NASA Goddard's Vision* for 10 Gigabit Ethernet

J. Patrick Gary

Network Projects Leader

Earth and Space Data Computing Division

NASA Goddard Space Flight Center

pat.gary@nasa.gov

301-286-9539

Presentation for informational seminar sponsored by Force10 Networks, Inc.

NASA Goddard's Vision* for 10 Gigabit Ethernet

***Caveat: actually <u>merely</u> from Pat's point of view**

- NASA
  - » Earth Science Enterprise (Code Y)
    - – Goddard Space Flight Center
      - Earth Sciences Directorate
        - » Earth and Space Data Computing Division

Presentation/Discussion Topics

- Mission and Goals of NASA, its Earth Science Enterprise, & GSFC's Earth Science Directorate
- Challenges & Initiatives in NASA's Earth Science Information Technology Program
- NASA/GSFC's Various Networks
- Some Multi-GE Network R&D involving GSFC
- 10 GE Testing by Bill Fink (GSFC) & Paul Lang (ADNET)

NASA'S VISION

To improve life here,
To extend life to there,
To find life beyond.

NASA'S MISSION

To understand and protect our home planet
To explore the Universe and search for life
To inspire the next generation of explorers
… as only NASA can.

ESE Fundamental Science Questions

How is the Earth changing and what are the consequences of life on Earth?

- How is the global Earth system *changing*?
- What are the primary *forcings* of the Earth system?
- How does the Earth system *respond* to natural and human-induced changes?
- What are the *consequences* of changes in the Earth system for human civilization?
- How well can we *predict* future changes in the Earth system?

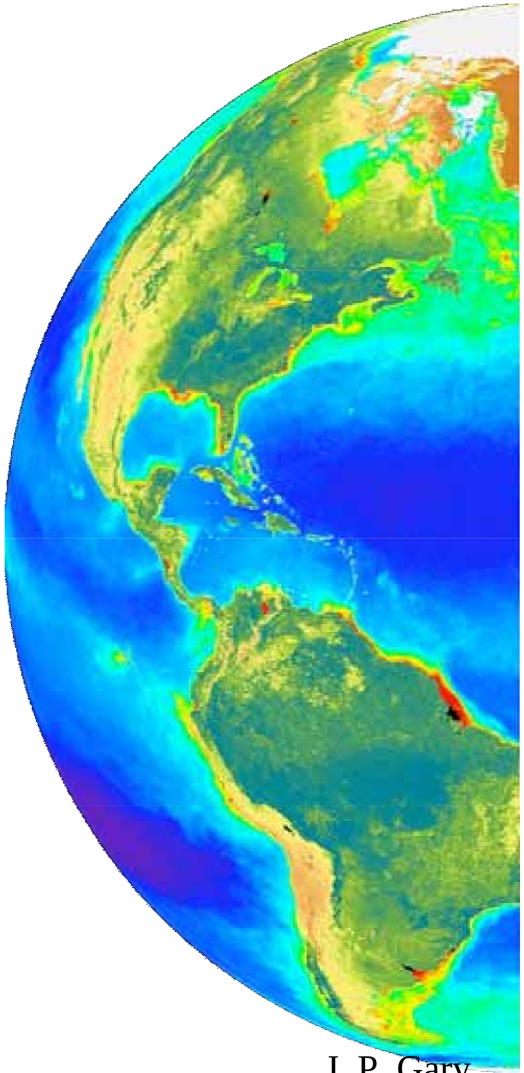

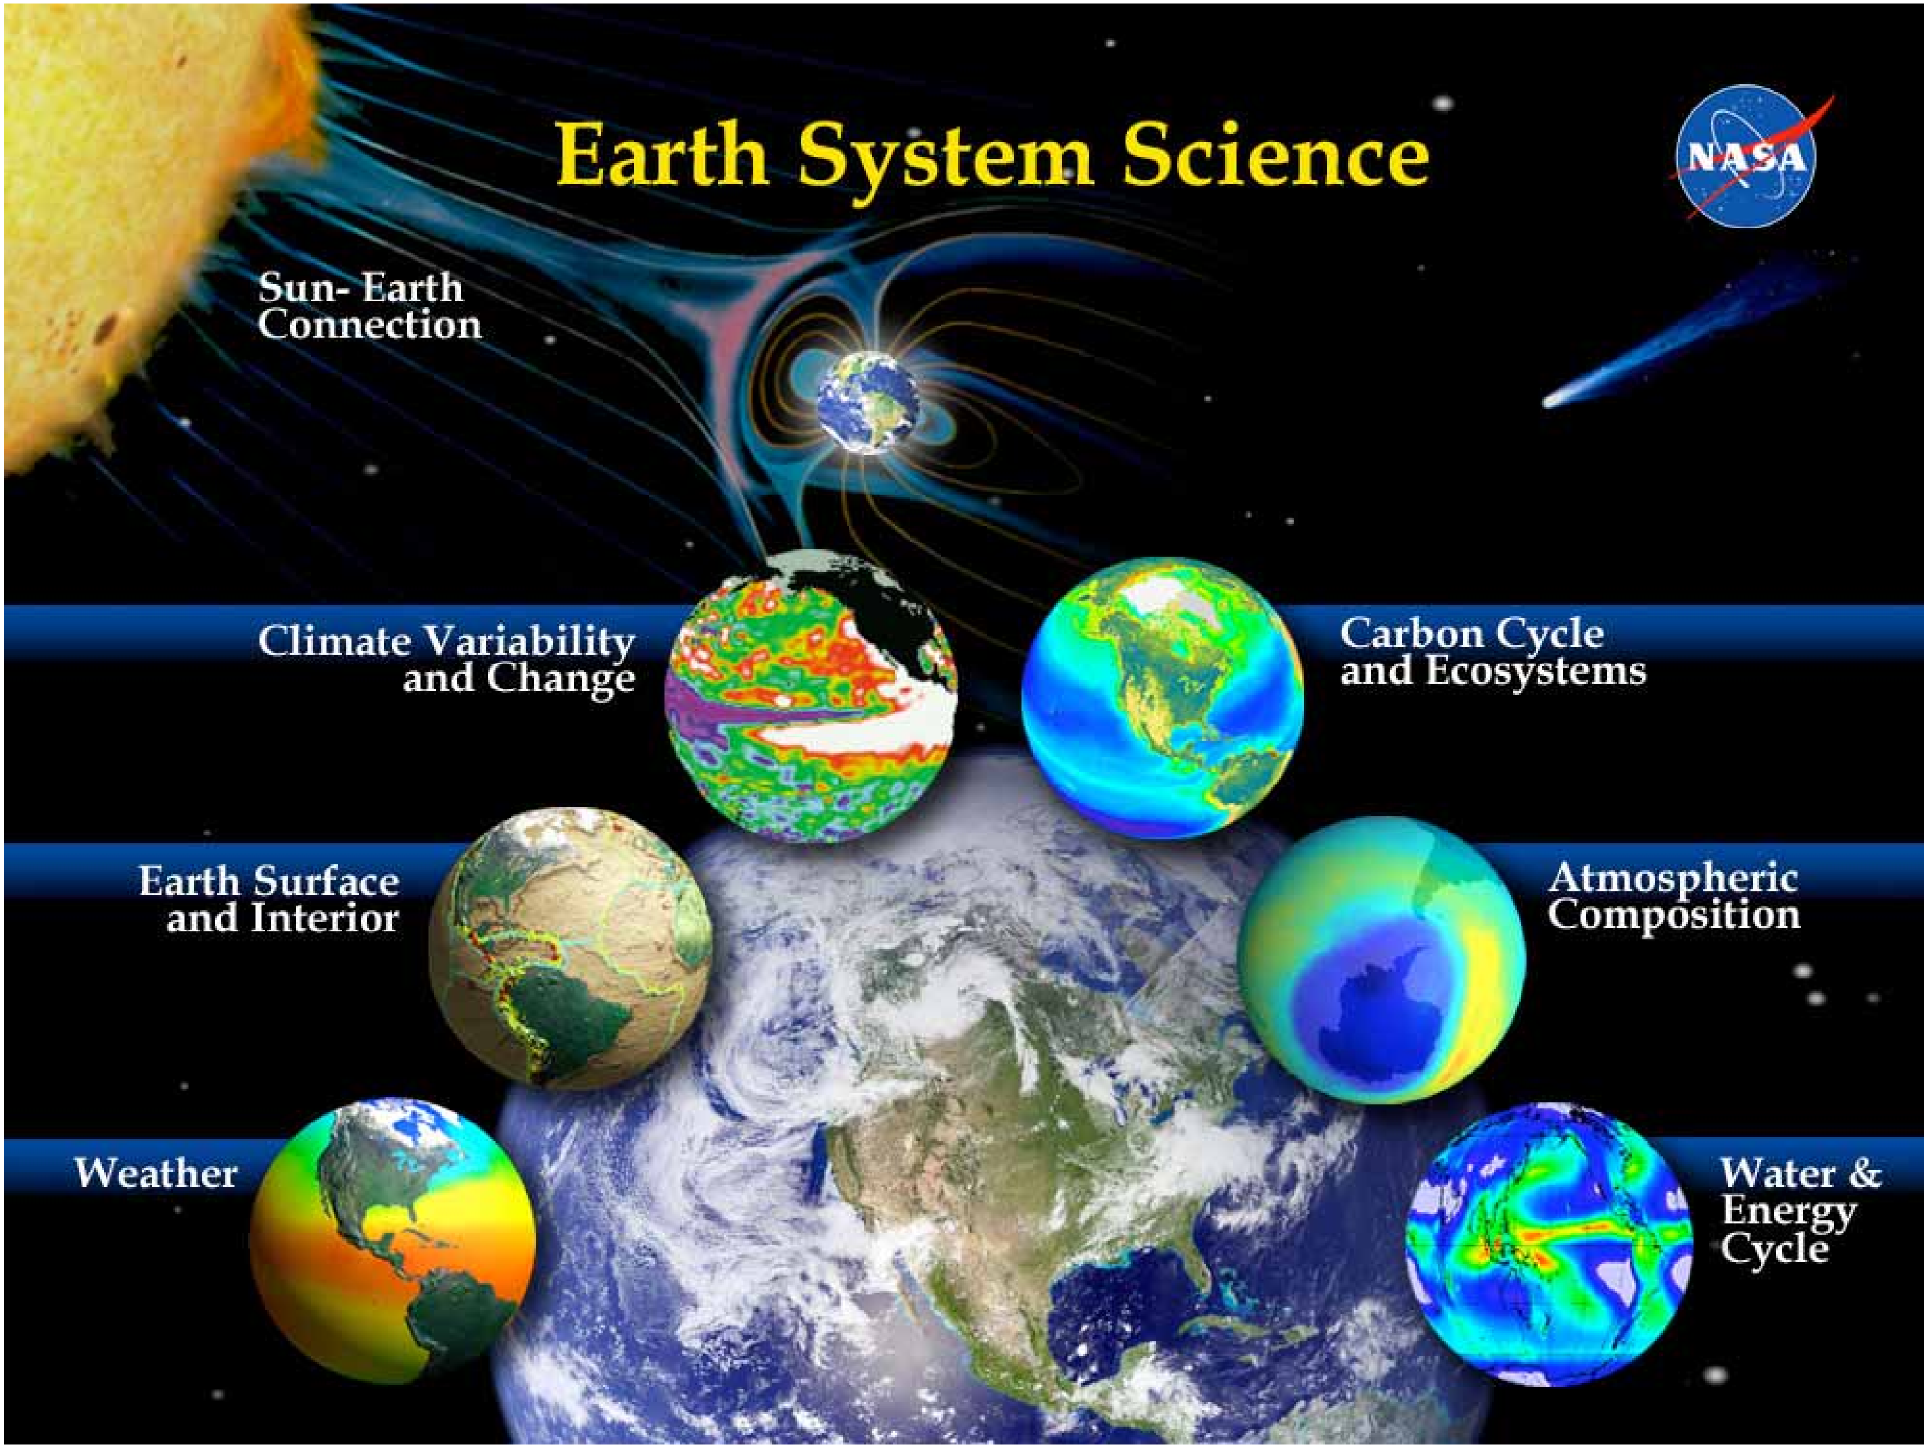

Earth System Science
NASA
Sun- Earth Connection
Climate Variability and Change
Carbon Cycle and Ecosystems
Earth Surface and Interior
Atmospheric Composition
Weather
Water & Energy Cycle

NASA Earth Science Research Satellites

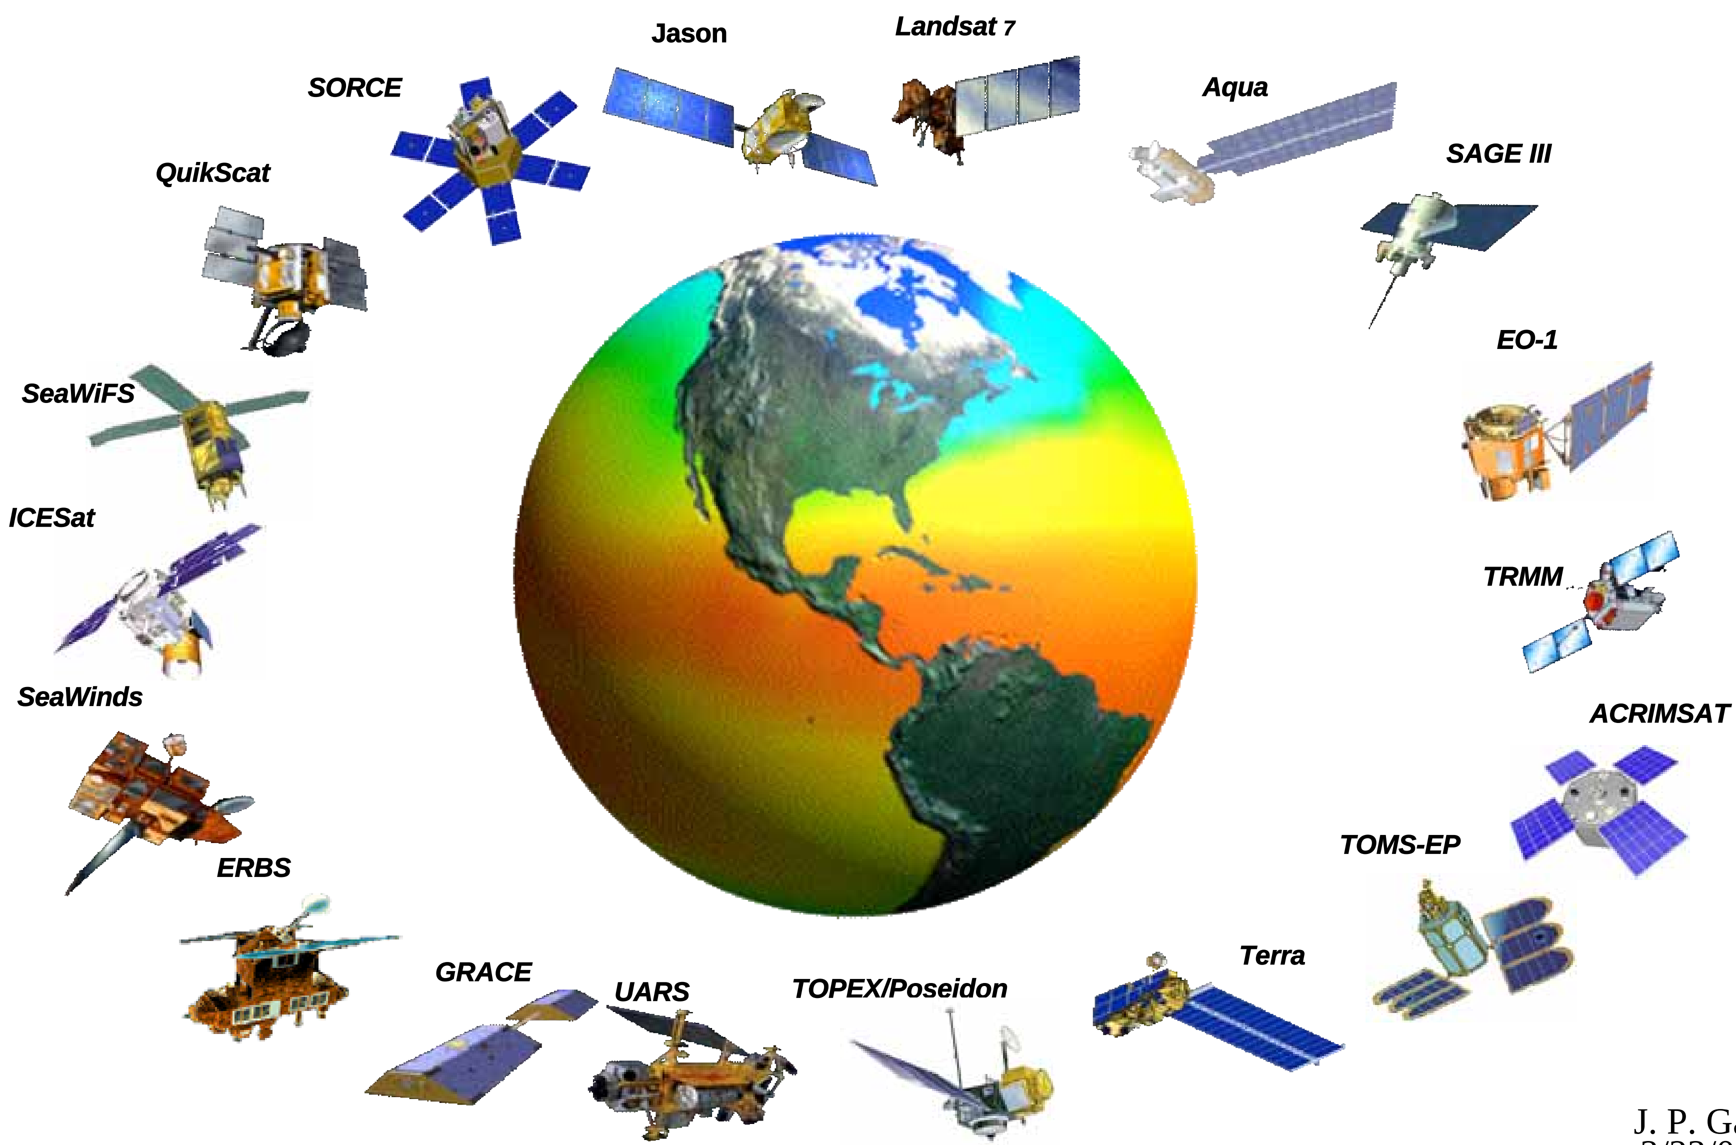

J. P. Gary
3/23/04

Next Generation Missions

- The Earth Sensorweb Concept Involves Satellites Working In
- Intelligent Constellations, Adapting To Observed And Modeled Changes
- And Delivering Tailored Information Products From Space To Science Users

Components of a Future Global System for Earth Observation

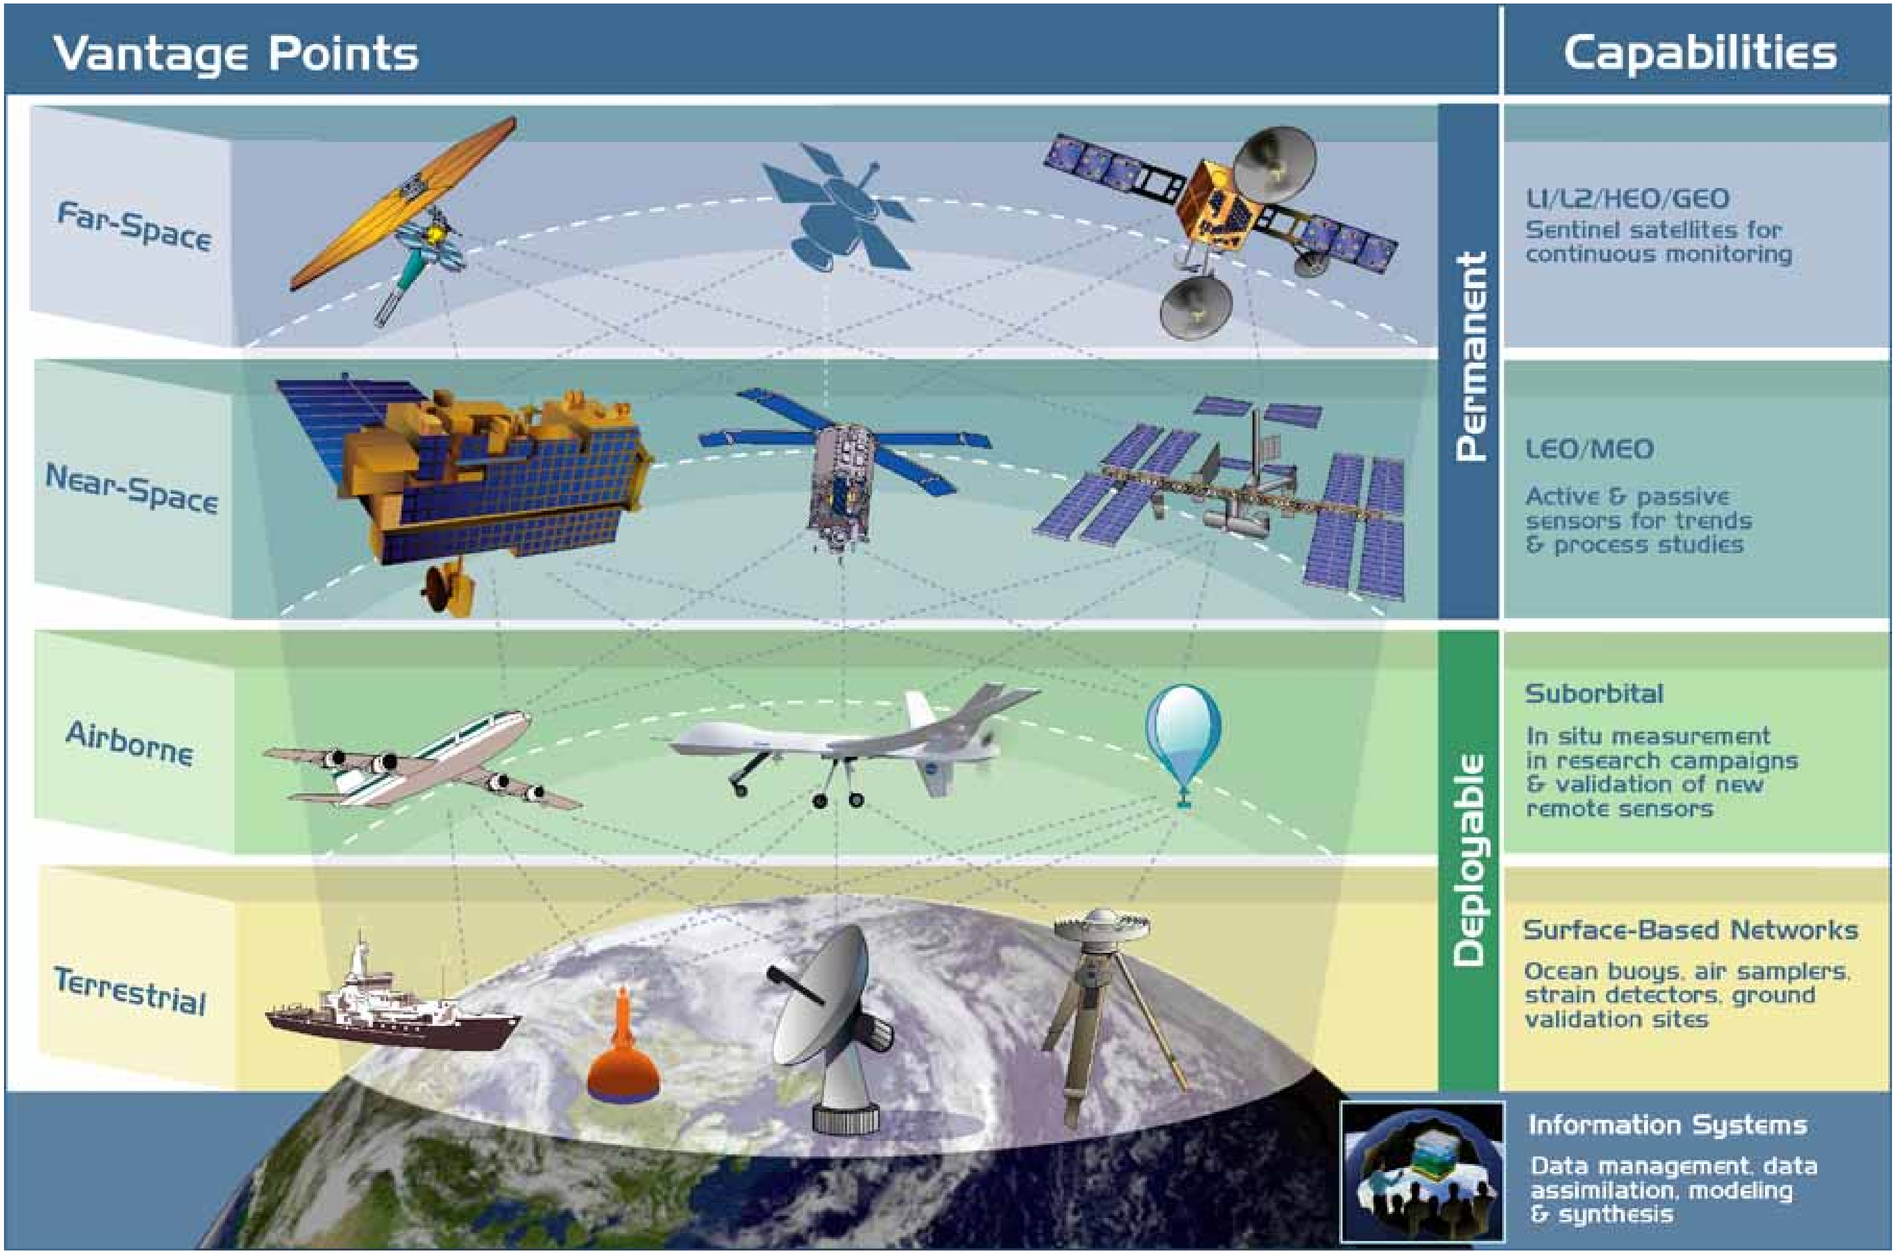

EARTH SCIENCES DIRECTORATE

Goddard Space Flight Center

J. P. Gary
3/23/04

Earth Science Themes

Earth Science Enterprise Themes

- *Water and Energy Cycle*
- *Carbon Cycle and Ecosystem*
- *Climate Variability and Change*
- *Atmospheric Composition*
- *Weather*
- *Earth Surface and Interior*

→

Earth Sciences Directorate Themes

- **Global Water and Energy Cycle**
- **Carbon Cycle**
- **Weather and Short Term Climate Forecasting**
- **Long Term Climate Change**
- **Atmospheric Chemistry**
- **Aerosols**
- **Geodynamics and Other Solid Earth Activities**

GSFC Earth Science Directorate Vision and Mission

VISION

To develop and acquire knowledge of the Earth through discovery leading to the improvement of life

MISSION

- Provide leadership and serve as a resource in Earth system science and technology
- Improve predictions of the Earth system through new observational and modeling capabilities
- Establish partnership with agencies with operational responsibility to promote Earth science applications
- Advance understanding of the evolution of the Earth System through the exploration of planets
- Enhance the Nation's scientific and technological literacy

What characterizes our activities and what are our functions

The Earth Sciences Directorate deals with large, sustained, multi-year projects that require significant collaborative efforts. Such activities are:

- Development, design, and implementation of new satellite missions and suborbital science campaigns
- Instrument algorithm development and data analysis
- Model development and data assimilation
- Distribution of geophysical and model data products

  Examples:

  - Atmospheric ozone
  - Land use and land cover change
  - Global precipitation
  - Ocean biology
  - Aeronet and aerosols data sets
  - International Satellite Cloud Climatology Project (ISCCP)

- Provide a resource for environmental assessment and policy decisions

Sensor Webs Could Link *in situ* and Remotely-Sensed Observations With Model Outputs & Federated Information Repositories

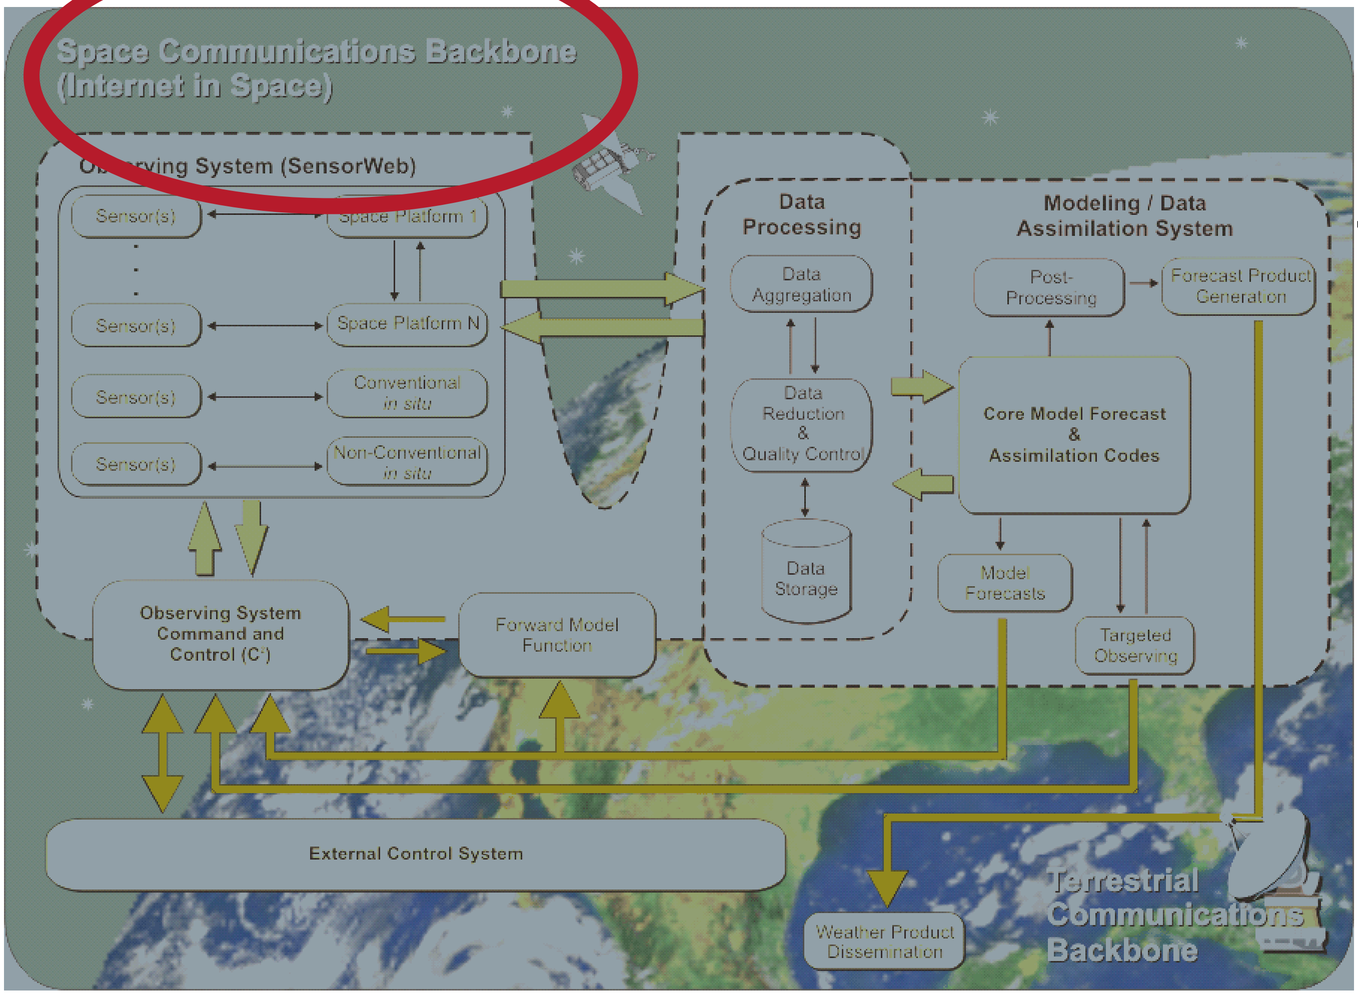

Sensor Webs Enable The Use Of Dynamic Targeting -- Potentially Reducing Error Growth and Improving Forecast Skill

Current GSFC Activities Are Focused on the Simulation of a Dynamic Sensor Web

Enhanced Sensor Web Architecture Enables Assets to Process Data and For The Assimilation System to Command the Observing System

InterPlaNetary Internet
Defining a New NASA Space Communications Architecture

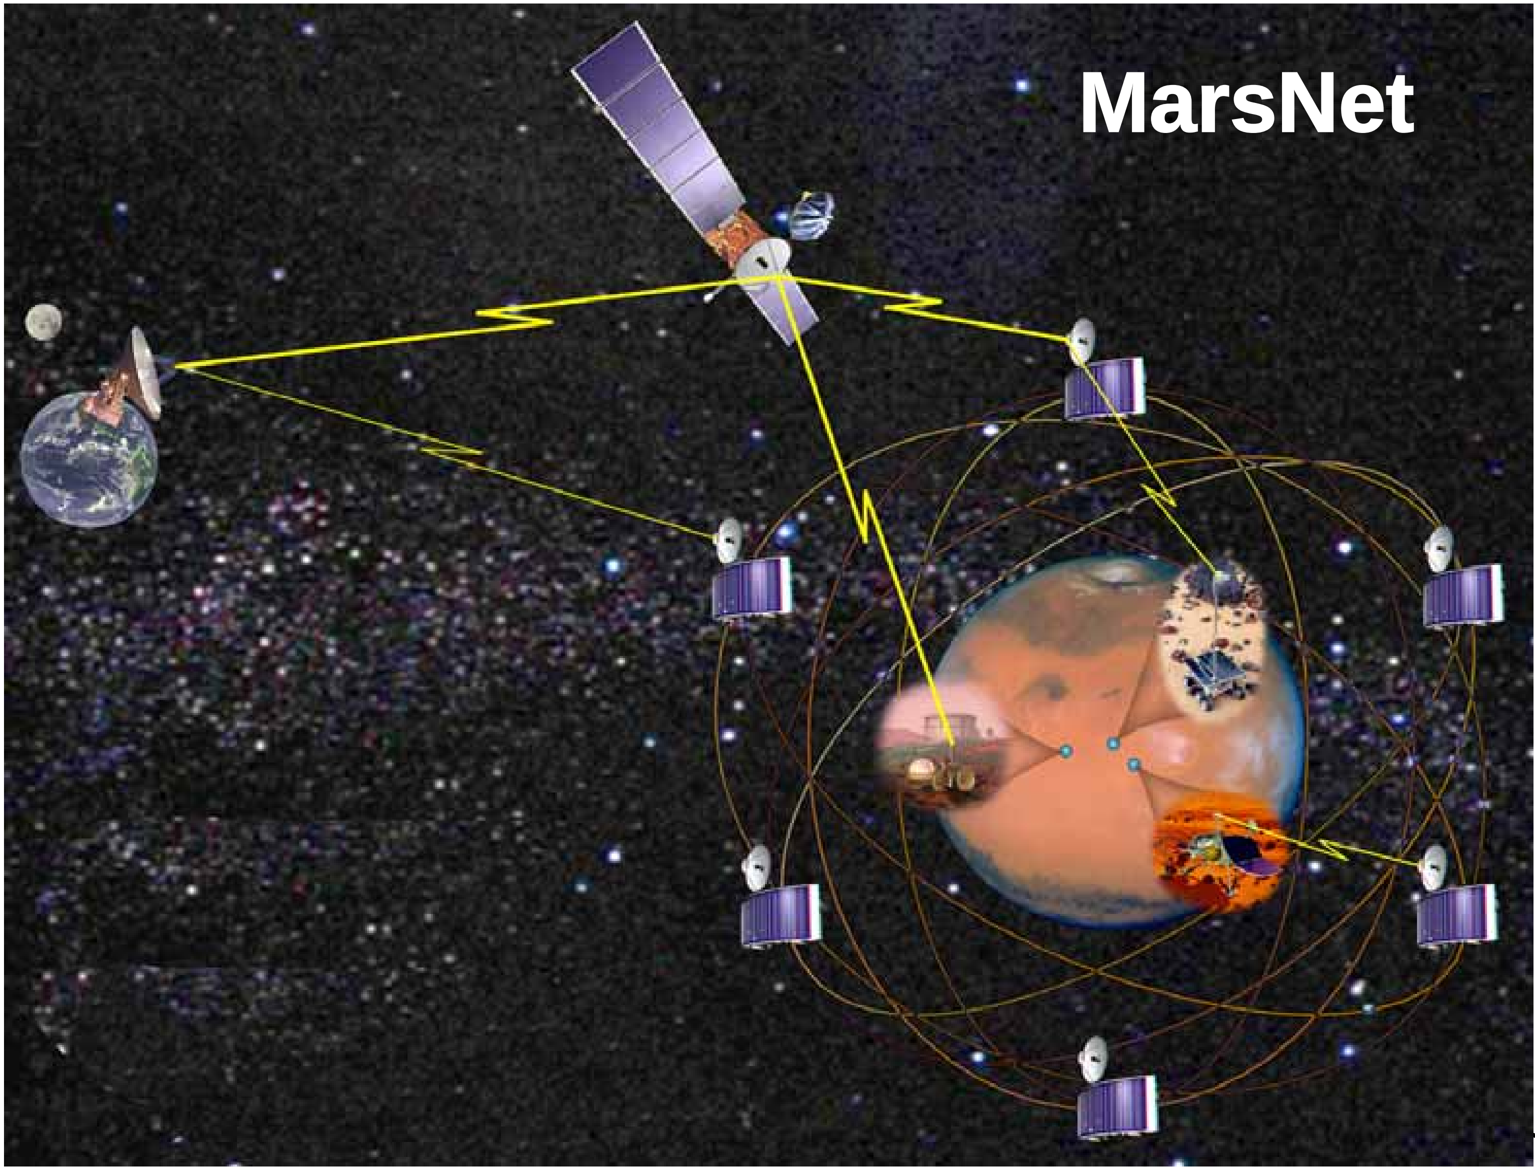

Source: JPL, Vint Cerf, MCI

Earth System Enterprise-Data Lives in Distributed Active Archive Centers (DAAC)

ASF (256 TB)
SAR Products
Sea Ice
Polar Processes

NSIDC (67 TB)
Cryosphere
Polar Processes

LPDAAC-EDC (1143 TB)
Land Processes
& Features

SEDAC (0.1 TB)
Human Interactions in
Global Change

GES DAAC-GSFC (1334 TB)
Upper Atmosphere
Atmospheric Dynamics, Ocean Color, Global Biosphere, Hydrology, Radiance Data

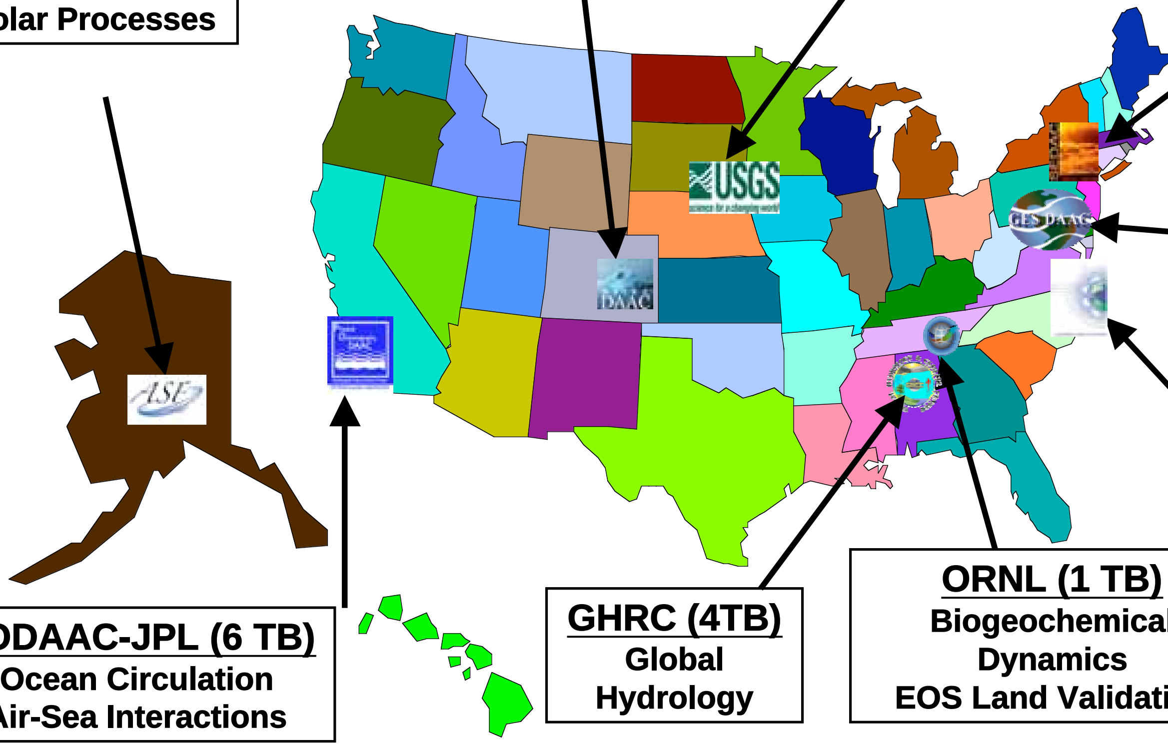

PODAAC-JPL (6 TB)
Ocean Circulation
Air-Sea Interactions

GHRC (4TB)
Global
Hydrology

ORNL (1 TB)
Biogeochemical
Dynamics
EOS Land Validation

ASDC-LaRC (340 TB)
Radiation Budget,Clouds
Aerosols, Tropospheric
Chemistry

EOS Aura Satellite Will Be Launched Soon
Challenge is How to Evolve to New Technologies

NASA Earth System Science IT Challenges

- EOSDIS Currently:
  - » Ingests Nearly 3 Terabytes of Data Each Day
  - » In 2003 it Delivered Over 25 Million Data Products
  - » In Response to Over 2.3 Million User Requests
  - » Making It the Largest "e-Science" System in the World
- This Capability Must Evolve To Handle Still Larger Data Volumes As Well As New Data Types (e.g. Laser-LIDAR Data)
- Earth System Modeling is a Driving Requirement for High-End Computing, and will Continue to be so as Models:
  - » Increase in Resolution and
  - » Are Further Coupled
    - – (e.g., Atmosphere-Ocean-Land Processes)

❑ *Other Agencies are Learning from EOSDIS and are Moving Beyond. As NASA Lays Out the Evolution of its Information Infrastructure to Meet its Earth Science Challenges Over The Next Decade, it will Again Need to Move to The Leading-Edge.*

Removing Barriers to Earth Observing & Simulation

- One Current Barrier: The Low Throughput of Today's Internet

- Even Though Internet2 Backbone is 10 Giga bits per second
  - » Network is Shared Using TCP/IP Protocol
- A Remote NASA Earth Observation System User Only Sees:
  - » 10-50 Mbps (May 2003) Throughput to Campuses
    - – Typically Over Abilene From Goddard, Langley, or EROS
  - » Best FTP with Direct Fiber OC-12: Goddard to UMaryland
    - – 123 Mbps
  - » UCSD's SIO to Goddard (ICESAT, CERES Satellite Data)
    - – 12.4 Mbps—1/1000 of the Available Backbone Speed!

http://www.evl.uic.edu/cavern/rg/20030817_he

Additional Factors Affecting Throughput Performance (very partial list)

…Other than Layers 1 (Physical) and 2 (Data Link) Network Infrastructure at Core or Edge

- Layer 7/Application: Many with non-optimal I/O designs
- Layer 6/Presentation: Huge reformatting requirements
- Layer 5/Session:
- Layer 4/Transport: Standard TCP with multiple parallel streams
  TCP-Mods: FAST, XCP, HSTCP
  TCP-Alternates: TSUNAMI, SABUL
- Layer 3/Network: IPv4 best effort vs. with DiFFServ or MPLS
  IPv6 with per-user-flow QoS features
- NIC's, I/O bus, and CPU capabilities

Agency High End Networking

Motivation

- » NASA has fallen significantly behind the state of the art in advanced networks as indicated in figure 1
- » With the introduction of NASA's newest supercomputer the lack of bandwidth is a significant barrier to collaboration and data sharing-2TByte per day data set cannot be effectively transferred between research teams
- » Ames in conjunction with JPL and GSFC has completed a study on options for solving the problem
- » Eventually the agency must solve this for all the centers and a preliminary analysis has been completed for the agency

ECCO Ocean Modeling

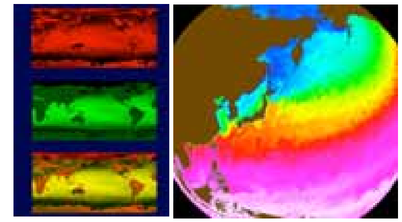

Run Requirements: (Ames – JPL)

–Nov 2003 = 340 GBytes / day

–Feb 2004 = 2000 GBytes /day

Conclusion

– Not enough bandwidth for distributed data intensive applications

– Opportunities exists to work with emerging NLR high bandwidth systems but Agency Infrastructure will not support this

Approach

- – Ames High End Computers have been upgraded to 10Gbps capability
- – Consortium formed and negotiations underway to extend Dark Fiber to Ames Site from local POP
- – Cenic/National Lambda Rail NLR membership investigated-budget and plan developed. Operating costs after upgrade (x50) anticipated to be the same as current OC3-12 charges
- – Design and Estimates for Router and Switch upgrades completed.
- – Estimated costs at both ARC and GSFC $1M-$1.5M. Team working to identify funding.
- – ISSUES-Short term-no natural owner of this problem in previous years HPCC or CICT program attempted to solve, Long Term-Maintaining balanced Computing, Network and Storage systems requires capital upgrades to the agency research networks

Research Network Capacity

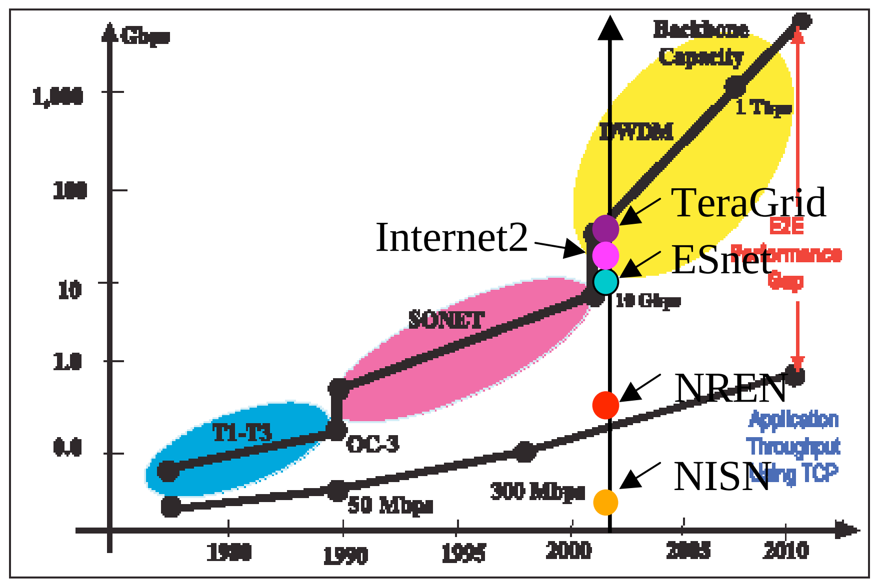

Figure 1

DoE Network Challenge, 2000

Technology Emphasis Areas

Earth System Science in the future will leverage three ongoing technology revolutions:

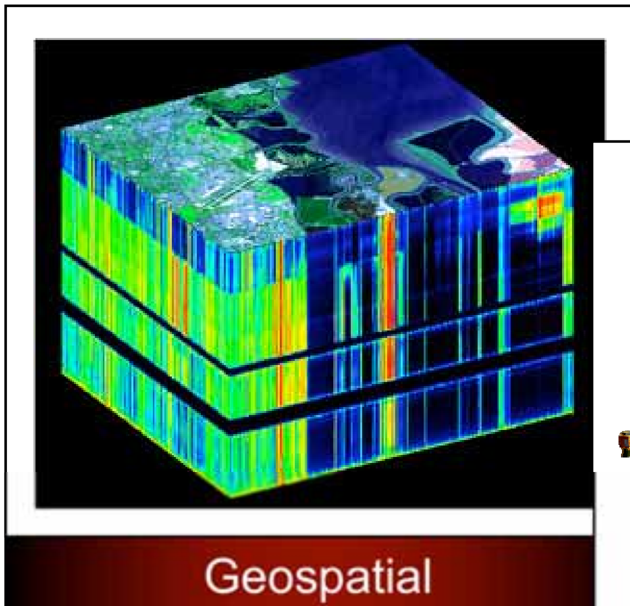

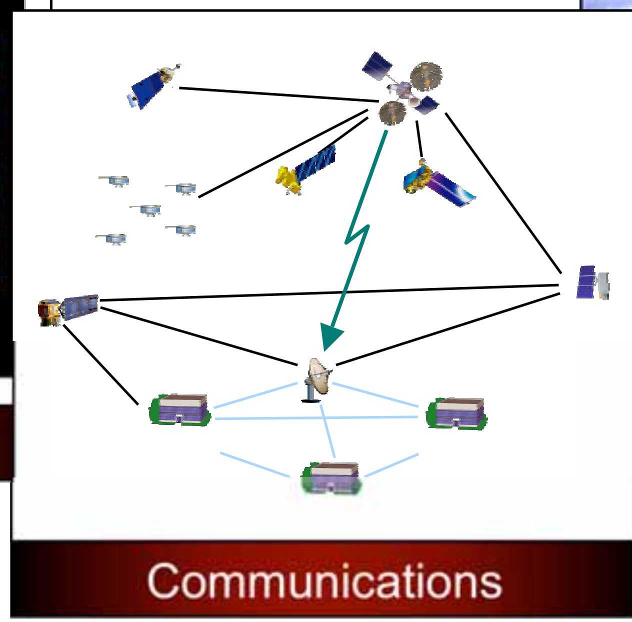

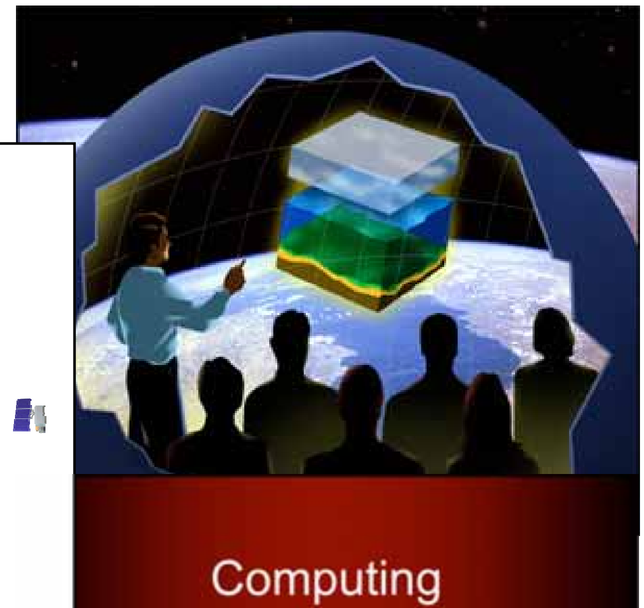

...to enable timely and affordable delivery of Earth Science data and information to users

Difficulties (Murphy’s Law)

Technology doesn’t always work exactly as you hope it will!

Data System Technology Evolution Cycle

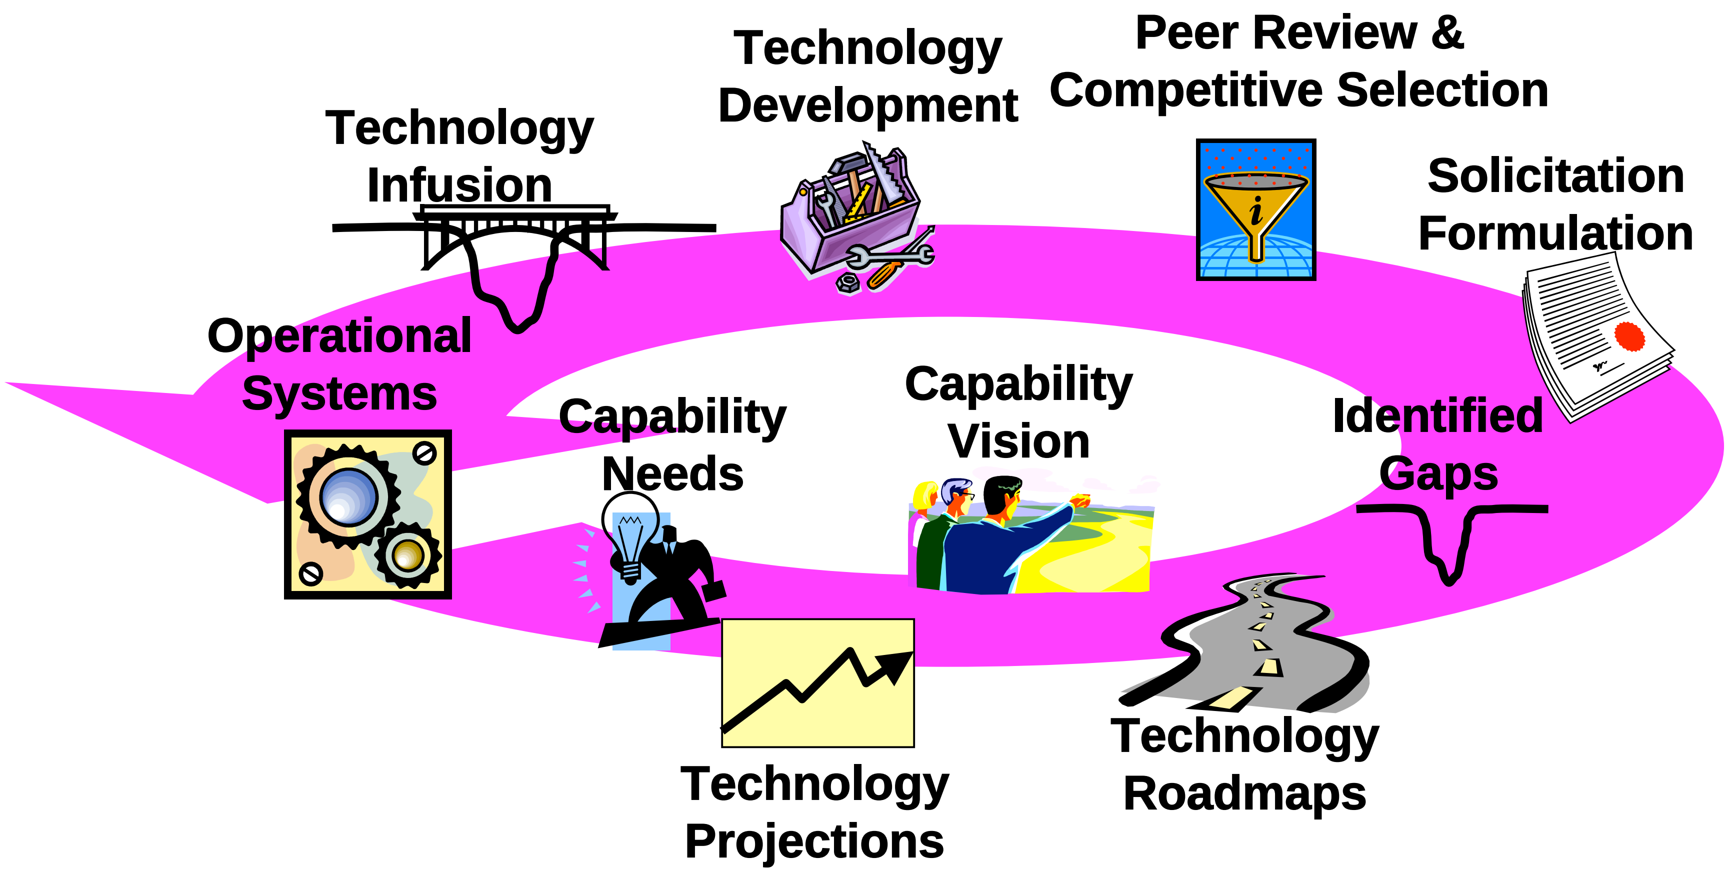

NASA ESE Technology Planning Process

Utilizes & Extends Current ESTO Processes for ESE Data Systems

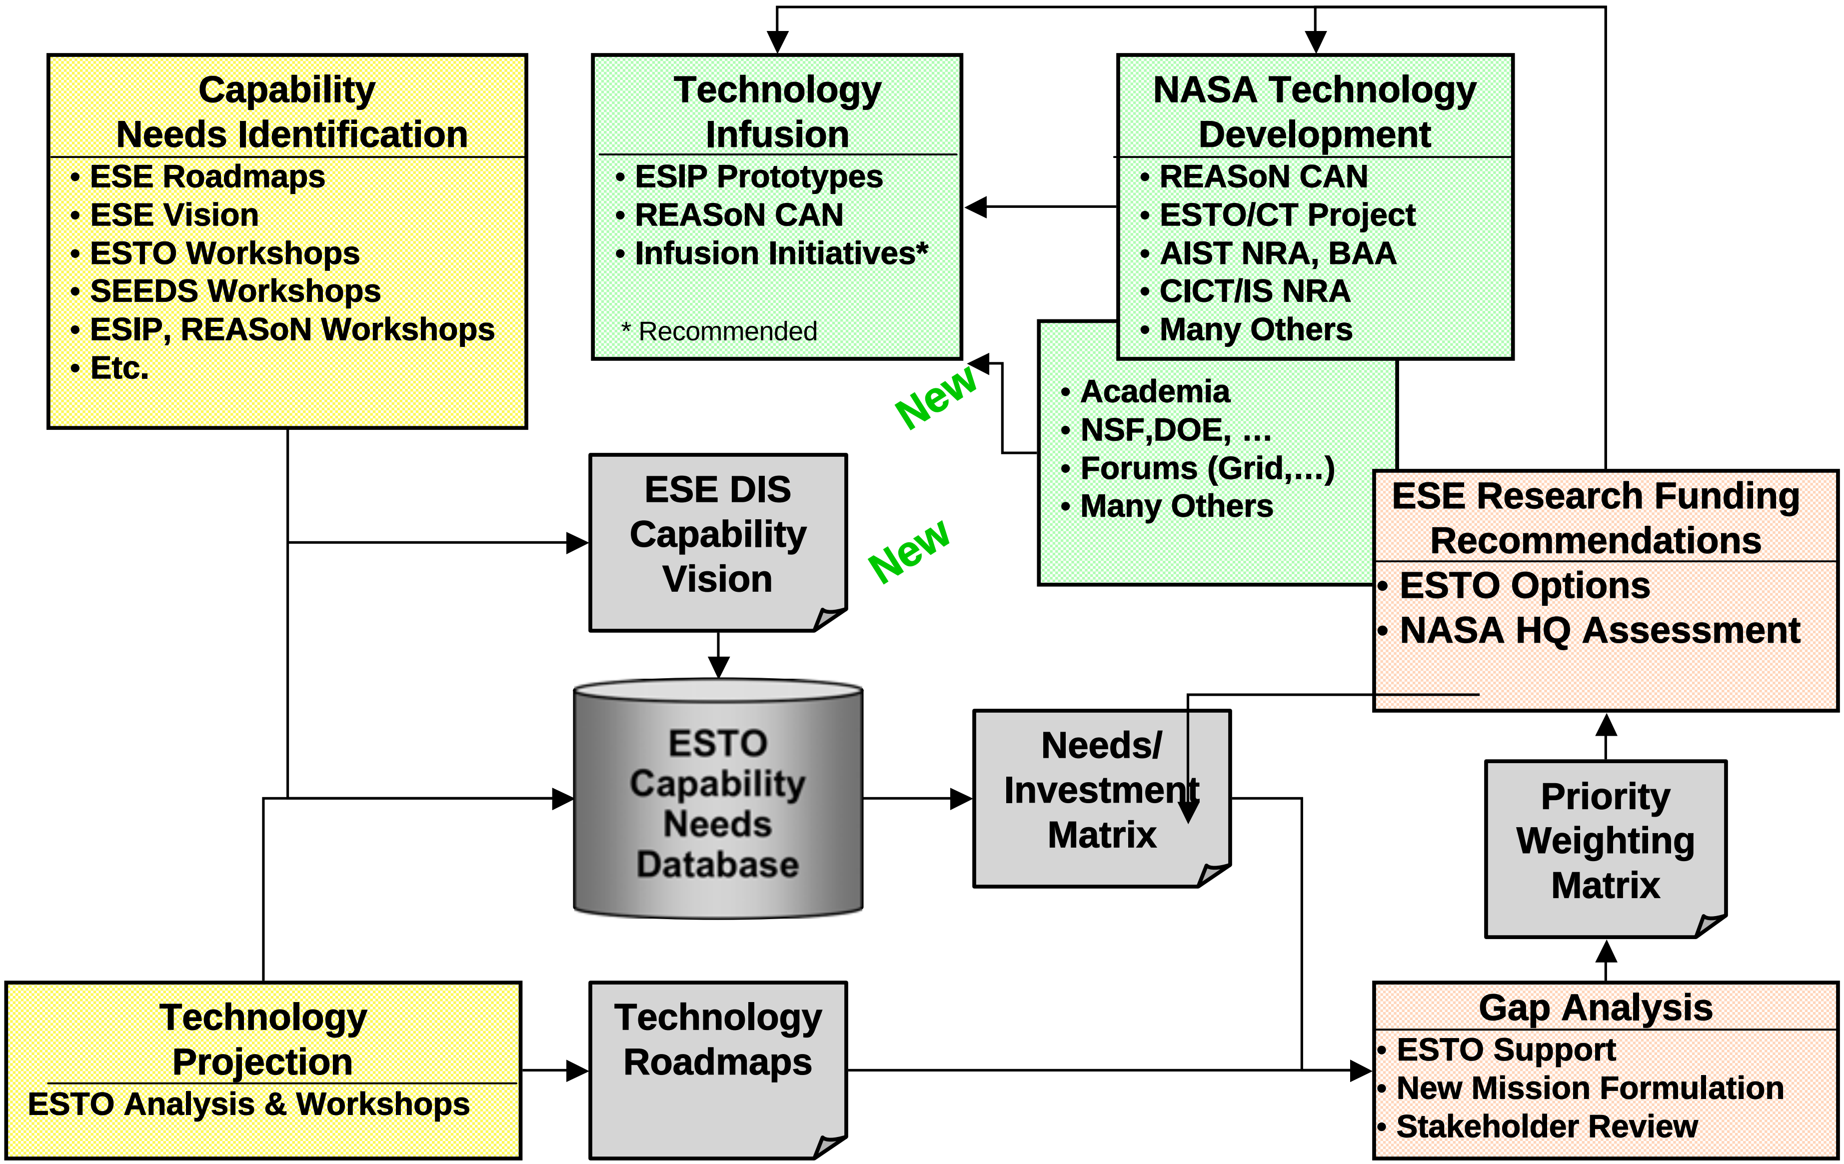

NASA is Developing Grid Technologies to Enable Users to Easily Fuse Distributed Data

Committee on Earth Observation Satellites (CEOS) Grid Testbed

Goals And Objectives:

•Establish A Data Grid Between USGS EDC, NASA Goddard ADG And IPG.

Achievements:

• Phase One, A Simulated Data Fusion Algorithm Is Initiated at a Scientist's Workstation With Processing Taken Place at Another Site

•Data Will Reside In Two Locations And Be Moved Using Grid Technology

Future Plans:

•Phase Two-- Use Real Data Fusion Algorithms Using The Grid To Demonstrate Distributed Processing of Data Sets and to Experiment With Grid Workflow Techniques

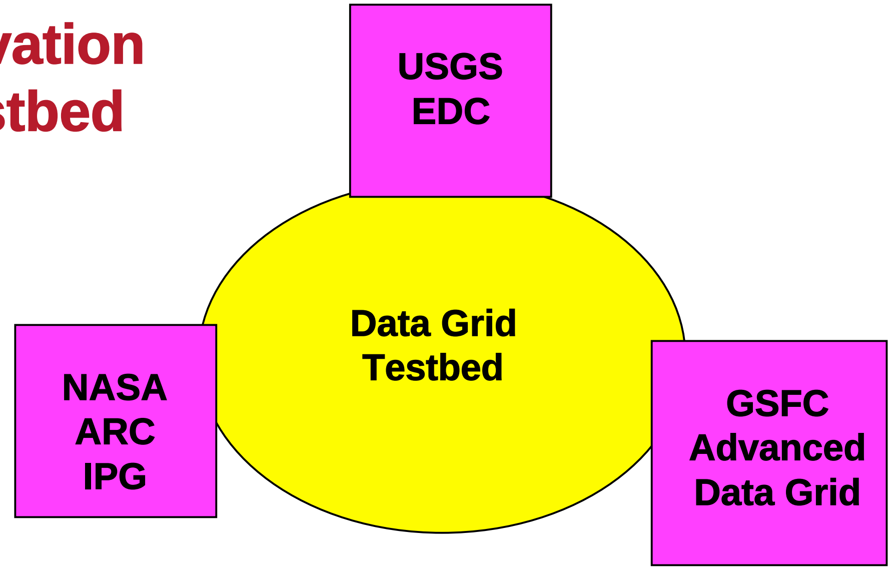

Participants: CEOS Member sites

- EOSDIS & George Mason University (GMU)
- European Space Agency (ESA)
- DutchSpace
- NOAA Operational Model Archive & Distribution System (NOMADS)
- University of Alabama – Huntsville (UHA)
- United States Geological Survey – EROS Data Center (EDC)
- NASA Advanced Data Grid (ADG)
- China Spatial Information Grid (SIG)
- ARC

GSFC Networks

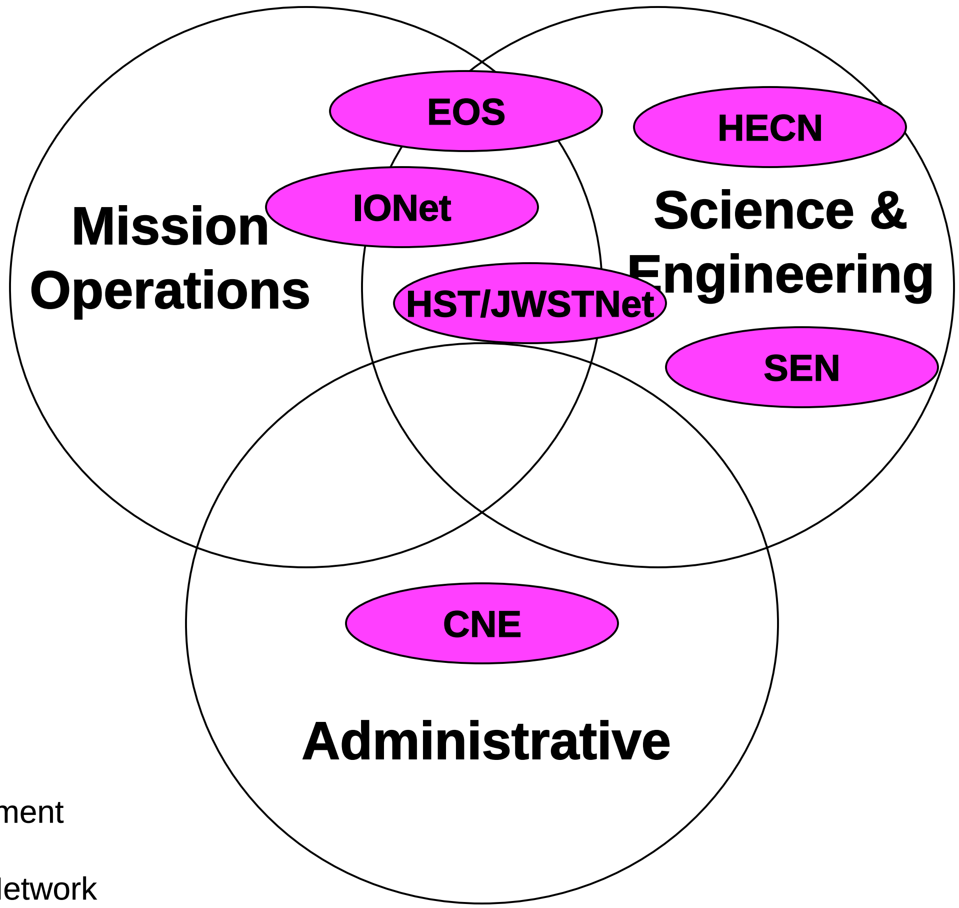

CNE: Center Network Environment
EOS: Earth Observing System
HECN: High End Computing Network
HST/JWSTNet:: Hubble Space Telescope/James Webb Space Telescope Network
IONet: IP Operational Network
SEN: Science & Engineering Network

Notional Key Characteristics of GSFC's Scientific and Engineering Network (SEN) and High End Computer Network (HECN)

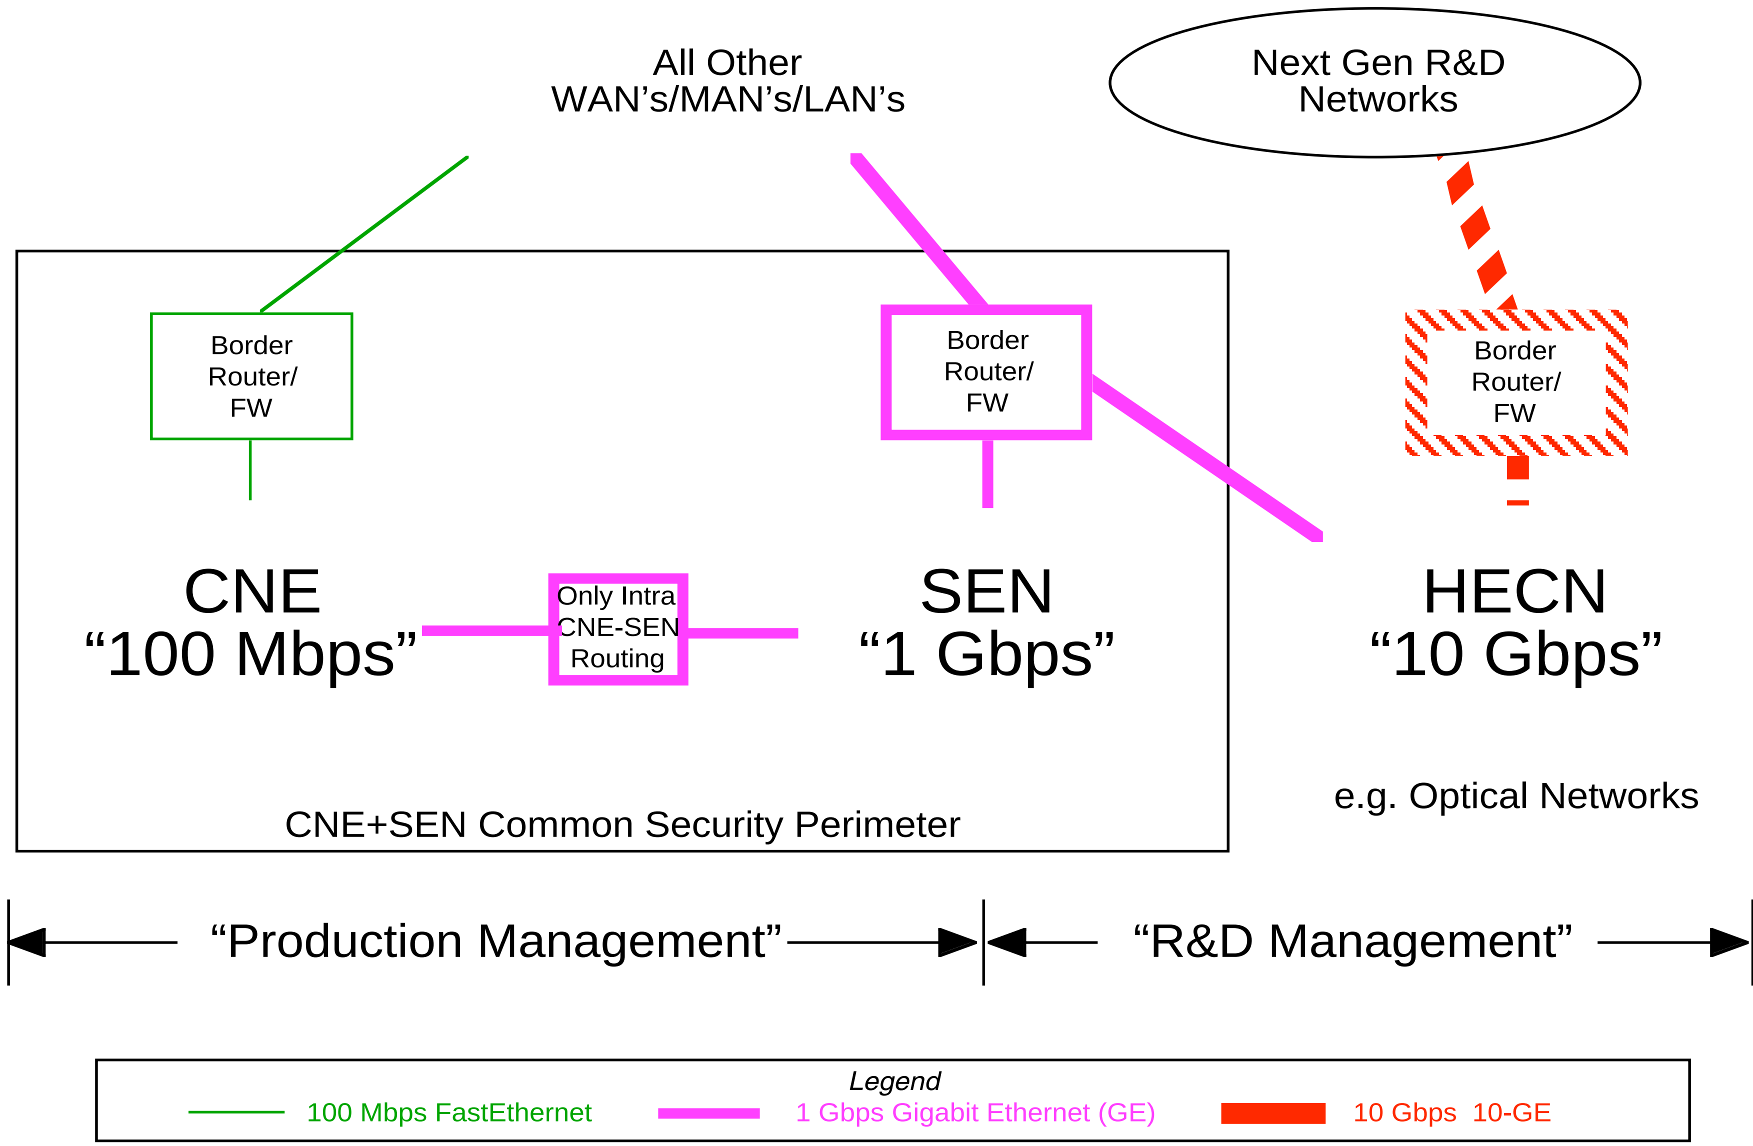

JPG 02/25/04

Schematic of Gbps e-VLBI Demonstration Experiment

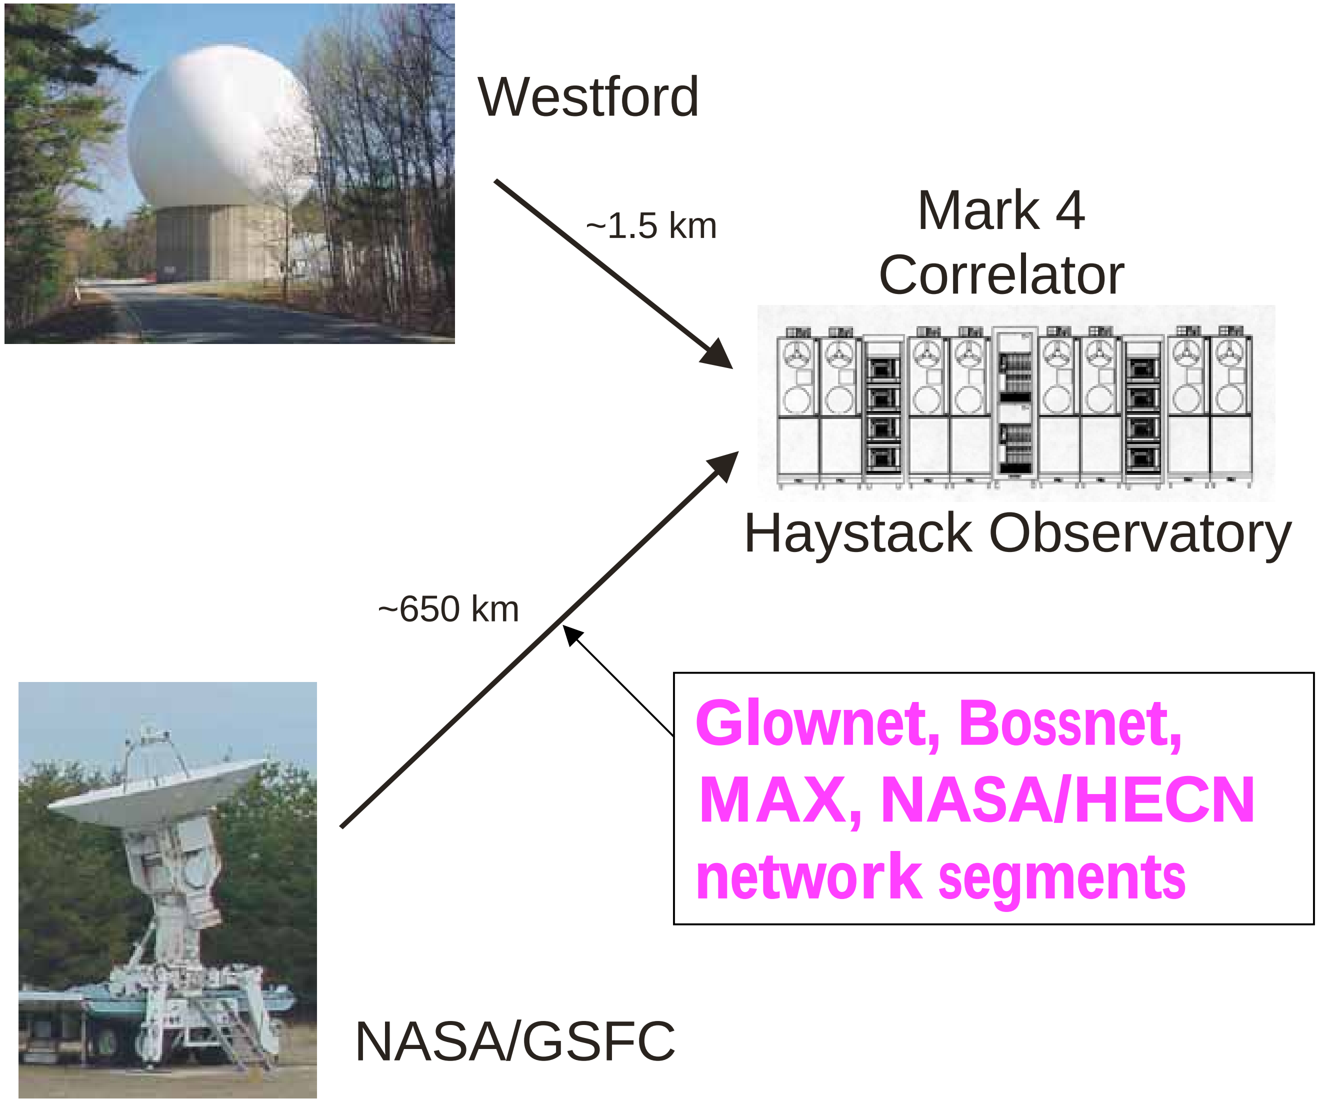

e-VLBI GGAO-Haystack Data Rates Sustained During a 16-Hour-Long Evaluation Test

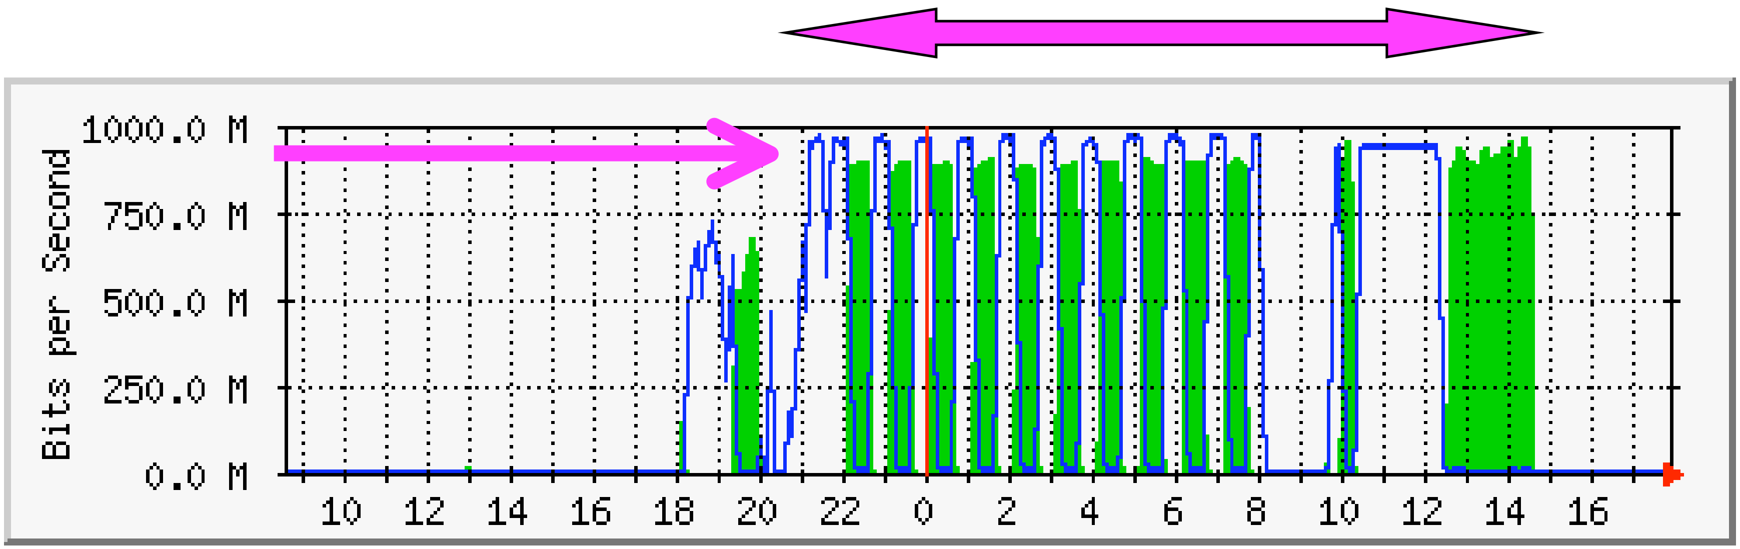

Max In:970.5 Mb/s (97.1%) Average In:210.8 Mb/s (21.1%) Current In:168.0 b/s (0.0%)
Max Out:978.1 Mb/s (97.8%) Average Out:263.6 Mb/s (26.4%) Current Out:216.0 b/s (0.0%)

GSFC SAN Pilot

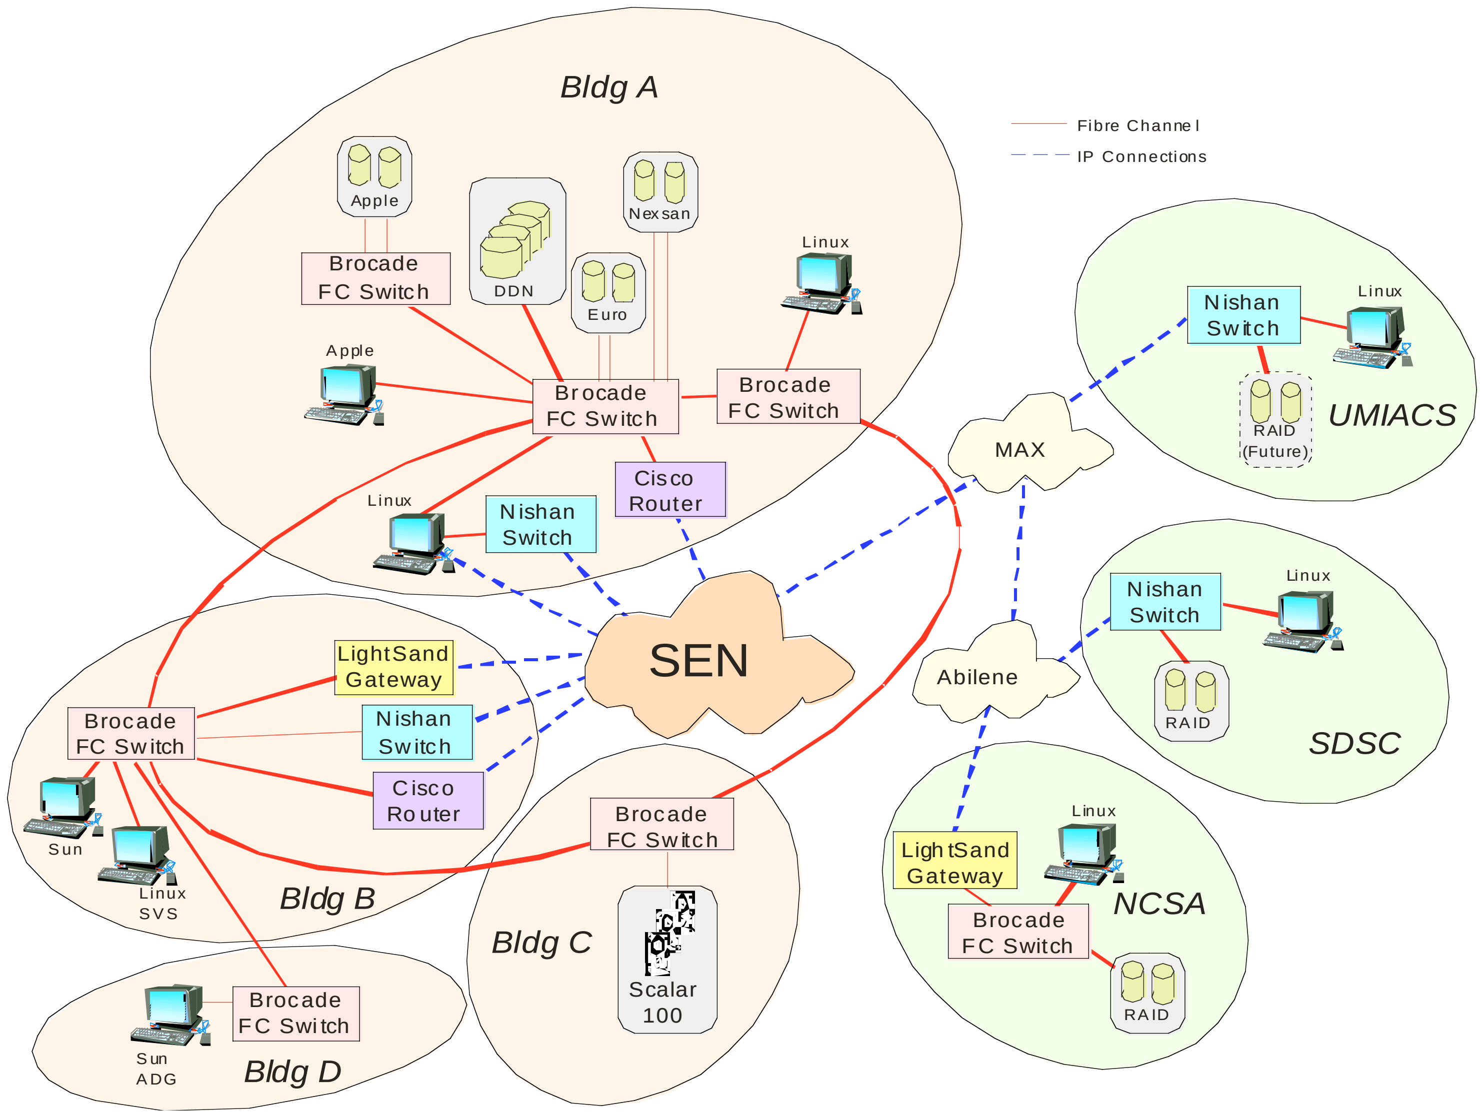

J. P. Gary
3/23/04

DRAGON - Complete Network by Year 3

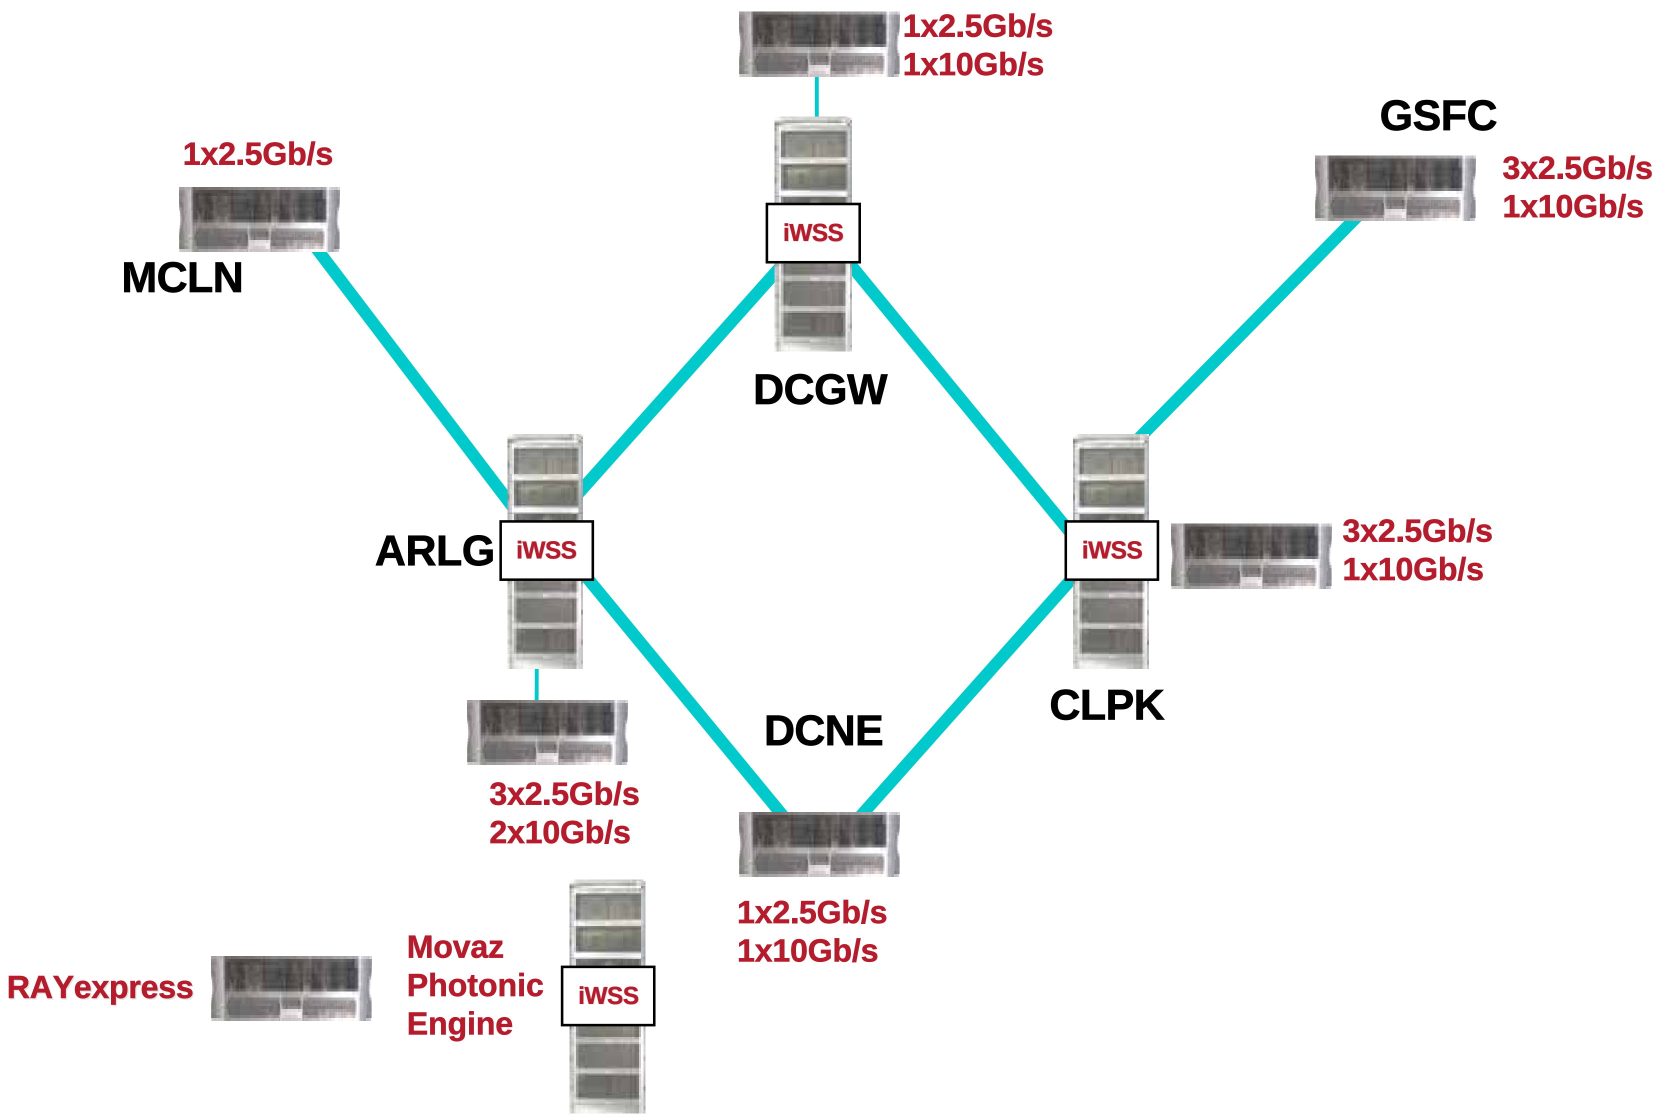

Considerations for Transcontinental Backbone Network

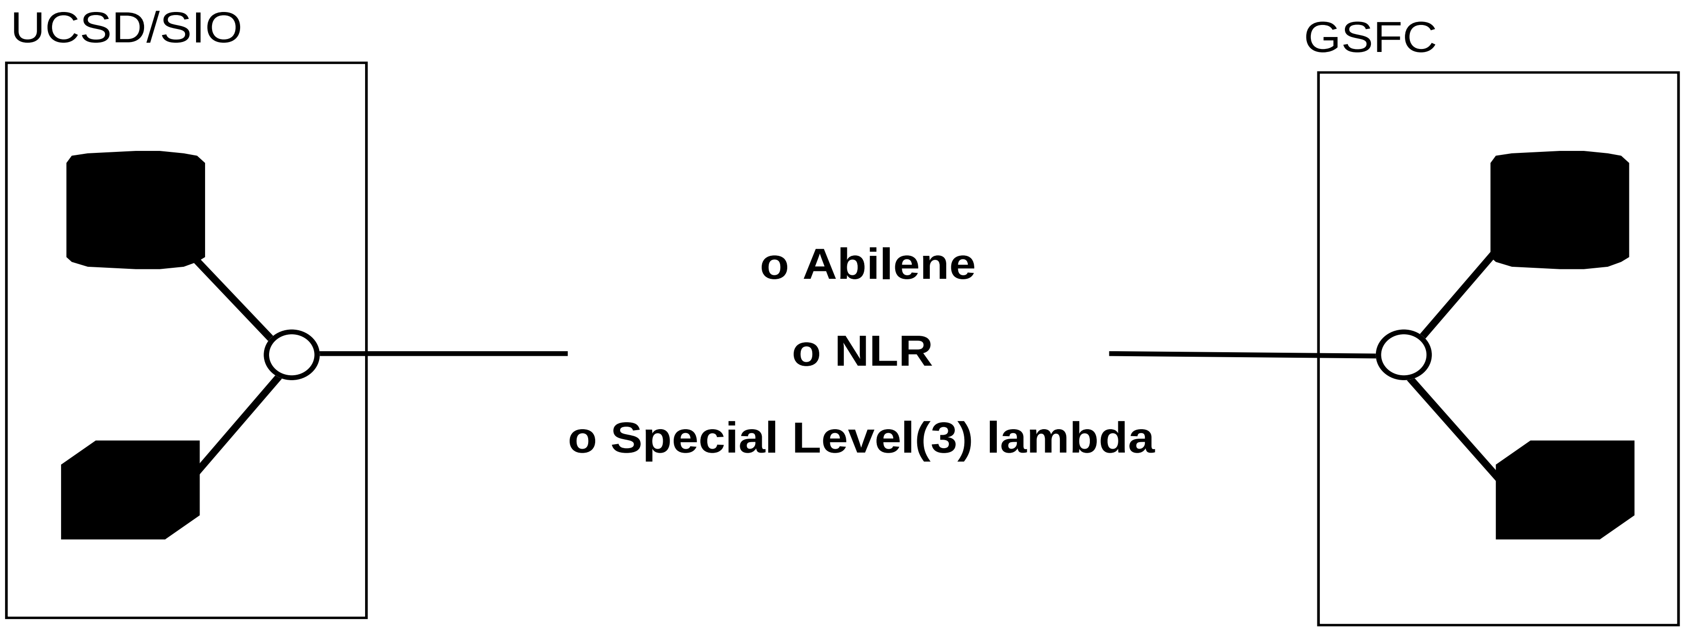

completed connections:
238 participants
51 connectors + 3 NGIXs + STAR TAP, StarLight, AMPATH, PacificWave ManLan IXs
48 connections to 32 peer networks

The Abilene Network

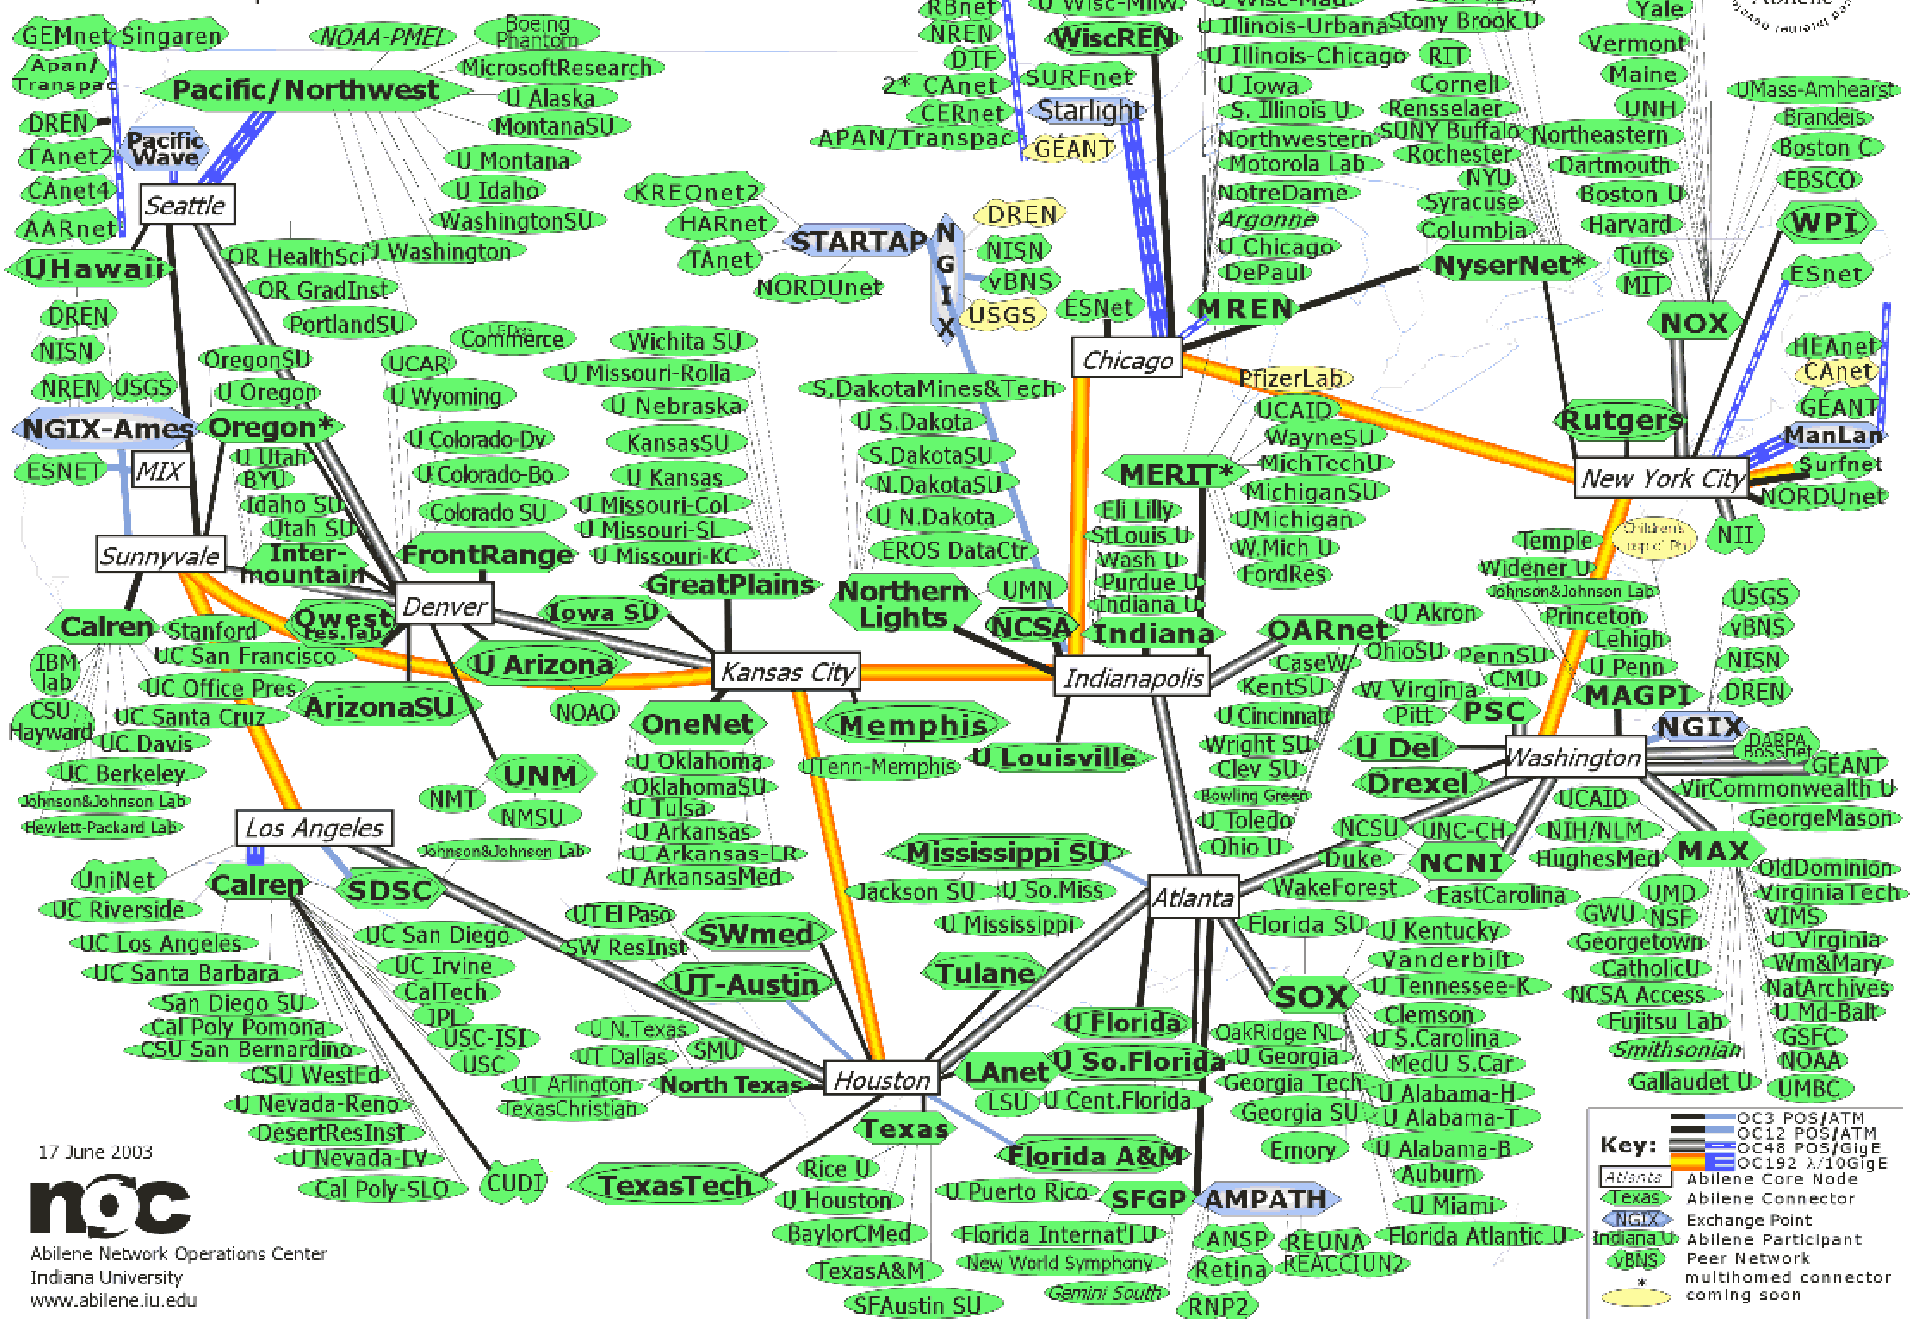

17 June 2003

Abilene Network Operations Center
Indiana University
www.abilene.iu.edu

National LambdaRail (http://www.nationallambdarail.org/)

- Provide an enabling network infrastructure for new forms and methods for research in science, engineering, health care, and education as well as for research and development of new Internet technologies, protocols, applications and services.
- Provide the research community with direct control over a nationwide optical fiber infrastructure, enabling a wide range of facilities, capabilities and services in support of both application level and networking level experiments and serving diverse communities of computational scientists, distributed systems researchers and networking researchers.

Network-based Limitations of Abilene Removed with NLR

- Applications traffic must be IP-based
- 1 GE present limits at access POP's
- Shared 10 GE backbone
- Typically 13 store-and-forward router hops between GSFC and SIO; ~75 msec RTT
- Private addresses of UCSD's OptIPuter not advertised via Abilene

R&D Test: Move to Internet Protocol Over Dedicated Optical Lightpaths

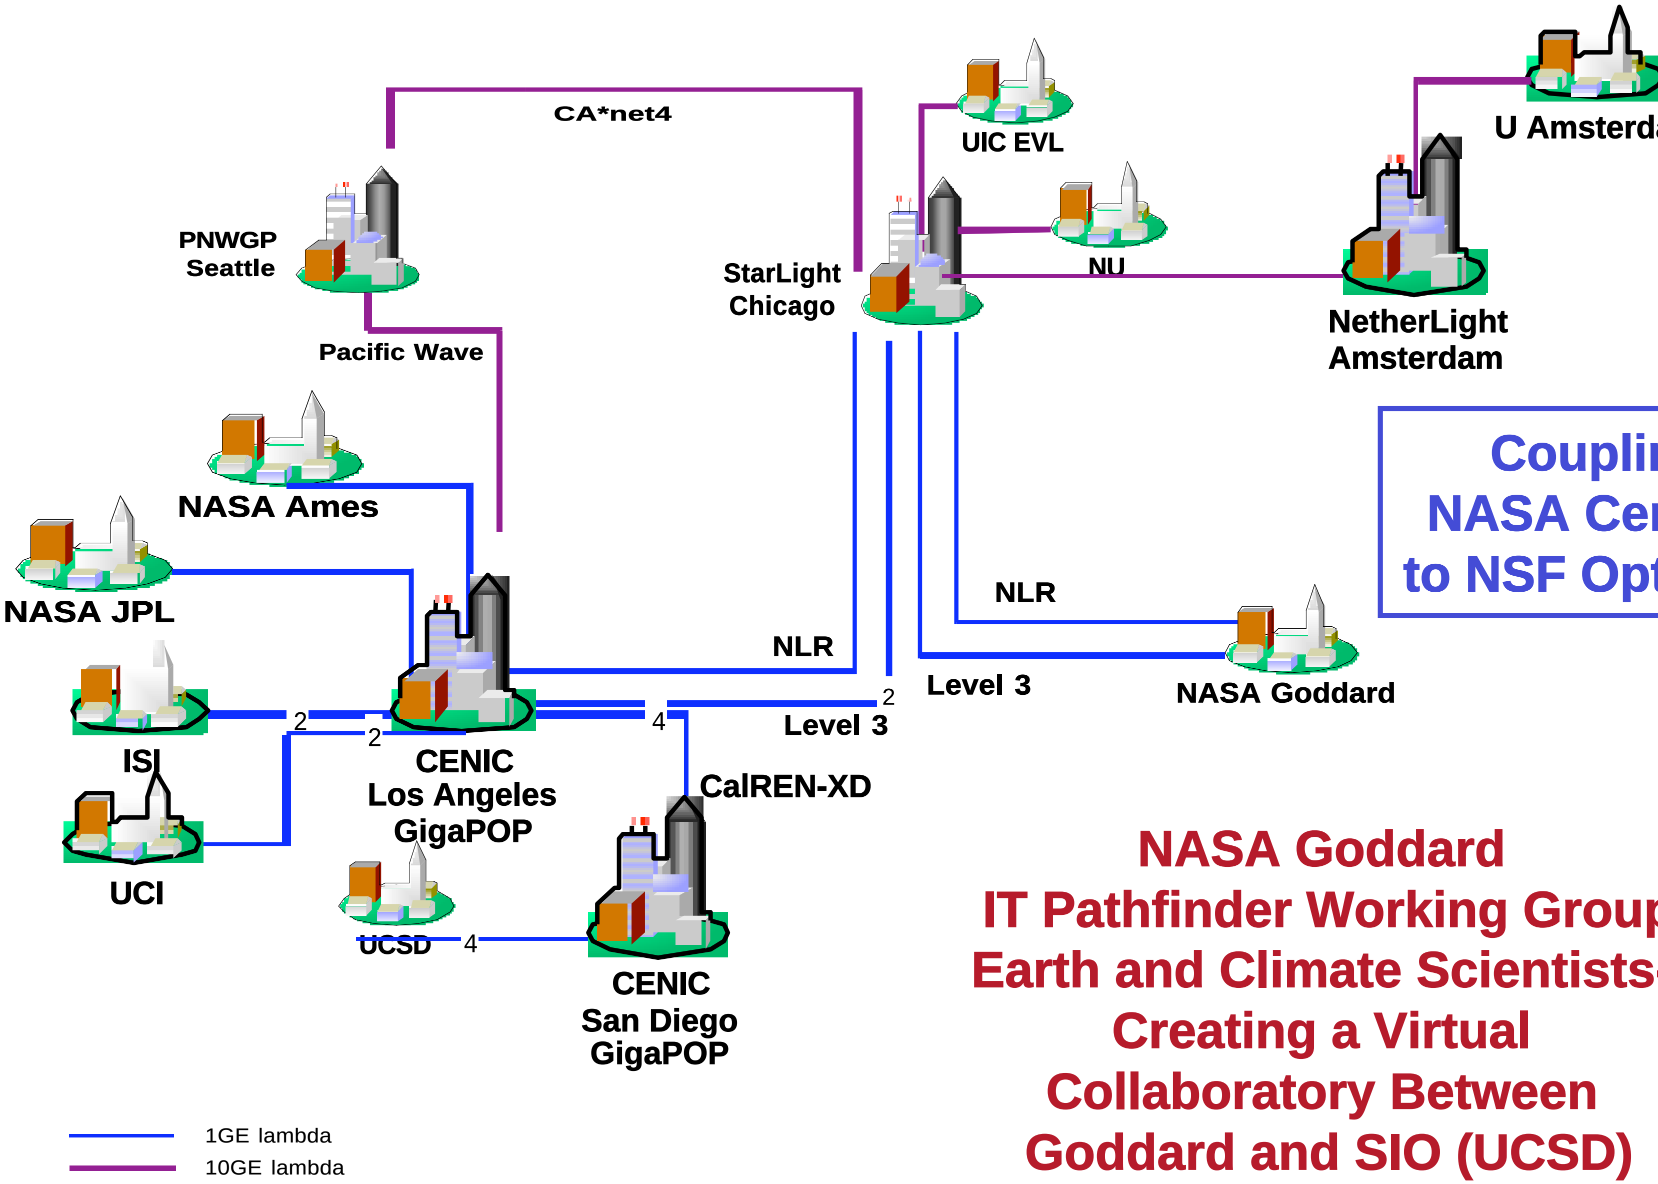

Coupling NASA Centers to NSF OptIPuter

NASA Goddard IT Pathfinder Working Group Earth and Climate Scientists- Creating a Virtual Collaboratory Between Goddard and SIO (UCSD)

J. P. Gary
3/23/04

Force10 E300 10 GE switch/router being readied…

For use in or as:

- Test upgrade for SEN's inter- and intra-building GE switch infrastructure
- Multiple 1-GE up/downlink multiplexor between Beowulf clusters
- Switch/router-host for testing 10 GE NIC-based host connections
- Switch/router-host for testing 10 Gbps-capable firewall
- Test upgrade for SEN's link with MAX/Abilene
- GSFC CPE connection for proposed 10 Gbps Lambda Network connection with UCSD/Scripps

10 GE Testing by Bill Fink (GSFC) & Paul Lang (ADNET)