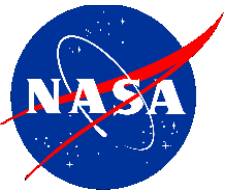

The MODIS online archive and on-demand processing

Edward Masuoka

NASA Goddard Space Flight Center, Greenbelt, MD, USA

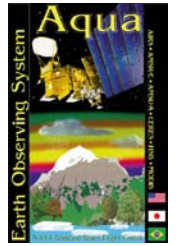

Production – Driven by Science

- Over the Terra and Aqua mission lifetimes, better calibration and characterization of the instruments have been performed and new and improved algorithms have been developed
- To extract the maximum value from the investment in the these missions, multiple reprocessing campaigns are needed
- For the MODIS instruments, the schedule of the reprocessing is also driven by the complex interdependencies between algorithms
- Terra and Aqua, like many of the NASA missions, have lasted well beyond their design lifetimes of 5 years, and are expected to survive long enough to each produce more than 15 years of high quality data
- These reprocessing activities are typically driven by the science team and community
- The reprocessing cycle involves three phases
  - Development of the algorithm improvements
  - Testing the improvements
  - (Then) the actual reprocessing

Contents of the Data Archive

- A reprocessing of the entire MODIS product suite occurs every 3 years driven by the science team's schedule for algorithm improvement. Additional reprocessing of selected products may occur more often, e.g. aerosols.
- The archive contains at least two complete versions of the reprocessed data, the current reprocessing and the one immediately before it, both of which cover the entire data record from launch to the current day.
- For earlier reprocessing campaigns a "Golden Month" is stored in the archive in the event that products from earlier collections need to be compared with those from the most current.

Reproducing Products

- Reprocessing vs. Reproducing Products
- Why not just archive products?
- When do we reproduce products?
- What we need to in order to do this
- How we use on-demand processing with our online archive

Why not just archive products?

- In some cases it is more cost effective to produce products on-demand rather than store them, e.g. Level 1a.
- We have many instances of data products produced in testing connected with improving algorithms that shouldn't be archived permanently but may need to be recreated at a later date
- If we only archive products then there is no guarantee that we or a third party can reproduce them again
- We distribute the MODIS science processing software and its documentation. A benefit of reproducing a particular version of a product is that in doing so we have demonstrated that enough information is available for us to reproduce it in the future.

When do we reproduce products?

- Versions of products deleted from the test data archive are needed again as baselines for testing improvements
- Products not archived due to size and limited demand, e.g. L1A, L2g, are ordered by users
- Versions of products are replaced in the archive by newer versions, e.g. Collection 3, but an older version is ordered
- Products missing due to hardware failure or human error from our archive or other MODIS archives need to be regenerated.

What we need to do this

- Correct versions of:
  - Basic input files for processing. Level 0, telemetry and ancillary/auxiliary files
  - Science software, look-up-tables and production rules
  - Science Data Processing Toolkit (SDPTK)
- Operating system and compilers
- MODAPS processing system
- Information in our database and product metadata that ties the items above together to produce a given version of the products.

How we use on-demand processing with our archive

- Ordering a data product that doesn't exist on disk will caused an order to be sent to one of the 20 (24 core) compute servers that handle processing on-demand (POD)
- If the product was generated earlier, version information in the database is used to reproduce the product. Default is the current version in the archive
- If the product is a custom product based on transforming existing products then the POD system applies the requested transformation(s), i.e. subset, subsample, reprojection, mosaic, mask and format conversion
- At the completion of the POD request, a link to the files is placed in the user's order directory. Files are available for download for a limited time.

Other points

- Reproducing a product is necessary but not sufficient. It does not ensure that software and documentation are complete and that an end-user can understand how a complex algorithm works.
- Reproducing a suite of products is not always easy if your system did not make the products, e.g. Level 1 MODIS products from Collection 4 made in GES DISC.
- Over the years computers and their operating systems change, we keep samples from earlier Collections, the "Golden Months", to ensure that as we change underlying hardware we are still able to reproduce products though not always exactly. How close is close enough is determined by the Science lead for a given product.

Backup Slides on MODIS Processing, Archiving and Distribution

Production – Reprocessing Rates

- MODIS team is in the process of preparing for the fourth major reprocessing for which the production phase is schedule to begin later this year, about 11 years after launch
- Average time between reprocessing is about 3 years with the first one taking place after the algorithms were stabilized after launch
- With a three year cycle, the reprocessing is done in less than a year with the majority of time spent in algorithm development and testing phases
- For current reprocessing campaign of 11 mission-years of MODIS/Terra data and 9 mission years of MODIS/Aqua, the reprocessing capacity is 100 data-days per day (a conservative estimate based on ingest capabilities at the archives) and network bandwidth

MODIS Adaptive Processing System

- Scalable system for MODIS, AVHRR, VIIRS and Landsat processing
  - Designed so that processing resources can be easily moved where needed
- Built with commodity hardware
  - Easier to scale, cost savings and easier technology refresh
- Built with open source components
  - Linux, Apache, Perl, Postgres, Subversion, FUSE
- All data products are online
  - Facilitates reprocessing of Level 2 and Level 3 products
- Designed to run with limited staff
  - “Lights out” processing outside of normal business hours
  - Months of processing can be queued up and execute w/o human intervention
  - Alerts emailed to system administrators when hardware components generate warnings or fail
  - Easy to use tools for monitoring the system (Ganglia) and investigating failed jobs
- Rapid updating or provisioning of servers with science processing software, the operating system and applications (Depot and SATE)
- Capability to separate data products into archive sets to accommodate storing a large variety of test results in an online archive

MODIS Processing System

- Middle-tier servers and disk storage
- Open-source software
- Uniform system configurations
- Automated updates daily from Depot
- Automated H/W problem reporting
- >1,000 servers, 5PB of storage
- Maintained by 6 staff members (Security, Database, Linux Sys Admin, Facility Management, Documentation, Property Management and Purchasing) who are shared with 8 other projects in our computing facility

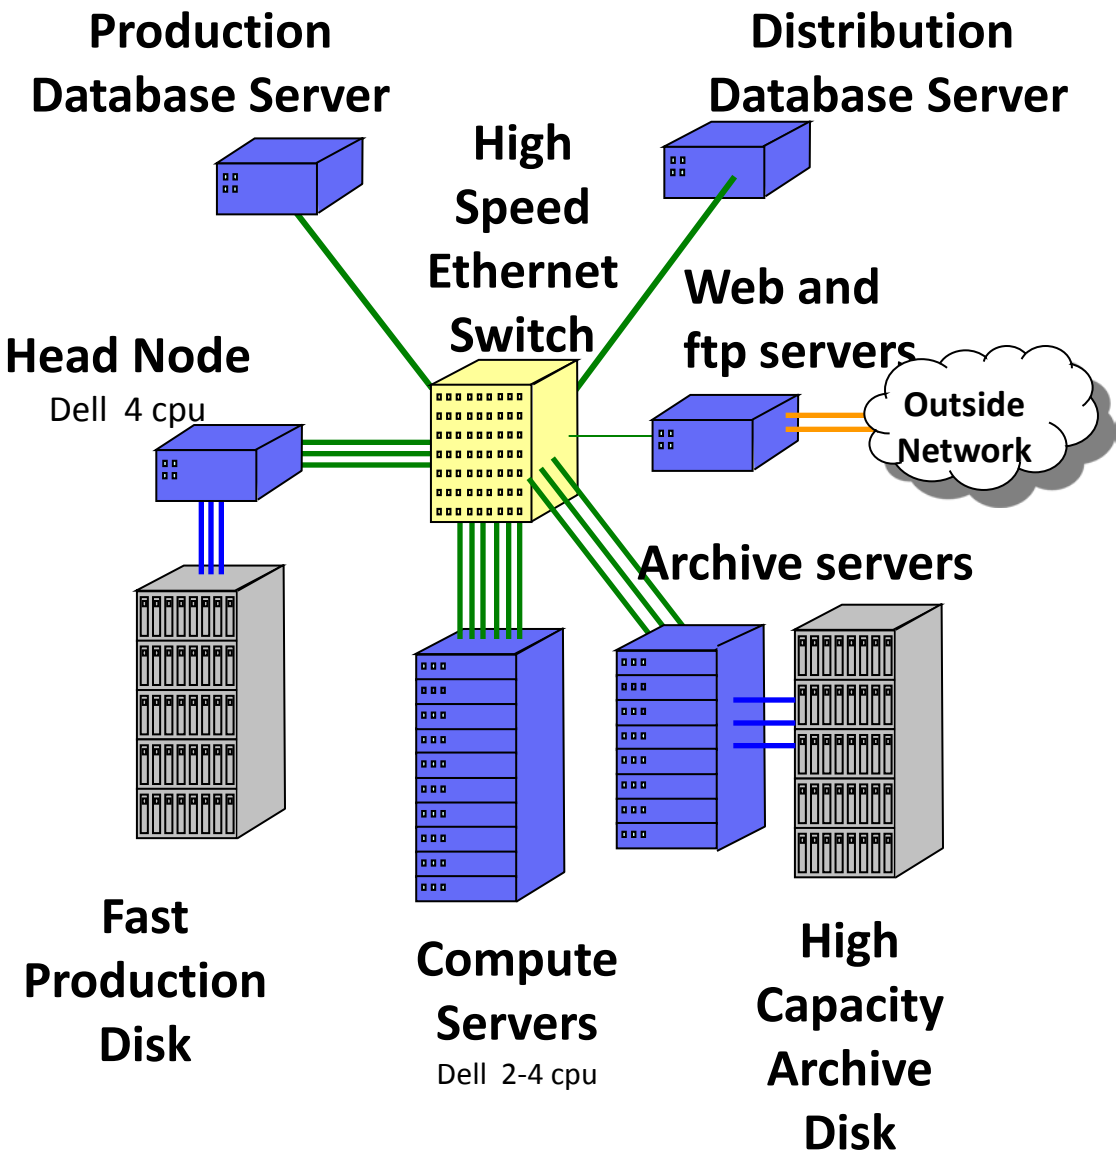

Some of the 1,000 servers

Level 1 and Atmosphere Archive and Distribution System

- Web-based search and order at http://ladsweb.nascom.nasa.gov/
- Online archive of MODIS Level 1, Atmosphere and Land products can also be accessed using ftp though a directory structure organized as:

  /allData/Reprocessing Collection #/Data Product/Year/Day
- Built upon the MODAPS framework to support post processing of products including the following operations: subset, sub-sample, mosaic, mask, parameter selection, geographic reprojection and format conversion)
- The archive also includes on-demand products which are produced when ordered by the end-user
- Web services allow machine to machine access for functions on web site
- Separating data products into archive sets to accommodate storing the large number of test results in the online archive as well as multiple reprocessing campaigns
- Other interfaces to the online archive include: an iRODS server providing access to the MODIS atmosphere products for users of the NASA Center for Climate Simulation (NCCS) and a server for EPA scientists that works in concert with visualization and analysis software running on their desktop systems to subset, sub-sample and combine MODIS atmosphere products.

Half of the 5PB archive