System & Network Engineering research UvA

Cees de Laat

EU

COMMIT

UvA

NWO

PID/EFRO

SURFnet

TNO

NCF

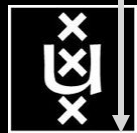

Science Faculty @ UvA

Informatics Institute

- CSA: Computer Systems Architecture (dr. A.D. Pimentel)
- FCN: Federated Collaborative Networks (Prof. dr. H. Afsarmanesh)
- IAS: Intelligent Autonomous Systems (Prof. dr. ir. F.C.A. Groen)
- ILPS: Information and Language Processing Systems (Prof. dr. M. de Rijke)
- ISIS: Intelligent Sensory Information Systems (Prof. dr. ir. A.W.M. Smeulders)
- SCS: Section Computational Science (Prof. dr. P.M.A. Sloot)
- SNE: System and Network Engineering (Prof. dr. ir. C.T.A.M. de Laat)
- TCS: Theory of Computer Science (Prof. dr. J.A. Bergstra)

Internet developments

... more data!

... more realtime!

... more users!

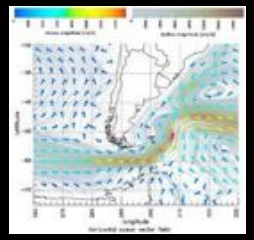

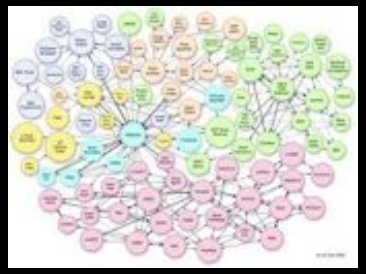

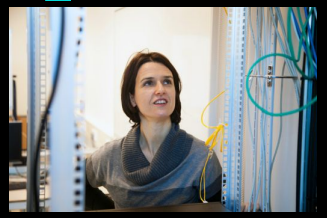

Internet developments

... more data!

Speed

Volume

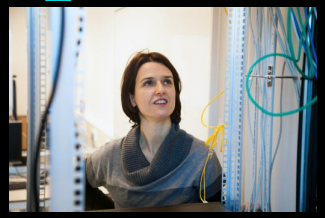

Deterministic

Real-time

... more realtime!

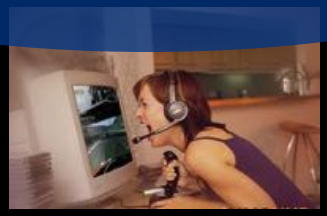

Scalable

Secure

... more users!

GPU cards are distruptive!

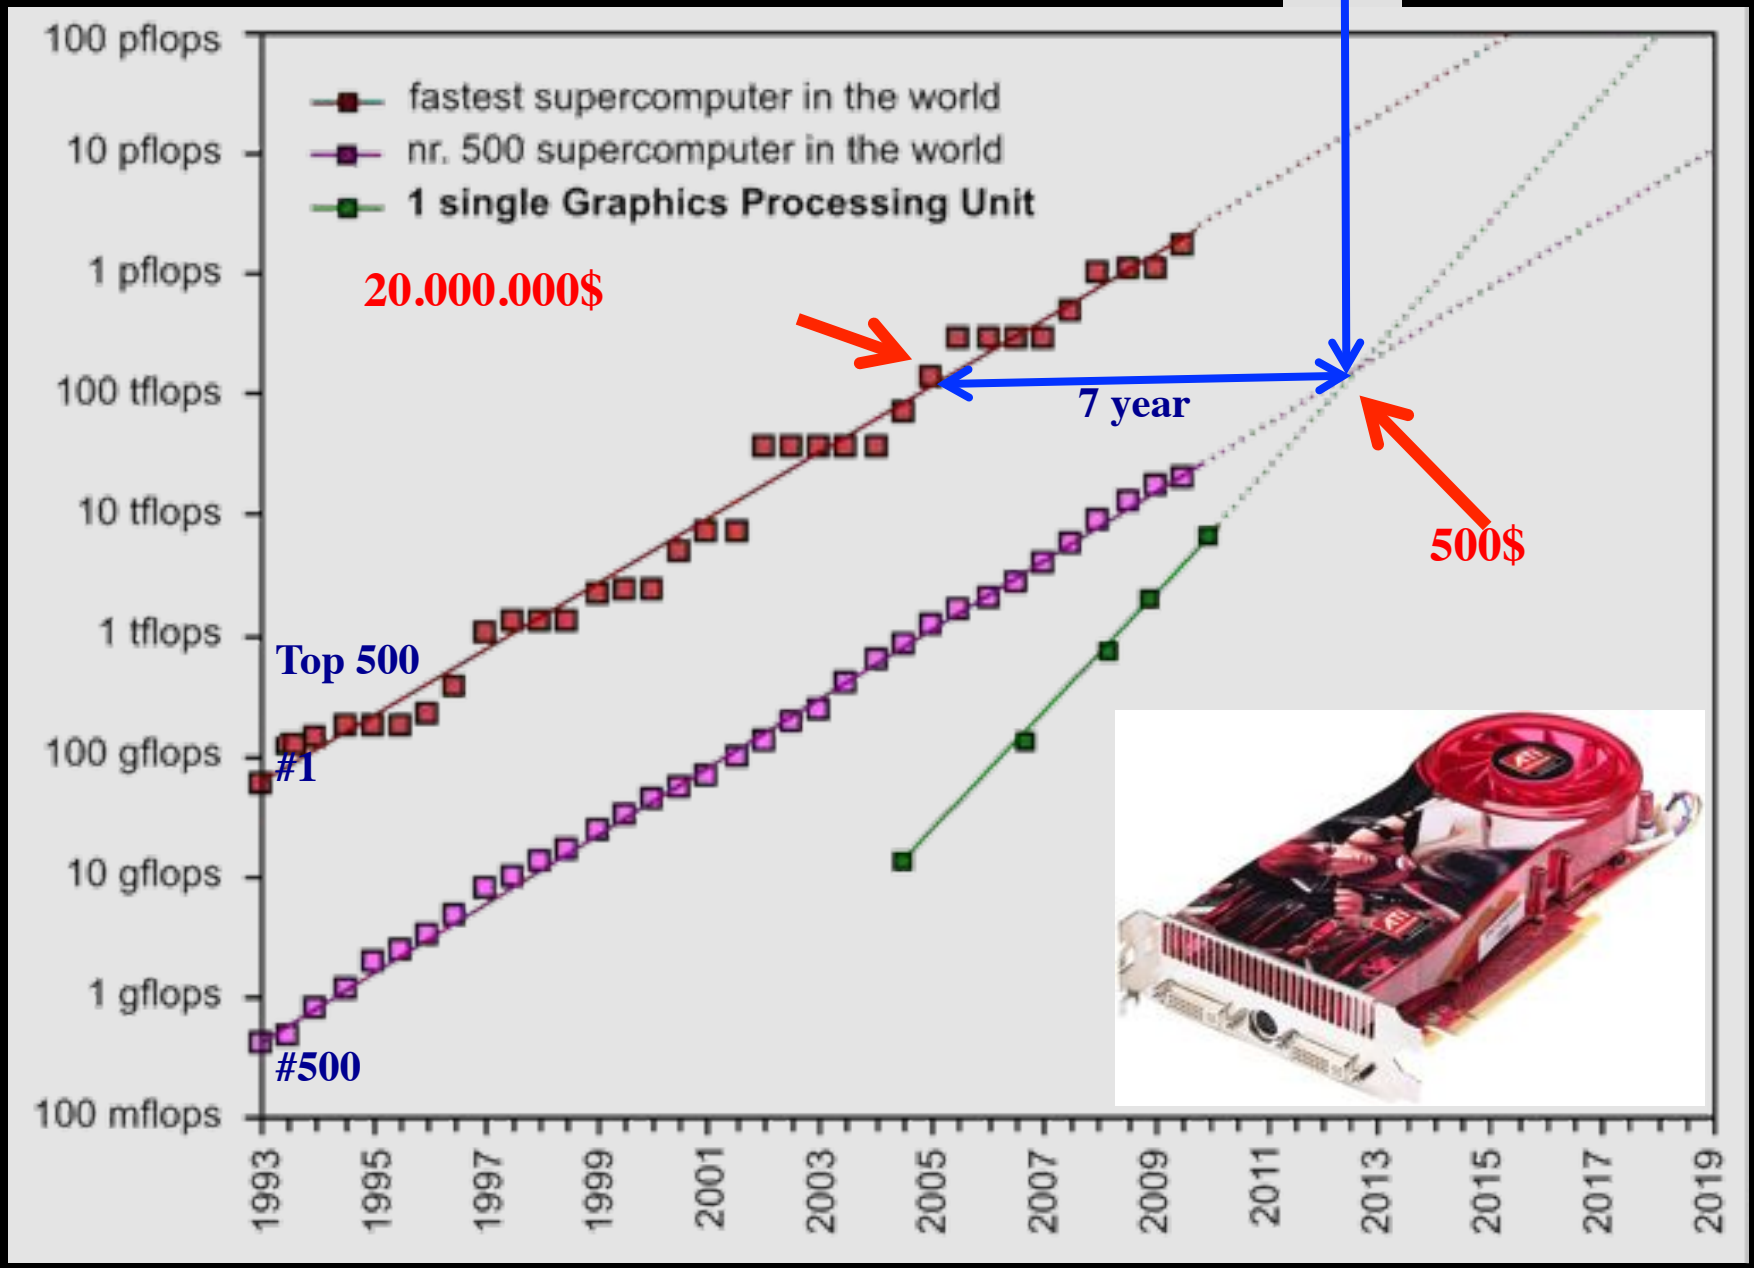

Data storage: doubling every 1.5 year!

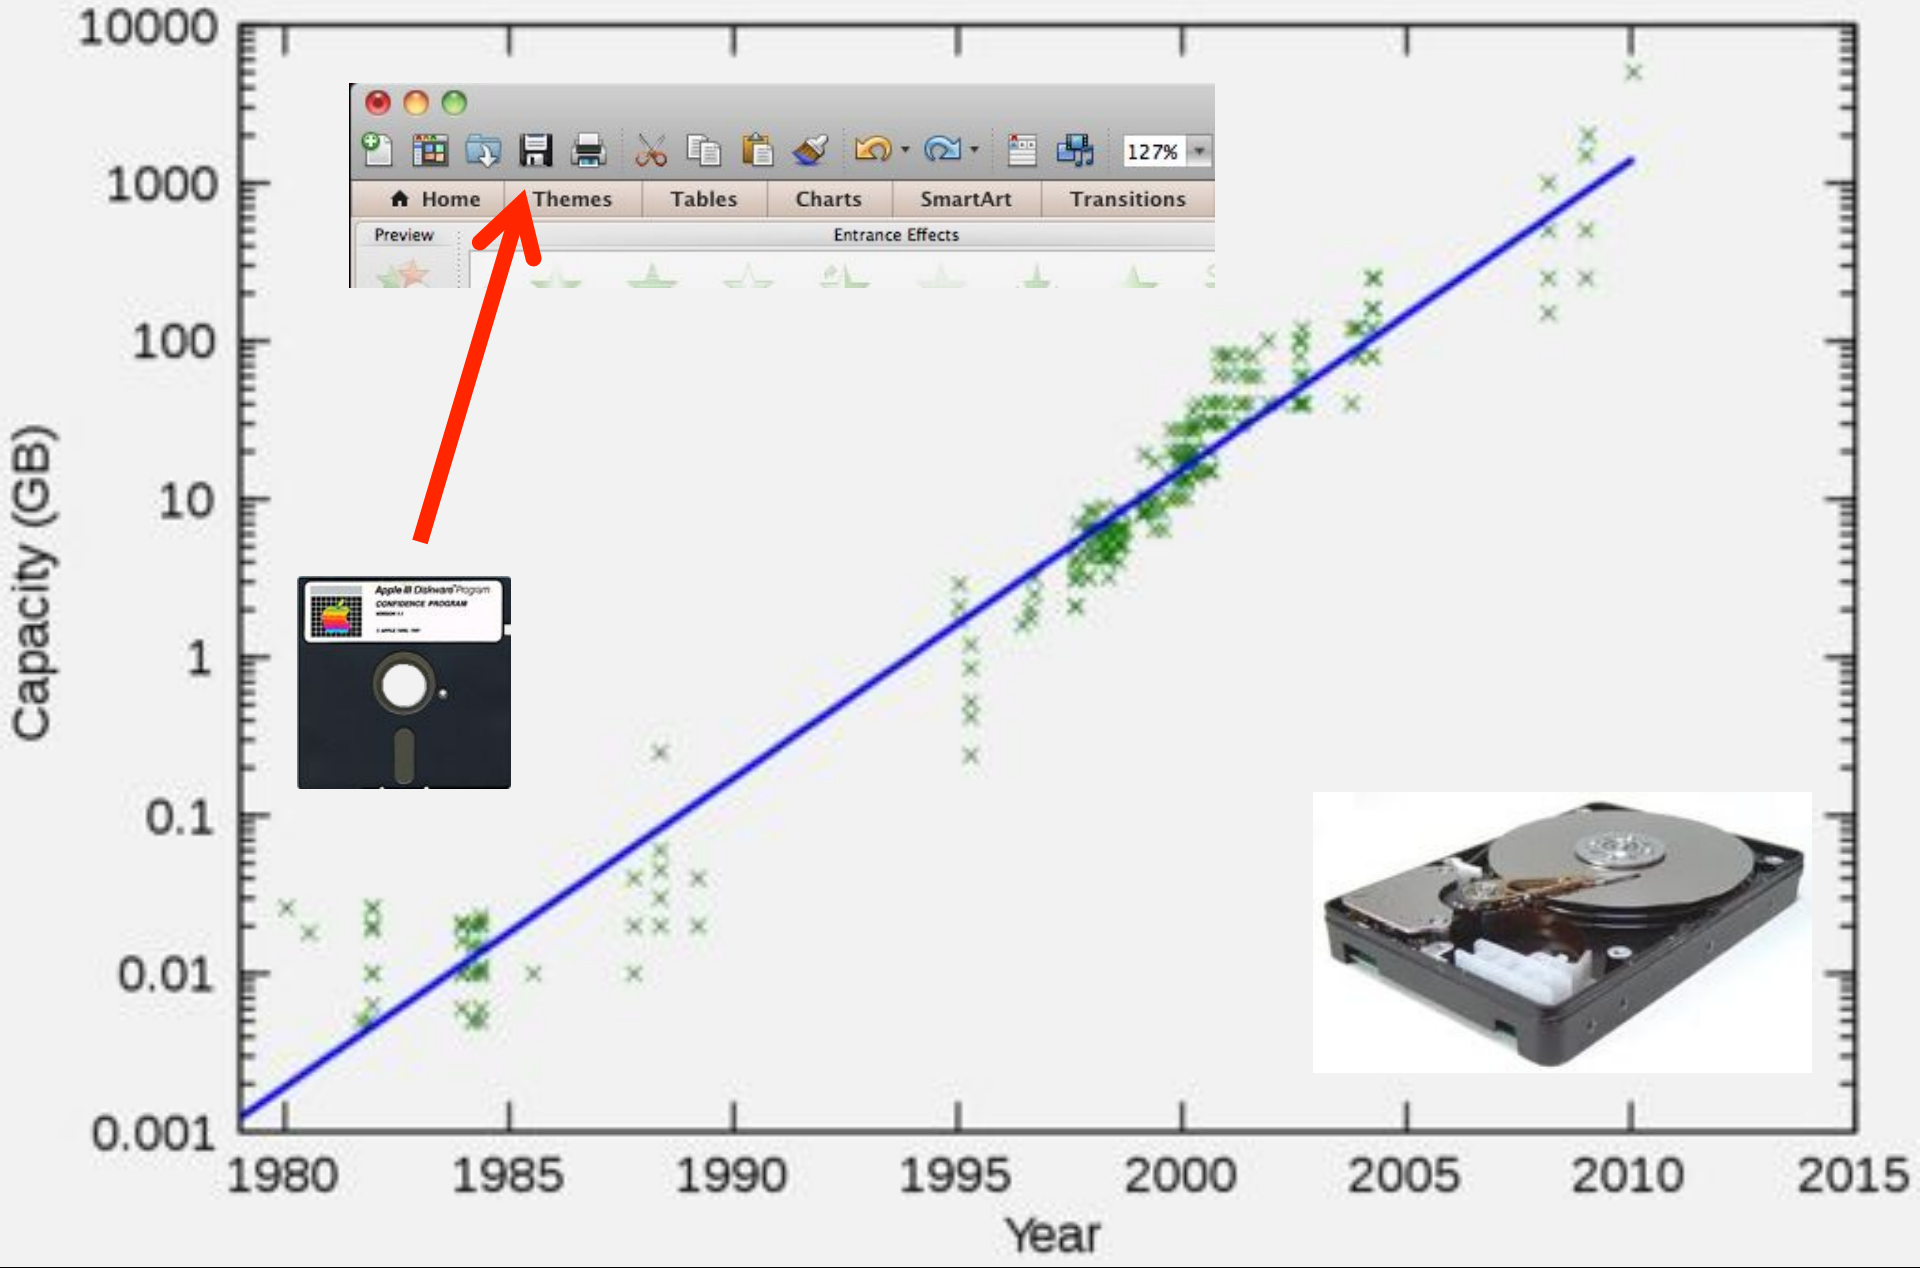

Multiple colors / Fiber

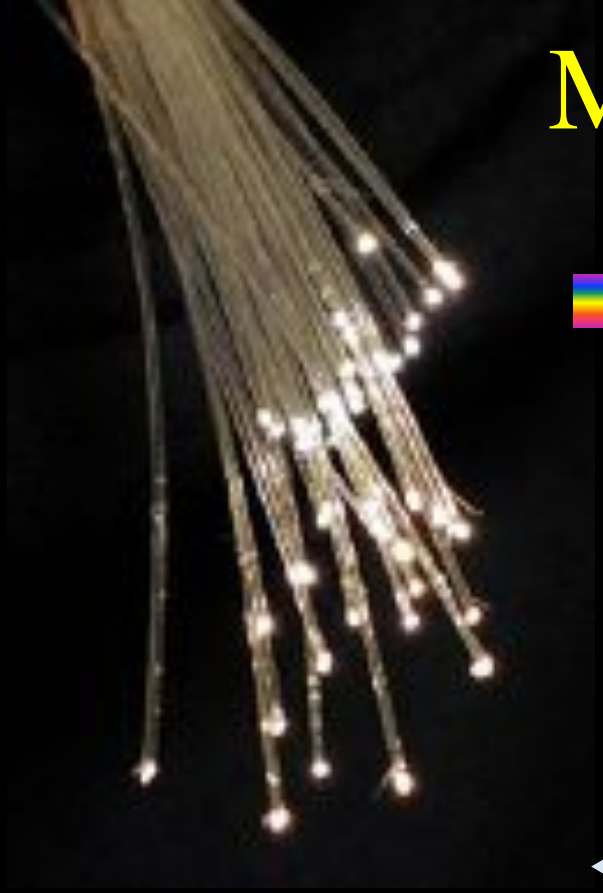

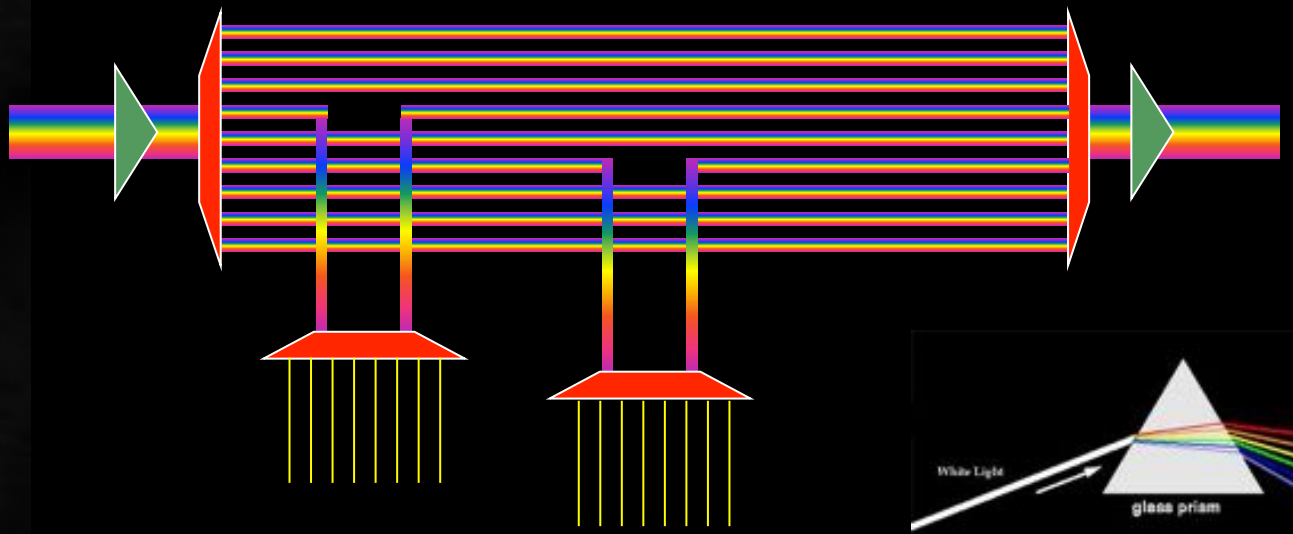

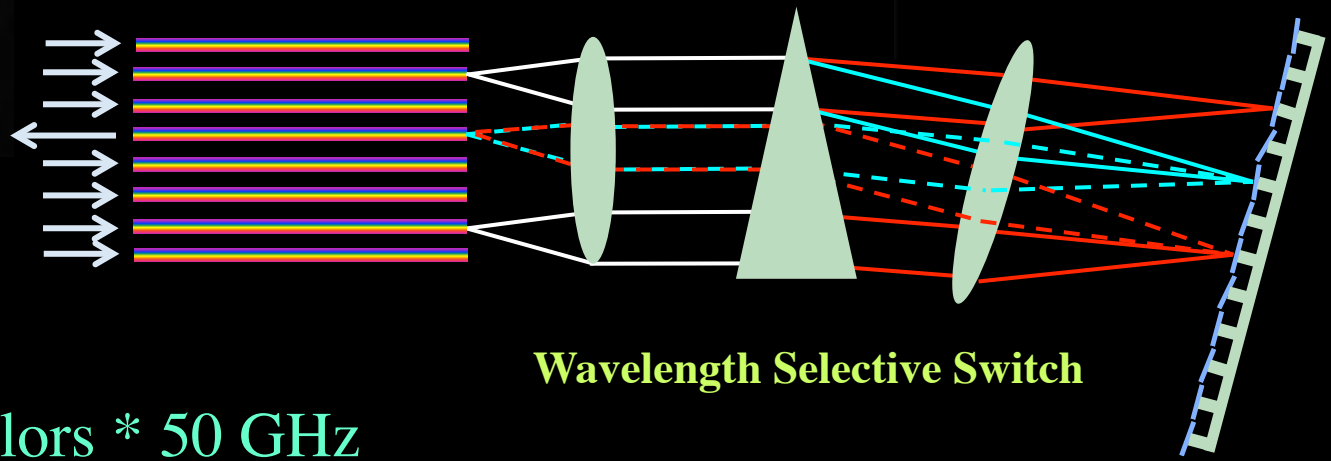

Wavelength Selective Switch

Per fiber: ~ 80-100 colors * 50 GHz

Per color: 10 – 40 – 100 Gbit/s

BW * Distance ~ $2*10^{17}$ bm/s

New: Hollow Fiber!

➔ less RTT!

Optical transmission

... more possibilities

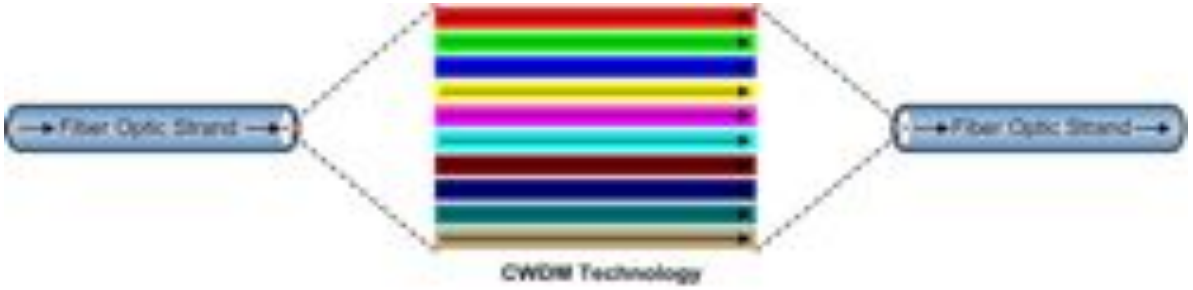

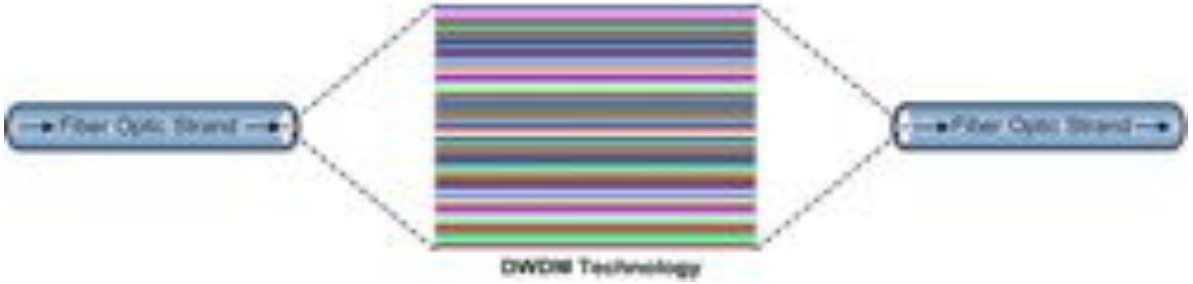

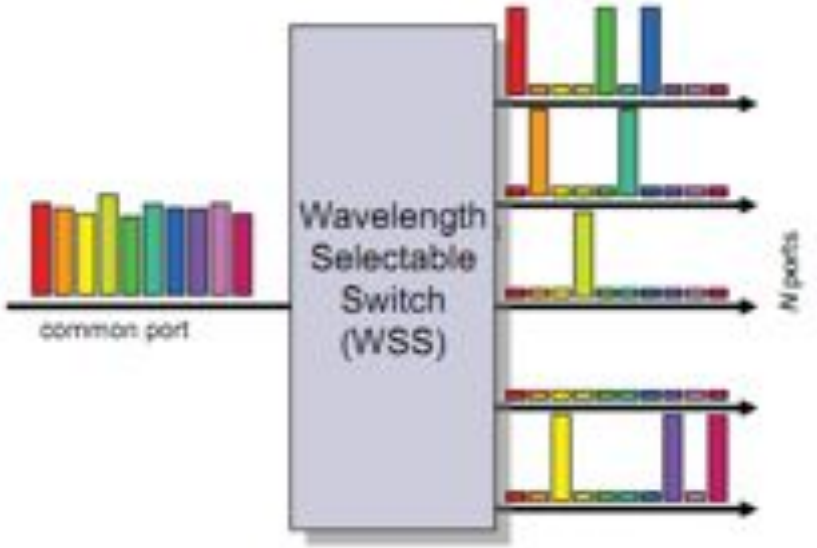

Virtualization

SNE @ UvA

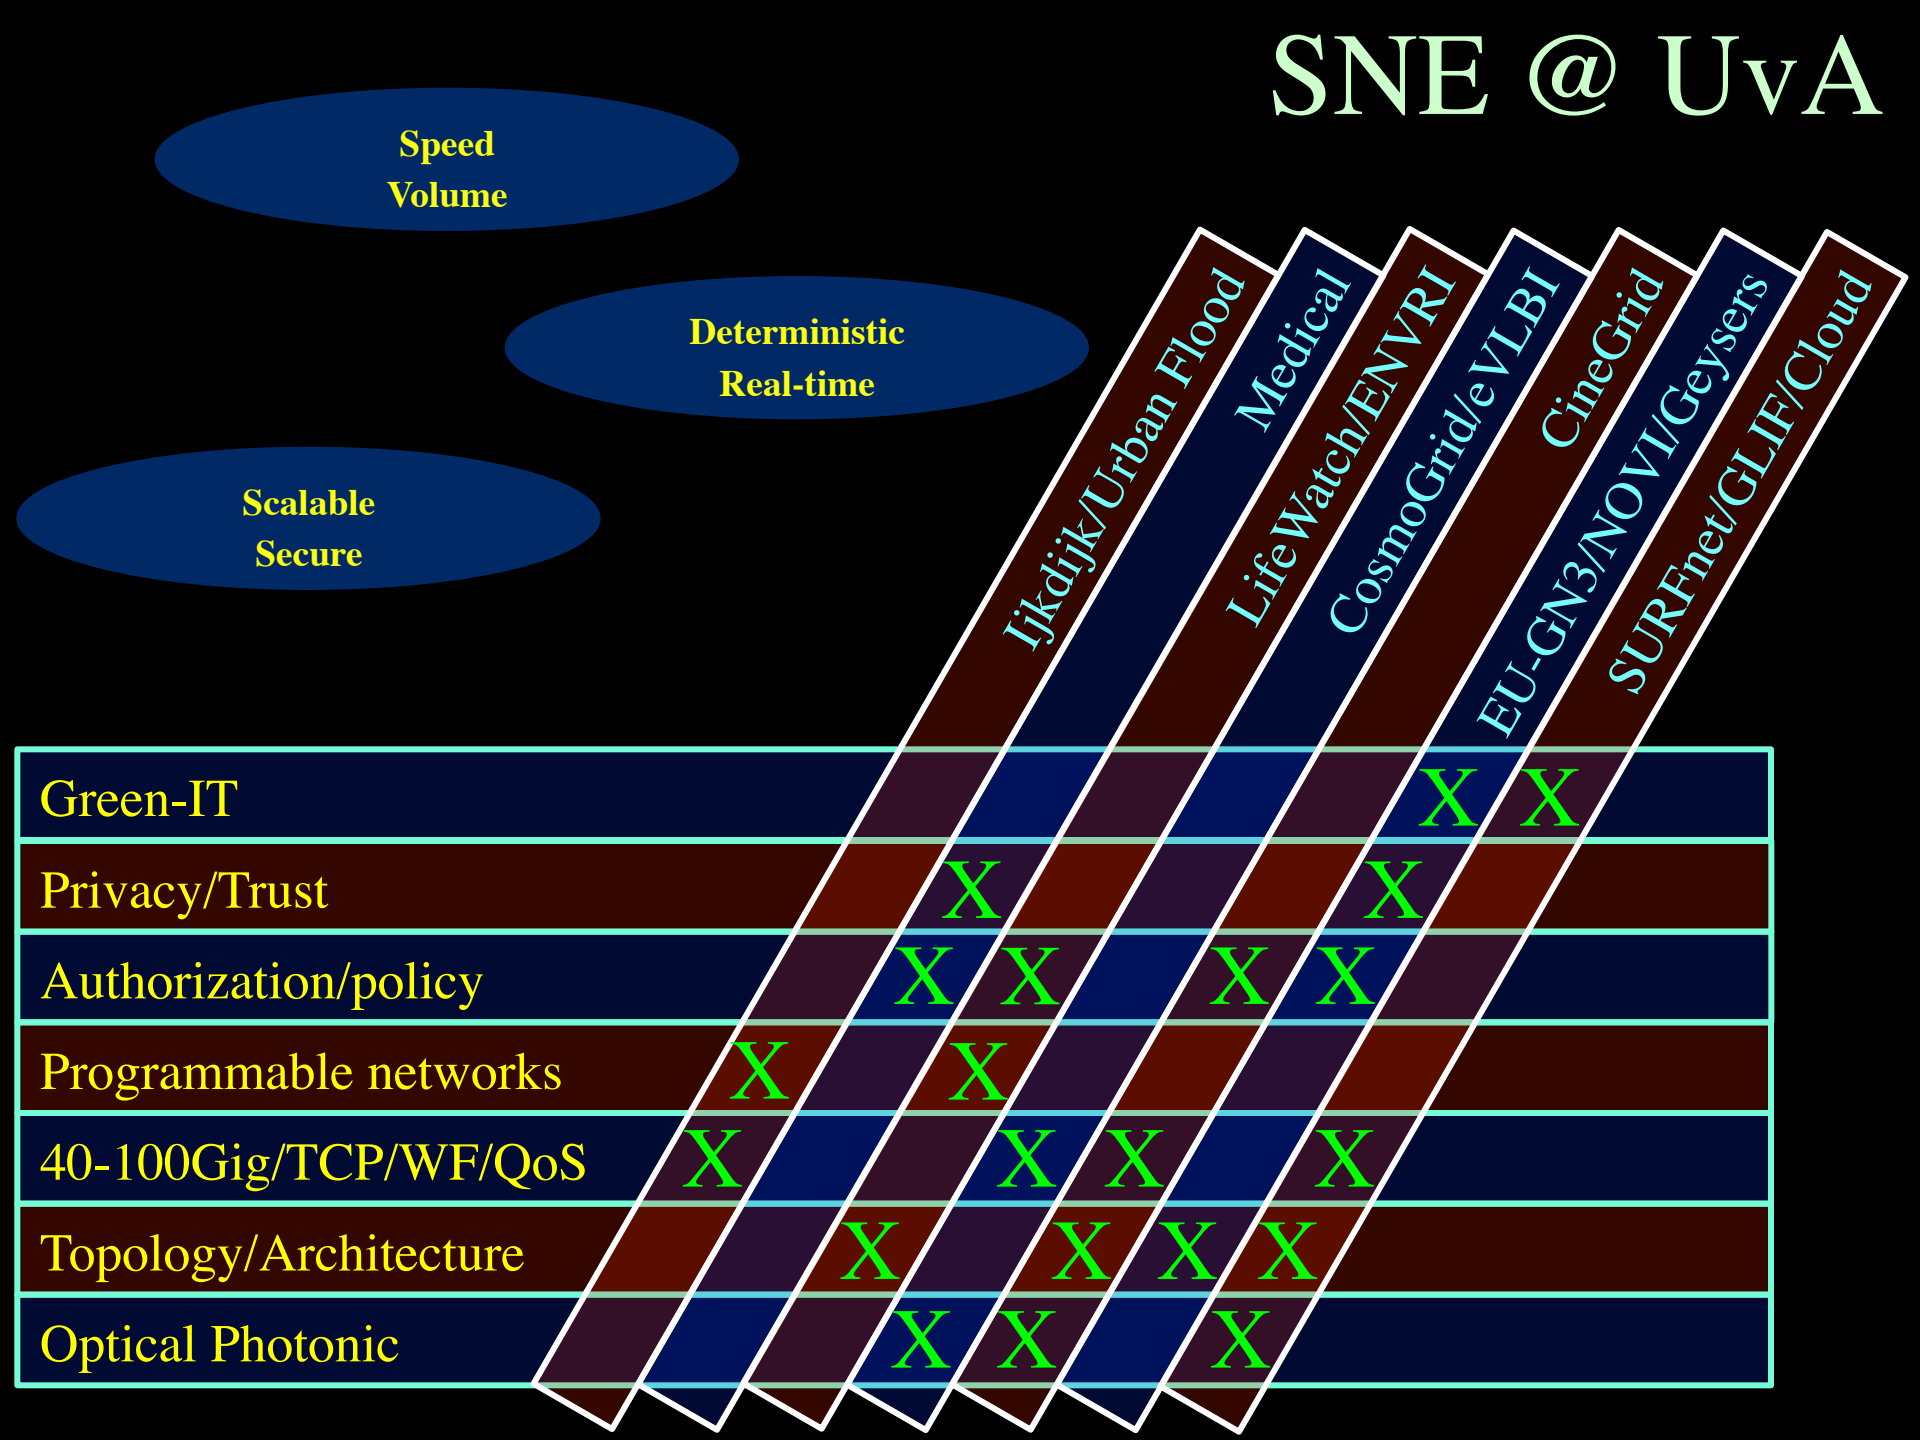
Speed
Volume
Deterministic
Real-time
Scalable
Secure
Ijkdijk/Urban Flood
Medical
LifeWatch/ENVRI
CosmoGrid/eVLBI
CineGrid
EU-GN3/NOVI/Geysers
SURFnet/GLIF/Cloud
Green-IT
Privacy/Trust
Authorization/policy
Programmable networks
40-100Gig/TCP/WF/QoS
Topology/Architecture
Optical Photonic

SNE @ UvA

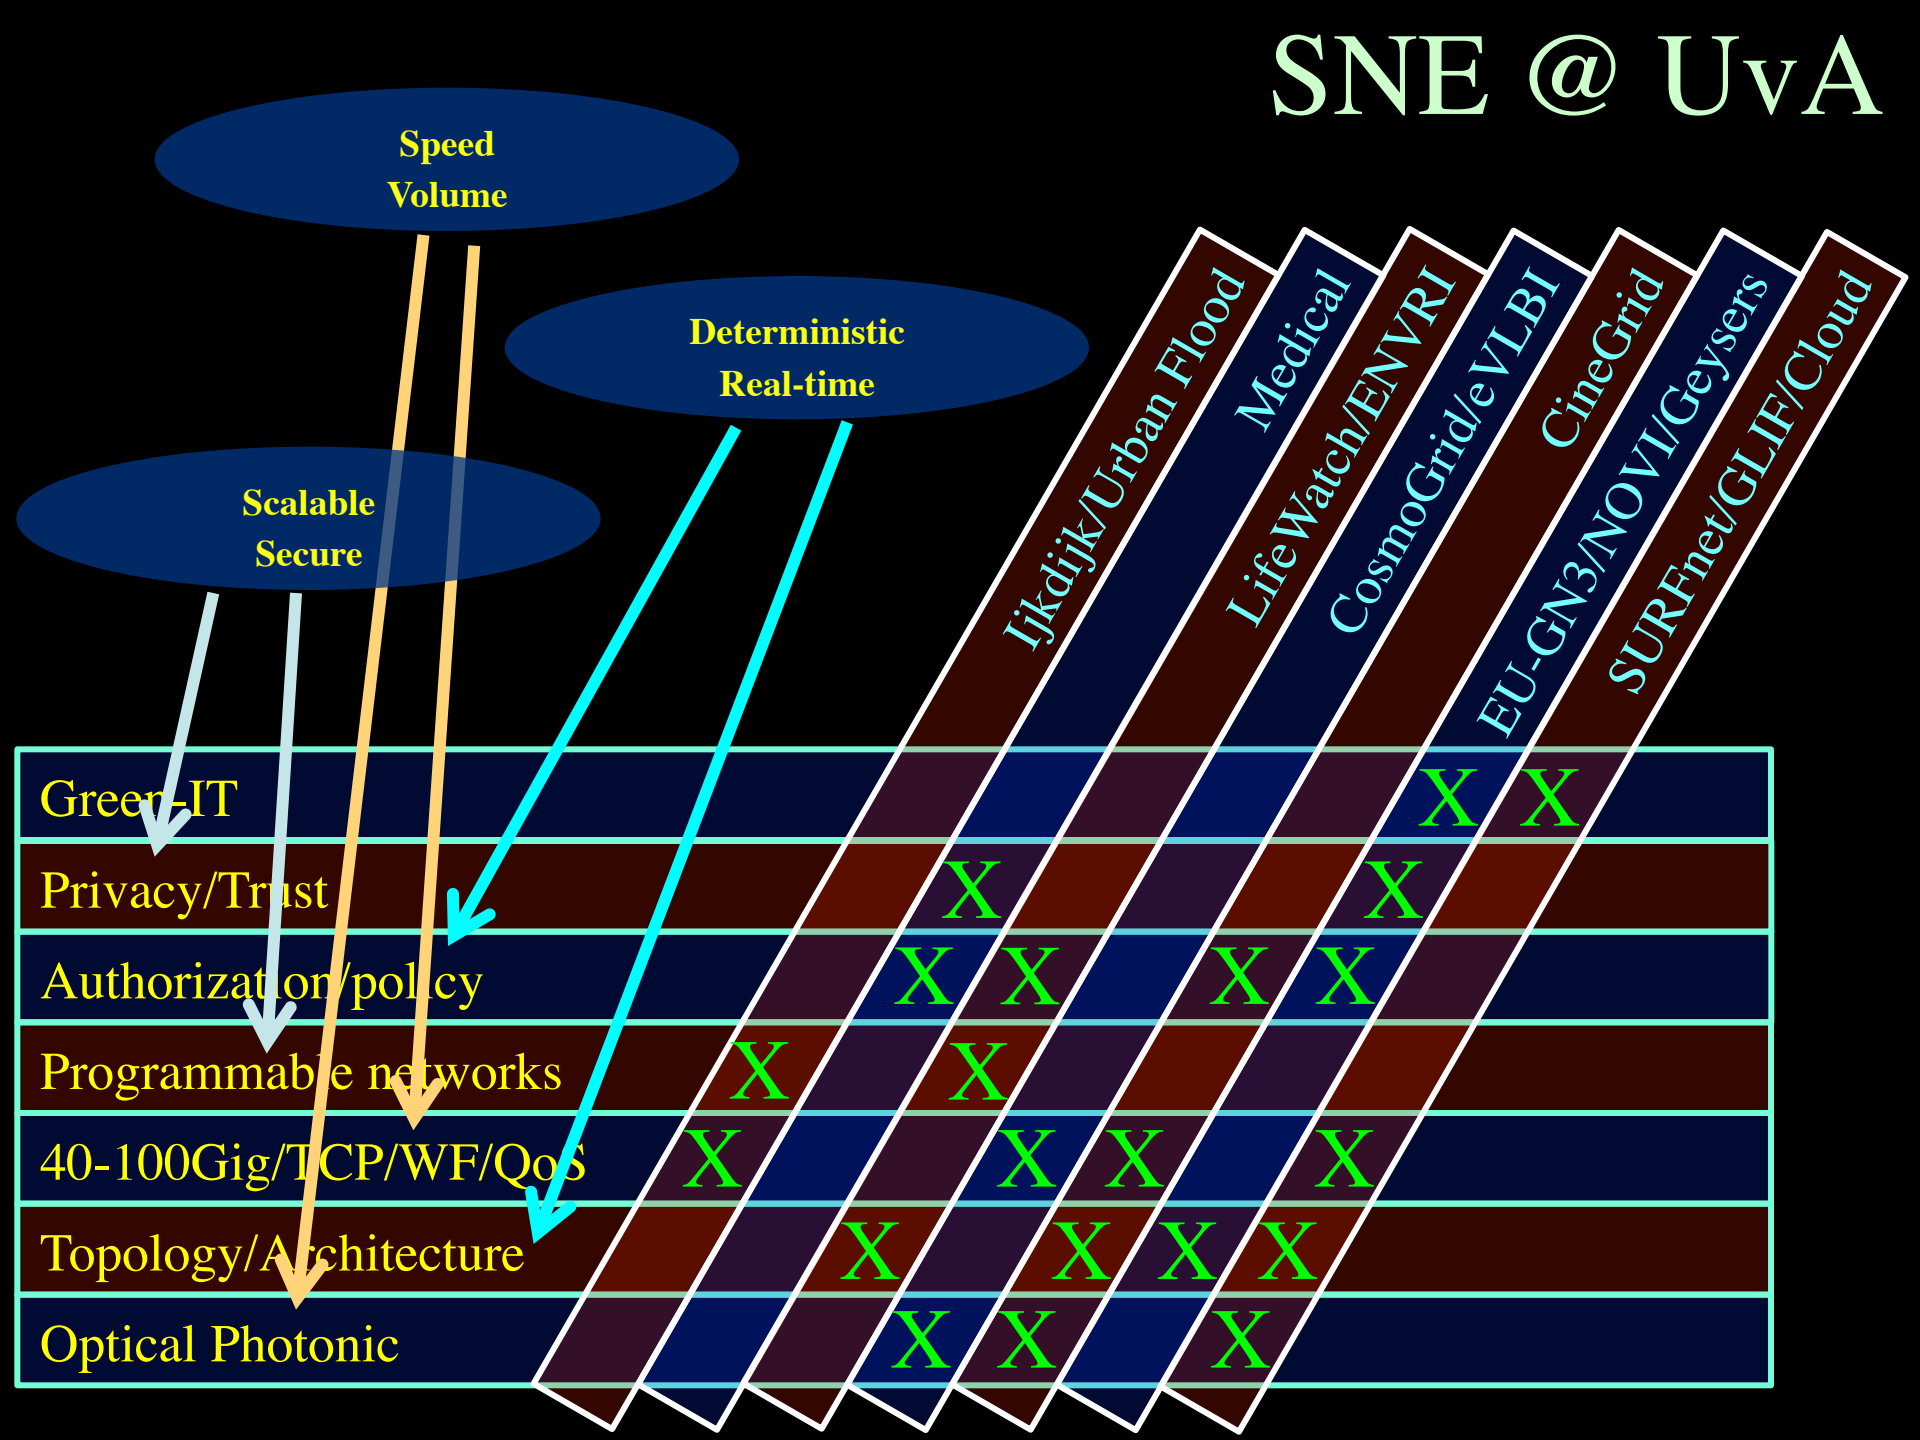
Speed
Volume
Deterministic
Real-time
Scalable
Secure
Ijkdijk/Urban Flood
Medical
LifeWatch/ENVRI
CosmoGrid/eVLBI
CineGrid
EU-GN3/NOVI/Geysers
SURFnet/GLIF/Cloud
Green-IT
Privacy/Trust
Authorization/policy
Programmable networks
40-100Gig/TCP/WF/QoS
Topology/Architecture
Optical Photonic

SNE @ UvA

	Ijkdijk/Urban Flood	Medical	LifeWatch/ENVRI	CosmoGrid/eVLBI	CineGrid	EU-GN3/NOVI/Geysers	SURFnet/GLIF/Cloud
Green-IT						X	X
Privacy/Trust			X		X		
Authorization/policy		X	X		X	X	
Programmable networks	X		X				
40-100Gig/TCP/WF/QoS	X		X	X		X	
Topology/Architecture		X		X	X	X	
Optical Photonic		X	X		X		

e -Very Large Base Interferometer

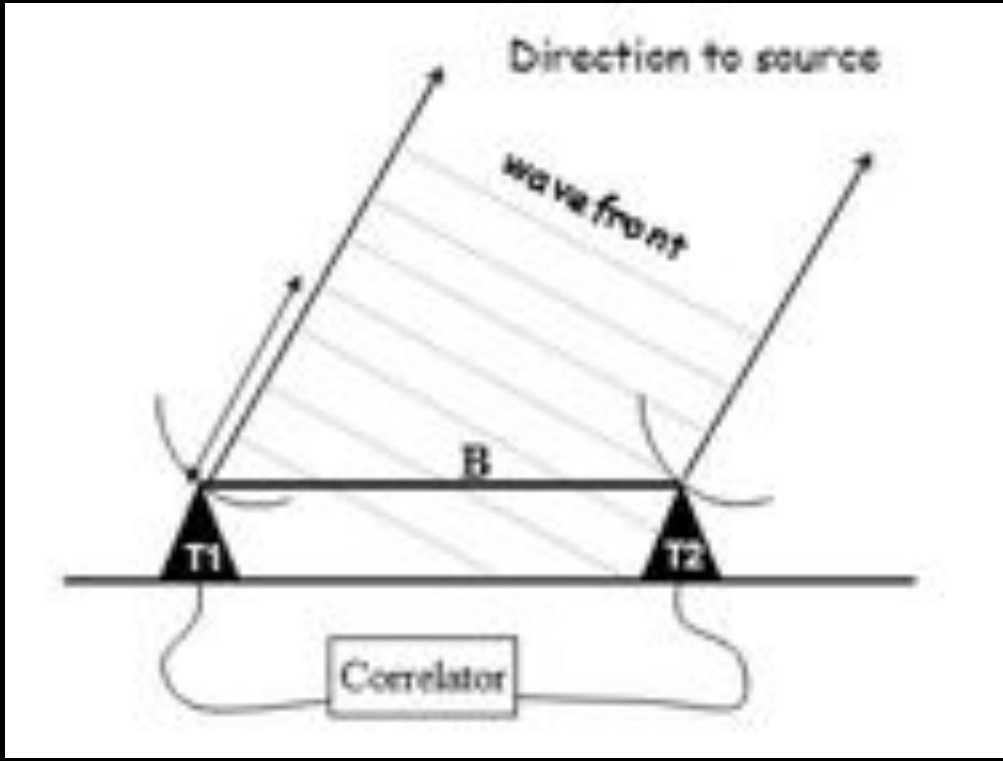

eEVN: European VLBI Network

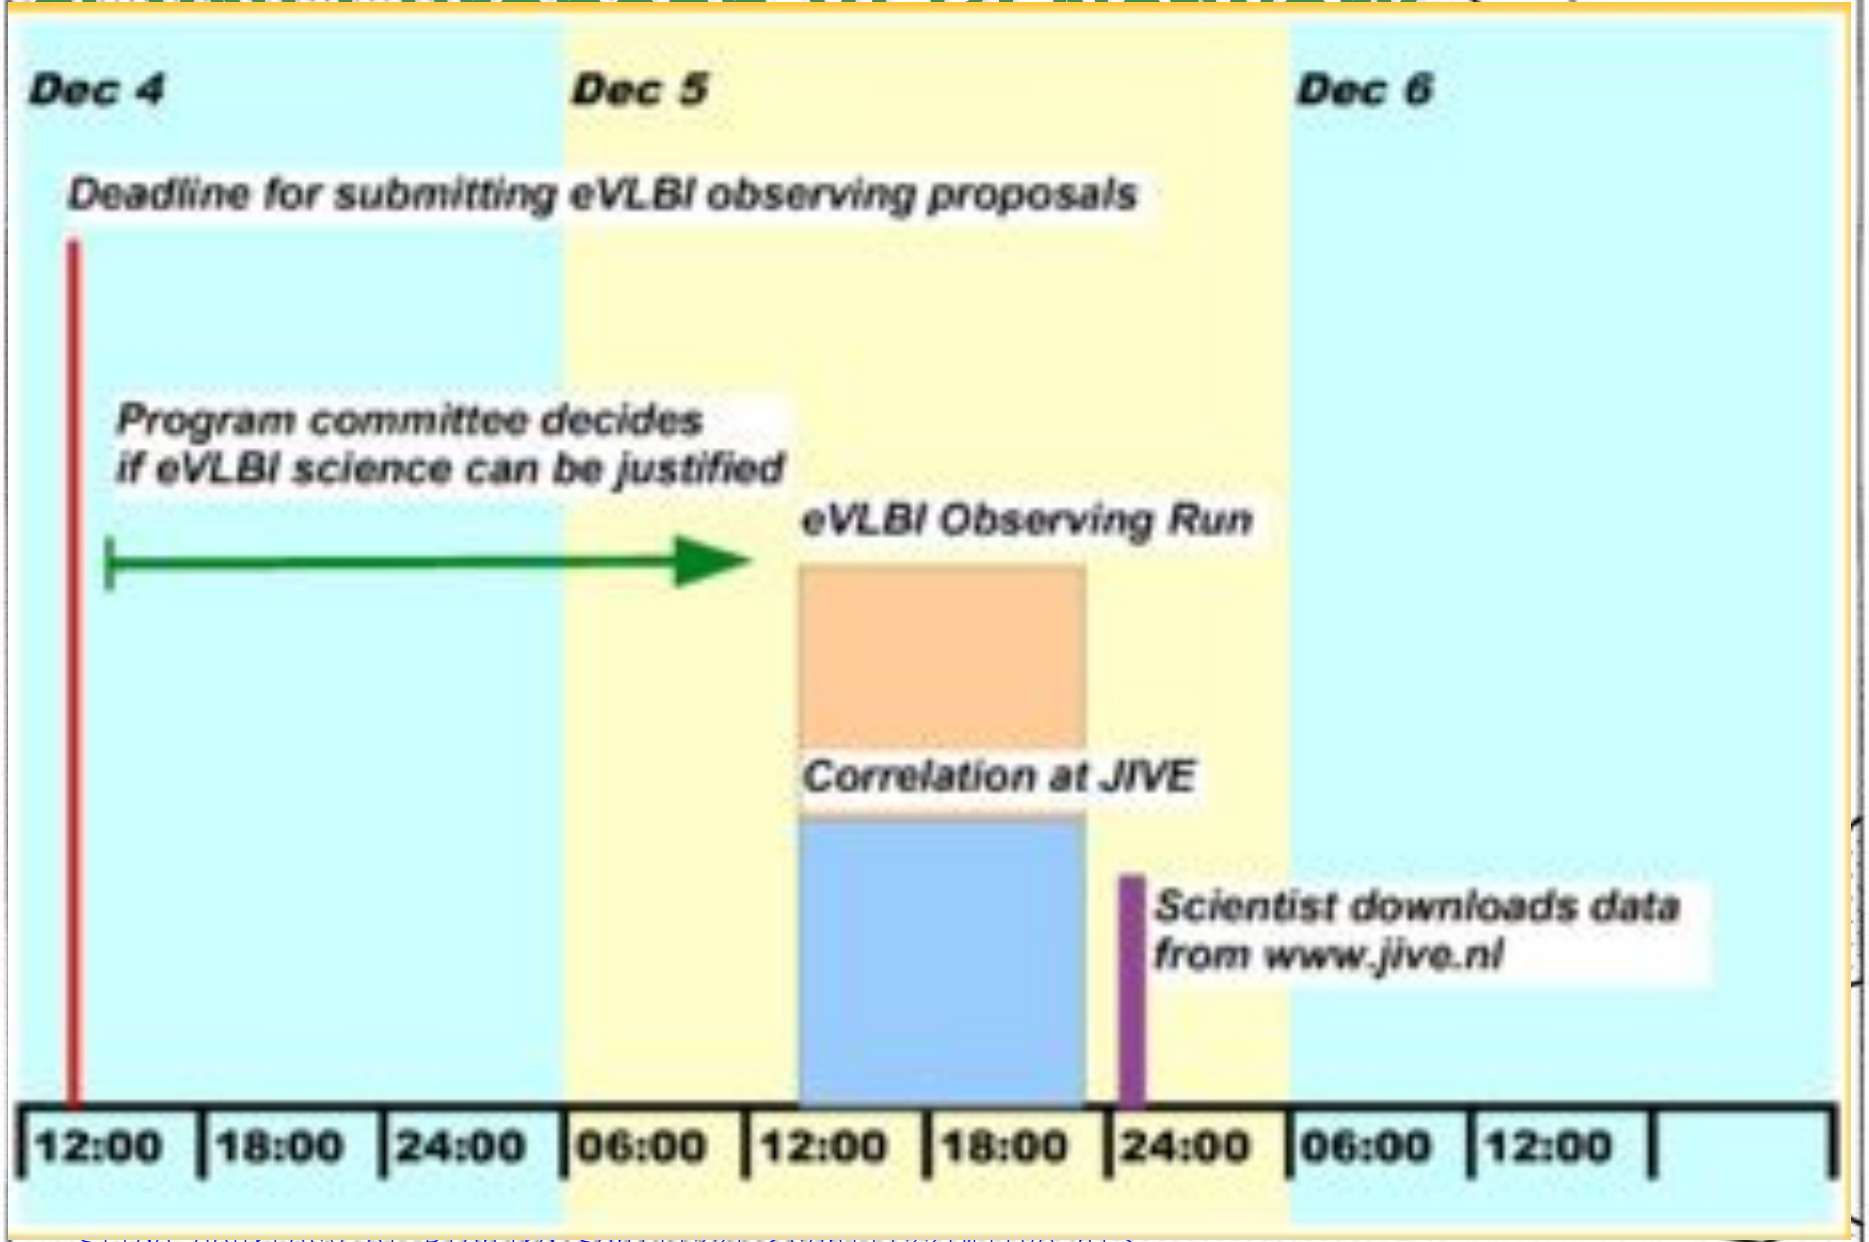

Slide courtesy of Richard Schilizzi <schilizzi@jive.nl>

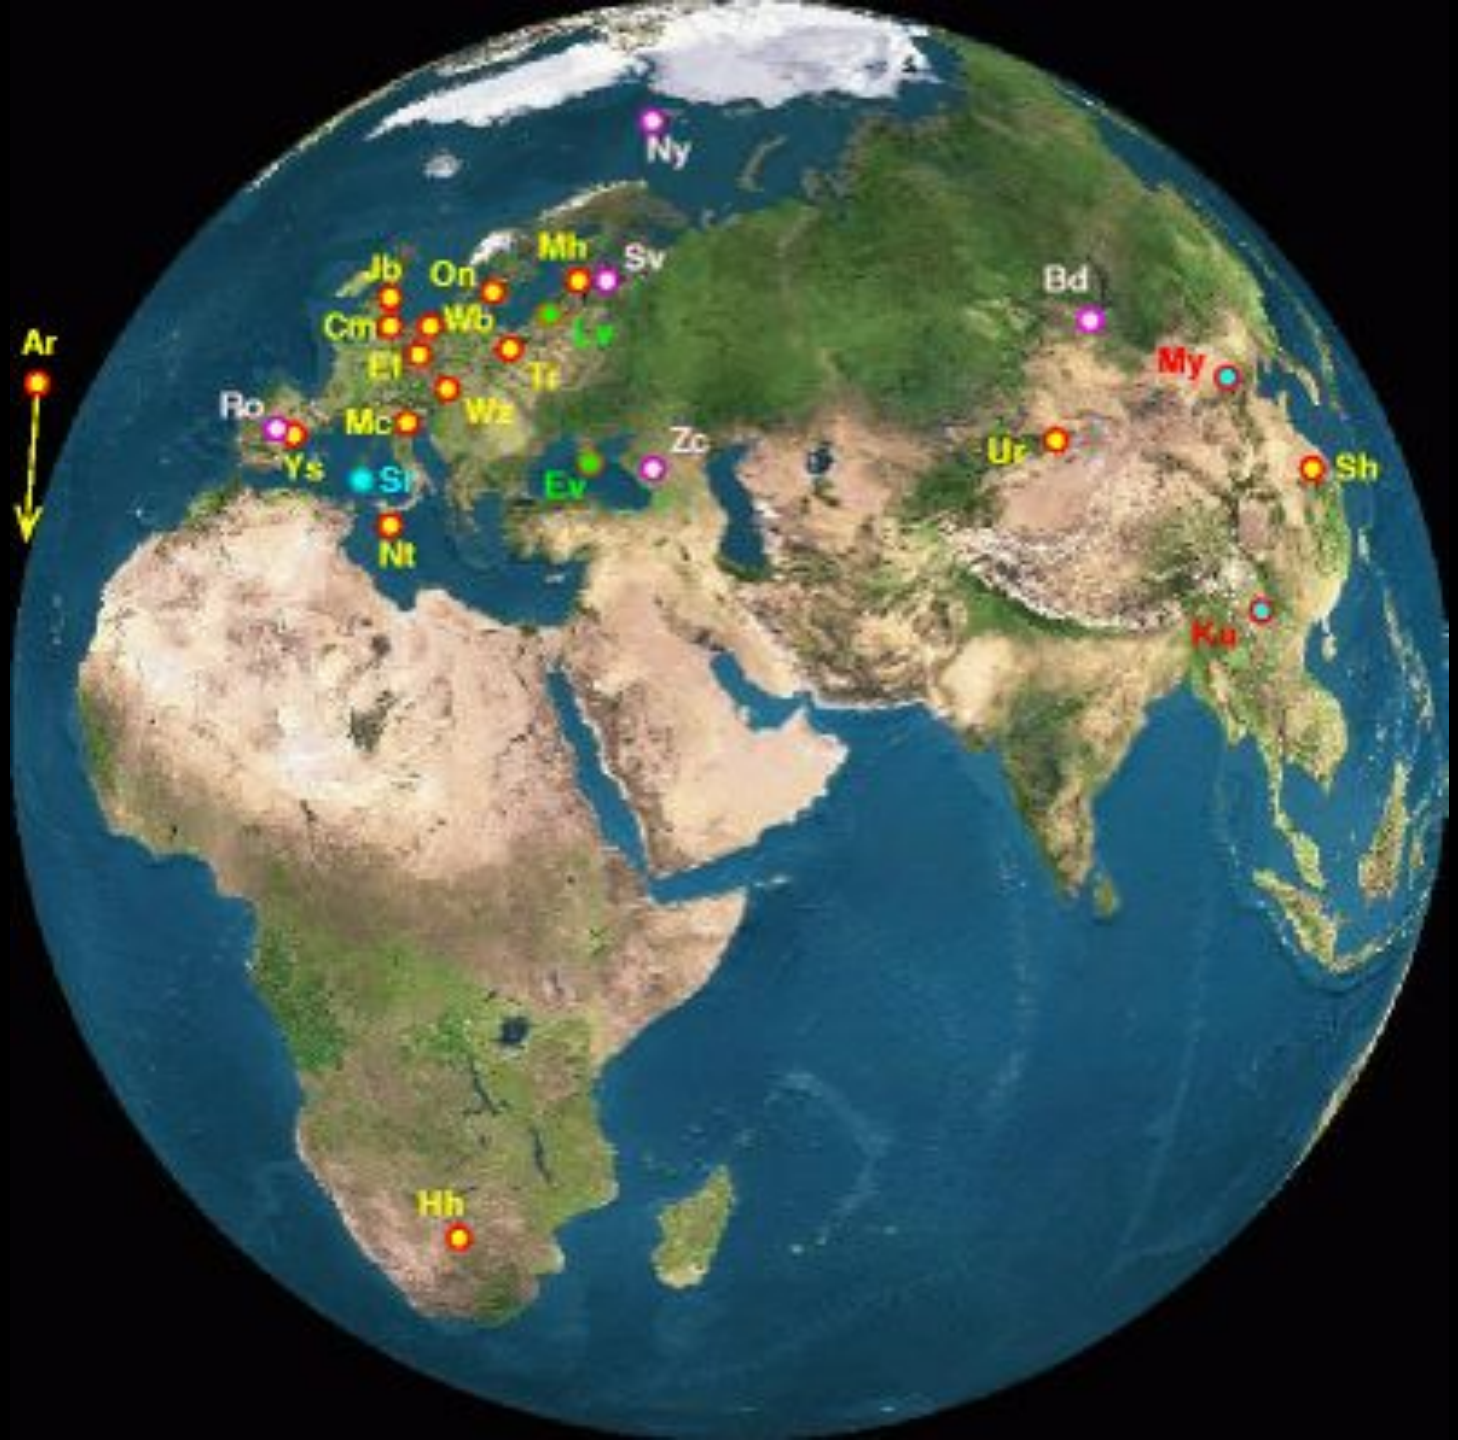
Ny
Mh
Sv
Jb
On
Cm
Wb
Lv
El
Tr
Wz
Ro
Mc
Ys
Sl
Ev
Zc
Nt
Ar
Bd
My
Ur
Sh
Ka
Hh

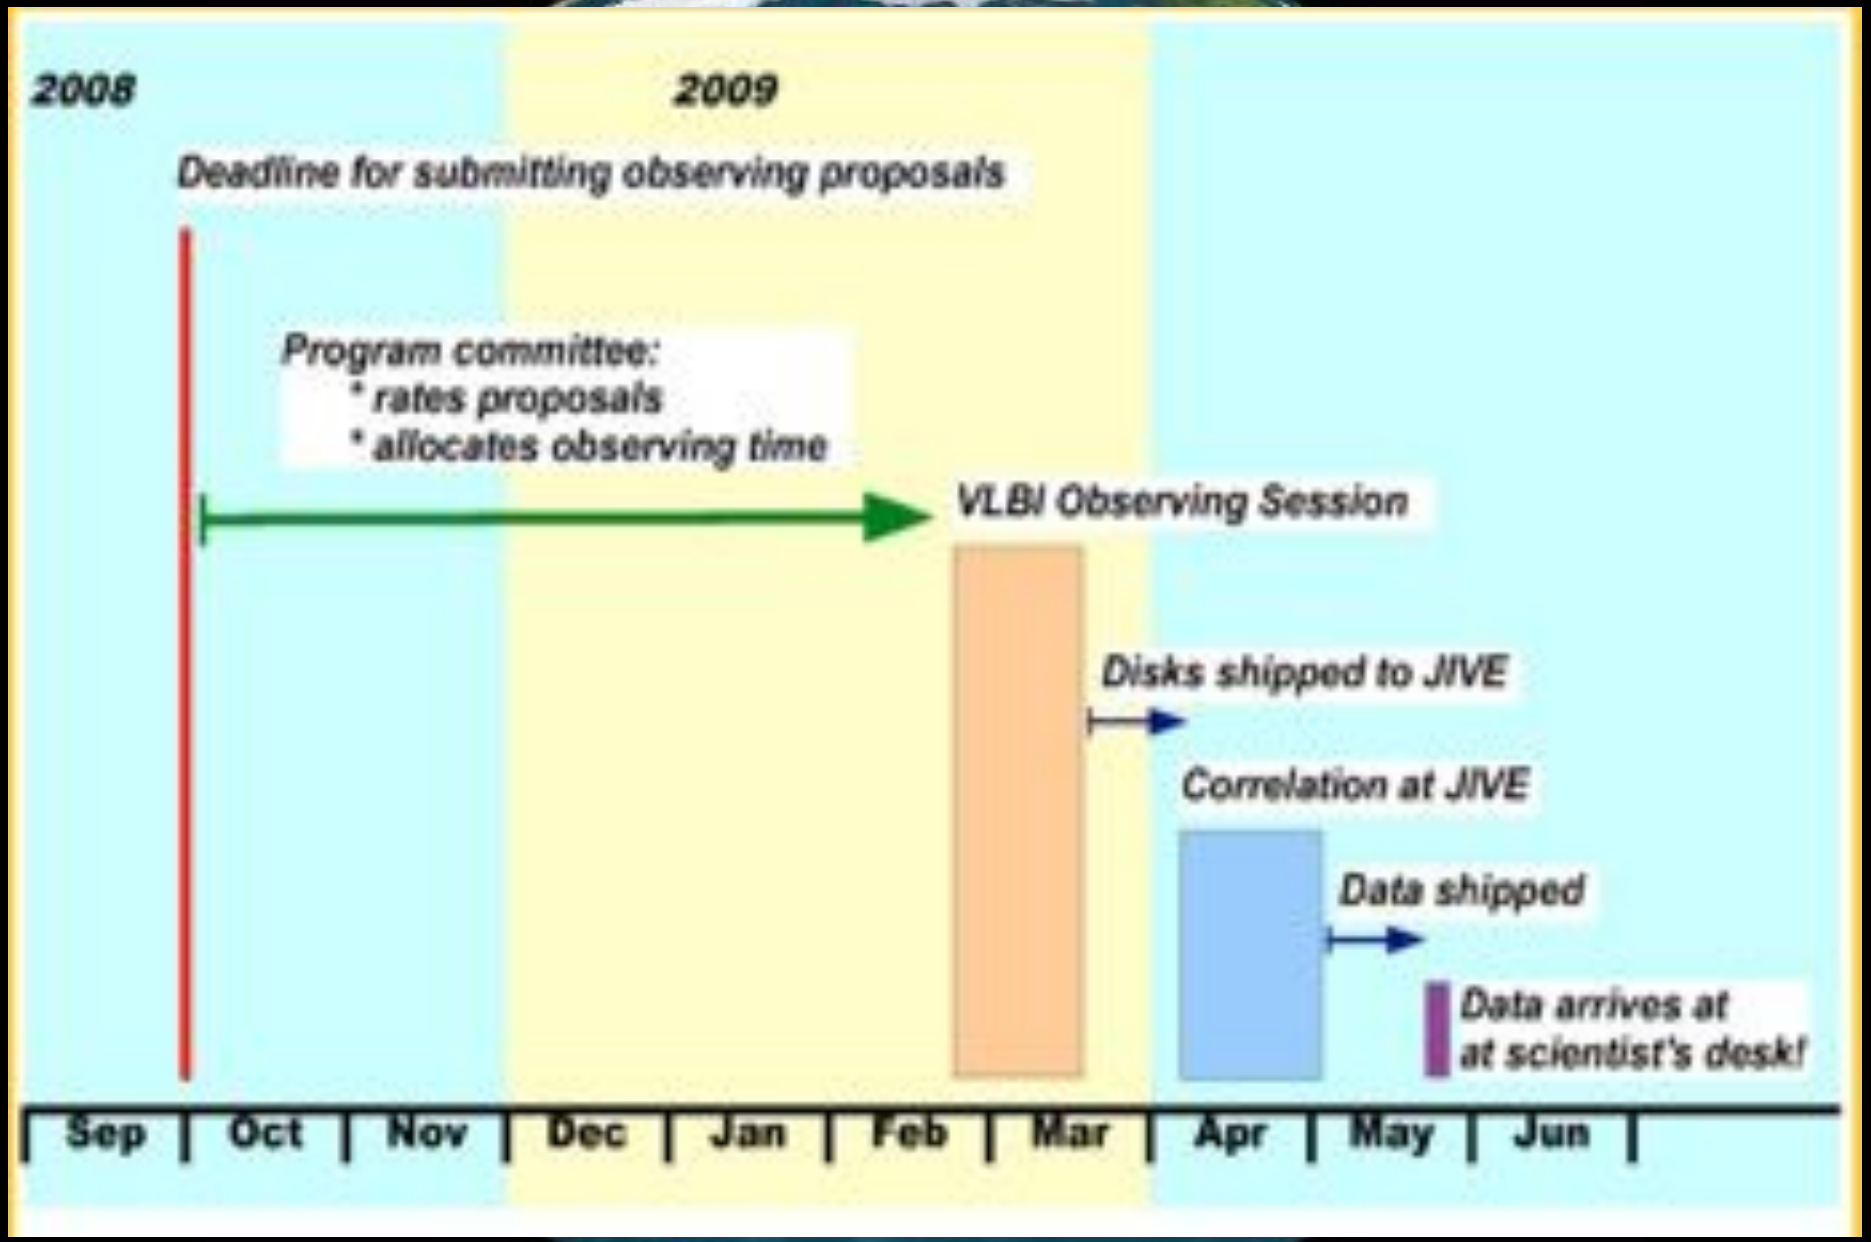
2008
2009
Deadline for submitting observing proposals
Program committee:
* rates proposals
* allocates observing time
VLBI Observing Session
Disks shipped to JIVE
Correlation at JIVE
Data shipped
Data arrives at
at scientist's desk!
Sep
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun

2008

2009

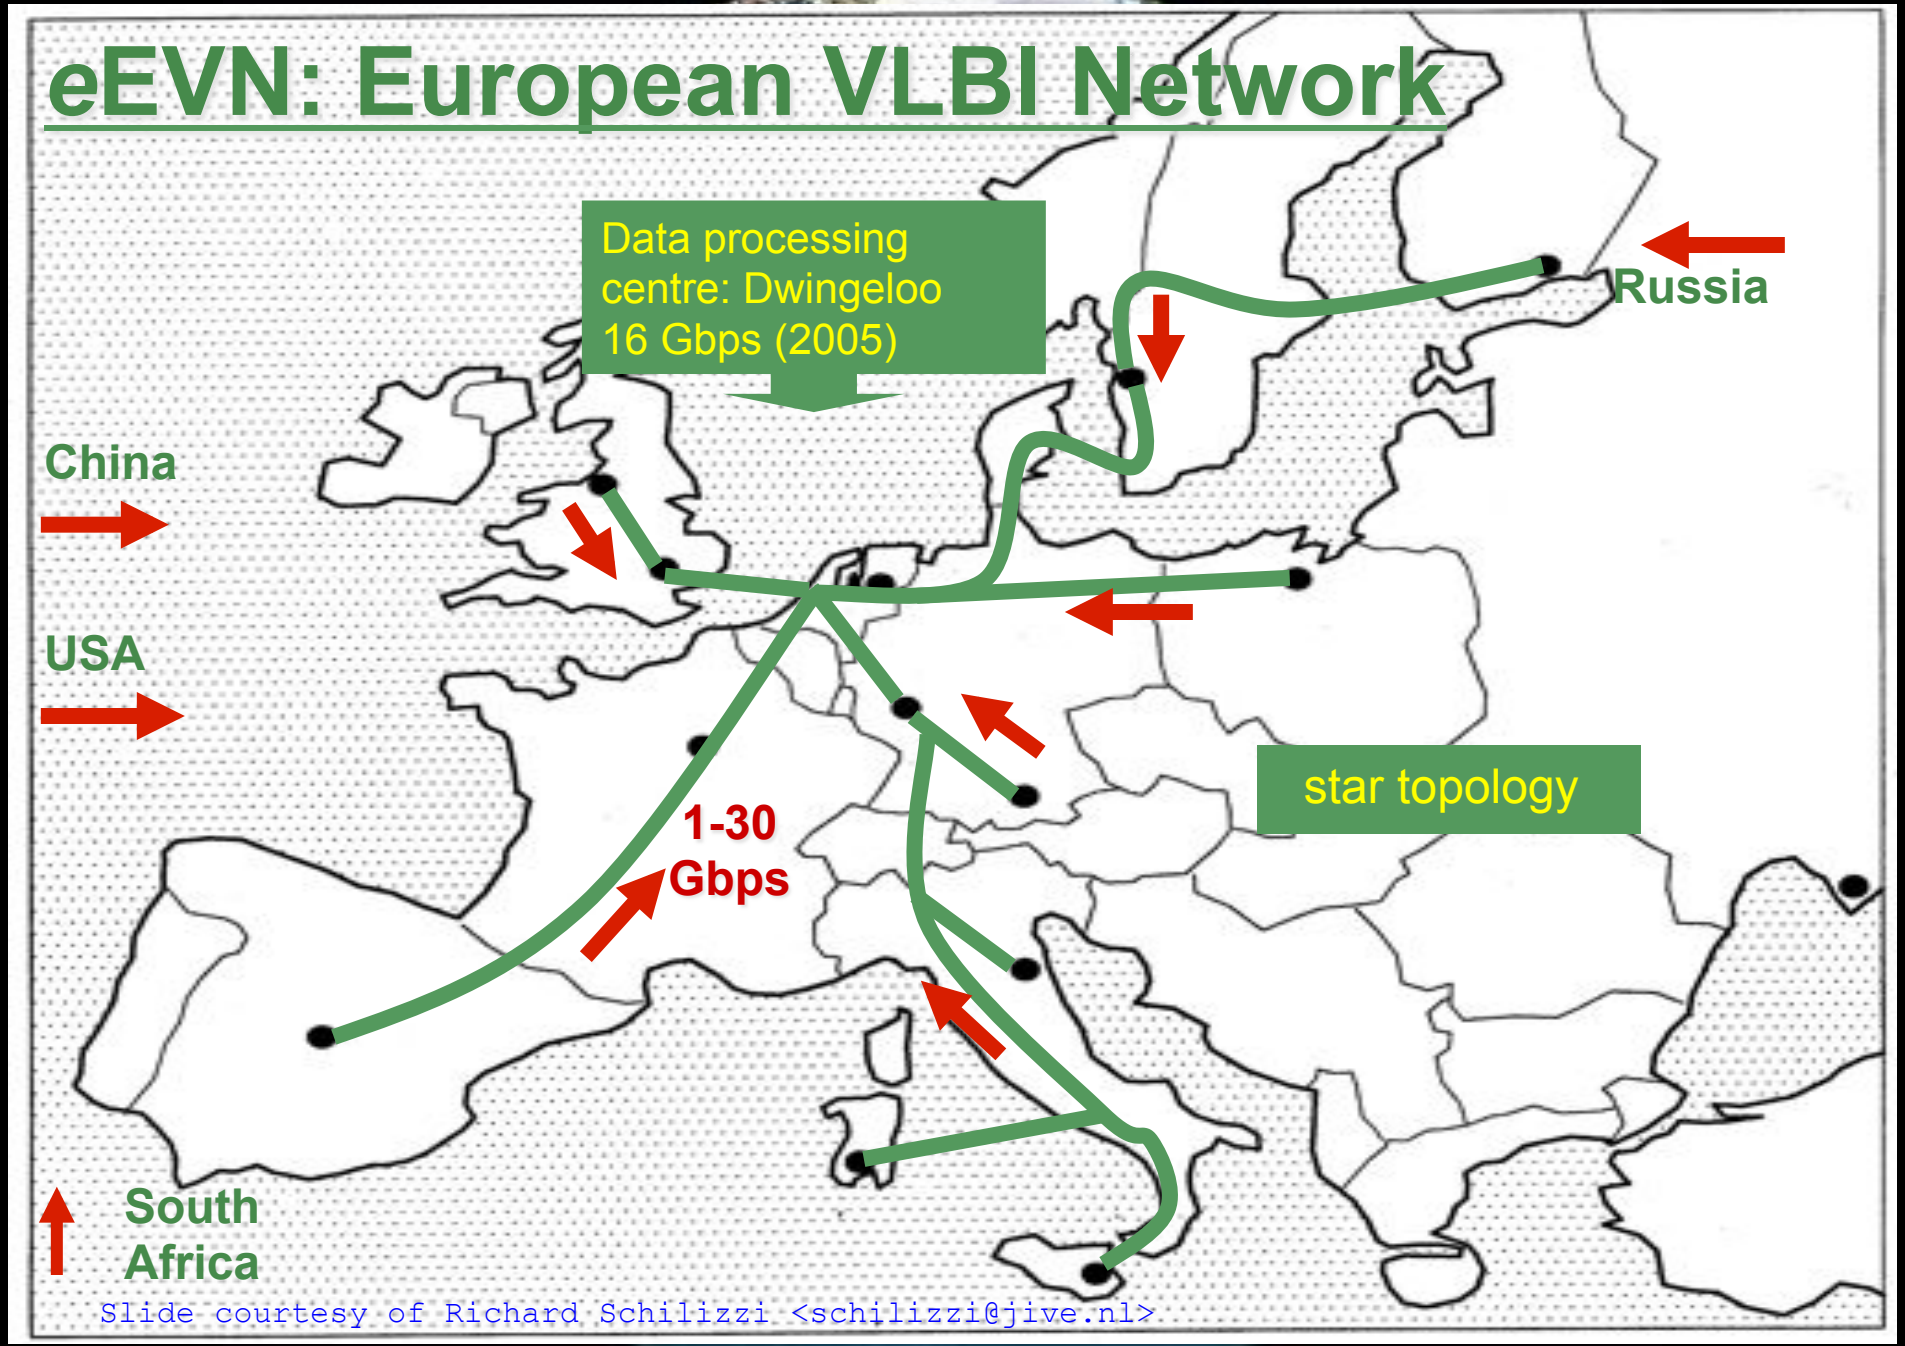
eEVN: European VLBI Network
Data processing centre: Dwingeloo
16 Gbps (2005)
Russia
China
USA
1-30 Gbps
star topology
South Africa
Slide courtesy of Richard Schilizzi <schilizzi@jive.nl>

eEVN: European VLBI Network

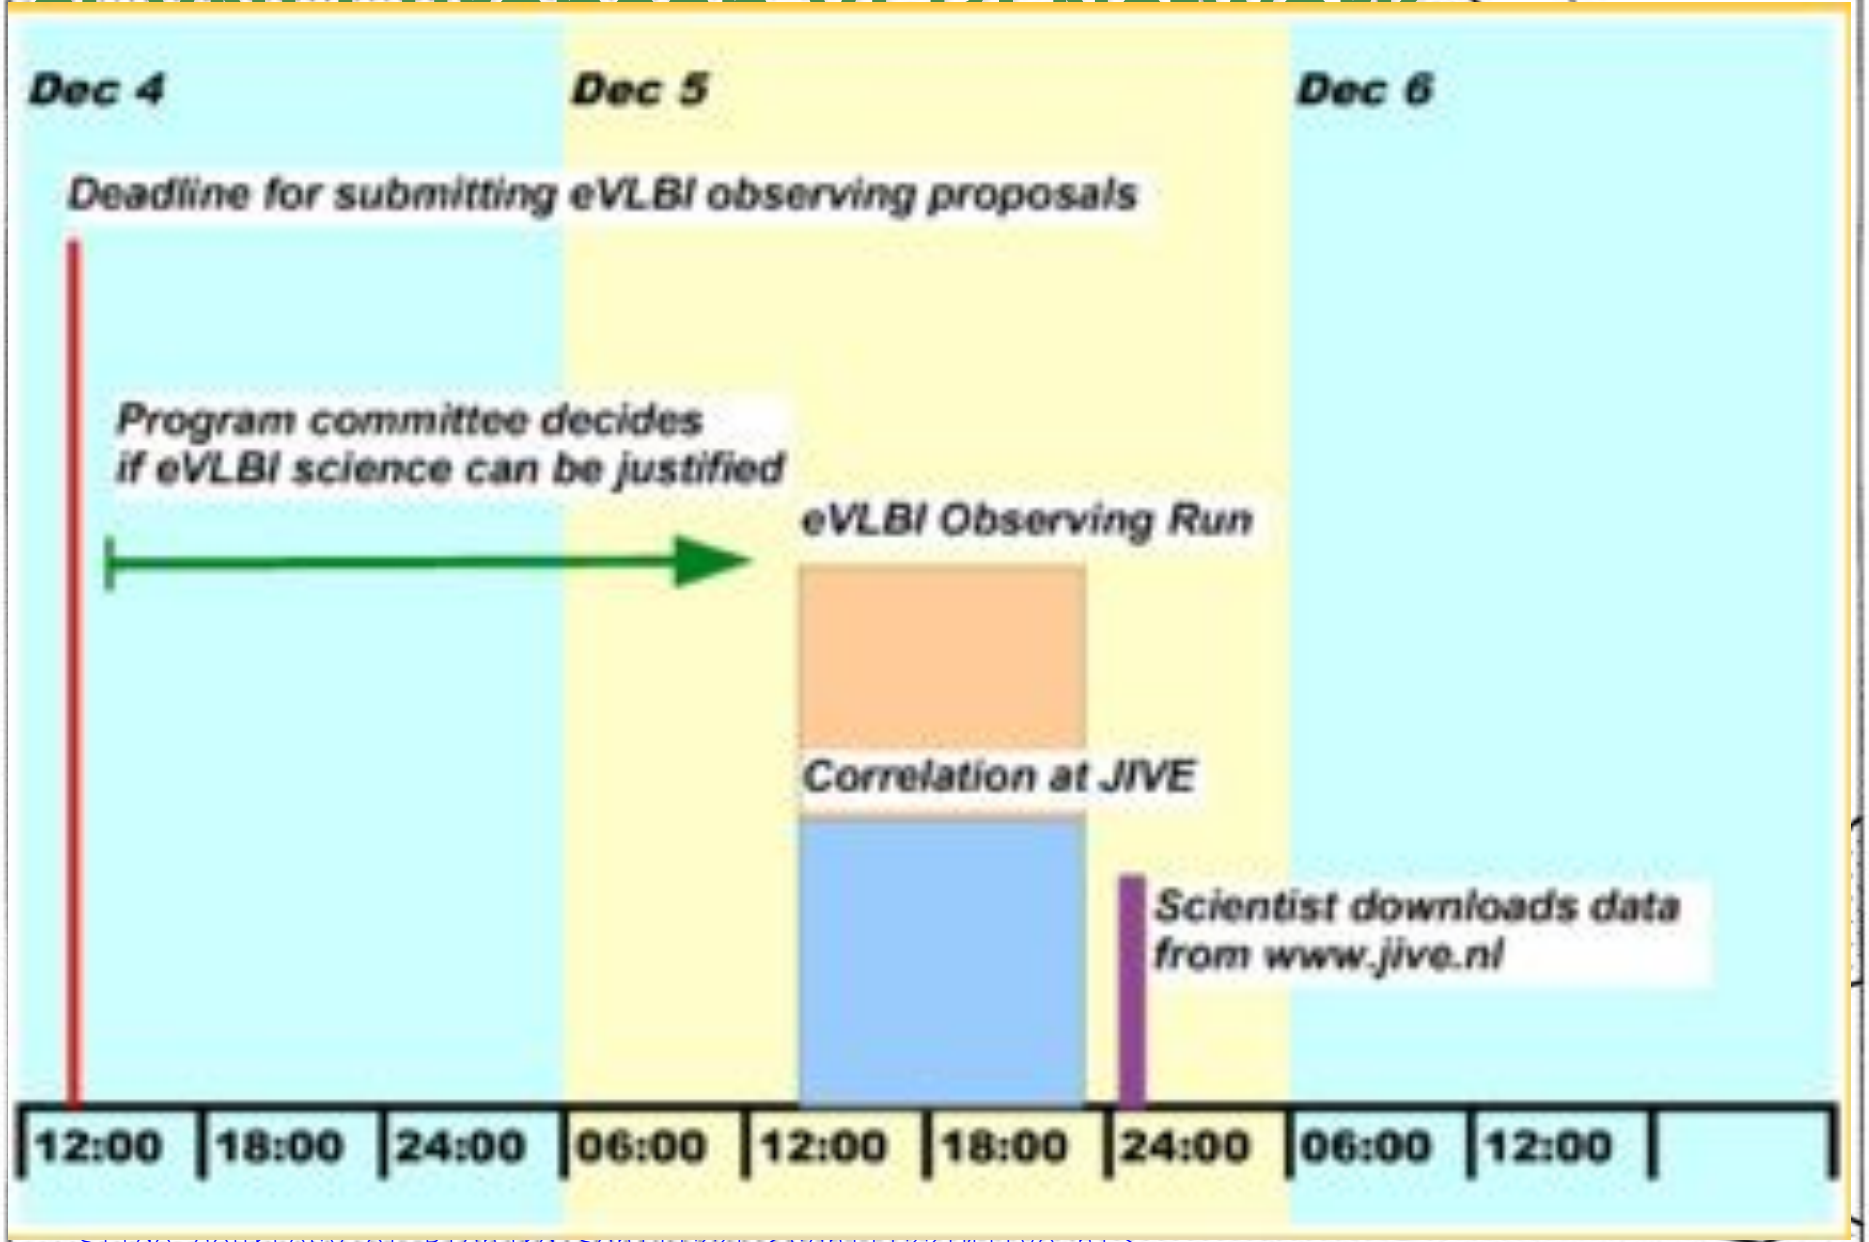

Slide courtesy of Richard Schilizzi <schilizzi@jive.nl>

The SCARIe project

SCARIe: a research project to create a Software Correlator for e-VLBI.
VLBI Correlation: signal processing technique to get high precision image from spatially distributed radio-telescope.

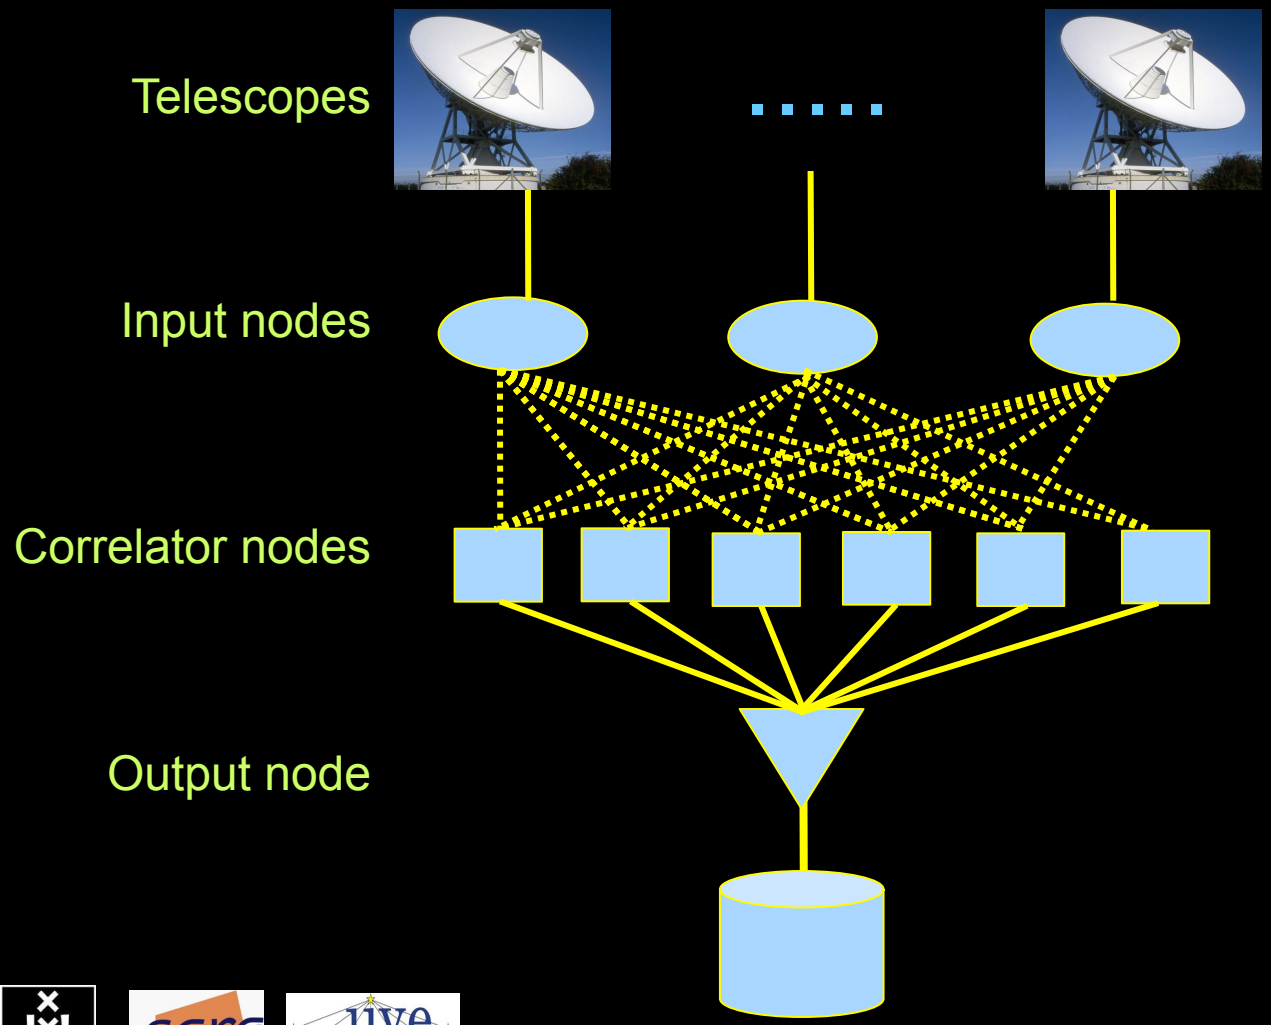

16 Gbit/s - 2 Tflop ➔ THIS IS A DATA FLOW PROBLEM !!!

Research:

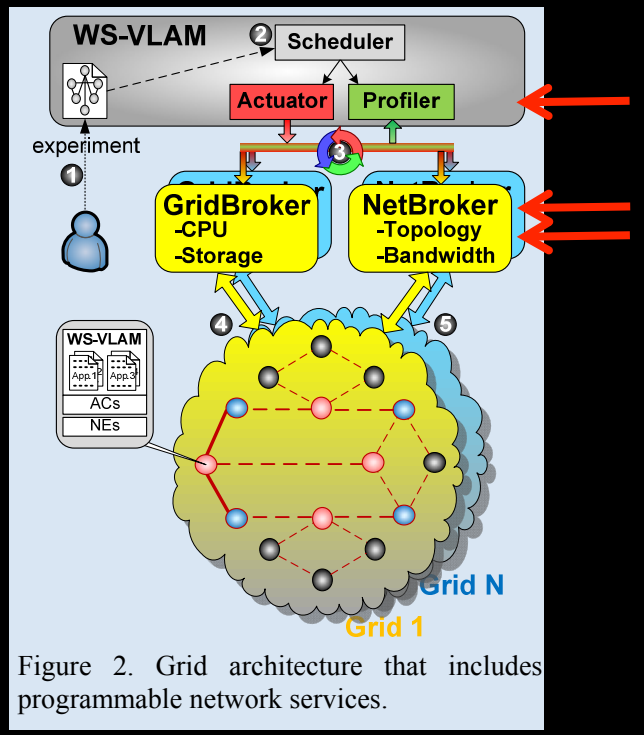

Figure 2. Grid architecture that includes programmable network services.

SNE @ UvA

	Ijkdijk/Urban Flood	Medical	LifeWatch/ENVRI	CosmoGrid/eVLBI	CineGrid	EU-GN3/NOVI/Geysers	SURFnet/GLIF/Cloud
Green-IT						X	X
Privacy/Trust		X			X		
Authorization/policy		X	X		X	X	
Programmable networks	X		X				
40-100Gig/TCP/WF/QoS	X		X	X		X	
Topology/Architecture		X		X	X	X	
Optical Photonic		X	X		X		

IJKDIJK

Sensors: 15000km* 800 bps/m ->12 Gbit/s to cover all Dutch dikes

Sensor grid: instrument the dikes

First controlled breach occurred on sept 27th ‘08:

Tera-Thinking

- What constitutes a Tb/s network?
- think back to teraflop computing!
  - MPI turns a room full of pc's in a teraflop machine
- massive parallel channels in hosts, NIC's
- TeraApps programming model supported by
  - TFlops -> MPI / Globus / Cloud
  - TBytes -> DAIS / MONETdb …
  - TPixels -> SAGE
  - TSensors -> LOFAR, LHC, LOOKING, CineGrid, ...
  - Tbit/s -> ?

  - ? -> Programmable Networks

User Programmable Virtualized Networks.

- The network is virtualized as a collection of resources
- UPVNs enable network resources to be programmed as part of the application
- Mathematica interacts with virtualized networks using UPVNs and optimize network + computation

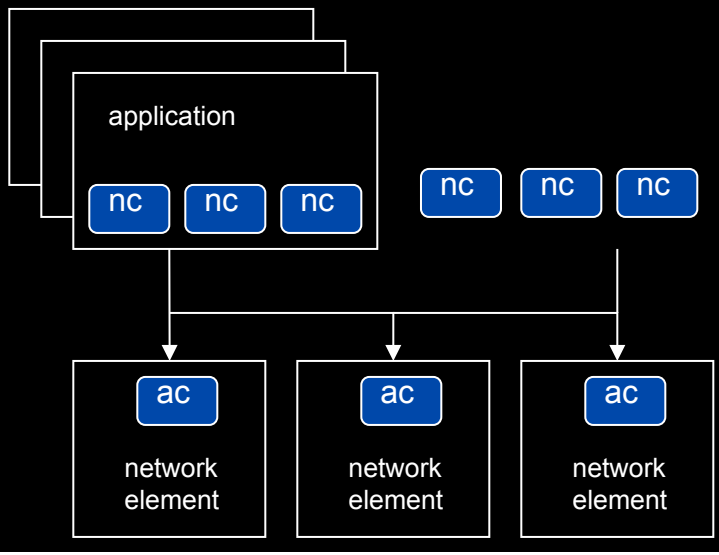

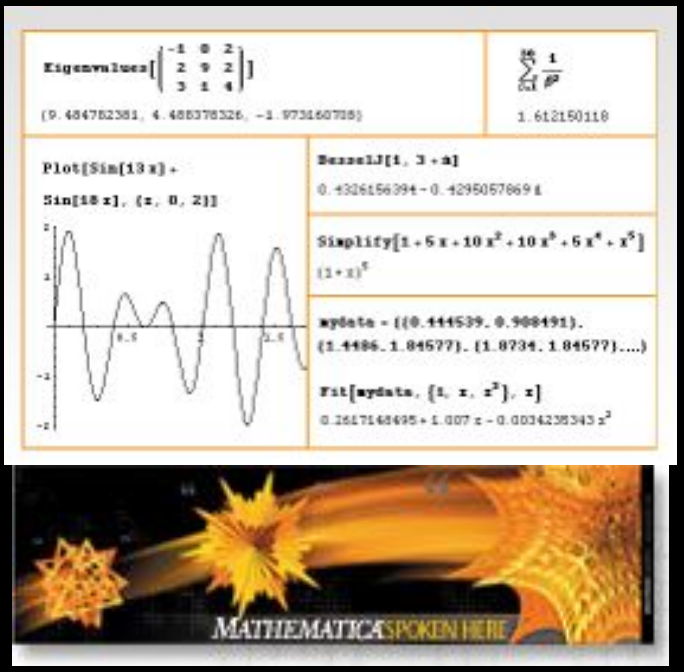

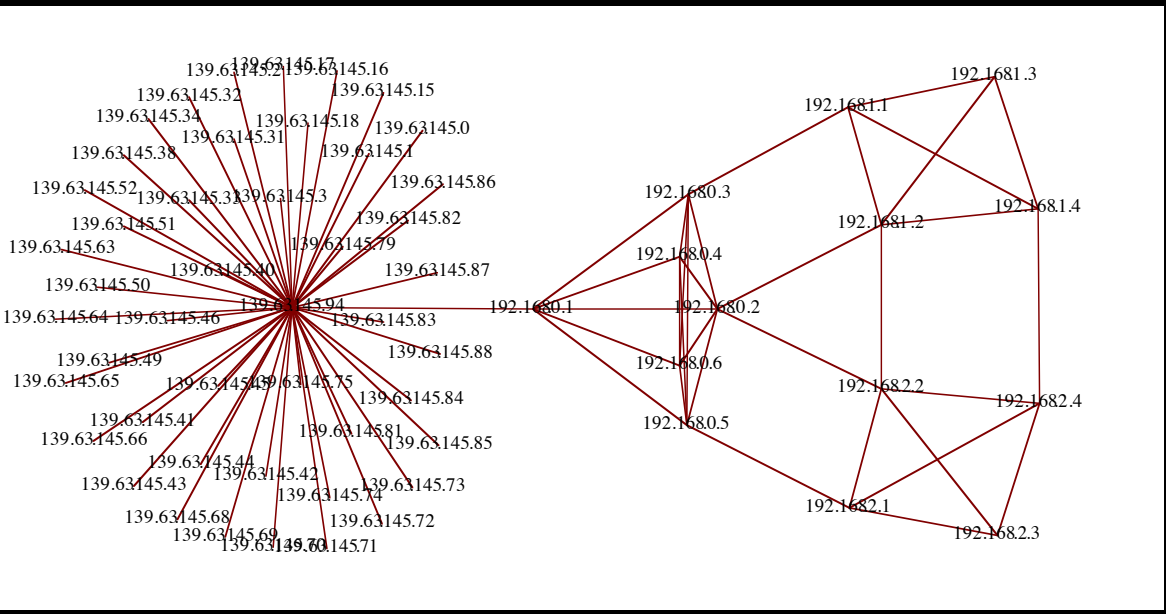

ref: Robert J. Meijer, Rudolf J. Strijkers, Leon Gommans, Cees de Laat, User Programmable Virtualiized Networks, accepted for publication to the IEEE e-Science 2006 conference Amsterdam.

TouchTable Demonstration @ SC08

SNE @ UvA

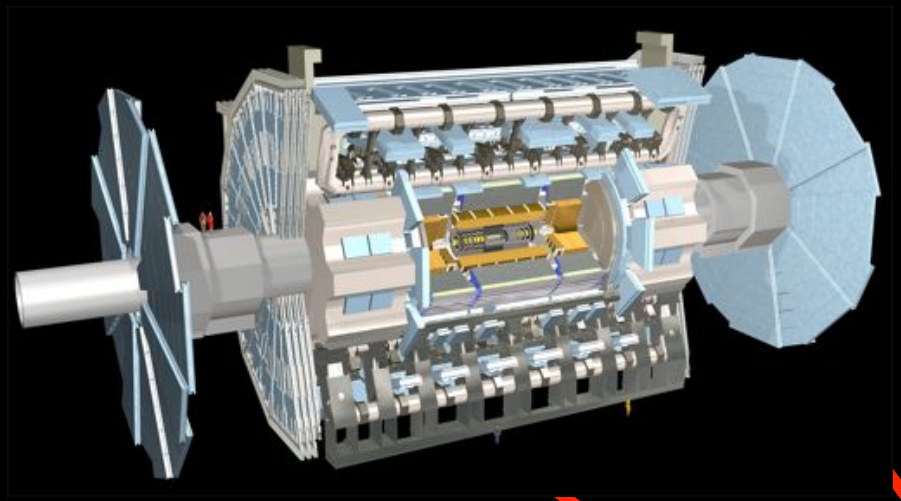

	Ijkdijk/Urban Flood	Medical	LifeWatch/ENVRI	CosmoGrid/eVLBI	CineGrid	EU-GN3/NOVI/Geysers	SURFnet/GLIF/Cloud
Green-IT						X	X
Privacy/Trust			X			X	
Authorization/policy		X	X		X	X	
Programmable networks	X		X				
40-100Gig/TCP/WF/QoS	X		X	X		X	
Topology/Architecture		X		X	X	X	
Optical Photonic		X	X		X		

ATLAS detector @ CERN Geneve

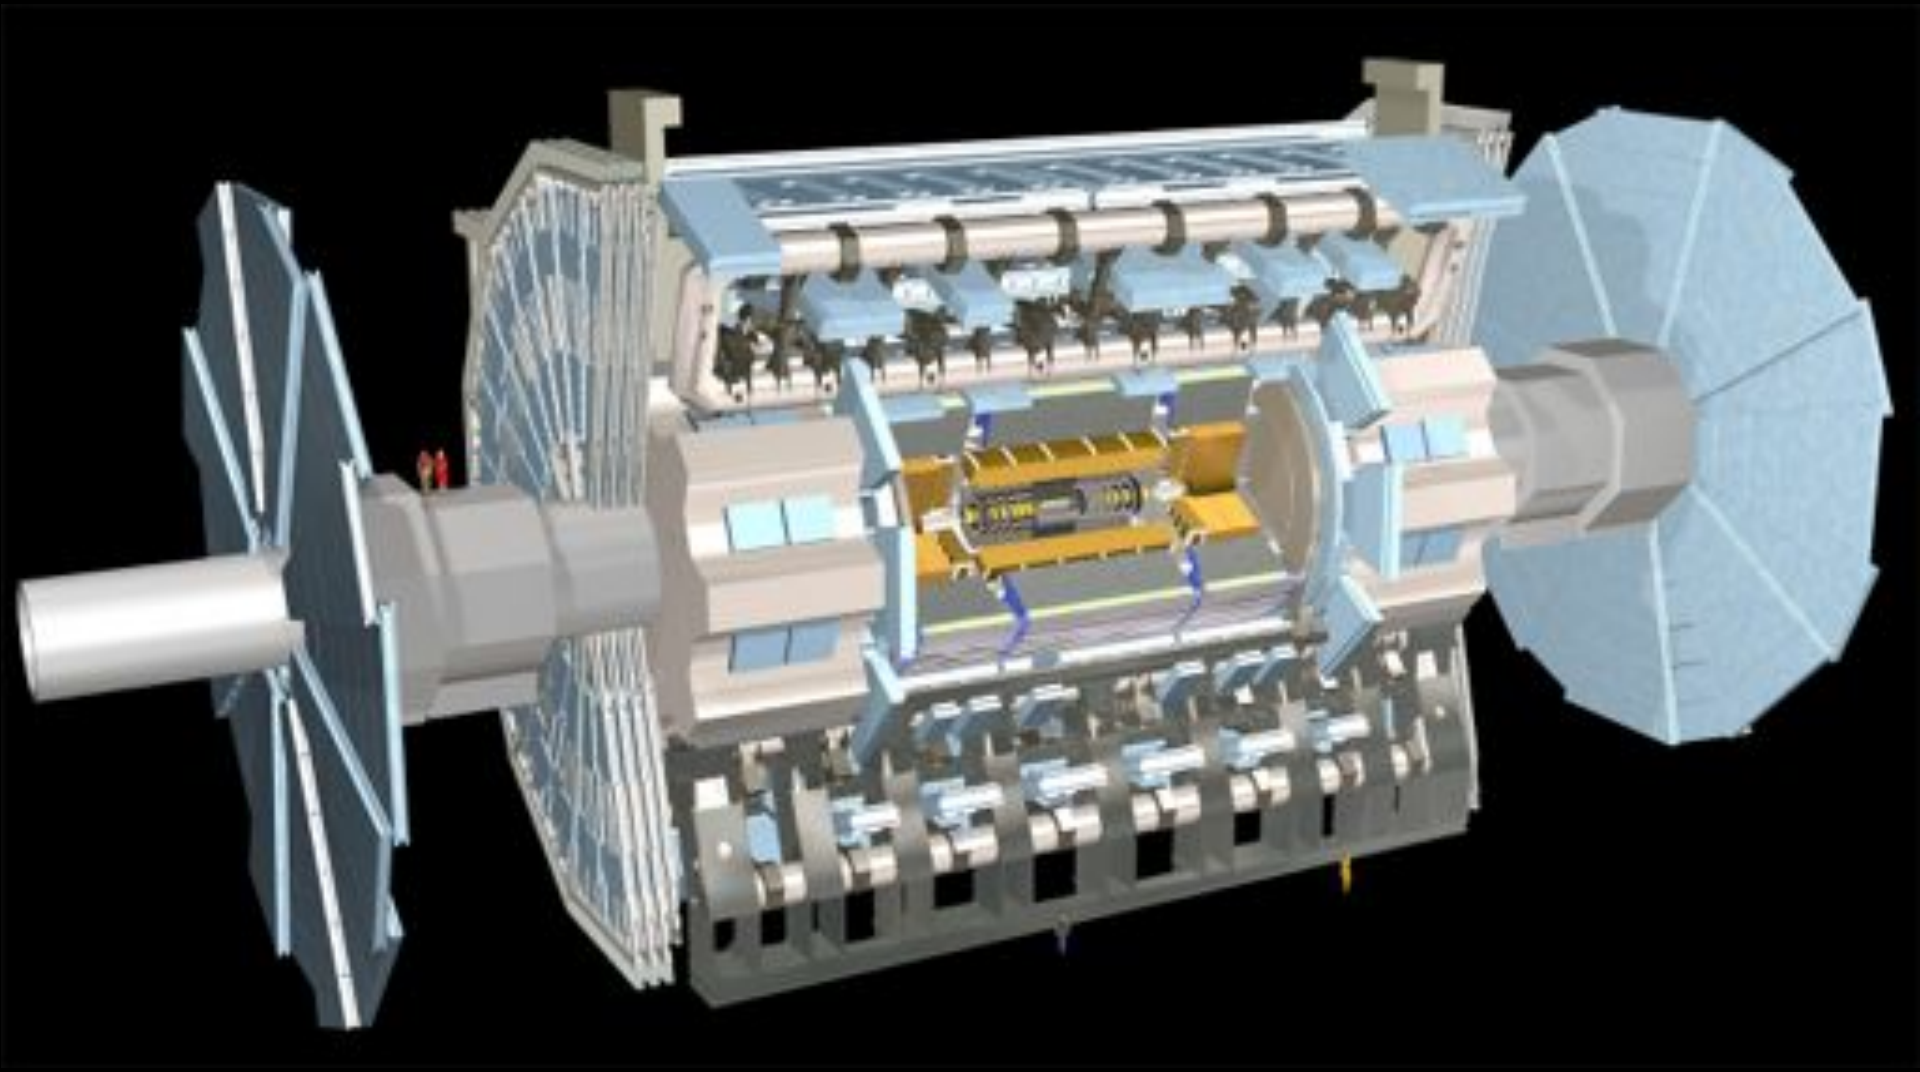

ATLAS detector @ CERN Geneve

One Heavy Ion Collision in Atlas!

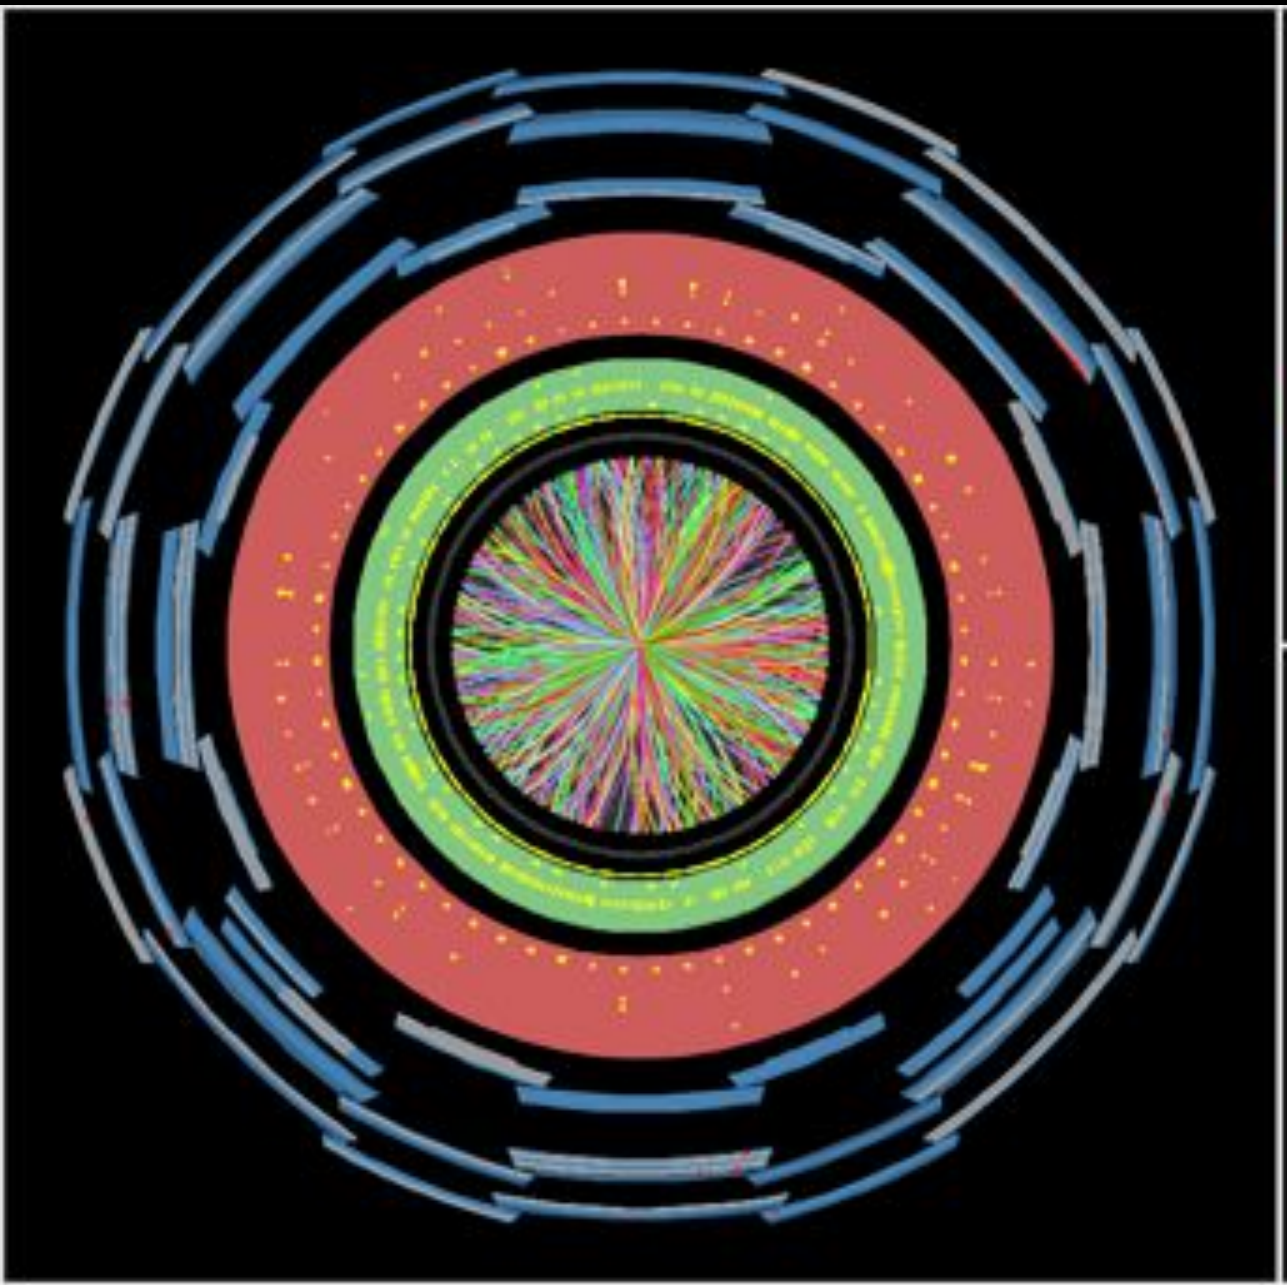

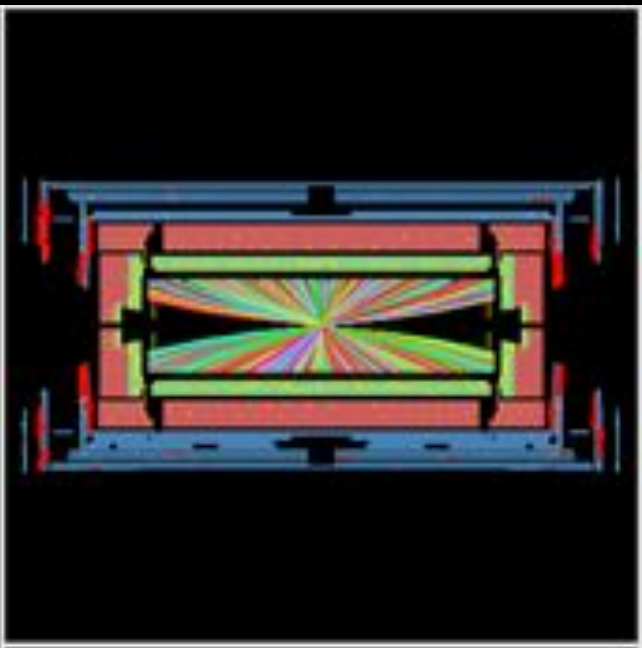

ATLAS EXPERIMENT

Run Number: 170482, Event Number: 3936308

Date: 2010-12-06 17:21:31 CET

Snapshot of a heavy ion collision directly from the ATLAS experiment

LHC Data Grid Hierarchy

CMS as example, Atlas is similar

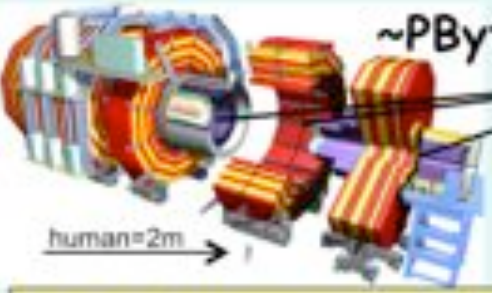

~PByte/sec

Online System

~100 MBytes/sec

100000 flops/byte

10 Pflops/s

event simulation

Tier 0 +1

HPSS

event reconstruction

CMS detector: 15m X 15m X 22m

12,500 tons, $700M.

~2.5 Gbits/sec

Status 2002!

Tier 1

Italian Regional Center

German Regional Center

NIKHEF Dutch Regional Center

FermiLab, USA Regional Center

...

~0.6-2.5 Gbps

analysis

Tier2 Center

Tier 2

~0.6-2.5 Gbps

Tier 3

Institute ~0.25TIPS

Physics data cache

100 - 1000 Mbits/sec

Tier 4

Workstations

Courtesy Harvey Newman, CalTech and CERN

CERN/CMS data goes to 6-8 Tier 1 regional centers, and from each of these to 6-10 Tier 2 centers.

Physicists work on analysis "channels" at 135 institutes. Each institute has ~10 physicists working on one or more channels.

2000 physicists in 31 countries are involved in this 20-year experiment in which DOE is a major player.

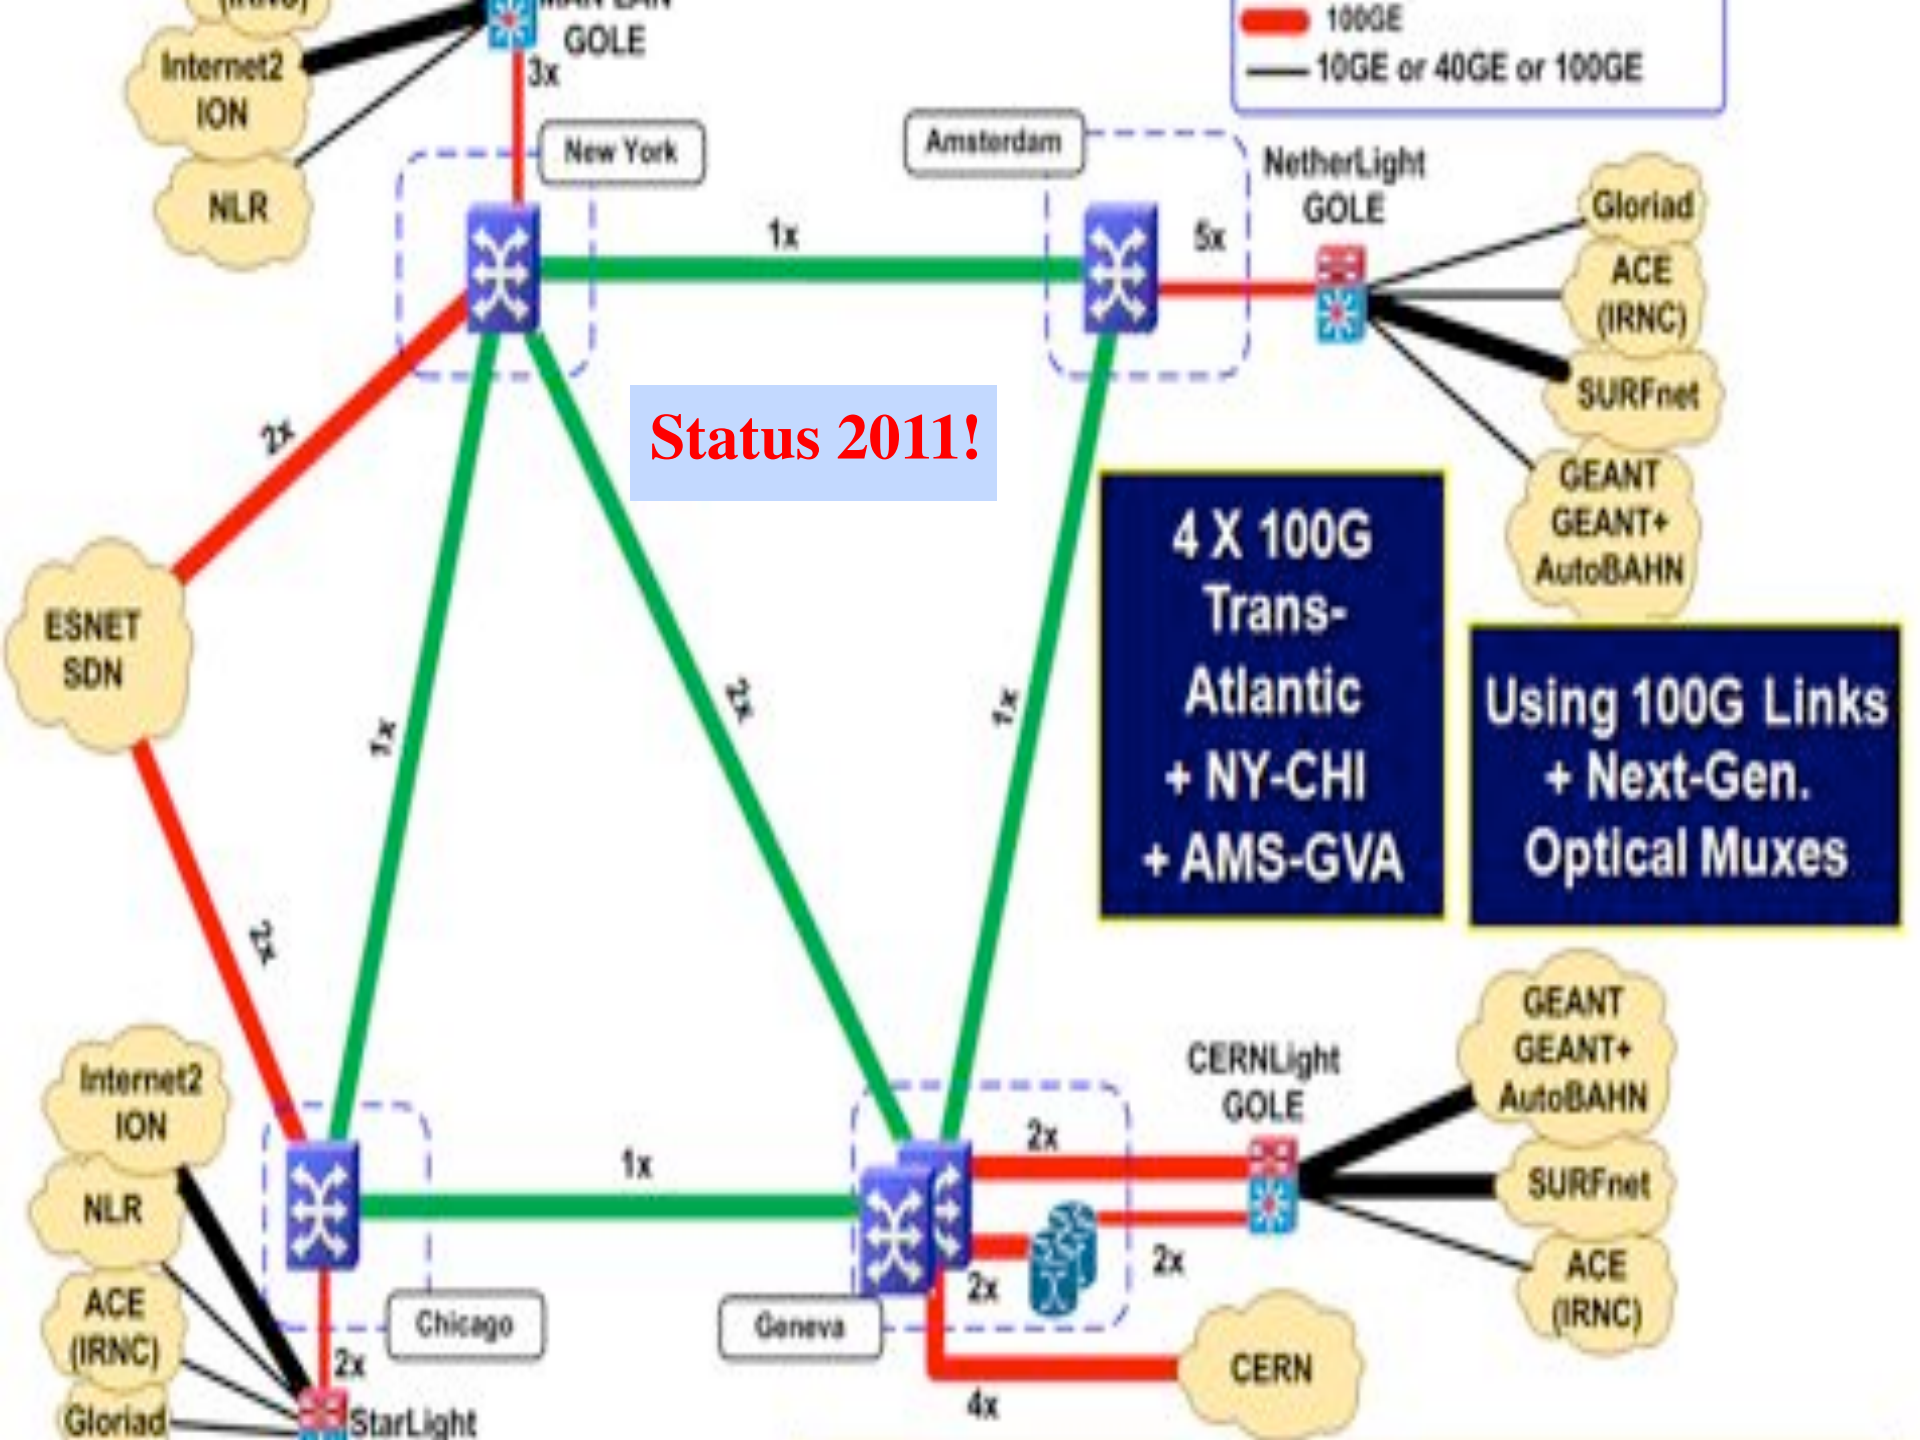
100GE
10GE or 40GE or 100GE
Internet2 ION
NLR
GOLE
3x
New York
Amsterdam
NetherLight GOLE
1x
5x
Gloriad
ACE (IRNC)
SURFnet
GEANT GEANT+ AutoBAHN
Status 2011!
2x
ESNET SDN
1x
2x
1x
4 X 100G Trans-Atlantic + NY-CHI + AMS-GVA
Using 100G Links + Next-Gen. Optical Muxes
4x
Internet2 ION
NLR
ACE (IRNC)
Gloriad
2x
Chicago
StarLight
1x
Geneva
2x
2x
2x
CERNLight GOLE
GEANT GEANT+ AutoBAHN
SURFnet
ACE (IRNC)
CERN
4x

Big and small flows don’t go well together on the same wire! ☹

Diagram for SAGE video streaming to ATS

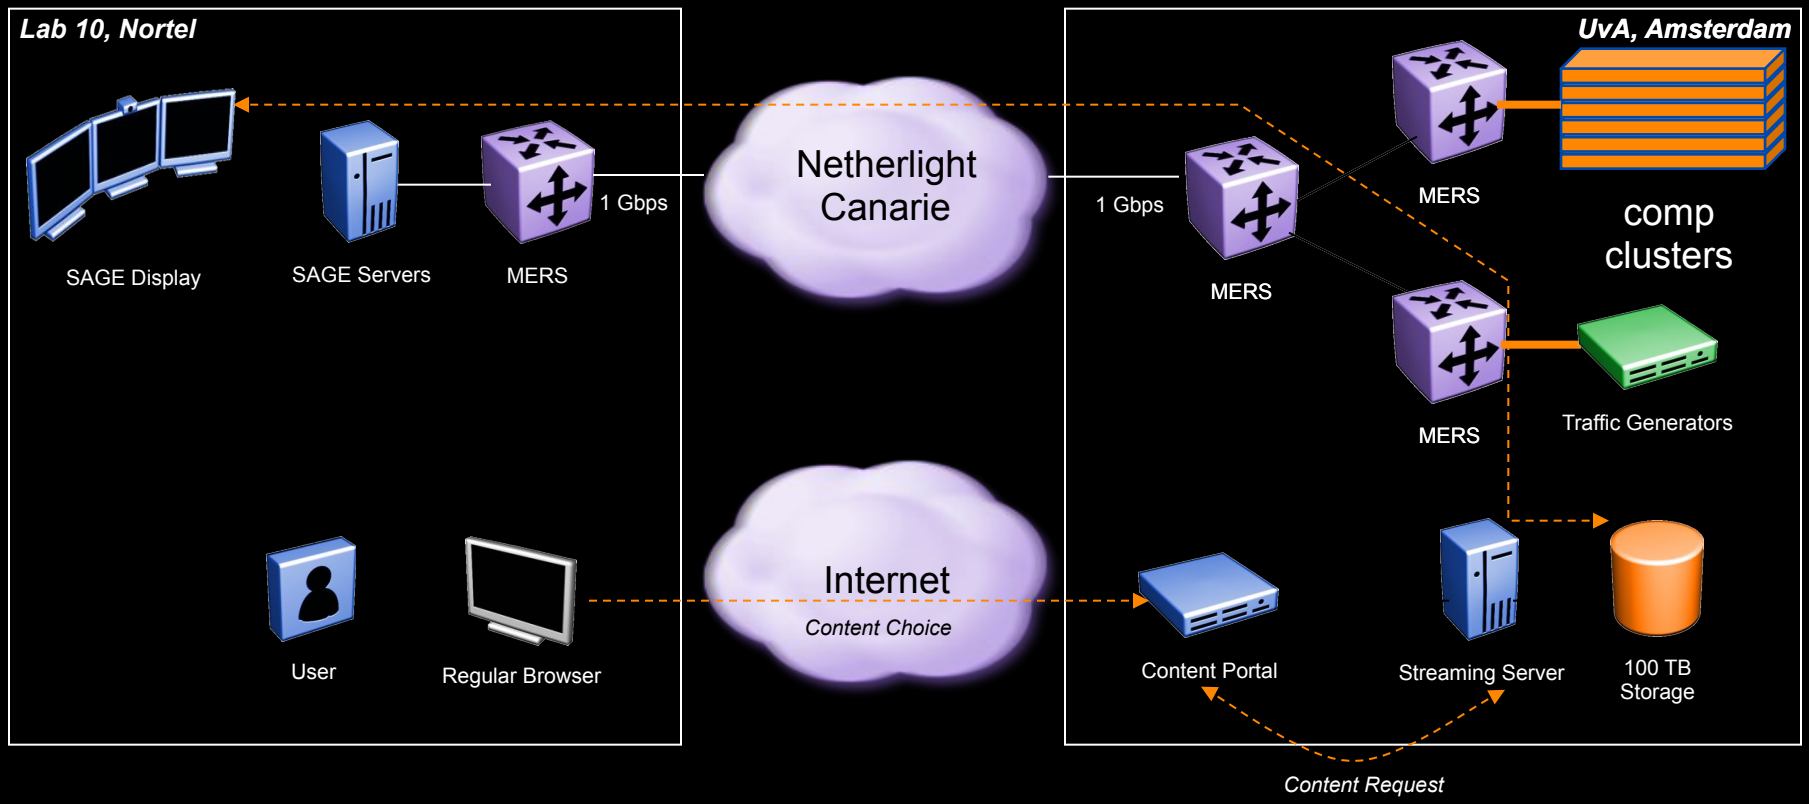

Experimental Data

Sage without background traffic

Sage with background traffic

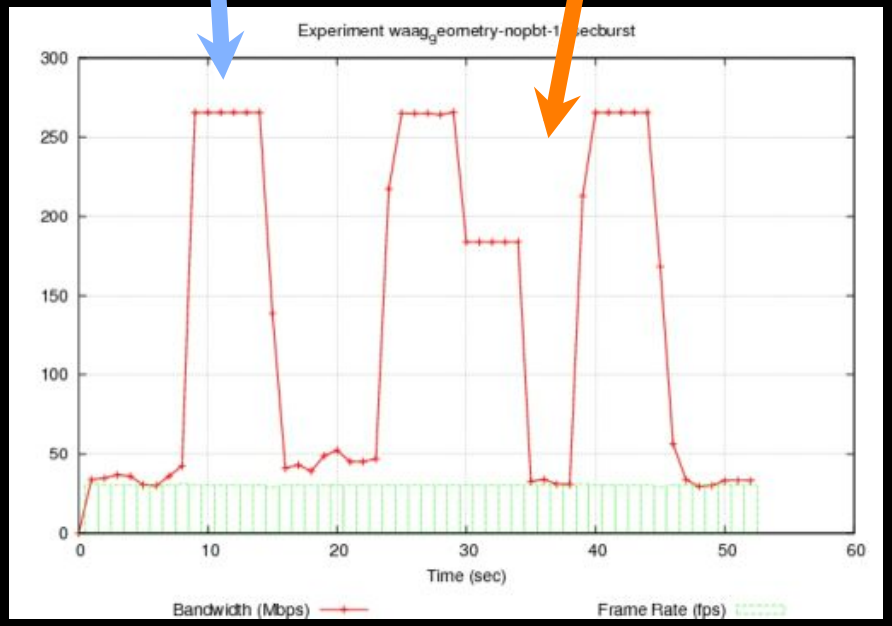

10 Second Traffic bursts with No PBT

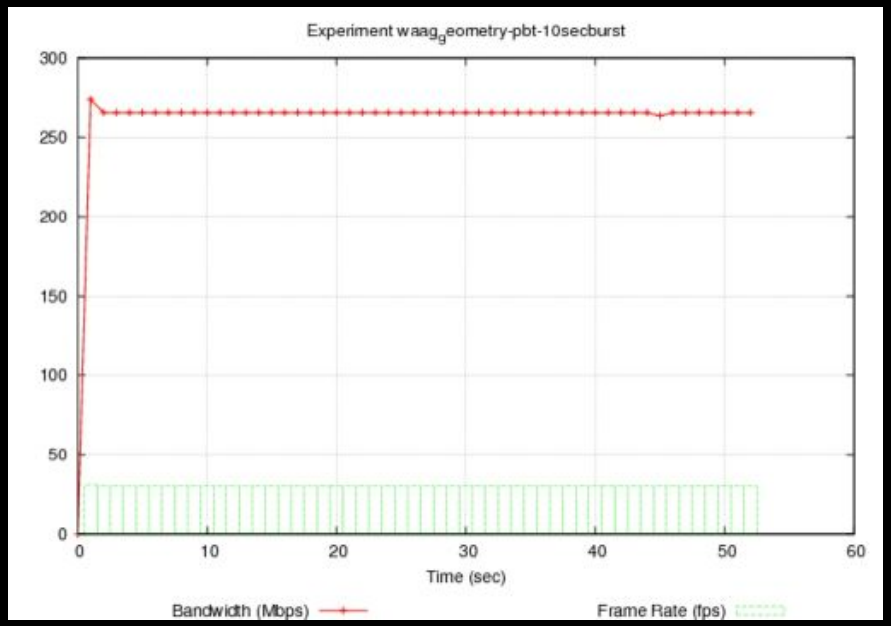

10 Second Traffic bursts with PBT

PBT is *SIMPLE* and *EFFECTIVE* technology to build a shared Media-Ready Network

Alien light From idea to realisation!

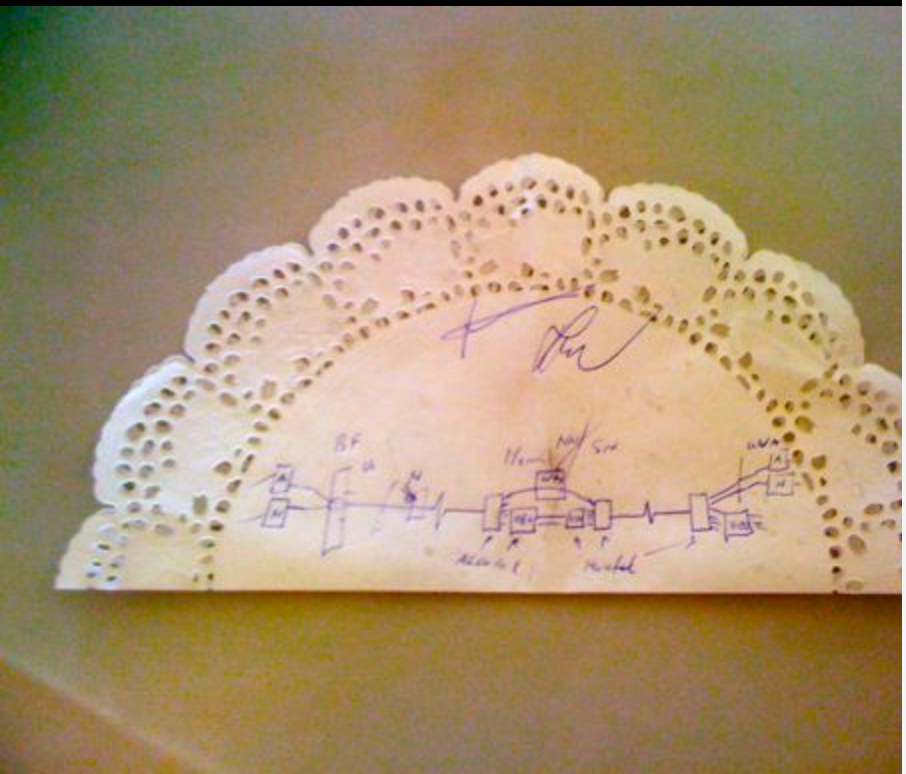

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

NCF

Alien wavelength advantages

- Direct connection of customer equipment[1] → cost savings
- Avoid OEO regeneration → power savings
- Faster time to service[2] → time savings
- Support of different modulation formats[3] → extend network lifetime

Alien wavelength challenges

- Complex end-to-end optical path engineering in terms of linear (i.e. OSNR, dispersion) and non-linear (FWM, SPM, XPM, Raman) transmission effects for different modulation formats.
- Complex interoperability testing.
- End-to-end monitoring, fault isolation and resolution.
- End-to-end service activation.

In this demonstration we will investigate the performance of a 40Gb/s PM-QPSK alien wavelength installed on a 10Gb/s DWDM infrastructure.

New method to present fiber link quality, FoM (Figure of Merit)

In order to quantify optical link grade, we propose a new method of representing system quality: the FOM (Figure of Merit) for concatenated fiber spans.

$$FOM = \sum_{j=1}^{N} 10^{\left(\frac{L_j}{10}\right)}$$

Lj, span losses in dB
N, number of spans

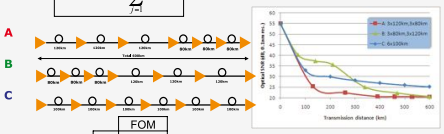

	FOM
A	5504
B	5504
C	1897

Easy-to-use formula that accurately quantifies transmission system performance

Transmission system setup

JOINT SURFnet/NORDUnet 40Gb/s PM-QPSK alien wavelength DEMONSTRATION.

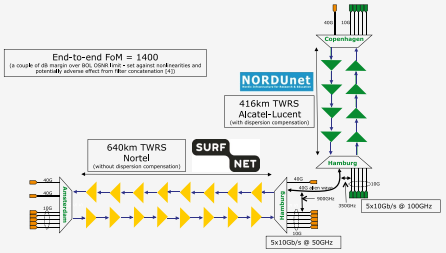

Test results

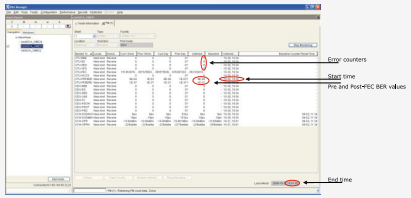

Error-free transmission for 23 hours, 17 minutes → BER < 3.0 10^{-16}

Conclusions

- We have investigated experimentally the all-optical transmission of a 40Gb/s PM-QPSK alien wavelength via a concatenated native and third party DWDM system that both were carrying live 10Gb/s wavelengths.
- The end-to-end transmission system consisted of 1056 km of TWRS (TrueWave Reduced Slope) transmission fiber.
- We demonstrated error-free transmission (i.e. BER below 10-15) during a 23 hour period.
- More detailed system performance analysis will be presented in an upcoming paper.

NORTEL NORDUnet telindus together with ISIT SURF NET

REFERENCES [1] "OPERATIONAL SOLUTIONS FOR AN OPEN DWDM LAYER", O. GERSTEL ET AL, OFC'2009 | [2] "AT&T OPTICAL TRANSPORT SERVICES", BARBARA E. SMITH, OFC'09 [3] "OPEX SAVINGS OF ALL-OPTICAL CORE NETWORKS", ANDREW LORD AND CARL ENGINEER, ECOC2009 | [4] NORTEL/SURFNET INTERNAL COMMUNICATION

ACKNOWLEDGEMENTS WE ARE GRATEFUL TO NORDUNET FOR PROVIDING US WITH BANDWIDTH ON THEIR DWDM LINK FOR THIS EXPERIMENT AND ALSO FOR THEIR SUPPORT AND ASSISTANCE DURING THE EXPERIMENTS. WE ALSO ACKNOWLEDGE TELINDUS AND NORTEL FOR THEIR INTEGRATION WORK AND SIMULATION SUPPORT

Alien light From idea to realisation!

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

NCF

Alien wavelength advantages

- Direct connection of customer equipment[1] → cost savings
- Avoid OEO regeneration → power savings
- Faster time to service[2] → time savings
- Support of different modulation formats[3] → extend network lifetime

Alien wavelength challenges

- Complex end-to-end optical path engineering in terms of linear (i.e. OSNR, dispersion) and non-linear (FWM, SPM, XPM, Raman) transmission effects for different modulation formats.
- Complex interoperability testing.
- End-to-end monitoring, fault isolation and resolution.
- End-to-end service activation.

In this demonstration we will investigate the performance of a 40Gb/s PM-QPSK alien wavelength installed on a 10Gb/s DWDM infrastructure.

New method to present fiber link quality, FoM (Figure of Merit)

In order to quantify optical link grade, we propose a new method of representing system quality: the FOM (Figure of Merit) for concatenated fiber spans.

$$FOM = \sum_{j=1}^{N} 10^{\left(\frac{L_j}{10}\right)}$$

Lj, span losses in dB
N, number of spans

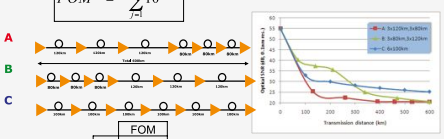

	FOM
A	5504
B	5504
C	1897

Easy-to-use formula that accurately quantifies transmission system performance

Transmission system setup

JOINT SURFnet/NORDUnet 40Gb/s PM-QPSK alien wavelength DEMONSTRATION.

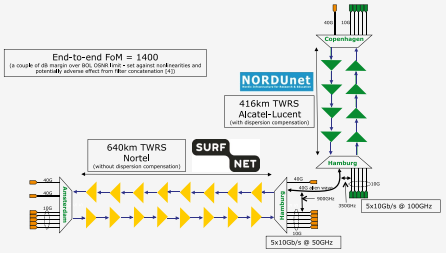

Test results

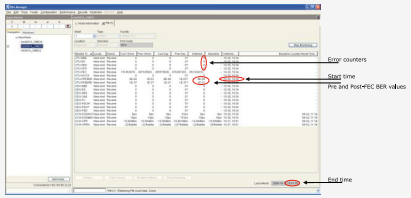

Error-free transmission for 23 hours, 17 minutes → BER < 3.0 10^{-16}

Conclusions

- We have investigated experimentally the all-optical transmission of a 40Gb/s PM-QPSK alien wavelength via a concatenated native and third party DWDM system that both were carrying live 10Gb/s wavelengths.
- The end-to-end transmission system consisted of 1056 km of TWRS (TrueWave Reduced Slope) transmission fiber.
- We demonstrated error-free transmission (i.e. BER below 10-15) during a 23 hour period.
- More detailed system performance analysis will be presented in an upcoming paper.

NORTEL NORDUnet telindus together with ISIT SURFNET

REFERENCES [1] "OPERATIONAL SOLUTIONS FOR AN OPEN DWDM LAYER", O. GERSTEL ET AL, OFC'2009 | [2] "AT&T OPTICAL TRANSPORT SERVICES", BARBARA E. SMITH, OFC'09
[3] "OPEX SAVINGS OF ALL-OPTICAL CORE NETWORKS", ANDREW LORD AND CARL ENGINEER, ECOC2009 | [4] NORTEL/SURFNET INTERNAL COMMUNICATION

ACKNOWLEDGEMENTS WE ARE GRATEFUL TO NORDUNET FOR PROVIDING US WITH BANDWIDTH ON THEIR DWDM LINK FOR THIS EXPERIMENT AND ALSO FOR THEIR SUPPORT AND ASSISTANCE DURING THE EXPERIMENTS. WE ALSO ACKNOWLEDGE TELINDUS AND NORTEL FOR THEIR INTEGRATION WORK AND SIMULATION SUPPORT

ClearStream @ TNC2011

Setup codename:
FlightCees

UvA

iPerf
I7 3.2 GHz Q-core

iPerf
Amd Ph II 3.6 GHz HexC

Mellanox

40G E

Copenhagen

iPerf
2* dual 2.8 GHz Q-core

iPerf

Mellanox

CIENA
OME
6500

CERN

CIENA DWDM

Hamburg

Alcatel DWDM

CIENA
OME
6500

CIENA
OME
6500

17 ms RTT

LH

27 ms RTT

LH

Amsterdam – Geneva (CERN) – Copenhagen – 4400 km (2700 km alien light)

Demo setup codename: FlightCees

Ciena ActiveFlex(OME) 6500

Broadcom 40GE 18 port L2 Ethernet Switch

Supermicro Intel Server

Dell R815 Server

Visit CIENA Booth
surf to http://tnc11.delaat.net

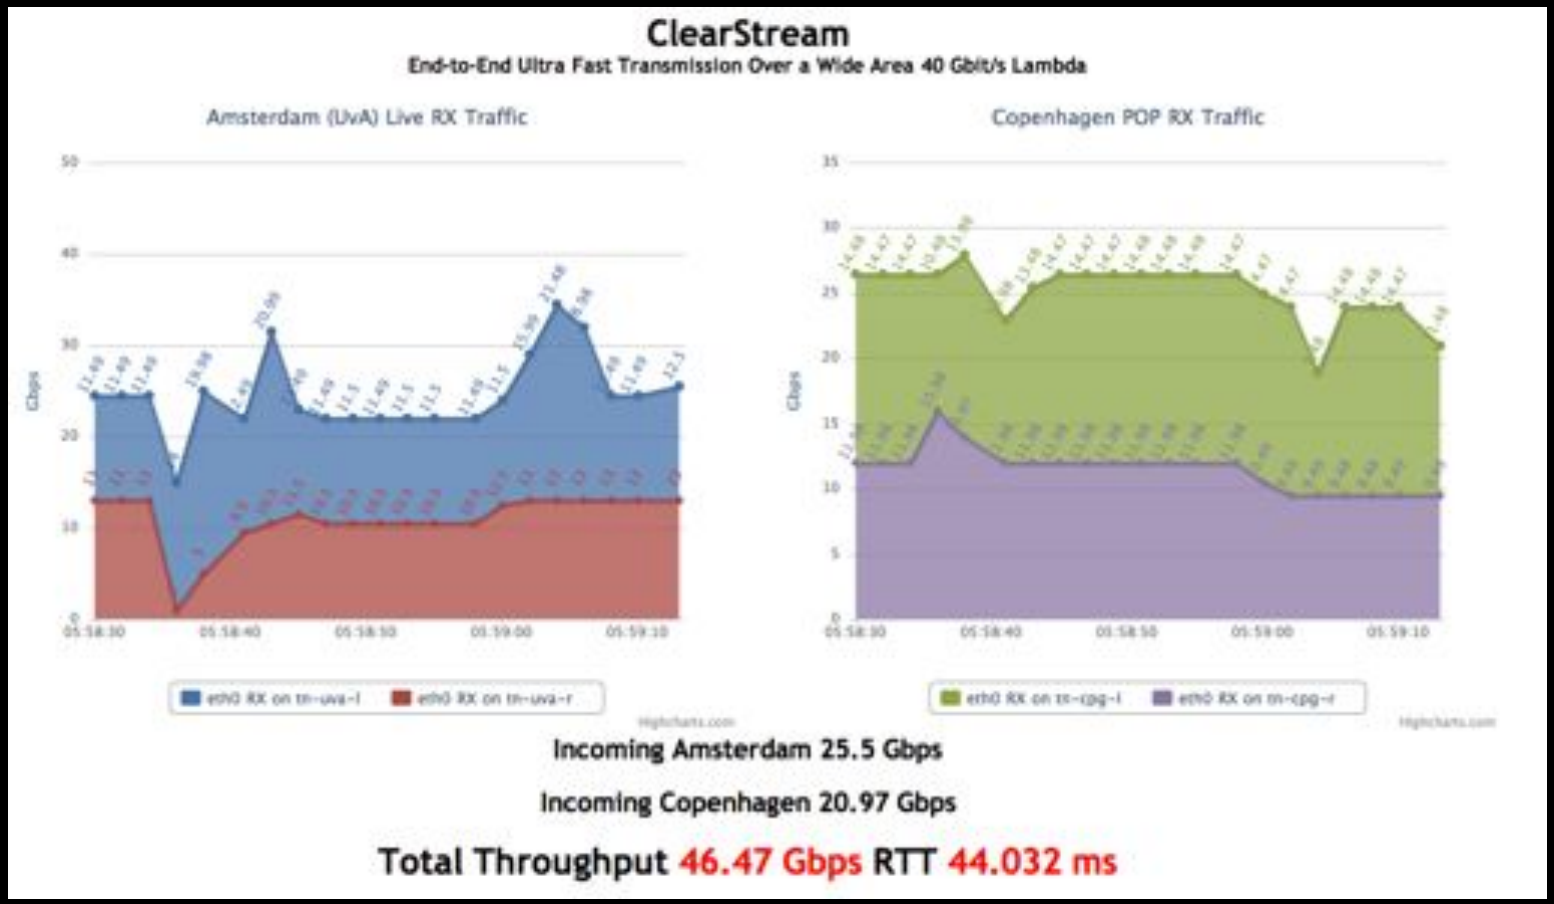

From GLIF October 2010 @ CERN

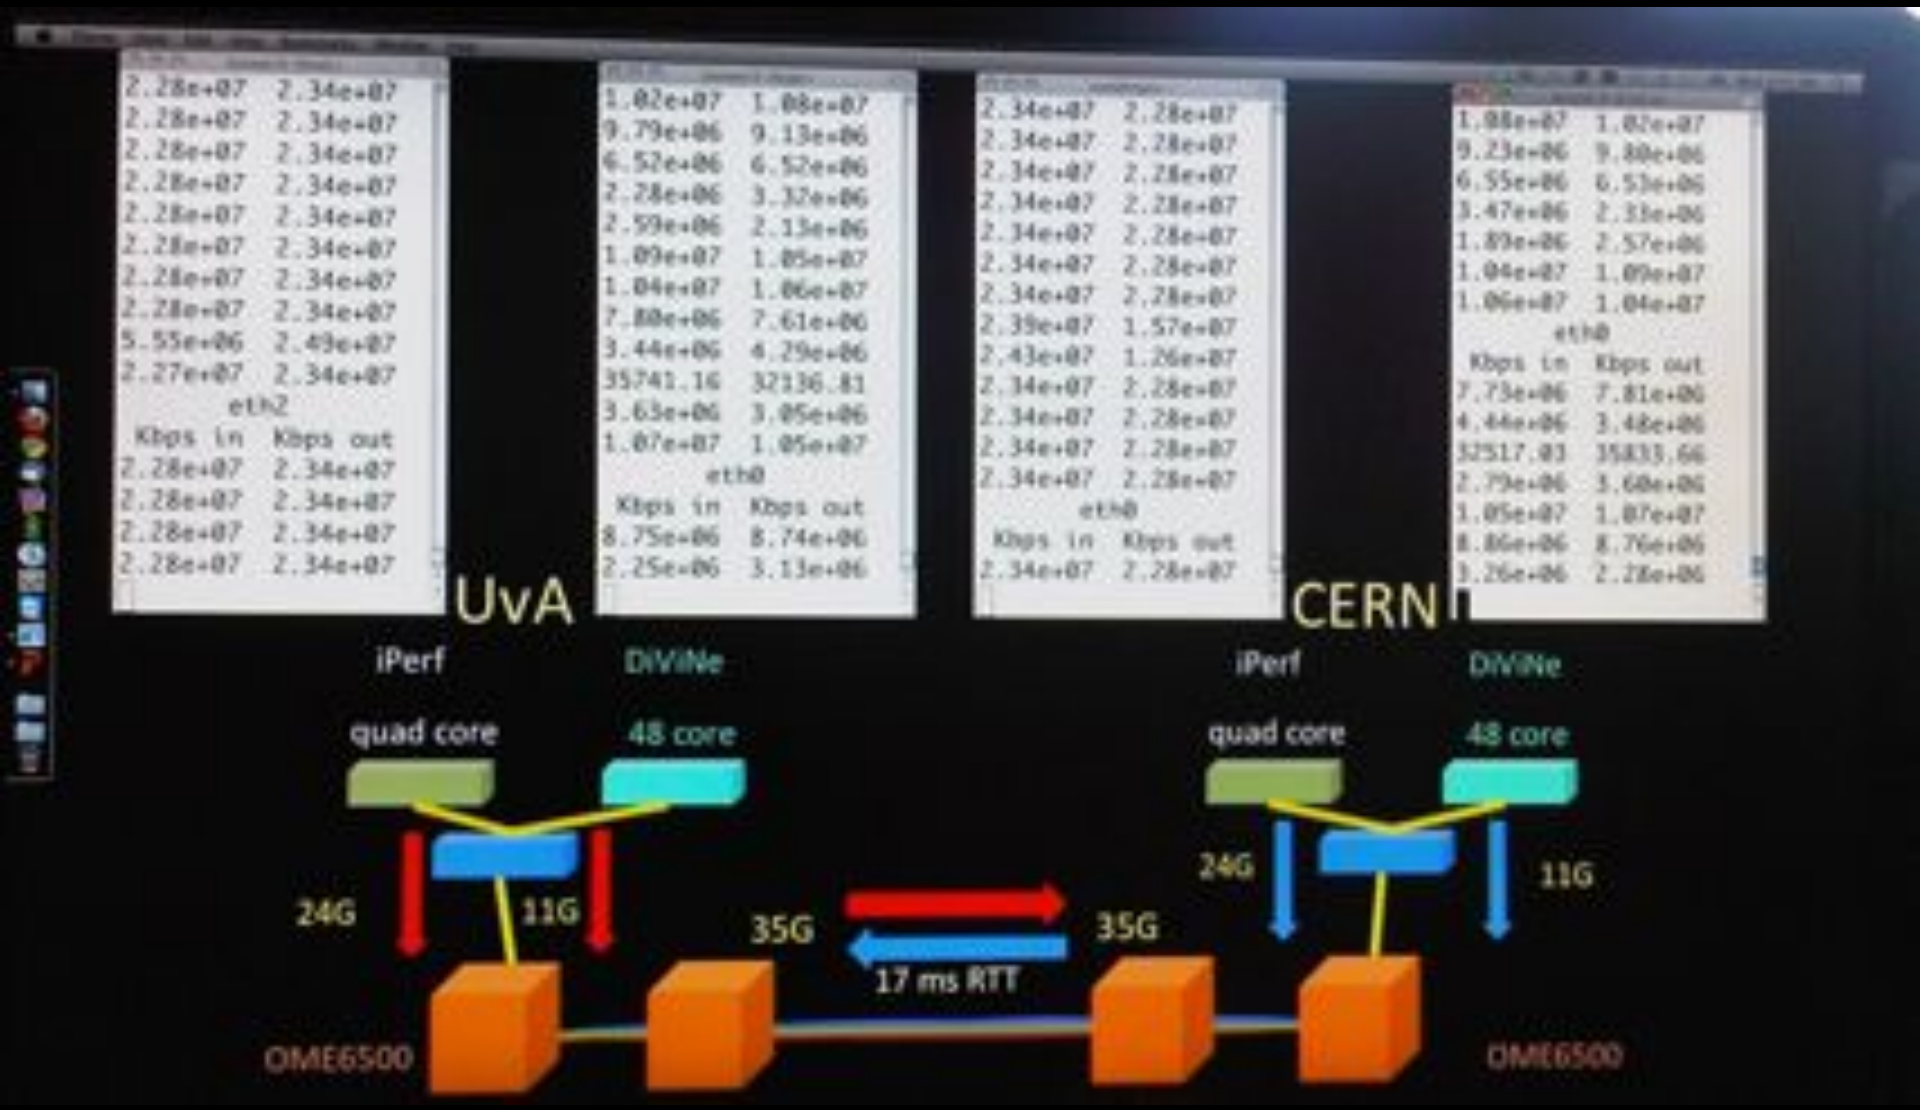

Results (rtt = 17 ms)

- Single flow iPerf 1 core -> 21 Gbps
- Single flow iPerf 1 core <> -> 15+15 Gbps
- Multi flow iPerf 2 cores -> 25 Gbps
- Multi flow iPerf 2 cores <> -> 23+23 Gbps
- DiViNe <> -> 11 Gbps
- Multi flow iPerf + DiVine -> 35 Gbps
- Multi flow iPerf + DiVine <> -> 35 + 35 Gbps

Performance Explained

- Mellanox 40GE card is PCI-E 2.0 8x (5GT/s)
- 40Gbit/s raw throughput but ….
- PCI-E is a network-like protocol
  - 8/10 bit encoding -> 25% overhead -> 32Gbit/s maximum data throughput
  - Routing information
- Extra overhead from IP/Ethernet framing
- Server architecture matters!
  - 4P system performed worse in multithreaded iperf

Server Architecture

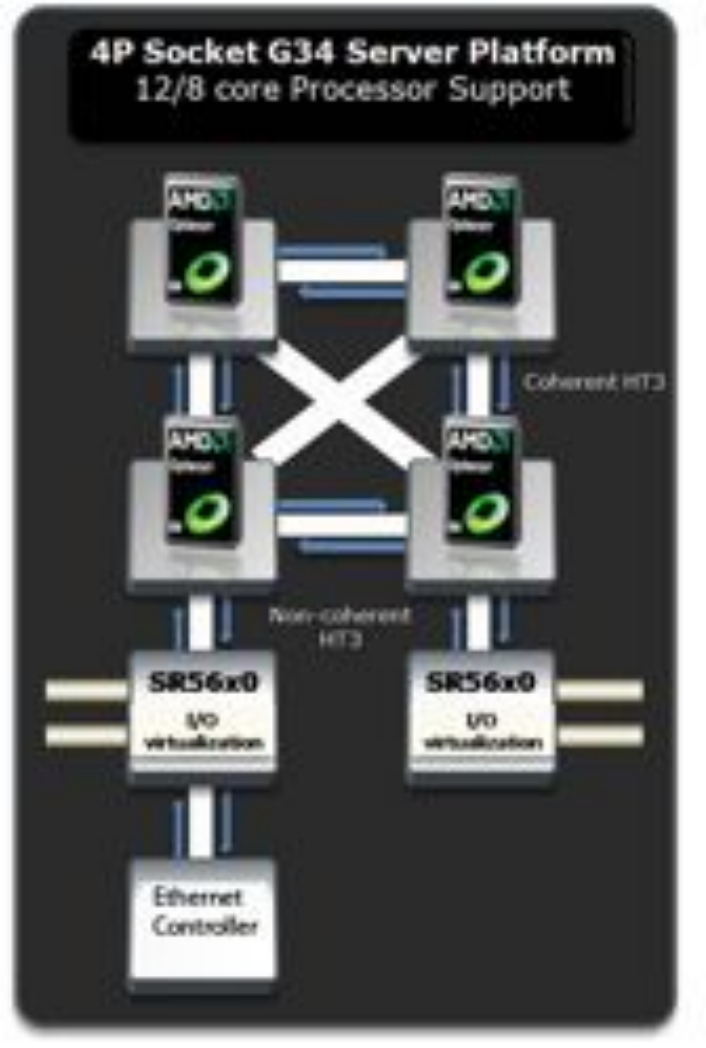

DELL R815

4 x AMD Opteron 6100

Supermicro X8DTT-HIBQF

2 x Intel Xeon

CPU Topology benchmark

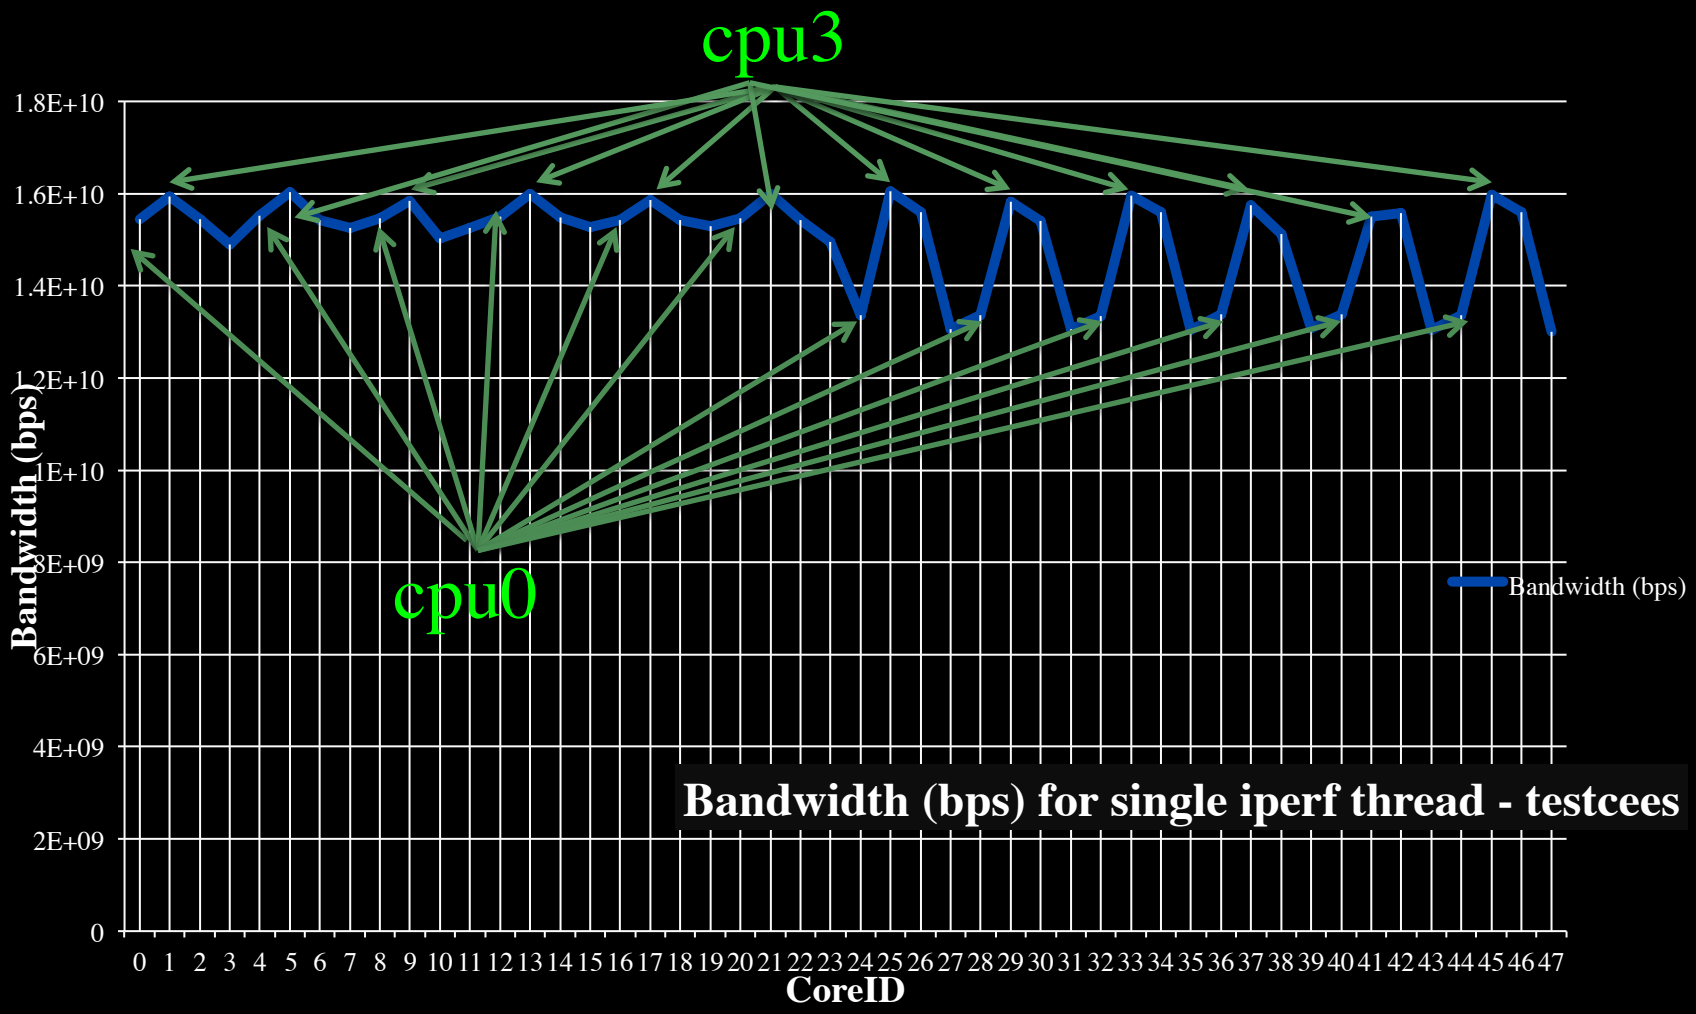

We used numactl to bind iperf to cores

We investigate: TOMTOM for complex networks!

LinkedIN for Infrastructure

- From semantic Web / Resource Description Framework.
- The RDF uses XML as an interchange syntax.
- Data is described by triplets (Friend of a Friend):

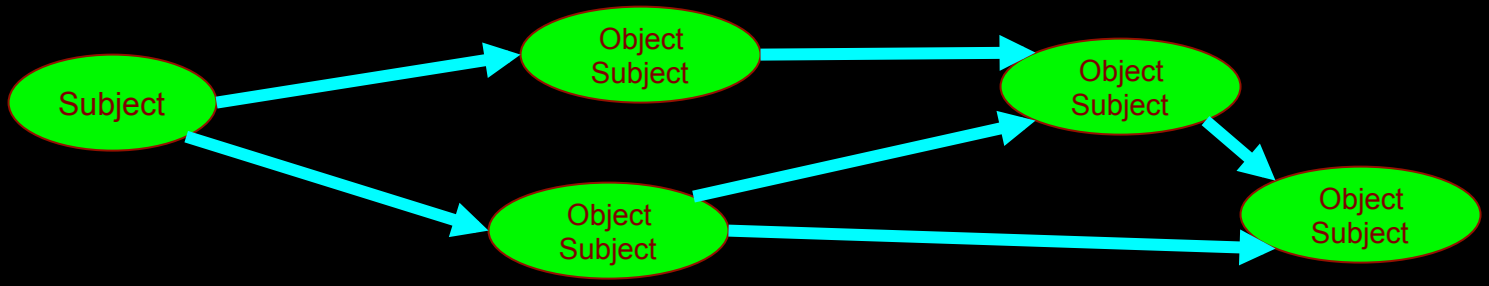

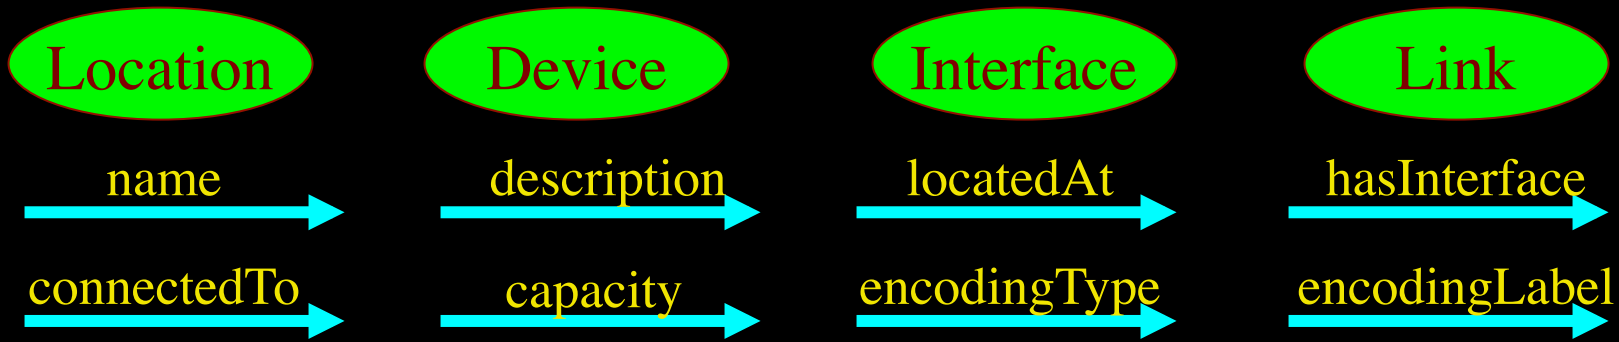

NetherLight in RDF

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
      xmlns:ndl="http://www.science.uva.nl/research/air/ndl#">
<!-- Description of Netherlight -->
<ndl:Location rdf:about="#Netherlight">
      <ndl:name>Netherlight Optical Exchange</ndl:name>
</ndl:Location>
<!-- TDM3.amsterdam1.netherlight.net -->
<ndl:Device rdf:about="#tdm3.amsterdam1.netherlight.net">
      <ndl:name>tdm3.amsterdam1.netherlight.net</ndl:name>
      <ndl:locatedAt rdf:resource="#amsterdam1.netherlight.net"/>
      <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/1"/>
      <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/3"/>
      <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/4"/>
      <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:503/1"/>
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
      <ndl:hasInterface rdf:resourc
```

```
<!-- all the interfaces of TDM3.amsterdam1.netherlight.net -->

<ndl:Interface rdf:about="#tdm3.amsterdam1.netherlight.net:501/1">
            <ndl:name>tdm3.amsterdam1.netherlight.net:POS501/1</ndl:name>
            <ndl:connectedTo rdf:resource="#tdm4.amsterdam1.netherlight.net:5/1"/>
 </ndl:Interface>
<ndl:Interface rdf:about="#tdm3.amsterdam1.netherlight.net:501/2">
            <ndl:name>tdm3.amsterdam1.netherlight.net:POS501/2</ndl:name>
            <ndl:connectedTo rdf:resource="#tdm1.amsterdam1.netherlight.net:12/1"/>
</ndl:Interface>
```

Multi-layer descriptions in NDL

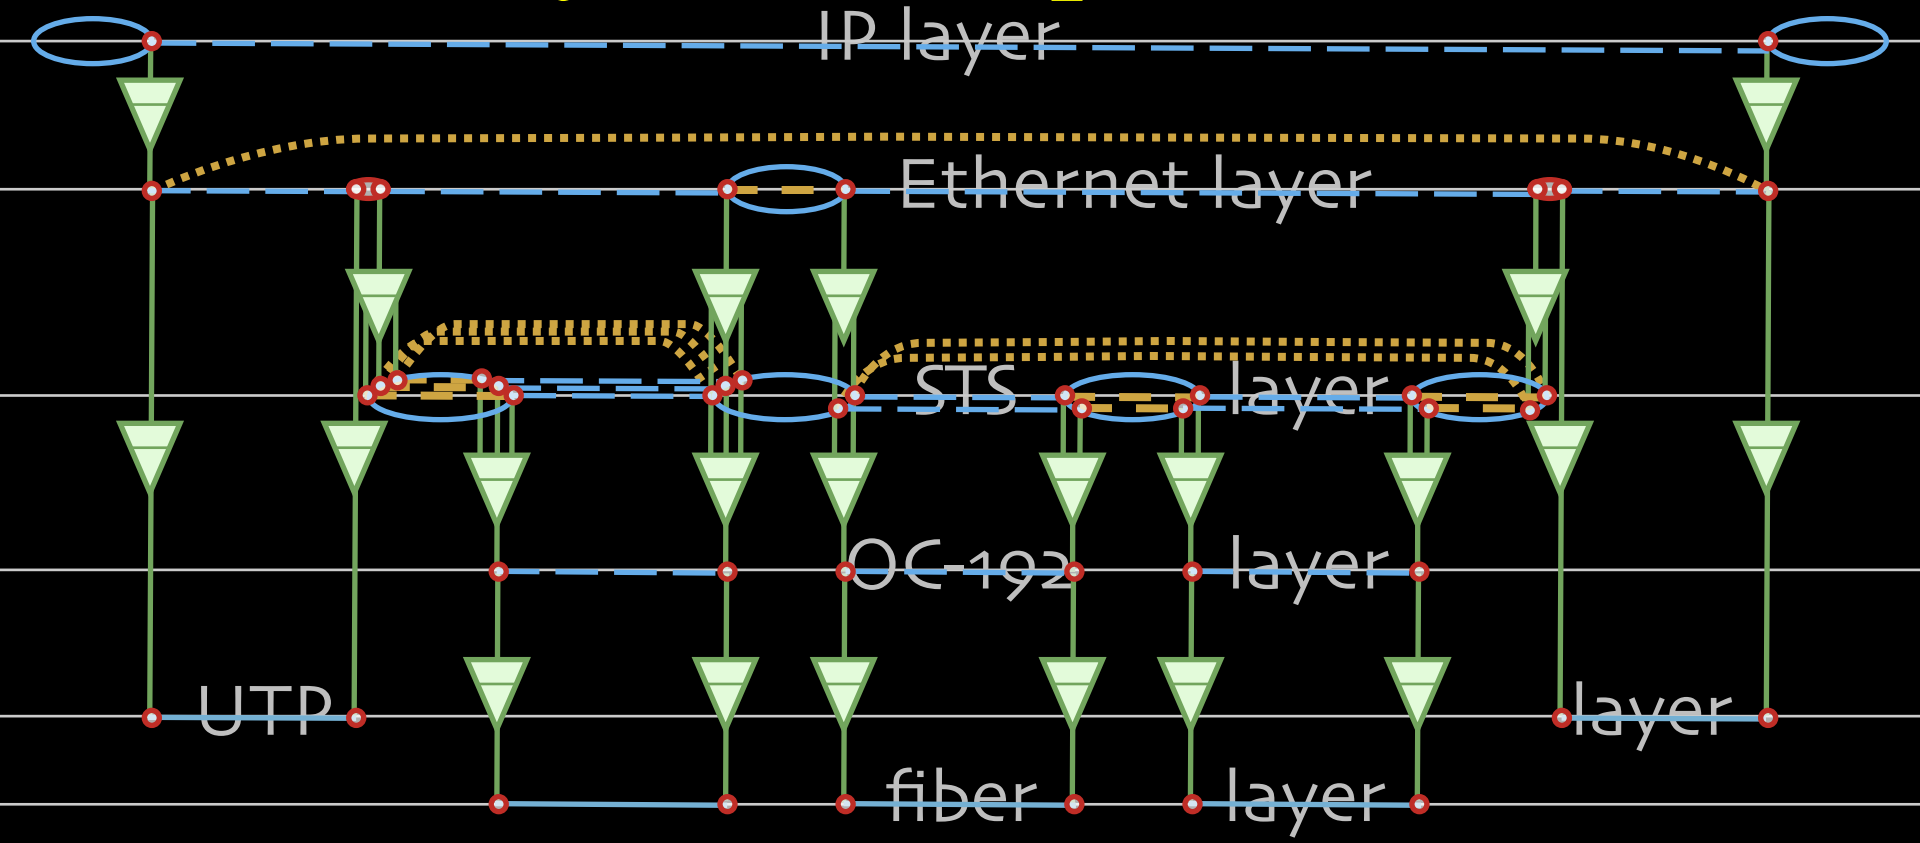

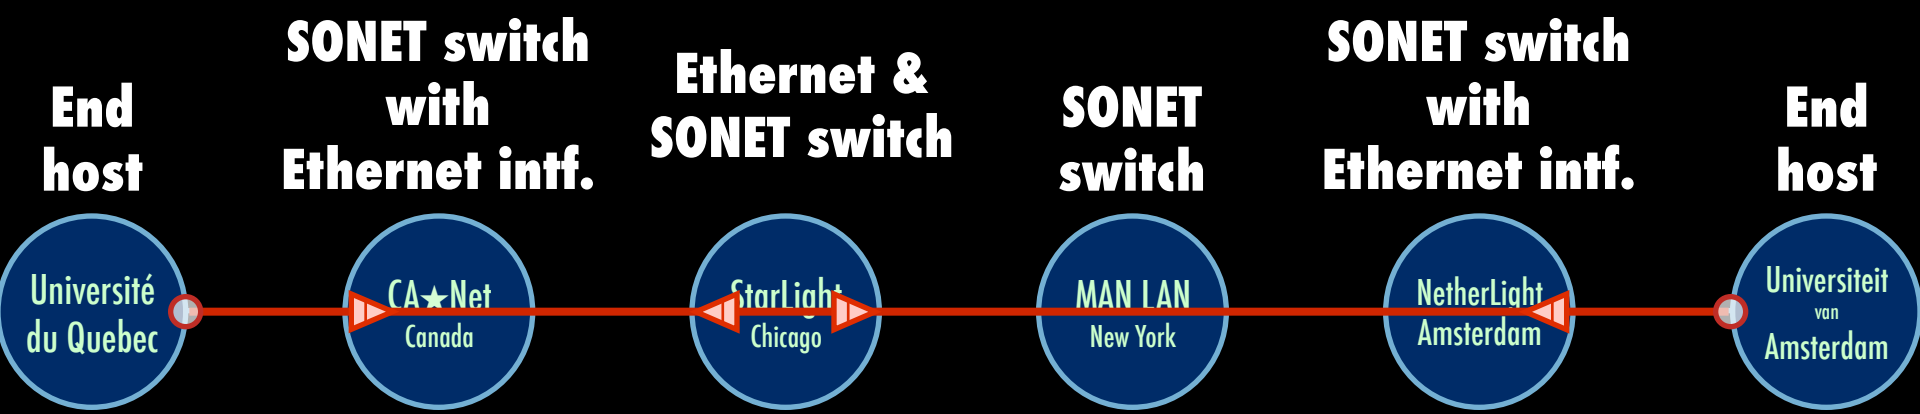

Multi-layer Network PathFinding

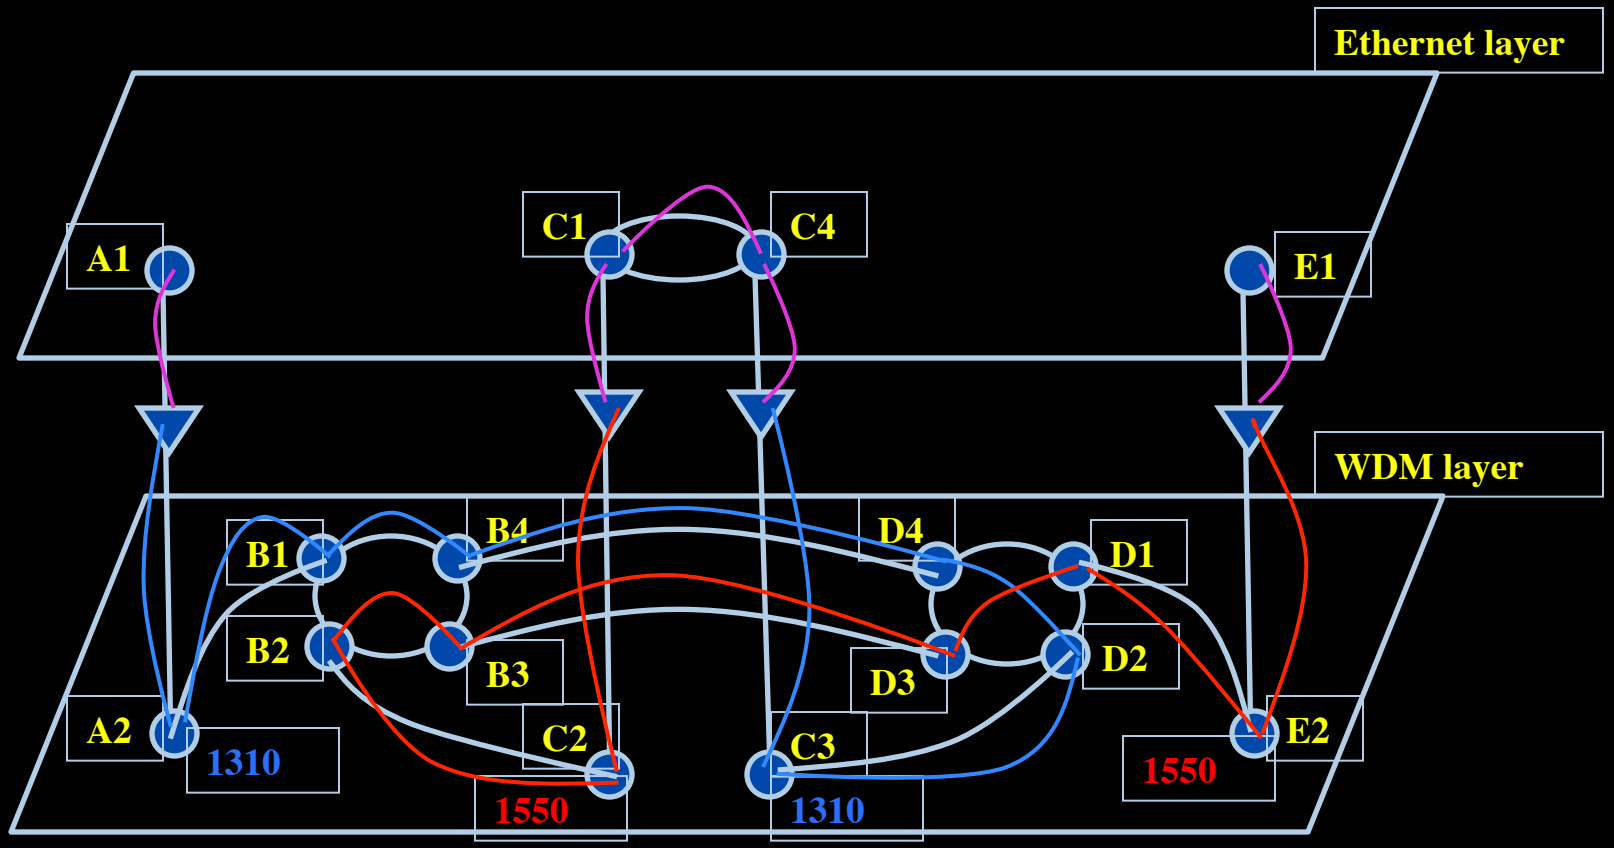

Path between interfaces A1 and E1:

A1-A2-B1-B4-D4-D2-C3-C4-C1-C2-B2-B3-D3-D1-E2-E1

Scaling: Combinatorial problem

Virtualisatie van infrastructuur & QoS

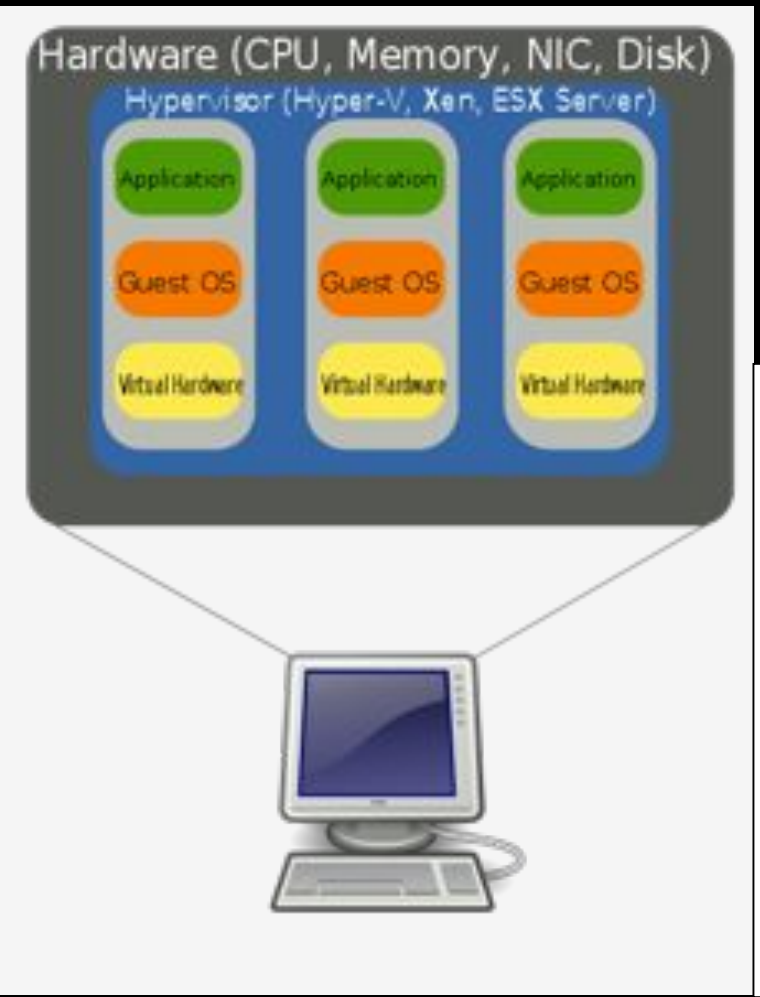

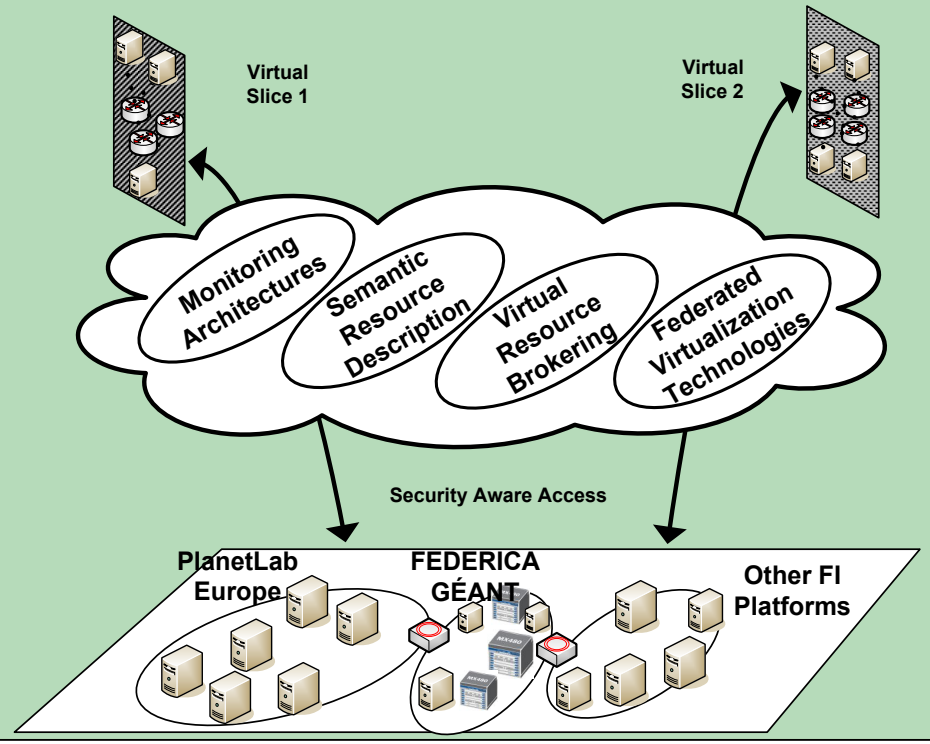

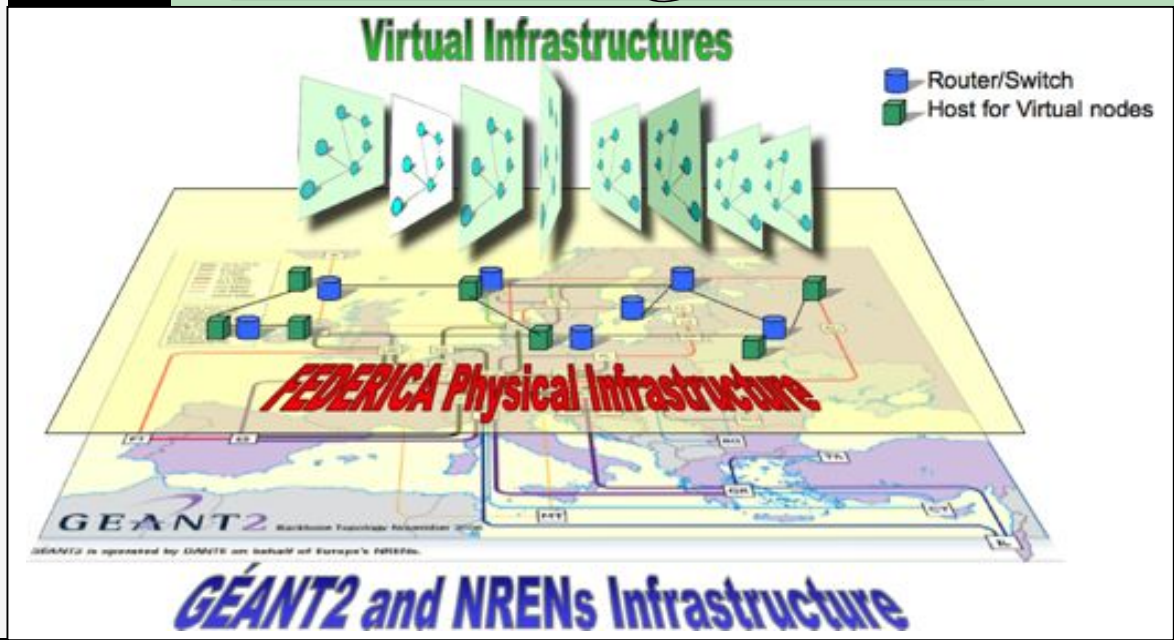

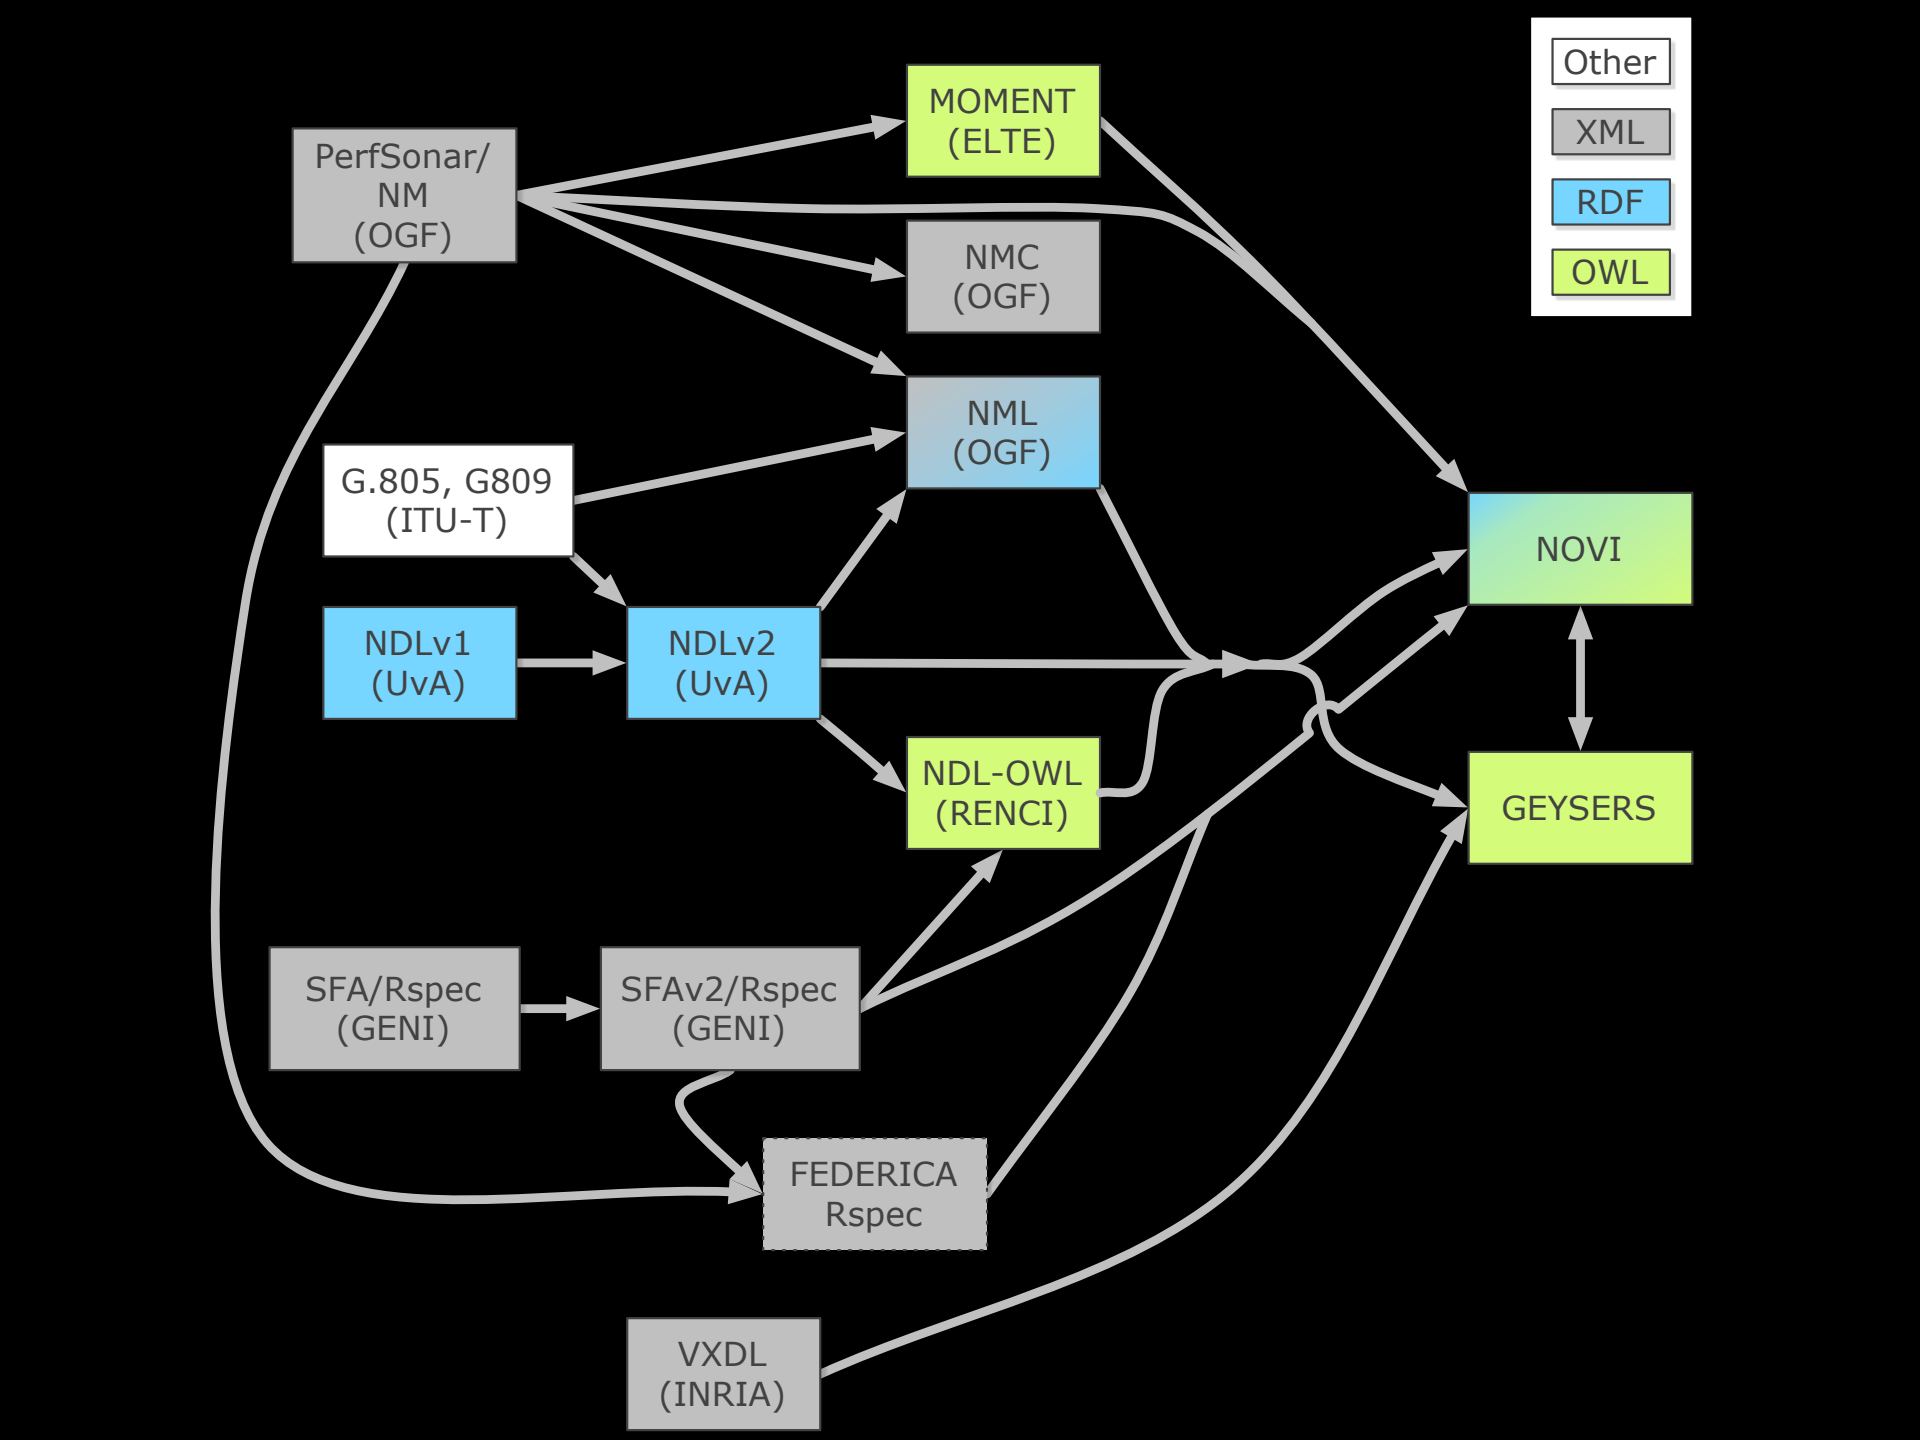
Other
XML
RDF
OWL
MOMENT (ELTE)
PerfSonar/ NM (OGF)
NMC (OGF)
NML (OGF)
G.805, G809 (ITU-T)
NOVI
NDLv1 (UvA)
NDLv2 (UvA)
NDL-OWL (RENCI)
GEYSERS
SFA/Rspec (GENI)
SFAv2/Rspec (GENI)
FEDERICA Rspec
VXDL (INRIA)

Information Modeling

Define a common information model for ***infrastructures*** and ***services***. Base it on Semantic Web.

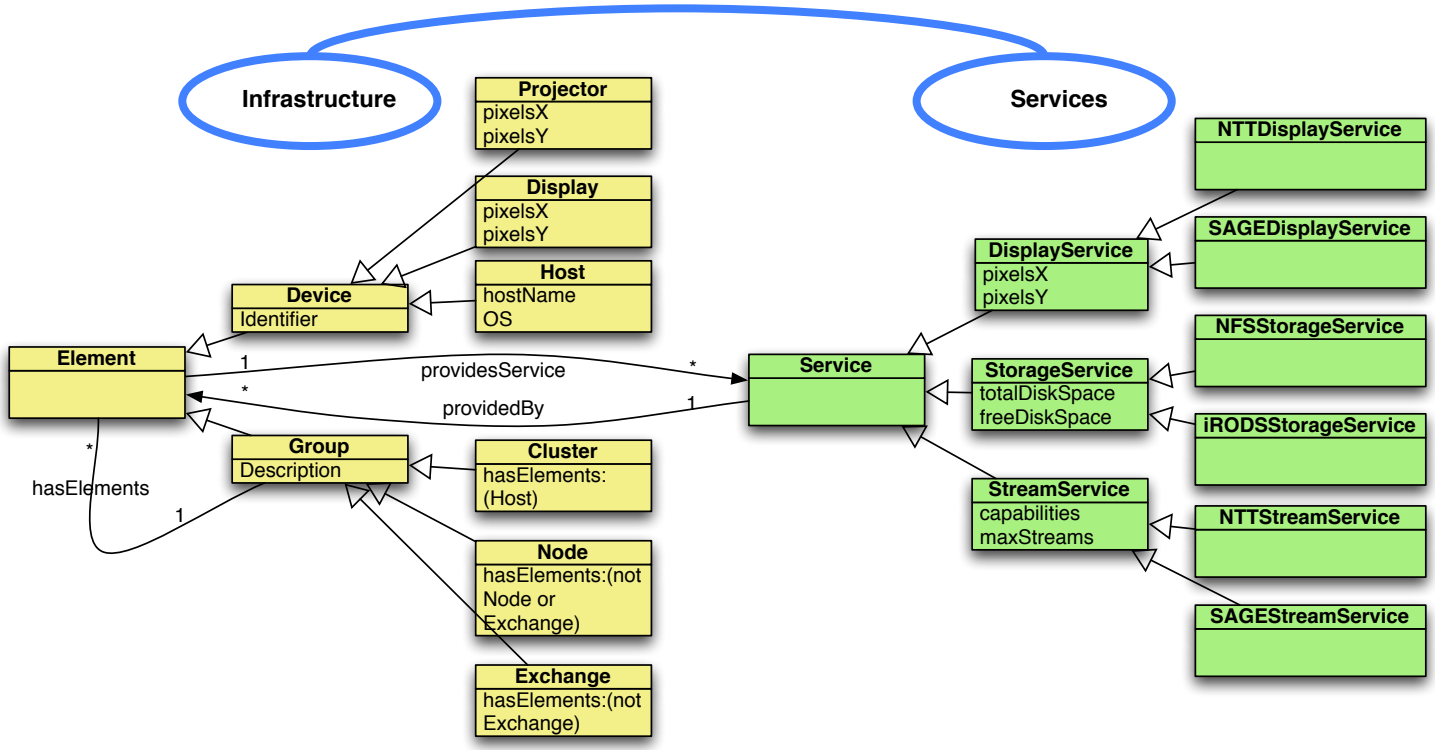

J. van der Ham, F. Dijkstra, P. Grosso, R. van der Pol, A. Toonk, C. de Laat
A distributed topology information system for optical networks based on the semantic web,
In: Elsevier Journal on Optical Switching and Networking, Volume 5, Issues 2-3, June 2008, Pages 85-93

R.Koning, P.Grosso and C.de Laat
Using ontologies for resource description in the CineGrid Exchange
In: Future Generation Computer Systems (2010)

SNE @ UvA

	Ijkdijk/Urban Flood	Medical	LifeWatch/ENVRI	CosmoGrid/eVLBI	CineGrid	EU-GN3/NOVI/Geysers	SURFnet/GLIF/Cloud
Green-IT						X	X
Privacy/Trust		X			X		
Authorization/policy		X	X		X	X	
Programmable networks	X		X				
40-100Gig/TCP/WF/QoS	X		X	X		X	
Topology/Architecture		X		X	X	X	
Optical Photonic		X	X		X		

Why is more resolution is better?

1. More Resolution Allows Closer Viewing of Larger Image
2. Closer Viewing of Larger Image Increases Viewing Angle
3. Increased Viewing Angle Produces Stronger Emotional Response

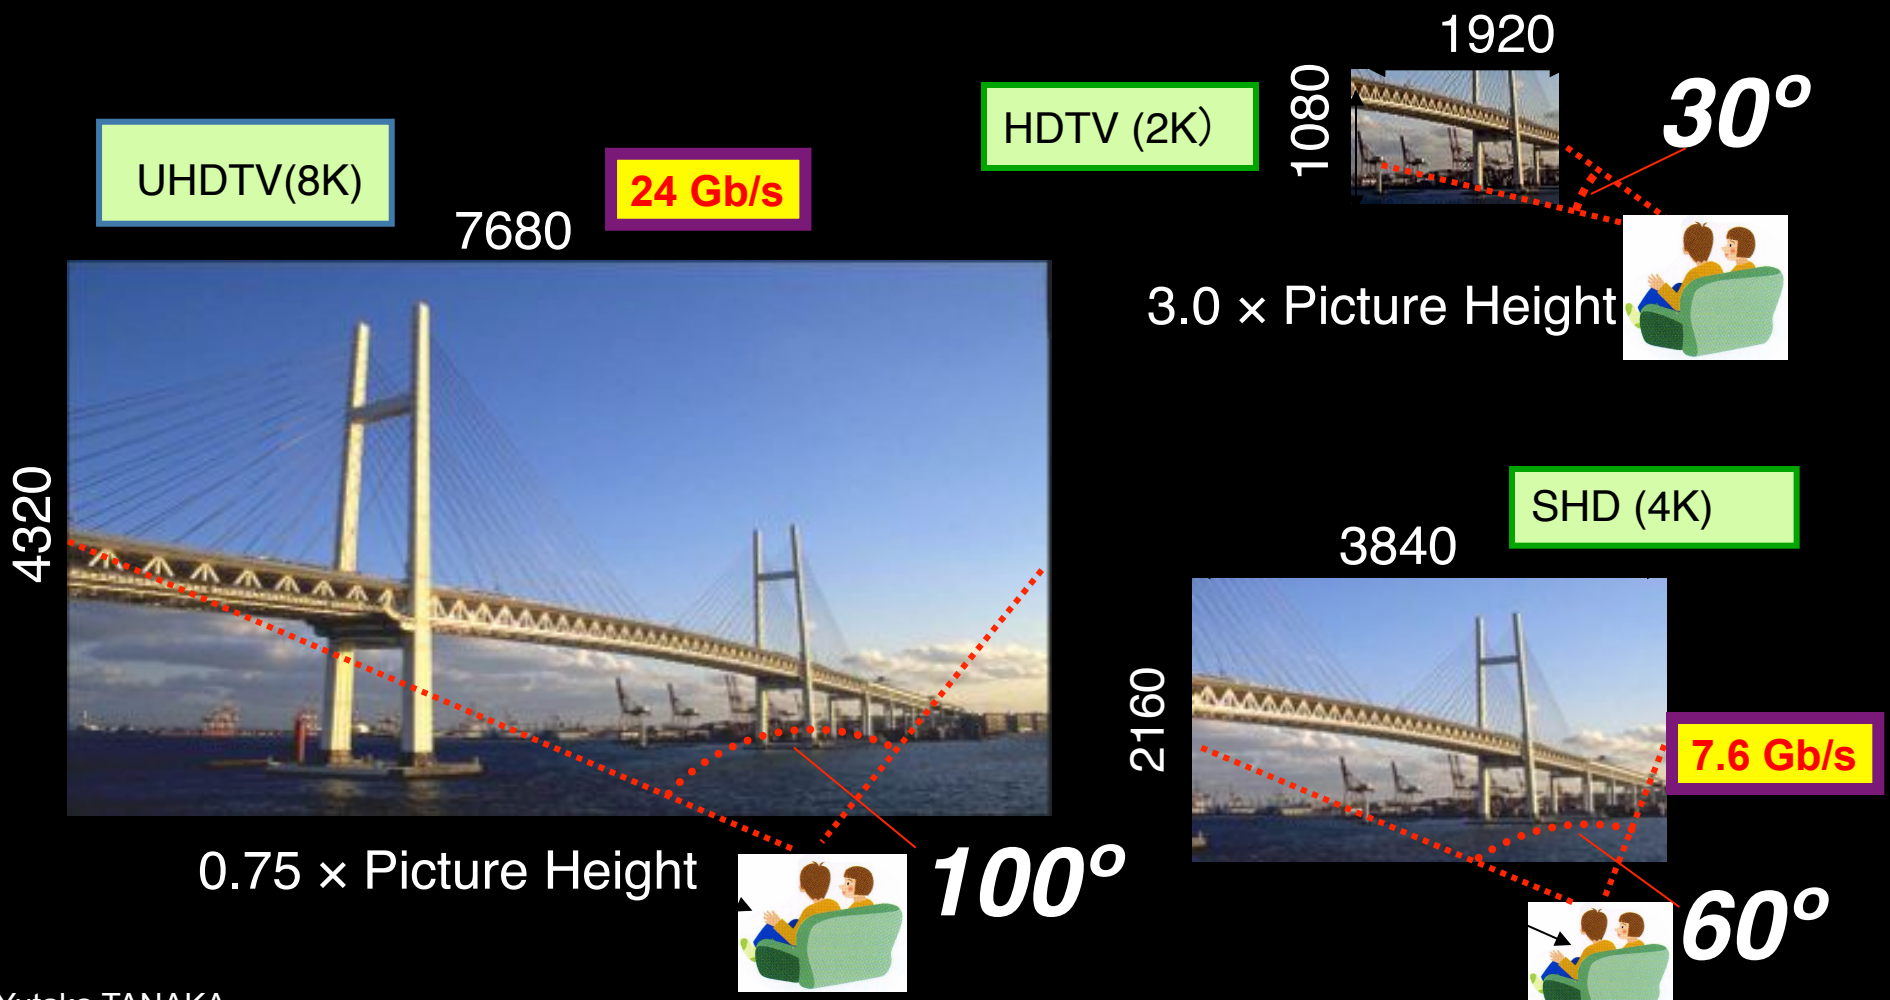

Yutaka TANAKA
SHARP CORPORATION
Advanced Image Research Laboratories


Red End
Robin Noorda & Bethany de Forest


STEREO 4K Recording
Viktoria Mullova
Holland Festival 2010


Hey at all
We're almost done. Shhh..

Why?

I want to:

"Show Big Bug Bunny in 4K on my Tiled Display using green Infrastructure"

- Big Bugs Bunny can be on multiple servers on the Internet.
- Movie may need processing / recoding to get to 4K for Tiled Display.
- Needs deterministic Green infrastructure for Quality of Experience.
- Consumer / Scientist does not want to know the underlying details.
  ➔ His refrigerator also just works.

RDF describing Infrastructure

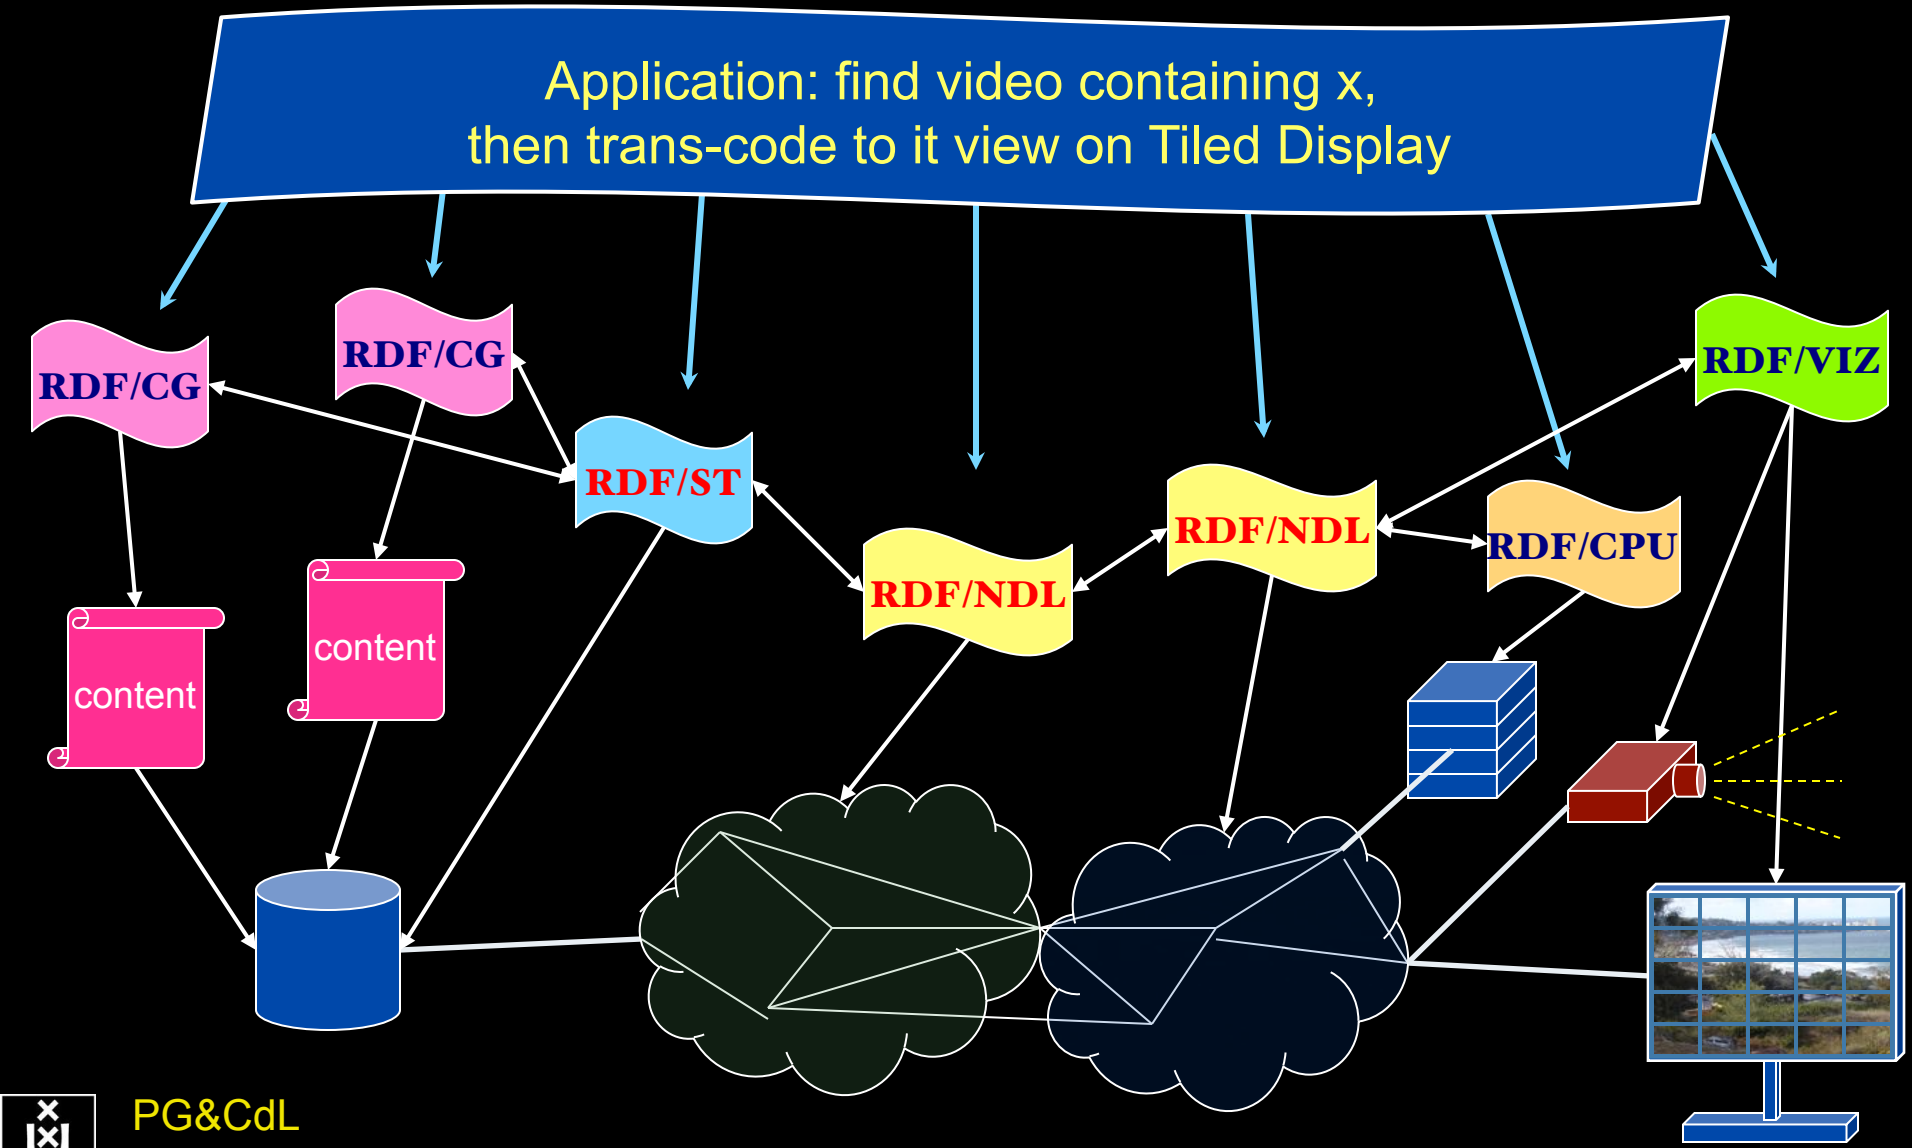

PG&CdL

Applications and Networks become aware of each other!

CineGrid Description Language

CineGrid is an initiative to facilitate the exchange, storage and display of high-quality digital media.

The CineGrid Description Language (CDL) describes CineGrid resources. Streaming, display and storage components are organized in a hierarchical way.

CDL has bindings to the NDL ontology that enables descriptions of network components and their interconnections.

With CDL we can reason on the CineGrid infrastructure and its services.

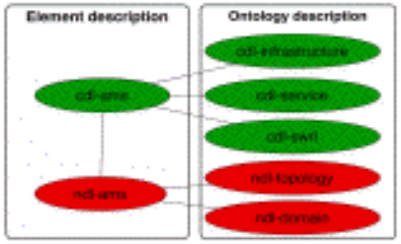

SQWRL is used to query the Ontology.

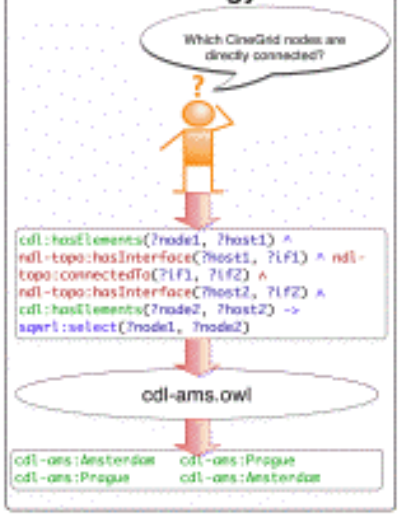

UML representation of CDL

CDL links to ***NDL*** using the ***owl:SameAs*** property. ***CDL*** defines the services, ***NDL*** the network interfaces and links. The combination of the two ontologies identifies the host pairs that support matching services via existing network connections.

GigaPort

CineGrid portal

100 Tbyte
Cache & Store & Forward

CineGrid distribution center Amsterdam

Home | About | Browse Content | cinegrid.org | cinegrid.nl

Amsterdam Node Status:

node41:
Disk space used: 8 GiB
Disk space available: 10 GiB

Search node:

Search

Browse by tag:

amsterdam animation antonacci blender boat bridge bunny cgi delos holland hollandfestival leidschestraat muziekgebouw nieuwmarkt opera prague ship train tram trams waag

CineGrid Amsterdam

Welcome to the Amsterdam CineGrid distribution node. Below are the latest additions of super-high-quality video to our node.

For more information about CineGrid and our efforts look at the about section.

Latest Additions

Wypke

Wypke

Available formats:
4k dxt (4.0 kB)
Duration: 1 hour and 8 minutes
Created: 1 week, 2 days ago
Author: Wypke
Categories:

Prague Train

Steam locomotive in Prague.

Available formats:
4k dxt (3.9 kB)
Duration: 27 hours and 46 minutes
Created: 1 week, 2 days ago
Author: CineGrid
Categories: delos prague train

VLC: Big Buck Bunny

(c) copyright Blender Foundation | http://www.bigbuckbunny.org

Available formats:
1080p MPEG4 (1.1 GB)
Duration: 1 hour and 0 minutes
Created: 1 month, 1 week ago
Author: Blender Foundation
Categories: animation blender bunny cgi

CineGrid Workflow Planner

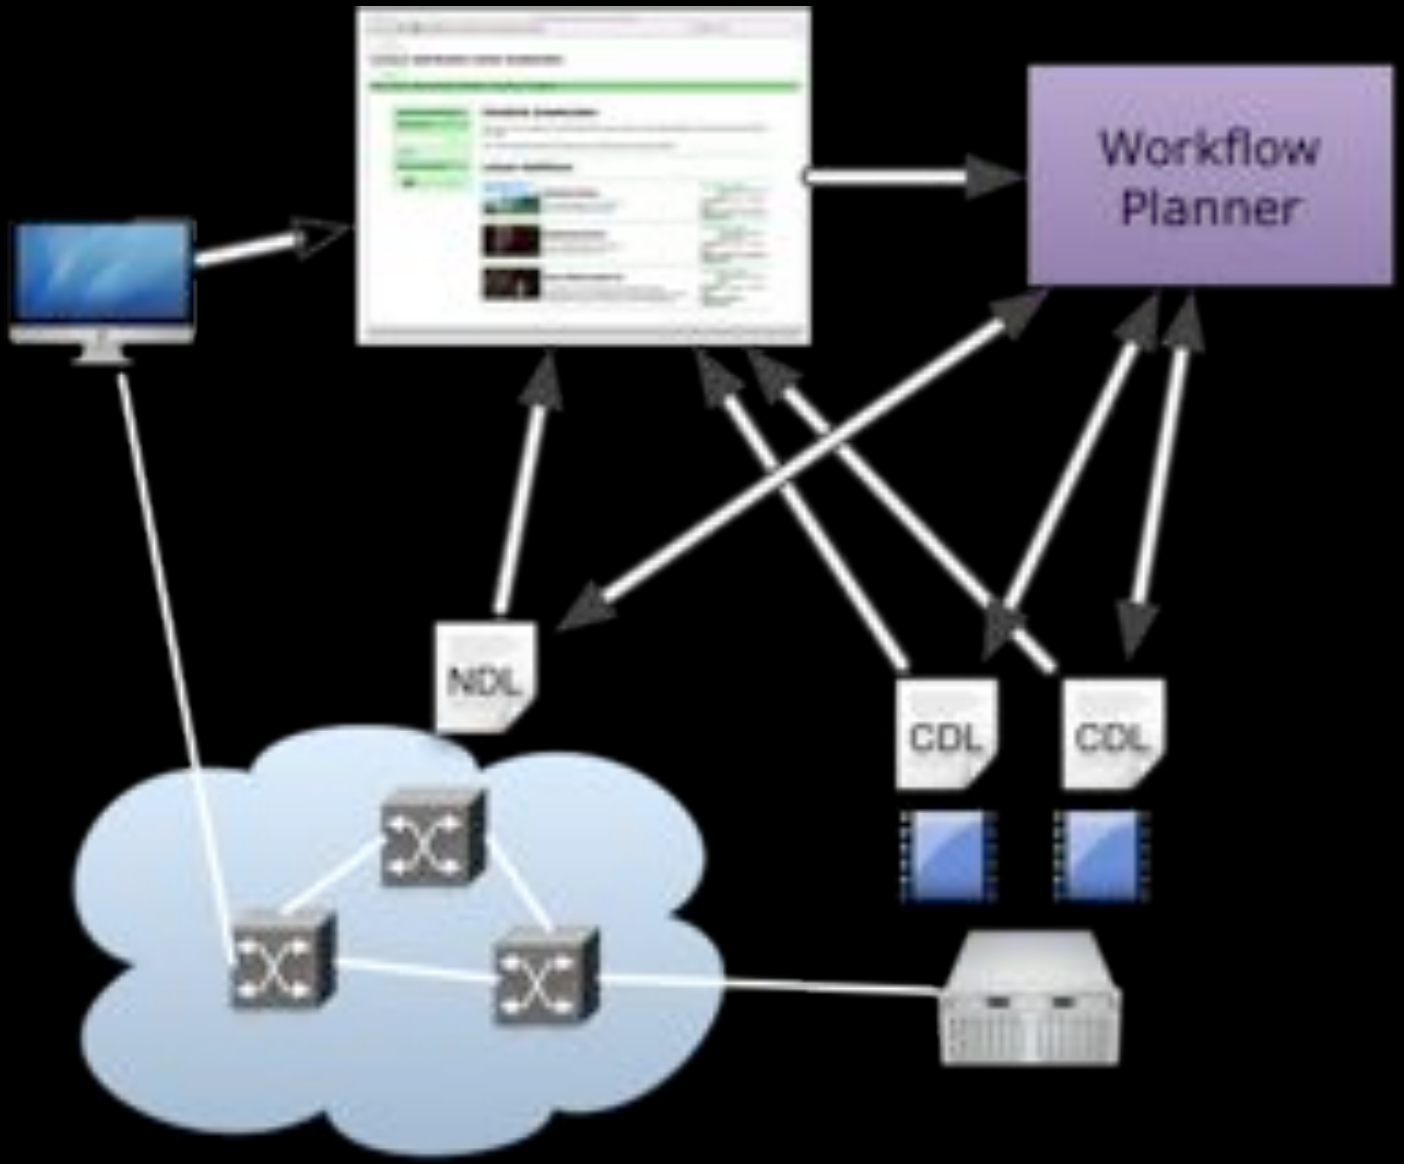

SNE @ UvA

	Ijkdijk/Urban Flood	Medical	LifeWatch/ENVRI	CosmoGrid/eVLBI	CineGrid	EU-GN3/NOVI/Geysers	SURFnet/GLIF/Cloud
Green-IT					X	X	
Privacy/Trust		X			X		
Authorization/policy		X	X		X	X	
Programmable networks	X		X				
40-100Gig/TCP/WF/QoS	X		X	X		X	
Topology/Architecture		X		X	X	X	
Optical Photonic		X	X		X		

The GLIF – lightpaths around the world

The GLIF – lightpaths around the world

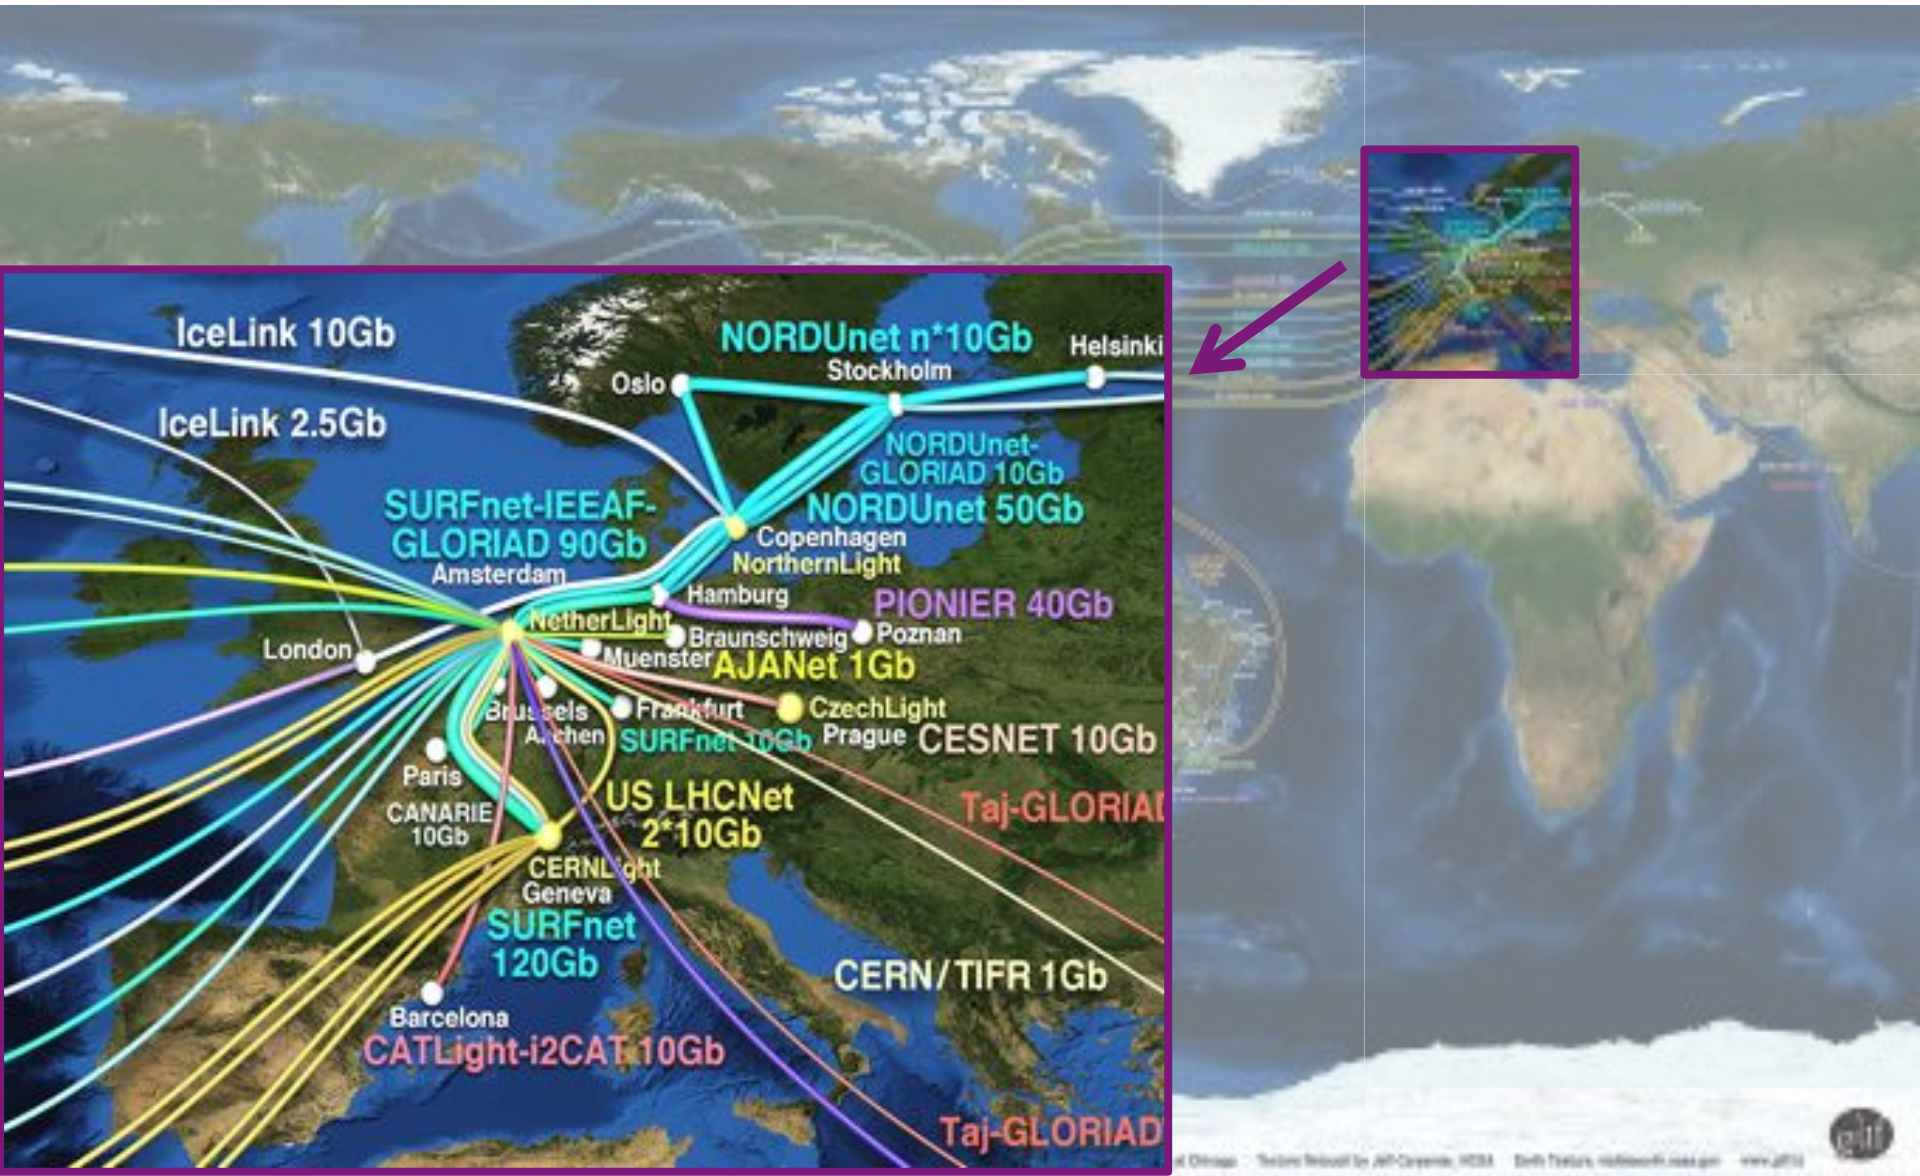

The Ten Problems with the Internet

1. Energy Efficient Communication
2. Separation of Identity and Address
3. Location Awareness
4. Explicit Support for Client-Server Traffic and Distributed Services
5. Person-to-Person Communication
6. Security
7. Control, Management, and Data Plane separation
8. Isolation
9. Symmetric/Asymmetric Protocols
10. Quality of Service

Nice to have:

- Global Routing with Local Control of Naming and Addressing
- Real Time Services
- Cross-Layer Communication
- Manycast
- Receiver Control
- Support for Data Aggregation and Transformation
- Support for Streaming Data
- Virtualization

ref: Raj Jain, "Internet 3.0: Ten Problems with Current Internet Architecture and Solutions for the Next Generation", Military Communications Conference, 2006. MILCOM 2006. IEEE

The Ten Problems with the Internet

1. **Energy Efficient Communication**
2. Separation of Identity and Address
3. Location Awareness
4. **Explicit Support for Client-Server Traffic and Distributed Services**
5. Person-to-Person Communication
6. Security
7. **Control, Management, and Data Plane separation**
8. **Isolation**
9. Symmetric/Asymmetric Protocols
10. **Quality of Service**

Nice to have:

- Global Routing with Local Control of Naming and Addressing
- **Real Time Services**
- **Cross-Layer Communication**
- Manycast
- Receiver Control
- Support for Data Aggregation and Transformation
- **Support for Streaming Data**
- **Virtualization**

ref: Raj Jain, "Internet 3.0: Ten Problems with Current Internet Architecture and Solutions for the Next Generation", Military Communications Conference, 2006. MILCOM 2006. IEEE

TimeLine

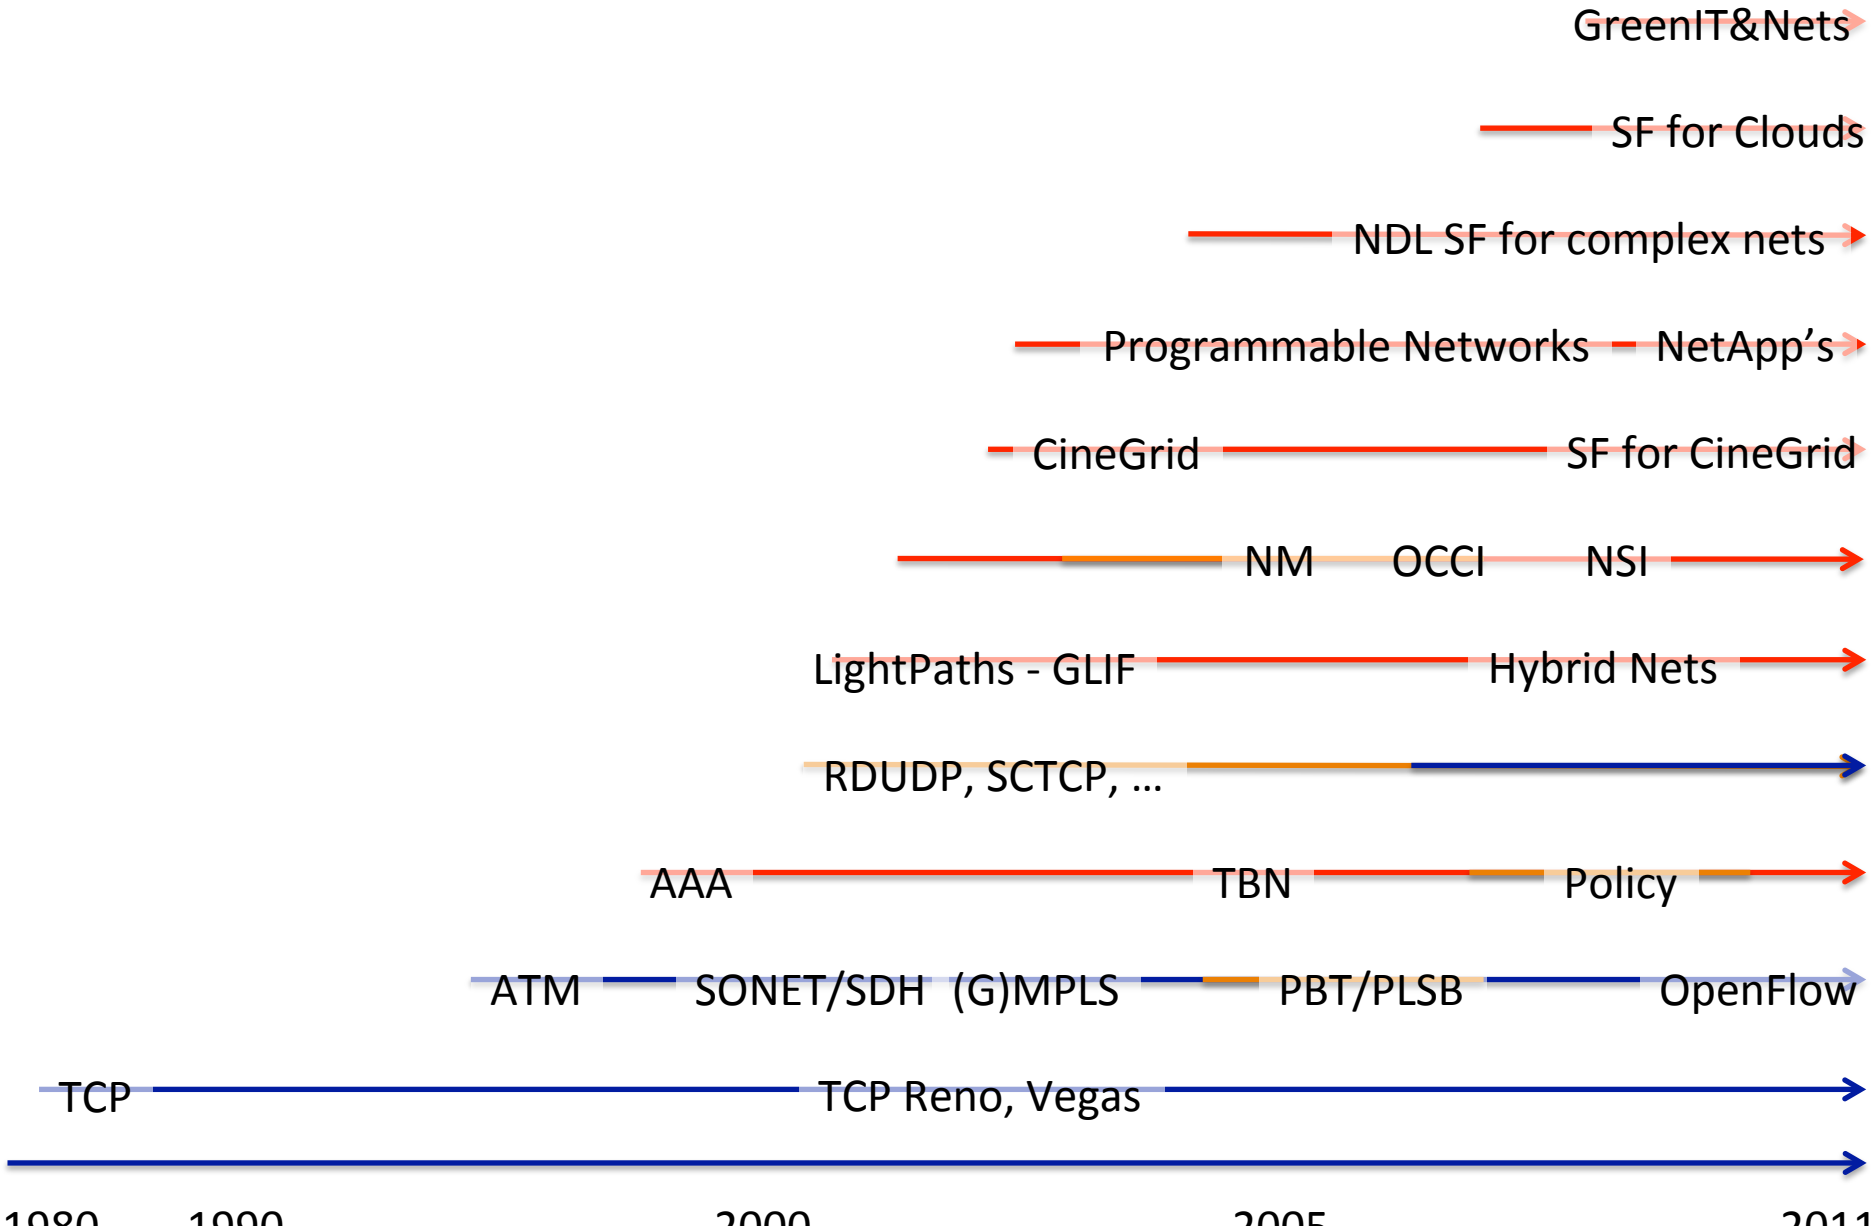
GreenIT&Nets
SF for Clouds
NDL SF for complex nets
Programmable Networks
NetApp's
CineGrid
SF for CineGrid
NM
OCCI
NSI
LightPaths - GLIF
Hybrid Nets
RDUDP, SCTCP, …
AAA
TBN
Policy
ATM
SONET/SDH
(G)MPLS
PBT/PLSB
OpenFlow
TCP
TCP Reno, Vegas
1980
1990
2000
2005
2011

TimeLine

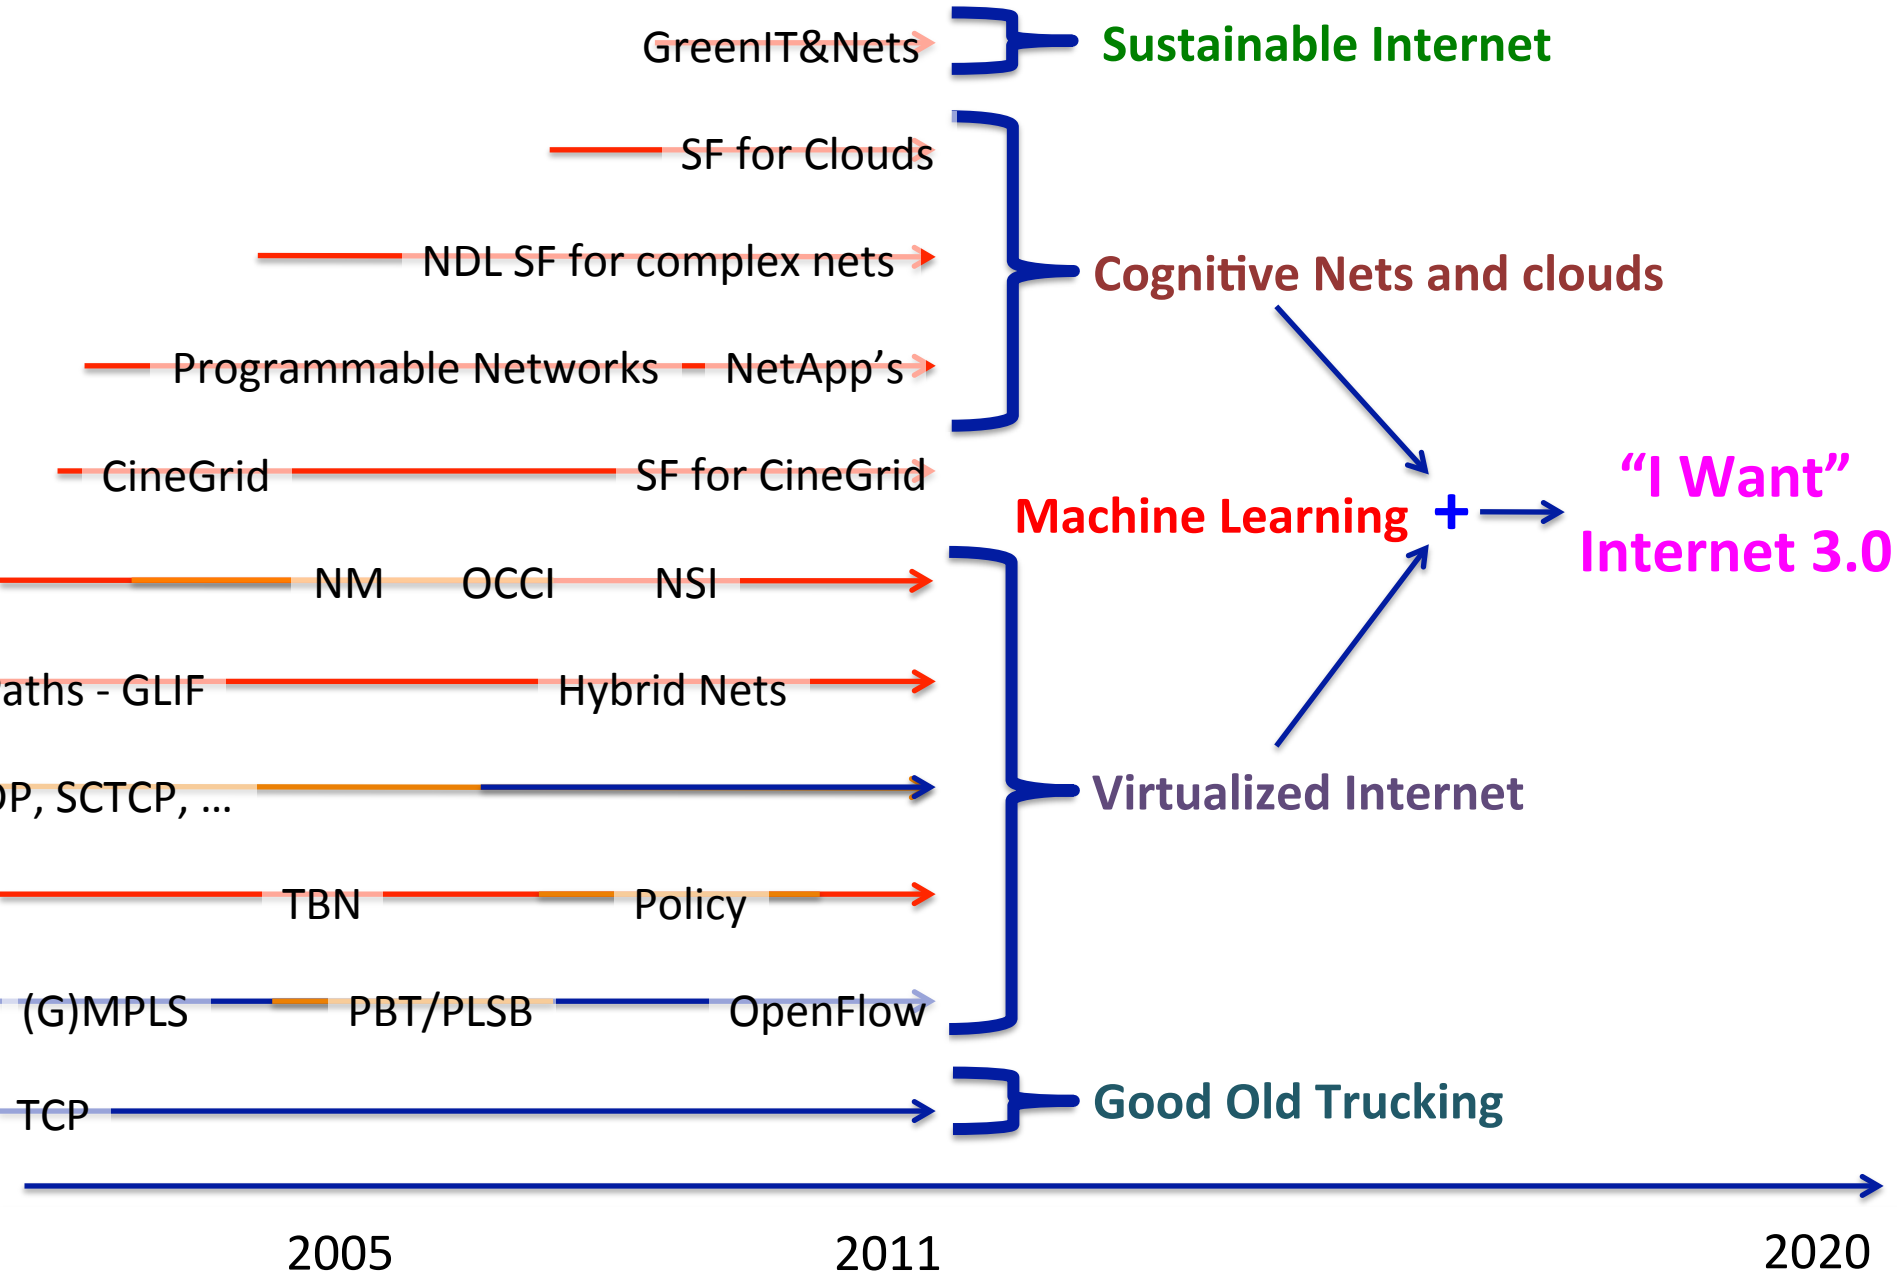
GreenIT&Nets
Sustainable Internet
SF for Clouds
NDL SF for complex nets
Programmable Networks
NetApp's
Cognitive Nets and clouds
CineGrid
SF for CineGrid
Machine Learning
+
"I Want"
Internet 3.0
NM
OCCI
NSI
aths - GLIF
Hybrid Nets
P, SCTCP, …
Virtualized Internet
TBN
Policy
(G)MPLS
PBT/PLSB
OpenFlow
TCP
Good Old Trucking
2005
2011
2020

TimeLine

- Sustainable Internet
- Cognitive Nets and clouds
- Machine Learning + → "I Want" Internet 3.0
- Virtualized Internet
- Good Old Trucking

I retire

2020

2040

Complex e-Infrastructure!

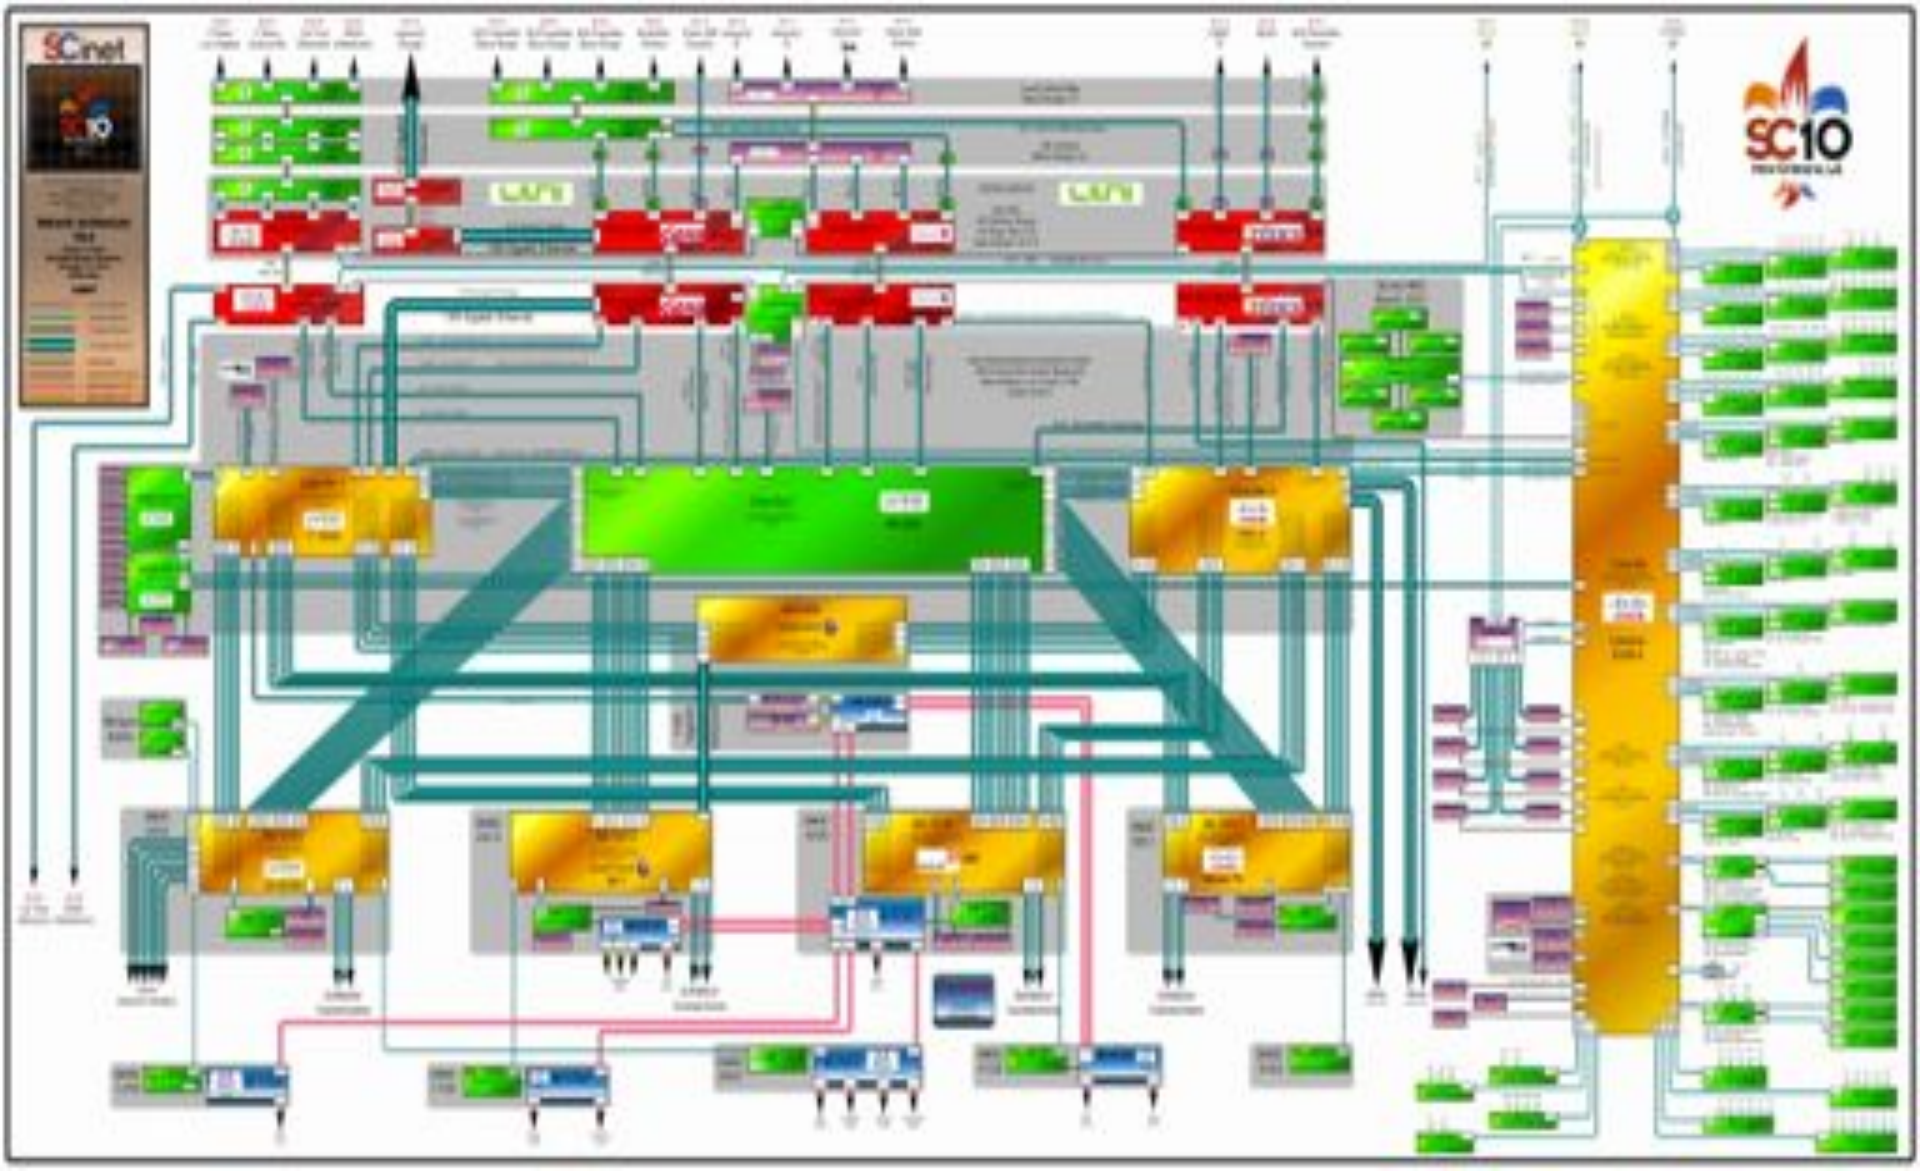

Complex e-Infrastructure!

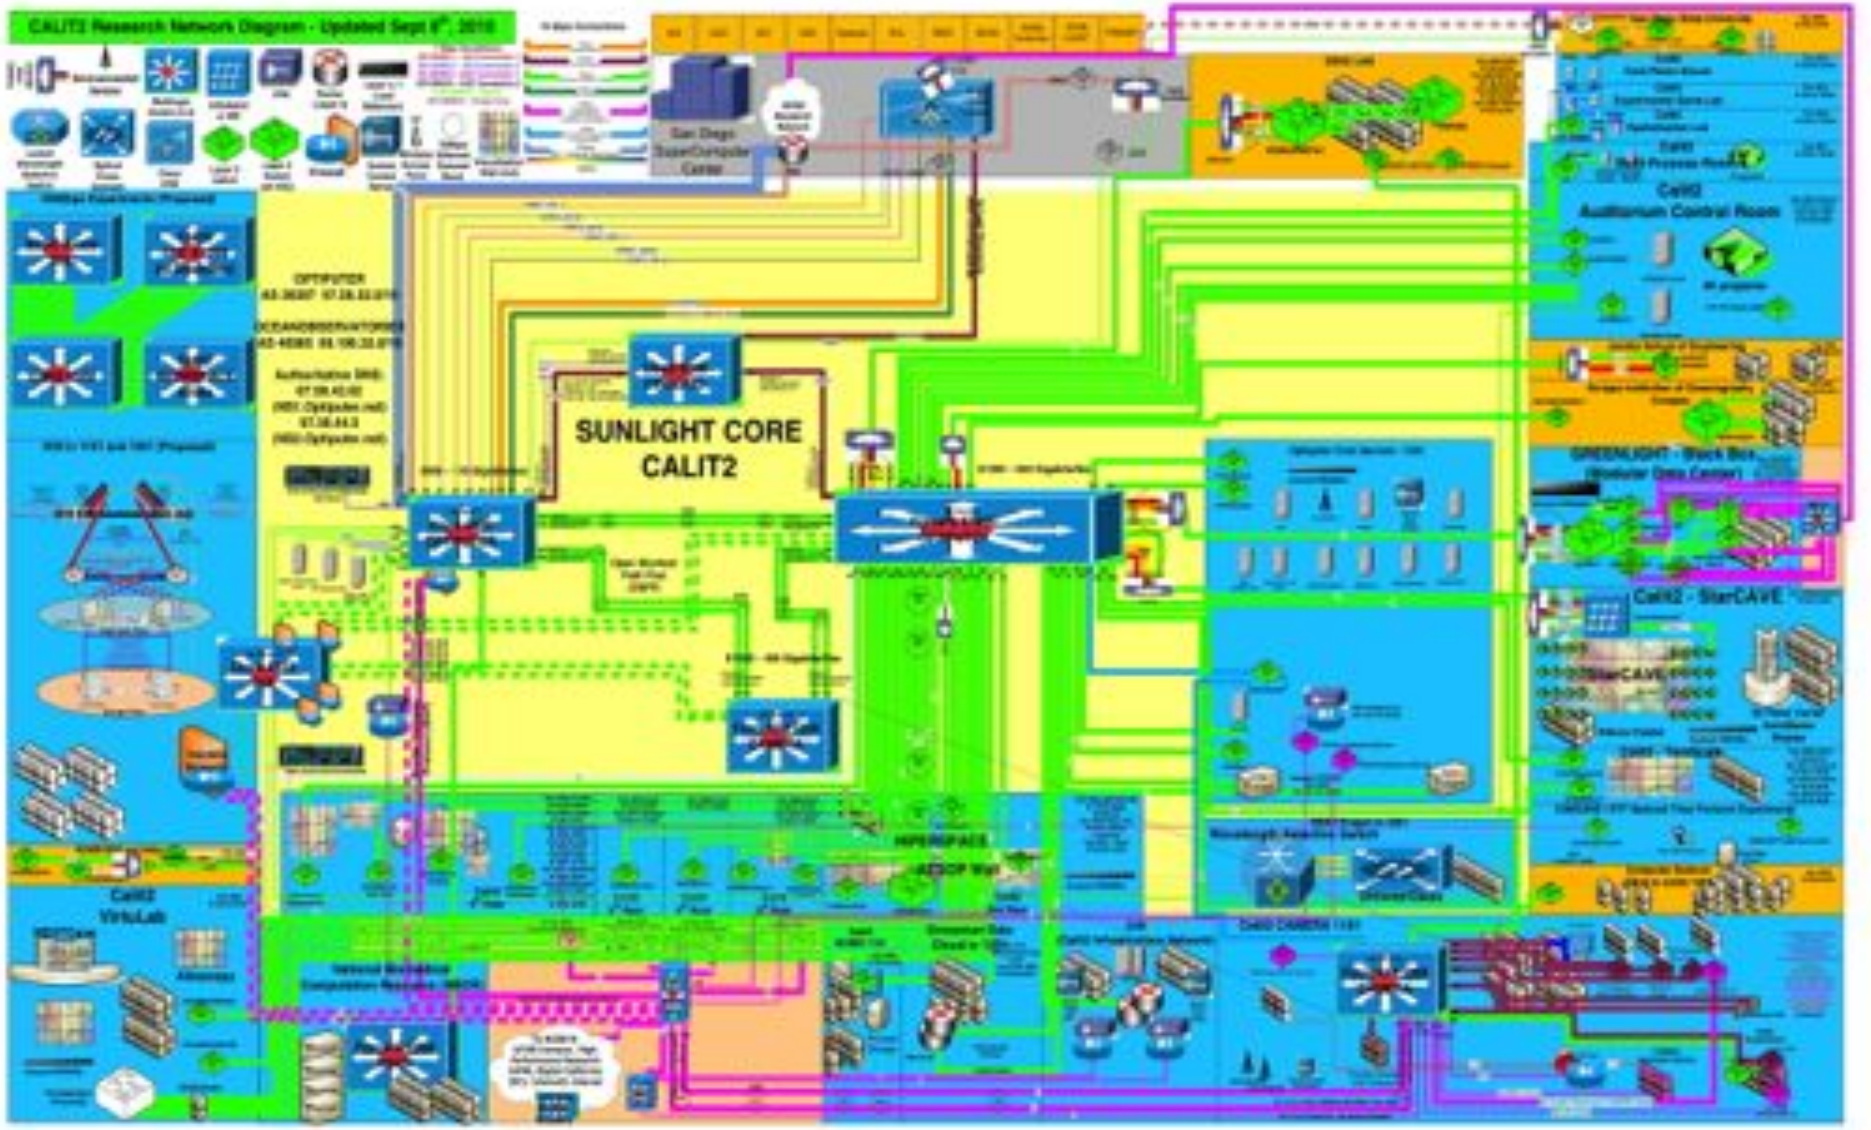

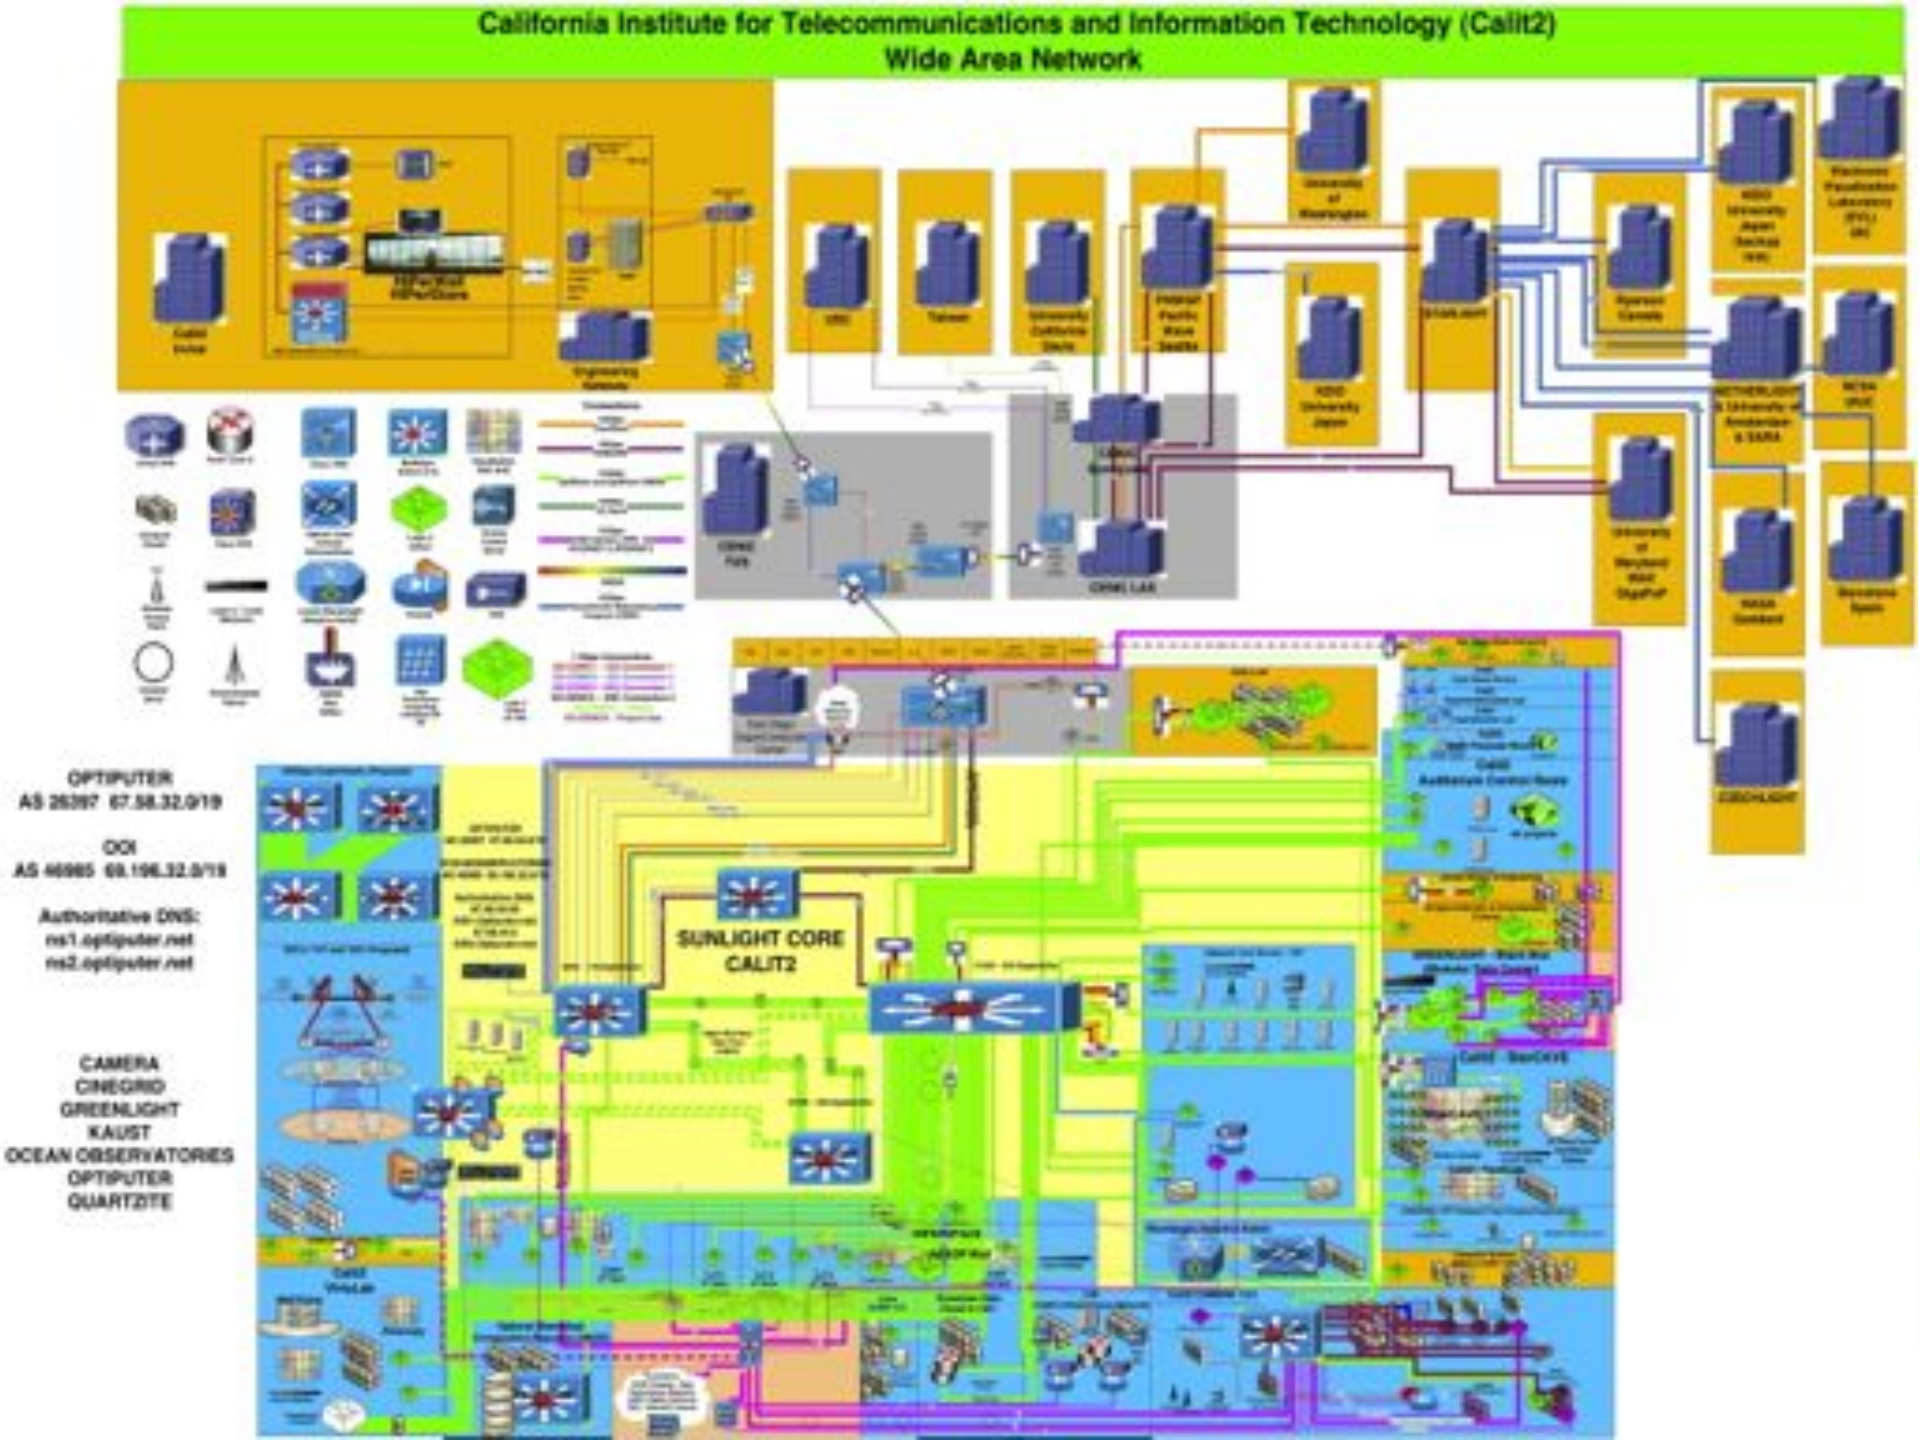
California Institute for Telecommunications and Information Technology (Calit2)
Wide Area Network
OPTIPUTER
AS 25397 67.58.32.0/19
OOI
AS 46985 69.196.32.0/19
Authoritative DNS:
ns1.optiputer.net
ns2.optiputer.net
CAMERA
CINEGRID
GREENLIGHT
KAUST
OCEAN OBSERVATORIES
OPTIPUTER
QUARTZITE
SUNLIGHT CORE
CALIT2

Semantics in Multiple Spaces

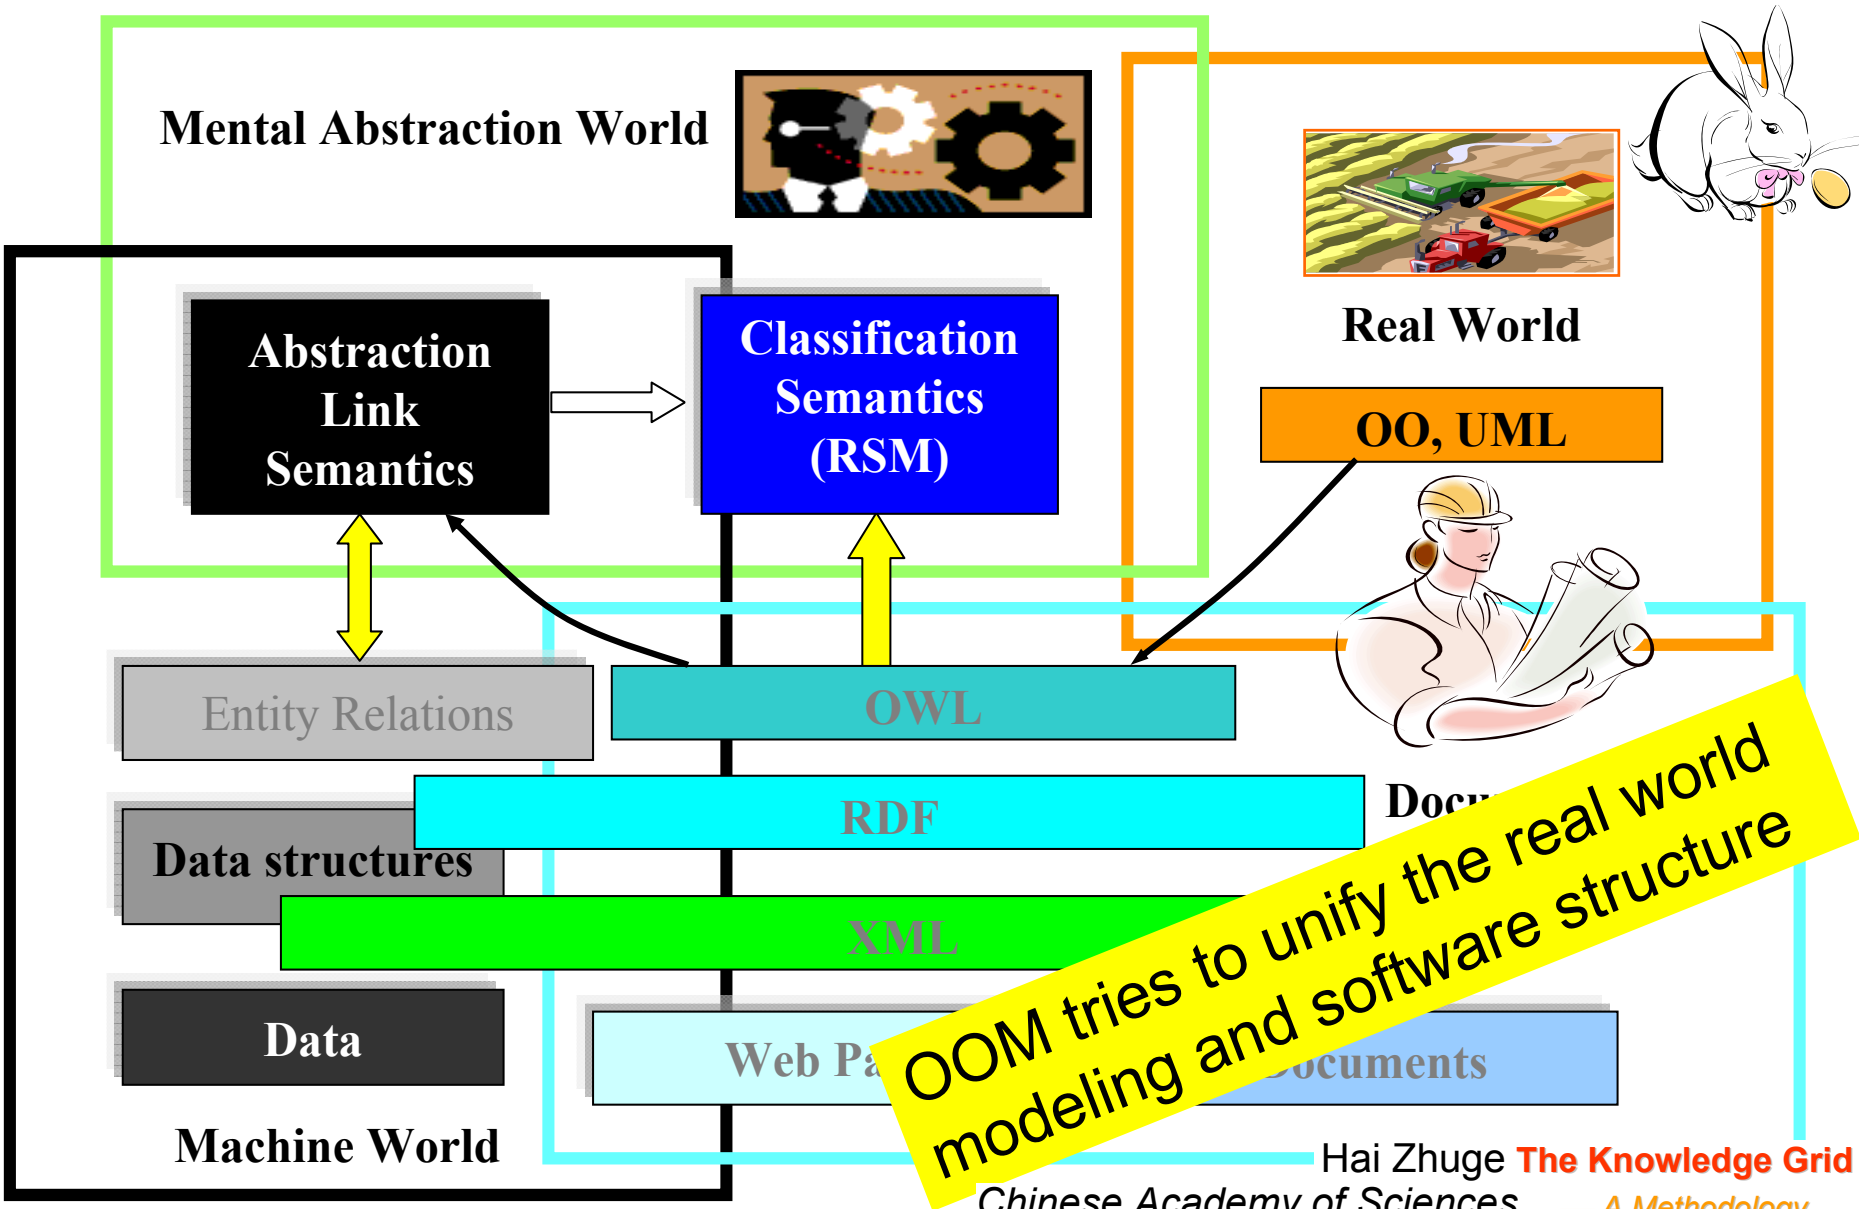

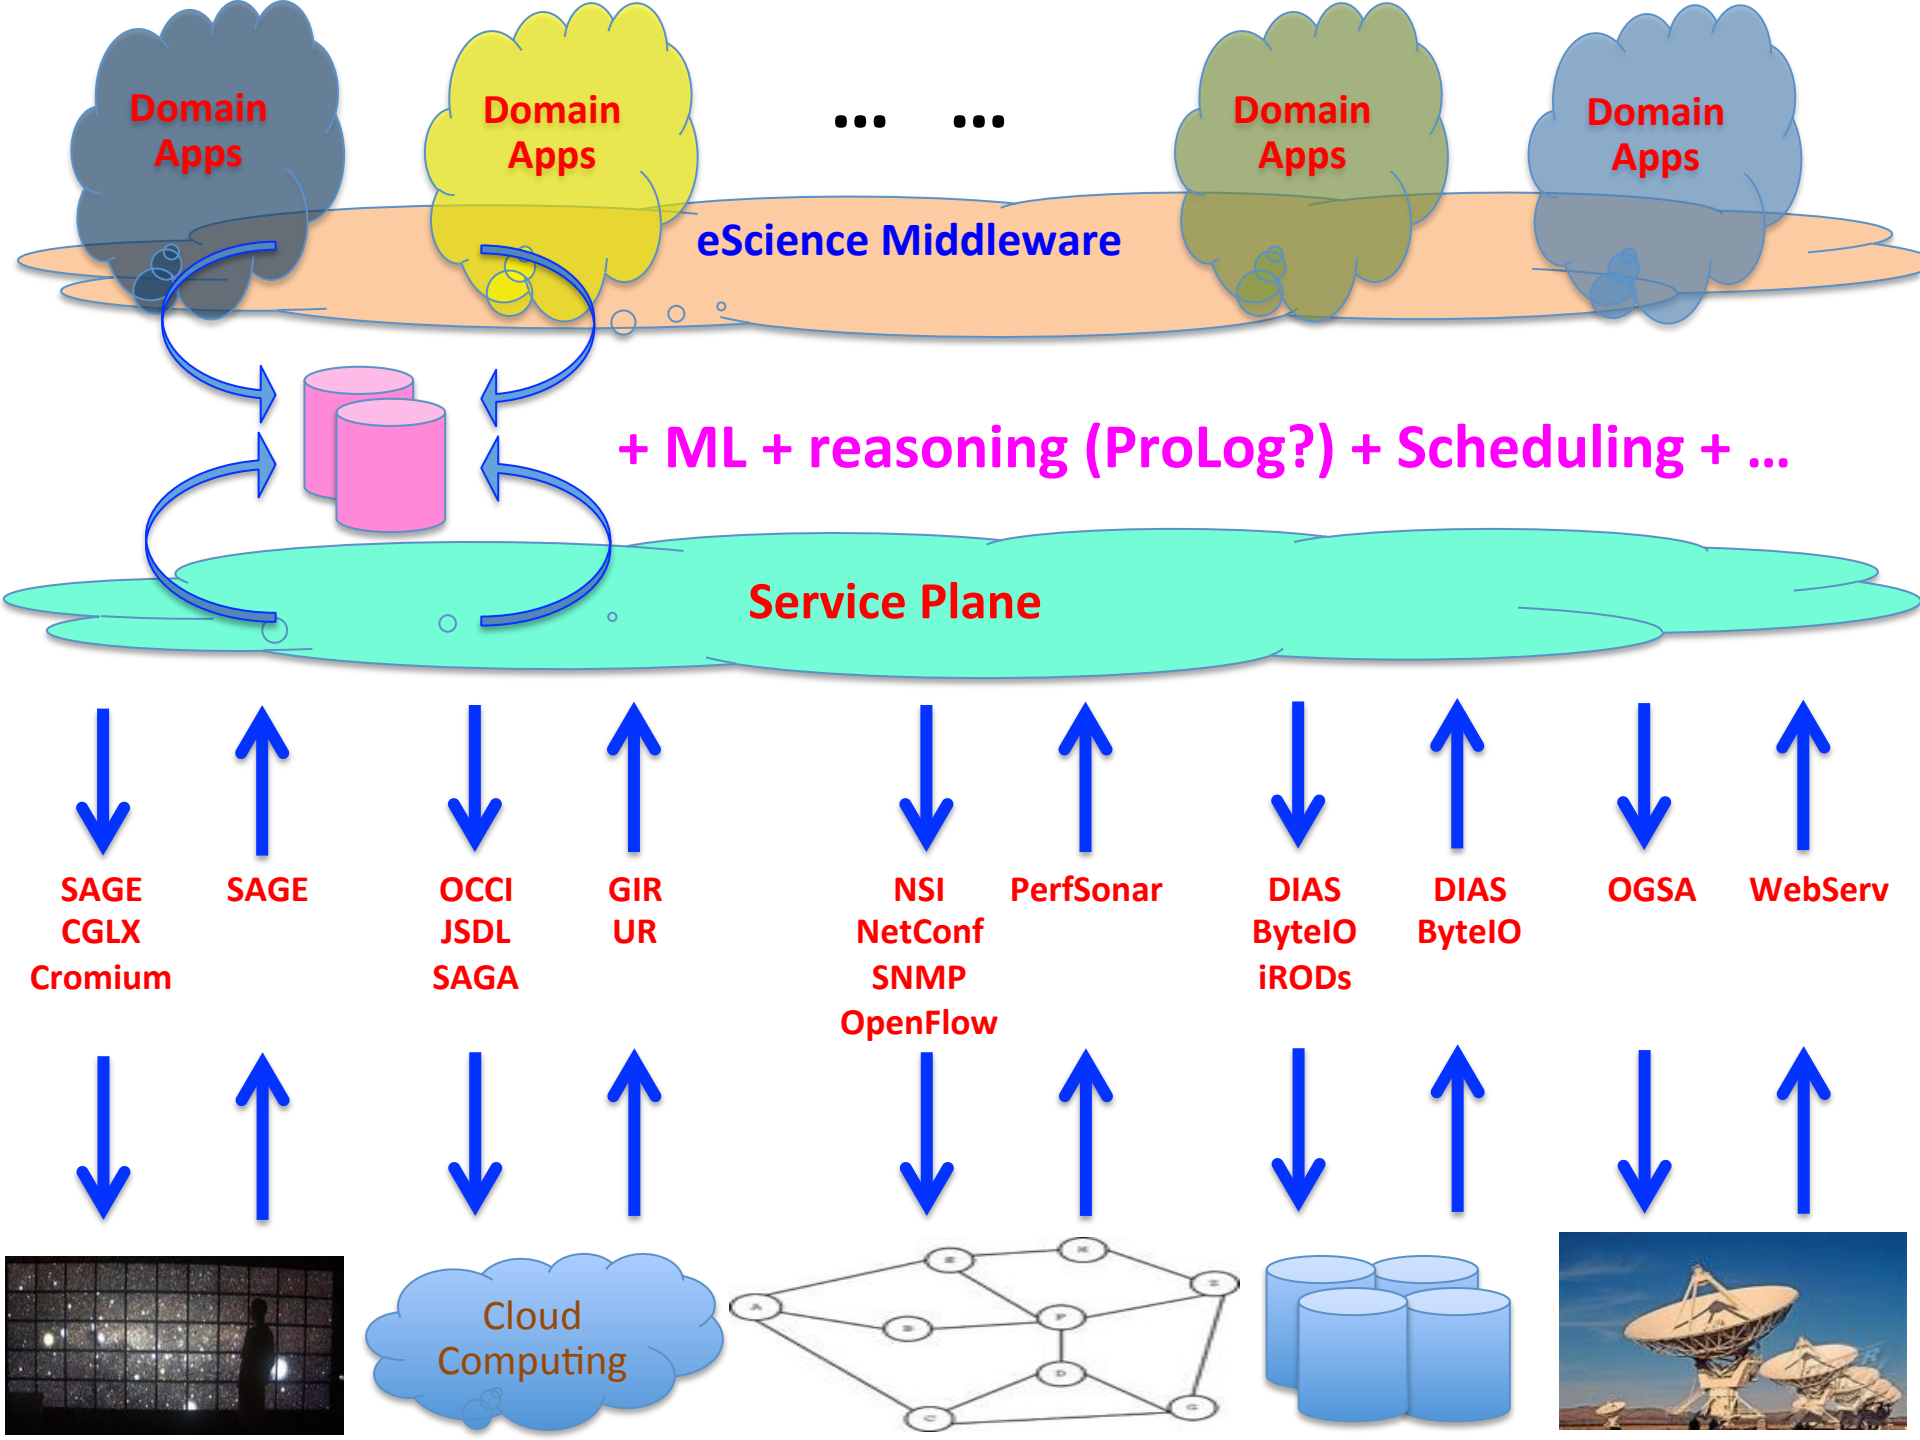
Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
+ ML + reasoning (ProLog?) + Scheduling + …
Service Plane
SAGE
CGLX
Cromium
SAGE
OCCI
JSDL
SAGA
GIR
UR
NSI
NetConf
SNMP
OpenFlow
PerfSonar
DIAS
ByteIO
iRODs
DIAS
ByteIO
OGSA
WebServ
Cloud Computing

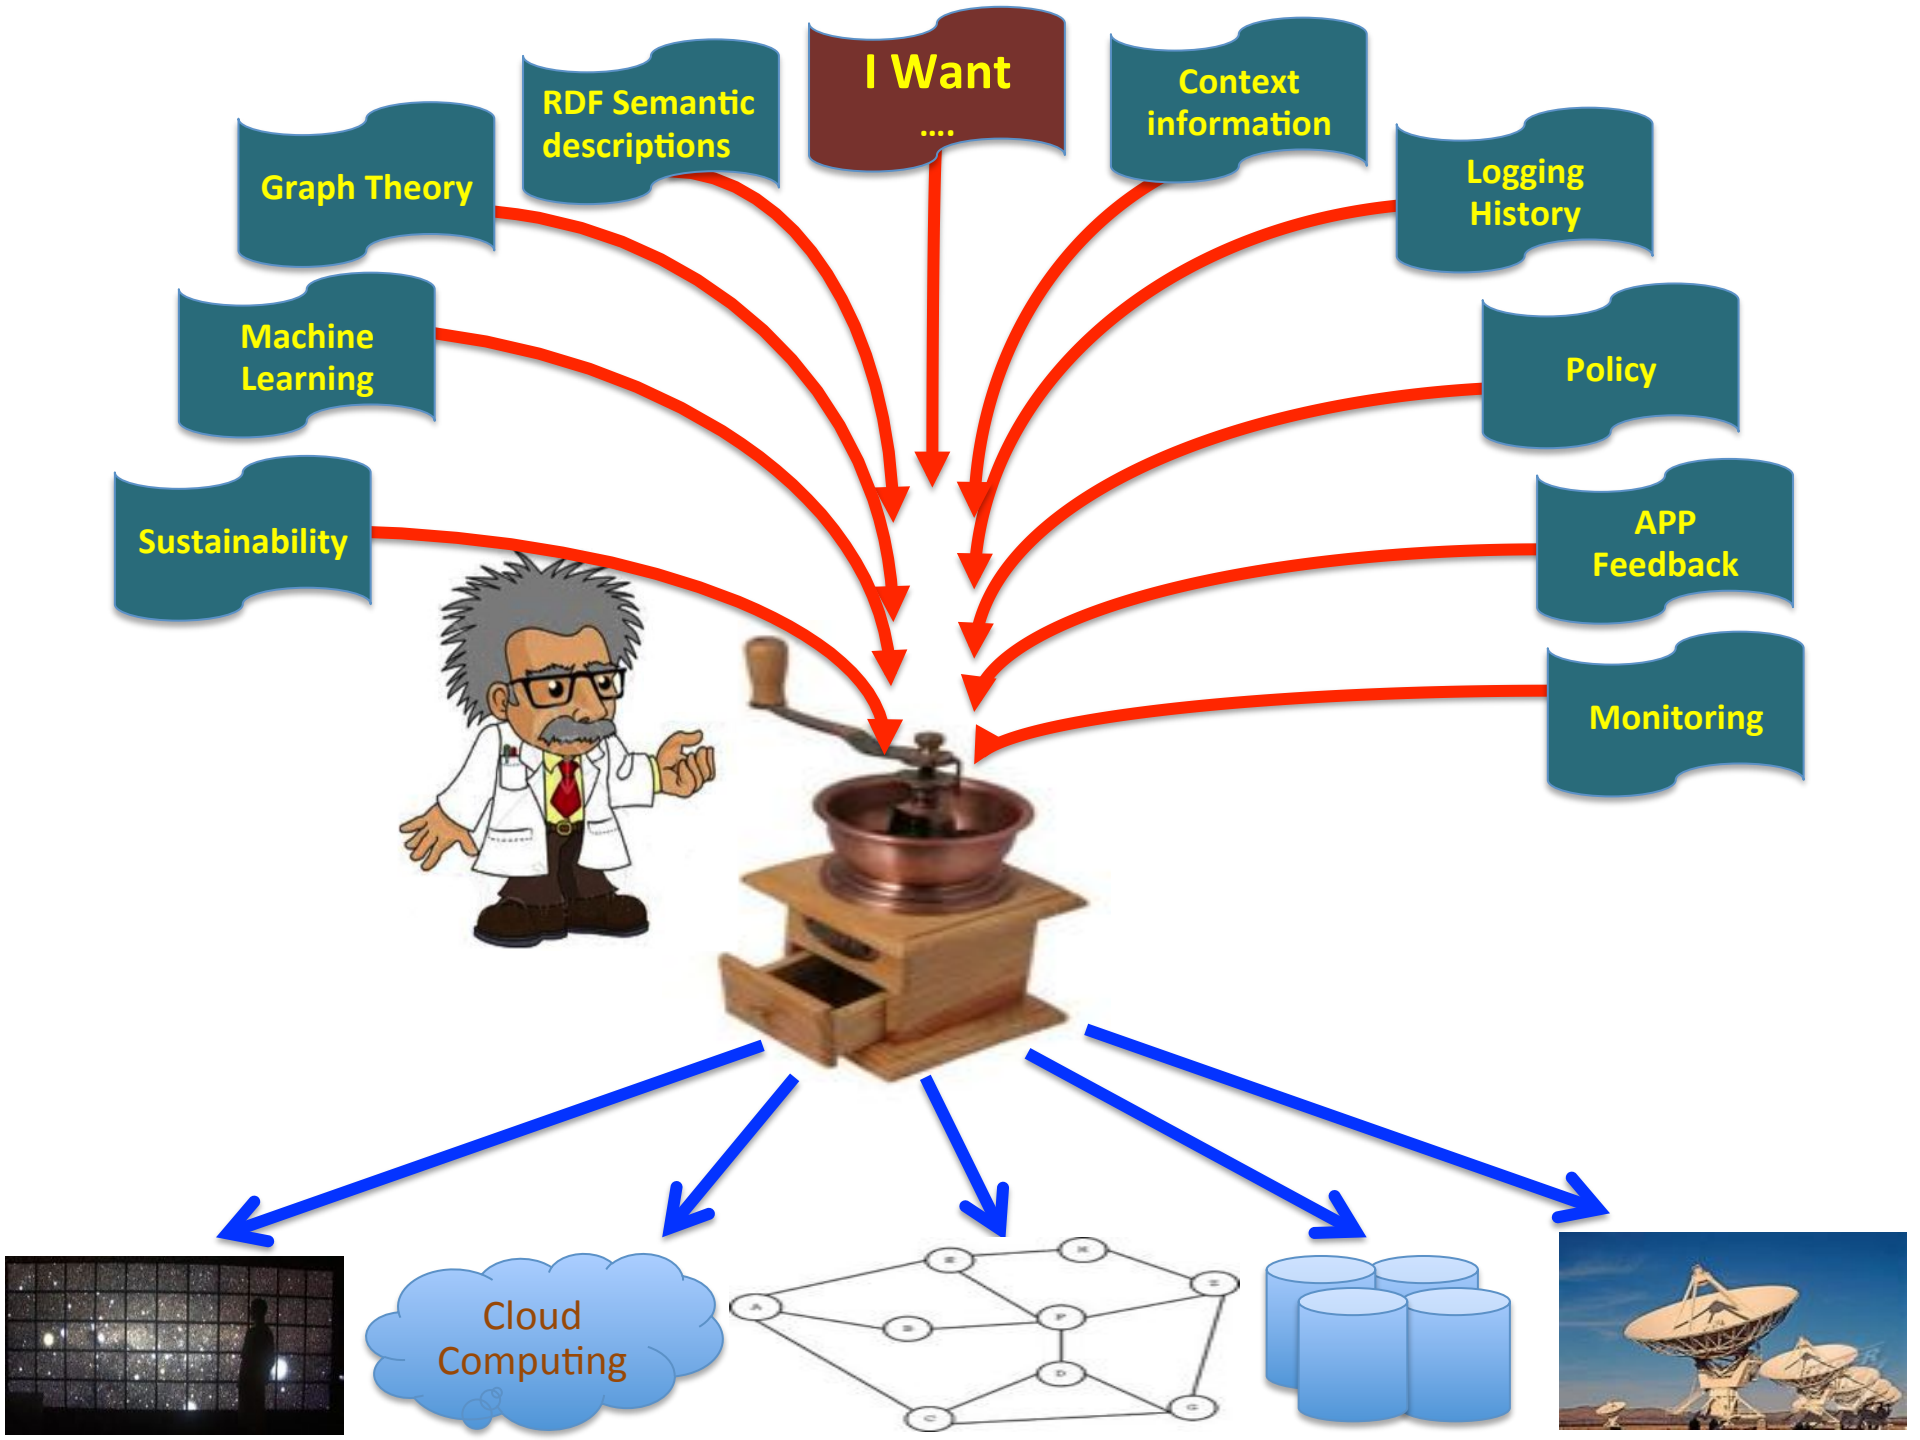
I Want
....
Graph Theory
RDF Semantic descriptions
Context information
Logging History
Machine Learning
Policy
Sustainability
APP Feedback
Monitoring
Cloud Computing

ECO-Scheduling

Hybrid Networking <-> Computing

Routers ⇐ ⇒ Supercomputers

Ethernet switches ⇐ ⇒ Grid & Cloud

Photonic transport ⇐ ⇒ GPU's

What matters:

- Energy consumption/multiplication
- Energy consumption/bit transported

Challenges

- Data – Data – Data
  - Archiving, publication, searchable, transport, self-describing, DB innovations needed, multi disciplinary use
- Virtualisation
  - Another layer of indeterminism
- Greening the Infrastructure
  - e.g. Department Of Less Energy: http://www.ecrinitiative.org/pdfs/ECR_3_0_1.pdf
- Disruptive developments
  - BufferBloath, Revisiting TCP, influence of SSD's & GPU's
  - Multi layer Glif Open Exchange model
  - Invariants in LightPaths (been there done that ☺)
    - X25, ATM, SONET/SDH, Lambda's, MPLS-TE, VLAN's, PBT, OpenFlow, ….
  - Authorization & Trust & Security and Privacy

The Way Forward!

- Nowadays scientific computing and data is dwarfed by commercial & cloud, there is also no scientific water, scientific power.
  - Understand how to work with elastic clouds
  - Trust & Policy & Firewalling on VM/Cloud level
- Technology cycles are 3 – 5 year
  - Do not try to unify but prepare for diversity
  - Hybrid computing & networking
  - Compete on implementation & agree on interfaces and protocols
- Limitation on natural resources and disruptive events
  - Energy becomes big issue
  - Follow the sun
  - Avoid single points of failure (aka Amazon, Blackberry, …)
  - Better very loosly coupled than totally unified integrated…

Q & A

http://ext.delaat.net/

Slides thanks to:

- Paola Grosso
- Sponsors see slide 1. ☺
- SNE Team & friends, see below