NASA Supercomputing, Big Data, and a Vision for the Future Compute Services

Presented at
NASA **Big Data Task Force Meeting**

January 16, 2016

Tsengdar Lee – tsengdar.j.lee@nasa.gov, NASA Headquarters
High Performance Computing Program Manager
with HECC and NCCS teams

Agenda

- **NASA Supercomputing Program**
- **Changing Workload from Big Compute to Big Data**
- **Big Compute + Big Data = New Compute Service**

High-end Computing at NASA

- **The mission of NASA's High-End Computing (HEC) Program is to:**
  - Plan and provision HEC systems and services to support NASA's mission needs.
  - Operate and manage these HEC resources for the benefit of agency users, customers, and stakeholders.

•The HEC Program manages two projects, which serve all three NASA Mission Directorates: Aeronautics (ARMD), Human Exploration and Operations (HEOMD), Science (SMD), and Space Technology (STMD).

HECC provides world-class high-end computing, storage, and associated services to enable scientists and engineers supporting NASA missions to broadly and productively employ large-scale modeling, simulation, and analysis to achieve successful mission outcomes. It supports 1,200 users from NASA and around the U.S. HECC is run by the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

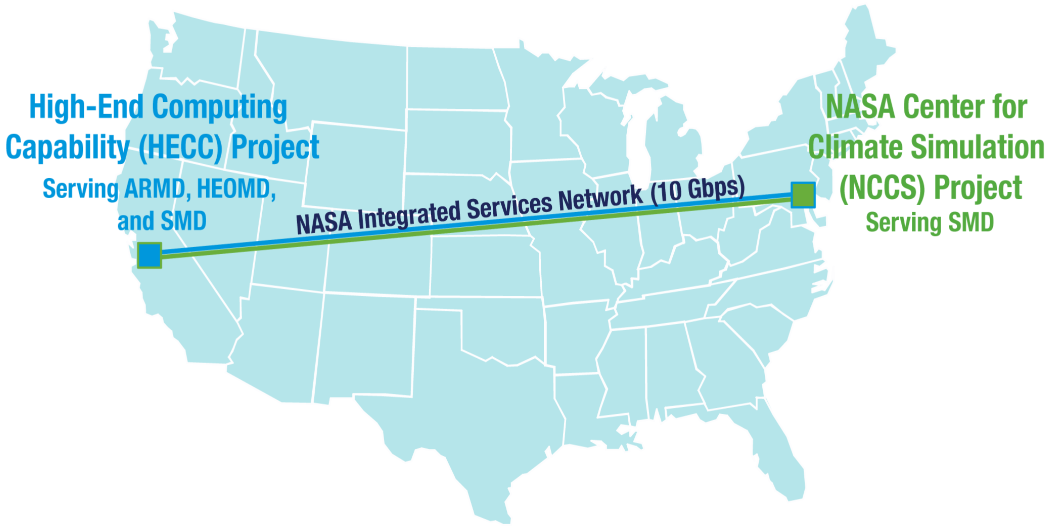

NCCS offers an integrated set of supercomputing, visualization, and data interaction technologies to enhance NASA capabilities in weather and climate prediction and enable future scientific discoveries that will benefit humankind. It serves hundreds of users at NASA centers, laboratories, and universities across the U.S. NCCS is run by the Computational and Information Sciences and Technology Office (CISTO) at Goddard Space Flight Center.

Strategic Support for NASA Programs

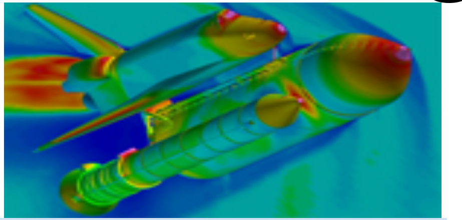

External tank redesign

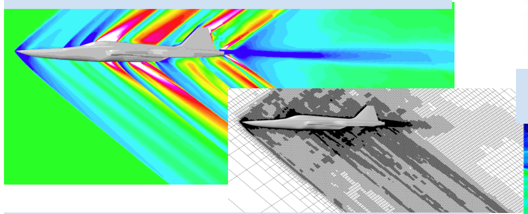

Sonic boom optimization

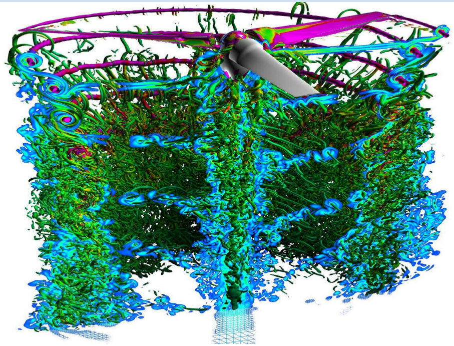

Rotary wing aerodynamics

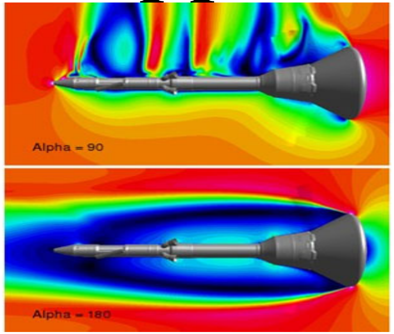

Launch abort system

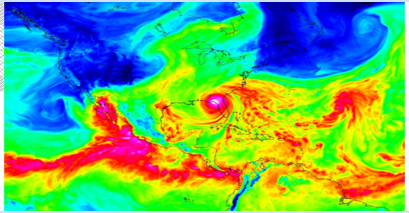

Hurricane prediction

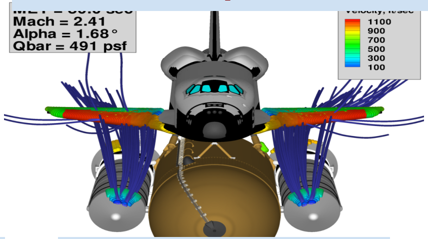

Debris transport

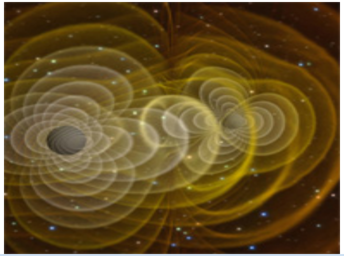

Merging black holes

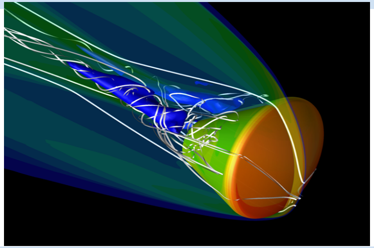

Orion/MPCV reentry

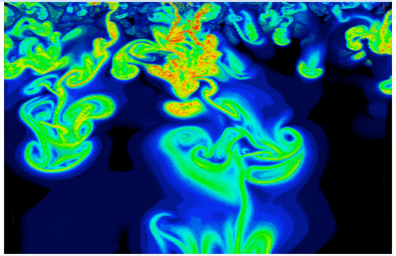

Solar magnetodynamics

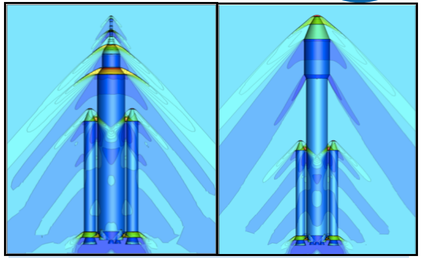

SLS vehicle designs

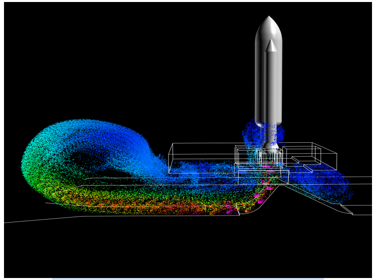

Flame trench

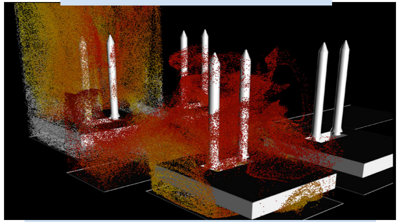

SRB burn in VAB

Pleiades Surpasses a Half Billion SBUs/ $300M* Delivered to Users

- NASA's flagship supercomputer, Pleiades, delivered 16,535,458 Standard Billing Units (SBUs) in August, 2015, bringing the Pleiades total to 505,389,818 SBUs; nearly $300M over 8 years.

- HECC deployed multiple expansions to Pleiades, increasing its computational capability by about 8x since its installation in 2008. Pleiades now delivers in 6 weeks the same amount of processing that its predecessor, Columbia, delivered in its entire 9-year life span, which ended in 2013.

- Pleiades already delivered ~25x the processing Columbia delivered. Columbia delivered approximately 10x the combined processing of all NAS supercomputers before it, going back to the establishment of the NAS facility in 1987.

- Among recent projects using the most computer time on Pleiades: flow computations for the Space Launch System configuration; validation of Kepler mission planet candidates; investigations to reduce the environmental impact of aircraft; and global simulations to help monitor ocean, sea ice, and atmospheric systems.

** 1 SBU equals 1 hour of a Pleiades Westmere 12-core node and costs $0.28/SBU in FY15, $0.26/SBU in FY16 (avg $.59/SBU per FY.)*

Mission Impact: The HECC project's continuous increase in computational capability and capacity is essential for all mission directorates to meet the high demand for supercomputing resources to accomplish their goals and objectives.

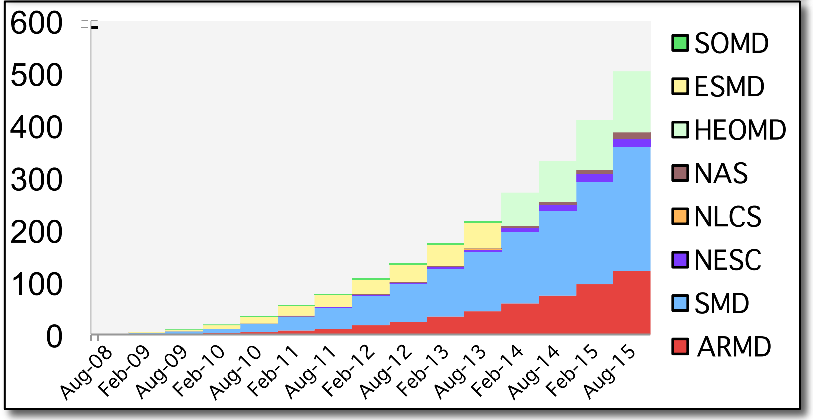

Cumulative chart in millions since the first mission directorate jobs ran on Pleiades from August 2008 to August 2015. One hour of computer time is defined as one hour of computing on a 12-core Westmere node. Therefore, Pleiades' 500-million hour milestone is equivalent to delivering 6-billion Westmere core hours.

All Missions HEC Capacity Shares in SBUs

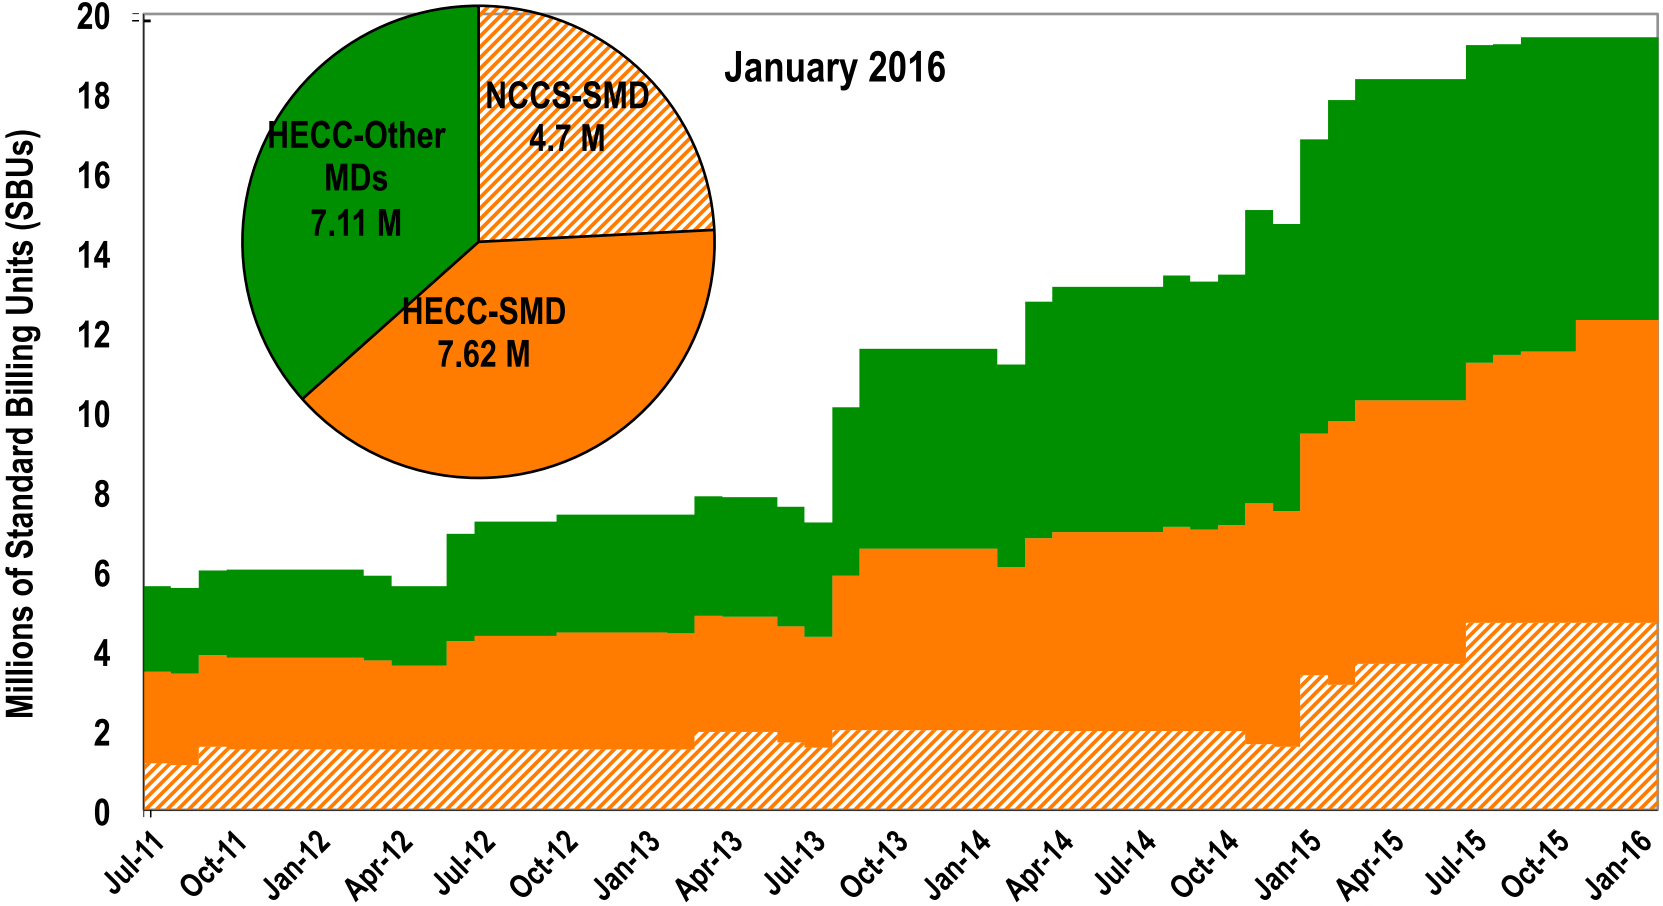

Agenda

- NASA Supercomputing Program
- **Changing Workload from Big Compute to Big Data**
- Big Compute + Big Data = New Compute Service

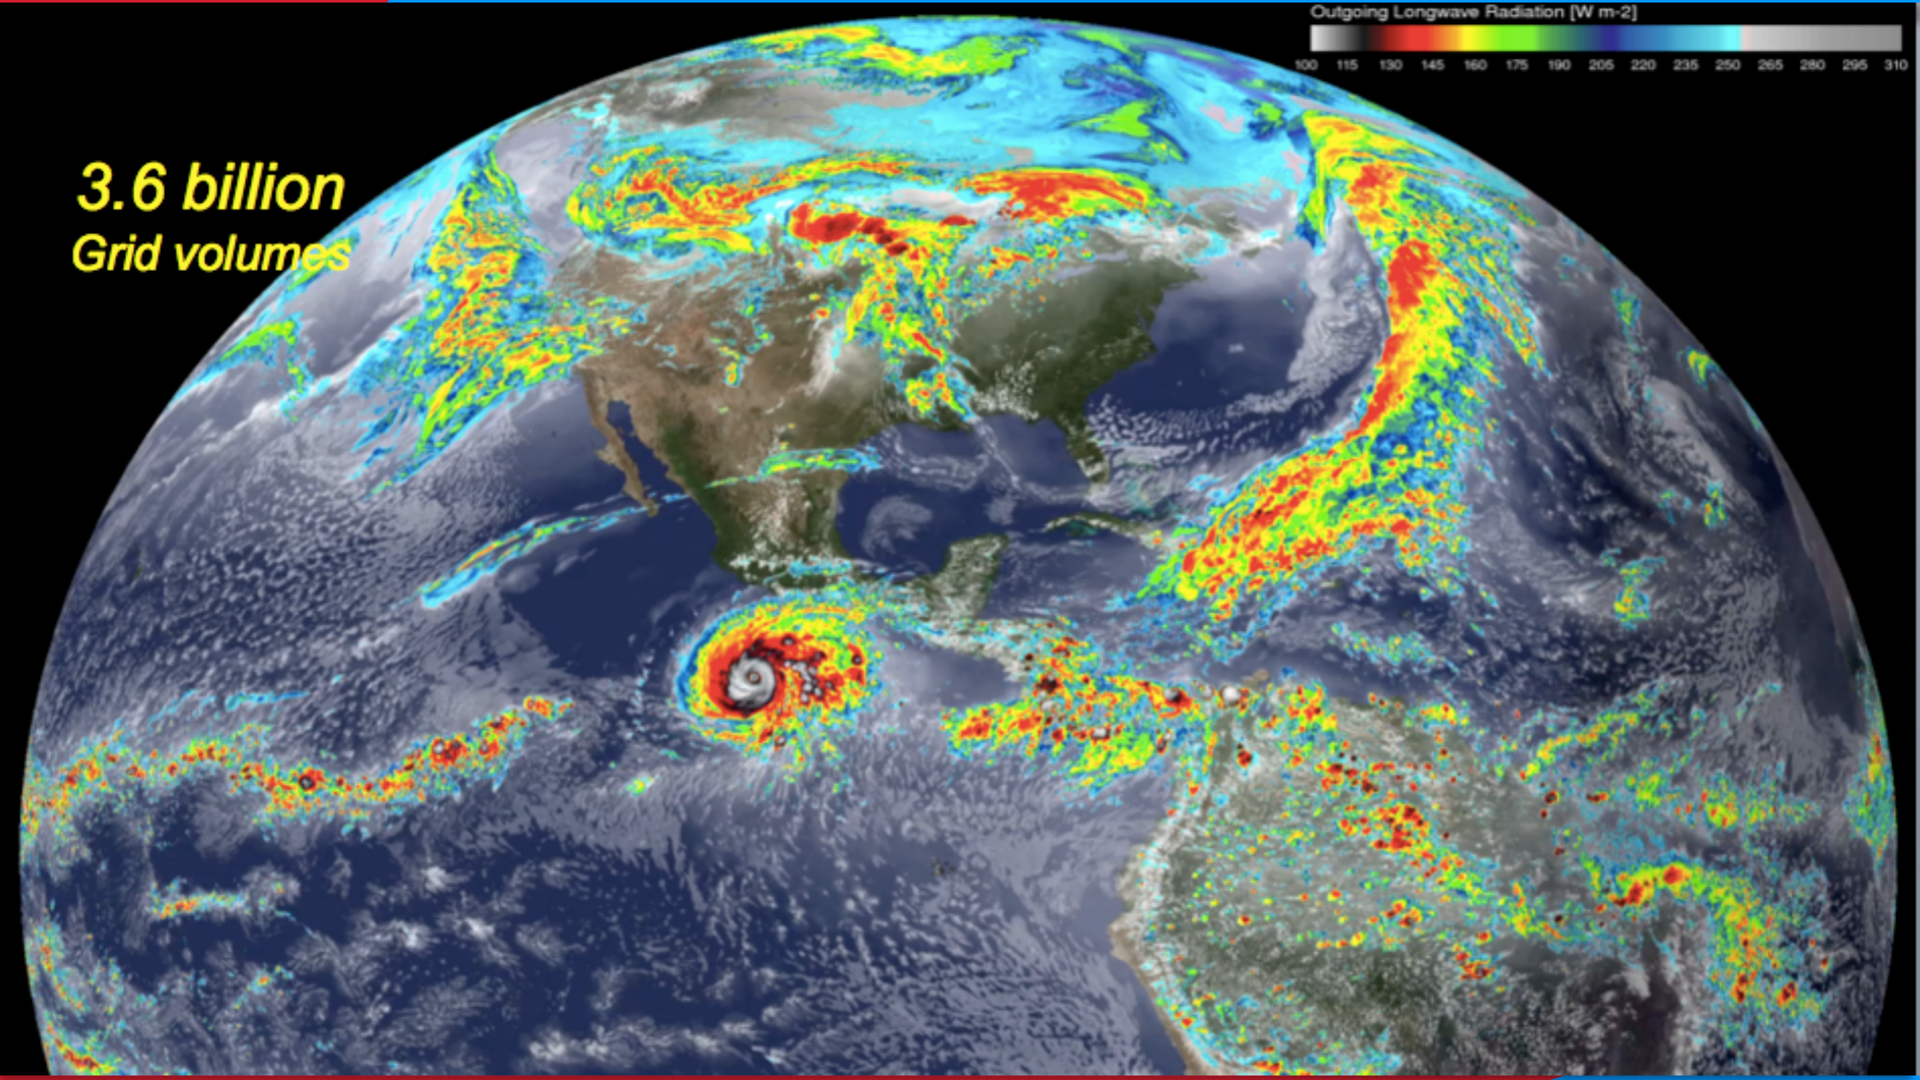
3.6 billion
Grid volumes
Outgoing Longwave Radiation [W m-2]
100 115 130 145 160 175 190 205 220 235 250 265 280 295 310

HECC Experts Deliver Crucial Support for Validation of New Kepler Planets

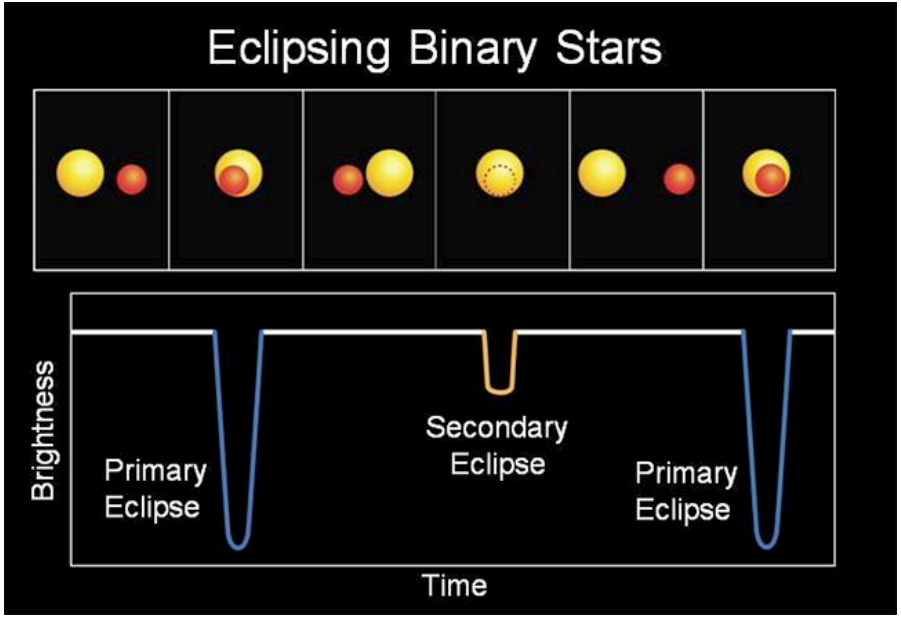

Diagram illustrating how eclipsing binary stars vary in brightness. When seen from certain angles, one star may eclipse the other, causing a reduction in brightness. Wikipedia

- HECC data analysis, visualization, and computer science experts provide invaluable code optimization and infrastructure support for validating the Kepler Mission's discoveries of Earth-size planets:
  - Enhanced and automated Kepler science codes, enabling rapid turnaround of analyses between the Kepler science processing/data analysis pipeline and Pleiades to run transiting planet searches and data validation.
  - Made major enhancements to the JAVA process management code to improve the efficiency of some modules.
  - Significantly optimized arithmetic-intensive pipeline modules to work with Kepler's MATlab equation solver.

Pleiades is Critical to Meet Kepler Resource Demands and Turnaround

NASA's supercomputer, Pleiades, with 211,872 cores, 11,312 nodes, and 724 terabytes of memory, and 4 generations of Intel processors, is currently ranked as the seventh most powerful system in the U.S. (Nov. 2015 TOP500 List).

- HECC's Pleiades supercomputer capacity and capability provides crucial resources required to analyze Kepler data that has grown 10x over the life of the project. Project usage, as of of May 6, 2015:
  - Total number of SBUs to date since the project began in 2011: 20M
  - Total Number of CPU hours: 193,045,109
  - Working Storage: 165 terabytes
  - Archive storage: 7 terabytes of data
- Pleiades is essential for completing the enormous, increasingly complex planetary transit searches of 200,000 stars observed by Kepler in a matter of weeks. This timeframe includes not only pure processing, but data transfers, network adjustments, handling job failures and reruns, and more.

Earth Observations

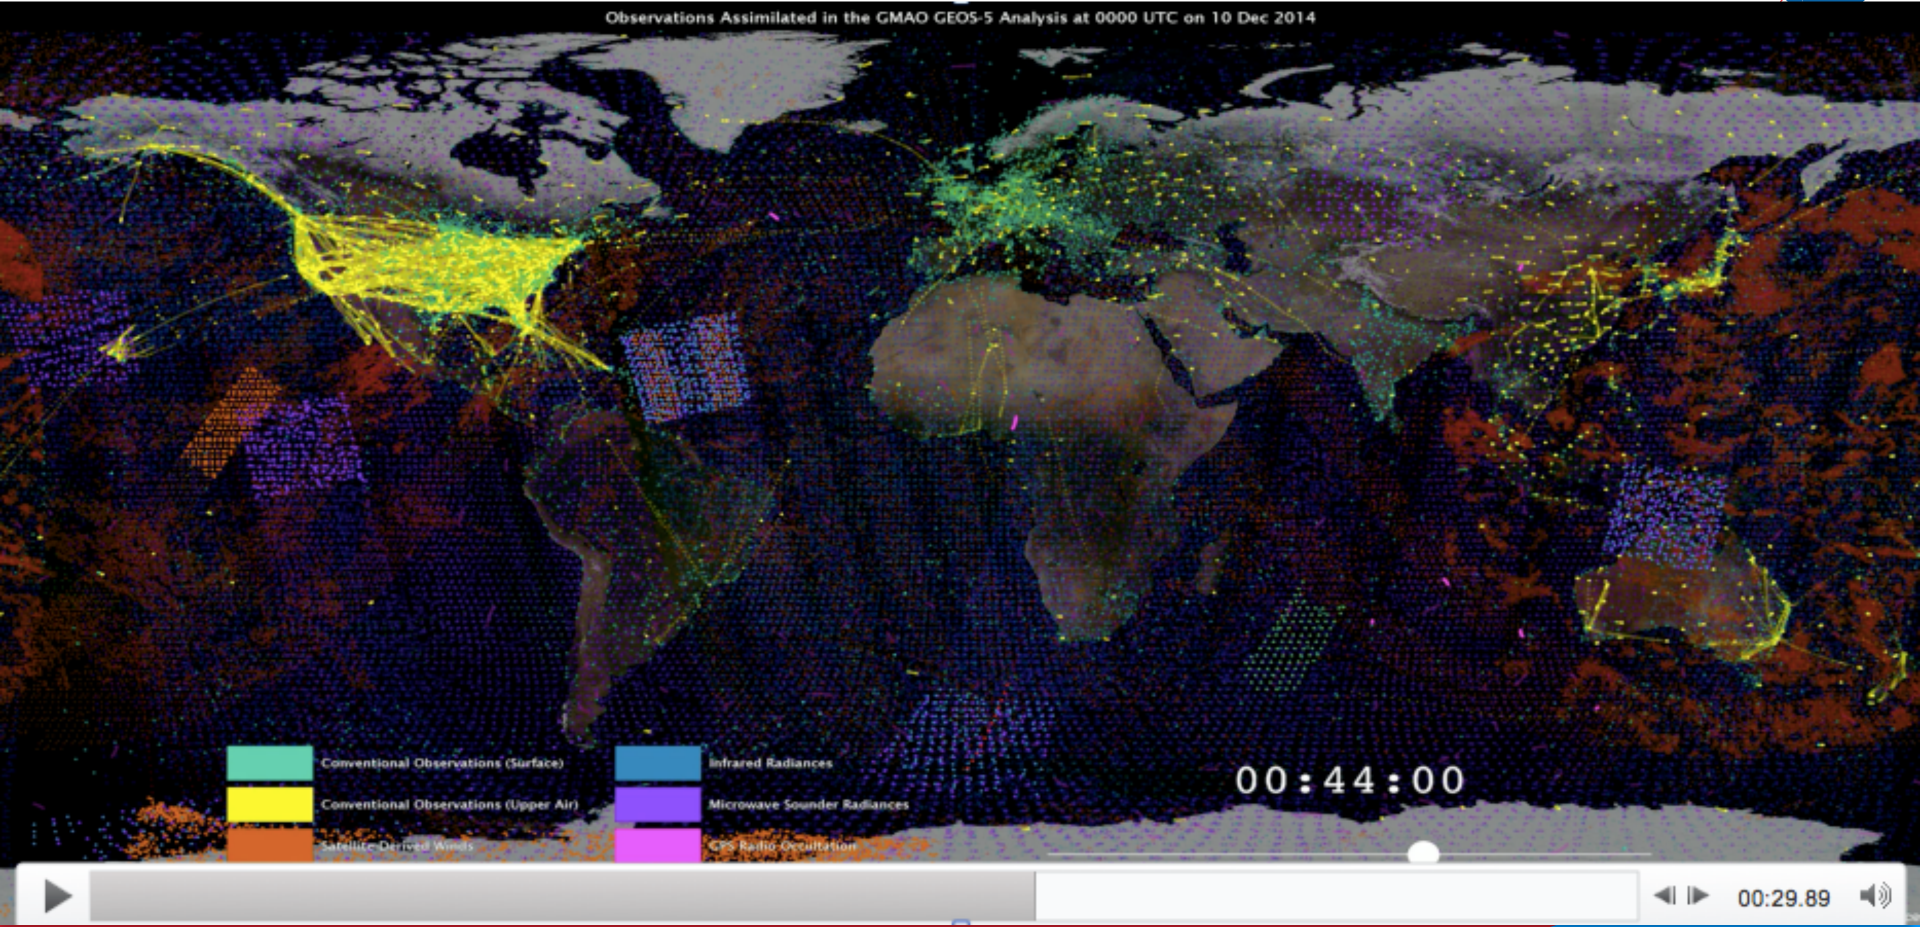

OCO-2 Data Processing Challenges

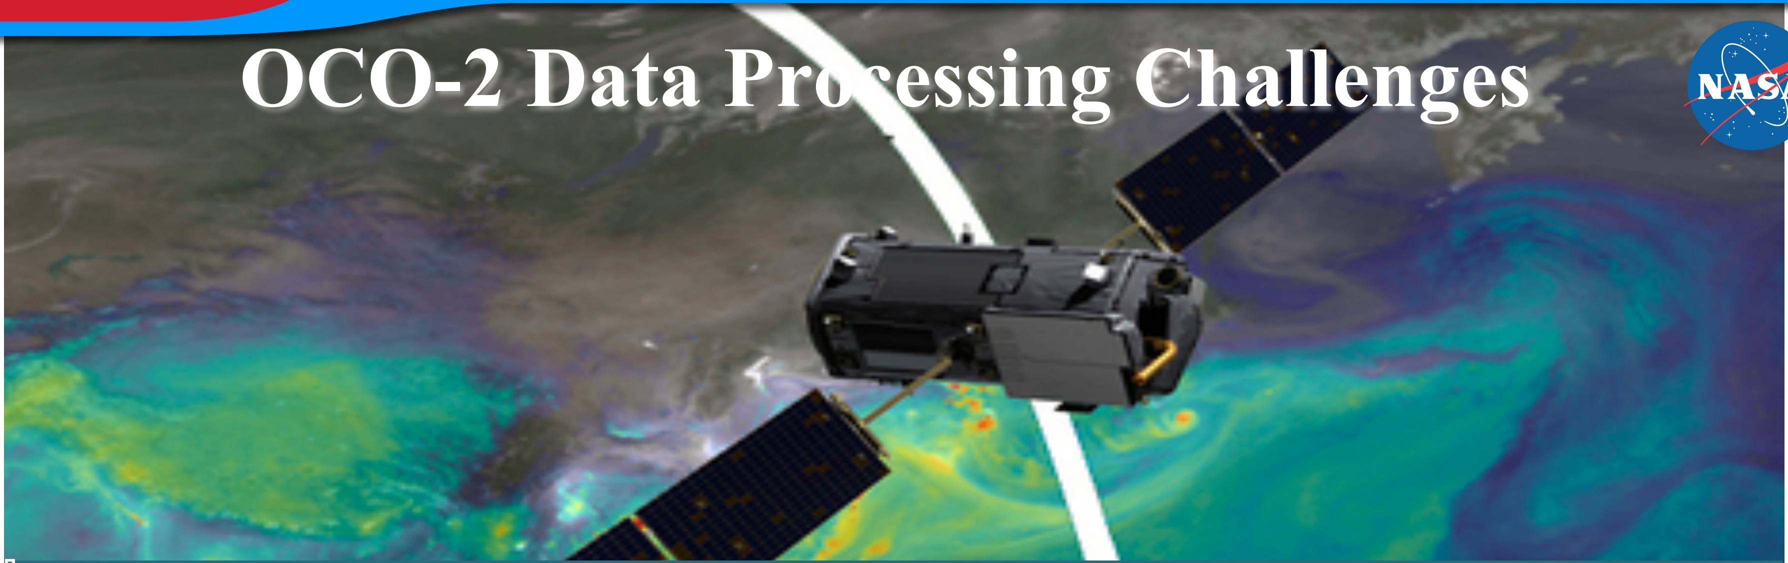

Mission Phase	Scope	Local cluster usage	Pleiades usage
Initial routine processing	6% of 1M soundings/day	Daily routine processing plus science/algorithm work	
First reprocessing campaign	6% of 7 months of data in 3 weeks	Process 5 months plus daily routine processing	Initial Operations involvement: process 6% of two months
Algorithm update validations (each system update)	Up to 6 M soundings	20% of scope plus daily routine processing, science/algorithm work	80% of scope for timely analysis
Final reprocessing campaign	800 M soundings - all cloud-free (~30%) in 2 mo.	8% of scope plus daily routine processing, science/algorithm work	92% of scope for timely closeout

Pleiades-Enabled Reprocessing by OCO-2 Provides a Consistent Ten-month Data Record

- The OCO-2 team at JPL recently completed a major reprocessing campaign covering ten months of the mission, enabled by the extensive use of NASA's Pleiades supercomputer. About 60% of the campaign was pushed through Pleiades, including 18.6 million instrument soundings (using 1.4 million CPU-hours).
- The effort utilized the highly-reliable transfer tools and up to 500 Haswell nodes concurrently. Because the Full Physics retrieval process operates on the soundings independently to extract X_{CO2}, it is well suited to the Pleiades architecture.
- The added computing resources allowed the team to complete the campaign much faster than initially planned and get a consistent, well-calibrated and comprehensive data set out to the public quickly.
- Science team members and the larger community are now mining the data record at both global (see maps at right) and regional scales.

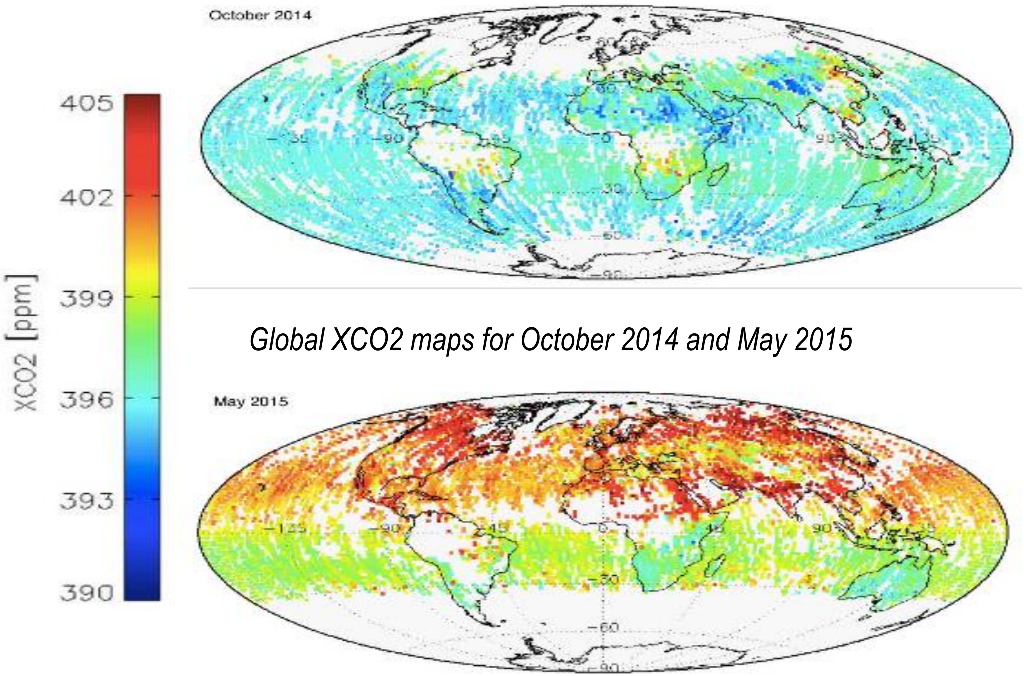

Global XCO2 maps for October 2014 and May 2015

HECC is Used to Clear 5 Years of UAVSAR Data Processing Backlog and Reduce Latency

Impact: Enabled by NASA's Pleiades supercomputer, UAVSAR was able to clear a 5-year processing backlog to deliver surface deformation products to scientists for research.

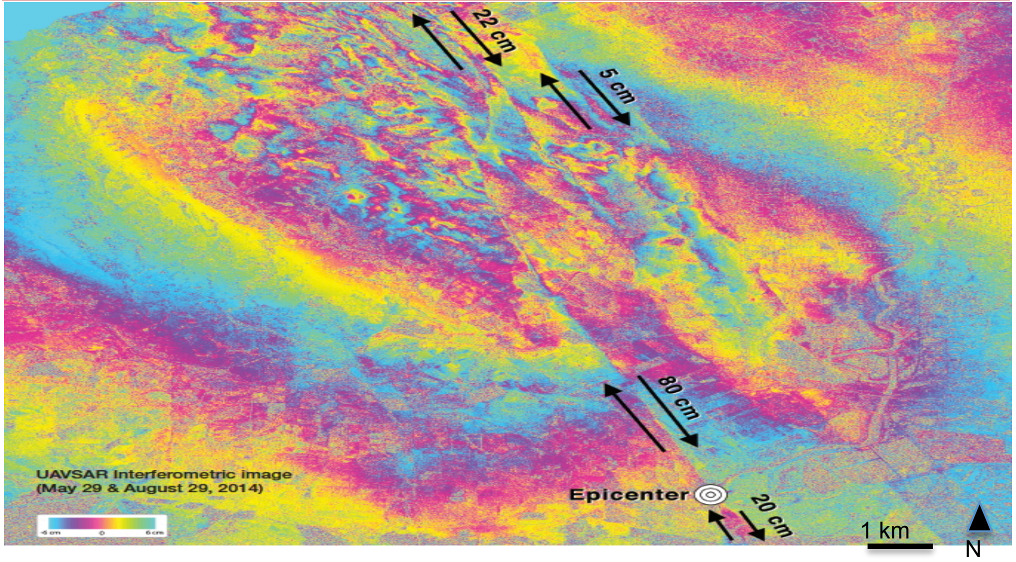

- Scientists are using repeat-pass interferometric (InSAR) data from Uninhabited Airborne Vehicle Synthetic Aperture Radar (UAVSAR) to study centimeter-scale surface deformation of earthquake faults, volcanoes, landslides, and glaciers.
- The computationally intensive InSAR processing code was ported to Pleiades Supercomputer at Ames to take advantage of the large number of processing nodes, each with more than 32 GB RAM, and ample data storage.
- With Pleiades and processor automation, HECC was able to clear a 5-year InSAR processing backlog in 6 months and reduce processing latency to 2 weeks.

UAVSAR Interferometric image (May 29 & Aug 29, '14) of the M6.0 South Napa Earthquake, CA. Colors in the image represent the amount of ground motion between the two flights from the radar's point of view. Linear discontinuities in the colors indicate locations where a surface rupture is highly likely, giving government agencies the exact location of the fault traces that shifted during the earthquake and how they relate to levees, buildings, and vital infrastructures, as well as to help provide a fundamental understanding of earthquakes processes. ***Andrea Donnellan, NASA/JPL***

Agenda

- NASA Supercomputing Program
- Changing Workload from Big Compute to Big Data
- **Big Compute + Big Data = New Compute Service**

Shifting Technologies Toward Data

High Performance Computing

- Shared everything environment
- Very fast networks; tightly coupled systems
- Cannot lose data
- Big data (100 PBs)
- Bring the data to the application
- Large scale applications (up to 100K cores)
- Applications cannot survive HW/SW failures
- Commodity and non-commodity components; high availability is costly; premium cost for storage

What technologies can be leveraged from both sides to create a more data centric and data analytic environment?

Large Scale Internet

- Examples: Google, Yahoo, Amazon, Facebook, Twitter
- Shared nothing environment
- Slower networks
- Data is itinerant and constantly changing
- Huge data (Exabytes)
- Bring the application to the data
- Very large scale applications (beyond 100Ks)
- Applications assume HW/SW failures
- Commodity components; low cost storage

HPC as a private cloud (IaaS, PaaS) service!

Data Centric HPC, Big Data and IT Environment

Data Sharing and Publication

- Capability to share data & results
- Supports community-based development
- Data distribution and publishing

Code Development

- Code repository for collaboration
- Environment for code development and test
- Code porting and optimization support
- Web based tools

User Services

- Help Desk
- Account/Allocation support
- Computational science support
- User teleconferences
- Training & tutorials

DATA Storage & Management

Global file system enables data access for full range of modeling and analysis activities

Analysis & Visualization

- Interactive analysis environment
- Software tools for image display
- Easy access to data archive
- Specialized visualization support

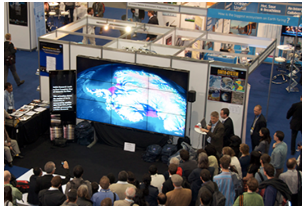

Data Transfer

- Internal high speed interconnects for HPC components
- High-bandwidth to data center users
- Multi-gigabit network supports on-demand data transfers

Security

HPC Computing

- Large scale HPC computing
- Comprehensive toolsets for job scheduling and system monitoring

Data Archival and Stewardship

- Large capacity storage
- Tools to manage and protect data
- Data migration support

New Architecture

Goal

- Demonstrate support for the scientific investigation and engineering optimization without moving data

Concept

- Unified data analysis platform that provides a co-location of data, compute, data management, and data services
- Data storage surrounded by a compute cloud
- Large amount of data storage, high performance compute capabilities, very high speed interconnects
- High speed network to internal **and external** communities

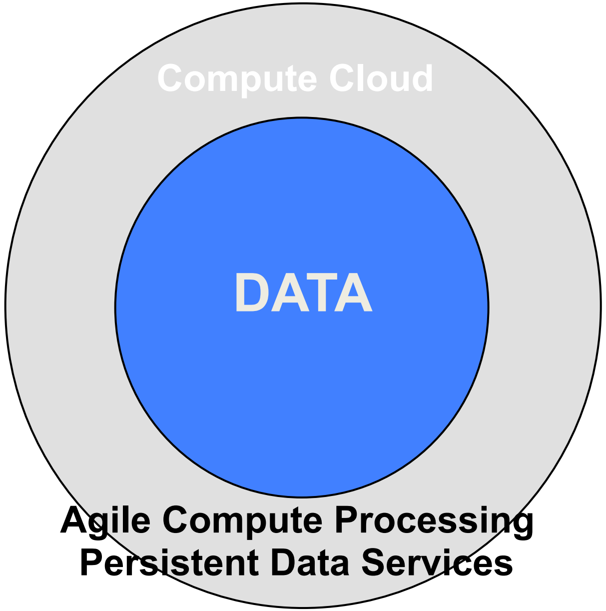

Science Cloud Architecture

- Agile, high level of support
- Storage is 90% full prior to use
- The system owns the data
- The users own their analysis
- Extensible storage; build and expand as needed
- Persistent data services built in VMs, Containers, or bare metal
- Create purpose build VMs for specific science projects
- Image management

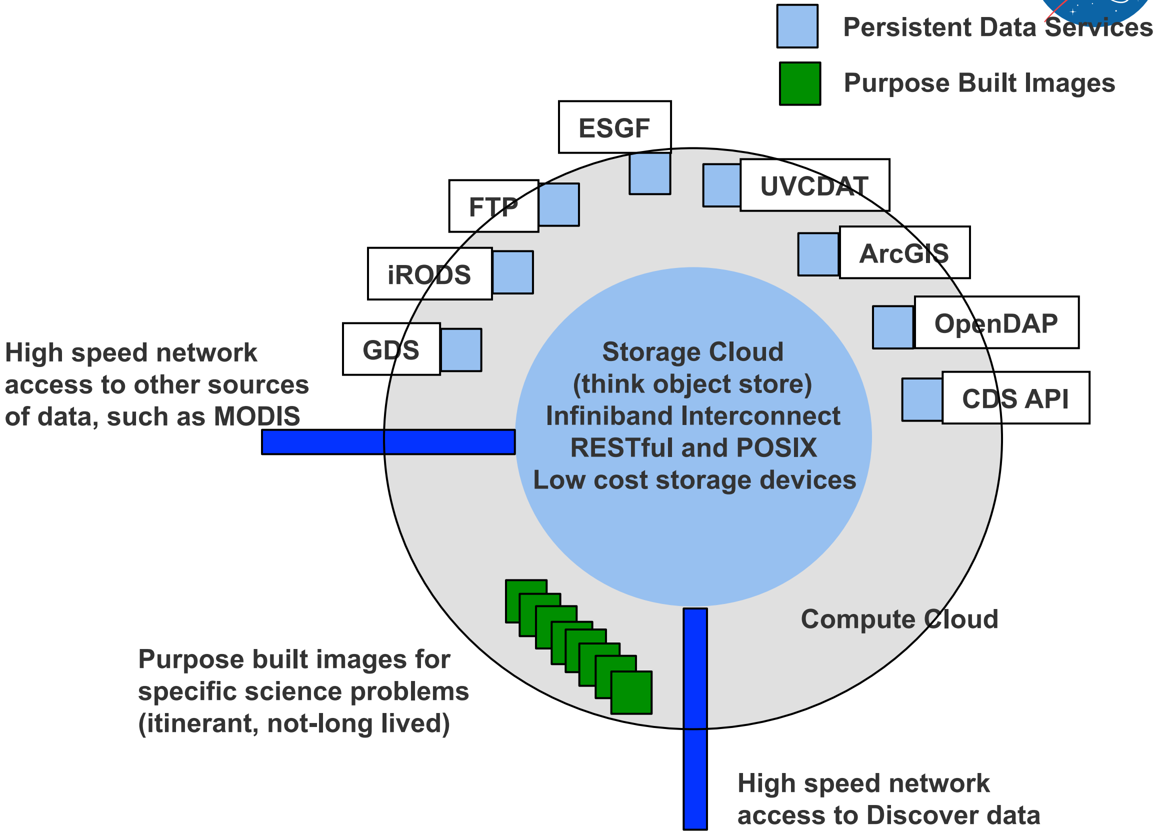

What Are We Doing to Build This?

Data Management System

- NASA Center for Climate Simulation Data Management System moving toward an iRODS-Based Approach to Scientific Data Services

Data Analytics Service

- Supports Map-reduce project for large scale data analysis projects

Multiple projects started such as CDaaS, Kepler, OCO-2, UAVSAR, ABoVE Science Cloud

Technologies of Interest

- **Adoption of Coprocessors in NASA Applications**
  - Tools still have a long way to go
  - Exploitation of concurrency requires applications work
- **Science Cloud and Virtualization**
  - Virtualized Infiniband for very high speed interconnects (near wire speed)
  - High speed object storage environments
- **High Performance Object Stores**
  - Shift from POSIX compliant to REST-ful file systems
- **MapReduce, Hadoop, and Spark**
  - Very large adoption throughout industry
  - Robust eco system built around Hadoop
  - High speed in-memory data processing with Spark
- **SciDB**
  - Scientific Database (array based badabase)
- **Data Management**
  - Use of the Integrated Rule-Oriented Data System (iRODS) for data management
  - Apache OODT
  - Apache Climate Work Bench
- **Data Services and Data Analytics**
  - Robust data services – shift to server side analysis
  - Remote visualization of large-scale data sets
  - Data Services API with Plug & Play Algorithms

Science Benefit of a Science Cloud

Loboda et al MODIS cloud dynamics and burned area mapping

- Dramatically reduces download time by utilizing multiple high speed connections simultaneously
- Data processing times will be reduced:
  - No need to cycle through inputs, all can stay online at once
  - Multiple instances can be run simultaneously
- Enables the option for reprocessing as the algorithm evolves
- Extra time for QA/verification of the products
- Enables completion of circumpolar data
- Solves data storage and data management issues freeing up personnel time
- Reduction of effort because Loboda will also use the Landsat from Carroll

NASA Releases Detailed Global Climate Change Projections

- CMIP5 archive downscaled using the Bias Correction and Spatial Disaggregation (BCSD) approach (Wood et al., 2002; Thrasher et al. 2012)
- Global daily temperature and precipitation scenarios from 1950-2100 at 0.25 degree spatial resolution (~25km x 25km)
- 21 coupled General Circulation Models (global climate models)
- 2 Representative Concentration Pathways (RCPs) - RCP 4.5 and RCP 8.5
- 12 terabyte dataset available via THREDDS at NCCS; visualizations developed by OpenNEX community and available at www.climateinternational.org
- Used ~50,000 SBUs on Pleiades
- More information available at:
    - https://cds.nccs.nasa.gov/nex-gddp/
    - https://nex.nasa.gov/nex/projects/1356/

Above: Chief Scientist Dr. Ellen Stofan presented a talk during launch of the Climate Services for Resilient Development Partnership describing NASA's NEX-GDDP release, "Visualization of Climate Services for Development." ***Below:*** NEX-GDDP provides forecast of how global temperature and precipitation might change up to 2100 (shown here for the RCP 8.5 scenario).

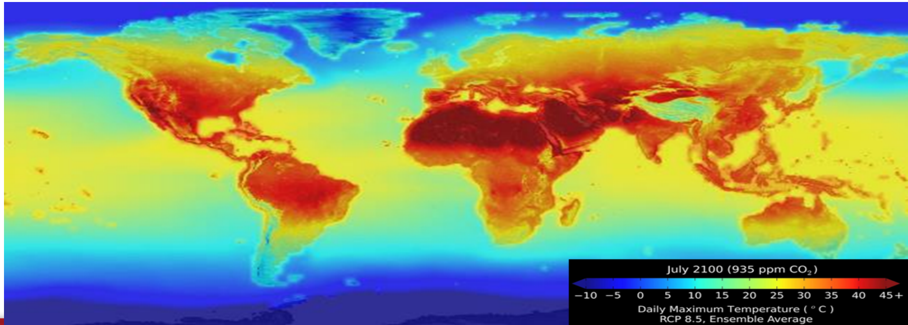

Global Monitoring of Fires in Tropical Forests

Fires in tropical forests during 2001-2014

1 million sq. km. burned area found in tropical forests

- ***more than three times the total area reported by state-of-art NASA products***.

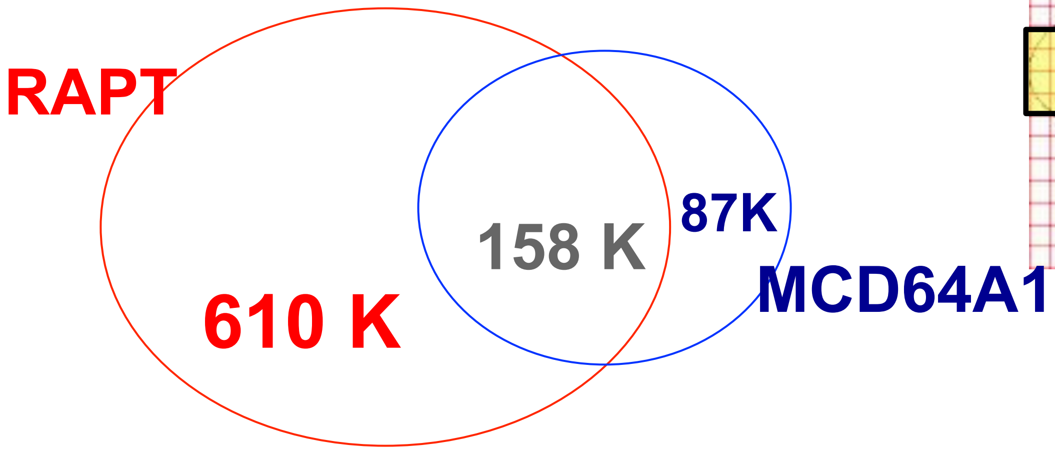

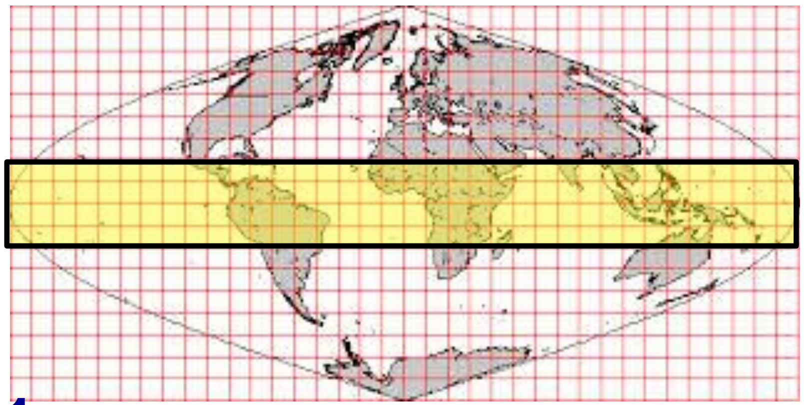

Validation: Multiple sources

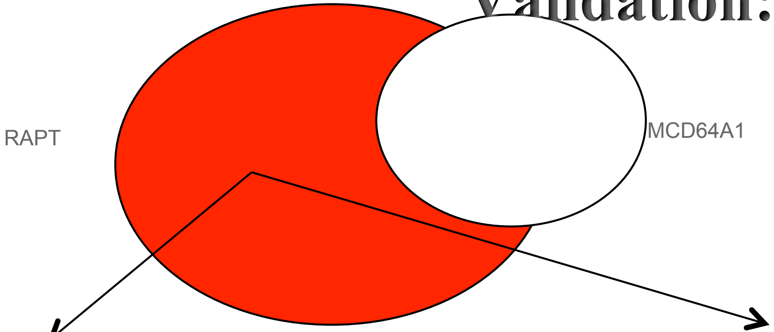

Change in Vegetation series

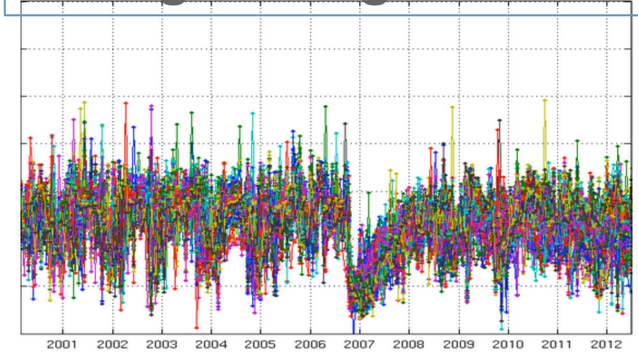

Burn scar in Landsat composite

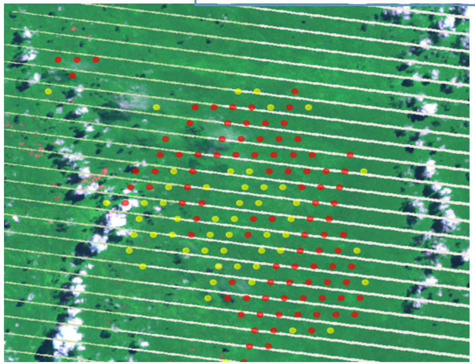

Before Fire Event

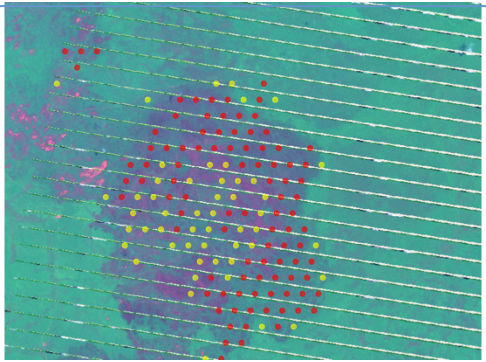

After Fire Event

Validation confirmed that the additional burned areas detected using RAPT are actual burns missed by state-of-art products

Summary

- High-end computing (HEC) supported all the NASA Mission Directorates.
- HEC can be viewed as PaaS cloud computing.
- Recent R&D efforts include satellite data processing and data analytics.
  - By coupling HEC and Cloud Computing, we are creating “best of breed” computing service environment.
  - There are a lot of benefit from the use of object stores to be harvested.
- Coupling HEC and data centers has produced significant scientific value.
- Scaling out data processing has significantly reduced the time to solution in data processing and data analytics.

Thank You!

Tsengdar Lee, Ph.D.
High-end Computing Program Manager
Weather Focus Area Program Scientist
NASA Headquarters
tsengdar.lee@nasa.gov

Intel® Head in Clouds Challenge Award to Estimate Biomass in South Sahara

Using NGA data to estimate tree and bush biomass over the entire arid and semi-arid zone on the south side of the Sahara

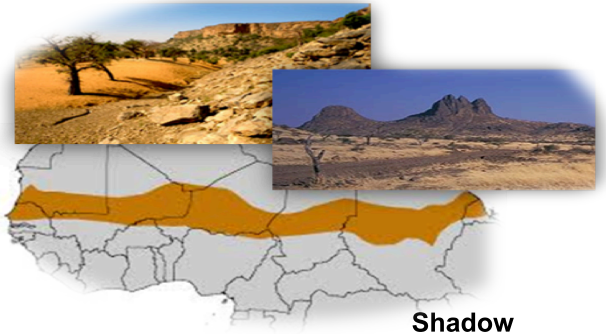

Project Summary:

- Estimate carbon stored in trees and bushes in arid and semi-arid south Sahara
- Establish carbon baseline for later research on expected CO_2 fertilization of photosynthesis will first be manifested in the arid and semi-arid zones, because bushes and trees will use less water to grow, thus growing more.

Principal Investigators:

- Dr. Compton J. Tucker, NASA Goddard Space Flight Center and Dr. Paul Morin, Univ. Minnesota

Resources From the Partners:

- Intel awards professional services
- AWS provides cloud computing resources
- Cycle Computing brings cloud resource management software

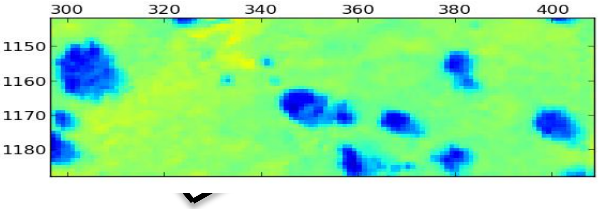

NGA Imagery representing tree & shrub automated recognition

Partners:

Existing Sub-Saharan Arid and Semi-arid Sub-meter Commercial Imagery

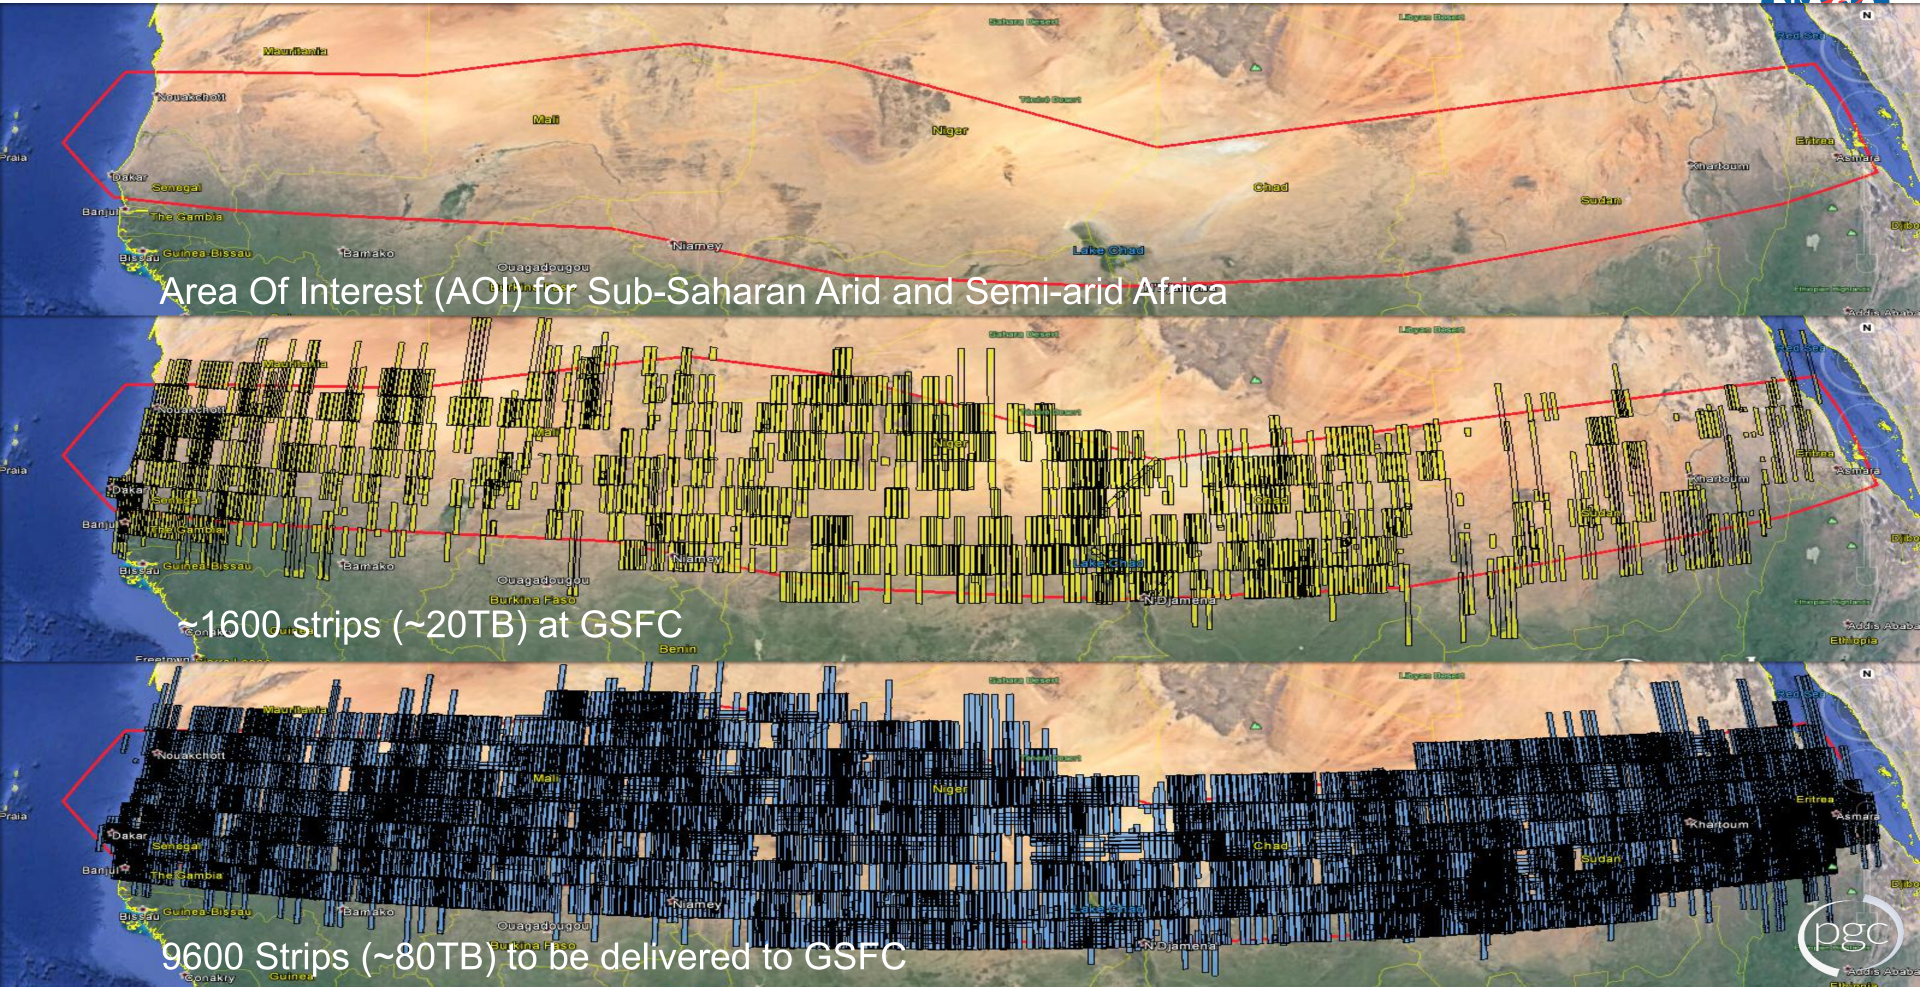

Panchromatic & Multi-spectral Mapping at the 40 & 50 cm scale

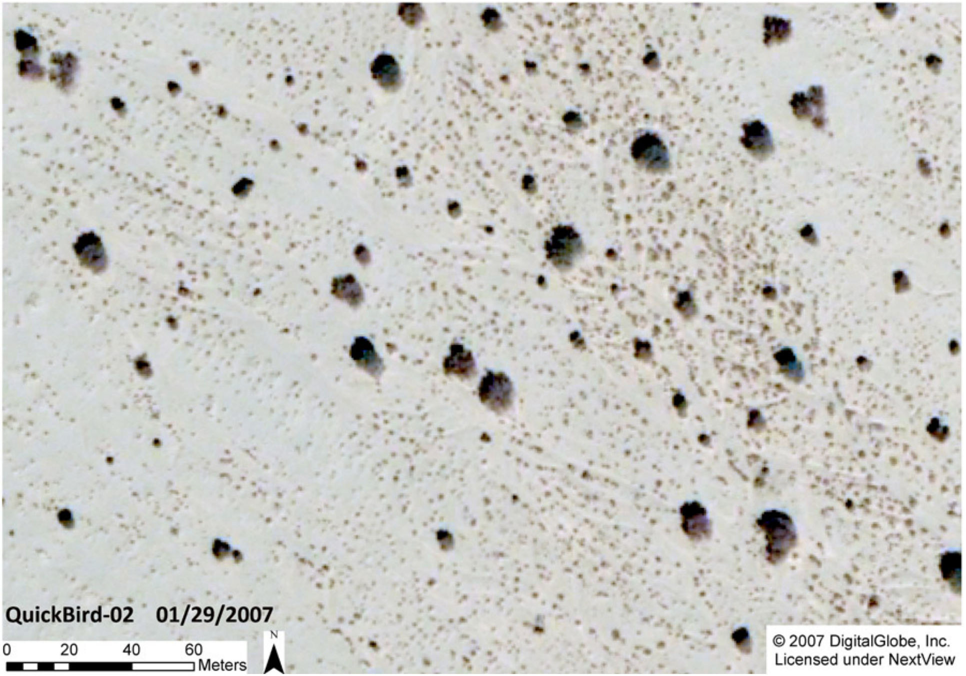

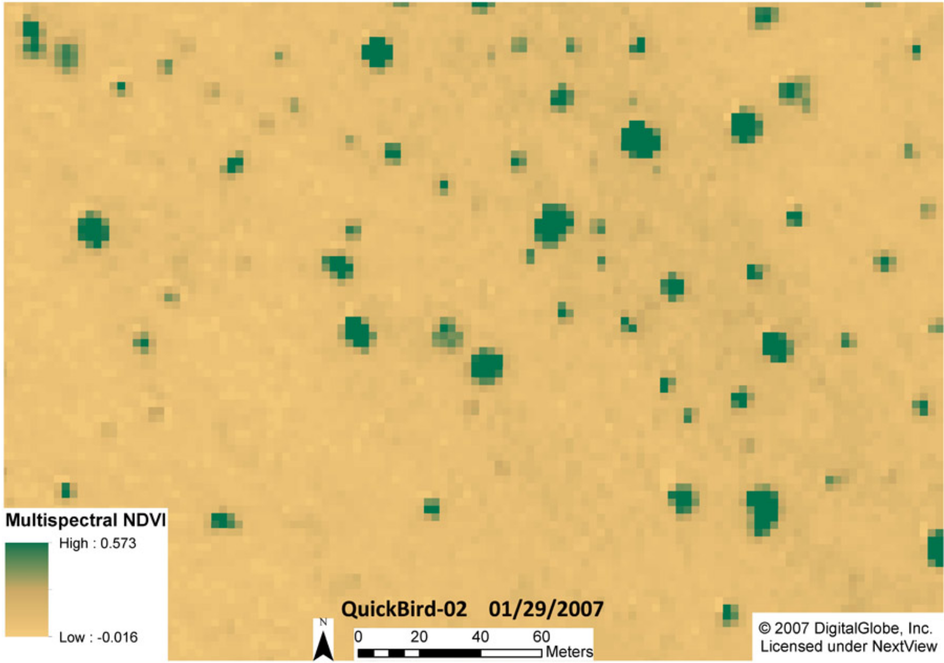

30