Incident Report #201712 Reported by: Marcos Rocha Incident Manager: Sam McCullough	Incident Date: Dec. 28, 2017 PST Report Date: Dec. 28, 2017 PST

Outage Incident Report

Summary

A virtual hard disk failed during a routine operation to expand the size. This operation failure render the sites within the hosting environment offline.

Timeline

- 11:33AM - Operation to expand size of virtual hard disk of hosting environment failed
- 11:33AM - All shared LMS sites in the environment went offline
- 11:35AM - Operations team begins recovery effort on hard disk
- 11:39AM - Support sends out notification to affected clients
- 11:48AM - Operations team fails at initial attempt to recover hard disk
- 11:52AM - Operations team enacts disaster recovery, beings to download recovery images
- 1:10PM - First of site images completes download
- 1:20PM - First site is fully restored
- 1:58PM - Second site images completes download
- 2:07PM - Second site is fully restored
- 2:20PM - Third site image completes download
- 2:31PM - Third site is fully restored
- 2:49PM - Original hard disk is recovered and used
- 2:51PM - All sites operational

Statistical evaluations

Number of Incidents	**Recovery/Resolution time**	**Impact on Uptime SLA**
1	3.5 Hours	3.5 Hours

Resolution and recovery

In preparation for the introduction of new client sites to this shared environment, the Operations team was expanding the virtual hard disk. An operational error during this procedure lead to the disk to crash.

The operations team made an attempt to recover the disk initial, but it failed. The team enacted Disaster Recovery protocols to restore LMS functionality. The team was able to recover three of the five sites on the environment via Disaster Recovery, before the original disk came back online.

Corrective and Preventative Measures

The Operations team is looking to make improvements to the disaster recovery protocol. The goal is to reduce the time required to recover (download) the backup and restore it into the hosting environment. Currently Lambda utilizes a service called Duplicity. This stores images outside of Amazon Web Services.

This service creates a bottleneck where Lambda is required to download the image files to use. This can take a significant amount of time. Some client sites are in the are of 1 Terabyte in size. It could take up to 16-24 hours depending on connection to recover this image file.

The plan is to upgrade to a Amazon Web Service that stores the images outside the affected hosting environment, but within AWS. This can lead to a significantly reduced recovery time. Internal estimates at Lambda suggest we could see Disaster Recovery times of 30 minutes regardless of image file size.