

Reason For Outage (RFO) report - SUNET Drive

Affected Services:

- S3-Storage hosted in Sto4
- Sunet Drive users with data in Sto4

Outage time:

- 2021-06-17 - 2021-10-09

Sequence of significant events in UTC time:

June 17th, 2021:	Reports of high error rate coming from storage cluster
June 18th:	Cluster unavailable due to high load, cluster was taken offline
June 19th:	Default restart and recovery of storage cluster failed
June 20th:	Development of customized restore solution was initiated
July 15th:	First individual files restored
August 17th:	First complete bucket restored
September 6th:	Restoration of data finalized, with ongoing verification
September 10th:	Incident closed

Root cause and corrective actions

A single customer overloaded the S3 storage cluster in Sto4 by uploading 680 million objects within the course of a few days. This overloaded the metadata-part of the S3-cluster, leading to its shutdown. This resulted in a state of the cluster, where metadata was stored on one side and files which have been split by the erasure encoding on the other side. The restoration required custom development by technical experts and took several weeks.

Affected customers were immediately referred to the S3 storage cluster in Sto3 to ensure continuous operations of services. This required customers to restore data from their own backups to Sto3. Customers who reportedly had no backups of data stored in Sto4 (i.e., unique data) were identified and restoration of their data was initiated.

Corrective actions:

Rate-limiting has been implemented on both Sto3 and Sto4 to prevent further incident and ensure quality of service. - **Done**

Implemented temporary mirroring of data from Sto3 to Sto2 via TSM - **Done**

Migration to Sto3<-->Sto4-mirroring when the cluster has been rebuilt and will offer this as an option to customers. - **Pending**

Documentation of services will be reviewed and will put extra focus on the importance of backups and disaster protection. - **Pending**

Other notes:

A detailed technical problem report has also been written and is available from the SUNET Drive project lead, Richard Freitag (freitag@sunet.se).