Lambda-Grid developments

Global Lambda Integrated Facility

www.science.uva.nl/~delaat

Cees de Laat

GigaPort EU

University of Amsterdam

SARA

NCF

Contents

This page is intentionally left blank

- Ref: www.this-page-intentionally-left-blank.org

Four LHC Experiments: The Petabyte to Exabyte Challenge

- **ATLAS, CMS, ALICE, LHCB**

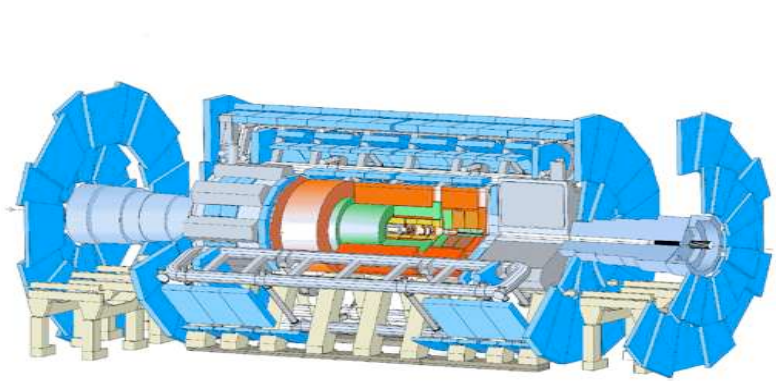

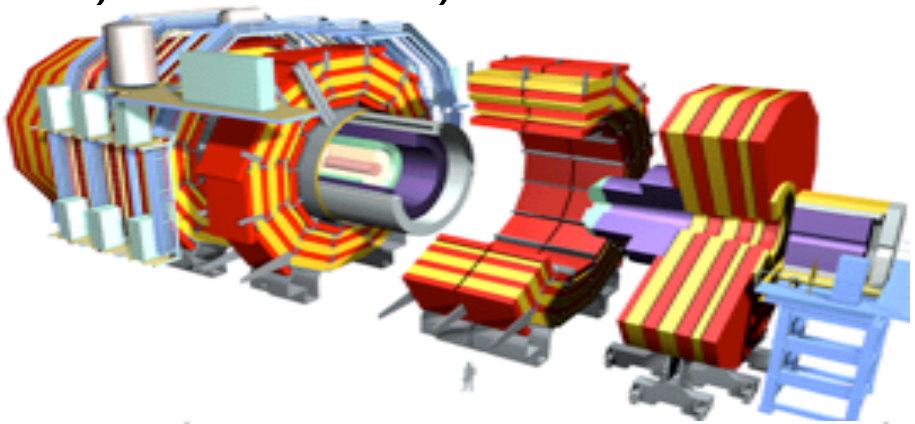

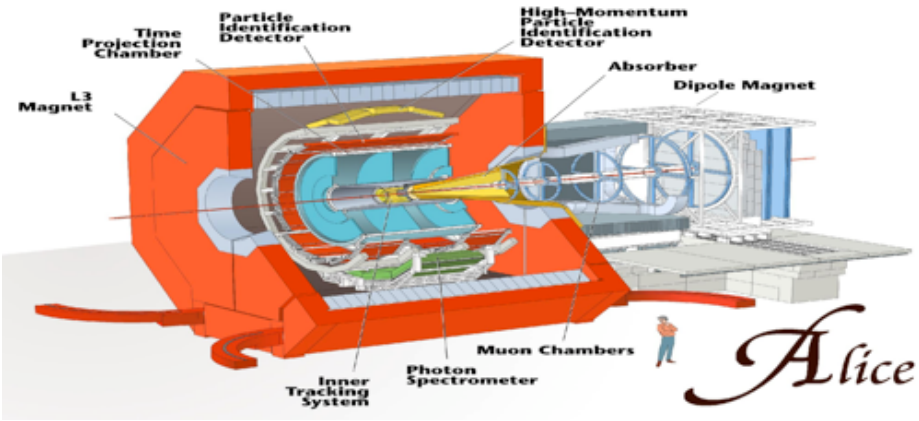

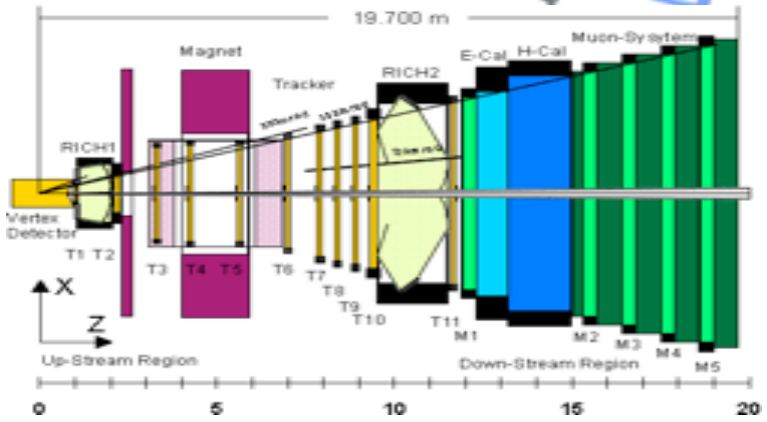

6000+ Physicists & Engineers; 60+ Countries; 250 Institutions

Tens of PB 2008; To 1 EB by ~2015

Hundreds of TFlops To PetaFlops

VLBI

ger term VLBI is easily capable of generating many Gb of data per
The sensitivity of the VLBI array scales w
(data-rate) and there is a strong push to
. Rates of 8Gb/s or more are entirely feasibl
der 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

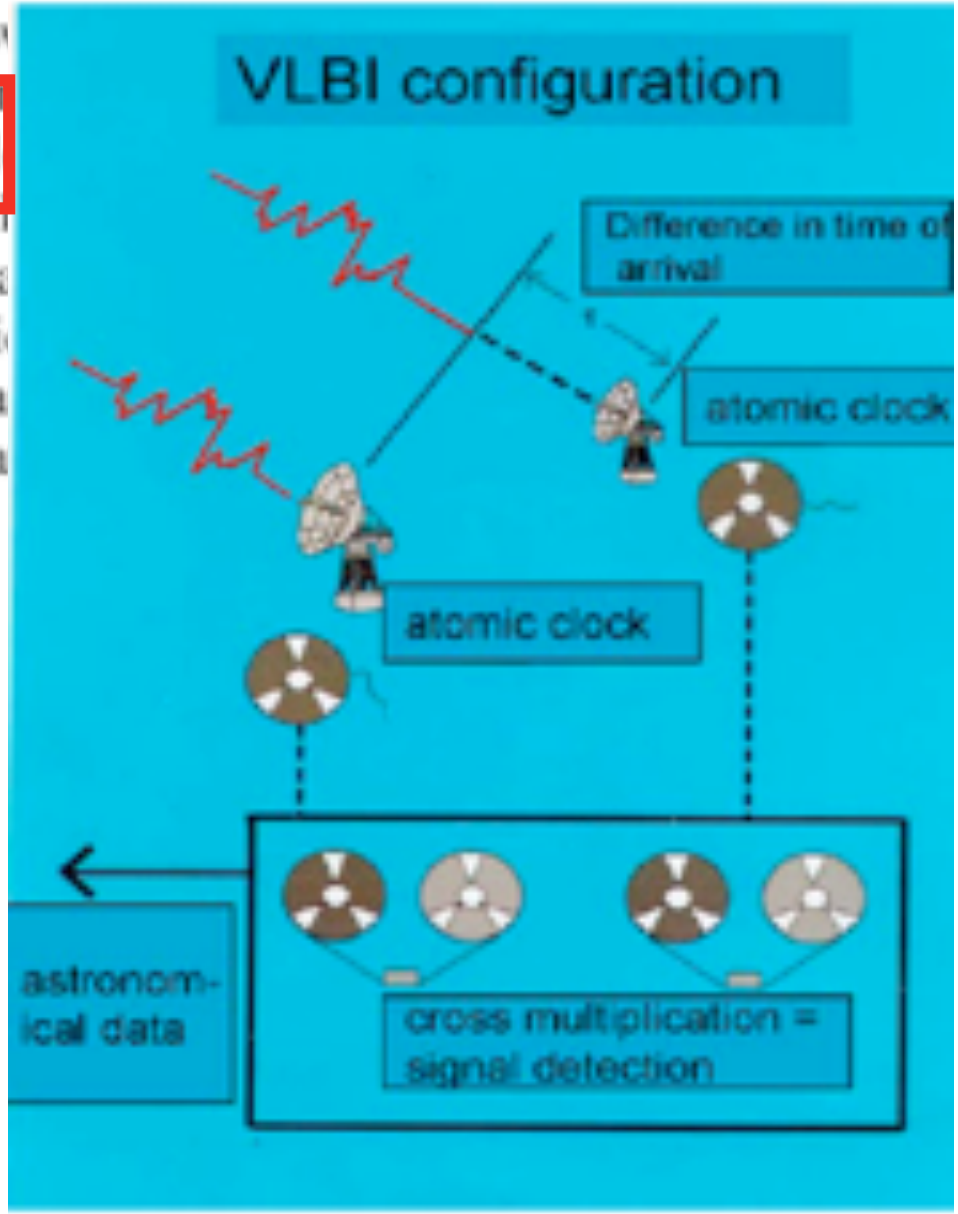

Lambdas as part of instruments

www.lofar.org

37 Tbit/s - 116 T-ops/s

http://www.lofar.org/p/systems.htm

http://web.haystack.mit.edu/lofar/technical.html

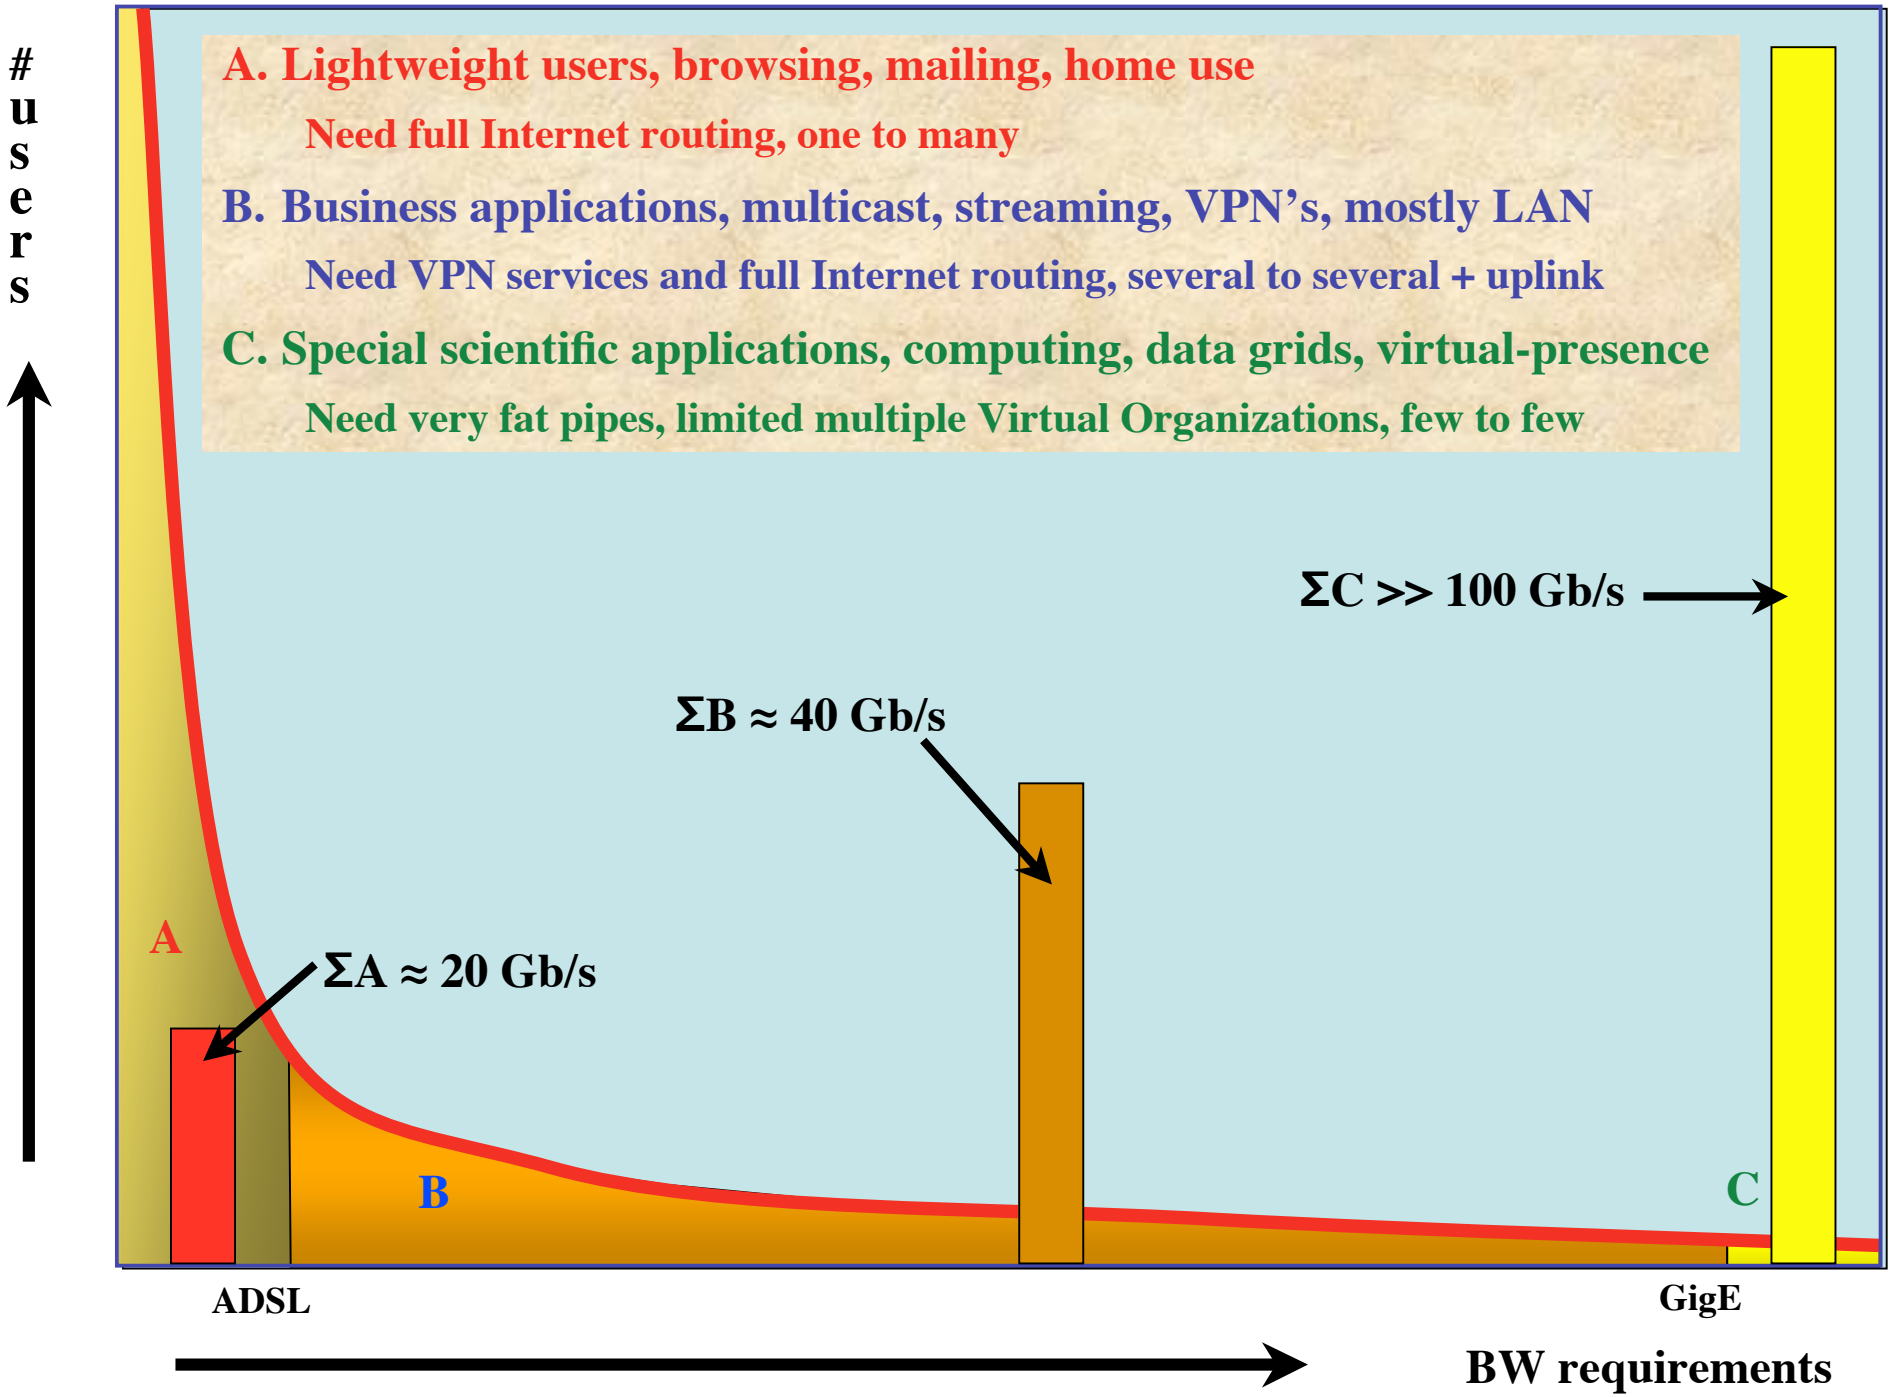
A. Lightweight users, browsing, mailing, home use
Need full Internet routing, one to many
B. Business applications, multicast, streaming, VPN's, mostly LAN
Need VPN services and full Internet routing, several to several + uplink
C. Special scientific applications, computing, data grids, virtual-presence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

So what?

- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
  - 10G routerblade -> 100-300 k$, 10G switch port -> 10-20 k$, MEMS port -> 0.7 k$
  - DWDM lasers for long reach expensive, 10-50k$ (???)
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way (A -> L3 , B -> L2 , C -> L1)
- Give each packet in the network the service it needs, but no more

L1 - 0.7 k$/port

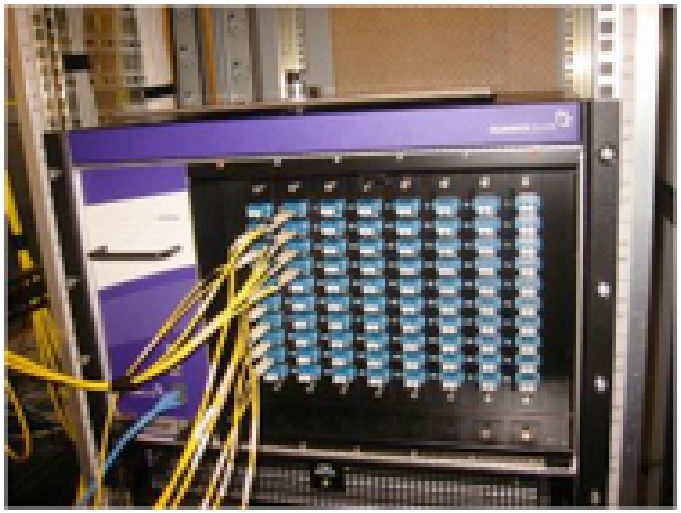

L2 - 10-20 k$/port

L3 - 100-300 k$/port

How low can you go?

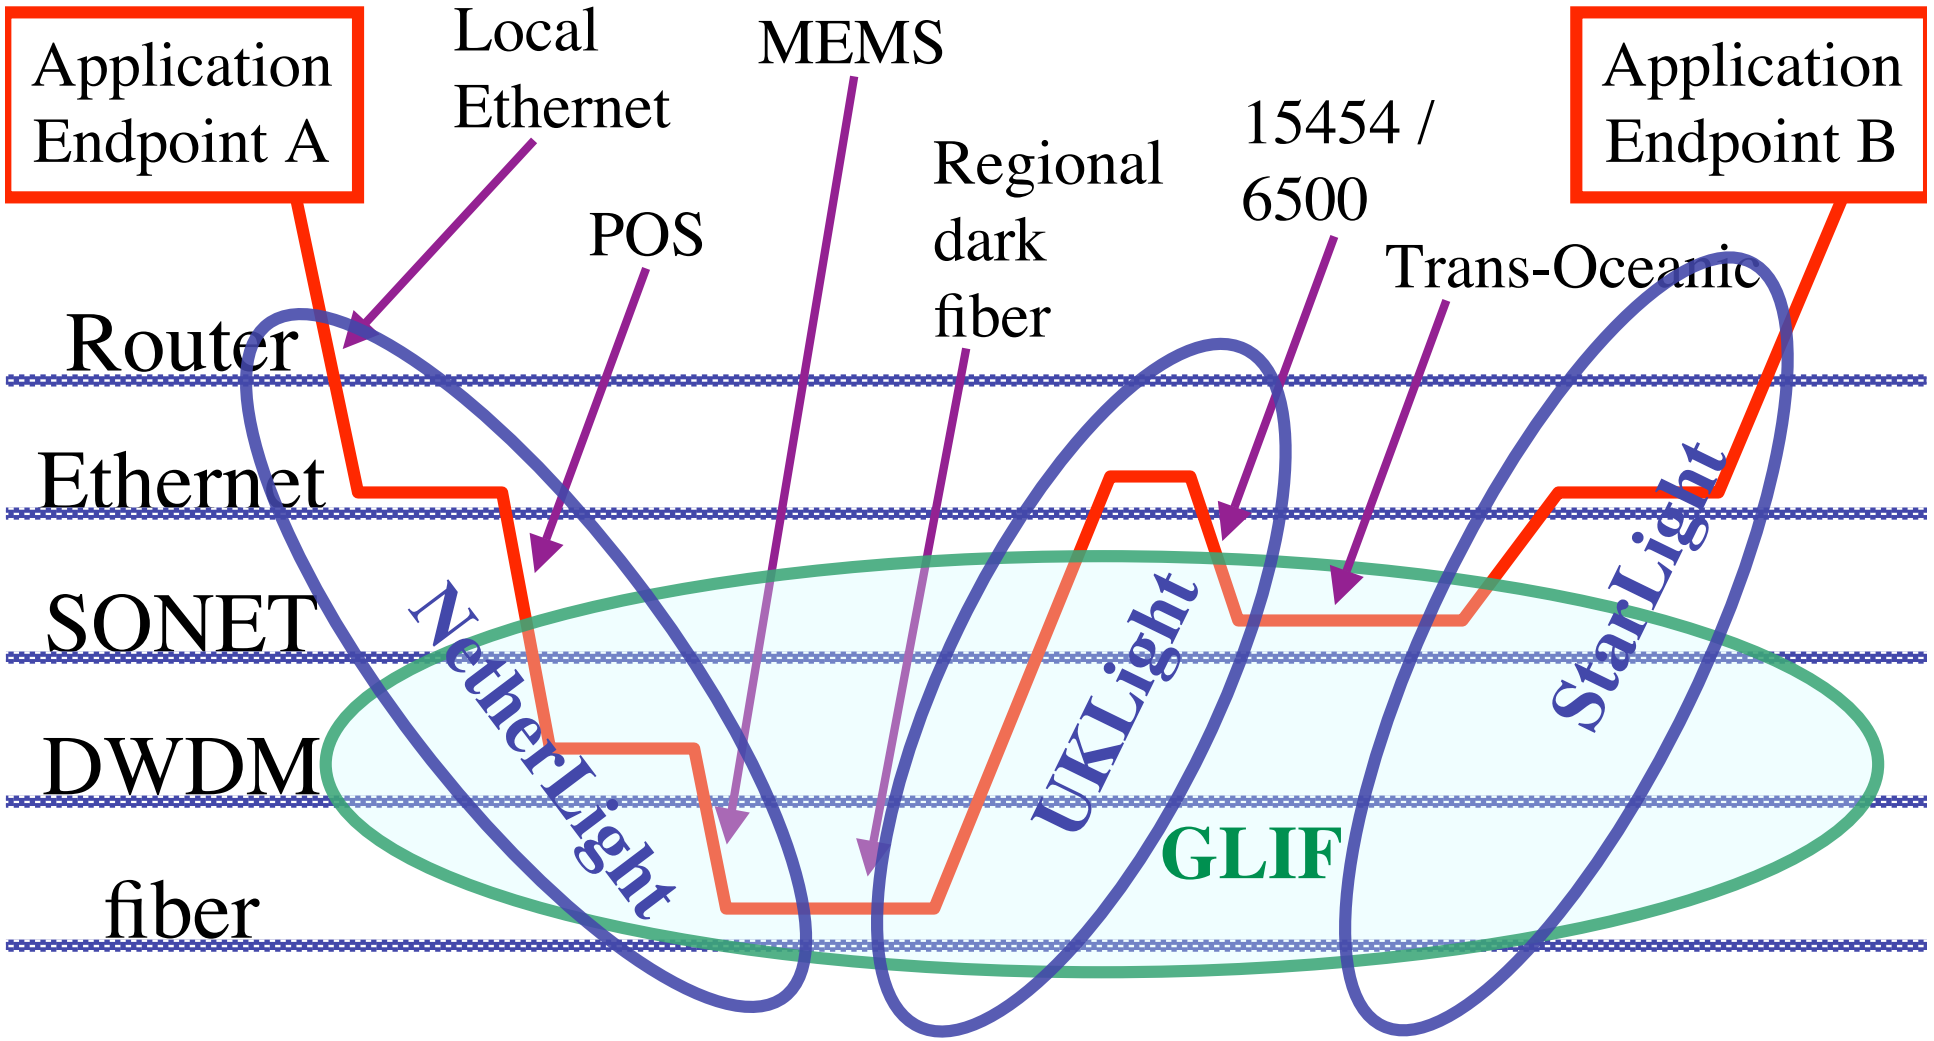

Optical Exchange as Black Box

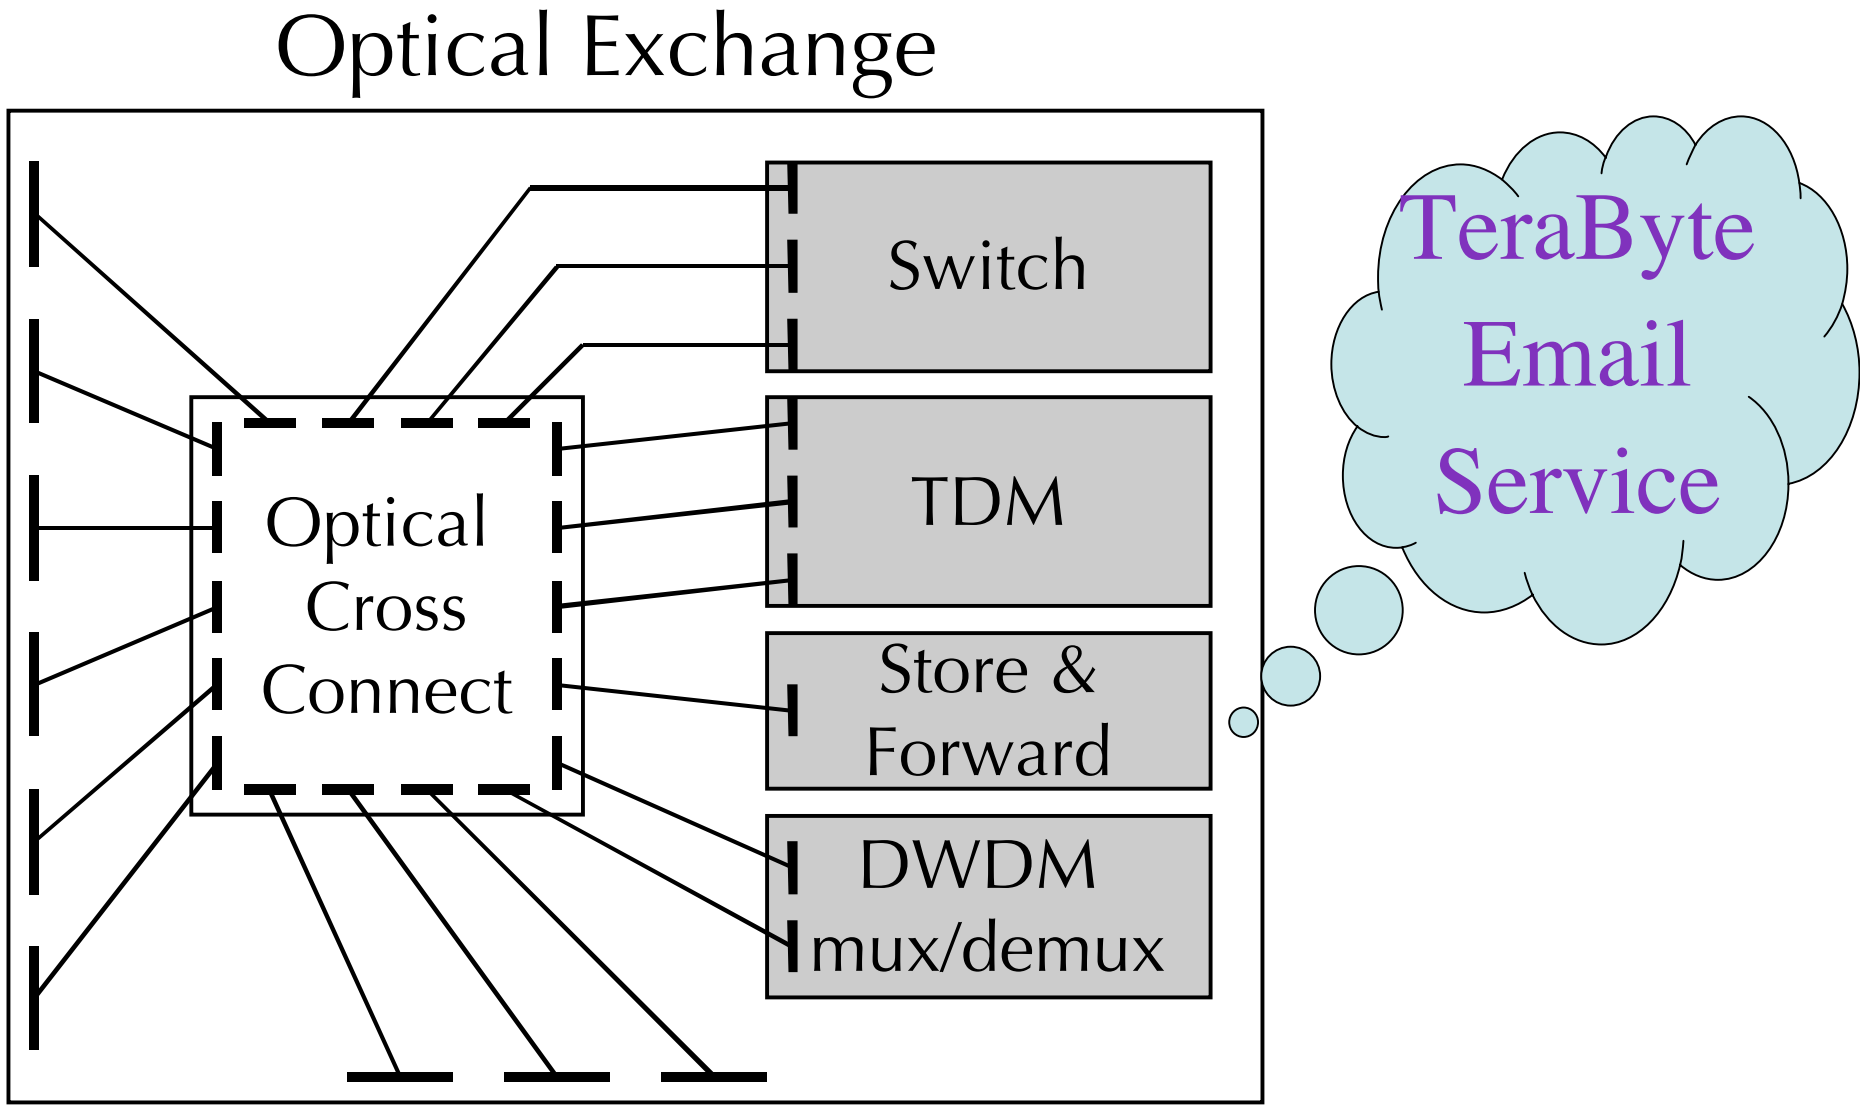

Discipline Networks

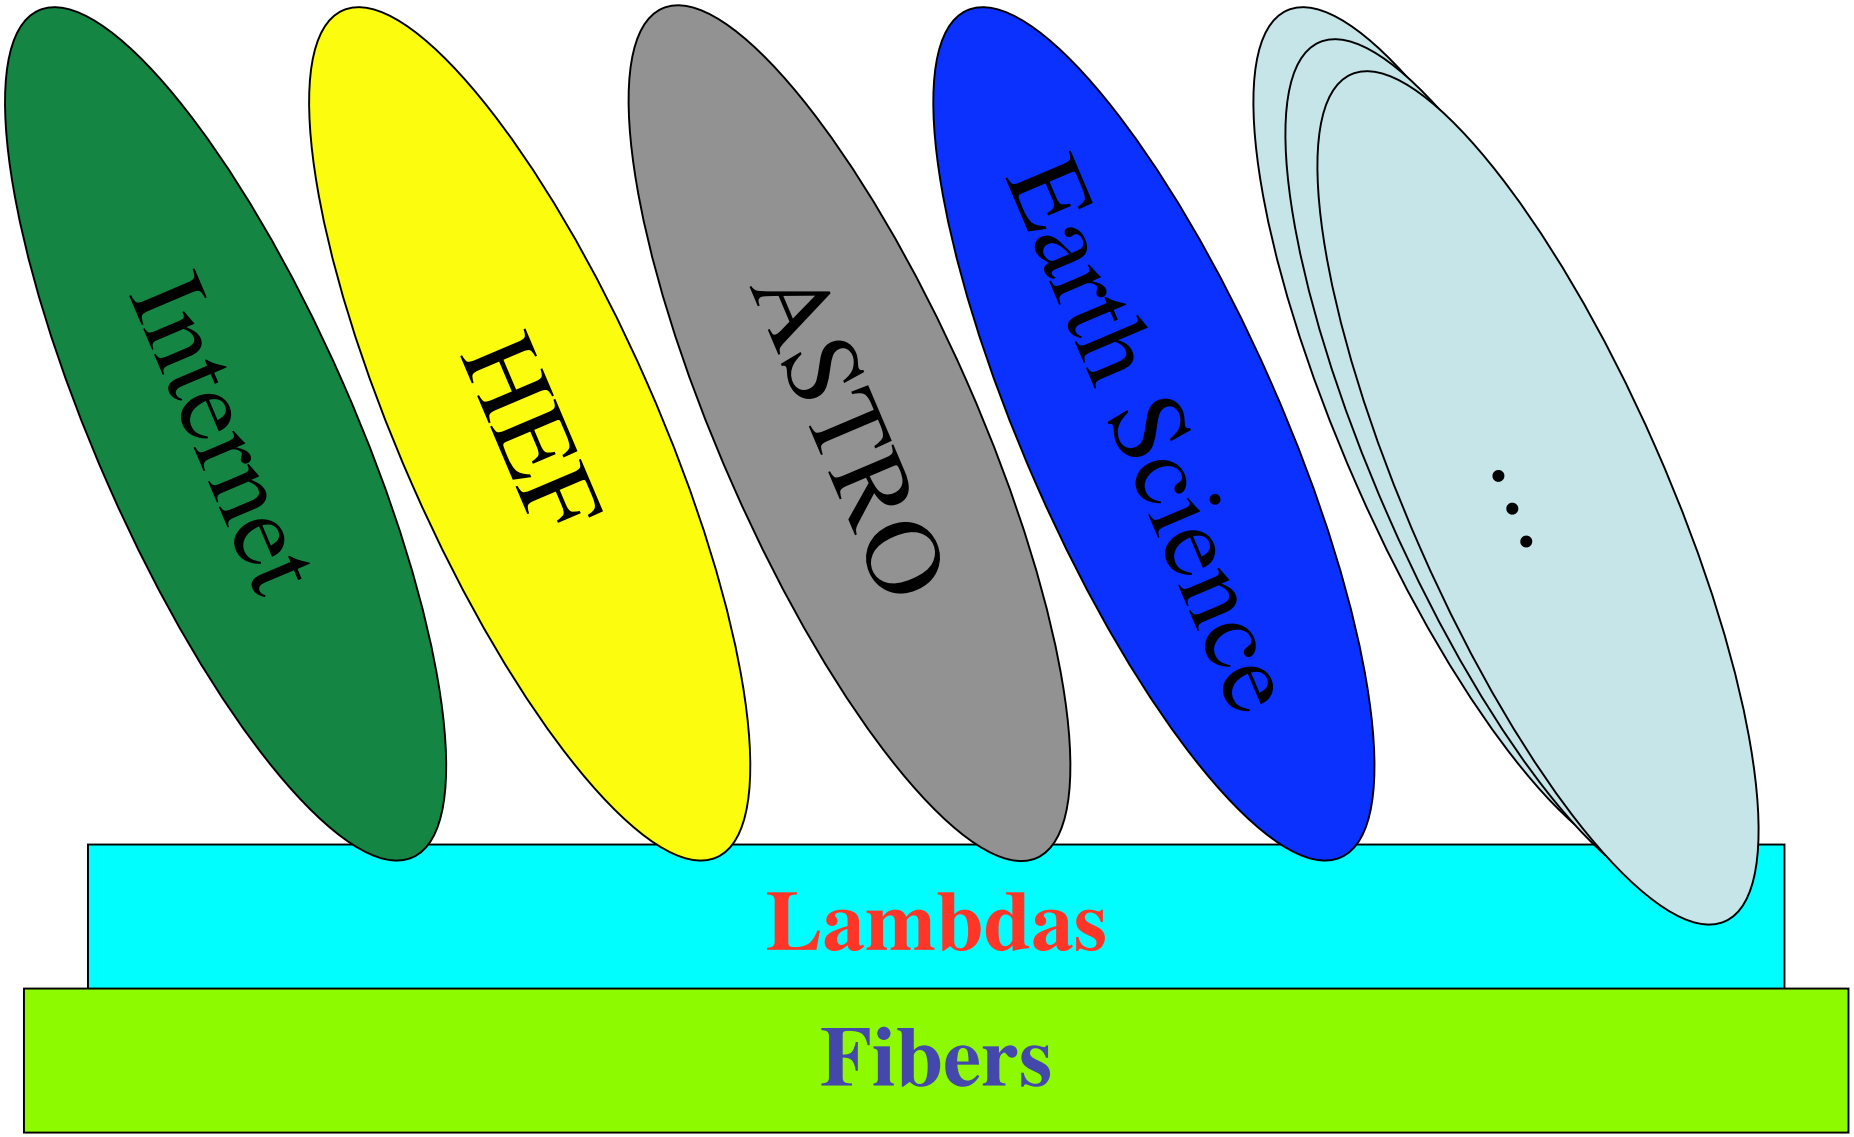

History of GLIF

- **Brainstorming in Antalya at Terena conf. 2001**
- **1th meeting at Terena offices 11-12 sep 2001**
  - **On invitation only (15) + public part**
  - **Thinking, SURFnet test lambda Starlight-Netherlight**
- **2nd meeting appended to iGrid 2002 in Amsterdam**
  - **Public part in track, on invitation only day (22)**
  - **Core testbed brainstorming, idea checks, seeds for Translight**
- **3th meeting Reykjavik, hosted by NORDUnet 2003**
  - **Grid/Lambda track in conference + this meeting (35!)**
  - **Brainstorm applications and showcases**
  - **Technology roadmap**
  - **GLIF established**

GLIF - 4 meeting

- Invitation only
- Nottingham 3 September (+preparatory afternoon on 2 September)
- 60 participants
- Attendance from China, Japan, Netherlands, Switzerland, US, UK, Taiwan, Australia, Tsjech, Korea, Canada, Ireland, Russia, Belgium, Denmark
- www.glif.is

GLIF Mission Statement

- **GLIF is a world-scale Lambda-based Laboratory for application and middleware development on emerging LambdaGrids, where applications rely on dynamically configured networks based on optical wavelengths**

- **GLIF is an environment (networking infrastructure, network engineering, system integration, middleware, applications) to accomplish real work**

Nottingham 2004

The main objectives of this year's (2004) meeting

GLIF Governance and policy

Our small-scale Lambda Workshop is now turning into a global activity. TransLight and similar projects contribute to the infrastructure part of GLIF. A good and well understood governance structure is key to the manageability and success of GLIF. Our prime goal is to decide upon and agree to the GLIF governance and infrastructure usage policy.

GLIF Lambda infrastructure and Lambda exchange implementations

A major function for previous Lambda Workshops was to get the network engineers together to discuss and agree on the topology, connectivity and interfaces of the Lambda facility. Technology developments need to be folded into the architecture and the expected outcome of this meeting is an agreed view on the interfaces and services of Lambda exchanges and a connectivity map of Lambdas for the next year, with a focus on iGrid 2005 and the emerging applications.

Persistent Applications

Key to the success of the GLIF effort is to connect the major applications to the Facility. We, therefore, need a list of prime applications to focus on and a roadmap to work with those applications to get them up to speed. The demonstrations at SC2004 and iGrid 2005 can be determined in this meeting.

Control Plane and Grid Integration

The GLIF can only function if we agree on the interfaces and protocols that talk to each other in the control plane on the contributed Lambda resources. The main players in this field are already meeting, almost on a bi-monthly schedule. Although not essential, this GLIF meeting could also host a breakout session on control plane middleware.

GLIF Q3 2004

Visualization courtesy of
Bob Patterson, NCSA.

UCLP intended for projects like National LambdaRail

CAVEwave partner acquires a separate wavelength between San Diego and Chicago and wants to manage it as part of its network including add/drop, routing, partition etc

NLR Condominium lambda network

CA*net 4 Architecture

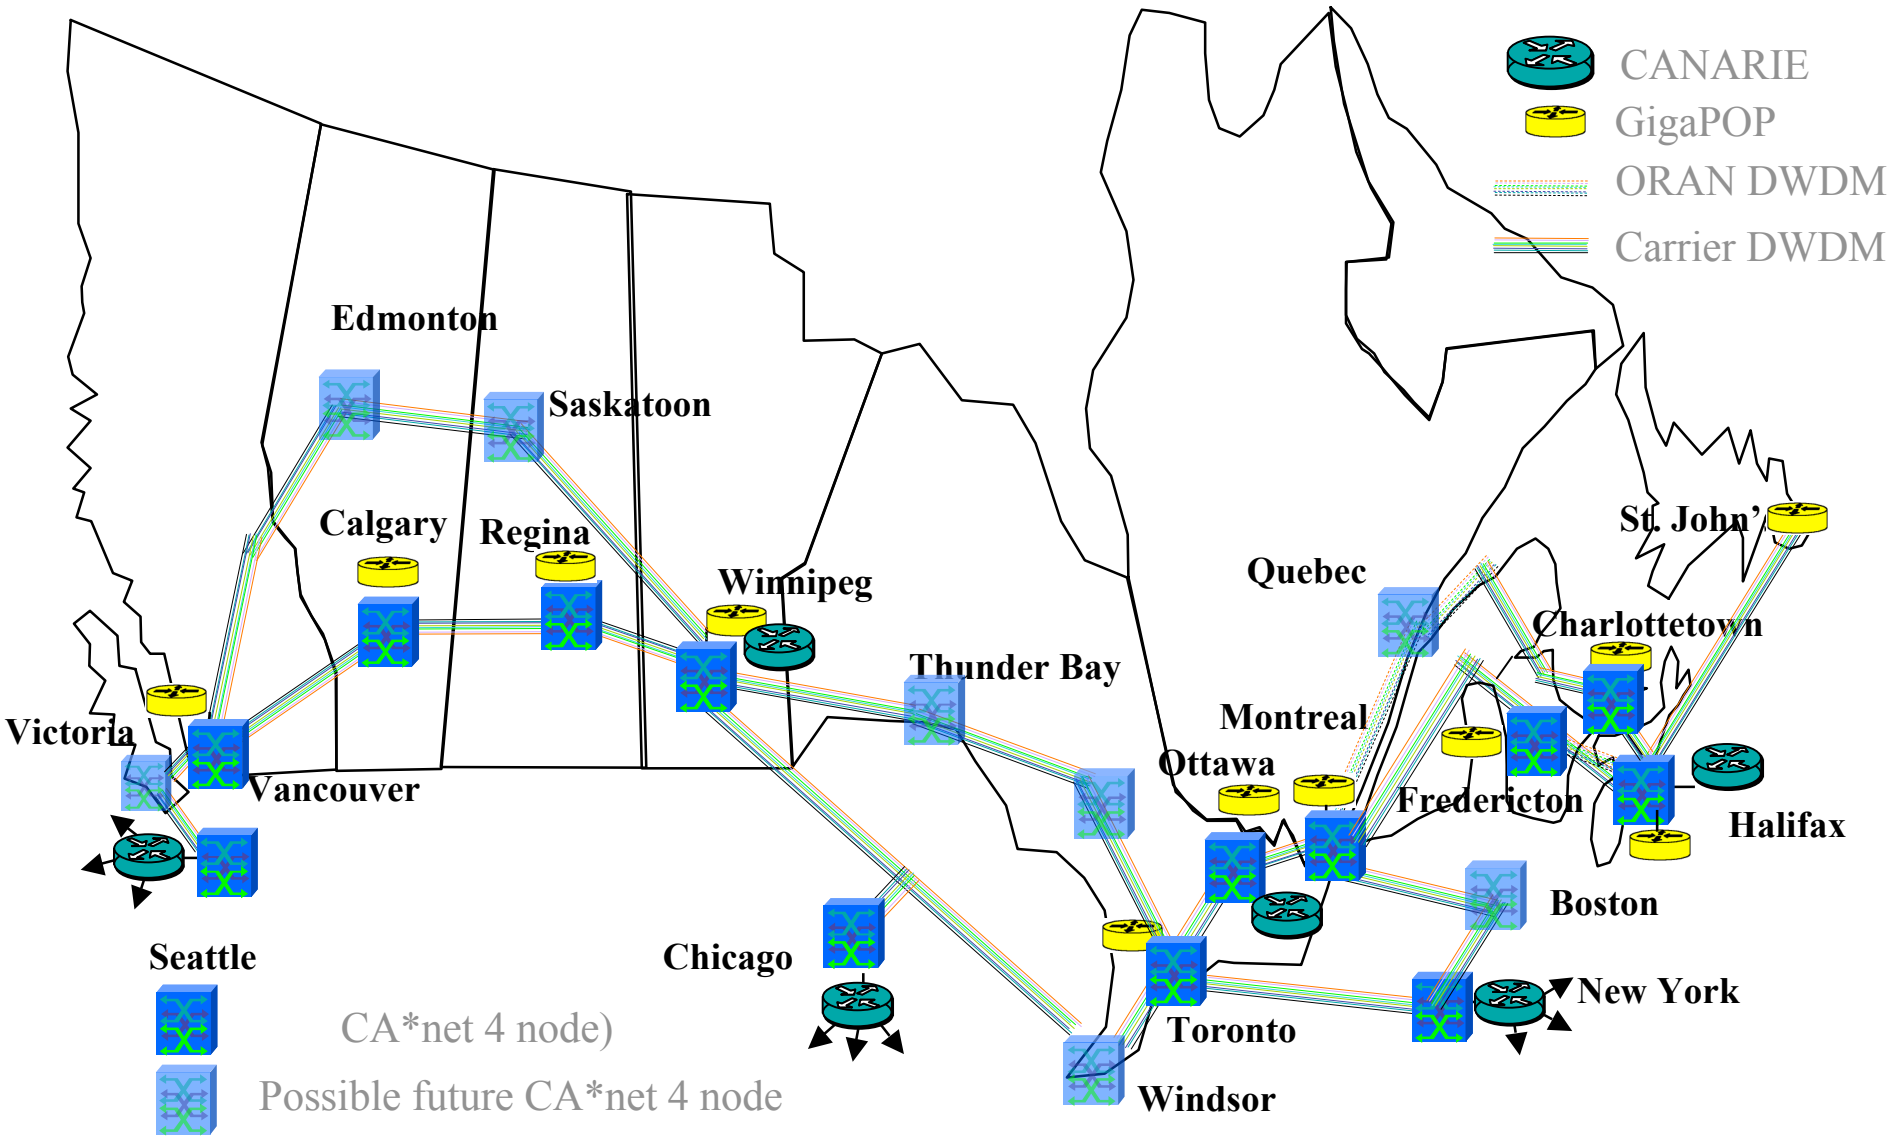

UltraLight Network: **PHASE III**

- Move into production
- Optical switching fully enabled amongst primary sites
- Integrated international infrastructure

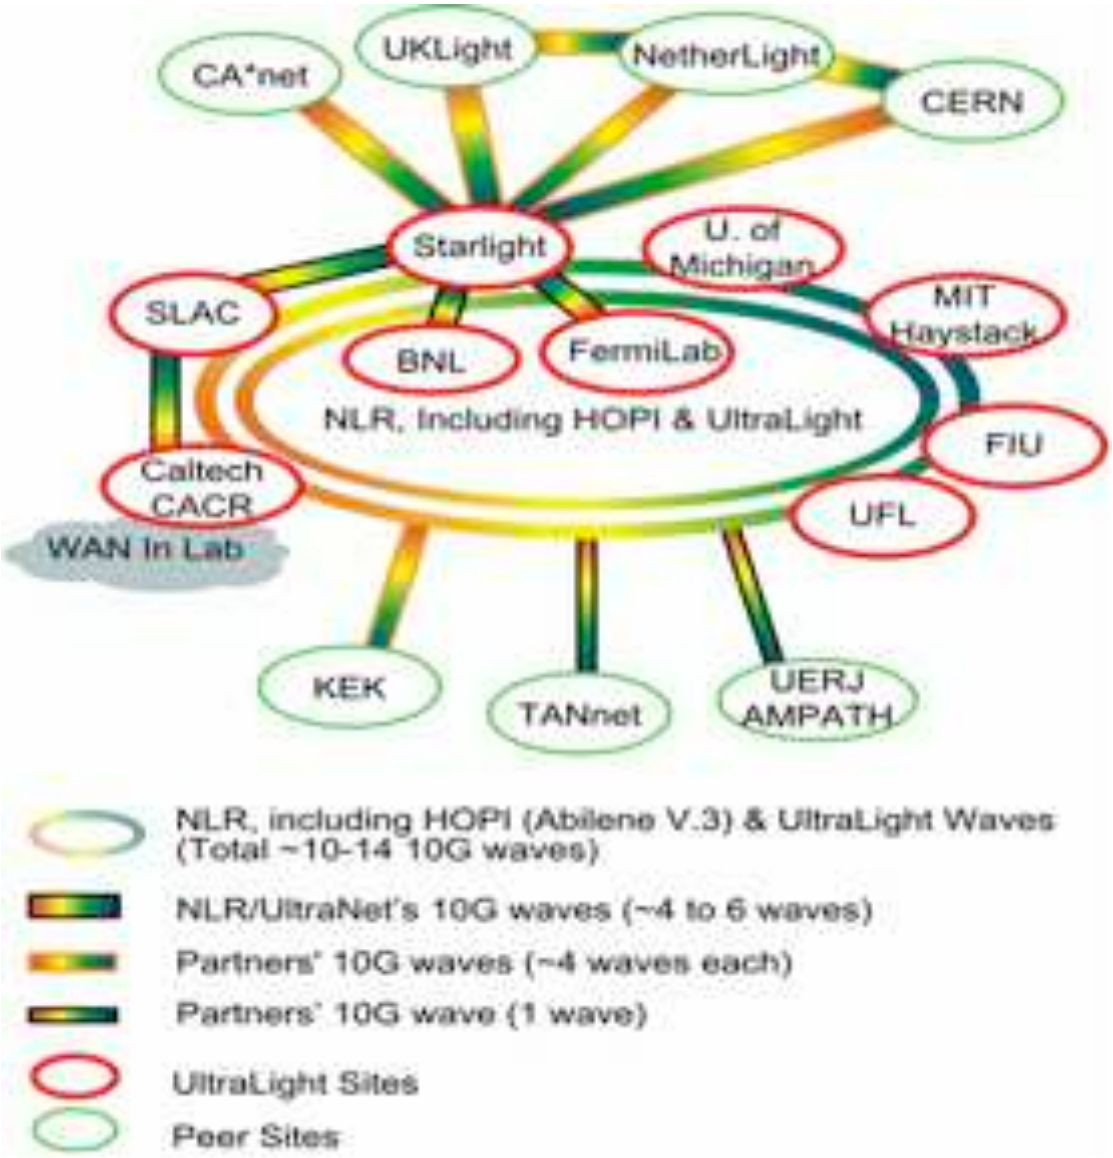

Important notes

- Dark fiber based Multi Lambda infrastructures are emerging
- Networks are build ON TOP OF THEM !
- Need for a hybrid (L1, L2, L3) networks to support GRIDs, communities and science disciplines
- Need for open policy exchanges
- Need for a persistent photonic infrastructure
- Need for a European FiberCo
- Participation in the worldwide infrastructure

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!!!**

Last slide

Thanks