Itinerary of 3.2 PB

Don Petravick
Senior Project Manager
NCSA/University of Illinois at Urbana-Champaign

What (Instrumentation -- Terra):

Earth Observing Satellite in Low Earth Orbit. Launched1999. Four instruments.

Consistent and comprehensive long baseline record of Earth.

Particularly suitable for climate science. Crosses the equator at the same time each day. Same pattern of orbital paths.

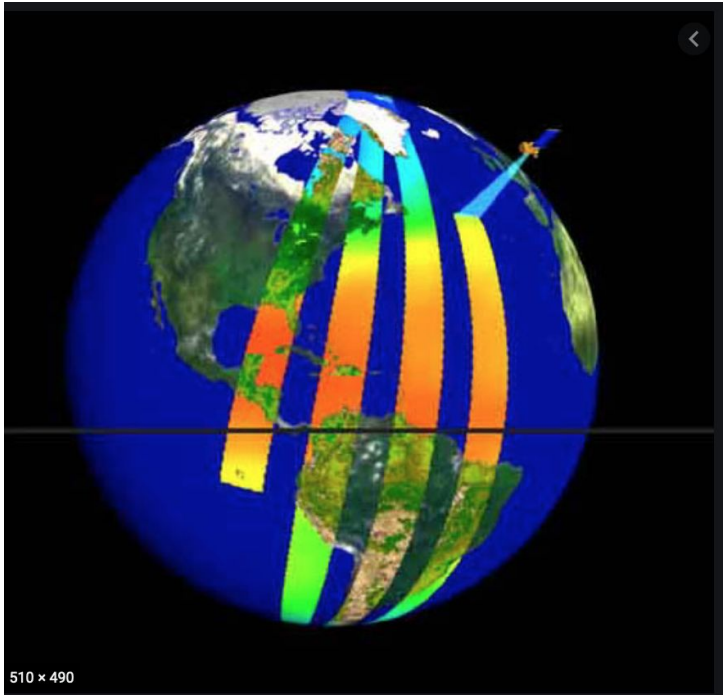

Credit: history.nasa.gov

What --Mission-scale science opportunities

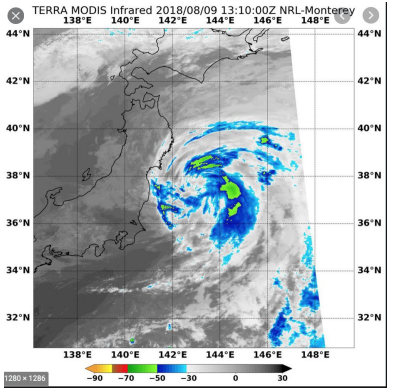

A good deal of science based on Terra is based on data the satellite captures for specific phenomena occurring at a definite location in space and time.

However, the satellite has also recorded a 20-year detailed record of the earth's climate, include those requiring large spatial areas over decades.

The records consist of millions of files, with differences in packaging and concept for each of the instruments.Gathering data for one climate study was onerous.

Gathering and Fusing the data once was seen as enabling Mission Scale Science.

The Gather Problem

NASA Awarded the Terra Access proposal, in mid 2010's, The project would build a fused data set from all Terra instruments to enable mission scale data fusion project. The award anticipated the resulting data set would be hosted by a NSF-Supported National Data Service.

The requisite input files were on a variety of NASA DAACs, with varying access methods for the files. Acquisition and ingest of > 1 PB of input files was time consuming; **NCSA copies were made to tape.**

- **Acquisition was enhanced by border controls optimized for connectivity at NCSA.**
- **Acquisition was by a variety of protocols, -- not everything was Grid FTP.**

The Fuse Problem

-- Fusing

- 1PB, > 10*6 files -> 3.2 PB, 80K ~35 GB HDF5 files.
  - Several runs, on
    - ROGER (optimized for Earth Science)
    - BW (Powerful Petascale) computers at NCSA.
  - Production machines expiring, now what?

The New neighborhoods

NSF-Awarded computing has a five-year lifetime, with no reliable prospect of renewal. The hoped-for National Data Service did not materialize as a hosting vehicle for the fusion product.

Transfer workflows:

- Due to an interest at NCSA in large scale earth science, NCSA Provided 3 PB of Disk-based storage as an interim archive measure, and to support local science use.
- NASA provided a deep archive, understood not for science use.

“Samplers” -- partial replicas

Mission scale use for the community still remains a vision.

Samplers: ~150 TB subsets of paths (e.g all time, specific orbits).

1. AWS Public Dataset Program, https://registry.opendata.aws
2. Open Storage Network, https://www.openstoragenetwork.org
3. An AWS bucket on glacier managed by NASA
4. National Center for Supercomputing Applications (NCSA), http://www.ncsa.illinois.edu

See: https://digirolamo.web.illinois.edu/projects/terra-fusion/

High-Level Design of a Data Carousel for the Basic Fusion Files

Idea: exploit AWS deep archive.

- “Reprocessing” is an envisioned use case
- Concern is movement costs.
- Approach is to gather requests in a “tour”
  - Carousel operator has fix maximum yearly cost.
- Analogous to train that runs on time.

Whitepaper at : http://hdl.handle.net/2142/107186

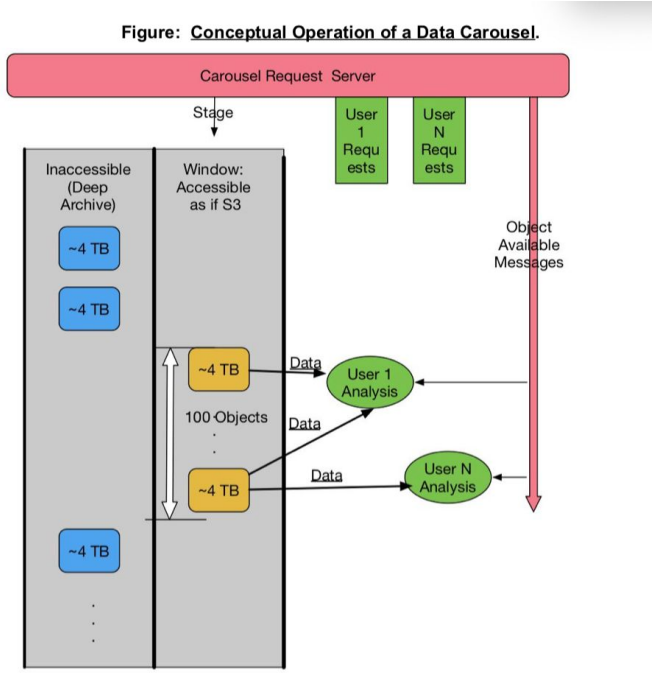

A low cost cloud home?

Characteristic	Value
Storage cost/year	$28,000
Restores/day	96
Restores # objects	480 @ 5TB
Min duration of full tour.	5 days
Movement cost for tour	$13,000
Duration of useful access once stage is complete	24 hour

Annual

S3 cost : $604,000

Carousel (26 tours) $185,000

Carousel (no interest) $28,000

Conclusion

- Finding homes for large data set which enable science analogous to mission scale science is is problematic.
- The Basic Fusion projects had challenges in gather operations and to find a way for community exploitation of its total data product
    - OSN is important
- The lowest-cost cloud-resident hosting method we could find provided access patterns requiring batch techniques.
    - We were unable to find a funding mechanism to sustain the data in this low cost mode. This is coupled with scientists' concerns for funding their computations and storage under commercial conditions.