International Perspective:

Facilities supporting research and development with Lightpaths

www.science.uva.nl/~delaat

Cees de Laat

SURFnet

EU

University of Amsterdam

SARA

NIKHEF

NCF

International Perspective:
Facilities supporting research and development with Lightpaths

www.science.uva.nl/~delaat

Cees de Laat

SURFnet

EU

University of Amsterdam

SARA

NIKHEF

NCF

Contents of this talk

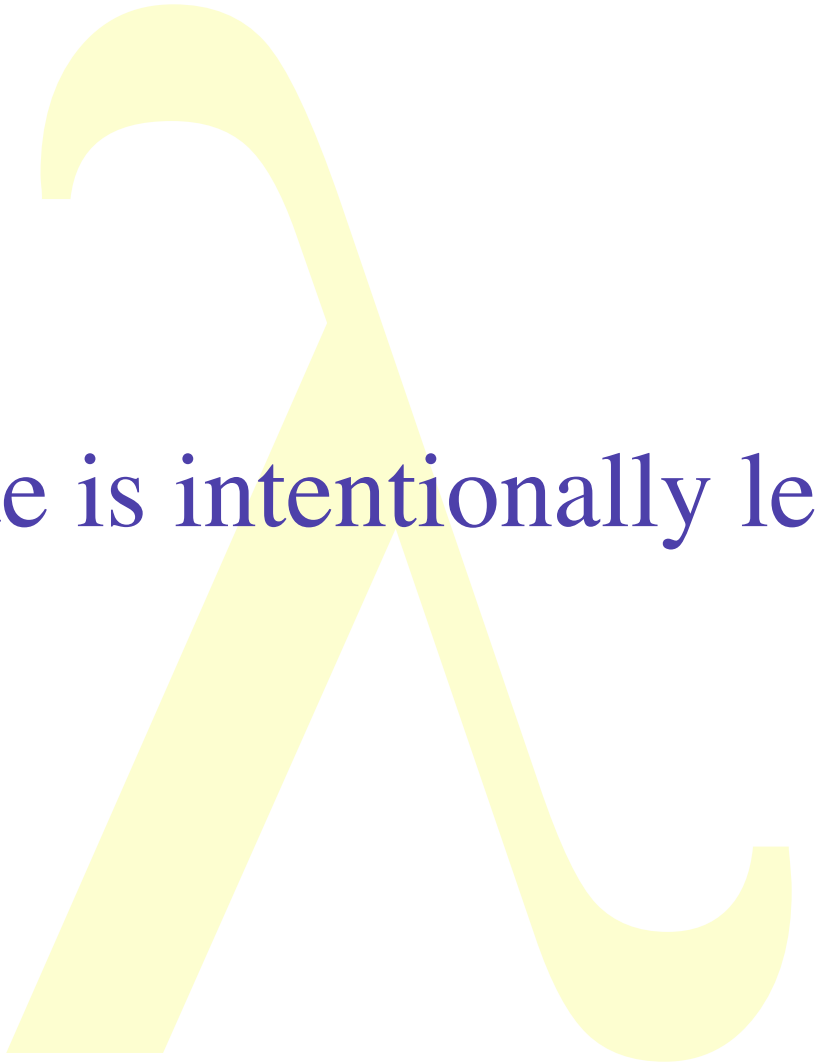

This slide is intentionally left blank

eVLBI

VLBI

ger term VLBI is easily capable of generating many Gb of data per
The sensitivity of the VLBI array scales v
(data-rate) and there is a strong push to
. Rates of 8Gb/s or more are entirely feasibl
ider development. It is expected that paral
orrelator will remain the most efficient approa
s distributed processing may have an appli
lti-gigabit data streams will aggregate into la
or and the capacity of the final link to the da
tor.

Westerbork Synthesis Radio Telescope - Netherlands

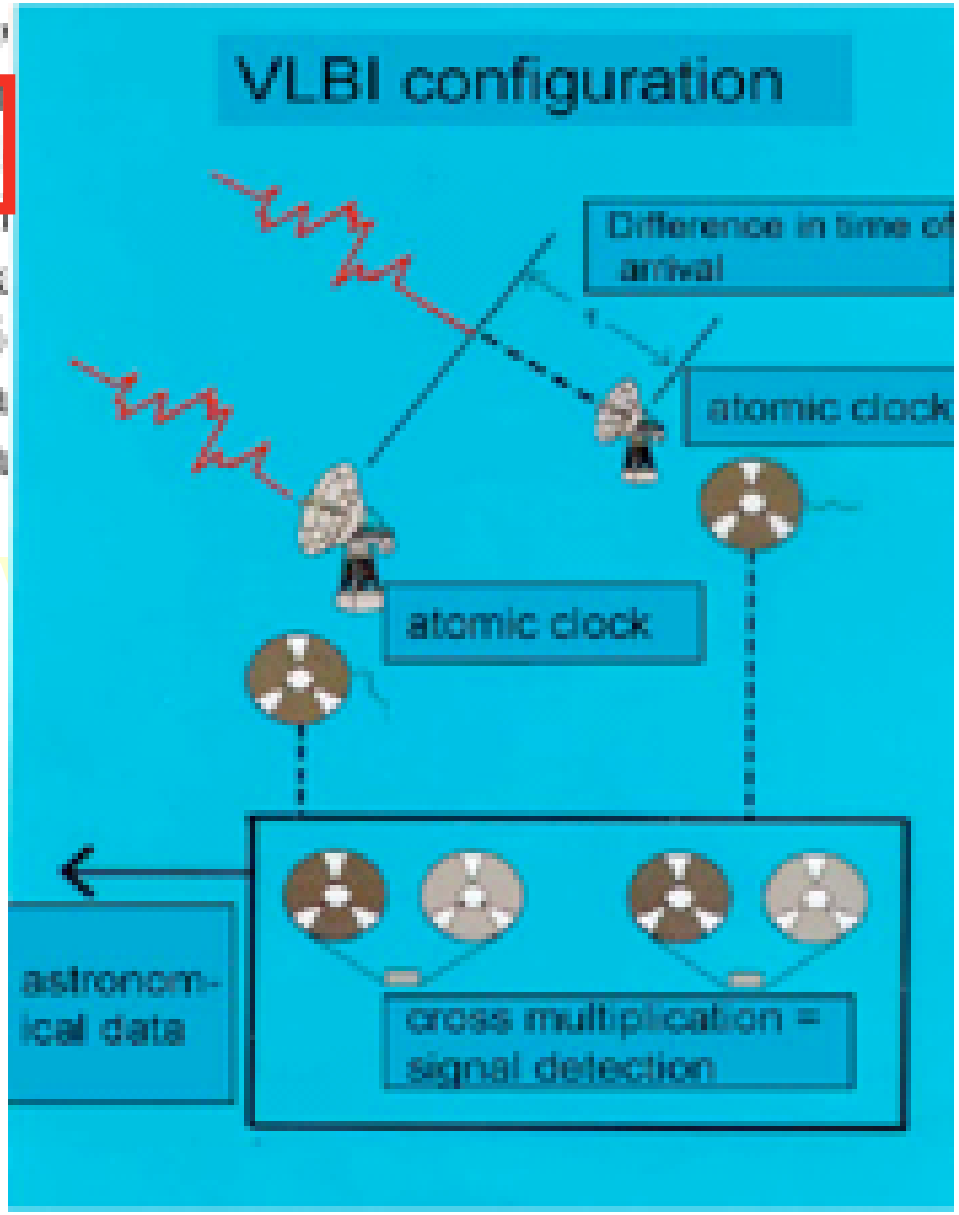

iGrid 2002

September 24-26, 2002, Amsterdam, The Netherlands

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

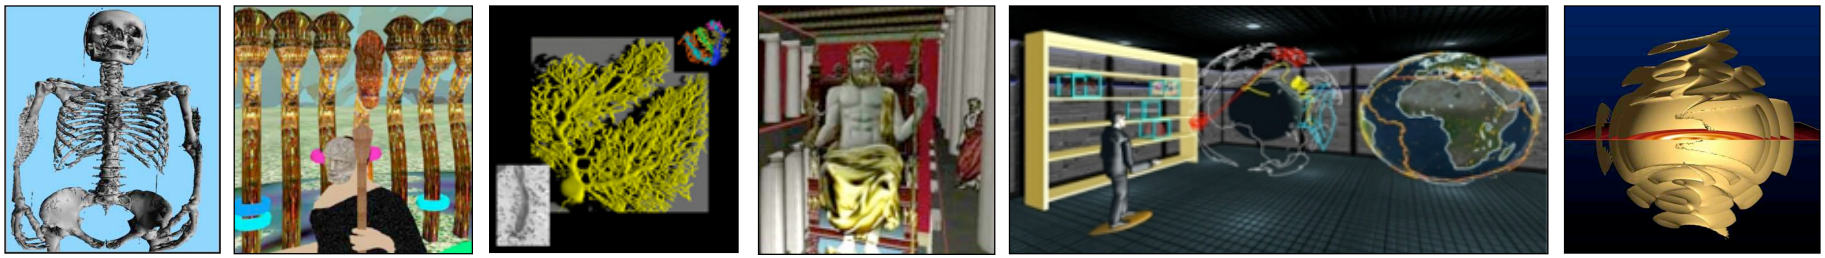

- Grid technologies demonstrated: Major emphasis on grid middleware, data management grids, data replication grids, visualization grids, data/visualization grids, computational grids, access grids, grid portals
- 25Gb transatlantic bandwidth (100Mb/attendee, 250x iGrid2000!)

www.igrid2002.org

iGrid 2002

Sept 24-26, 2002, Amsterdam, The Netherlands

Conference issue
FGCS
Volume 19 (2003)
Number 6 august
22 refereed papers!

THESE
ARE
THE
APPLICATIONS!

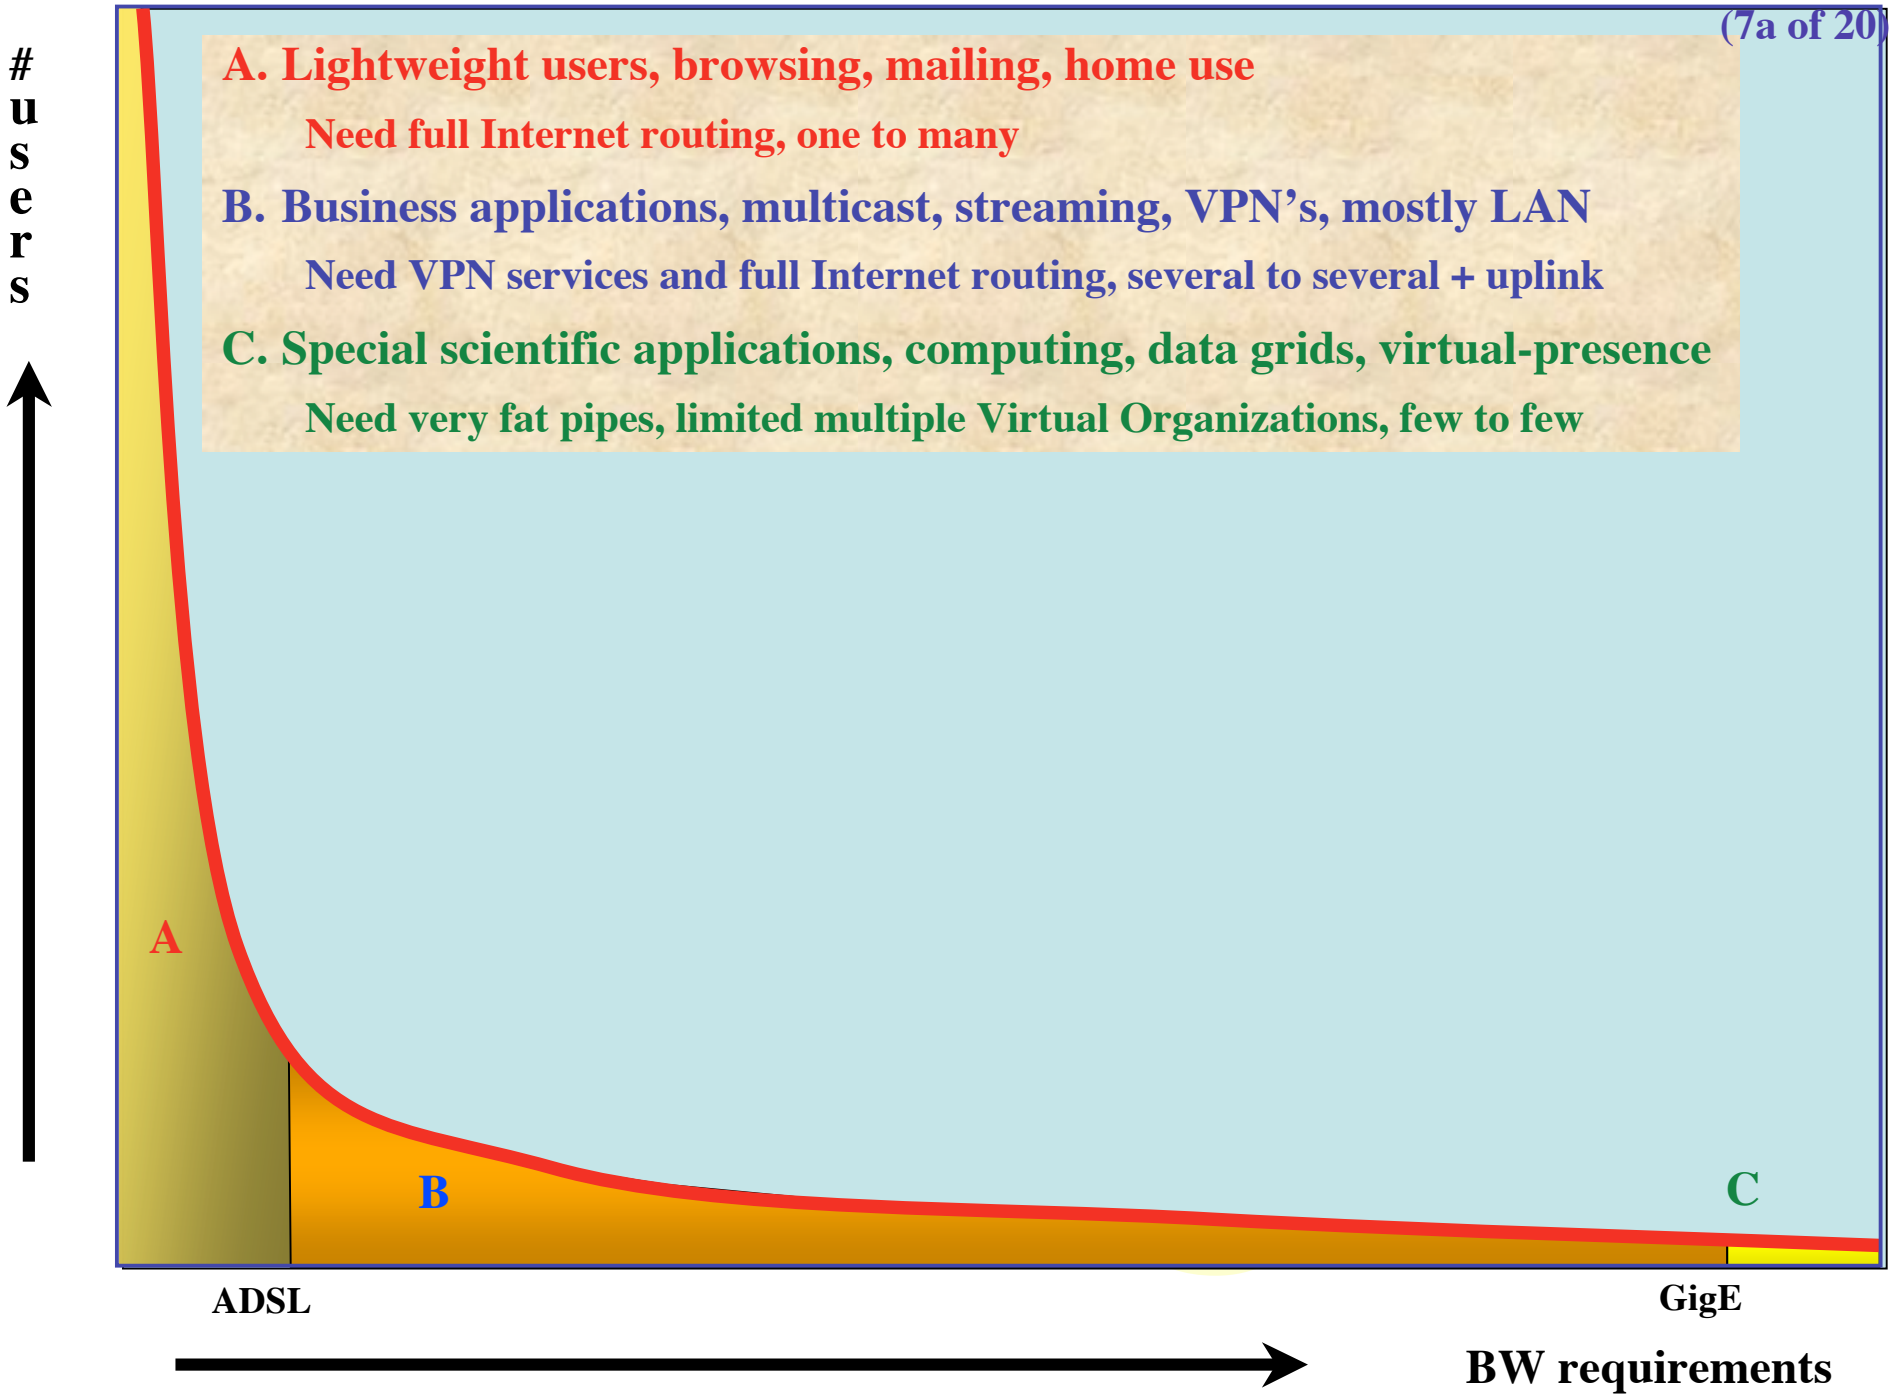
A. Lightweight users, browsing, mailing, home use
Need full Internet routing, one to many
B. Business applications, multicast, streaming, VPN's, mostly LAN
Need VPN services and full Internet routing, several to several + uplink
C. Special scientific applications, computing, data grids, virtual-presence
Need very fat pipes, limited multiple Virtual Organizations, few to few
users
A
B
C
ADSL
GigE
BW requirements

The Dutch Situation

- **Estimate A**
  - **17 M people, 6.4 M households, 25 % penetration of 0.5 Mb/s ADSL, 40 times under-provisioning ==> 20 Gb/s**
- **Estimate B**
  - **SURFnet has 10 Gb/s to about 12 institutes and 0.1 to 1 Gb/s to 180 customers, estimate same for industry (overestimation) ==> 20-40 Gb/s**
- **Estimate C**
  - **Leading HEF and ASTRO + rest ==> 80-120 Gb/s**

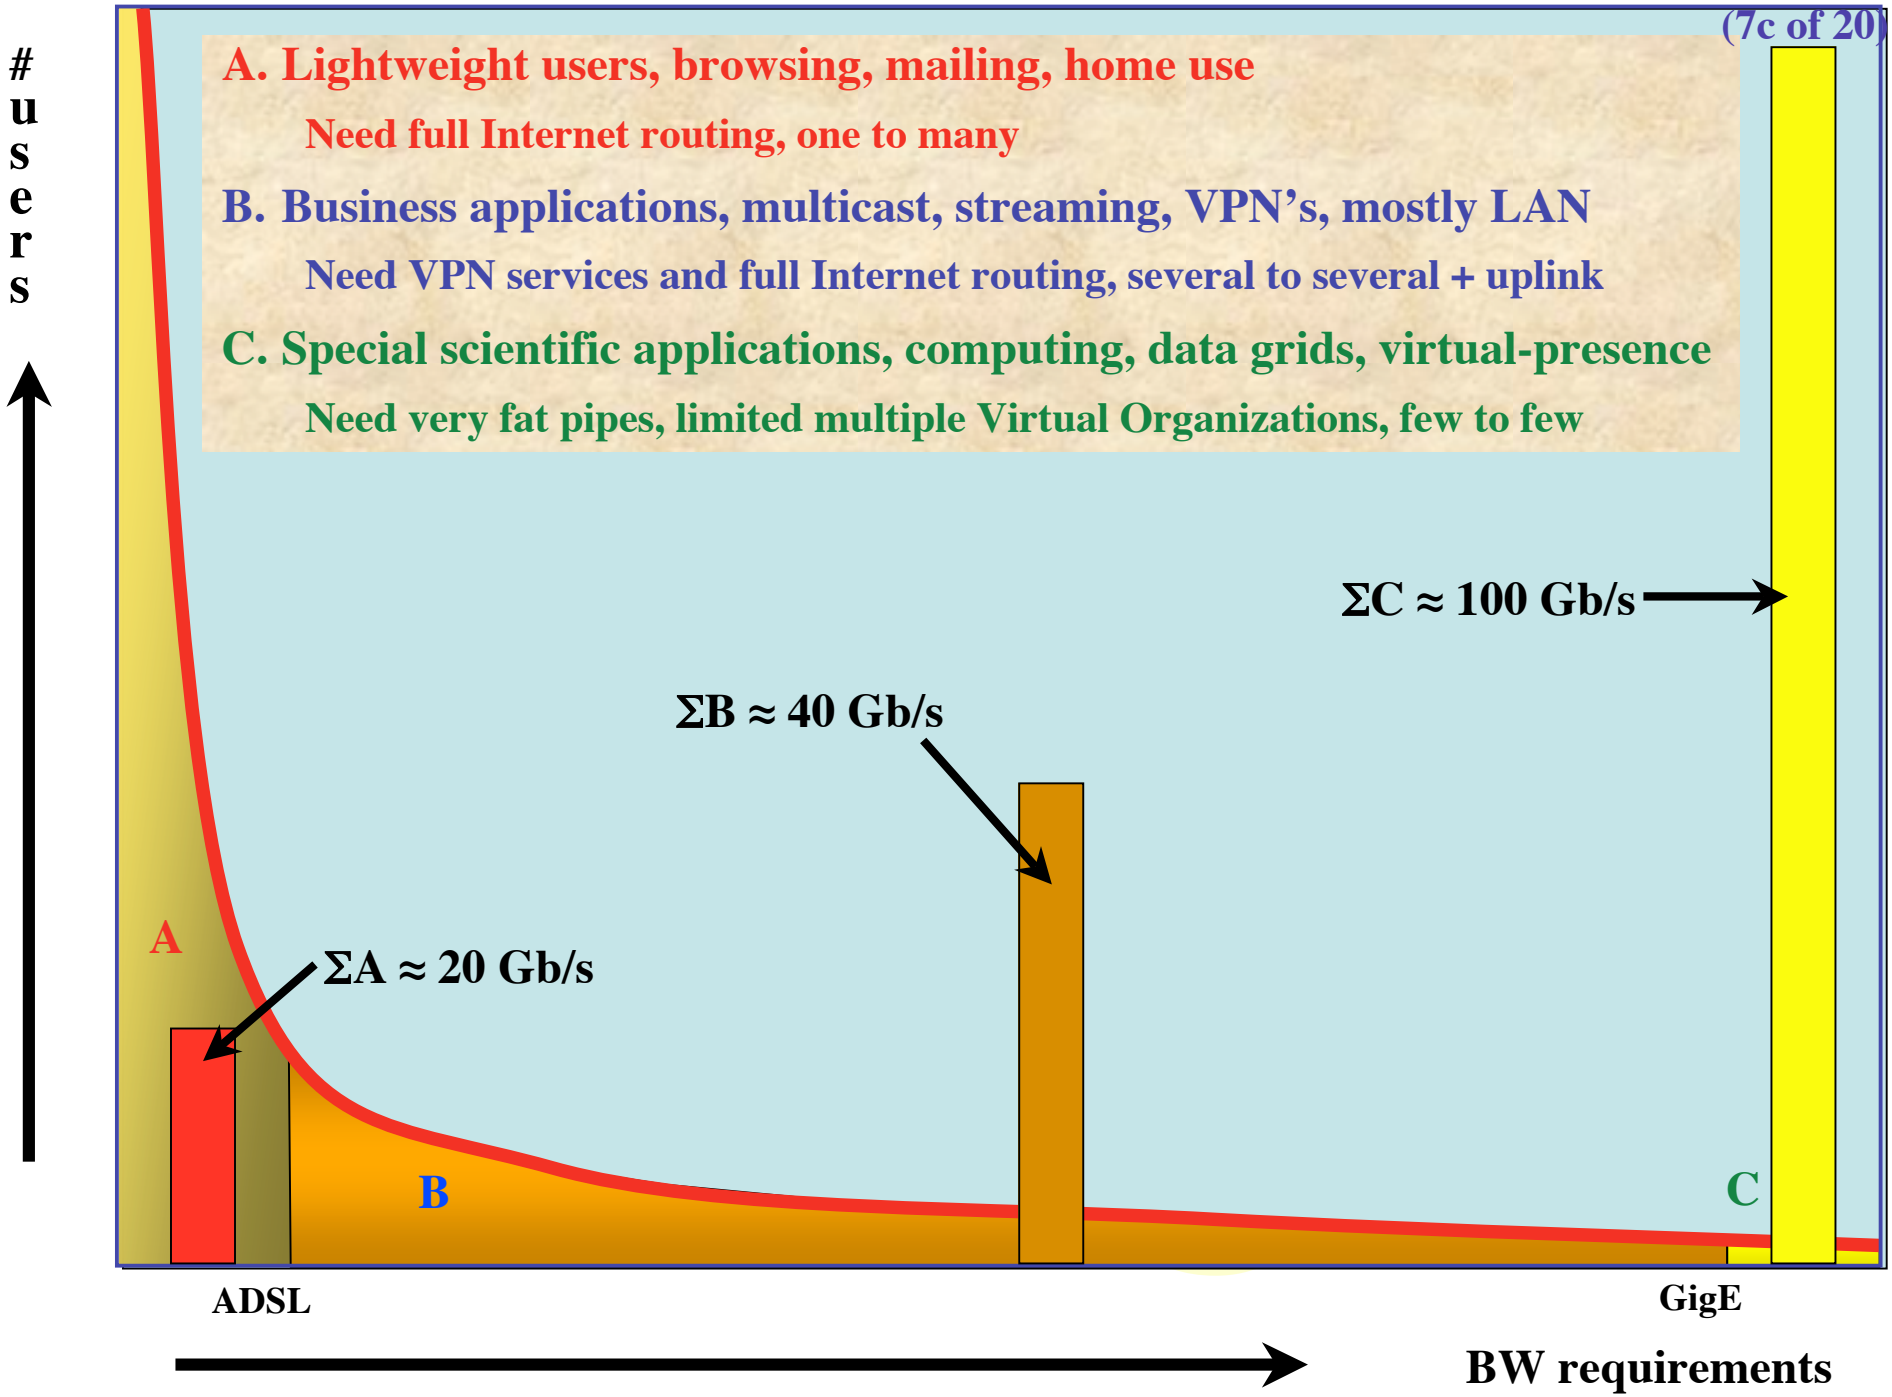
#users
A. Lightweight users, browsing, mailing, home use
Need full Internet routing, one to many
B. Business applications, multicast, streaming, VPN's, mostly LAN
Need VPN services and full Internet routing, several to several + uplink
C. Special scientific applications, computing, data grids, virtual-presence
Need very fat pipes, limited multiple Virtual Organizations, few to few
ΣC ≈ 100 Gb/s
ΣB ≈ 40 Gb/s
ΣA ≈ 20 Gb/s
A
B
C
ADSL
GigE
BW requirements

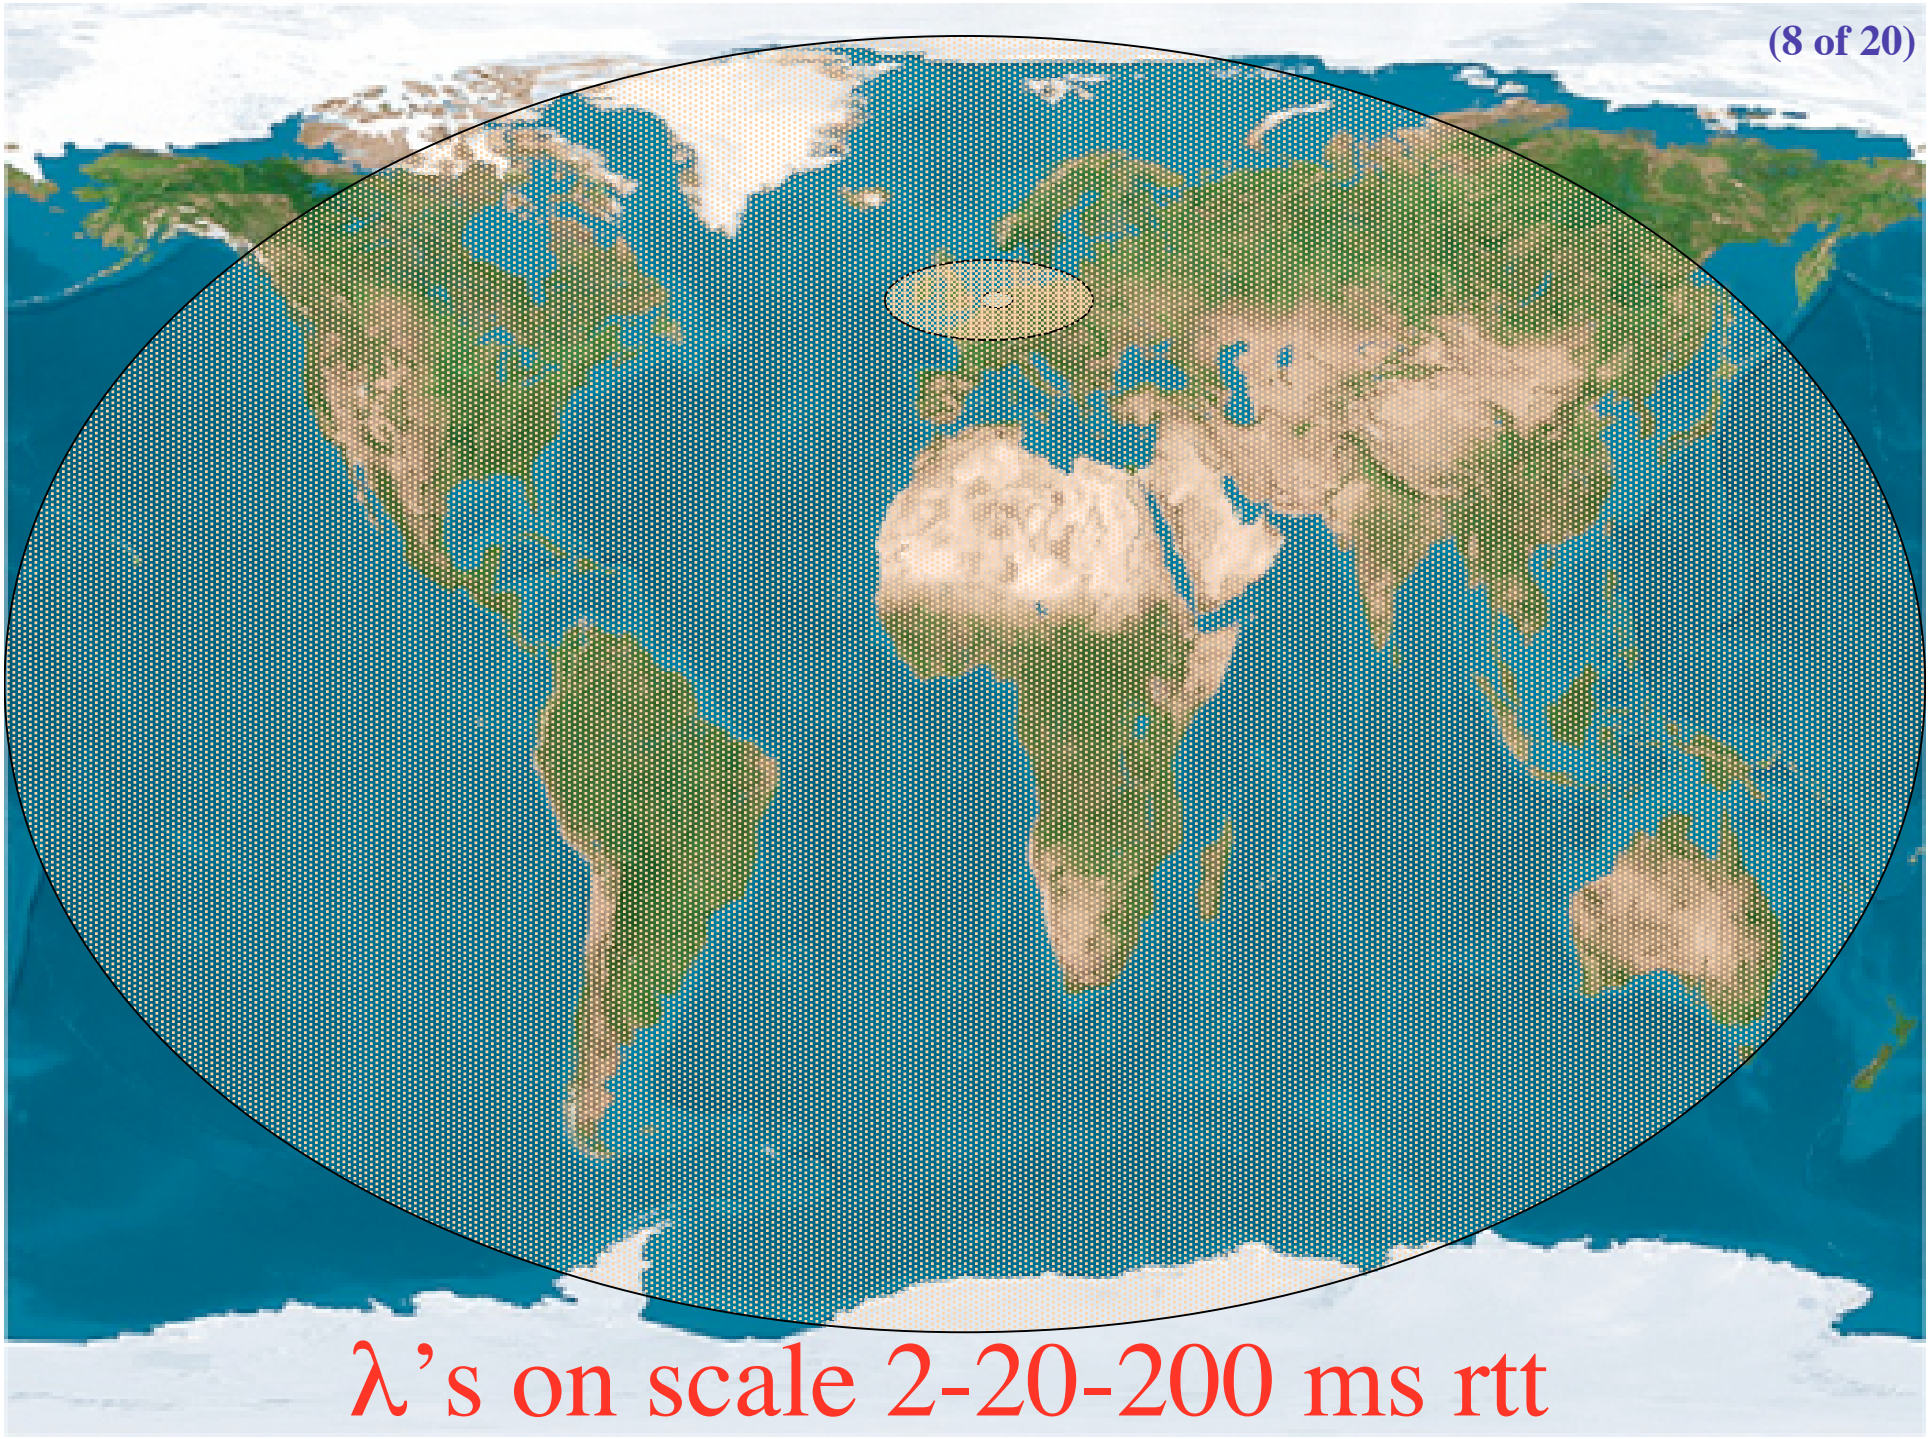

λ's on scale 2-20-200 ms rtt

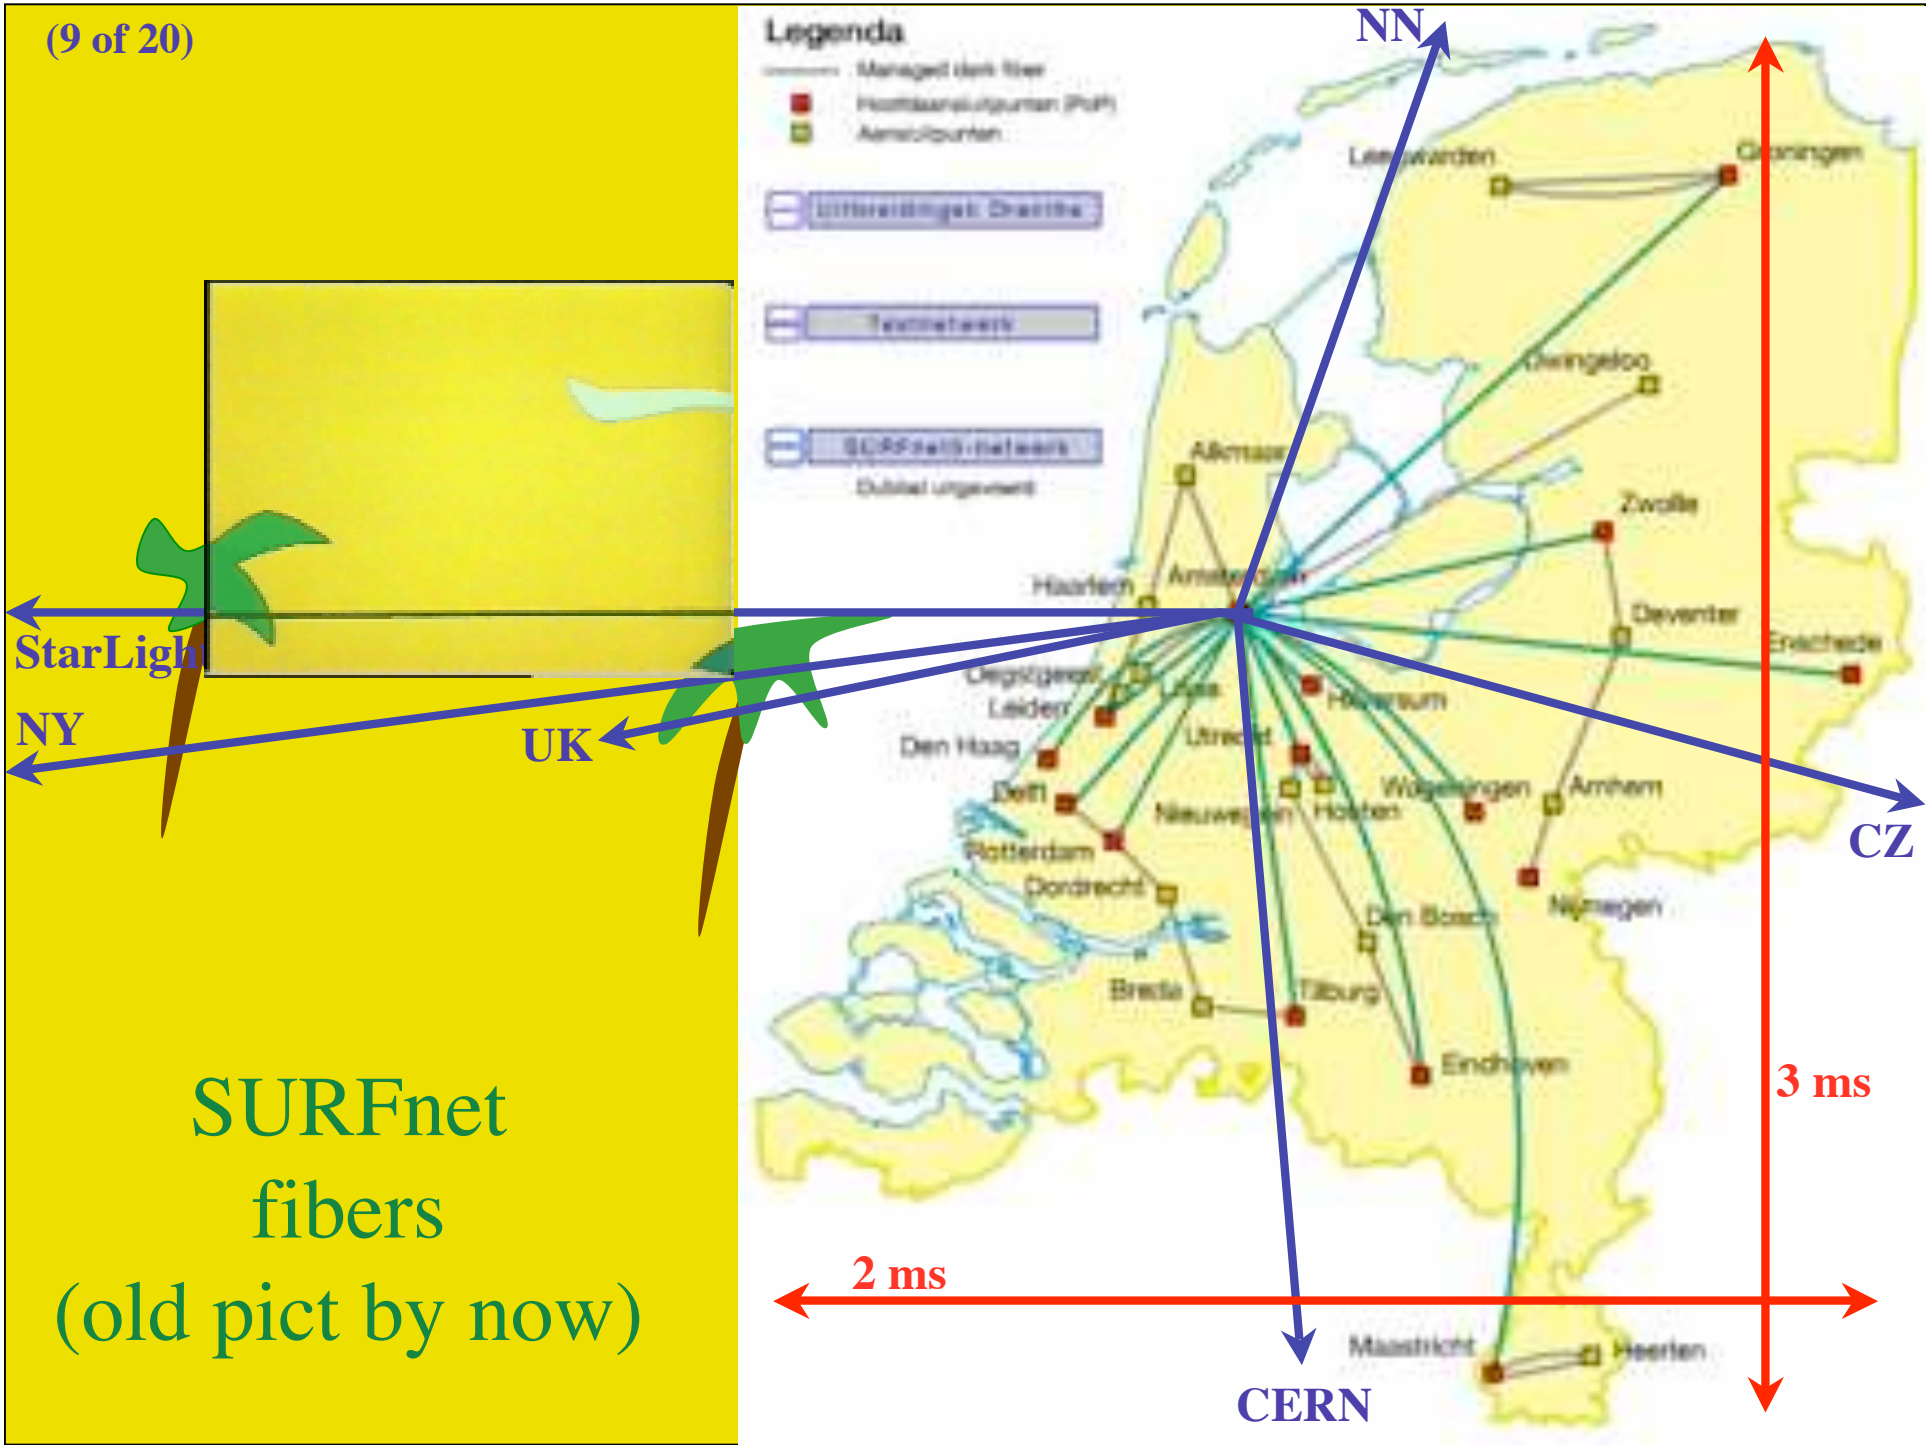

SURFnet
fibers
(old pict by now)

The only formula's

$$\#\lambda(rtt,t) \approx \frac{200 * e^{(t-2002)}}{rtt}$$

Now, having been a High Energy Physicist we set

c = 1

e = 1

$\bar{h}$ = 1

and the formula reduces to: $\#\lambda(rtt,t) \approx \frac{200 * e^{(t-2002)}}{rtt}$

So what are the facts

- Costs of fat pipes (fibers) are one/third of cost of equipment to light them up
  - Is what Lambda salesmen tell me
- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
  - 100 Byte packet @ 40 Gb/s -> 20 ns to look up in 140 kEntries routing table (light speed from me to you!)
- Big sciences need fat pipes
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way (A -> L3 , B -> L2 , C -> L1)
- Tested 10 gbps Ethernet WANPHY Amsterdam-CERN and Amsterdam-StarLight
  - http://www.surfnet.nl/en/publications/pressreleases/021003.html

Services

SCALE / CLASS	2 Metro	20 National/ regional	200 World
A	Switching/ routing	Routing	ROUTER$
B	Switches + E-WANPHY VPN's,	Switches + E-WANPHY (G)MPLS	ROUTER$
C	dark fiber Optical switching	Lambda switching	Sub-lambdas, ethernet-sdh

Lightwave Networking

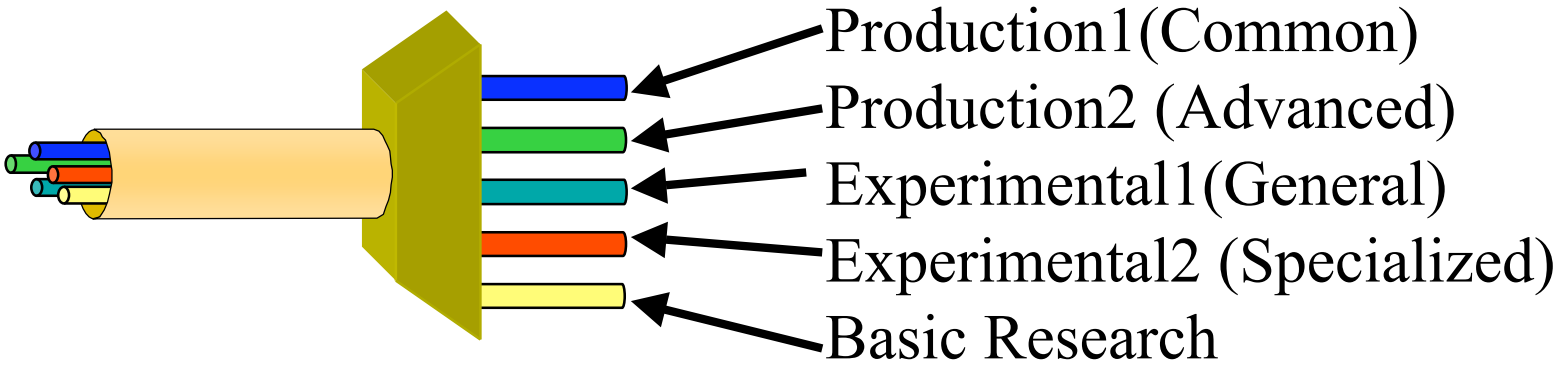

Eg, Experimental2 Could Be Provisioned From Remote Instruments In Other Countries Directly to University Labs –
Finally, Separate Networks On the Same Infrastructure!

Source: Joe Mambretti

Architectures - L1 - L3

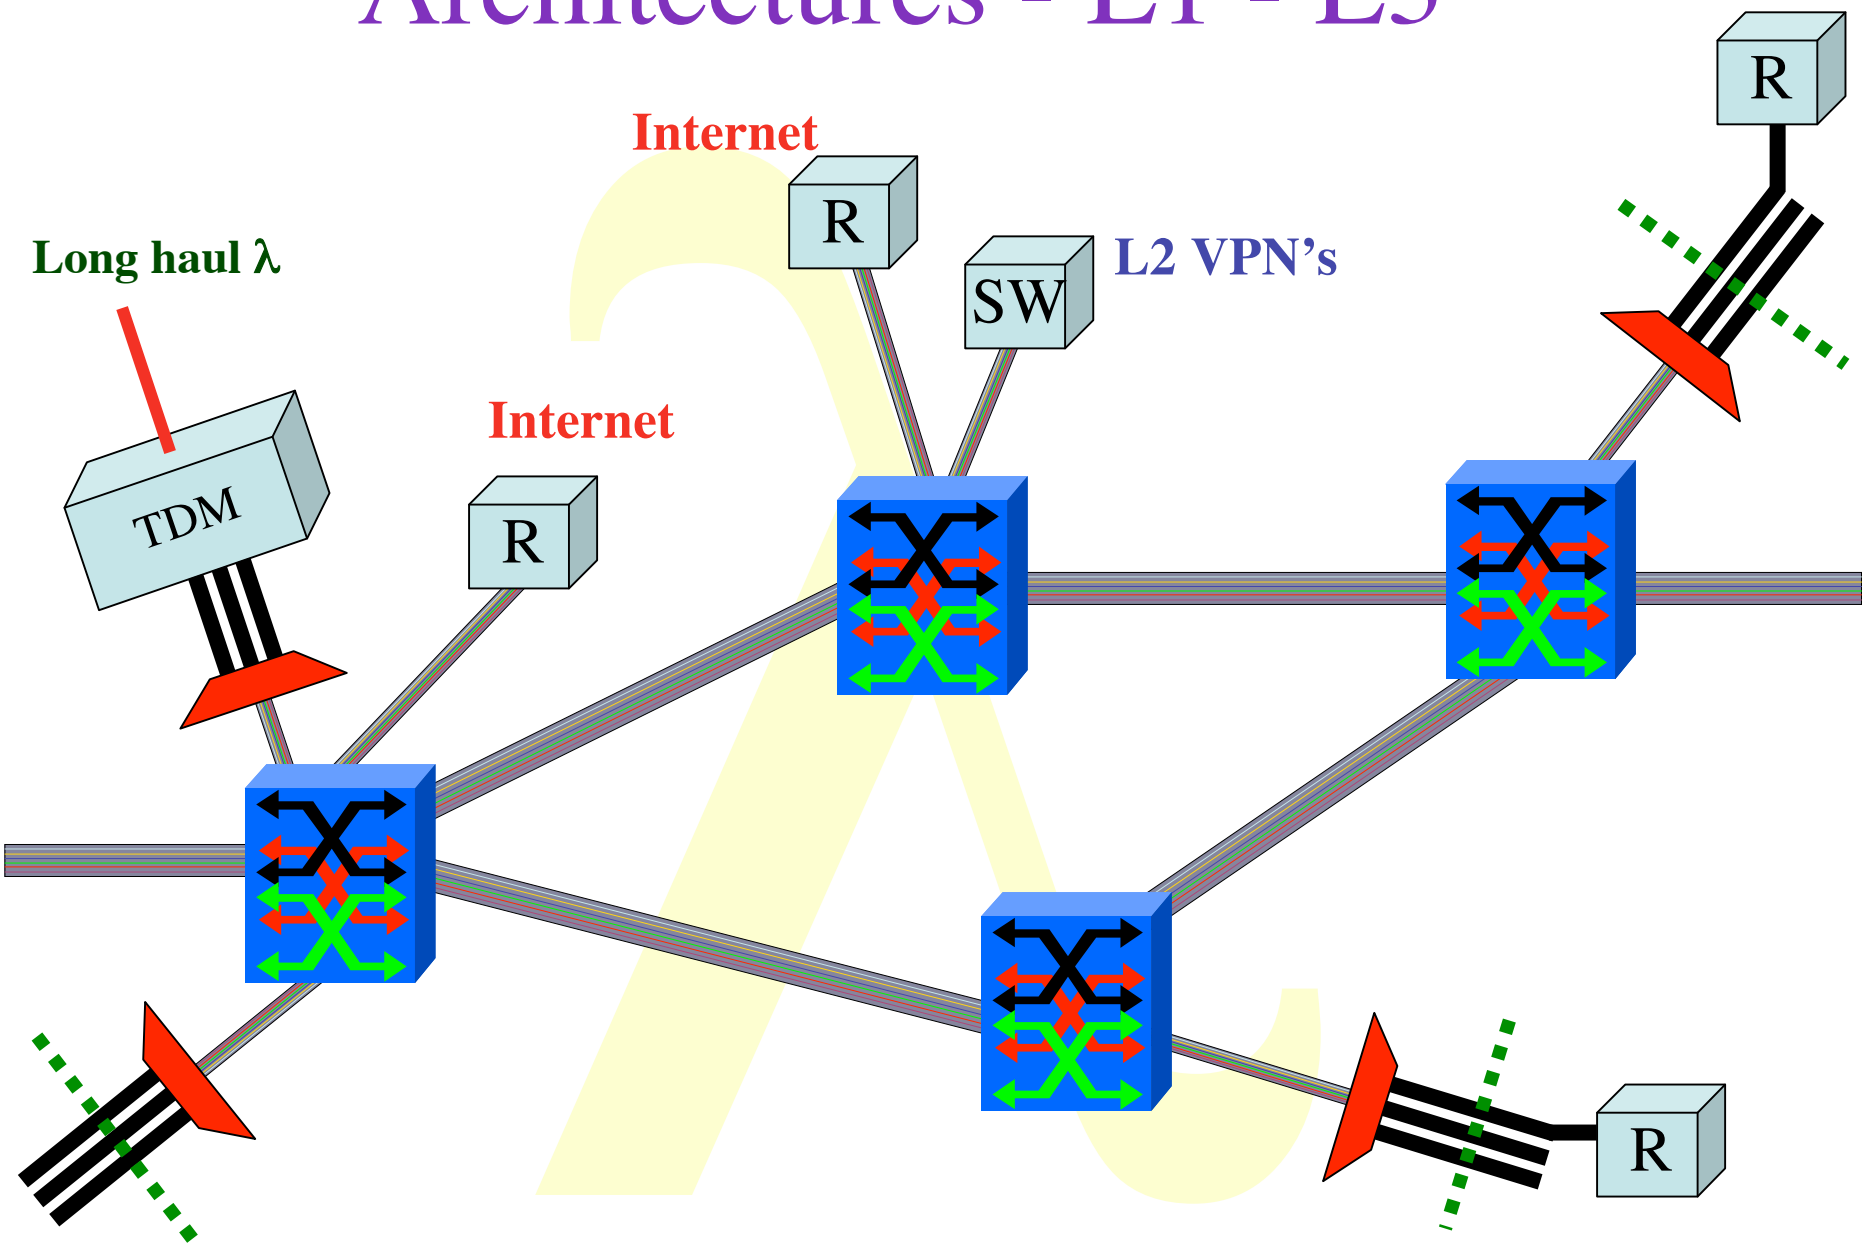

MREN
Metropolitan Research & Education Network

OMNInet Testbed Experiments
(MREN Used as Out-of-Band Control Channel)

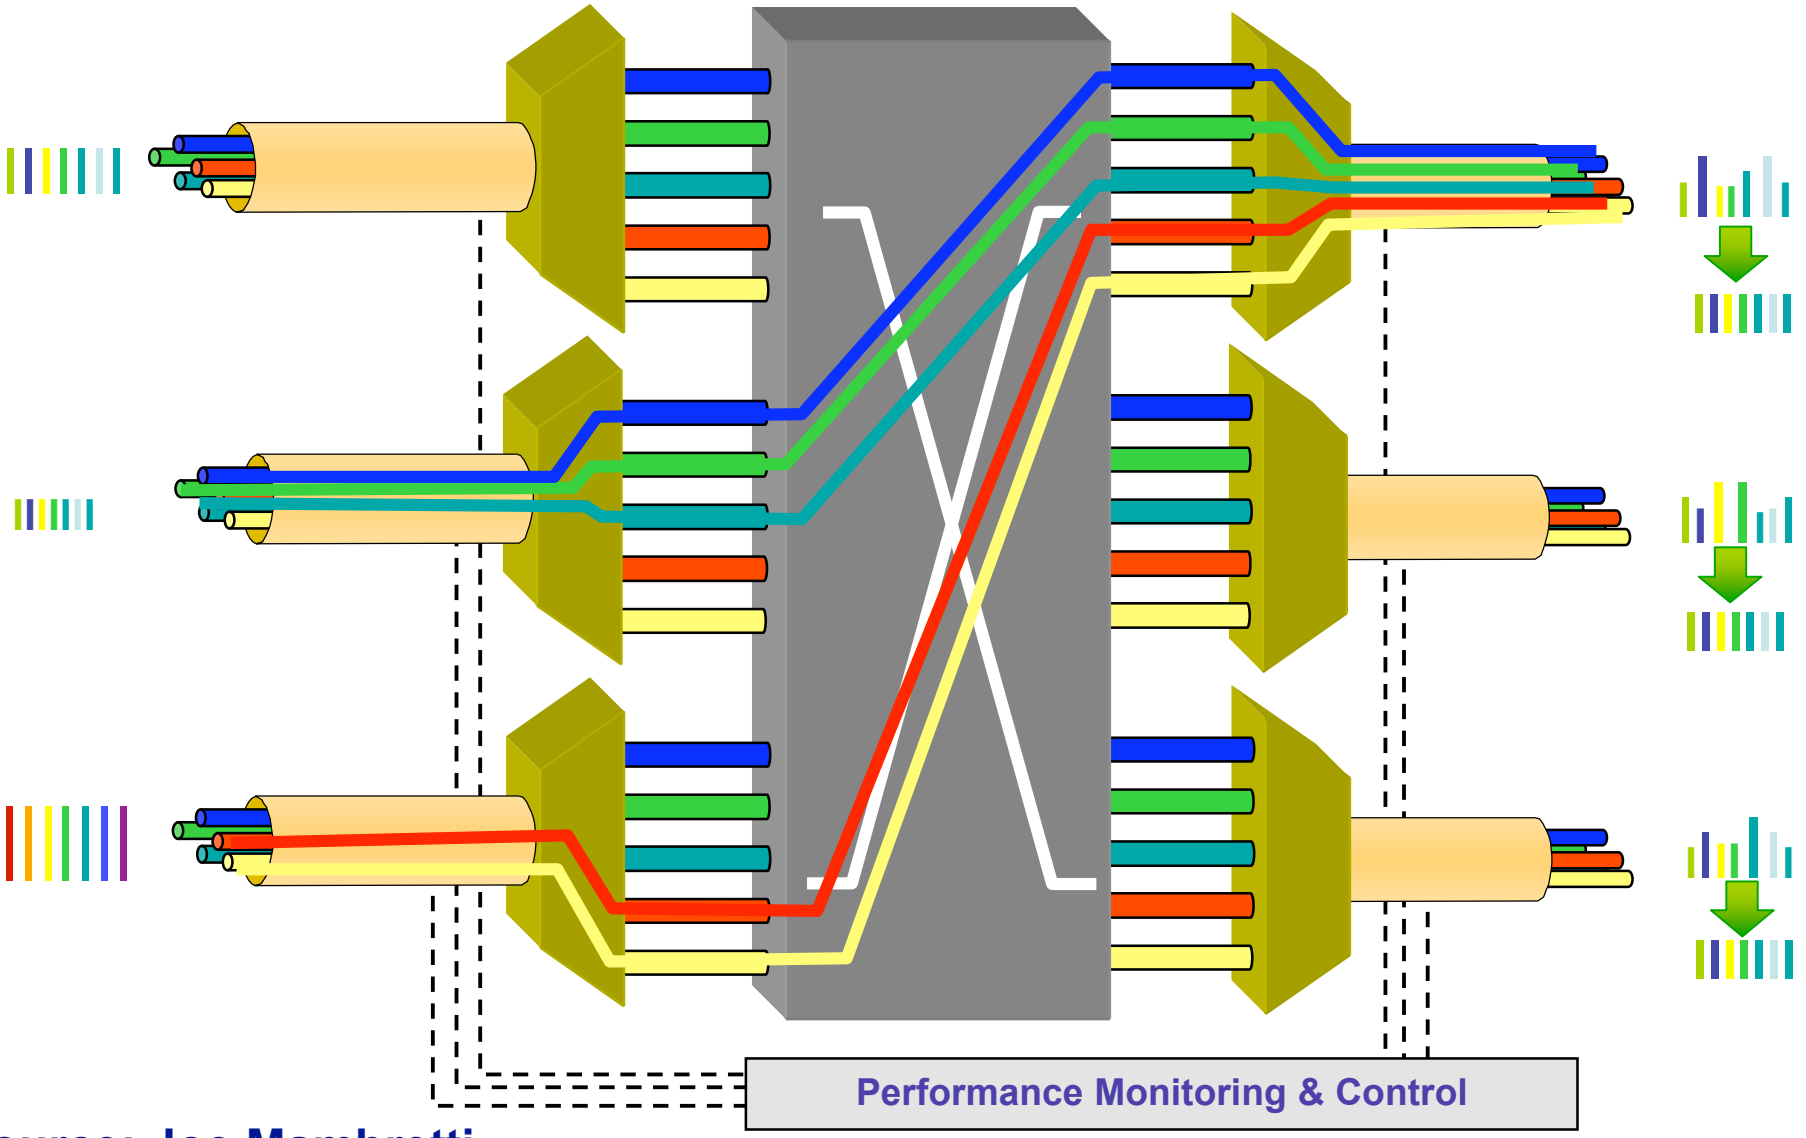

Source: Joe Mambretti

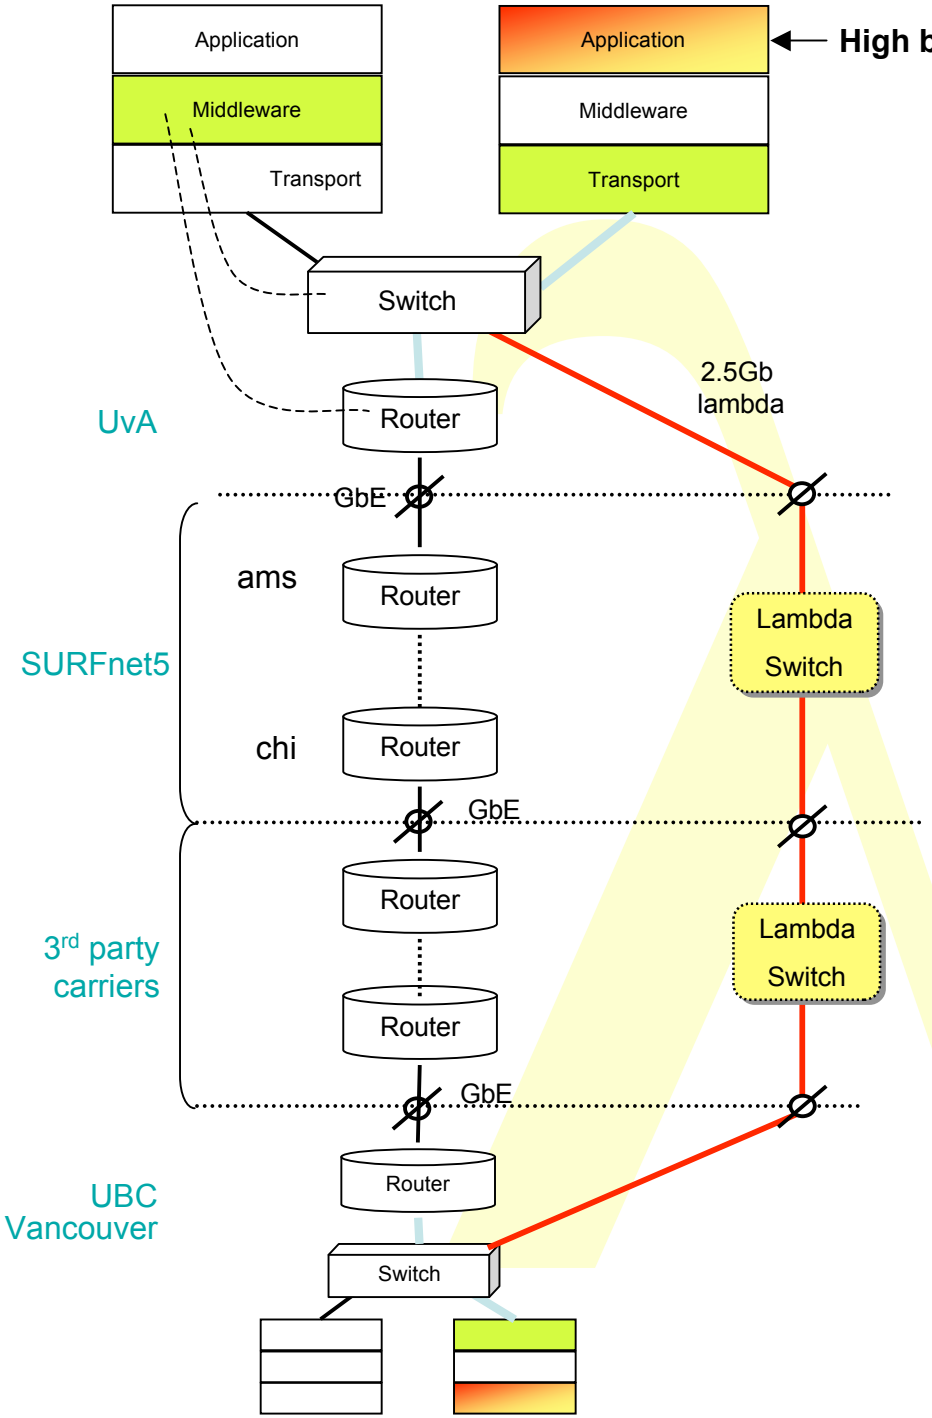

- lambda for high bandwidth applications
  - Bypass of production network
  - Middleware may request (optical) pipe
- RATIONALE:
  - Lower the cost of transport per packet
  - Use Internet as controlplane!

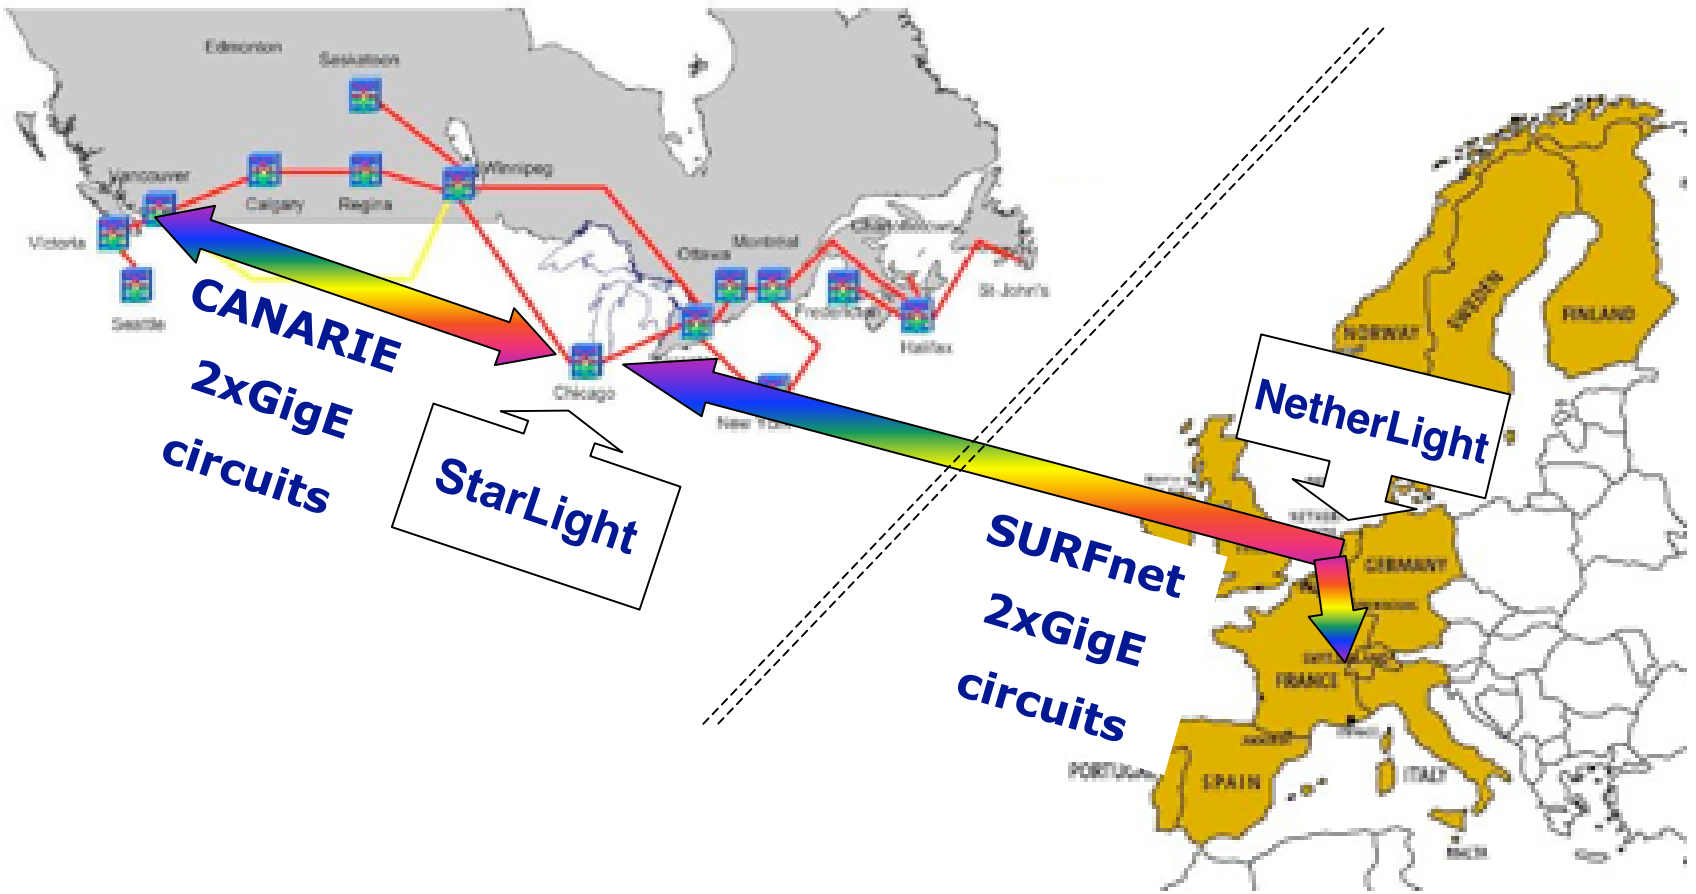

Source: Tom DeFanti

How low can you go?

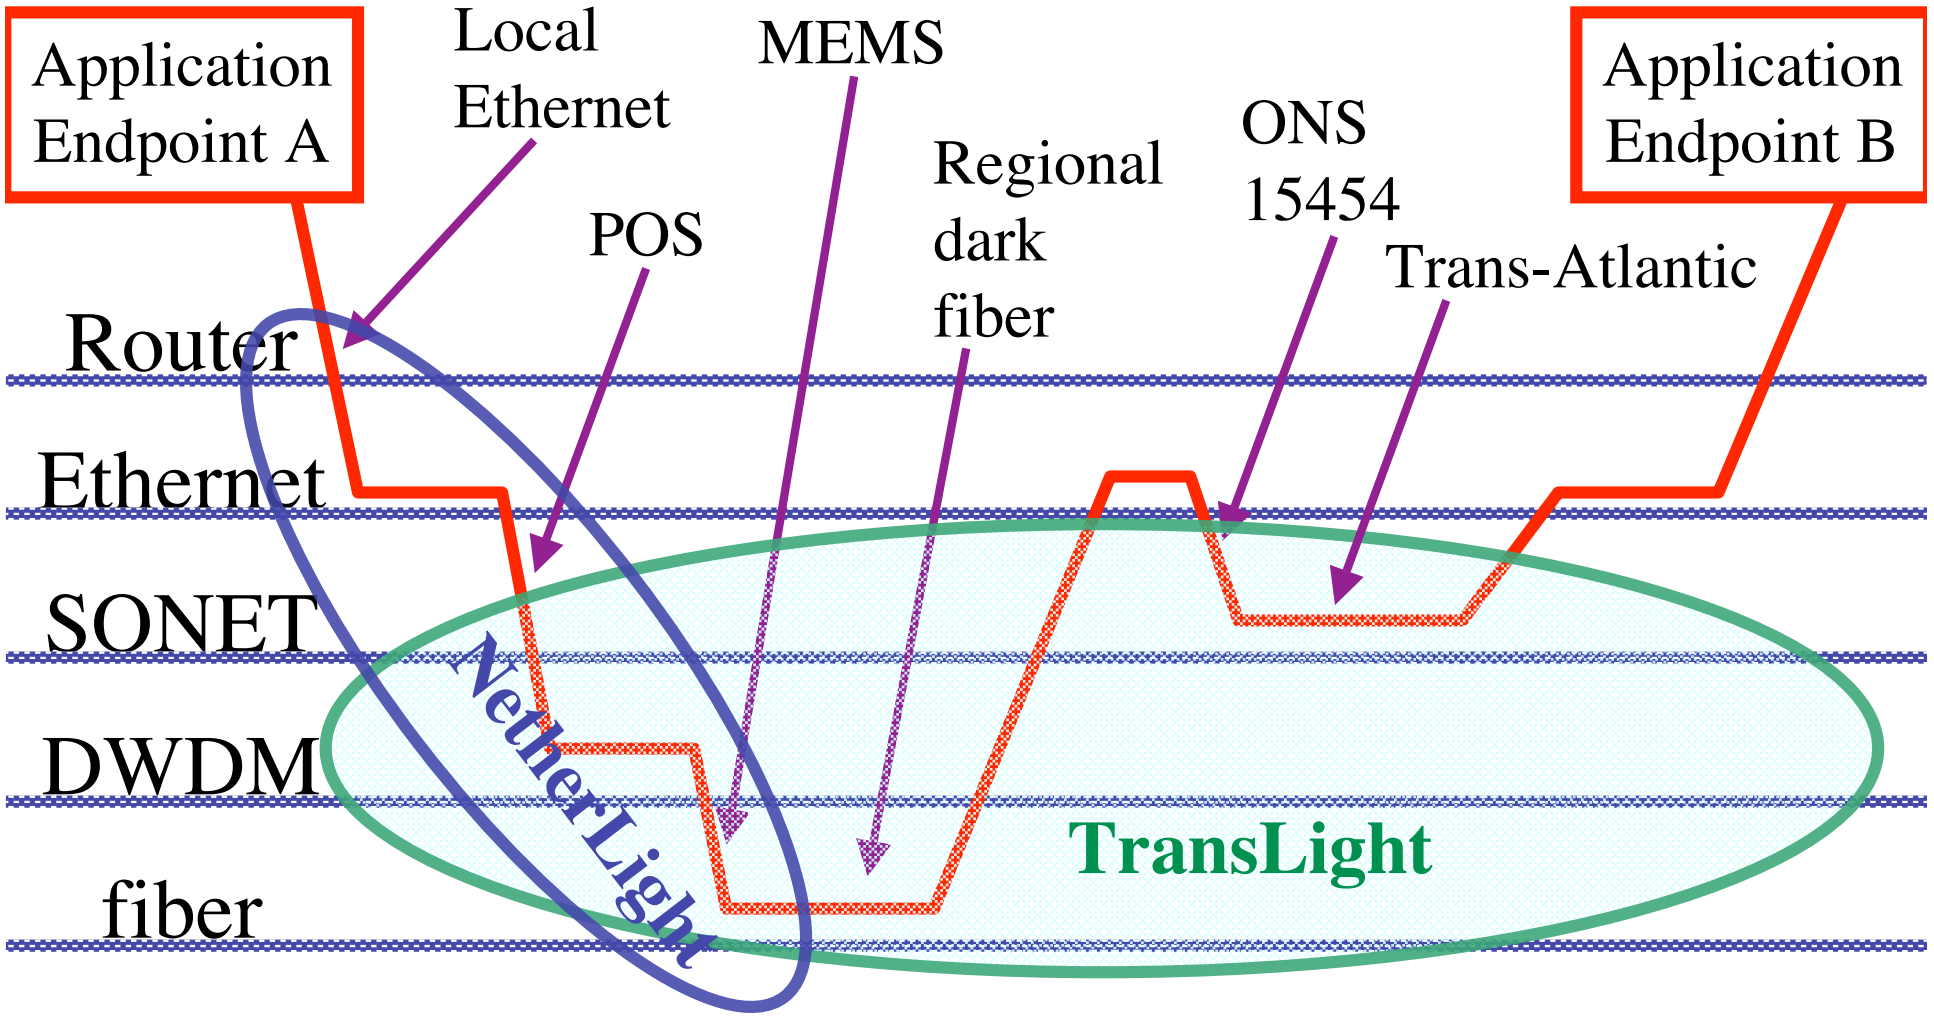

TransLight Lambdas

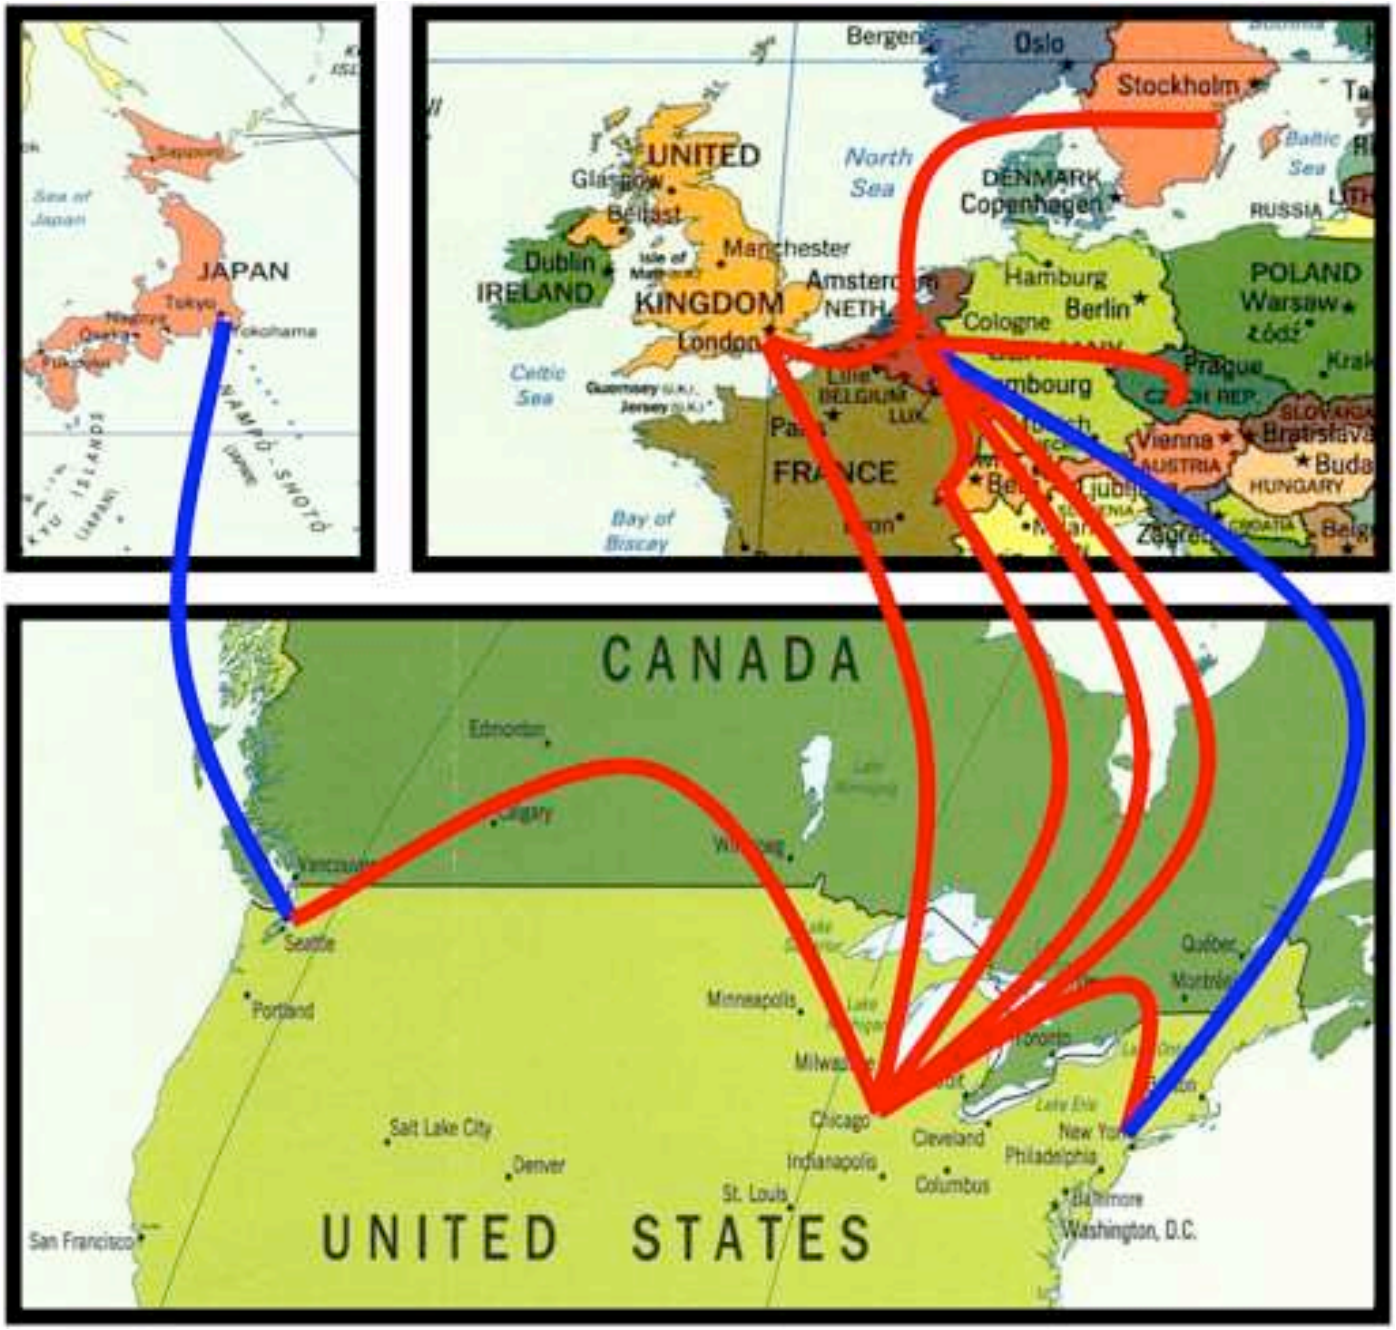

European lambdas to US
–6 GigEs Amsterdam—Chicago
–2 GigEs CERN—Chicago
–8 GigEs London—Chicago

Canadian lambdas to US
–8 GigEs Chicago—Canada—NYC

–8 GigEs
Chicago—Canada—Seattle

US lambdas to Europe
–4 GigEs Chicago—Amsterdam
–2 GigEs Chicago—CERN

European lambdas
–8 GigEs Amsterdam—CERN
–2 GigEs Prague—Amsterdam
–2 GigEs
Stockholm—Amsterdam
–8 GigEs London—Amsterdam

IEEAF lambdas (blue)
–8 GigEs Seattle—Tokyo
–8 GigEs NYC—Amsterdam

NetherLight UvA Setup

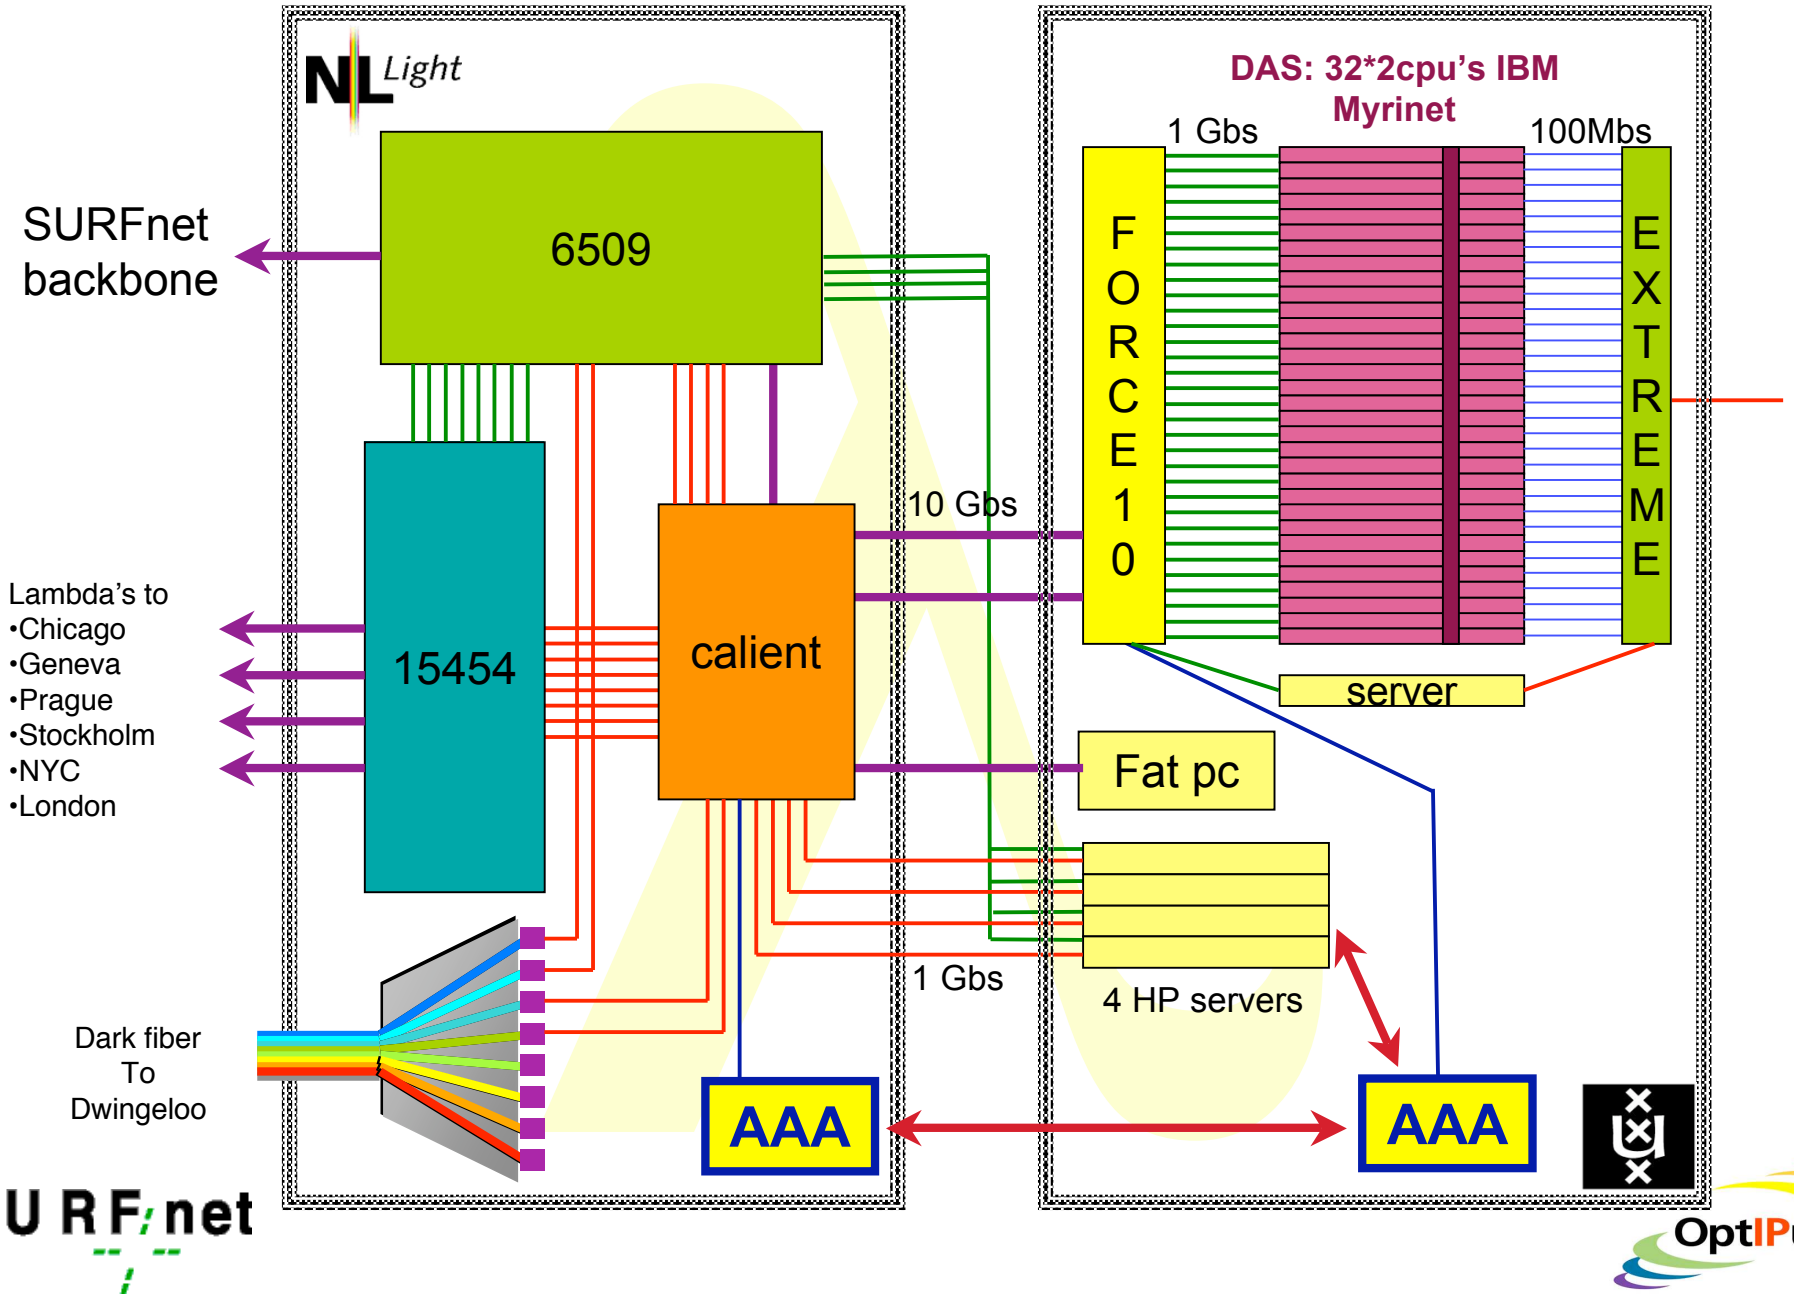

SURFnet

OptIPuter

UVA/EVL's
64*64
Optical Switch
@ NetherLight
in SURFnet POP
@ SARA

Costs 1/100th of
a similar
throughput router
but with specific
services!

Core Switch Technology

(Intermezzo)
Calient NETWORKS

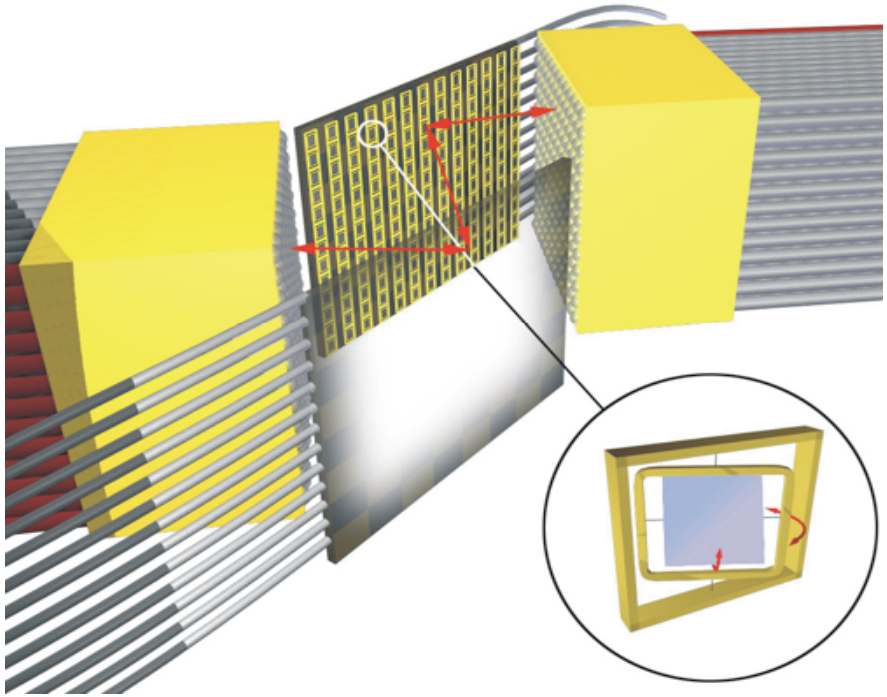

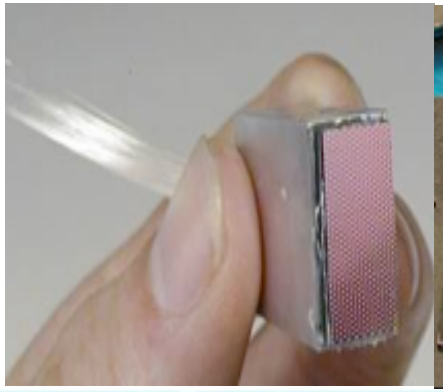

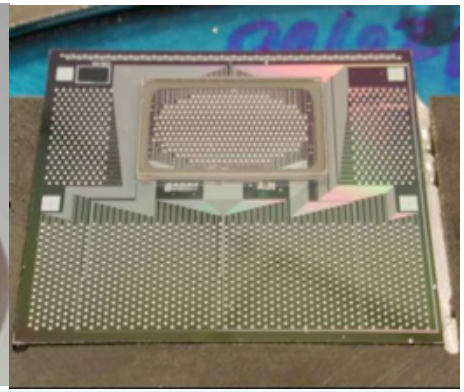

- **3D MEMS structure**
  - Bulk MEMS – High Density Chips
  - Electrostatic actuation
  - Short path length (~4cm)
  - <1.5 dB median loss
- **Completely Non-blocking**
  - Single-stage up to 1Kx1K
  - 10 ms switching time
- **Excellent Transparency**
  - Polarization
  - Bit rate
  - Wavelength

OMNInet Technology Trial: January 2002

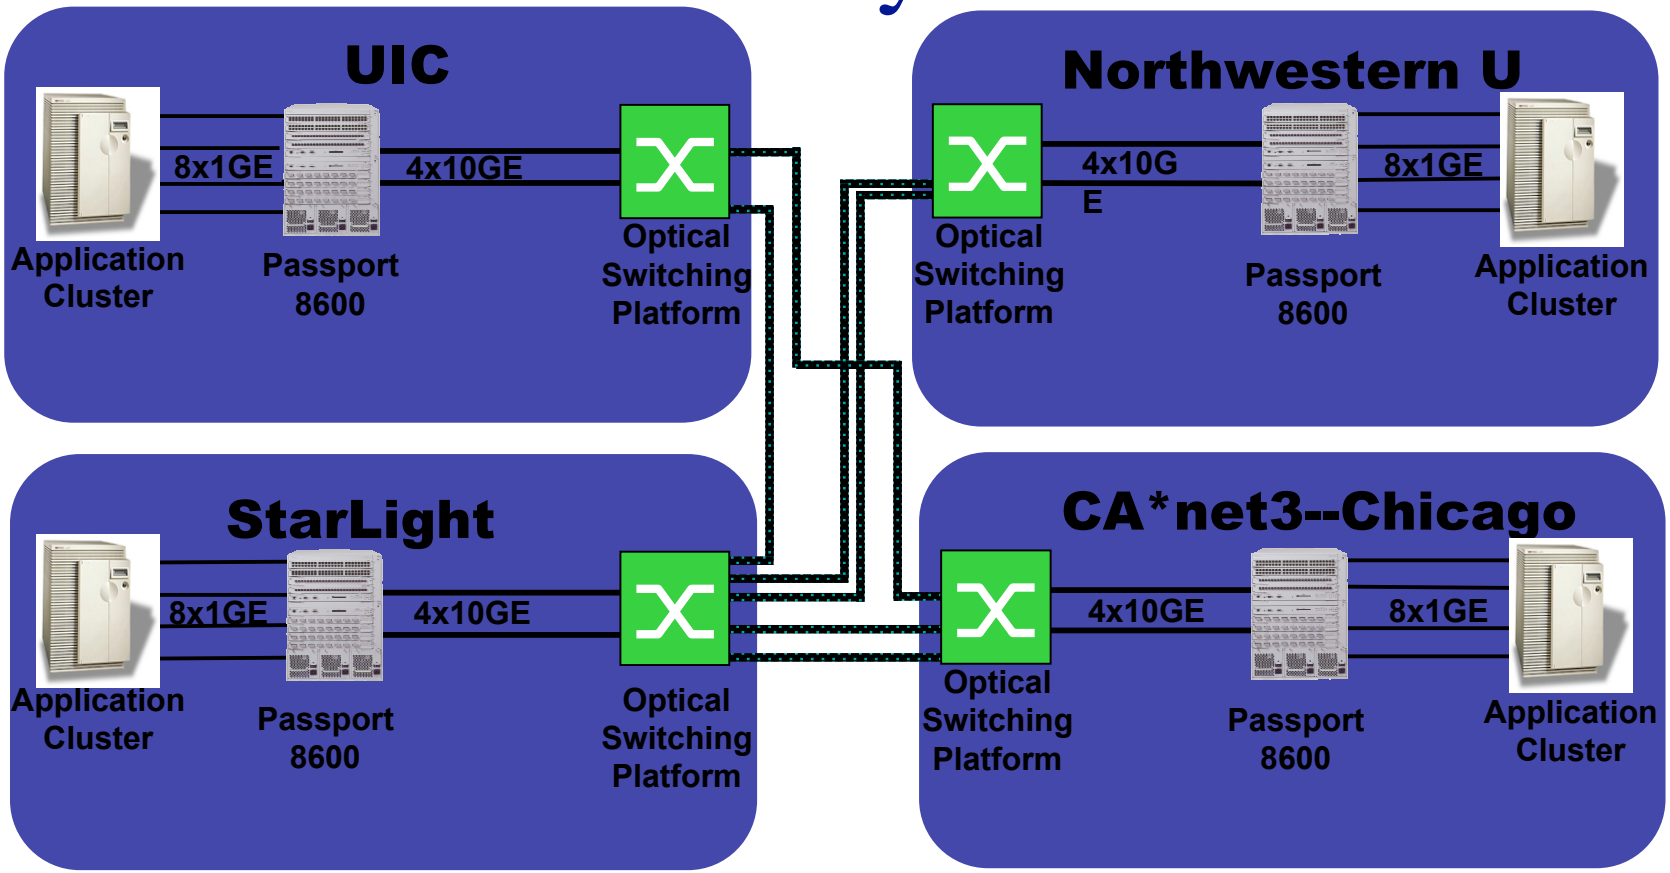

- A four-site network in Chicago -- the first 10GE service trial!
- A test bed for all-optical switching and advanced high-speed services
- Partners: SBC, Nortel, iCAIR at Northwestern, EVL, CANARIE, ANL

iCAIR

Source: Joe Mambretti

What is StarLight?

StarLight is an experimental optical infrastructure and proving ground for network services optimized for high-performance applications
GE+2.5+10GE Exchange
Soon: Multiple 10GEs Over Optics – World's "Largest" 10GE Exchange!

View from StarLight

Abbott Hall, Northwestern University's Chicago downtown campus

Source: Joe Mambretti

Where is StarLight?

- Located in Northwestern's Downtown Campus: 710 N. Lake Shore Drive
- Carrier POPs
- Chicago NAP

Source: Joe Mambretti

StarLight Infrastructure

StarLight is *a large research-friendly co-location facility* with space, power and fiber that is being made available to university and national/international network collaborator as a *point* of presence in Chicago

Source: Joe Mambretti

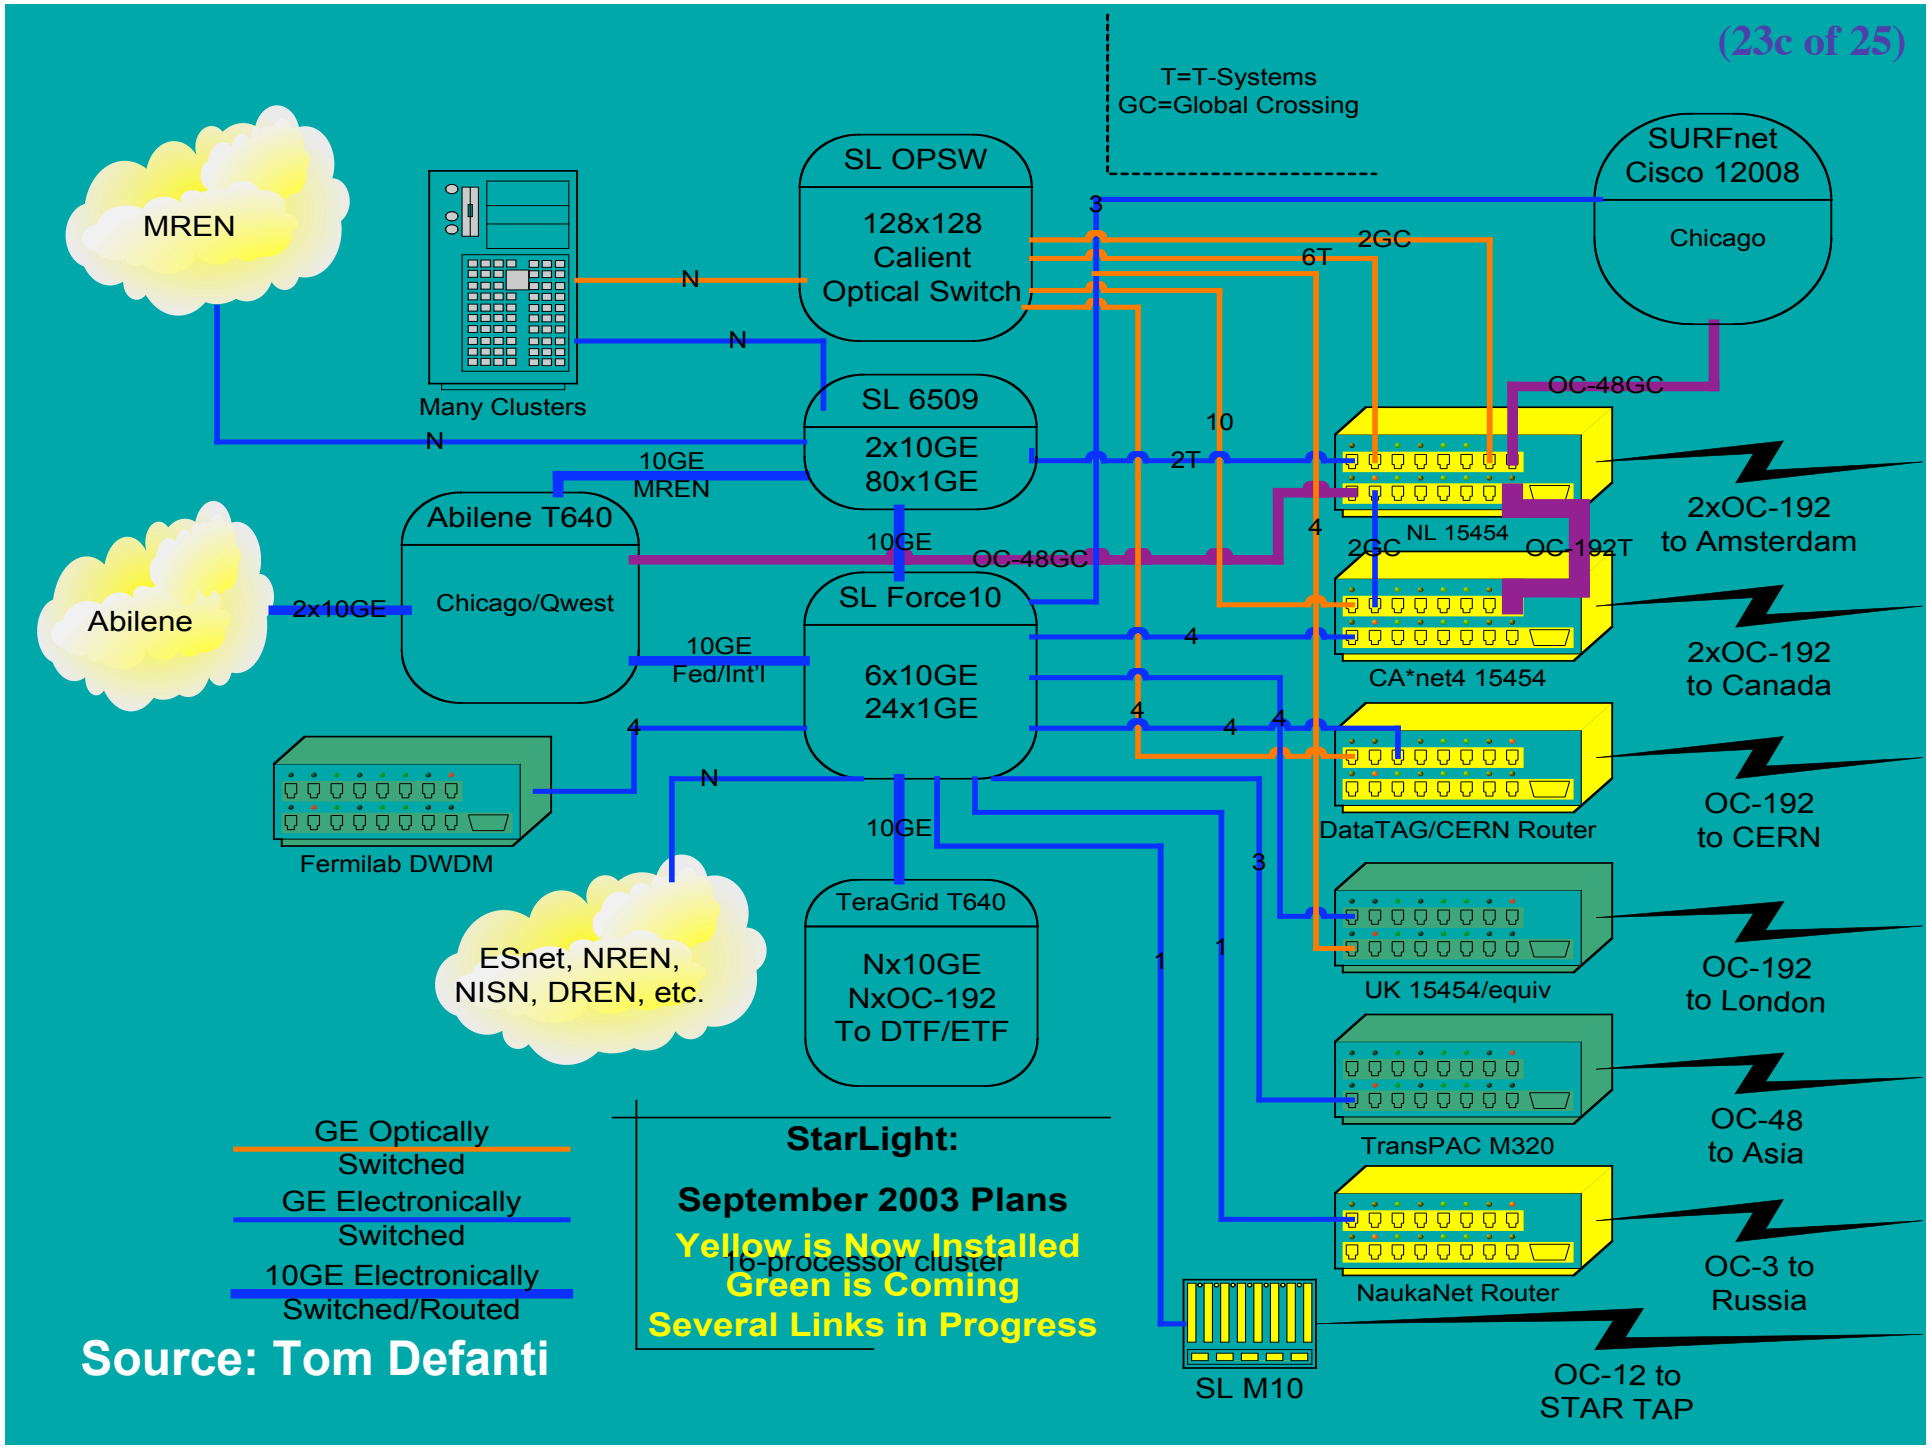
T=T-Systems
GC=Global Crossing
MREN
Many Clusters
SL OPSW
128x128
Calient
Optical Switch
SURFnet
Cisco 12008
Chicago
SL 6509
2x10GE
80x1GE
10GE
MREN
Abilene T640
Chicago/Qwest
Abilene
2x10GE
10GE
Fed/Int'l
SL Force10
6x10GE
24x1GE
OC-48GC
2GC
6T
2T
10
4
NL 15454
OC-192T
CA*net4 15454
DataTAG/CERN Router
UK 15454/equiv
TransPAC M320
NaukaNet Router
SL M10
16-processor cluster
Fermilab DWDM
ESnet, NREN,
NISN, DREN, etc.
TeraGrid T640
Nx10GE
NxOC-192
To DTF/ETF
2xOC-192
to Amsterdam
2xOC-192
to Canada
OC-192
to CERN
OC-192
to London
OC-48
to Asia
OC-3 to
Russia
OC-12 to
STAR TAP
GE Optically
Switched
GE Electronically
Switched
10GE Electronically
Switched/Routed
StarLight:
September 2003 Plans
Yellow is Now Installed
Green is Coming
Several Links in Progress
Source: Tom Defanti

CA*net 4 Physical Architecture

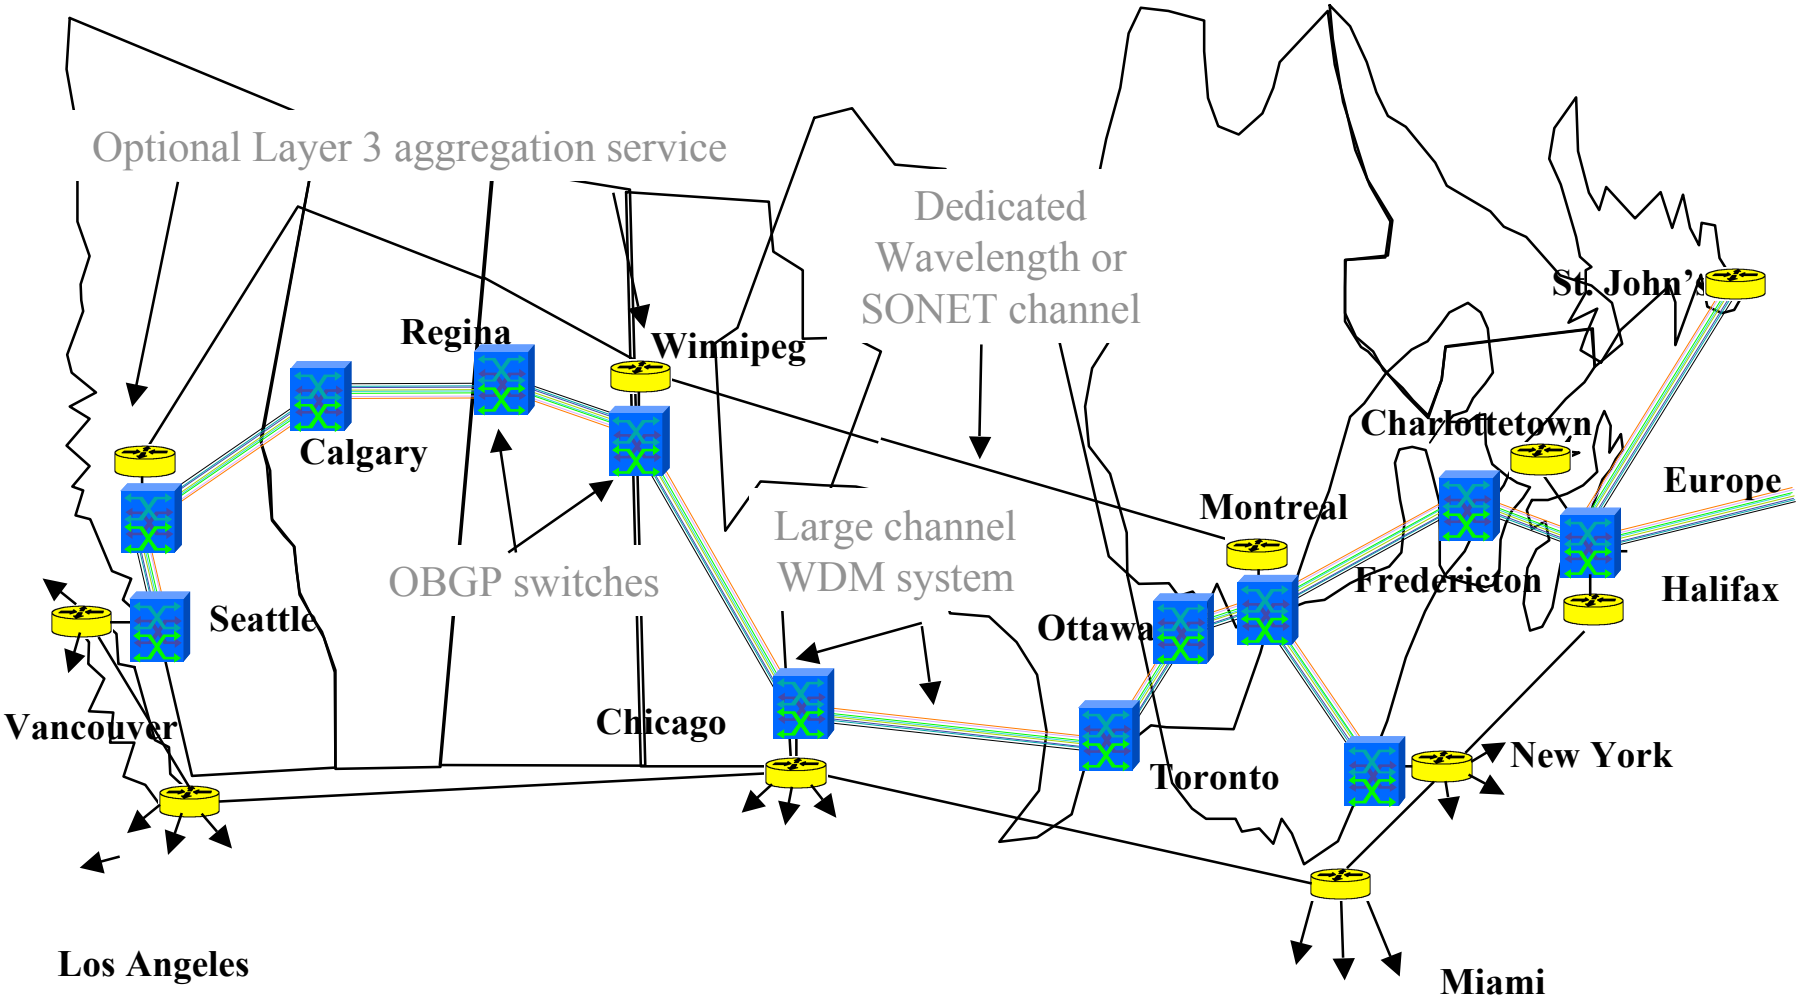

Source: Bill St Arnaud

LambdaGrid Control Plane Paradigm Shift

Traditional Provider Services:
Invisible, Static Resources,
Centralized Management

OptIPuter:
Distributed Device, Dynamic Services,
Visible & Accessible Resources,
Integrated As Required By Apps

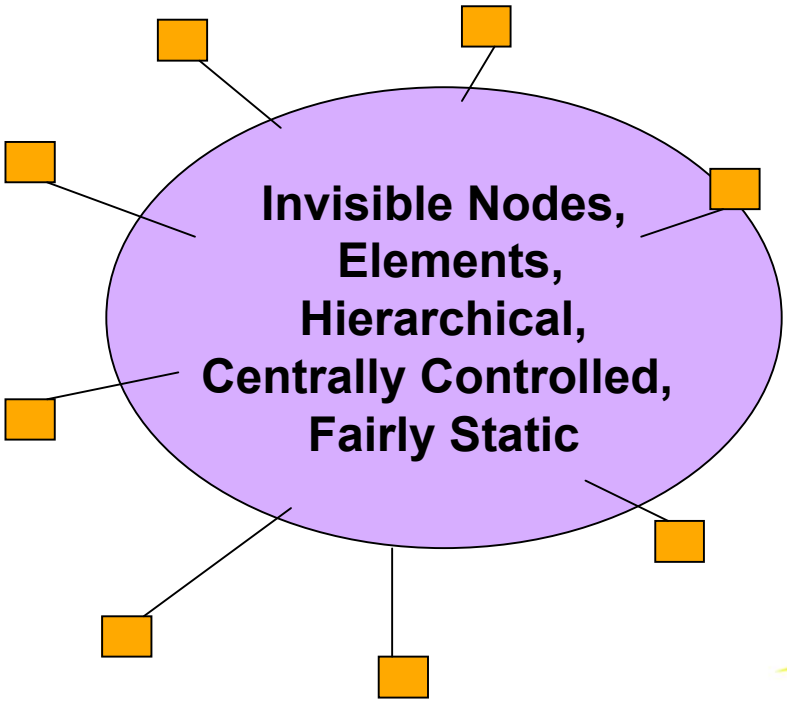

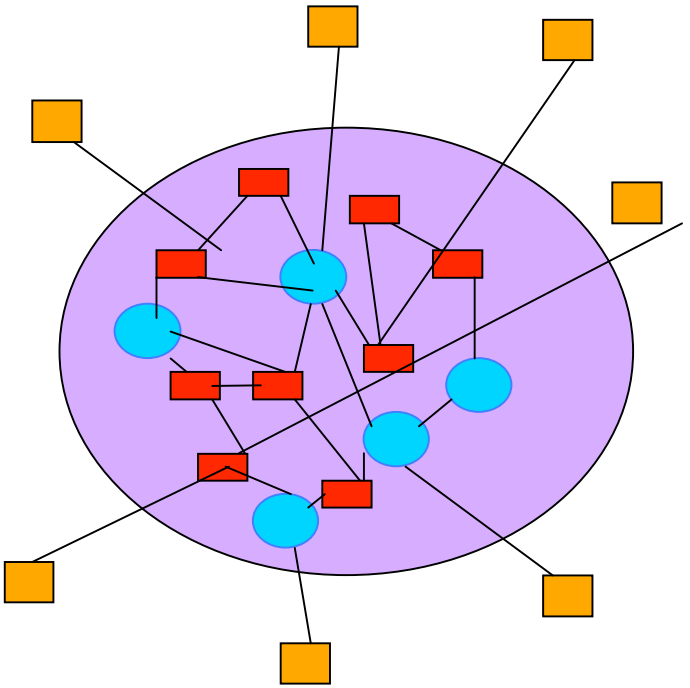

Limited Functionality,
Flexibility

OptIPuter

Unlimited Functionality,
Flexibility

Source: Joe Mambretti, Oliver Yu, George Clapp

Multi Domain Lambda setup

- AAA based on RFC 2903-2906
- OGSI wrapper
- Interface to CALIENT optical switch, layer 2 switches
- Interface to PDC
- Broker for path searching, selection
- Web and application interface
- Demonstration on SC2003

Multi Domain Lambda setup

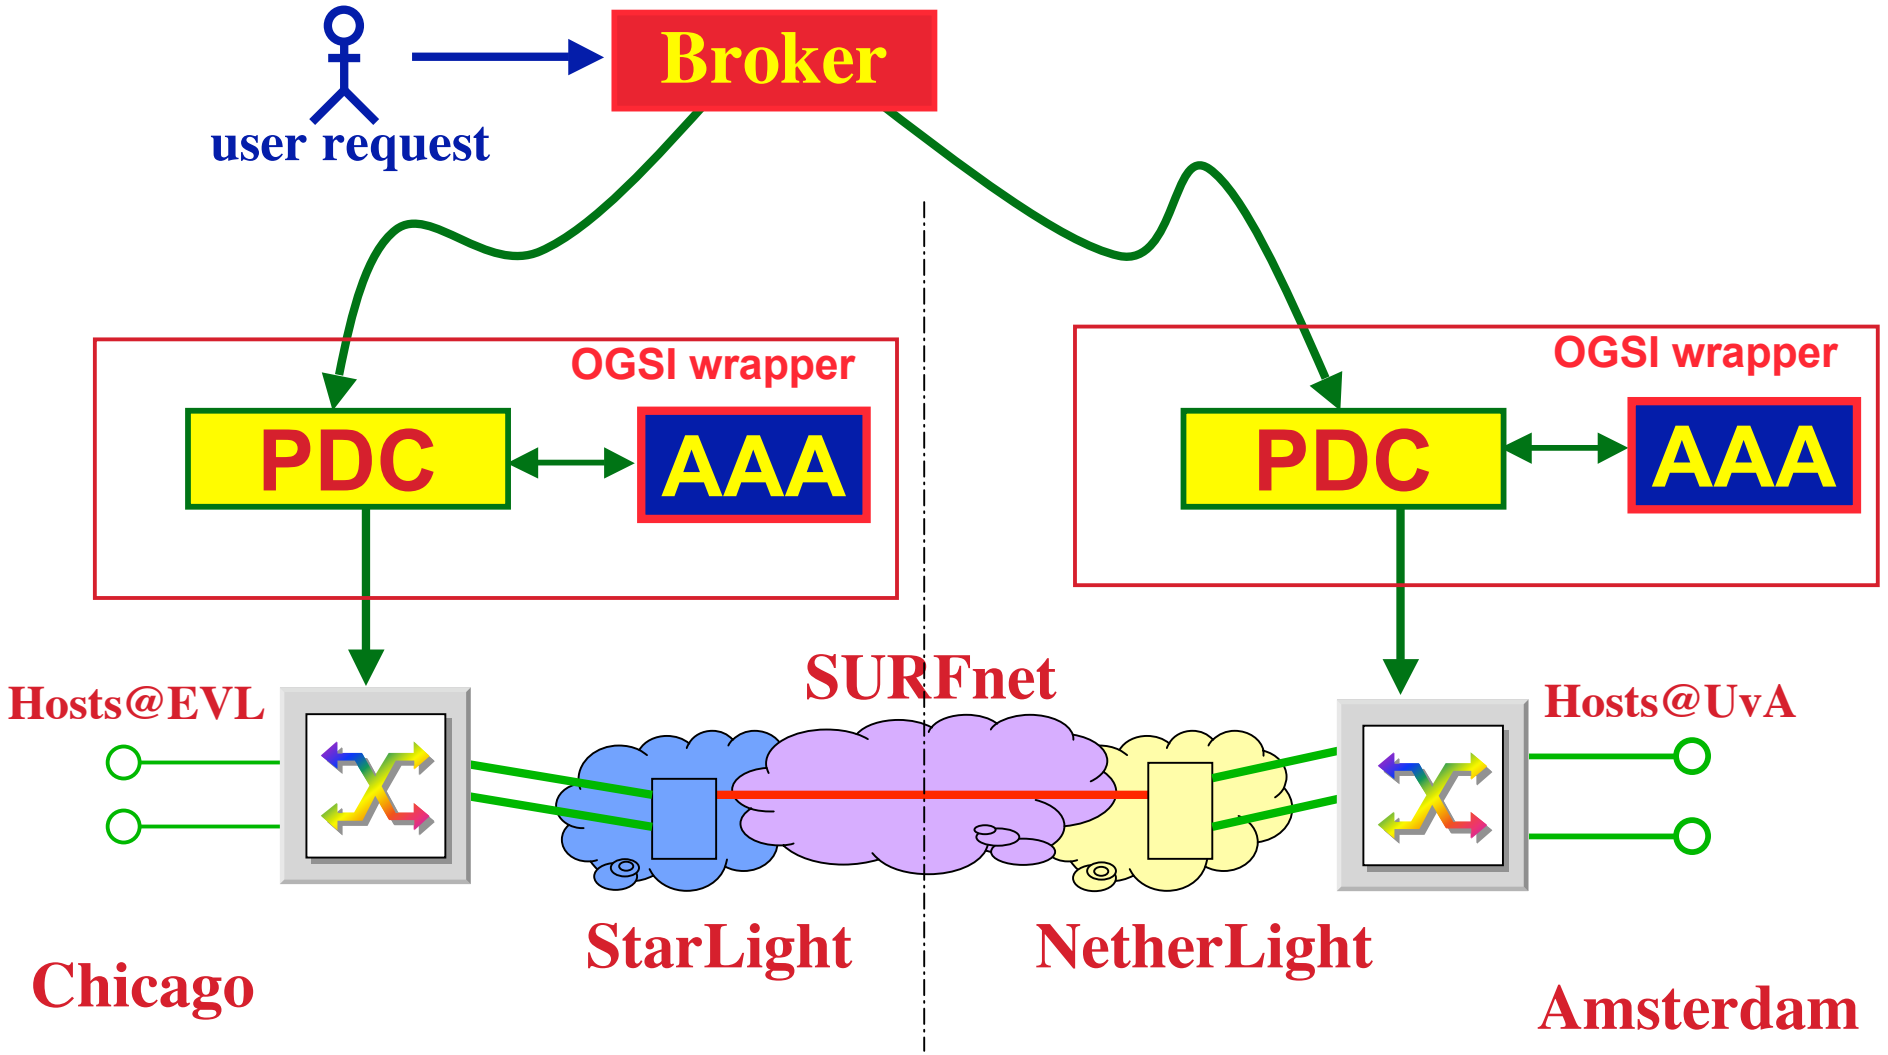

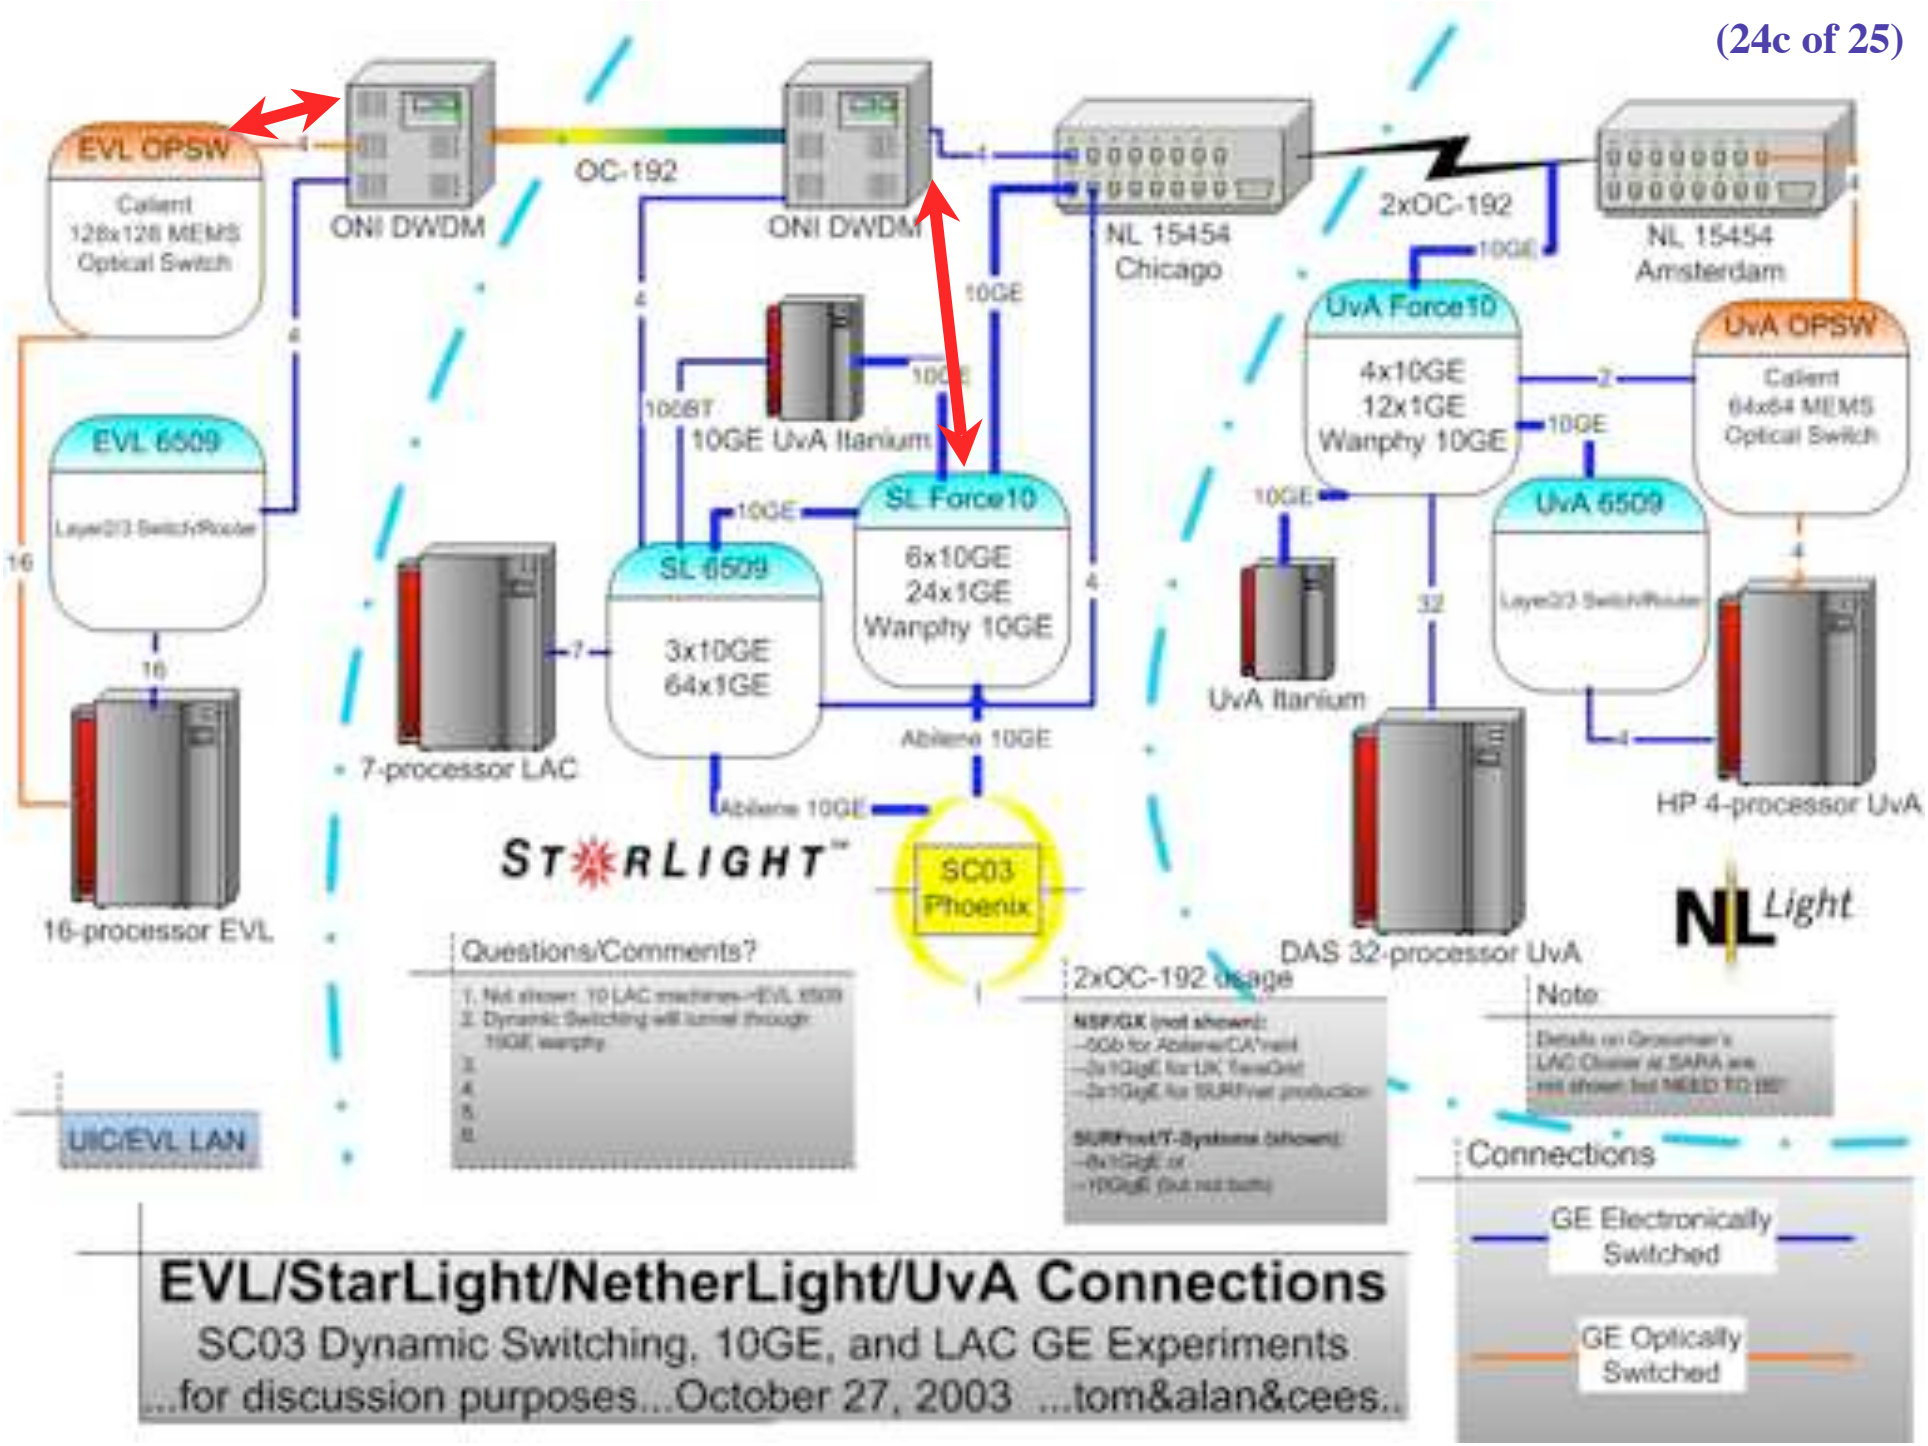
EVL OPSW
Calient
128x128 MEMS
Optical Switch
ONI DWDM
OC-192
ONI DWDM
NL 15454
Chicago
2xOC-192
NL 15454
Amsterdam
UvA Force10
4x10GE
12x1GE
Wanphy 10GE
UvA OPSW
Calient
64x64 MEMS
Optical Switch
10GE UvA Itanium
EVL 6509
SL Force10
6x10GE
24x1GE
Wanphy 10GE
SL 6509
3x10GE
64x1GE
UvA 6509
UvA Itanium
7-processor LAC
HP 4-processor UvA
16-processor EVL
DAS 32-processor UvA
STARLIGHT
SC03
Phoenix
NL Light
Abilene 10GE
Questions/Comments?
2xOC-192 usage
Note
Connections
GE Electronically
Switched
GE Optically
Switched
UIC/EVL LAN
EVL/StarLight/NetherLight/UvA Connections
SC03 Dynamic Switching, 10GE, and LAC GE Experiments
...for discussion purposes...October 27, 2003 ...tom&alan&cees..

Research topics

- Optical networking architectures and models for usage
- Transport protocols for massive amounts of data
- Authorization of complex resources in multiple domains
- Embedding in Grid environments

3th Lambda workshop @ NORDUnet 2003

Lambda workshop

- **Amsterdam - Terena**
  - **Concepts**
  - **Initial testbed (SURFnet Lambda to StarLight)**
- **Amsterdam - iGrid2002**
  - **Rechecking concepts models**
  - **Initial experiences and measurements**
  - **Expansion of Lambda testbed**
- **Reykjavik - NORDUnet**
  - **Towards persistent demonstrations and applications**

Transport in the corners

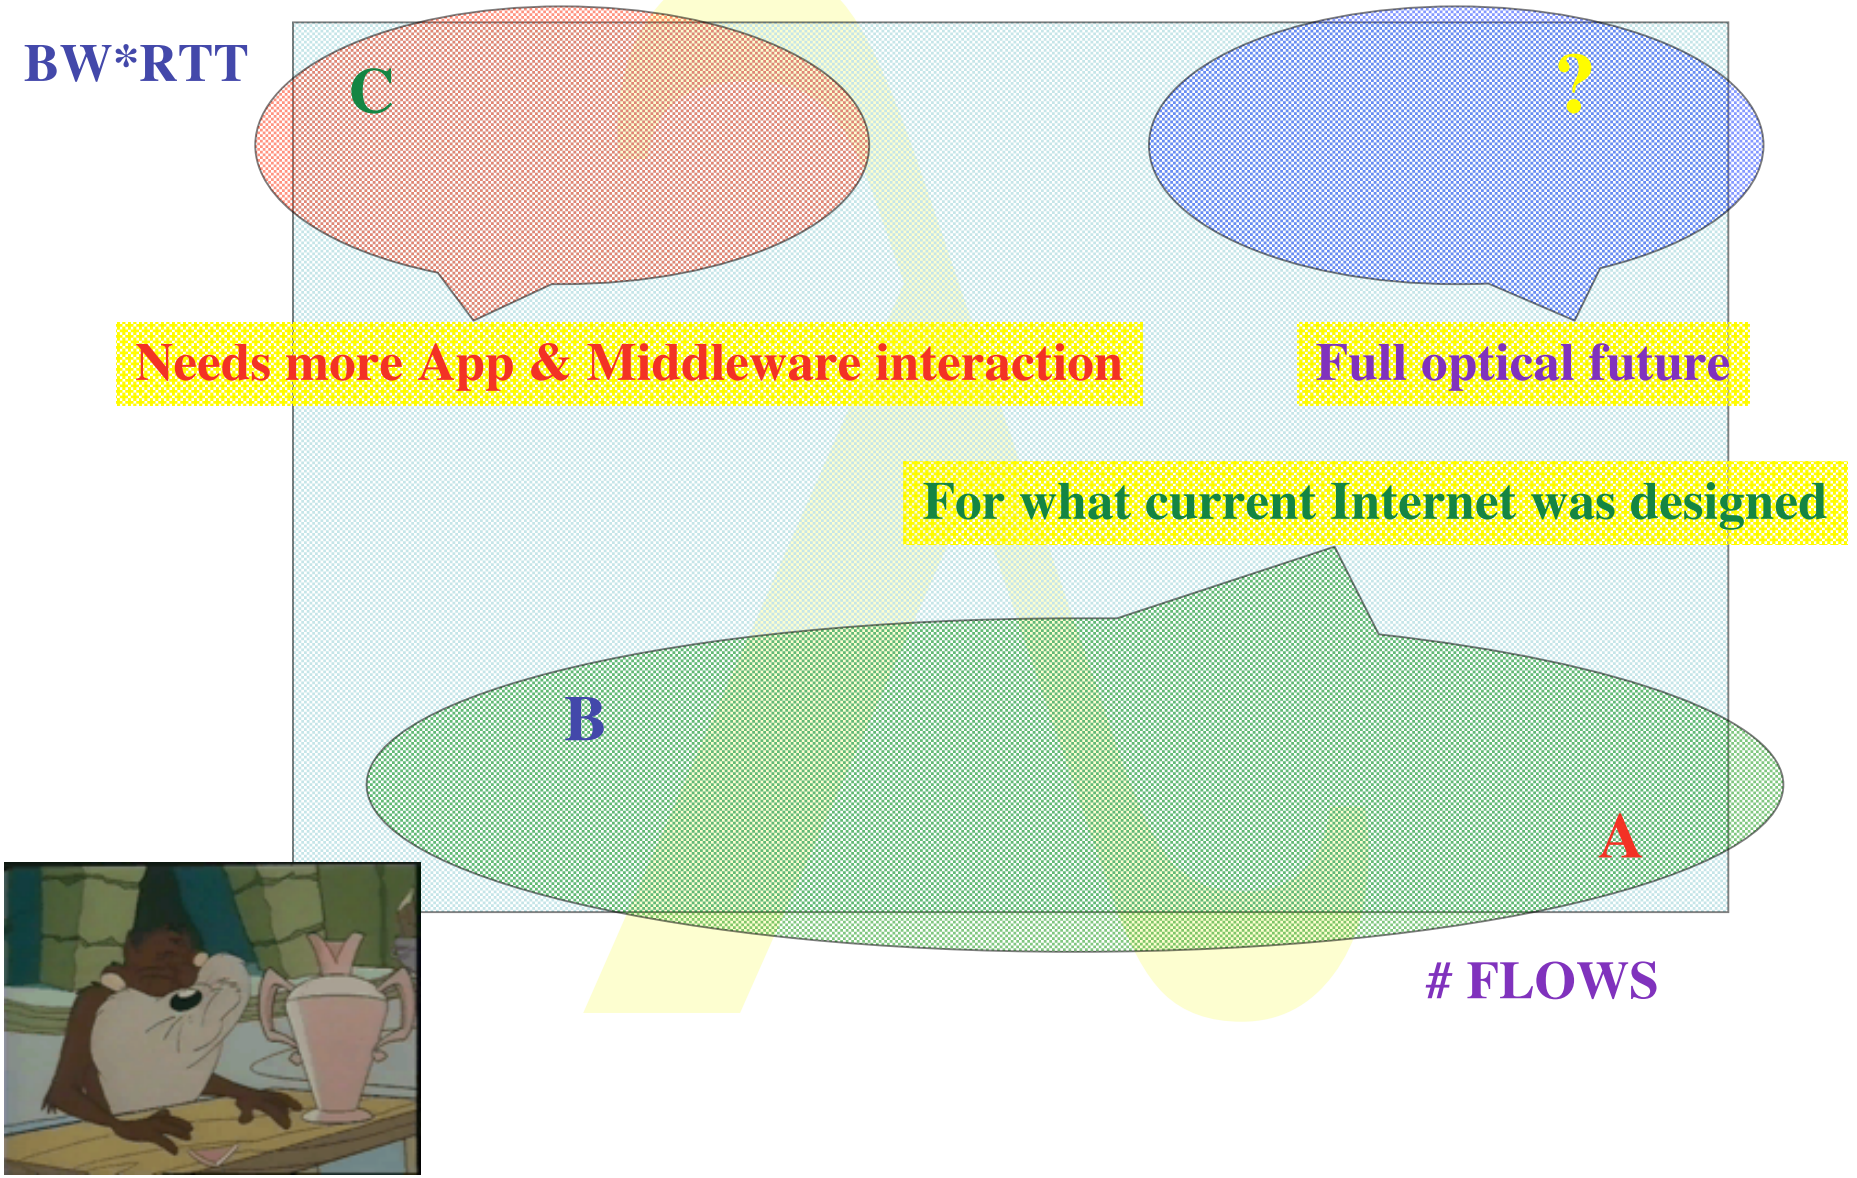

Revisiting the truck of tapes (one but last)

Consider one fiber

- **Current technology allows 320 λ in one of the frequency bands**
- **Each λ has a bandwidth of 40 Gbit/s**
- **Transport: 320 * 40*10^9 / 8 = 1600 GByte/sec**

- **Take a 10 metric ton truck**
  - **One tape contains 50 Gbyte, weights 100 gr**
  - **Truck contains (10000 / 0.1) * 50 Gbyte = 5 PByte**

- **Truck / fiber = 5 PByte / 1600 GByte/sec = 3125 s ≈ one hour**
- **For distances further away than a truck drives in one hour (50 km) minus loading and handling 100000 tapes the fiber wins!!!**

The END

Thanks to

SURFnet: Kees Neggers,UIC&iCAIR: Tom DeFanti, Joel Mambretti, CANARIE: Bill St. Arnaud

Freek Dijkstra, Hans Blom, Leon Gommans, Bas van oudenaarde, Arie Taal, Pieter de Boer, Bert Andree, Martijn de Munnik, Antony Antony, Rob Meijer, VL-team.