Global Lambda Networking with GLIF

Serving Demanding Applications Worldwide

Lars Fischer

SUNET TREFpunkt

15 October 2008

The Rise of e-Science

LHC: A large, shared scientific instrument with distributed (global) use and vast resource requirements.

Not all users are equal

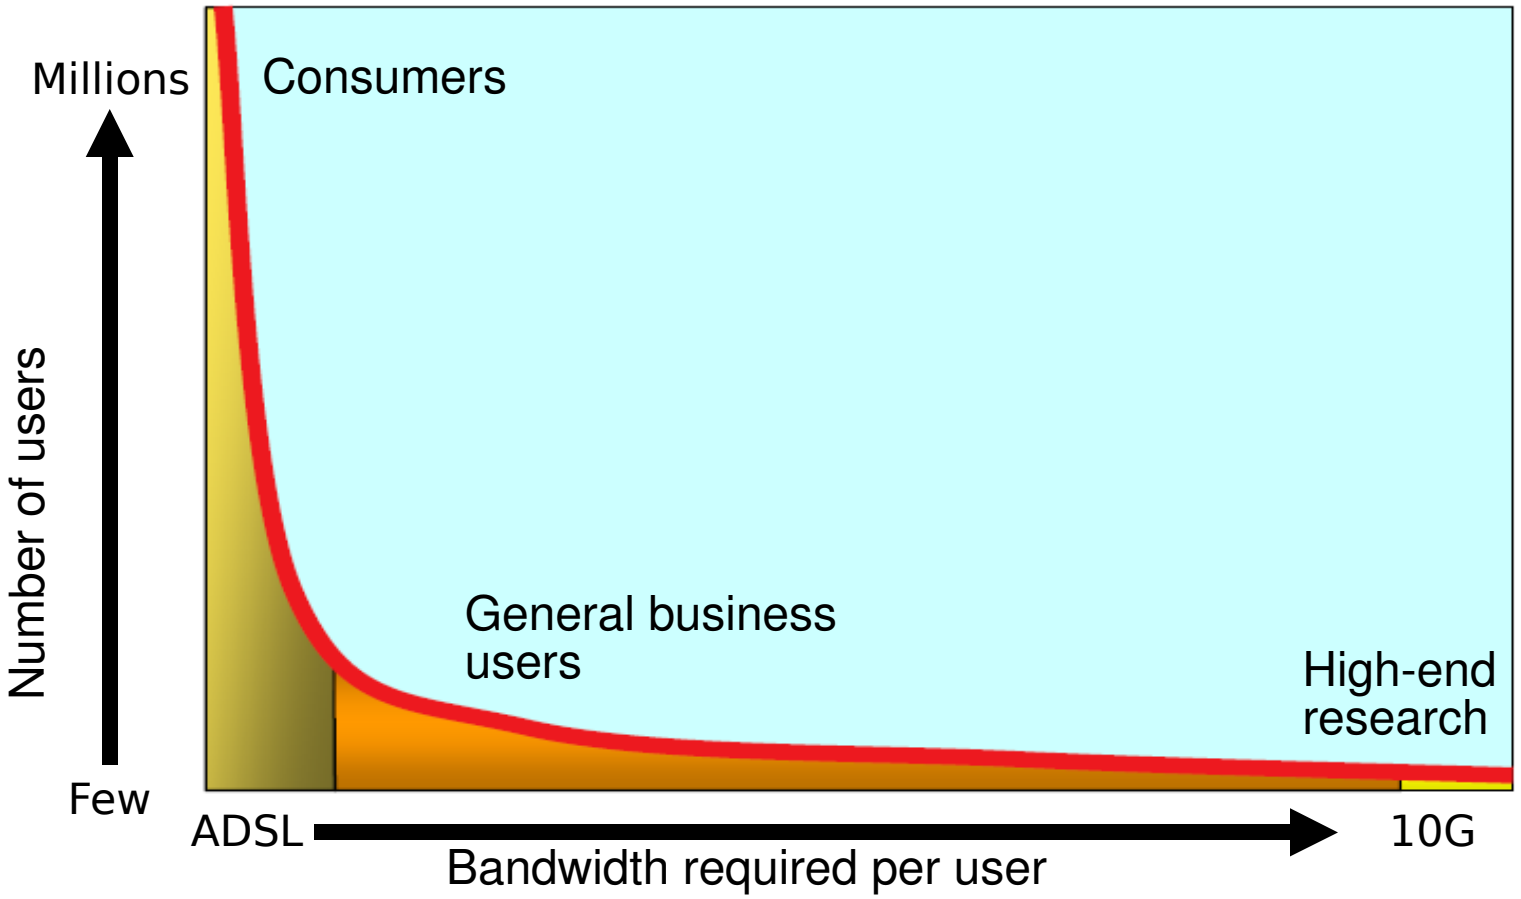

SERENATE Study Final Report, 2003, *Cees De Laat, David Williams* et. al.

Cost Concerns

- Growth: demands, capacity, cost
  - In particular cost & power for routing
- Observations
  - Costs of optical port is 10% of switching port is 10% of router port with same capacity
  - Many high-volume traffic flows are community or project specific, with few & known sites
- Idea
  - Dedicated circuits for massive users
  - Go as low (in the protocol stack) as possible
  - Give each packet in the network the service it needs, but no more

Lambda Networking

- Offload heavy users onto (simple) private circuit-switched networks
  - High performance for demanding users
  - Better service for average users on shared IP
  - Save capacity and cost on routed networks
  - As a bonus, private networks provide predictable network properties (i.e., true QoS)
- Challenges
  - Capability to set up, manage lambdas
  - Coordination for multi-domain lambdas
  - International connectivity
- Early experiments, collaborations ~2000-2001

NORDUnet
Nordic infrastructure for Research & Education

2003 Lambda Grid Workshop

Founding of GLIF

- August 2003: third Lambda Workshop, in Reykjavik, hosted by NORDUnet and attached to the NORDUnet 2003 Conference
- 33 participants from Europe, Asia and North America agree to form forum for collaboration, coordination, and promotion for lambda networking:

  GLIF: Global Lambda Integrated Facility

GLIF vision

Linking the World with Light

- Facilitate large-scale application through sharing of network resources
- Optical networks are the central architectural element in support of the most demanding e-science applications.
- Hybrid networks are the next-generation networks
  - Dedicated lightpaths for guaranteed high-speed few-to-few usage.
  - High-speed packet-switched, routed, shared Internet for regular many-to-many usage
- Research should not have any geographical boundaries.

What is GLIF

- GLIF is a collaborative initiative among worldwide NREN's and institutions.
  - GLIF enables and coordinates collaborative splicing of (federated) networks from individual resources
  - “it's not a network”
- GLIF is a world-scale lambda-based laboratory to facilitate application and middleware development
  - Enable experiments and collaboration
  - Promote middleware development
  - Facilitate novel applications
- Forum for discussion & presentation of progress on major challenges
  - Example: Control plane development

GLIF Organization

- Open to any organization sharing a vision of optical interconnection of facilities, who voluntarily contribute network resources (e.g. equipment, lambdas) or actively participate in relevant activities
- Lightweight governance structure.
  GLIF is managed as a cooperative activity
- GLIF has participants, not members
- GLIF is an open community
- Secretariat functions provided by TERENA with contributions from sponsors

Activities in GLIF

- Resource coordination
  - Open lightpath exchanges - like internet exchange points for IP
  - Glueing together infrastructure
  - Coordination lists and meeting
- Short-lived Task Forces
  - Develop procedures, document best practice
  - Propose standards (but not a standards body)
  - Tools development & trials
  - Example: PerfSONAR taskforce
- Workshops every 6 months

1st Generation (2004)

- First NORDUnet lambda platform
- Services
  - 10G / 2.5 G shared IP
  - 1G (subrate) Ethernet lightpath service
- Design
  - Based entirely on leased lines
  - Cisco ONS 15454 SONET lambda service platform, OC48 carrier
  - Lambda service platform and IP service platform separate infrastructures (w/ shared PoP's)
  - Leased OC48 to NetherLight

(e-)VLBI

- Very long baseline virtual radiotelescope, created by correlating outputs of multiple actual radiotelescopes (VLBI)

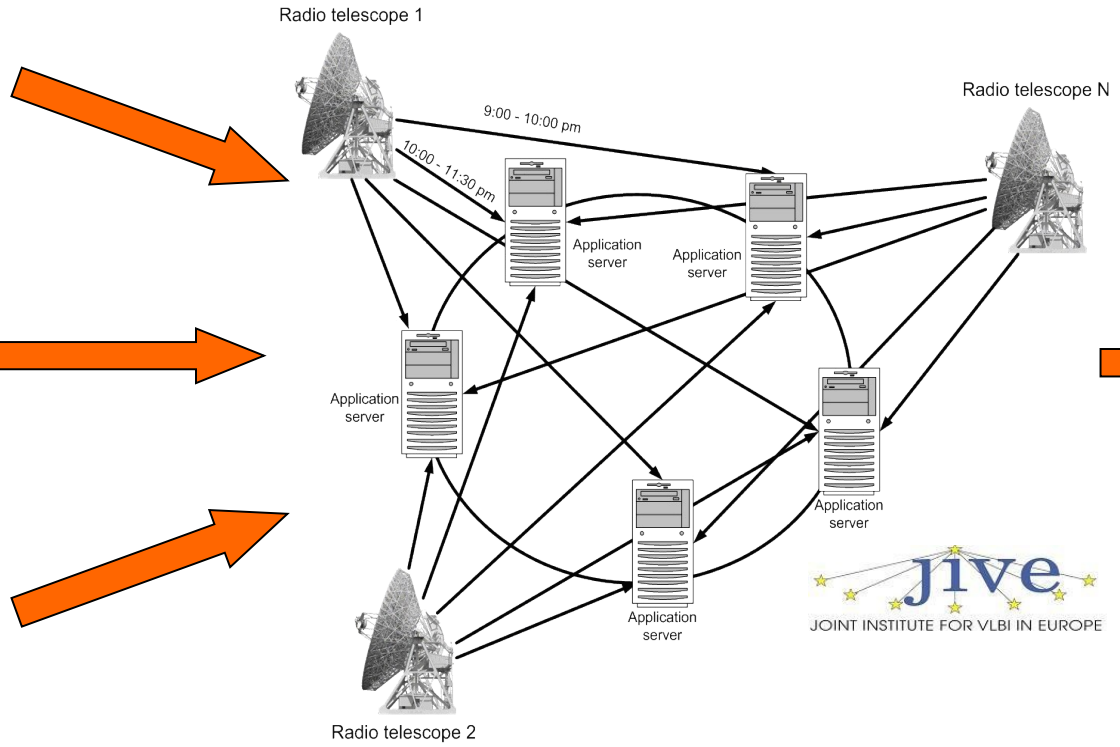

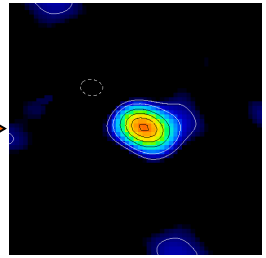

iGrid 2005 e-VLBI demo

GLIF terminology

- **GOLE**: GLIF open lightpath exchange, comprised of one or more network devices performing lightpath switching.
- **lambda**: a high-capacity (optical) circuit which terminates on a GOLE
- **lightpath** [service]: a high capacity circuit or QoS-supported virtual circuit (or the concatenation of several of these to form an end-end lightpath). A lightpath can take up a portion or all of the capacity of a lambda.
- **GLIF network resources**: GOLEs and lambdas

A GOLE

- The key infrastructure of the GLIF collaboration
- Open Exchange Points for circuit-switched networks – in the tradition of IP Internet Exchanges
- Everyone can bring a circuit and exchange traffic with everyone else
- A GOLE is typically a SDH/SONET switch matrix, possibly with (next-generation) Ethernet and photonic switching
- Each GOLE will have lambdas connecting it to other GOLEs

Topology of a GOLE

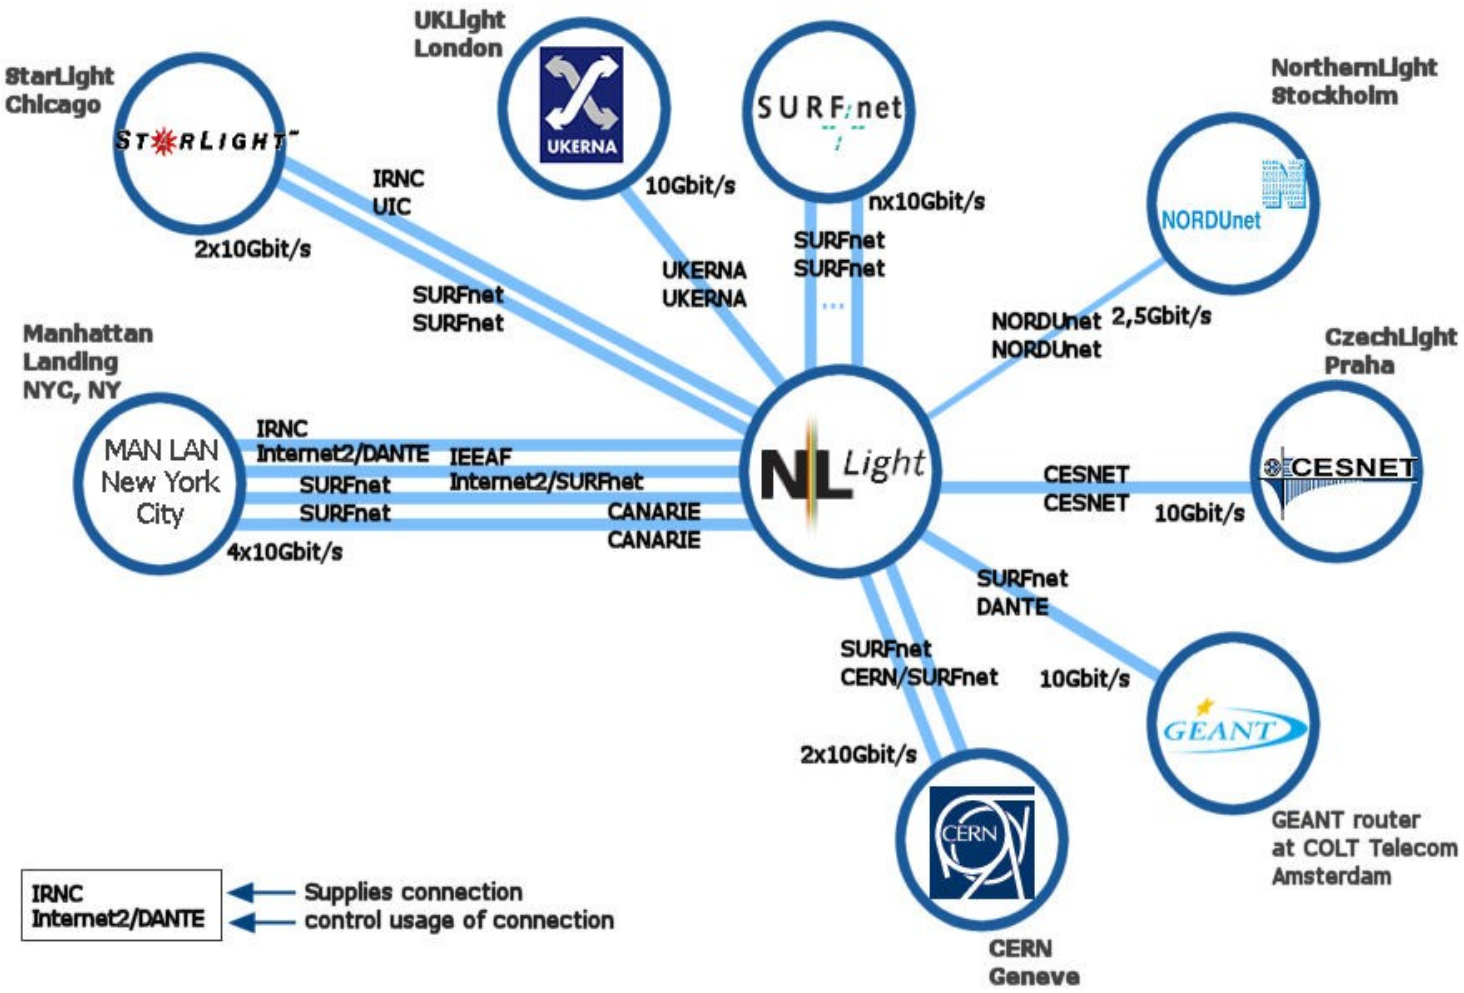

An Example GOLE

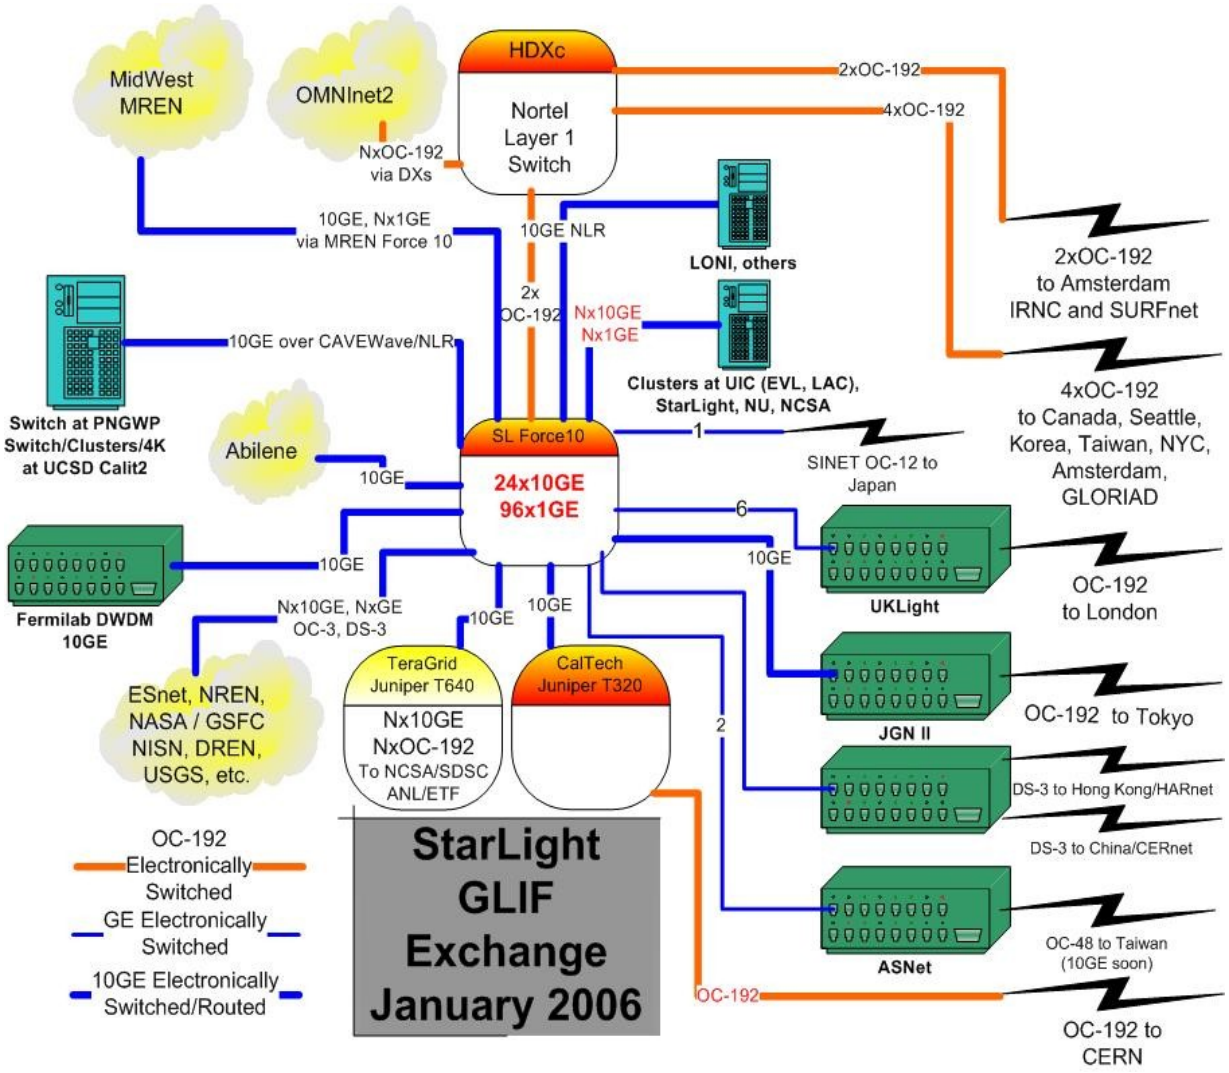

GOLE to GOLE networking

- In GLIF, end-to-end connectivity is constructed by piecing together circuits from GOLE to GOLE, eventually connecting end-sites
- There are GOLEs all over the globe
- GOLEs are typically operated by NRENs
- Coordination through GOLE operator meetings, mailing lists, personal network
- Resource sharing – no inter-GOLE payment system established

NORDUnet 2nd Generation: Hybrid (2007)

Nordic Infrastructure for Research & Education

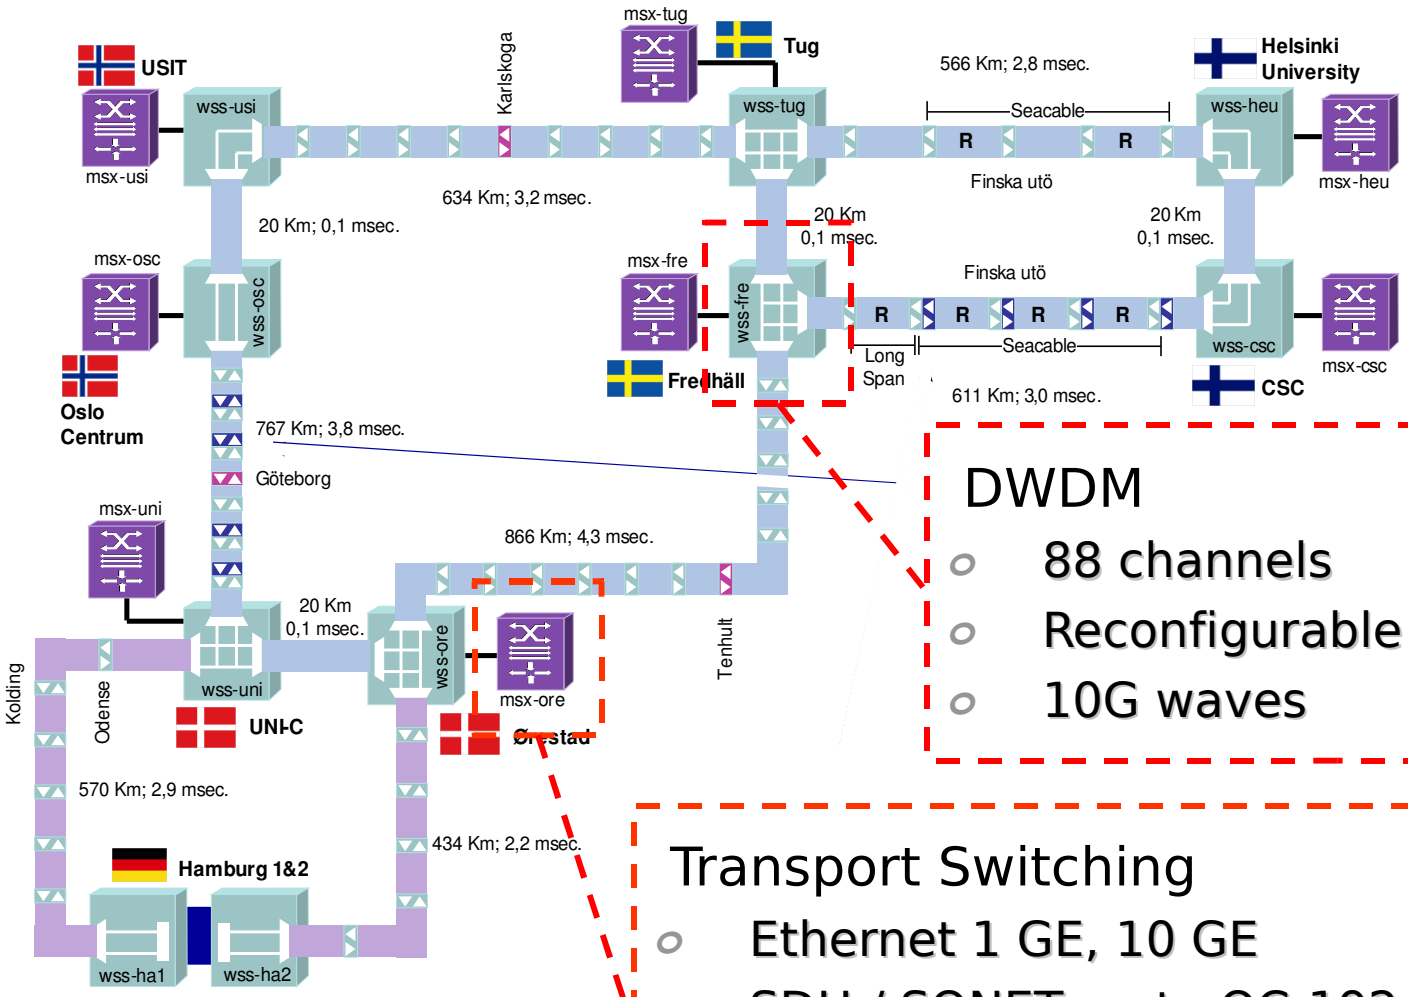

DWDM

- 88 channels
- Reconfigurable and tunable
- 10G waves

Transport Switching

- Ethernet 1 GE, 10 GE
- SDH / SONET up to OC-192
- Subrate services
- Bandwidth sharing, dynamic capacity allocation

NORDUnet
Nordic infrastructure for Research & Education

Lambdas across the world

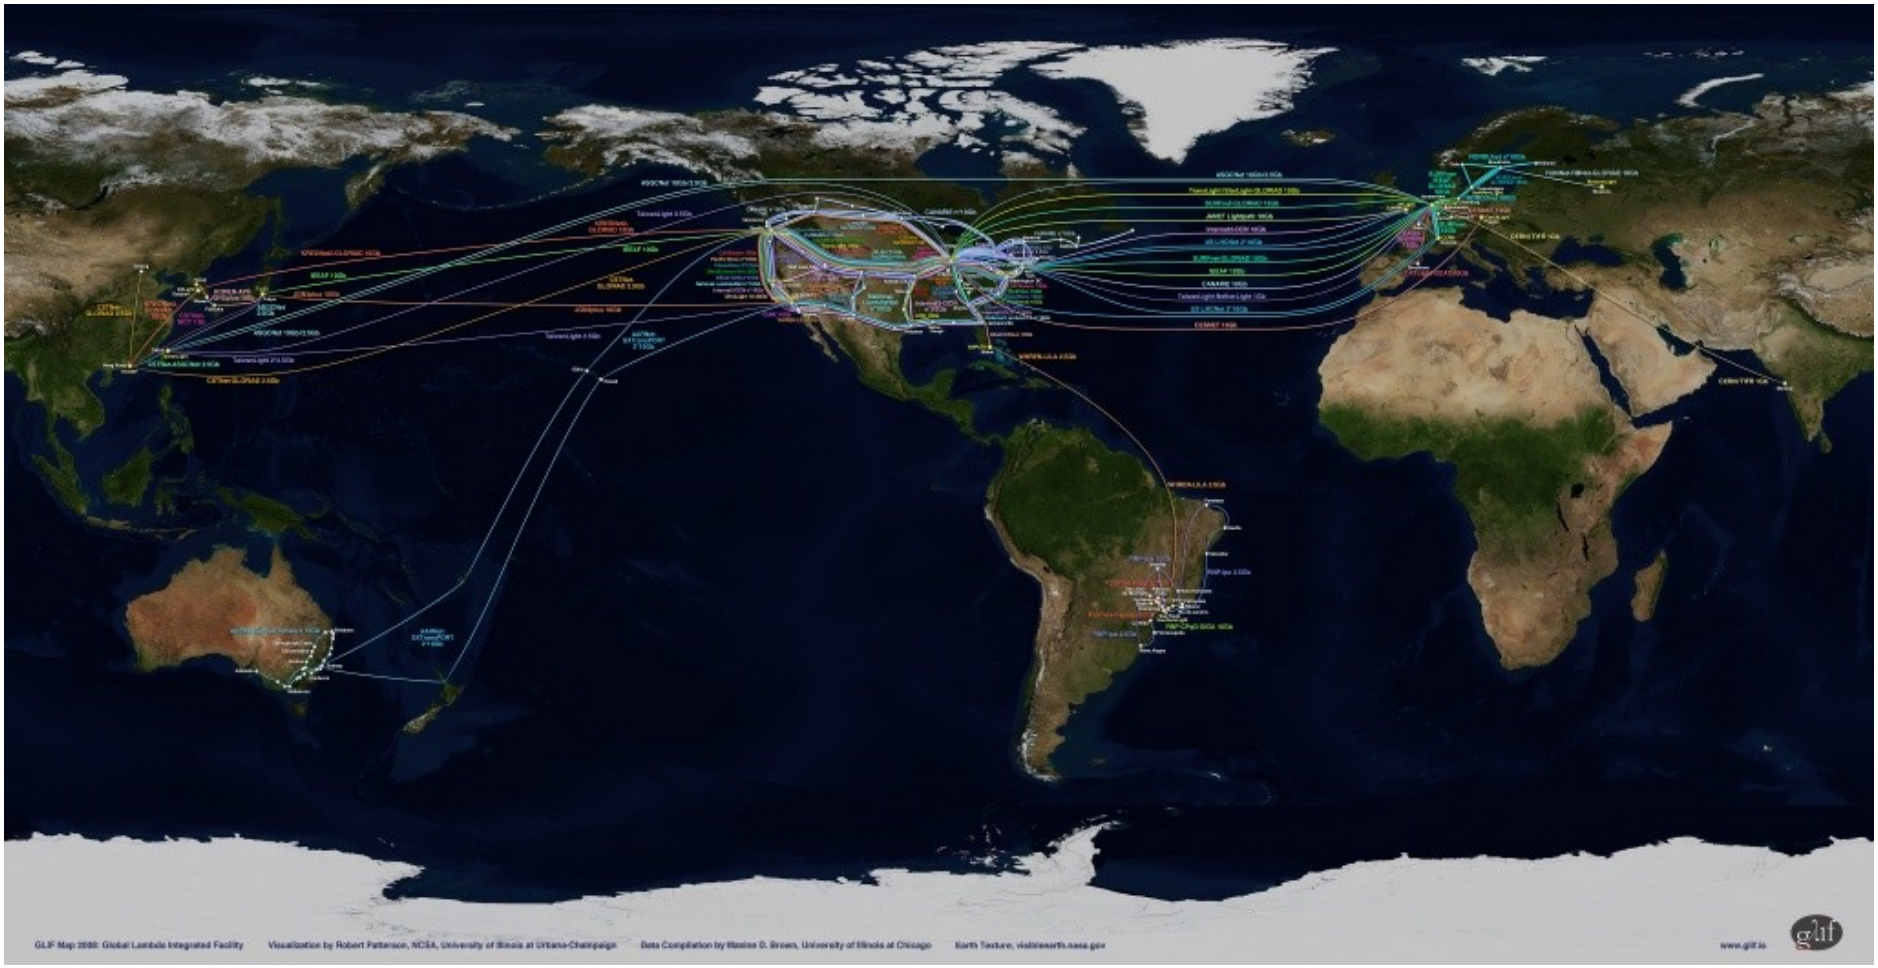

A European view

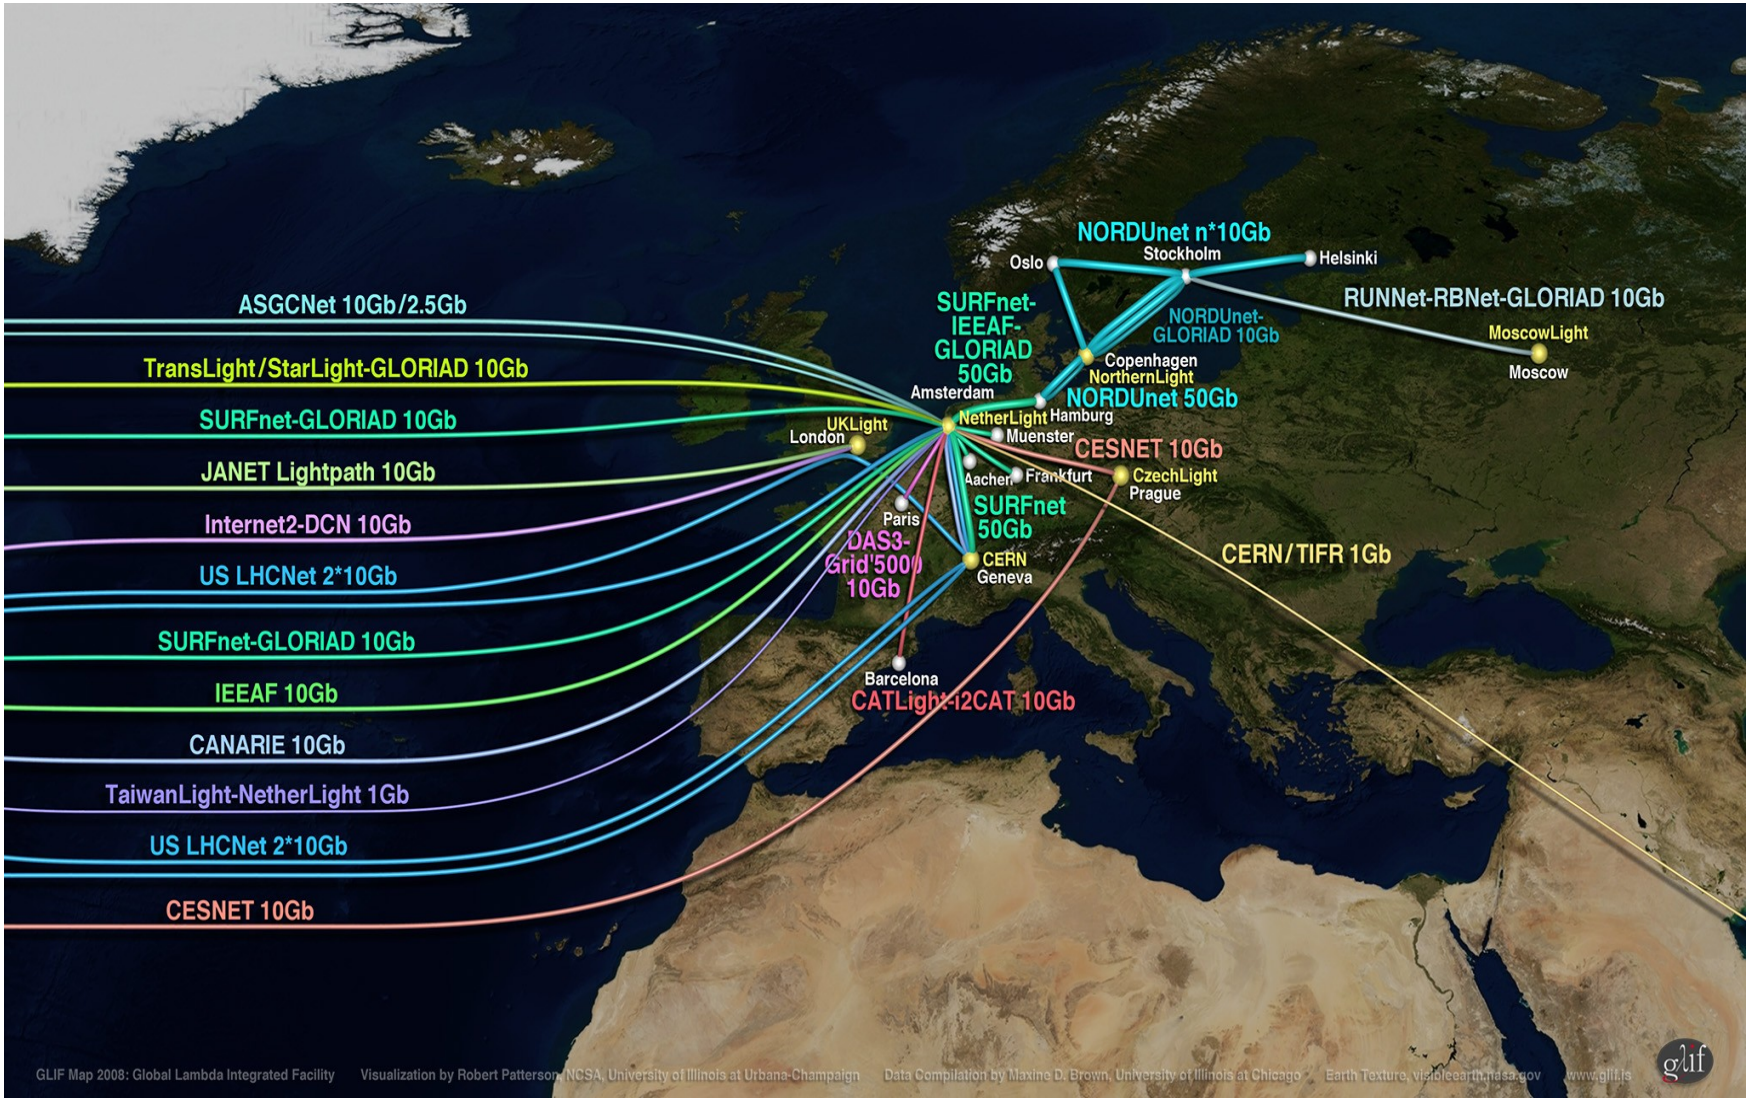

Live 4k video across the globe

- Collaboration between KTH and Keio University, Japan
- Challenge
  - Stream the Kyoto price ceremony, live, from Kyoto to Stockholm, using 4k video (Nov. 2007)
  - Use Sony projector, NTT encoder / decoder
- 4k video – cinema quality
  - 4096 x 2160 pixels – 4x HDTV 1080p
  - Uncompressed 4k = 8.6 Gbps
- 10G circuit provisioned in two weeks
  - GLIF resources, GOLE-to-GOLE
  - Across Pacific, North America, Atlantic, N. Europe

Korea – Norway Medical

- Medical Media HD Live Transmission
  - From Seoul, Korea to Trondheim, Norway
  - Test technologies for HD live medical surgery transmission
  - Viewing laparoscopy surgery in High Definition Video enable doctors to collaborate, learn, and see details normally not available
- Video Streaming
  - 800 Mbps, low jitter required
  - Dedicated 1 GE link Norway – Korea provide by collaboration in GLIF

The Network

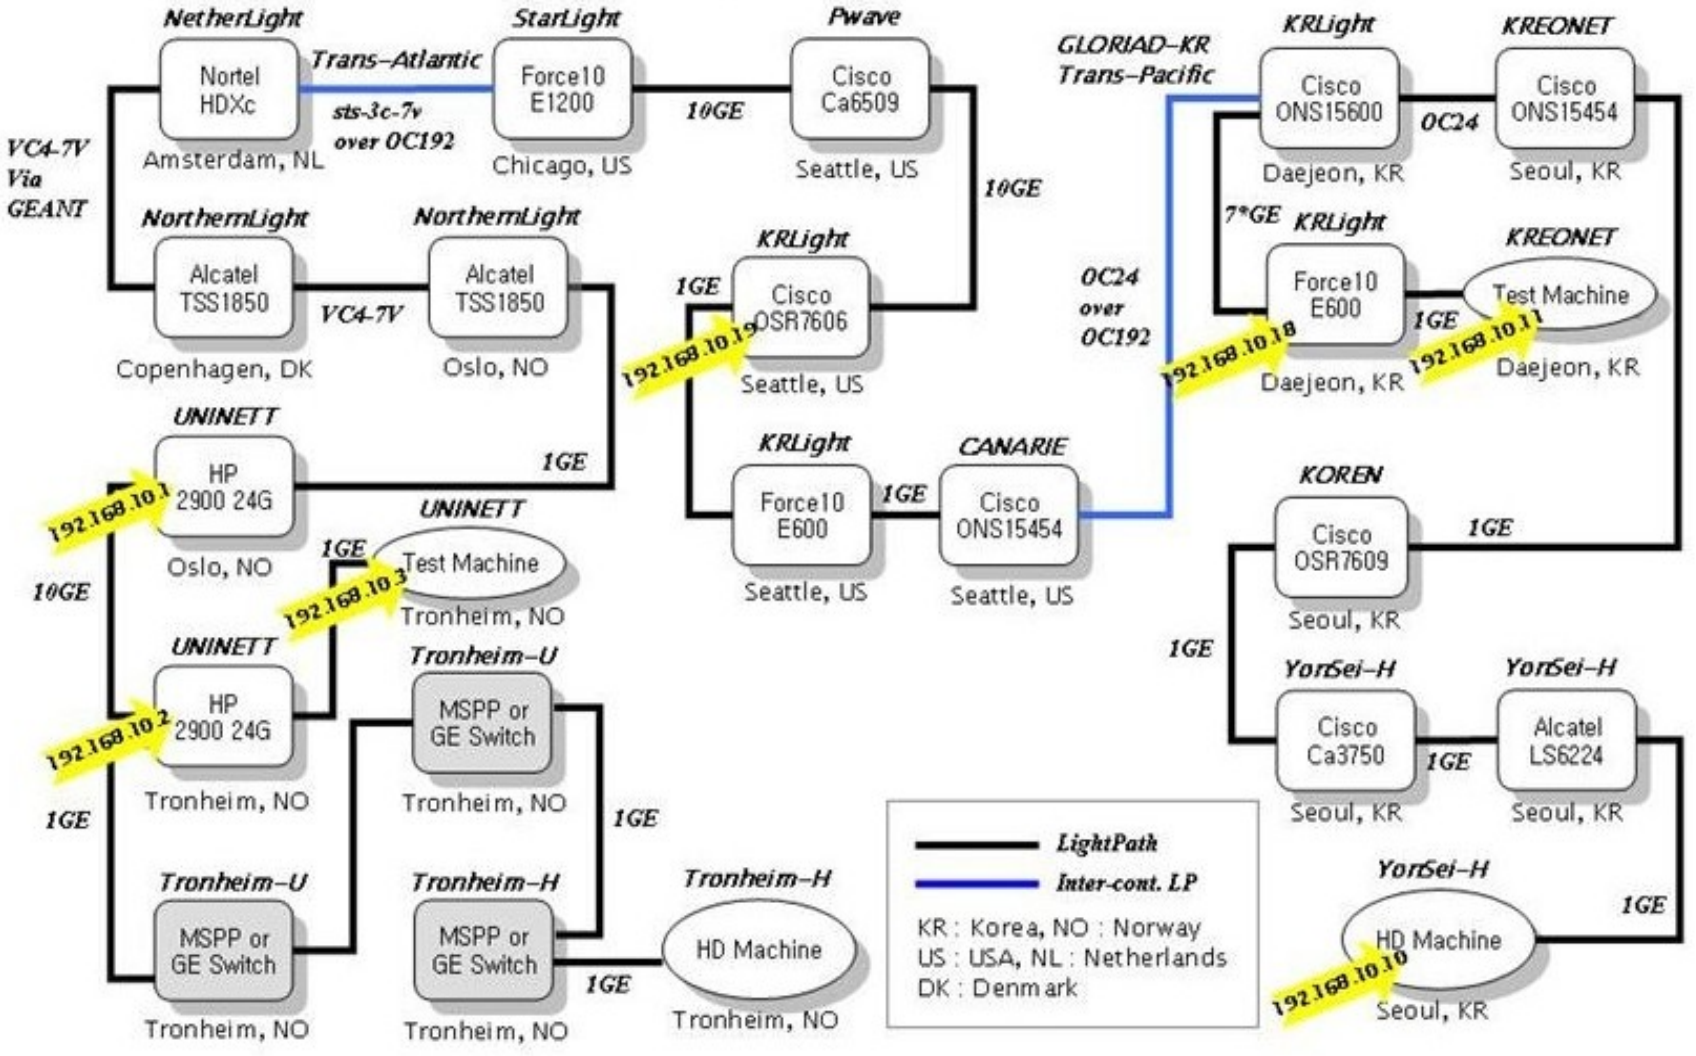
NetherLight
Nortel HDXc
Amsterdam, NL
Trans-Atlantic
sts-3c-7v over OC192
StarLight
Force10 E1200
Chicago, US
Pwave
Cisco Ca6509
Seattle, US
10GE
VC4-7V Via GEANT
NorthernLight
Alcatel TSS1850
Copenhagen, DK
VC4-7V
NorthernLight
Alcatel TSS1850
Oslo, NO
KRLight
Cisco OSR7606
Seattle, US
1GE
192.168.10.19
UNINETT
HP 2900 24G
Oslo, NO
192.168.10.1
1GE
10GE
UNINETT
Test Machine
Tronheim, NO
192.168.10.3
UNINETT
HP 2900 24G
Tronheim, NO
192.168.10.2
1GE
Tronheim-U
MSPP or GE Switch
Tronheim, NO
1GE
Tronheim-U
MSPP or GE Switch
Tronheim, NO
Tronheim-H
MSPP or GE Switch
Tronheim, NO
1GE
Tronheim-H
HD Machine
Tronheim, NO
KRLight
Force10 E600
Seattle, US
1GE
CANARIE
Cisco ONS15454
Seattle, US
GLORIAD-KR Trans-Pacific
OC24 over OC192
KRLight
Cisco ONS15600
Daejeon, KR
KREONET
Cisco ONS15454
Seoul, KR
OC24
7*GE
KRLight
Force10 E600
Daejeon, KR
192.168.10.18
KREONET
Test Machine
Daejeon, KR
1GE
192.168.10.11
KOREN
Cisco OSR7609
Seoul, KR
1GE
1GE
YonSei-H
Cisco Ca3750
Seoul, KR
1GE
YonSei-H
Alcatel LS6224
Seoul, KR
YonSei-H
HD Machine
Seoul, KR
192.168.10.10
1GE
LightPath
Inter-cont. LP
KR : Korea, NO : Norway
US : USA, NL : Netherlands
DK : Denmark

The Results

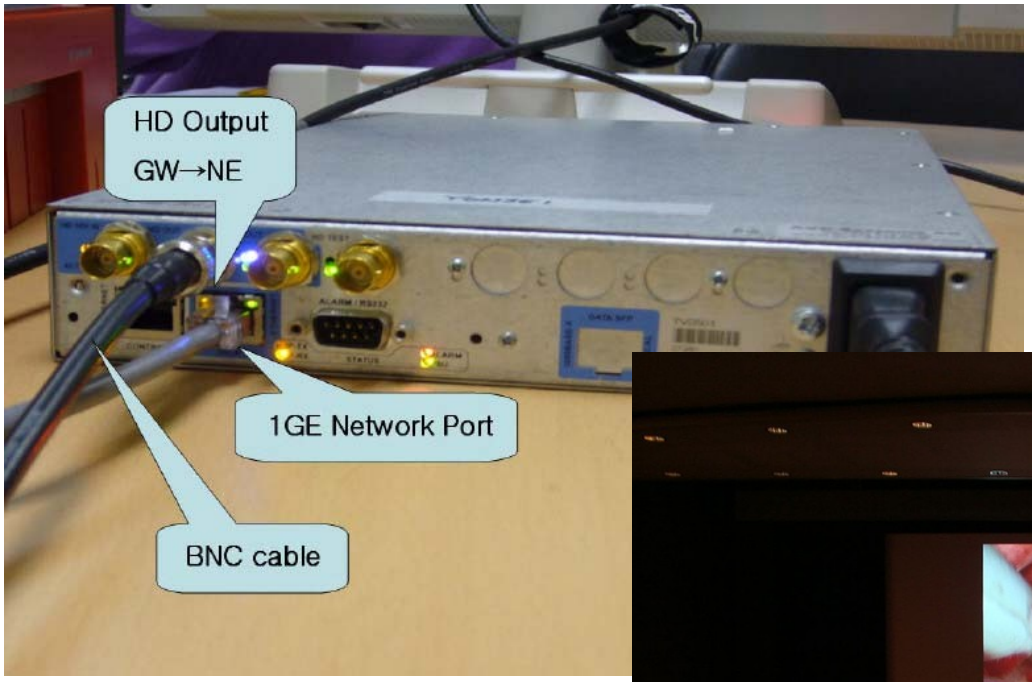

From Experiments to Production

- Demand of e-Science is growing
  - Initial experiments were successful - now users expect a service – predictable, ubiquitous
  - New advanced user communities are appearing, expecting permanent infrastructures with local and global reach
- Advanced infrastructure is taken for granted
  - Vast network capacity (as well as massive storage and computing resources) is expected by e-Science projects; massive projects depend on this infrastructure.
  - It's the way science is done – not an exciting experiment

e-VLBI

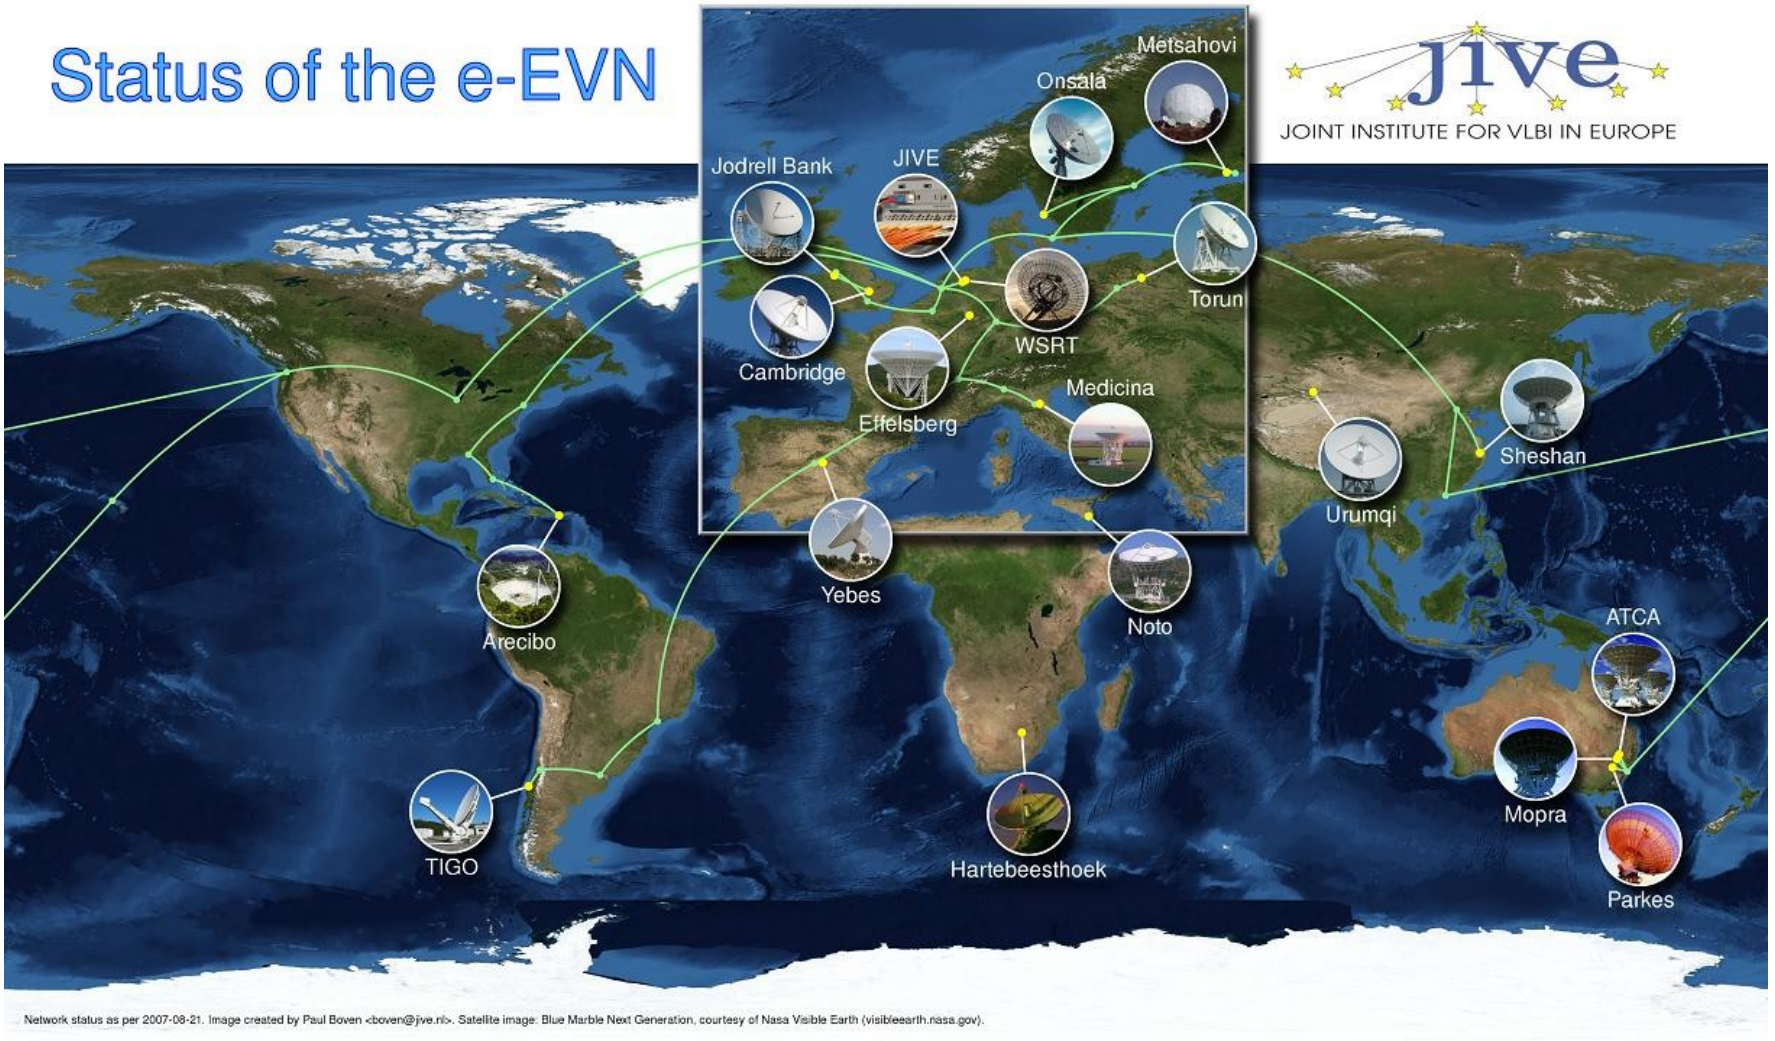
Status of the e-EVN
JIVE
JOINT INSTITUTE FOR VLBI IN EUROPE
Metsahovi
Onsala
Jodrell Bank
JIVE
Torun
WSRT
Cambridge
Effelsberg
Medicina
Sheshan
Urumqi
Yebes
Noto
Arecibo
ATCA
Mopra
TIGO
Hartebeesthoek
Parkes
Network status as per 2007-08-21. Image created by Paul Boven <boven@jive.nl>. Satellite image: Blue Marble Next Generation, courtesy of Nasa Visible Earth (visibleearth.nasa.gov).

NORDUnet
Nordic infrastructure for Research & Education

Nordic Astronomy OPN

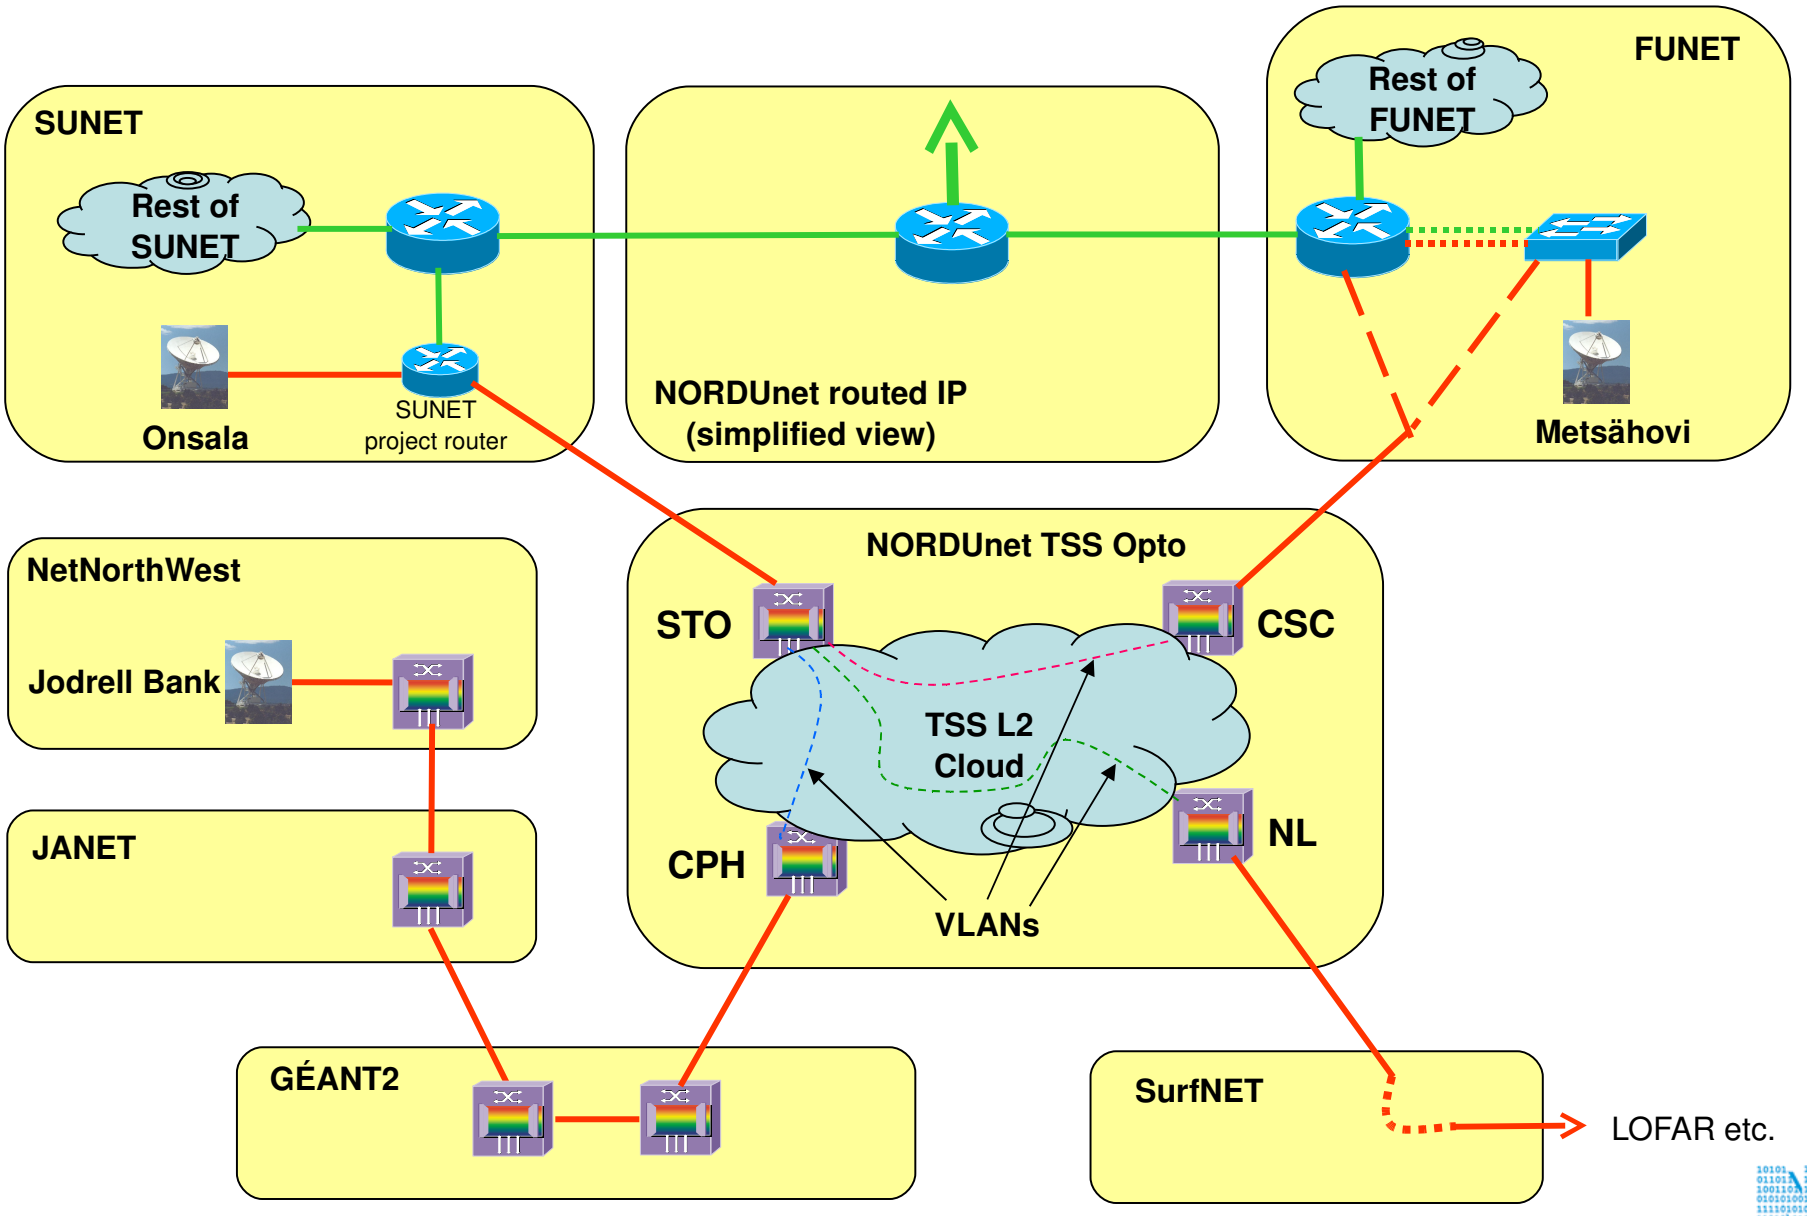

NDGF Tier-1 for WLCG

- A virtual HPC centre, made of resources from major Nordic HPC sites
- Resources (Storage and Computing) are scattered
- 10G Private Network + 10G connections to CERN, Amsterdam, Ljubljana

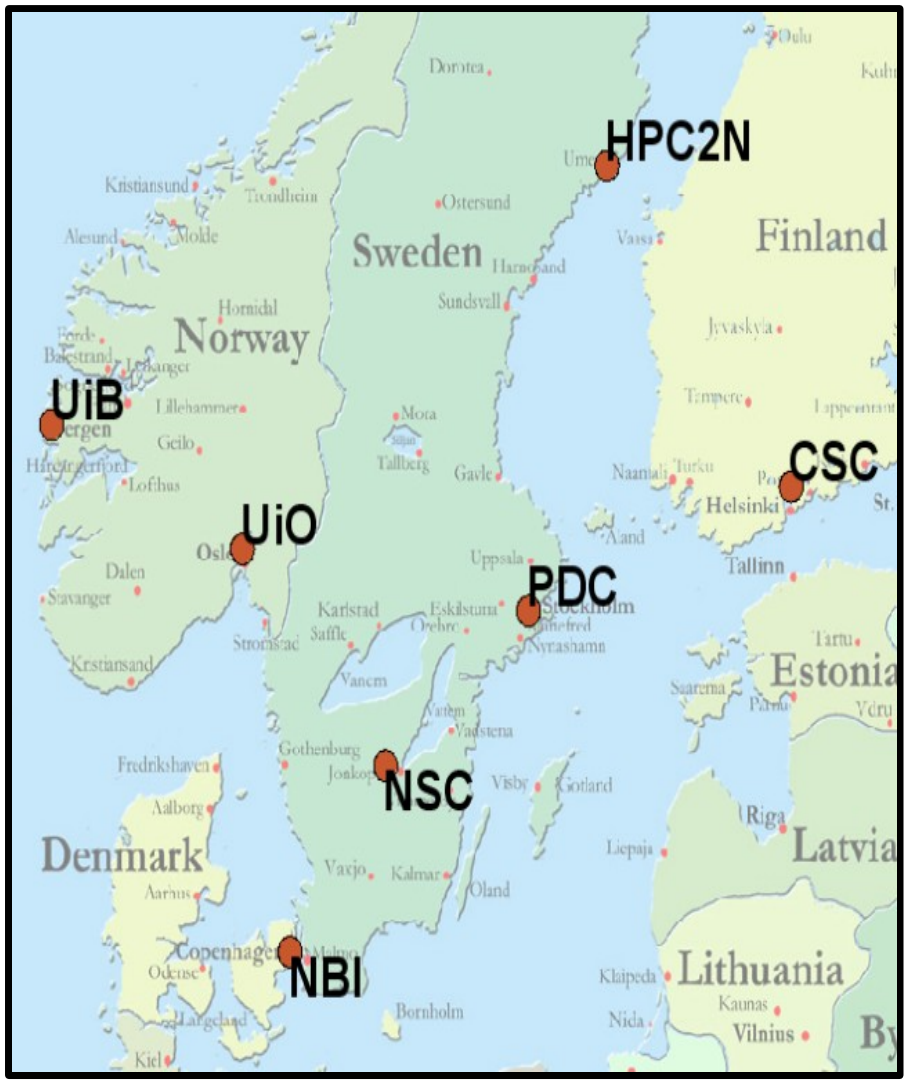

NORDUnet
Nordic infrastructure for Research & Education

NDGF Tier-1 Network

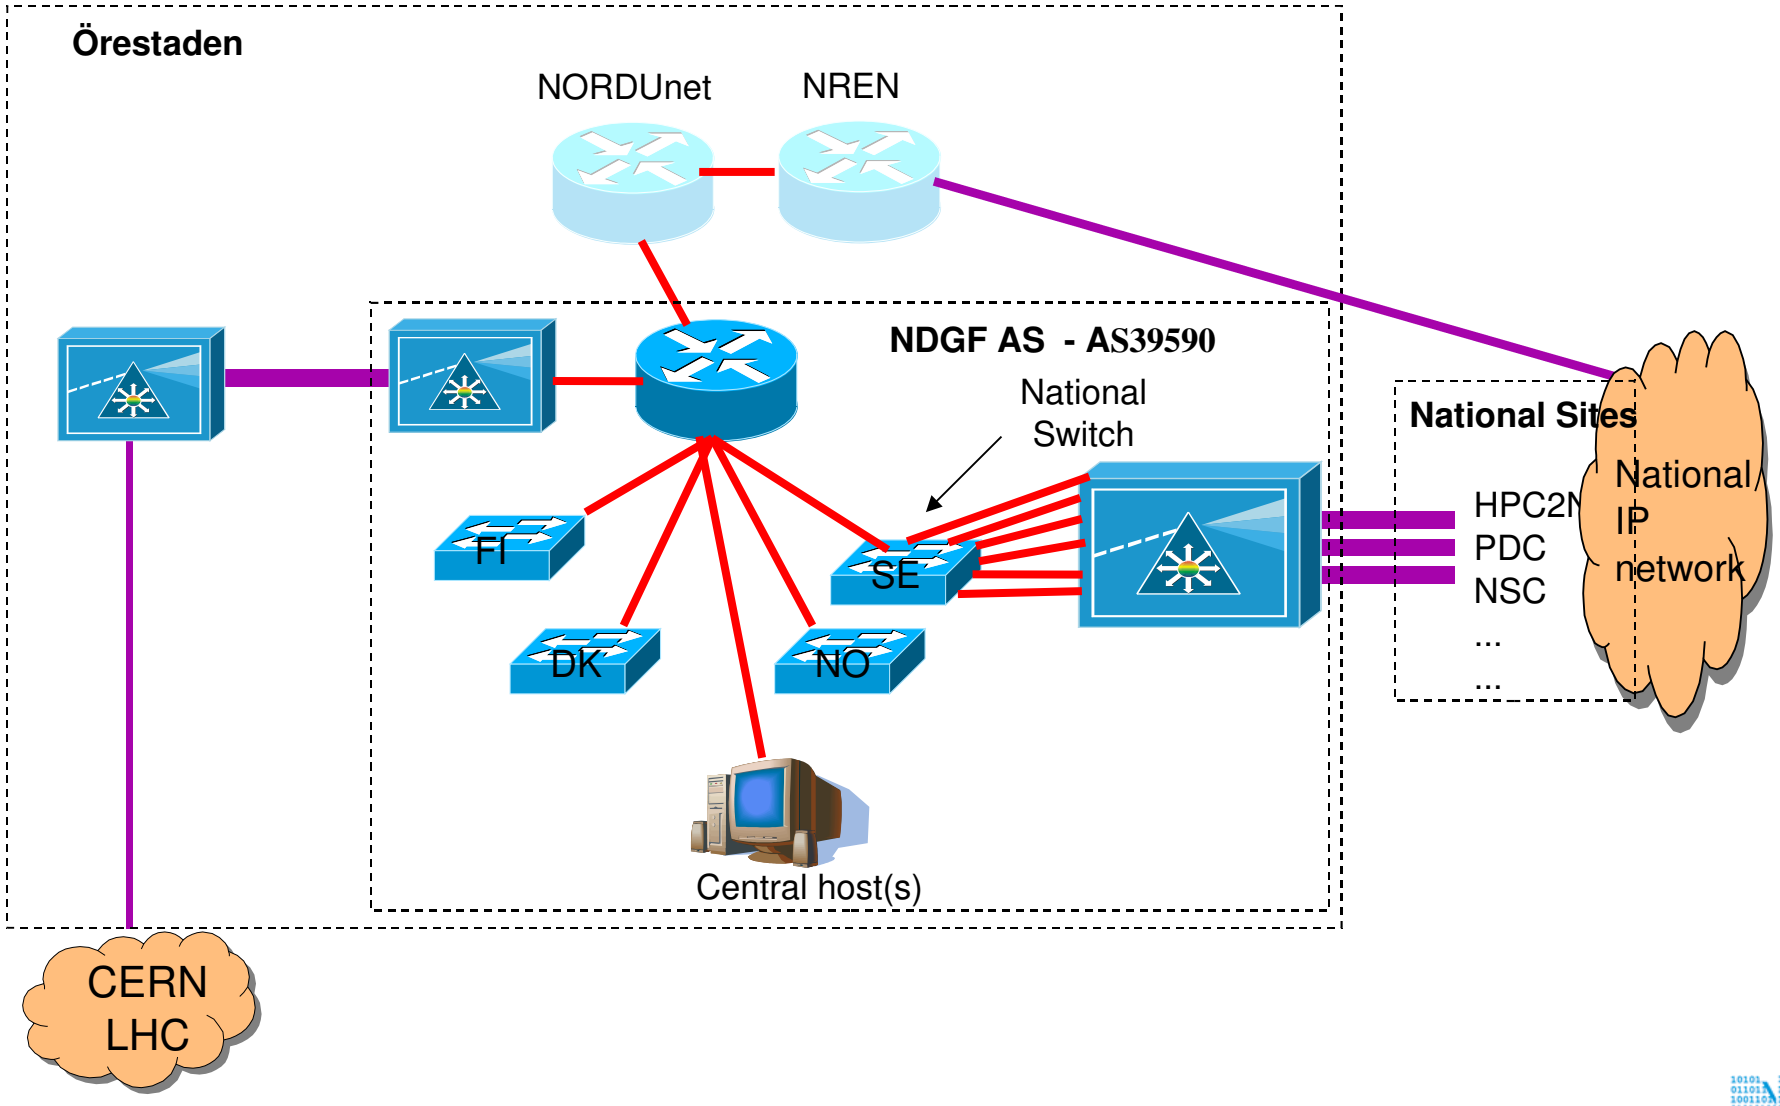

Optical Networking trends

- More bandwidth at lambda level
  - 10G now standard
  - 40G here, 100G coming soon
- More flexible handling of lambdas
  - From static configurations to tunable lasers and filters
  - Wavelength Selector Switches for flexible routing of entire lambdas in the optical domain
  - Alien waves – for inter-domain lambdas
- Dynamic configuration
  - Allows control plane systems to alter lambda routes “on the fly”

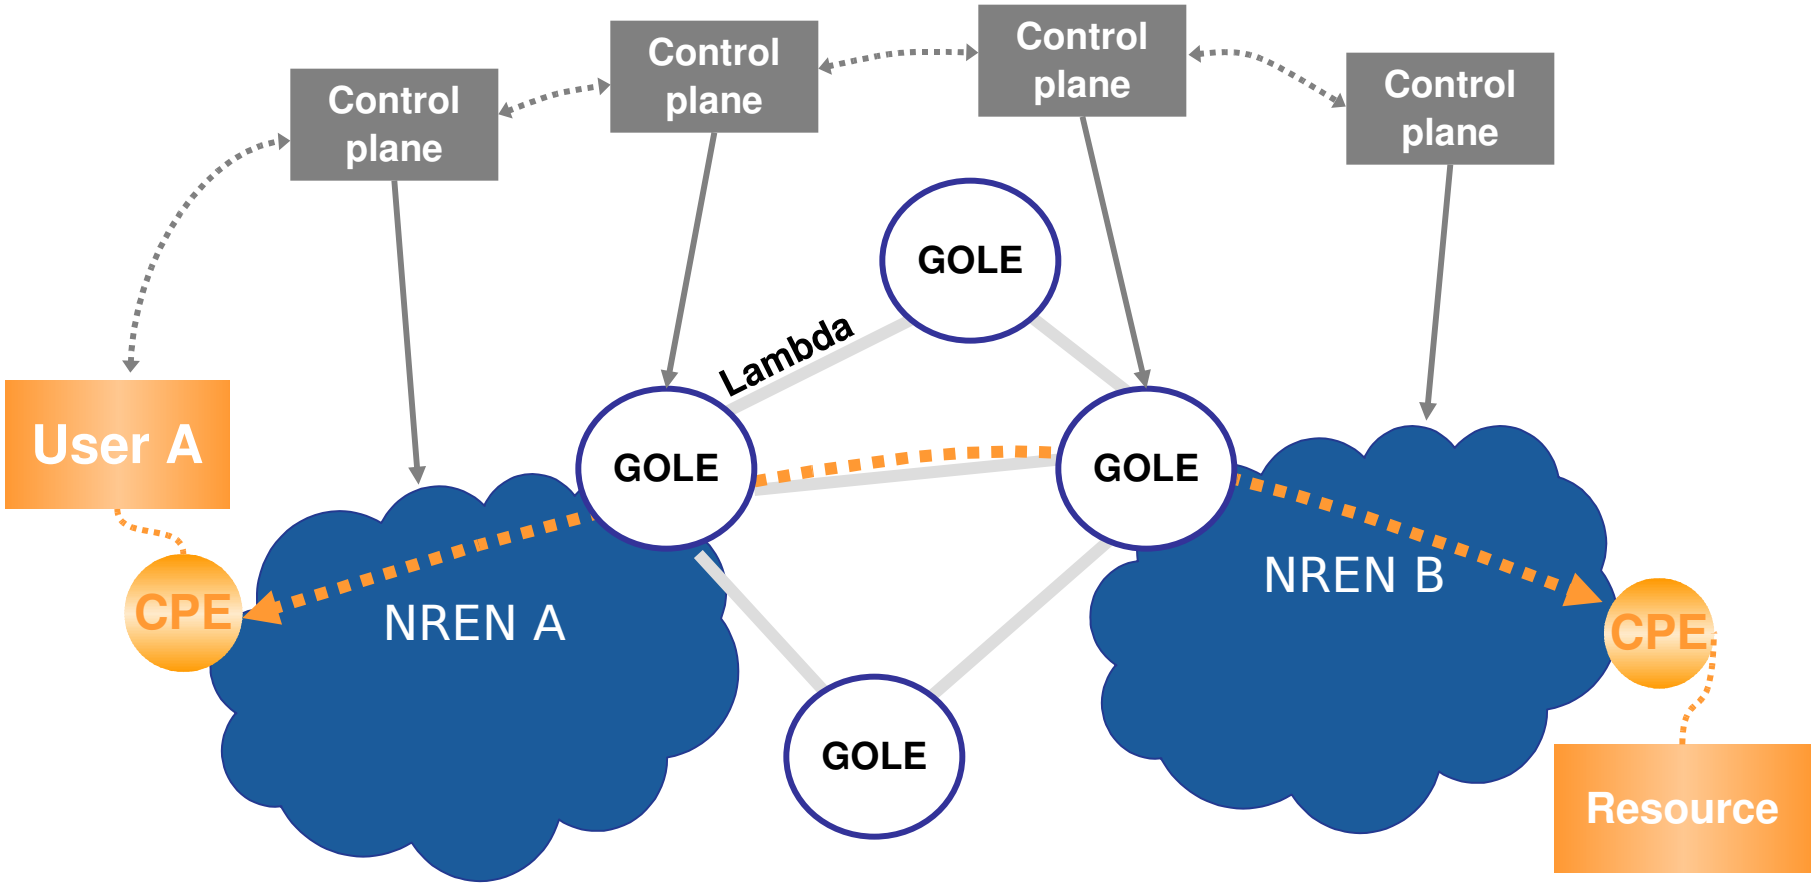
Control plane
Control plane
Control plane
Control plane
GOLE
Lambda
User A
GOLE
GOLE
CPE
NREN A
NREN B
CPE
GOLE
Resource

Multi-domain control planes

- Users need ubiquitous end-to-end lightpath connectivity over a multi-domain infrastructure
- How to glue things together
  - Cooperation between networks requires a shared control plane
  - Centralized models will not scale and will have role and authority issues
- Create a loose cooperation between domains
  - Standardized interfaces between domains
  - Create user friendly AAA features
  - Promote middleware developments (and embedding)

NORDUnet
Nordic Infrastructure for Research & Education

3rd Generation: WSS (2008+)

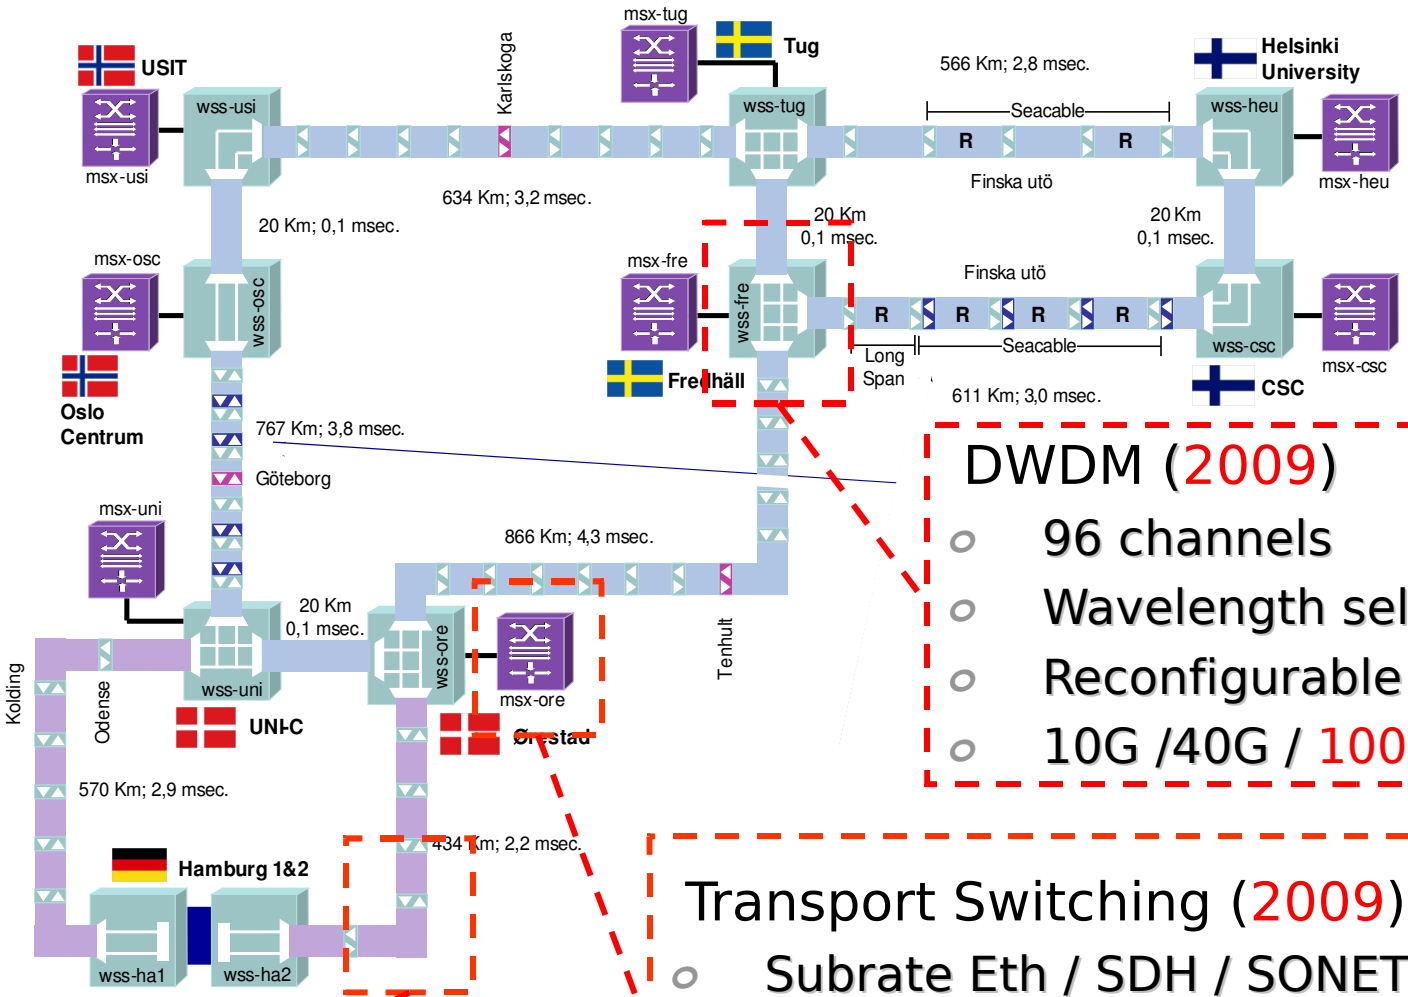

DWDM (2009)

- 96 channels
- Wavelength selector switches
- Reconfigurable and tunable
- 10G /40G / 100G waves

Transport Switching (2009)

- Subrate Eth / SDH / SONET - partial 10G / 40G / 100G waves
- Next-generation Ethernet
- T-MPLS

Cross-Border Fibre

ALU 1626 Light Manager

- Wavelength Selector Switches
  - LH/ULH up to 96 channels
  - Add/Drop 72 channels
- Automatic system alignment
- Enhanced functionality
  - Tuneable filters
  - 50 GHz filters
- Long Haul:
  - Fully Tuneable ROADM
- Fully C band tuneable interfaces over any interface
  - 10 G Universal Transponder
    - STM-64
    - 10GE WAN & LAN
  - 40G transponder
    - PSBT (50 GHz - SH)
    - qDPSK (50 GHz – LH/ULH)
- Embedded Optical Protection
- Raman Amplification

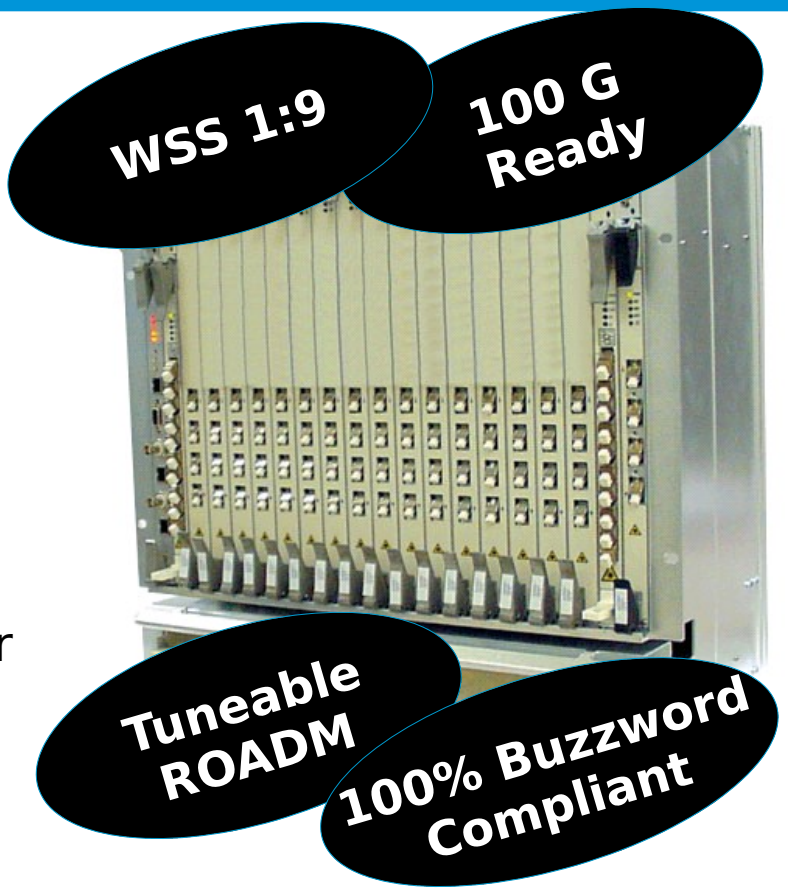

-
- Same modules for OADM and ILA
  - Shelf
  - Amplifier

NORDUnet
Nordic infrastructure for Research & Education

ALU 1850 TSS-320

SDH/SONET

- STM-1, 4, 16 & 64
- Cross-Connection
- Termination
- ETH Mapping over SDH
- Complete Scope of SDH/SONET Features

Ethernet

- 10GE LAN/WAN - Optical
- GE - Optical
- 10/100/1000 - Electrical
- ETH Traffic Classification
- Complete Scope of Ethernet Features

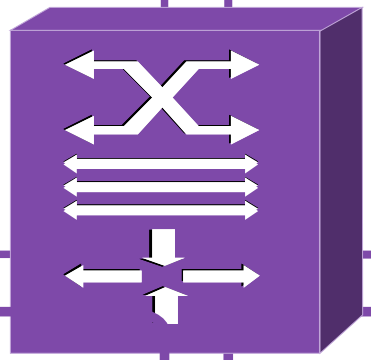

CWDM

- Terminal, Hub, OADM Ring
- Stacked C-WDM Rings

MSPP

- CLI, SNMP and TL1
- 2008 GMPLS feature set

1850 TSS-320

Square kilometre Array

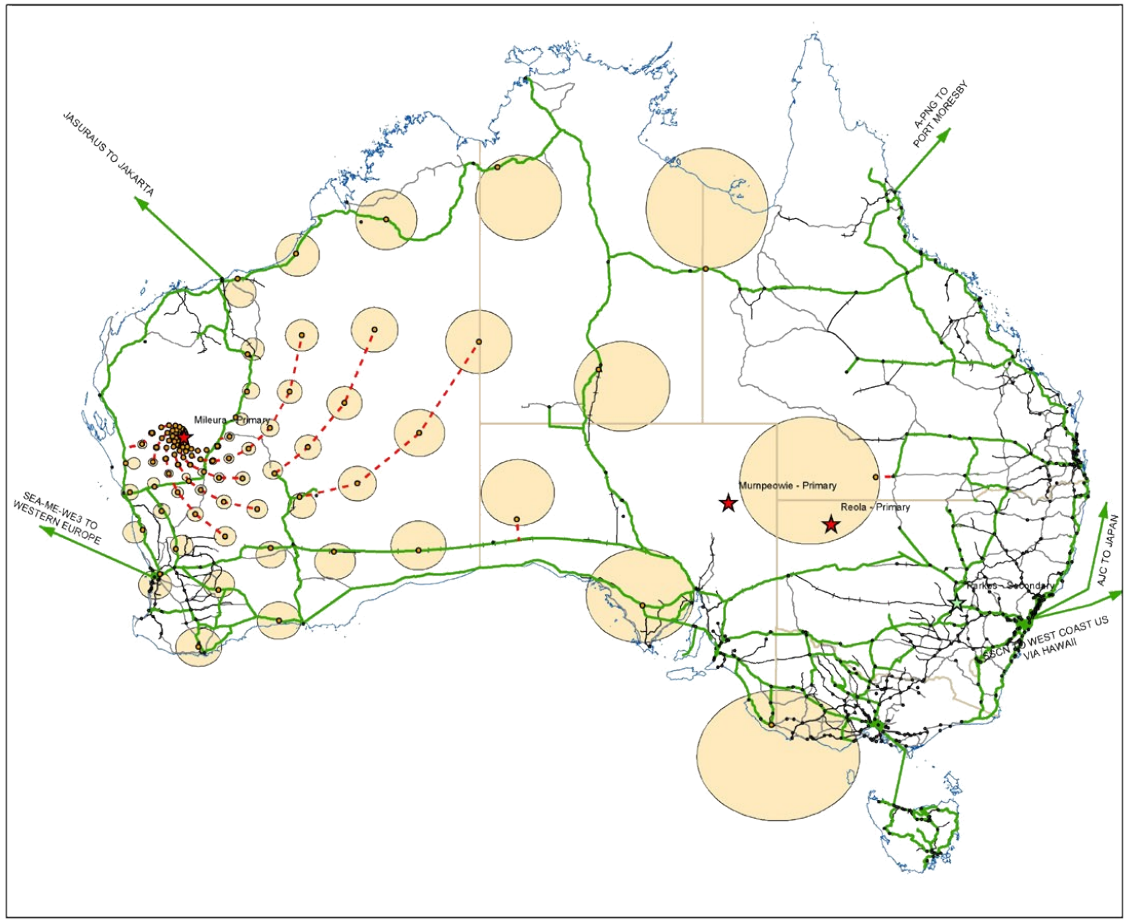

- $2Bn investment in infrastructure
- Real-time data analysis at petabits per second
- Storage >40 years
- Building the world's largest computational & data facility in one of the world's most isolated locations

Heterogeneous Networking

- Don’t attempt to build a single, complex network from different access, backhaul and core networks
- Hide the differences between technologies inside the relevant layers
- Give users freedom to use the network at each layer in any way they want

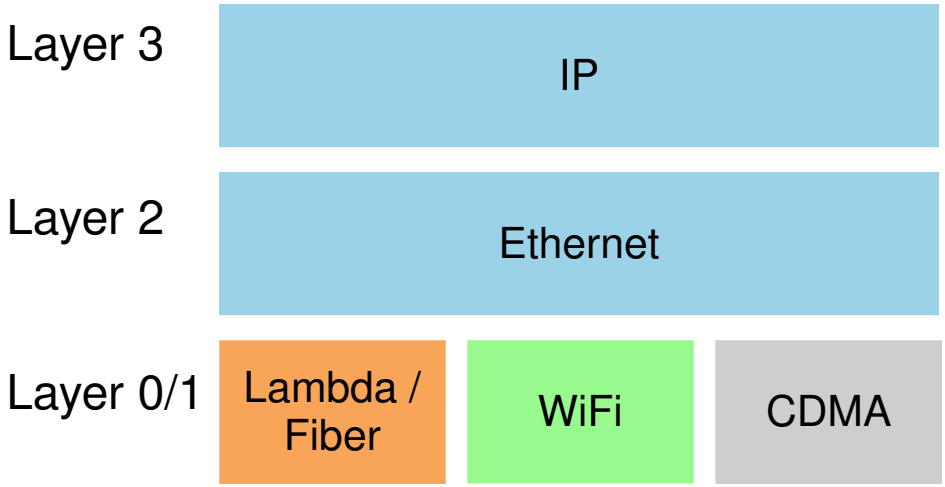

Reaching users

- So far innovative applications have to come to the emerging GLIF infrastructure
- The challenge is to bring GLIF to the desktop of researchers and to scientific instruments
  - Bringing dark fiber to remote instruments
  - Need for global collaboration to ensure global network reach
  - Bringing hybrid networking and lambda switching capability to campus LANs
- Network virtualization as a means to create heterogeneous networks of pure lambda resources and shared IP infrastructure

ALMA

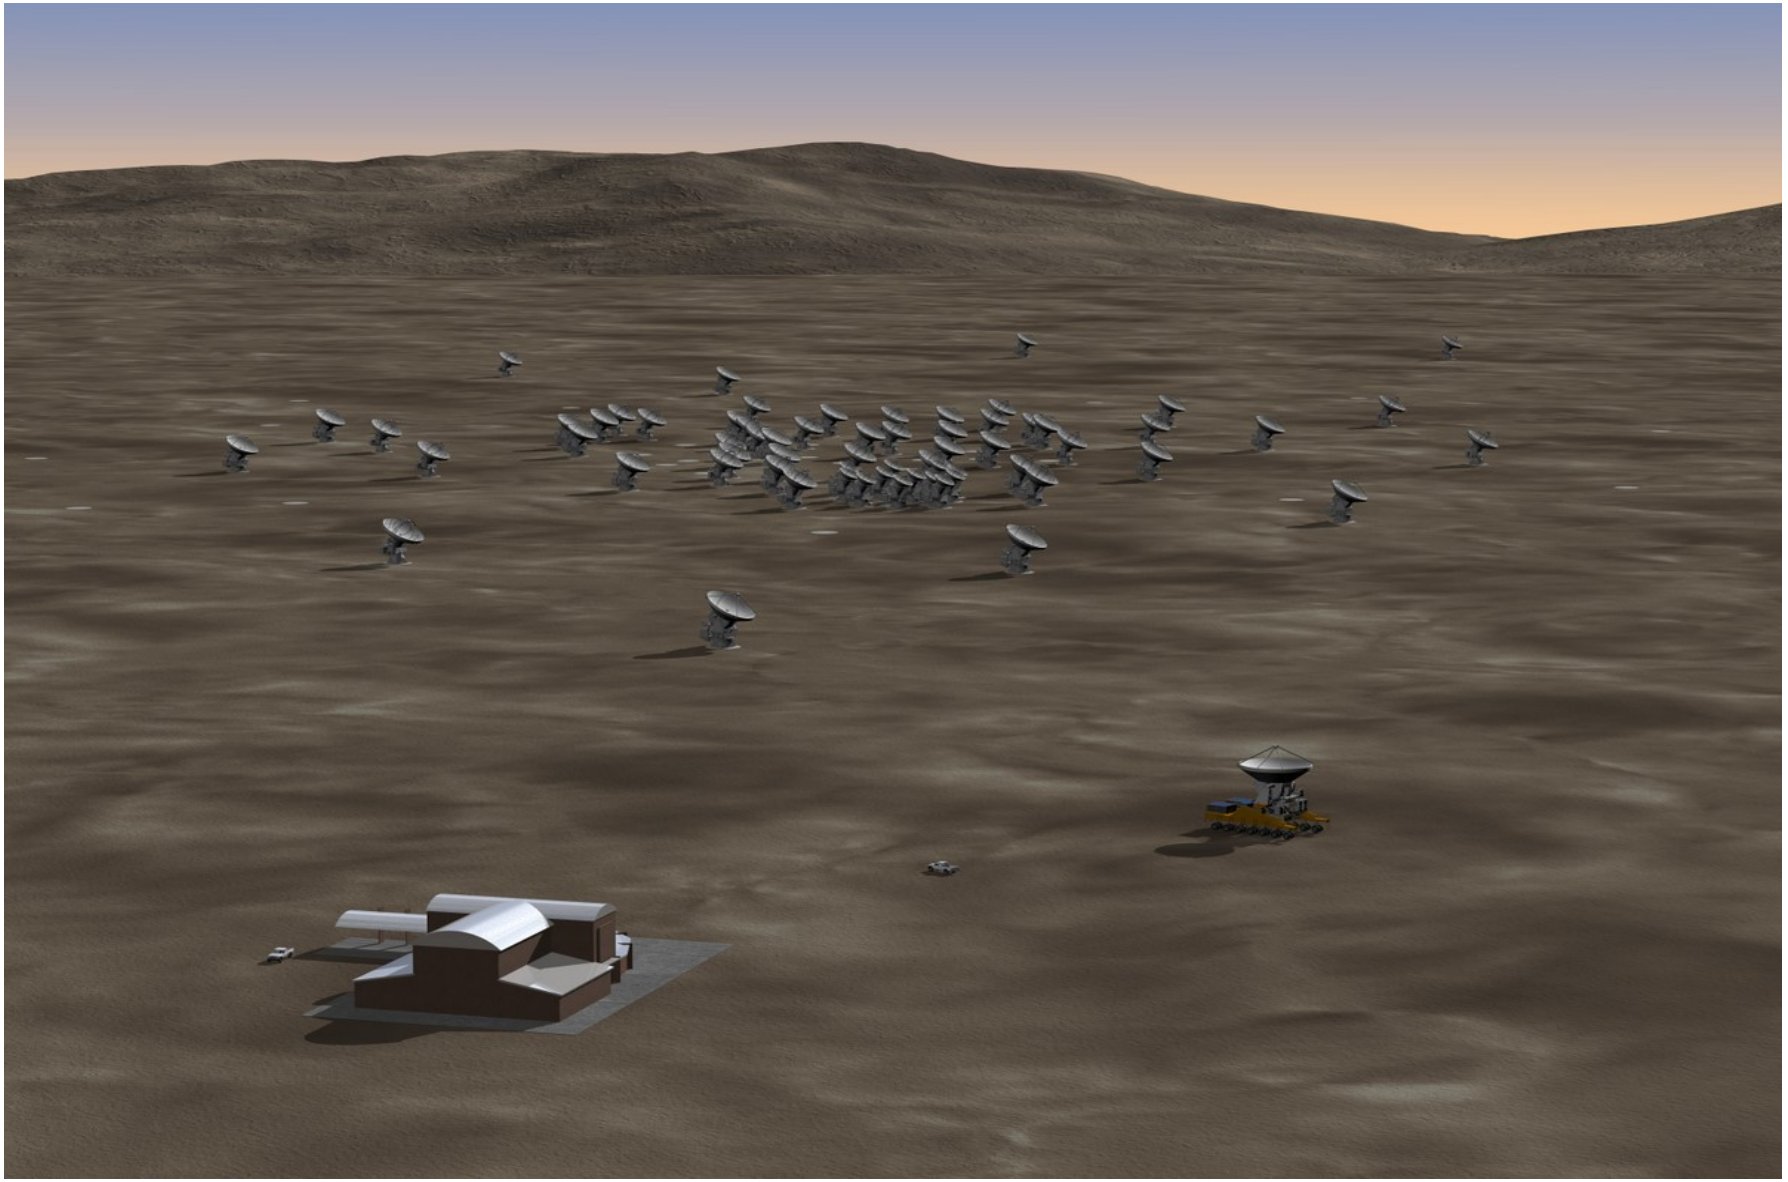

ALMA

NORDUnet
Nordic infrastructure for Research & Education

ALMA transporter

NORDUnet
Nordic infrastructure for Research & Education

NorthernLight 2008

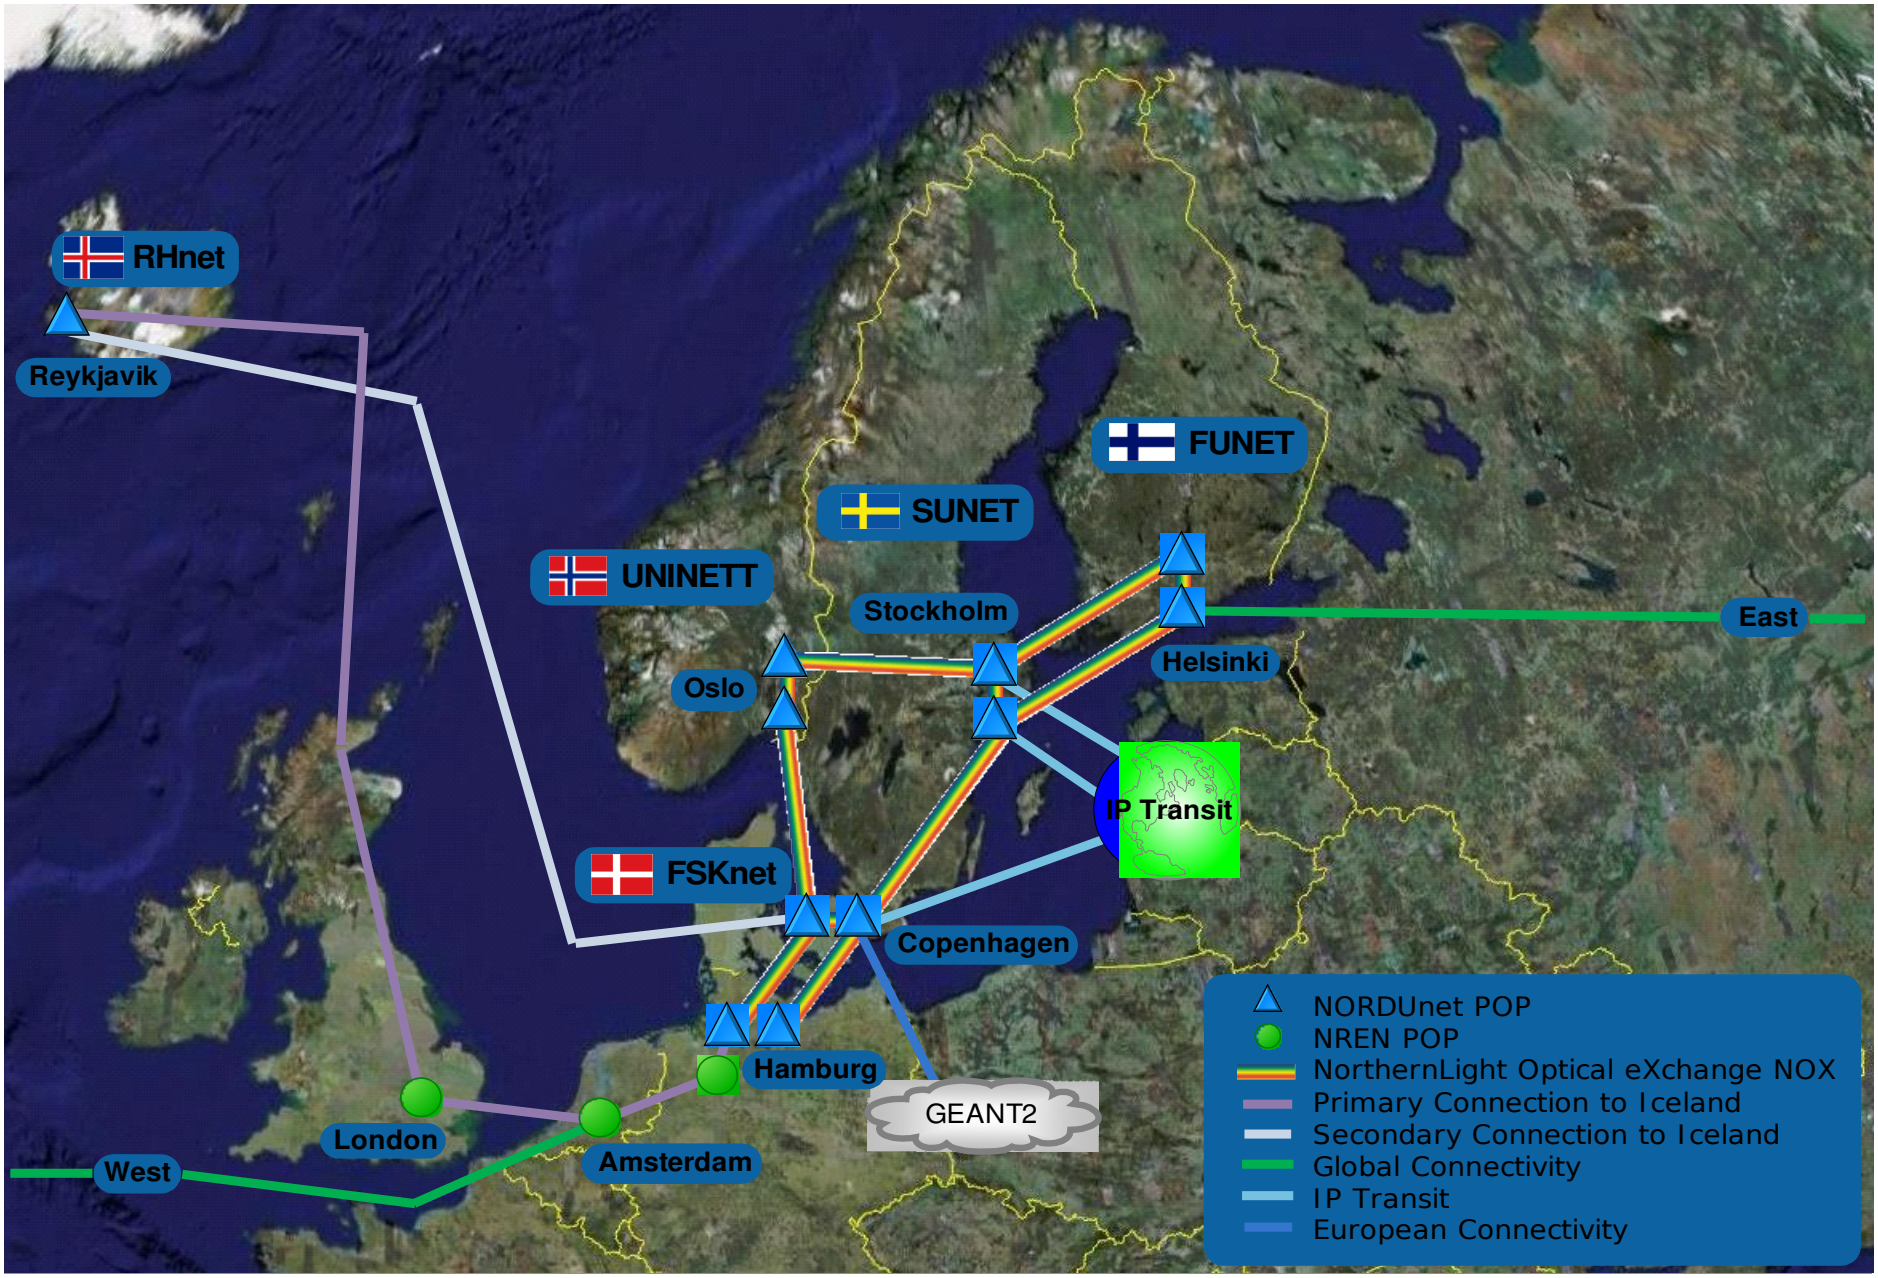

NORDUnet
Nordic infrastructure for Research & Education

Nordic NREN Optical Networks

- Sweden: SUNET
  - Dark Fibre & DWDM acquired
  - Ligthpath service in operation
  - Ciena
- Denmark: Forskningsnet
  - Dark Fibre & DWDM acquired
  - Lightpath Service being deployed
  - AlcatelLucent
- Norway: UNINETT
  - Dark Fibre & DWDM acquired
  - Ligthpath service in operation
  - Siemens
- Finland: FUNET
  - Dark Fibre & DWDM acquired
  - Lightpath service being deployed
  - NokiaSiemens Networks
- Iceland: RHnet
  - Dark fibre network in Reykjavik area

Some Next Steps

- Dynamic Circuit Networking
  - Nordic DCN solution – NORDUnet and Nordic NRENs Interoperable with major DCN approaches
  - NORDUnet will push for standardization of interfaces and development of demonstration implementations
  - Participate in interoperability trials
- (Network) Virtualization
  - FEDERICA: European testbed for network and service virtualization
  - MANTICORE: Logical IP networks, integration of network virtualization across L1 - L3
- Transmission
  - Cross-border alien waves, multi-domain WSS
  - 40G, 100G trials

NORDUnet
Nordic infrastructure for Research & Education

Thank you!

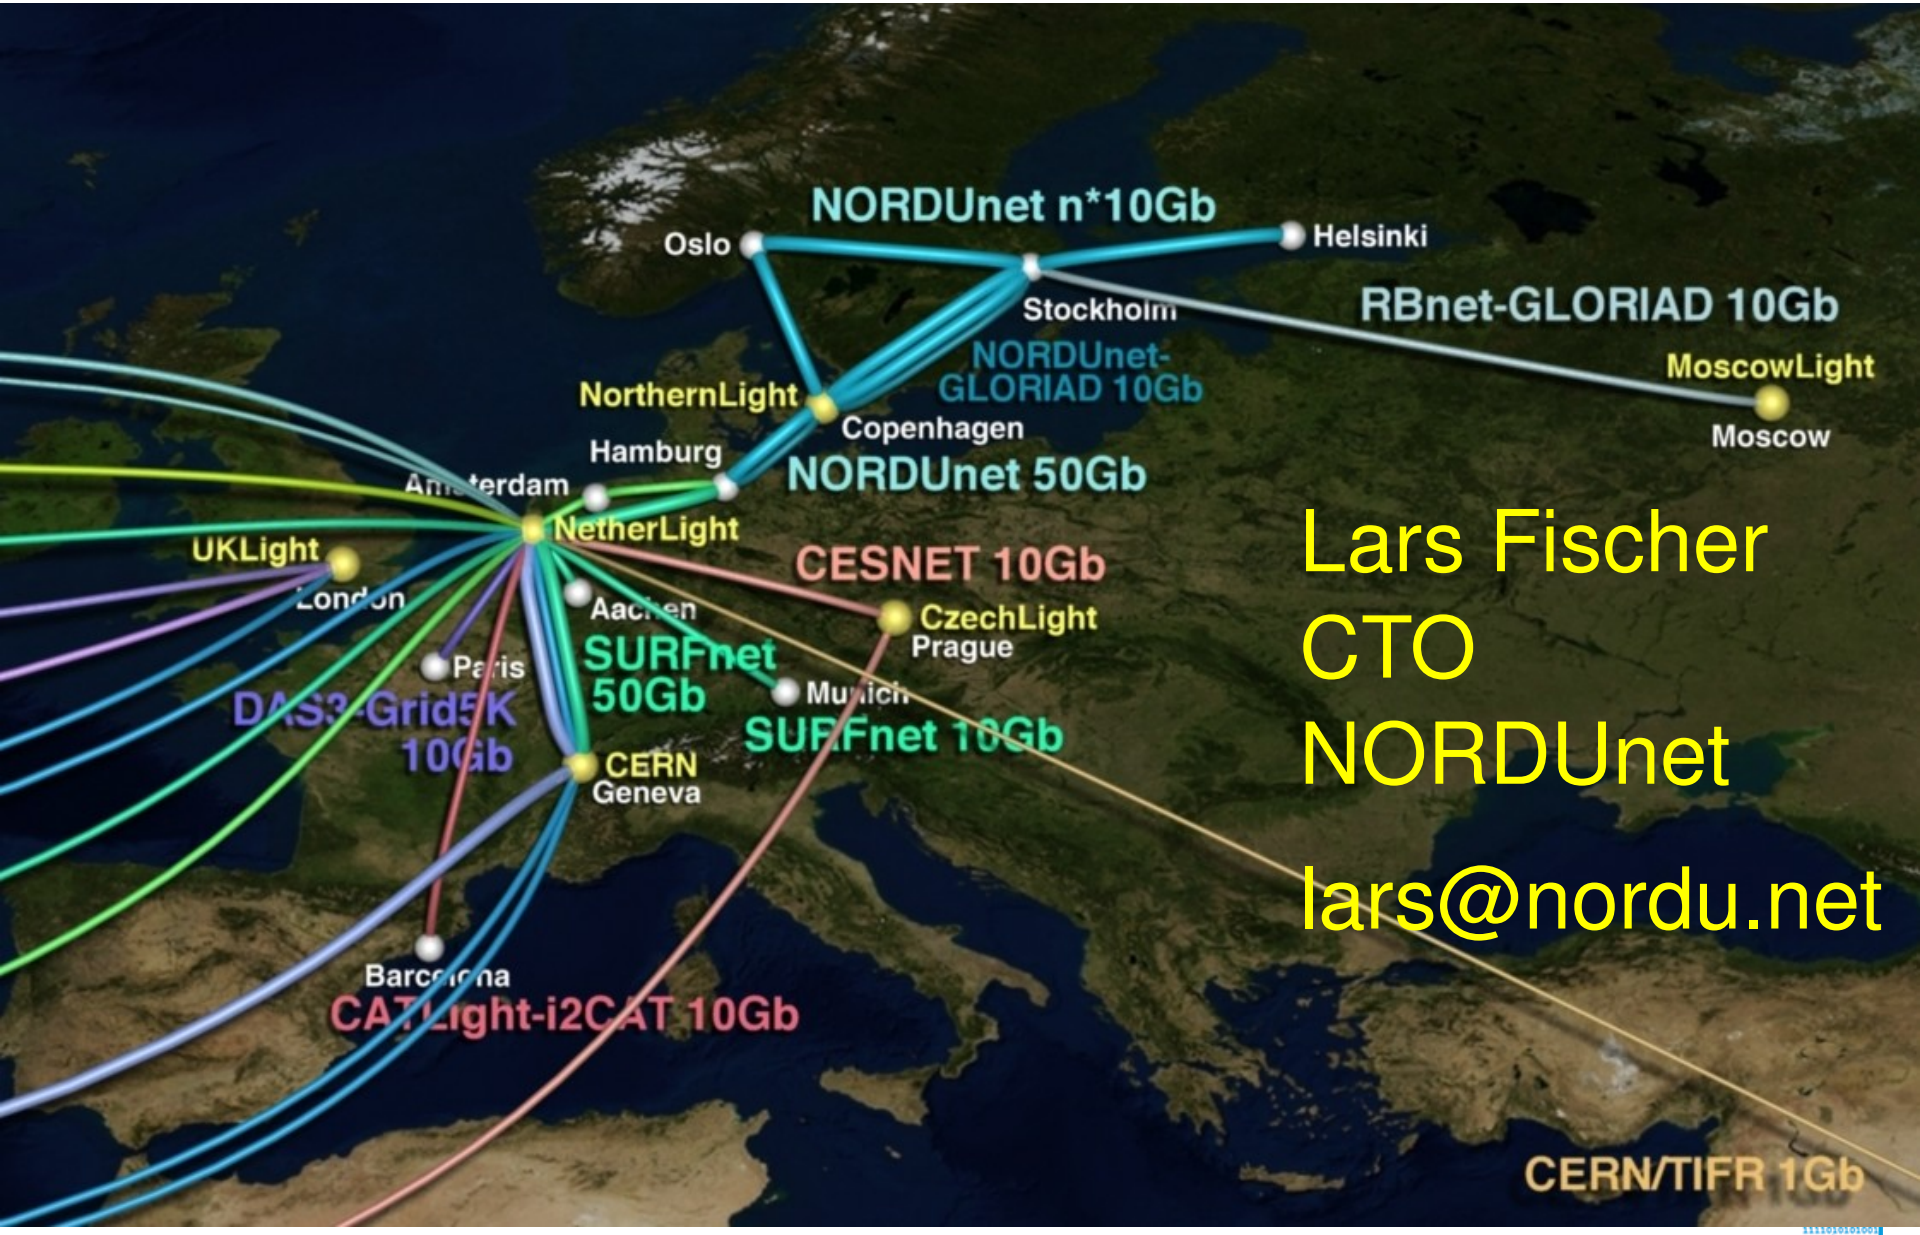

Lars Fischer
CTO
NORDUnet

lars@nordu.net