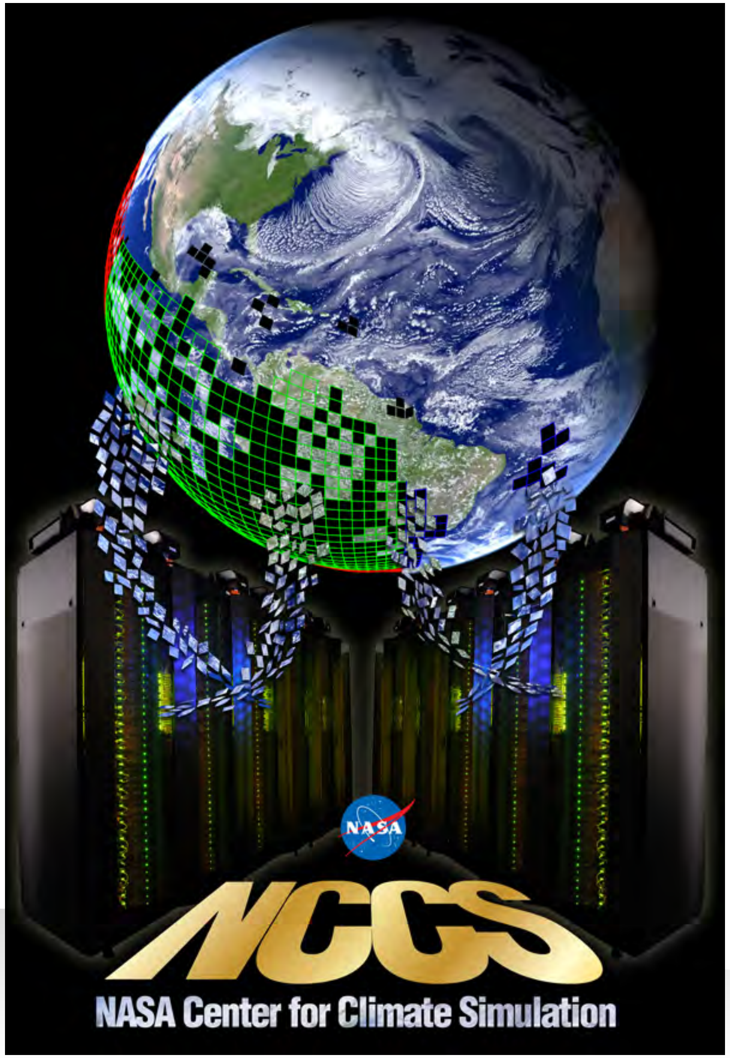

Evolving Storage and Cyber Infrastructure at the NASA Center for Climate Simulation

Ellen Salmon

Slides and content from Daniel Duffy, Carrie Spear, Scott Sinno, Garrison Vaughan, Michael Bowen
NASA Center for Climate Simulation (NCCS)

NASA Goddard Space Flight Center

Greenbelt, Maryland

Standard Disclaimers and Legalese Eye Chart

- **All Trademarks, logos, or otherwise registered identification markers are owned by their respective parties.**
- Disclaimer of Liability: With respect to this presentation, neither the United States Government nor any of its employees, makes any warranty, express or implied, including the warranties of merchantability and fitness for a particular purpose, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.
- Disclaimer of Endorsement: Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government. In addition, NASA does not endorse or sponsor any commercial product, service, or activity.
- The views and opinions of author(s) expressed herein do not necessarily state or reflect those of the United States Government and shall not be used for advertising or product endorsement purposes.
- All errors in this presentation are inadvertent and are the responsibility of the primary author.

AGENDA:

- About the NASA Center for Climate Simulation (NCCS)
- Recent History & Today
- Science Drivers
- Recent Developments
  - Advanced Data Analytics Platform (ADAPT)
  - Data Analytics Storage System (DASS)
  - Cloud-Bursting Project: Semi-Arid Carbon Sink
- Looking Toward Exascale

NASA Science Mission Directorate

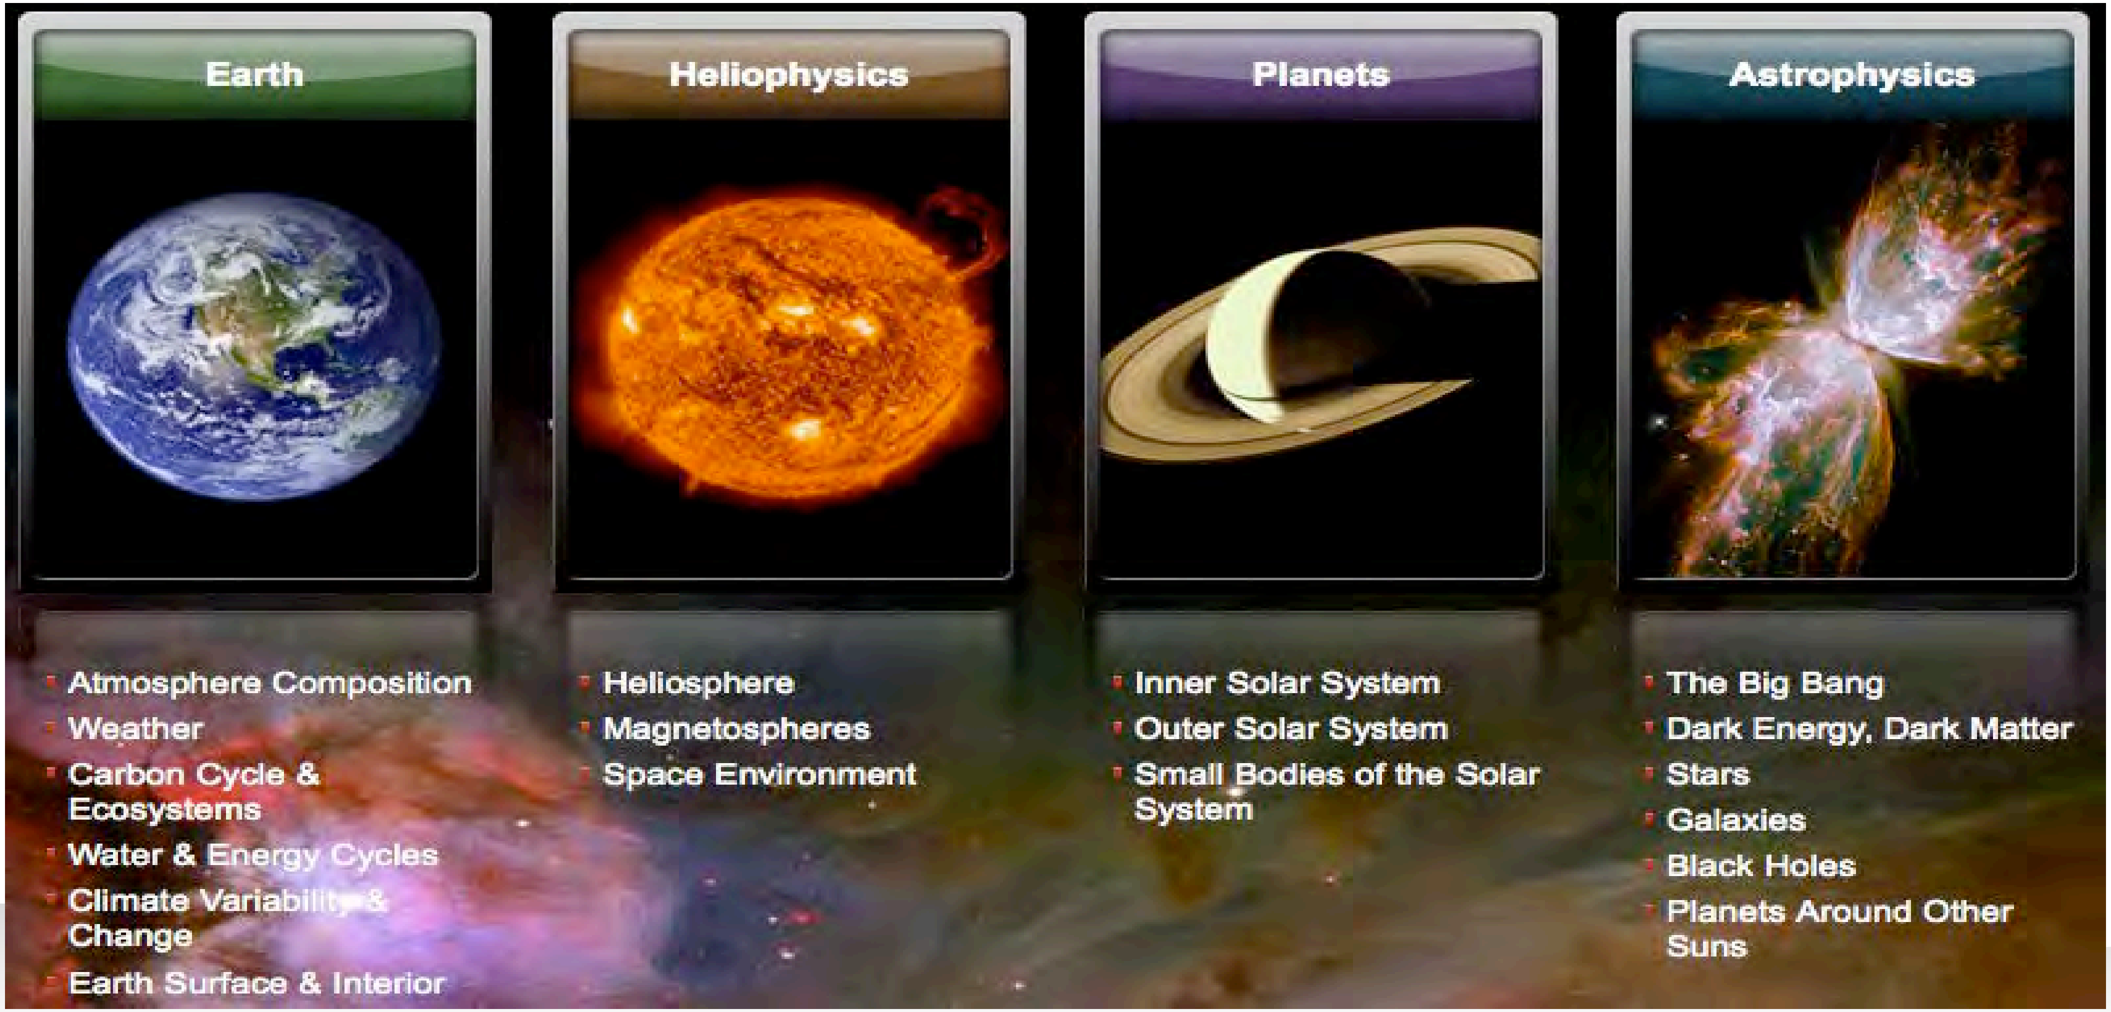

NASA Center for Climate Simulation (NCCS)

Provides an integrated high-end computing environment designed to support the specialized requirements of Climate and Weather modeling.

- High-performance computing, cloud computing, data storage, and networking technologies
- High-speed access to petabytes of Earth Science data
- Collaborative data sharing, publication, and analysis services

Primary Customers (NASA Science Mission Directorate)

- NASA-funded science projects can get access to these resources
- Global Modeling and Assimilation Office (GMAO)
- Land Information Systems (LIS)
- Goddard Institute for Space Studies (GISS)
- Variety of other Research and Development (R&D) and Engineering
  - ABoVE, HiMAT, CALET, WFIRST

Enabling High-Performance Science

- http://www.nccs.nasa.gov
- Funded by the High End Computing (HEC) program under Science Mission Directorate
  - Dr. Tsengdar Lee, Program Manager
- Code 606.2 at NASA Goddard Space Flight Center in Greenbelt, MD.

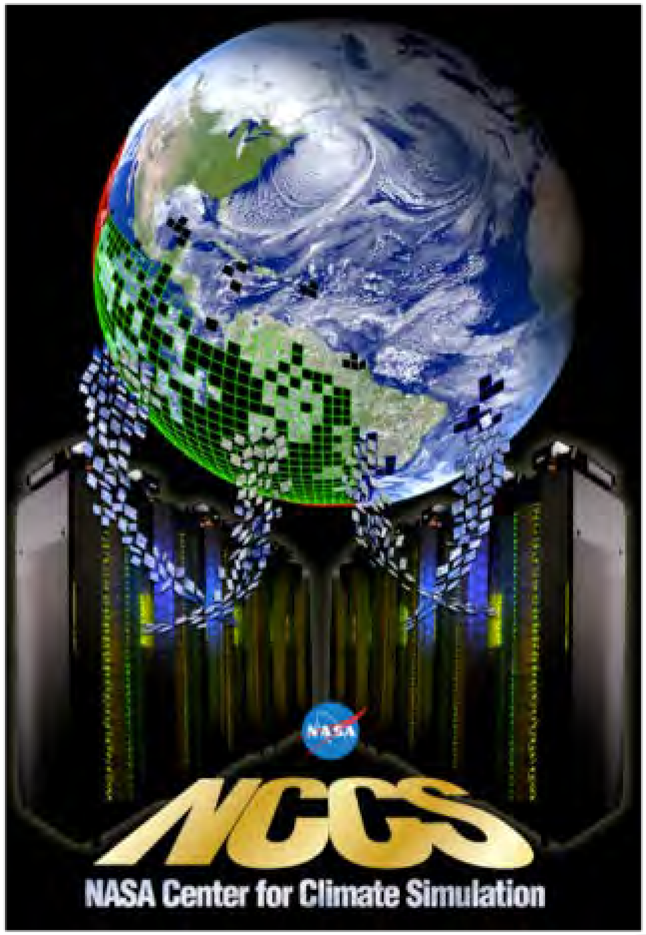

NCCS: Weather, Ocean Climate Simulation Research Supported

- Global Modeling and Assimilation Office
  - Goddard Earth Observing System (GEOS) model
- Goddard Institute for Space Studies research includes:
  - Climate forcings, climate model development, Earth observations, atmospheric radiation, atmospheric chemistry, climate impacts, planetary atmospheres and astrobiology, paleoclimate, …
- Land Information System (LIS)

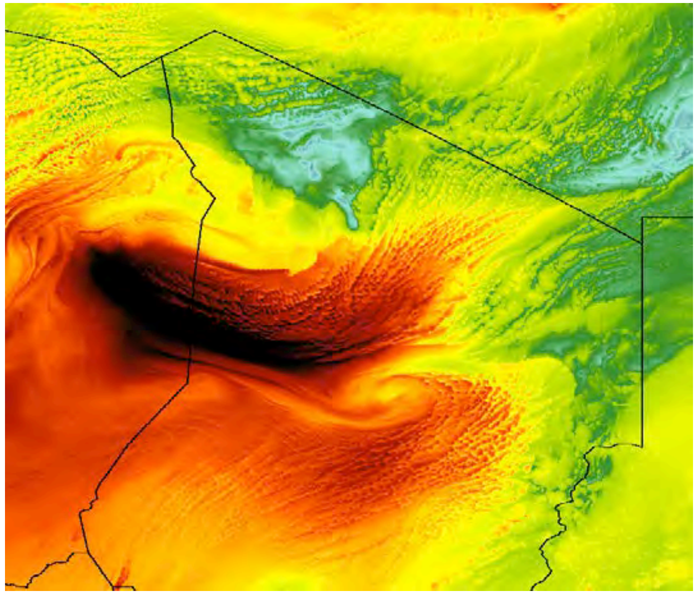

Detail of the Total Aerosol Optical Depth for a dust storm over Chad from a 1.75-km GEOS global simulation for 1200 GMT 15 June 2012 (forecast hour 15). Image source: William Putman/GMAO

NCCS: Support for NASA Carbon Research

Both on Discover (HPC), increasingly on ADAPT (virtualization environment)

- ABoVE – Arctic Boreal Vulnerability Experiment
- HiMAT – High Mountain Asia Team
- Semi-Arid Carbon Sink Study (some via Amazon Web Services & “cloud bursting”)

Many data sets & new services

- ArcGIS geospatial information system
- High-res imagery from National Geospatial-Intelligence Agency (NGA) & other sources
- Connection to MODIS data in B32 via Lightweight Virtualized File System (LVFS)

Rivers and snow in the Himalayas as seen from the International Space Station on April 8, 2015 (Image credit: NASA JSC)

NCCS: Some Additional Projects Supported on ADAPT

- Planetary Defense (near-Earth asteroids)
- CALET (CALorimetric Electron Telescope)
  - On International Space Station (ISS); searches for signatures of dark matter and makes direct measurements of the cosmic ray electron spectrum in our local region of the Galaxy.

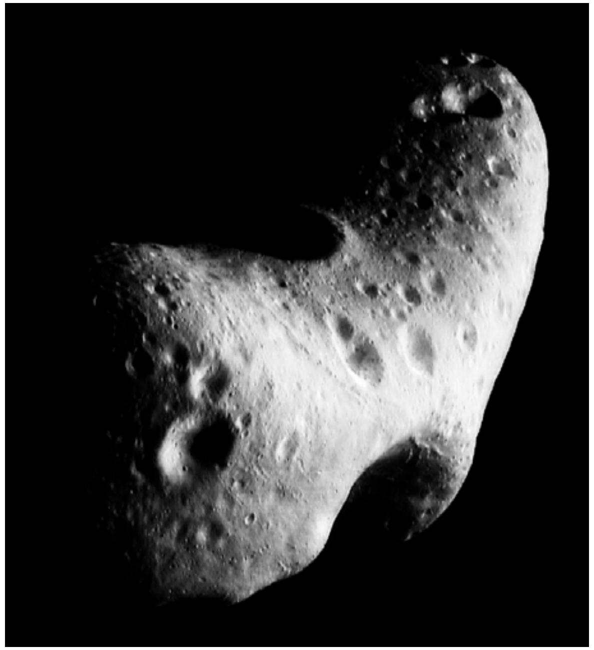

Mosaic of asteroid Eros as observed by NEAR (Near Earth Asteroid Rendezvous - Shoemaker) spacecraft on February 29, 2000 (image credit: Johns Hopkins Applied Physics Laboratory and NASA JPL).

NCCS Recent Evolution of Major Systems

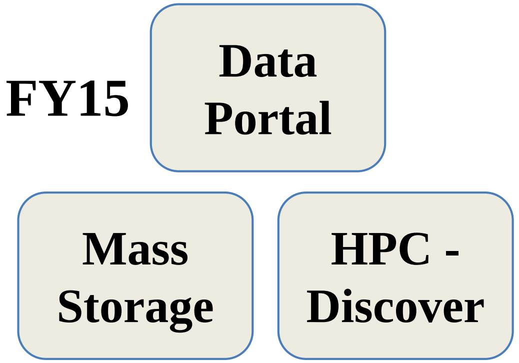

FY16: Creation of the Advanced Data Analytics Platform (**ADAPT**), a High Performance Science cloud (virtual environment) designed for traditional data services, data analytics, and web services:

move the data to the analysis.

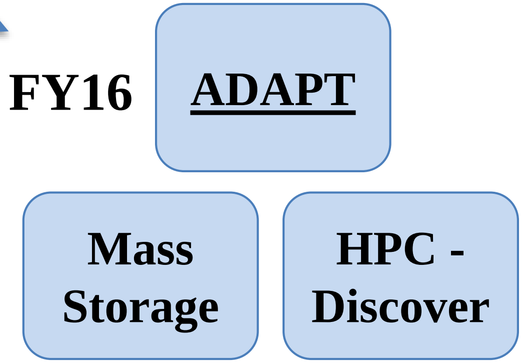

FY17: Creation of the Data Analytics Storage Service (**DASS**), a combined High Performance Computing and Data environment to enable emerging analytics:

move the analysis to the data.

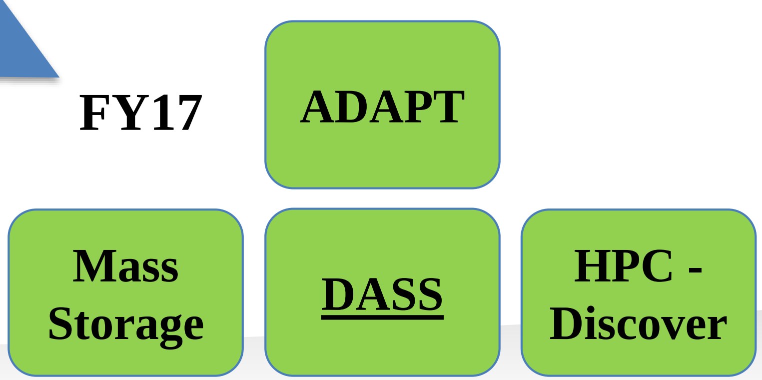

NCCS Major NCCS Systems/Services

ADAPT

High Performance Science Cloud

- Designed for large scale data analytics
- Loosely coupled applications
- Low barrier to entry for scientists
- Agile virtual environment
- Mixture of compute and storage
- Data and analytics services

Dirac

Mass Storage Environment

- ~65 PB of stored data
- ~5 PB disk front end
- Tape back end
- HPE/SGI DMF System

DASS

Data Analytics Storage System

- 1,000's of cores
- TFLOPS of compute
- PBs of storage
- High Speed Networks
- ***Operational Late Spring 2017***

Discover

High Performance Computing Cluster

- ~3.5 PFLOPS of peak computing
- Almost 90K cores; 3,400 nodes
- 42 PB (usable) shared storage (GPFS)
- High-speed networks
- Tightly coupled applications

NCCS Discover High Performance Computing and Mass Storage Growth, 2009 – Early 2017

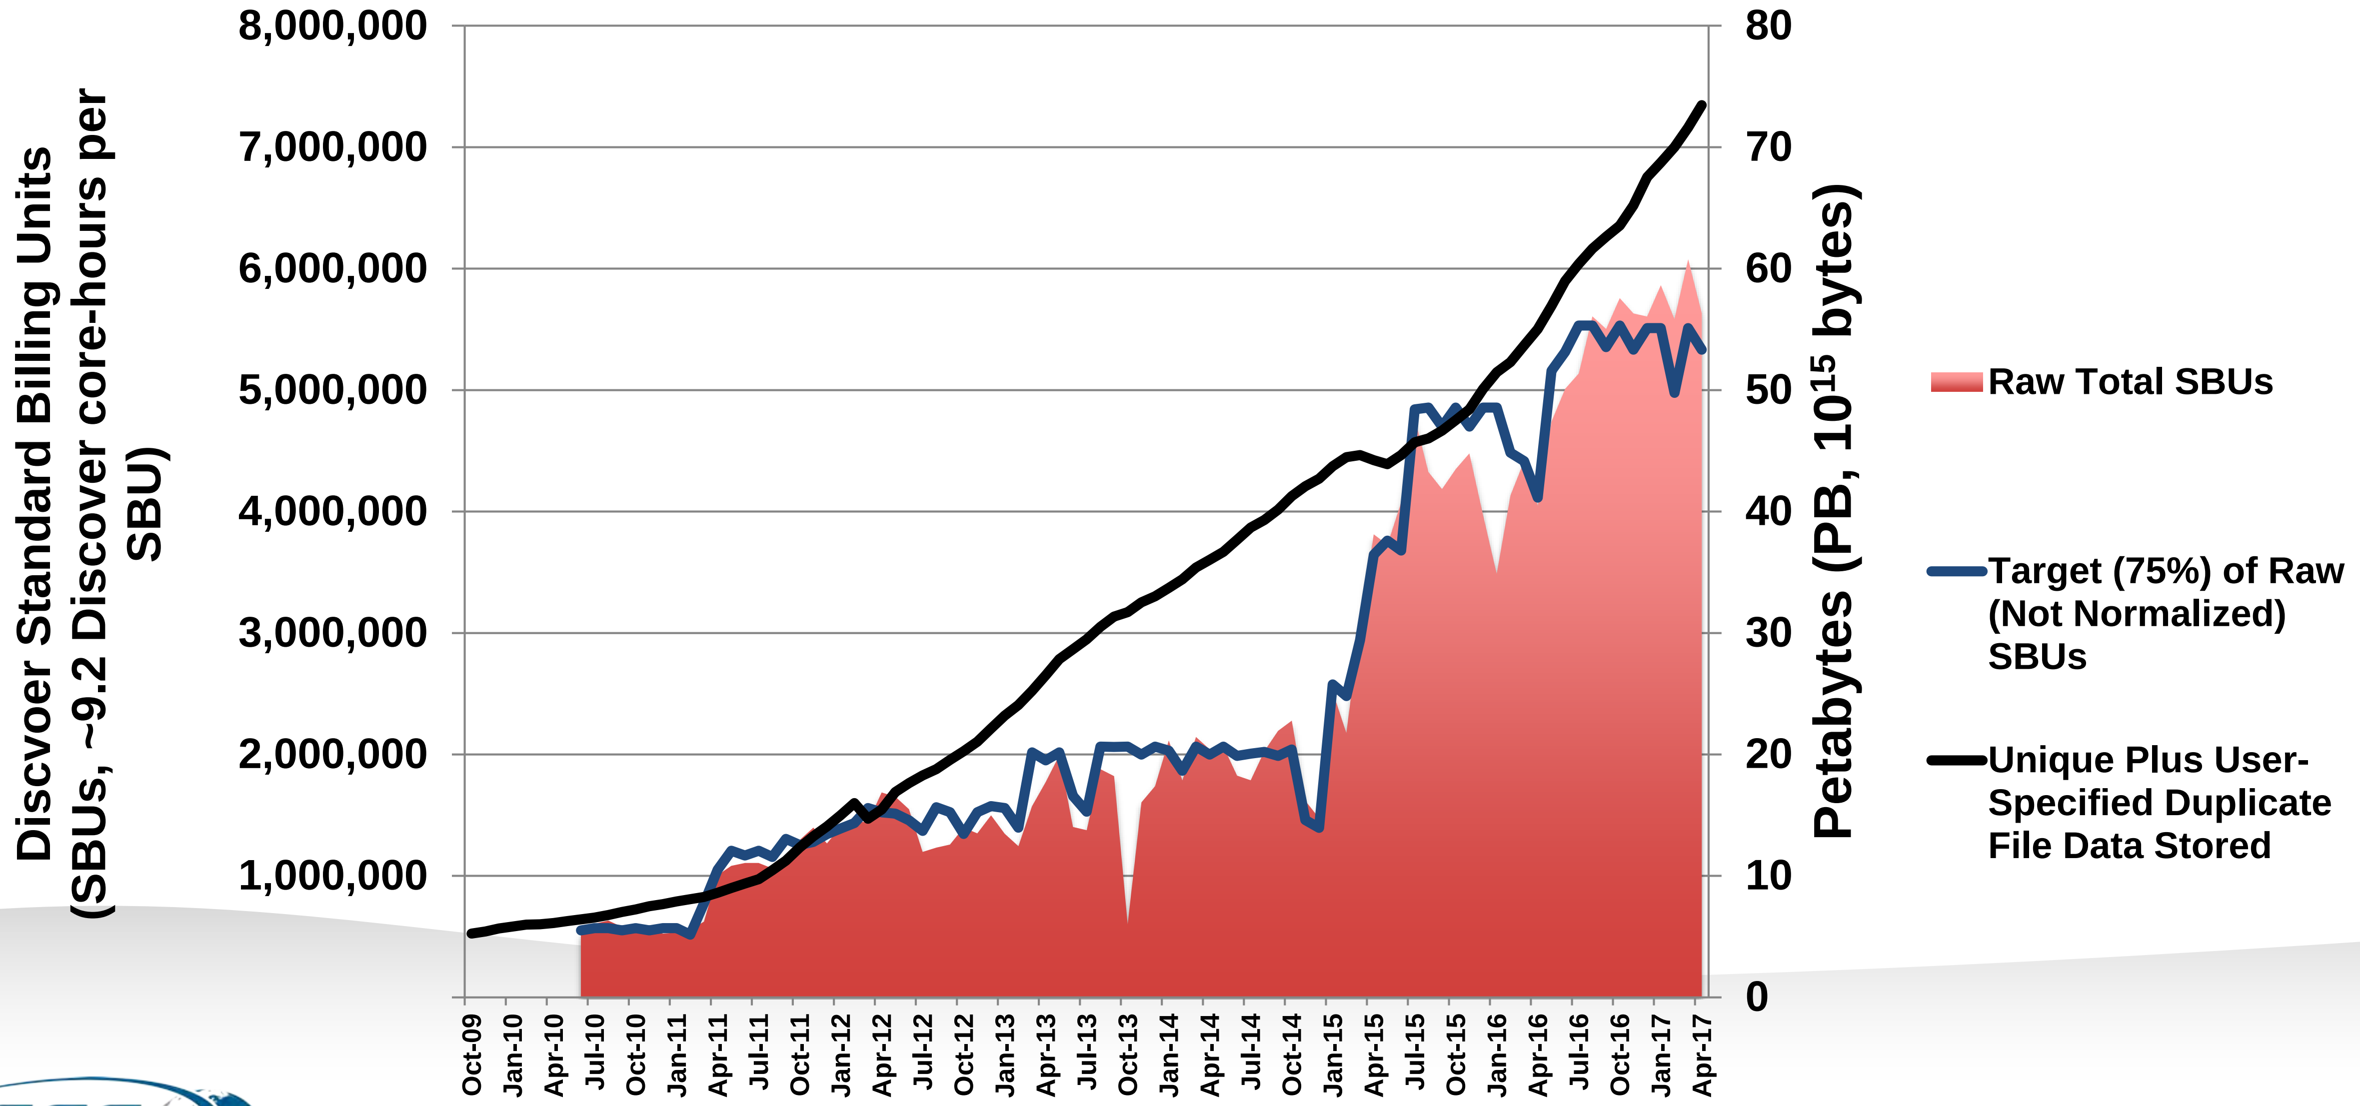

NCCS Science Drivers Example: Increasing the GEOS Model Resolution for Research

The following table shows the requirements to simulate 1-year of the Earth's atmosphere.

Year	Resolution (meters)	X and Y Grid	Total Grid Points	Cores	PB
2017	1,736	5,760	26×10^9	30,000	60
2019	868	11,520	105×10^9	240,000	480
2021	434	23,040	420×10^9	960,000	1,920
2023	217	46,080	$1{,}682 \times 10^9$	3,840,000	7,680
2025	109	92,160	$6{,}727 \times 10^9$	30,720,000	61,440

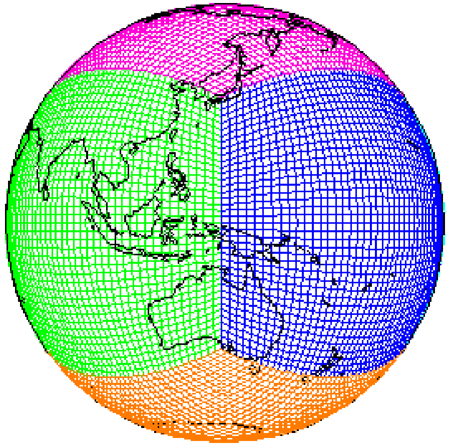

The GEOS fvCubed Sphere grid.

- Need the ability to run simulations to generate the data needed for specific science questions
- Data sets are getting too big to move or to simply serve through traditional data services.
- Cannot store all fields from each simulation!
- This is just for the atmosphere! Fully coupled models (dynamic ocean, ice, land, …) will require even more capabilities!

Advanced Data Analytics Platform (ADAPT) High Performance Science Cloud

ADAPT-wide technologies

- 100's of nodes, capable of 1,000's of virtual machines

"Data Lake" concept – data available as needed to all VMs

- ~15 PB lower cost, commodity storage, easily expandable
- High performance file systems using IBM Spectrum Scale/GPFS

High speed external networks

- 10 and 40 GbE
- remote mounts to external data sources (e.g., NCCS HPC, GES DISC, MODIS)

High speed internal networks

- repurposed HPC InfiniBand switches
- Ethernet switches

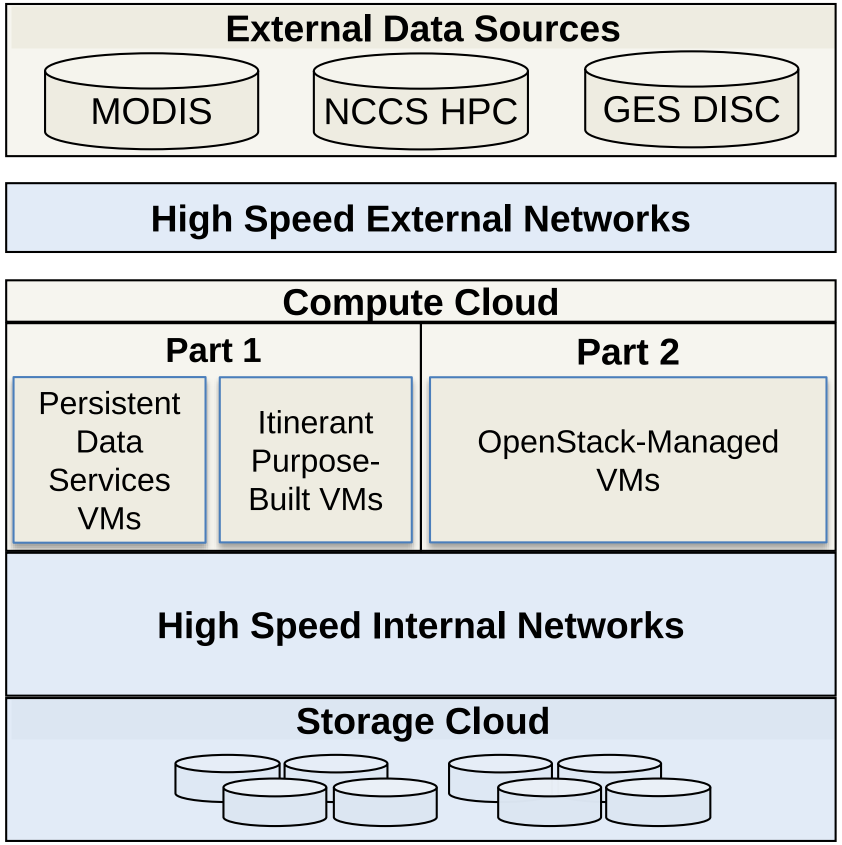

Part 1: compute resources: older re-purposed HPC nodes, providing

- "Data Portal" Persistent Data Services:
  - Long-lived virtual machines (VMs) specifically designed for data or web services (including ESRI ArcGIS)
- Itinerant purpose-built virtual machines
  - Customized for each user/project, spun up and down as needed.

Part 2: modular container compute resources:

- OpenStack-managed virtual machines (in development)

NCCS ADAPT Part 2: Modular Container

Data Analytics Storage System (DASS) - Motivation

- Data movement and sharing of data across services within the NCCS is a continuing challenge
- Large data sets created on Discover (HPC)
  - On which users perform many analyses
  - And much of this data is not stored in a NASA Distributed Active Archive Center (DAAC)
- Approach: create a true centralized combination of storage & compute capability
  - Capacity to store many PBs of data for long periods of time
  - Architected to be able to scale both horizontally (compute and bandwidth) and vertically (storage capacity)
  - Can easily share data to different services within the NCCS
  - Free up high speed disk capacity within Discover
  - Enable both traditional and emerging analytics
  - No need to modify data; use native scientific formats

Data Analytics Storage System (DASS) Concept

Read access for Advanced Data Analytics Platform (**ADAPT**) nodes and Virtual Machines (VMs)

- Serve to **Dataportal** (persistent data services VMs: ftp, HTTP, OPeNDAP, ESGF, etc.)
- Serve to **Itinerant purpose-built VMs**, and **OpenStack-managed VMs** for additional processing
- Allows integration/comparison of simulation results and observations

Analysis request is sent to a service.

ADAPT

Climate Analytics as a Service

Answer is returned.

HPC – Discover Cluster

Mass Storage

DASS (~20 PB)

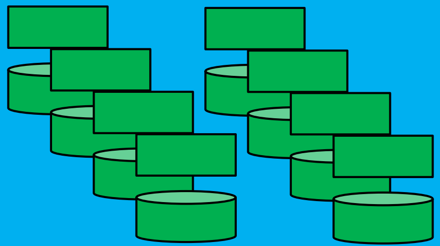

Read/write access from the **Mass Storage**

- Stage data into and out of the centralized storage environment as needed

Discover is the source of much of the data:

- Write and Read access to DASS from all network-accessible nodes in Discover – models write data into Discover GPFS which can then be staged into the DASS centralized storage (burst buffer like).

Note that all systems will still have local filesystems to enable local optimized writes and reads as needed within their respective security domains.

DASS: Analyze Weather, Ocean, Climate Datasets

MPI, Open, Read, Write, etc.	***Traditional HPC***	***Big Data Analytics***	MapReduce, Spark, ML
Network, IB, RDMA	**Classical Usage Patterns:** **Data is moved to the process**	**Hadoop-Like Usage:** **Analytics moved to the data**	Cloudera, Horton, BDAS
GPFS	**POSIX Interface**	**RESTful Interface**	**SIA** / Hadoop Connector
IBM Spectrum Scale (GPFS)	**Object Store/Posix Parallel File System** Very large, scaling both horizontally (throughput) and vertically (capacity); permeated with compute capability at all levels		IBM Spectrum Scale (GPFS)
	Traditional HPC Storage	**Server & JBOD Commodity-Based Hardware**	

Spatiotemporal Indexing Approach (SIA)

Collaboration with GMU

Spatiotemporal Index Approach (SIA) and Hadoop

- Use what we know about the structured scientific data
- **Create a spatiotemporal query model to connect the array-based data model with the key-value based MapReduce programming model using grid concept**
- Built a spatiotemporal index to
  - Link the logical to physical location of the data
  - Make use of an array-based data model within HDFS
  - Developed a grid partition strategy to
  - Keep high data locality for each map task
  - Balance the workload across cluster nodes

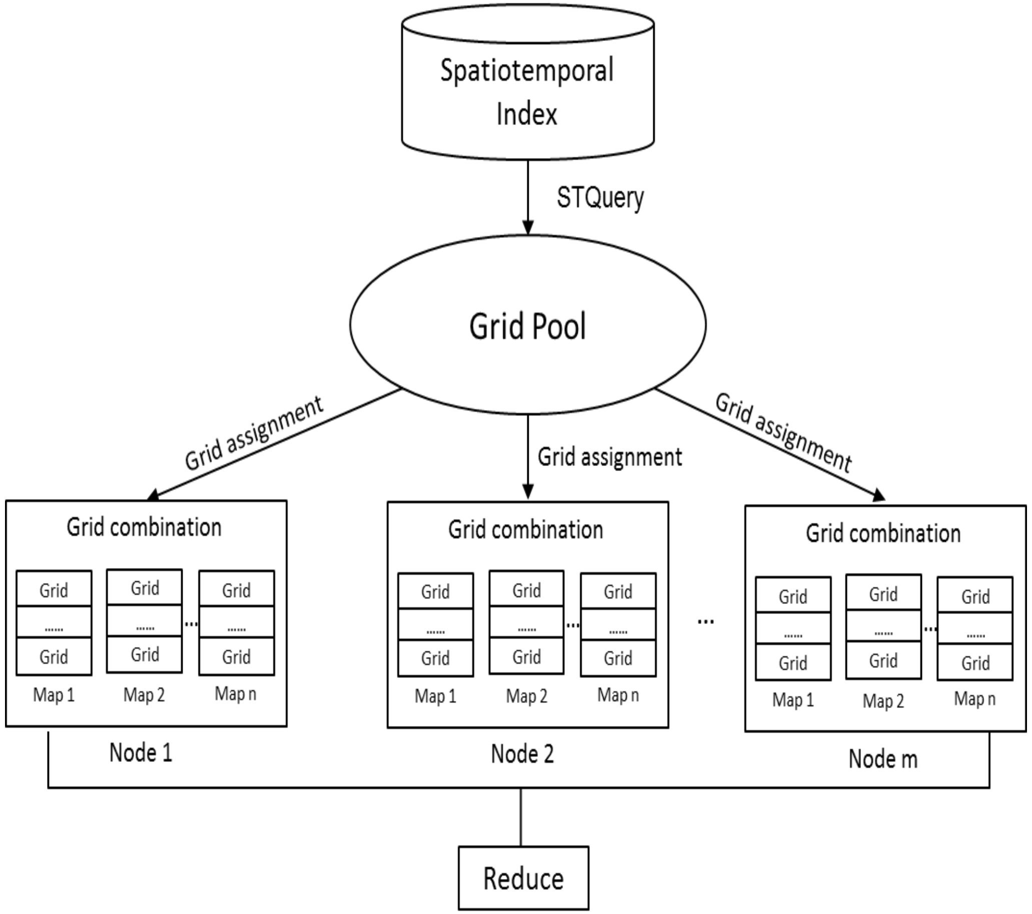

A spatiotemporal indexing approach for efficient processing of big array-based climate data with MapReduce
Zhenlong Lia, Fei Hua, John L. Schnase, Daniel Q. Duffy, Tsengdar Lee, Michael K. Bowen and Chaowei Yang
International Journal of Geographical Information Science, 2016
http://dx.doi.org/10.1080/13658816.2015.1131830

Estimating Woody Biomass on South Side of the Sahara at the 40-50 cm Scale Using Cloud Bursting to Amazon Web Services

Project Science Goal – Using high-resolution satellite imagery data, estimate tree and bush biomass over the entire arid and semi-arid zone on the south side of the Sahara to:

- Estimate carbon stored in trees and bushes.
- Establish carbon baseline for later research on expected CO_2 uptake.

Cloud Bursting into Amazon Web Services (AWS) Cloud

- Satellite imagery preprocessed to remove distortions on NCCS ADAPT.
- Batch system (with Cycle Computing) in ADAPT communicates with elastic twin in AWS to move data and initiate processing.
- Biomass calculations run on AWS virtual machines (VMs) for 6-7 hours each, via low-cost AWS spot instances; work scales linearly.
- AWS returns results (~25% size of input) to an ADAPT file system.
- Easily scale up the number of virtual machines using the Cycle Computing software and the AWS resources.

Principal Investigators

- Dr. Compton J. Tucker, Goddard Space Flight Center (Code 610)
- Dr. Paul Morin, University of Minnesota

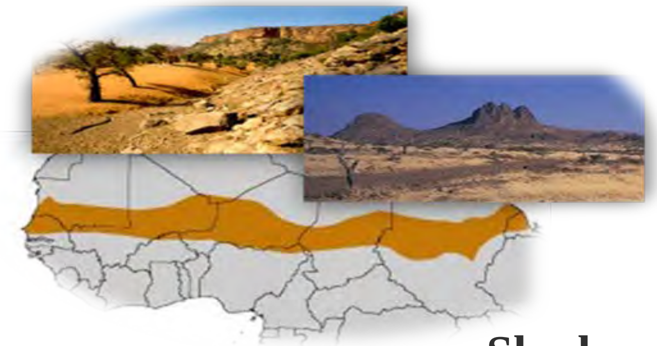

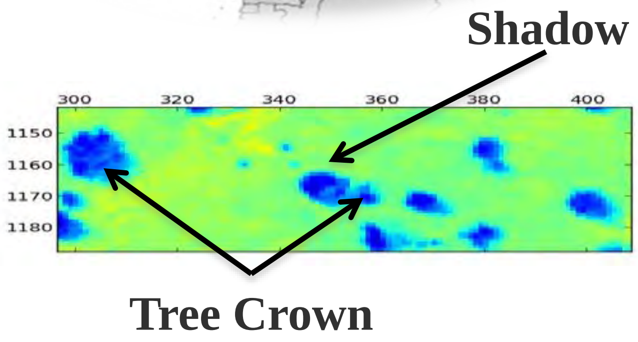

40-cm imagery representing tree & shrub automated recognition. © DigitalGlobe, Inc., licensed under NextView.

Looking Toward Exascale

ADAPT
Virtualization Environment
HPC and Cloud
Existing Size:
~1,000 cores
~10 PB storage

Designed for Big Data Analytics

Mass Storage
Tiered Storage
Existing Size:
~75 - 100 PB storage

Designed for longer term storage and retrieval, not compute

DASS
Tiered Storage
Memory, SSD, Disk
Existing Size:
~10 PB storage

Designed for compute, analytics, and longer term storage

Future Exascale Environment
Merging of HPC and Big Data Analytics Capabilities

Ability for in-situ analytics throughout the environment … both known analytics and machine learning

HPC/Discover
HPC Cluster
Existing Size:
~100,000 cores
~50 PB storage

Designed for Large-Scale Weather, Ocean Climate Simulations

Analytics Intensity

Computational Intensity

Thank You


Dave Van
Lee
Glenn
Reitz Palm
Hoot Settel Laura
Sinno Pfaff
Zed
Ben Lyn
Clune Probe
Luczak
Mason
Al Tamkin Duffy
Dewan
Gerner
Gooden
Craig Rumney Thompson
Bruce
Chang Taylor Schardt Fitzgerald Gene
Schnase
George Maharaj Nancy
Bhat Tsengdar
Webster
Savannah
Adina Tom
Tarshish Phil Woodford John Ripley
Mike
Chyatte Greg Sorathia Laubenthal Mcinerny
Pelissier
Donovan
Kameza Williams Scott
Bill
Gyorgy Strong McHale
Salmon Carrier Bledsoe
Aartsen Nadeau Parnell
Debbie
Edward
Lynn
Kareem Daniel Jasen
Jasaun
Acks
Denis Guillen Ellen Mark Nicko Burch
Fekete
Max Tim Fred
Jim Neff Heidi
David