# Incident Report - CUBE-DS-1477 drive failure

## Summary

17 April 2016 - During routine maintenance to attach a network storage device to the server to allow you to manage your own backups and increase data retention of automated control panel backups the server failed to reboot.

Ultimately this led to a new server being built and data being restored from backup.

In short

- Data lost - none
- Mail lost - none
- Total downtime 8 working hours (24hrs total)

## Incident Timeline

**2245** on 17 April server CUBE-DS-1477 failed to reboot during routine maintenance.

The engineers identified the cause as corrupt metadata on the system drive and were unable to correct the issue.

**2255** the engineers attempted to restore the most recent snapshot from the primary backup routine.

**0435** on 18 April  engineers reported attempts failed and moved to restore the next most recent snapshot..

**0845** engineers reported the second and third snapshots failed to restore due to filesystem errors.

This was entirely unexpected as the snapshots undergo an integrity check immediately after being taken which did not identify any errors.

At this point the engineers began building a new server to restore the secondary backups.

**1225** the new server was built and ready for required software and configuration to be applied.

**1515** the new server was online, the secondary backup restoration process was initiated and estimated to take 3 hours.

No progress indication is provided during the restore process to indicate time remaining and this took far longer than expected..

**0035** on 19 April the secondary backup restoration process completed. The Backup Manager reported a number of email account passwords were changed during the process.

No advance warning of this was given. Nor did the Backup Manager report it also changed over 1000 control panel, database and other system passwords.

This behaviour was caused by an undocumented change to the password security policy introduced by a micro update to the server control panel software.

When the majority of subscriptions were created in the system the security policy was to enforce *Normal* strength passwords. The micro update changed this to *Very Secure*.

When the restoration process found original passwords that didn't meet the *Very Secure* policy it changed them.

This prevented affected users sending and receiving mail as their devices did not have the new passwords. Incoming mail was accepted so no data was lost.

The password change also prevented some websites communicating with their database.

All affected passwords were updated manually and the changes communicated to affected users.

**1420** a report was made that the mailserver was applying the wrong timestamp to new mail. This was corrected at 1500.

# Data continuity

## Files

The restored backup was from Friday 15 April 2016 at 0015.

This means any changes made to websites during 15 April were not automatically restored.

There was no data loss as the tertiary backup procedures run daily and store data in multiple off site locations.

These were used to manually restore data for a single client that reported a change made on 15 April was not restored to their website.

### Email

Clients using their own on site mailserver or cloud email including Google Apps or Microsoft Exchange were unaffected.

Clients who use the basic email facility included with website hosting as their primary email system were affected.

Delivery of email sent during the time the server was inaccessible (Monday 18 April 2016) meant messages would fail to be delivered and would typically be retried for 24hrs before being returned to sender.

Any messages sent during working hours on Monday 18 April should have been successfully delivered within the 24hr period as the mailserver came back online well in advance of 0800 the following day.

One client found historic email failed to display in the Webmail control panel. The mail was on the server and forcing Webmail to rescan the inbox restored access to the mail.

## Backup procedure overview

The backup systems in place are

- Primary
    - Daily snapshot of server, data stored on site and retained for 3 days
- Secondary
    - Weekly subscriptions backup by control panel software, data stored on site and retained for 30 days
- Tertiary
    - Nightly website files backup to Amazon EU datacentre, files retained for 14 days
    - Nightly database dumps, stored both on and off site, files retained for 14 days

### Post incident analysis

- Communications - ability to communicate with clients during the incident was hampered because many use the email system that went down during the incident.
    - A secondary email list will be created with alternative email addresses email addresses where available
    - A Service Status channel has been established on a third party system to ensure it is accessible if any of our systems go down to provide details of what is happening. It can be accessed at https://goo.gl/d7lHg2 and links will be in all future email signatures so you always know where to find it
- Following discussion with the engineers minor adjustments will be made to the backup procedures. Neither of these would have allowed faster recovery if they had been in place at the time of the incident

    - Primary - will remain the same with the addition of a 6 monthly restore test to verify data integrity
    - Secondary - the planned addition of the attached network storage device (see summary) will go ahead
    - Tertiary - an additional on site backup of individual website files and databases will be introduced alongside the existing off site service

- Websites that use heavy caching were unaffected by database password change issue. This meant they were accessible up to 12hrs before those that required a manual password update. Details of how to implement such caching will be issued to those who could benefit will be issued in due course.

- When the new network storage device (see summary) is in place you will be offered the ability to manage your own backups. This means in the event of a catastrophic failure you have your own copy of all the files required to host your website with another provider

## Can this be more robust?

Options to develop a more robust solution are noted below, the first is strongly recommended, the second is unlikely to be financially viable, the third is recommended as it can keep your website online during an incident and brings additional benefits

- Separate email and web hosting - if you use the email accounts provided with your hosting package you should consider moving to a cloud solution. This introduces a layer of redundancy as the system is separate to your website hosting and brings features required for modern business such as synchronising across multiple devices, enterprise level anti-spam and backups held by the datacentre rather than your own machine. Two options are available with further detail on request.

- Colocation/Failover - it is possible to operate two servers with the secondary server mirroring the first. In the event of an issue traffic can be routed to the secondary server however this is is a complex setup.

  You need to keep the backup server update with any changes made to the main server, you need a method to detect failure of the first server and re-route traffic to the backup and you need to sync any changes made during an incident on the second server back to the first once the problem is resolved eg. online store orders or changes made with content management.

  All of this means the cost of implementing two servers like this is not simply double the cost of the first, it’s at least 3 or 4 times more expensive.

- Content Distribution Network - the leading CDN provider offers a service that will serve the most popular pages of your website from it’s cache if your primary server

fails to respond.  This will keep you partially online in the event of an incident.  The cost of implementing this is usually just a few hours to cover configuration.