

Table of Contents

Ceph Block Store incident and associated service incidents.....	1
Description.....	1
Impact.....	1
What was affected (Grid-facing services).....	1
What was not affected (Grid-facing services).....	1
Time line of the incident.....	1
Analysis.....	2
Root Cause.....	2
Follow up.....	2
Links.....	2

Ceph Block Store incident and associated service incidents

Description

As noted on [OTG:0054944](#), on Thursday 20th February 2020, around 10am, some parts of the Ceph infrastructure went down.

Ceph is used as the backing store for VMs providing attached disks and shared filesystems to many services - consequently several CERN site and VO services were affected.

Impact

What was affected (Grid-facing services)

- The CERN Batch Service was unavailable for job submission. Longer running jobs eventually failed due to losing touch with the batch `schedds` for too long. ([OTG:0054948](#))
- The CVMFS stratum-0 was unavailable for some hours ([OTG:005496](#))
- HammerCloud was degraded ([OTG:0054991](#))
- The WLCG monitoring infrastructure was degraded, recovering the backlog some hours later ([OTG:0054946](#))
- Multiple VO services, both local and those supporting the experiments' distributed production and analysis were unavailable.
 - ◆ Impact to VO services was summarised in the weekly WLCG operations meeting.
- Other non-Grid services were affected and can be found on the [IT-SSB Service Incident page](#)
 - ◆ Notably `lxplus`, `AIADM`, `acron`, some AFS volumes, Gitlab and file access in Indico.

What was not affected (Grid-facing services)

- The EOS Grid services were not affected
- The BDII, MyProxy and VOMS services were not affected
- `aiadm-homeless`, a server for administering agile infrastructure machines with minimum external dependencies

Time line of the incident

LXBATC

- 10:10: `lxbatch` submissions blocked
- 18:00: `lxbatch` job submissions open again
- 19:30: `lxbatch` back to normal levels

Analysis

- In most cases, the direct cause of secondary service degradation was due to a direct dependency on Ceph.
- In a few cases (e.g. Indico file access) the effect was tertiary, due to affected AFS volumes hosted on VMs.
- The services' dependencies (on Ceph) were reviewed by IT management at the C5 Service Meeting [\(PDF\)](#) and generally found to be reasonable and a required component of the service delivery.

Root Cause

- The root cause of the Ceph outage was determined after one week of investigations.
- Versions of the LZ4 compression library earlier than 1.8.2 have an encoding bug when compressing fragmented memory buffers.
 - ◆ CentOS 7 ships version 1.7.5.
 - ◆ The outage occurred when a central cluster metadata file (the osdmap) was incorrectly compressed, corrupting 4 bits, which then caused the ceph OSD daemons to throw a CRC exception and exit.
- On Tuesday Feb 25th the LZ4 library was first suspected and was subsequently disabled. LZ4 will remain disabled.
- Block storage data previously compressed is not expected to be at risk, because it is normally compressed from a consolidated memory buffer.
- We are highly confident in the root cause analysis:
 - ◆ We are able to reproduce the exact bit corruption in a new Ceph unit test.
 - ◆ Ceph upstream has developed a workaround for the buggy LZ4, and is getting input from the LZ4 developer.
- Links:
 - ◆ Post mortem presentation (WLCG Ops Coord): <https://codimd.web.cern.ch/p/SkII0nc4I>
 - ◆ Technical post mortem presentation (IT-ASDF): <https://codimd.web.cern.ch/p/HJGnq3HVI>
 - ◆ Ceph bug tracker: <https://tracker.ceph.com/issues/39525>
 - ◆ Github pull request: <https://github.com/ceph/ceph/pull/33584>

Follow up

Actions:

- Reconfigure aiadm servers, which provide hosts for infrastructure tools, to have no dependencies on external storage. SSB at OTG:0055070 [\(PDF\)](#)
- Ceph: disable LZ4 compression until an upstream fix comes available.

Links

Ceph OTG:

- Ceph Block Storage degraded OTG:0054944 [\(PDF\)](#)

Linked OTGs:

This topic: LCG > CERNProdIncident200220

Topic revision: r5 - 2020-03-02 - DanielVanDerSter



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)