Recent Advancements in High Resolution Climate Modeling

14th Workshop on the use of HPC in Meteorology
1 November 2010

Dr. Richard Loft
loft@ucar.edu
Technology Development Division
Computational and Information Systems Laboratory
National Center for Atmospheric Research

CISL Computational & Information Systems Laboratory

NCAR

Why High Resolution? Resolving Ocean Mesoscale Eddies

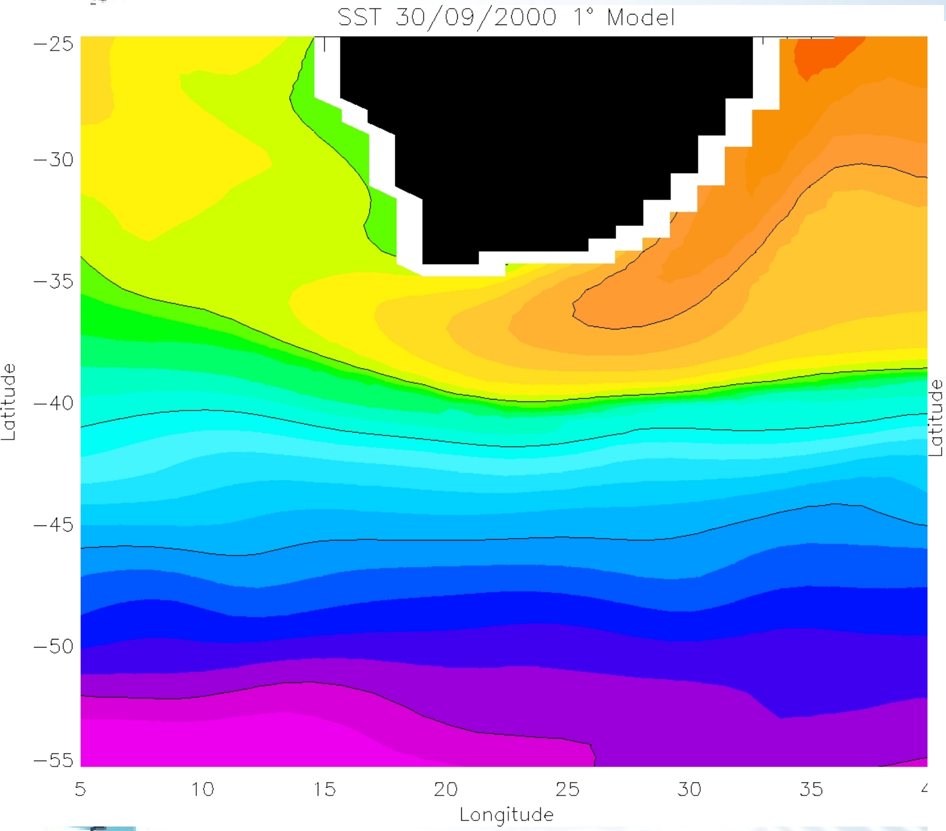

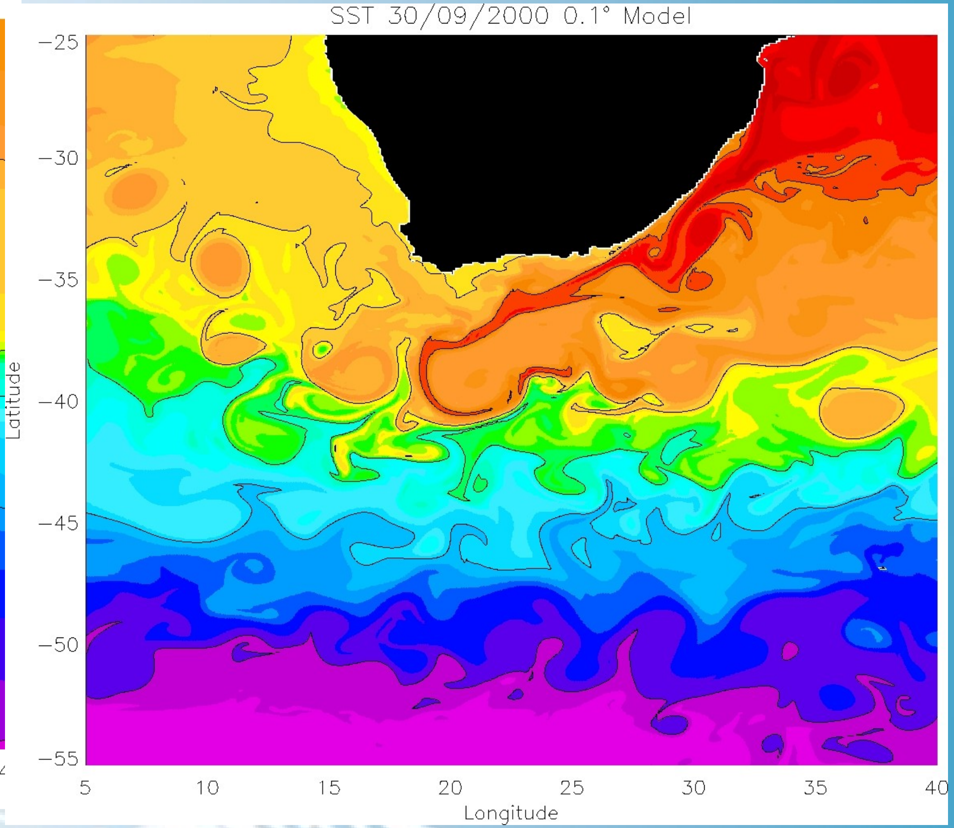

1° Ocean component of CCSM (Collins et al, 2006)

0.1° Eddy Resolving (Maltrud & McClean, 2005)

NCAR

Ocean-Atmosphere Interactions: North Atlantic Winter Storm Track

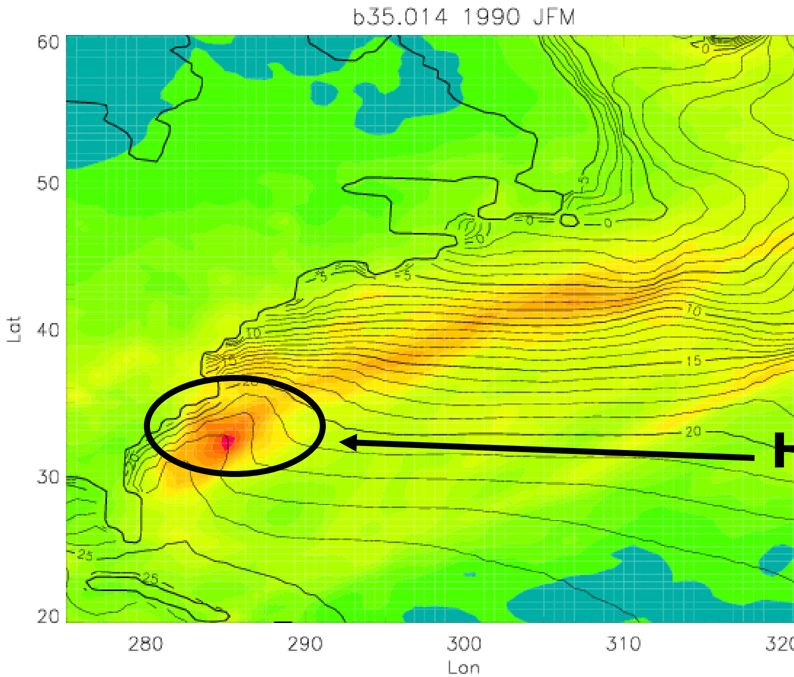

0.5° atm + 1° ocn

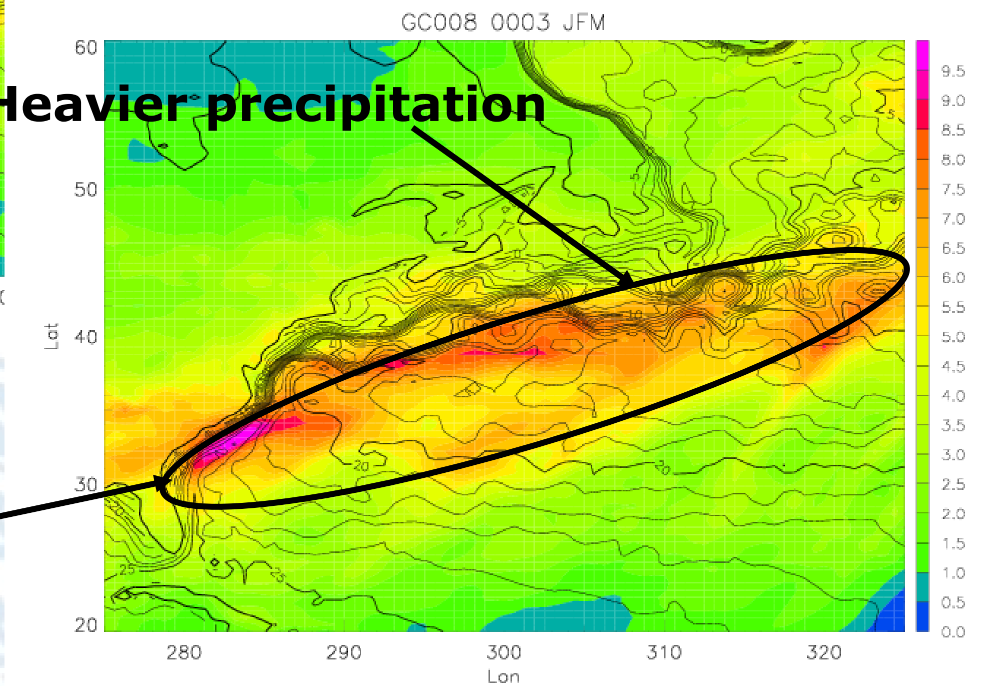

0.5° atm + 0.1° ocn

Heavier precipitation

Stronger SST gradient

NCAR

Engineering Petascale Software

a different result

CESM: New Modeling Capabilities for Ultra-High Resolution Simulations

Computational & Information Systems Laboratory

Including…

- **Flexible coupling infrastructure**
- **Memory scalability of all components**
  - *Minimize global arrays*
- **Performance scalability of all components**
  - *Hybrid MPI and OpenMP for multicore architecture*
  - *ALL active components - CAM, CLM, CICE and POP2 - now meet this requirement*
- **Parallel I/O throughout system**
- **Scalable Dynamical Core Option (HOMME)**

NCAR

CESM1 “Hub and Spoke” Coupling Architecture

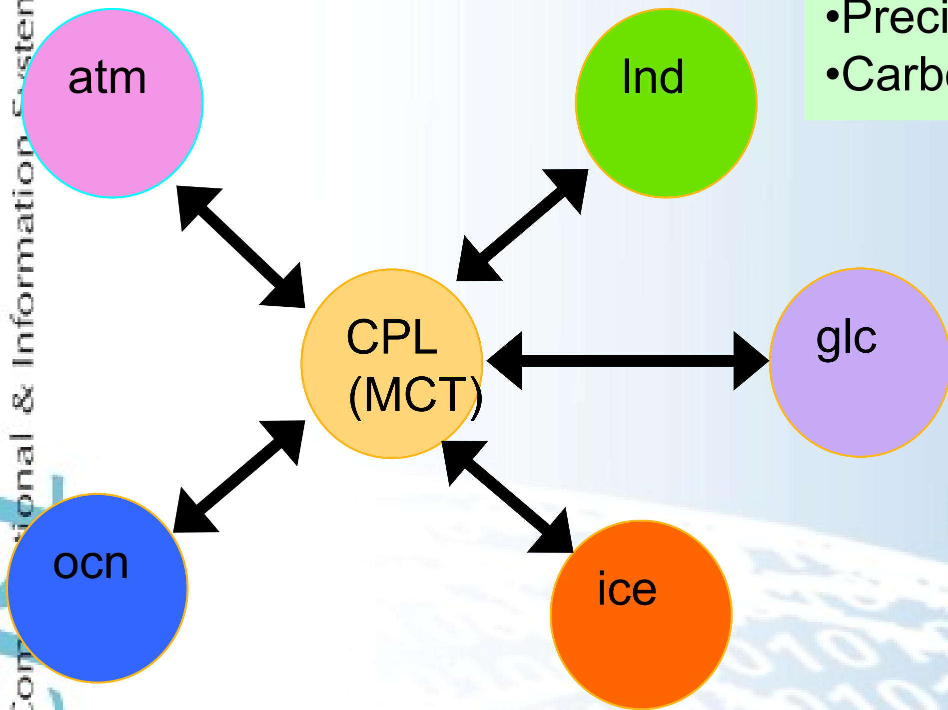

Atm -> Coupler

- Bottom level temperature, pressure, wind…
- Downward Shortwave (vis, nir)
- Precipitation
- Carbon and Dust fluxes

Coupler -> Atm

(merged from lnd, ice and ocn)

- Latent, sensible heat fluxes
- Surfaces Stresses
- Upward long wave
- Evaporative water flux
- Surface Albedos

Note: Glacier (glc) component is new with CESM1 and not benchmarked here.

CESM: Coupler-7 Architecture

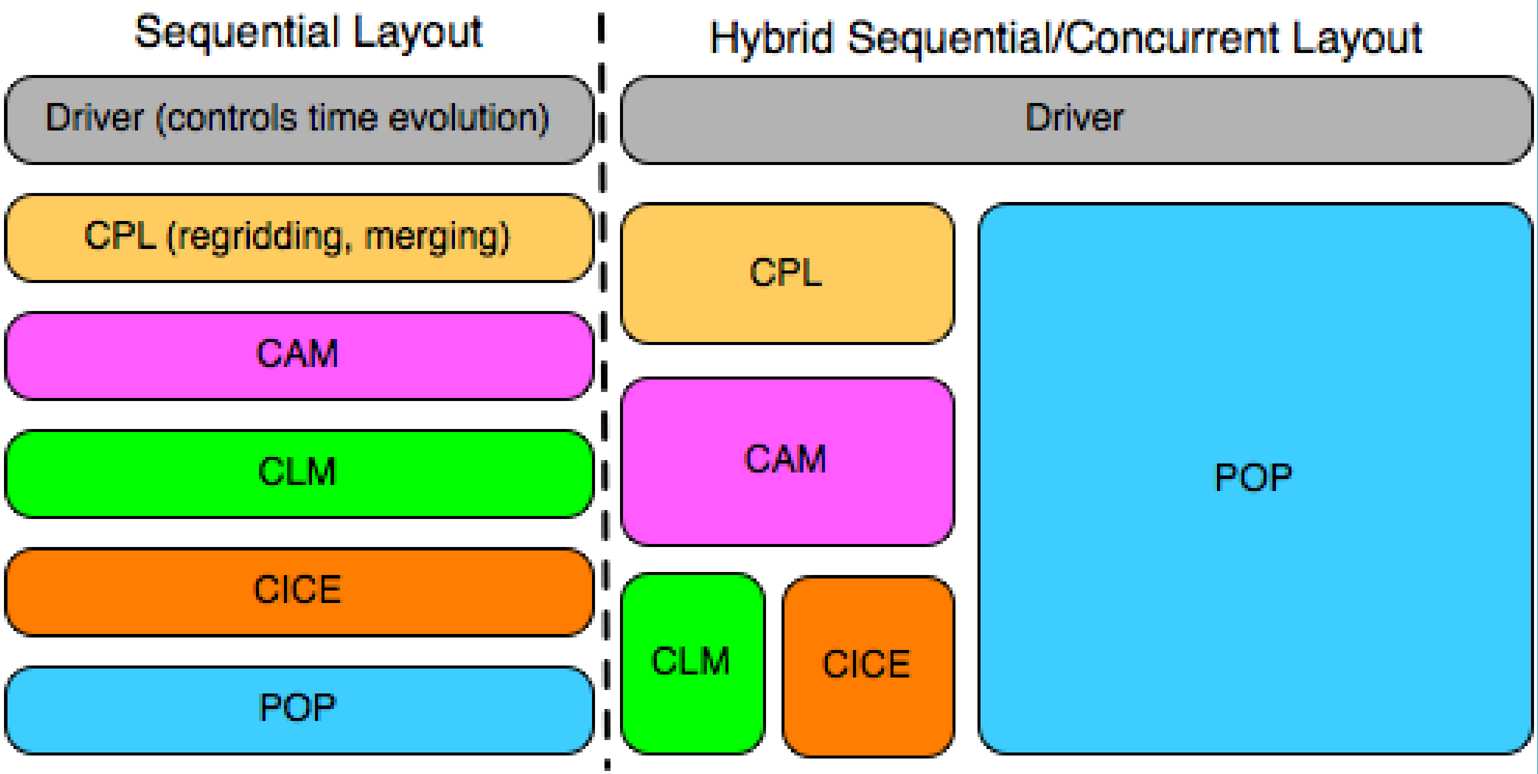

Climate Model Nomenclature/Details

- **CCSM renamed CESM1 in June 2010**
- **Configuration Nomenclature- N° x M°**
  - ***N°*** *Atmosphere/Land models*
  - ***M°*** *Ocean/Sea Ice models*
- **Example: component grids – (0.5° x 0.1°)**
  - *0.50° ATM [576 x 384 x 26] - CAM*
  - *0.50° LND [576 x 384 x 17] - CLM*
  - *0.1° OCN [3600 x 2400 x 42] – POP2*
  - *0.1° ICE [3600 x 2400 x 20] – CICE*
- **SYPD – simulated years per day**
  - *Throughput measure of performance*
  - *You won't see flops mentioned here*

CESM: Parallel I/O (PIO) Library

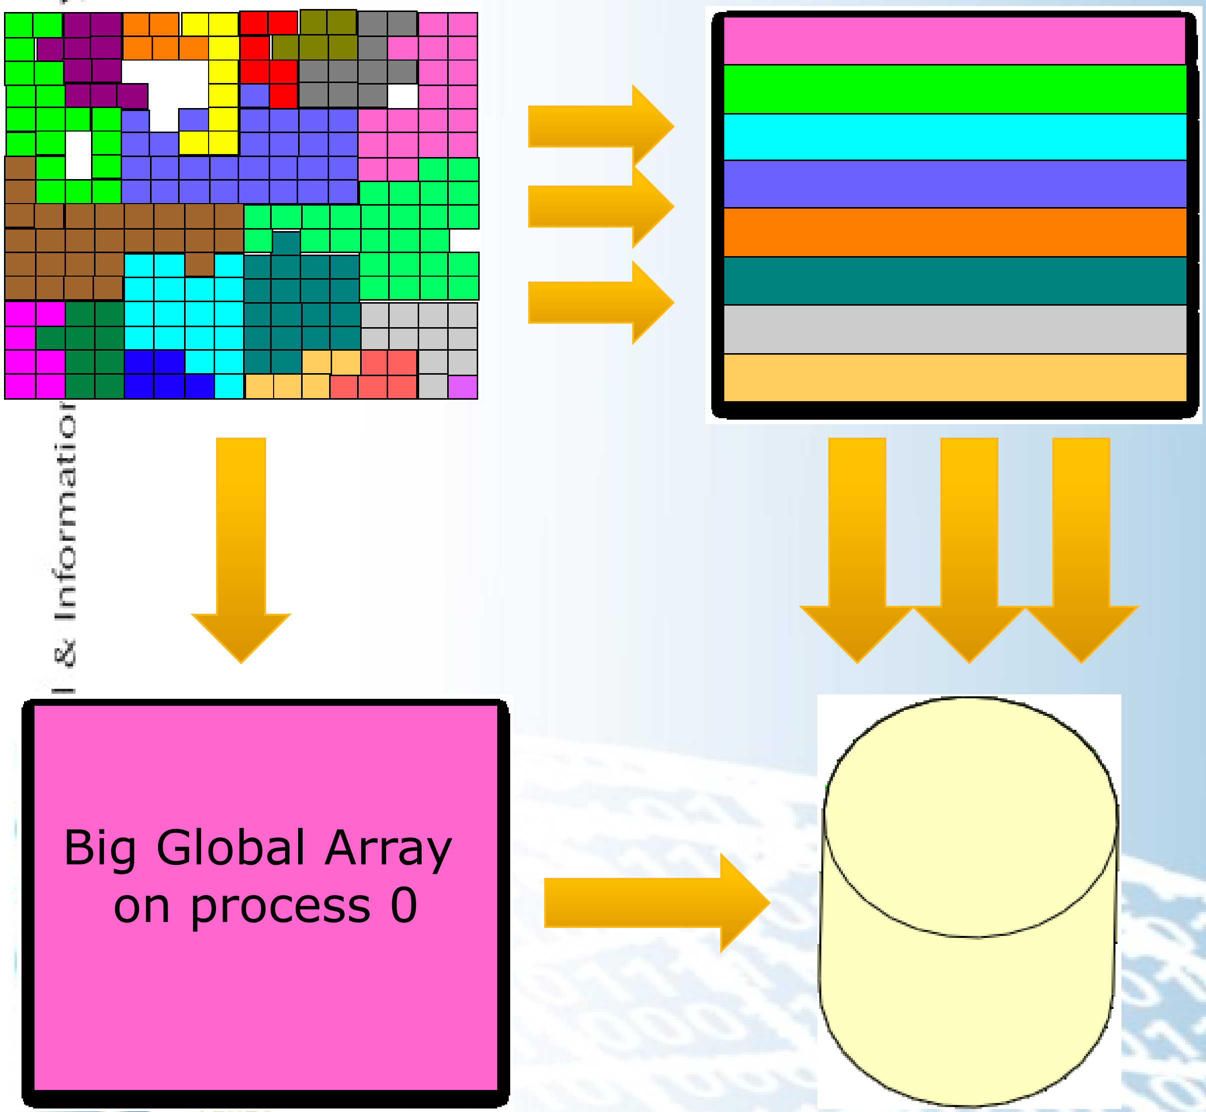

Rearranges data from model decomp to I/O friendly decomp

Interface between the model and the I/O library. Supports

- Binary
- NetCDF3 (serial netcdf)
- Parallel NetCDF (pnetcdf) (MPI/IO)
- NetCDF4

How do we Load Balance Multi-component models?

Optimize throughtput and decrease idle cycles

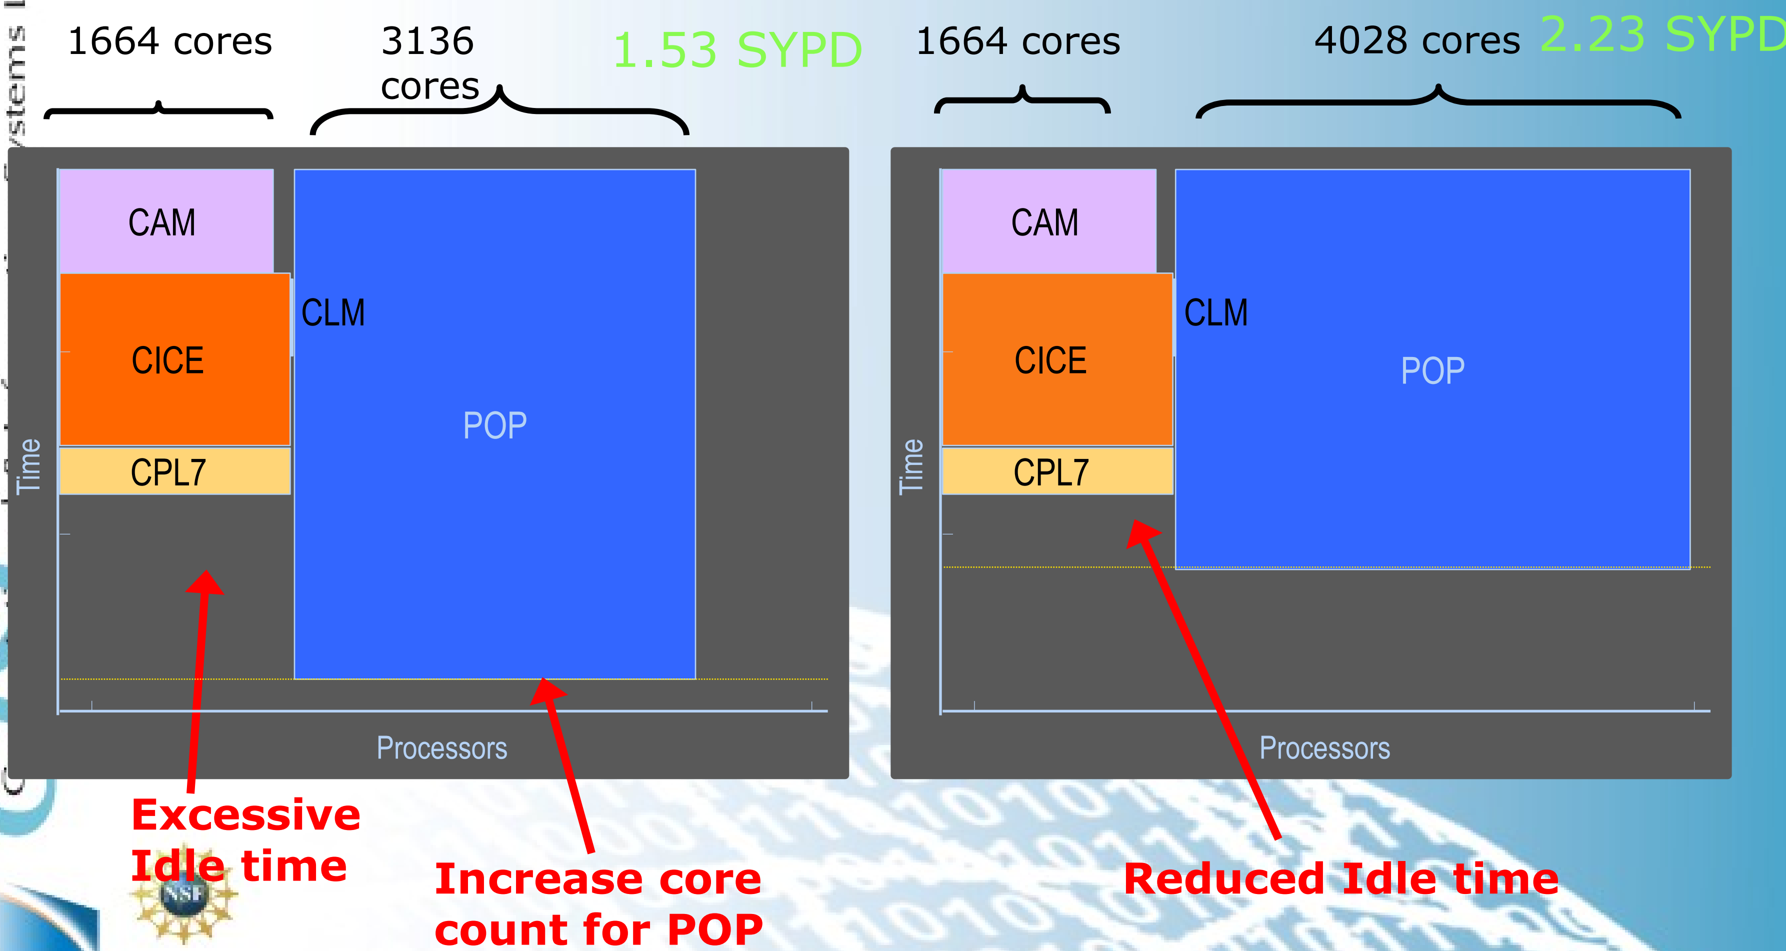

Courtesy of John Dennis

CESM1 0.5° x 0.1° Scalability on Cray XT systems

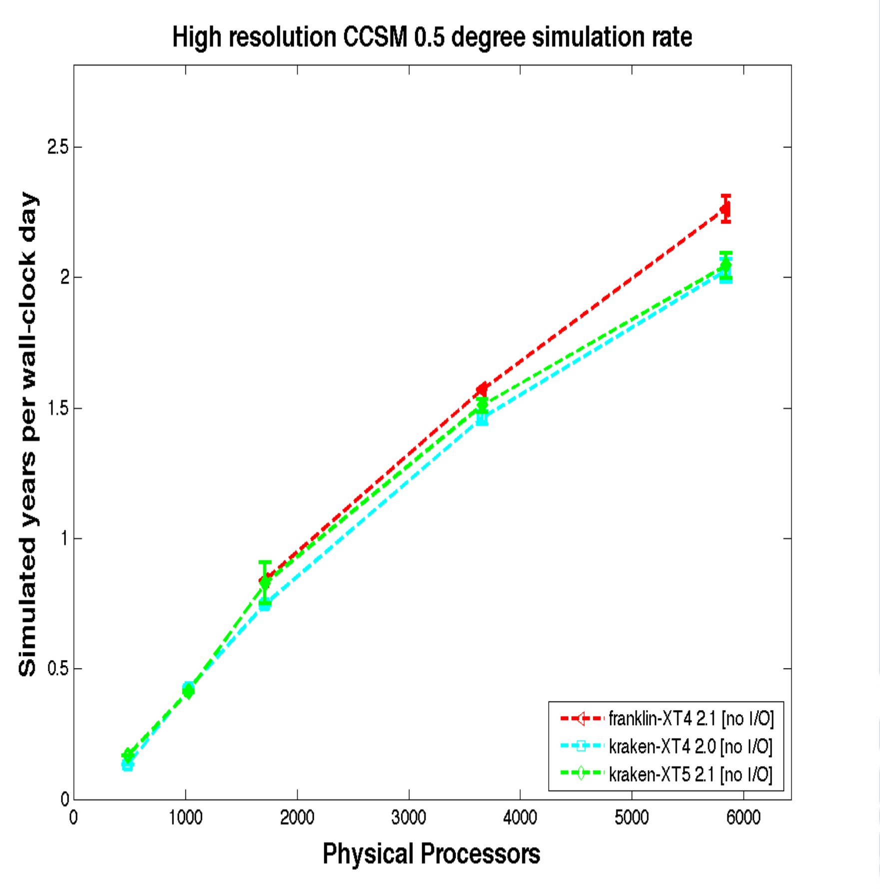

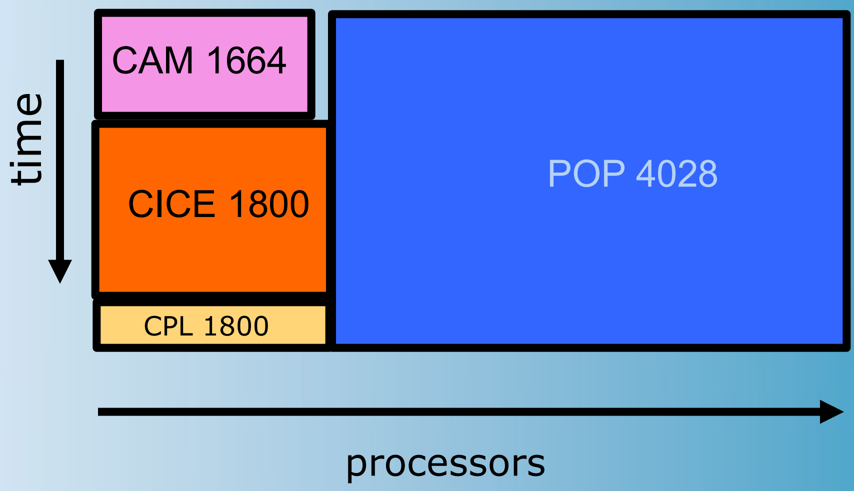

1.9 years/day on 5844 cores with I/O on kraken hex-core XT5 (no threading)

(Courtesy of John Dennis)

CESM1 0.5° x 0.1°
Tuning Computational Efficiency

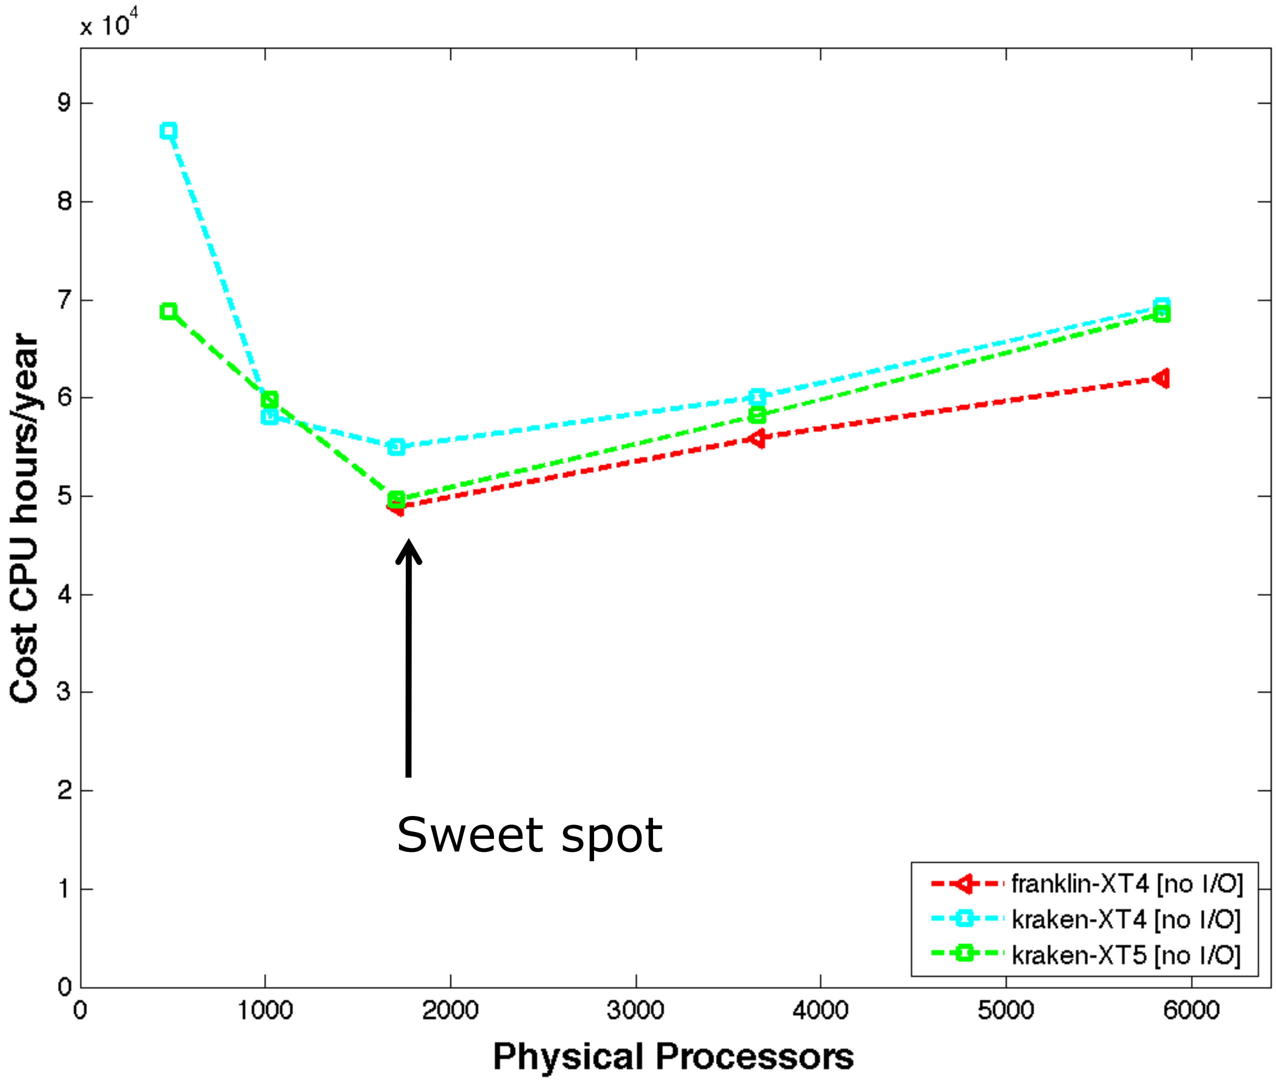

CESM1 OpenMP/MPI Cray XT5 Hex-core Scalability 0.25° x 0.1° case

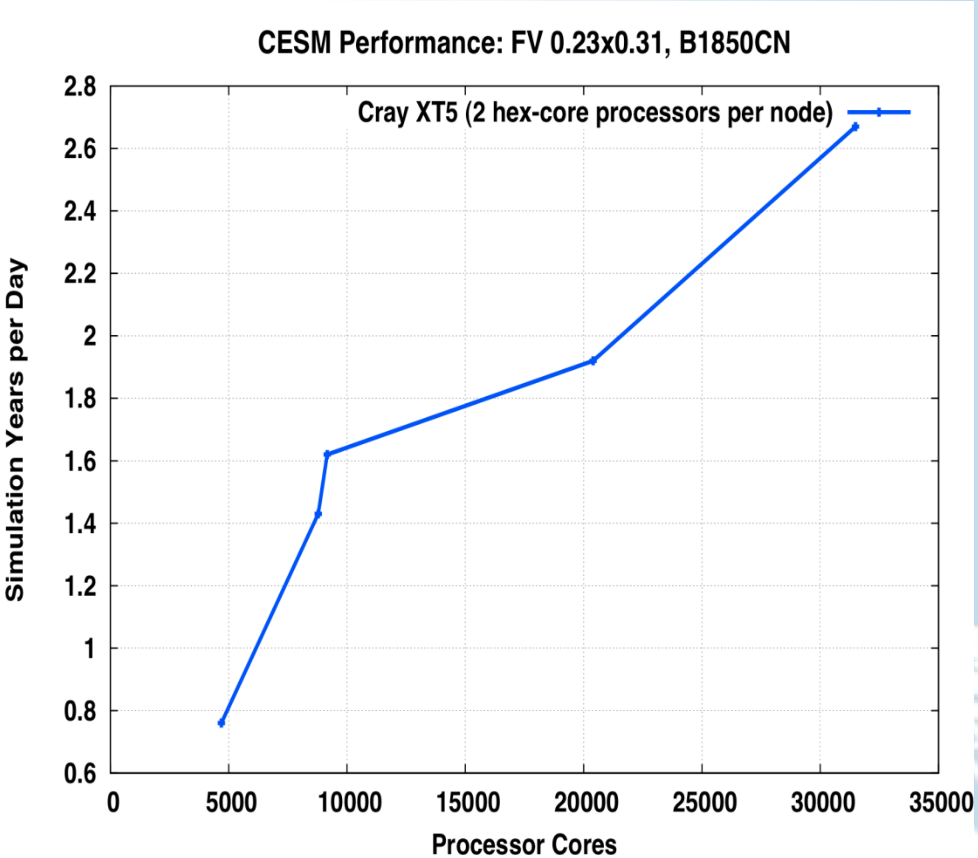

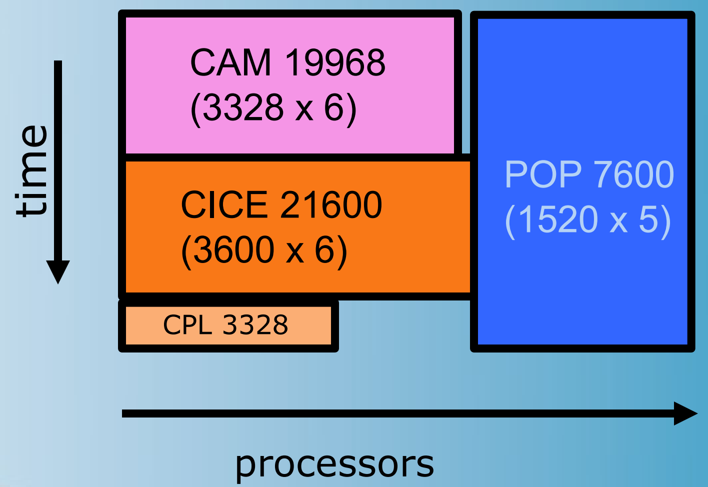

2.6 sypd on 30K cores with I/O (x N threads)

(Courtesy of Pat Worley)

HRC06 Production Run
0.5° x 0.1° Details

- **155 year control run**
- **~18M CPU hours**
- **5844 cores for 4-5 months**
- **~100 TB of data generated**
- **0.5 to 1 TB per wall clock day generated**

Computational & Information Systems Laboratory

NCAR

CESM1 0.1° ocean/sea-ice visualization

Note: atmospheric physics has not been retuned to its higher 0.5° resolution

This mov file was removed from pdf version to reduce size.

CESM1 Scalable HOMME Dycore: High-Order Methods Modeling Environment

- **Algorithmic Advantages of High Order Methods**
  - *h-p element-based method on quadrilaterals (N_e x N_e)*
  - *Exponential convergence in polynomial degree (N)*

- **Computational Advantages of High Order Methods**
  - *Naturally cache-blocked N x N computations*
  - *Nearest-neighbor communication between elements (explicit)*
  - *Well suited to parallel µprocessor systems*

HOMME: Quasi-uniform Cube-Sphere Grid

- Sphere is decomposed into 6 identical regions using a central projection (Sadourny, 1972) with equiangular grid (Rancic et al., 1996).
- Avoids pole problems, quasi-uniform.
- Non-orthogonal curvilinear coordinate system with identical

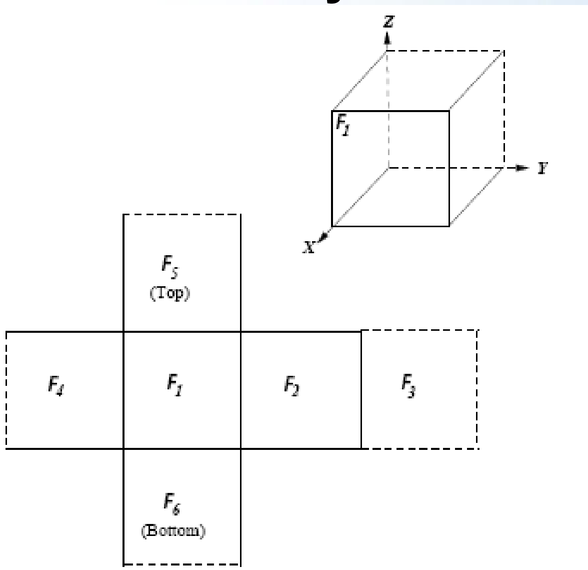

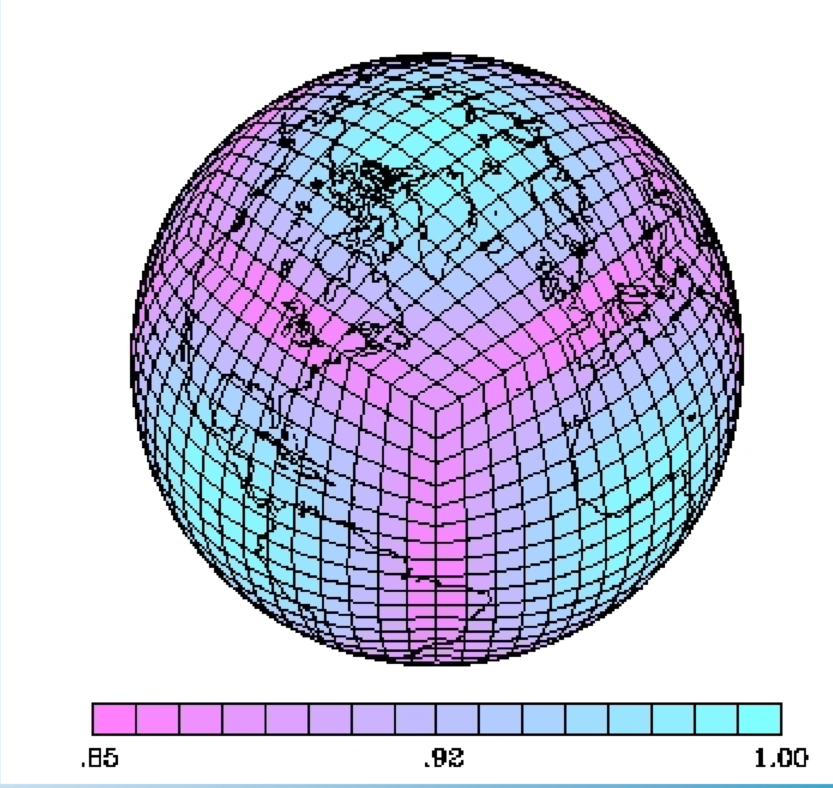

Ne=16 Cube Sphere Showing degree of non-uniformity

CCSM/HOMME Scalability
0.125° atm / 0.25° land / 0.1° ocean

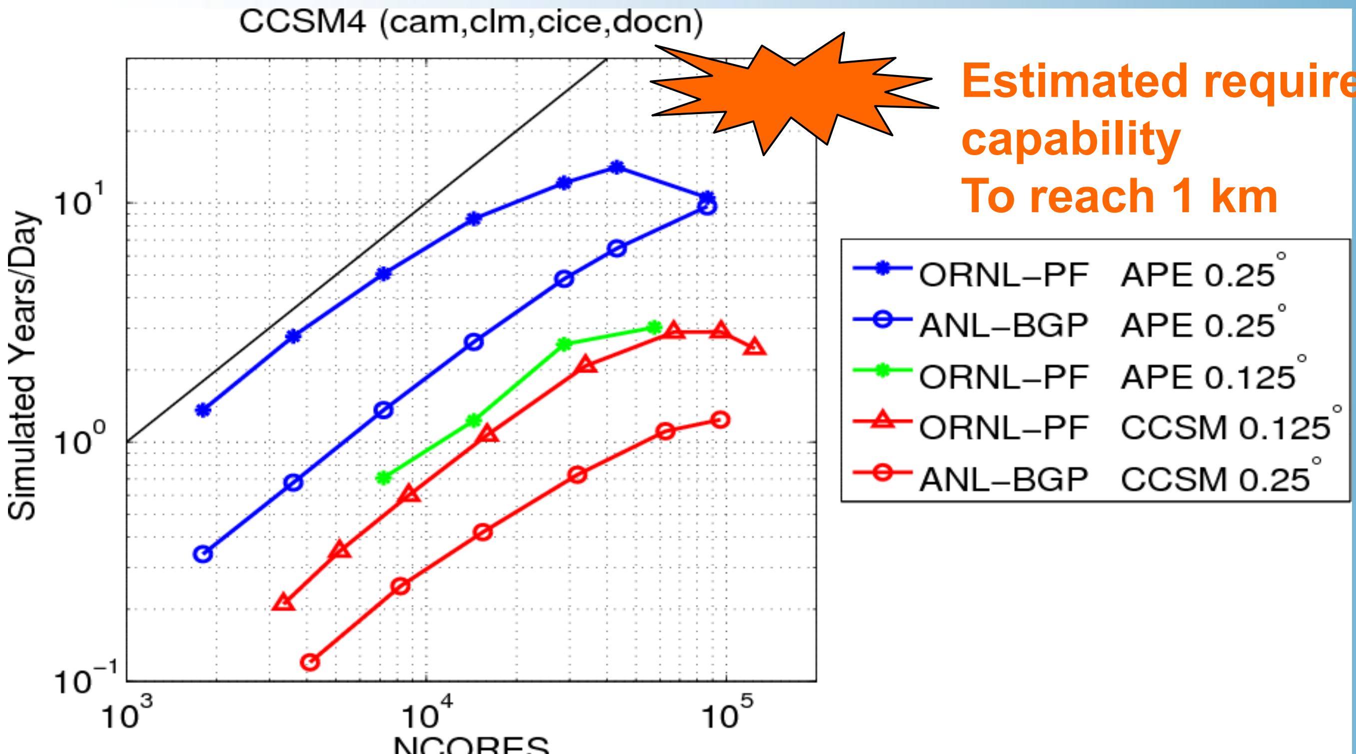

Work of Mark Taylor, Jim Edwards and Brian Eaton

- **CCSM times include ALL CCSM components (PIO use was critical)**
- **Scalability of the dynamical core is preserved by CAM *and* scalability of CAM is preserved by CCSM**
- **Scale out to 86000 cores (BGP) and get 3 SYPD (Jaguarpf)**

Next Generation NCAR/CISL Infrastructure

(Planning for 2012)

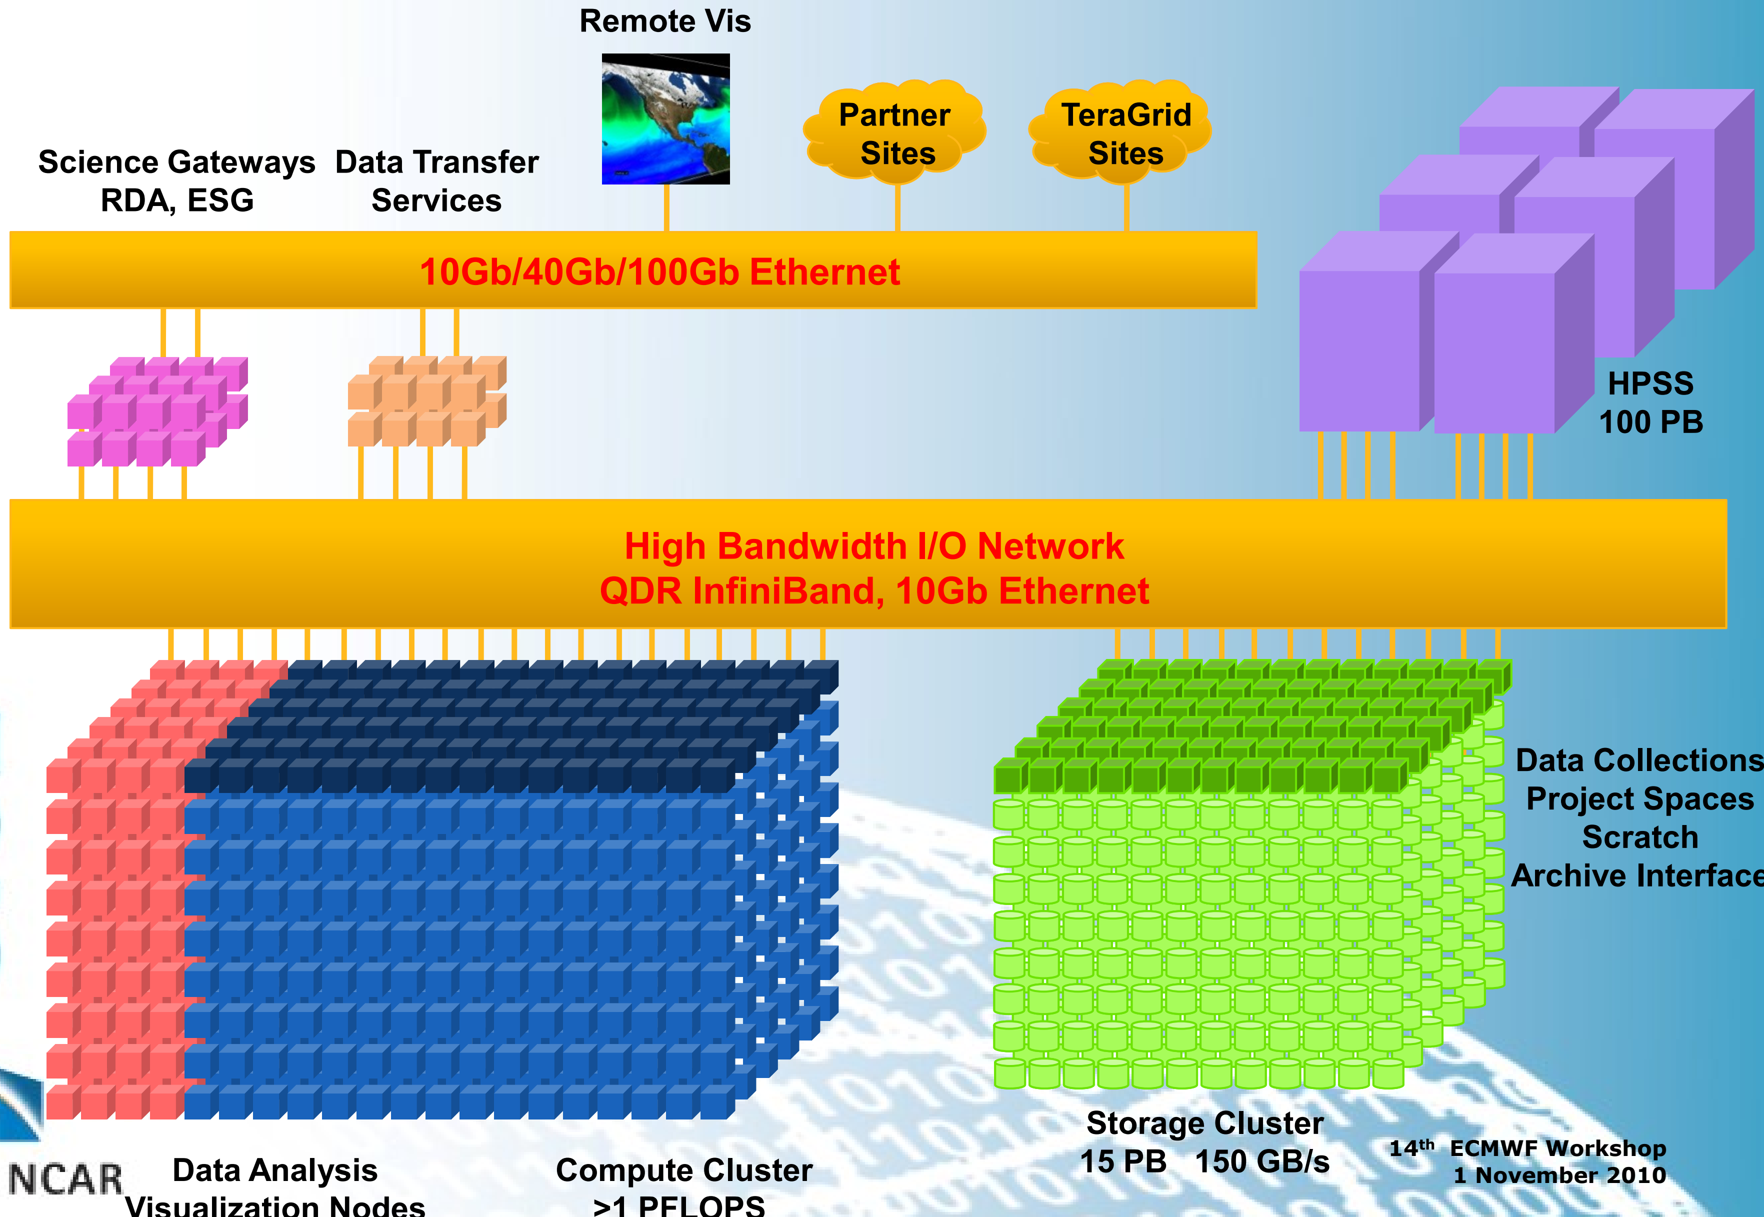

Computational & Information Systems Laboratory

CISL

NCAR

New Systems and Facilities

The NCAR Wyoming Supercomputing Center:

A Petascale Facility
Dedicated to the Atmospheric Sciences

boratory

NWSC: Timeline

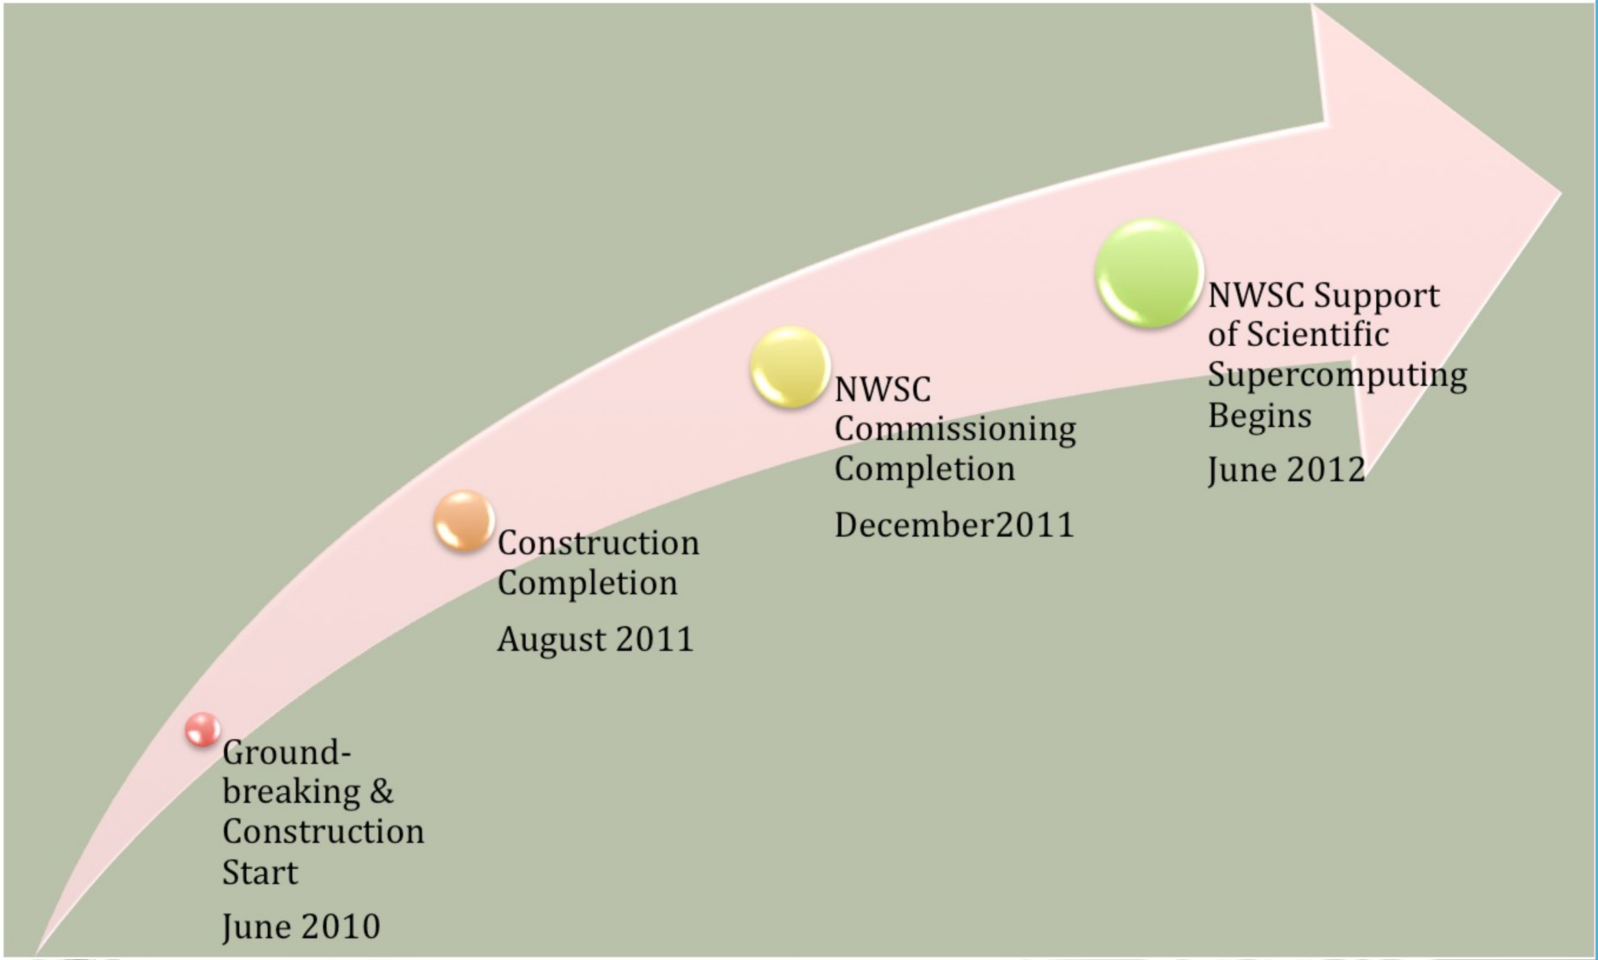

Computational & Information Systems Laboratory

CISL

NCAR

NWSC Fact Sheet

- **Project cost: ~$70M**
- **HPC system install begins January 2012**
- **Two, 1,115 m^2 raised floor areas**
  - *Power density – 6727 W/m^2*
  - ***Rack power day 1 – 4.5 MW***
- **221-595 m^2 archive space**
- **Total floor area: 15,885 m^2**
  - *Main floor: 14,242 m^2*
  - *Upper floor: 1,643 m^2*
- **Facility/site expandable to 24 MW**
  - *With additional funding of course!*
- **Distance from NCAR: 170 km**

NWSC: Power Efficiency

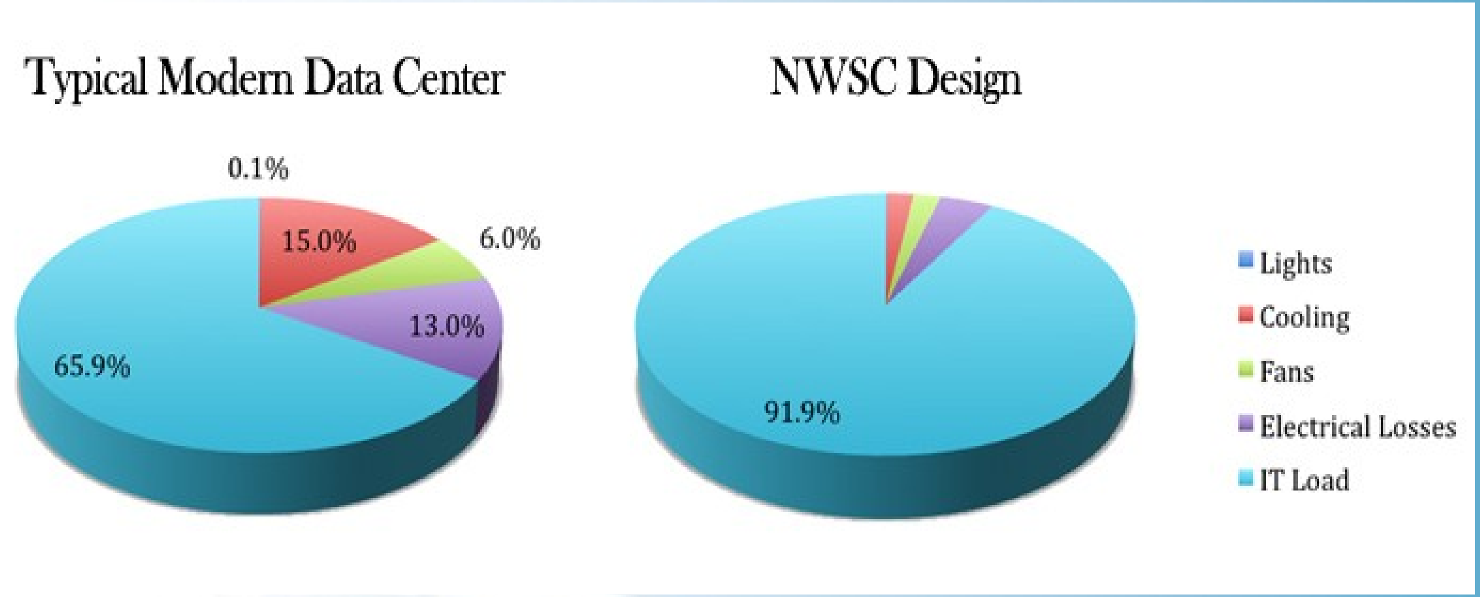

NWSC PUE target: 1.10

NWSC Facility Progress

Acknowledgements

- **NCAR:**
  - *D. Bailey*
  - *F. Bryan*
  - *T. Craig*
  - *B. Eaton*
  - *J. Edwards [IBM]*
  - *N. Hearn*
  - *K. Lindsay*
  - *N. Norton*
  - *M. Vertenstein*
- **COLA:**
  - *J. Kinter*
  - *C. Stan*
- **U. Miami**
  - *B. Kirtman*
- **U.C. Berkeley**
  - *W. Collins*
  - *K. Yelick (NERSC)*
- **U. Washington**
  - *C. Bitz*

- NICS:
  - M. Fahey
  - P. Kovatch
- ANL:
  - R. Jacob
  - R. Loy
- LANL:
  - E. Hunke
  - P. Jones
  - M. Maltrud
- LLNL
  - D. Bader
  - D. Ivanova
  - J. McClean (Scripps)
  - A. Mirin
- ORNL:
  - P. Worley

and many more…

- **Grant Support:**
  - *DOE*
    - DE-FC03-97ER62402 [SciDAC]
    - DE-PS02-07ER07-06 [SciDAC]
  - *NSF*
    - Cooperative Grant NSF01
    - OCI-0749206 [PetaApps]
    - CNS-0421498
    - CNS-0420873
    - CNS-0420985
- **Computer Allocations:**
  - *TeraGrid TRAC @ NICS*
  - *DOE INCITE @ NERSC*
  - *LLNL Grand Challenge*
- **Thanks for Assistance:**
  - *Cray, NICS, and NERSC*

CISL Computational & Information Systems Laboratory

NCAR


Thanks!
Any Questions?

NCAR/CISL HPC-Data Infrastructure

(Summer 2010)

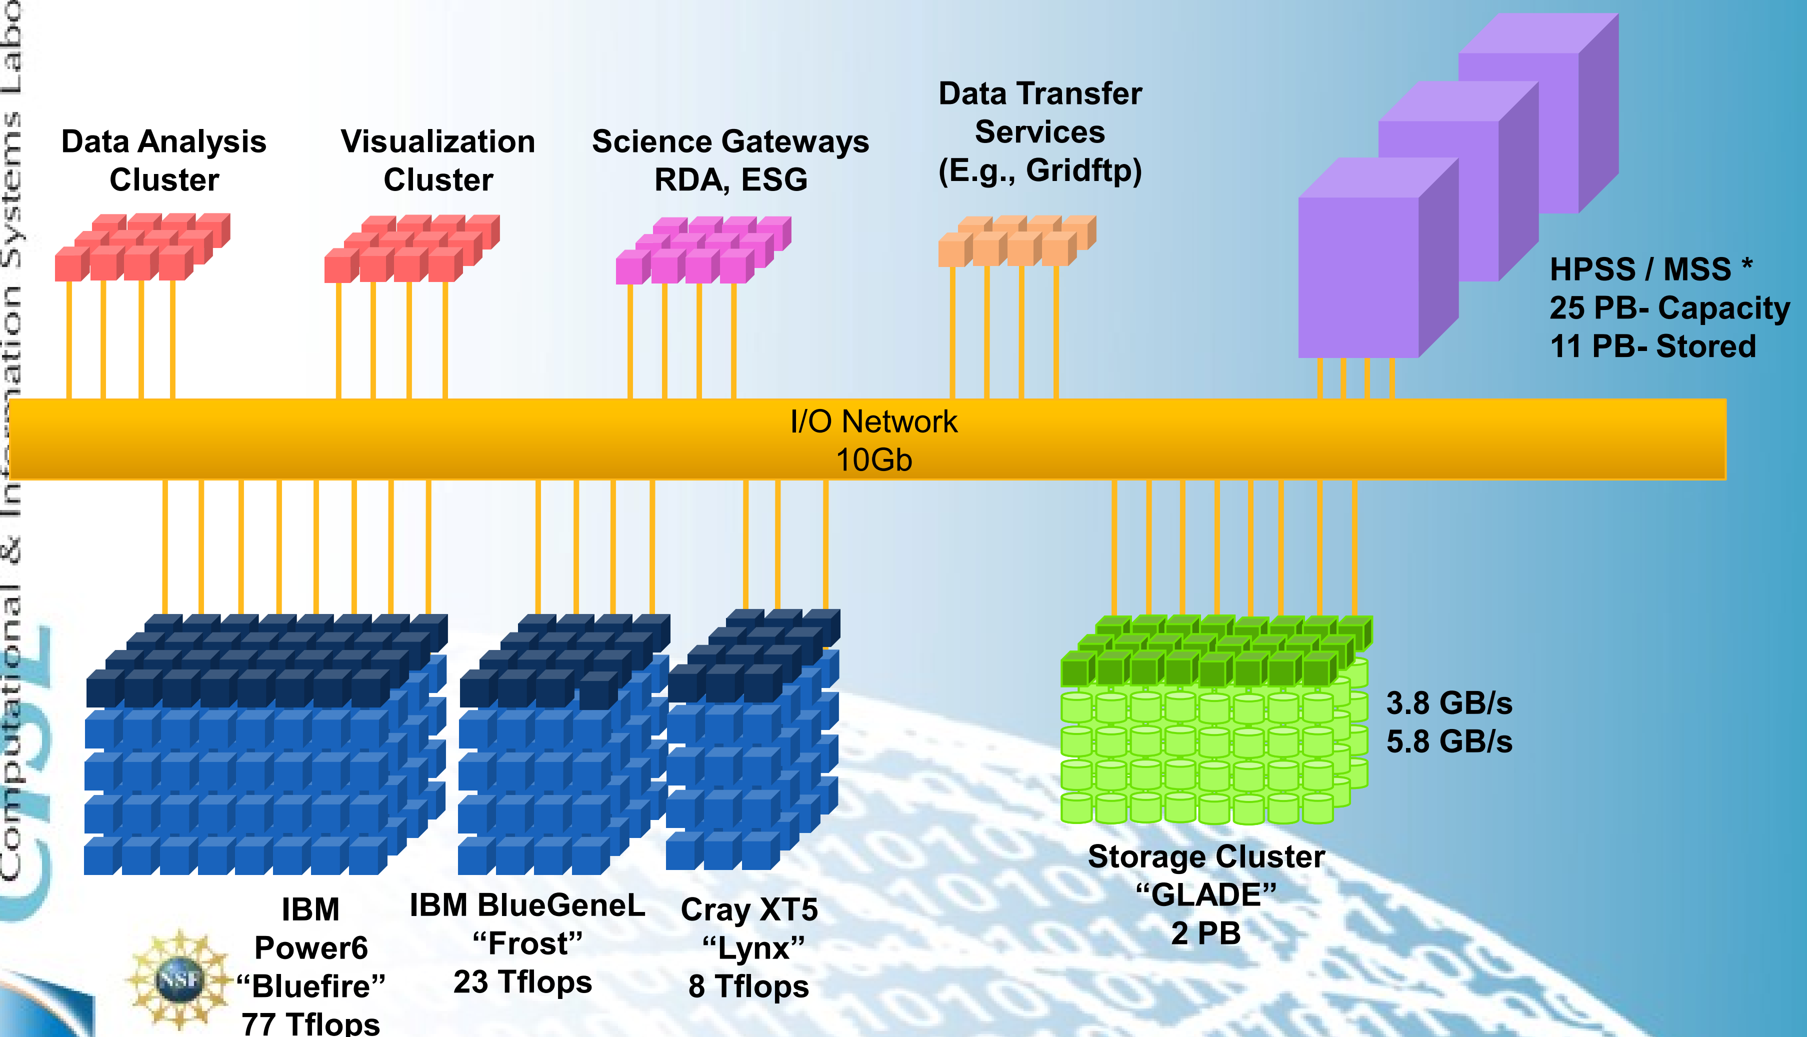

* Migrating from NCAR Mass Storage System to HPSS during 2010

CESM History File Sizes (GB)

	IPCC 1 x1	0.5 x 0.1	0.25 x 0.1	0.125 x 0.1	Exascale 1 km
ATM	0.2	0.9	3.6	14.4	2250
LND	0.1	0.2	0.9	3.6	560
ICE	0.7	4.4	4.4	4.4	440
OCN	1.2	19	19	19	1900
Total	2.2	24.5	28.9	41.4	5150

Extrapolation Estimate