StarLight

Tom DeFanti, Maxine Brown, Jason Leigh, Alan Verlo, UIC
Joe Mambretti, David Carr, Tim Ward, NU
Linda Winkler, Bill Nickless, Caren Litvanyi, ANL

StarLight: an Experimental Optical Infrastructure

Perhaps the World's Largest 1GigE and 10GigE Exchange

Operational since summer 2001, StarLight is a 1GigE and 10GigE switch/router facility for high-performance access to participating networks and is becoming a true optical switching facility for wavelengths.

Abbott Hall, Northwestern University's Chicago downtown campus

STARLIGHT℠

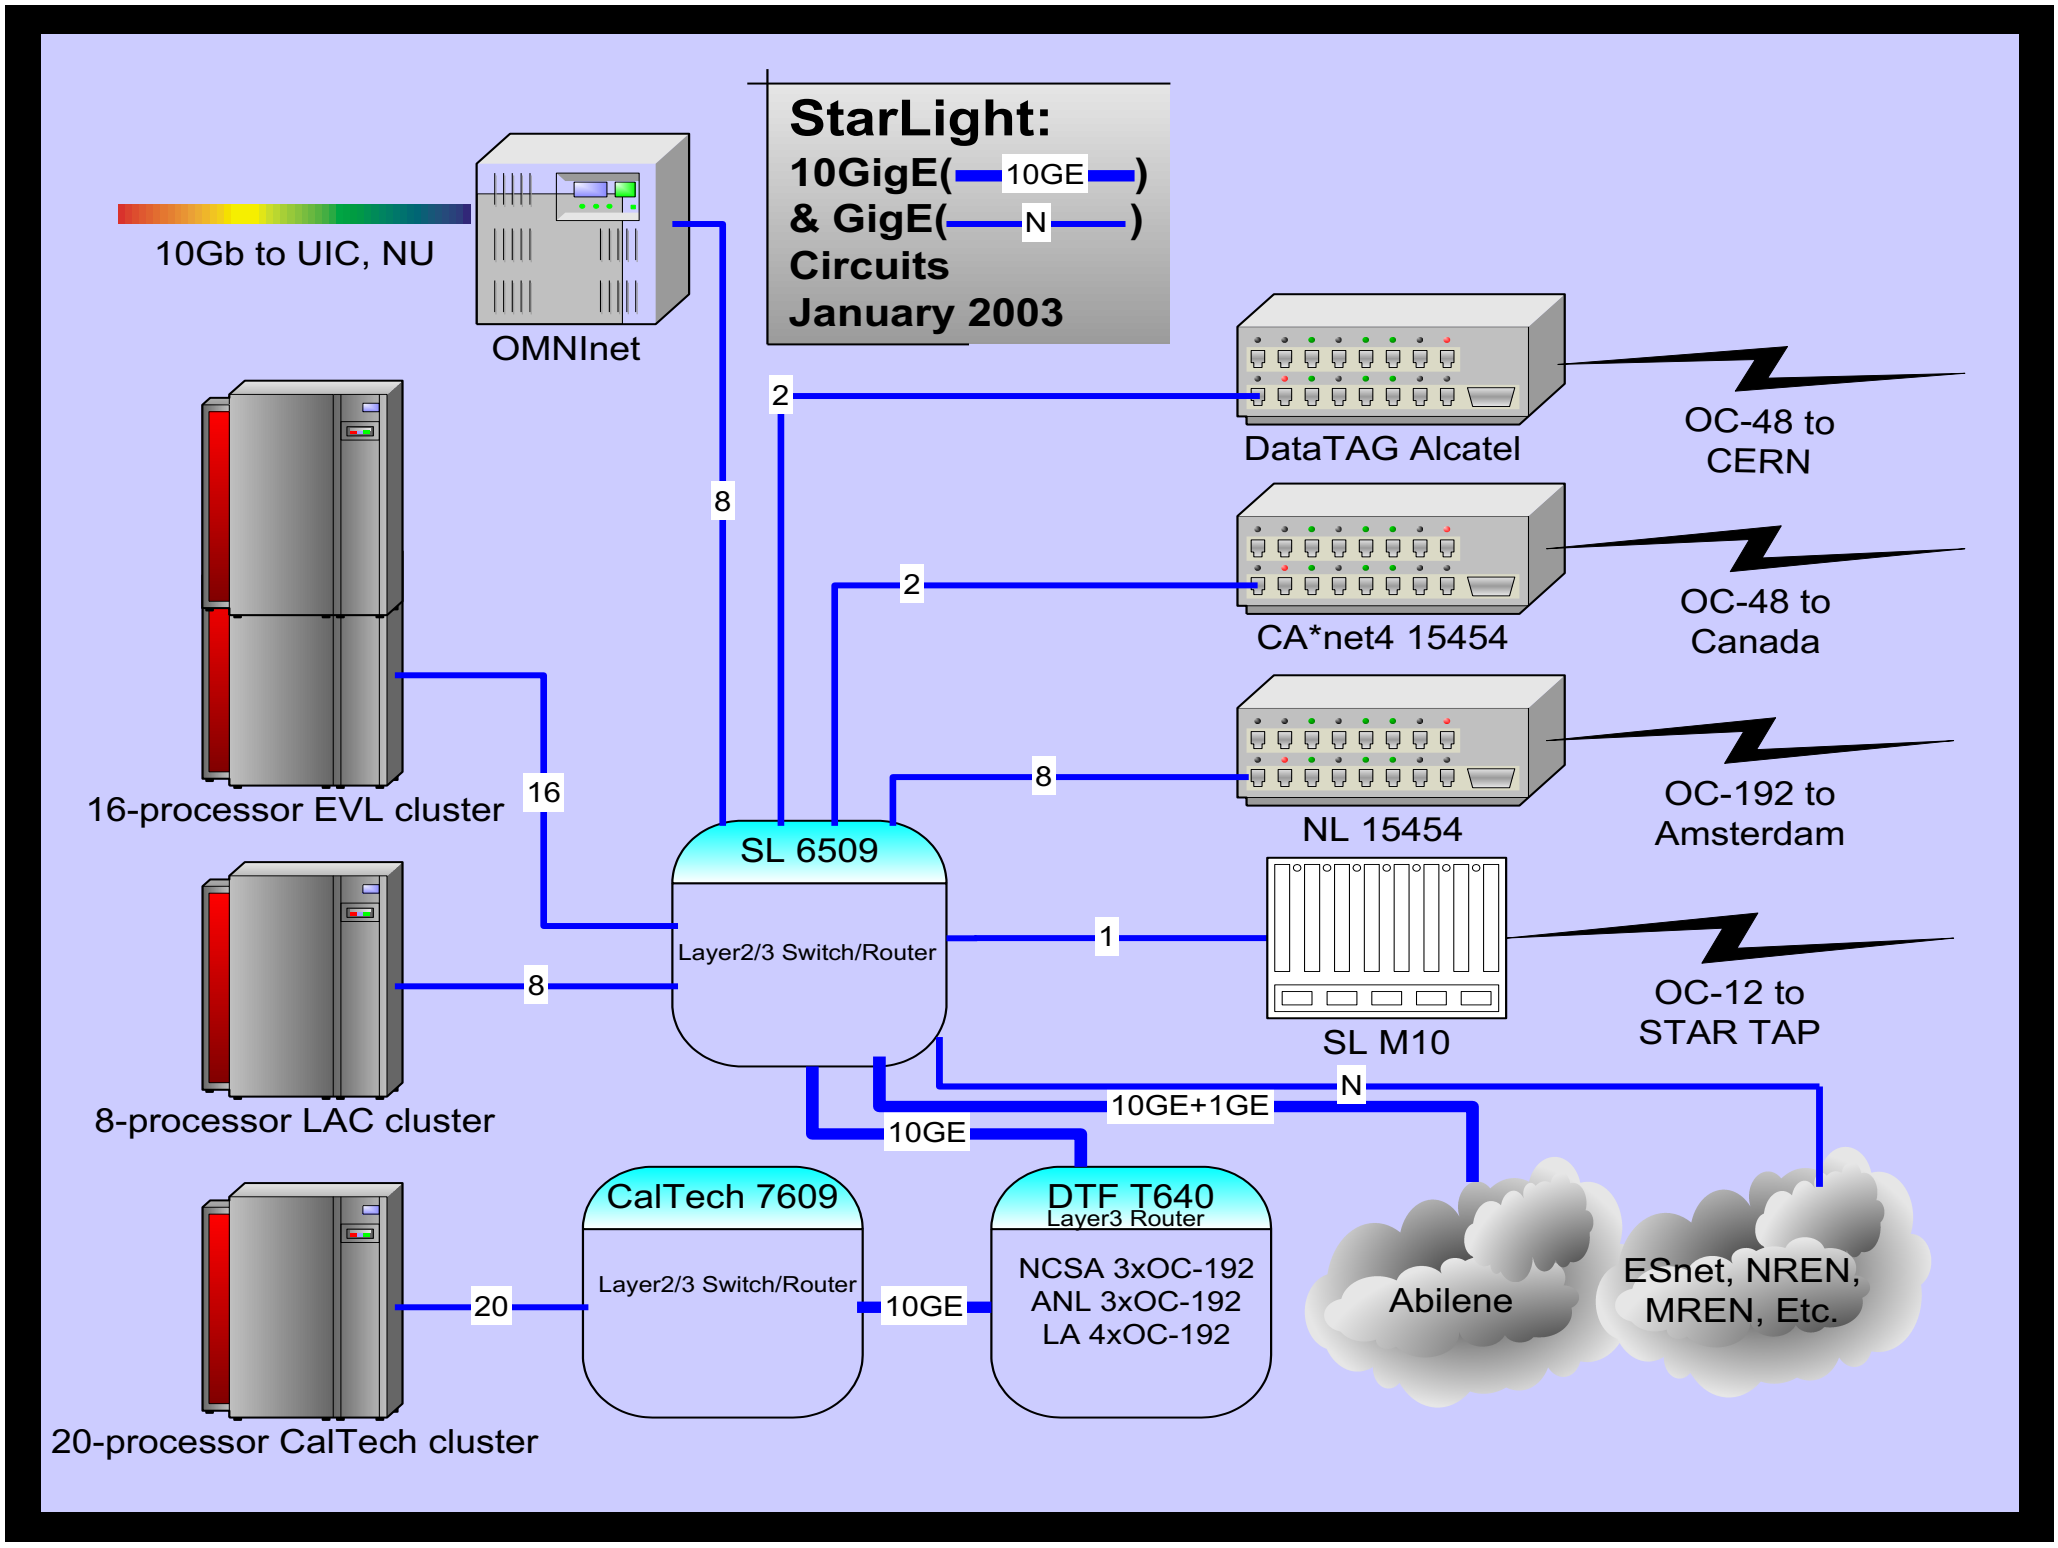
StarLight:
10GigE(10GE)
& GigE(N)
Circuits
January 2003
10Gb to UIC, NU
OMNInet
8
2
DataTAG Alcatel
OC-48 to CERN
2
CA*net4 15454
OC-48 to Canada
8
NL 15454
OC-192 to Amsterdam
16-processor EVL cluster
16
SL 6509
Layer2/3 Switch/Router
1
SL M10
OC-12 to STAR TAP
8
8-processor LAC cluster
N
10GE+1GE
10GE
CalTech 7609
Layer2/3 Switch/Router
DTF T640
Layer3 Router
NCSA 3xOC-192
ANL 3xOC-192
LA 4xOC-192
Abilene
ESnet, NREN, MREN, Etc.
20
10GE
20-processor CalTech cluster

StarLight Wants to be Everyone's *Other End*

- A **Production Network** 1GigE and 10GigE exchange
- An **Experimental Network** lambda exchange
- A **Research Network** 1GigE and 10GigE MEMS-switched exchange
- Host to DTFnet, the TeraGrid's 4x10Gb T640-based Experimental Network
- A co-location space with 66 racks for networking *and* computing and data-management equipment
- An OIX with fiber and/or circuits from SBC/Ameritech, Qwest, AT&T, Global Crossing, Looking Glass Networks, Level 3, RCN, Deutsche Telekom/T-Systems, I-WIRE
- A **lambda-switching** facility, with links coming from NetherLight, DataTAG, CA*net4, and proposed from UK-Light and APAN forming ***Trans-Light***

STARLIGHT℠

What is a Lambda?

- A *lambda*, in networking, is a fully dedicated wavelength of light in an optical network, typically used today for 1-10Gbps.
- Lambdas are circuit-based technology, but can carry packet-based information.
- We are now mostly working with 1Gb dedicated Layer 2 circuits that act like lambdas

STARLIGHT℠

A Proposed Global *Trans-Light* Project to Prototype the LambdaGrid

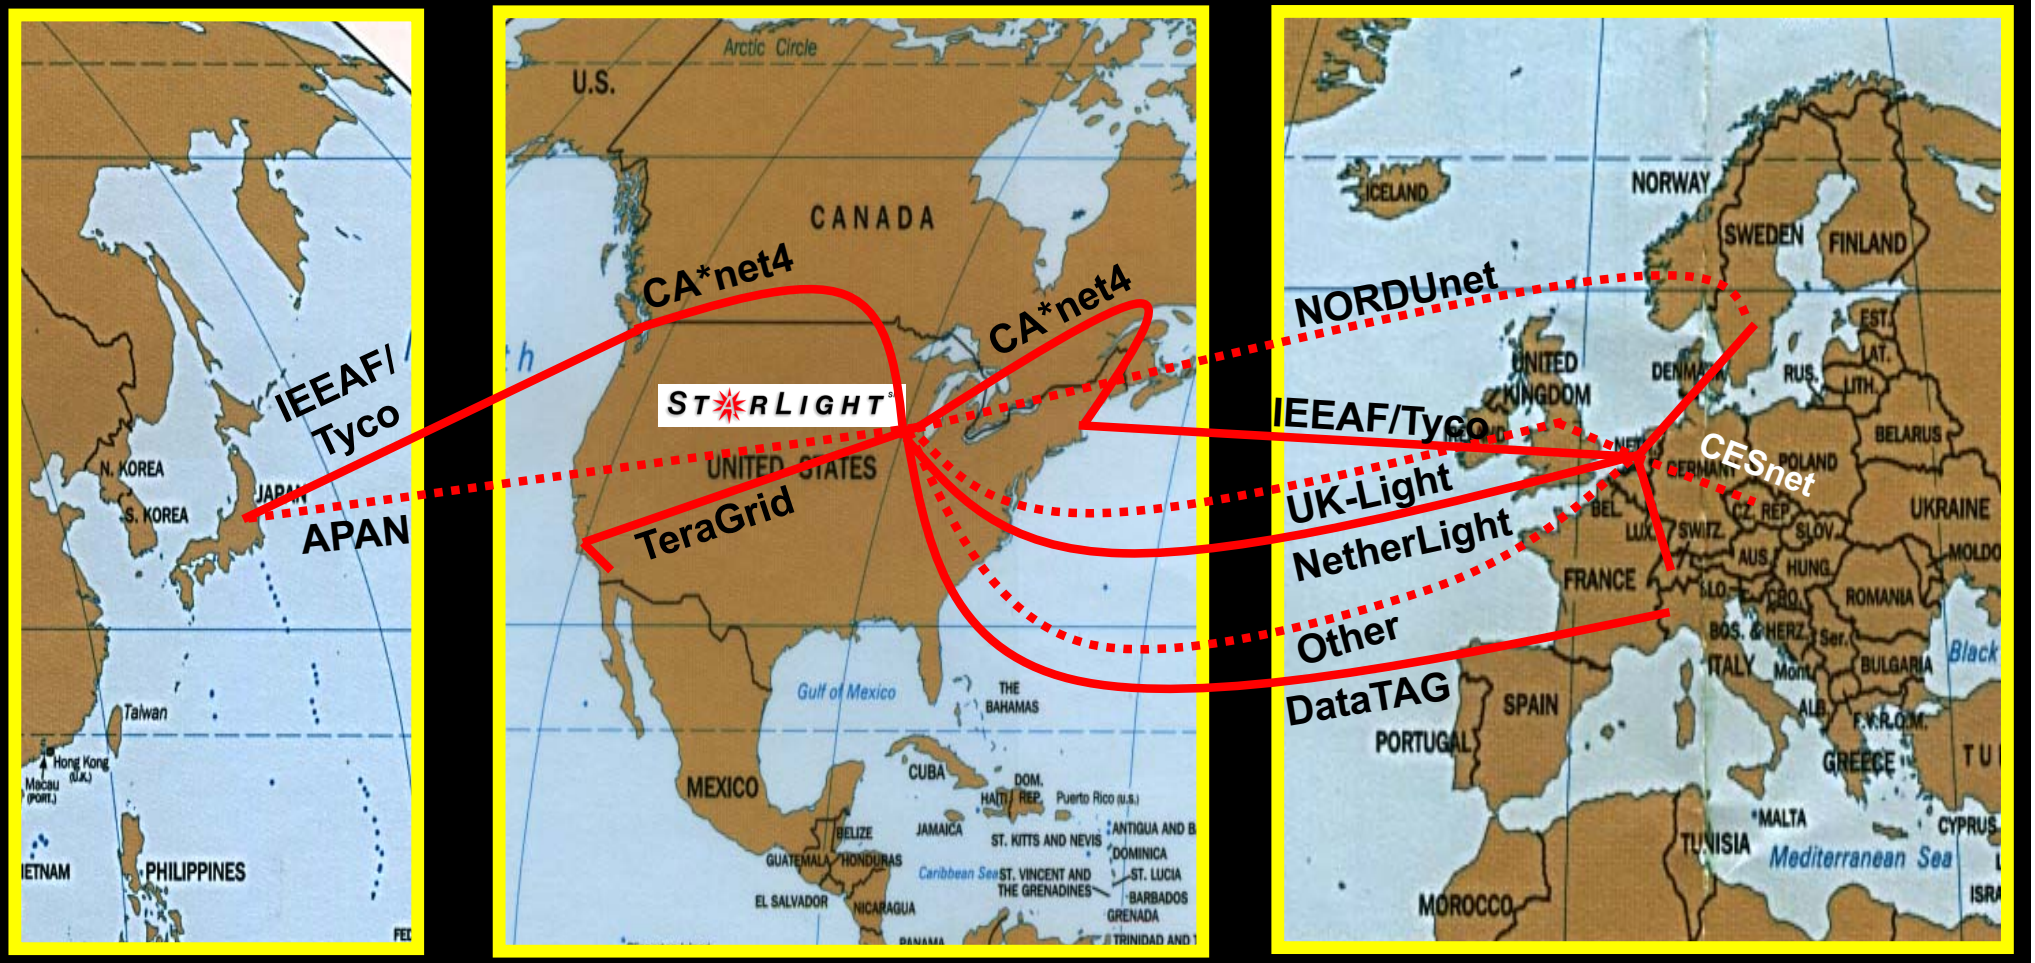

2.5Gb and 10Gb Links

2.5Gb and 10Gb Links (Proposed)

STARLIGHT℠

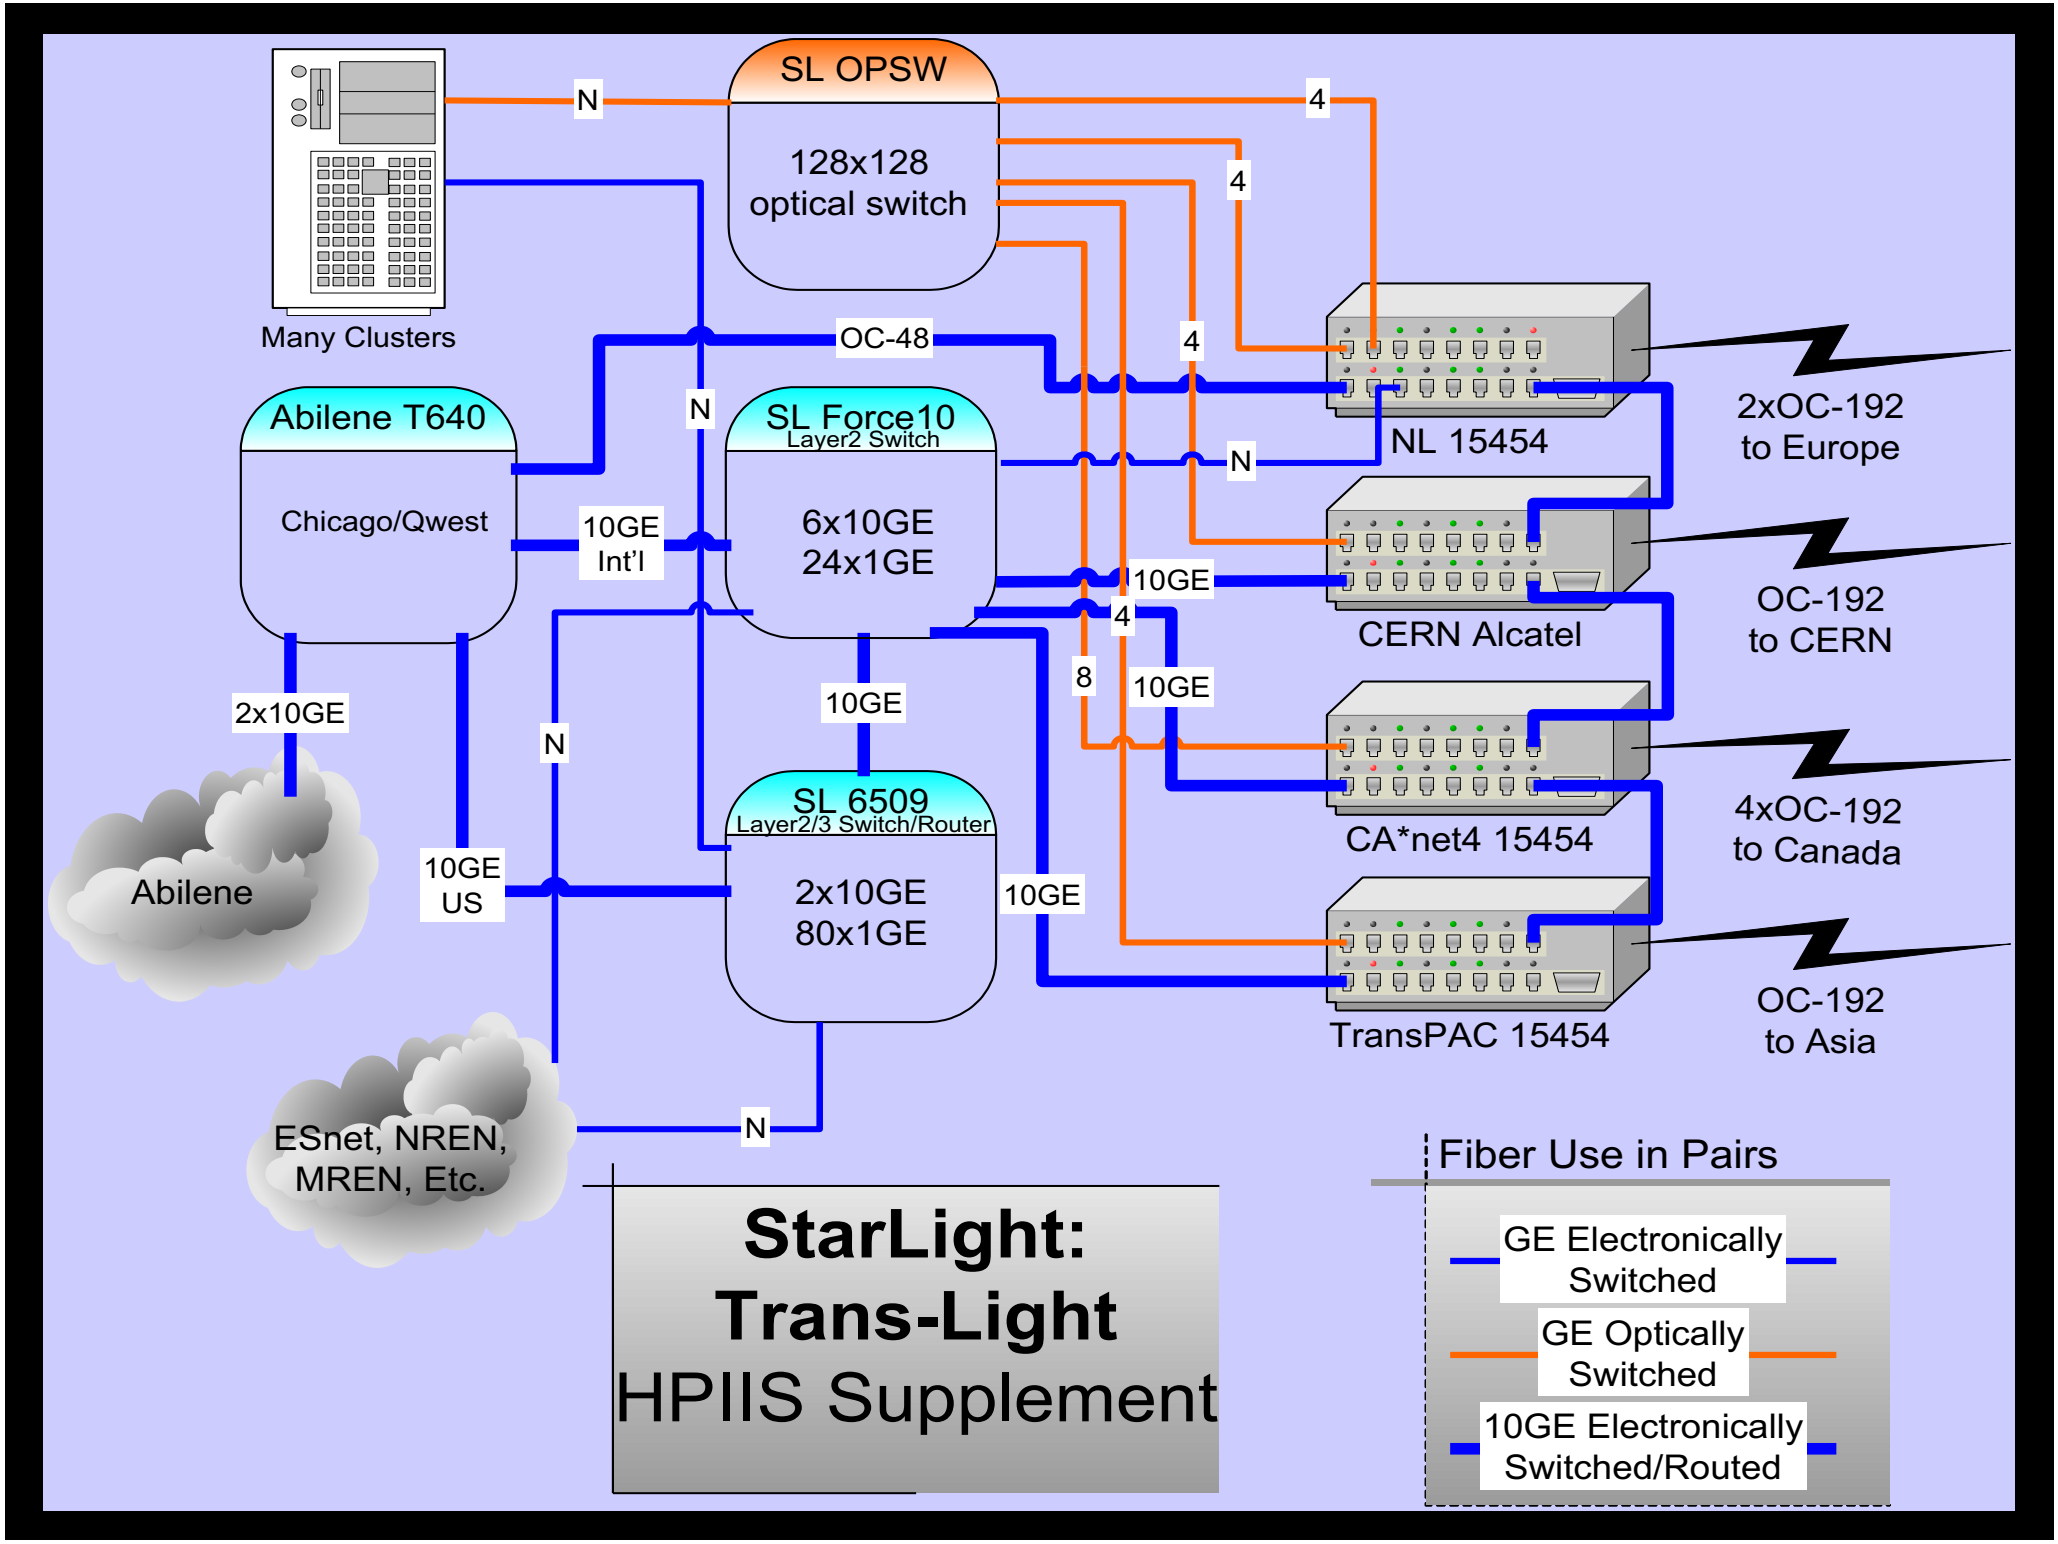
Many Clusters
N
SL OPSW
128x128
optical switch
4
4
4
OC-48
N
Abilene T640
Chicago/Qwest
10GE
Int'l
SL Force10
Layer2 Switch
6x10GE
24x1GE
N
10GE
4
8
10GE
2x10GE
10GE
N
SL 6509
Layer2/3 Switch/Router
2x10GE
80x1GE
10GE
10GE
US
Abilene
N
ESnet, NREN,
MREN, Etc.
NL 15454
2xOC-192
to Europe
CERN Alcatel
OC-192
to CERN
CA*net4 15454
4xOC-192
to Canada
TransPAC 15454
OC-192
to Asia
StarLight:
Trans-Light
HPIIS Supplement
Fiber Use in Pairs
GE Electronically
Switched
GE Optically
Switched
10GE Electronically
Switched/Routed

iGrid 2002
September 23-26, 2002, Amsterdam, The Netherlands

- **28 demonstrations from 16 countries:** Australia, Canada, CERN/Switzerland, France, Finland, Germany, Greece, Italy, Japan, Netherlands, Singapore, Spain, Sweden, Taiwan, the United Kingdom and the United States
- **Applications** demonstrated: art, bioinformatics, chemistry, cosmology, cultural heritage, education, high-definition media streaming, manufacturing medicine, neuroscience, physics, tele-science

- **Grid technologies** demonstrated: Major emphasis on grid middleware, data management grids, data replication grids, visualization grids, data/visualization grids, computational grids, access grids, grid portals
- **25Gb transatlantic bandwidth (100Mb/attendee, 250x iGrid 2000!)**

www.startap.net/igrid2002

STARLIGHT℠

Reliable Blast UDP (RBUDP)

- At IGrid 2002 *all* applications which were able to make the most effective use of the 10G link from Chicago to Amsterdam used **UDP**
- RBUDP[1], SABUL[2] and Tsunami[3] are all similar protocols that use UDP for bulk data transfer. All are based on NETBLT- RFC969
- RBUDP has fewer memory copies & a prediction function to let applications know what kind of performance to expect.

  - [1] J. Leigh, O. Yu, D. Schonfeld, R. Ansari, et al., Adaptive Networking for Tele-Immersion, Proc. Immersive Projection Technology/Eurographics Virtual Environments Workshop (IPT/EGVE), May 16-18, Stuttgart, Germany, 2001.
  - [2] Sivakumar Harinath, Data Management Support for Distributed Data Mining of Large Datasets over High Speed Wide Area Networks, PhD thesis, University of Illinois at Chicago, 2002.
  - [3] http://www.indiana.edu/~anml/anmlresearch.html

Source: Jason Leigh, UIC

STARLIGHT℠

Large-Scale International Application Development

	Guaranteed Latency	Guaranteed Scheduling	Guaranteed Bandwidth
Access Grid USA, Canada, The Netherlands, UK, Italy, Germany, Russia, Australia, China, Korea, Thailand, Taiwan, Japan, Brazil	•		
BABAR USA and internationally			•
The D0 Experiment USA, CERN, Germany, France, Japan and other worldwide collaborators			•
GiDVN: Global Internet Digital Video Network CCIRN DVWG worldwide membership	•	•	
Hubble Space Telescope USA, France and others worldwide			•
SC Global USA and International	•		
Sloan Digital Sky Survey (SDSS) USA, France and worldwide			•
Virtual Room Videoconferencing System (VRVS) CERN, Switzerland; Caltech, USA; Others	•		
vlbiGrid USA, The Netherlands, Finland, UK and worldwide			•

STARLIGHT℠

Large-Scale International Middleware and Toolkit Development

	Guaranteed Latency	Guaranteed Scheduling	Guaranteed Bandwidth
EU DataGrid CERN, France, Italy, The Netherlands, UK, Czech Republic, Finland, Germany, Hungary, Spain, Sweden (in cooperation with US grid projects, notably GriPhyN, PPDG and iVGL	•	•	•
EU DataTAG USA and Europe			•
Globally Interconnected Object Databases (GIOD) USA and CERN			•
Globus USA, Sweden, others internationally	•	•	•
MONARC for LHC Experiments CERN, Switzerland; Caltech, USA; Others			•
UK e-Science Programme UK and USA	•	•	•

STARLIGHT℠

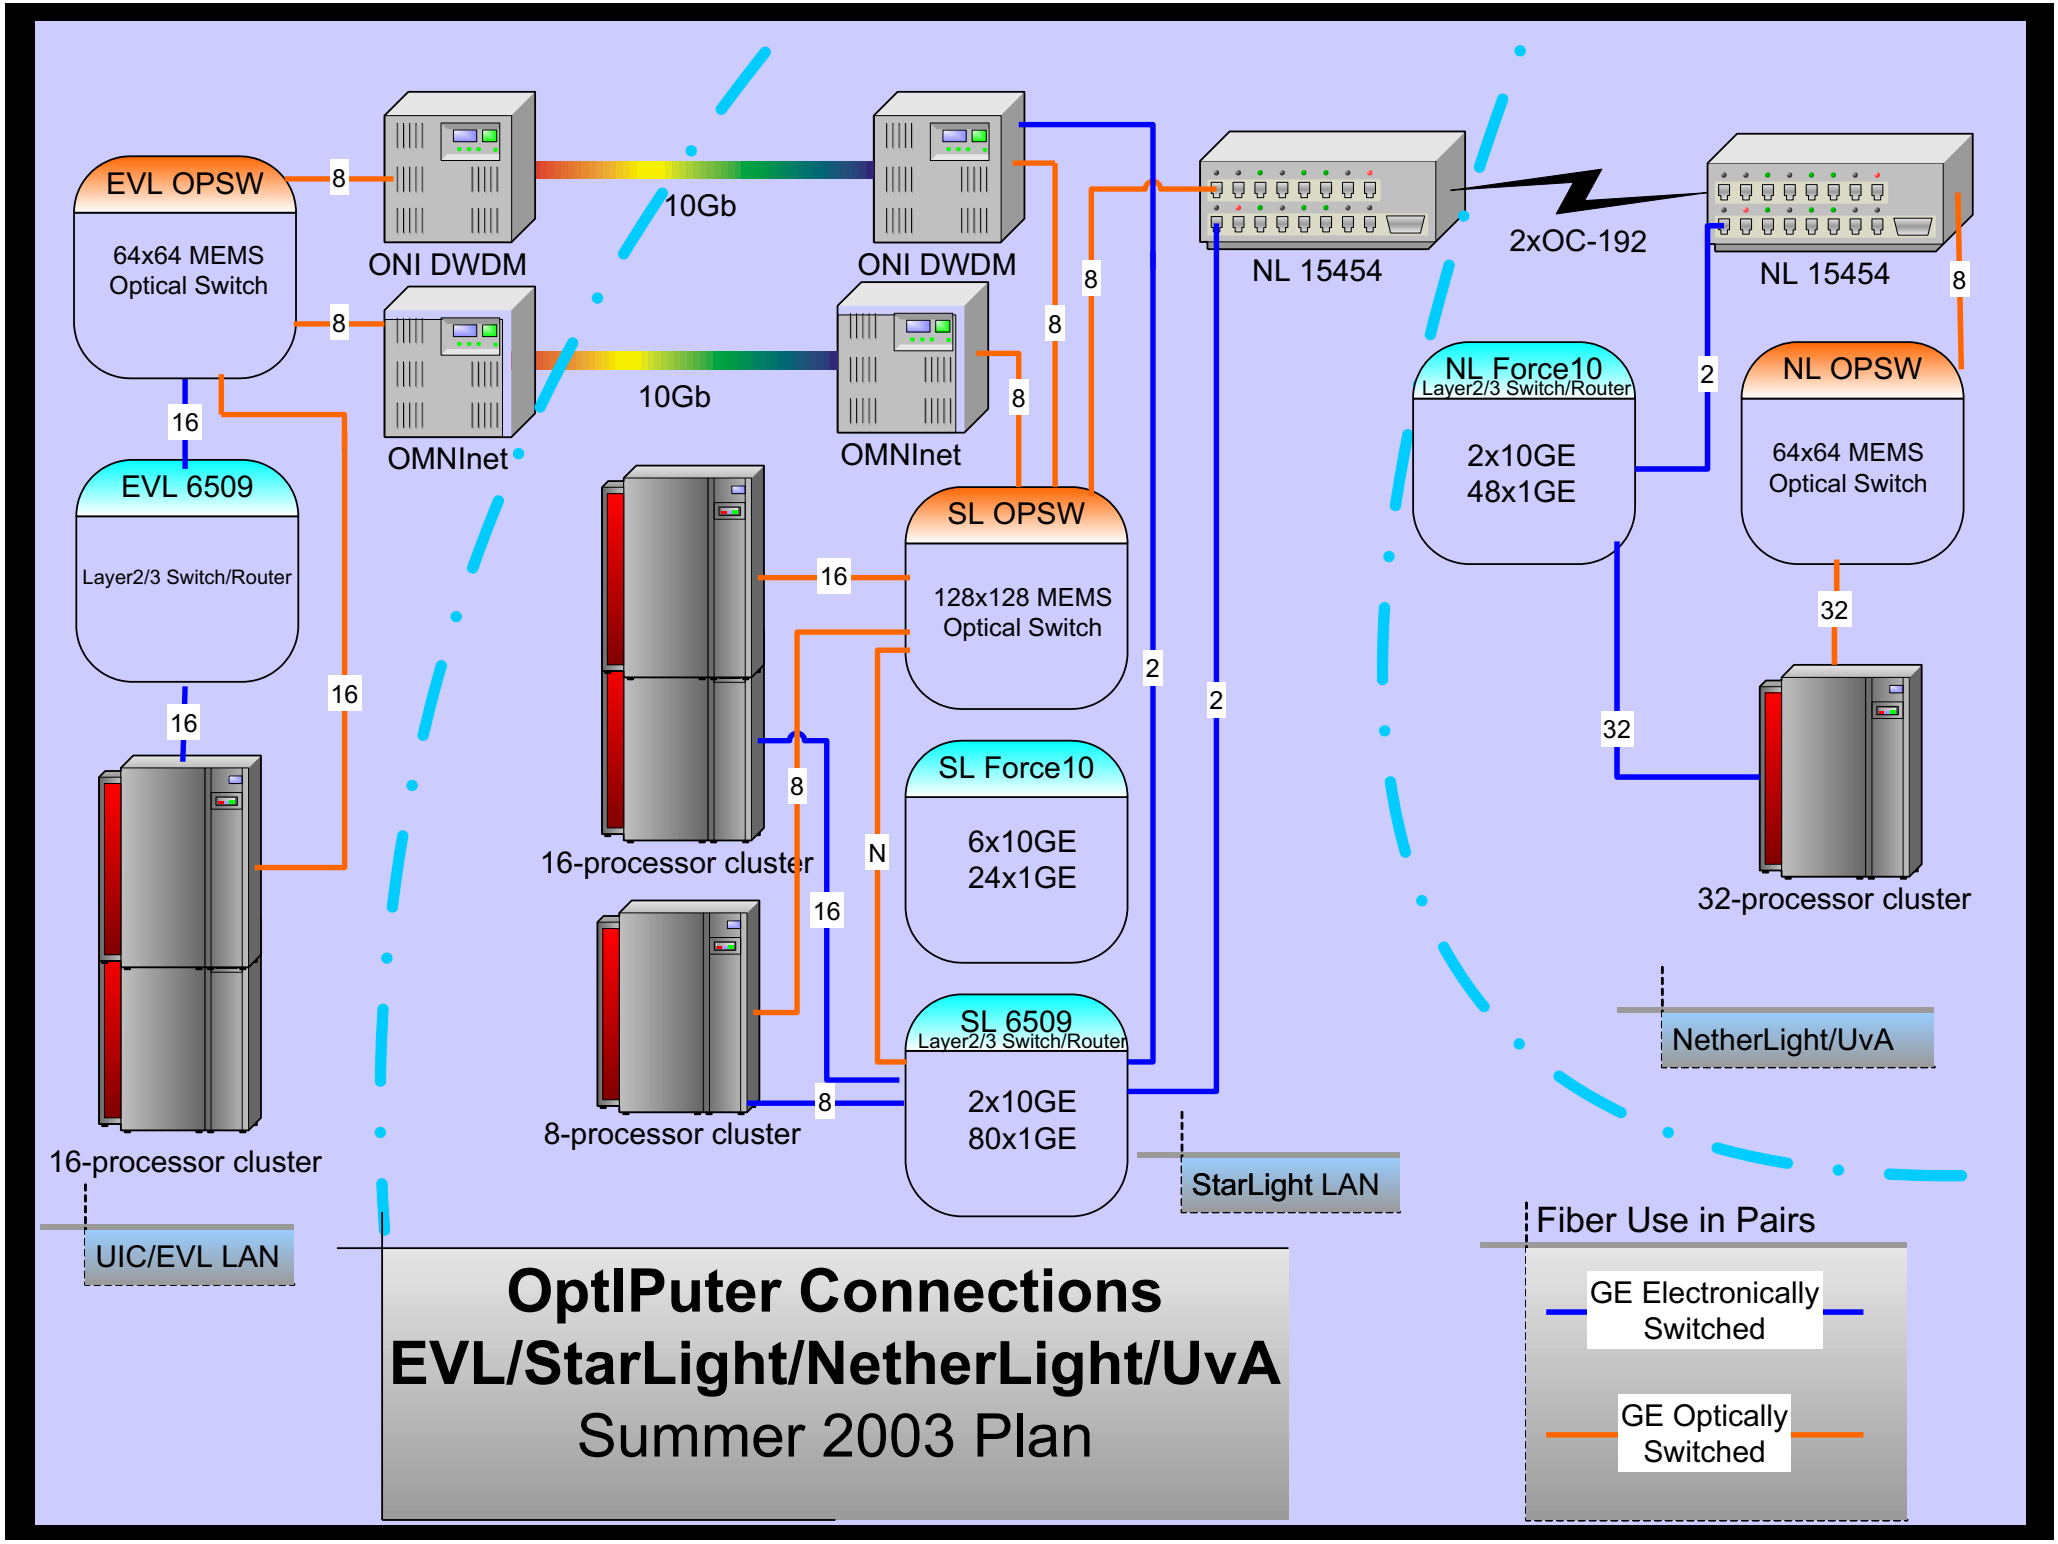
EVL OPSW
64x64 MEMS
Optical Switch
8
8
16
16
16
EVL 6509
Layer2/3 Switch/Router
16
16-processor cluster
UIC/EVL LAN
ONI DWDM
10Gb
ONI DWDM
OMNInet
10Gb
OMNInet
8
8
8
16
16-processor cluster
SL OPSW
128x128 MEMS
Optical Switch
2
2
8
N
16
8
8-processor cluster
SL Force10
6x10GE
24x1GE
SL 6509
Layer2/3 Switch/Router
2x10GE
80x1GE
StarLight LAN
NL 15454
2xOC-192
NL 15454
8
NL Force10
Layer2/3 Switch/Router
2x10GE
48x1GE
2
NL OPSW
64x64 MEMS
Optical Switch
32
32
32-processor cluster
NetherLight/UvA
Fiber Use in Pairs
GE Electronically
Switched
GE Optically
Switched
OptIPuter Connections
EVL/StarLight/NetherLight/UvA
Summer 2003 Plan

Why Optical Switching?

- No need to look at every packet when transferring a terabyte of information
  - 1% the cost of routing
  - 10% the cost of switching
  - 64x64 10Gb:
    - $100,000 O-O-O switched
    - $1,000,000 O-E-O switched
    - $10,000,000 O-E-O Routed
- Spend the savings on computing and collaboration systems instead!
- Replaces patch panels; allows rapid reconfiguration of 1 and 10Gb experiments

STARLIGHT℠

Thanks to…

- **StarLight planning, research, collaborations, and outreach efforts are made possible, in major part, by funding from:**
  - National Science Foundation (NSF) awards ANI-9980480, ANI-9730202, EIA-9802090, EIA-9871058, ANI-0225642, and EIA-0115809
  - NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to NCSA
  - State of Illinois I-WIRE Program, and major UIC cost sharing
  - Northwestern University for providing space, $$, engineering and management
- **NSF/CISE/ANIR and DoE/Argonne National Laboratory for StarLight and I-WIRE network engineering and planning leadership**
- **NSF/CISE/ACIR and NCSA/SDSC for DTF/TeraGrid/ETF opportunities**
- **UCAID/Abilene for Internet2 and ITN transit; IU for the GlobalNOC**
- **Bill St. Arnaud of CANARIE, Kees Neggers of SURFnet, Olivier Martin of CERN, Michael McRobbie of IU, and Harvey Newman of CalTech for networking leadership**

STARLIGHT℠

StarLight:
"Bring Us Your Lambdas"

STARLIGHT℠