Science and Technology Facilities Council

Echo

A Ceph-backed storage service for particle physics data.

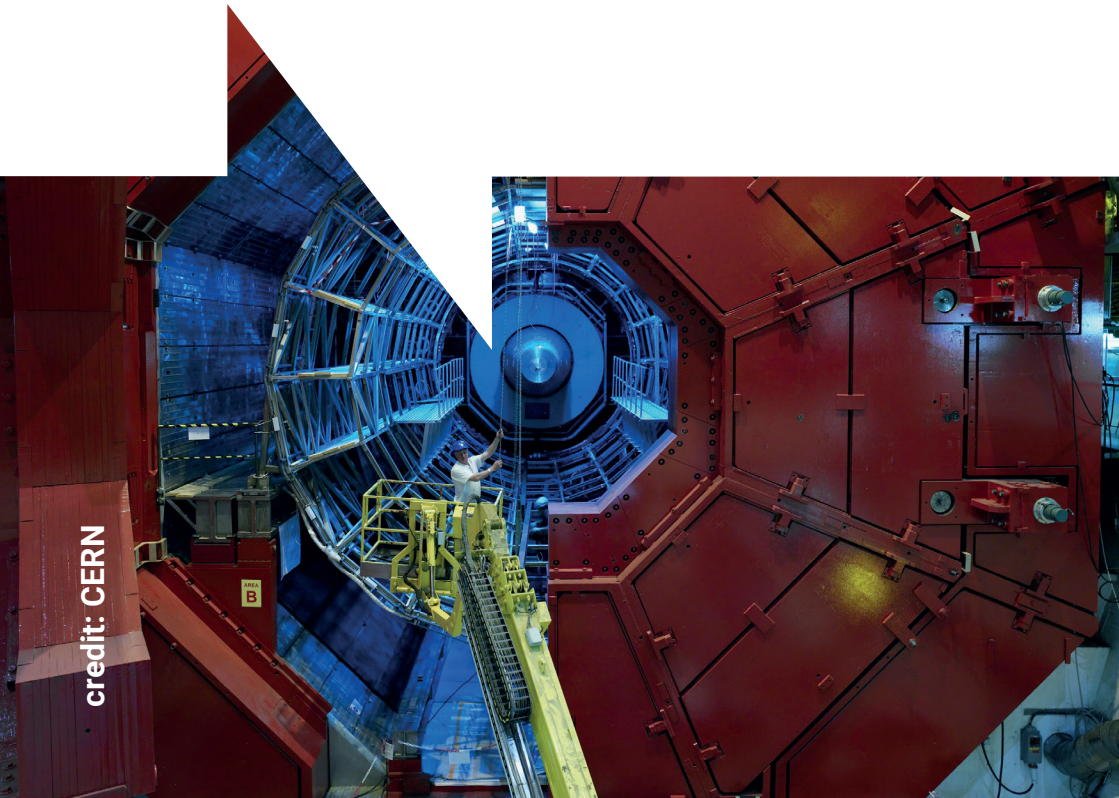

About Echo

Echo is the Ceph-backed storage service developed at STFC to meet the UK's data storage commitments to the Large Hadron Collider (LHC) experiments at CERN. The LHC experiments produce vast quantities of data which are stored across a worldwide grid of Tier 1 computer centres, including the UK Tier 1 centre at STFC's Rutherford Appleton Laboratory. Echo uses erasure-coding to protect the data it stores against disk failure while also reducing the storage space needed to implement such protection.

Successful external transfer throughput, by country of data origin

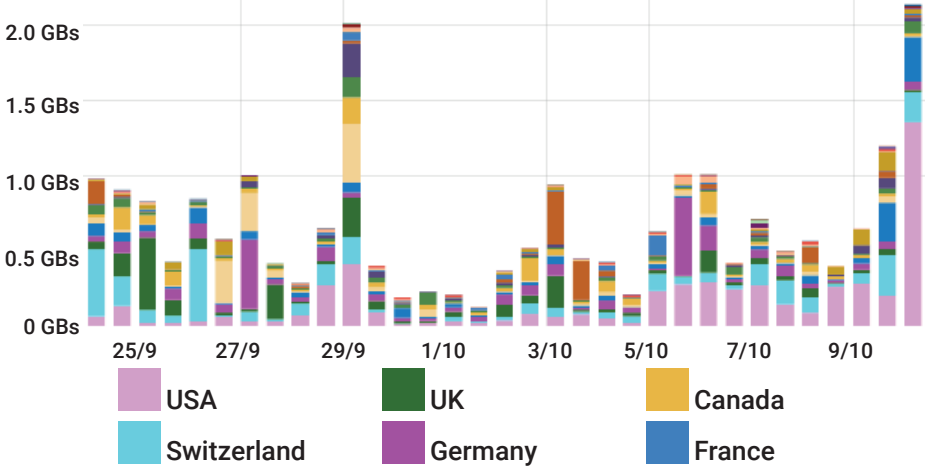

Impact studies

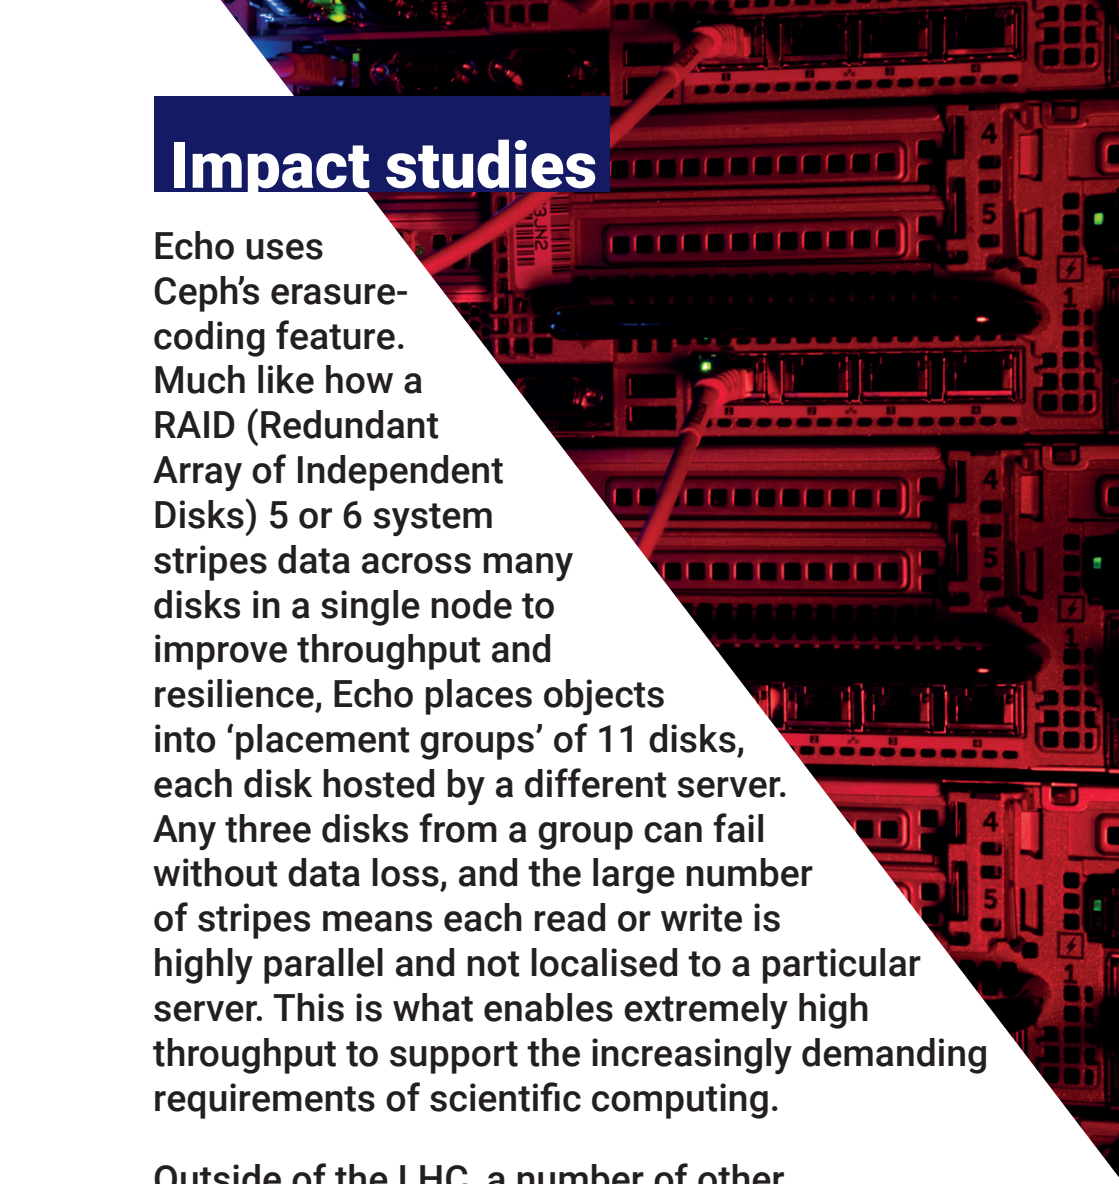

Echo uses Ceph's erasure-coding feature. Much like how a RAID (Redundant Array of Independent Disks) 5 or 6 system stripes data across many disks in a single node to improve throughput and resilience, Echo places objects into 'placement groups' of 11 disks, each disk hosted by a different server. Any three disks from a group can fail without data loss, and the large number of stripes means each read or write is highly parallel and not localised to a particular server. This is what enables extremely high throughput to support the increasingly demanding requirements of scientific computing.

Outside of the LHC, a number of other experiments have a small allocation of storage on Echo. These include Diamond Light Source, the UK's national synchrotron, and a prototype of the US-based Deep Underground Neutrino Experiment (DUNE).

50PB
raw storage capacity

over
1PB data
transferred daily

310
storage nodes

23PB
of LHC data stored

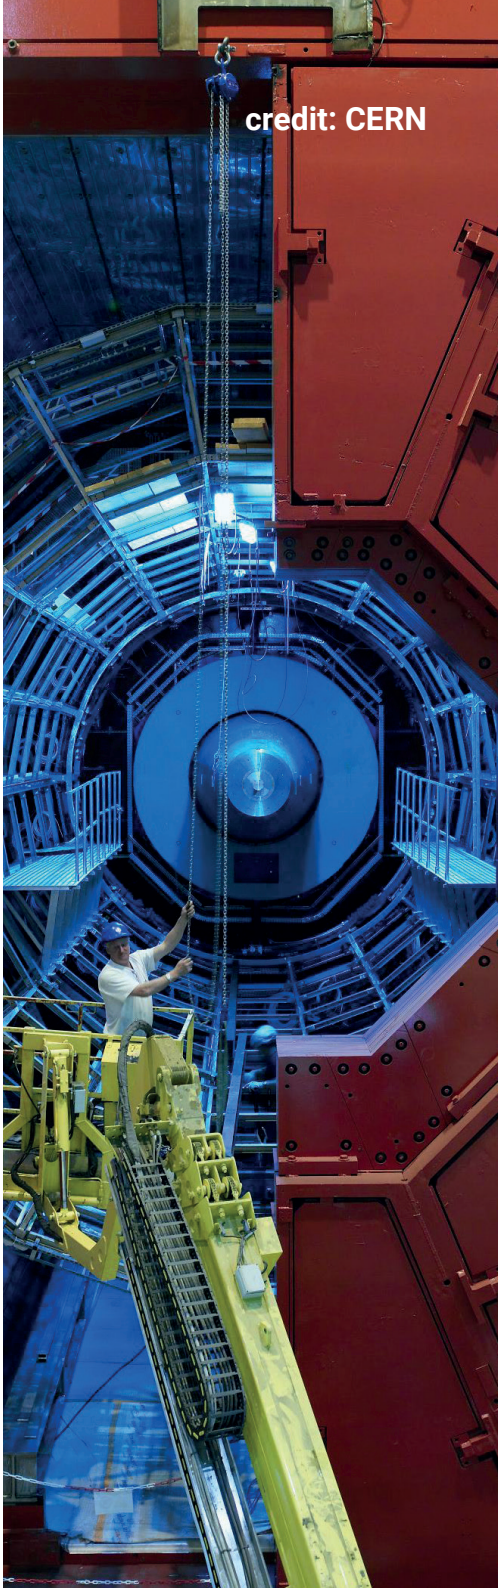
credit: CERN

GridPP
UK Computing for Particle Physics