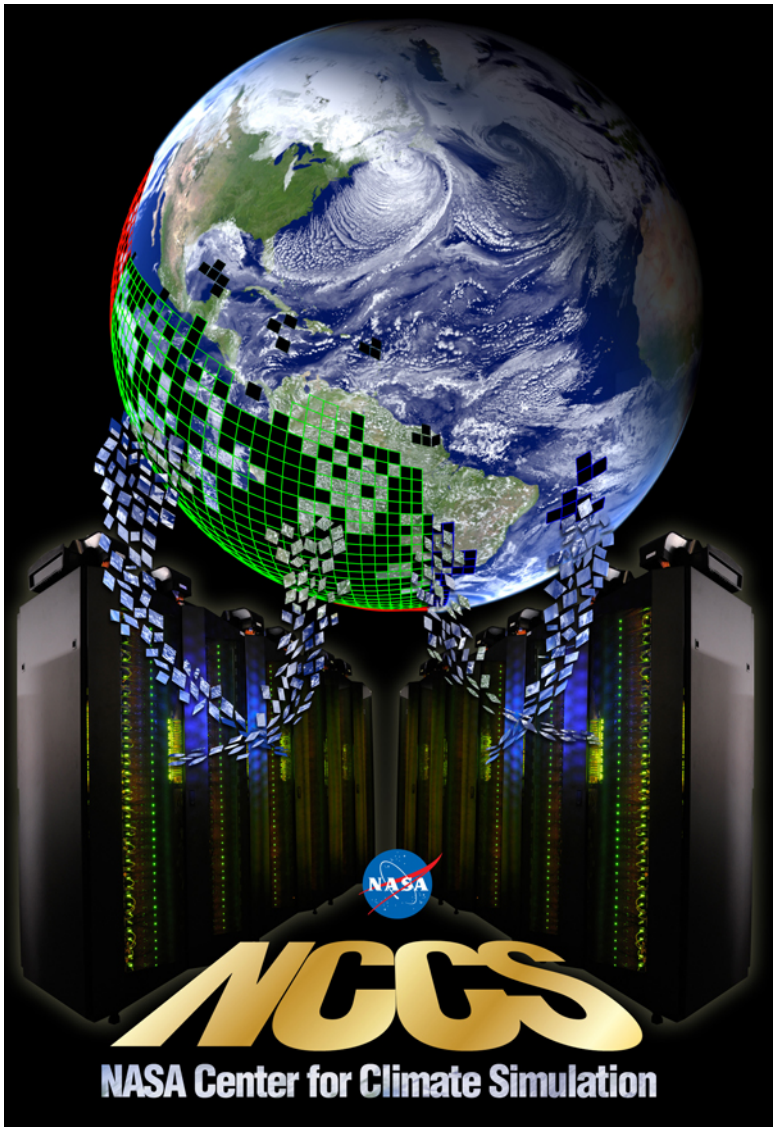

Climate SuperComputing at NASA Goddard Space Flight Center

Fifteenth ECMWF Workshop on the Use of High Performance Computing in Meteorology

Dr. W. Phillip Webster

3 October 2012

ECMWF

Agenda

- **NASA Context**
- **Computing Environment**
- **Data Services Model**

NASA Mission Structure

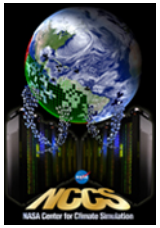

To implement NASA's Mission, NASA Headquarters is organized into Three *Mission Directorates*.

- **Aeronautics**: Pioneers and proves new flight technologies that improve our ability to explore and which have practical applications on Earth.
- **Human Exploration and Space Operations**: Creates new capabilities and spacecraft for affordable, sustainable human and robotic exploration. Provides critical enabling technologies for much of the rest of NASA through the space shuttle, the International Space Station, and flight support.
- **Science**: Explores the Earth, moon, Mars, and beyond; charts the best route of discovery; and reaps the benefits of Earth and space exploration for society.

Science Mission Directorate

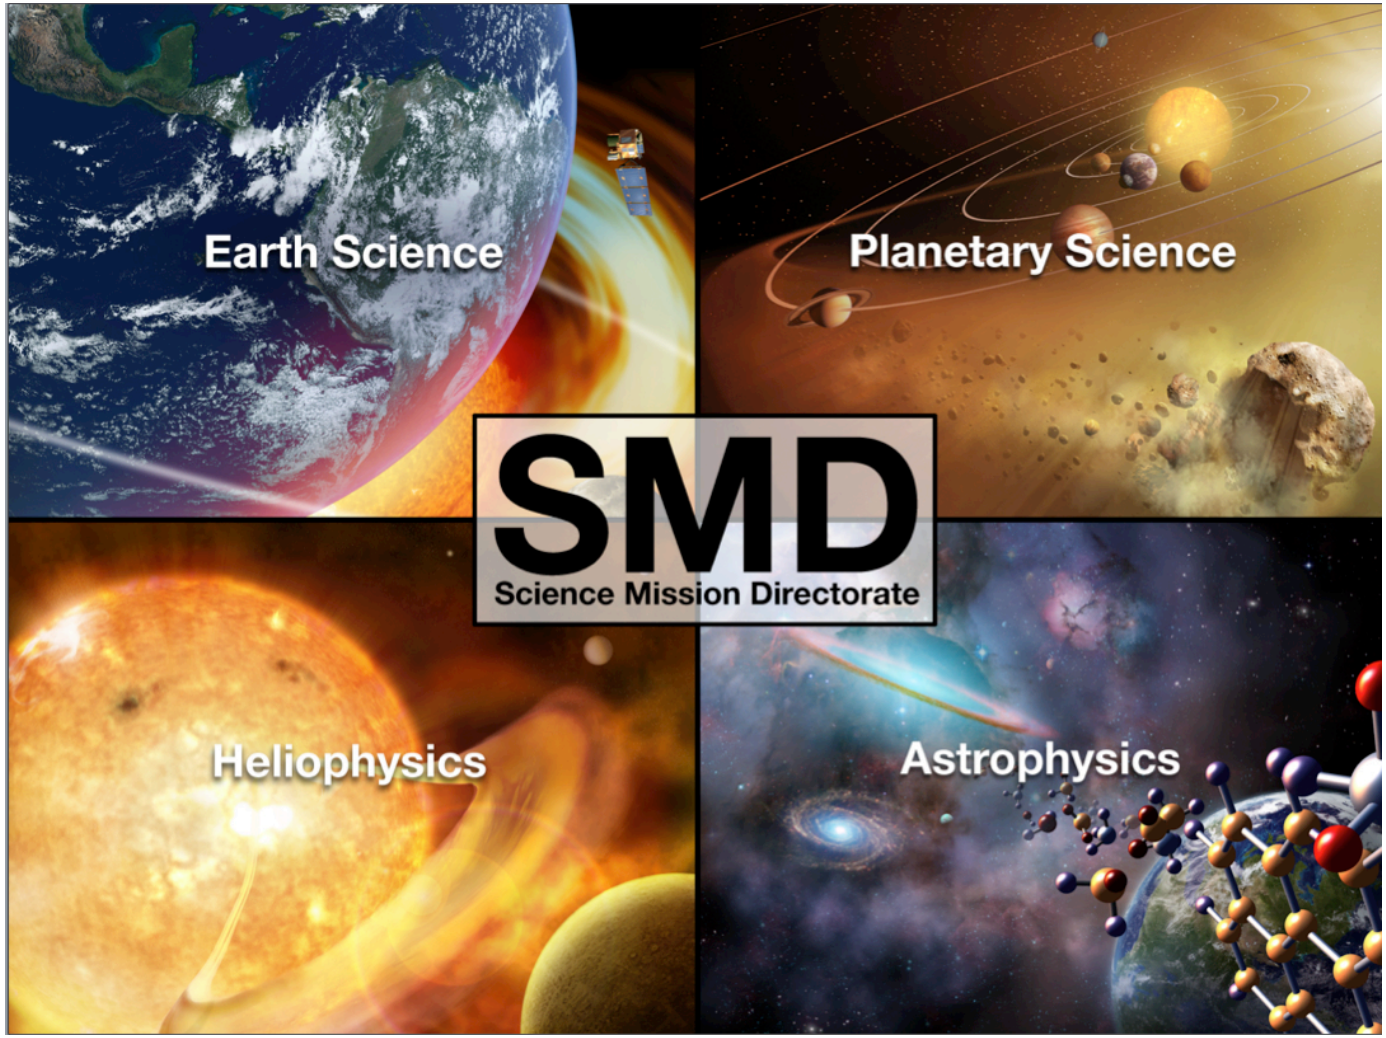

Earth Science Division Overview

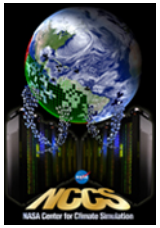

Overarching Goal: to advance Earth System science, including climate studies, through spaceborne data acquisition, **research and analysis, and predictive modeling**

Six major activities:

- Building and operating Earth observing satellite missions, many with international and interagency partners
- Making high-quality data products available to the broad science community
- **Conducting and sponsoring cutting-edge research in 6 thematic focus areas**
  - Field campaigns to complement satellite measurements
  - **Modeling**
  - Analyses of non-NASA mission data
- **Applied Science**
- **Developing technologies to improve Earth observation capabilities**
- Education and Public Outreach

ECMWF

HEC in 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), and Science (SMD).

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.

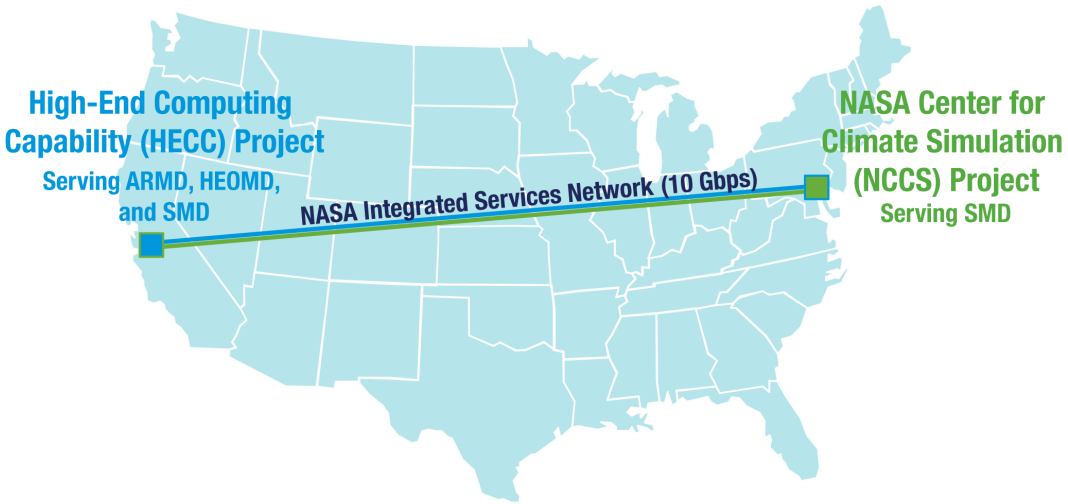

NCCS offers an integrated set of supercomputing, visualization, and data interaction technologies to enhance NASA capabilities in weather and climate prediction and enable future scientific discoveries that will benefit humankind. It serves hundreds of users at NASA centers, laboratories, and universities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

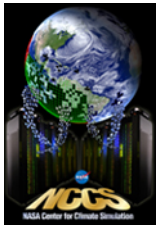

Agenda

- NASA Context
- **Computing Environment**
- Data Services Model

Climate Science Computing

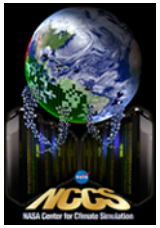

Climate Science is "Data Intensive". . . . . .

- Our understanding of Earth processes is based in the observational data record and is expressed as mathematical models.
- Data assimilation combines observational data with model prediction.
- Climate models produce large data sets (100s of terabytes) for the scientific community as well as decision makers.
- Reanalysis with improved models results in vast data sets (100s of terabytes) for the scientific community.

. . . . . And Requires "Data Centric" Computing

- Designed for effective manipulation of large data sets.
- Global file system makes data available to all services. Effective data management tools are required.
- Efficient data analysis needs to have "supercomputing" capability with data sets online.
- Data sets must be made easily accessible to "external users" with analysis and visualization capability.

NCCS Data Centric Climate Simulation Environment

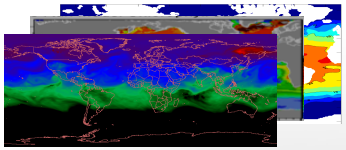

Data Sharing and Publication

- 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 NCCS for GSFC users
- Multi-gigabit network supports on-demand data transfers

HPC Computing

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

Data Archival and Stewardship

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

*Joint effort with SSSO

*Joint effort with SEN

NCCS Data Centric Climate Simulation Environment

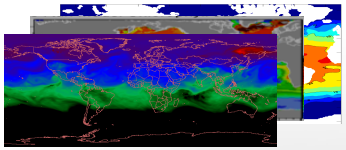

Data Sharing and Publication

- 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 NCCS for GSFC users
- Multi-gigabit network supports on-demand data transfers

HPC Computing

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

Data Archival and Stewardship

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

*Joint effort with SSSO

*Joint effort with SEN

NCCS Architecture

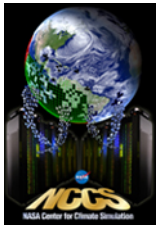

Existing | Planned for FY13

NCCS LAN (1 GbE and 10 GbE)

Data Portal

Login Nodes

Data Gateways

Data Management

Viz Wall

Discover-Fabric 1
139TF, 13K Cores

Dali Analysis Nodes

GPFS I/O Nodes

Discover-Fabric 2
650TF, 31K Cores
39K GPU Cores
14K MIC Cores

Dali-GPU Analysis Nodes

GPFS I/O Nodes

FY13 Upgrade

EST. ~160TF

GPFS I/O Nodes

ARCHIVE

Disk
~970 TB

Tape
~30 PB

GPFS Disk Subsystems
~ 8.0 PB

Management Servers | License Servers | GPFS Management | PBS Servers | Other Services

Internal Services

ECMWF

NCCS Architecture Standards

- Intel Xeon processors
  - 8265 Nahalem cores
  - 26,784 Woodcrest cores
    - 64 NVIDIA GPUs
  - 7680 Stony Bridge
    - 240 Intel MIC this fall
- SUSI Linux
  - V11 SP1
- IBM GPFS
- Altair PBS Pro

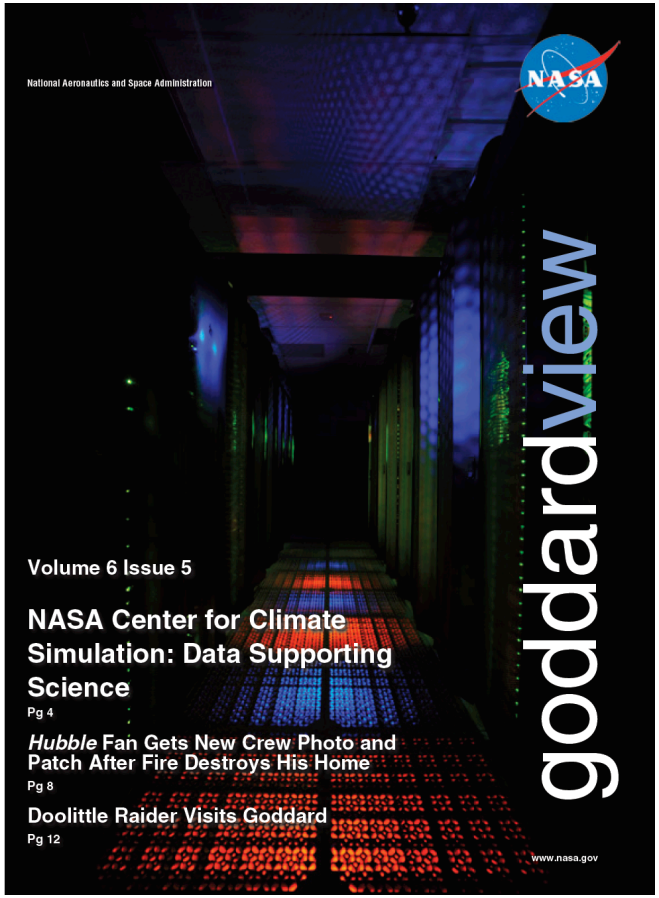

NCCS Compute Capacity Evolution 2007-2013

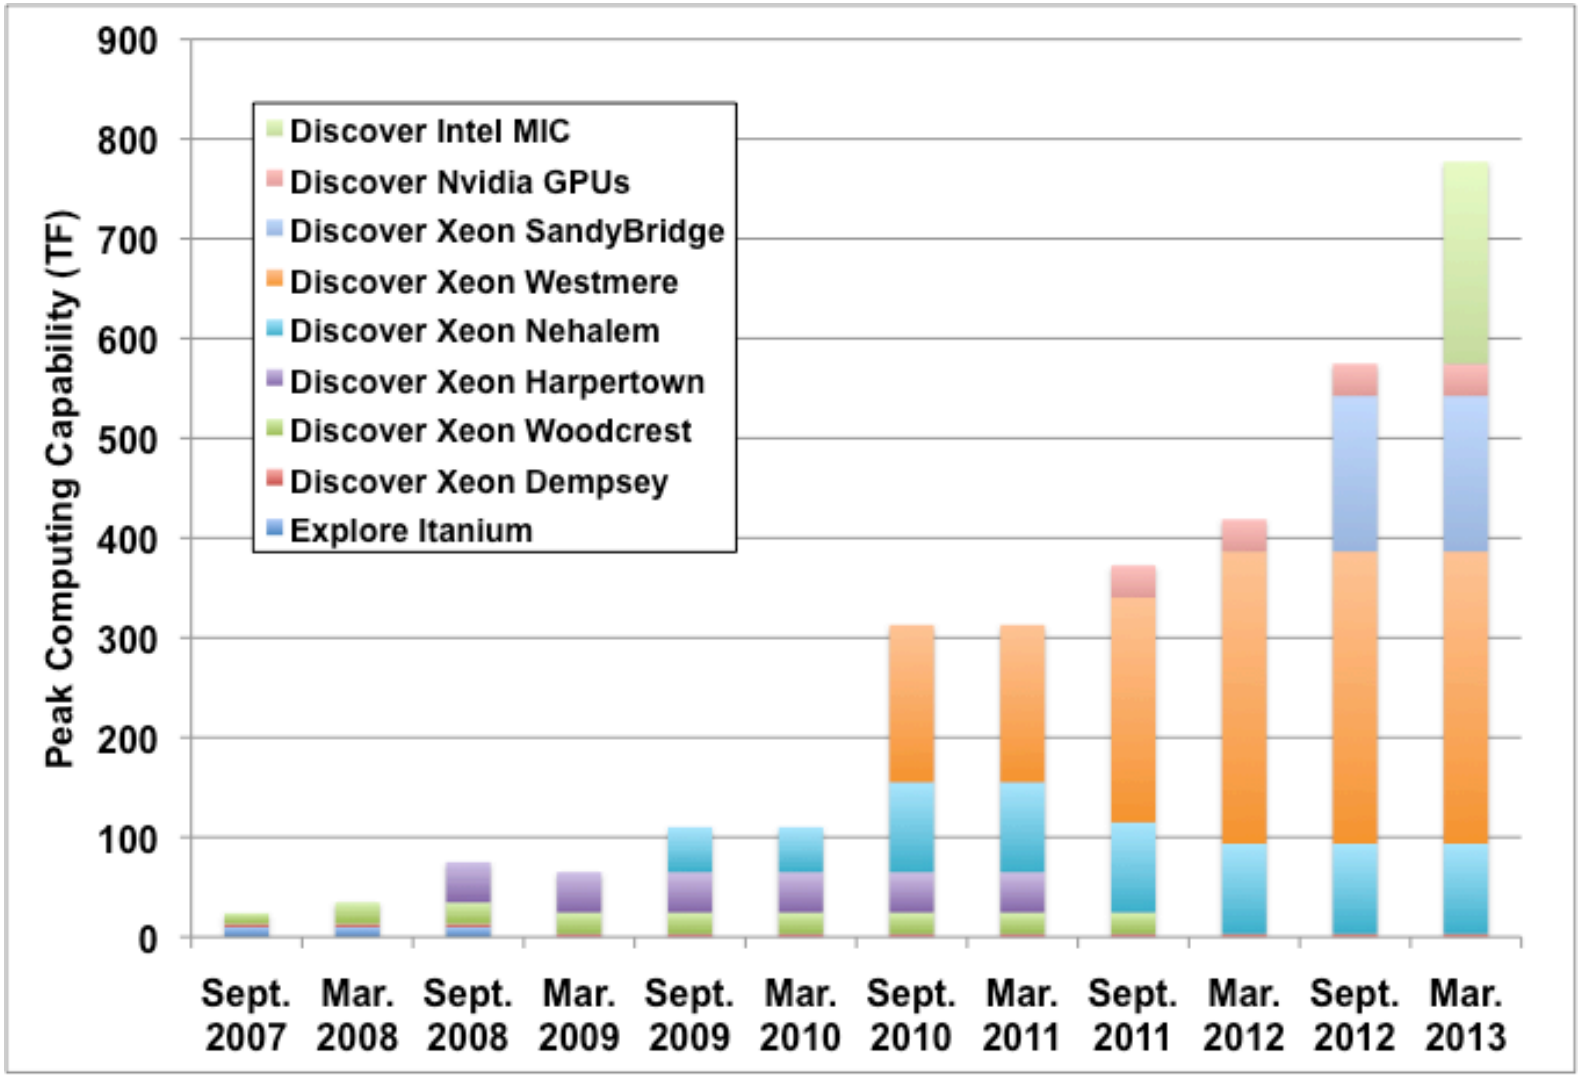

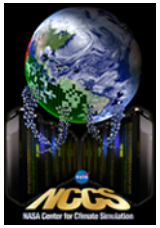

Agenda

- NASA Context
- Computing Environment
- **Data Services Model**

Climate Simulation Data Communities

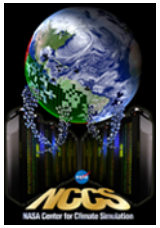

NASA Scientific Community
- Simulation data consumers
- Advance scientific knowledge
- Direct access to systems
- *Supercomputer* capability required for effective analysis

NASA Modeling Community
- Model development, testing, validation, and execution
- Data creation
- Largest HPC usage
- Requires observational data as input

Simulation Data

External Applications Community
- Huge opportunity for impact from climate change data
- Simulation data consumers
- Limited Earth Science data expertise
- Web-based access to systems

External Scientific Community
- Simulation data consumers
- Advance scientific knowledge
- Web-based access to data

Each community has different capabilities and data usage requirements.

Evolution of Data Services Centric Environment

Data

HPC Models
- GEOS 5
- ModelE
- WRF

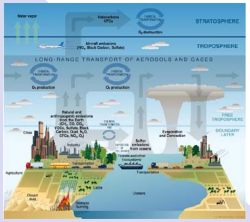

Observations
- Ground Based
- Satellite
- In Situ

Reanalysis
- MERRA
- NOAA
- Others

HPC Computing and Storage
- NASA NCCS
- NOAA
- Others

Analytics

Data Services

Moving beyond just a file system and a storage repository.

Technologies
- Dali Analysis Nodes
- vCDS
- Hadoop (HDFS)
- Merra Analytic Service
- Earth System Grid
- Web Portals

Discovery

Modelers/ Scientists

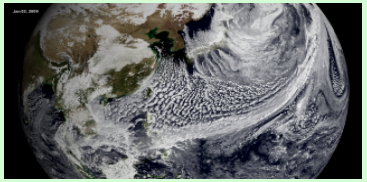

Downstream Users
- Agriculture
- Water Management
- Public Health
- Famine Prediction

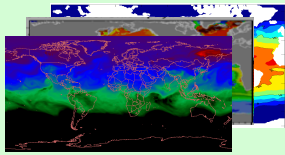

Commercial
- Insurance/Reinsurance
- Commodity Trading

Public/Citizen Scientists

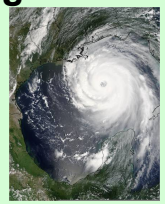

Data Management System
iRODS based management of federated data sets

iRODS Architecture and Interfaces

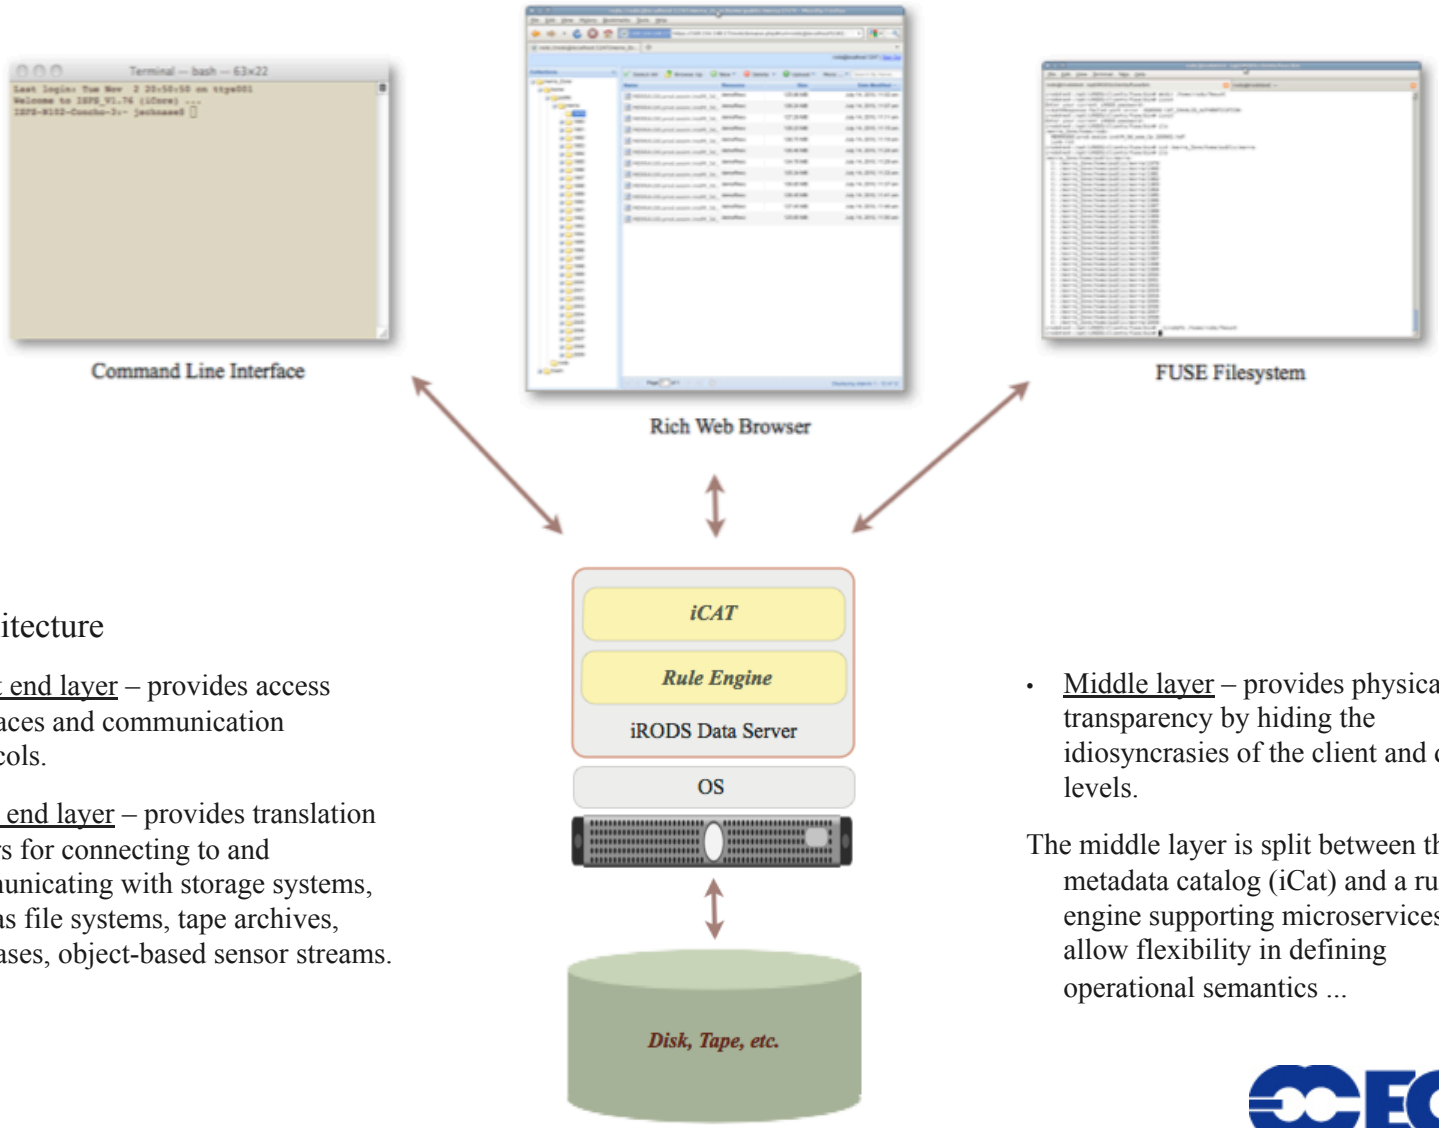

Architecture

- Front end layer – provides access interfaces and communication protocols.
- Back end layer – provides translation drivers for connecting to and communicating with storage systems, such as file systems, tape archives, databases, object-based sensor streams.
- Middle layer – provides physical transparency by hiding the idiosyncrasies of the client and diver levels.

The middle layer is split between the metadata catalog (iCat) and a rule engine supporting microservices that allow flexibility in defining operational semantics ...

iRODS Architecture and Interfaces

Command Line Interface

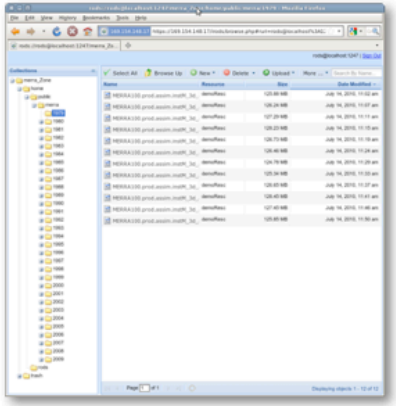

Rich Web Browser

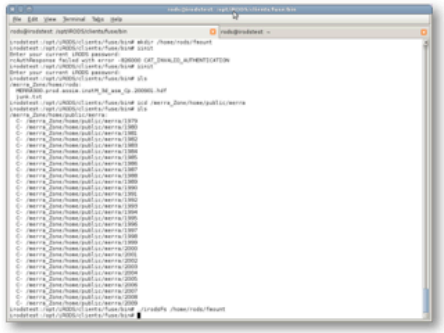

FUSE Filesystem

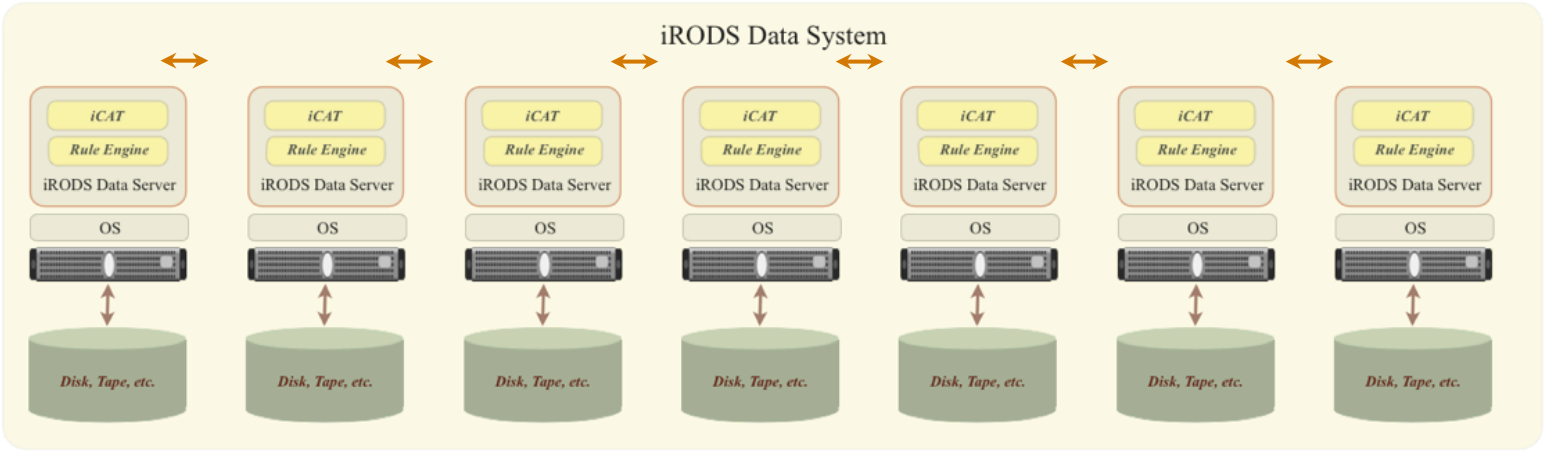

vCDS (Virtual Climate Data Services) Anatomy

Our approach has been to build a core suite of general purpose scientific kits – such as NetCDF, HDF, and GeoTIF – that sit in the vertical stack above iRODS and below application-specific climate kits such as IPCC, MERRA, and SMOS ...

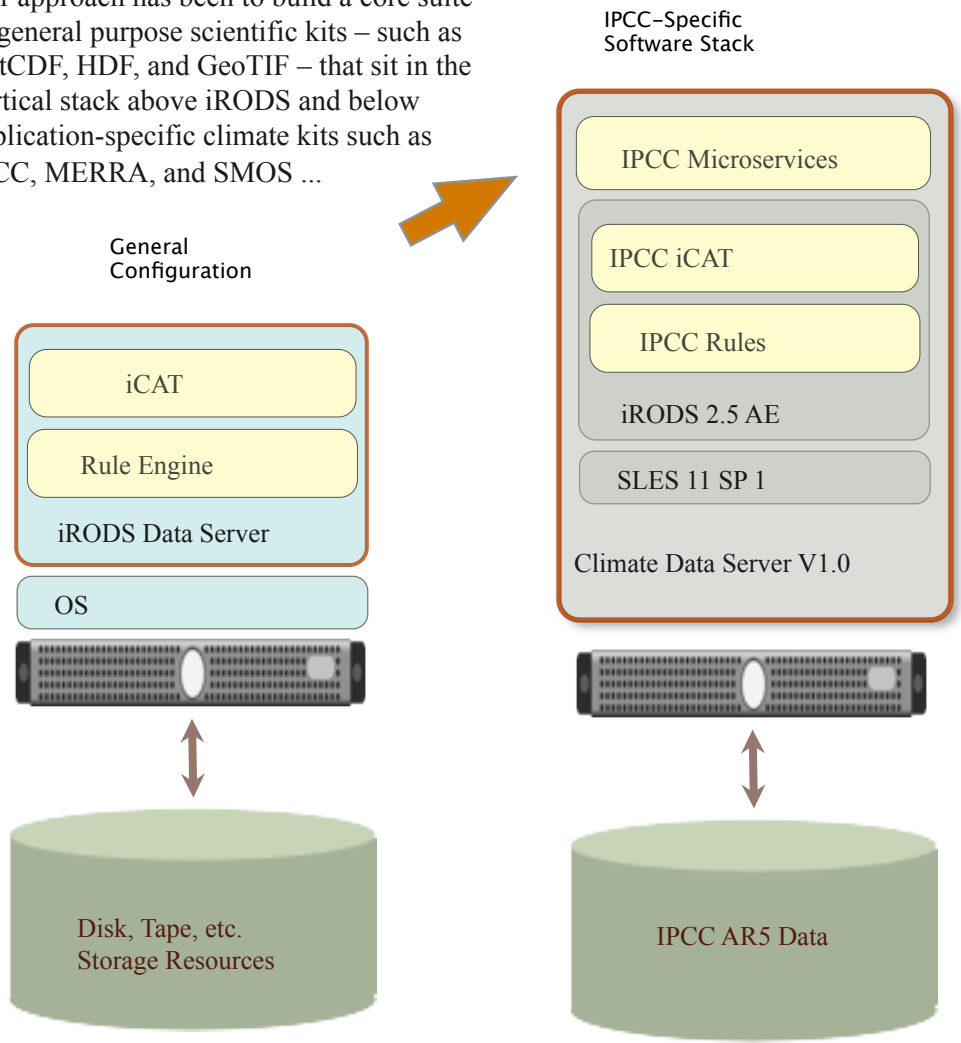

IPCC Kit + NetCDF Kit

CDS 1.0 Products

1.IPCC / NetCDF Module

iRODS microservices, rules, configuration settings, and software utilities required to implement canonical CRUD operations for IPCC/NetCDF system kernel ...

•Administrative Extensions

iRODS Postgres extensions and utilities to log system-level object provenance and provide QA for OAIS metadata compliance (plus associated Rich Web Browser GUI extensions) ...

•Repetitive Provisioning

RPM scripts to build software stacks for the SLES 11 SP1 (IaaS), iRODS AE (PaaS), and CDS/IPCC (SaaS) virtual images ...

•Deployment and Distribution

Product library, documentation, and SLA infrastructure for distribution, deployment, and help desk support ...

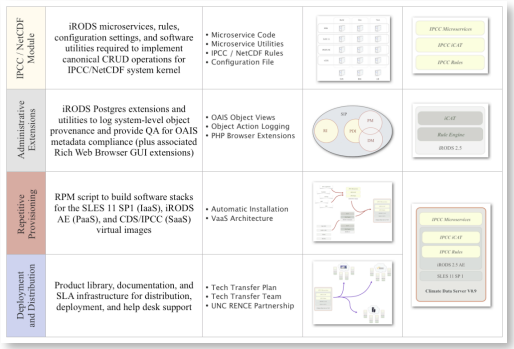

MERRA/AS

MERRA Analytic Services

•Modern Era-Retrospective Analysis for Research and Applications

- Integrates observing systems with numerical models to produce a temporally and spatially consistent synthesis of difficult-to-observe climate variables

- Covers time period of 1979 through the present

- Special focus of the atmospheric assimilation is the hydrological cycle.

•MERRA/AS = MapReduce Service

- 1PB HDFS integrated into an iRODS-, cloud-, and vCDS-based MERRA Analytic Server

- Intermediates/supports data-proximal hosting of exogenous MR analytic codes

- Implemented as suite of NetCDF, MapReduce, and MERRA Kits ...

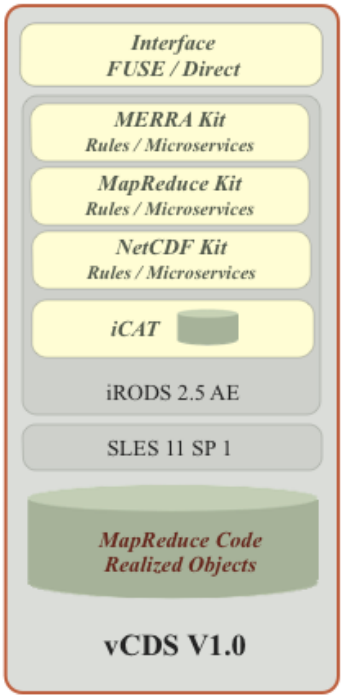

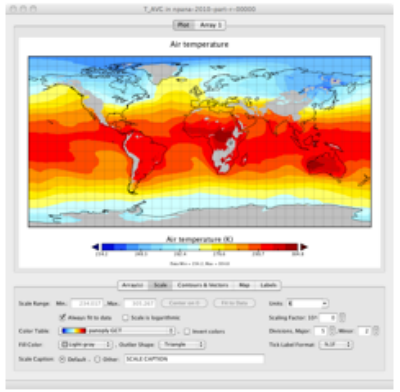

ECMWF

CISTO Data Vision

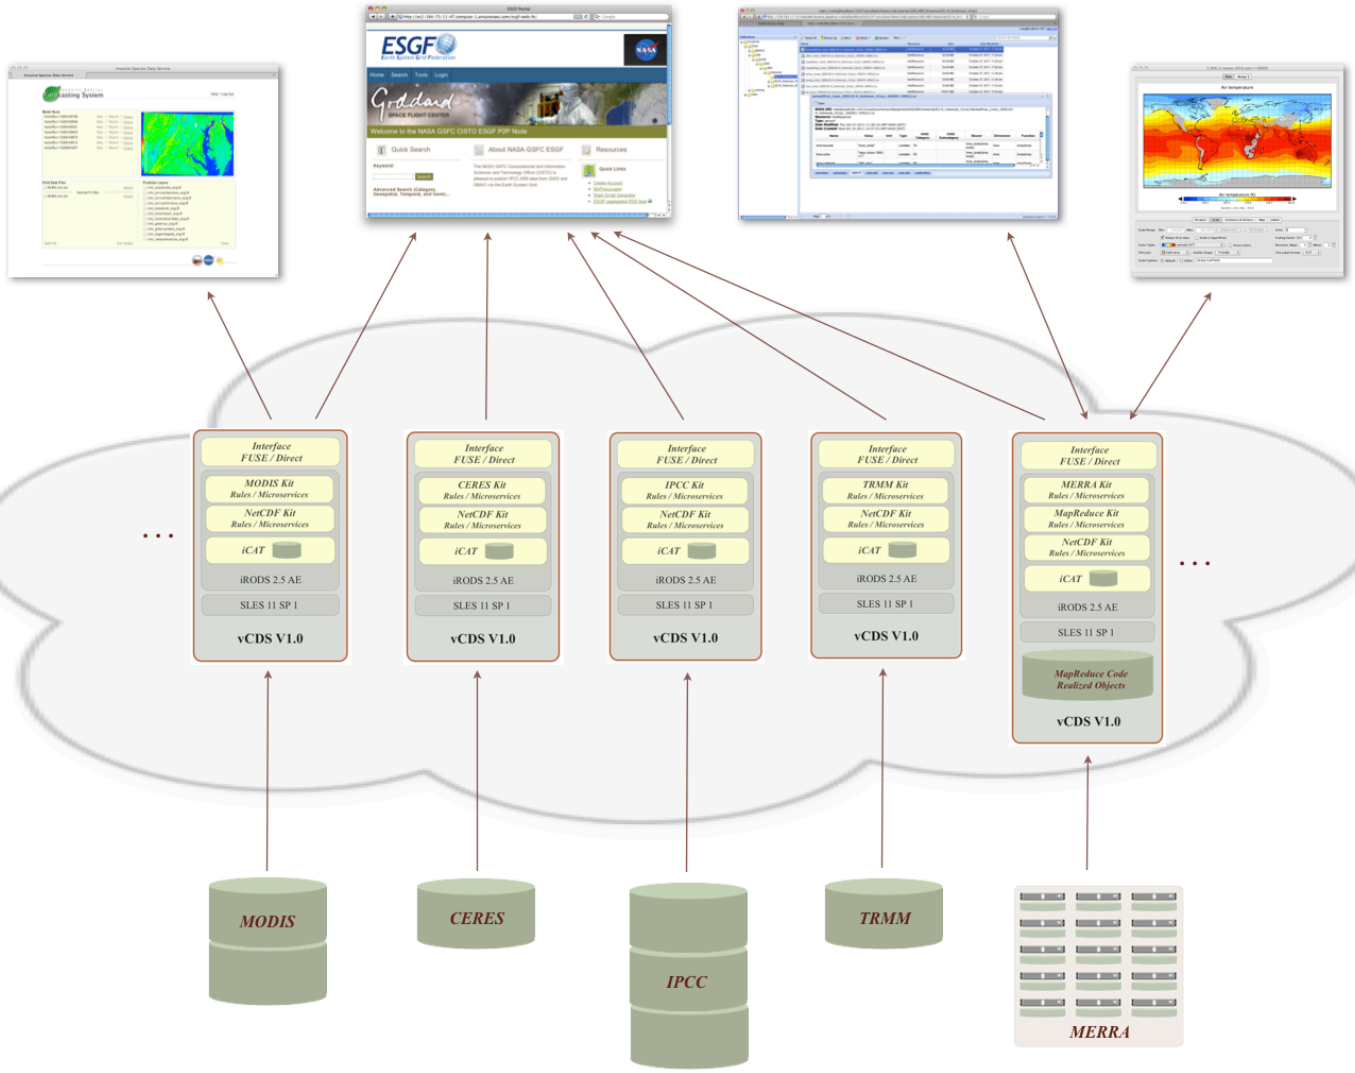

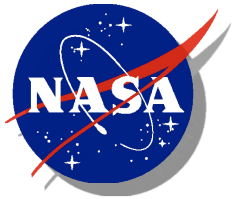

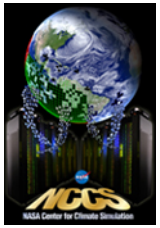

Thank you

ECMWF