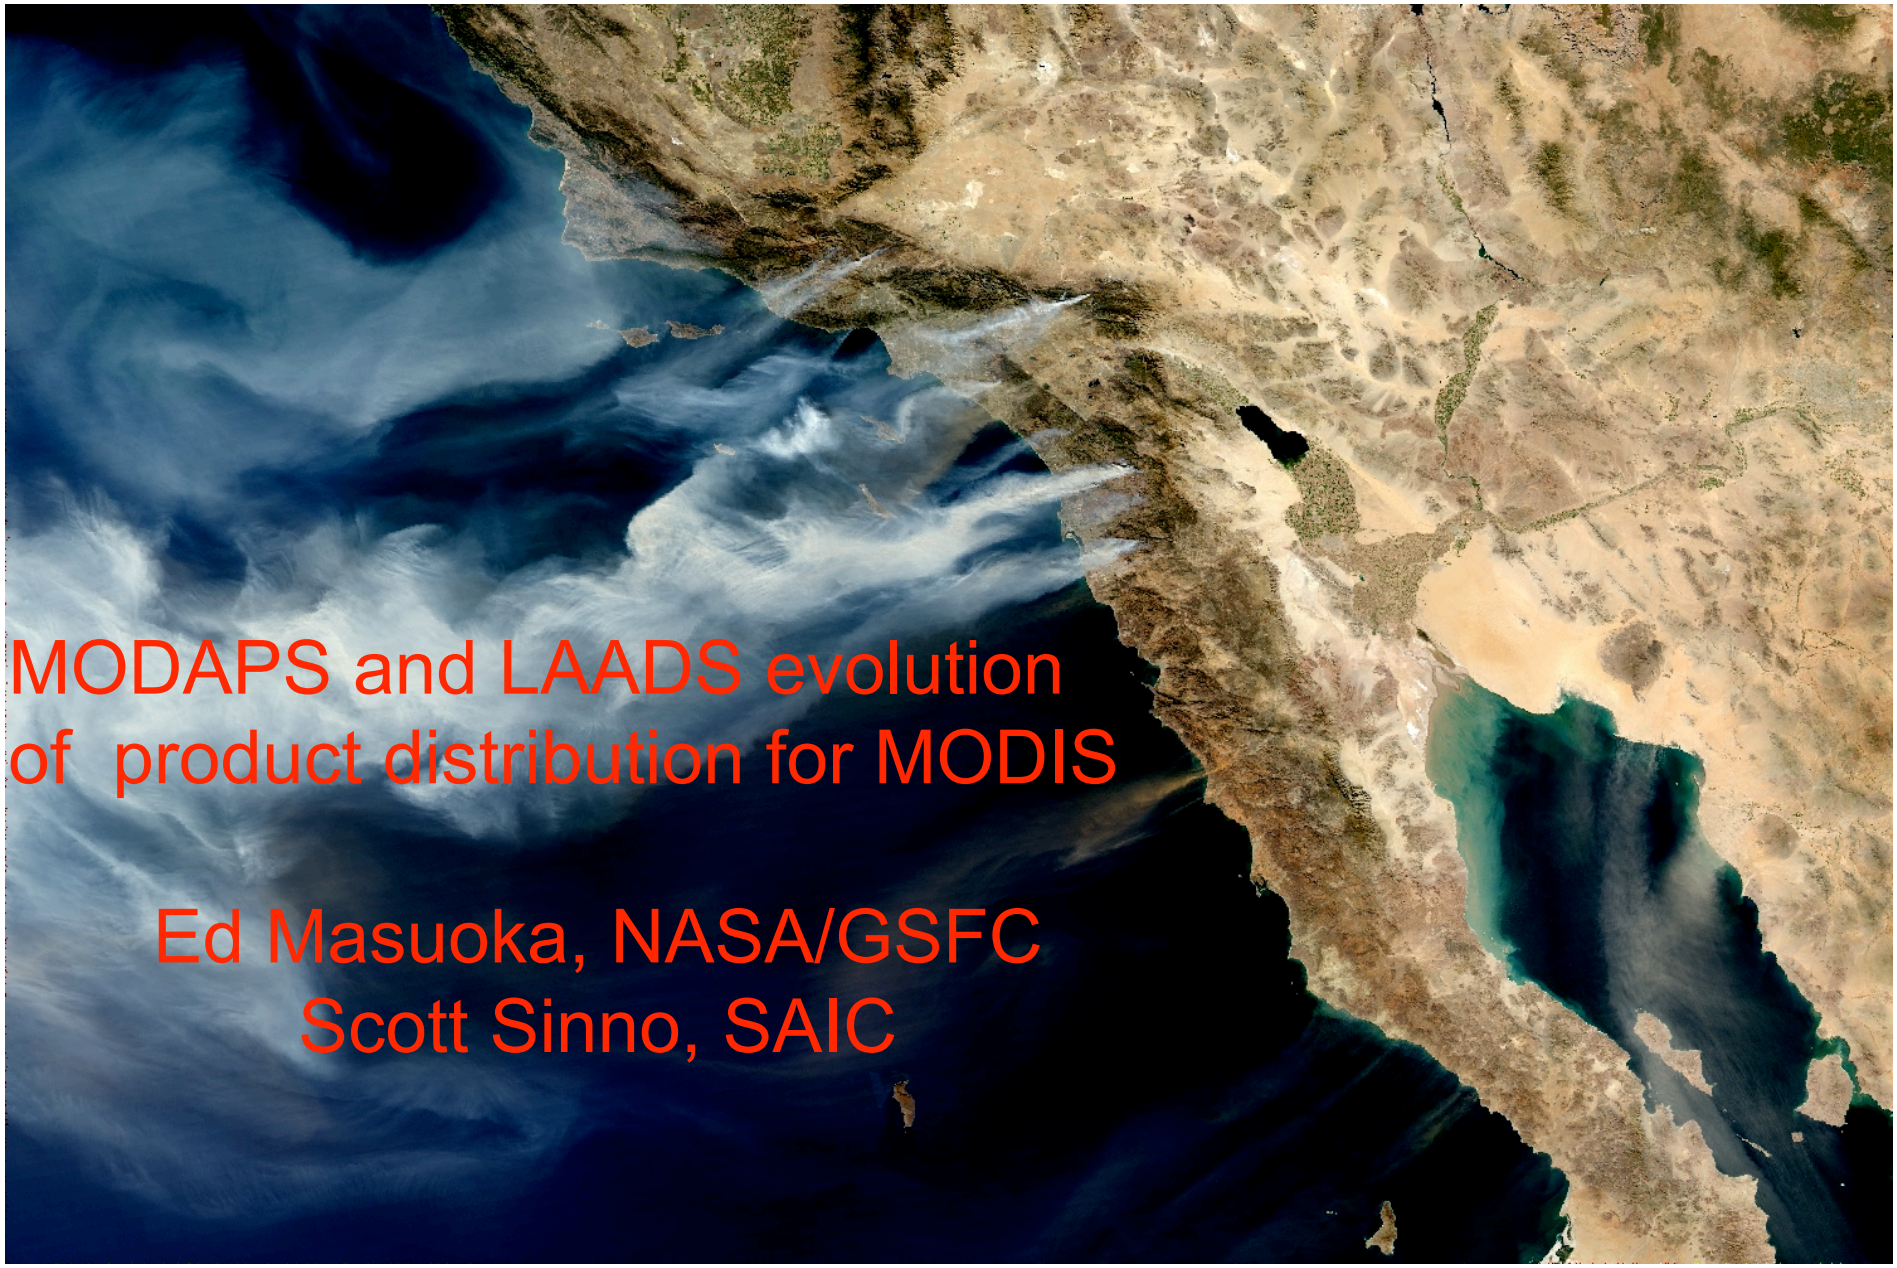

MODAPS and LAADS evolution of product distribution for MODIS

Ed Masuoka, NASA/GSFC
Scott Sinno, SAIC

MODAPS/LAADS Overview

- MODIS Adaptive Processing System
  - Produces and distributes science data products from the Moderate-resolution Imaging Spectroradiometer (MODIS) instruments on EOS-Terra and EOS-Aqua spacecraft
- LAADS
  - Archives and distributes MODIS products
  - Customers are DAACs (Land Processes, National Snow and Ice and Goddard Earth Sciences), MODIS Science Team and the Earth Science and applications communities

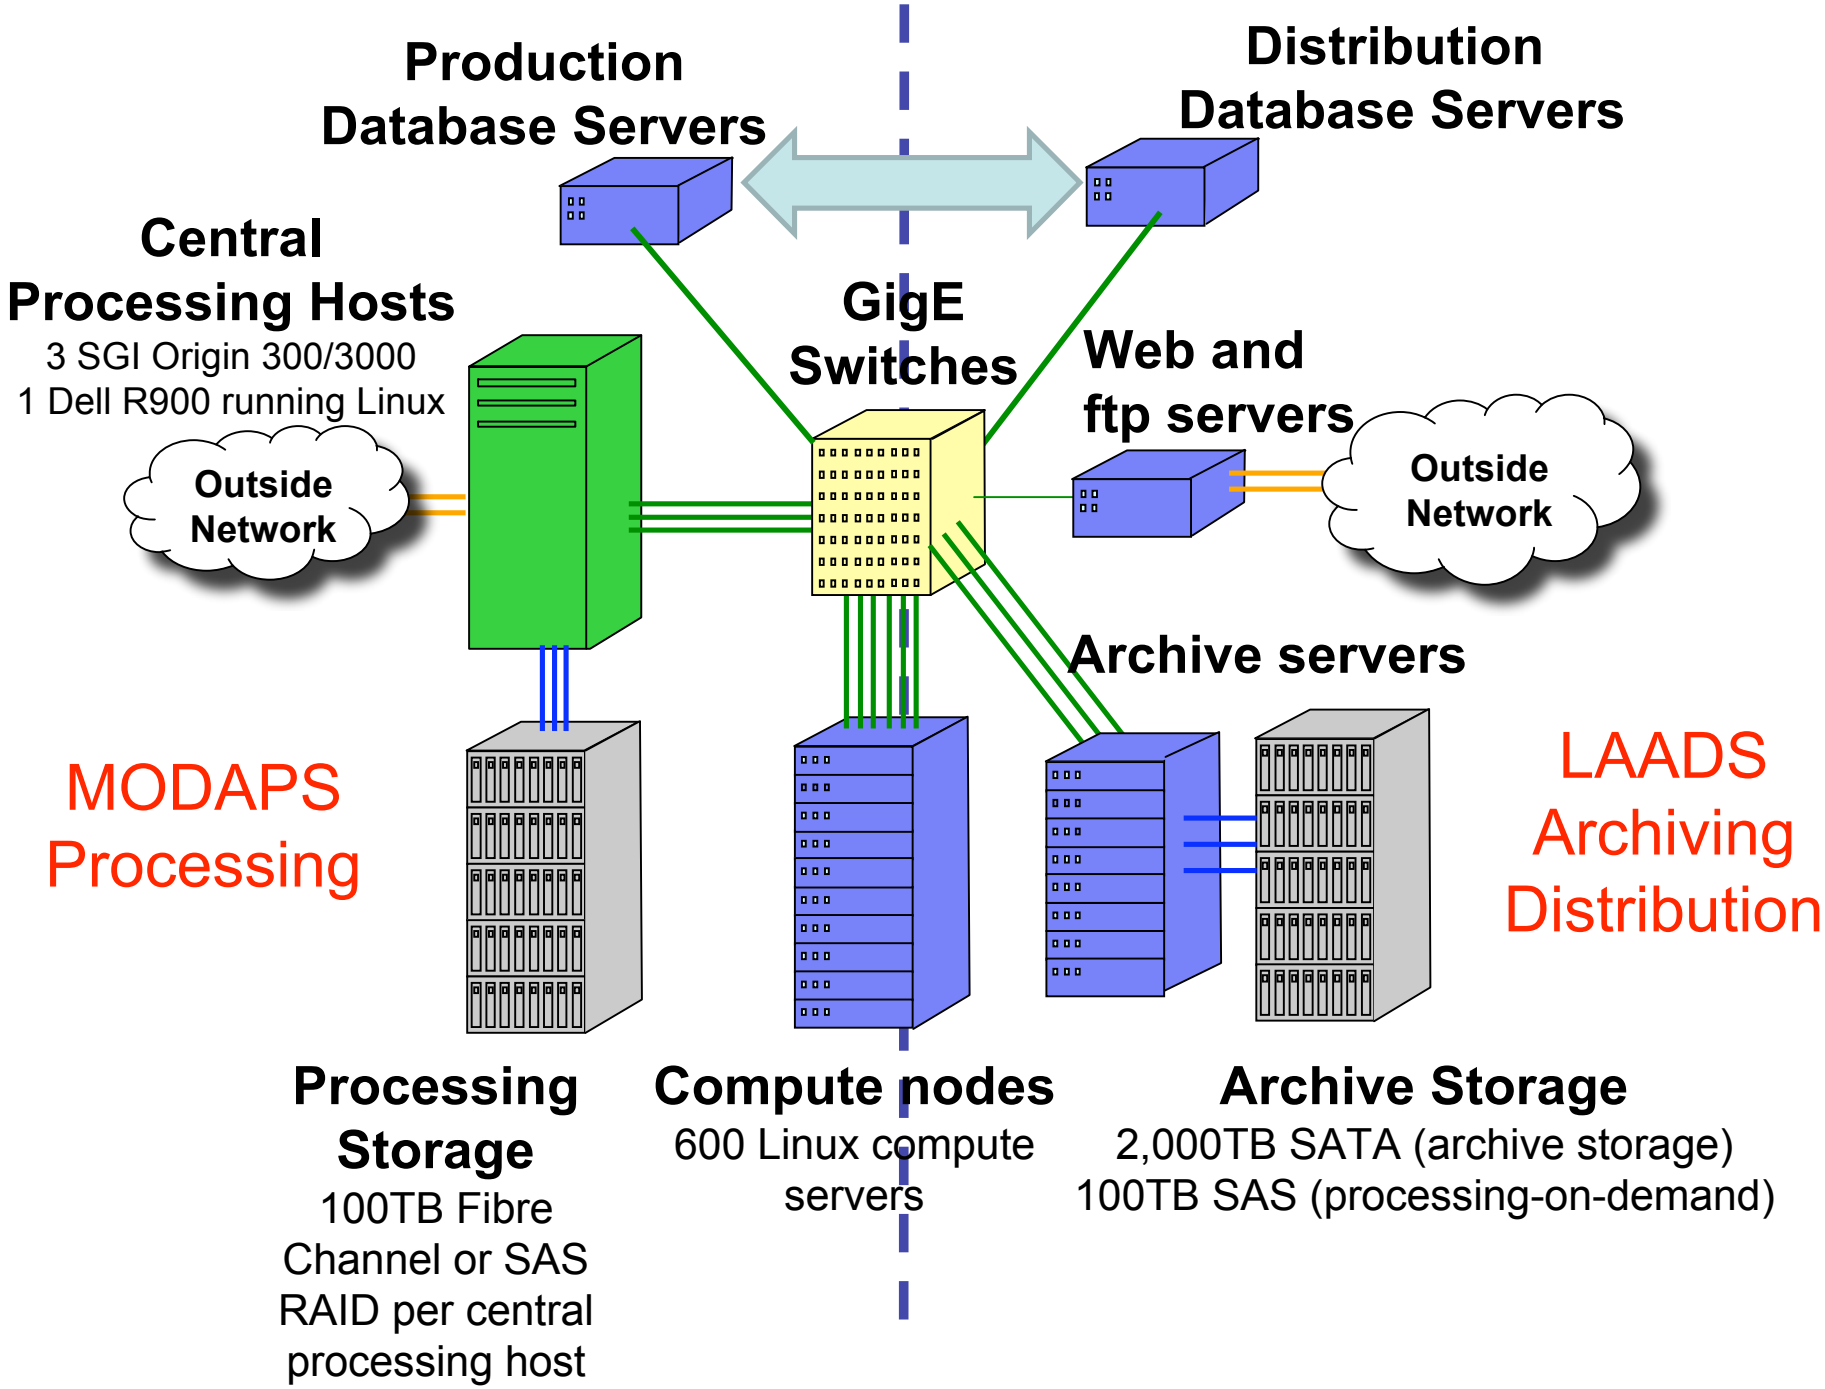
Production
Database Servers
Distribution
Database Servers
Central
Processing Hosts
3 SGI Origin 300/3000
1 Dell R900 running Linux
GigE
Switches
Web and
ftp servers
Outside
Network
Outside
Network
Archive servers
MODAPS
Processing
LAADS
Archiving
Distribution
Processing
Storage
100TB Fibre
Channel or SAS
RAID per central
processing host
Compute nodes
600 Linux compute
servers
Archive Storage
2,000TB SATA (archive storage)
100TB SAS (processing-on-demand)

Level 2 Products

- Retrieved geophysical parameters at the same location and in the same format as the MODIS Level 1 calibrated instrument data
  - 288 5-minute files
  - Each cover roughly 2340 x 2030 km
  - 250m, 500m and 1km spatial resolution

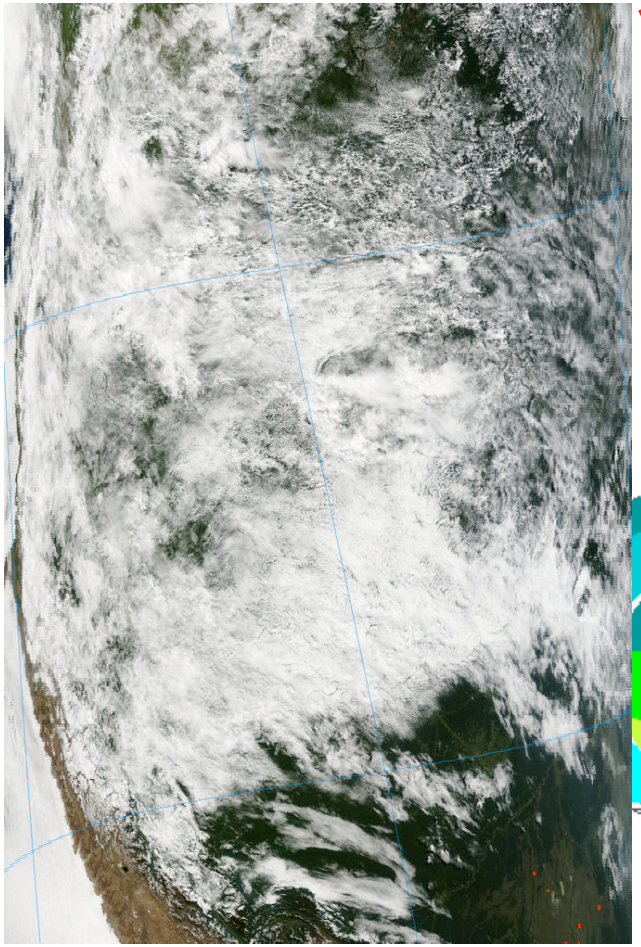

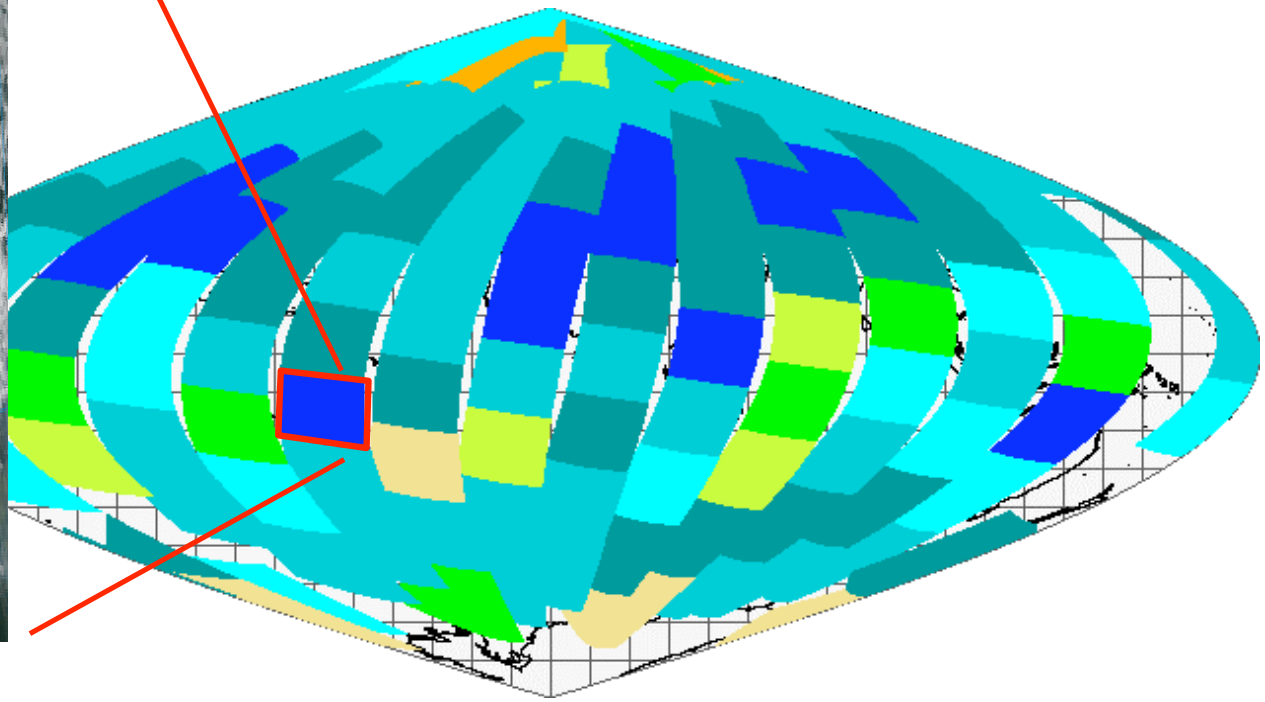

Level 2G, 3 and 4 Products

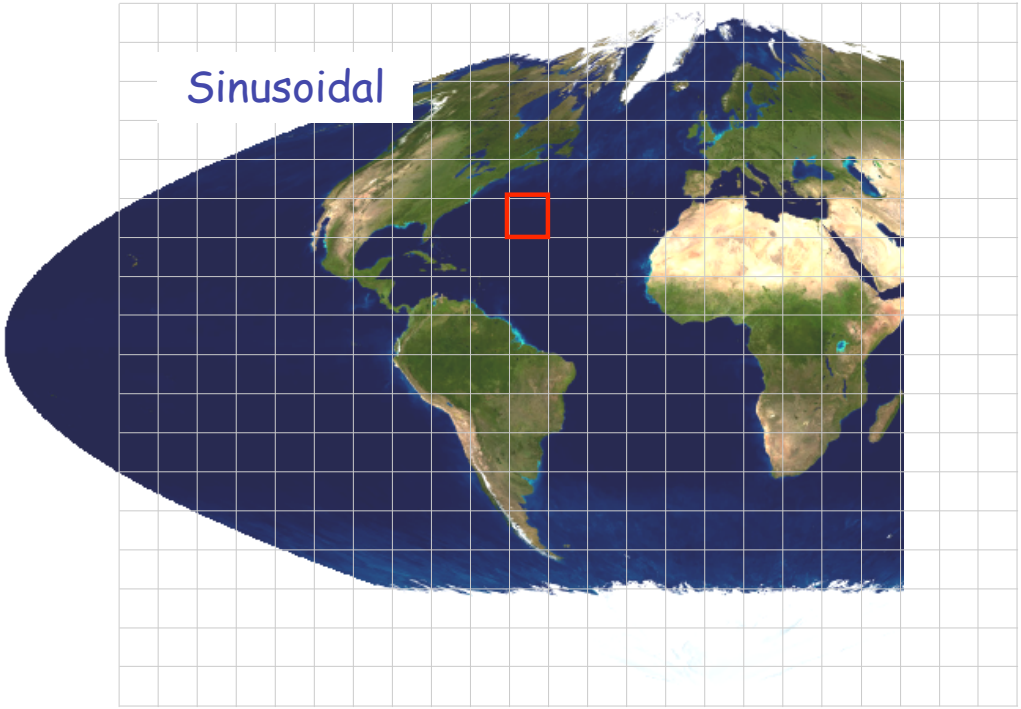

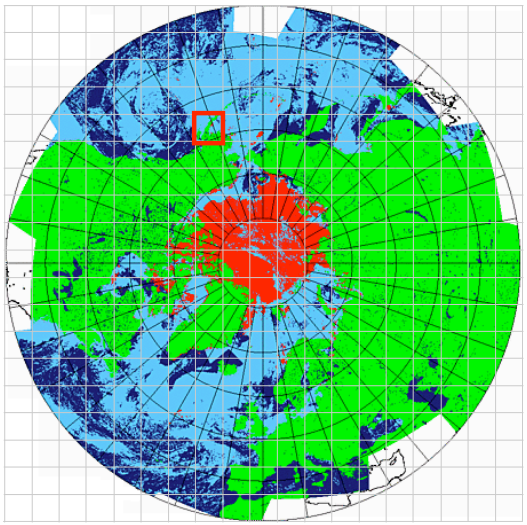

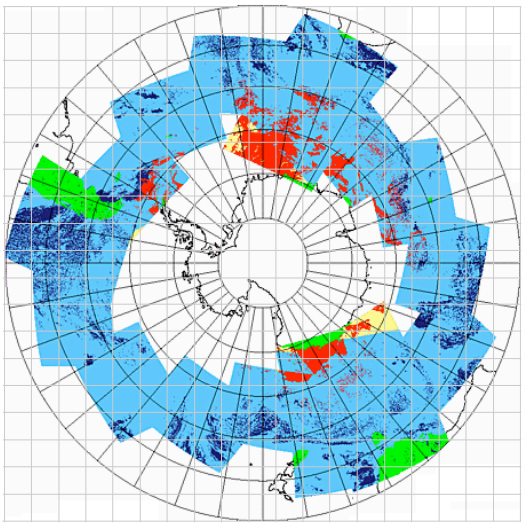

- Level 2G/3: earth-gridded geophysical parameters
- Level 4: earth-gridded model outputs
- Daily, 8-day, 16-day, 32-day, monthly and yearly products
- 10º x 10º Tiles Sinusoidal (equatorial); 7.5, 15 and 30 arcsec. resolution (roughly 250m, 500m and 1 km)
- LAEA (sea-ice products, polar projection)

Lambert Azimuthal Equal Area (LAEA)

Climate modeling grid products

- Resolution: 0.05º (now) and 0.25º (previous) degrees
- All products are lat/long grid
  - except sea-ice in polar grid

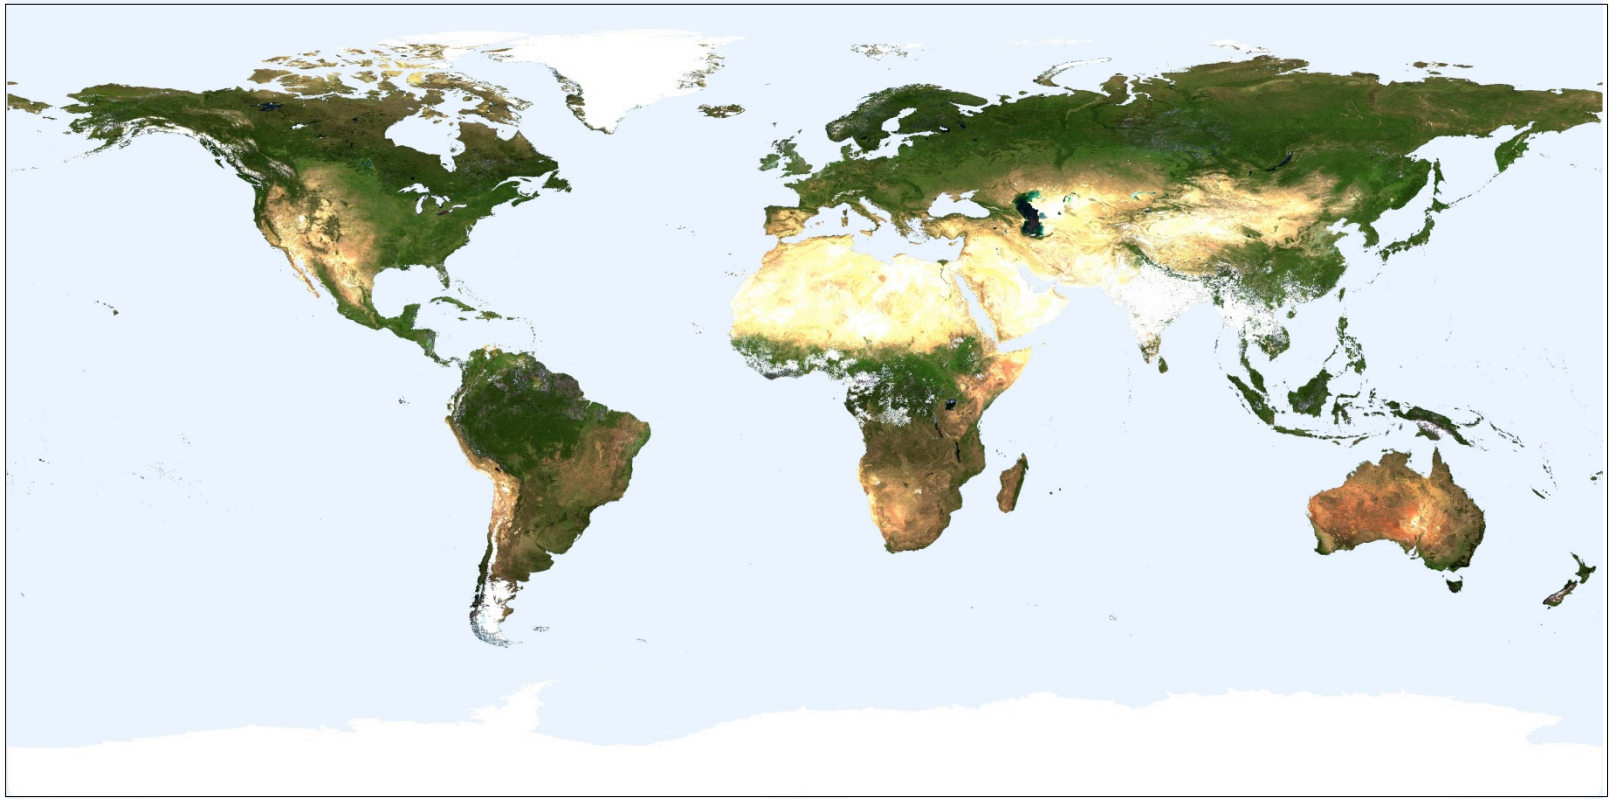

(from BU – NBAR CMG – days 193-208, 2001)

MODAPS – Product Generation

- MODIS Adaptive Processing System (MODAPS)
- Location: NASA/GSFC, Greenbelt, MD
- Produces: 51 global MODIS science products
- Daily distribution: 1.5 TB/day in 300K files/day
- Ships products to:
  - Land Processes DAAC,
  - National Snow and Ice Data Center
  - LAADS (Level 1/ Land and Atmosphere Archive and Distribution System)
  - Individual research and applications groups via subscriptions
  - MODIS near real-time processing systems

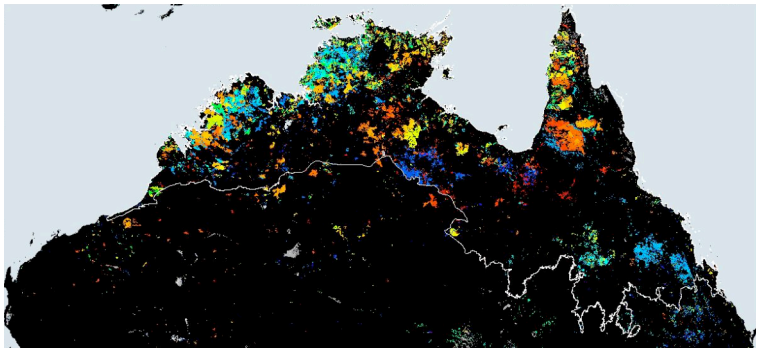

Burned areas over the 2003 dry season in Australia (March-November) from the new MODIS burned area product

Land Surface Reflectance with dynamic aerosol model improves correction for aerosol

LAADS – Archive and Distribution

- Level 1/Land and Atmosphere Archive and Distribution System (LAADS)
- Location: NASA/GSFC, Greenbelt, MD
- Archives : MODIS Level 0, Level 1, Land and Atmosphere products
- Archive: 350TB Level 0, 670 TB science products
- Daily ingest: 0.2TB to 2TB (during reprocessing)
- Daily distribution: 3.7TB in 160K files
- 17,500 customers September 2006 – August 2007
- Products distributed include:
  - Standard science products in HDF-EOS
  - Post-processed science products that have been transformed through subsetting, sub-sampling, mosaicing, reprojection, masking and reformatting

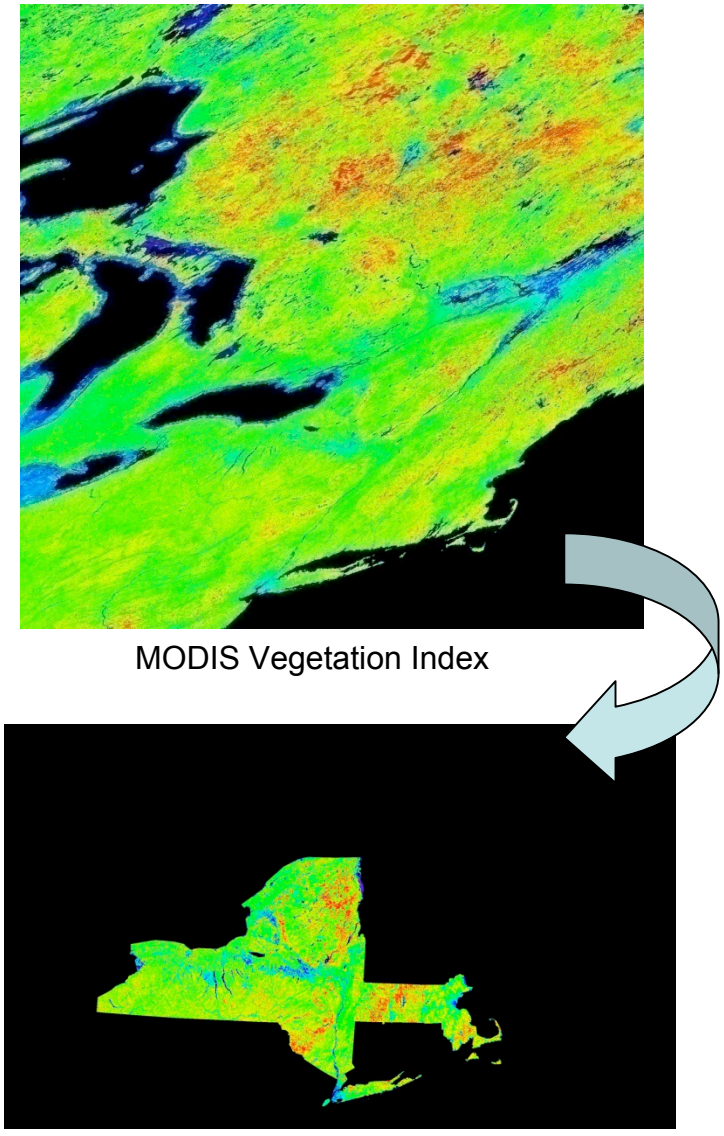

MODIS Vegetation Index

Subset, Masking and Reprojection to Albers Equal Area

LAADS Web

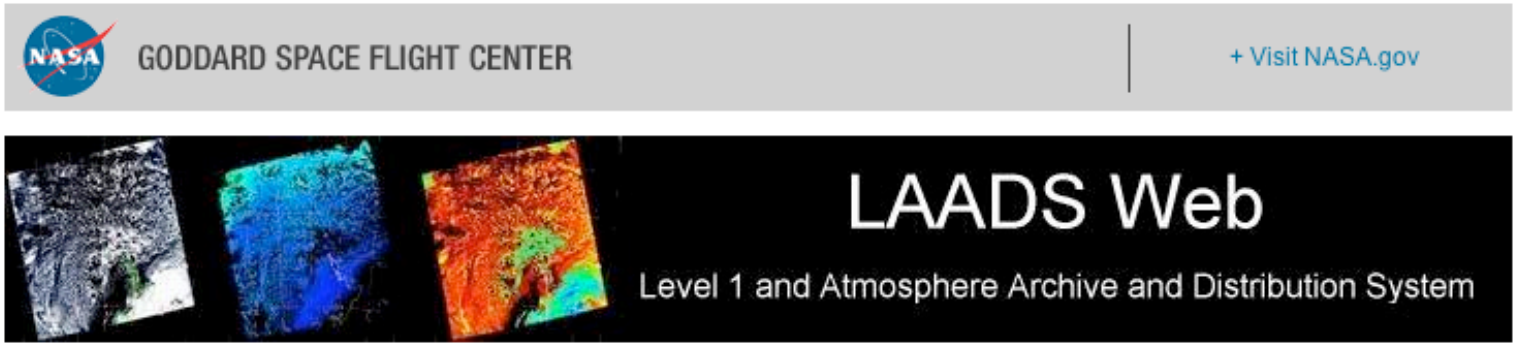

http://ladsweb.nascom.nasa.gov/

How end-users get products

- Use web site to search and place orders
  - Granule, tile-based views and metadata search
  - Transforming standard products
  - Files appear under: /orders/order # with directory path sent to end-user
- Pull products from online archive via ftp
  - Links to products stored under: /allData/Collection#/Product /Year/Day_of_year
- Arrange for products to be pushed to their systems during MODAPS processing

Files Distributed/Day

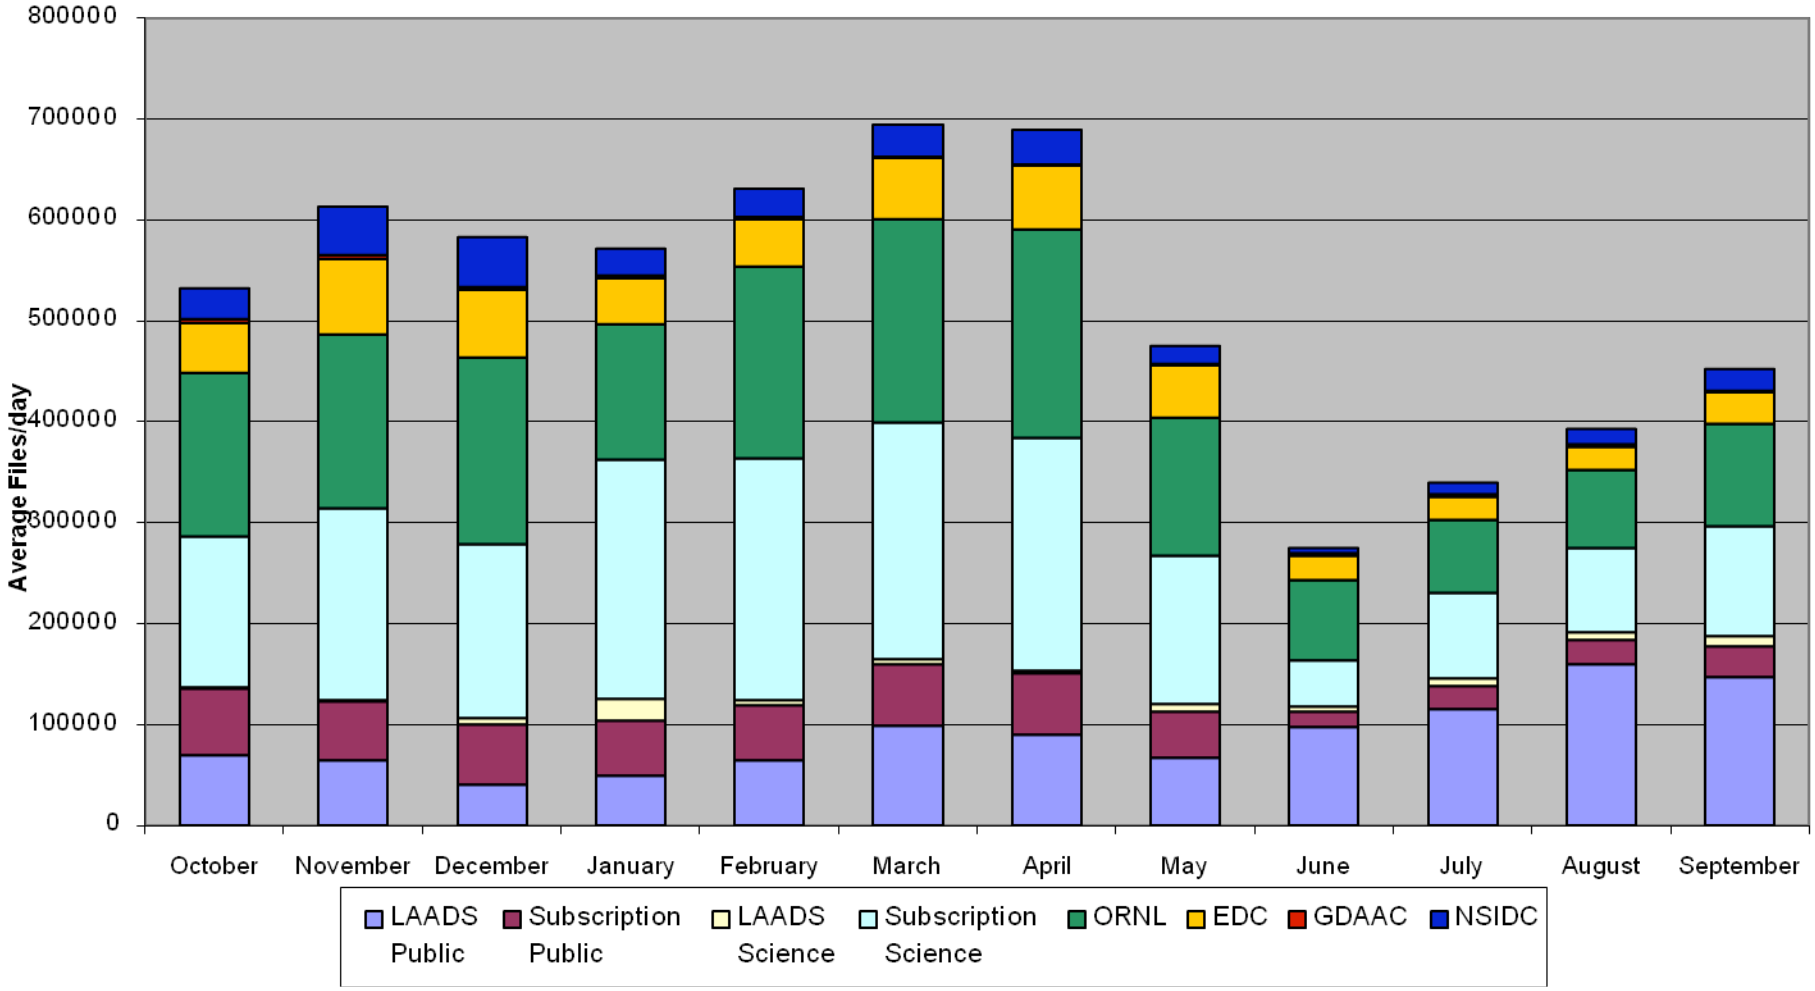

Volume Distributed/Day

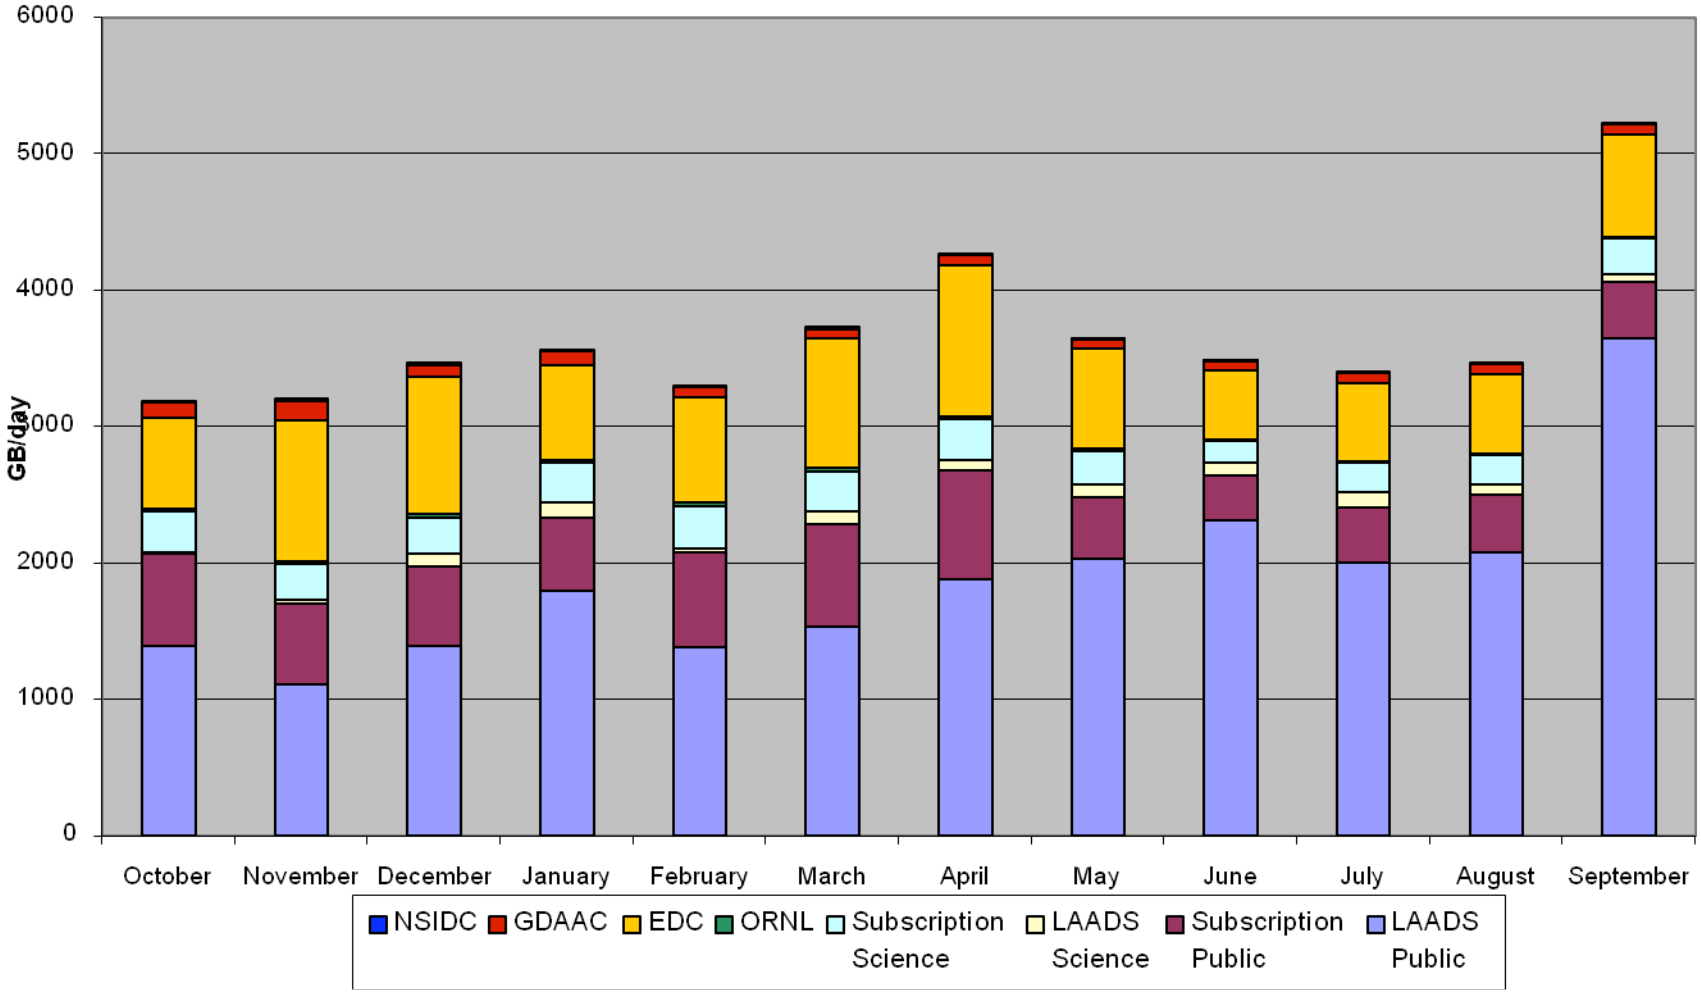

Order Response Time

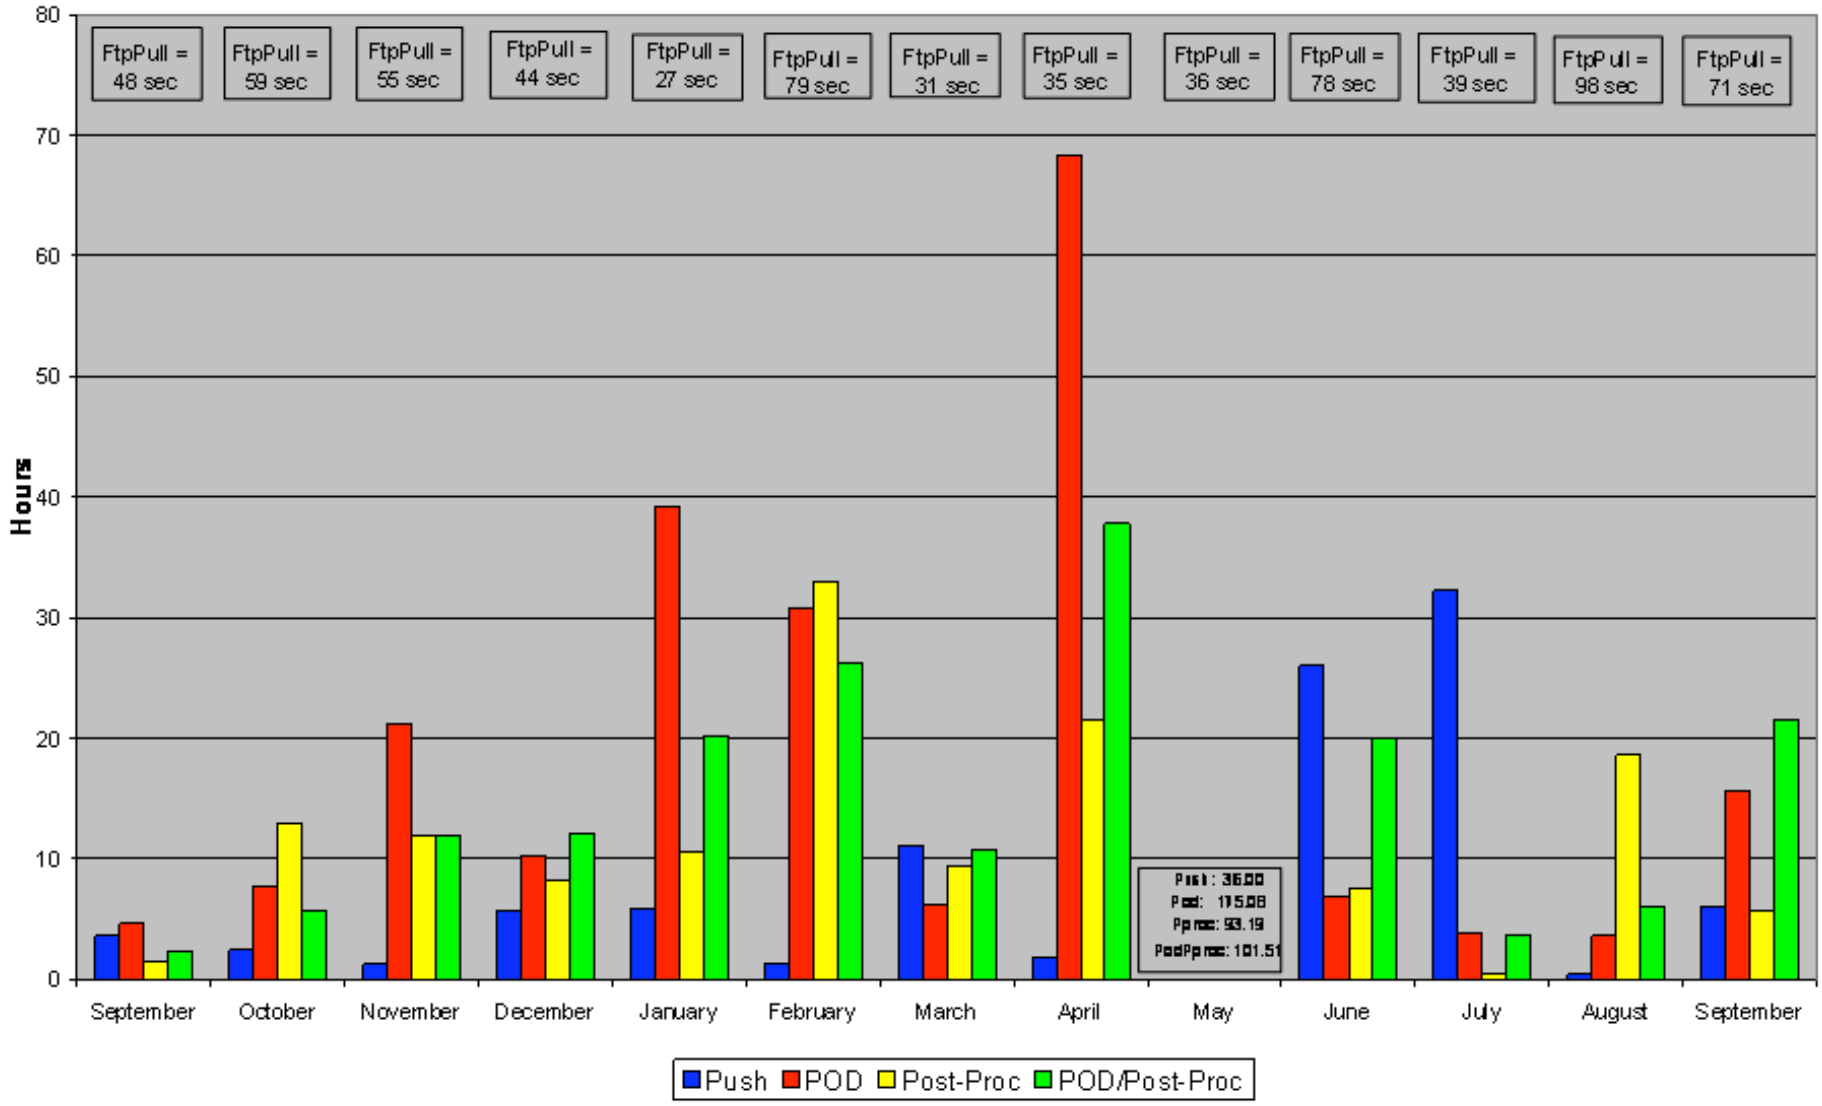

Archive Approach

- Protect original instrument data (Level 0) against corruption or loss
- Store frequently ordered, small and costly to produce products in online archive.
- Large infrequently ordered products are kept online in a data pool for 30 days. Products are produced through on-demand processing after expiration from data pool.

Products in the archive

- Level 0 – raw instrument counts from which all other products are derived.
  - 3 copies: Tapes at White Sands, NM and two copies on RAID 6 disk at two processing facilities at GSFC
- Level 1 Calibrated radiances
  - 1 copy of these products over North America and western Europe, for other regions of the world products are re-made if they have expired from 30-day data pool

Products in the archive

- Atmosphere Science Products
  - Small product volume. Quickly remade
  - 1 copy of all daily products and 2 copies of monthly products are in online archive
- Land Science Products
  - Large daily product volume. Land daily products are produced on-demand once they have expired from the data pool
  - 2 copies of multi-day (8, 16 and 32 day) products are in online archive
  - Archive limited to Land Science Team

Verifying file integrity

- Check_consistency – verifies that records in the production and archive databases match files that are stored online in terms of file_name and file_size
- Check_checksum – verifies that checksum stored for a file in the databases matches the checksum from the file, random check of 30,000 files/day. Followed up by a file system check if problems are found.
- Level 0 archive auditor – unzips compressed files (bzip2) containing L0 data from LAADS. If OK, stores files in archive and links them into L0 distribution tree. Also verifies integrity of current archive files by unzipping them.

Monitor servers with Ganglia

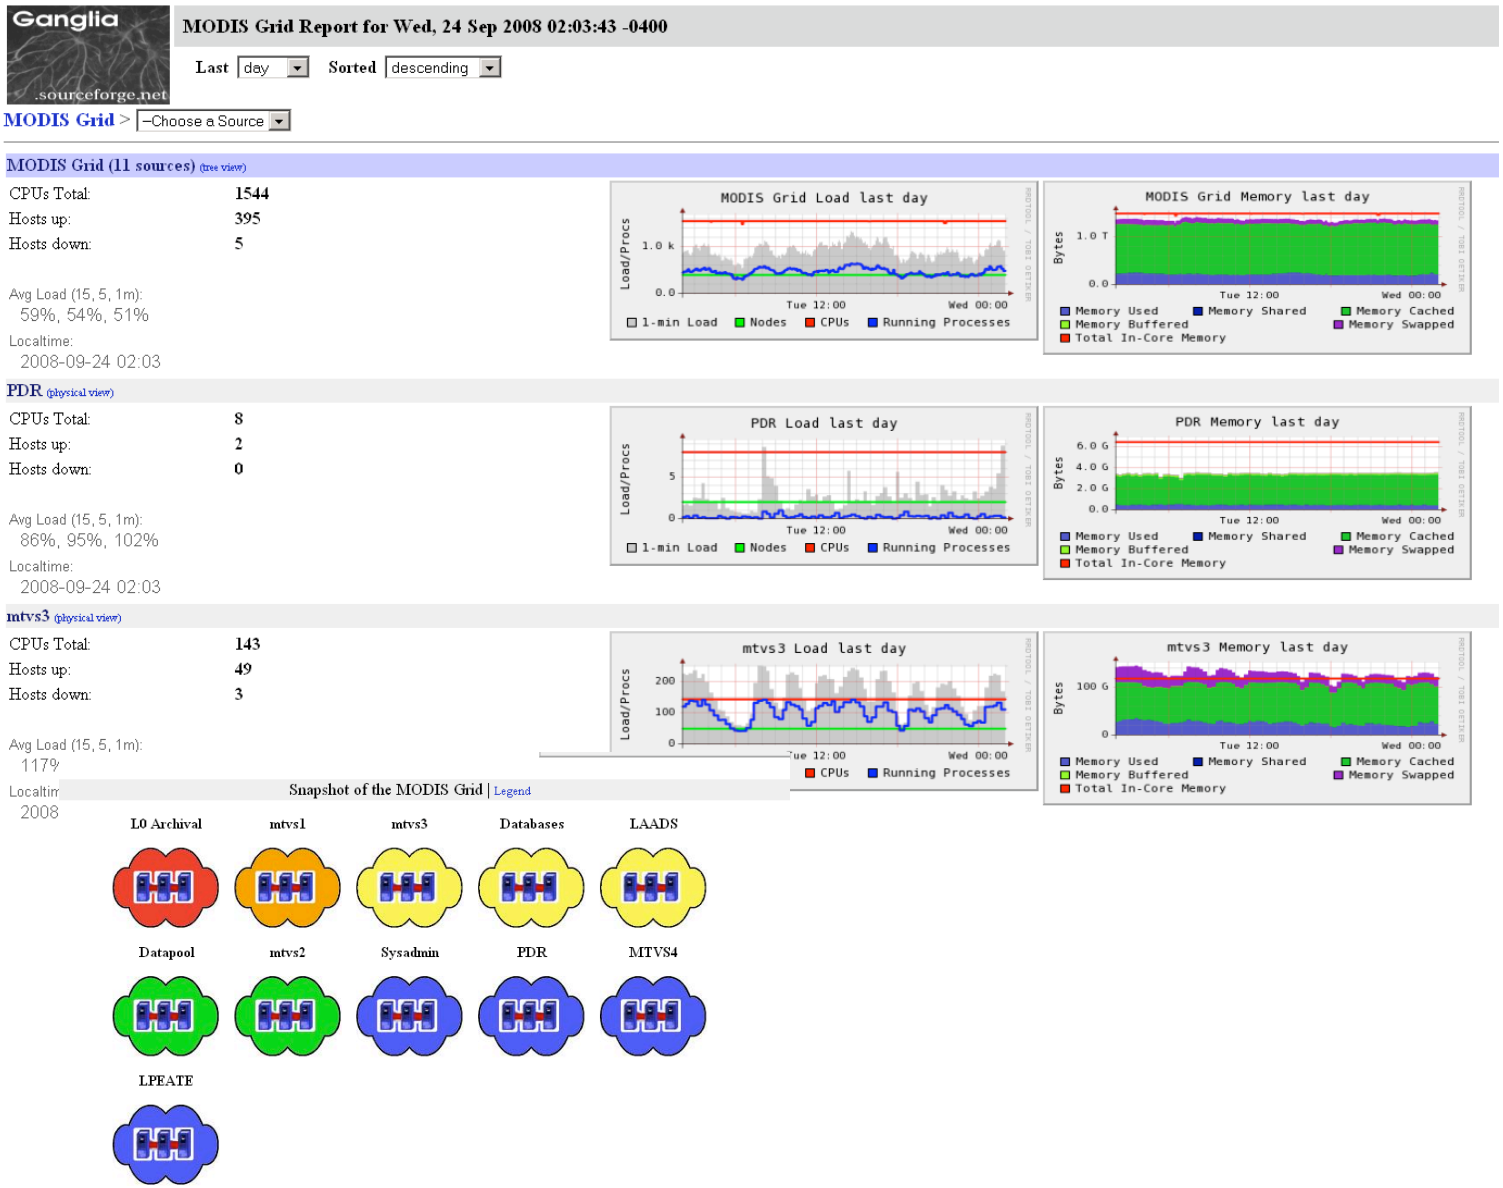

Monitoring RAID systems

Scripts checking Data Direct Networks RAID:

1. Parse logs from each controller
   - produces human readable version
   - emails alerts to problem queue and system administrators
2. On Monday verify 25% of available LUNs at random, performing parity calculation to check for bad blocks
3. Twice daily, check each controller to see if a drive has been failed but is rebuilding.

Monitoring RAID systems

- RAID Inc/MPAK (Infotrend controllers) – On Monday 1/4 of all RAID arrays have parity verified for all blocks.
- Mandriva with software RAID 6 – No special verification, only MODAPS check of checksums for files
- Open Solaris with ZFS – scrub to locate problems

Time to rebuild/repair drives

Storage System	Drives/File system	Rebuild Hours	Comments
DDN S2A 6000,8000,9500 >4,000 drives (14)	8+1 8+2	<2 - 30	Time to rebuild depends on drive size. Can repair bad blocks if drive has not failed.
RAID Inc (RAID 5) (RAID 5) (RAID 6) 1,700 drives (36)	16 bay (7+1), (6+1), HS 24 bay 3*(5+1), (4+1), HS 24 bay 2*(10+2)	30	Rebuild entire disk
XFS JBOD(RAID 6) 72 drives (3)	10+2	TBD	Rebuild entire disk
ZFS JBOD(RAID 6) 120 drives (5)	10+2	TBD*	Can rebuild only the portions of a failed drive that contain files.

* Partially filled volume took 4 hours to rebuild

In 2009

- Increase outbound network bandwidth
- Retire SGI Origin hosts
- Store Level 1B calibrated radiance products rather than produce on demand
- Portions of archive spin down when idle
- Continue testing Open Solaris servers with JBOD as alternative to hardware RAID