



## **Incident report: 2017-10-16**

### ***Description***

A storage cluster issue caused blocked I/O for disk operations for several customer VMs. The fault was first reported by a customer at 05:57 on the morning of Monday 10 October. The root cause was rapidly diagnosed and corrective measures were taken within minutes of the incident being reported. The problem had occurred some 5 hours prior to the first report, with the same blocked I/Os preventing the automated monitoring process from raising an appropriate alarm. A related issue became apparent a few hours later that prevented the management of VMs on one of our compute nodes. This related issue necessitated a reboot of the compute node in question, resulting in the restarting of certain customer VMs.

### ***Analysis***

The Ceph cluster utilized for the storage backing isoho.st VMs is a redundant multi-terabyte array. The space is subdivided into several 4TB Object Storage Daemons (OSDs), each of which receives a portion of the data. The array is designed to withstand the failure of one or more OSDs, provided all the OSDs reside within the same storage node. In some cases, the array can survive multiple distributed failures, but usually not without affecting a subset of customer VMs.

Four of these OSDs (three in one storage node and one in another) simultaneously experienced a crash due to a software assertion failing. These software assertions are designed to protect the integrity of the array in the event of failure scenarios that might corrupt the data in the array. A software bug in the version the Ceph code we are running is the most probable cause.

### ***Impact***

Data are stored across a subset of nodes in accordance with a hash distribution function. All disk reads/writes for those customer VMs accessing parts of their data contained exclusively within the 4 OSDs in question were affected for the duration of the incident.

One of our compute nodes was also fully reliant on data within these failed OSDs hanging the management of VMs on that node in such a way as to require a reboot. Additionally, our fault reporting infrastructure was similarly affected.

### ***Actions***

The crashed OSDs were brought back online and the storage array was allowed to re-sync. Later, a compute node was administratively rebooted in order to recover from the fault affecting the management of VMs.

In the short term, we will be reducing our reliance on our own storage infrastructure for the reporting of administrative alarms, with the intent of being able to respond to such a scenario more rapidly.

In the medium term, we will be scheduling upgrades of the underlying Ceph software. Such upgrades may necessitate some scheduled down time in the future.

**Please direct any questions or comments to [support@isoho.st](mailto:support@isoho.st)**