

Trends in Large Scale Data Processing

Liam Gumley
CIMSS Science Symposium, Dec 12, 2012

My perspective

- > 20 years working with NASA Earth Observing System (Terra, Aqua).

- > 5 years working with NASA PEATE Project.

Experience with applications running on systems as small as iPhone and as large as S4.

Why do we need large scale data processing?

To apply an enhanced calibration or retrieval algorithm to historical data.

To apply a common retrieval algorithm to multiple sensors (e.g., VIIRS, MODIS, AVHRR).

To create time-series of climatological data (e.g., global frequency of high cloud).

Trends in data processing

Data Center hosts $> 4,000$ CPU cores. About 85% of the available floor space is in use.

Several different clusters (e.g., S4, zara, PEATE).

Data Center hosts about 220 non-cluster servers.

Clusters are configured either for high throughput (lots of jobs & data), or high performance (few jobs, lots of computations).

Power and cooling are finite resources.

Recent results from PEATE

33 year record of AVHRR GAC data (14 TB)
reprocessed on PEATE cluster in 7 days (rate =
1720X).

3 month test of VIIRS and CrIS data:

- VIIRS SDR = 100X
- CrIS SDR = 200X
- VIIRS CLAVR-X = 45X

Trends in data volume

Data Center holdings are approaching 4 PB.


GOES archive growing at 50 TB/yr, will grow at 125 TB/yr in GOES-R era.

PEATE archive growing at 350 TB/year.

Long term datasets with increasing spatial/spectral/temporal resolution.

Need data to be close to your compute servers.

DC archive growth



GOES archive
700 TB by end
of 2012

PEATE archive
350 TB by end
of 2012

Trends in software

64-bit Intel Linux.

Grid Engine for running jobs.

PostgreSQL for relational database.

Python for scripting.

GlusterFS and Lustre for filesystems.

Matlab, C, Fortran, IDL for algorithm development.

Subversion and Trac for source code management.

The historical data processing paradigm

We (CIMSS/ASPB/SSEC) develop and test software to implement or refine an algorithm.

We deliver the software to NOAA or NASA.

They implement the software in their own processing system.

They decide when reprocessing is done (if ever) and they distribute the data.

Reprocessing of MODIS

MODIS Ocean team officially reprocesses MODIS historical data 1 or 2 times per year.

MODIS Atmosphere and Land teams have officially reprocessed 5 times in 12 years (#6 is happening now).

There is no reason why reprocessing couldn't be done every 6 months.

A new data processing paradigm

CIMSS/ASPB/SSEC as a community hub for

- algorithm development and refinement,
- operational real-time forward processing,
- reprocessing to create climate datasets,
- data search, order, and distribution.

Like a combined NASA DAAC/MODAPS,
but located at UW.

Why will it be better?

Close cooperation between science experts who develop the algorithms and the computing experts who process the data.

New algorithms can be tested quickly.

Multiple datasets can be used easily (e.g., CrIS and VIIRS, MODIS and AIRS).

Any requirements regarding formats or code design will be internally driven.

It will be cheaper for the government.

Clusters and Clouds

Do you really need to build your own cluster?

Comment to PEATE Team at AGU: “Wow, you guys build your own clusters”.

Rent one from Amazon Web Services (AWS) for as little as 6.5 cents/core/hour.

Spin up a new server in AWS in < 60 seconds.

Spin up a Grid Engine cluster on AWS in < 10 minutes using MIT StarCluster.

Who uses AWS?

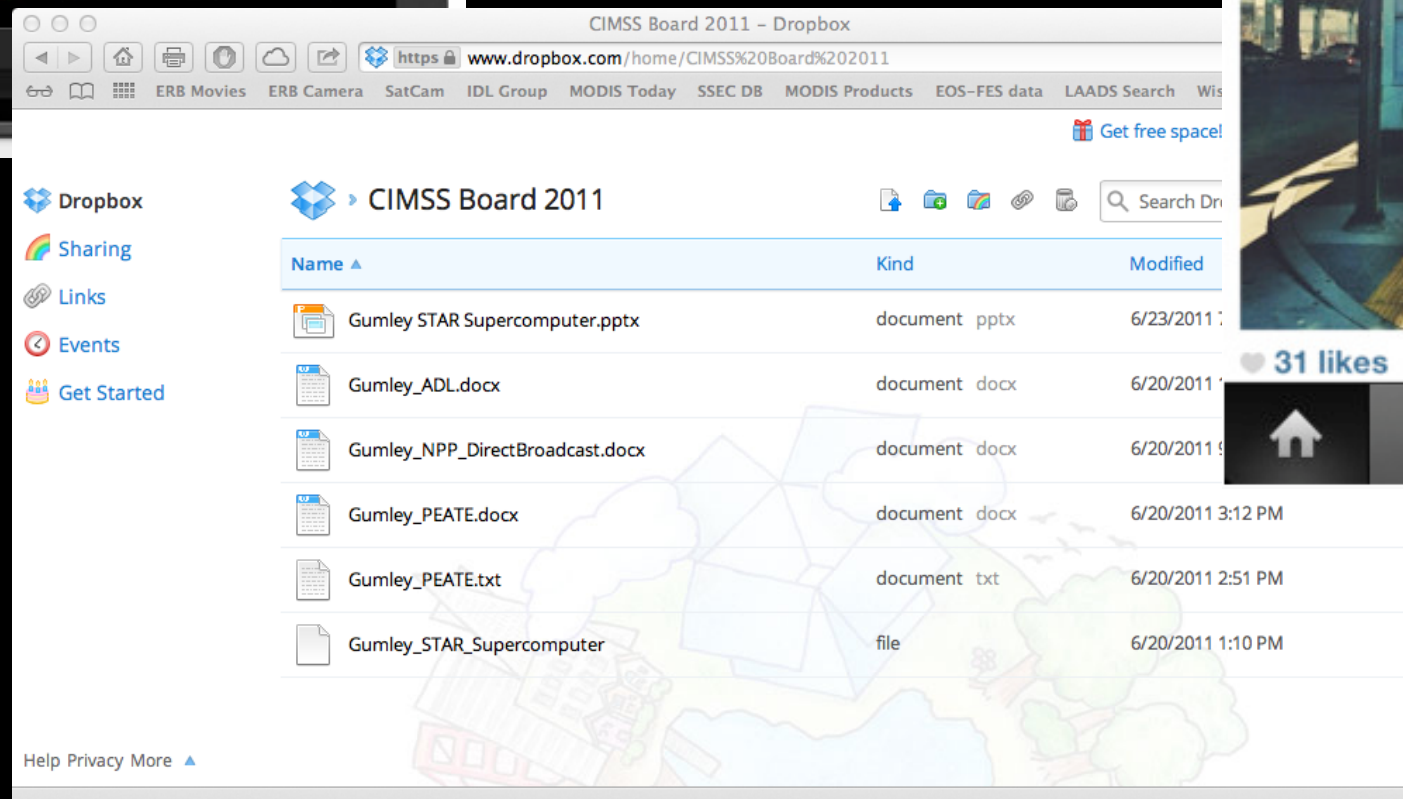
Netflix



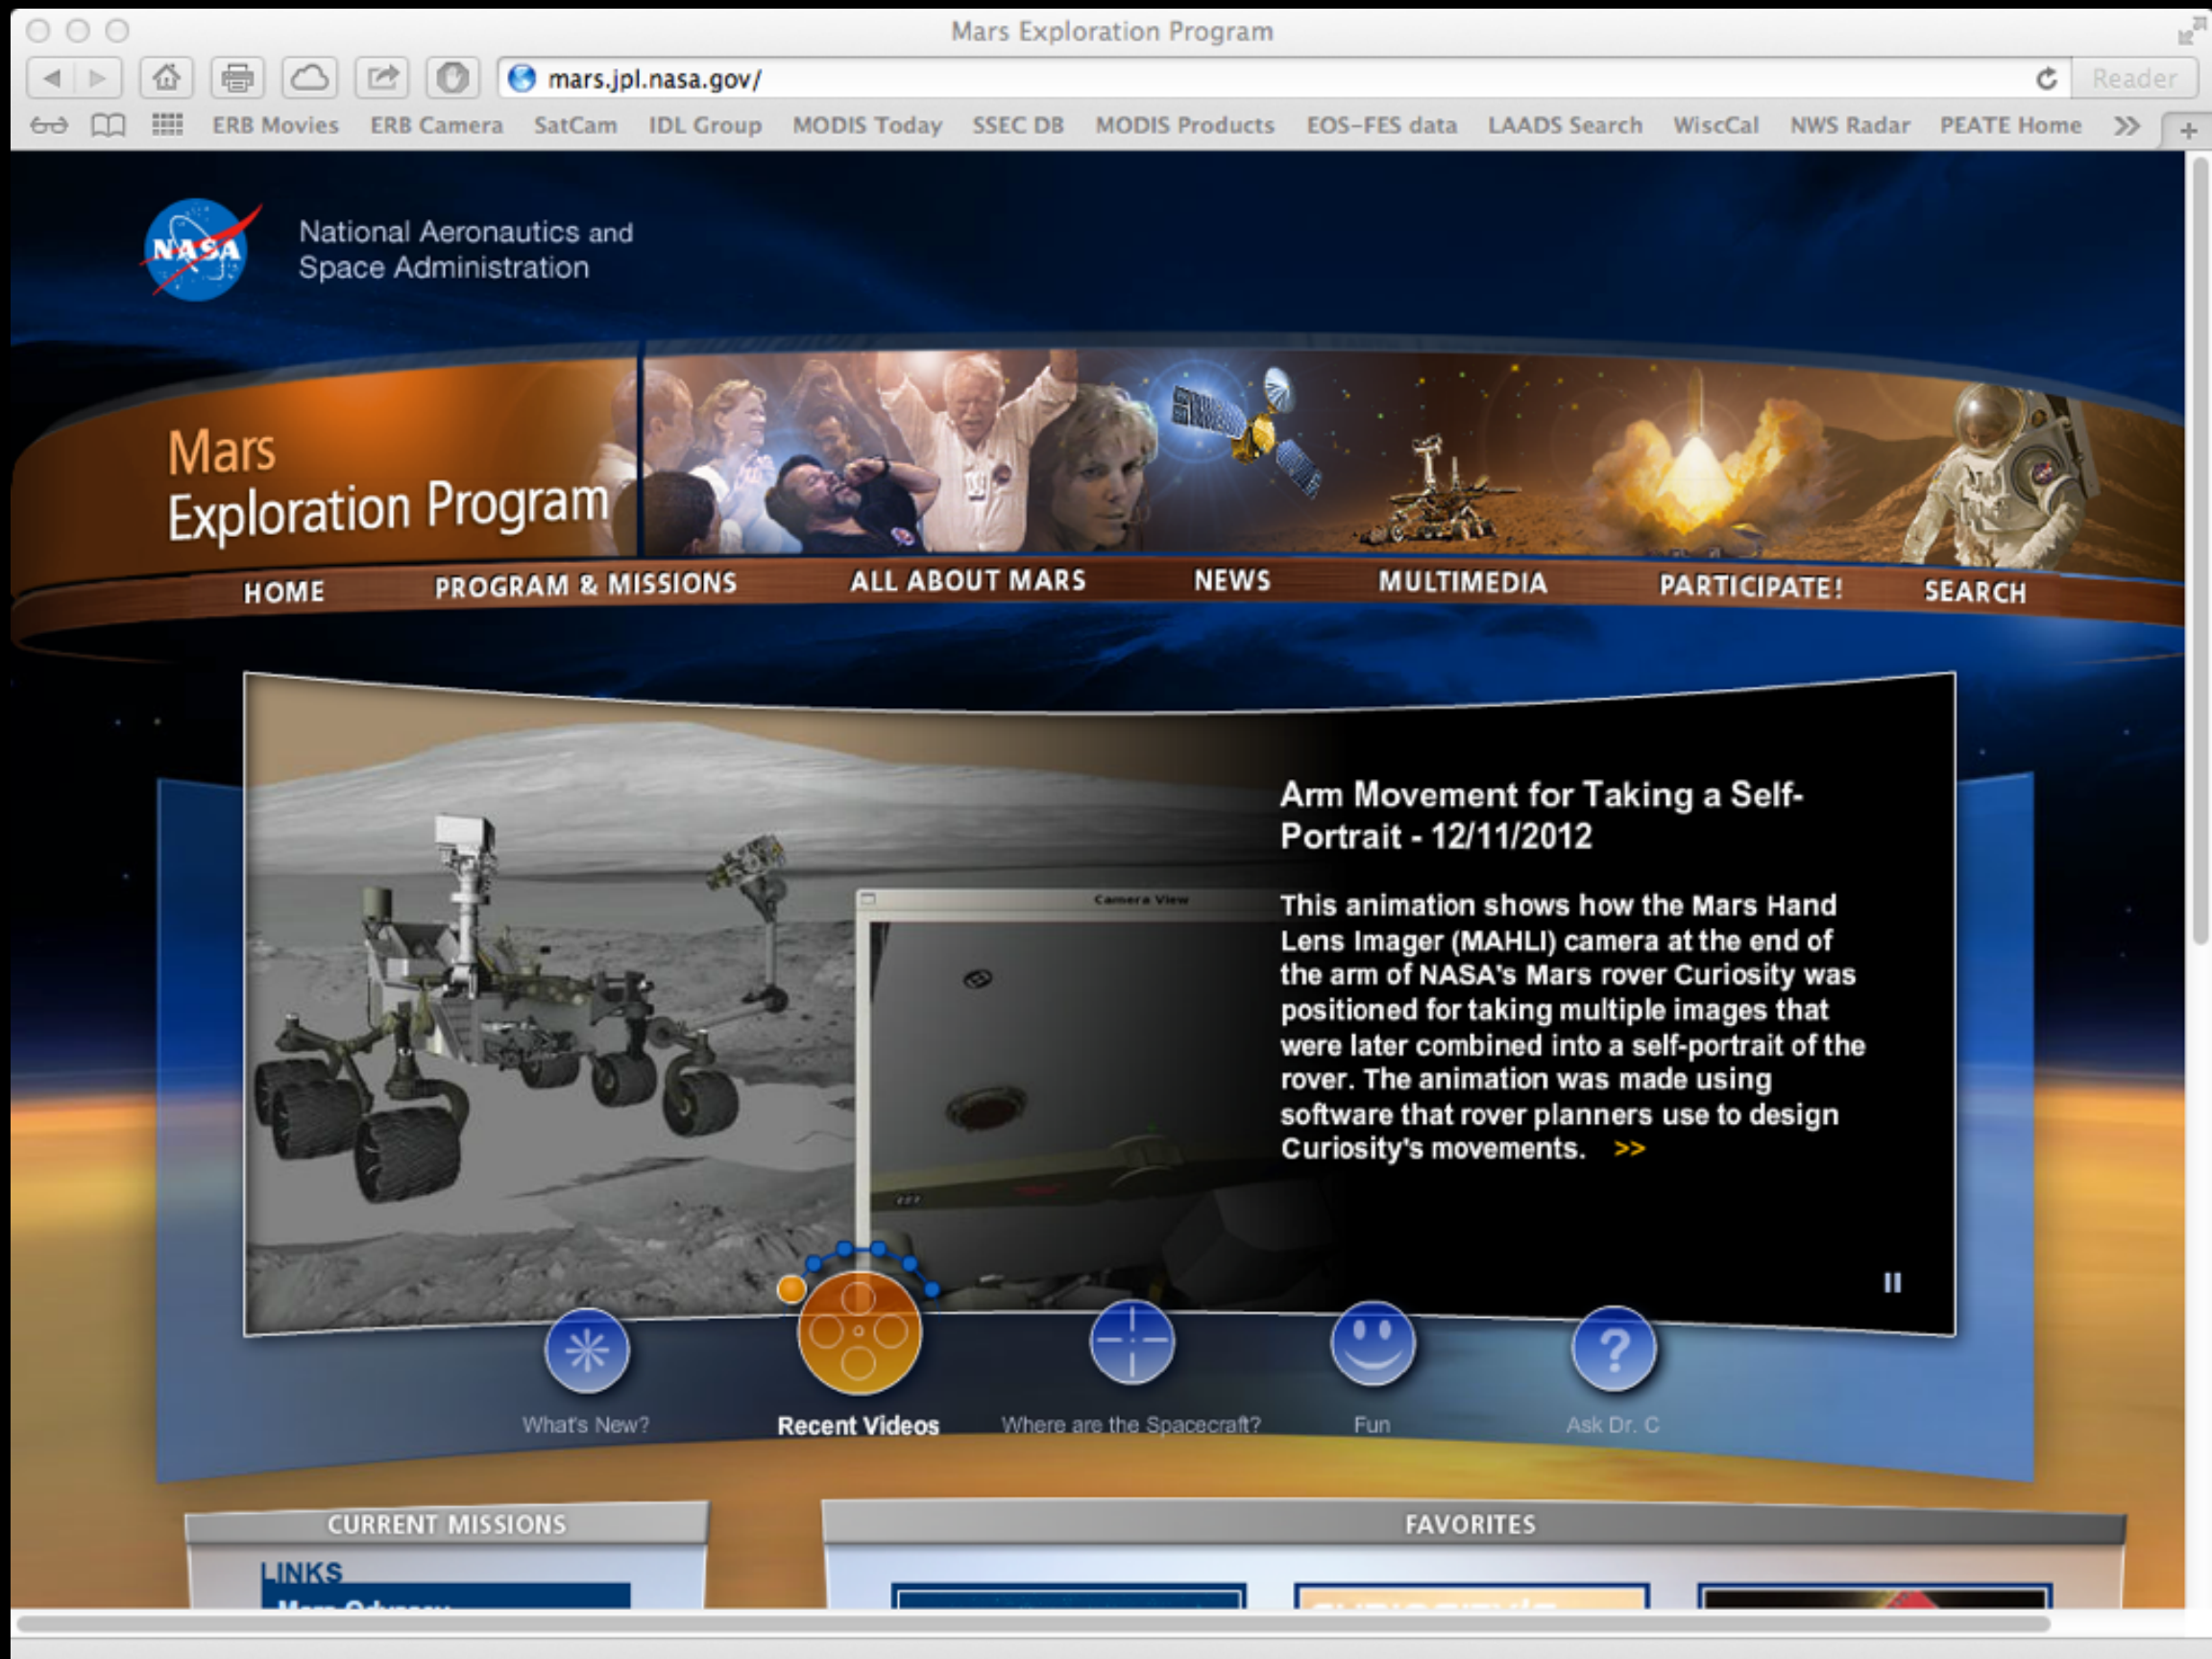
Instagram



Dropbox



Also hosted on AWS...



Cloud Use Case

CIMSS/ASPB/SSEC maintains a core set of servers either in-house or in AWS for routine daily processing.

When needed for reprocessing, servers are added on demand from AWS to handle the reprocessing load.

Less work on server purchasing, installation, setup, maintenance, and disposal; more work on daily and climate data processing.



What is she doing?

Analyzing 10 years of Aqua MODIS data for her PhD (about 130 TB).

Tools needed: Laptop, Wifi, Python, AWS.

Budget < \$200 (10 sec/granule = 2920 hr)

Summary

- We will be tasked and funded to do large scale data processing routinely (\$3-5M/yr?).
- LEO and GEO.
- AVHRR/MODIS/VIIRS, HIRS/AIRS/IASI/CrIS.
- The expertise (Science & Computing) and software will be in-house at CIMSS/ASPB/SSEC.
- The hardware (compute and storage) may or may not be in-house. It shouldn't matter.