



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

# Climate Science Responds to “Big Data” Challenges: Accessing Analyzing Model Output and Observations

October 3, 2012

**Dean N. Williams**

**On behalf of Multiple Earth System Communities and Projects**



This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Security, LLC, Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

# Outline: diverse research interests and coverage

---

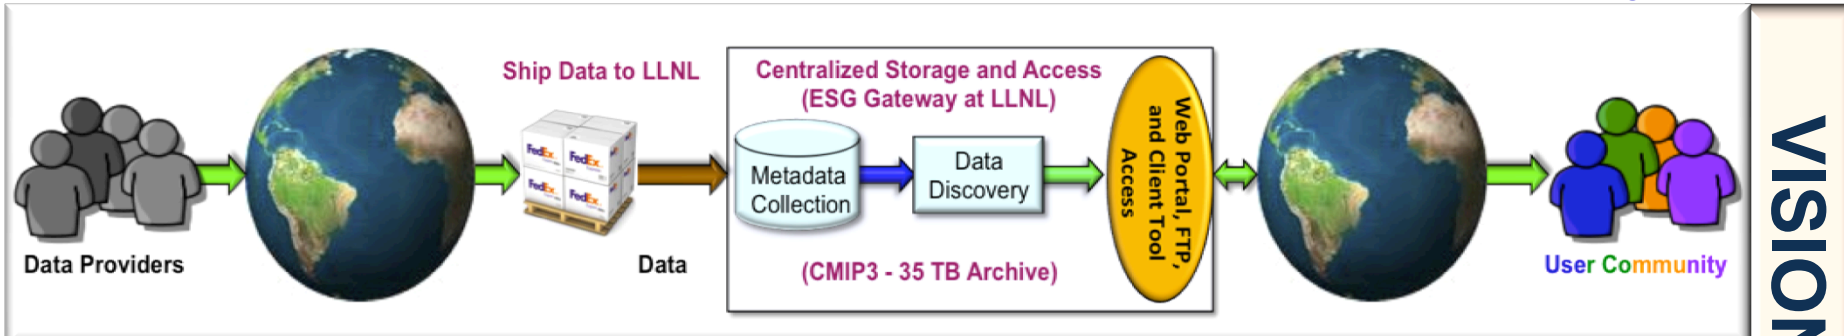
- **Enterprise Software & Systems**
  - Earth System Grid Federation
  - Data & Management (Game-Changer)
    - CMIP3 and CMIP5 – Vision for the future
  - Distributed and collaborative workflows, databases, and groupware
    - Large-scale scientific data management
    - Data publication and quality control
- **Large-scale analysis and visualization**
- **End-to-end simulation and analysis workflow**
- **Storage, Movement, and Computational Resources**
  - Computer hardware
  - Computer networks
  - High-performance compute cluster
- **Other associated projects**
- **Current status and future directions**

# A brief history of the Earth System Grid (ESG): ESG-I, ESG-II, ESG-CET, ESGF

- **ESG-I** funded under DOE' s Next Generation Internet (NGI) to address the emerging challenge of climate data 1999 – 2001 (ANL, LANL, LBNL, LLNL, NCAR, USC/ISI)
  - Data movement and replication; Prototype climate “data browser”; Hottest Infrastructure” Award at SC’ 2000.
- **ESG-II** funded under DOE’ s Scientific Discovery through Advanced Computing (SciDAC), turning climate data sets into community resources 2001-2006 (ORNL addition)
  - Web-based portal, metadata, access to archival storage, security, operational services, 2004 first operational portal CCSM (NCAR), IPCC CMIP3/AR4 (LLNL); 200 TB of data, 4,000 users, 130 TB served.
- **ESG-CET** funded under DOE’ s Offices of ASCR and BER to provide climate researchers worldwide with access to: data, information, models, analysis tools, and computational resources required to make sense of enormous climate simulation and observational data sets 2006 – 2011 (PMEL addition)
  - 2010 Awarded by American Meteorological Society (AMS) for leadership which led to a new era in climate system analysis and understanding.
  - CMIP3, CMIP5, CCSM, POP, NARCCAP, C-LAMP, AIRS, MLS, Cloudsat, etc.
  - 25,000 users, 500-800 users active per month, over 1 PB served
- **ESGF P2P** under the DOE’s Office of BER, it is an open consortium of institutions, laboratories and centers around the world that are dedicated to supporting research of climate change, and its environmental and societal impact. (Additional U.S. funding from NASA, NOAA, NSF.) The federation includes: multiple universities and institution partners in the U.S., Europe, Asia, and Australia.

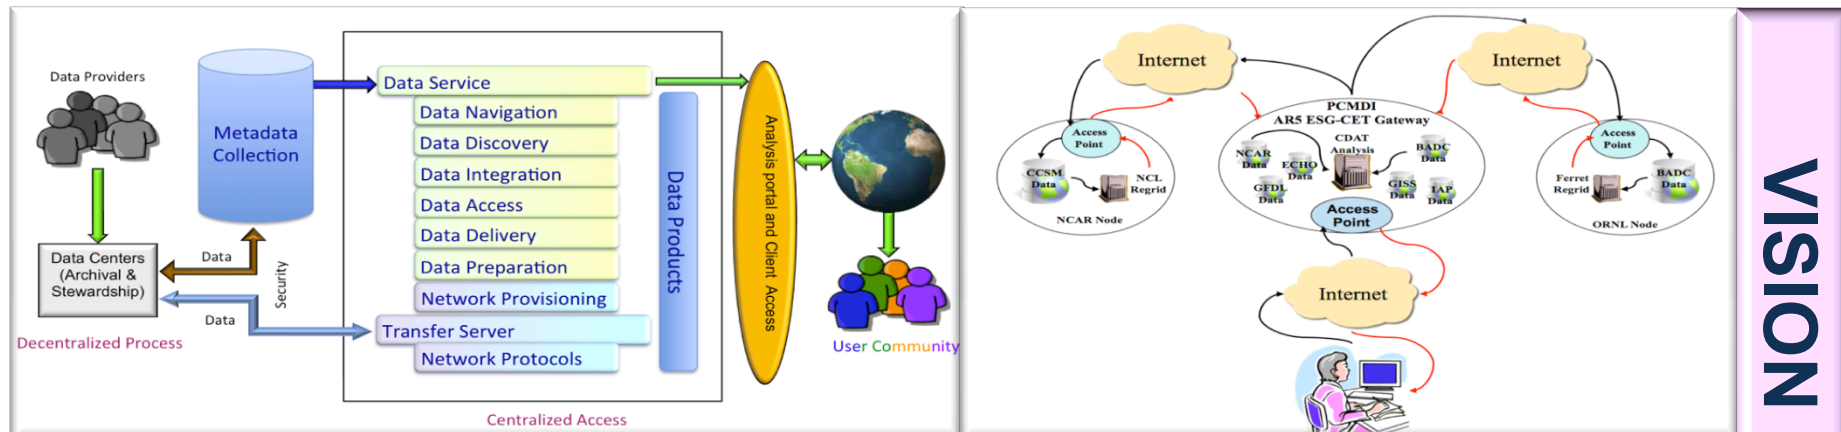
# Climate research requirements challenge data intensive science and results in a game-changing approach

- 2003 centralized processes for CMIP3. CMIP3, data was shipped to LLNL and stored in a single



“For leadership in implementing, maintaining, and facilitating access to the CMIP3 multi-model data set archive, which led to a new era in climate system analysis and understanding.”—AMS Award 2010

- 2006 decentralized process for CMIP5, a new cloud type approach for distributed data was necessary



“The new portal is wonderful! In one afternoon work achieved more than several weeks work with the old system.”— Most recent quote from a user 2012

VISION

VISION

# CMIP: experiment design

## CMIP5: 47 models **available** from 21 centers

- **CMIP = Coupled Model Intercomparison Project**
  - Phase 1: Idealized simulations of present-day climate ( ~1 Gigabyte (GB))
  - Phase 2: Idealized simulations of future climate changes (~500 GB: **CMIP2/CMIP1=500**)
  - Phase 3: More realistic simulations (2004 – present) (~35 Terabytes (TB): **CMIP3/CMIP2 = 70**)
- **CMIP 5 multi-model archive expected to include (3.5 Petabytes (PB) **CMIP5/CMIP3 = 100**):**
  - 3 suites of experiments
  - 24 modeling centers in 19 countries
  - 60 models
  - Total data, ~3.5 PB
  - Replica 1 – 2 PB
  - Derived data ~1 PB
- **Global distribution**
- **Timeline fixed by IPCC (2012 - 2013)**
- **LLNL organizes, manages and distributes** the CMIP/IPCC (Intergovernmental Panel on Climate Change) database of climate model output
- **CMIP6 (350 PB – 3 EB ?)**

kilobyte (kB)	$10^3$
megabyte (MB)	$10^6$
gigabyte (GB)	$10^9$
terabyte (TB)	$10^{12}$
<b>petabyte (PB)</b>	$10^{15}$
exabyte (EB)	$10^{18}$
zettabyte (ZB)	$10^{21}$
yottabyte (YB)	$10^{24}$

# Data challenge of CMIP3 archive vs. CMIP5 archive

CMIP3 Modeling Centers		volume (GB)
BCCR	Norway	862
CCCma	Canada	2,071
CNRM	France	999
CSIRO	Australia	2,088
GFDL	USA	3,843
GISS	USA	1,097
IAP	China	2,868
INGV	Italy	1,472
INMCM3	Russia	368
IPSL	France	998
MIROC3	Japan	3,975
MIUB	Germany/Korea	477
MPI	Germany	2,700
MRI	Japan	1,025
CCSM	USA	9,173
UKMO	UK	973
<b>Totals</b>		<b>34,989 (TB)</b>

**Archive size: 35 TB**

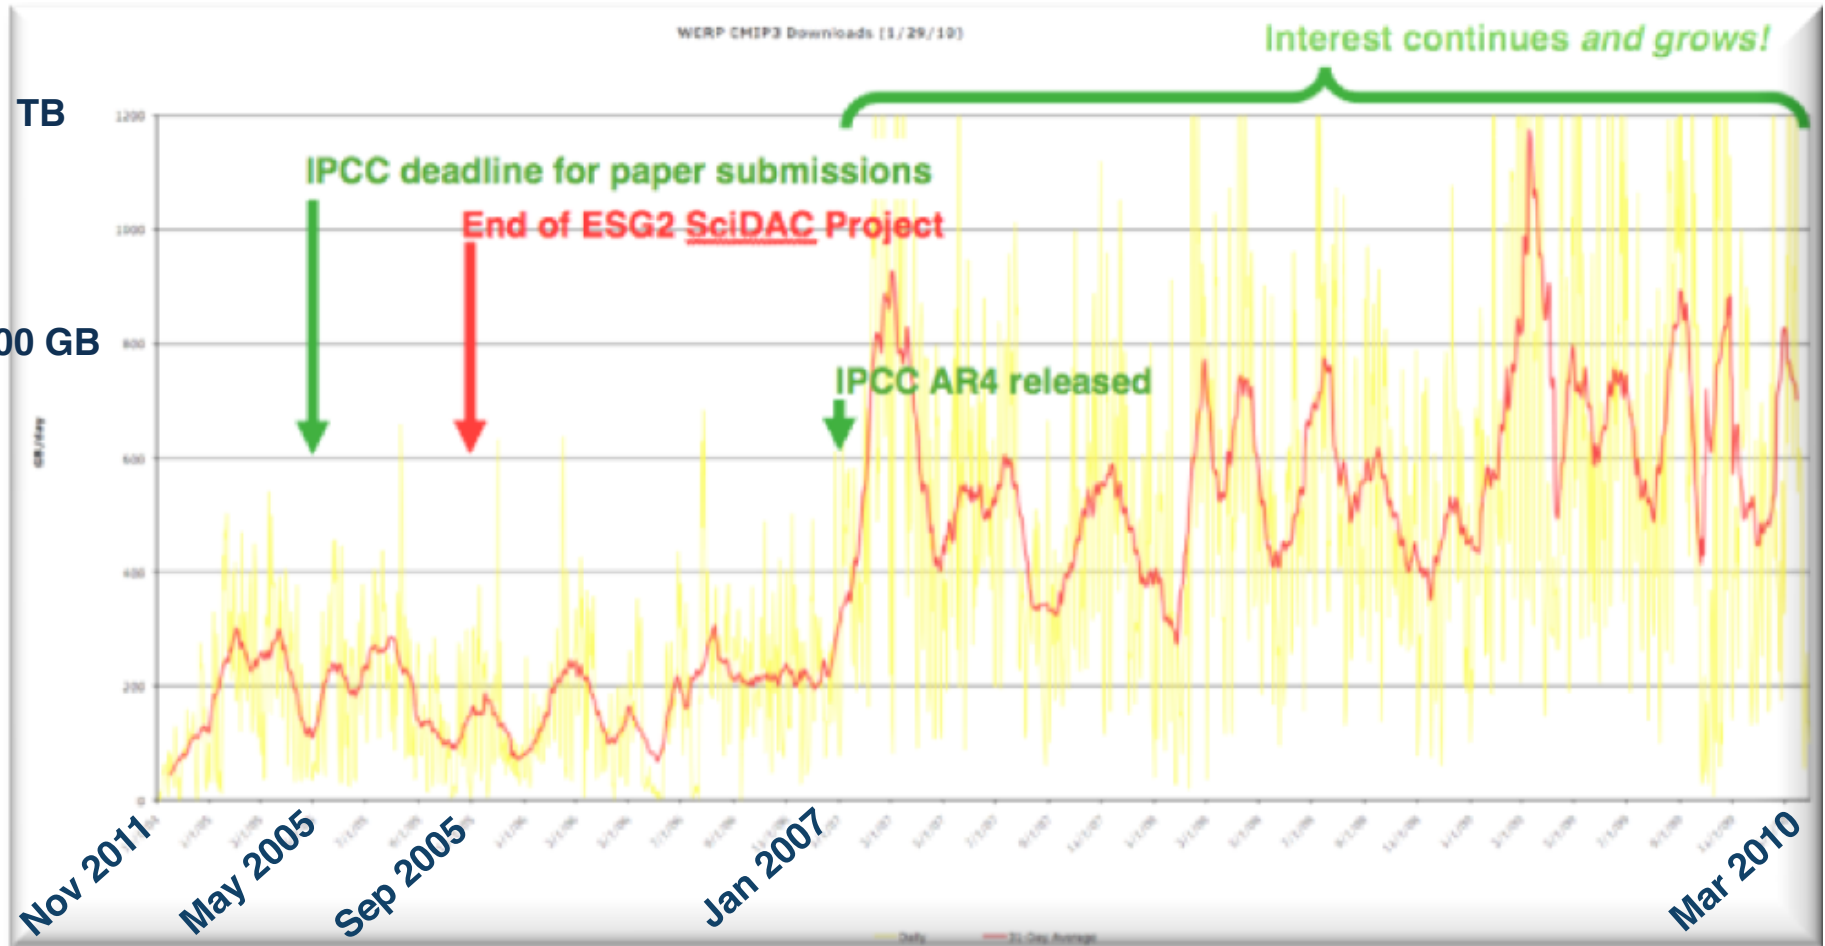
9/5/12

CMIP5 Modeling Centers		volume (TB)
BCC	China	51
CCCma	Canada	51
CMCC	Europe (Italy)	158
CNRM	France	71
CSIRO	Australia	81
EC-EARTH	Europe (Netherland)	97
GCESS	China	24
INM	Russia	30
IPSL	France	121
LASG	China	100
MIROC	Japan	350
MOHC	UK	195
MPI	Germany	166
MRI	Japan	269
NASA	USA	375
CESM	USA	739
NCC	Norway	32
NCEP	USA	26
NIMR/KMA	Korea	14
NOAA GFDL	USA	158
<b>Totals</b>		<b>3,108 (PB)</b>

**Archive size:  
currently: 1.4 PB  
total: 3.1 PB by 2013**

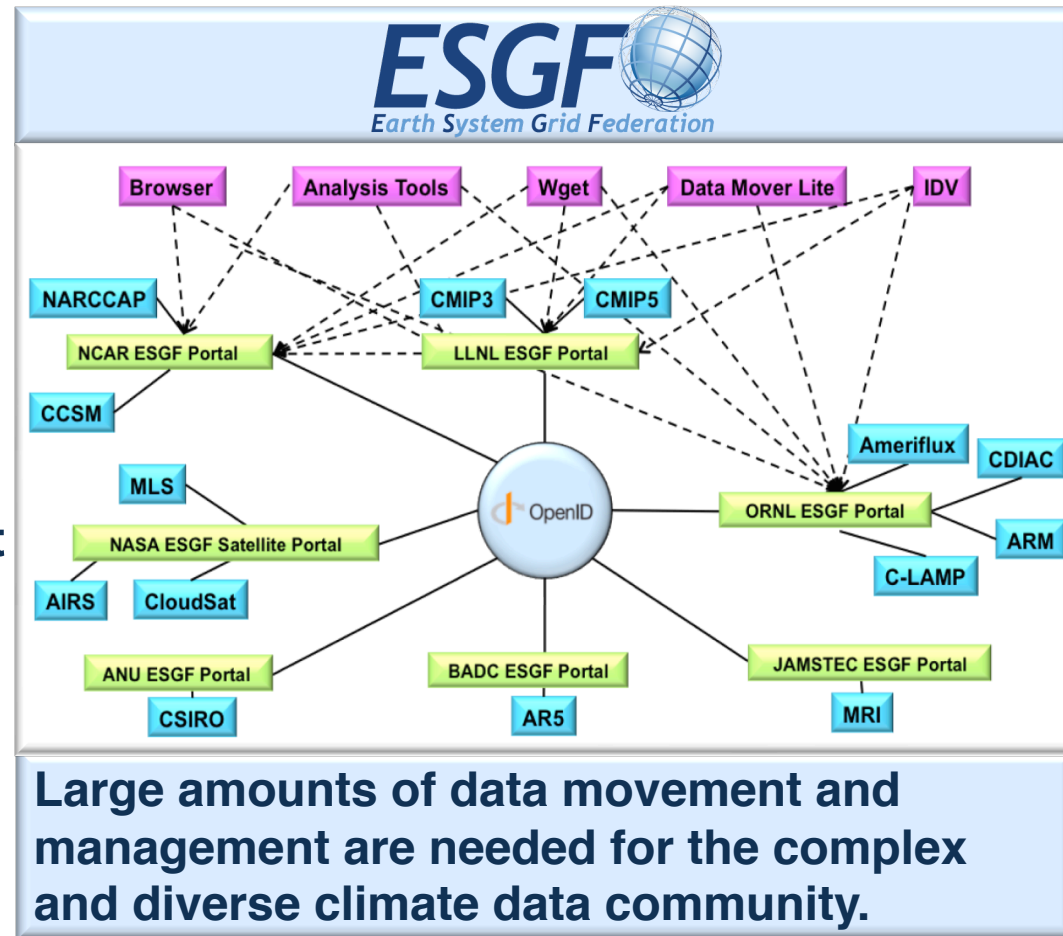
**CMIP5/CMIP3 = 10<sup>2</sup>**

# CMIP3 (IPCC AR4) Download Rates in Gigabytes per day



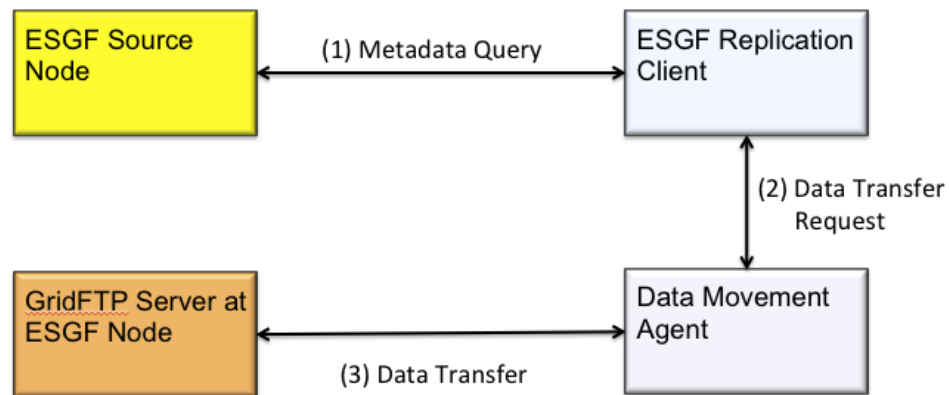
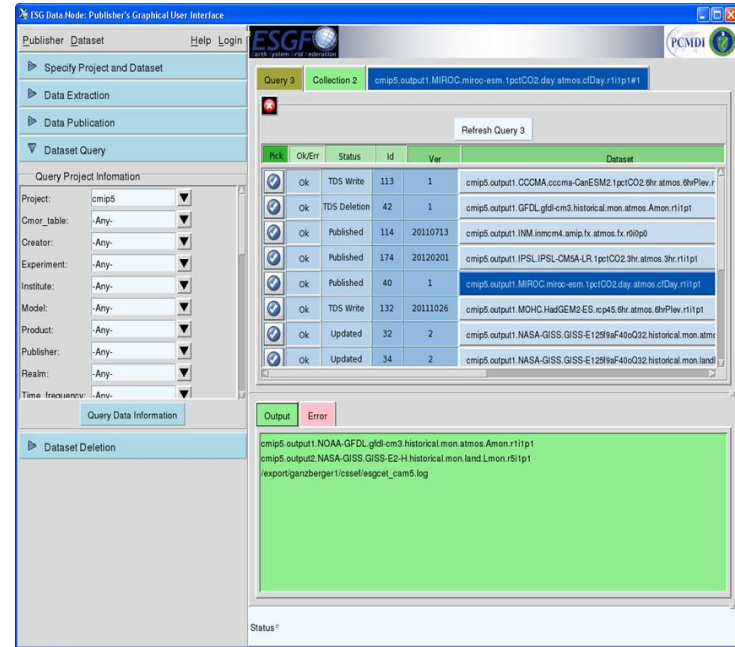
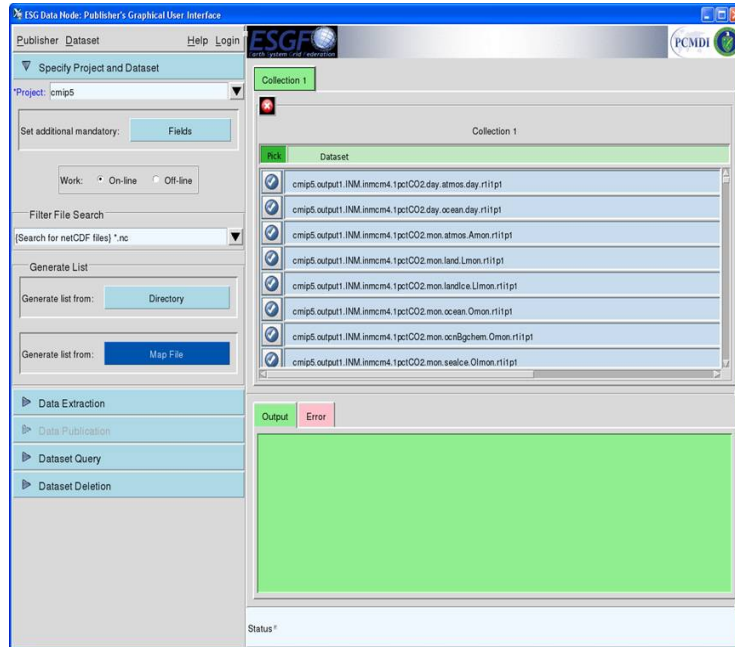
# The ESGF distributed data archival and retrieval system

- Distributed and federated architecture
- Support discipline specific portals
- Support browser-based and direct client access
- Single Sign-on
- Automated script and GUI-based publication tools
  - A collection of files, usually ordered by simulation time, that can be treated as a single file for purposes of data access, computation, and visualization
- User notification service
  - Users can choose to be notified when a data set has been modified






# Publication and smart replication



# Data quality control check operations end in digital object identifiers (DOIs)

## 3-Layer Quality Assurance Concept

QC Level 1 at ESGF Data Nodes



All CMIP5 requested data (CMOR-2)

QC Level 2 at Primary Sites



Replicated CMIP5 data

QC Level 3 at PubAgency



Replicated CMIP5 data and metadata



Receives DOI  
CMIP5 replicated data and metadata in long-term archive (IPCC DDC/WDC)

- **Publishing data to an ESGF portal performs QC Level 1 (QCL1) check**
  - QCL1 data are visible to users and are identified as QCL1 on the UI
- **DKRZ (MPI) quality control code is run on data to perform QC Level 2 (QCL2) check**
  - QCL2 data are visible to users and are identified as QCL2 on the UI
- **Visual inspections are performed for inconsistencies and metadata correctness at QC Level 3 (QCL3) check**
  - QCL3 data are visible to users and are identified as QCL3 on the UI
  - **Digital Object Identifiers (DOIs)** are given to data sets that pass the QCL3 check

# CMIP5 QC status

---

- **QC Status of CMIP5 as of 08/2012:**

- **Quality Control 1:** 54,269 data sets
- **Quality control 2:** 10,986 data sets (finalized 6535)
- **Quality Control 3:** 1468 data sets
- **DataCite DOI:** 1468 data sets

For more information on QC Status, see the following URLs:

<https://redmine.dkrz.de/collaboration/projects/cmip5-qc/wiki>

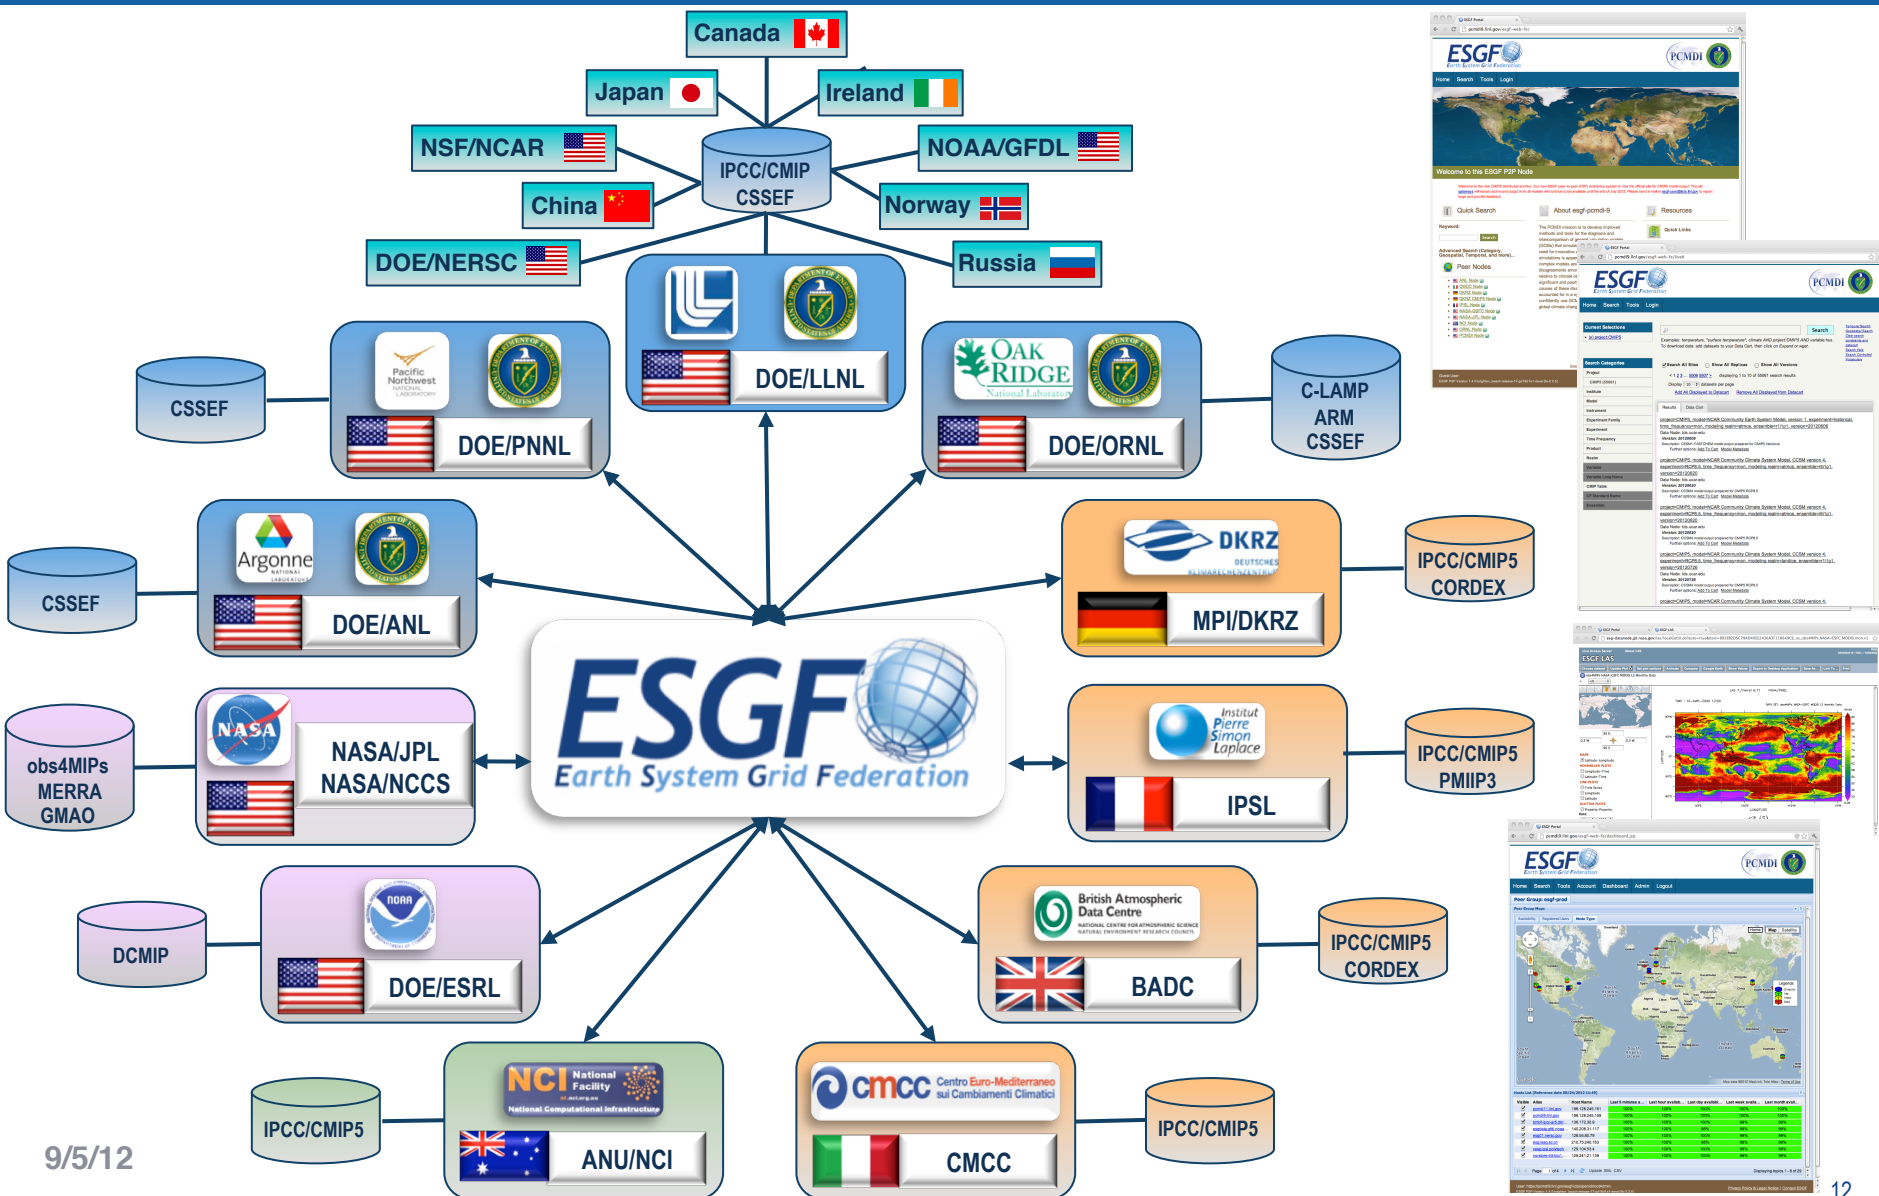
<http://cera-www.dkrz.de/WDCC/CMIP5/QCResult.jsp>

- **Replicated CMIP5**

- **LLNL replicated data sets** 23,745 data sets
- **DKRZ replicated data sets** 23,745 data sets

# ESGF is more than CMIP: federated and integrated data from multiple sources

<http://pcmdi9.llnl.gov/esgf-web-fe/>



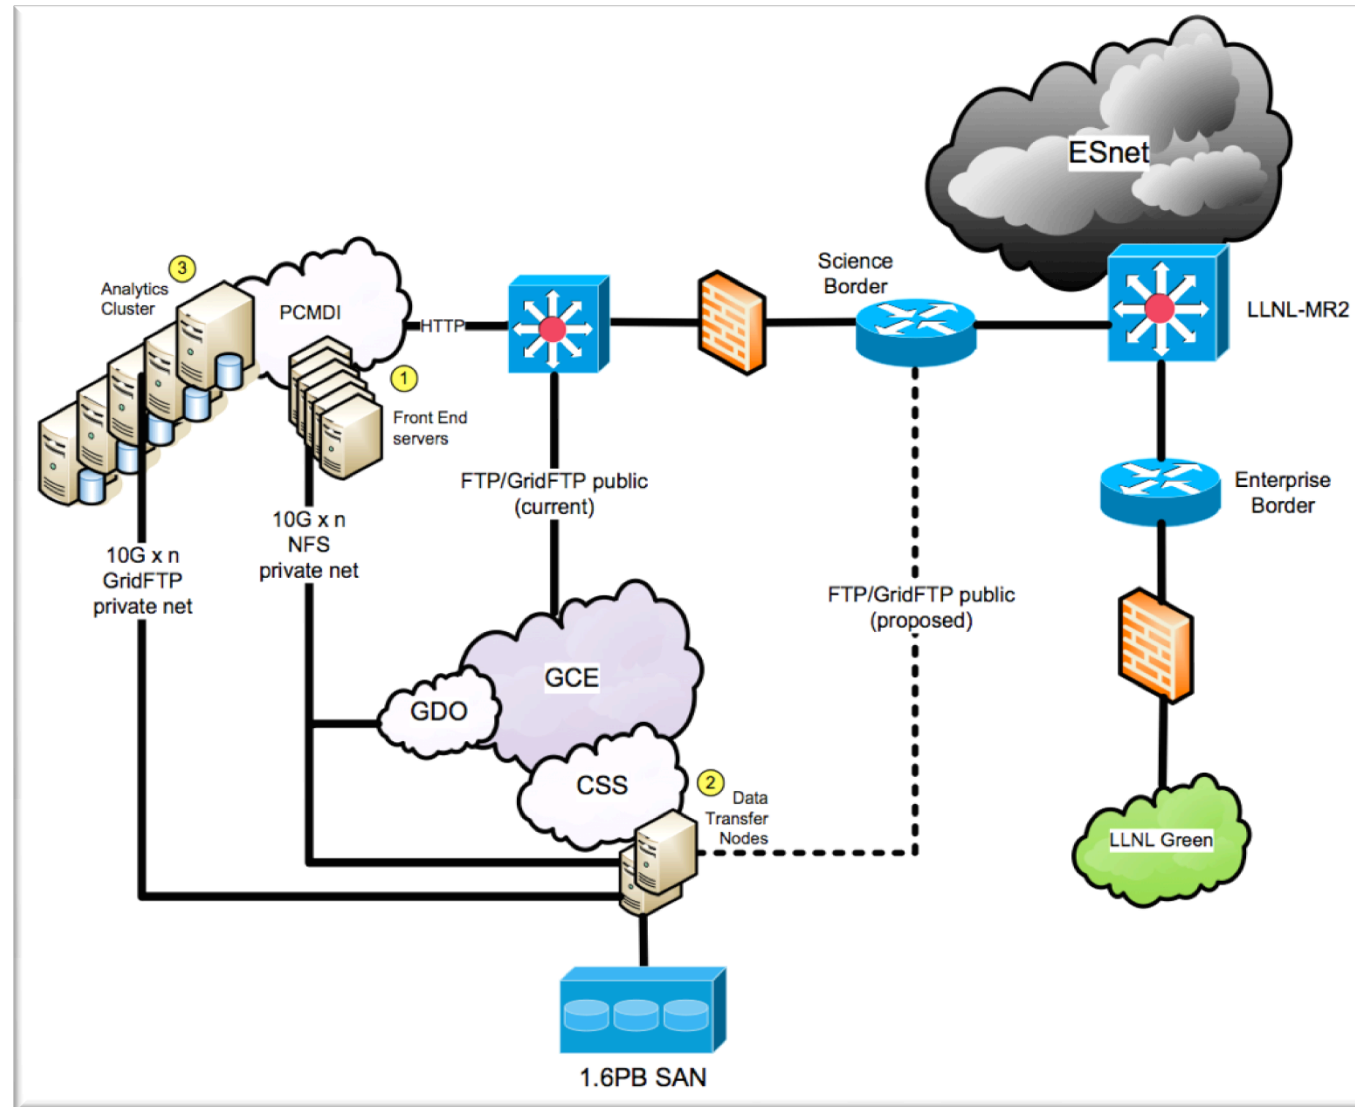
# ESGF Governance

---

- **Governance – critical ideas of fairness, transparency, responsibility, accountability**
- **The ESGF Review Board (ERB) consist of two committees:**
  - **Executive Committee**
    - General guidance and overall high-level decisions in directing the course of the ESGF project in correspondence with multiple sponsor and stakeholders needs—ultimate responsibility for ensuring that ESGF meets the needs of customers and stakeholders.
  - **Technical Committee**
    - General guidance and overall high-level decisions in directing the course of the ESGF project in correspondence with multiple sponsor and stakeholders needs—ultimate responsibility for ensuring that ESGF meets the needs of customers and stakeholders.

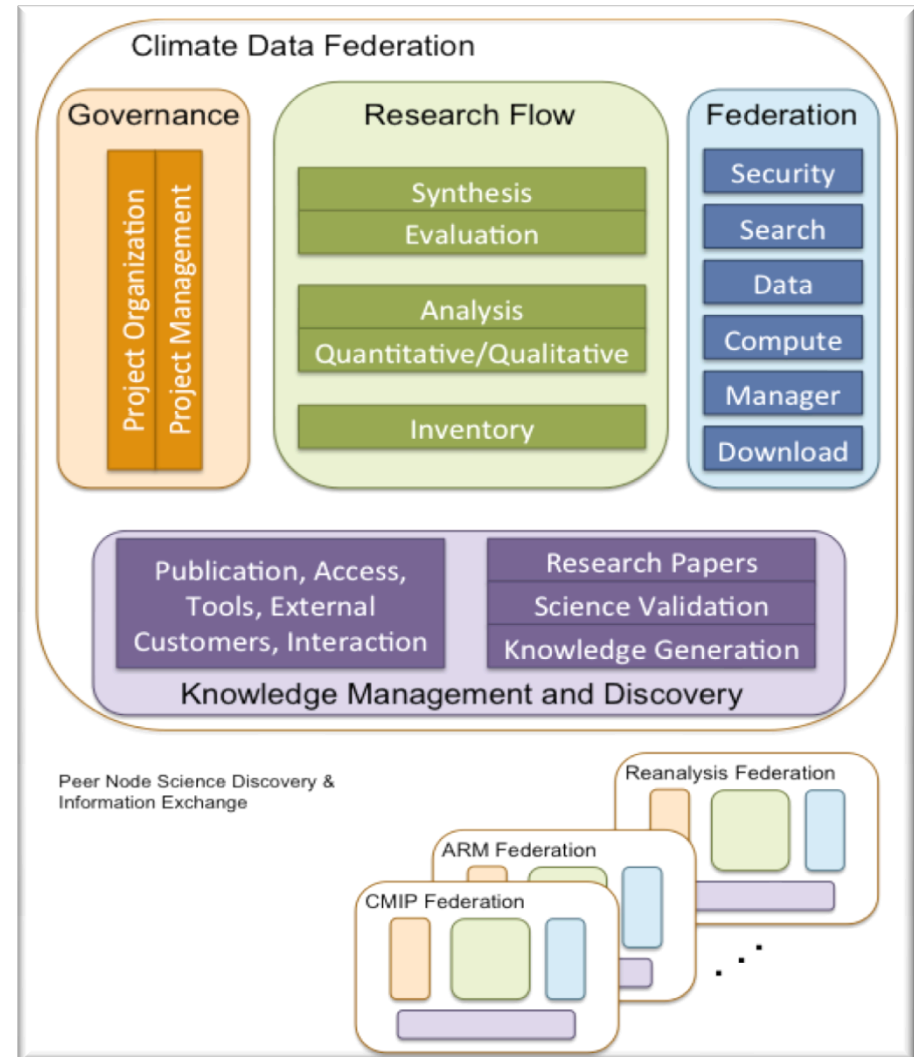
# ESGF integrated with the hardware and network

- 1 Users communicate with ESGF front-end servers via HTTP
- 2 Large data sets are made available to users directly from the Climate Storage System (CSS) via vsftp and GridFTP
- 3 Through UV-CDAT, ESGF will perform analysis of raw data if requested by users through the front-end servers to the analysis (hadoop) cluster






# ESGF software system integrates data federation services

- NetCDF Climate and Forecast (CF) Metadata Convention
  - (LibCF)
  - Mosaic
- Climate Model Output Rewriter 2 (CMOR-2)
- Regriders: GRIDSPEC, SCRIP, & ESMF
- Publishing
- Search & Discovery
- Replication and Transport
  - GridFTP, OPeNDAP, DML, Globus Online, ftp, BeSTMan (HPSS)
  - Networks
- Data Reference Syntax (DRS)
- Common Information Model (CIM)
- Quality Control
  - QC Level 1, QC Level 2, QC Level 3, Digital Object Identifiers (DOIs)
- Websites and Web Portal Development
  - Data, Metadata, Journal Publication Application
- Notifications, Monitoring, Metrics
- Security
- Product Services
  - Live Access Server, UV-CDAT



# Advanced analytics, informatics, and visualization for scientists

- Analysis and visualization is a key aspect of scientific analysis and discovery
- Advanced interactive visualization is rarely used by scientists
- Interfaces too complex, pickup too costly
- Interactive climate visualization requires:
  - Intuitive interfaces
  - Seamless integration with high performance analysis workbenches
  - Parallel streaming visualization pipelines



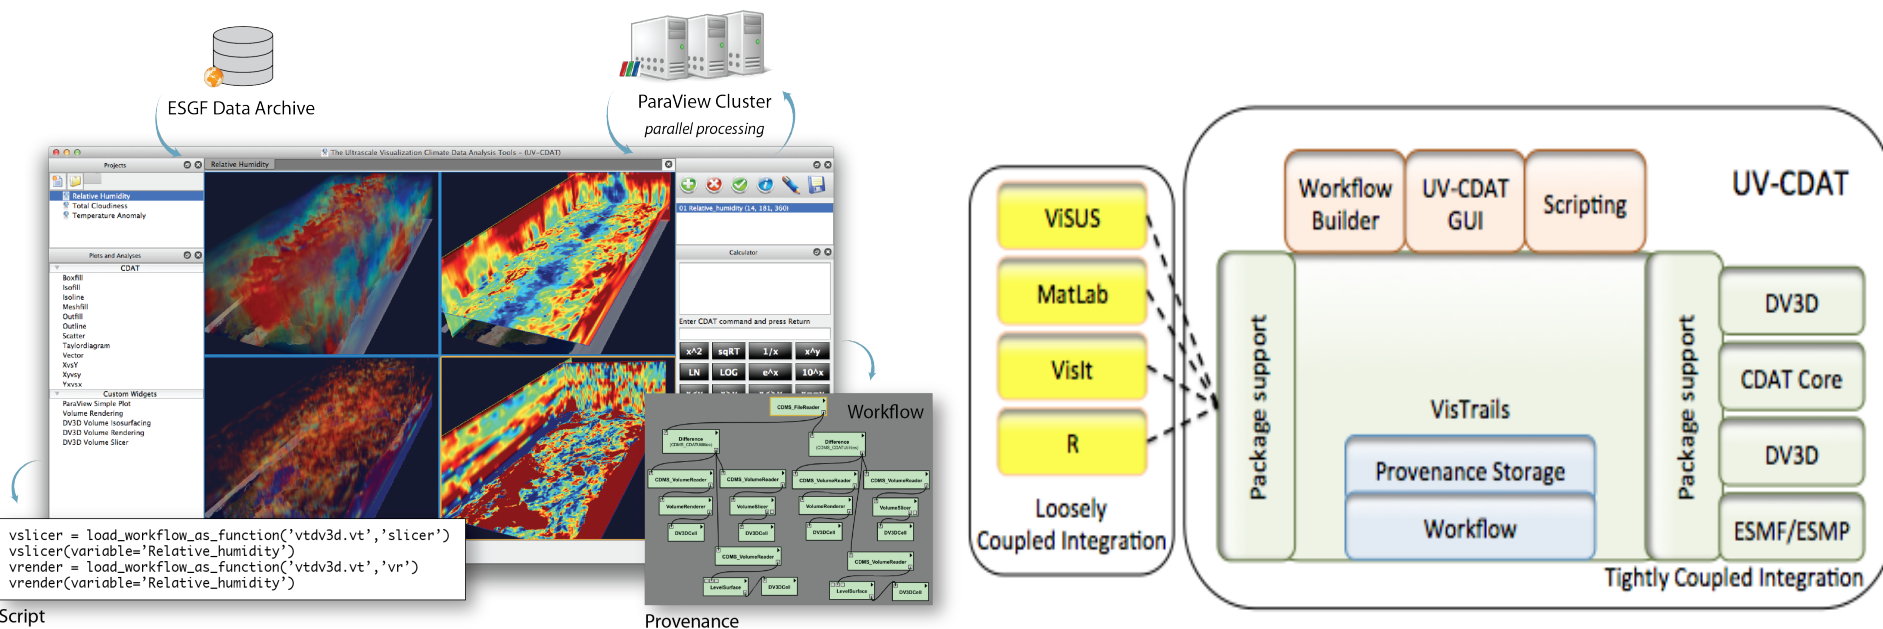


# Supported data activities (complementary BER funded project)

<http://uvcdat.llnl.gov>

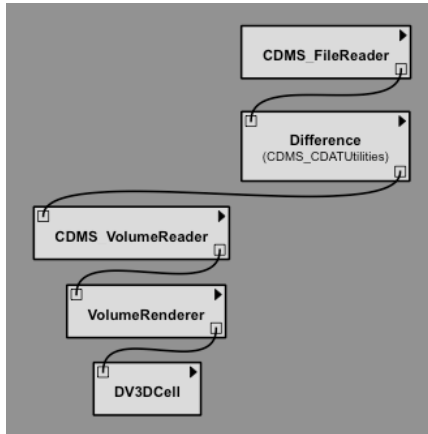
## Ultra-scale Visualization Climate Data Analysis Tools (UV-CDAT):

- Integrate DOE's climate modeling and measurements archives
- Develop **infrastructure** for national and international model/data comparisons
- Deploy a wide-range of climate data visualization, diagnostic, and analysis tools with familiar interfaces for very large, high resolution climate data sets (**CDAT, VTK, R, VisIt, ParaView, DV3D, ...**)
- Workflow** – data flows are directed graphs describing computational tasks
- Takes advantage of **ESGF** data management




# Exploitation of hardware and system parallelism for climate data analysis

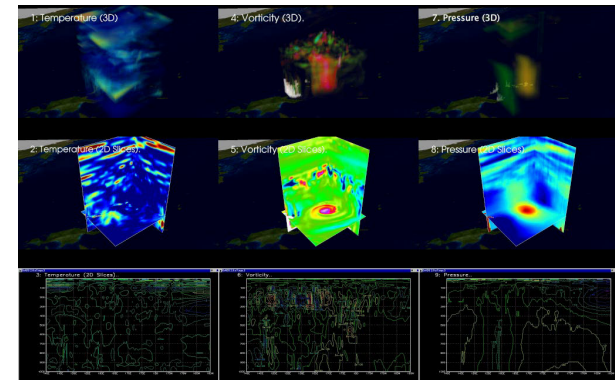
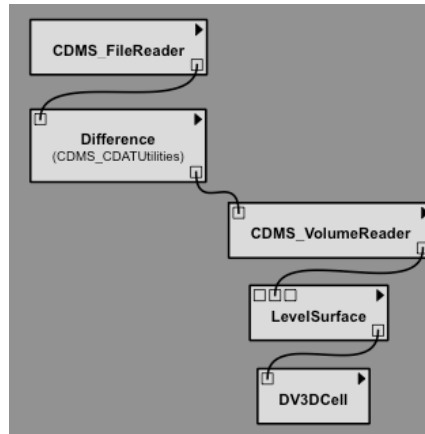
Client-0



Client-1

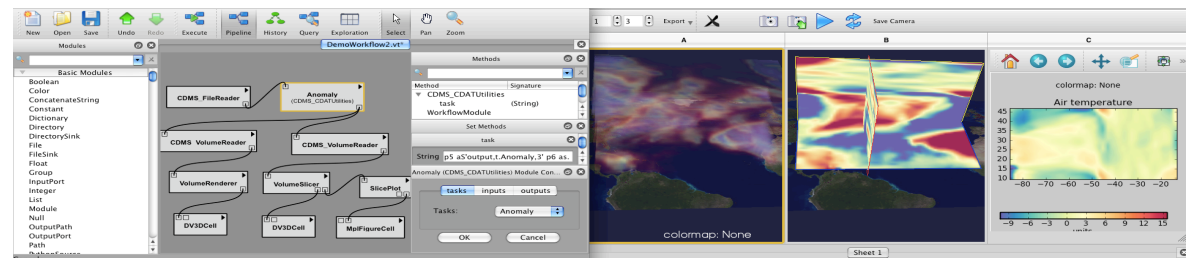
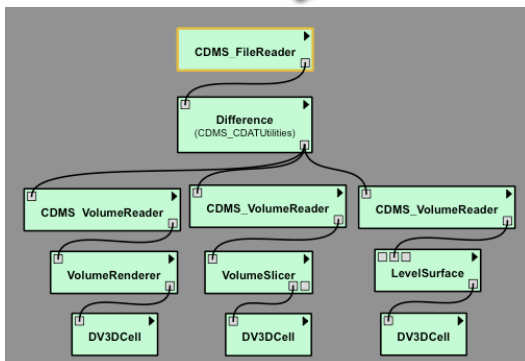


Client-2



- Data Parallelism
- Task Parallelism
- Time Parallelism

Server



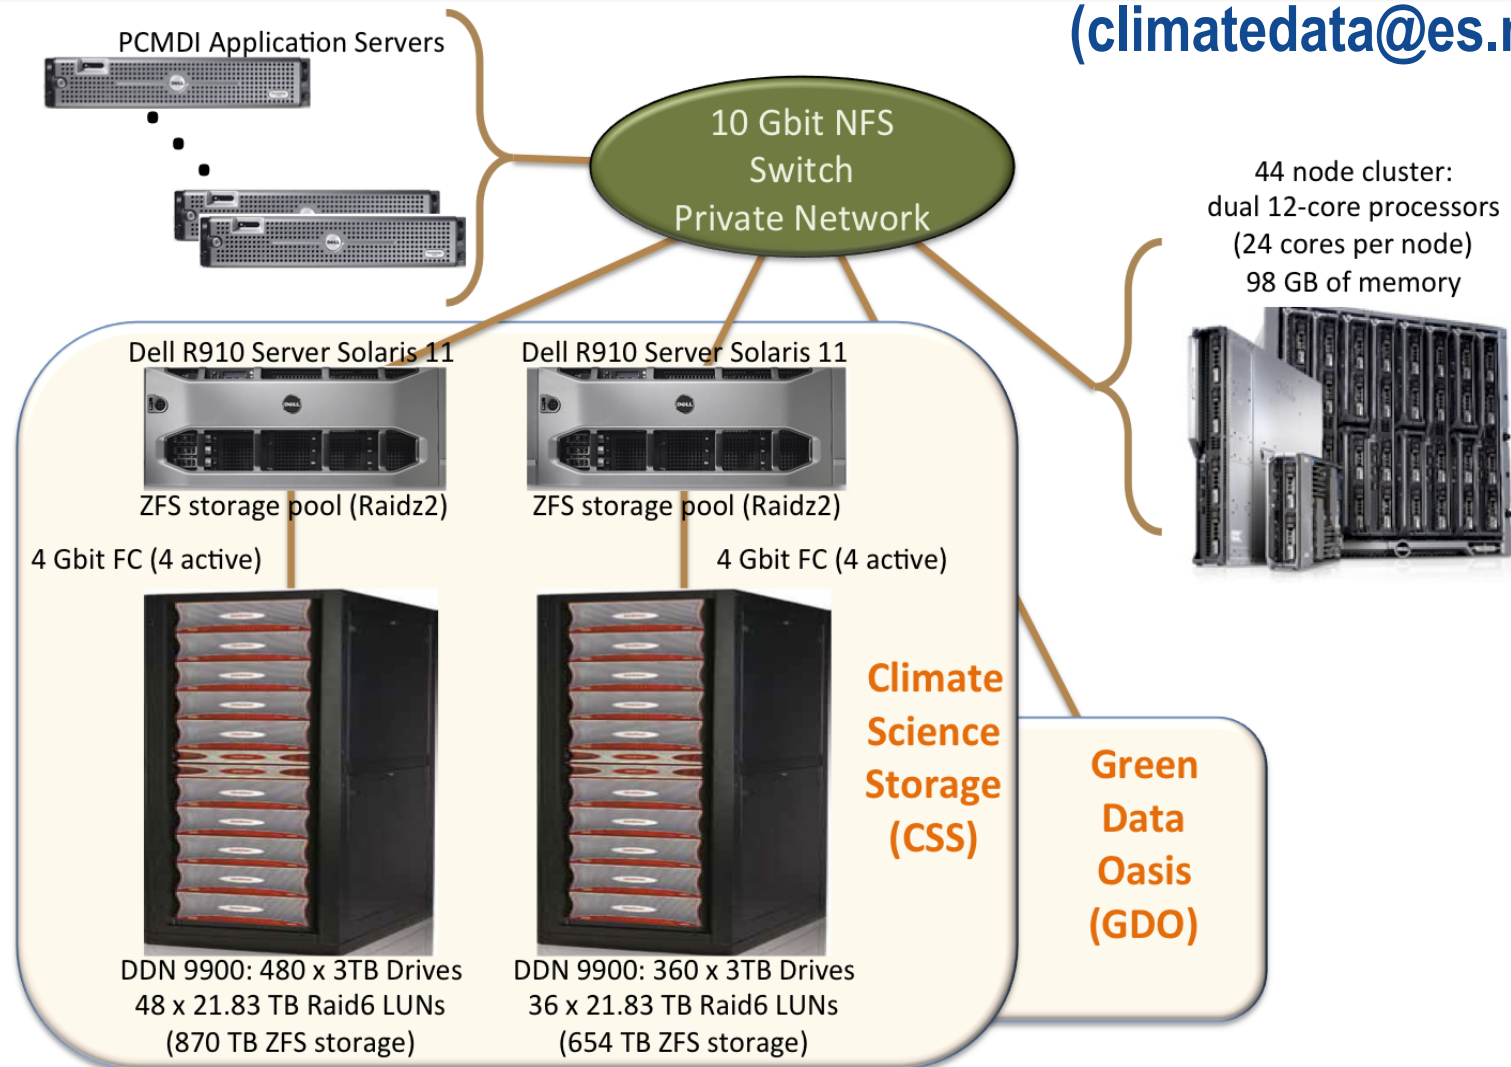
9/5/12

D. N. Williams, IGIM



# ASCR Climate 100 project: connecting our system to the largest network in the world and organizing the climate network community

## ESnet 100 Gbps Internet

(climatedata@es.net)



# ESGF's multiple collaborations



# Notable national and international science data systems interested in using ESGF

- The NASA Global Change Master Directory (GCMD) is a well-established clearinghouse for geoscientific metadata records based on the popular DIF metadata standard.
- The British Atmospheric Data Center (BADC) is a Python-based infrastructure that shares many technologies with ESGF, and is interoperable based on common APIs and trust relations.
- The NOMADS network of NOAA sites is serving a plethora of heterogeneous Earth datasets via TDS servers and OpenDAP, sometimes enhanced by a Live Access Server.
- The USGS Geo Data Portal provides access to climate data sets and derivative products through standard OGC protocols.
- The Mercury system is a web portal deployed at ORNL that provides extensive metadata capabilities and data access for DOE observations and models.
- The Planetary Data System (PDS) is deployed at NASA and other international Space centers, is based on a federation model similar to ESGF, and uses some of the same search technology such as Solr.
- The Regional Climate Model Evaluation System (RCMES) provides a common data model for accessing model and observational climate data sets at regional scale.
- The European Infrastructure for the European Network for Earth System Modeling II (IS-ENES II), developers from BADC, DKRZ, IPSL, and CMCC will continue to funding to work on ESGF and related model data archiving issues

# Bring Together Large Volumes of Diverse Data at any Velocity to Generate New Insights

## Data integrating enterprise system

Insight into big data reveals three very significant challenges:

- **Variety:** managing complex data, including storage and retrieval, from multiple regional and non-regional data indices, types and schemas
- **Velocity:** distributing live data streams and large volume data movement quickly and efficiently
- **Volume:** analyzing large-volume data (from terabytes to zettabytes) in-place for big data analytics

BER and ASCR invests in:

- **Accessing Global Information:** Accessing climate data and content information from everywhere via the web, sensors, and applications in an integrated and federated environment
- **Flexible Infrastructure:** Flexible automated administration, easy-to-use analytics, and virtualization at every level
- **Scalable Framework:** Big data analytics in a scalable environment with efficient parallelism, workload-optimization, and real-time streaming process


Simulation



Observation



Reanalysis



# ESGF Roadmap: Vision for the Future

*Testbed*

*Production*

*Interoperability Across Science Domains*

## ESGF Data System Evolution

**2013**

### Foundation Development

- ESGF architecture refinement for science domain use case studies for diverse data sets
  - Ontology & Provenance
- ESGF collaborative distributed analytics infrastructure using UV-CDAT
  - Local and remote analysis
  - Enabling reproducibility via workflow
- GIS Integration
- Training and documentation

**2014 - 2015**

### Integration and Release

- Expanded to other science domains
- Full suite of server-side analysis and visualization
- Machine learning for pattern discovery and prediction
- Decision analytics based quantifying uncertainties
- Streaming analysis, visualization and sensors
- Model intercomparison metrics
- Training and documentation

**2016**

### Evaluation and Deployment

- Evaluation and Iterative science domain community feedback and upgrade
- Debugging
- Continued user feedback
- Operational transitions support by domain
- Extended community training and documentation

**Climate**

**ESGF Science Domains**

**Petabytes ( $10^{15}$ )**

**Exabytes ( $10^{18}$ )**

**Astrophysics,  
Biology, Chemistry,  
Climate, Combustion,  
Energy, Fusion, Materials,  
Nuclear Energy**

# Questions and discussion

---

