GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Summary Status of

GSFC 10 Gbps "Lambda Network" (L-Net)

J. Patrick Gary

Network Projects Leader

Earth and Space Data Computing Division

NASA Goddard Space Flight Center

pat.gary@nasa.gov

301-286-9539

(Courtesy of Dr. Larry Smarr/UCSD)
Goddard Space Flight Center IRAD

- Establish a "Lambda Network"
  - » Use optical wavelength technology (10Gbps Ethernet per wavelength)
  - » Connect GSFC's Earth Science Greenbelt facility in Maryland to Scripps Institution of Oceanography in through the University of California at San Diego (UCSD)
  - » Ride the National LambdaRail
- Synergies
  - » Falls in with reinvigorating the nation's cyber infrastructure
  - » Takes advantage of next generation networking technologies (Lambda-Nets)
  - » Makes use of NSF funded compute, storage and visualization resources being implemented at UCSD (OptIPuter and GEON)
- Benefits
  - » Develop real-time interactive collaborations with leading Earth and space science academic institutions
  - » Enable scientists at both institutions to share and use compute intensive community models, complex data base mining and multi-dimensional streaming visualization
  - » Creates a virtual laboratory and SIO wing within GSFC's Building 33
  - » Several NASA missions benefit from the high bandwidth connection, eg:
    - – HPC between ARC/Project Columbia and GSFC/NCCS
    - – CEOP data analyses between Rhodes (SIO) and Bosilovich (GSFC)

(Courtesy of Dr. Larry Smarr/UCSD)
GSFC FY04 IRAD Proposal

- "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"
- Transcontinental, Regional, and Local Networking
  - » Became a member of the NLR, with NREN Project assistance, through a Mid-Atlantic Terascale Partnership membership arrangement
  - » Deploying GMPLS-managed Movaz Optical Switches and Add Drop Muxs at GSFC and in DRAGON regional multiwavelength network
  - » Utilize NLR lambdas between GSFC and UCSD/SIO
  - » Interconnect GSFC's Thunderhead and other clusters to UCSD/SIO's OptIPuter network at 10GigE via this Lambda Network
- Application Development
  - » Integrate Earth System Modeling Framework software with GRID middleware by constructing prototype interfaces between the components
  - » Identify requirements for new methods and/or messages that would be desirable for supporting GSFC models and data assimilation

NLR – Optical Infrastructure - Phase 1

NLR Phase 1 - Installation Schedule
Will Complete Aug 2004

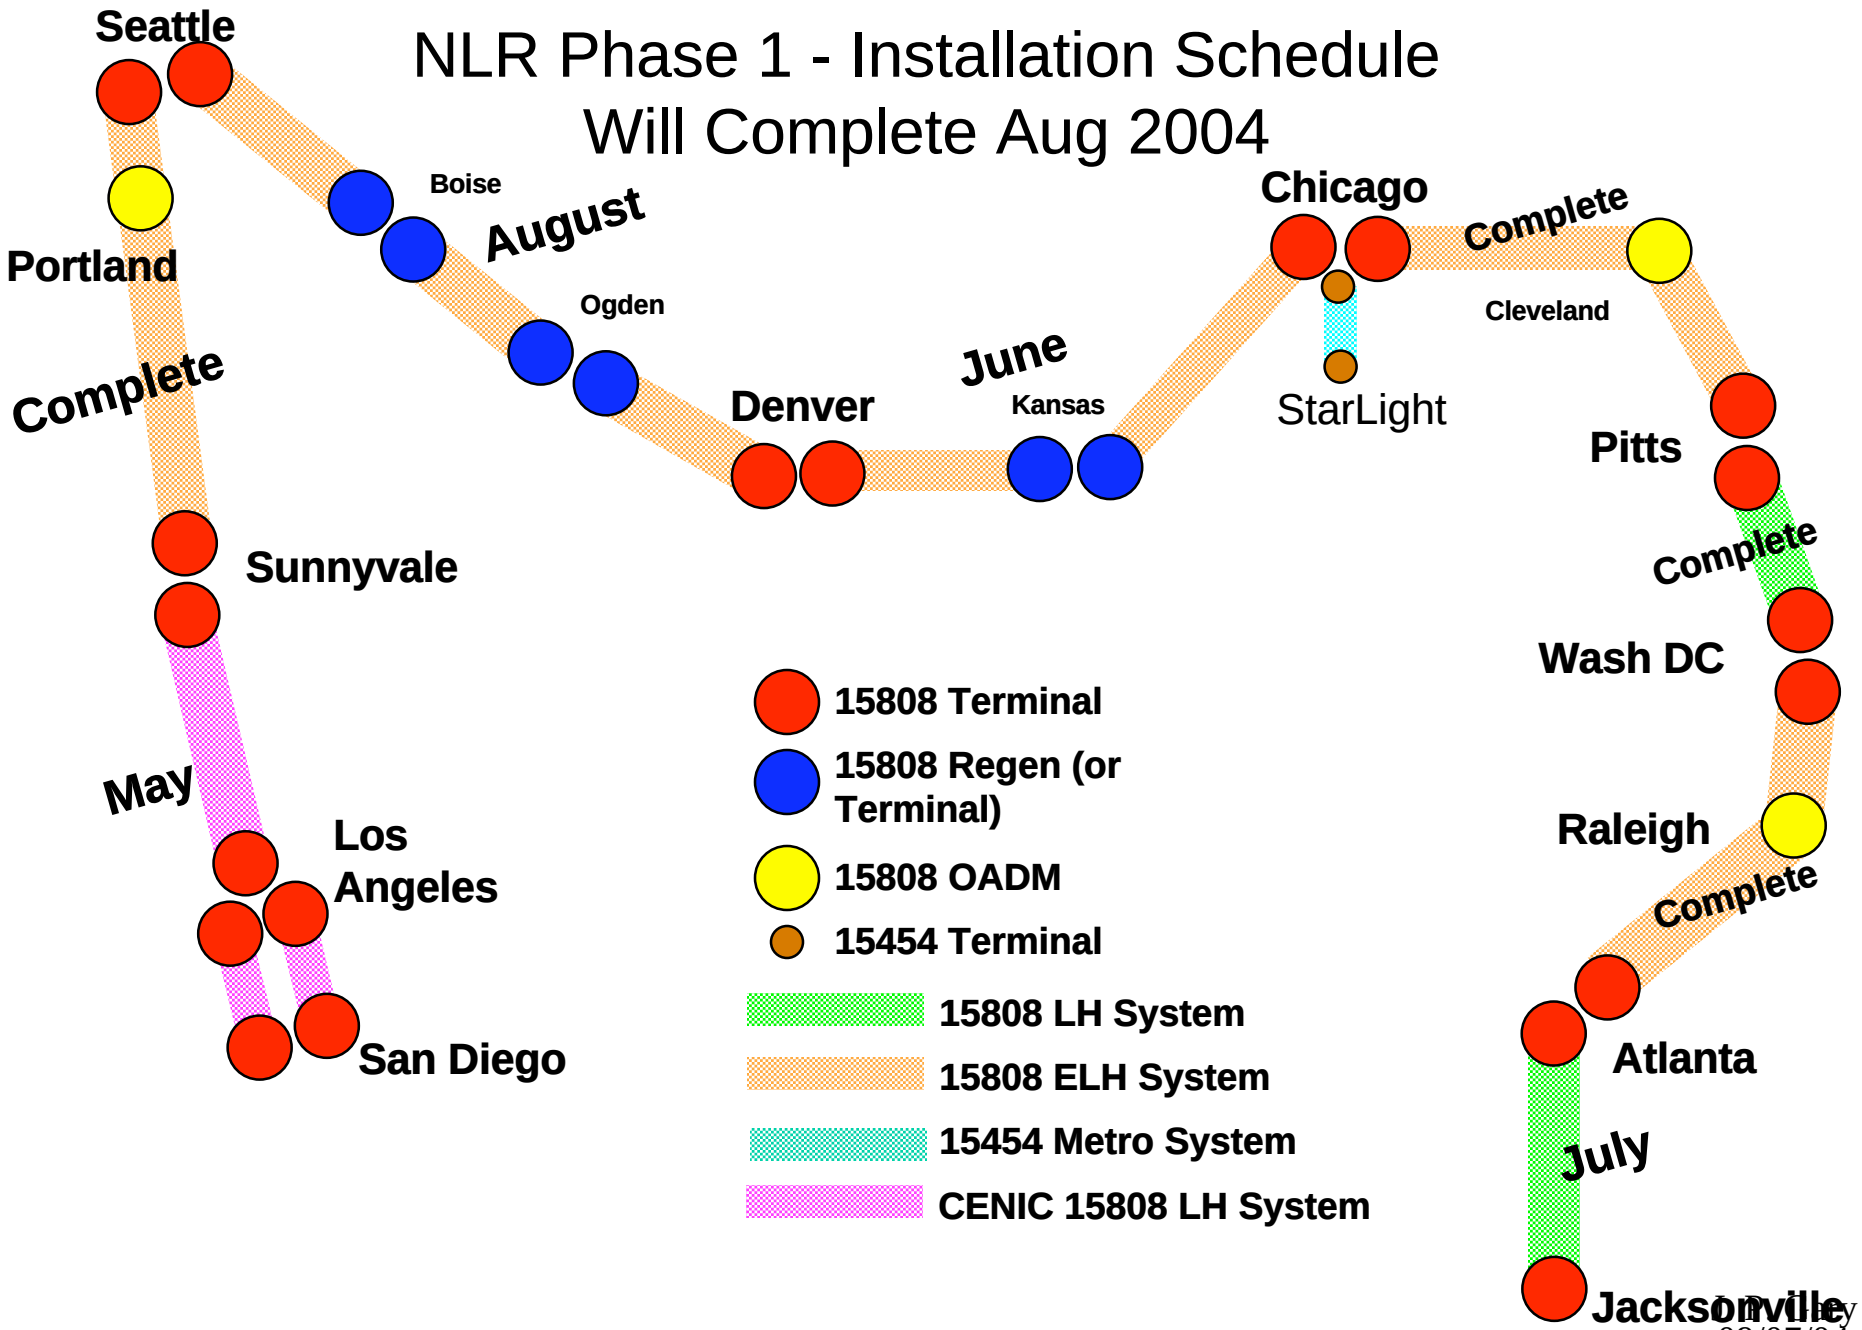

NLR – Phase 1 +
Route to NY & ORNL Waves & Salt Lake City

NLR – Proposed Southern Route

Wiltel Fiber in the West -- AT&T Fiber in the Southeast

Qwest Fiber North-South

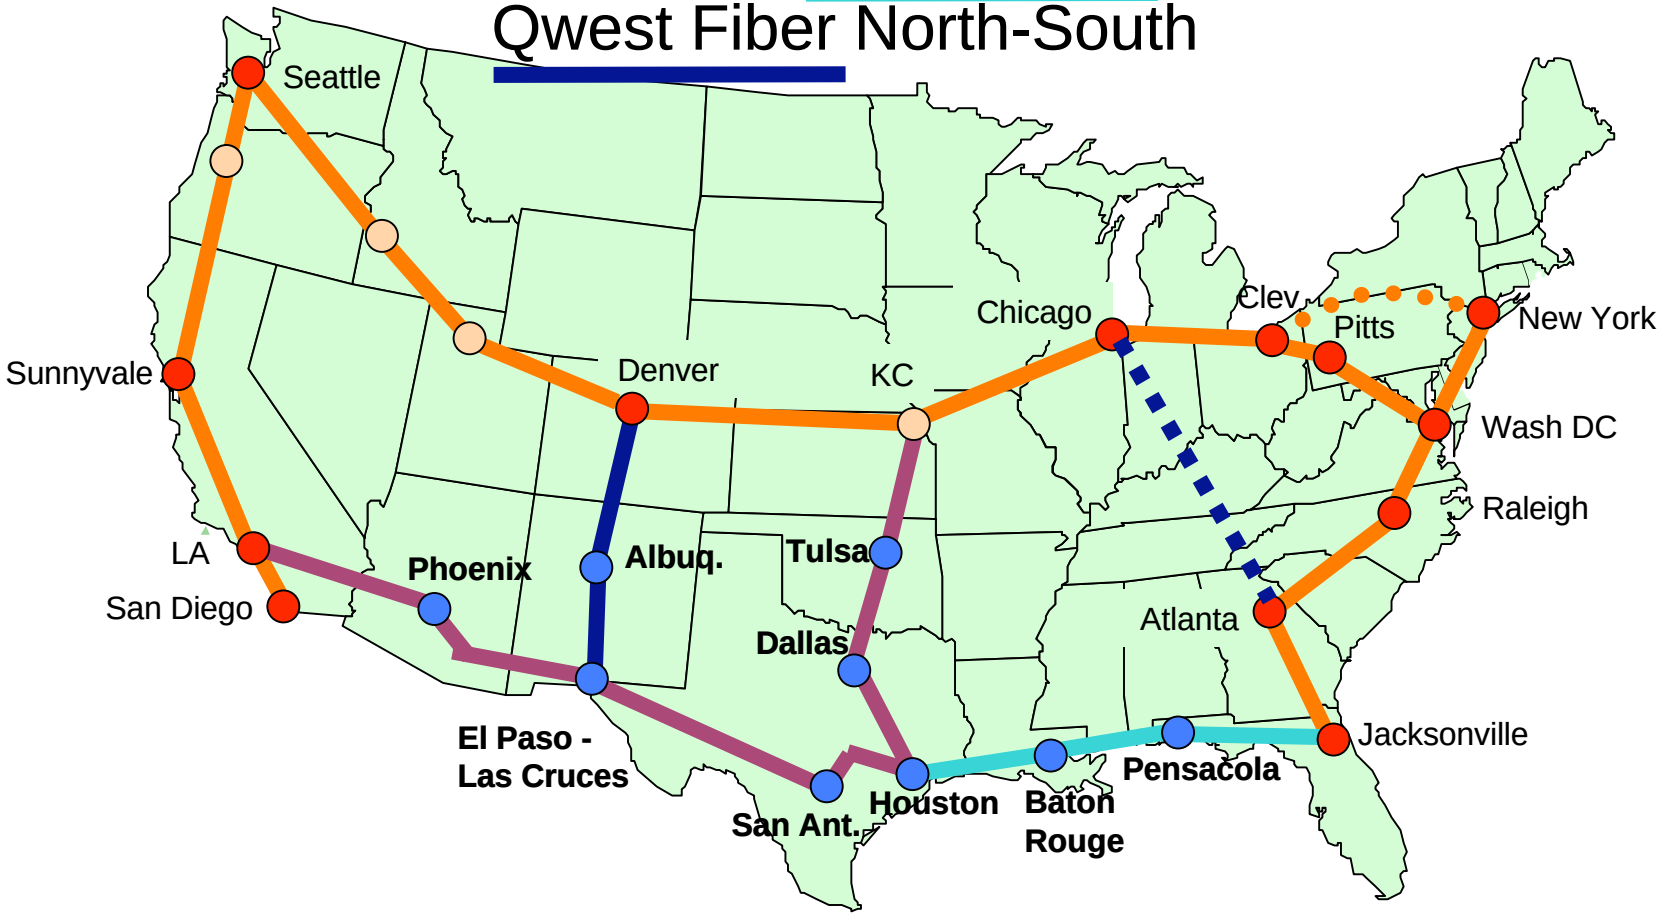

Using Cisco 15454 DWDM gear

NLR Wavelengths

- **Initial complement of 4 λs installed and available at outset**
  - One λ for national switched Ethernet experimental network
  - Another λ for national 10 Gbps IP network to support internetworking and end-to-end transport protocol experiments
    - Similar to Internet2's Abilene except routers will be available for measurement and experimentation
  - Third λ will serve as a quick start facility for new research projects
  - Fourth λ will be used by Internet2's HOPI testbed

  More λs will be activated as needed to support the research and operational objectives of the community

NLR Members as of 19Jul04

- CENIC
- Pacific Northwest GigaPOP
- Pittsburgh Supercomp. Center
- Duke (coalition of NC univers.)
- Mid-Atlantic Terascale Partnership
- Cisco Systems
- Internet2
- Florida LambdaRail
- Georgia Institute of Technology
- Committee on Institutional Cooperation (CIC)
- Texas / LEARN
- Cornell
- Louisiana Board of Regents
- University of New Mexico
- Oklahoma State Regents
- UCAR/FRGP

Plus Agreements with:

- SURA (AT&T fiber donation)
- Oak Ridge National Lab (ORNL)

MATP Aggregation Facility Architecture

DRAFT

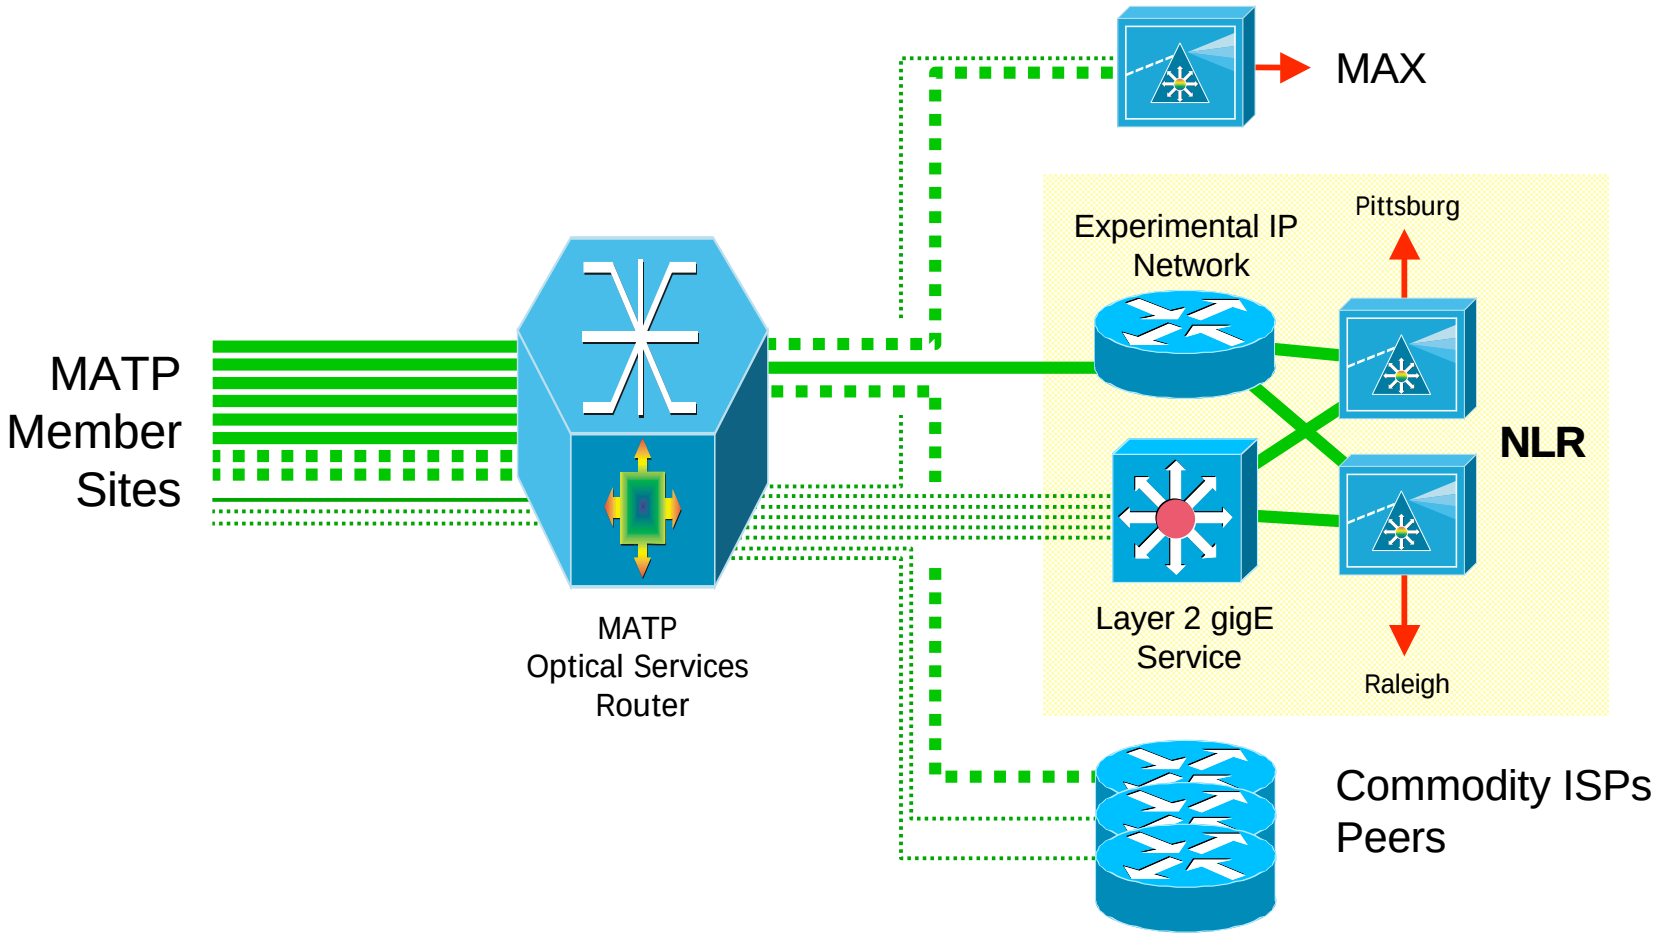

HOPI Node

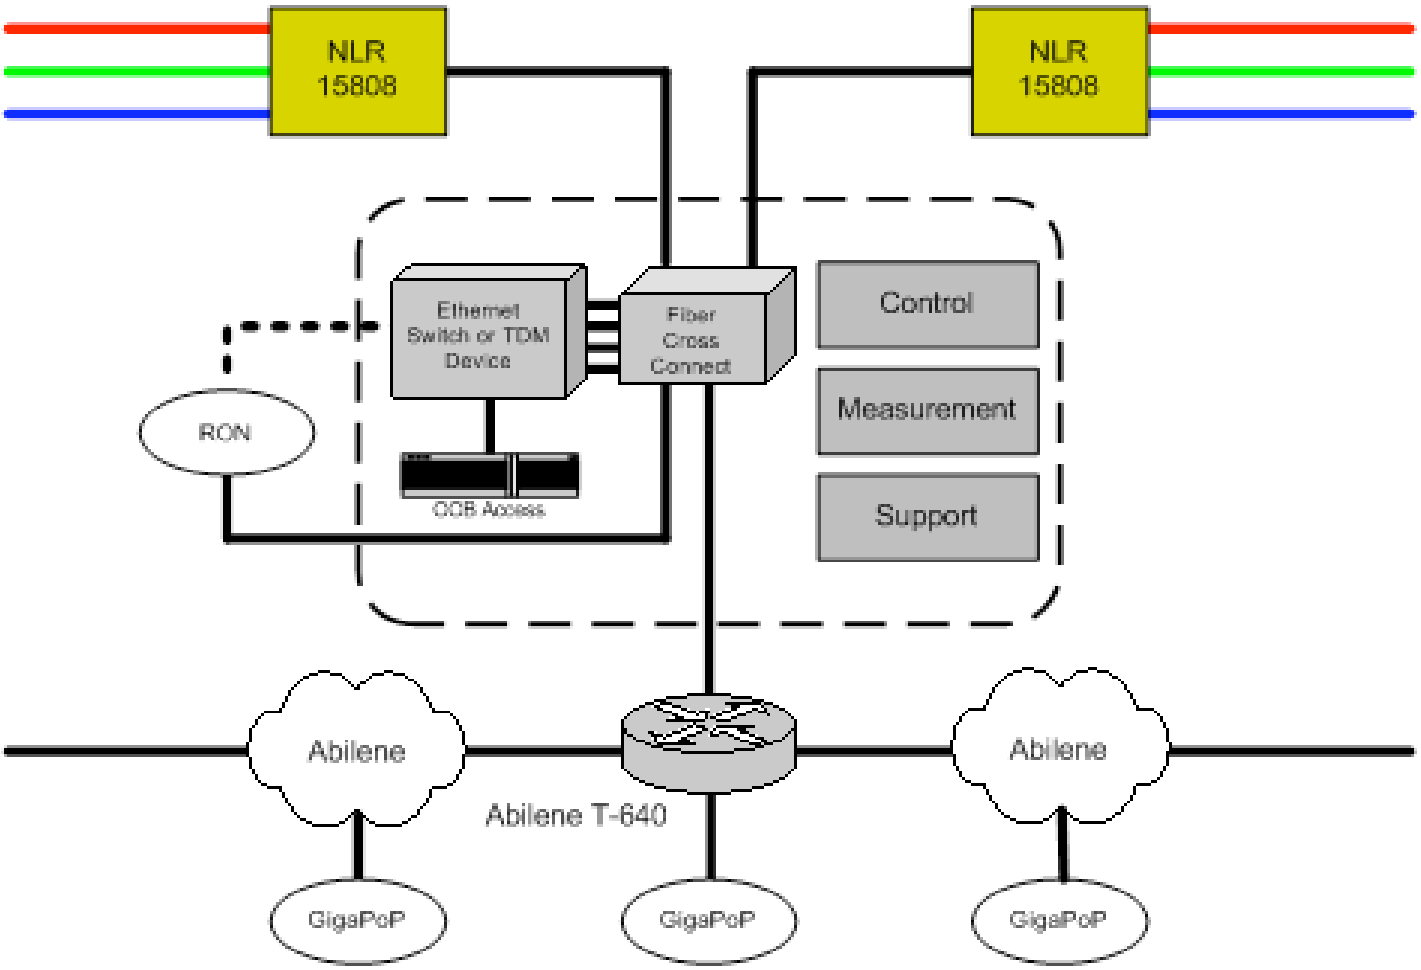
NLR
15808
NLR
15808
Ethernet
Switch or TDM
Device
Fiber
Cross
Connect
Control
Measurement
Support
RON
OOB Access
Abilene
Abilene
Abilene T-640
GigaPoP
GigaPoP
GigaPoP

Dynamic Resource Allocation with GMPLS on Optical Networks (DRAGON) Configuration

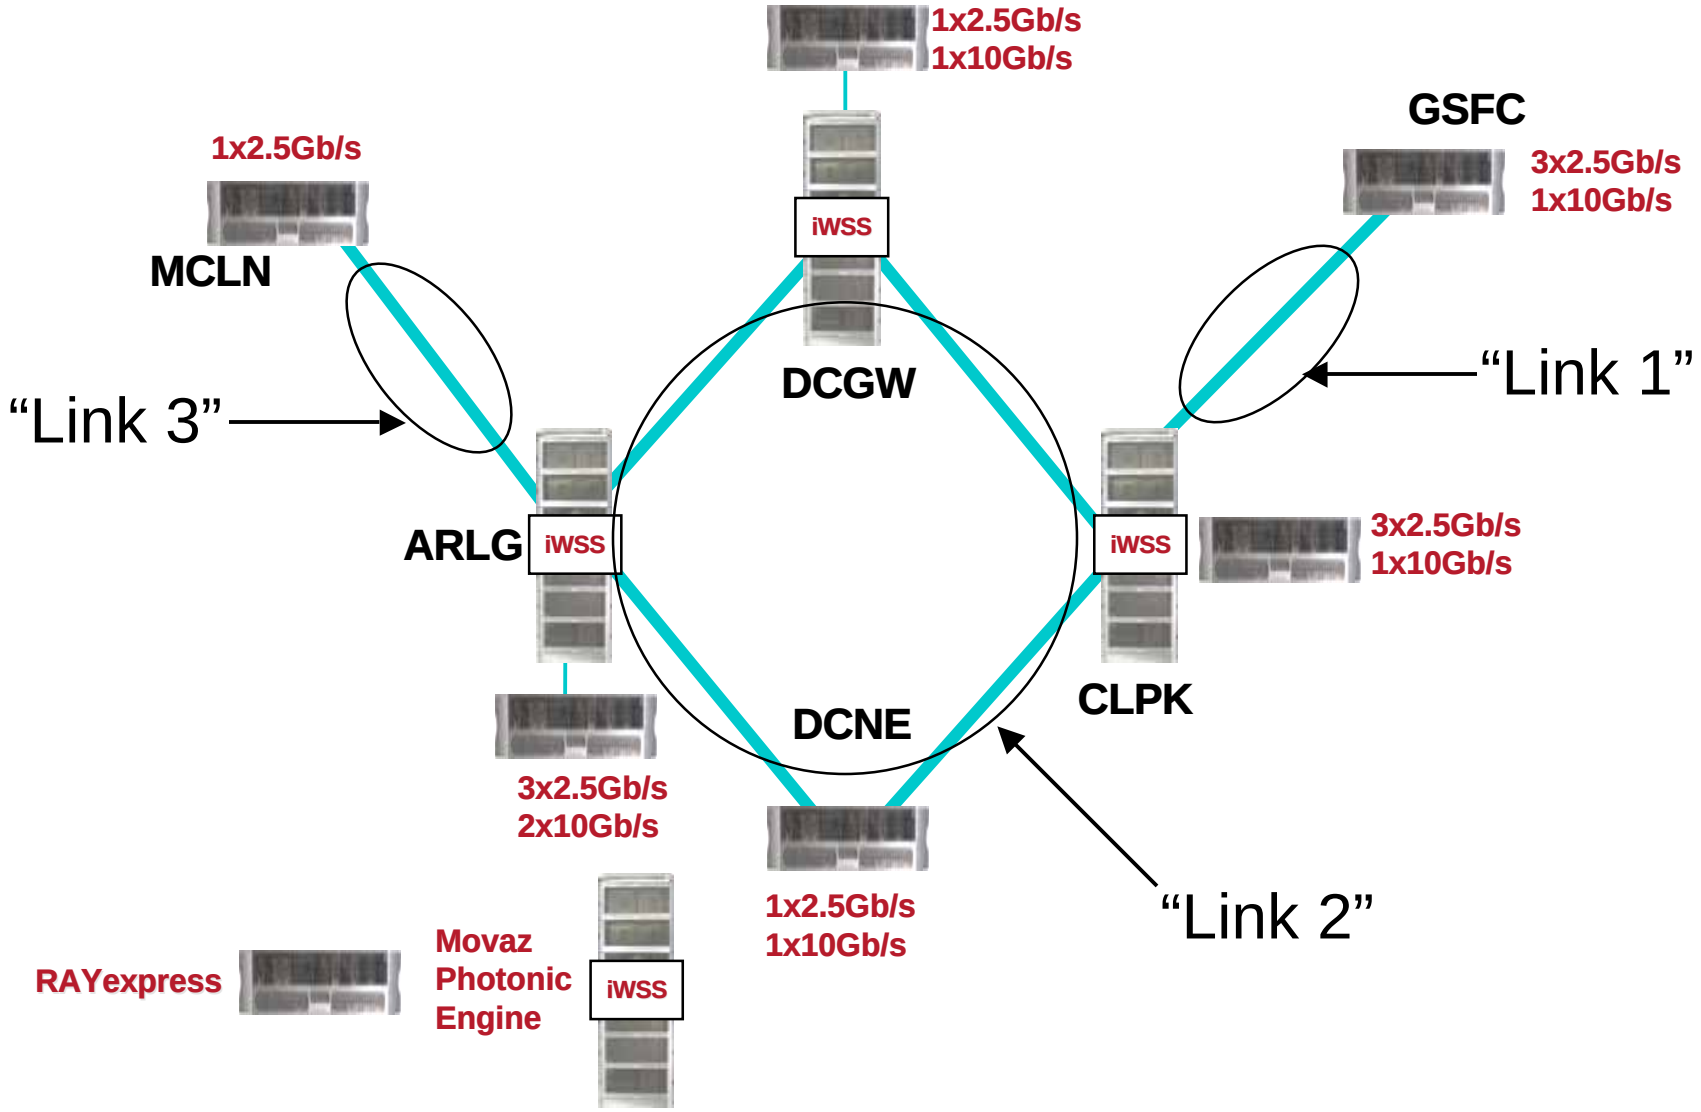

GSFC L-Net Configurations at McLean and Greenbelt

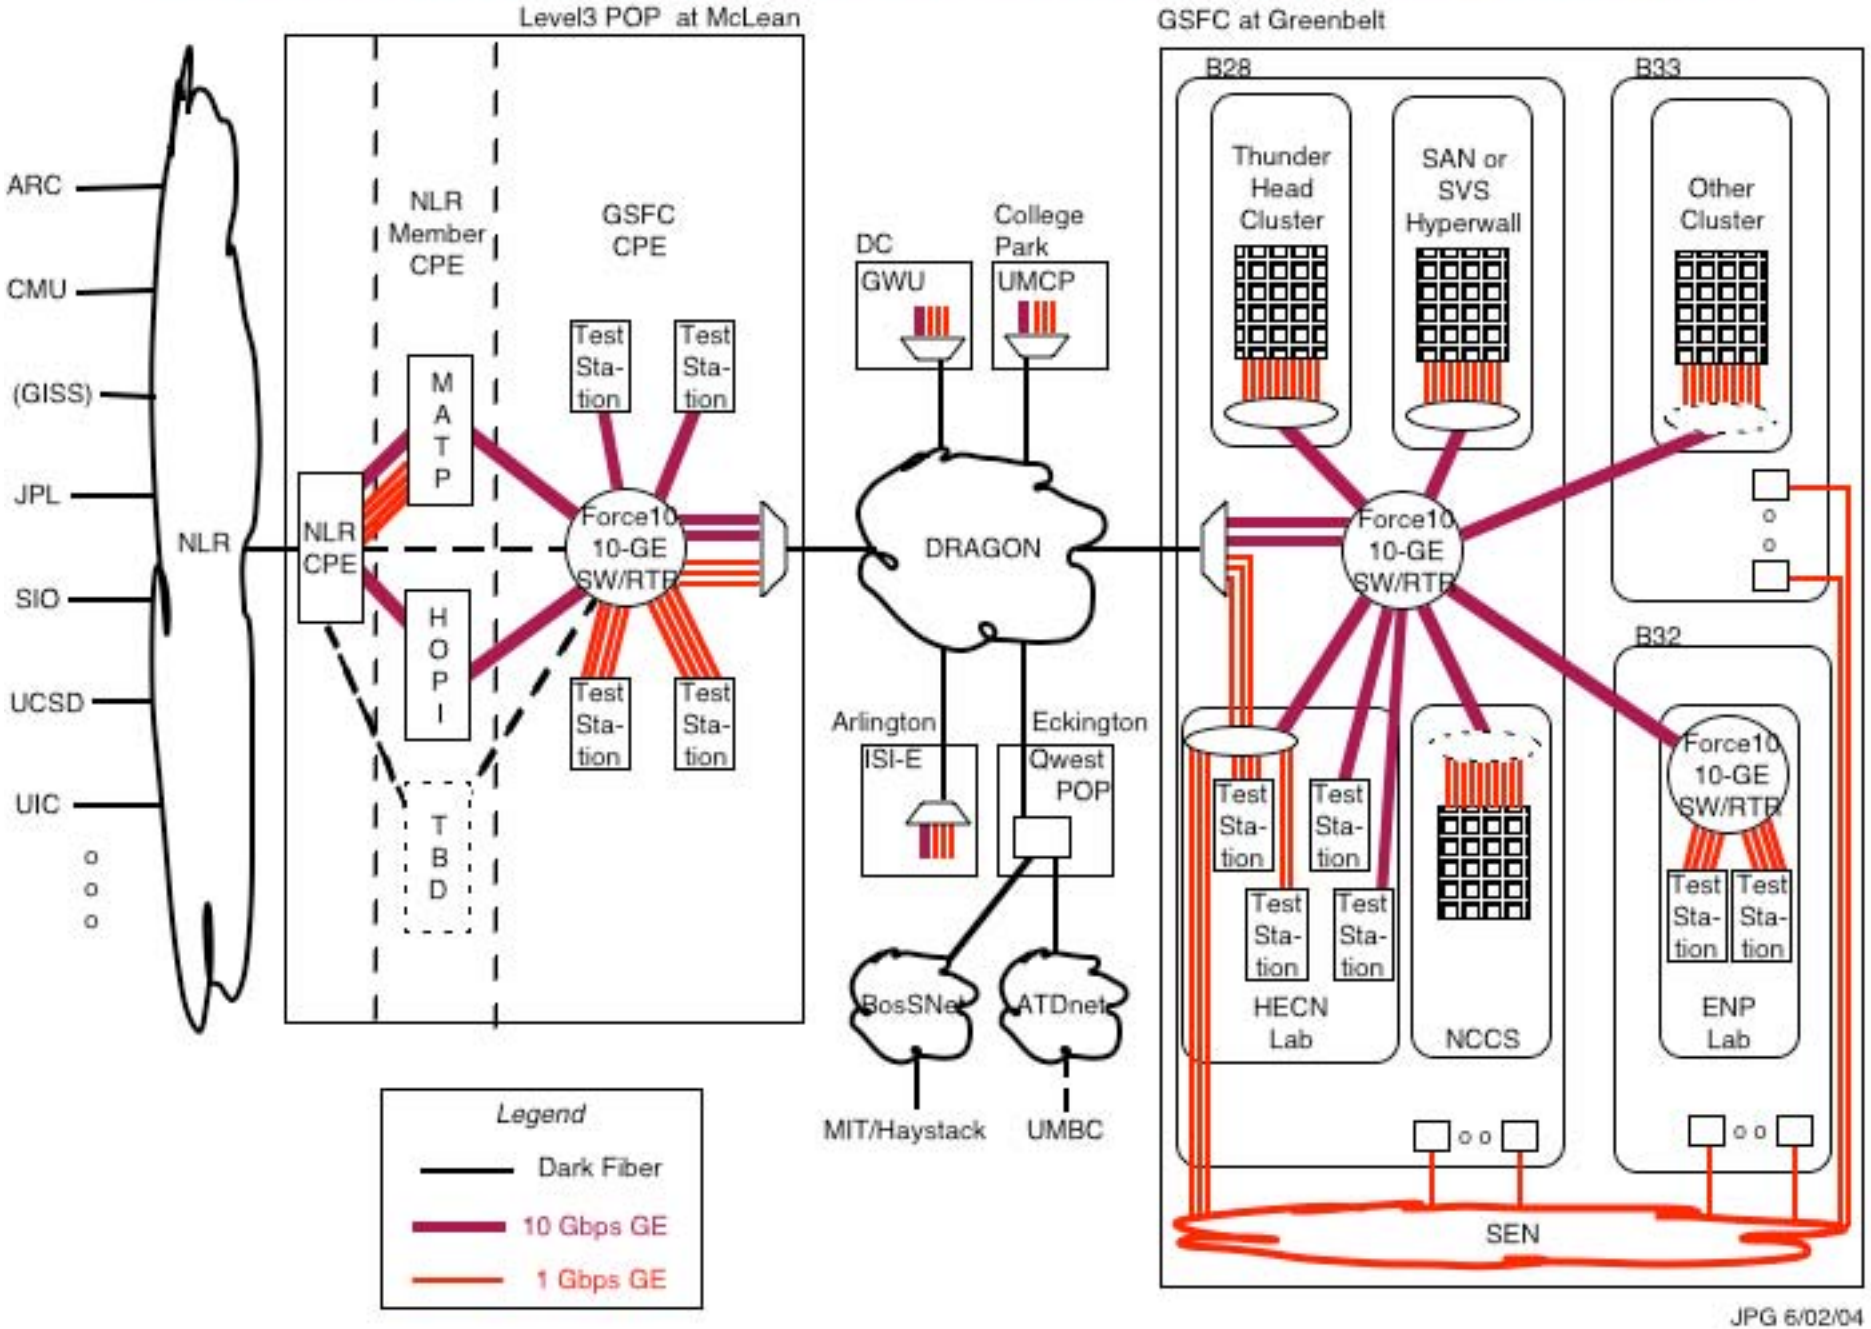

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Key Initial Features of GSFC L-Net Design

- GSFC Local Network Part
  - » 10-GE (Gigabit Ethernet) connections for:
    - – Thunderhead cluster in B28
    - – SVS Hyperwall and/or GSFC SAN Pilot interface in B28
    - – Network test stations in HECN lab in B28 and ENP lab in B32
    - – Optical switch from UMBC/Ray Chen
  - » 10-GE ports for future connections:
    - – NCCS in B28
    - – One other cluster such as Houser's in B33
    - – Others can be easily planned
  - » Four 1-GE connections with GSFC's Science and Engineering Network

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Key Features of GSFC L-Net Design (continued)

- Regional Network Part
  - » Two 10-GE connections with DRAGON at GSFC in Greenbelt
  - » Two 10-GE connections with DRAGON at Level3 POP in McLean
  - » Two 10-GE and multiple 1-GE connections for network test stations at Level3 POP in McLean
- Transcontinental Network Part
  - » 10-GE connection with NLR/MATP's IP Backbone and Switched Ethernet lambdas
  - » 10-GE connection with NLR/Internet2's HOPI lambda

GSFC L-Net Milestone Schedule

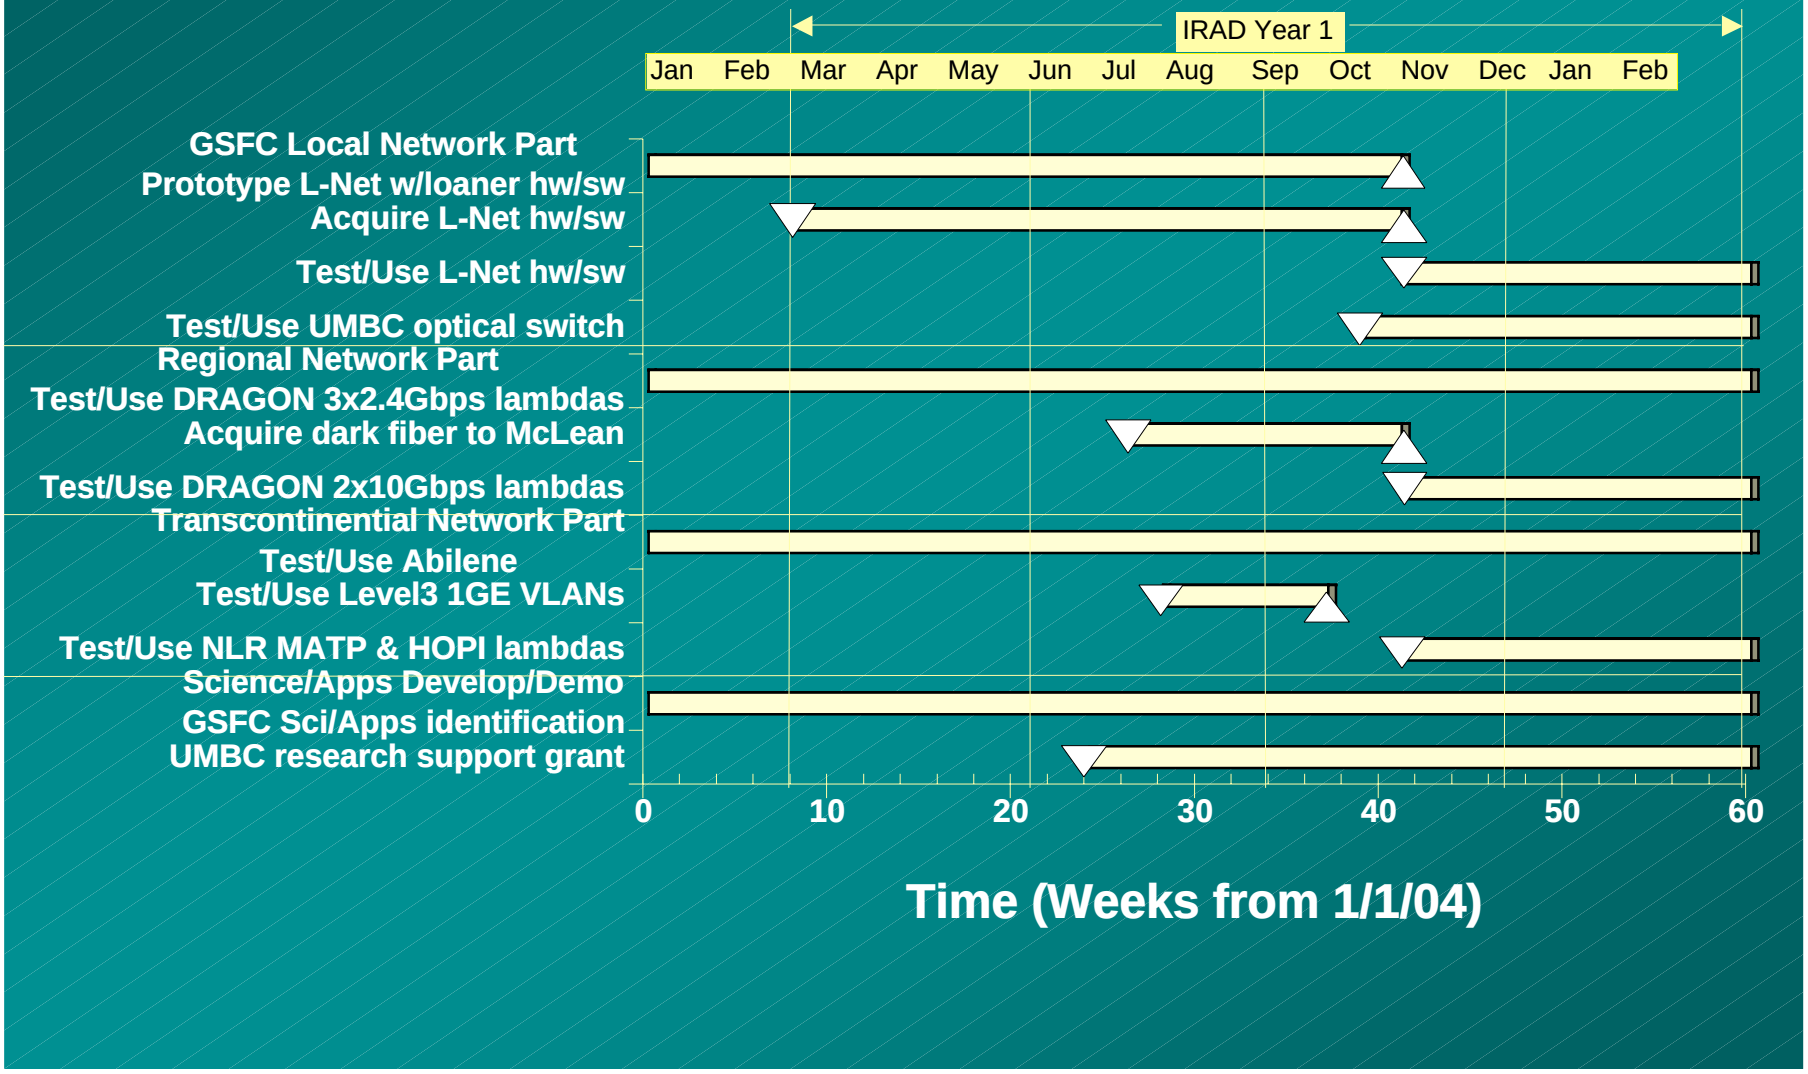

High Performance Remote Data Access Via GSFC L-Net Follow-on

FY05 L-Net Follow-on Considerations

- Enable high performance remote data access for all GSFC Earth and space science efforts via 10-GE lambda networking and SAN-over-IP technologies via creation of inter-facility virtual SANS
- Leverages
  - » Design and testing of the initial GSFC 10-GE L-Net funded by GSFC's FY04 IRAD program
    - – WAN: NLR's four 10 Gbps lambdas (via MATP membership, HOPI experimenting, OptIPuter participant, etc)
    - – MAN: DRAGON's two 10 Gbps and three 2.4 Gbps lambdas
    - – LAN: Force10's 10-GE switch/router with 12 "backbone" ports, and Extreme Network's 48x1-GE/2x10-GE "user uplink" switches
  - » Experience gained through GSFC SAN Pilot effort
    - – Especially SAN-over-IP via the new IETF protocols: iSCSI, FCIP, and iFCP

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

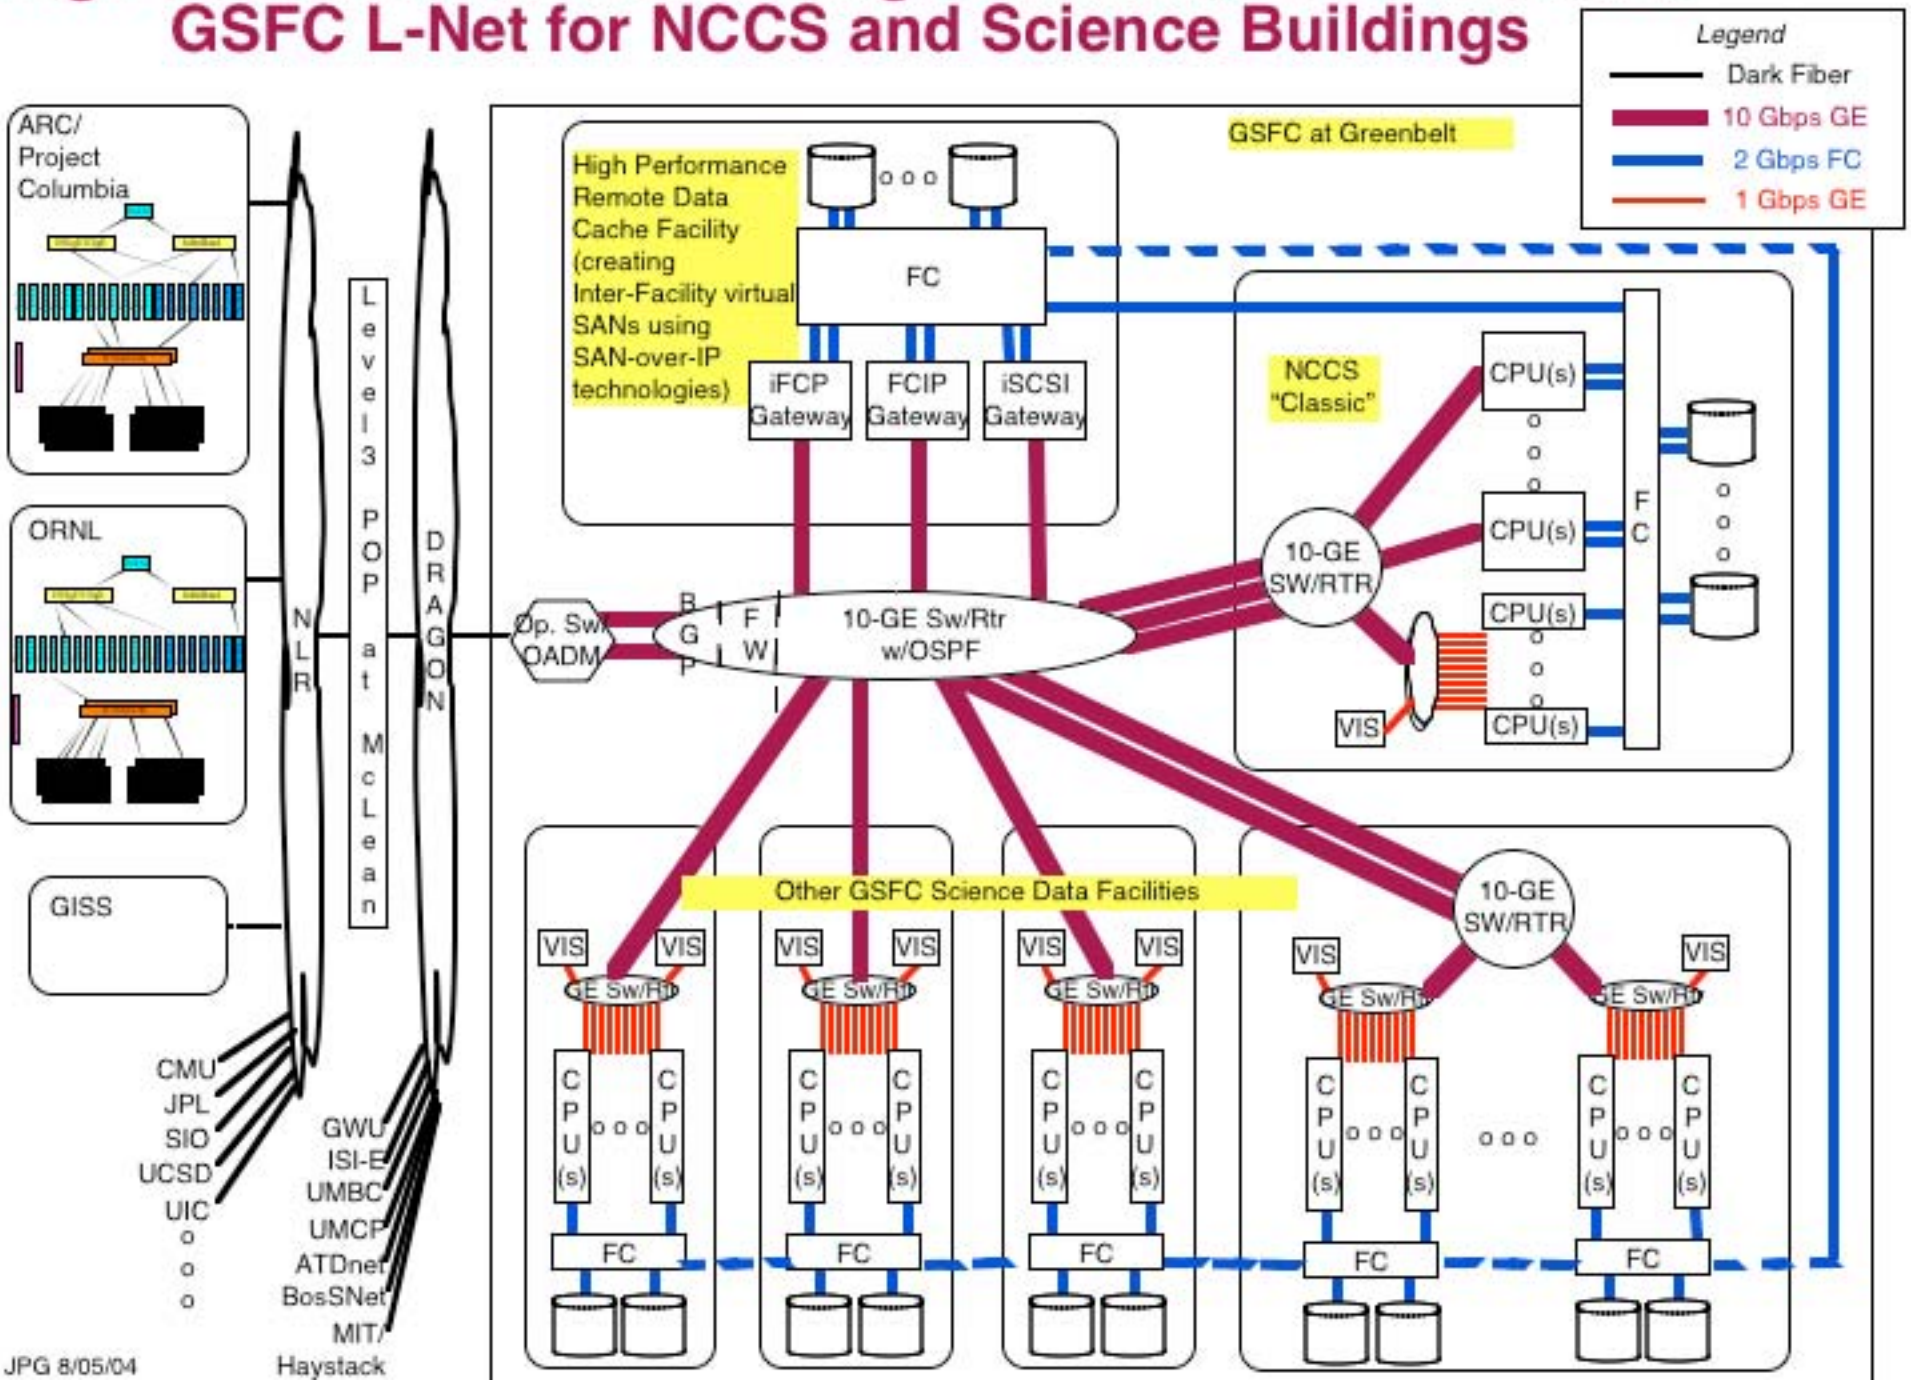

High Performance Remote Data Access Via GSFC L-Net Follow-on

Initial L-Net Follow-on Capability Possibility

- Wide Area SAN: CXFS-SGI between NAS and NCCS
  - » Demonstrate >15 MByte per second sustained data flows between ARC/Project Columbia and GSFC/NCCS within 9 months of full requested funding approval
- Dependencies or Drivers
  - » Funding required asap for GSFC L-Net's current support contractors
  - » Requires full cooperation by ARC/Project Columbia
  - » May need "security waivers" versus standard NASA EA proposals

High Performance Remote Data Access Via GSFC L-Net Follow-on

Other L-Net Follow-on Capability Possibilities

- Extending 10-GE L-Net within GSFC to more science buildings/clusters
- Dedicated 10-GE NLR lambda(s) between GSFC and:
  - NASA ARC
  - UCSD/SIO & OptIPuter
  - ORNL
  - UIC/OptIPuter
- GISS on shared or dedicated 10-GE NLR lambda
- Optical switch for both GSFC's East and West campuses

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

For More Info Related to the GSFC L-Net

- DRAGON - Dynamic Resource Allocation via GMPLS Optical Networks
  - http://dragon.east.isi.edu/
  - http://www.itrd.gov/iwg/lsn/jet/conferences/20040413/jetroadmapworkshop03.pdf
  - http://duster.nren.nasa.gov/workshop7/pps/04.Sobieski.DRAGON.ppt
- HOPI - Hybrid Optical and Packet Infrastructure
  - http://networks.internet2.edu/hopi/
  - http://www.itrd.gov/iwg/lsn/jet/conferences/20040413/jetroadmapworkshop10.pdf
  - http://duster.nren.nasa.gov/workshop7/pps/11.Winkler.HOPI.ppt
- Level3 - Level 3 Communications, Inc
  - http://www.level3.com/
  - http://www.itrd.gov/iwg/lsn/jet/conferences/20040413/jetroadmapworkshop31.pdf

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

For More Info Related to the GSFC L-Net (continued)

- MATP - Mid-Atlantic Terascale Partnership
  - http://www.midatlantic-terascale.org/
- Movaz Networks
  - http://www.movaz.com/
  - http://www.itrd.gov/iwg/lsn/jet/conferences/20040413/jetroadmapworkshop14.pdf
- NLR - National Lambda Rail
  - http://www.nlr.net/
  - http://www.itrd.gov/iwg/lsn/jet/conferences/20040413/jetroadmapworkshop43.pdf
  - http://duster.nren.nasa.gov/workshop7/pps/10.Farber.LambdaRail.ppt

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

For More Info Related to the GSFC L-Net (continued)

- OptIPuter
  - http://www.optiputer.net/
  - Smarr, L., Chien, A., DeFanti, T., Leigh, J., and Papadopoulos, P., "The OptIPuter", Communications of the ACM, November 2003, Vol. 46, No. 11, pp. 59-67; http://delivery.acm.org/10.1145/950000/948410/p58-smarr.pdf?key1=948410&key2=4972772901&coll=ACM&dl=ACM&CFID=25879183&CFTOKEN=50598646
  - http://duster.nren.nasa.gov/workshop7/pps/07.Smarr.OptIPuter.ppt

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Backup Slides

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Summary of Purpose

- "…establish a "Lambda Network" (in this case using optical wavelength technology and 10 Gbps Ethernet per wavelength) from GSFC's Earth science Greenbelt facility in MD to the Scripps Institute of Oceanography (SIO) through the University of California, San Diego (UCSD) facility over the National Lambda Rail (NLR), a new national dark optical fiber infrastructure."
- "…make data residing on Goddard's high speed computer disks available to SIO with access speeds as if the data were on their own desktop servers or PC's."
- "…enable scientists at both institutions to share and use compute intensive community models, complex data base mining and multi-dimensional streaming visualization over this highly distributed, virtual working environment."

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.**

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

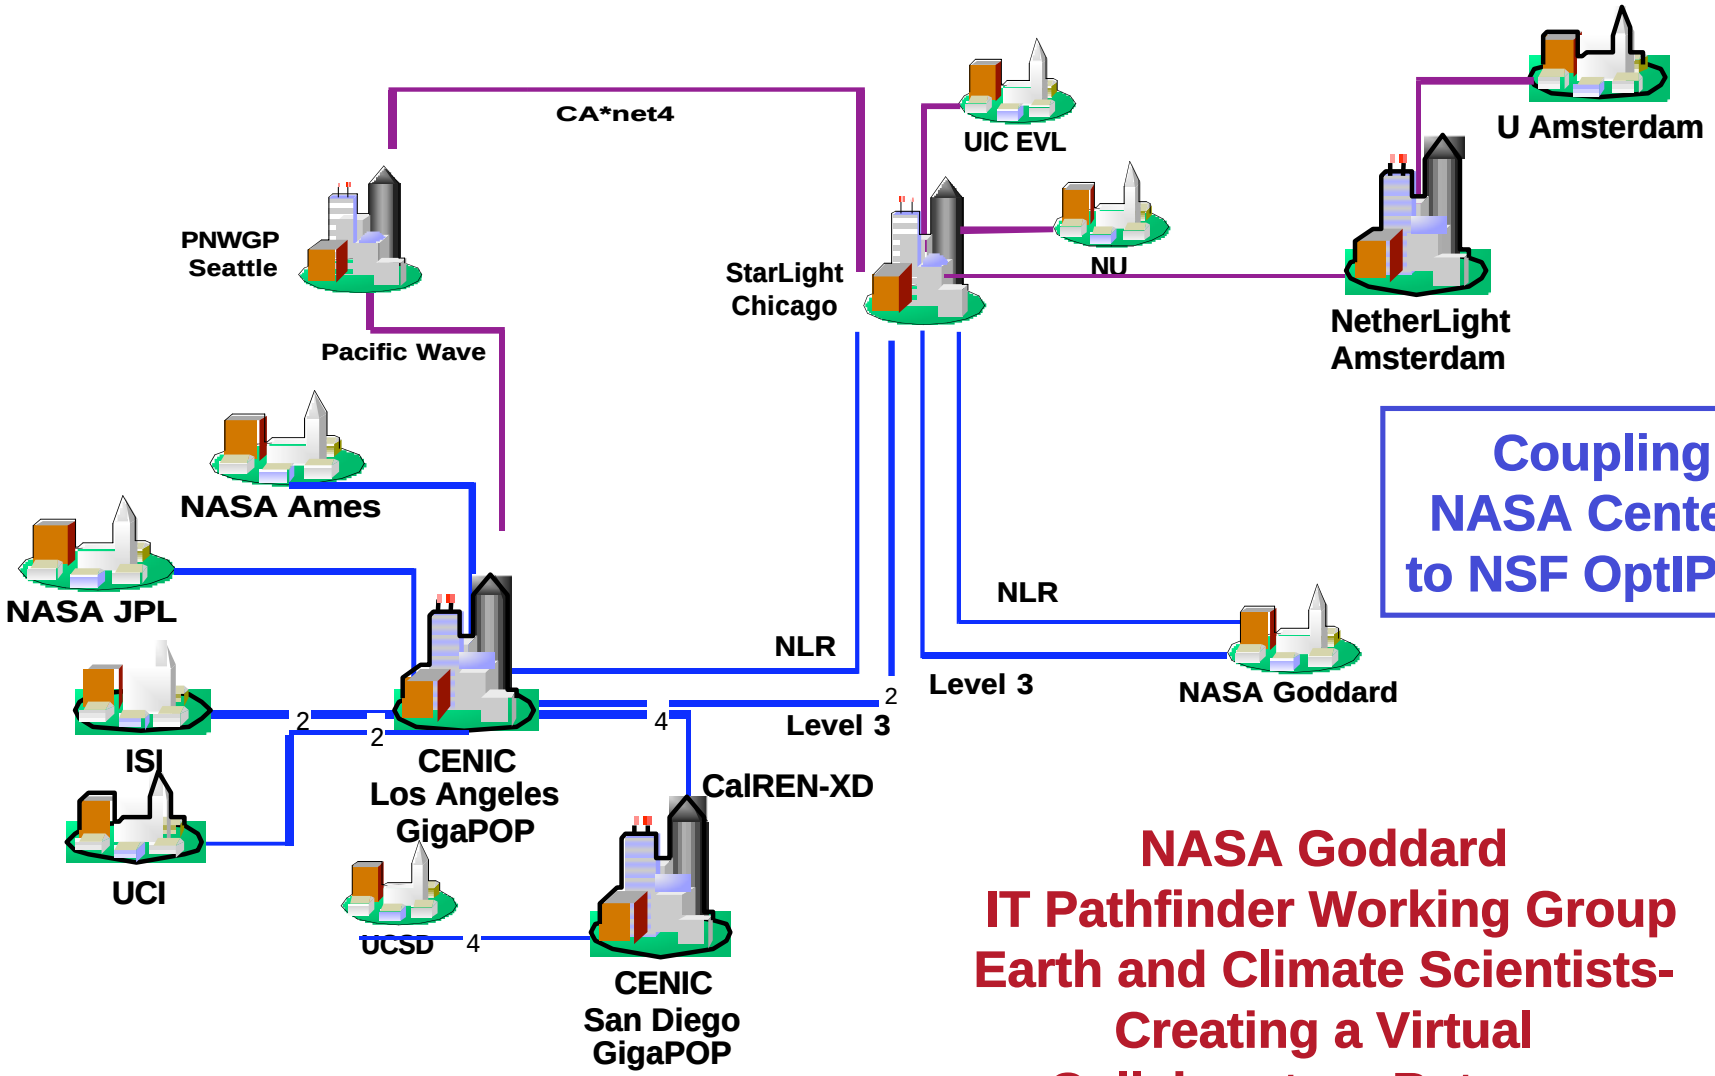

Coupling NASA Centers to NSF OptIPuter

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

Southern California CalREN-XD Build Out

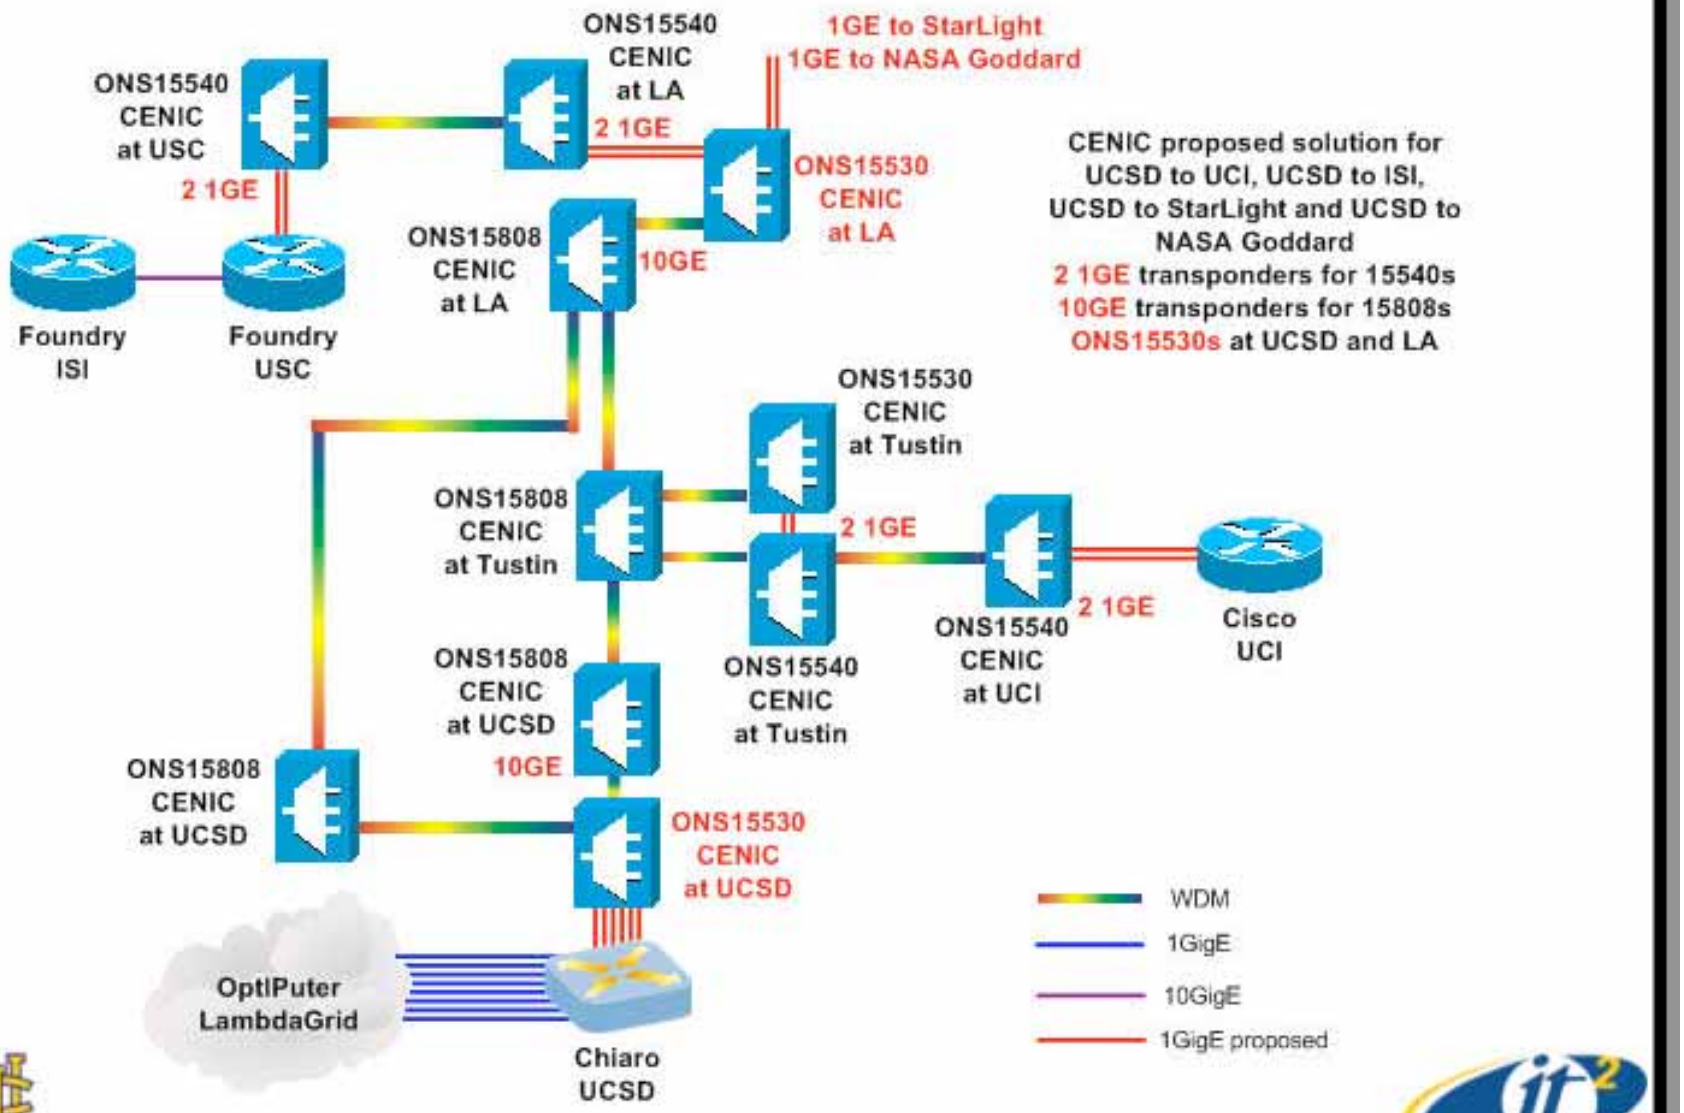

UCSD OptIPuter Nodes 2004

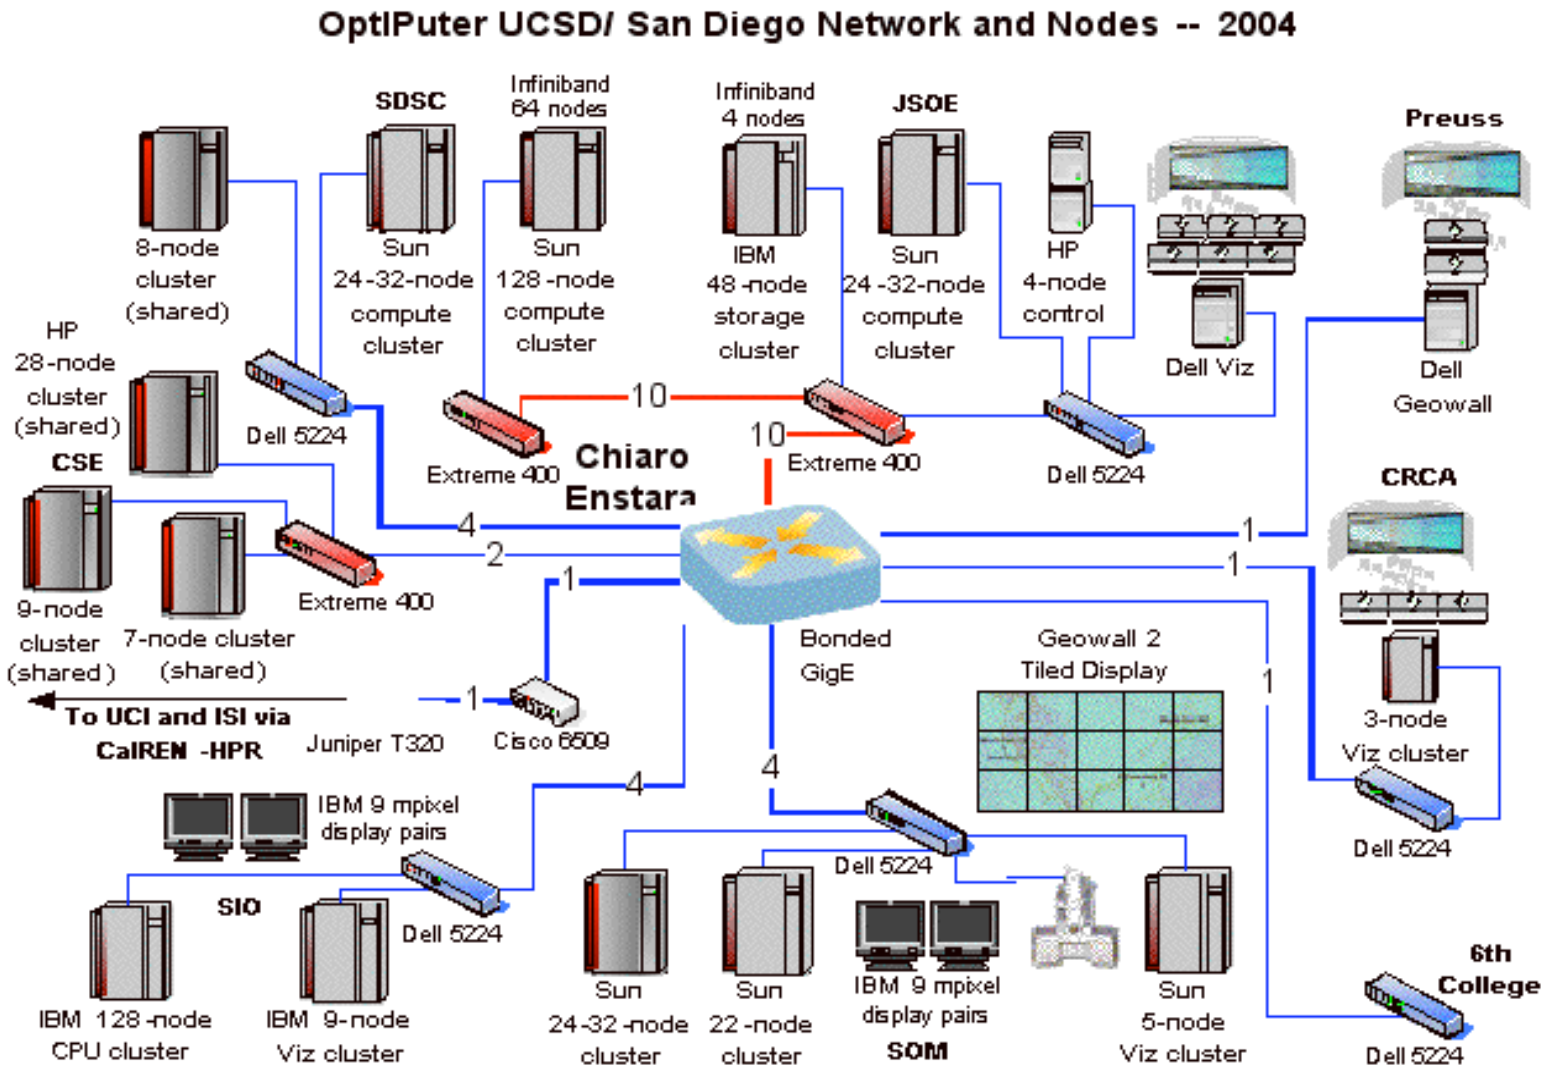

Fibers Considered in DRAGON-based Approach

- Fibers already “deployed” for all needed DRAGON links; and usage contracts in various states of progress
  - » “Link 1” GSFC -- UMD at College Park (UMCP): Fibers leased from FiberGate became available to be characterized on 10Jun04, were extended to GSFC building 28 on 22Jun04, and were planned to be characterized on 13Jul04
  - » “Link 2” UMCP -- ISI/E at Arlington & GWU in DC: Fibers leased from Qwest became available to be characterized on 28Jun04, and were planned to be characterized the week of 12Jul04
  - » “Link 3” ISI/E (& GWU) -- Level(3) POP in McLean: Have various quotes from Level3 including 20 year Indefeasible Right to Use (IRU). Negotiations are expected to be completed in Aug04

DRAGON Network

Washington, DC Metro Area

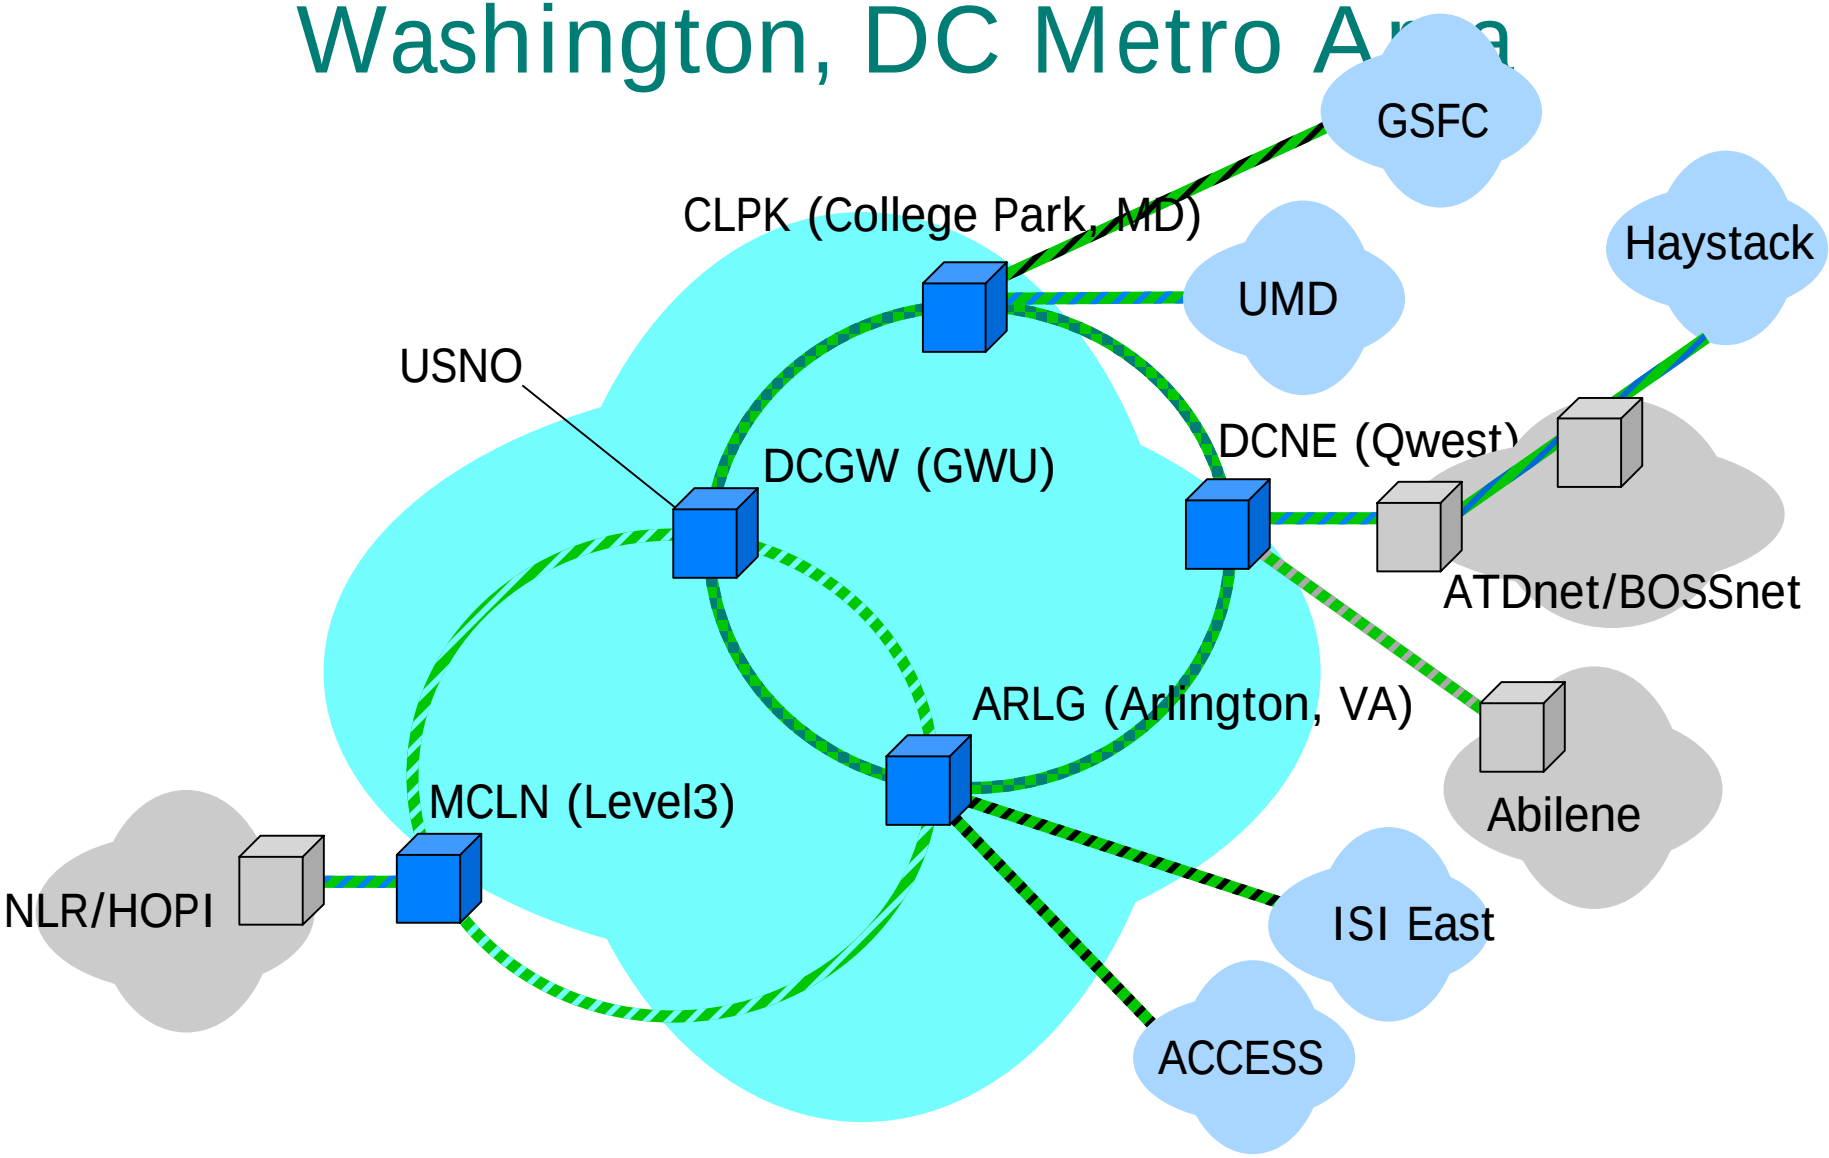

DRAGON eVLBI Experiment Configuration

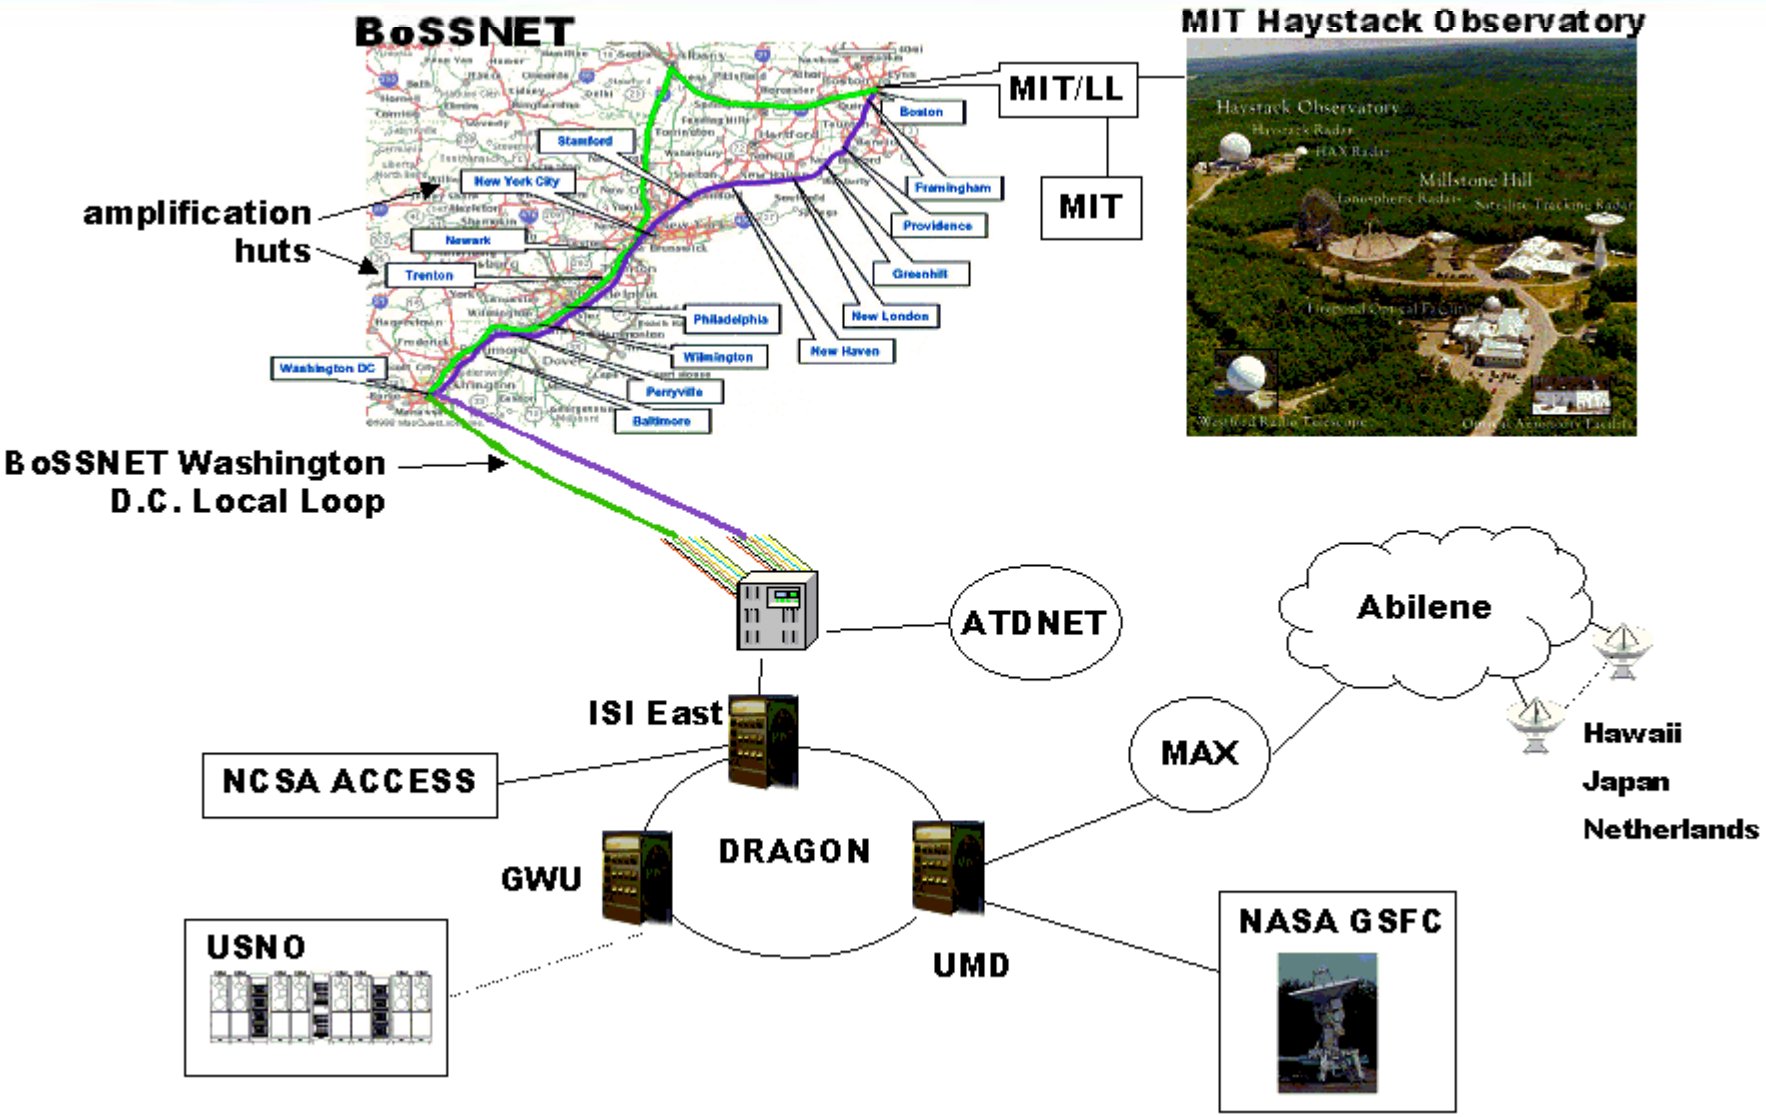

Testing of Force10 E300 10-GE Switch Capability Completed by Bill Fink (930) and Paul Lang (ADNET)

- Throughput testing
- Link Aggregation (two 10GigE links, similar physical and VLAN as above)
- Initial QoS testing (port based)
- VLAN Stacking

Force10 10GigE Testing

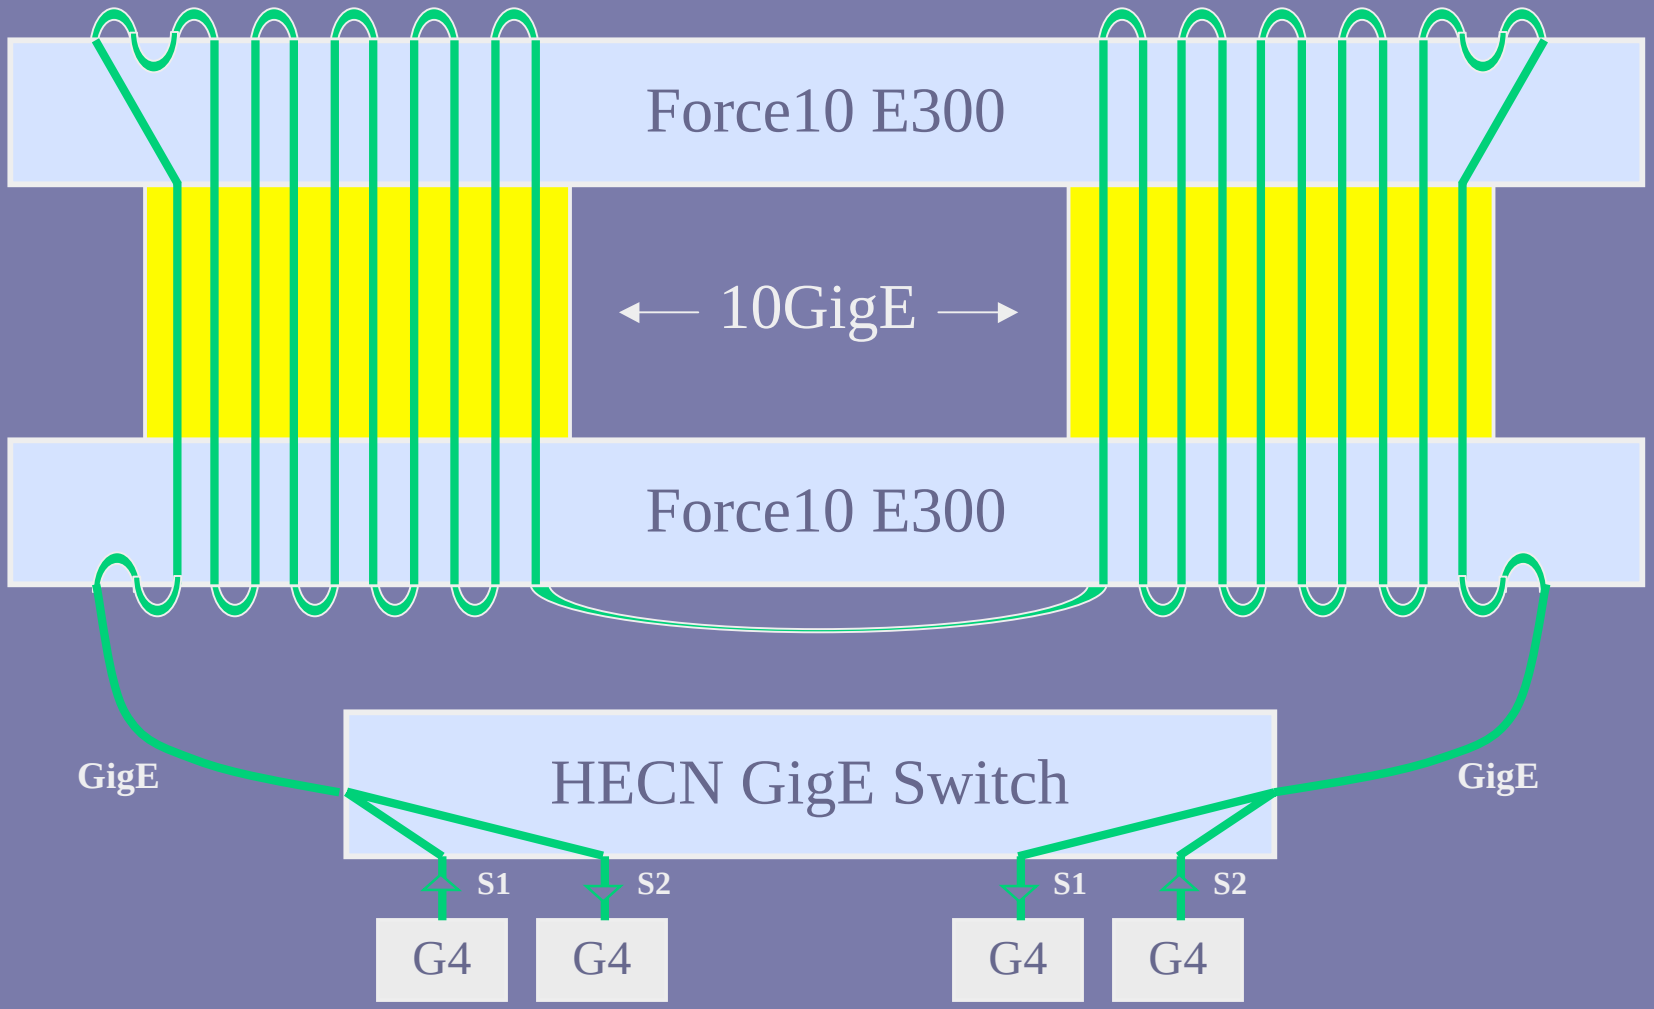

Test data generated with nuttcp and VLAN looping

Force10 10GigE nuttcp Stress Test (Over 1 Petabyte in 3 days *)

S1: porthos (2x800 MHz PowerMac G4) -> clifford (867 MHz PowerMac G4)
S2: underdog (867 MHz PowerMac G4) -> bigdog (2x1 GHz G4 XServe)

All systems running YellowDog Linux with Intel Pro/1000 Server Adapter or NetGear GA620T NIC, and using a 9K Jumbo Frame MTU

```
clifford% nuttcp -u -r -T72h -w2048 -Iporthos2clifford porthos & \
          nuttcp -u -T72h -w2048 -Iunderdog2bigdog underdog bigdog
underdog2bigdog:  30687801.2734 MB / 259772.86 sec =  990.9732 Mbps 74 %TX 29 %RX
                  1950222 / 3929988785 drop/pkt 0.05 %loss
porthos2clifford: 30695954.5781 MB / 259773.58 sec =  991.2337 Mbps 65 %TX 17 %RX
                  6056720340 / 9985802526 drop/pkt 60.65 %loss
```

(30687801.2734 MB + 30695954.5781 MB)*10*2/1024/1024/1024 = 1.14336 PB

* Over 0.25 PB transferred bidirectionally across each 10GigE link

Force10 Petabyte Challenge

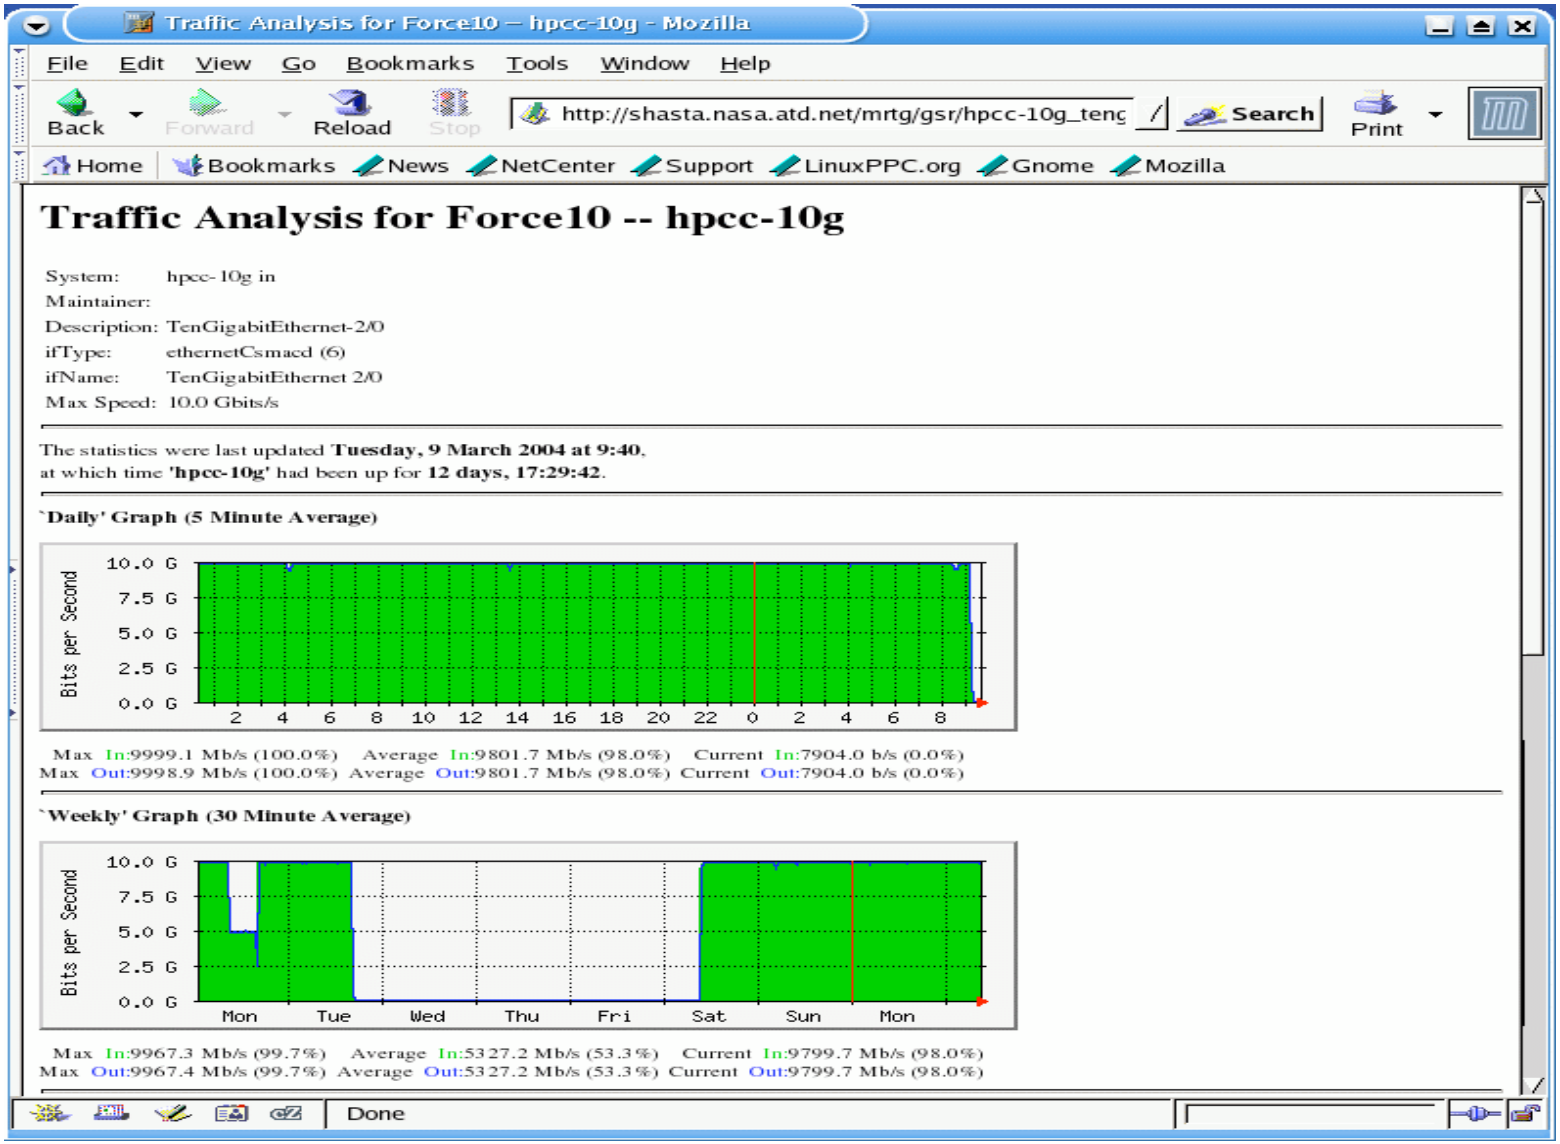

Graph for TenGigabitEthernet-3/0 basically identical

Abilene / NLR Map

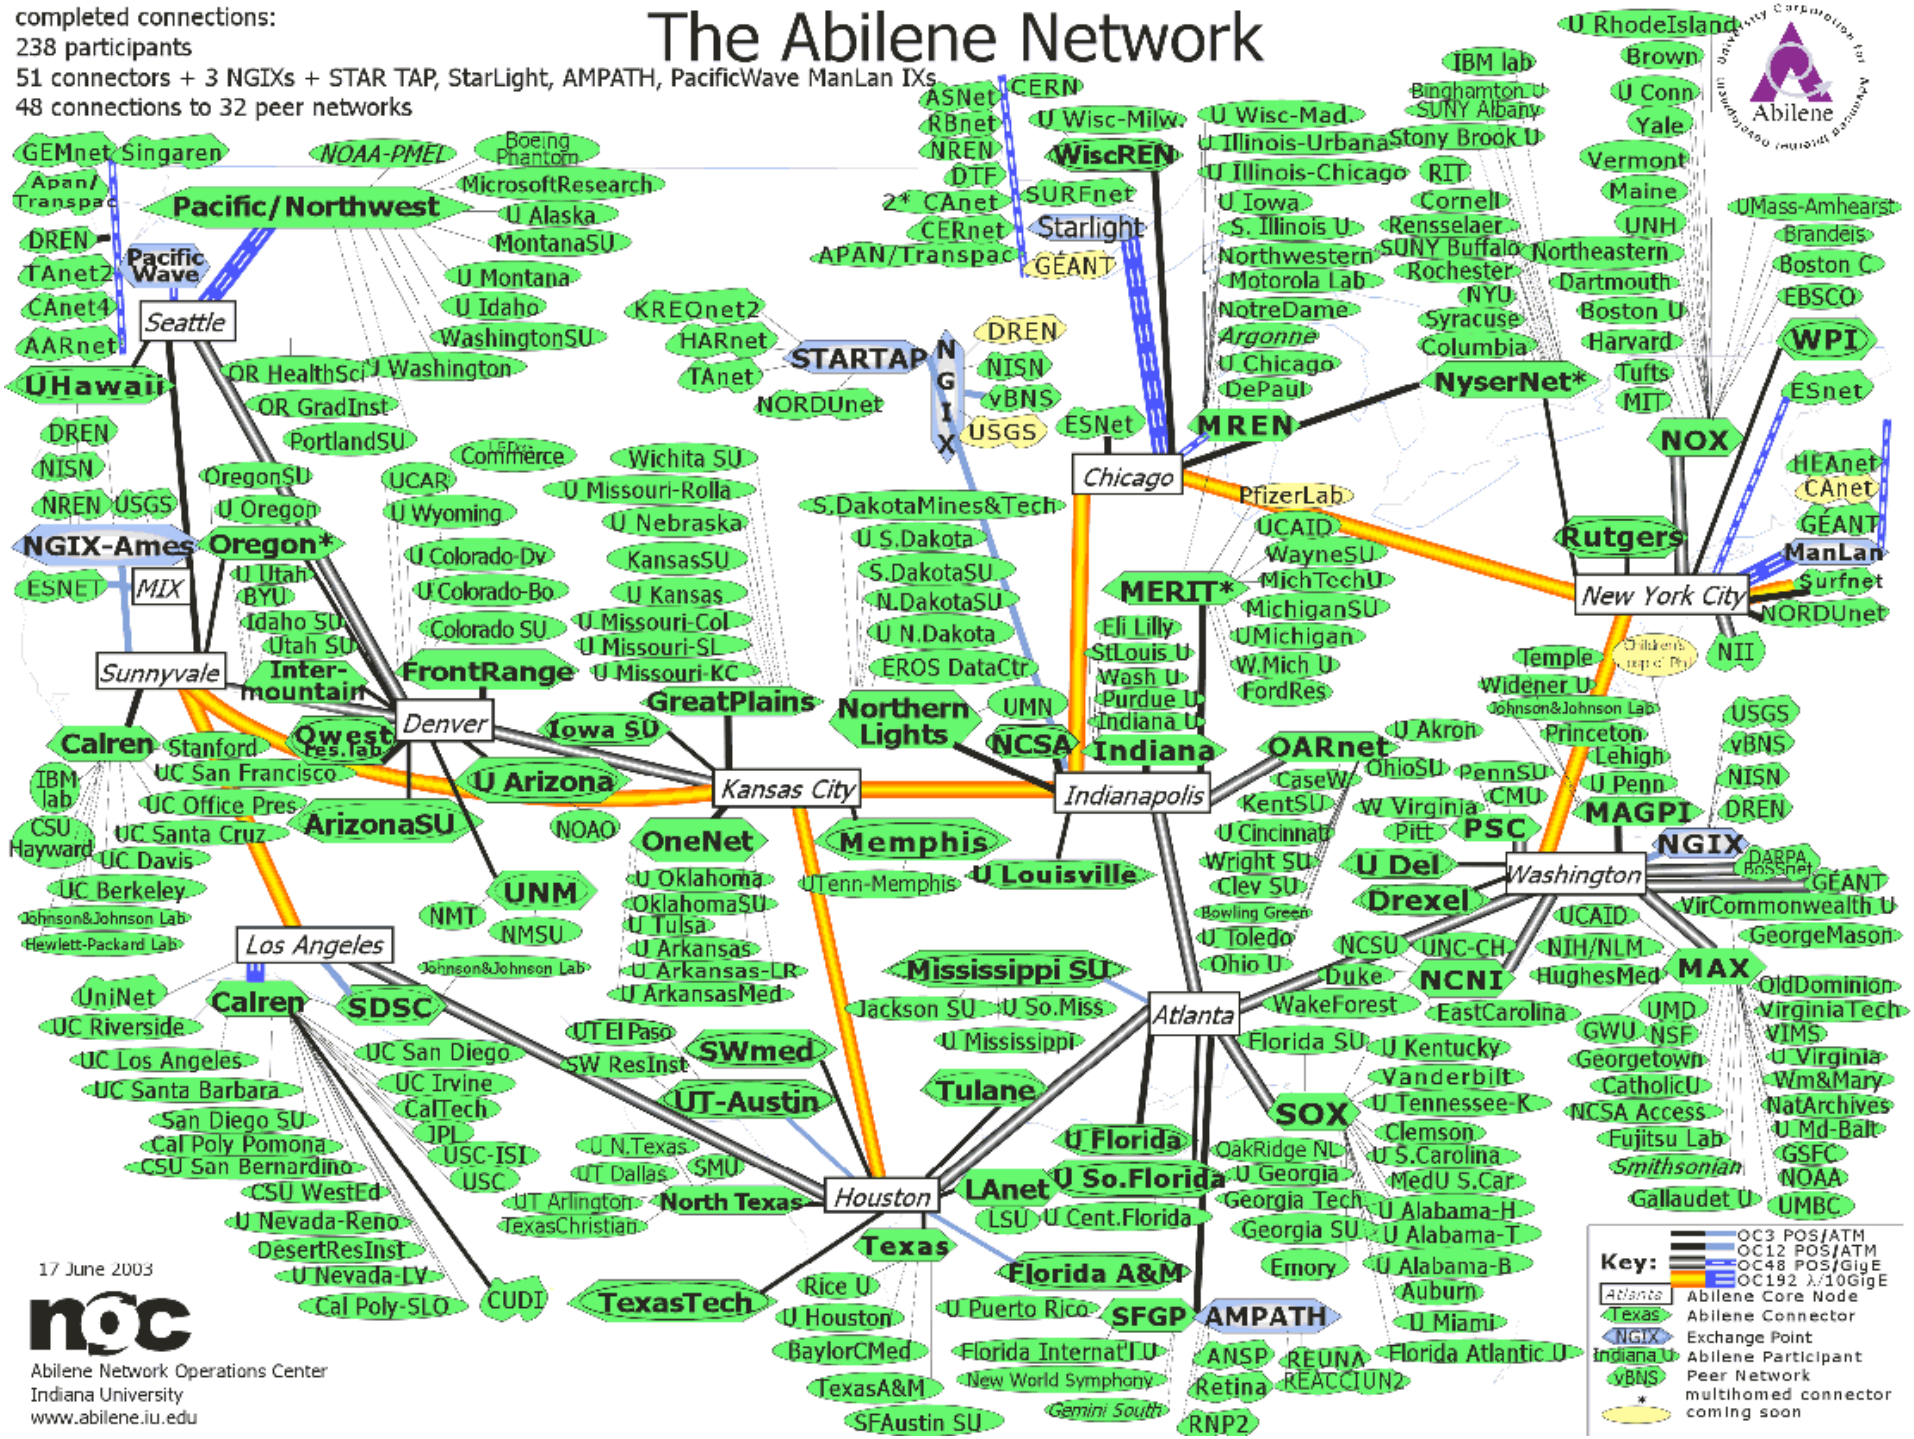
completed connections:
238 participants
51 connectors + 3 NGIXs + STAR TAP, StarLight, AMPATH, PacificWave ManLan IXs
48 connections to 32 peer networks
The Abilene Network
Abilene
GEMnet
Singaren
Apan/ Transpac
DREN
TAnet2
CAnet4
AARnet
Pacific Wave
Pacific/Northwest
NOAA-PMEL
Boeing Phantom
MicrosoftResearch
U Alaska
MontanaSU
U Montana
U Idaho
WashingtonSU
U Washington
OR HealthSci
OR GradInst
PortlandSU
Seattle
UHawaii
DREN
NISN
NREN
USGS
NGIX-Ames
ESNET
MIX
OregonSU
U Oregon
Oregon*
U Utah
BYU
Idaho SU
Utah SU
Inter-mountain
Sunnyvale
Commerce
UCAR
U Wyoming
U Colorado-Dv
U Colorado-Bo
Colorado SU
FrontRange
Denver
Wichita SU
U Missouri-Rolla
U Nebraska
KansasSU
U Kansas
U Missouri-Col
U Missouri-SL
U Missouri-KC
GreatPlains
Iowa SU
Kansas City
KREQnet2
HARnet
TAnet
NORDUnet
STARTAP
NGIX
DREN
NISN
vBNS
USGS
ESNet
CERN
ASNet
RBnet
NREN
DTF
2* CAnet
CERnet
APAN/Transpac
U Wisc-Milw
WiscREN
SURFnet
Starlight
GÉANT
Chicago
S.DakotaMines&Tech
U.S.Dakota
S.DakotaSU
N.DakotaSU
U N.Dakota
EROS DataCtr
Northern Lights
UMN
NCSA
Eli Lilly
StLouis U
Wash U
Purdue U
Indiana U
Indiana
Indianapolis
MERIT*
U Wisc-Mad
U Illinois-Urbana
U Illinois-Chicago
U Iowa
S. Illinois U
Northwestern
Motorola Lab
NotreDame
Argonne
U Chicago
DePaul
MREN
PfizerLab
UCAID
WayneSU
MichTechU
MichiganSU
UMichigan
W.Mich U
FordRes
IBM lab
Binghamton U
SUNY Albany
Stony Brook U
RIT
Cornell
Rensselaer
SUNY Buffalo
Rochester
NYU
Syracuse
Columbia
NyserNet*
U RhodeIsland
Brown
U Conn
Yale
Vermont
Maine
UNH
Northeastern
Dartmouth
Boston U
Harvard
Tufts
MIT
NOX
UMass-Amhearst
Brandeis
Boston C
EBSCO
WPI
ESnet
HEAnet
CAnet
GÉANT
ManLan
Surfnet
NORDUnet
NII
Rutgers
New York City
Temple
Widener U
Johnson&Johnson Lab
Princeton
Lehigh
U Penn
OARnet
U Akron
OhioSU
CaseW
KentSU
U Cincinnati
Wright SU
Clev SU
Bowling Green
U Toledo
Ohio U
PennSU
CMU
W Virginia
Pitt
PSC
MAGPI
USGS
vBNS
NISN
DREN
NGIX
DARPA BoSSnet
GÉANT
VirCommonwealth U
GeorgeMason
MAX
OldDominion
VirginiaTech
VIMS
U Virginia
Wm&Mary
NatArchives
U Md-Balt
GSFC
NOAA
UMBC
Gallaudet U
Washington
U Del
Drexel
UCAID
NIH/NLM
HughesMed
NCSU
UNC-CH
Duke
WakeForest
Florida SU
NCNI
EastCarolina
U Kentucky
Vanderbilt
U Tennessee-K
Clemson
U S.Carolina
MedU S.Car
U Alabama-H
U Alabama-T
U Alabama-B
Auburn
U Miami
Florida Atlantic U
SOX
OakRidge NL
U Georgia
Georgia Tech
Georgia SU
Emory
Atlanta
Mississippi SU
Jackson SU
U So.Miss
U Mississippi
Tulane
U Florida
U So.Florida
LAnet
LSU
U Cent.Florida
Florida A&M
SFGP
AMPATH
Florida Internat'l U
New World Symphony
Gemini South
ANSP
REUNA
REACCIUN2
Retina
RNP2
Houston
Texas
Rice U
U Houston
BaylorCMed
TexasA&M
SFAustin SU
TexasTech
U Puerto Rico
North Texas
U N.Texas
UT Dallas
SMU
UT Arlington
TexasChristian
SWmed
UT El Paso
SW ResInst
UT-Austin
Los Angeles
Calren
SDSC
UniNet
UC Riverside
UC Los Angeles
UC Santa Barbara
San Diego SU
Cal Poly Pomona
CSU San Bernardino
CSU WestEd
U Nevada-Reno
DesertResInst
U Nevada-LV
Cal Poly-SLO
CUDI
UC San Diego
UC Irvine
CalTech
JPL
USC-ISI
USC
Calren
Stanford
UC San Francisco
UC Office Pres
UC Santa Cruz
UC Davis
UC Berkeley
Johnson&Johnson Lab
Hewlett-Packard Lab
IBM lab
CSU Hayward
Qwest Res.lab
ArizonaSU
U Arizona
NOAO
UNM
NMT
NMSU
Johnson&Johnson Lab
OneNet
U Oklahoma
OklahomaSU
U Tulsa
U Arkansas
U Arkansas-LR
U ArkansasMed
Memphis
UTenn-Memphis
U Louisville
17 June 2003
noc
Abilene Network Operations Center
Indiana University
www.abilene.iu.edu
Key:
OC3 POS/ATM
OC12 POS/ATM
OC48 POS/GigE
OC192 λ/10GigE
Atlanta Abilene Core Node
Texas Abilene Connector
NGIX Exchange Point
Indiana U Abilene Participant
vBNS Peer Network
multihomed connector
coming soon

GSFC FY04 IRAD Proposal "Preparing Goddard for Large Scale Team Science in the 21st Century: Enabling an All Optical Goddard Network Cyberinfrastructure"

Principal Investigator & Co-Investigators

- Name: Pat Gary (930) & Jeff Smith (585) Co-PI's & GSFC's Information Technology Pathfinder Working Group (ITPWG) as Co-I's
- Organizations: Code 420, Code 580, Code 920, & Code 930
- Telephone: 301-286-9539 & 301-614-5038 for Co-PI's
- E-mail: Pat.Gary@nasa.gov, Jeff.Smith@nasa.gov for Co-PI's