Report from the
Earth Science Enterprise

Computational Technology Requirements Workshop

April 30 - May 1, 2002

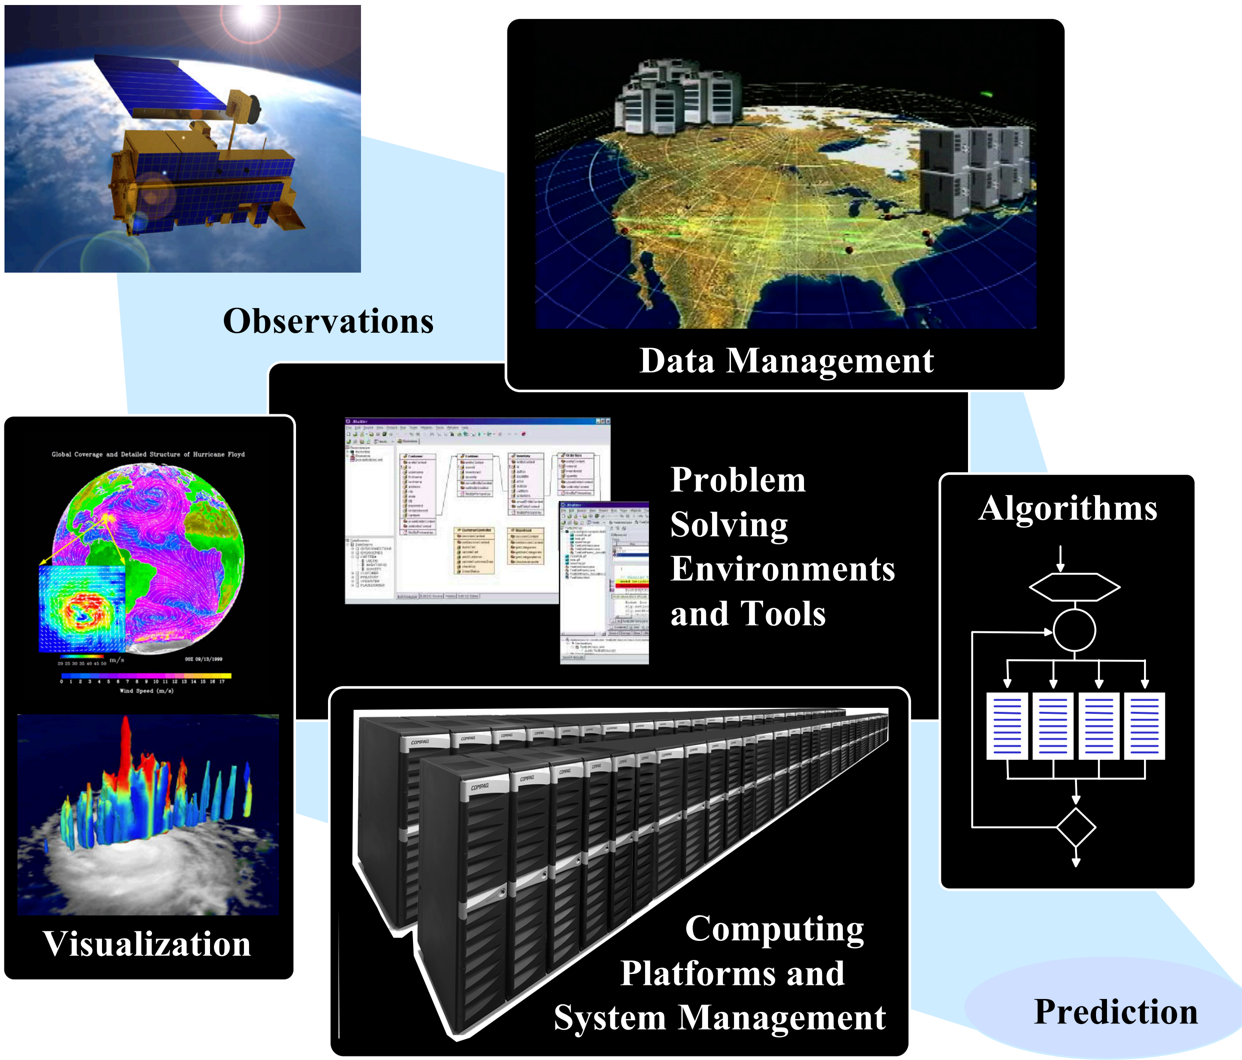

"The overarching purpose of Earth system science is to develop the knowledge basis for predicting future changes in the coupled physical, chemical, geological, biological, and social state of the Earth and assessing the risks associated with such change".

Earth Science Research Strategy for 2000-2010, p.17

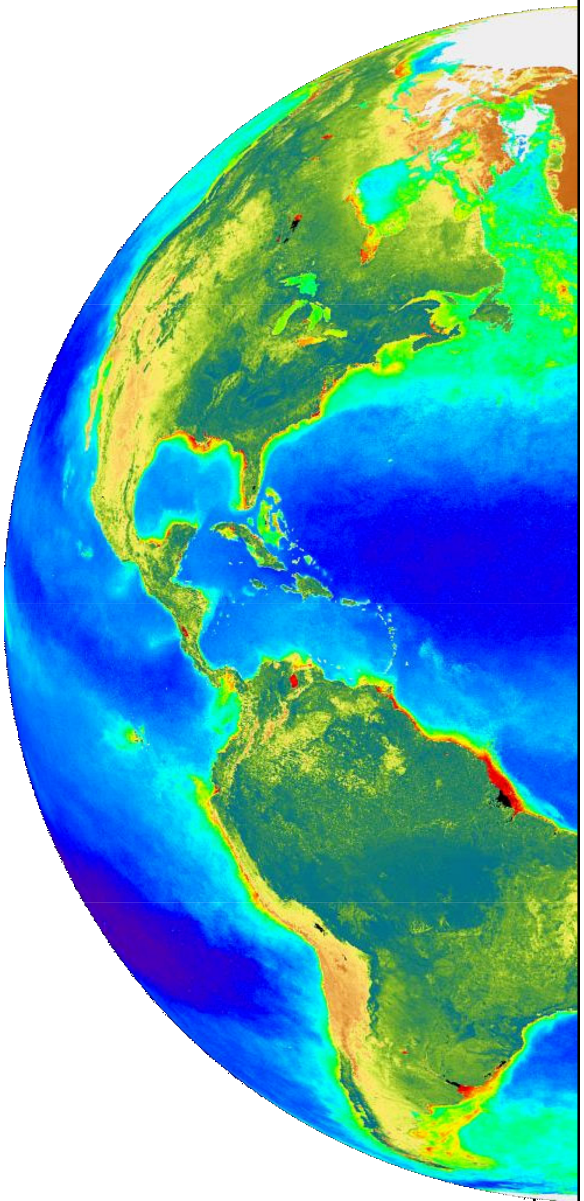

The First Step Toward Developing a Vision for the Future of High End Computing Program:

- Organized computational technologies requirement workshop April 30-May 2
- Discussions driven by ESE prediction goals around science themes of **Weather**, **Climate**, and **Solid Earth**

Workshop Charter:

- Identify ESE computing capabilities required to achieve ESE prediction goals for 2010 and beyond
- Evaluate current ESE capabilities against these requirements
- Determine and quantify what gaps in capabilities exist
- Determine which gaps can be addressed by advances in computational technologies
- Identify and quantify advancements in computational technologies that require NASA investment in order to bridge these gaps.
- Prioritize these capability advancements in terms of their likelihood to enable the Earth system prediction goals for 2010.
- Create a roadmap for each unique capability advancement.
- Create a final report.

Science Drivers

Weather - Robert Atlas/GSFC

Climate - Jim Kinter/COLA

Solid Earth - Andrea Donnellan/JPL

Capability Requirements

Cross-cut of panel reports - Robert Ferraro/JPL

Weather Science
Computational Requirements

Presented by Robert Atlas

Panel Members:

Robert Atlas/GSFC co-chair
Ricky Rood/GSFC co-chair
Arlindo da Silva/GSFC
Jim Dodge/HQ
Tim Hogan/NRL
Bill La Penta/MSFC
Shian-Jiann Lin/GSFC
Jeff McQueen/NOAA
Wei-Kuo Tao/GSFC

ESE Prediction Goals for Weather Prediction

Today's Capability

- 3-day forecast at 93%
- 7 day forecast at 62%
- 3 day rainfall not achievable
- Hurricane landfall+/- 400 Km at 2-3 days
- Air Quality day by day

2010+ Capability

- 5 day forecast at >90%
- 7-10 day forecast at 75%
- 3 day rainfall forecast routine
- Hurricane landfall+/- 100 Km at 2-3 days
- Air Quality forecast at 2 days

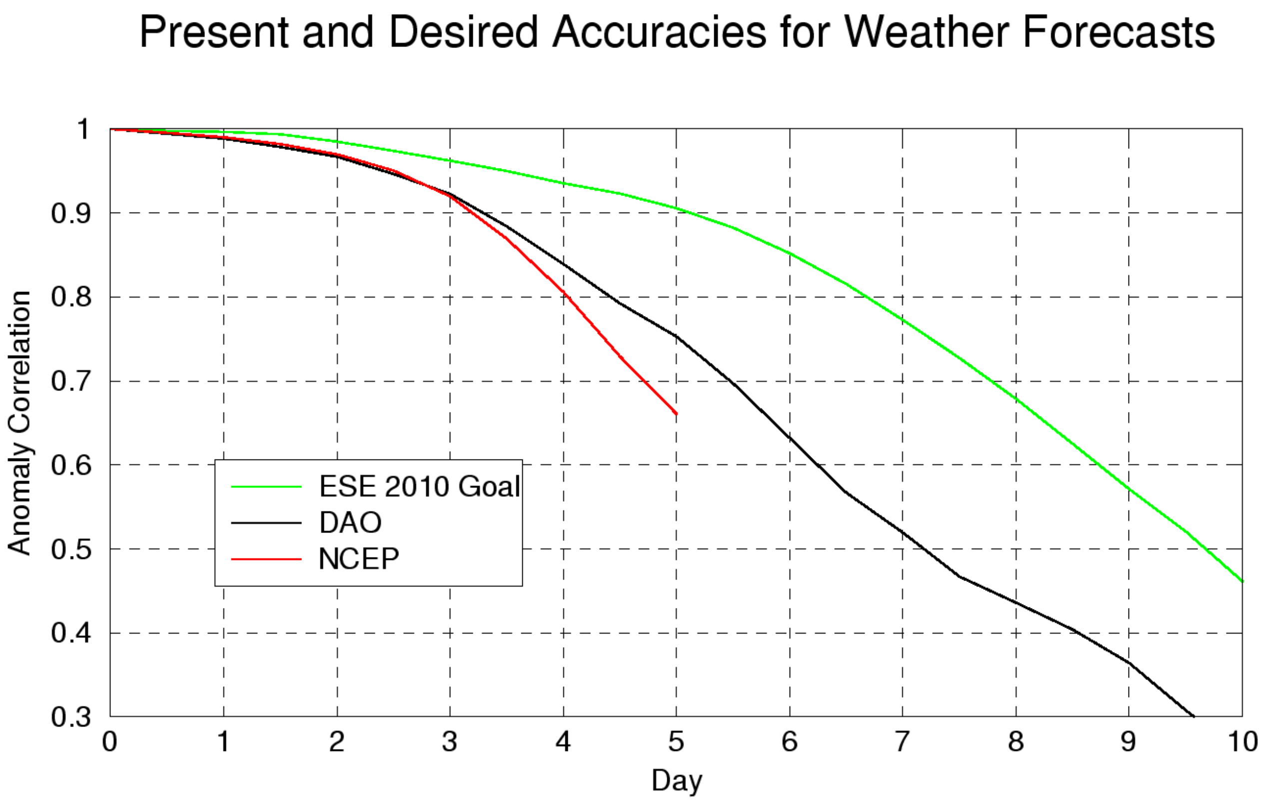

Development Areas (Systems Investment)

- Models
- Analysis
- Validation
- Optimization of Observing System
- Data Management
- Computational Systems
- Interfaces to Operational Agencies & Customers

Forecasting System for 2010

Forecast
model

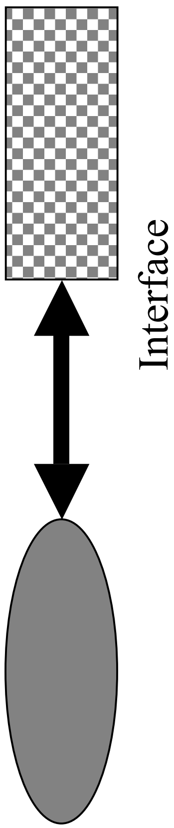

Assimilation/
Analysis
Algorithm

Global Resolution
(10 km, 100L)
Deterministic

Validation

Observations
used

(5 X 10^{10} daily)

Comprehensiveness/ Coupling

Upper Ocean

Land surface

Sea Ice

Chemistry

Physics

High Resolution Applications

Coastal

Air Quality

Meso/Cloud-scale

(Environmental Security)

Reanalysis

Regional Climate

Weather prediction:

10 day forecast (< 2 hours wall clock)

Climate reanalysis:

30 assimilated days (24 hours wall clock)

Operational Agency

End User

Impact on Storm Central Pressure With Increased Horizontal Resolution

132 hr GEOS forecasts

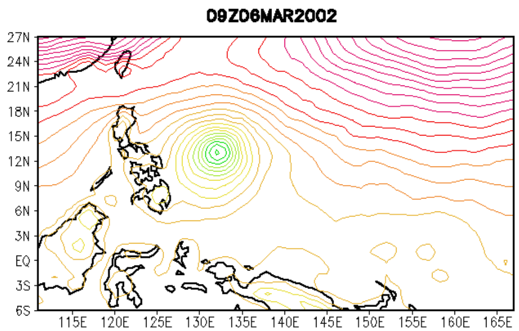

Min slp
1 x 1 :
998.754 hPa

Uniform grid

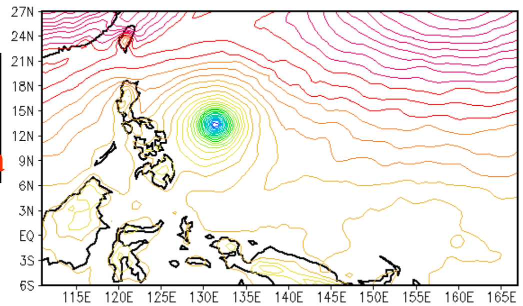

Min slp
0.5x0.5:
988.114 hPa

Stretched grid

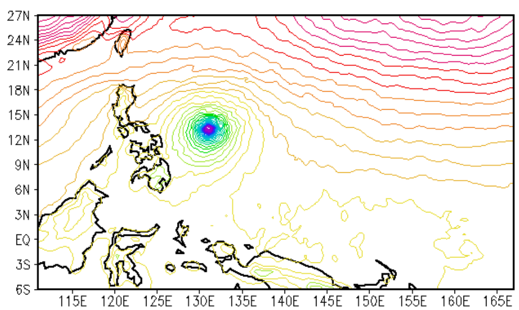

Min slp
0.25x0.25:
976.481 hPa

Stretched grid

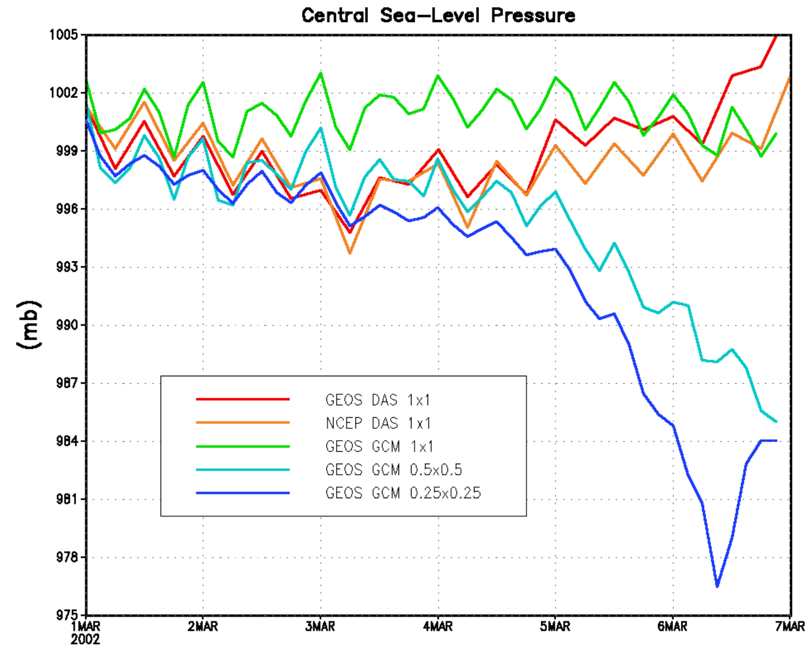

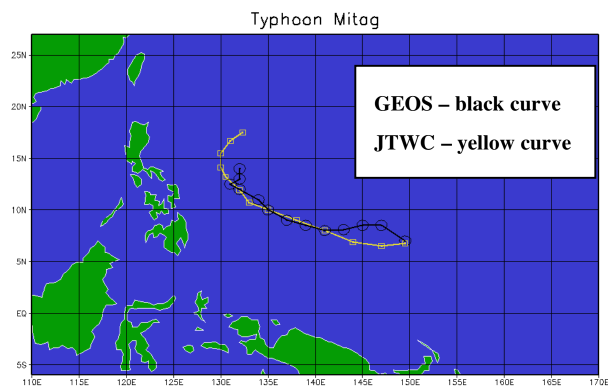

Super Typhoon Mitag: storm track

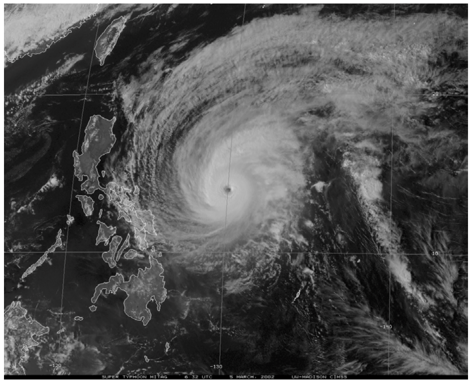

Super Typhoon Mitag: 06:32 20020305

http://cimss.ssec.wisc.edu/tropic/archive/2002/storms

Computing Requirements for Weather

	2002 System	2010+ System	
Resolution • Horizontal • Vertical levels • Time step • Observations o Ingested o Assimilated	 100 km 55 30 minutes 10^7 / day 10^5 / day	 10 km 100 6 minutes 10^{11} / day 10^8/ day	
System Components:	Atmosphere Land-surface Data assimilation	Atmosphere, Land-surface, Ocean, Sea-ice, Next-generation data assimilation Chemical constituents (100)	
Computing: • Capability (single image system) • Capacity (includes test, validation, reanalyzes, development)	 10 GFlops 100 GFlops	Must Have 20 TFlops (2000x) 400 TFlop (4000x)	Important 50 TFlops 1 PFlops
Data Volume: • Input (observations) • Output (gridded)	 400 MB / day 2 TB / day	 1 PB / day 10 PB / day	
Networking/Storage • Data movement o Internal o External • Archival	 4 TB / day 5 GB / day 1 TB / day	 20 PB / day 10 TB / day 10 PB / day	

Climate Science

Computational Requirements

Presented by Jim Kinter/COLA

Panel Members:

Michele Rienecker/GSFC co-chair
Al Kellie/NCAR co-chair
Dave Bader/DOE
Cecilia DeLuca/NCAR
Brian Gross/GFDL
Ming Ji/NOAA
Jeffrey Jonas/GISS
Jim Kinter/COLA
Tong Lee/JPL
Tsengdar Lee/HQ
C.R. Mechoso/UCLA
Max Suarez/GSFC

NASA ESE Short-term Climate Goals for 2010
6-12 month routine seasonal prediction
12-24 month experimental prediction

HARDWARE
Improve scaling: shared memory utilization, I/O, internode communication
single processor optimization
Network bandwidth

Single Image 2.5TF — aggregate throughput 80TF

Req'd turnaround: 1000 days/day

Required Scientific Capabilities
- **Coupled atmosphere-land-ocean-sea-ice**
- **Coupled (ocean and land) initialization**
- **Influences beyond ENSO ⇒ high latitudes, stratosphere, land cover/use changes**
- **Forecast reliability ⇒ multi-model ensembles**
- **Regional applications ⇒ more skill at intraseasonal timescales; resolving the mesoscale with non-hydrostatic models**
- **Prediction of extreme events ⇒ large ensembles**

2002
Difficult to share model modules
Slow data access & exchange

2°X30L atm
1/2°X30L ocean
10 member ensembles

2.5GF — 40GF

2005
SOFTWARE TOOLS
- 1st phase ESMF
- Interchangeable compts
- ODA support

DATA MANAGEMENT
- Distributed archives
- Manually generated catalogs & metadata

2010
SOFTWARE TOOLS
- ESMF extensions
- Flexible model components
- Full support for DA
- Integrated model environ.
- Analysis/visualization tools
- Data management S/W conforming to community code & data structure standards

DATA MANAGEMENT
- Virtual data archive
- 25PB distributed archive
- Fast access to remote 30TB

1/4°X40L atm
1/10°X40L ocean
100-member ensembles
Generate 500TB/day

6 mo seasonal mean forecast

12 mo seasonal forecast, include intraseasonal timescales
Extreme events

NASA ESE Medium-to-Long-Term Climate Goals for 2010: 10-year experimental prediction

Single Image	aggregate throughput
1.3TF	65TF
5GF	13GF

Req'd turnaround: 1000 days/day

Required Scientific Capabilities

- **Coupled atmosphere-land-ocean-ice-airchem-carbon models**
- **Resolve interannual-to-decadal variability**
- **use the same model, coarser resolution for longer global change integrations (~100 years)**
- **many small ensembles to explore parameter space**
- **include stratosphere, land cover/use changes**

HARDWARE

Improve scaling: shared memory utilization, I/O, internode communication
single processor optimization
Network bandwidth

2002

Lack of computer resources for required turnaround

2°X53L atm 4°X13L ocn
10 member ensembles

2005

SOFTWARE TOOLS
- Resources for code optimization

DATA MANAGEMENT
- Manually generated catalogs & metadata

2010

SOFTWARE TOOLS
- Transition to ESMF
- Analysis and visualization
- Data management software conforming to community code & data structure standards

DATA MANAGEMENT
- Fast access to remote 30TB

1°X100L atm, 1/2°X50L ocn; 40 tracers (10X computer time; 5-member ensembles
Generate 50TB/day

Decade-century integrations → Experimental decadal predictions; Global change assesments

NASA ESE CLIMATE: SCIENCE REQUIREMENTS → PERFORMANCE GOALS

The science requirements map to performance goals assuming a conservative rate of processor improvement (Moore's Law ~2X / 30 months)

-- IF --

Scalability is maintained through comparable improvements in switch & shared memory utilization H/W & S/W

Impediments to achieving science goals:

- **scaling limitations of H/W & S/W → can't reach anticipated performance**
- **lack of integrated tools to facilitate model use, cataloging of output**
- **lack of network bandwidth & S/W for sharing, analyzing & visualizing data**

Impediments to achieving performance goals: Machine balance

• highest priority	communication (internode & shared memory utilization)
• high priority	I/O bandwidth & storage
	Fortran compilers (improve processor performance)
• little influence	processor speed

THE NEXT GRAND CHALLENGE

Data Management & Distributed Access to Virtual Multi-Model Archive

The 2010 problem - using S-I prediction as an example:

Anticipated output: 500TB per image (multi-year output; low estimate)
5PB per "experiment": e.g., 10-member ensemble

Problem is probabilistic, so multi-model strategy required

Anticipated aggregate: ~ 5 groups sharing data
25 PB per experiment

Example analysis goal: explore extreme event for 1 season
need to share 30TB data volumes for analysis

Solid Earth Science
Computational Needs

Presented by Andrea Donnellan/JPL

Panel Members:
John Ries, UT Austin
John Rundle, UC Davis
Geoffrey Fox, Indiana U
Jay Parker, JPL
Robert Crippen, JPL
Eric DeJong, JPL
Ben Chao, GSFC
Weijia Kuang, GSFC

Solid Earth Science Questions

1. What is the nature of deformation at plate boundaries and the implications for earthquake hazards?
2. How is the land surface changing and producing natural hazards?
3. What are the interactions among ice masses, oceans, and the solid earth and their implications for sea level change?
4. How do magmatic systems evolve and under what conditions do volcanoes erupt?
5. What are the dynamics of the mantle and crust and how does the earth's surface respond?
6. What are the dynamics of the earth's magnetic field and its interactions with the earth system?

The Solid Earth is:

Complex, Nonlinear, and Self-Organizing

Computational technologies can help answer these questions:

1. How can the study of strongly correlated solid earth systems be enabled by space-based data sets?
2. What can numerical simulations reveal about the physical processes that characterize these systems?
3. How do interactions in these systems lead to space-time correlations and patterns?
4. What are the important feedback loops that mode-lock the system behavior?
5. How do processes on a multiplicity of different scales interact to produce the emergent structures that are observed?
6. Do the strong correlations allow the capability to forecast the system behavior in any sense?

Processes Take Place on Many Space and Time Scales

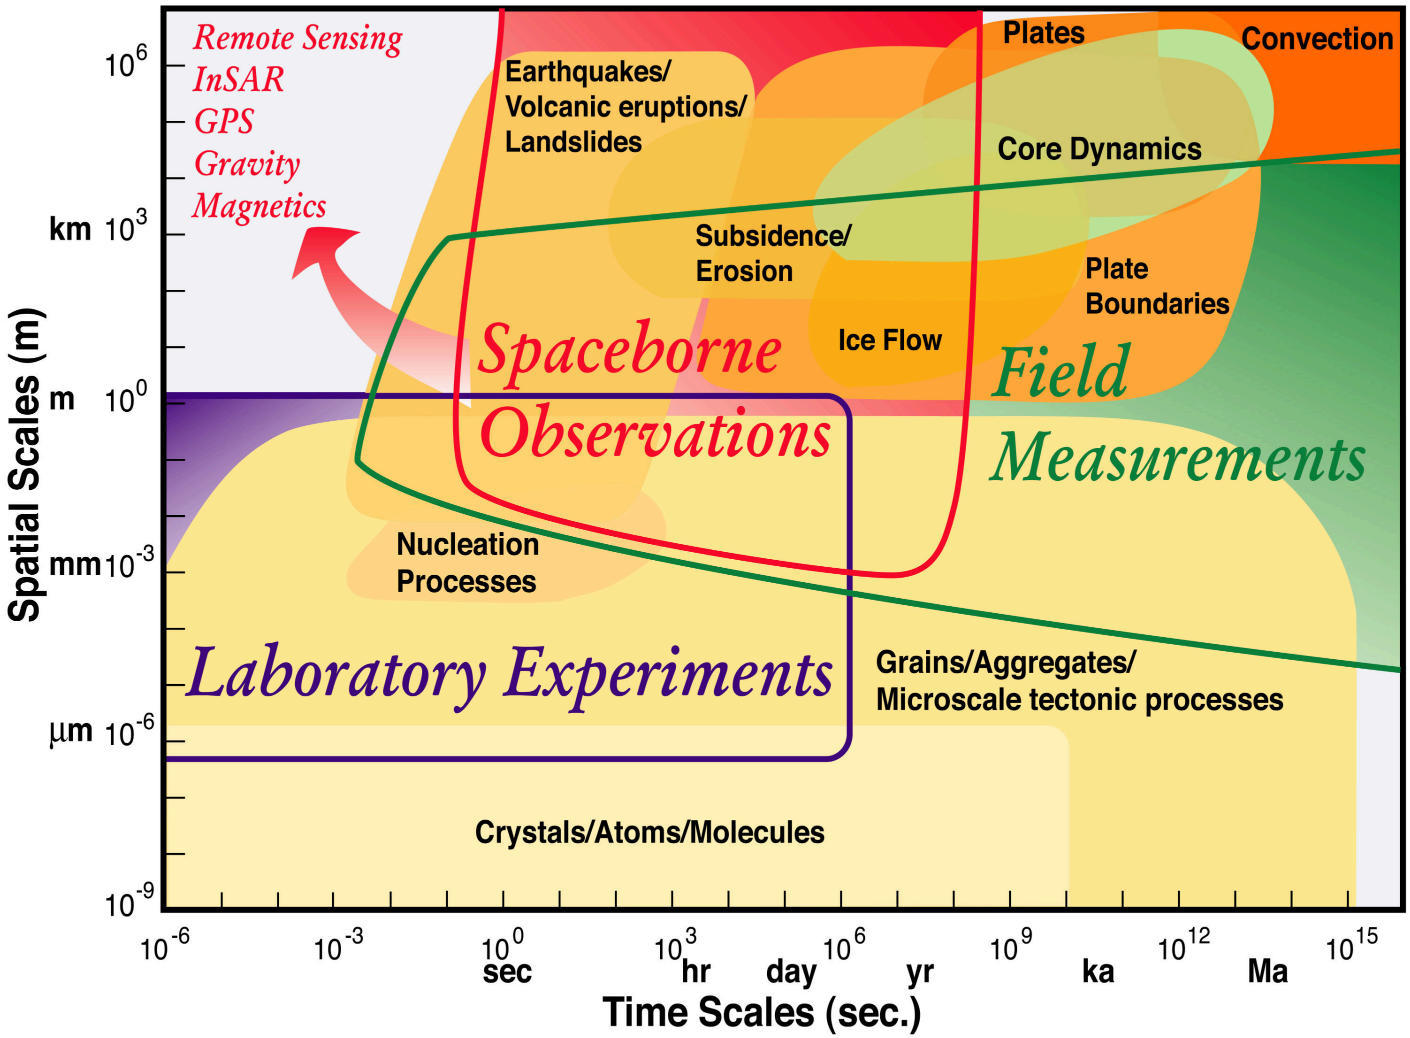

Recommendations

1. Create a Solid Earth Research Virtual Observatory (SERVO)
   - Numerous distributed heterogeneous real-time datasets
   - Seamless access to large distributed volumes of data
   - Data handling and archiving part of framework
   - Tools for visualization, datamining, pattern recognition, and data fusion
2. Develop an Solid Earth Science Problem Solving Environment (PSE)
   - Addresses the NASA specific challenges of multiscale modeling
   - Model and algorithm development and testing, visualization, and data assimilation
   - Scalable to workstations or supercomputer depending on size of problem
   - Numerical libraries existing within a compatible framework
3. Improve the Computational Environment
   - PetaFLOP computers with Terabytes of RAM
   - Distributed and cluster computers for decomposable problems
   - Development of GRID technologies

Solid Earth Research Virtual Observatory (SERVO)

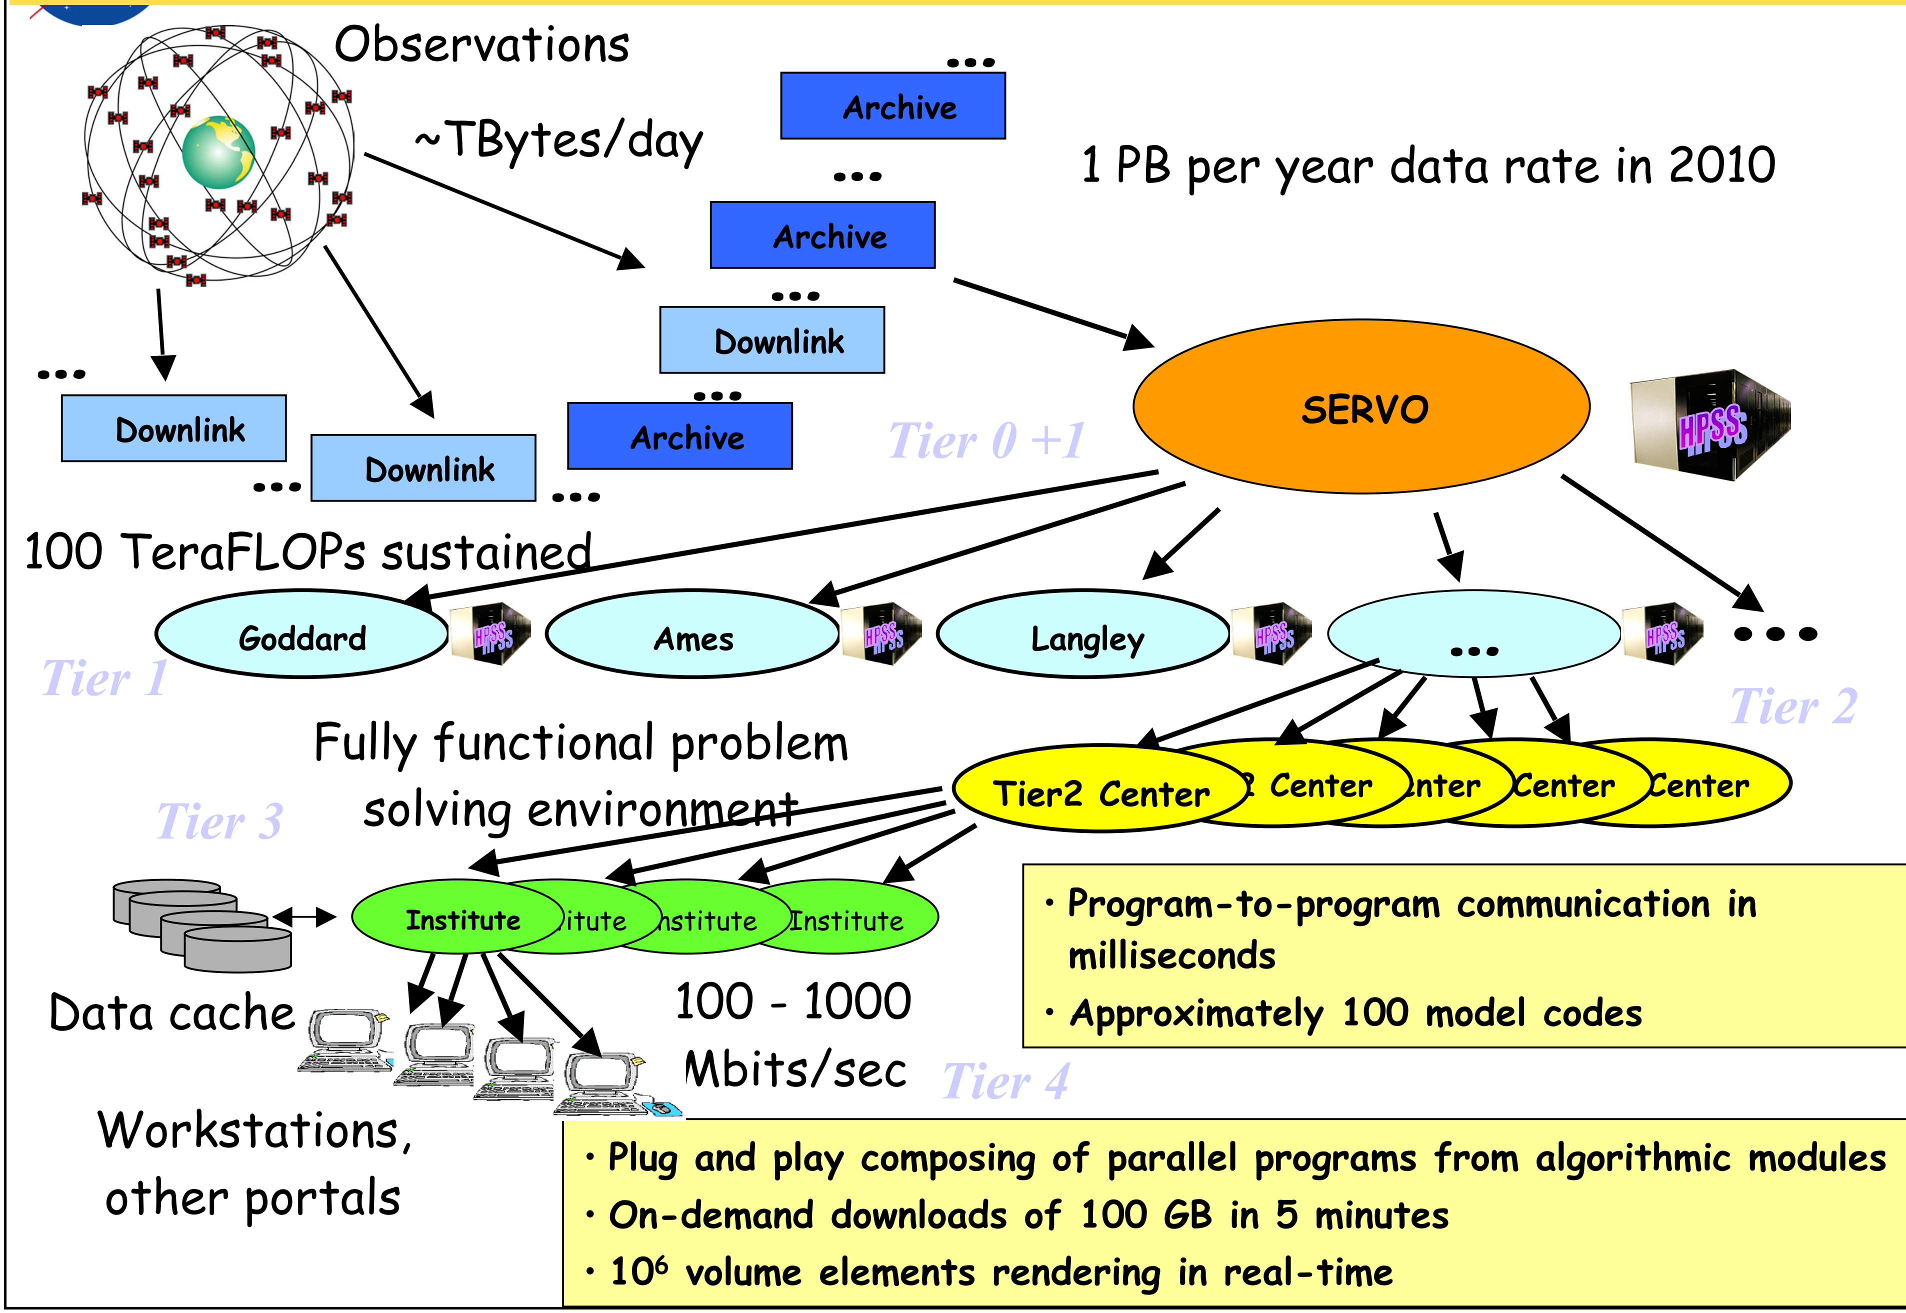

Virtual Observatory Project

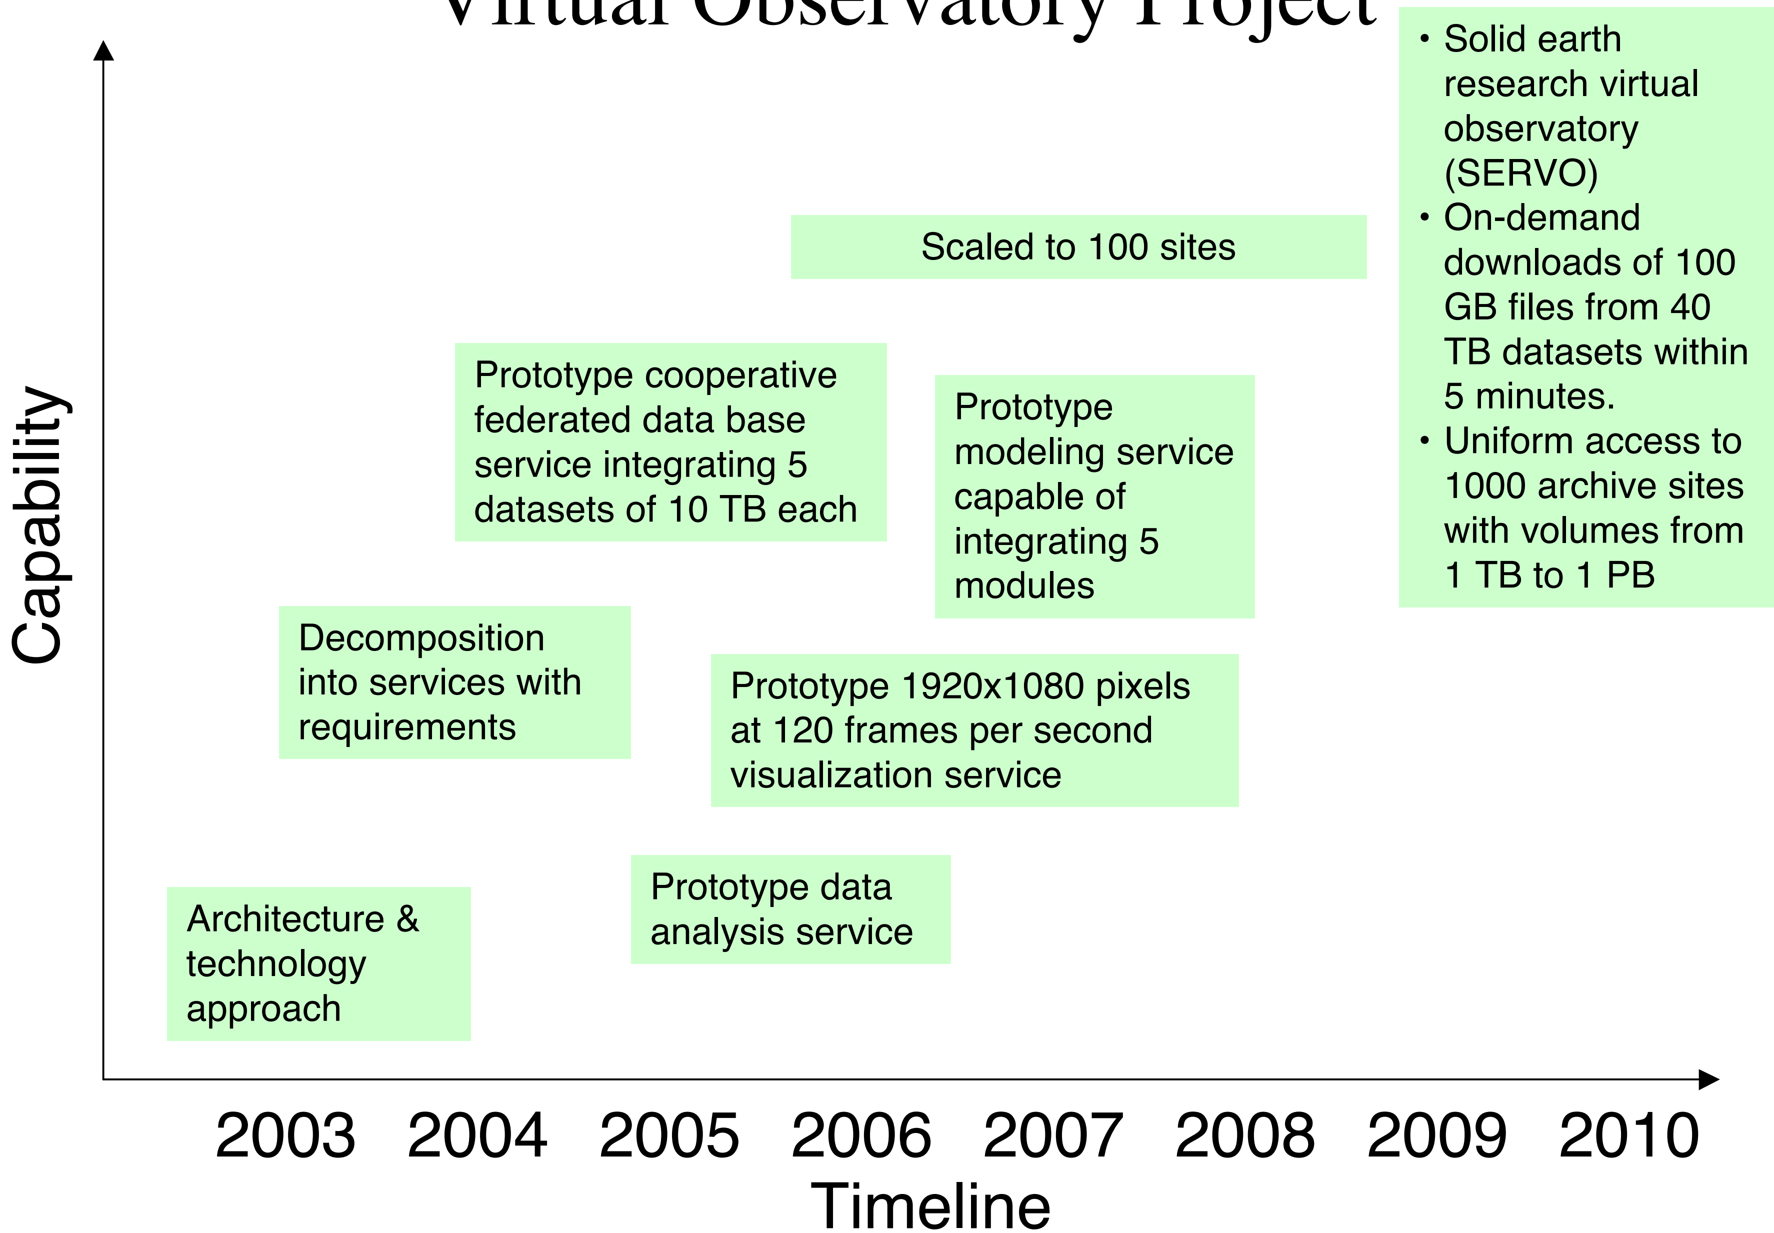

Problem Solving Environment Project

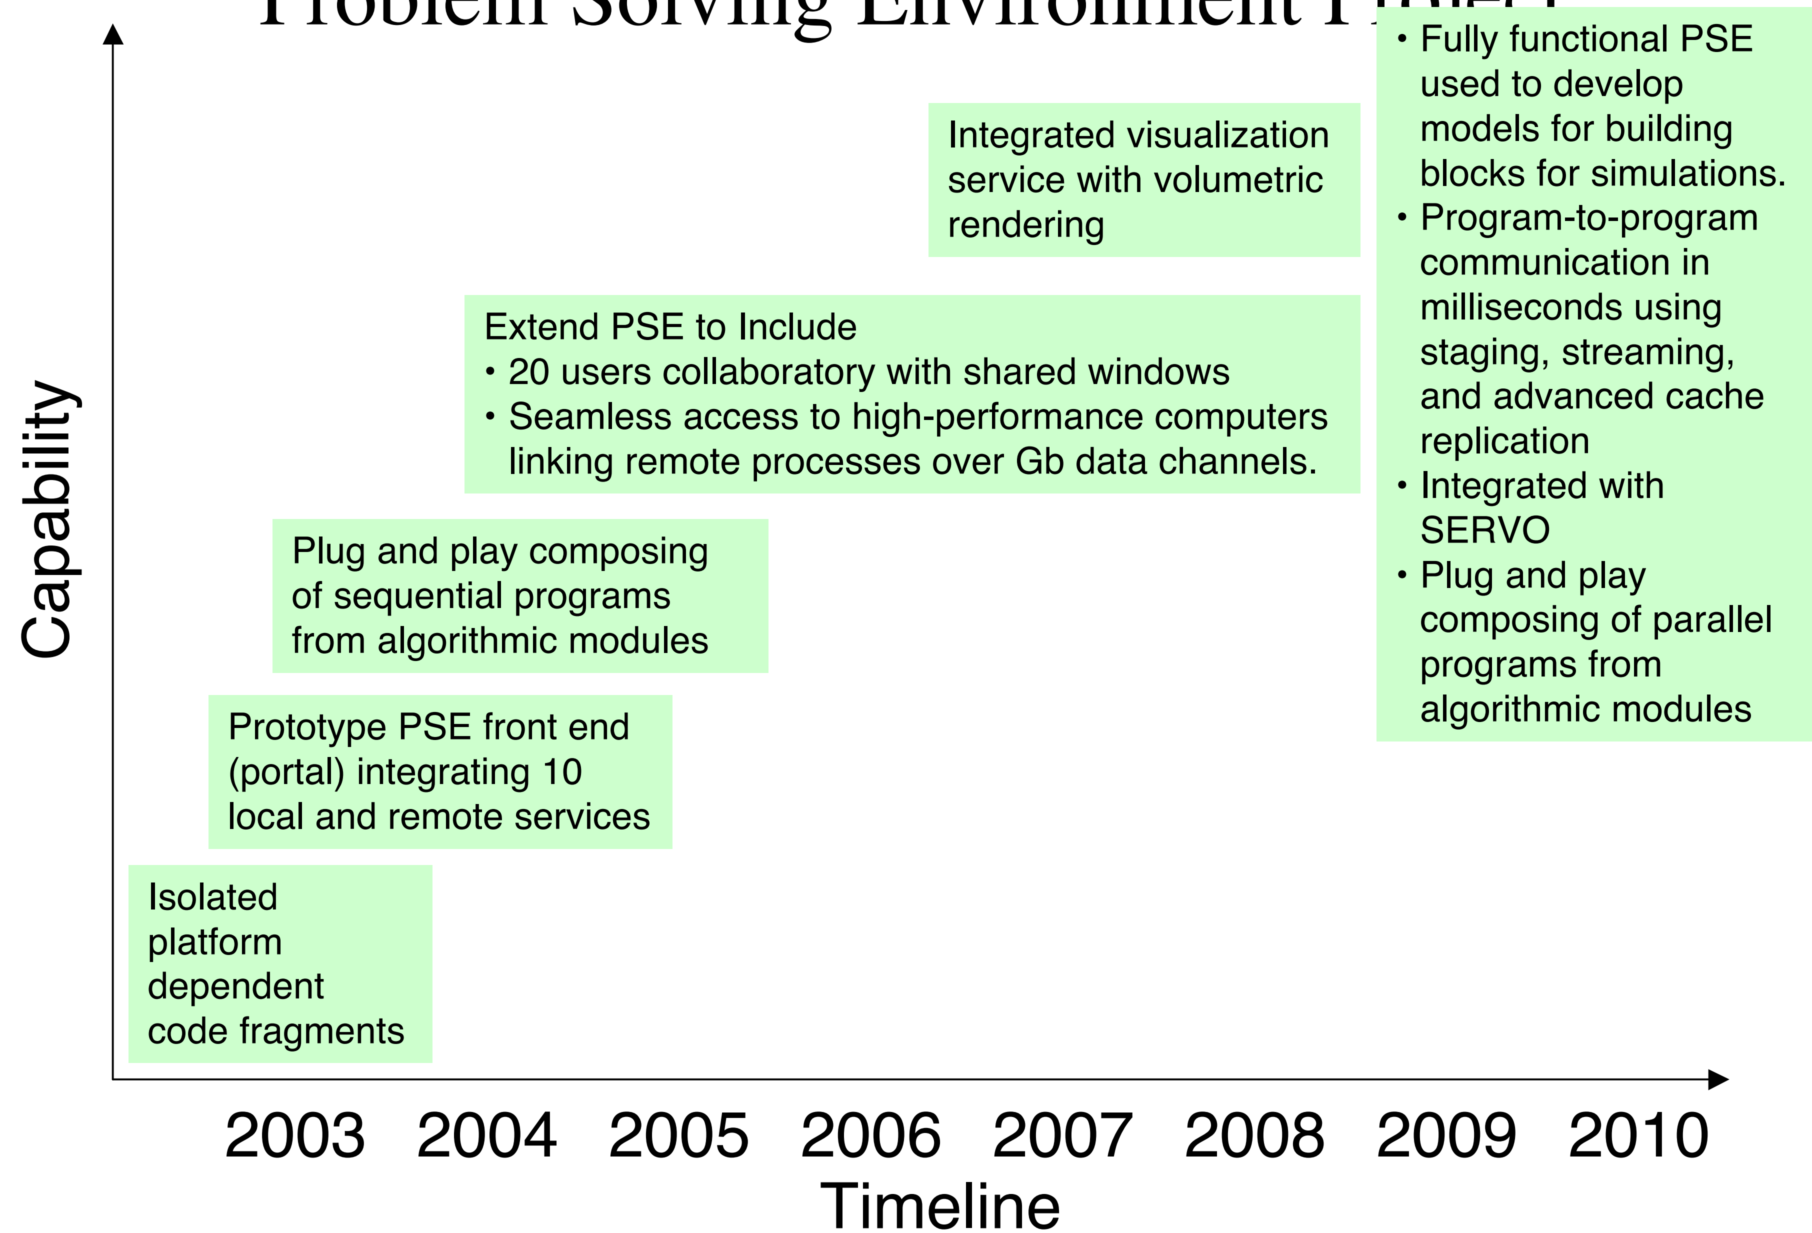

Computational Environment

~100 model codes with parallel scaled efficiency of 50%

~10^4 PetaFLOPs throughput per subfield per year

~100 TeraFLOPs sustained capability per model

~10^6 volume elements rendering in real time

Access to mixture of platforms low cost clusters (20-100) to supercomputers with massive memory and thousands of processors

100's GigaFLOPs
40 GB RAM
1 Gb/s network bandwidth

Capability

2003 2004 2005 2006 2007 2008 2009 2010

Timeline

Capability Requirements

Cross-cut of Panel Reports

Presented by Robert Ferraro/JPL

Cross-cut of Panel Capabilities Requirements

- Each Panel defined capabilities needed to achieve prediction goals in 2010
- Weather, Climate, and Solid Earth 2010 systems analyzed for stressing requirements
  - Strawman environments proposed by each panel
- Common Capability Requirements Identified
- Requirements Quantified where possible
- Use Scenarios presented where quantification is not possible
- Specific Technologies were not addressed by panels
  - Panels determined "what" was needed
  - Left it to technology experts to determine "how" to provide it
- The future path of computing technology is hard to predict
- The best approach is to drive development with specific science-required capabilities

Technology Cross-Cut of Gaps Identified

A. Computing Platforms

B. Data Management

C. Programming Environment and Tools

D. Distributed Computing

E. Other Requirements

A. Computing Platform Throughput Required

	Stressing Model	Single Image Throughput	Estimated Capacity Required
Weather	10 Day Forecast Atmosphere:10 km horizontal, 100 levels vertical 10^{11} observations	20 Tflops	400 Tflops
Climate	S-I Prediction Atmosphere: 25 km horizontal Ocean: 6 km horizontal	5 Tflops	100s Tflops
Solid Earth	Earthquake Fault Slip 16M finite elements 100k boundary elements	2 Tflops	10s – 100 Tflops
Sustained Throughput and Capacity Requirements			

- Single application requirements derived from current performance extrapolated by required resolution increase
- Capacity requirements are based on current experience scaled up to the 2010 strawman environments

Technology Issues

Vendors are expected to offer 20 - 50 Gflops processors, platforms with 10,000 processors, plenty of memory and storage

Gaps are in achievable applications performance

Needs:

- Single processor application performance at a significant fraction of peak
- Application scalability to thousands of processors
- I/O performance that scales with the application performance

Technology required:

- Performance optimization tools
- Compilers that achieve a significant fraction of peak architecture performance
  - Programming language/paradigm continues to impact this ability
- Scalable, portable algorithms
- Scalable operating systems
- Low latency, high bandwidth interprocessor communications
- Low latency, high bandwidth parallel I/O

B. Data Management Requirements

	Observational Data	Access Modes Rates	Output Data	Storage Term/Re-access Mode
Weather Forecast	1 TB/day Multiple Sources Continuous	Streamed input 20 GB/s	10 PB/day – Archival 10 TB/day – external distribution	Medium – Long Catalogued
Climate Modeling	10s of GB from archival sources	Data archive request 2 GB/s (latency tolerant)	100s TB/day	50% Short term - Immediate analysis 50% Medium term - Catalogued
Solid Earth Research	100s of GB/day Distributed sources	Distributed archives – low latency access	1 PB/day – ingested into distributed archives	Medium – Long Catalogued access

- Data volume is expected to be overwhelming and heterogeneous in format

  Model output data management is the problem

- Current practice does not scale to these volumes
- Data storage expected to be geographically distant from data consumers
- Uniform, seamless identification, indexing, and access methods required

Technology Issues

Vendors expected to provide physical storage solutions

Gaps are in management, distribution of data volume

Needs:

- Uniform, location independent service for identifying, managing, and accessing metadata and raw data
- Data transport performance that scales to consumer requirements
- Low latency random access

Technology required:

- Internal bandwidths @ 300 Gbytes/s
- WAN bandwidths that scale to provide 150 Mbytes/s for *every producer/consumer simultaneously*
- Reliable, location independent data transport services
- Intelligent Data caching
- New data organization & management applications
  - Next generation database tools for science data
  - Location independent view and control of data

C. Programming/Problem Solving Environment Requirements

- Applications to become much more complex
  - Composed of many separate models, data sources
  - Componentization necessary for maintainability, upgradability
  - Applications can consist of geographically distributed pieces
  - *No single person will understand all of the details*
- New applications need to be implemented in a month instead of a year
  - The research requirements demand agility in the model implementation
  - Reusable software is a must
- Performance (efficiency) must be maintained without heroic efforts
- Ensemble executions, distributed application executions must be transparently manageable

Technology Issues

No discipline specific vendor offerings expected in this area

Needs:

- Application frameworks/composable component architectures
- Platform independent program design and execution environment
- Highly efficient applications that scale to 1,000s of processors without heroic effort

Technology required:

- Discipline specific frameworks and components (e.g. ESMF, SERVO)
- Performance tuning tools
- Portable, scalable components from which to build applications
  - Numerical libraries
  - PDE component toolkits
  - Data, information, and knowledge components
- Standardized program execution environment independent of execution assets

D. Distributed Computing Requirement

- Integration of geographically distributed data servers, computing assets, and users will be the norm
  - Thousands of data servers
  - Thousands of application servers
  - Thousands of users
- Assets need to be unified in a seamless environment for maximum productivity
  - Scheduling of assets
  - Low latency request servicing
  - Uniform and universal naming [metadata] and access schemes
  - High bandwidth data transport
  - Programming Model needs to hide/optimize access latency penalties
- Transparent, reliable data transport layer for interservice communication is required

Technology Issues

Unpredictable vendor offerings in this area in 2010

Currently, there is a multi-agency investment (NSF, DOE, NASA) in this area (GRID computing)

Needs:

- Uniform, seamless, transparent access and programming environment

Technology required:

- Universal standard middleware
- Transparent data caching mechanisms
- Reliable high bandwidth, standardized data transport layer
  With automatic data format translation
- User single entry point with global application execution control
- Distributed application program composition tools
- Distributed performance evaluation tools spanning the data sources, application services, clients, and networks

E. Other Requirements

- Scalable, Efficient Implementations of Common Algorithms
  - Current practice is user implemented and application specific - generally not reusable
  - New implementation paradigms required that:
    - Allow transparent incorporation into applications
    - Maintain efficient execution on 1,000s of processors
    - Automate latency tolerance
    - Are completely portable
- Real time visualization of Terabytes of data
  - Current commercial offerings will not scale to these data volumes
  - Must use vendor provided display technology
    - Expect 10x display resolution over 10 years
  - Visualization Applications must;
    - Ingest Terabytes of data from geographically distributed sources
    - Render arbitrary combinations of diverse data sets in real time
    - Deliver rendered product to the end user interactively at their local site

Other Requirements

- Scalable Data Mining Applications
  - Currently in its infancy
  - New data mining applications for extracting science information require development
    - Pattern analysis over order 10^6 features
    - Wavelet analysis on order 10^6 scales
    - Inversion techniques for order 10^6 parameters
    - Distributed data sources
- Computing Platform Systems Management
  - Currently Ad Hoc
  - System engineering practice needs to integrated into the computing center infrastructure design and operation
    - Performability analysis tools
  - Building discipline specific computing centers needs to be supported by integration tools and systems management best practices

Technology Issues

- No vendor offerings expected for algorithms
- Vendor improvements in display technology will bound end-user viewing resolution
- "Last Mile" network bandwidth a factor that needs to be considered
- Commercial Data mining applications do not apply to science and model data

Summary

- NASA science requires major advances in computational technology
  - NASA's unique driver is the data
- Achievement of the prediction goals will require coordinated investments in science advancement and computational technologies
  - Key stressing science applications are already identified
  - Data management, application throughput, and problem solving environments are common across disciplines
  - Science applications will build upon industry best practices, standards, and commercial offerings - But industry will not provide key technologies required to enable the stressing applications
  - Coordinated investment in these identified technologies will benefit all the disciplines
- Continued, focused investment in a science driven technology development program is required for success in the ESE