ESG and Cloud Computing with an experience from Exploring Cloud for Parallelizing Tropical Storm Tracking

Alex Sim

Scientific Data Management Research Group
Computational Research Division
Lawrence Berkeley National Laboratory

Earth System Grid

- **Earth System Grid (ESG)**
  - To support the infrastructural needs of the national and international climate community, ESG is providing crucial technology to securely access, monitor, catalog, transport, and distribute data in today's grid computing environment.
  - ANL, LANL, LBNL, LLNL, NCAR, ORNL, PMEL, USC/ISI
- **ESG's mission is to provide climate researchers worldwide with access to**
  - Data,
  - Information,
  - Models,
  - Analysis tools, and
  - Computational resources

  required to make sense of enormous climate simulation datasets.
- **Project history**
  - ESG-I (1999-2001)
  - ESG-II (2001-2006)
  - ESG-CET (2006-present)
- **Production since 2004**

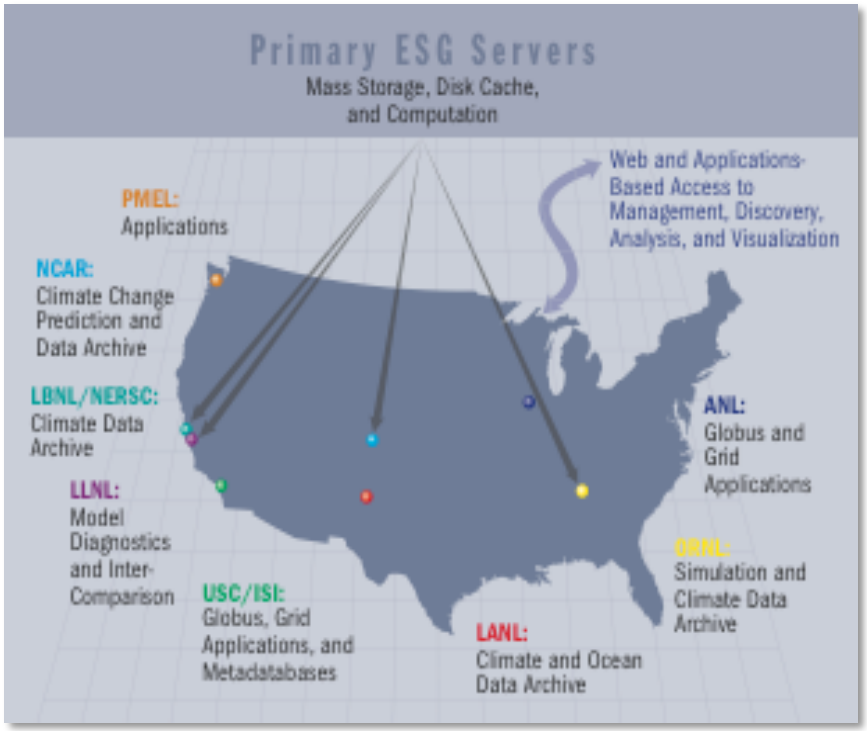

ESG current statistics

- **NCAR CCSM ESG portal**
  - 237 TB of data at four locations (NCAR, LBNL, ORNL, LANL) : 965,551 files
    - Includes the past 7 years of joint DOE/NSF climate modeling experiments
- **LLNL CMIP-3 (IPCC AR4) ESG portal**
  - 35 TB of data at one location
    - 83,337 files, model data from 13 countries
    - Generated by a modeling campaign coordinated by the Intergovernmental Panel on Climate Change (IPCC)
    - Over 565 scientific peer-review publications
- **Serving data to the community**
  - Coupled Model Intercomparison Project, Phase 3 (CMIP-3)
  - Community Climate System Model (CCSM)
  - Parallel Climate Model (PCM)
  - Parallel Ocean Program (POP)
  - The North American Regional Climate Change Assessment Program (NARCCAP)
  - Cloud Feedback Model Intercomparison Project (CFMIP)
  - Carbon-Land Model Intercomparison Project (C-LAMP)

ESG current download statistics

- **Geographic distribution of the users that downloaded data from ESG web portals**
  - Over 2,700 sites
  - 120 countries
  - 25,000 users
  - Over 1 PB downloaded

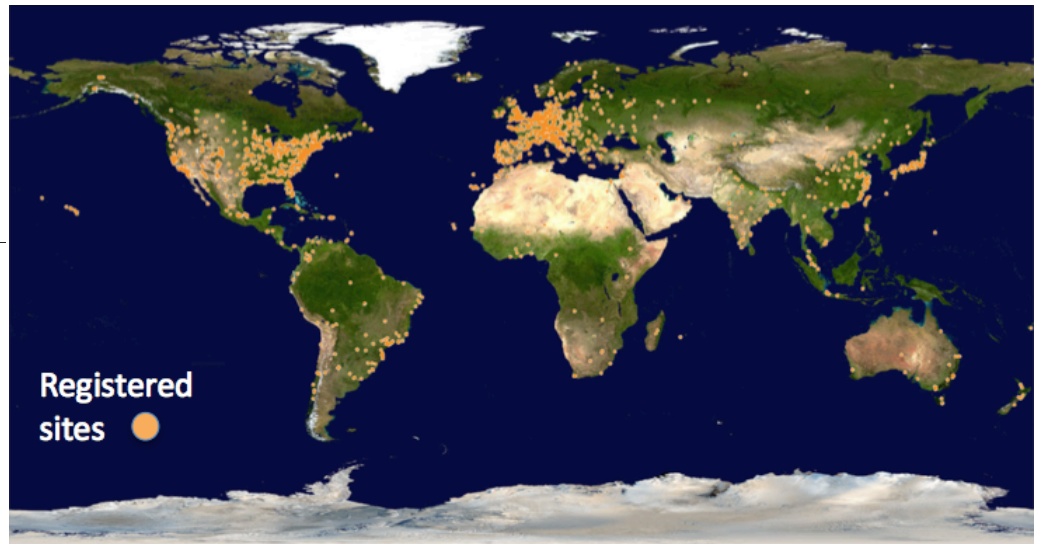

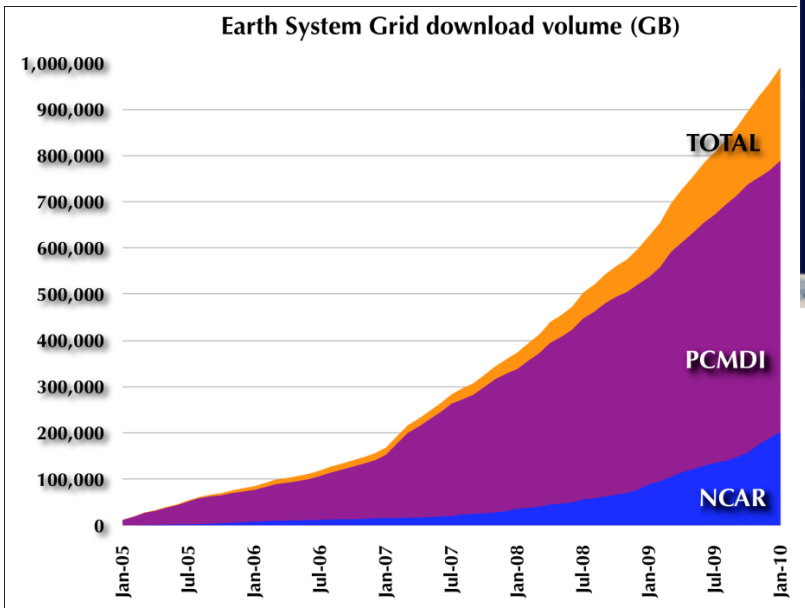

The Growing Importance of Climate Simulation Data

- **Broad investments in climate change research**
  - Development of climate models
  - Climate change simulation
  - Model intercomparisons
  - Observational programs
- **Climate change research is increasingly data-intensive**
  - Analysis and intercomparison of simulation and observations from many sources
  - Data used by model developers, impacts analysts, policymakers

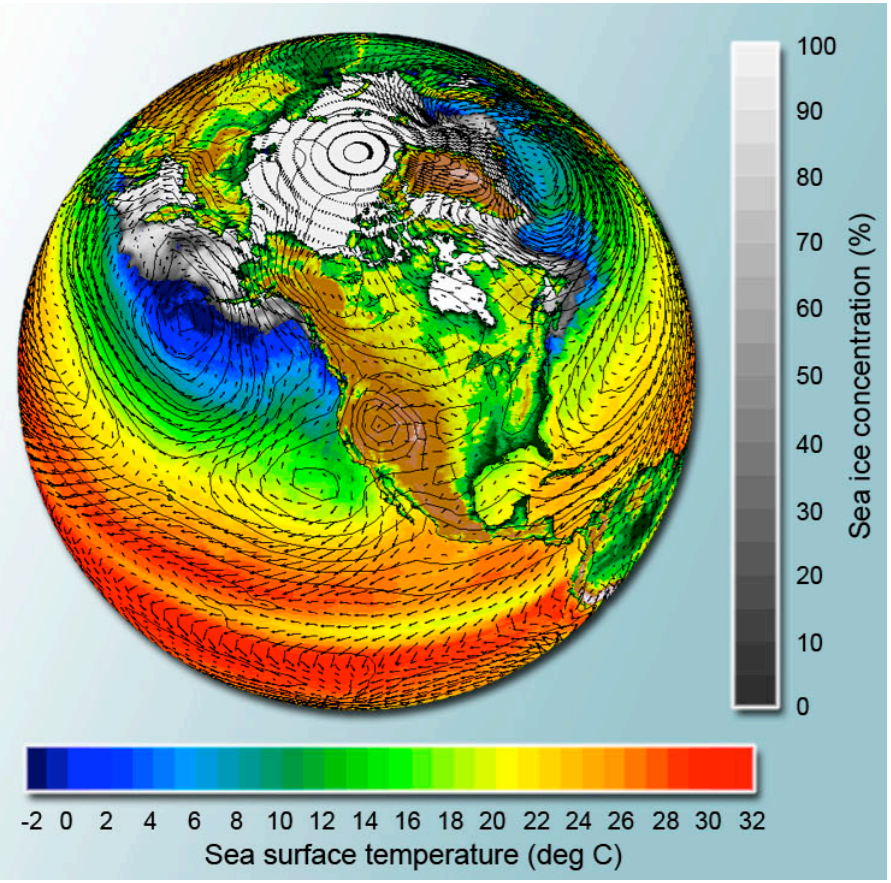

Results from the Parallel Climate Model (PCM) depicting wind vectors, surface pressure, sea surface temperature, and sea ice concentration. Prepared from data published in the ESG using the FERRET analysis tool by Gary Strand, NCAR.

The Growing Size of Climate Simulation Data

- Early 1990's (e.g., AMIP1, PMIP, CMIP1)
  - modest collection of monthly mean 2D files: ~1 GB
- Late 1990's (e.g., AMIP2)
  - large collection of monthly mean and 6-hourly 2D and 3D fields: ~500 GB
- In 2000's (e.g., IPCC/CMIP3)
  - fairly comprehensive output from both ocean and atmospheric components; monthly, daily, and 3 hourly: ~35 TB
- In 2011:
  - The IPCC 5th Assessment Report (AR5) in 2011: expected 5 to 15 PB
  - The Climate Science Computational End Station (CCES) project at ORNL: expected around 3 PB
  - The North American Regional Climate Change Assessment Program (NARCCAP): expected around 1 PB
  - The Cloud Feedback Model Intercomparison Project (CFMIP) archives: expected to be .3 PB
- CMIP5 is being defined now, available info neither complete nor final
  - Current estimates… 1.2 to 2 PB of "replica core archive" (RCA) results
    - In CMIP5, the RCA is expected to be 20% to 30% of total volume of data produced
- Climate model data is projected to exceed hundreds of Exabytes by 2020
  - (BES Science Network Requirements Workshop, 2007)

Emerging Requirements

- Data search, accessibility, versioning, replications, metrics…
- Data Analysis Requirements
  - Scalability
    - Data access from terabytes to petabytes to exabytes (many thousands of files)
    - Efficient handling on extreme variance in file sizes
    - Extensive data management such as packaging data into “containers”
    - Big coordination between storage and network, and between parallel and distributed processes
  - Reliability and Robustness
    - Asynchronous long-lasting operation
    - Recovery from transient failures and automatic restart
    - Work with local and remote security policies
    - Multiple network transfer protocol support
  - On-demand status
    - Estimation of analysis completion time
    - Statistics collection
  - Networking
    - Remote access to large datasets
    - Pre-staging required datasets based on the network performance

Cloud Computing in Climate Analysis

- Cloud computing is emerging as an alternative to dedicated computing clusters, likely to replace a majority of the computer clusters
  - No need for tightly coupled synchronization like MPI processes
  - Climate analysis typically processes one file at a time independently
- For climate analysis, the major attraction is it provides an alternative to rewrite existing sequential climate analysis programs in MPI
  - Less programming effort
- Virtualized computing node can replicate the environment required by the analysis codes and satisfy complex dependencies
- Needs alternative strategies of coordinating the jobs in virtual machines
- Reasons of not-using Hadoop
  - Hadoop typically duplicates data which increases storage requirements
  - Hadoop distributes data by size which may separate header from data under different nodes

Example: Finding Tropical Cyclones

- fvCAM (finite volume version of the Community Atmospheric Model) dataset (version 2.2)
  - 1/2° resolution, 6-hour interval, 3-year simulated time, 500 GB, 1000 netCDF files
- TSTORM code used to track tropical storms
  - Based on the criteria established by Knutson, et al., 2007 BAMS 88:10 1549-65
  - Searches for high vorticity, local pressure drop, and warm core
- Reference: "*Tropical Cyclones on a Cloud Computing Cluster: Using Parallel Virtualization for Large-Scale Climate Simulation Analysis*", D. Hasenkamp, A. Sim, M. Wehner, K. Wu, Proceedings of the 2nd IEEE International Conference on Cloud Computing Technology and Science, 2010

Tropical storms

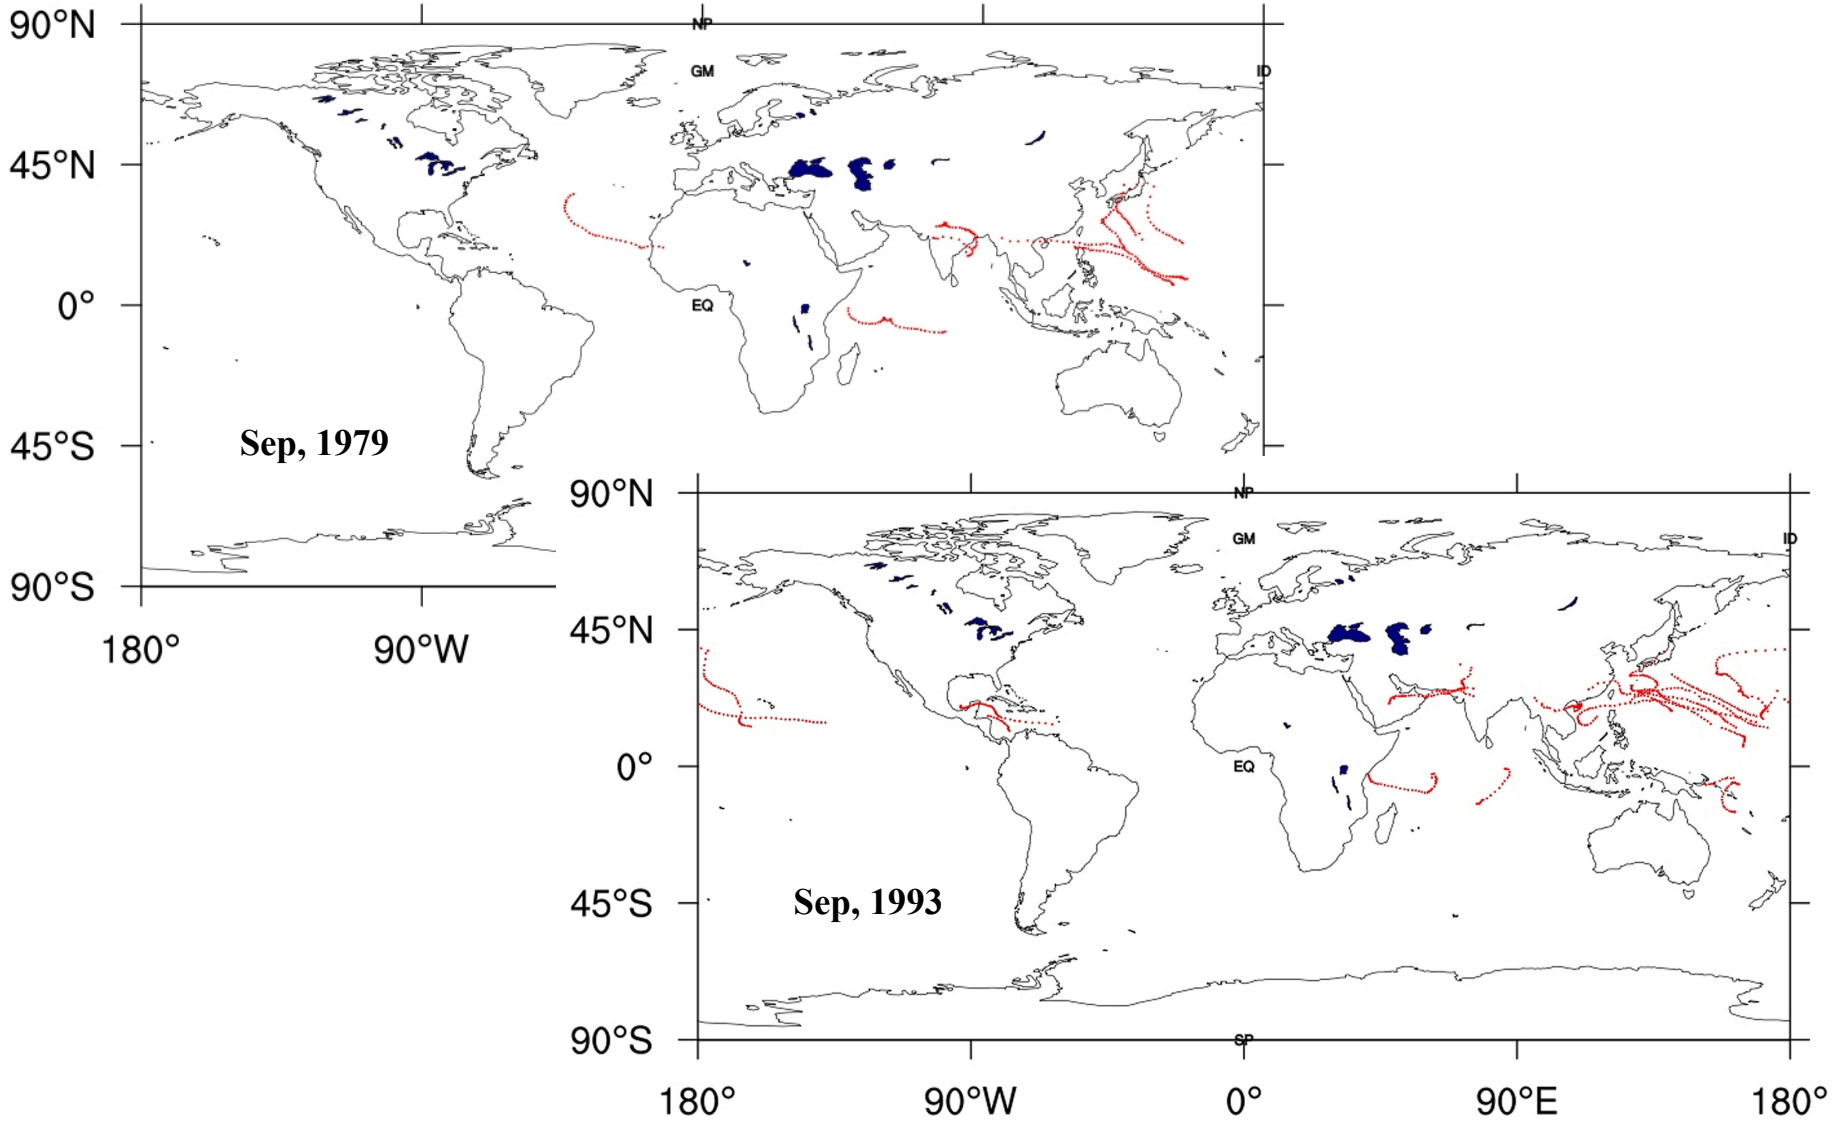

Analysis comparison

- Virtual machine on cloud computing
  - Eucalyptus VM submission
- Virtual machine on grid computing
  - Pre-loaded VMware image
- MPI parallel processing on cluster computing
  - Needed code re-write for MPI and local compilation

Analysis with virtual machines on cloud computing

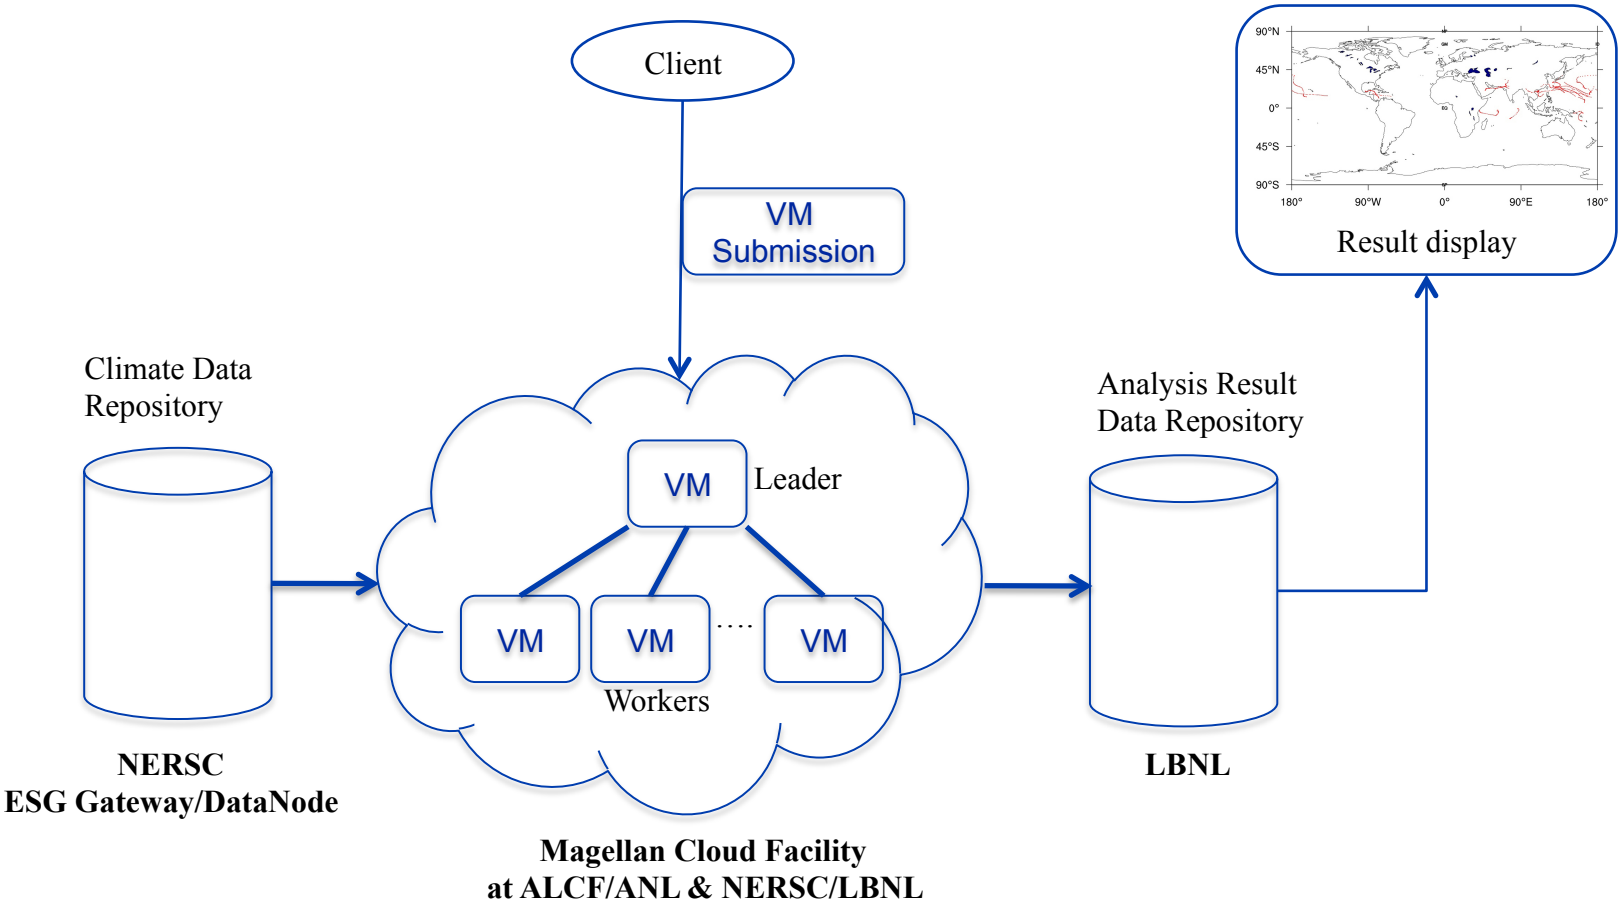

Analysis with Virtual Machines on Grid computing

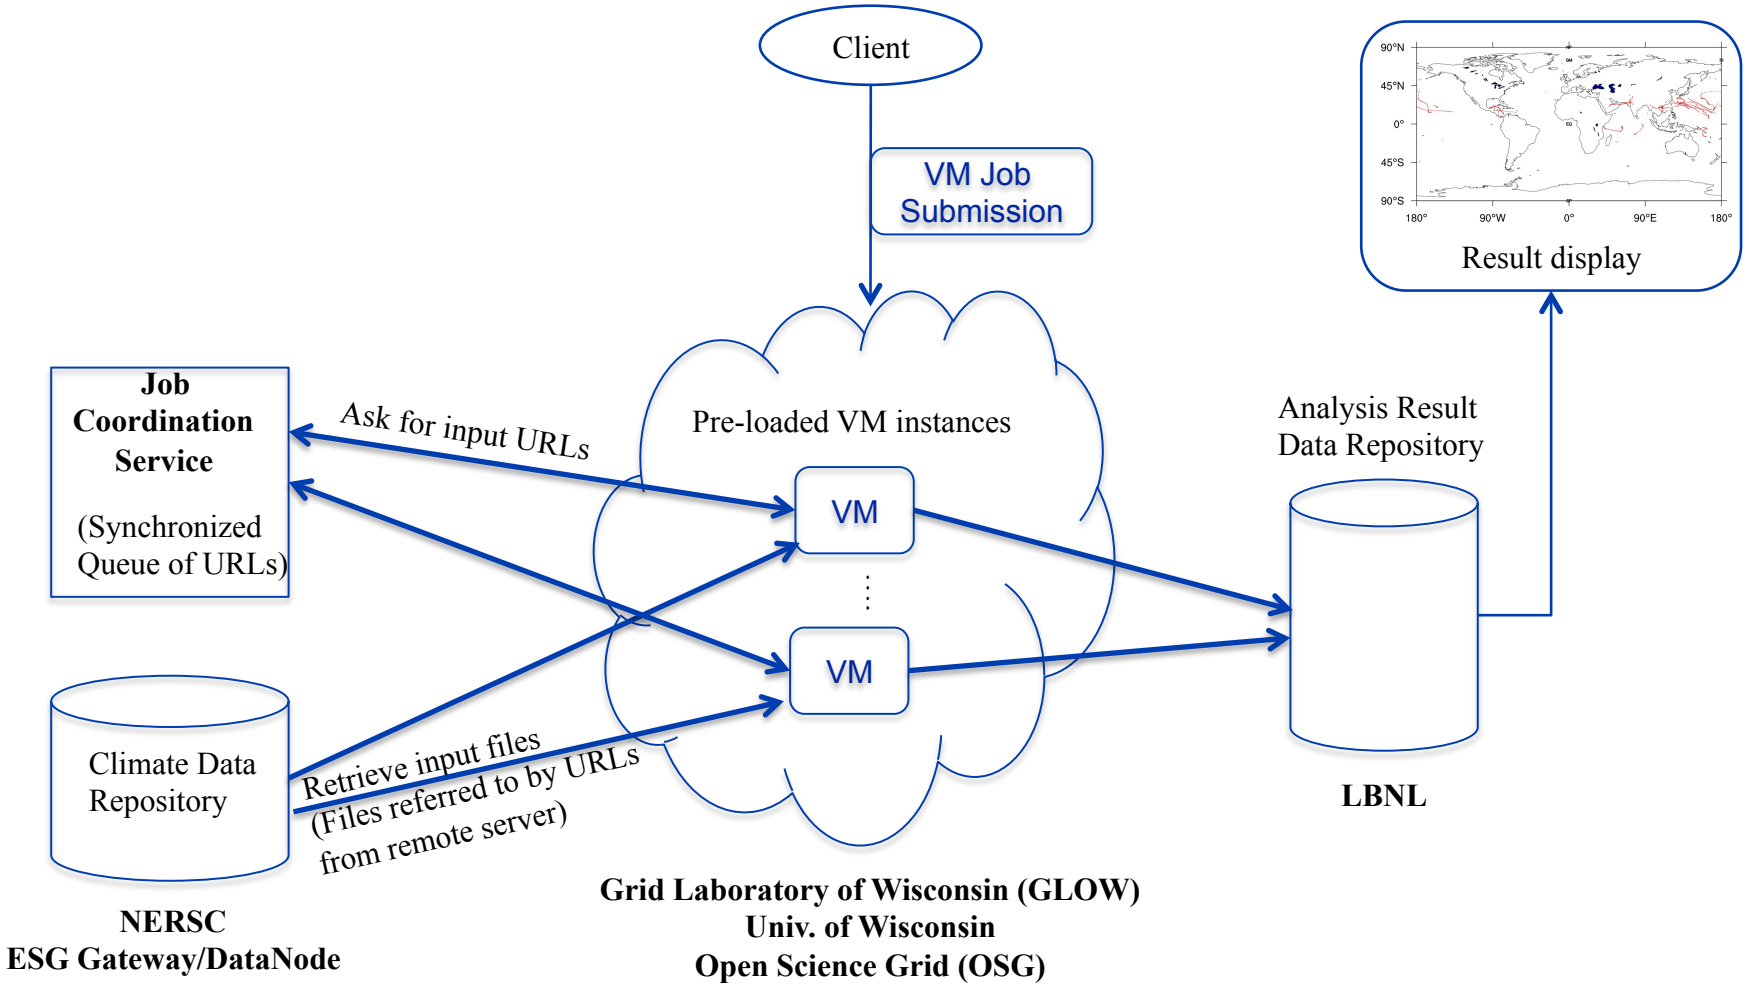

Analysis with MPI parallel processing on Clusters

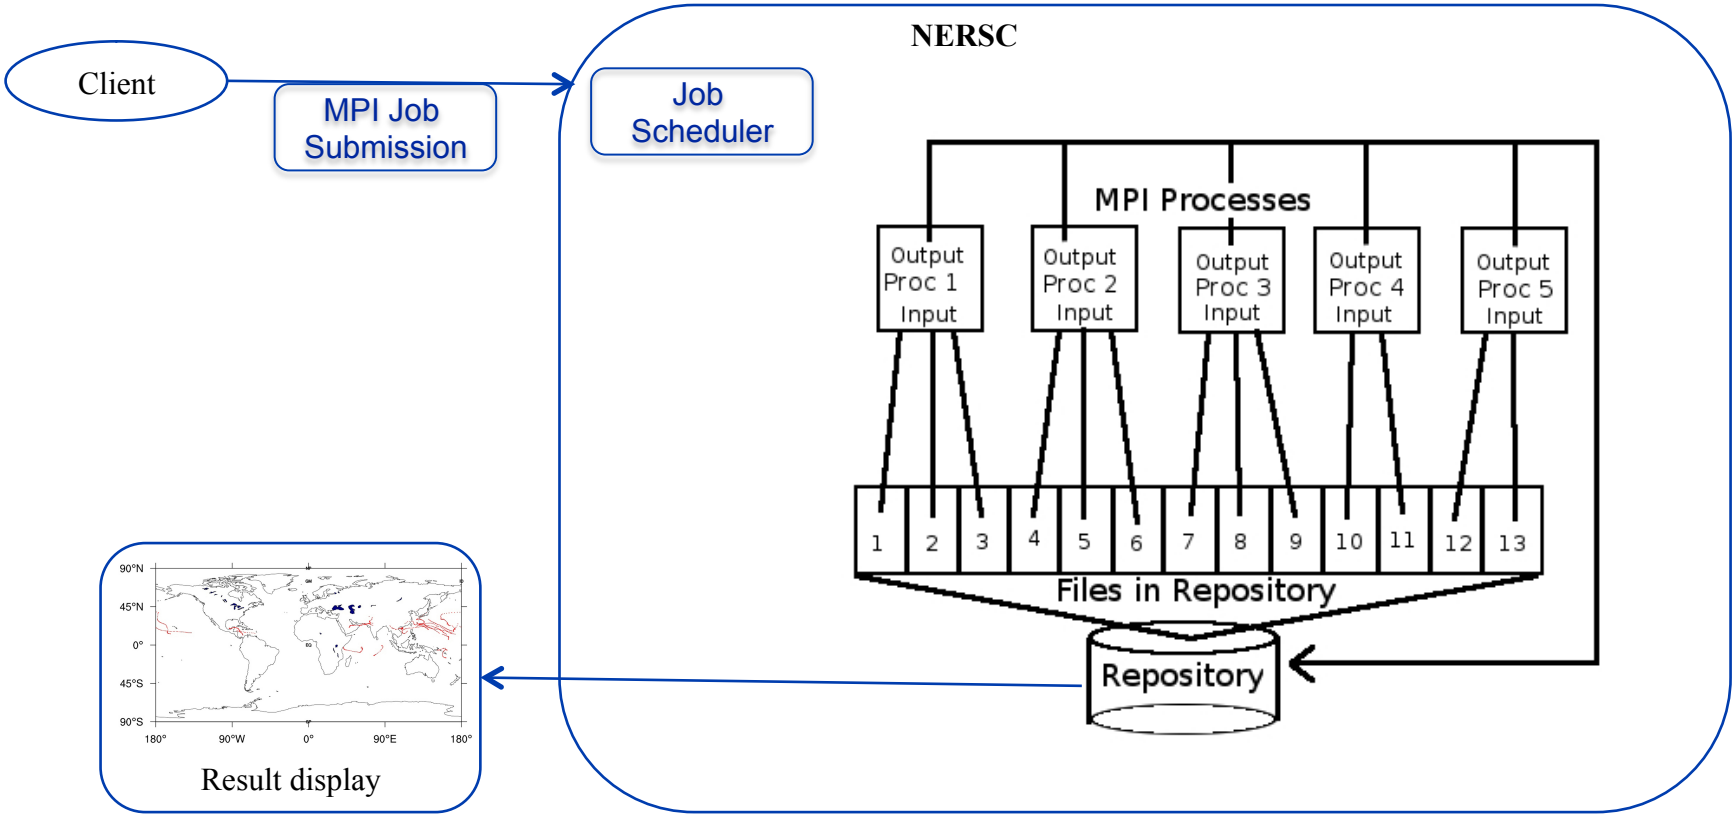

Analysis performance comparison

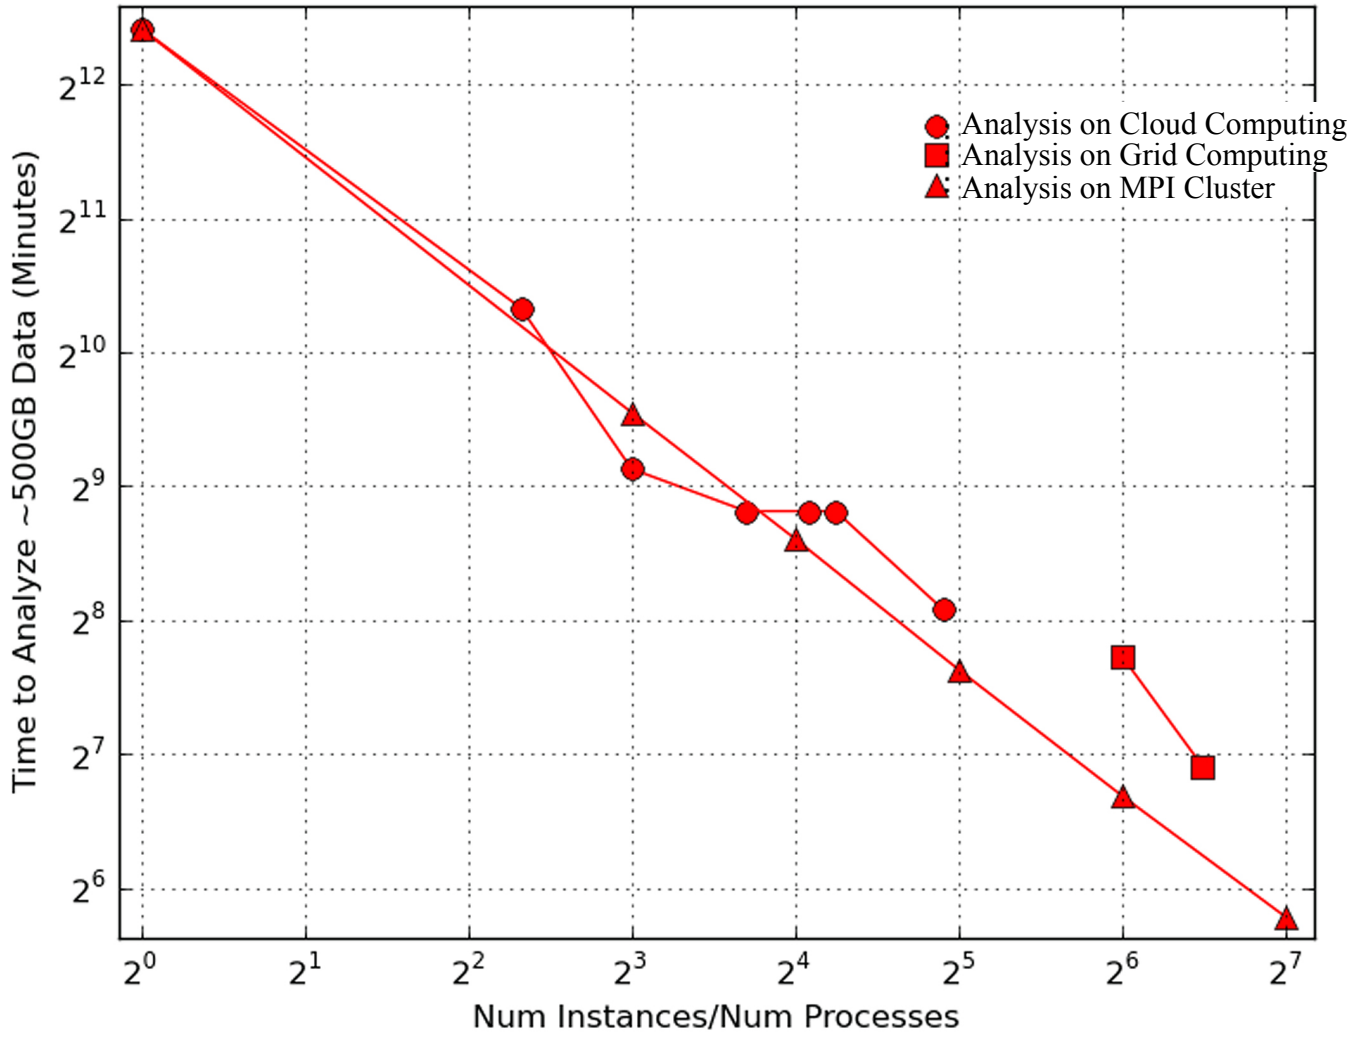

Lessons Learned

- Test analysis took 5-7 days on a workstation to 3 hours on 32 VMs on Cloud
- Cloud computing is just as efficient as MPI
  - MPI jobs are more predictable in performance
  - Variability on Cloud jobs is larger
    - Successful number of VM initialization varies
    - Network performance for remote data access
    - Storage capacity and performance

Envisioned Distributed Climate Analysis Framework

Analysis request

Distributed Extreme Climate Analysis (DECA)

Initialization & Planning

- Distributed resource verification
- Dataset Sampling & Dissemination planning
- Analysis planning
- Request planning

Distributed Analysis Execution & Management

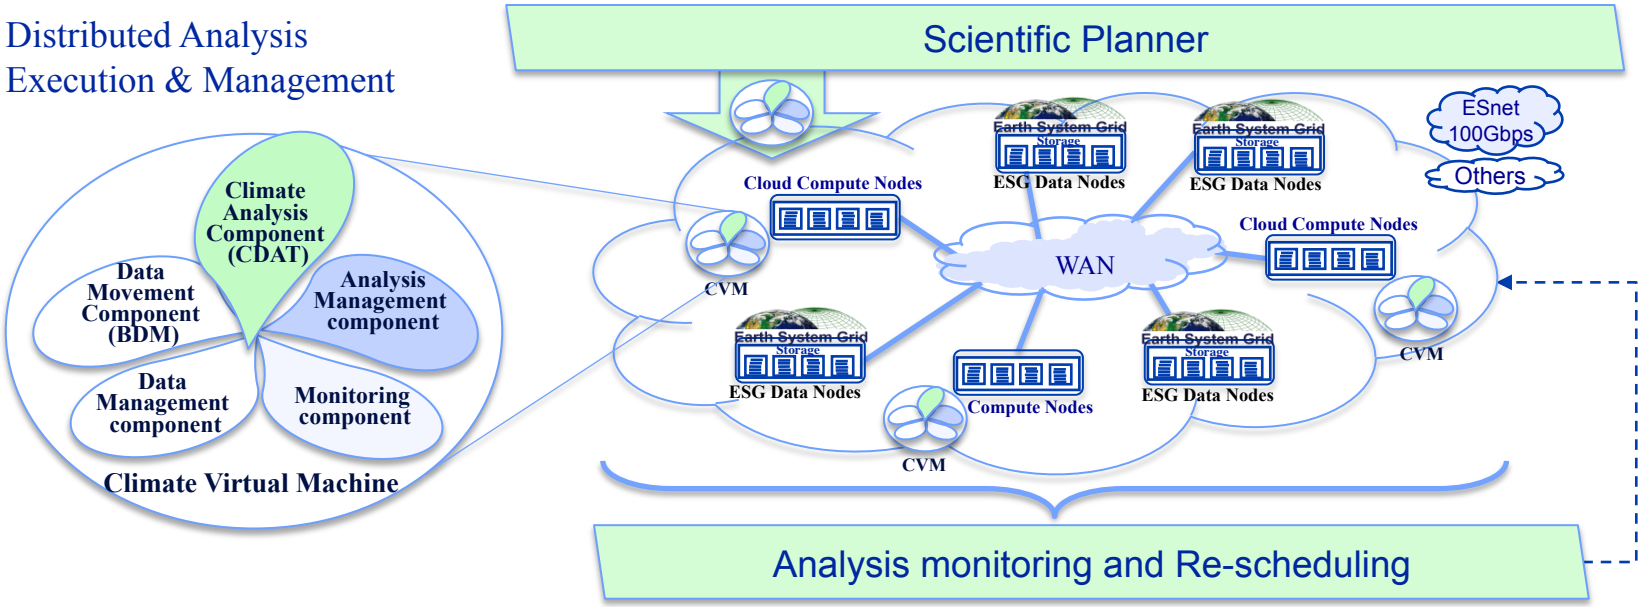

Post Analysis & Diagnostics

- Combination of results
- Ensemble
- Comparative Study
- Validation & Diagnostics
- Visibility & storage of analysis data

Future of ESG and Cloud

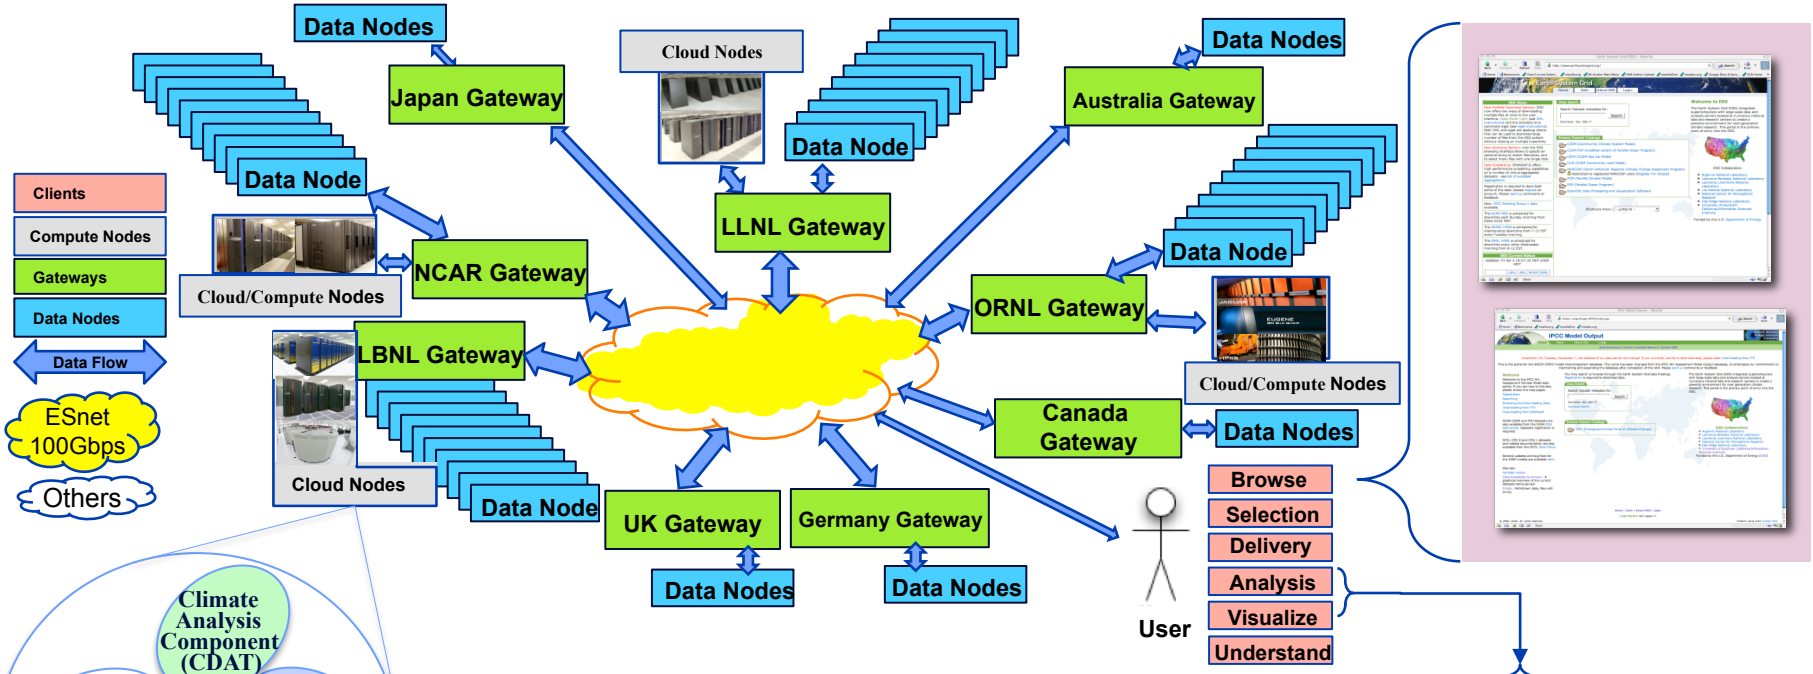

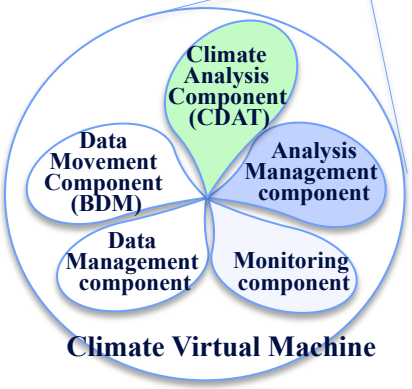

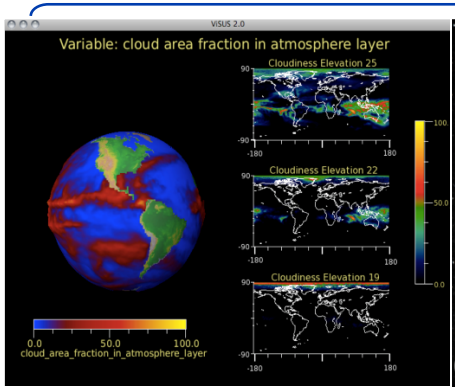

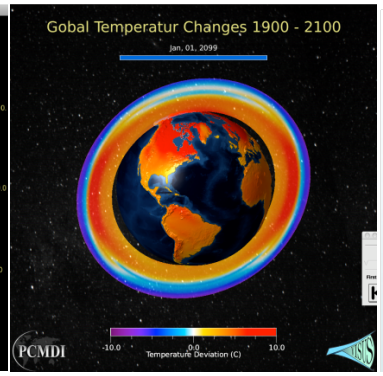

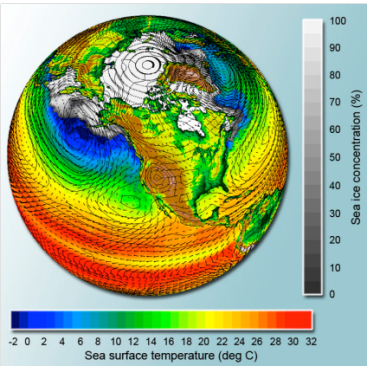

- **Enable new capabilities for analysis of data and visual exploration**
  - Visualization of uncertainty and ensemble data
- **Help scientists understand long-term climate impact**