# Large Scale File Migration

## INDUSTRY

Physical science research institute in accelerator physics and photon science

## RESULT

- 80% migration time savings
- Able to leverage existing storage, server hardware & network infrastructure investments
- Hands-off automation

## THE CHALLENGE

- Complex management and operational workloads with traditional tools
- Traditional tools inefficient in working with distributed storage
- Hardware upgrade not feasible as the two ends have different management and budget cycles

## THE SOLUTION

- Scale-out to leverage existing server hardware and network infrastructure
- Multi-level parallel processing for efficient utilization of distributed storage
- API-enabled automation for robust and very high-rate data transport

## BACKGROUND

Since 2015, the Zettar core engineering team has been engaged to support the highly ambitious data transfer requirements of a DOE's Light Source project hosted at a national lab. In late 2016 the project's Data Management informed Zettar that it wished to move 20PBs worth of data from the lab to National Energy Research Scientific Computing Center (NERSC) to be stored on tape. The entire transfer should be done at a rate ~70Gbps over the 100Gbps production network connection provided by the ESnet and all data must be moved in about a month.

## PROBLEM STATEMENT

The desired point-to-point data transfer rate and the transporting of massive amounts of data must be accomplished with a cluster of very modest hardware at both the lab and NERSC ends.

## LIMITATION OF CURRENT SOLUTIONS

Host-oriented data movers couldn't do anything with those modest hardware.  Other scale-out data movers did not have high enough efficiency.

## ZETTAR SOLUTION

Zettar first conducted pre-production trials transferring actual 1PB data a few times over a 5000-miles "loop"  provisioned by the U.S. DOE Energy Sciences Network (ESnet) and then formulated the solution in 2017.

## MAJOR FEATURES EMPLOYED

Scale-out; RESTful API; parallel processing of multiple storage systems and volumes. The design was presented in the U.S. DOE Supercomputing 2017 (SC17) booth in Salt Lake City, Utah as a Featured Talk for comments and it was well received. The following figure illustrates the 2017 setup. The cost-effectiveness and the target performance level definitely put this project in a category by itself. It is also a great example of effective application of the co-design principle.

Aggregate, Optimize, Divide and Conquer

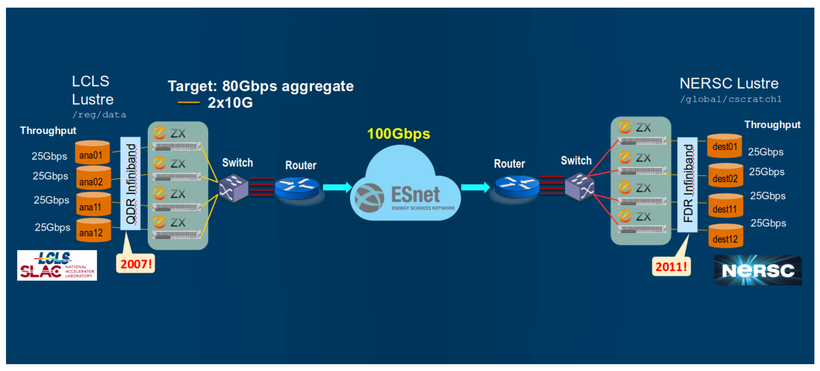


## ALSO APPLICABLE

- Data migration demanded by the facility location change of a data intensive business.
- Transporting of the aggregated data for large scale IoT deployments, smart cities, and fleets of autonomous vehicles.
- The detector output of any modern light source operation, nuclear accelerator, and large telescopes.