GSFC High End Computer Network (HECN) Availability

Topics

- **Brief Intro to HECN at GSFC**
- **Recent GSFC HECN Applications Support**
  - **SC06 Demos**
  - **Achieving Network Throughput Performance Gains**

J. Patrick (Pat) Gary
Network Projects Leader
Networks and Information Technology Security Group (Code 606.1)
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center

For IA@G January 22, 2007 Meeting at GSFC

GSFC Managed Networks

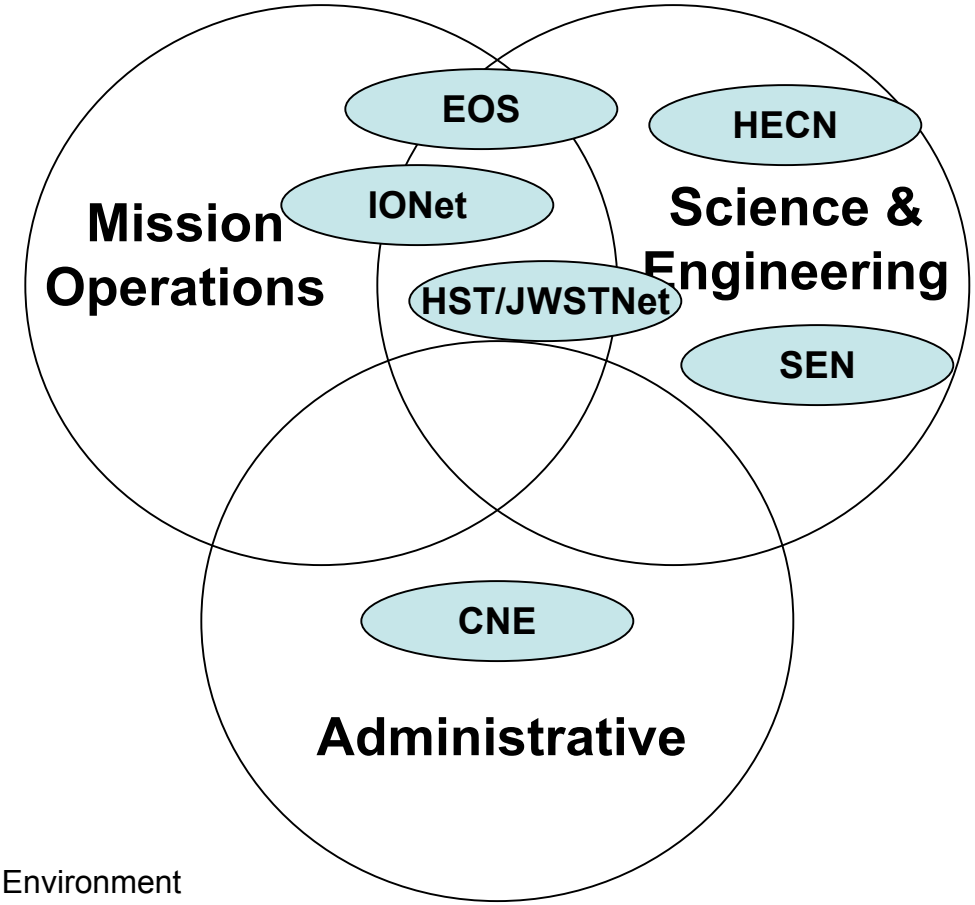

CNE: Center Network Environment
EOS: Earth Observing System
HECN: High End Computing Network
HST/JWSTNet:: Hubble Space Telescope/James Webb Space Telescope Network
IONet: IP Operational Network
SEN: Science & Engineering Network

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

[Note: Some data flows/paths are restricted by GSFC security policy and/or management agreement]

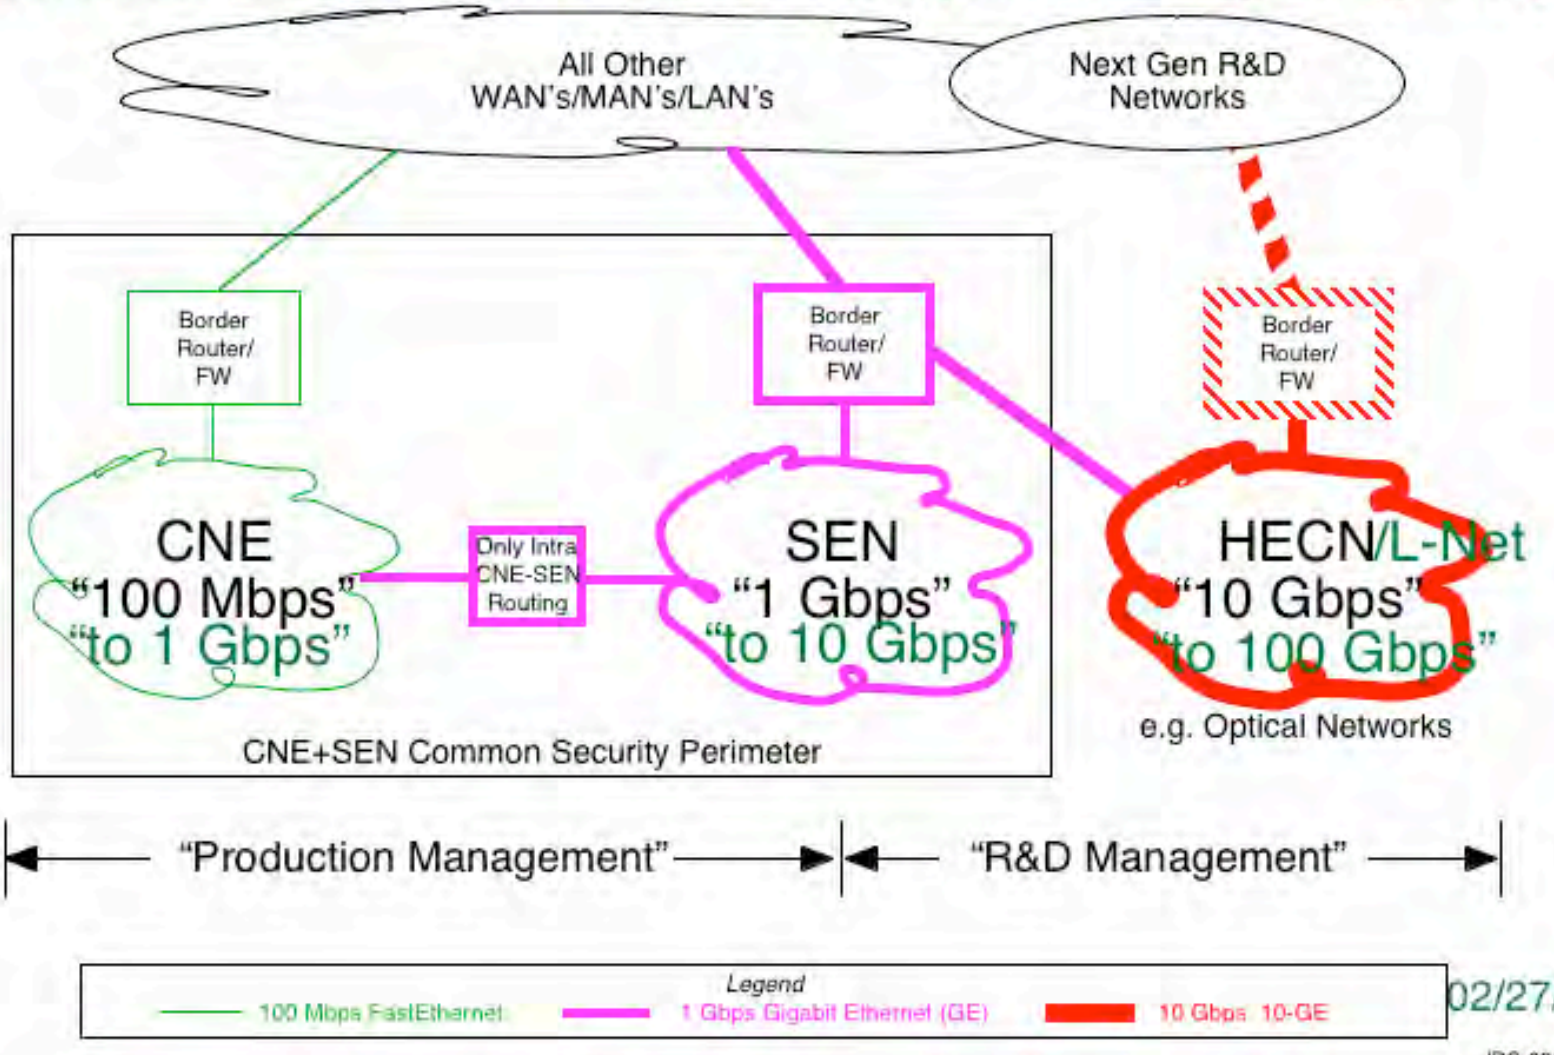

02/27/06

JPG 02/25/04

GSFC Scientific and Engineering Network (SEN) Major Links

Circa 1 ~~March~~ October 2006

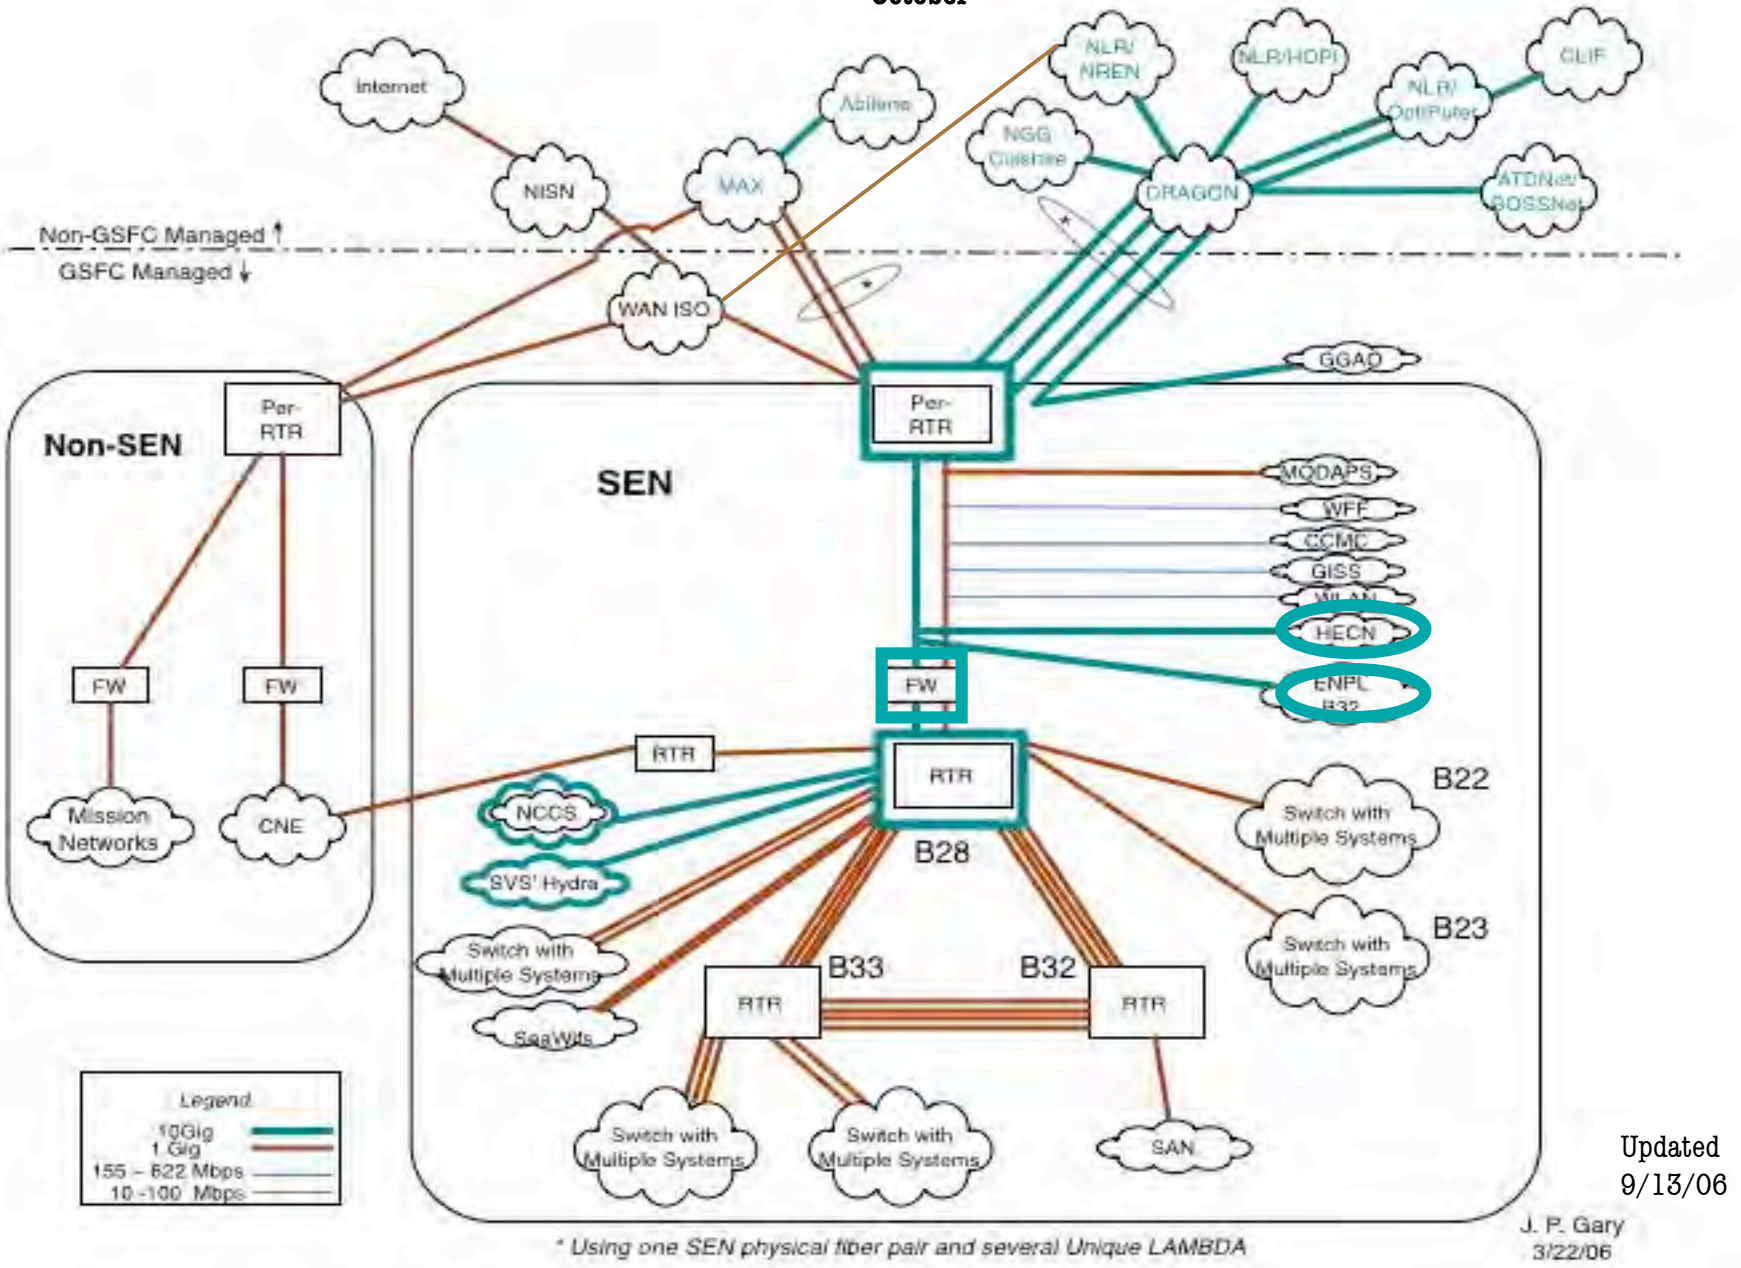

* Using one SEN physical fiber pair and several Unique LAMBDA

J. P. Gary
3/22/06

Updated
9/13/06

Security Services of GSFC SEN and HECN Summarized for GSFC DCIO for EA (7/7/06)

http://cisto.gsfc.nasa.gov/SENuserdocs/SEN_Network_Security_070706.pdf

- Infrastructure-Oriented
  - Perimeter Control
  - Hardware/Software Maintenance
  - Authentication and Configuration Control
  - Staff Training
- User-Oriented
  - Subnetting and VLANs
  - Firewall Management
  - Network Scans
  - Patch Monitoring via PatchLink
  - NIST SP 800-47 Compliant Interconnection Agreements in Testbeds for Advanced IP Services

Selected Security Services of GSFC SEN and HECN

- Firewalls and/or Access Control Lists (ACLs) for each GSFC-managed security domain/zone
  - Compliant with policies generated by GSFC's PCB
  - Moderately segmented network
    - ~15 subnets
    - ~50 VLANs
- Formal firewall waiver request process (similar to the CNE's)
- Separate subnets and/or VLANs for interfaces to console/management ports
- Kerberos V5 software-based authentication and access controls
- MOU with Code 700's Network Security Monitoring Team

Potential GSFC SEN and HECN User Advance Topics for a Future IA@G Meeting

Previous and/or On-Going Applications Support

- Using ARC/NAS/Columbia Supercomputer (w/NREN)
- Distributed ESMF Computing R&D (w/Code 610.3)
- eVLBI (w/MIT-Haystack, …)
- OptIPuter & Multi-channel Collaboration/Video Streaming Technologies(w/UCSD & UIC)
- 3D HDTV-over-IP R&D (w/Physical Optics Corporation)
- SAN-over-IP (w/UMIACS, NGC & NCCS/DICE)

SC06 Demos Supported By GSFC's HECN

- **DRAGON's XNET Demo**
  - Ability to dynamically establish application specific networks that exhibit deterministic, predictable, and repeatable performance characteristics
  - On demand provisioning of optical lambda and VLAN layer network services linking to facilities in Japan, Europe, and across the US to create a */dedicated/* distributed environments for scientific collaboration
  - http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/SuperComputingPlanning2006
- **TeraFlow Testbed Demo**
  - An international application testbed for exploring, integrating, analyzing, and detecting changes in massive and distributed data over wide area high performance networks
  - http://www.ncdm.uic.edu/ & http://sdss.ncdm.uic.edu/
- **OptIPuter Demo**
  - The California Institute for Telecommunications and Information Technology (Calit2), the Center for Earth Observations and Applications (CEOA), the National Center for Microscopy and Imaging Research (NCMIR), and the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago present collaborative research on sensor networks and instrument grids.
  - http://iebms.heiexpo.com/iebms/oep/oep_p2_details.aspx?sessionid=ejnff5ei1fb6fg7ei8&OrderNbr=1626&rescode=3101X62&newrestype=3101
- **DICE Demo**
  - Live data intensive computing environment between multiple booths
  - http://www.avetec.org/dice/SC06_overview.htm

GSFC HECN Team Supports Four Realtime Demonstrations at SC2006

- **GSFC's High End Computer Network (HECN) Team supported four realtime high performance networking data flow demonstrations into the showroom floor of the International Conference for High Performance Computing, Networking and Storage, a.k.a. SC2006, hosted in Tampa, FL, November 11-17, 2006.**
- **The demos supported were those of the following projects:**
  - **DRAGON: http://dragon.maxgigapop.net/**
  - **TeraFlow Testbed: http://www.teraflowtestbed.net/**
  - **OptIPuter: http://www.optiputer.net/**
  - **DICE: http://www.avetec.org/dice/**
- **The provided support was in the form of either HECN's physical network infrastructure used in the critical path of a demo's realtime data flows as illustrated in the top right figure, or HD video streaming and network engineering or troubleshooting expertise to help setup the demo as illustrated in the bottom right figure.**
- **Additional information about the projects supported, their SC06 demos, and the data flows across the relevant network infrastructure used is provided at http://cisto.gsfc.nasa.gov/private_implement/SC06_GSFC_netsupport.pdf.**

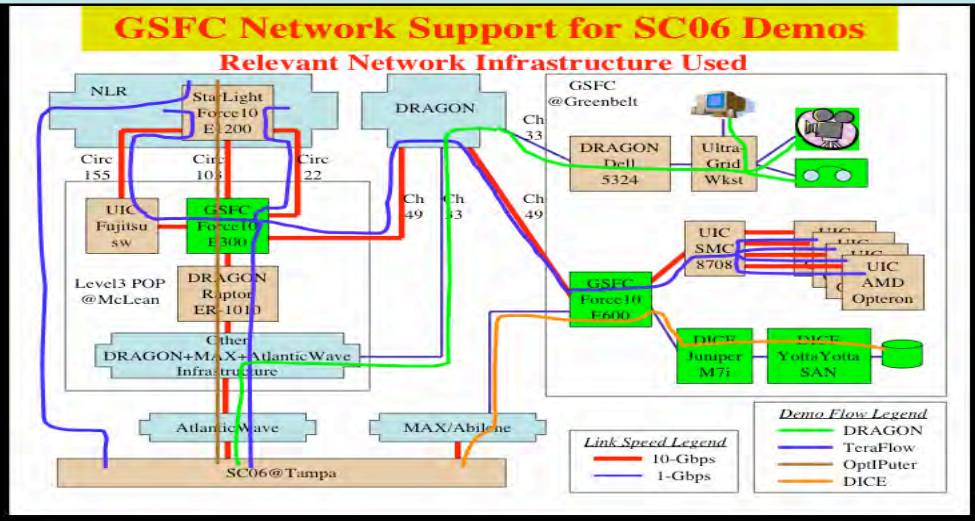

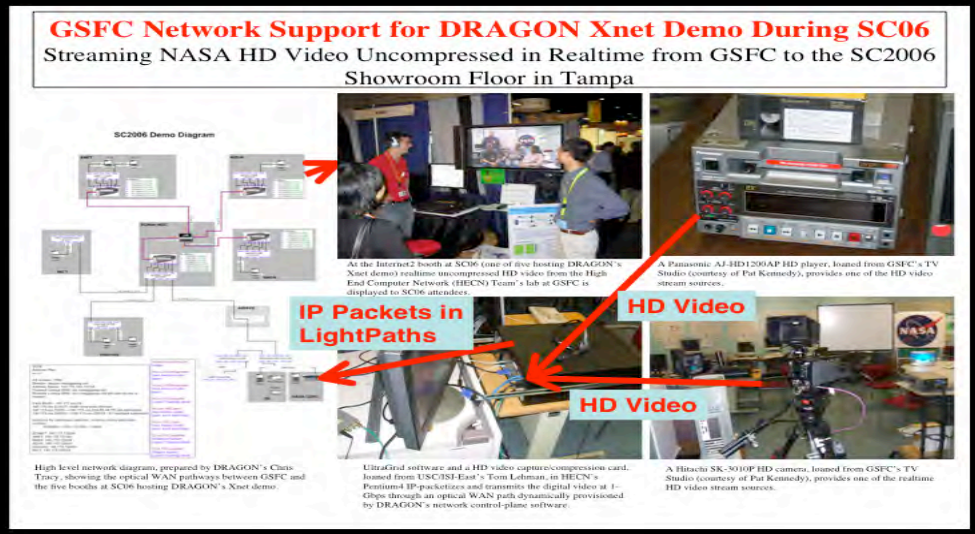

GSFC Network Support for DRAGON Xnet Demo During SC06

Streaming NASA HD Video Uncompressed in Realtime from GSFC to the SC2006 Showroom Floor in Tampa

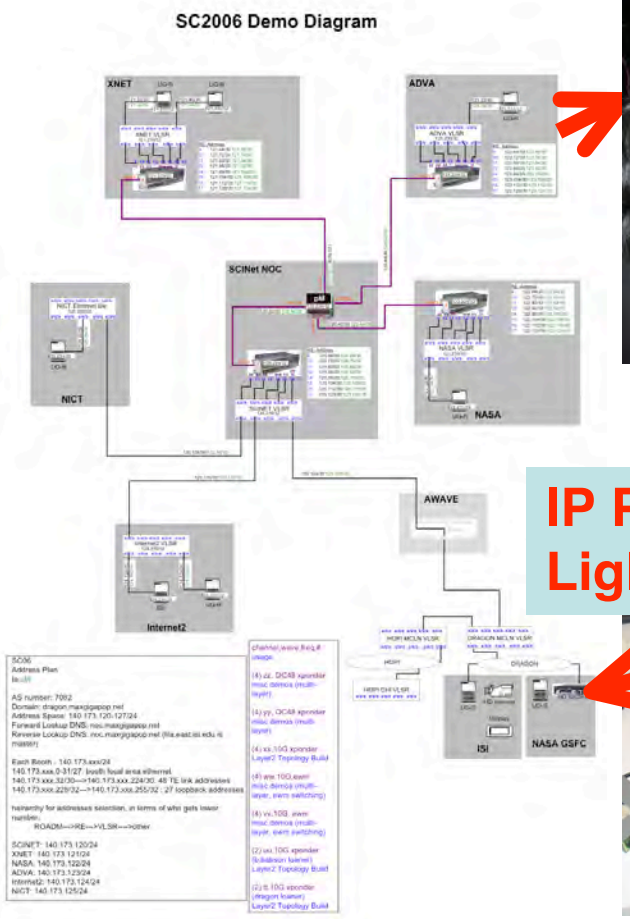

High level network diagram, prepared by DRAGON's Chris Tracy, showing the optical WAN pathways between GSFC and the five booths at SC06 hosting DRAGON's Xnet demo.

At the Internet2 booth at SC06 (one of five hosting DRAGON's Xnet demo) realtime uncompressed HD video from the High End Computer Network (HECN) Team's lab at GSFC is displayed to SC06 attendees.

A Panasonic AJ-HD1200AP HD player, loaned from GSFC's TV Studio (courtesy of Pat Kennedy), provides one of the HD video stream sources.

IP Packets in LightPaths

HD Video

HD Video

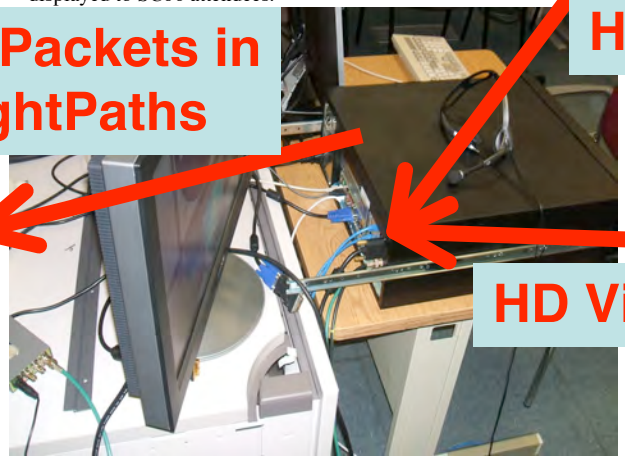

UltraGrid software and a HD video capture/compression card, loaned from USC/ISI-East's Tom Lehman, in HECN's Pentium4 IP-packetizes and transmits the digital video at 1-Gbps through an optical WAN path dynamically provisioned by DRAGON's network control-plane software.

A Hitachi SK-3010P HD camera, loaned from GSFC's TV Studio (courtesy of Pat Kennedy), provides one of the realtime HD video stream sources.

GSFC Network Support for SC06 Demos

Relevant Network Infrastructure Used

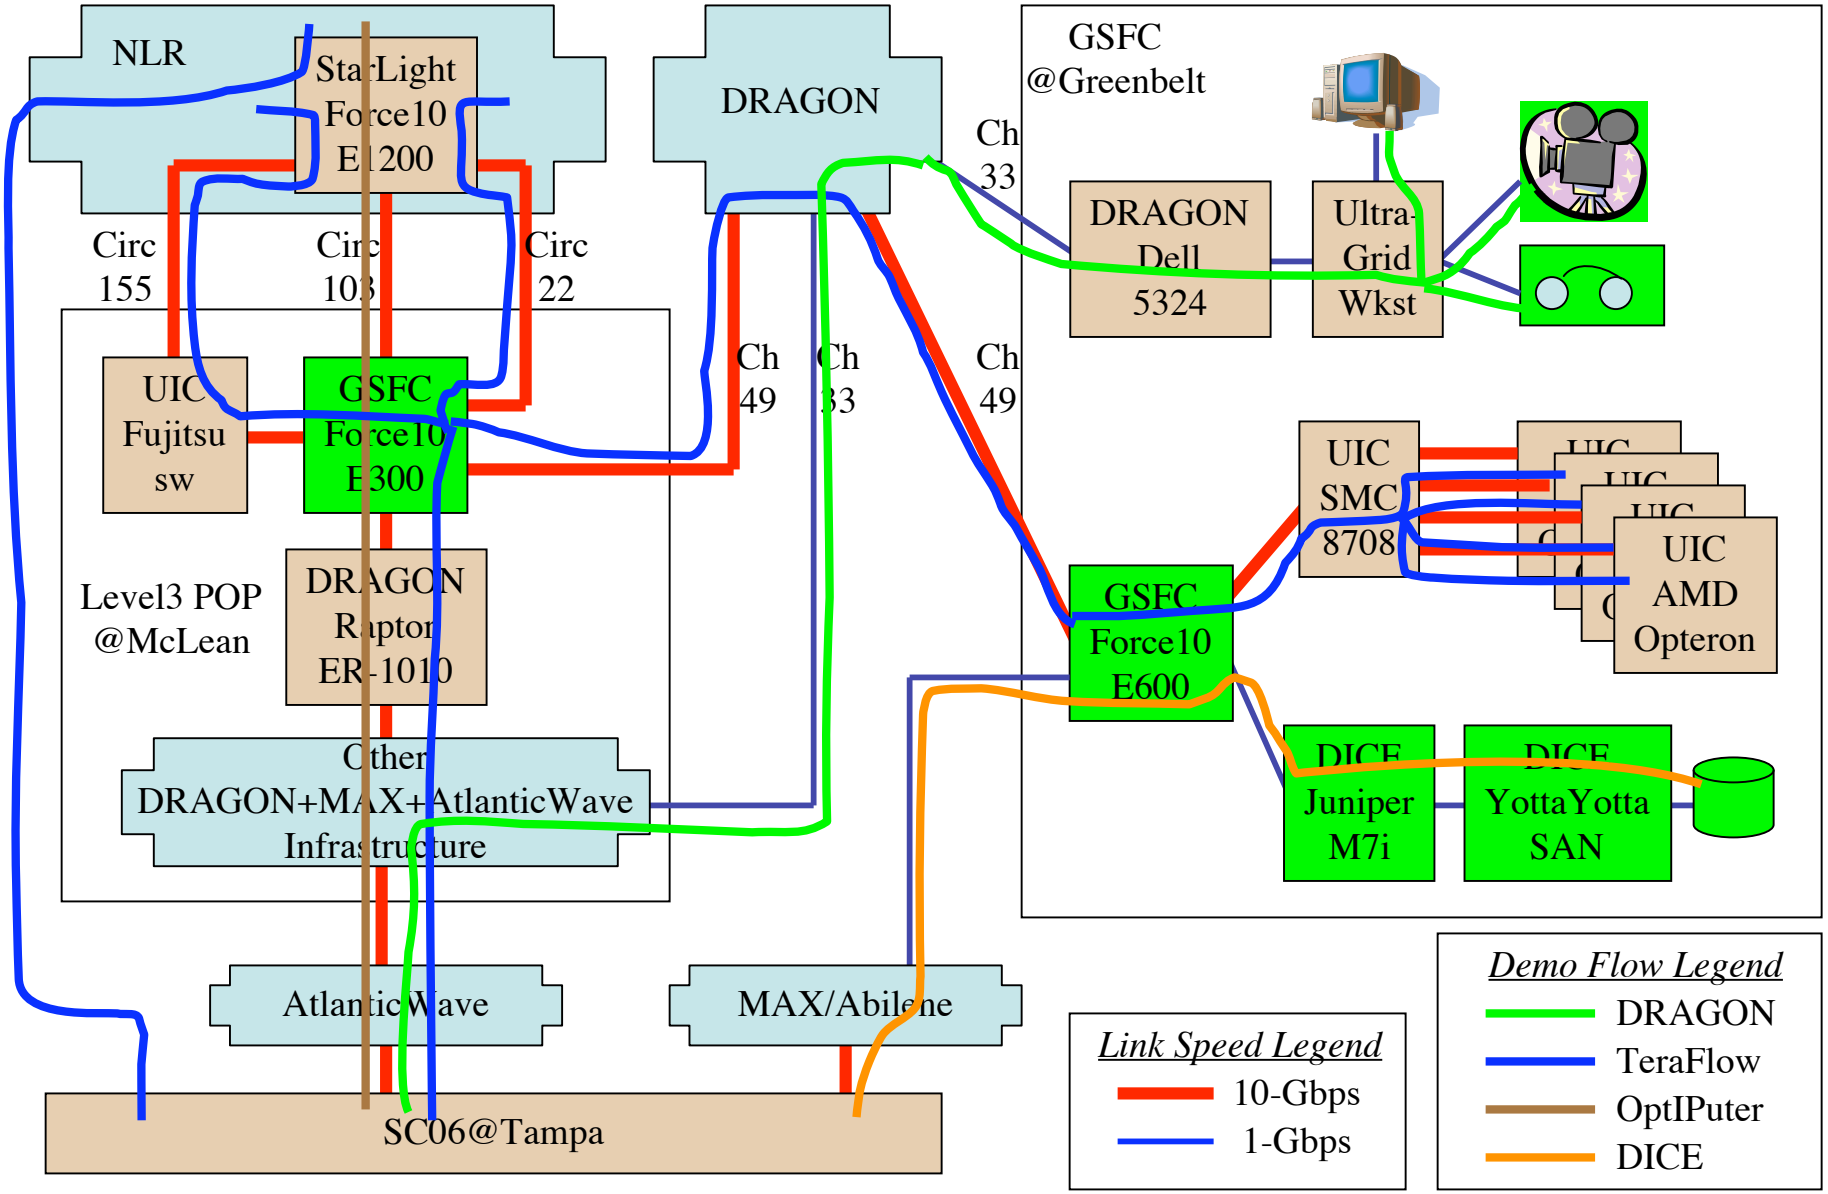

GSFC Network Support for SC06 Demos

Combined TeraFlow and OptIPuter Data Flows to/from SC06

GSFC High End Computer Network (HECN)
Mrtg-based Graphs
Bits per second In and Out
On Selected Interfaces

"Weekly"
30 Minute Averages
16 November 2006

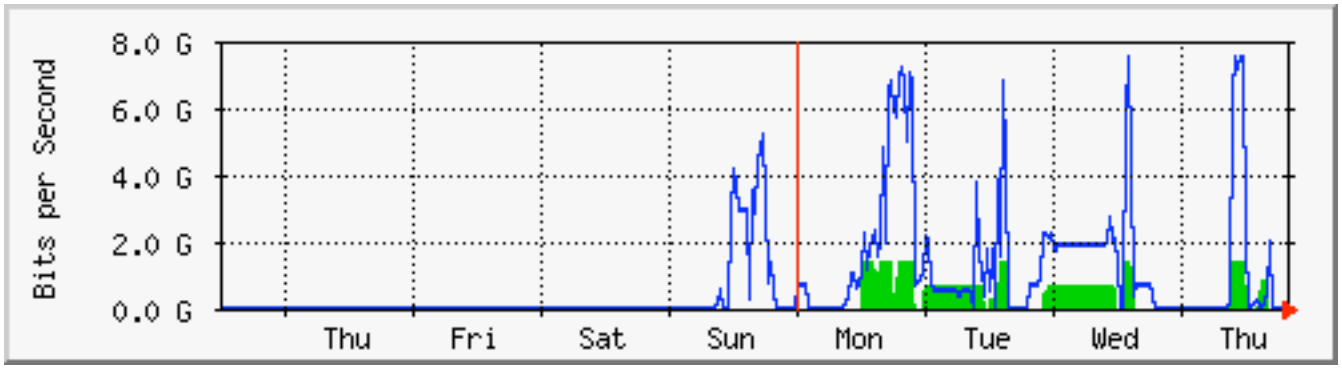

Max In:1515.8 Mb/s (15.2%) Average In:198.2 Mb/s (2.0%) Current In:56.0 b/s (0.0%)
Max Out:7533.7 Mb/s (75.3%) Average Out: 795.4 Mb/s (8.0%) Current Out:0.0 b/s (0.0%)

Measured At:

GSFC/HECN's Force10 E300 10-GE Interface with DRAGON's Raptor (and then AtlanticWave) in Level3 POP at McLean

54 Times Network Throughput Performance Gains Through Improved File Transfer Tool and End System Host/Network Adjustments

Team members from the NASA Research and Engineering Network (NREN) at Ames and the Science and Engineering Network (SEN) at Goddard worked together to help network users from Goddard's 3-D Cloud-Resolving Model project increase data transfer performance through improved file transfer tools and end system host adjustments.

The 3-D Cloud-Resolving Model team had been utilizing the standard file transfer application called Secure Copy (SCP), resulting in maximum sustainable data transfer rates of 5.7 Mbps, between Goddard and Ames. By utilizing the improved multi-stream file transfer application, BBFTP, these data transfer rates were improved to a maximum sustainable data transfer rate of 169 Mbps. Furthermore, by moving the user host to the Science and Engineering Network at Goddard allowed for the application of Jumbo Frames. This improvement resulted in an improved maximum sustainable data transfer rate of 308 Mbps, where the disk I/O speed of the user's local desktop machine is now suspected to be the limiting factor.

Source: Ken Freeman (ARC)

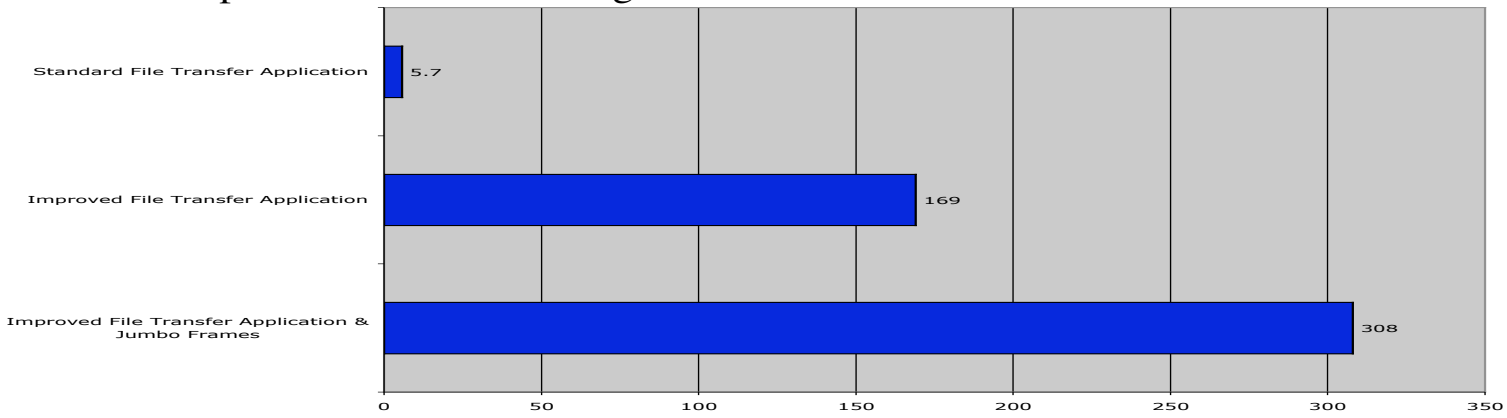

54 Times Network Throughput Performance Gains Through Improved File Transfer Tool and End System Host/Network Adjustments

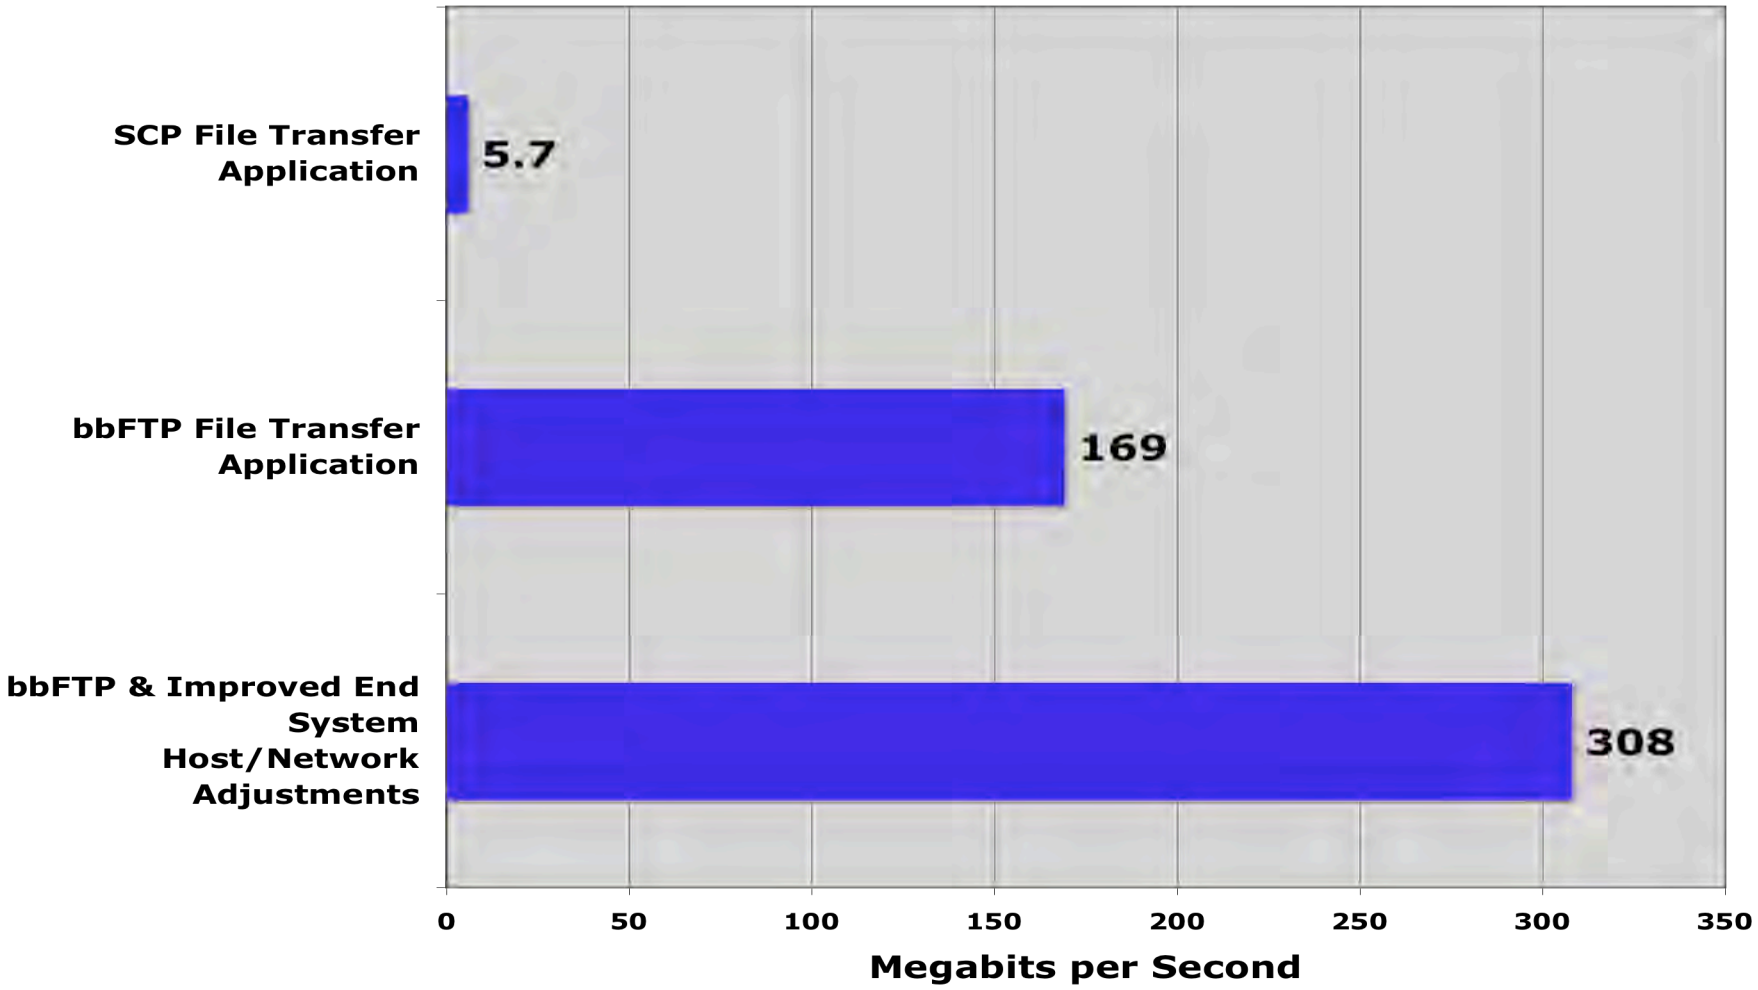

54 Times Network Throughput Performance Gains Through Improved File Transfer Tool and End System Host/Network Adjustments

From	***To***
Change 1	
• SCP	o bbFTP
– Encrypts all data, wasting cpu	- Encrypts only the user's password
– Small transport window hard-coded "defaultly"	- Large transport window via user-provided parameter (UPP)
– [Typically single stream flows]	- [Multi-stream flows via UPP]
• TCP window defaulted to 64KB	o Tuned to Bandwidth x Delay
Change 2	
• Max 1500 byte standard Ethernet frame size via Apple's 1-GE NIC	o Max 9000 byte jumbo frame size via new 1-GE NIC (Intel Pro/1000)
• CNE intra-building, inter-building & firewall infrastructure	o SEN intra-building, inter-building & firewall infrastructure

Excerpts from SEN's User-oriented Webpages at http://cisto.gsfc.nasa.gov/SENuserdocs/SENuser.html#Throughput

Throughput Performance Tuning Information

- System Specific Notes for System Administrators (and Privileged Users): http://www.psc.edu/networking/projects/tcptune/
- TCP Tuning Guide: http://www-didc.lbl.gov/TCP-tuning
- M. Mathis, et al, "NPAD/pathdiag unleashed", ESCC/Internet2 Joint Techs Workshop, Madison, July, 2006: http://events.internet2.edu/2006/jt-madison/sessionDetails.cfm?session=2753&event=253
- Phil Dykstra's tutorial at SC06: http://cisto.gsfc.nasa.gov/SENuserdocs/M07_tutorial.pdf
- Significant improvements in network throughput performance obtained by Code 613.1's Roger Shi's "hurricane" workstation/server: http://cisto.gsfc.nasa.gov/SENuserdocs/BBFTP-Jumbo-HL.110806.pdf
- NREN-provided info on Jumbo Frames, TCP Performance Tuning on End Systems, Multistream File Transfers (e.g., bbFTP), etc.: http://www.nren.nasa.gov/customer.php


NETWORK
BOTTLENECKS

Additional Information

- **SEN:**
  http://cisto.gsfc.nasa.gov/SENuserdocs/SENuser.html

- **HECN:**
  http://cisto.gsfc.nasa.gov/IRAD_Lambda.html

- **For further information about the SEN's or HECN's goals and present capabilities or about new users being attached to the SEN or HECN, please contact:**
  - **J. Patrick Gary (606.1), SEN Engineering Board Chair and HECN Project Manager/Leader**
  - **301-286-9539** **pat.gary@nasa.gov**

GSFC High End Computer Network (HECN) Availability

Backup Slides

The DRAGON Testbed

Washington, DC metro region

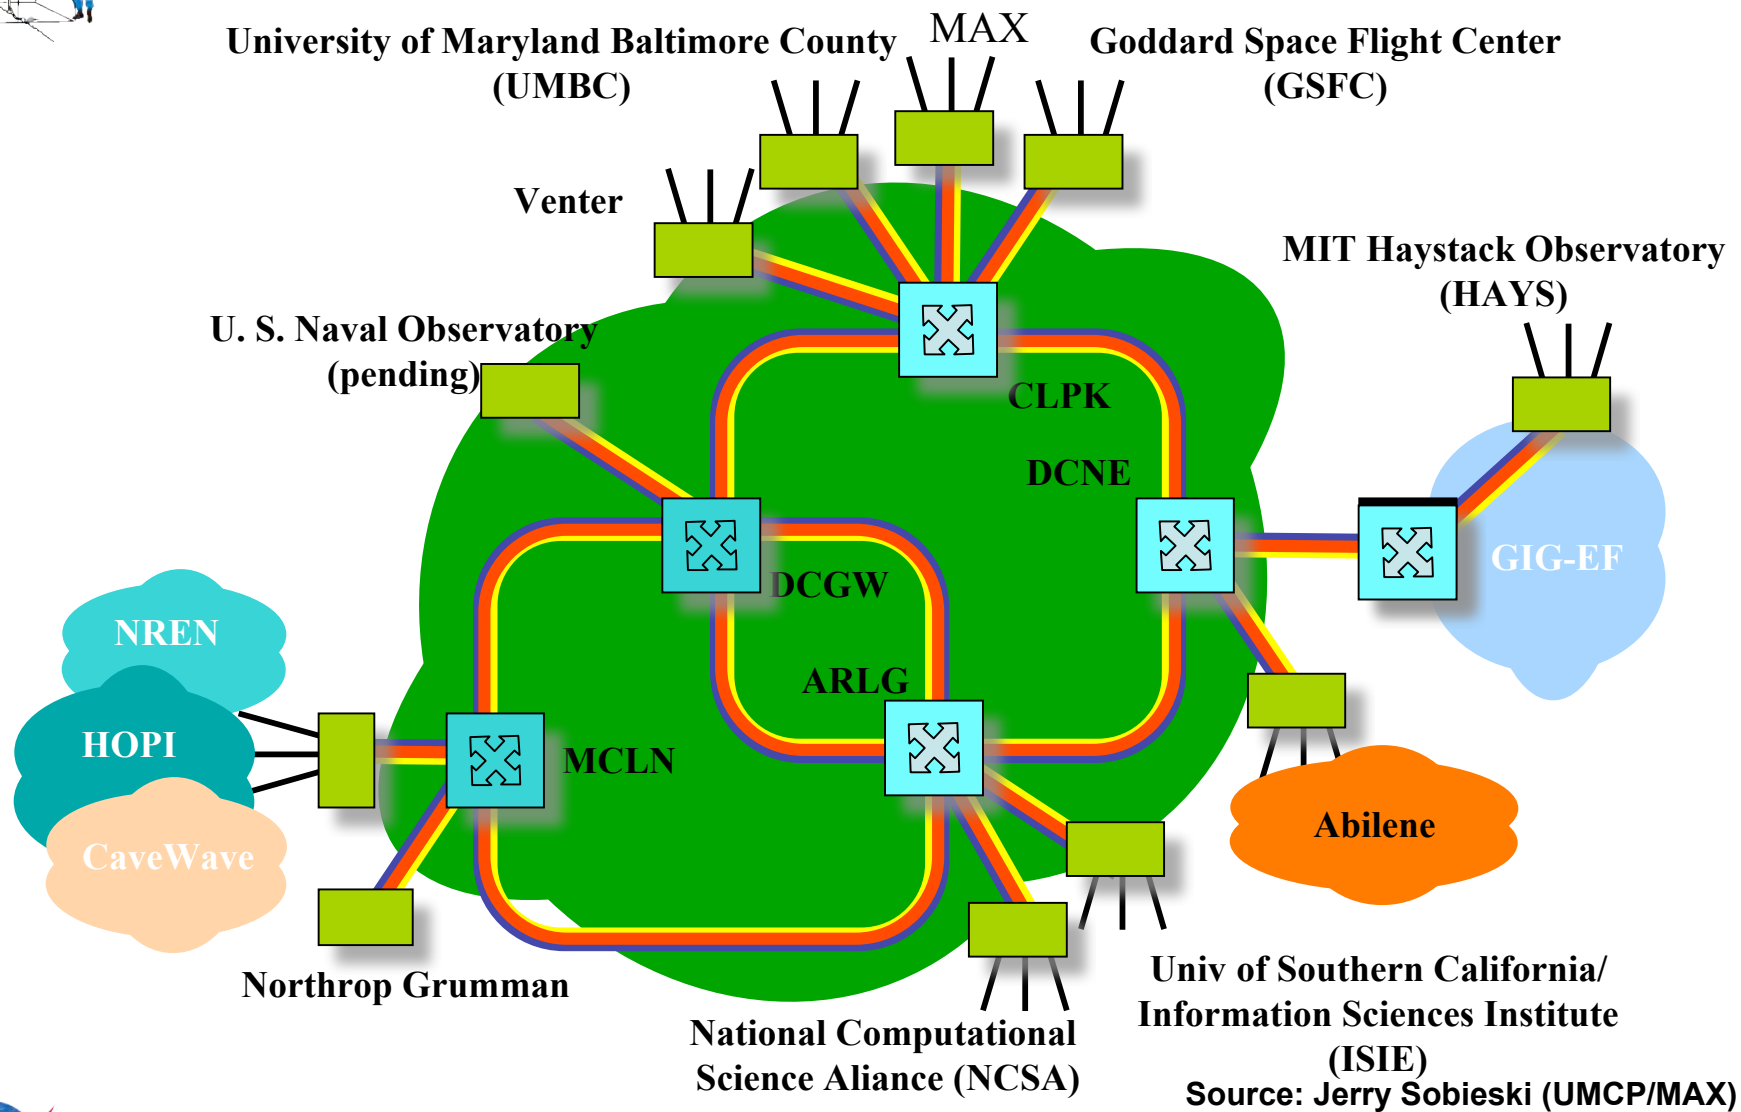

Source: Jerry Sobieski (UMCP/MAX)

National LambdaRail Architecture

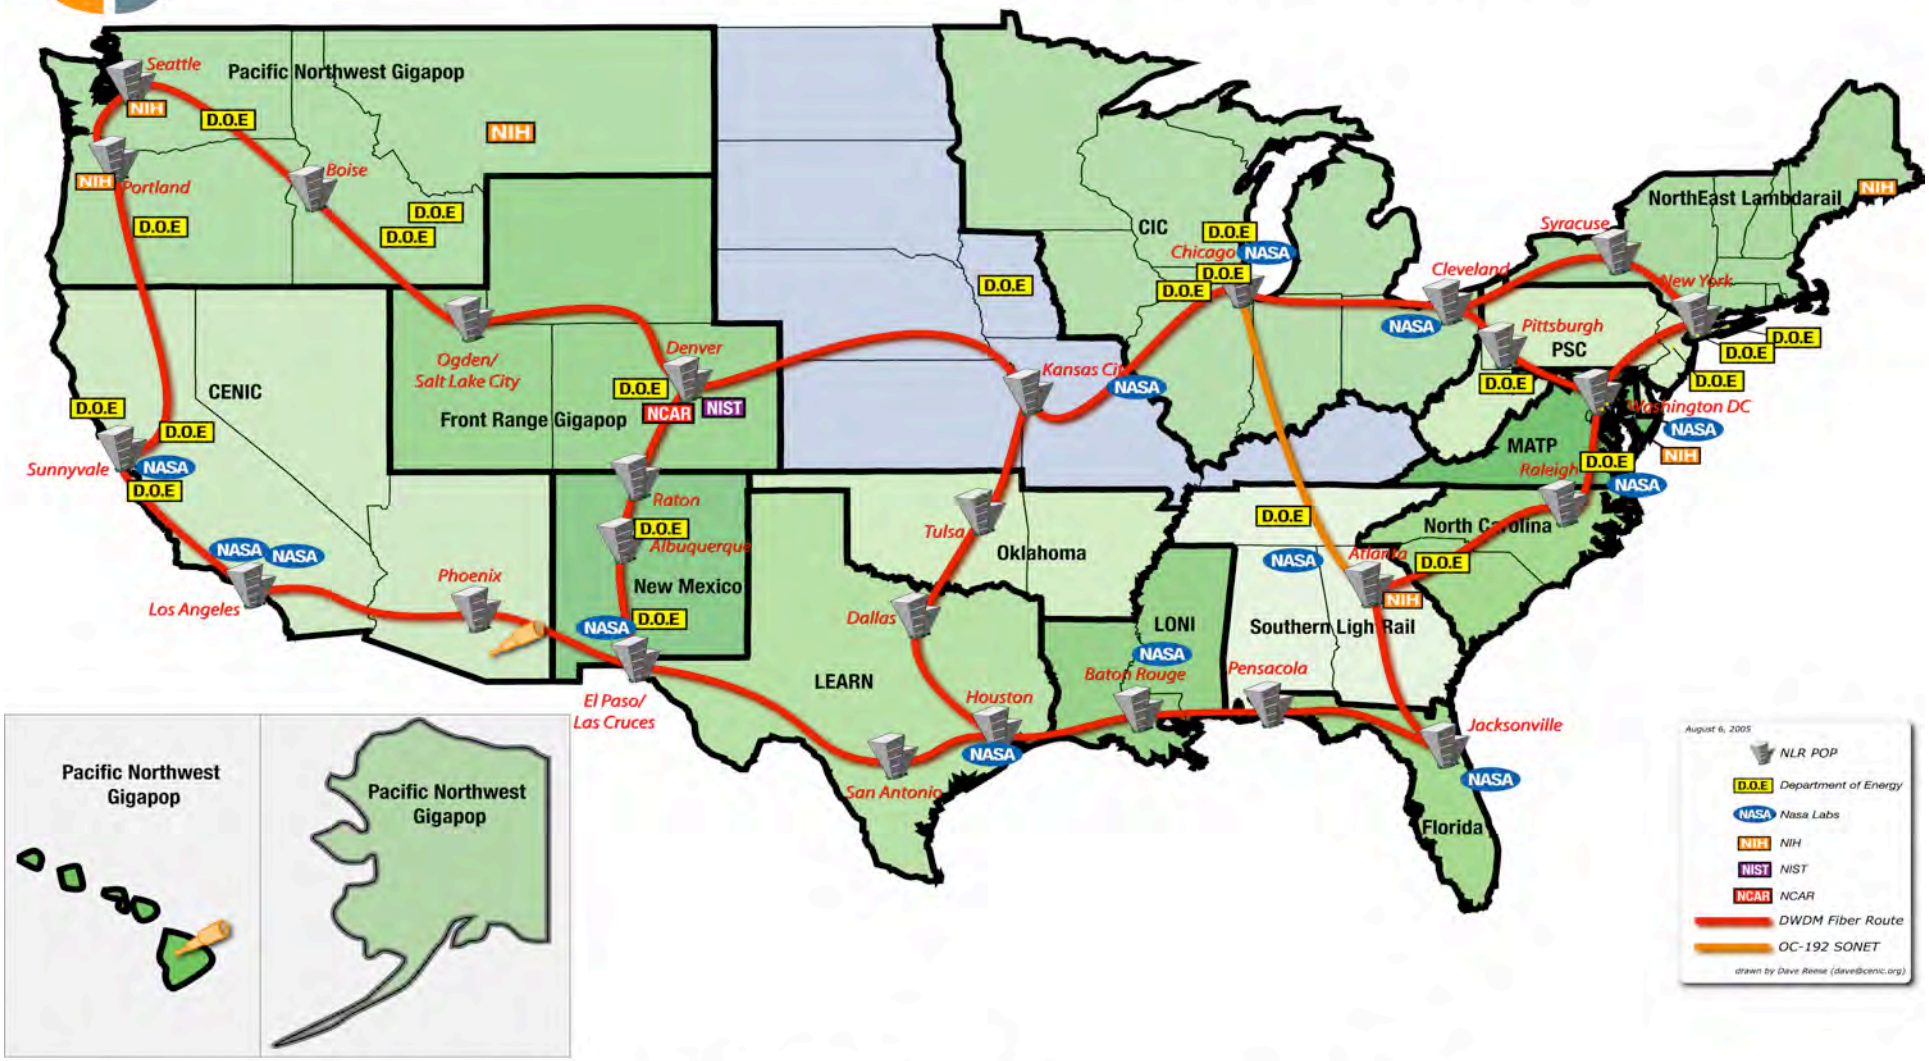

For more information regarding NLR see http://www.nlr.net or contact info@nlr.net

NASA GSFC Tests with OptIPuter Across the National LambdaRail

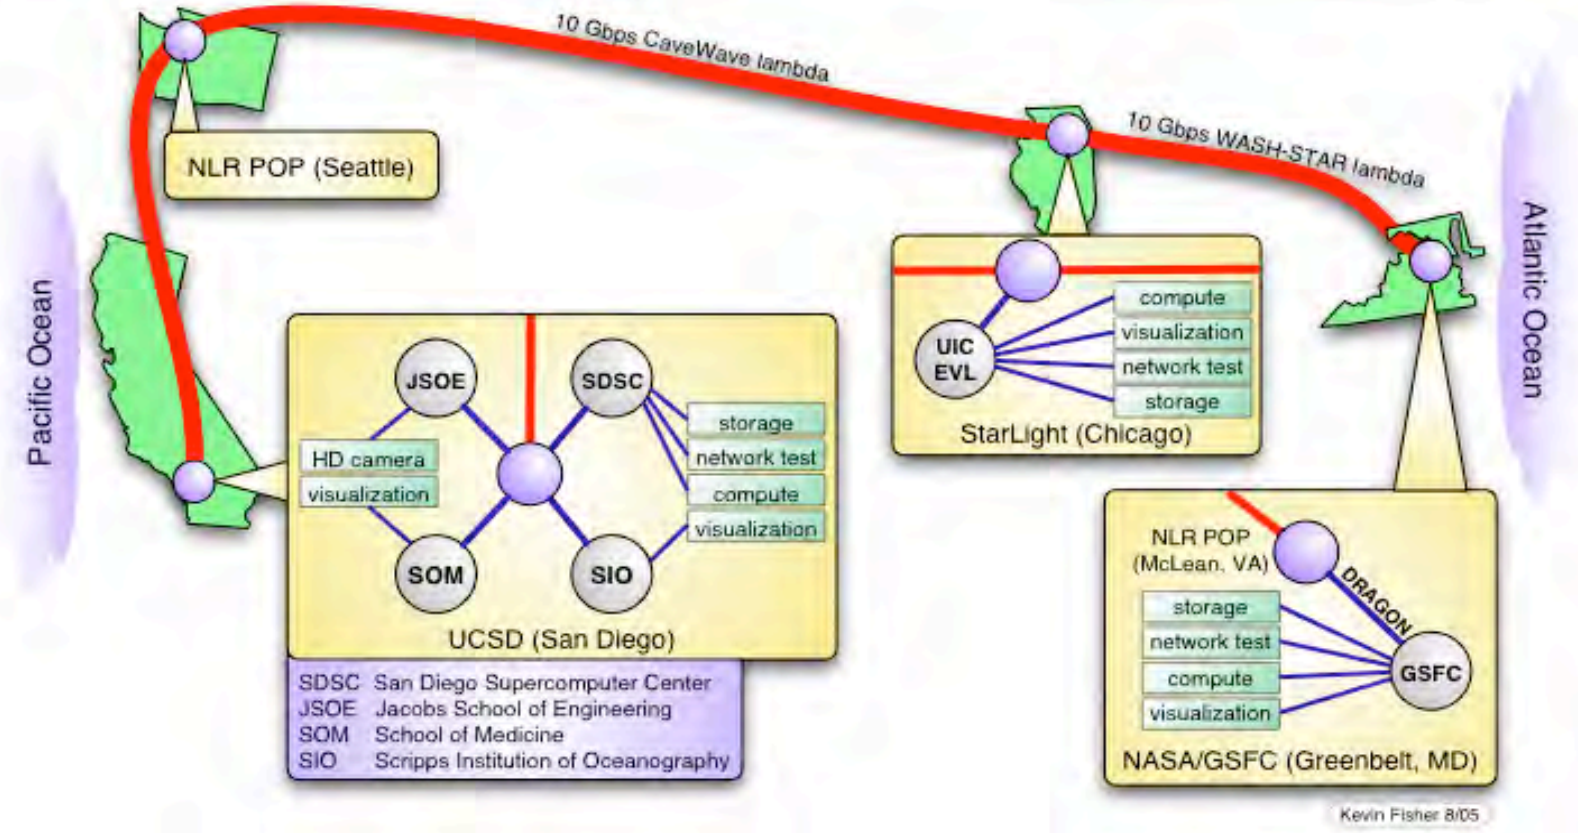

NREN Target FY06

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

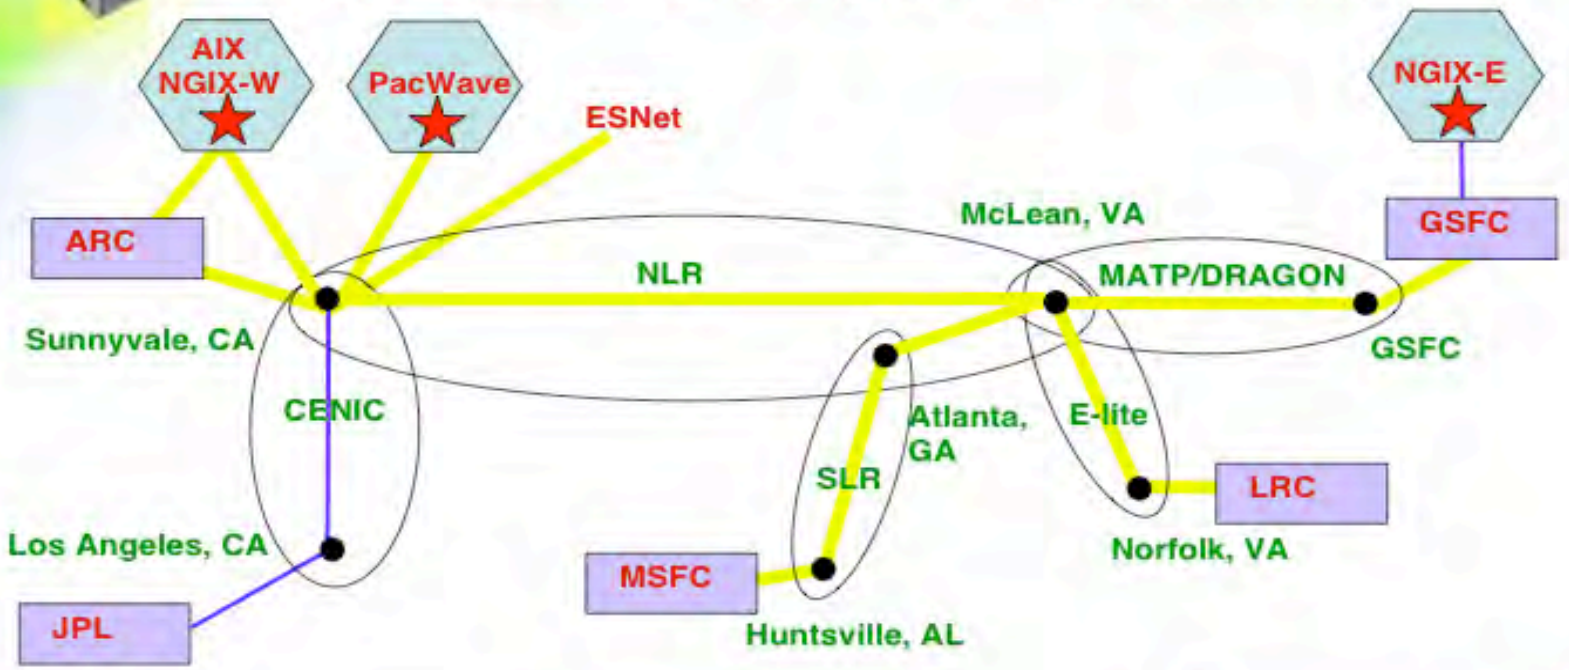

- National and Regional optical networks provide links over which 10 Gbps and 1 Gbps waves can be established.
- Distributed exchange points provide interconnect in metro and regional areas to other networks and research facilities

Source: Mark Foster (ARC)

Global Lambda Integrated Facility
World Map – December 2004

Predicted international Research & Education Network bandwidth, to be made available for scheduled application and middleware research experiments by December 2004.

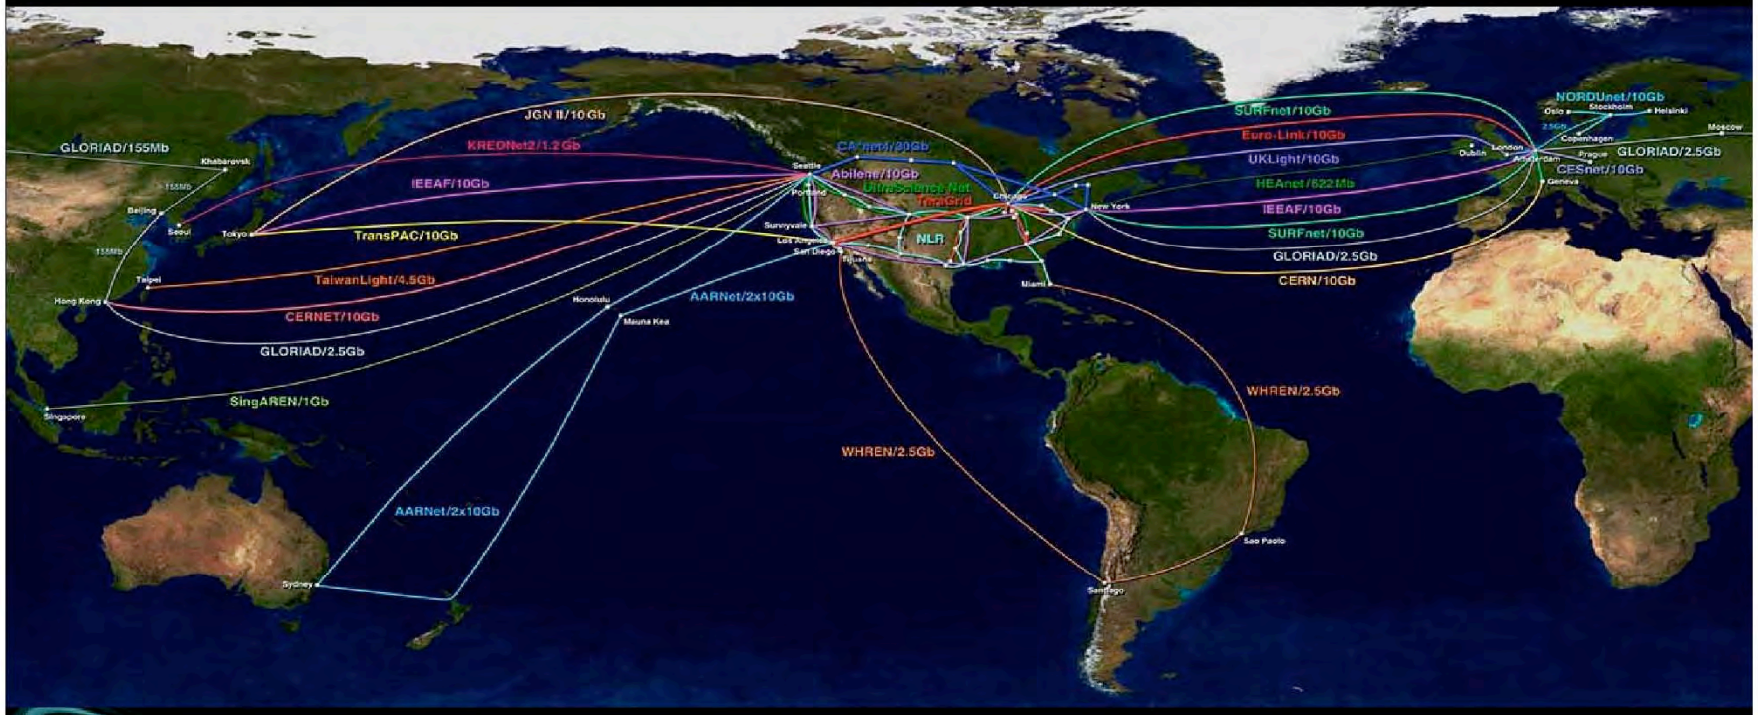

www.glif.is

Visualization courtesy of
Bob Patterson, NCSA.

STARLIGHT℠

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

EARTH SYSTEM MODELING FRAMEWORK

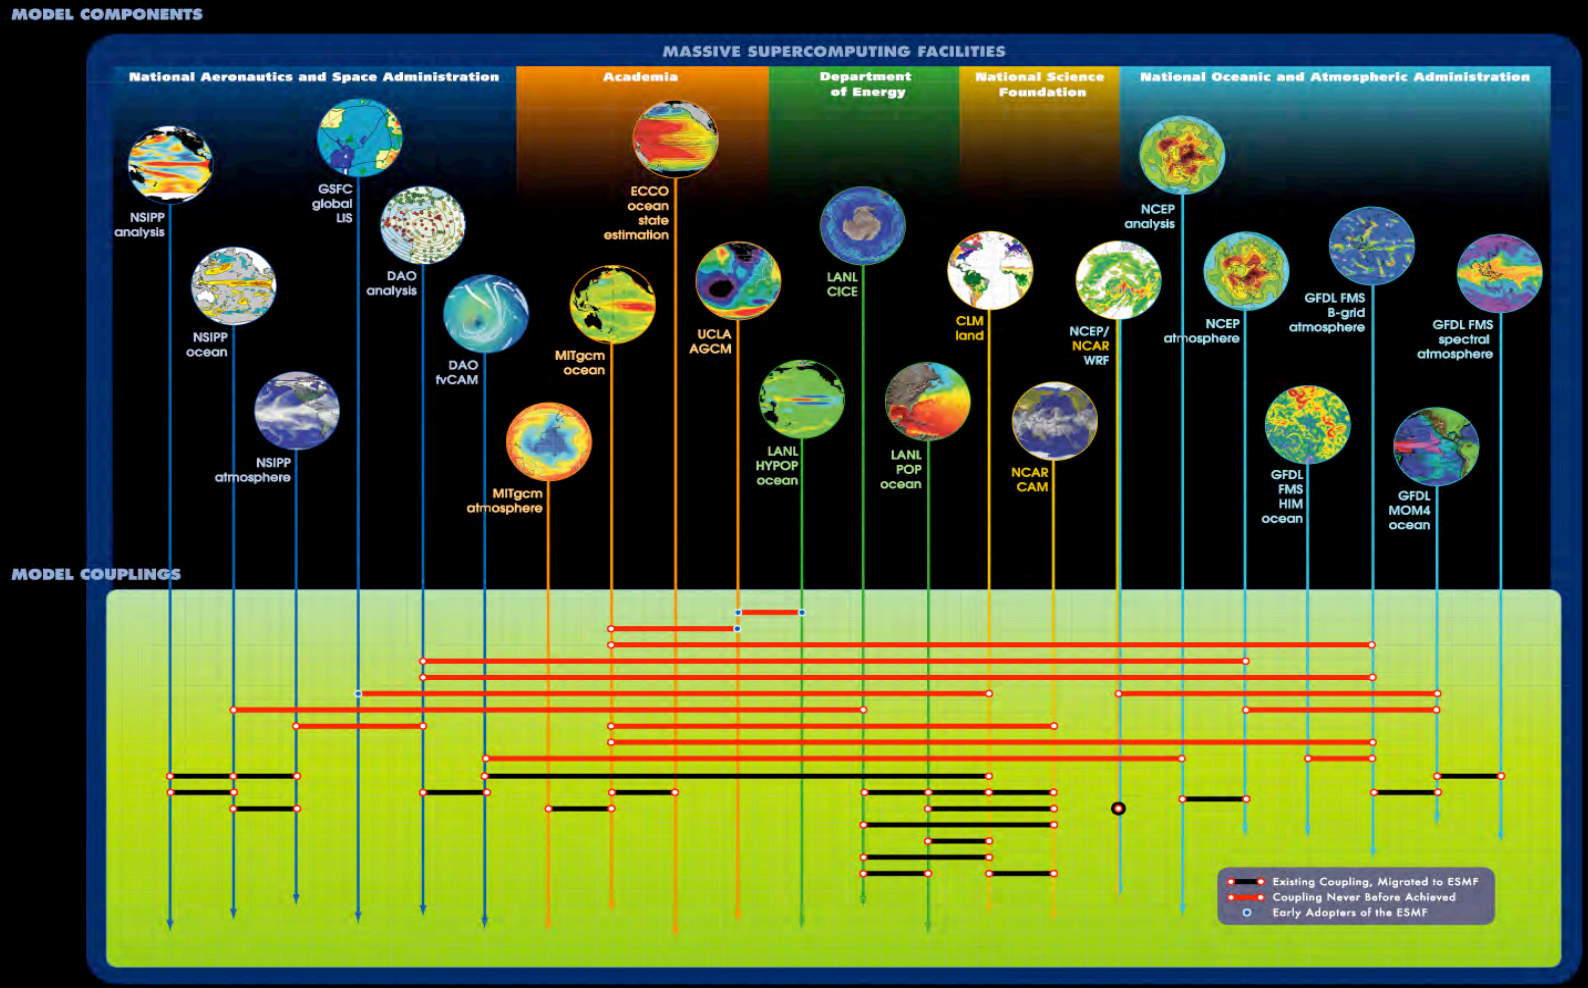

ESMF-Enabled Coupled Models

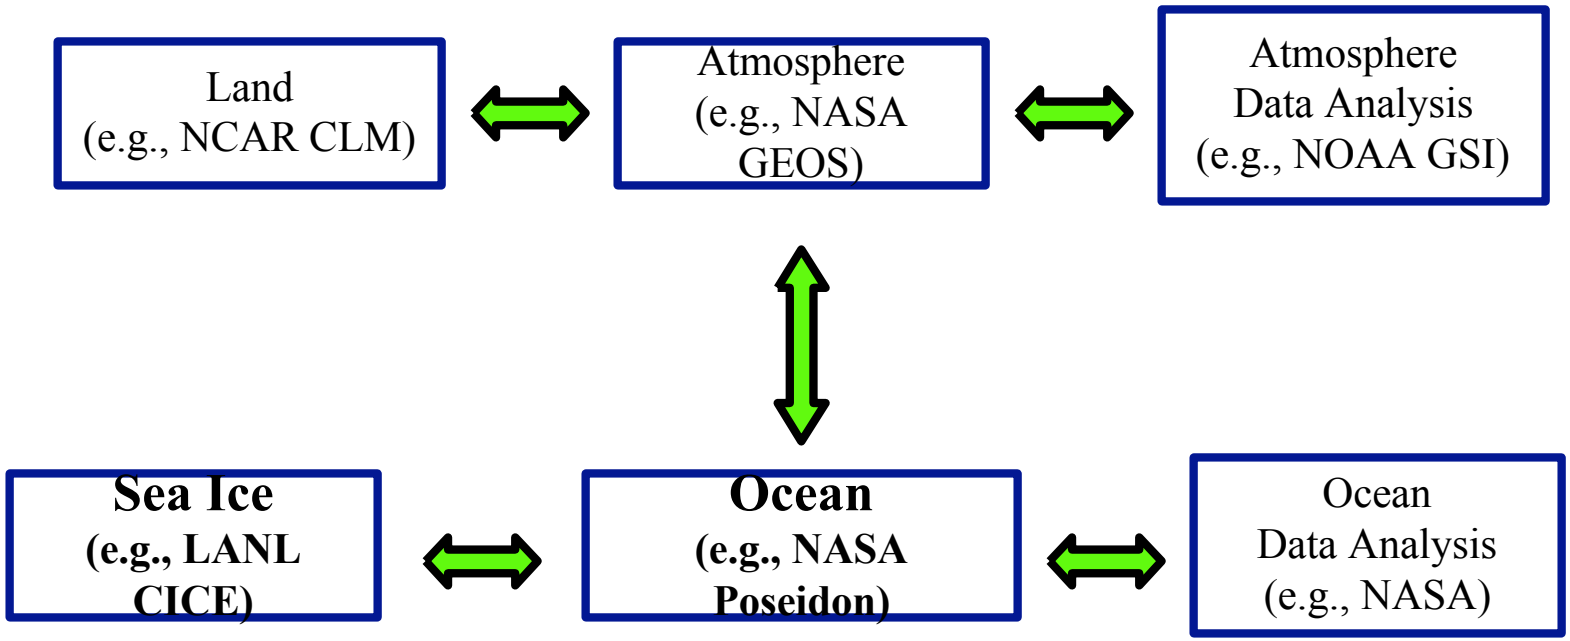

DRAGON eVLBI Experiment Configuration

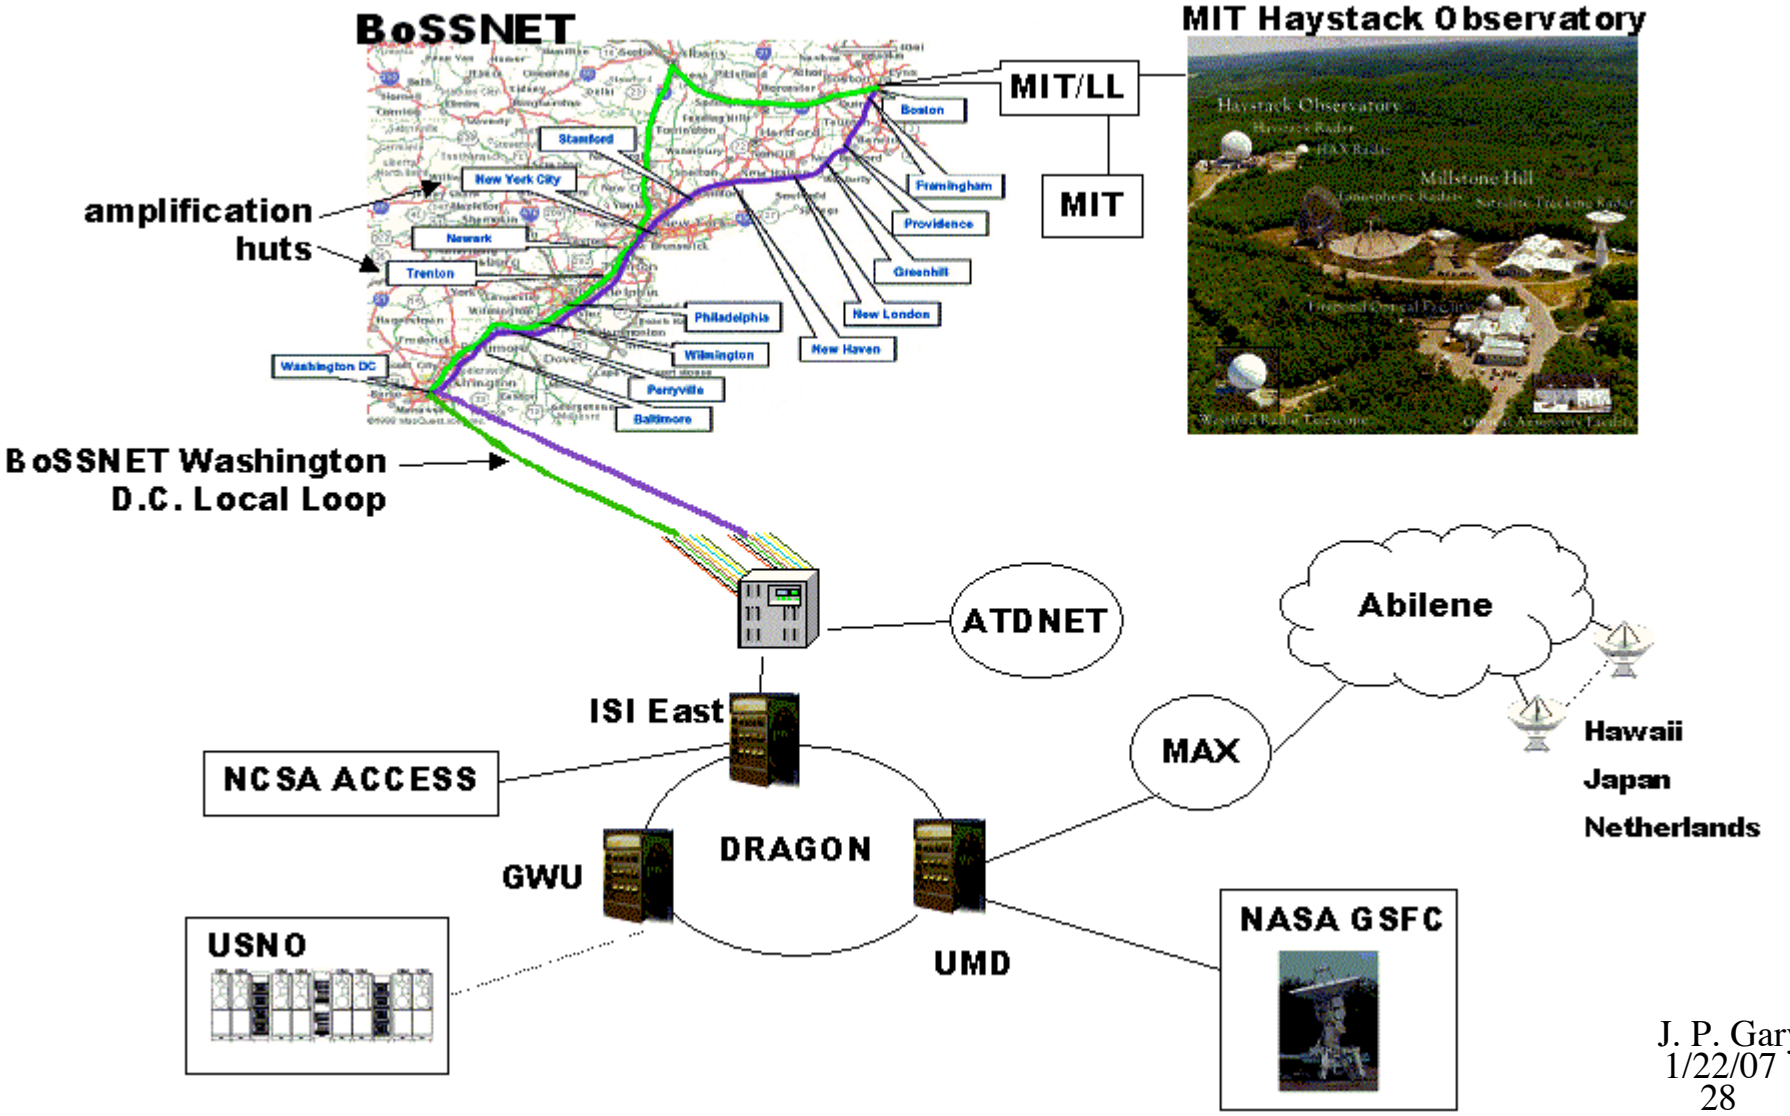

Combining Telepresence with Remote Interactive Analysis of Data Over NLR

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

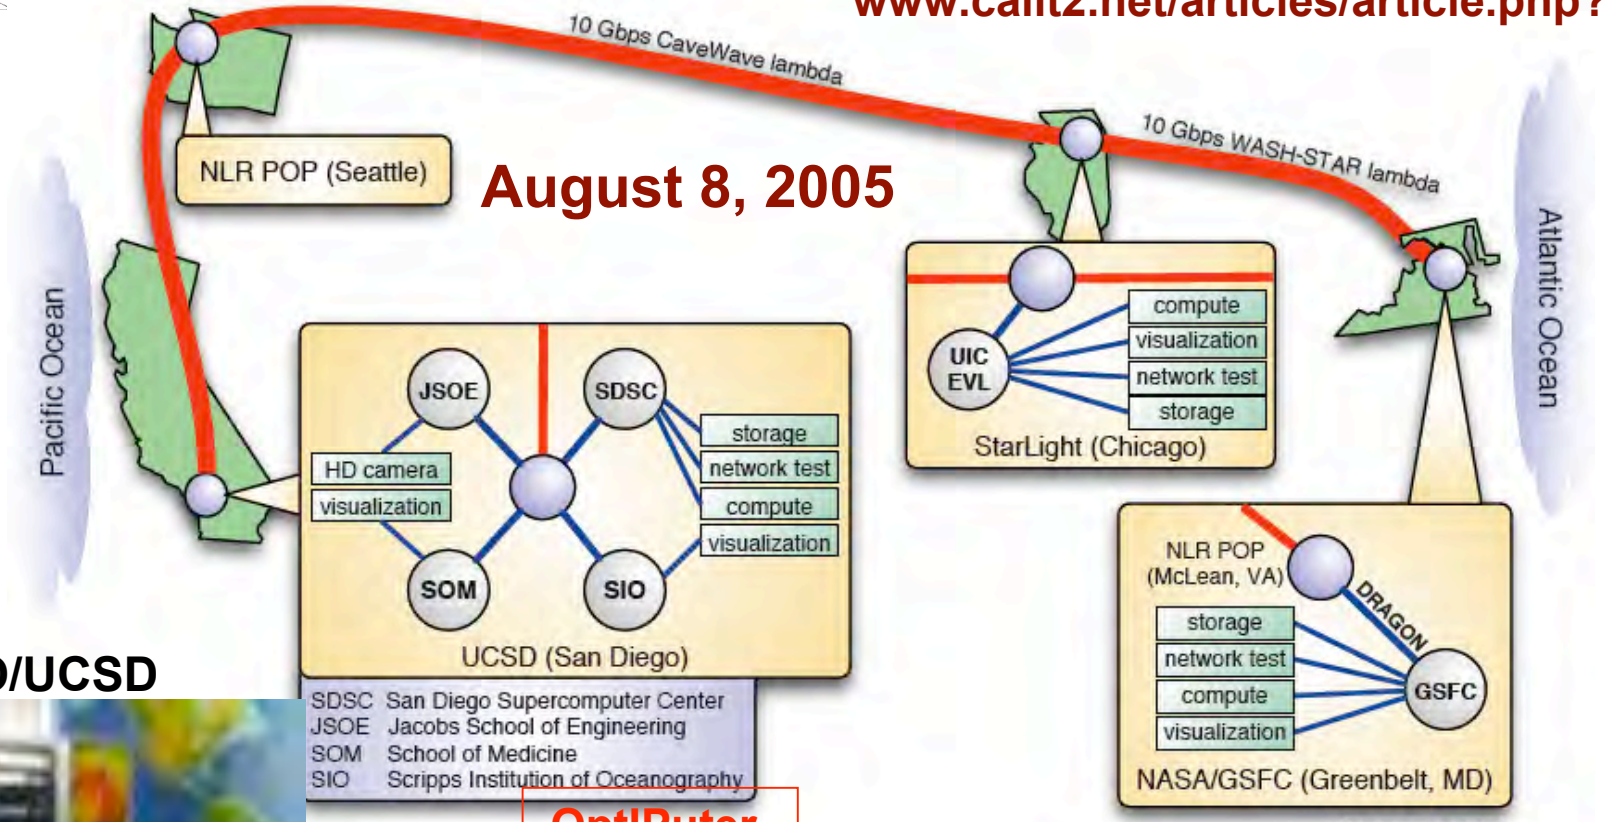

SIO/UCSD

OptIPuter Visualized Data

HDTV Over Lambda

NASA Goddard

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

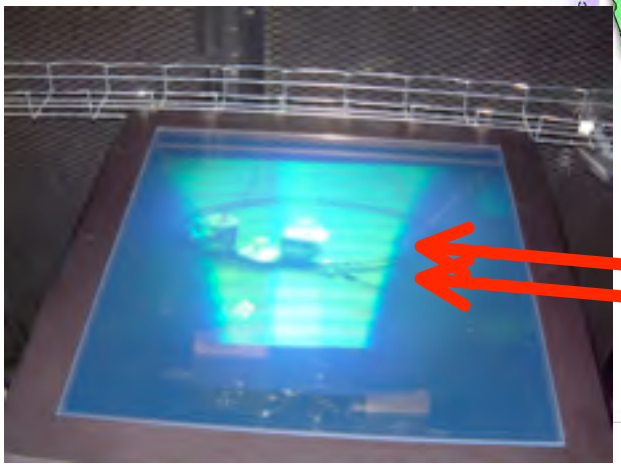

Only a non-stereo image of the True-3D display is captured in this photo of the real-time stereo-HDTV images transmitted from GSFC.

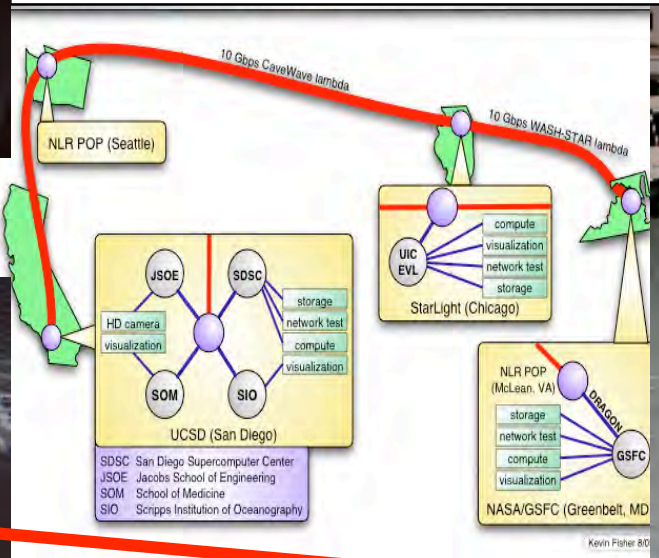

3D HDTV Over Lambda

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

Current SAN-over-IP Test-bed

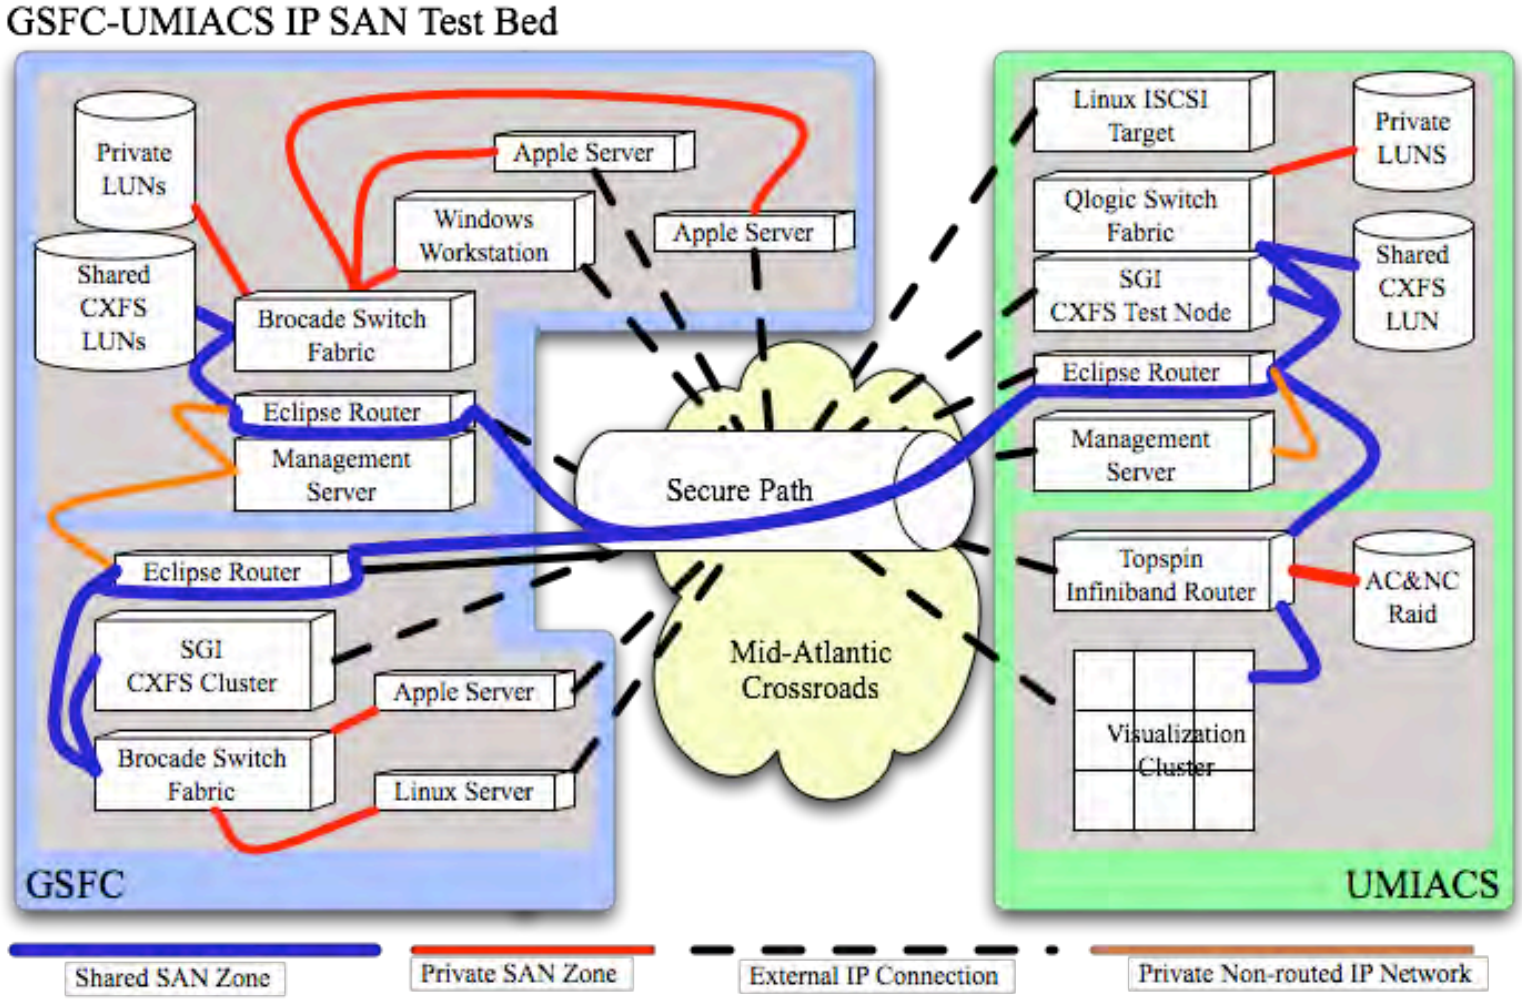

Source: Fritz McCall (UMIACS)