GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Outline

- **Current GSFC "Infrastructure" Using DRAGON**
- **Previous and/or On-Going Applications**
- **Upcoming and/or Future Applications**
- **Interest in Additional Collaborations**

J. Patrick (Pat) Gary
Network Projects Leader
Networks and Information Technology Security Group/
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center

August 30, 2006
For First DRAGON User Group (DUG-1) Meeting

GSFC High End Computer Network (HECN) Project's Research Partners and Collaborators

- *DRAGON Project: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/WebHome*
  - *PI: Jerry Sobieski (UMCP)*
  - *GSFC L-Net on DRAGON network diagram: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/Network*
- *e-VLBI Project: http://web.haystack.mit.edu/e-vlbi/evlbi.html*
  - *PI: Alan Whitney (MIT/Haystack)*
  - *GSFC L-Net on e-VLBI network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SC04_eVLBI_network.pdf*
- *GLIF: http://www.glif.is/*
  - *Chair: Kees Neggers (SURFnet)*
  - *GLIF network diagrams: http://www.glif.is/publications/#maps*
- *NGC IT Sector: http://www.it.northropgrumman.com/index.html*
  - *PI: Brice Womack (NGC)*
  - *GSFC L-Net on NGC IT Sector Colshire network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/DRAGON_NGC_030606.pdf*
- *NLR: http://www.nlr.net/*
  - *CEO: Tom West (NLR)*
  - *NLR network diagram: http://www.nlr.net/infrastructure/*
- *NREN Project: http://www.nren.nasa.gov/*
  - *PM: Ken Freeman (ARC)*
  - *GSFC L-Net/SEN on NREN network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/CENIC2006_13_mfoster_excerpts.pdf*
- *OptIPuter Project: http://www.optiputer.net/*
  - *PI: Larry Smarr (UCSD)*
  - *GSFC L-Net on OptIPuter network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SMARR-OptIPuter-AHM-gold.pdf*
- *TeraFlow Testbed Project: http://www.teraflowtestbed.net/*
  - *PI: Robert Grossman (UIC)*
  - *GSFC L-Net on TeraFlow Testbed network diagram: http://www.ncdm.uic.edu/maps/index.jpeg*

GSFC L-Net Configurations at McLean and Greenbelt

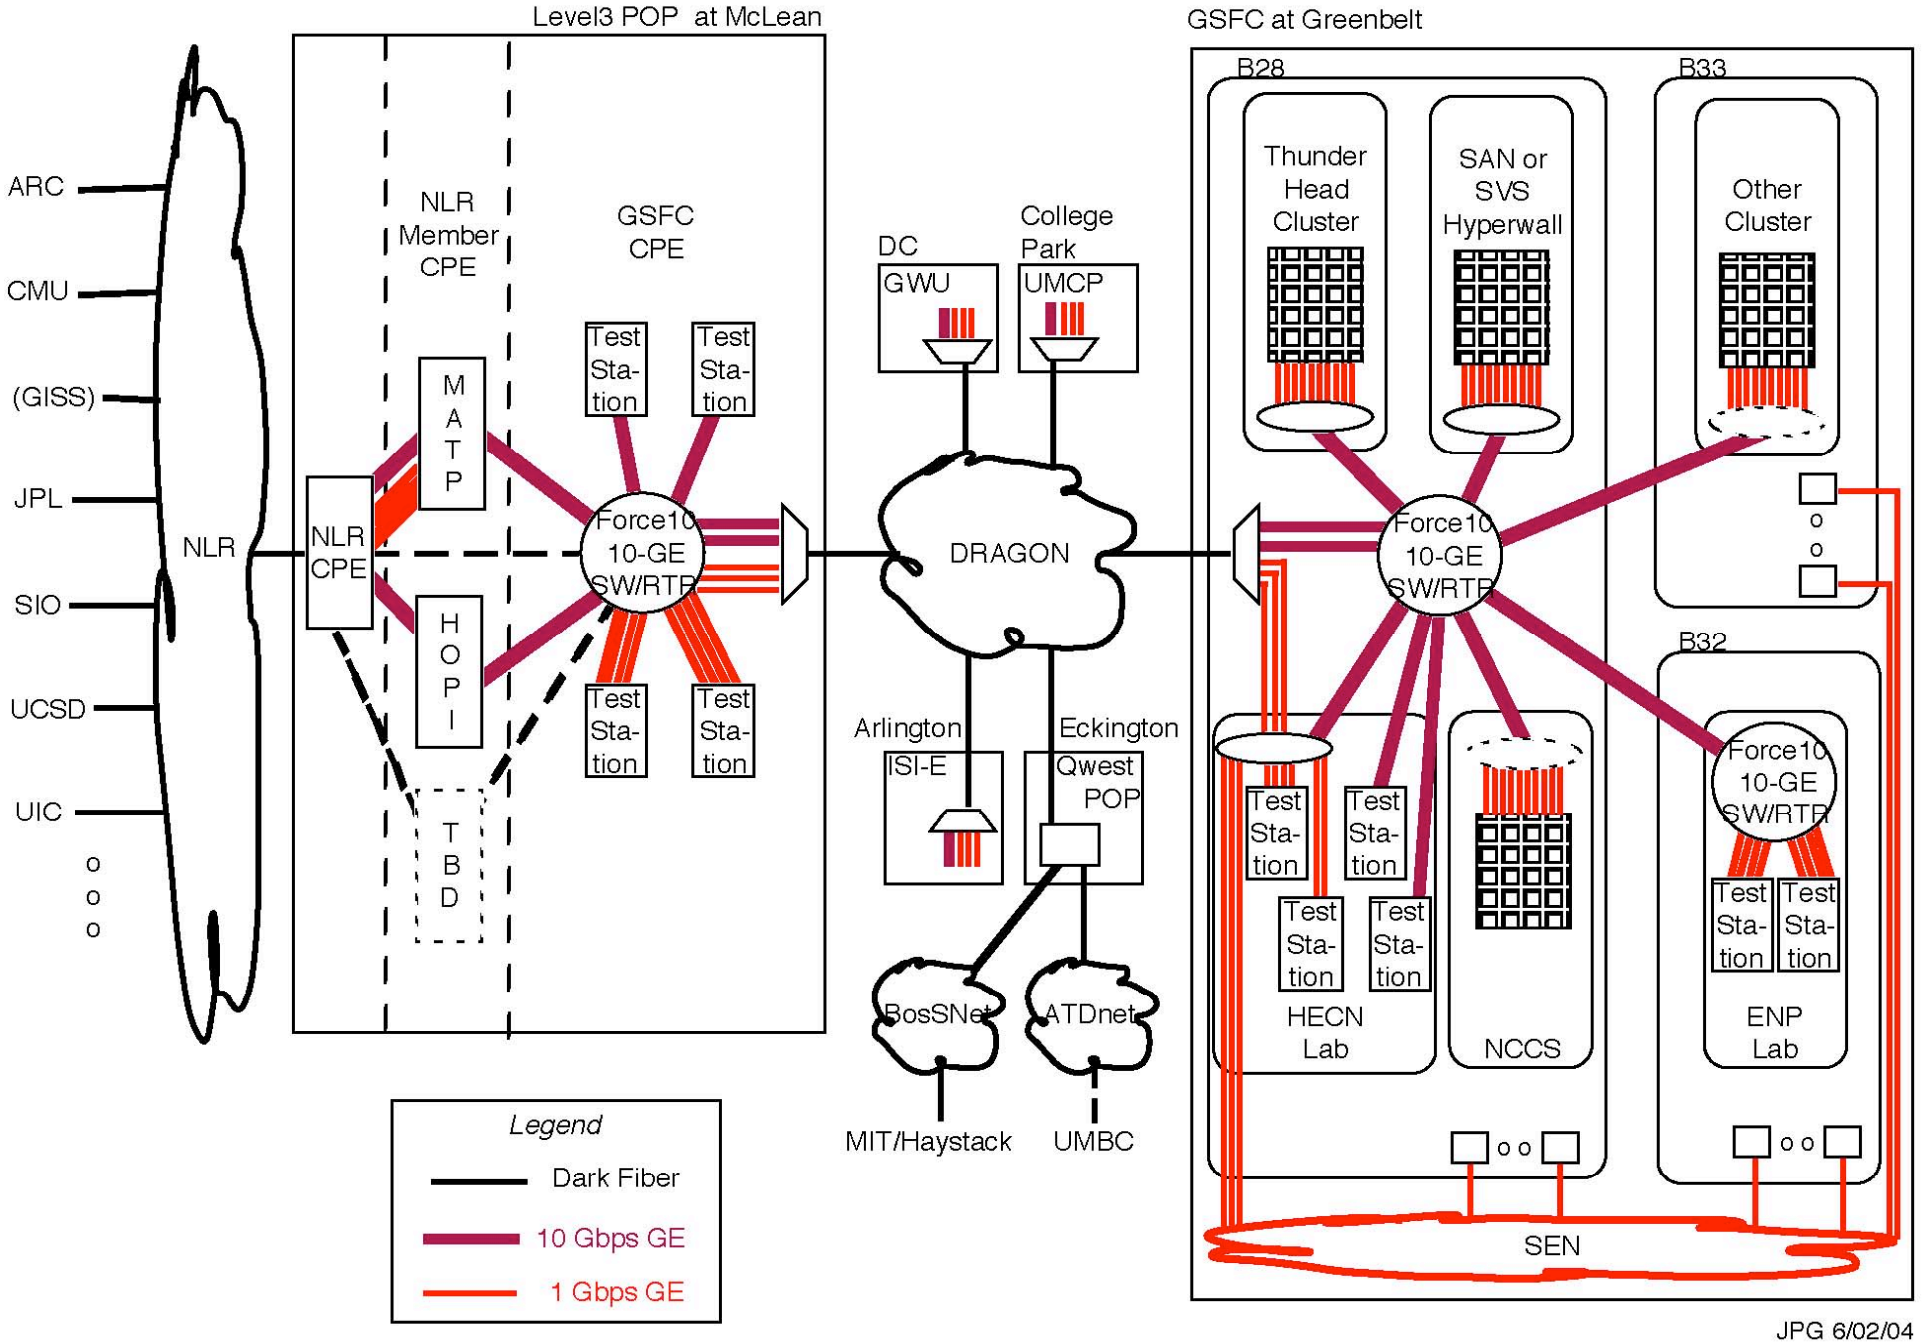

JPG 6/02/04

GSFC-Enabled High Performance Computer (HPC) Systems Using And/Or Accessible for Test Across DRAGON (1 of 4)

Project-Managed HPC Systems at MCLN

(All with own 10-GE connection to HECN's Force10 E300 10-GE switch/router)

HECN Project (PI: Pat Gary (GSFC))

- One Force10 E300 10-GE switch/router, w/eight "1-to-1 non-oversubscribed" 10-GE ports, eight "4-to-1 oversubscribed" 10-GE ports, and twenty-four 1-GE ports
  - Allocation of "1-to-1 non-oversubscribed" 10-GE ports
    - NLR-WASH-STAR-10GE-22 (Owner: Pat Gary courtesy of Tom West/NLR)
    - NLR-STAR-WASH-10GE-103 (Owner: Tom DeFanti/OptIPuter)
    - One connection with MCLN-based HECN/DRAGON Movaz-channel49 (to GSFC)
    - One connection with MCLN-based NREN Cisco 6506 (and indirectly its 10G NLR and DRAGON Movaz-channel26 (to GSFC) lambdas)
    - One connection with MCLN-based DRAGON Raptor ER-1010 (and indirectly its 10G NLR/HOPI and DRAGON Movaz-channel33 lambdas)
    - One connection with Colshire-based NorthropGrumman Cisco 6513
    - Two connections with MCLN-based HECN Intel Pentiums (see below)
  - Allocation of "4-to-1 oversubscribed" 10-GE ports
    - Eight connections TBD
- Two Intel Pentium 4 (dual 3.06 GHz Xeon) w/ PCI-X, each w/one 10-GE NIC (Intel PRO/10GbE LR Server Adapter) and one quad-GE NIC (Intel PRO/1000 MT Quad Port Server Adapter)

GSFC-Enabled High Performance Computer (HPC) Systems Using And/Or Accessible for Test Across DRAGON (2 of 4)

Project-Managed HPC Systems at GSFC

(All with own 10-GE connection to HECN's Force10 E600 10-GE switch/router)

HECN Project (PI: Pat Gary (GSFC))

- One Force10 E600 10-GE switch/router, w/twelve "1-to-1 non-oversubscribed" 10-GE ports and sixteen "4-to-1 oversubscribed" 10-GE ports
  - Allocation of "1-to-1 non-oversubscribed" 10-GE ports
    - One connection with GSFC-based HECN/DRAGON Movaz-channel49 (to MCLN)
    - One connection with GSFC-based NREN Cisco 6506 (and indirectly its DRAGON Movaz-channel26 (to MCLN) and 10G NLR lambdas)
    - One connection with GSFC-based DRAGON Movaz-channel33 lambda
    - Nine connections with other GSFC-based HPC systems (see below)
  - Allocation of "4-to-1 oversubscribed" 10-GE ports
    - Four connections with other GSFC-based HPC systems (see below)
    - Twelve connections TBD
- Two AMD Opteron (dual 2.8 GHz) w/ PCI-Express, each w/two 10-GE NIC (Myricom Myri-10G)
- Two AMD Opteron (dual 2.8 GHz) w/ PCI-Express, each w/one 10-GE NIC (Myricom Myri-10G)
- Two Intel Pentium 4 (dual 3.06 GHz Xeon) w/PCI-X, each w/one 10-GE NIC (Intel PRO/10GbE LR Server Adapter)
- Two Apple G5 w/ PCI-X, each w/10-GE NIC (Chelsio T110)
- One10-GE uplink from one Extreme Network Summit 400-48t 1/10-GE switch (at GGAO)
  - One Apple G5 w/ PCI-X, w/one 1-GE NIC (Intel PRO/1000 MT Server Adapter)
  - One eVLBI Mark5 system (Owner: Alan Whitney (MIT-Haystack))

GSFC-Enabled High Performance Computer (HPC) Systems Using And/Or Accessible for Test Across DRAGON (3 of 4)

Project-Managed HPC Systems at GSFC

(All with own 10-GE connection to HECN's Force10 E600 10-GE switch/router)

Cluster Computing R&D Project (PI: John Dorband (GSFC))

- Two 10-GE uplinks from one Extreme Network Summit 400-48t 1/10-GE switch
  - Two 5-node computer-clusters: Both clusters consist of five nodes with identical hardware specifications. Each node (each w/1-GE NIC (Intel Corp. PRO/1000 MT Server Adapters)) features two 2.8 GHz Intel Pentium 4 Xeon processors with HyperThreading technology on a Tyan Thunder i7500 motherboard, along with 1 gigabyte of system memory and an 80 gigabyte hard disk. Both clusters run the Rocks Cluster Distribution, a popular Linux-based cluster operating system. One cluster runs version 3.3 of Rocks (based on the Linux 2.4.21 kernel), while the other runs Rocks version 4.0 (based on the Linux 2.6.9 kernel).

Northrop Grumman IR&D Project (PI: Brice Womack (NGC))

- Two dual 2.2 GHz Opteron w/ PCI-Express, each w/one 10-GE NIC (Myricom Myri-10G)

Teraflow Testbed (TFT) Project (PI: Robert Grossman (UIC))

- One SMC SMC8708L2 10-GE switch w/eight "1-to-1 non-oversubscribed" 10-GE ports
  - Four dual-Opteron 265 servers, each w/one 10-GE NIC (Intel PRO/10GbE LR Server Adapter)

GSFC-Enabled High Performance Computer (HPC) Systems Using And/Or Accessible for Test Across DRAGON (4 of 4)

Project-Managed HPC Systems at GSFC

(All with own 10-GE connection to HECN's Force10 E600 10-GE switch/router)

EOSDIS Network Prototyping Lab (ENPL) Project (PI: George Uhl (SWALE))

- One Force10 E300 10-GE switch/router, w/four "1-to-1 non-oversubscribed" 10-GE ports and twenty-four 1-GE ports
  - One Intel Pentium 4 (dual 3.06 GHz Xeon) w/PCI-X, each w/one 10-GE NIC (Intel PRO/10GbE LR Server Adapter)

Scientific Visualization Studio (SVS) Project (PI: Horace Mitchell (GSFC))

- One 10-GE uplink from one Extreme Network Summit 400-48t 1/10-GE switch
  - One 10-node display-cluster: The cluster consists of ten nodes with identical hardware specifications. Each node (each w/1-GE NIC (Intel Corp. PRO/1000 MT Server Adapters) features a 3.2 GHz Intel Pentium 4. Each of nine nodes drives one display tile in a 3x3-tiled display. The 10th node serves as the display-cluster's controller.

NASA Center for Computational Sciences (NCCS) Project (PI: Phil Webster (GSFC))

- One Force10 E600 10-GE switch/router, w/four "1-to-1 non-oversubscribed" 10-GE ports, sixteen "4-to-1 oversubscribed" 10-GE ports, and twenty-four 1-GE ports
  - One SGI Origin 3800 w/PCI-X, w/one 10-GE NIC (Neterion)
  - Others in progress

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- eVLBI (w/MIT-Haystack, …)
- OptIPuter & Multi-channel Collaboration/Video Streaming Technologies(w/UCSD & UIC)
- 3D HDTV-over-IP R&D (w/Physical Optics Corporation)
- Distributed ESMF R&D
- SAN-over-IP (w/UMIACS & NGC)
- Using ARC/NAS/Columbia Supercomputer (w/NREN)

DRAGON eVLBI Experiment Configuration

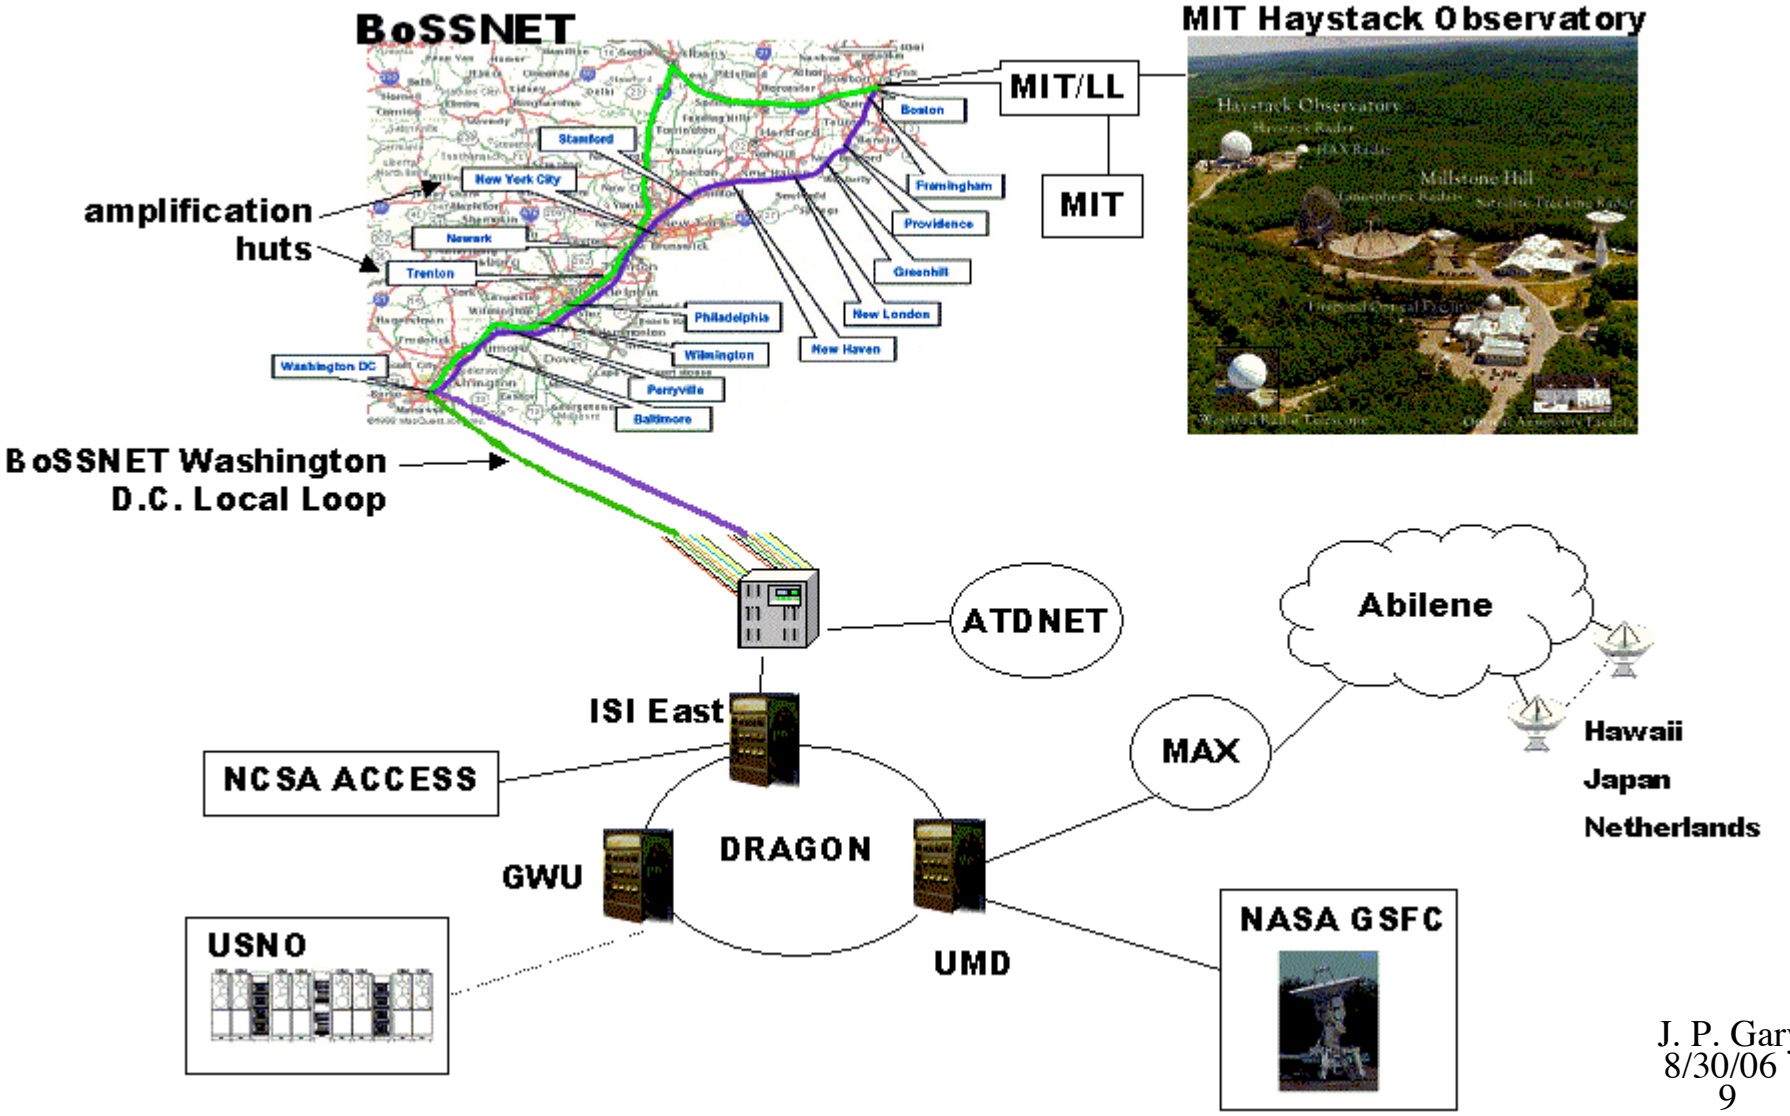

NASA GSFC Tests with OptIPuter Across the National LambdaRail

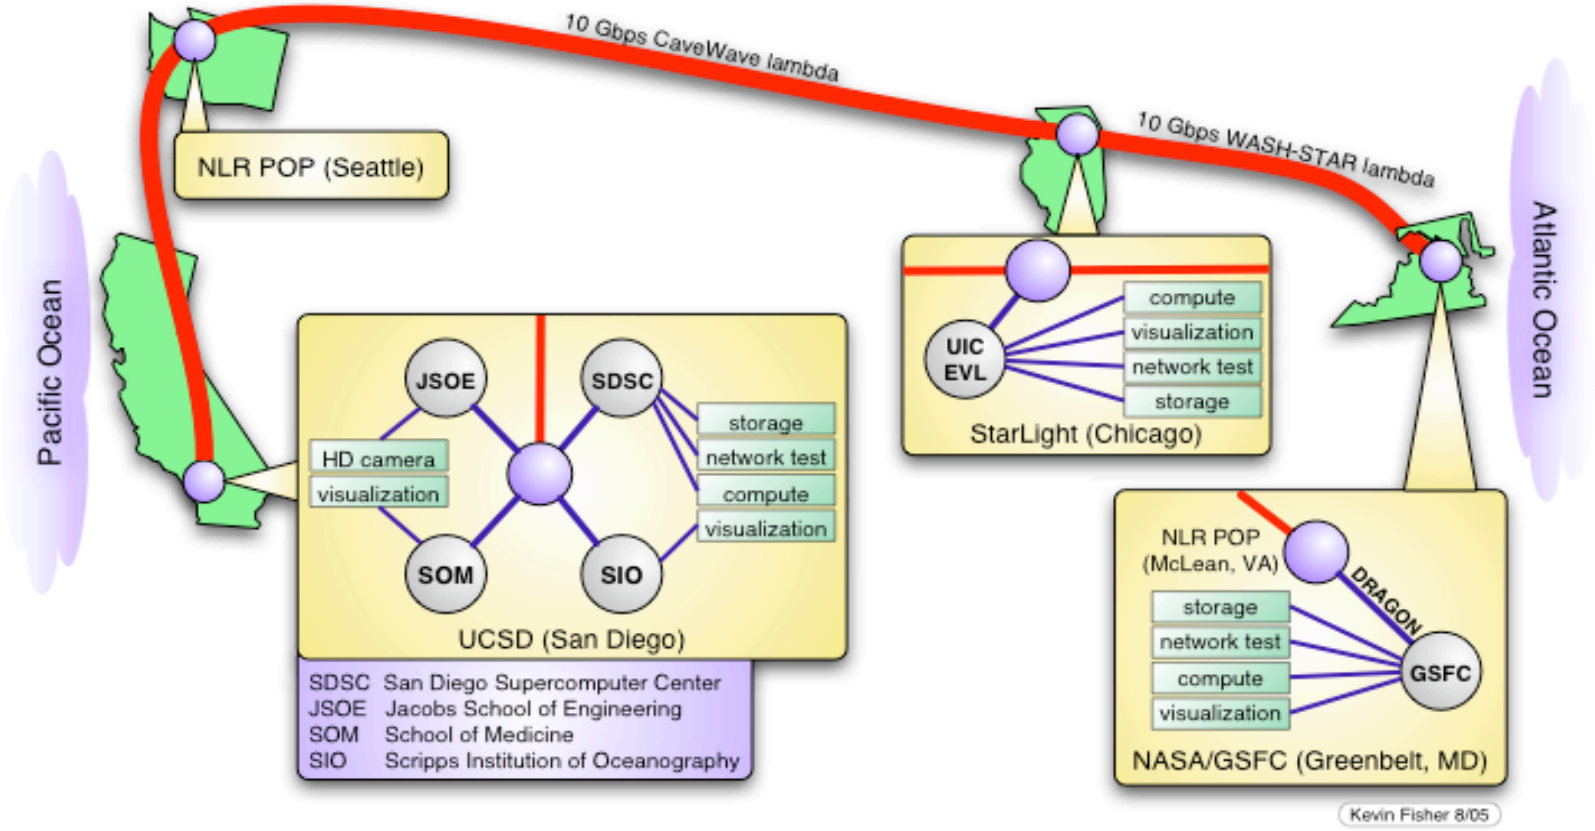

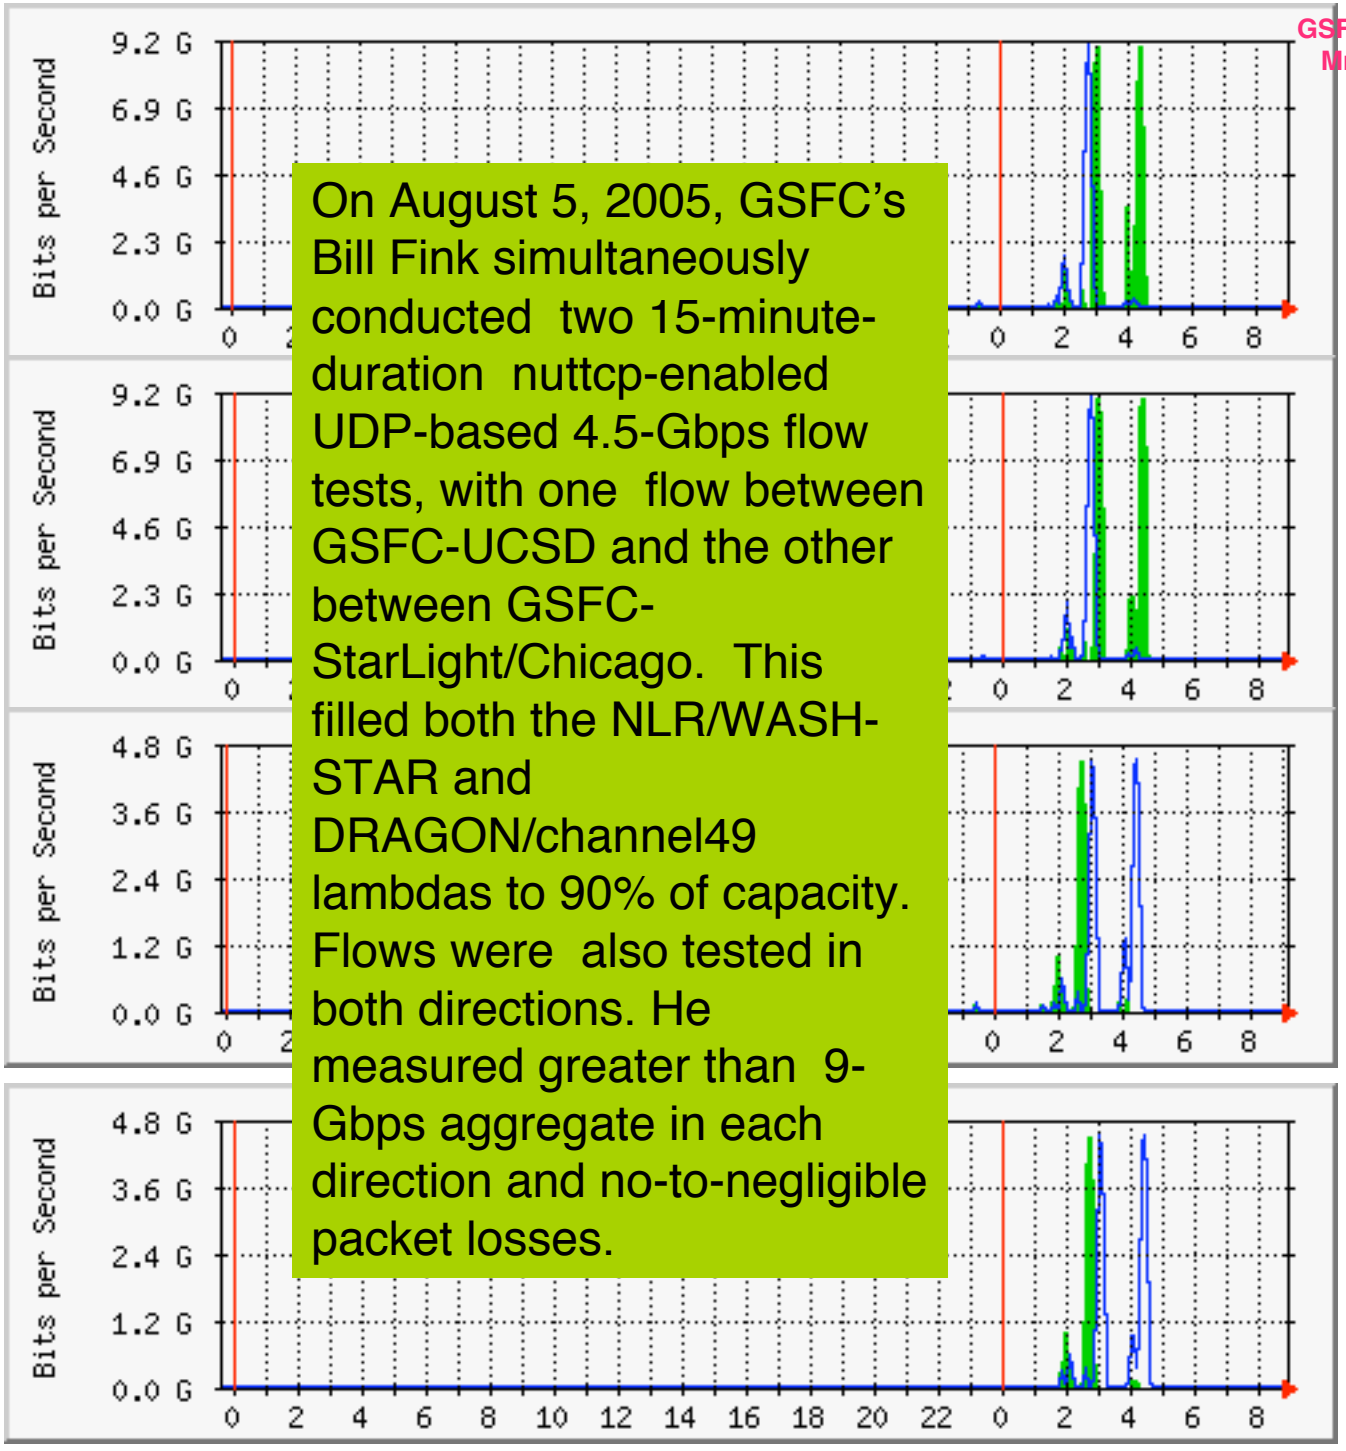

On August 5, 2005, GSFC's Bill Fink simultaneously conducted two 15-minute-duration nuttcp-enabled UDP-based 4.5-Gbps flow tests, with one flow between GSFC-UCSD and the other between GSFC-StarLight/Chicago. This filled both the NLR/WASH-STAR and DRAGON/channel49 lambdas to 90% of capacity. Flows were also tested in both directions. He measured greater than 9-Gbps aggregate in each direction and no-to-negligible packet losses.

10 GigE from McLean to Chicago via OptIPuter Lambda
5 August 2005

DRAGON 10Gig DWDM XFP
5 August 2005

chance1 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

chance2 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

Combining Telepresence with Remote Interactive Analysis of Data Over NLR

www.calit2.net/articles/article.php?id=660

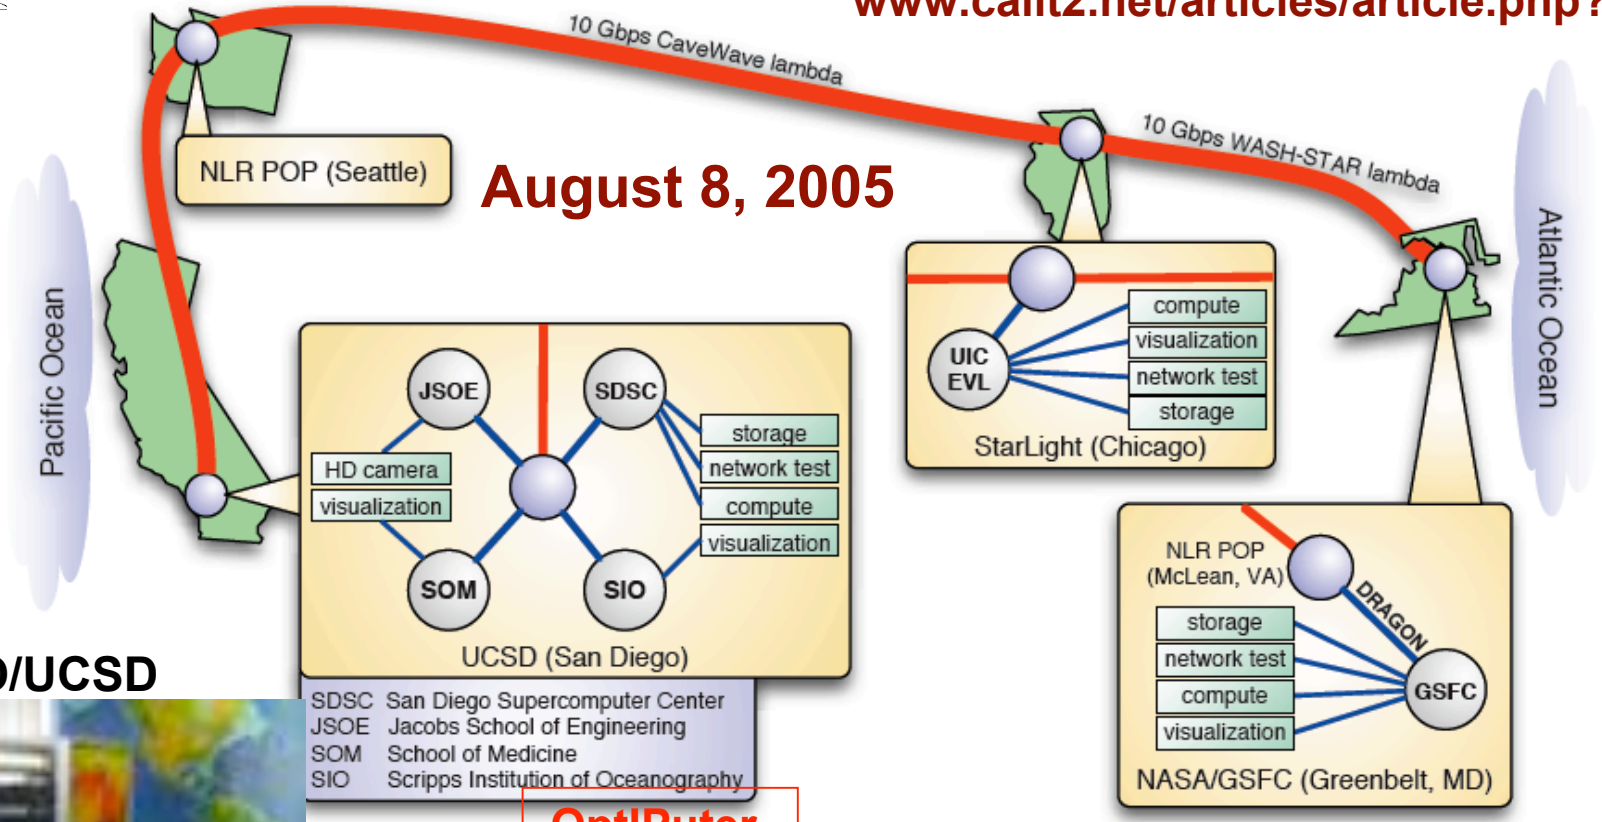

SIO/UCSD

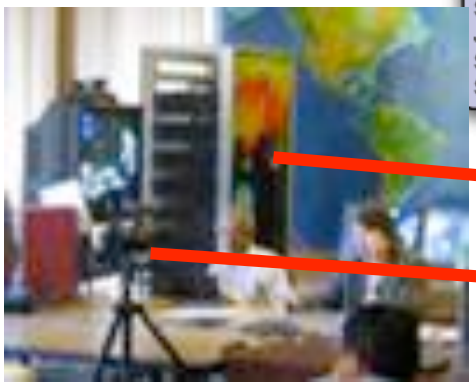

OptIPuter Visualized Data

HDTV Over Lambda

NASA Goddard

Gold Standard OptIPuter Initial Sites
Persistent 10G Services with SAGE End Appliances

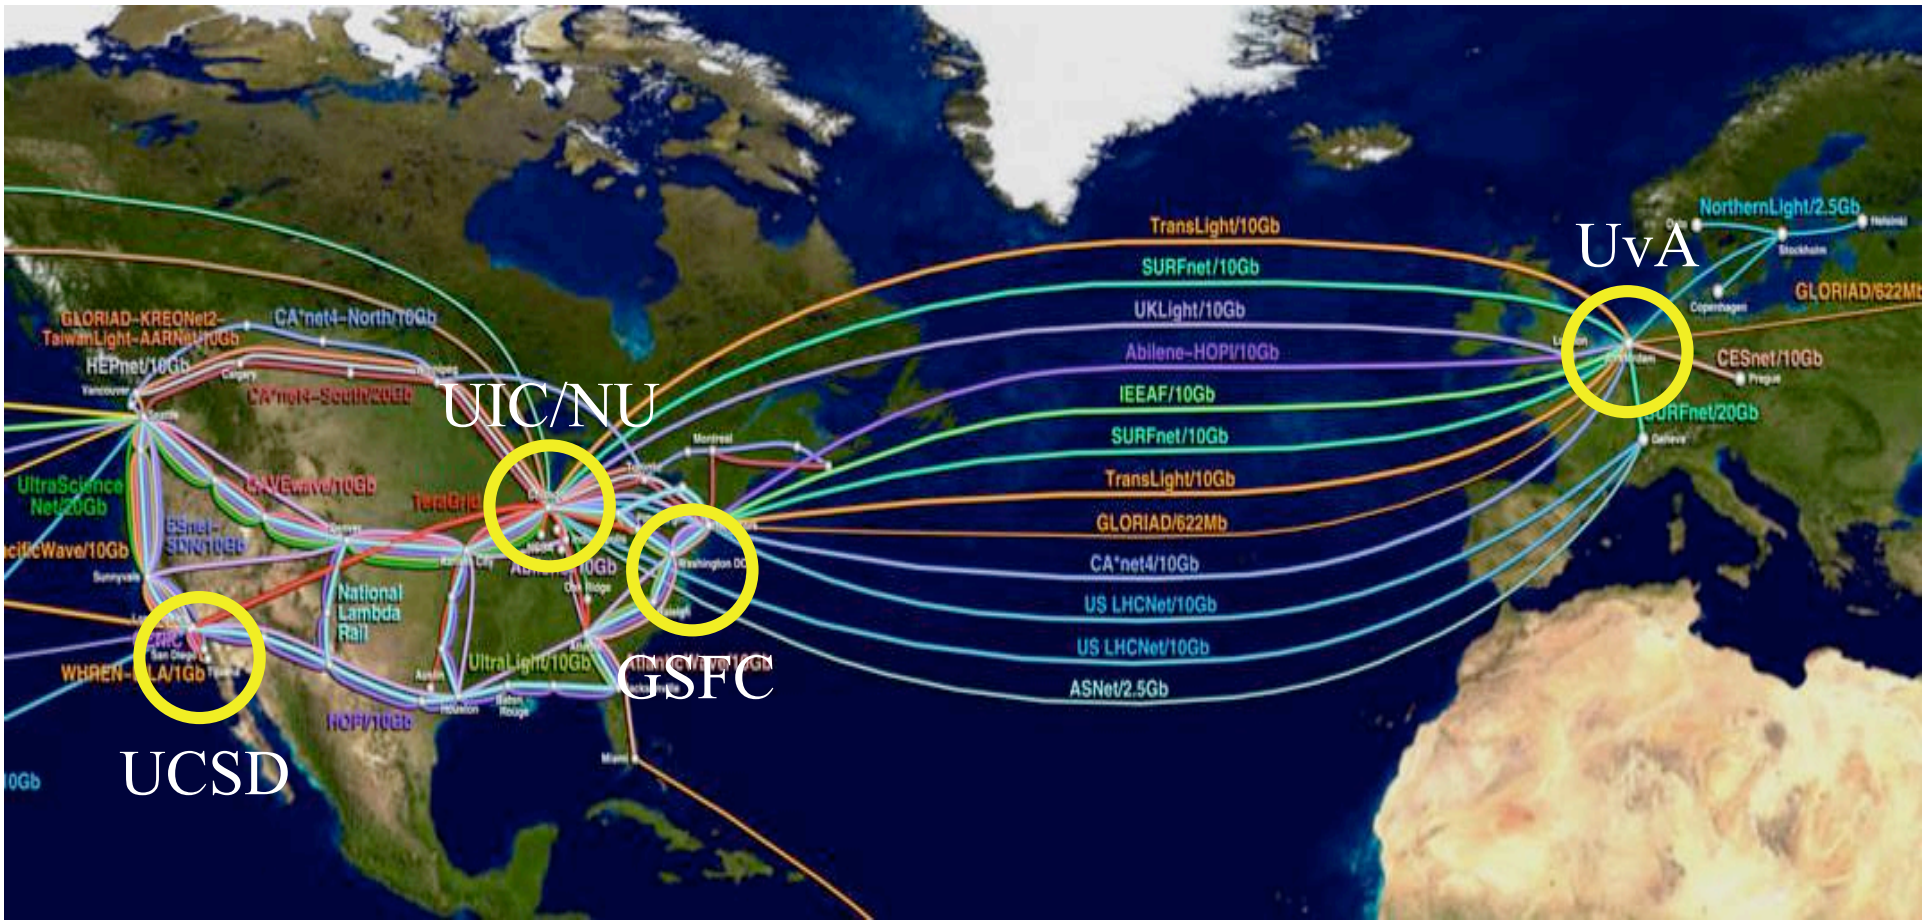

iGrid 2005 Workshop, 26-29Sep05, UCSD/CalIT2

Accelerating the Use of Multi-10Gigabit per Second International and National Networks: www.igrid2005.org

GSFC's Ben Kobler (left) and POC's Sookwang Ro and Kirill Kolesnikov (right) work to set up POC's 35" x 35" holographic 3D HDTV video display system (center) prior to the start of iGrid 2005.

US130: Real-Time True-3D/HDTV (No Goggles) Visualization Over the National LambdaRail

NASA and Physical Optics Corporation demonstrate a holographic 3D HDTV video display system that does not require goggles or other special head gear, using a live cross-country video feed from NASA Goddard Space Flight Center to the iGrid 2005 site in San Diego. POC is a NASA SBIR Phase 1 awardee, and worked with NASA GSFC on this project.

www.poc.com/emerging_products/3d_display/default.asp

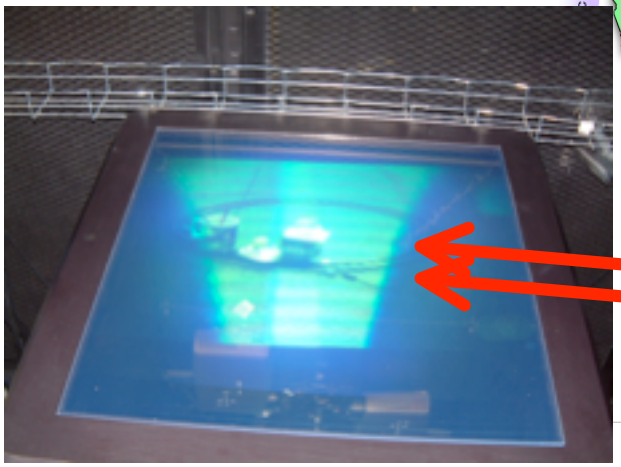

Only a non-stereo image of the True-3D display is captured in this photo of the real-time stereo-HDTV images transmitted from GSFC.

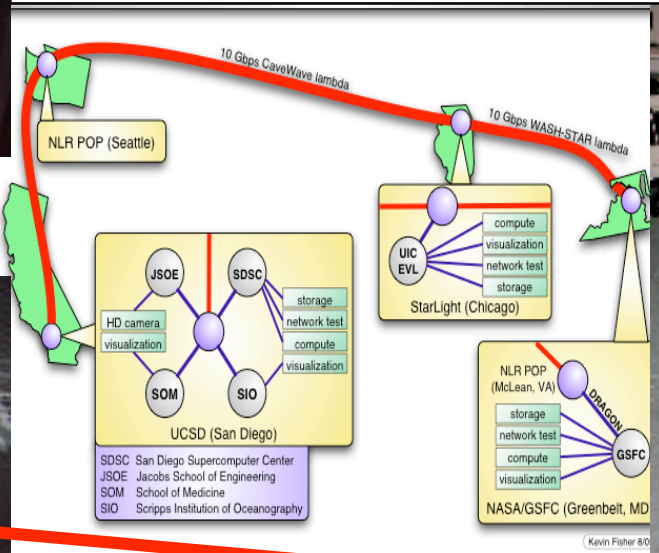

3D HDTV Over Lambda

Stereoscoptically-aligned Sony HDV 1080i HDR-FX1HDTV cameras and the viewed targets at GSFC.

Cross-Organization Coupling of Climate Models through ESMF

(A Prototype Over High-Speed Networks)

Shujia Zhou (Lead), C. Cruz, R. Burns, B. Womack, G. Higgins NASA SIVO/Northrop Grumman TASC

Collaborators:

- **High-speed network: P. Gary, B. Fink, P. Lang (NASA GSFC/ADNET)**
- **Cluster system admin: K. Fisher (NASA GSFC)**
- **XCAT/Proteus: M. Govindaraju, K. Chiu, M. Head (SUNY, Binghampton)**
- **Models: J. Spahr, C. Mechoso (UCLA), C. Hill (MIT), P. Jones (LANL)**

Presented at NASA Exhibit (booth 1810) at SC|05, November 14-18, 2005

EARTH SYSTEM MODELING FRAMEWORK

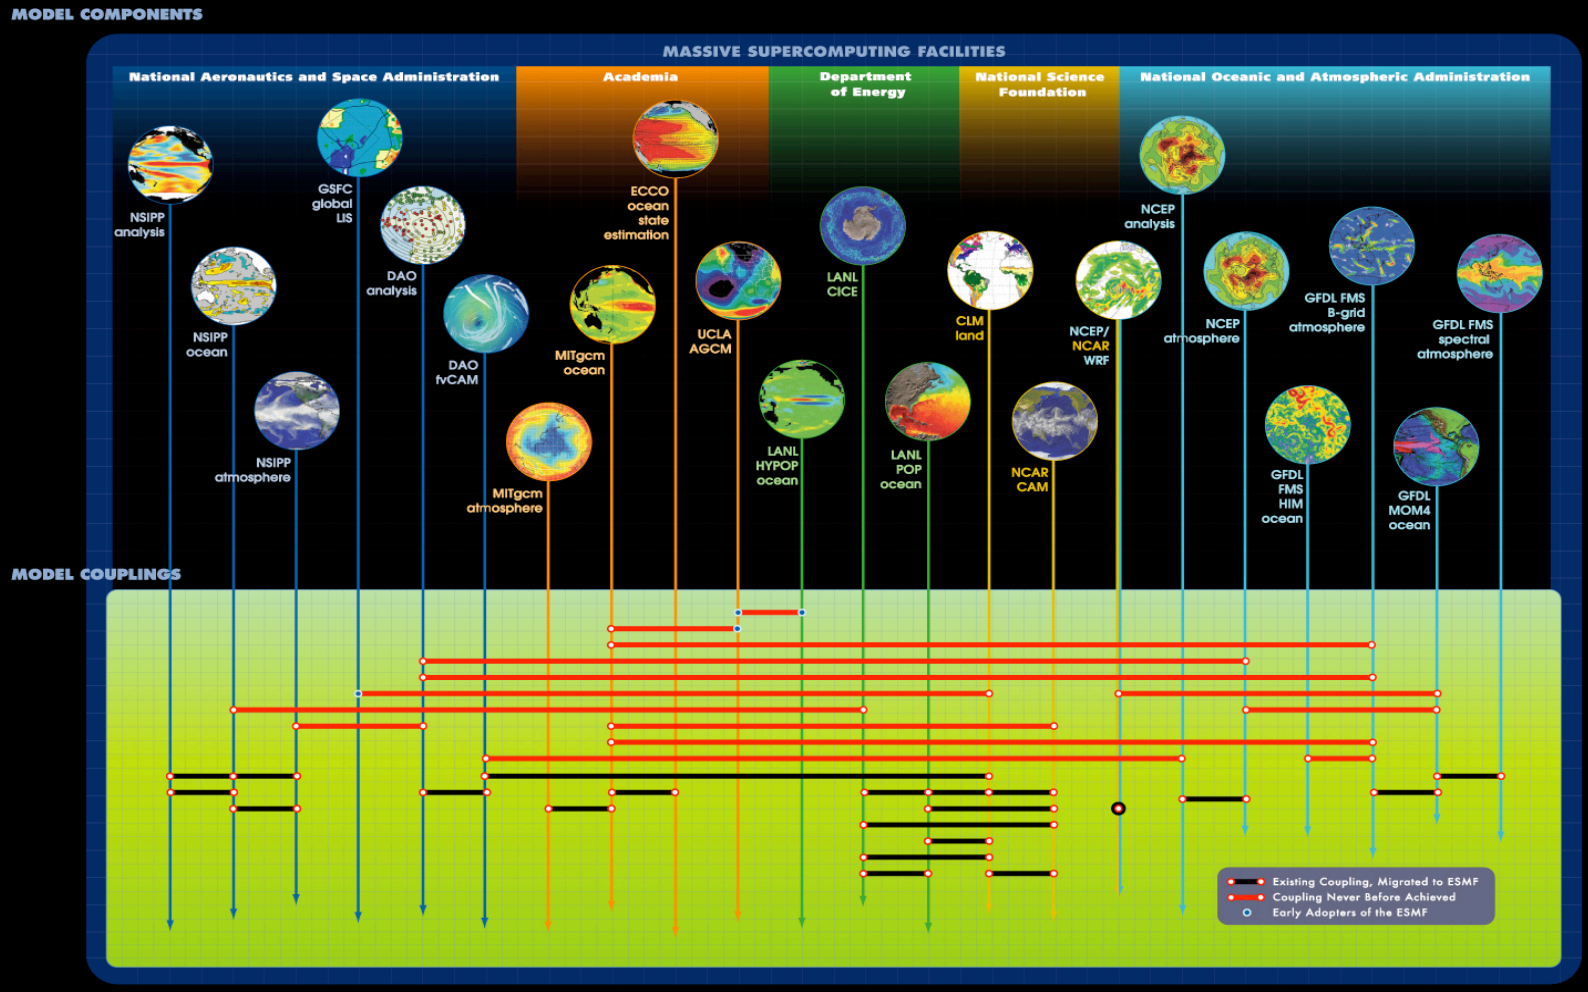

Coupled Atmosphere-Ocean Models

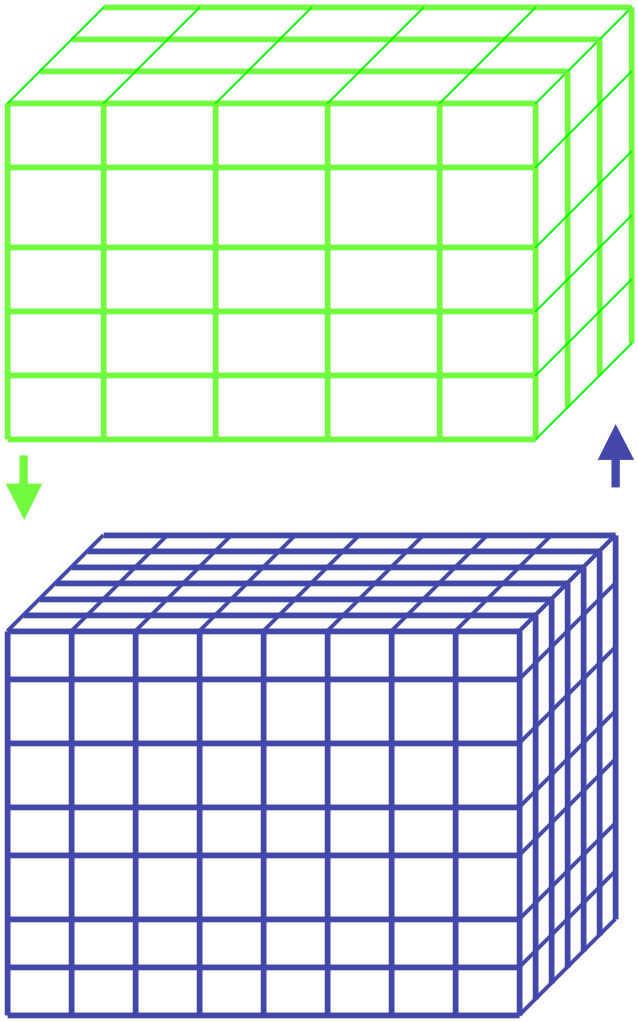

Atmosphere

Ocean

Different grid type, resolution

Thunder2-Thunder3 Looped Data Flows Across DRAGON Regional Optical Network

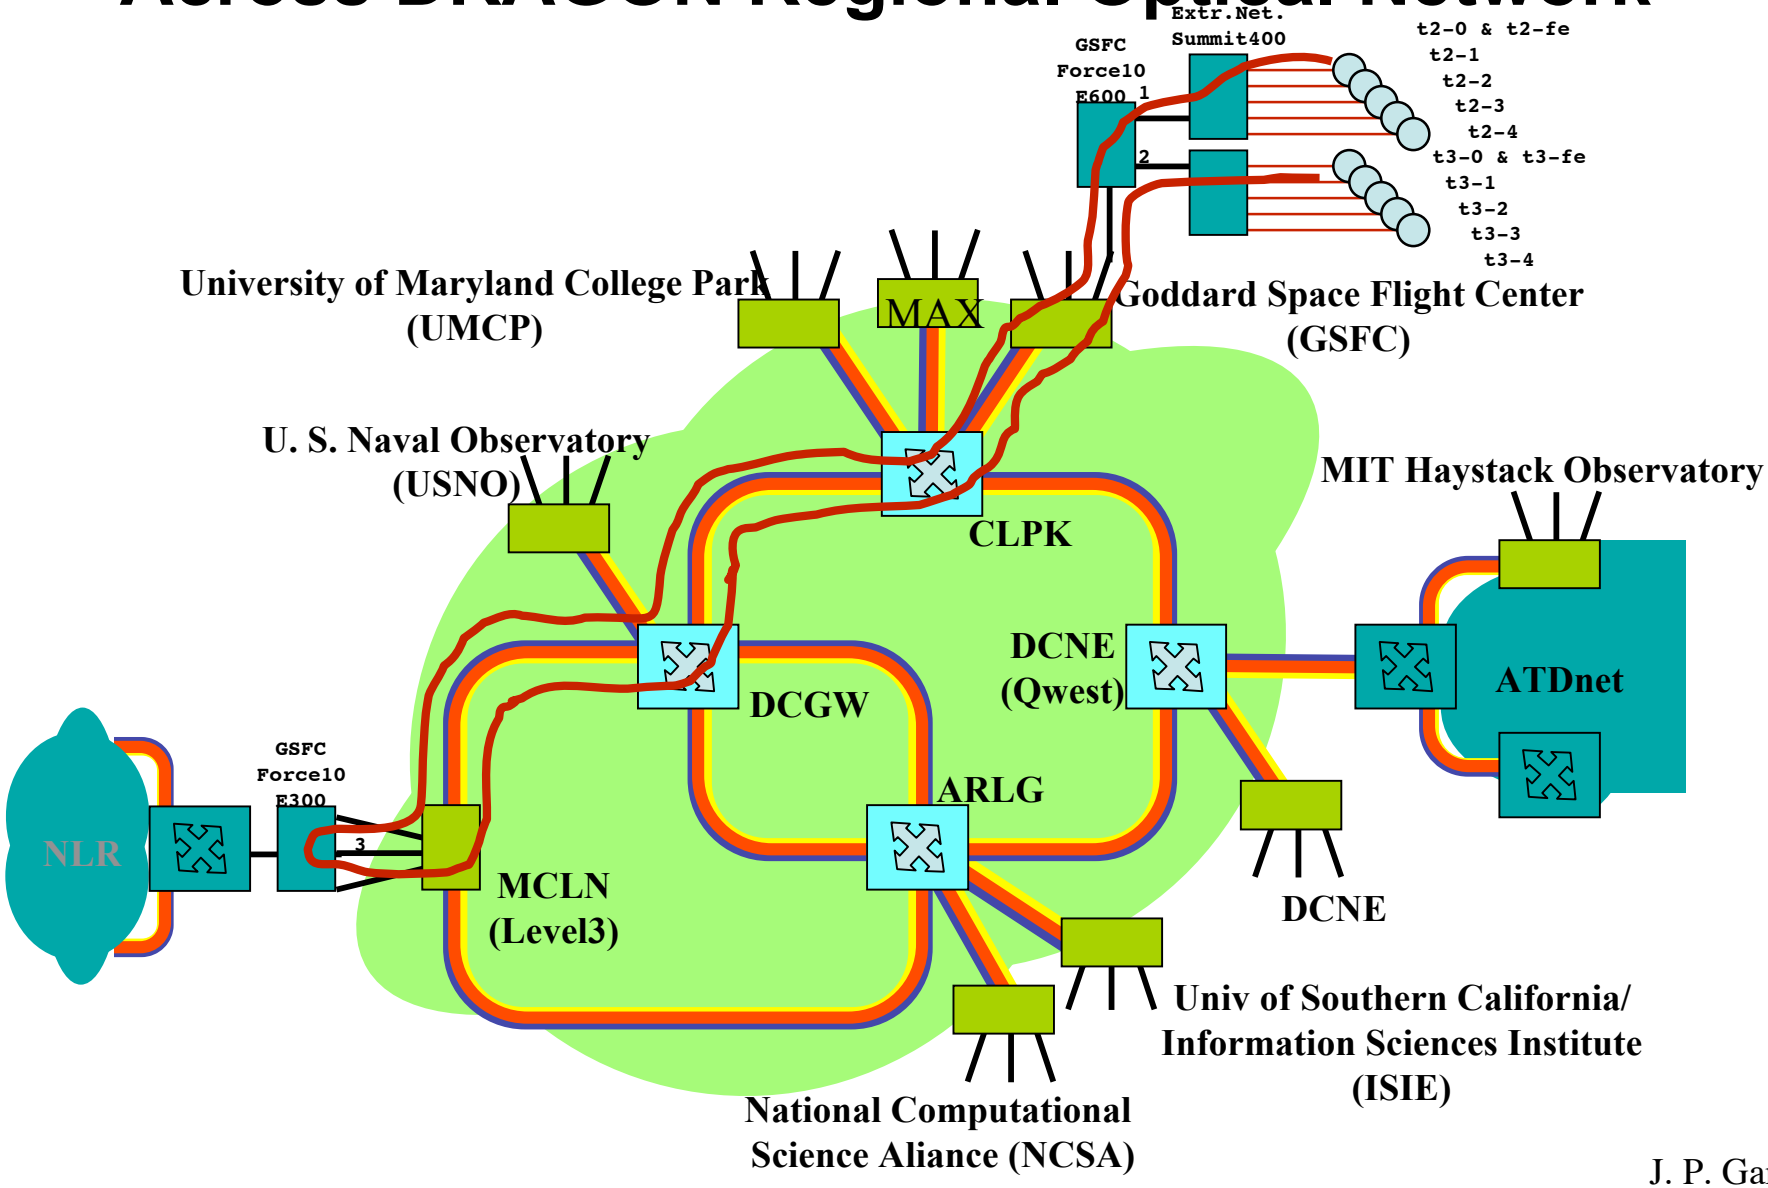

Current SAN-over-IP Test-bed

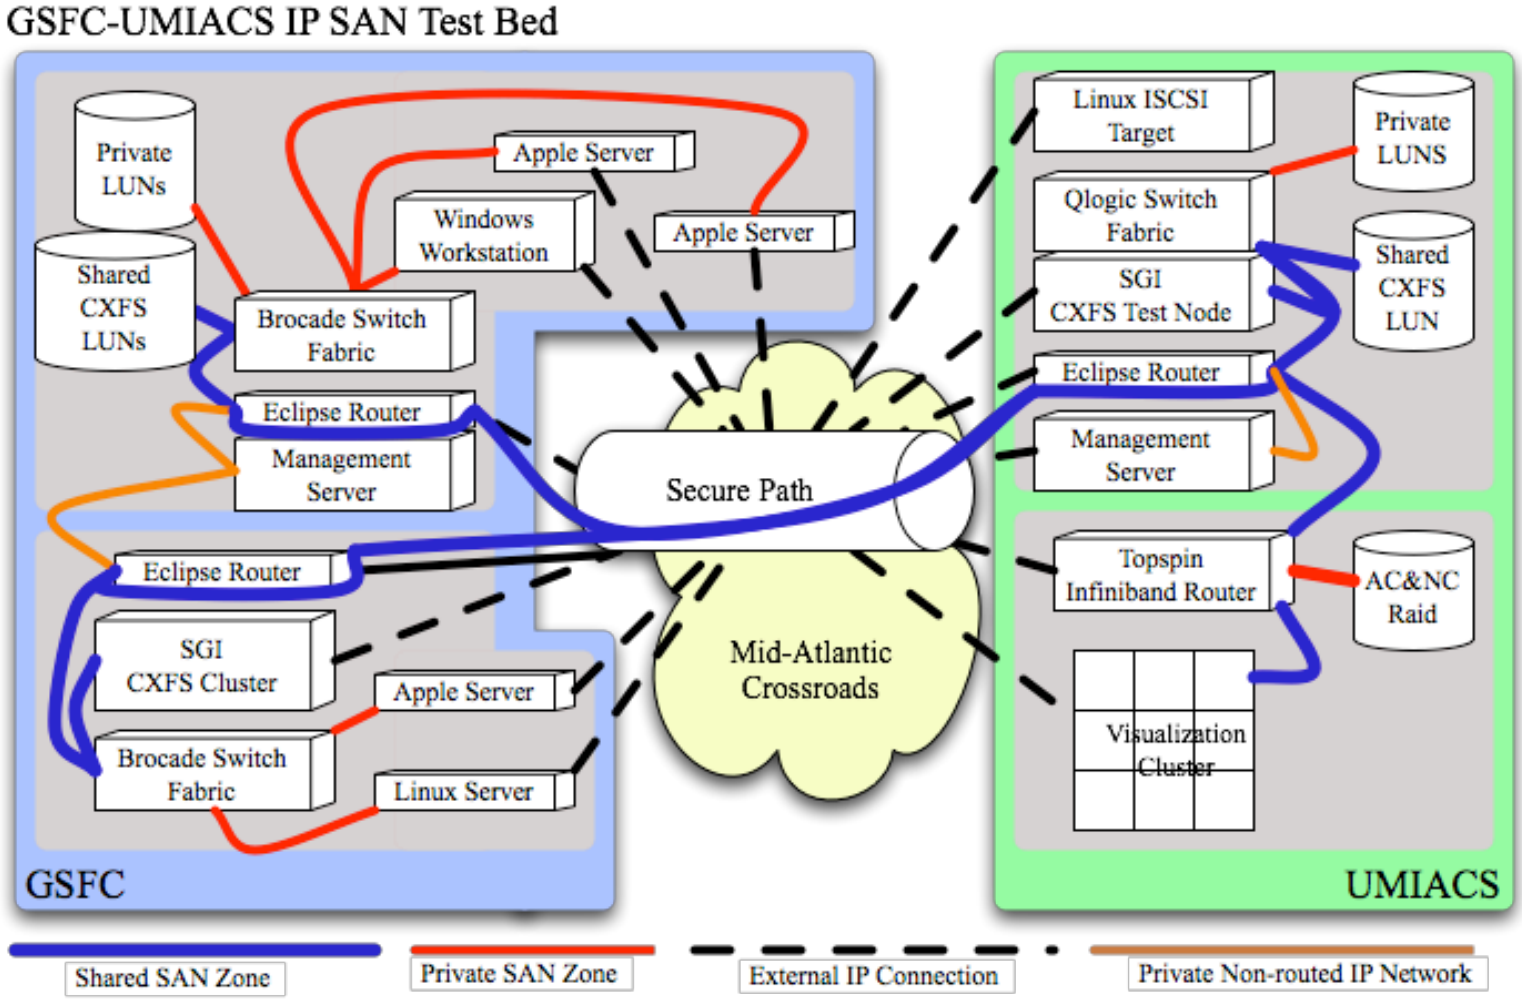

Source: Fritz McCall (UMIACS)

Wide Area Storage Configuration

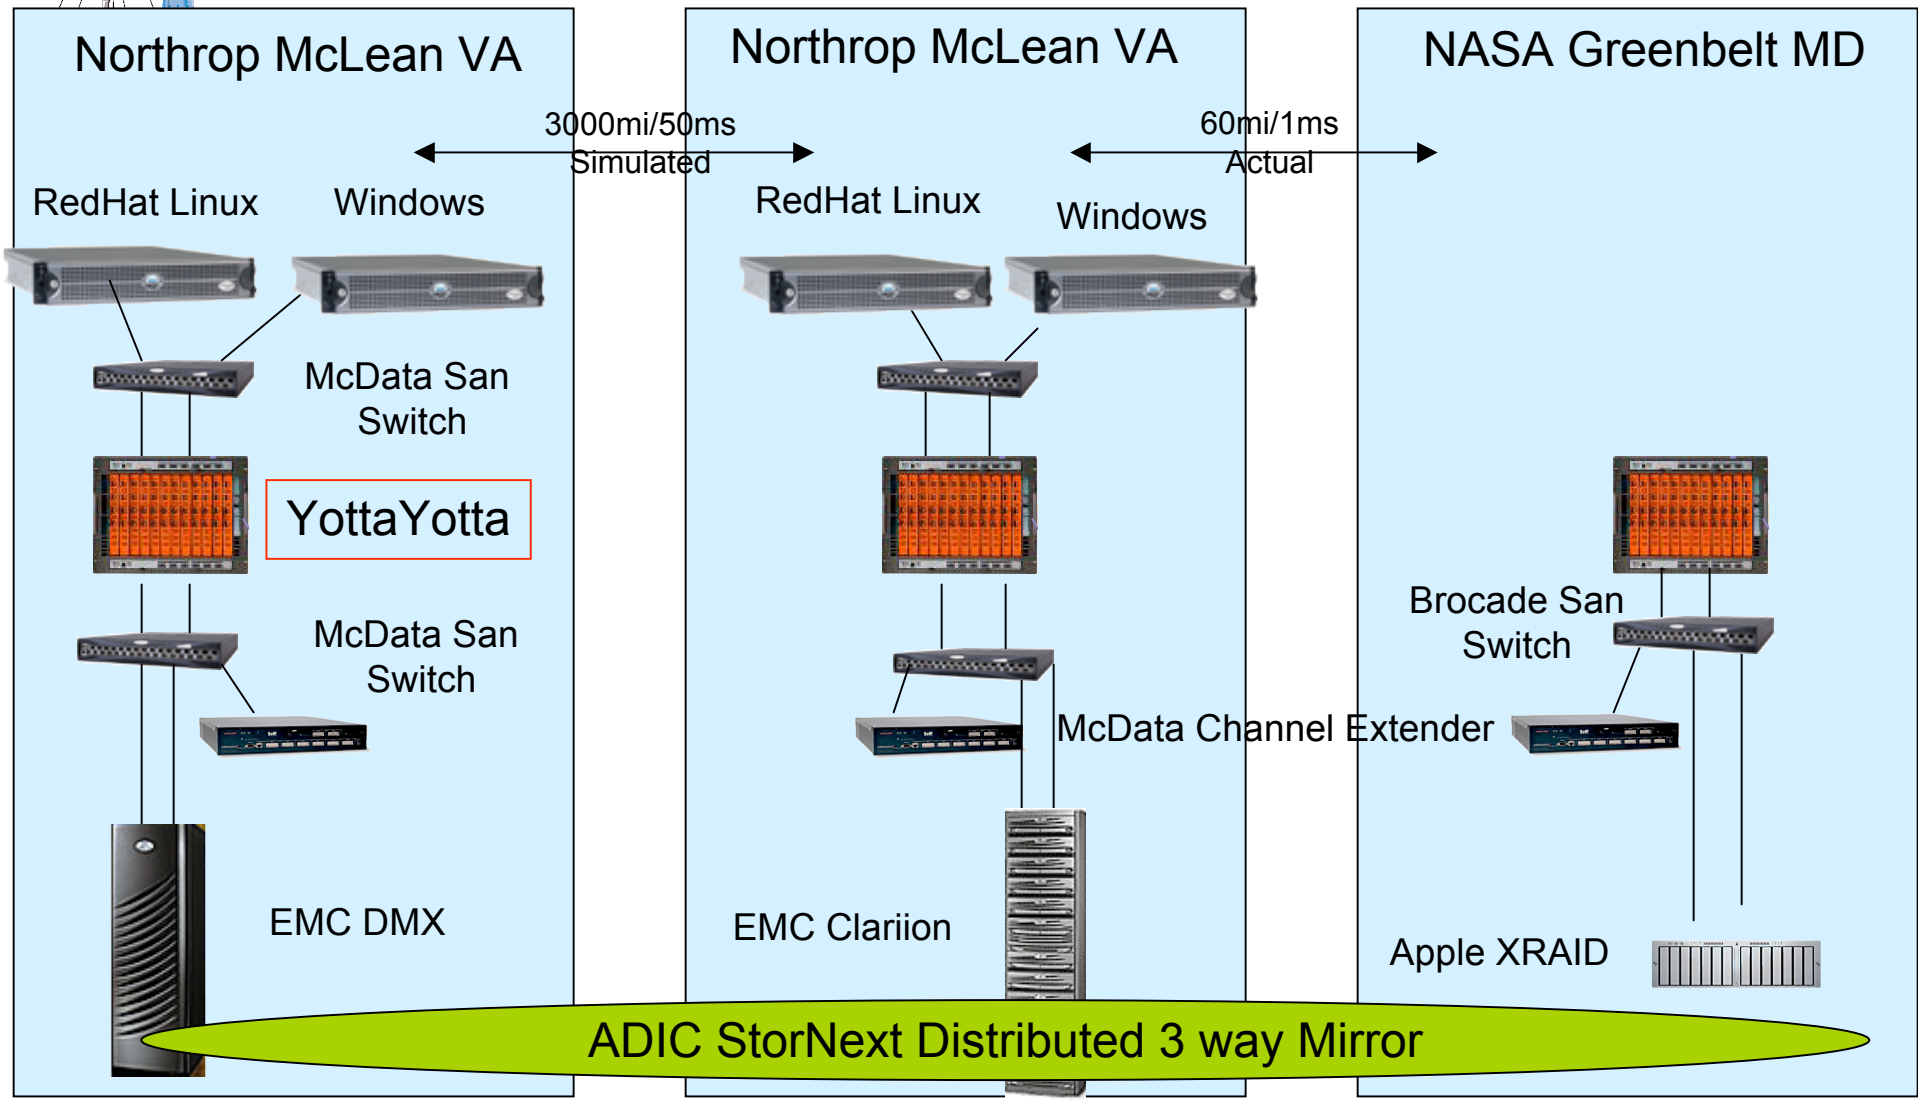

Source: Bob Bramow (YottaYotta)

Columbia Supercomputer

- 10,240 1.6 GHz CPUs
- Configured as twenty 512 CPU single-system image nodes via NUMA
- SGI Altix 3700 Architecture, runs Linux
- 1 Terabyte shared memory per node
- Over 500 terabytes of online disk space

WAN Infrastructure: NREN

Target CY06

10G waves at the core, dark fiber to end sites

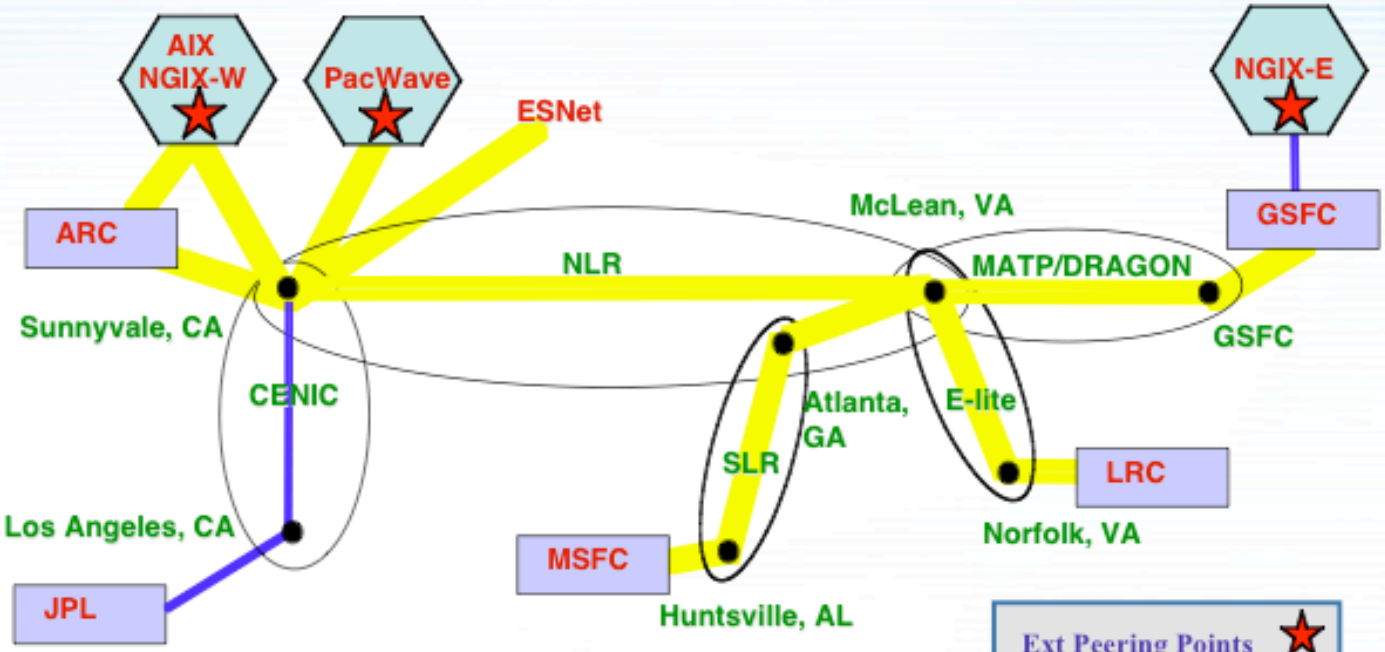

- National and Regional optical networks provide links over which 10 Gbps and 1 Gbps waves can be established.
- Distributed exchange points provide interconnect in metro and regional areas to other networks and research facilities

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Upcoming and/or Future Applications

- MAP'06 (w/NGC)
- Phoebus (w/Internet2, UDel & GSFC/ENPL)
- SOA+Brokering for ECHO (w/SIO, JPL & UAH)
- Dynamic Linking (w/ORNL, CUNY)
- Grid Computing (w/TBD: SURAGrid, UMBC, …)

2006 Hurricane Season - Global Modeling

NCEP
Inputs (1 GB)

INTEL Corp
Portland, Oregon

NASA Ames
Mt. View, California

Northrop Grumman
McLean, Virginia

NASA Goddard
Greenbelt, Maryland

DISTRIBUTED COMPUTING NODES

NEXT-GEN NETWORKS

Conventional Network
(600 Mb/s)

National Lambda Rail (10 - 40 Gb/s)

All model Outputs (500 GB per execution, ~15TB for entire season

Tape Backup

NASA Goddard
Greenbelt, Maryland

Main Server

NASA Goddard
Greenbelt, Maryland

DATA SERVERS / LONG TERM STORAGE

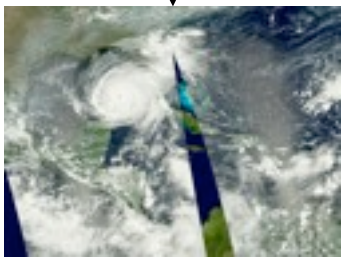

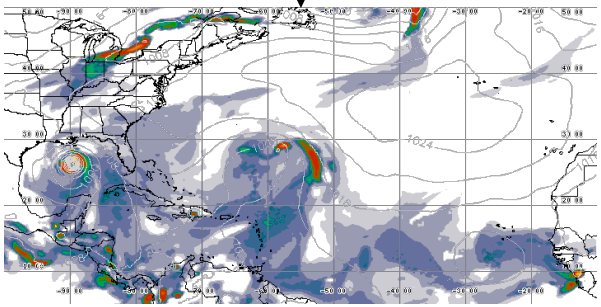

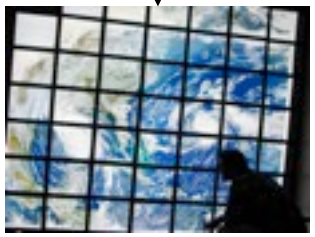

WEB SERVICES / ADVANCED VISUALIZATIONS

Source: Mike Seablom (GSFC/610.3)

ESMF-Enabled Coupled Models

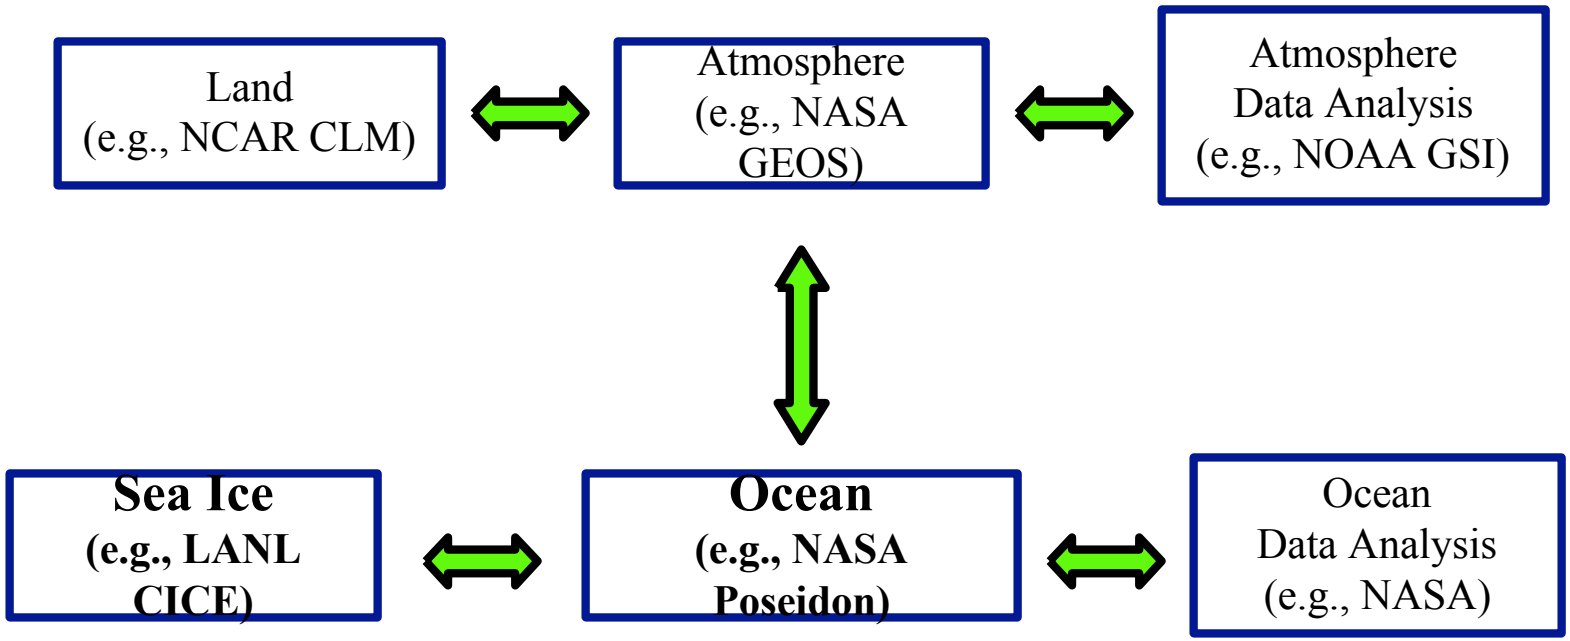

Example: Couple Different Ocean Models to NASA GEOS5

Source: Mike Seablom (GSFC/610.3)

DISTRIBUTED COMPUTING NODES

NASA Ames
Mt. View, California

Northrop Grumman
McLean, Virginia

NASA Goddard
Greenbelt, Maryland

NEXT-GENERATION NETWORKS

National Lambda Rail (10 - 40 Gb/s)

GEOS5/OGCM/ POSEIDON

GEOS5/OGCM/ MOM4

NCCS/NAS MODEL OUTPUT PORTALS

DAAC DATA PORTALS

SCIENCE ANALYSIS

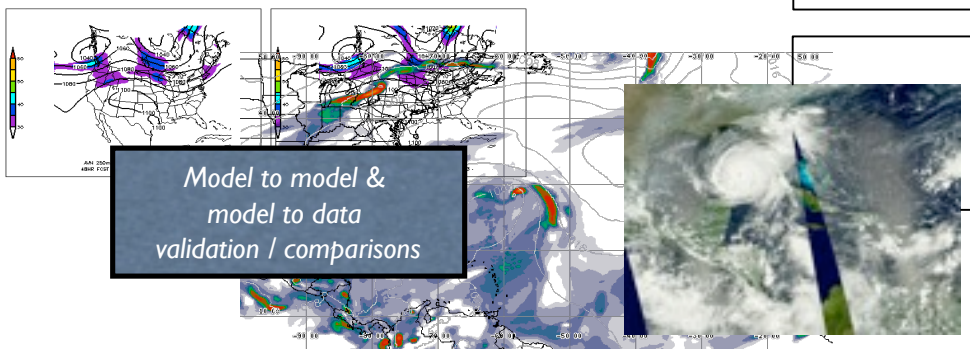

Model to model & model to data validation / comparisons

EXTERNAL COLLABORATORS

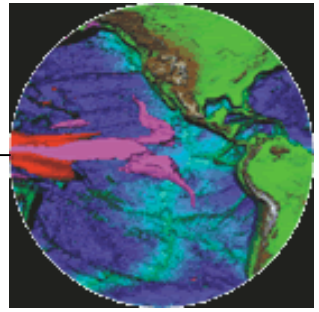

GFDL/NCDC MODEL OUTPUT PORTALS

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Phoebus

Date: Tue, 1 Aug 2006 16:38:41 -0400

To: almes@internet2.edu

From: Pat Gary <Pat.Gary@nasa.gov>

Subject: Follow-up to Pat Gary(GSFC)'s interest in Phoebus

Guy,

In follow-up to my interest in Phoebus that I expressed to you during the 17Jul06 Jt Tech Wksp, please asap reply to the following.

1. Could we work with your Phoebus team to soon arrange another Phoebus test between your test host at ColumbiaU and a test host we'd provide at GSFC in Greenbelt, MD, where the Phoebus/HOPI ingress/eqress POP's would be NYC and DC (actually McLean, to which we connect via DRAGON)?

(Note: GSFC via its GISS (http://www.giss.nasa.gov/) has both an organizational and physical presence on ColumbiaU's campus. Hence a good Phoebus/HOPI test between NYC and DC could have potential "real world" value to us.)

…

As I mentioned to you during the 17Jul06 Jt Tech Wksp, my longer-term application for Phoebus is between GSFC and the Columbia supercomputer at ARC (as partly described in http://www.internet2.edu/presentations/jt2006jul/20060718-columbiascc-jones.pdf) via either HOPI LA-DC initially or NREN Sunnyvale-DC once you allow us to copy Phoebus….

"Brokering and Chaining Distributed Services and Data Using OptIPuter and the National Lambda Rail" by Ramapriyan (GSFC) et al to NASA's ROSES NRA

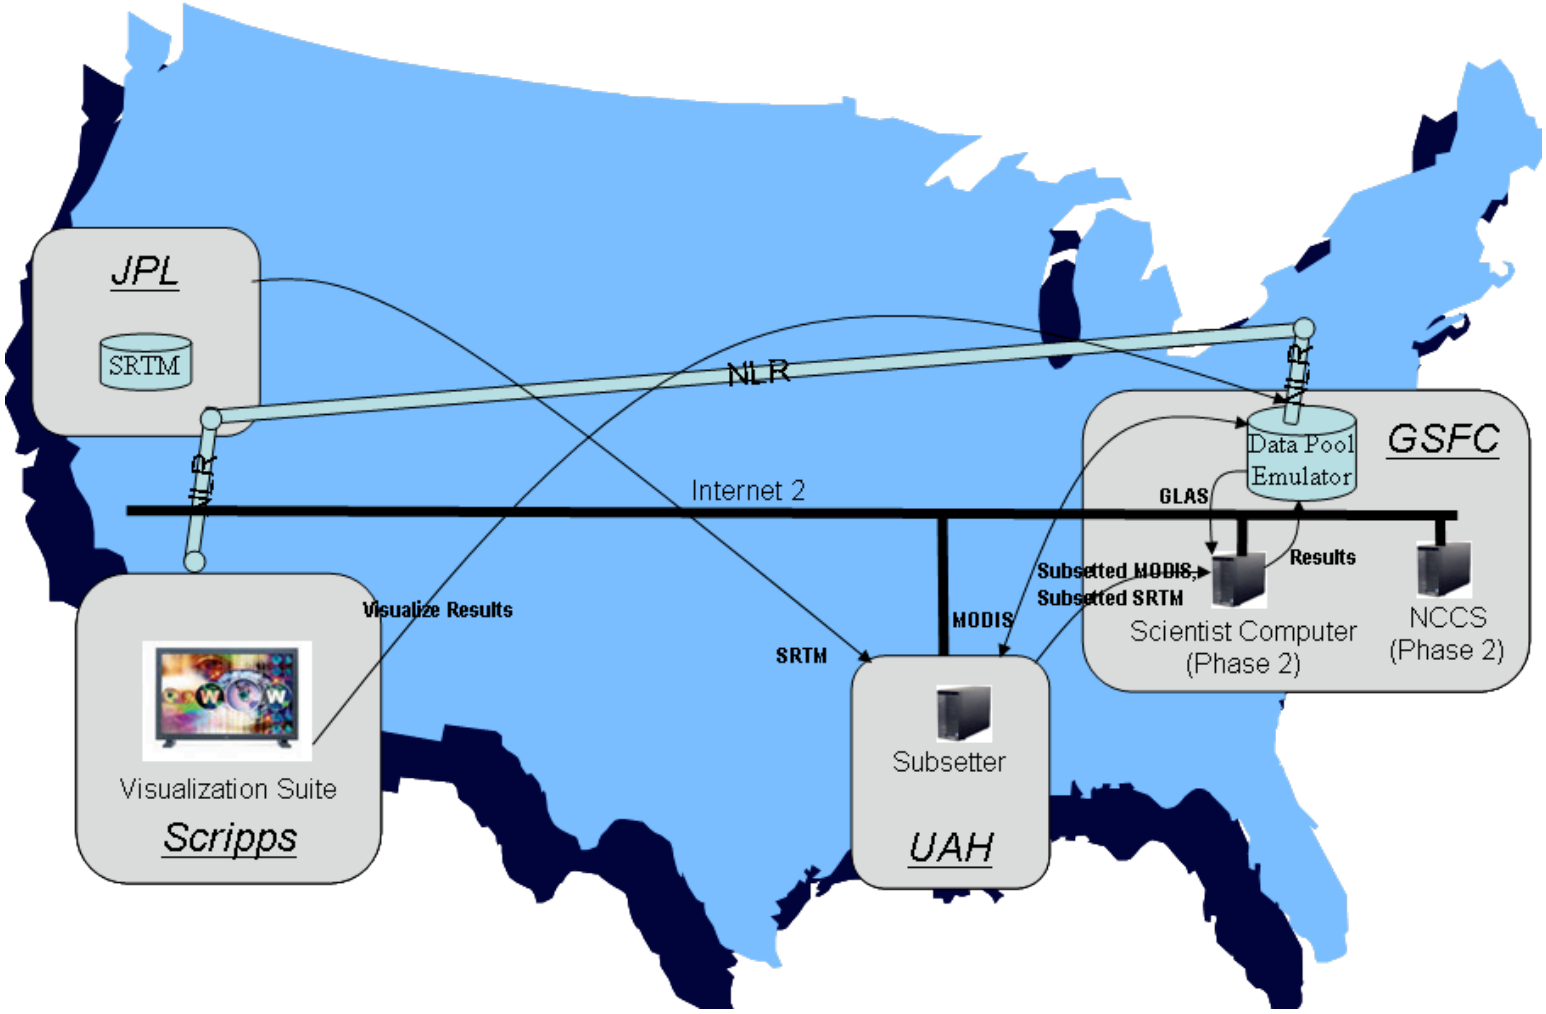

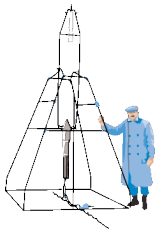

"Enabling NASA Applications Across Heterogeneous High Performance Networks" by Habib (CUNY) et al to NASA NNH05ZDA001N-Applied Information Systems Research (a.k.a. ROSES:D3)

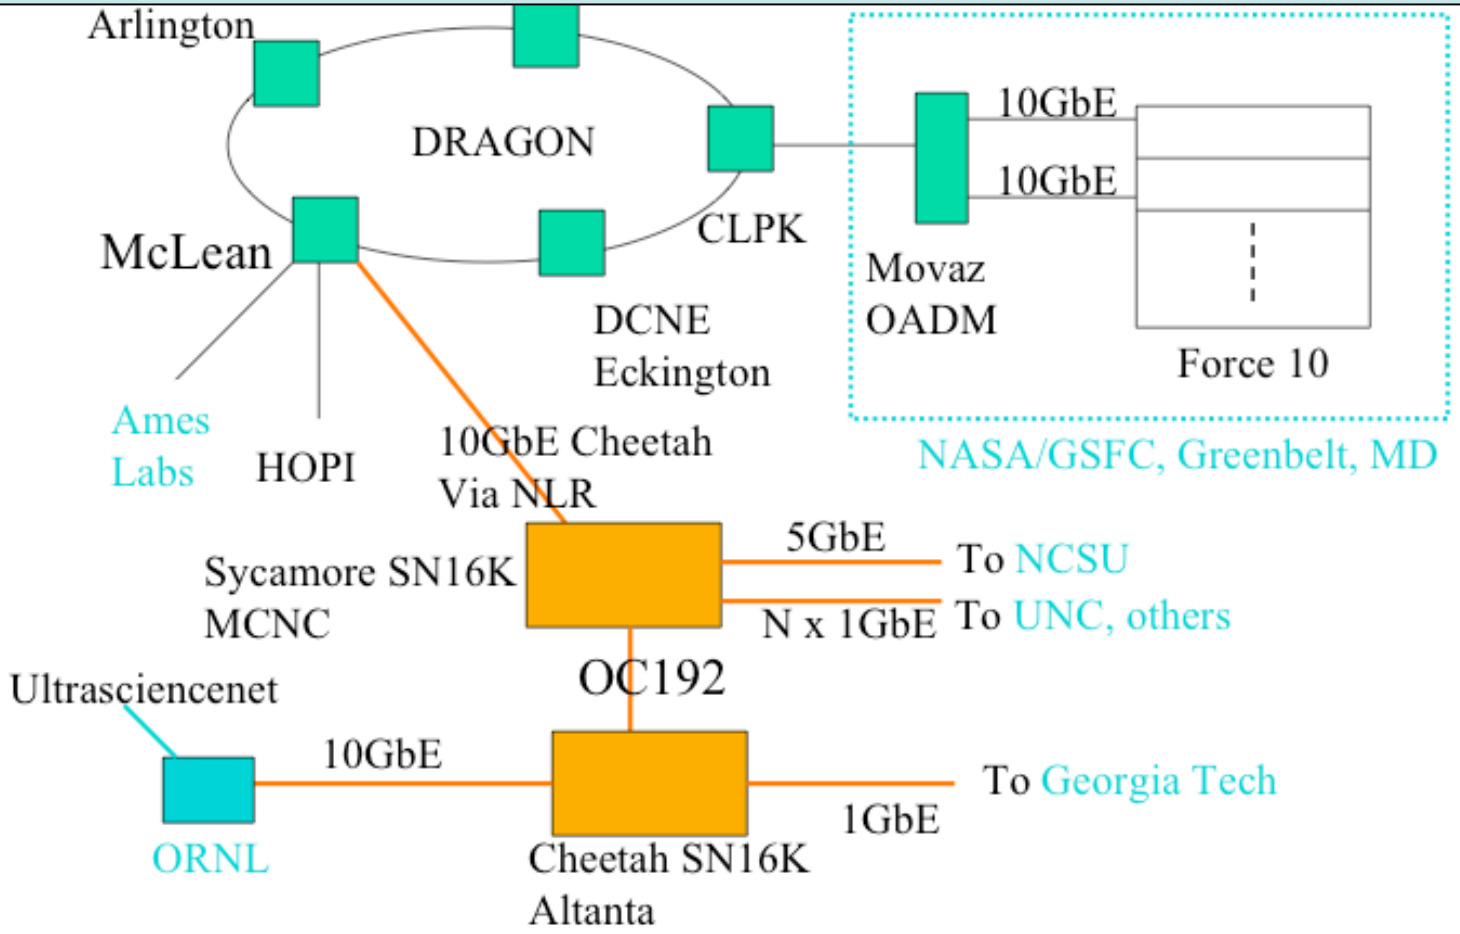

Fig. 1: Overall Proposed Network Connectivity to Cheetah

High Performance Networking and Remote Data Access
GSFC L-Net for NCCS and Science Buildings

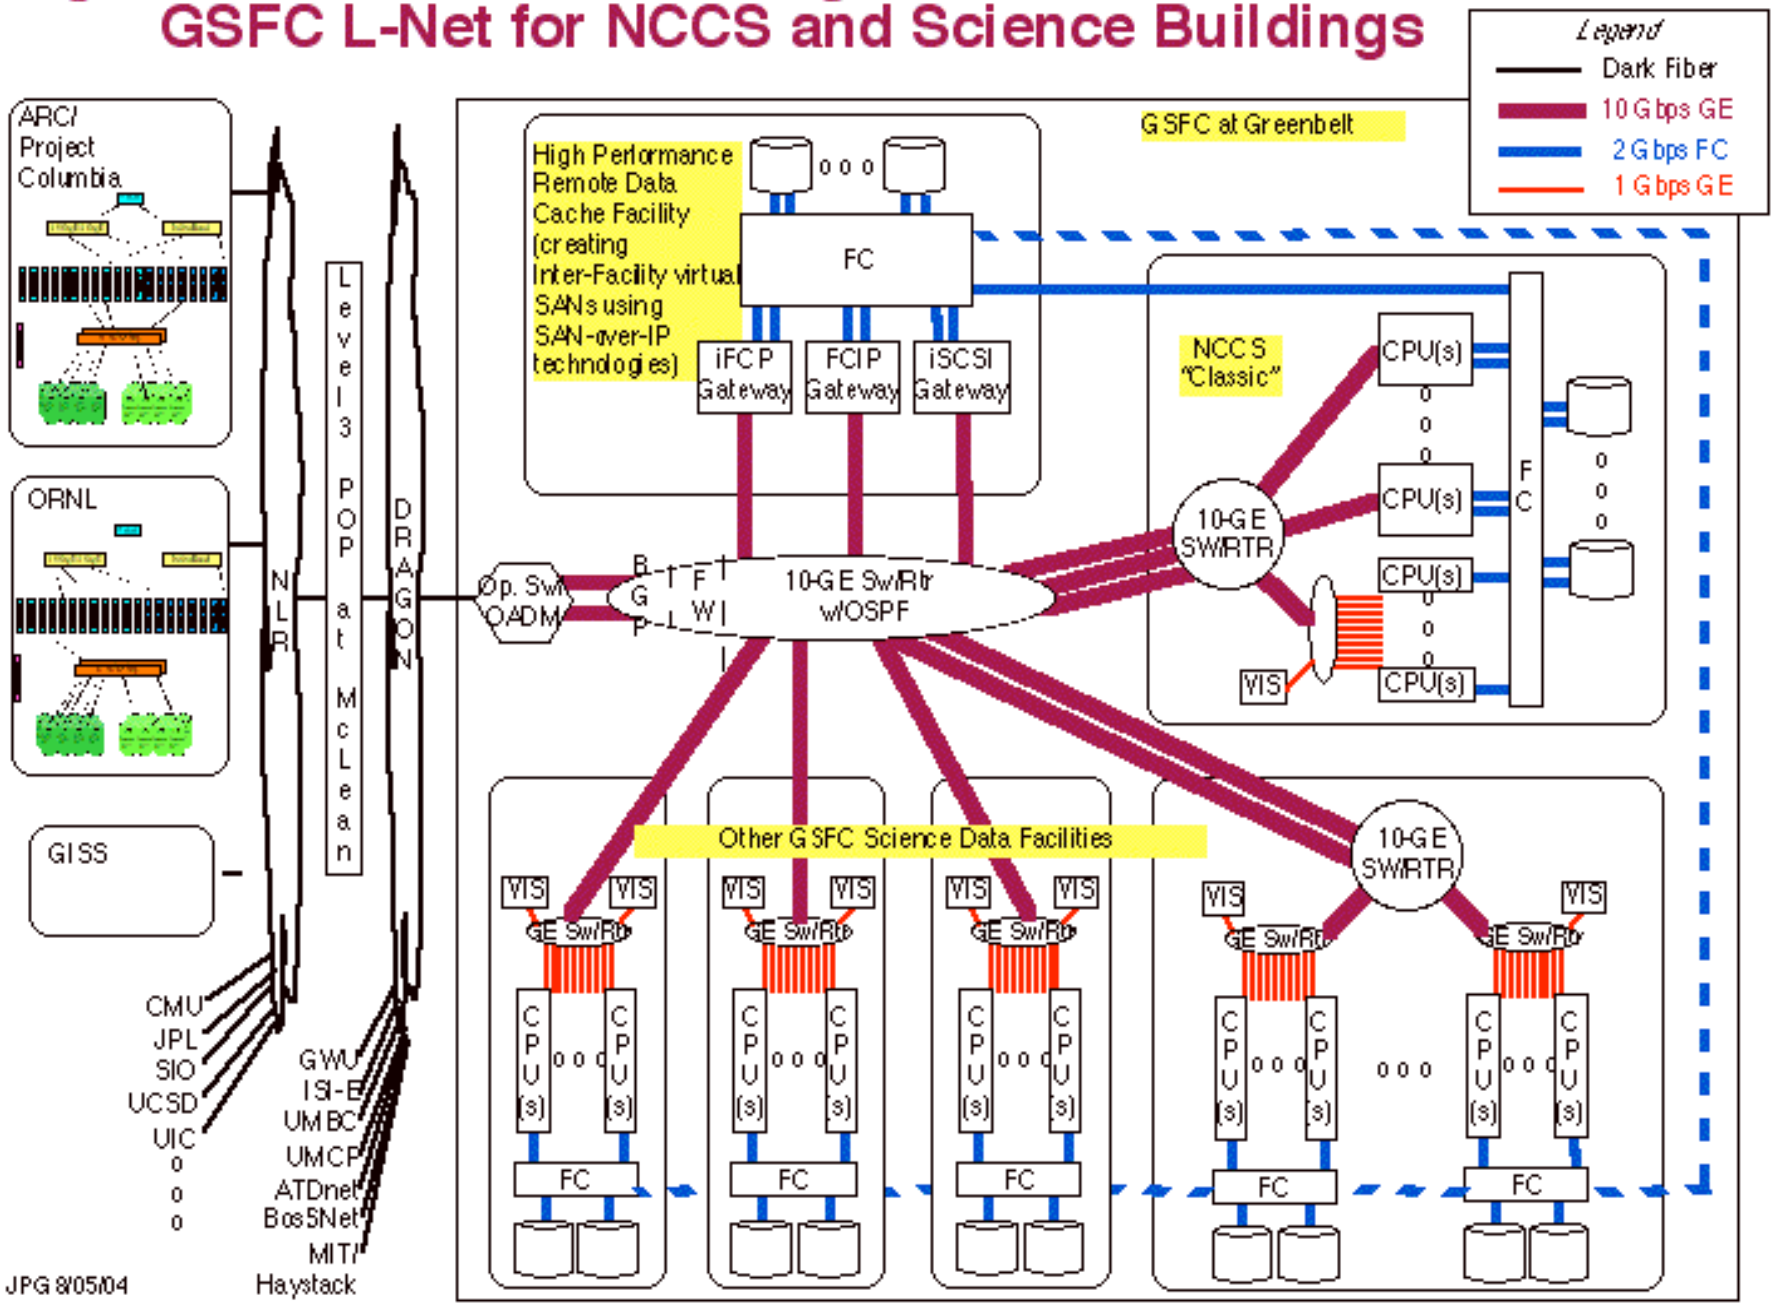

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Special Acknowledgements

GSFC Internal

- **High End Computer Network Team**
  - **Bill Fink/606.1**
  - **Kevin Kranacs/585**
  - **Paul Lang/ADNET/606.1**
  - **Aruna Muppalla/ADNET/606.1**
  - **Jeff Martz/CSC/606.2**
  - **Mike Steffenelli/CSC/606.2**
  - **Kevin Fisher/586/UMBC coop**
- **ESDIS Network Prototyping Lab**
  - **George Uhl/SWALES/423**
- **ESTC Computing Technology Project**
  - **PM: Jim Fischer/606**
- **IT Pathfinder Working Group**
  - **Chair: Dr. Milton Halem/Emeritus & UMBC**
- **Thunderhead Cluster**
  - **John Dorband/696**

GSFC External

- **National LambdaRail**
  - **CEO: Tom West**
  - **Net Eng Lead: Debbie Montano**
- **OptIPuter Project (NSF-funded)**
  - **PI: Dr. Larry Smarr/UCSD**
  - **Co-PI: Dr. Tom DeFanti/UIC**
  - **PM: Maxine Brown/UIC**
  - **UCSD Net Eng: Greg Hidley, Arron Chin, Phil Papodopolos**
  - **UIC Net Eng: Alan Verlo, Linda Winkler**
- **DRAGON Project (NSF-funded)**
  - **PI: Jerry Sobieski/UMCP**
  - **Co-I: Tom Lehman/USC-ISI/E**
  - **Net Eng: Chris Tracy/UMCP**
- **NASA Research and Education Network**
  - **DPM: Kevin Jones/ARC**


NETWORK
BOTTLENECKS

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Backup Slides

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- Multi-channel Collaboration/Video Streaming Technologies
  - Scalable Adaptive Graphics Environment (SAGE) (http://www.evl.uic.edu/cavern/sage)
  - HDTV-over-IP
  - Demonstrations of 21st Century National-Scale Team Science (http://www.calit2.net/newsroom/release.php?id=660)
- 3D HDTV-over-IP
  - 3D Multichannel Networked System via NASA SBIR FY06 Phase2 awardee Physical Optics Corporation
  - Live 3D HDTV multi-Gbps real-time data streaming from GSFC to holographic display at iGrid2005 as the US130/Real-Time_True-3D_Visualization exhibitor (http://www.igrid2005.org/program/applications/vizservices_3dviz.html)

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- Enabling e-VLBI real-time data flows from GGAO to MIT/Haystack (http://web.haystack.mit.edu/e-vlbi/evlbi.html)
- Prototyping of Earth System Modeling Framework (ESMF)-based cross-organization coupling of climate models over a high speed network (http://cisto.gsfc.nasa.gov/L-Netpdfs/sc05_esmf_demo_v5.pdf)
- Evaluating SAN-over-IP and distributed shared file system applicability to enhancing science data flows
  - NCCS' participation in the Data Intensive Computing Environment (DICE) Project (http://www.avetec.org/dice)
  - NCCS data portal environment
  - McCall et al, "A framework for Managing Inter-site Storage Area Networks using Grid Technologies" (http://romulus.gsfc.nasa.gov/msst/conf2006/Papers/2006-025-McCall.pdf)

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Future Plans (partial list)

- New NGC(Colshire) and UMBC DWDM connections to DRAGON
- Leverage existing DRAGON-provided 10-Gbps connection with Internet2's NLR/HOPI lambda
- Support plans identified in NASA NRA Proposals
  - "MAP Core Integration LambdaGrid Infrastructure" by Smarr (UCSD) et al to NASA's MAP NRA
  - "Brokering and Chaining Distributed Services and Data Using OptIPuter and the National Lambda Rail" by Ramapriyan (GSFC) et al to NASA's ROSES NRA
  - "Enabling NASA Applications Across Heterogeneous High Performance Networks" by Habib (CUNY) et al to NASA NNH05ZDA001N-Applied Information Systems Research (a.k.a. ROSES:D3)
- Extend GSFC's existing10 Gbps L-Net to additional GSFC buildings, computers, and users; increase the number and type of GSFC science/exploration research projects that benefit from the increased throughput performance that multi-wavelength optical networking can provide
- Expand SAN-over-IP testing: intra-GSFC, between GSFC-UMCP & GSFC-ARC & …

NLR/GSFC Applications: Hurricane Prediction

- The NASA Finite-Volume General Circulation Model (fvGCM) has been producing real-time, high-resolution (~25 km) weather forecasts focused on improving hurricane track and intensity forecasts.
- During the active 2004 Atlantic hurricane season, the fvGCM provided landfall forecasts with an accuracy of ~100 km up to 5 days in advance.
- The 50–100 Mbps throughput available between fvGCM users at GSFC and the Columbia supercomputer at ARC greatly hindered carrying out time-critical simulations of the hurricanes that devastated Florida.
- The 10 Gbps NLR access will enable remote, 3D visualization analysis as soon as forecast variables become available.
- Key Contacts: Ricky Rood, Bob Atlas, Horace Mitchell, GSFC; Chris Henze, ARC.

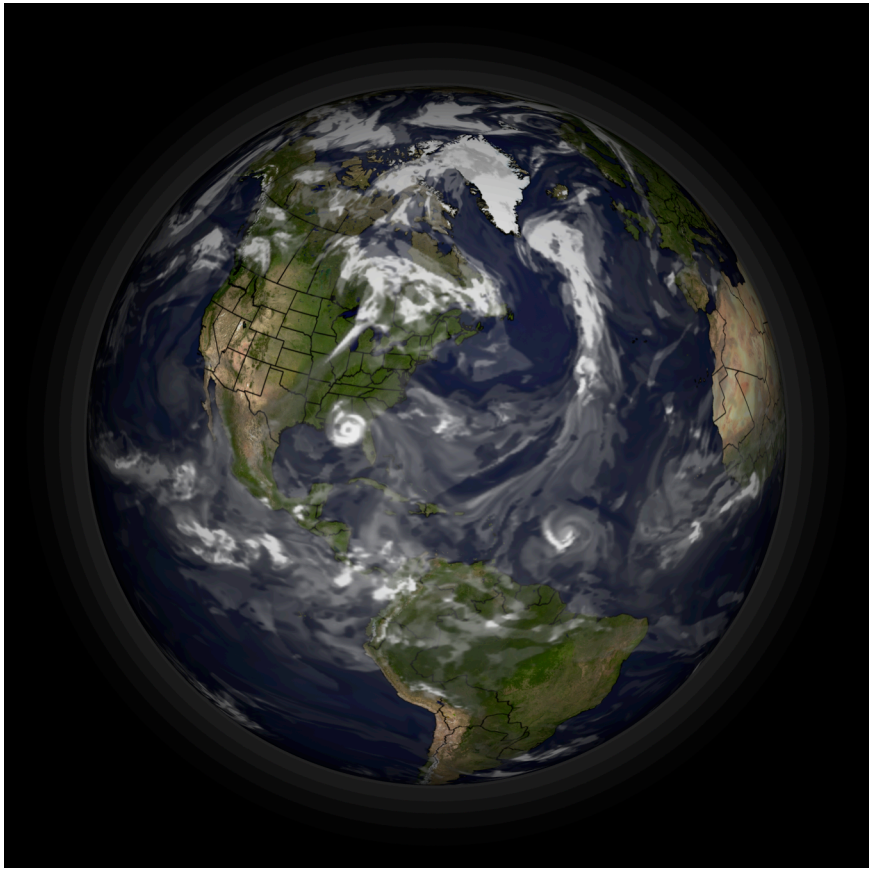

In an fvGCM forecast, Hurricane Frances makes landfall on the Gulf Coast of Florida while Hurricane Ivan intensifies in the tropical Atlantic. Visualization by J. Williams, GST.

http://fvnwp.gsfc.nasa.gov

NLR/GSFC Applications: Global Aerosols

- Project Atmospheric Brown Clouds (ABC) is an international effort to discover and analyze areas of brown colored atmosphere to learn how dust and pollution particles are transported and what impacts they have on the environment, climate, agricultural cycles, and quality of life.
- GSFC and the Scripps Institution of Oceanography (SIO) are planning a collaboration to predict the flow of aerosols from Asia across the Pacific to the U.S. on timescales of days to a week.
- GSFC will provide an aerosol chemical tracer model (GOCAR) embedded in a high-resolution regional model (MM5) that can assimilate data from Indo-Asian and Pacific ground stations, satellites, and aircraft.
- Remote computing and analysis tools running over the NLR will enable acquisition and assimilation of the Project ABC data.
- Key Contacts: Yoram Kaufman, William Lau, GSFC; V. Ramanathan, Chul Chung, SIO.

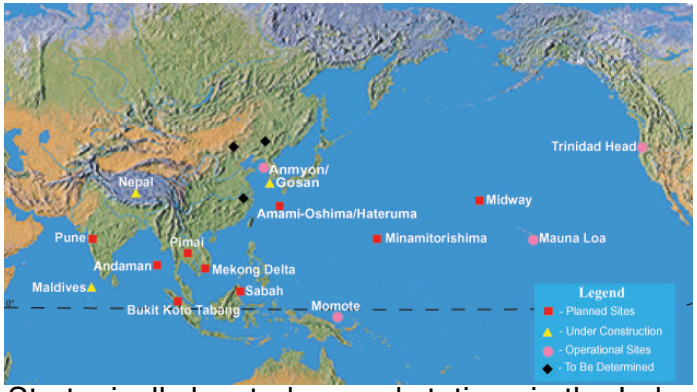

Strategically located ground stations in the Indo-Asian and Pacific regions monitor atmospheric pollution.

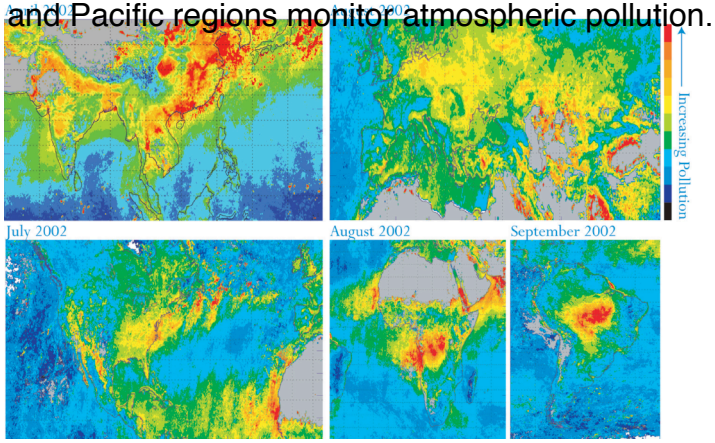

The global nature of brown clouds is apparent in analysis of NASA MODIS Data. Research by V. Ramanathan, C. Corrigan, and M. Ramana, SIO.

http://www-abc-asia.ucsd.edu

NLR/GSFC Applications: Remote Viewing and Manipulation of Large Earth Science Data Sets

- Remote viewing and manipulation of data sets at GSFC and JPL is needed to support EOSDIS and Earth system modeling.
- GSFC's EOSDIS Clearing House (ECHO) and JPL's GENESIS prototype science analysis system (iEarth) will become connected over the NLR. The link will enable comparison of hundreds of terabytes of data, generating large, multi-year climate records.
- Initial work will focus on the Estimating the Circulation and Climate of the Ocean (ECCO) modeling team. Besides ready access to the NLR, the team will need versatile subsetting and other data manipulation functions to reduce compute and bandwidth requirements as well as a set of Grid-accessible statistical analysis and modeling operators to refine and validate the ECCO models.
- Key Contacts: ECHO metadata gateway team, GSFC; GENESIS team, led by Tom Yunck, JPL.

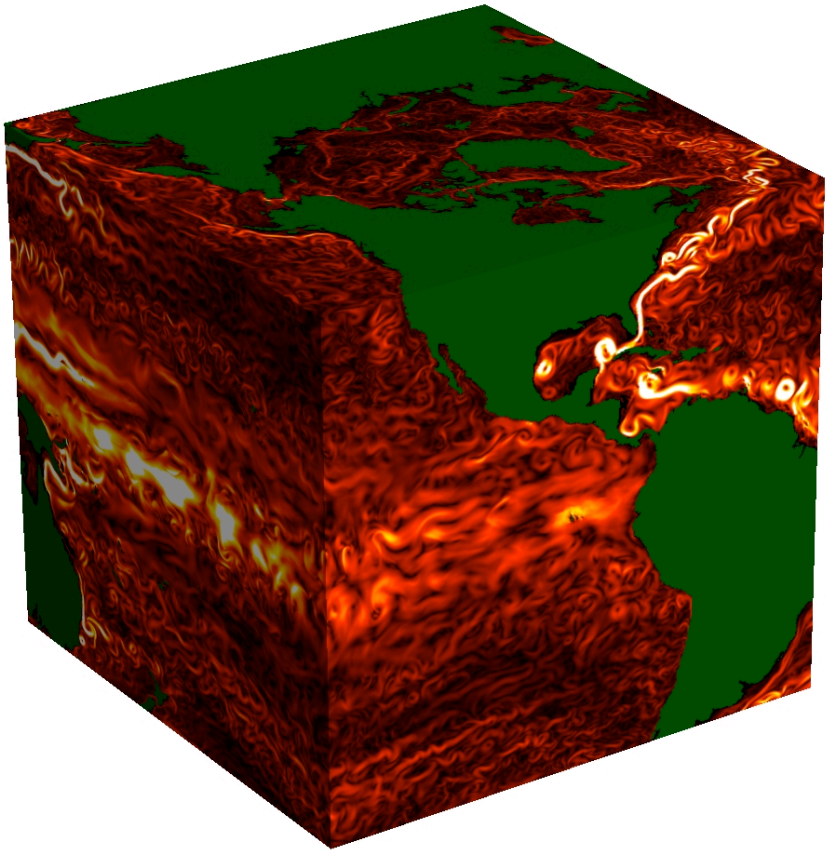

Near-surface (15-m) ocean current speed from an eddy-permitting integration of the cubed-sphere ECCO ocean circulation model. Research by JPL and MIT. Visualization by C. Henze, Ames.

http://www.ecco-group.org

NLR/GSFC Applications: Integration of Laser and Radar Topographic Data with Land Cover Data

- NASA has executed two advanced missions to create an accurate high-resolution topographic model of the Earth: the Shuttle Radar Topography Mission (SRTM) and ICESat, with its Geoscience Laser Altimeter System (GLAS).
- The agency now has the opportunity to merge the two data sets, using SRTM to achieve good coverage and GLAS to generate calibrated profiles. Proper interpretation requires extracting land cover information from Landsat, MODIS, ASTER, and other data archived in multiple DAACs.
- Use of the NLR and local data mining and subsetting tools will permit systematic fusion of global data sets, which are not possible with current bandwidth.
- Key Contacts: Bernard Minster, SIO; Tom Yunck, JPL; Dave Harding, Claudia Carabajal, GSFC.

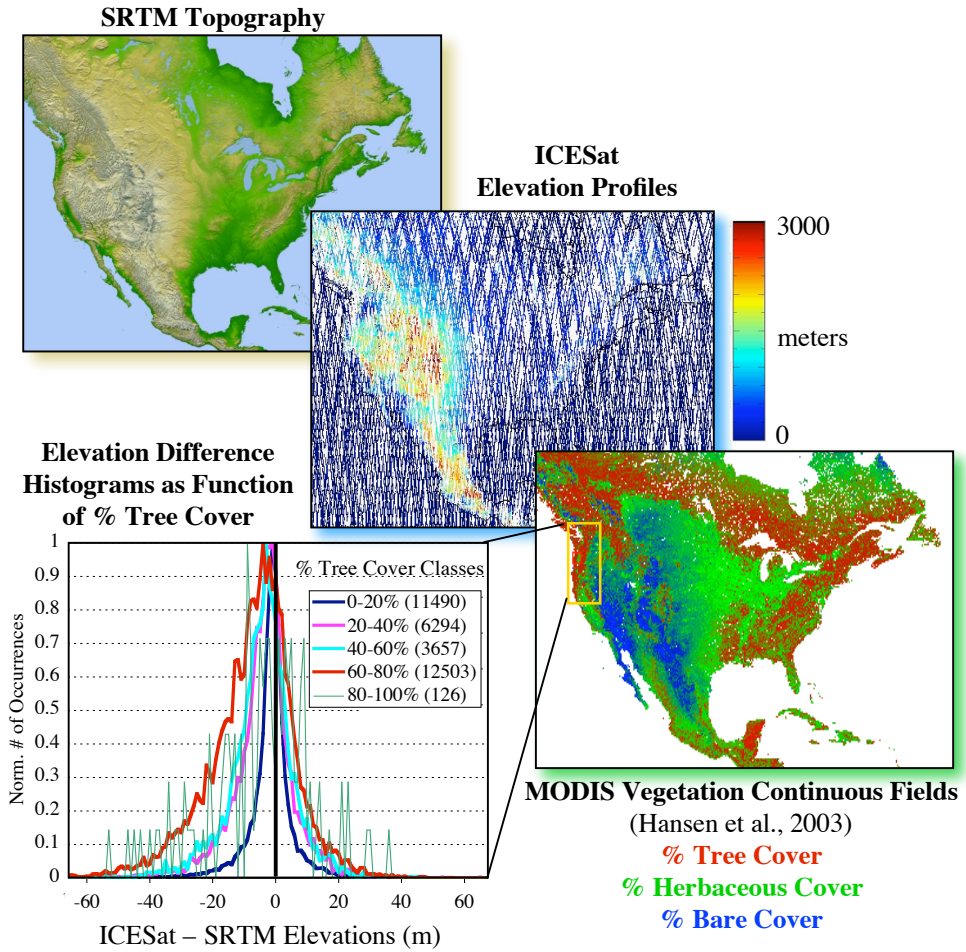

http://icesat.gsfc.nasa.gov
http://www2.jpl.nasa.gov/srtm
http://glcf.umiacs.umd.edu/data/modis/vcf

High speed networking and Grid computing for large-scale simulation in geodynamics

W. Kuang[1], W. Jiang[2], S. Zhou[3], P. Gary[1], M. Seablom[1], W. Truszkowski[1], J. Odubiyi[4], D. Liu[2], J. Palencia[5], G. Gardner[6]

[1] NASA Goddard Space Flight Center, [2] JCET, UMBC, [3] Northrop Grumman IT/TASC, [4] Bowie State University, [5] Raytheon ITSS, [6] INDUSCORP

Introduction

Now large-scale simulation has been wide-spread in many disciplines of solid Earth science research. A typical numerical test in the simulation can easily reach 10^{13} flops and beyond.

One such research problem that we are working on now is to establish a framework on predicting geomagnetic secular variation on decadal and longer time scales, utilizing surface geomagnetic/paleomagnetic records and our MoSST core dynamics model (Figure 1). In this approach, model forecast results and observations are weighted to provide initial state for assimilation (Figure 2). Typically 30 independent numerical tests are necessary for a reasonable ensemble size. This could easily require a computing cycle on orders of petaflops and larger.

A single super-computing facility for such studies is not an optimal choice, due to many limitations, in particular those on user management and administration. But it is relatively easy for users (researchers) to manage because of a unified system environment.

Grid computing can be a much better choice so that independent numerical tests can be carried out independently on different systems. However, researchers (users) have to deal with heterogeneous systems and other problems, such as those on network communication.

In this poster, we discuss our activities in GSFC on application of grid computation to geodynamics modeling.

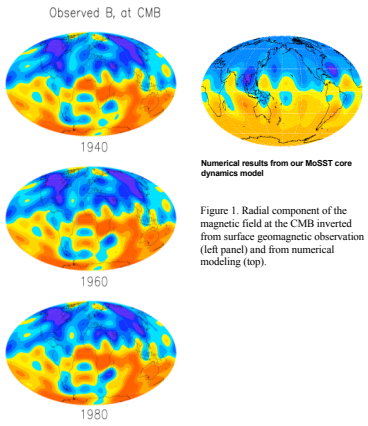

Numerical results from our MoSST core dynamics model

Figure 1. Radial component of the magnetic field at the CMB inverted from surface geomagnetic observation (left panel) and from numerical modeling (top).

Geomagnetic data assimilation

$$\mathbf{x}^a = \mathbf{x}^f + \mathbf{K}\left(\mathbf{x}^o - \mathbf{H}\mathbf{x}^f\right)$$

$\mathbf{X}^a$: Assimilation solution
$\mathbf{X}^f$: Forecast solution
$\mathbf{X}^o$: Observation data

Figure 2. Mathematical foundation of data assimilation. The common gain **K** depends on knowledge of error statistics of observations and of models. If ensemble Kalman-filter approach is applied. An ensemble size of at least 30 (i.e. independent tests) is required.

Prototype on MoSST simulation with independent systems

The objective of this prototype work is to test operability of executing our MoSST core dynamics model on independent computing systems. Individual computing units are slated out from selected components of our beowulf system to mimic independent computing environment. The prototype program for grid computing is built upon xcat3 framework (based on java/python). See Figure 3 for conceptual layout of our prototype experiment.

The sample script and the execution process are shown in Figure 4.

Our prototype experiment is very successful. With this experiment, we can proceed further our test on real remote systems. Also with this experiment, we can identify the needs from the user's considerations on supporting environment and other middleware that makes grid computing "friendly" .

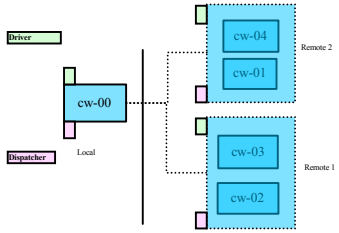

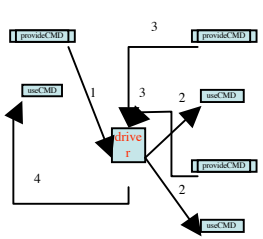

System configuration

OS: Fedora core 2; MPICH-1.2.5.2; Intel Fortran Compiler; Java 2
PE: Dual Intel Xeon, 2.4 Ghz, 1 GB, 1 GigEthernet

Figure 3. Prototype layout

```
import sys
import cca

from jarray import zeros

from java.lang import System
from java.lang import String, Object
# get the absolute location for XCAT
.
# create the TypeMap for the user component
# create component wrappers
geo1Provides = cca.createComponentWrapper("geo1Provider",
geo1ProviderMap)
geo2Provides = cca.createComponentWrapper("geo2Provider",
geo2ProviderMap)
#uses = cca.createComponentWrapper("user", userMap)
# assign a machine name
cca.setMachineName(uses, "cw-00")
cca.setMachineName(provides, "cw-00")
cca.setMachineName(geo1Provides, "cw-03")
cca.setMachineName(geo2Provides, "cw-04")

# set a creation mechanism to in-process
cca.setCreationMechanism(uses, "local")
cca.setCreationMechanism(provides, "local")
cca.setCreationMechanism(geo1Provides, "ssh")
cca.setCreationMechanism(geo2Provides, "ssh")

# connect their ports
cca.connectPorts(uses, "dispatchUsesPort", provides,
"dispatchProvidesPort")
cca.connectPorts(uses, "geo1UsesPort", geo1Provides,
"geo1ProvidesPort")
cca.connectPorts(uses, "geo2UsesPort", geo2Provides,
"geo2ProvidesPort")
# invoke the method
cca.invokeMethodOnComponent(uses,

                            portClassName,
                            portType,
                            methodName,
                            methodParams)
```

```
wjiang@coresun1:~/XCAT_HOME]wjiang@coresun1$ XCAT_HOME]$ xcat.jython.RunScript
/src/jython/samples/seesdigatchgeo2esh.py
.
Executing jython script : src/jython/samples/seesdigatchgeo2esh.py
. 20:00:11.504 main: xcat.coresvc.XCATServiceImpl.java:118 <init> | called with name: CompositionTool
. 20:00:12.501 main: xcat.jython.CompositionTool.java:353 createTypeMap | called
. 20:00:12.504 main: xcat.jython.ComponentWrapper.java:108 getInstanceName | called
. . . . . .
. 20:00:12.555 main: xcat.jython.CompositionTool.java:233 instantiate | called for component: user
. . . . . .
. 20:00:12.591 main: xcat.coresvc.XCATComponentIDServerImpl.java:110 <init> | exporting ComponentID as a
Grid service
.
. 20:00:12.597 main: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:42 setServices | adding
DispatchUseePort
. 20:00:19.228 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:152 startExecuting | Command
/home/wjiang/research/app/test1/run_main_mpi is sent to cluster cw-03
. 20:00:19.250 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:154 startExecuting | Command
/home/wjiang/research/app/test1/run_main_mpi is received!
. 20:00:19.358 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:158 startExecuting | Node
.
. 20:00:19.359 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:161 startExecuting | Command
/home/wjiang/research/app/test1/run_main_mpi is sent to cluster cw-04
. 20:00:19.380 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:163 startExecuting | Command
/home/wjiang/research/app/test1/run_main_mpi is received!
. 20:00:19.691 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:164 startExecuting | Node
.
. . . . . .
. 20:00:19.691 Thread-0: src.dispatch.CMCPortDispatch.java:91 useCMD |
CMCPortDispatch.java: useCMD(...)!
CMCPortDispatch: useCMD() geoCMD is Application is running on Cluster cw-04. Result is stored in
/data/test2
. 20:00:19.693 Thread-0: xcat.ports.BasicPortImpl.java:86 destroyImpl | called; going down in 5 seconds
.
. 20:00:19.579 Thread-0: samples.esmfcoa2.src.useegeo2.UseeDispatchGeo2.java:189 startExecuting | Your
simulation is successful!!!!!
```

Figure 4. Prototype Operation Script (left) and Screen Caption (right)

Discussions

1. Our research on geomagnetic data assimilation can greatly benefit from grid computing.
2. Our prototype experiment is successful and can be readily expanded to systems with identical settings and SSH communication protocol.
3. Our prototype experiment is limited in many areas, such as handling network communication between independent systems (e.g. instant feedback of remote systems to host systems), heterogeneous environment (e.g. prior knowledge on participating systems is necessary), authentication (e.g. prototype cannot handle high level access security requirement). Therefore, further experiment is needed to improve our work, such as integrating our work with other (developed and developing) middleware handling the problems.

Related work at GSFC

There are parallel, but related research going on in GSFC on networking and software development. These research activities are updated in http://esdcd.gsfc.nasa.gov/LNetImplement.html. Recent overview of GSFC research activities is given by Dr. M. Halem and can be found in http://esdcd.gsfc.nasa.gov/L-Netpdfs/ESSAAC_MHpres9904.pdf. Some of the activities listed in the report are shown in Figures, 5 and 6. These activities work towards establishing 21st century cyber infrastructure for large-scale scientific teamwork based on fast network.

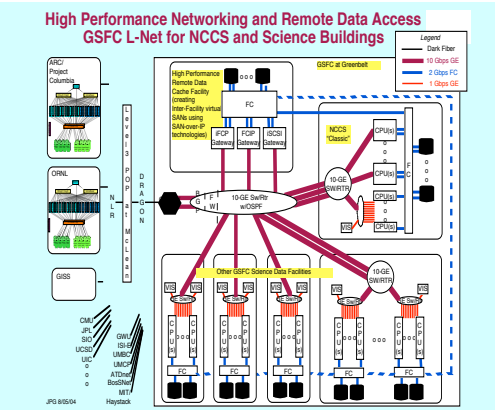

Figure 5. NASA GSFC IRAD work on regional fast network

An Example of Application Requiring L-NET

- Drs. Paul Houser and Mike Bosilovich of Code 970 are collaborating with Dr. John Roads of SIO on the **Coordinated Earth Observing Program (CEOP)** accessing 300TB's of Observational Data in Tokyo and 100TB's of Model Assimilation Data in MPI in Hamburg Germany and analyzing remote data using GRaD-DODS at these sites over the NLR and Starlight.

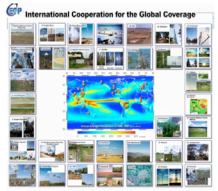

Figure 6.

APPLICATIONS -
Future GRID on 10-GE Network

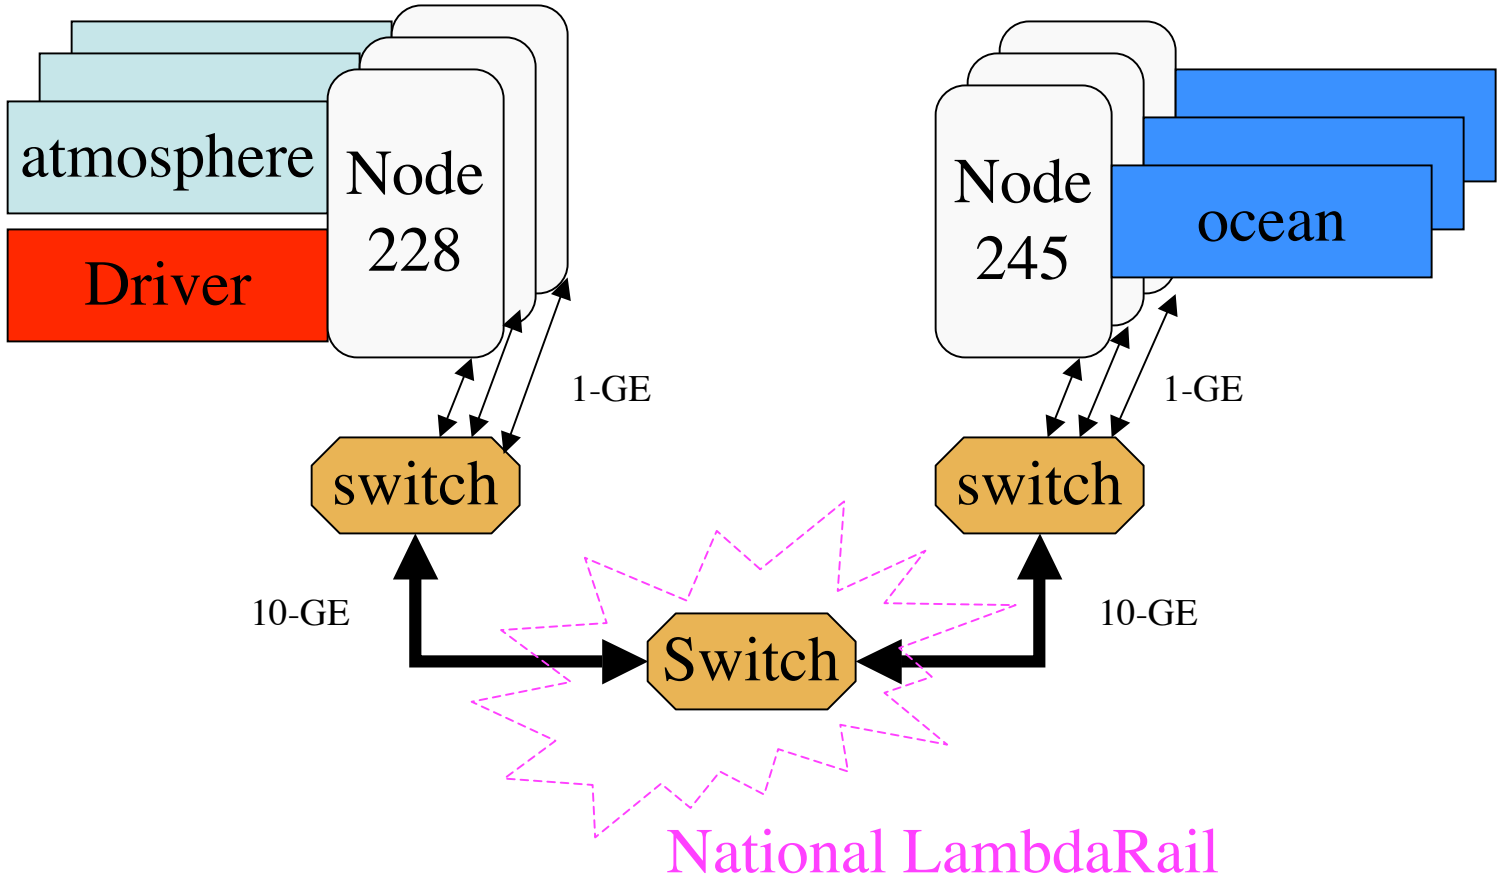

Dr. Zhou is working on applying Grid Computing and High-Speed Network to large-scale distributed computing in Earth and Space Science. More details can be found at http://esto.nasa.gov/conferences/estc2004/papers/a4p1.pdf.

OptIPuter: Empowering Global Data Intensive Services for Science

- Observations
  - The OptIPuter Project Has Moved Beyond “Proof of Concept”
  - Although the OptIPuter is a Technical Achievement, the Research Group Defines Success by Application Achievements
- Goals
  - Provide Persistent, Reliable Unique Services for Data Intensive Sciences and Applications
  - Design, Develop and Deploy Gold Standard OptIPuter Implementation, Which Will Provide Those Services
  - The “Gold Standard” OptIPuter Will Comprise a Distributed Contiguous, Foundation Facility
  - Key Focal Services – e.g., SAGE Distributed via Rocks
  - UltraHD Personal Collaborative Communications Directly Integrated With Extremely Large-Scale, Real-Time, Data Streaming
  - Flexible, Scalable OptIPuter Environments

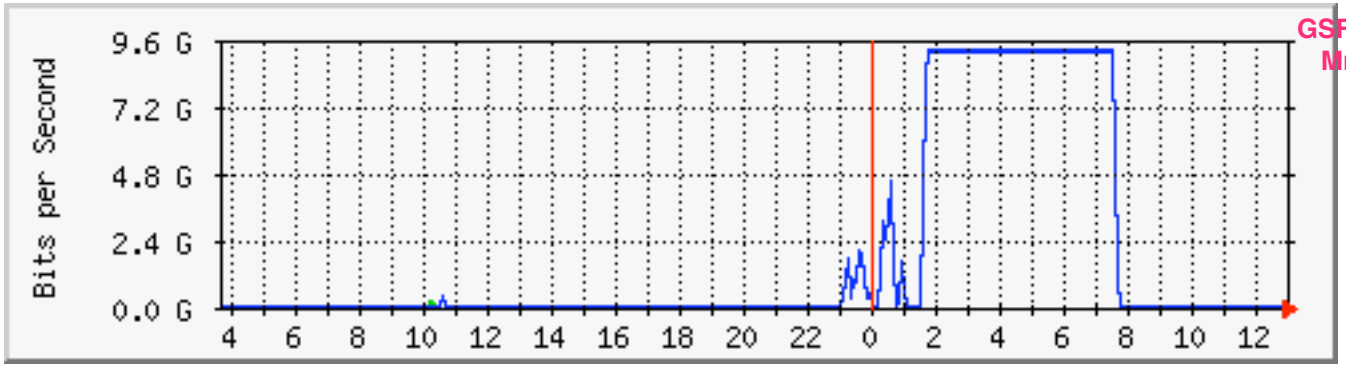

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

NLR/NREN TX 10 GigE

3 February 2006

Max In:	350.4 Mb/s (3.5%)	Average In:	1879.9 kb/s (0.0%)	Current In:	0.0 b/s (0.0%)
Max Out:	9305.3 Mb/s (93.1%)	Average Out:	1779.1 Mb/s (17.8%)	Current Out:	0.0 b/s (0.0%)

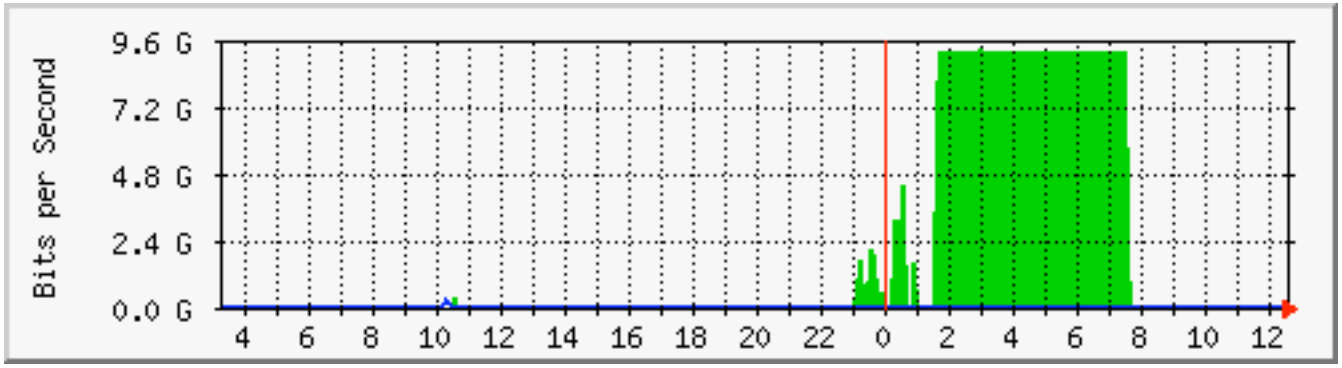

NLR/NREN RX 10 GigE

3 February 2006

Max In:	9316.8 Mb/s (93.2%)	Average In:	1779.1 Mb/s (17.8%)	Current In:	0.0 b/s (0.0%)
Max Out:	298.9 Mb/s (3.0%)	Average Out:	1879.9 kb/s (0.0%)	Current Out:	0.0 b/s (0.0%)

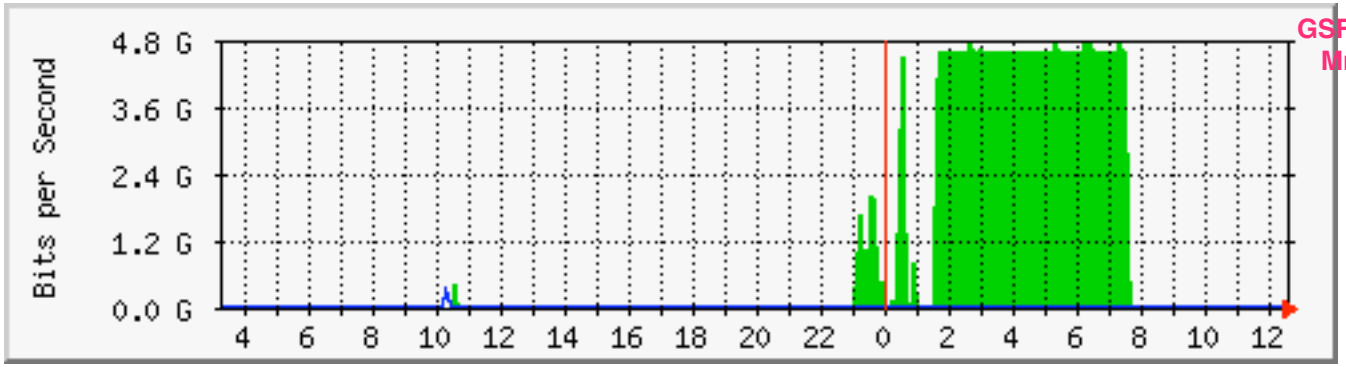

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

chance4 10 GigE

3 February 2006

Max In:	4772.6 Mb/s (47.7%)	Average In:	909.6 Mb/s (9.1%)	Current In:	0.0 b/s (0.0%)
Max Out:	379.1 Mb/s (3.8%)	Average Out:	1826.5 kb/s (0.0%)	Current Out:	8.0 b/s (0.0%)

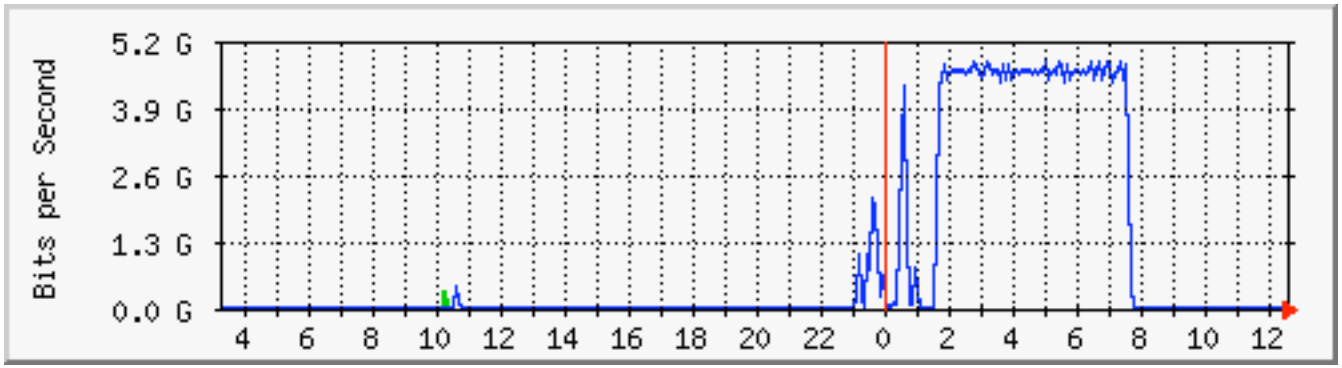

chance5 10 GigE

3 February 2006

Max In:	386.7 Mb/s (3.9%)	Average In:	1815.2 kb/s (0.0%)	Current In:	0.0 b/s (0.0%)
Max Out:	4817.9 Mb/s (48.2%)	Average Out:	907.1 Mb/s (9.1%)	Current Out:	8.0 b/s (0.0%)

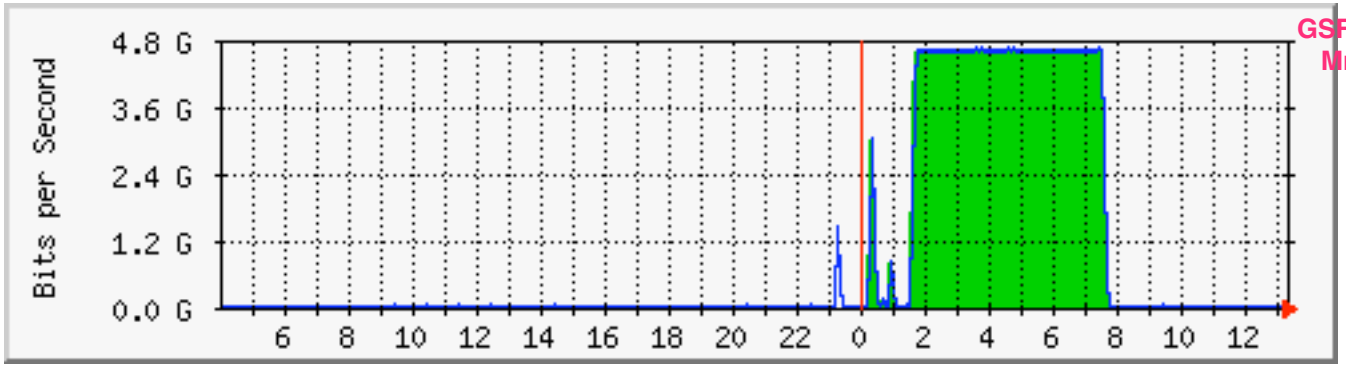

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

GSFC-McLean 10 GigE "Lambda49" across DRAGON

3 February 2006

Max In:	4690.0 Mb/s (46.9%)	Average In:	857.2 Mb/s (8.6%)	Current In:	552.0 b/s (0.0%)
Max Out:	4687.6 Mb/s (46.9%)	Average Out:	861.4 Mb/s (8.6%)	Current Out:	96.0 b/s (0.0%)

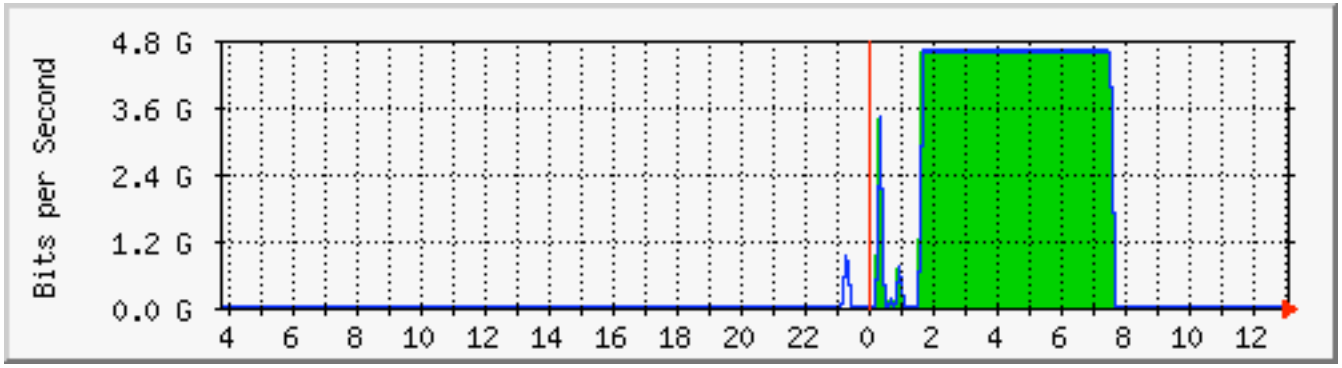

Both chance and chance2 via RAPTOR switch at GSFC 10 GigE

3 February 2006

Max In:	4653.9 Mb/s (46.5%)	Average In:	852.4 Mb/s (8.5%)	Current In:	0.0 b/s (0.0%)
Max Out:	4651.7 Mb/s (46.5%)	Average Out:	856.7 Mb/s (8.6%)	Current Out:	400.0 b/s (0.0%)

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

GSFC Lambda Network Project Website

- http://cisto.gsfc.nasa.gov/IRAD_Lambda.html
  - **Designs**
    - GSFC Local Network Part (i.e., within GSFC)
    - Regional Network Part (i.e., between GSFC in Greenbelt, MD, & Level3 POP in McLean, VA, typically involving the DRAGON optical network)
    - Transcontinental Network Part (i.e., use of NLR, GSFC 10-GE switch & workstations in the Level3 POP in McLean, VA, & remote end users/sites)
  - **Implementation Status**
    - GSFC Local Network Part
    - Regional Network Part
    - Transcontinental Network Part
  - **Presentations/Events in the News**
    - Eg: P. Gary's 18Feb05 presentation at GSFC's FY04 IRAD Colloquium <http://cisto.gsfc.nasa.gov/L-Netpdfs/FY04IRADGARY.pdf>
    - Live Demonstration of 21st Century National-Scale Team Science <http://www.calit2.net/articles/article.php?id=660>
  - **Related Links (e.g., DRAGON, HOPI, NLR, OptIPuter, …)**