TO BASH HISTORY
AWS_DEFAULT_REGION="eu-west-2"
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION

mkdir backup-testing
cd backup-testing

aws s3 ls s3://ocp-coalition-backup/
* Download a site backup and a database backup and test as normal
aws s3 ls s3://ocp-coalition-backup/database/
aws s3 cp s3://ocp-redmine-backup/database/$BACKUP .
zless $BACKUP.sql.gz | tail
# Look for Dump completed on
aws s3 ls s3://ocp-coalition-backup/site/
aws s3 cp s3://ocp-redmine-backup/site/$BACKUP .
tar -xzvf $BACKUP
# Compare against site files on server /home/redmine/public_html

cd
rm -rf backup-testing
REMOVE the backups you have downloaded!
```

## Linode

Each Linode server should have Backups enabled.

Linode takes the following:

Three backup slots are executed and rotated automatically: a daily backup, a 2-7 day old backup, and an 8-14 day old backup. Plans are priced according to the Linode plan selected above.

## Hetzner

The Hetzner servers do not have file backups. Currently these are not needed, all site data should be stored directly in GitHub.