Smart Cyber Infrastructure for Big Data Processing

Cees de Laat

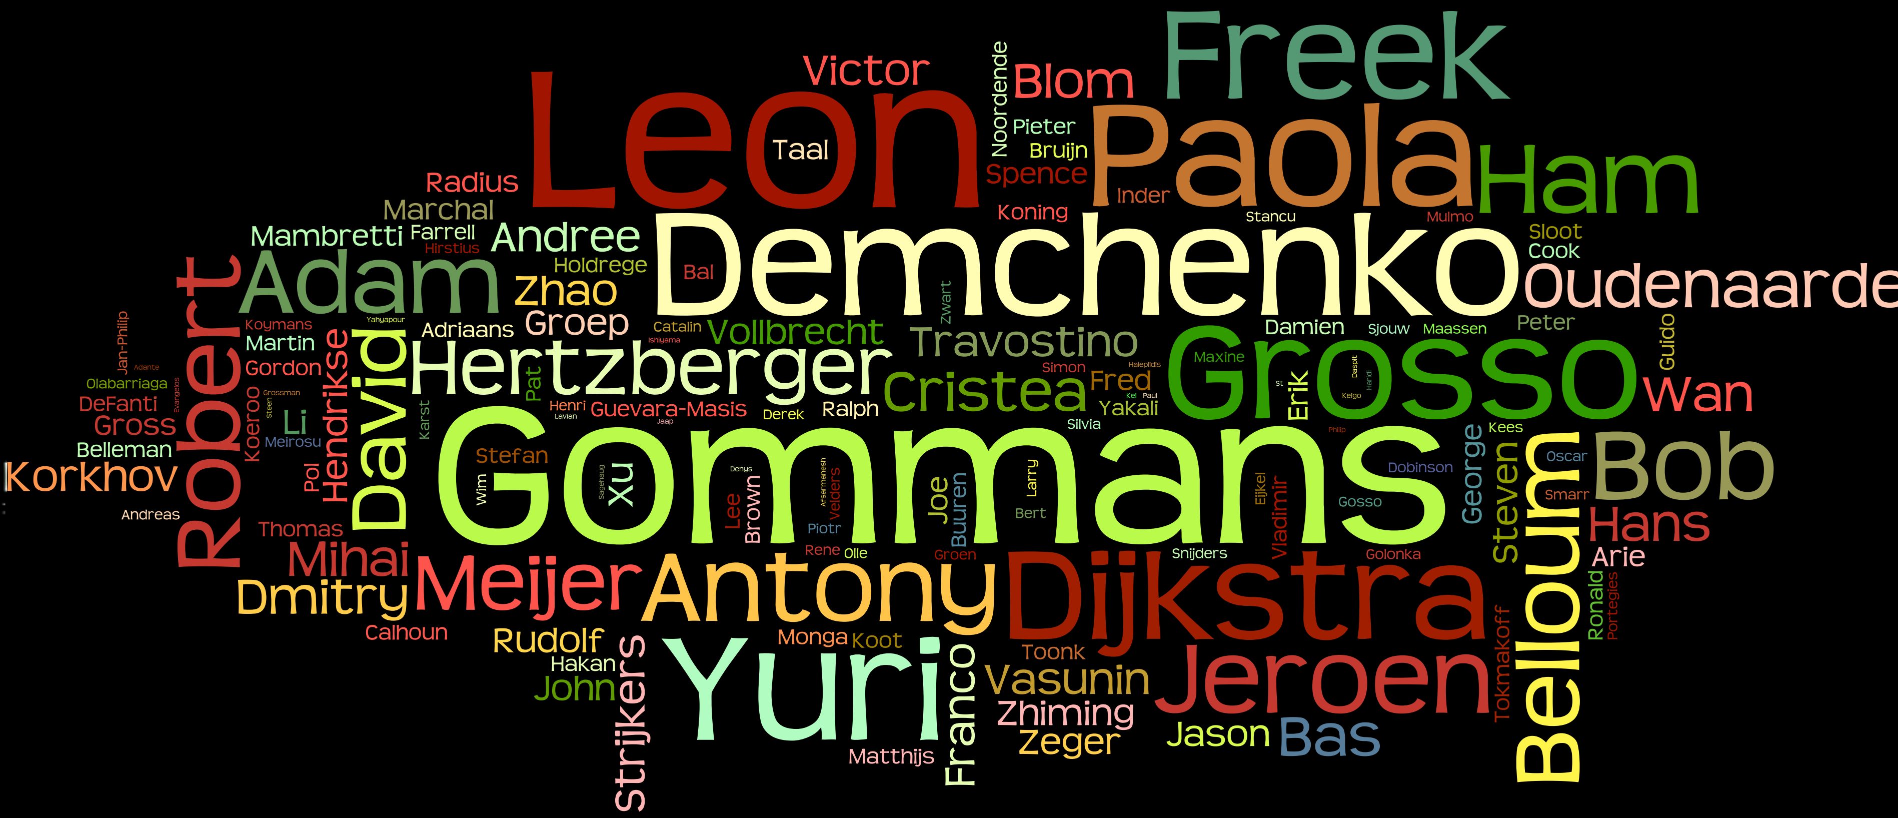

From King's Dutch Academy of Sciences

The Dutch Research Agenda

"Information technology (IT) now permeates all aspects of public, commercial, social, and personal life. bank cards, satnav, and weather radar... IT has become completely indispensable."

"But to guarantee the reliability and quality of constantly bigger and more complicated IT, we will need to find answers to some fundamental questions!"

https://www.knaw.nl/nl/actueel/publicaties/the-dutch-research-agenda/@@download/pdf_file/20111029.pdf

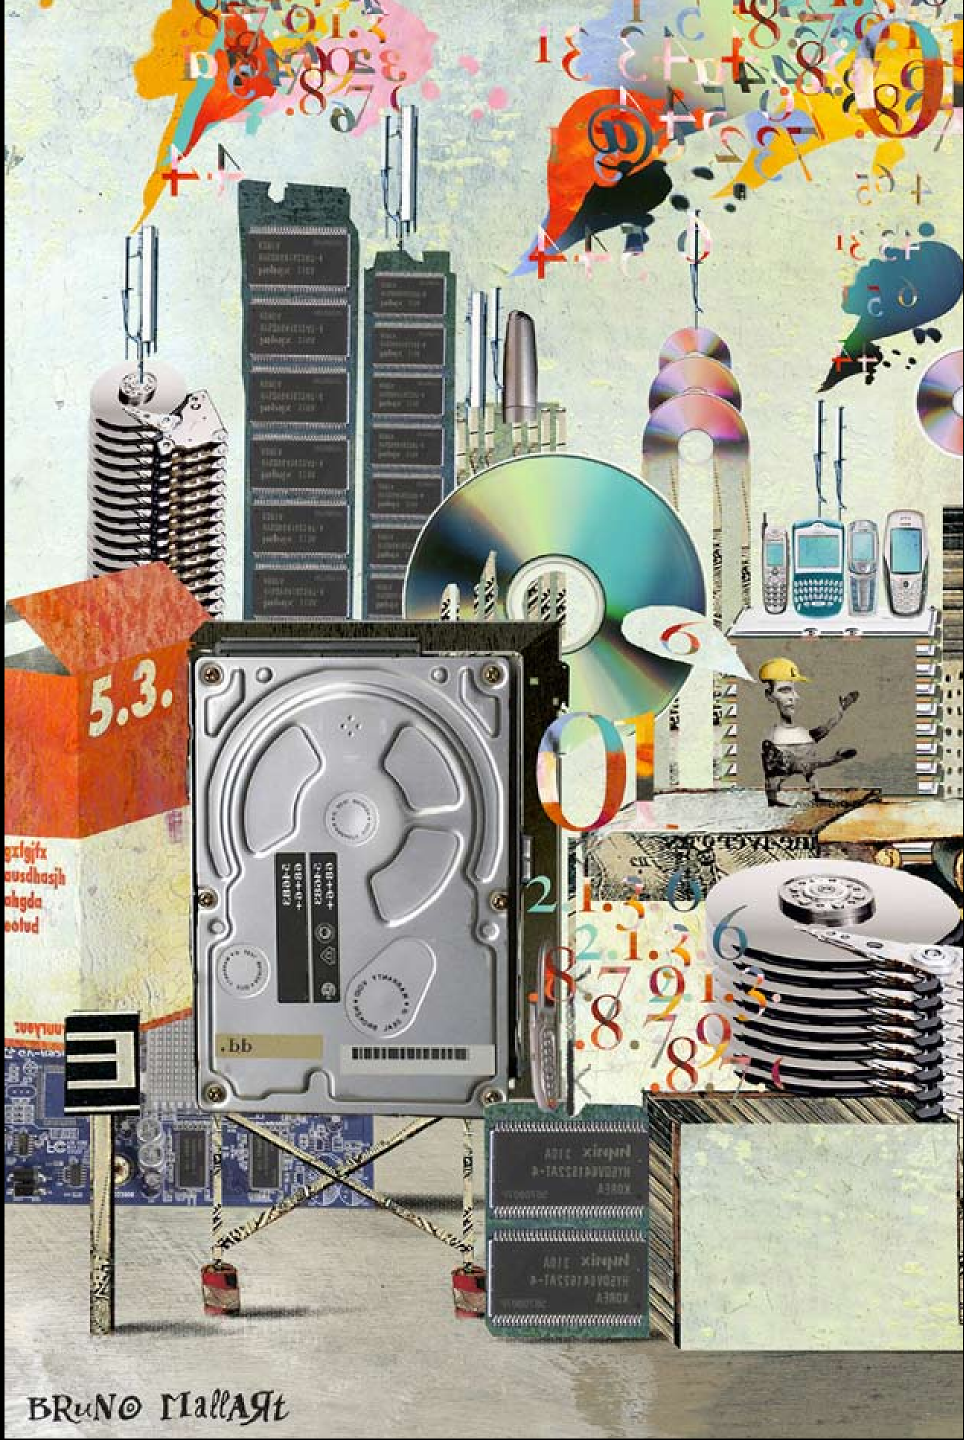

Reduction of Complexity by Integration

By combining services such as telephony, television, data, and computing capacity within a single network, we can cut down on complexity, energy consumption and maintenance.

- How can we describe and analyze complex information systems effectively?
- How can we specify and measure the quality and reliability of a system?
- How can we combine various different systems?
- How can we design systems in which separate processors can co-operate efficiently via mutual network connections within a much larger whole?
- Can we design information systems that can diagnose their own malfunctions and perhaps even repair them?
- How can we specify, predict, and measure system performance as effectively as possible?

SNE addresses a.o. the **highlighted** questions!

http://www.knaw.nl/Content/Internet_KNAW/publicaties/pdf/20111029.pdf

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- *Capacity*
  - *Bandwidth on demand, QoS, architectures, photonics, performance*
- *Capability*
  - *Programmability, virtualization, complexity, semantics, workflows*
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- *Sustainability*
  - *Greening infrastructure, awareness*
- *Resilience*
  - *Systems under attack, failures, disasters*

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- ***Capacity***
  - ***Bandwidth on demand, QoS, architectures, photonics, performance***
- *Capability*
  - *Programmability, virtualization, complexity, semantics, workflows*
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- *Sustainability*
  - *Greening infrastructure, awareness*
- *Resilience*
  - *Systems under attack, failures, disasters*

ATLAS detector @ CERN Geneve

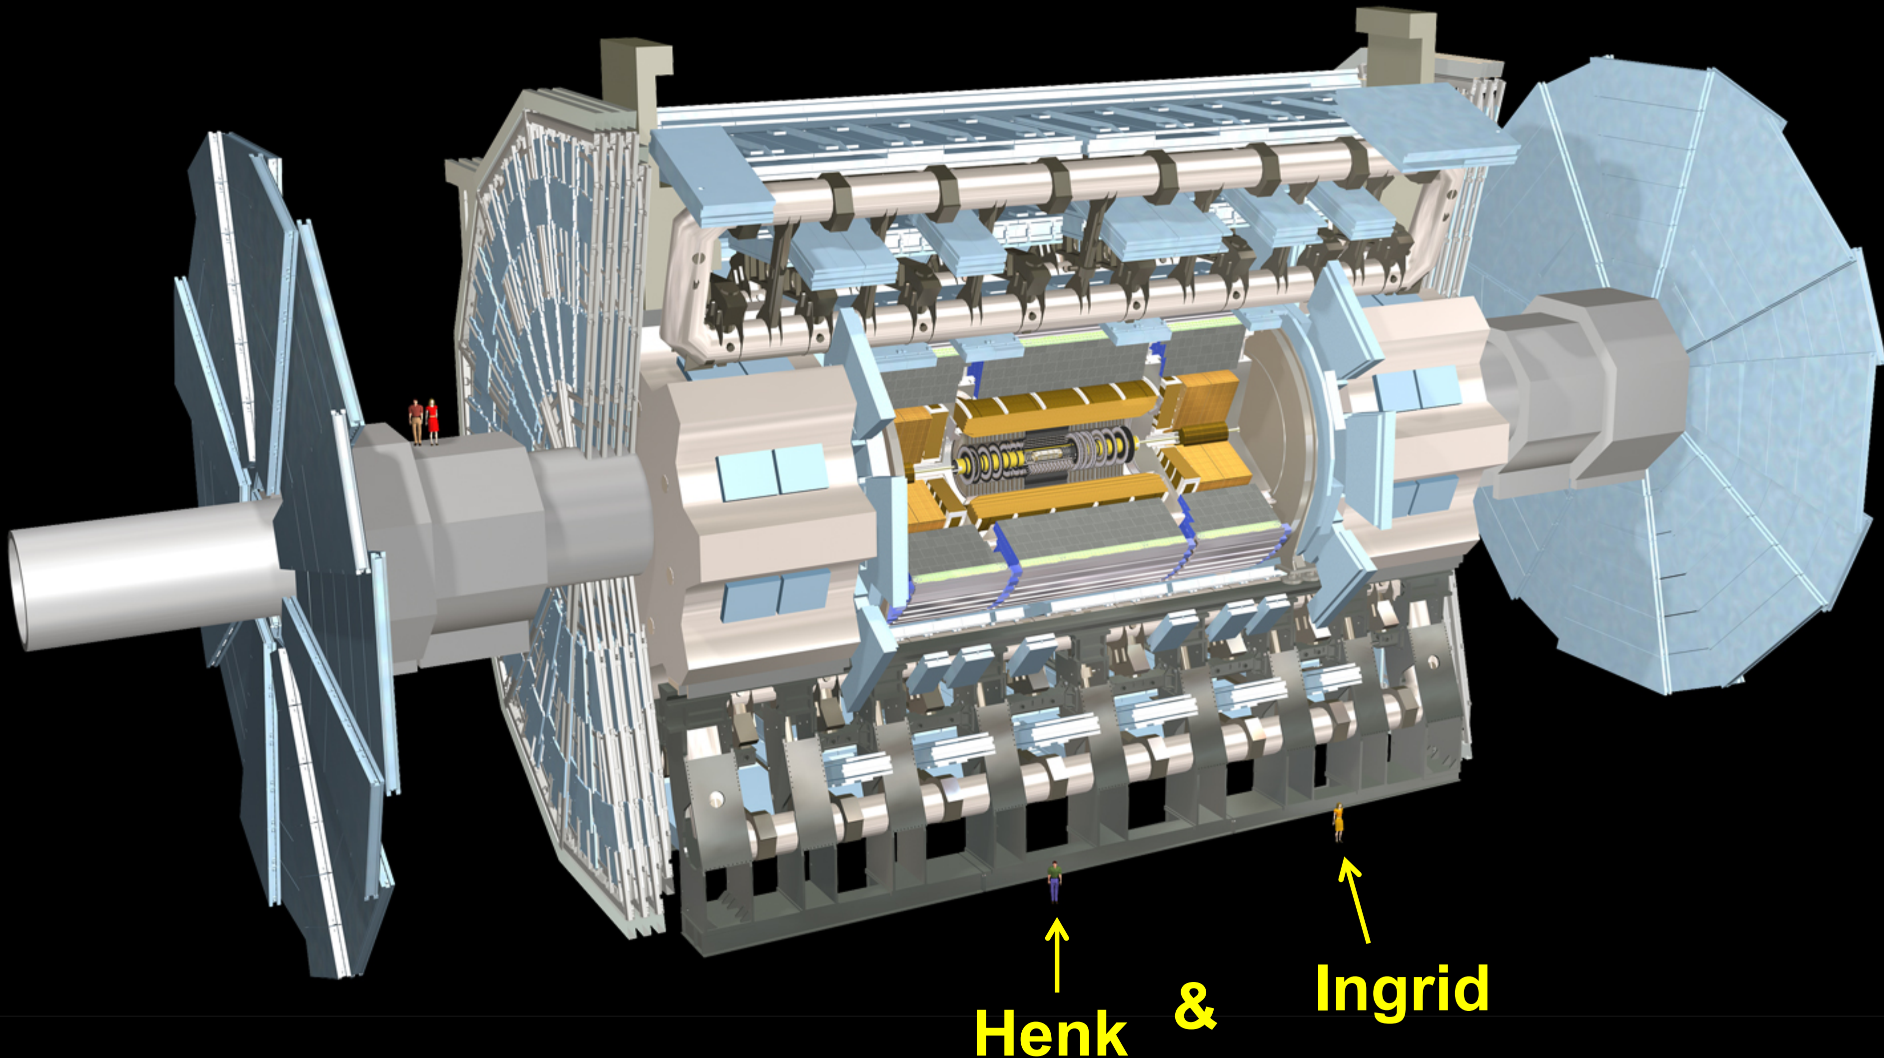

What Happens in an **Internet Minute?**

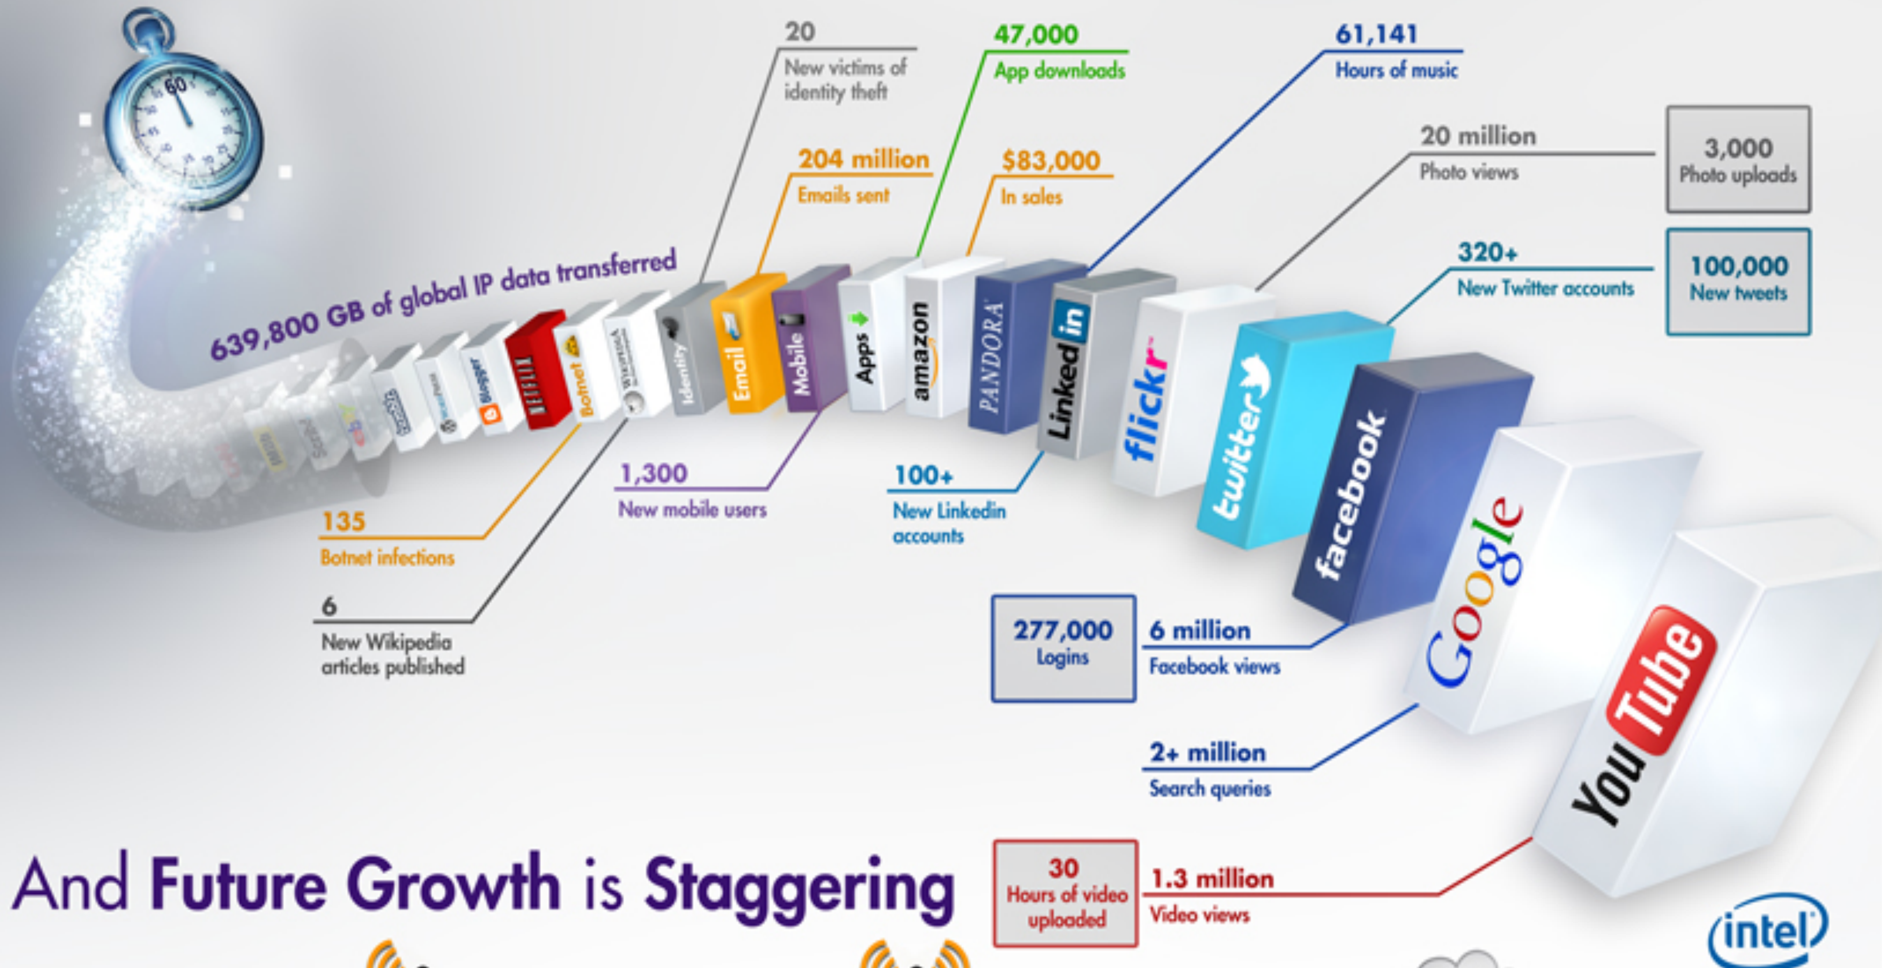

And **Future Growth** is **Staggering**

Today, the number of networked devices = the global population

By **2015**, the number of networked devices = 2x the global population

In **2015**, it would take you **5 years** to view all video crossing IP networks each **second**

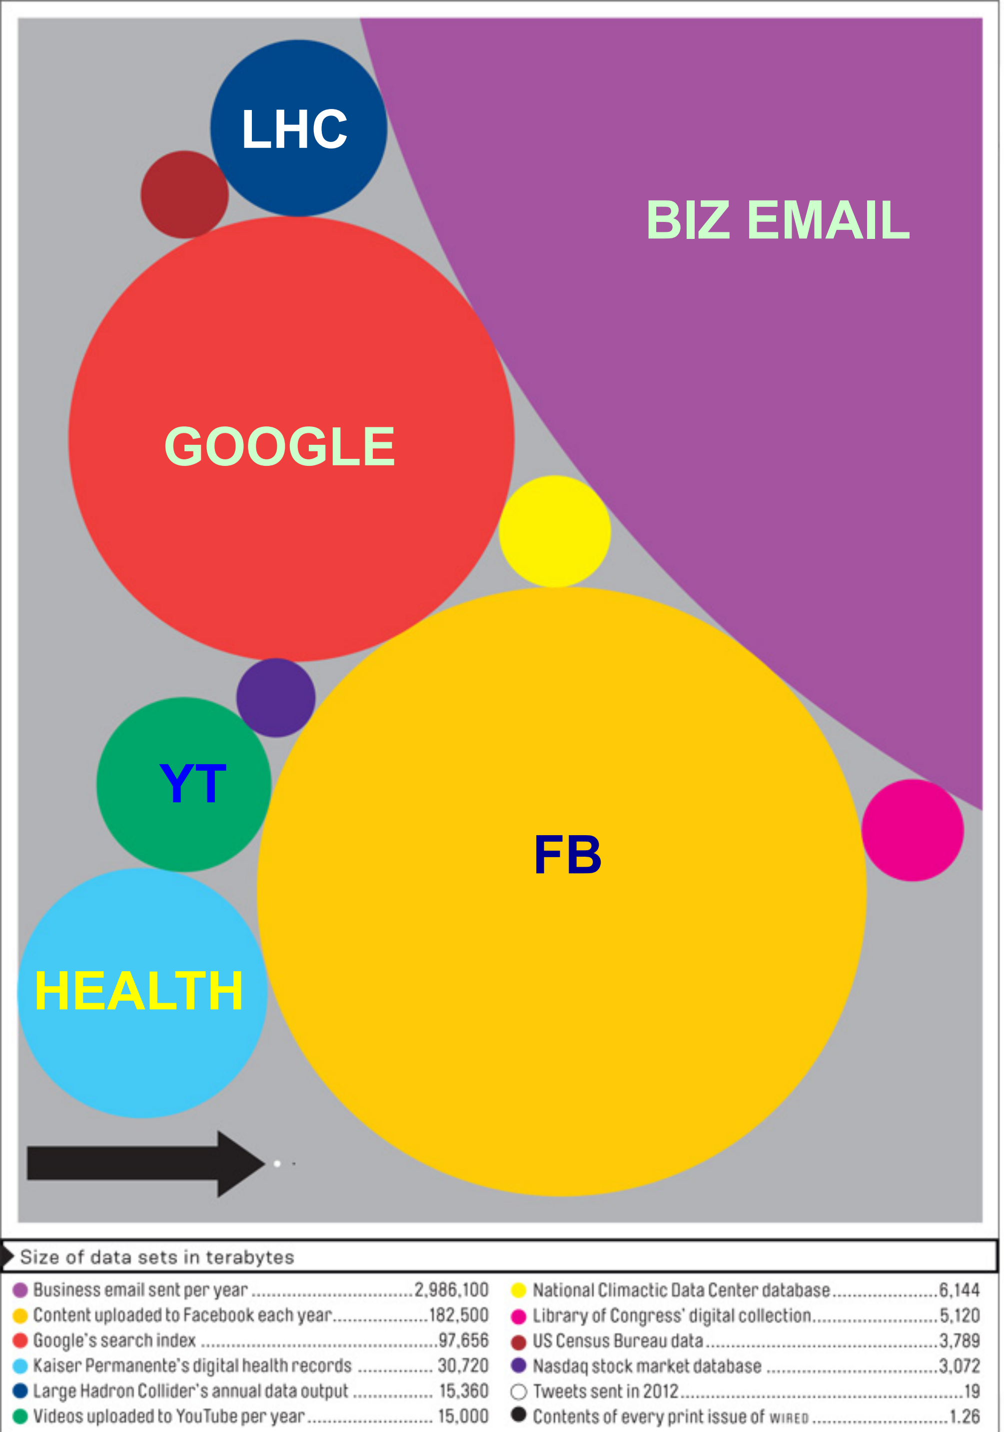

There is always a bigger fish

Internet developments

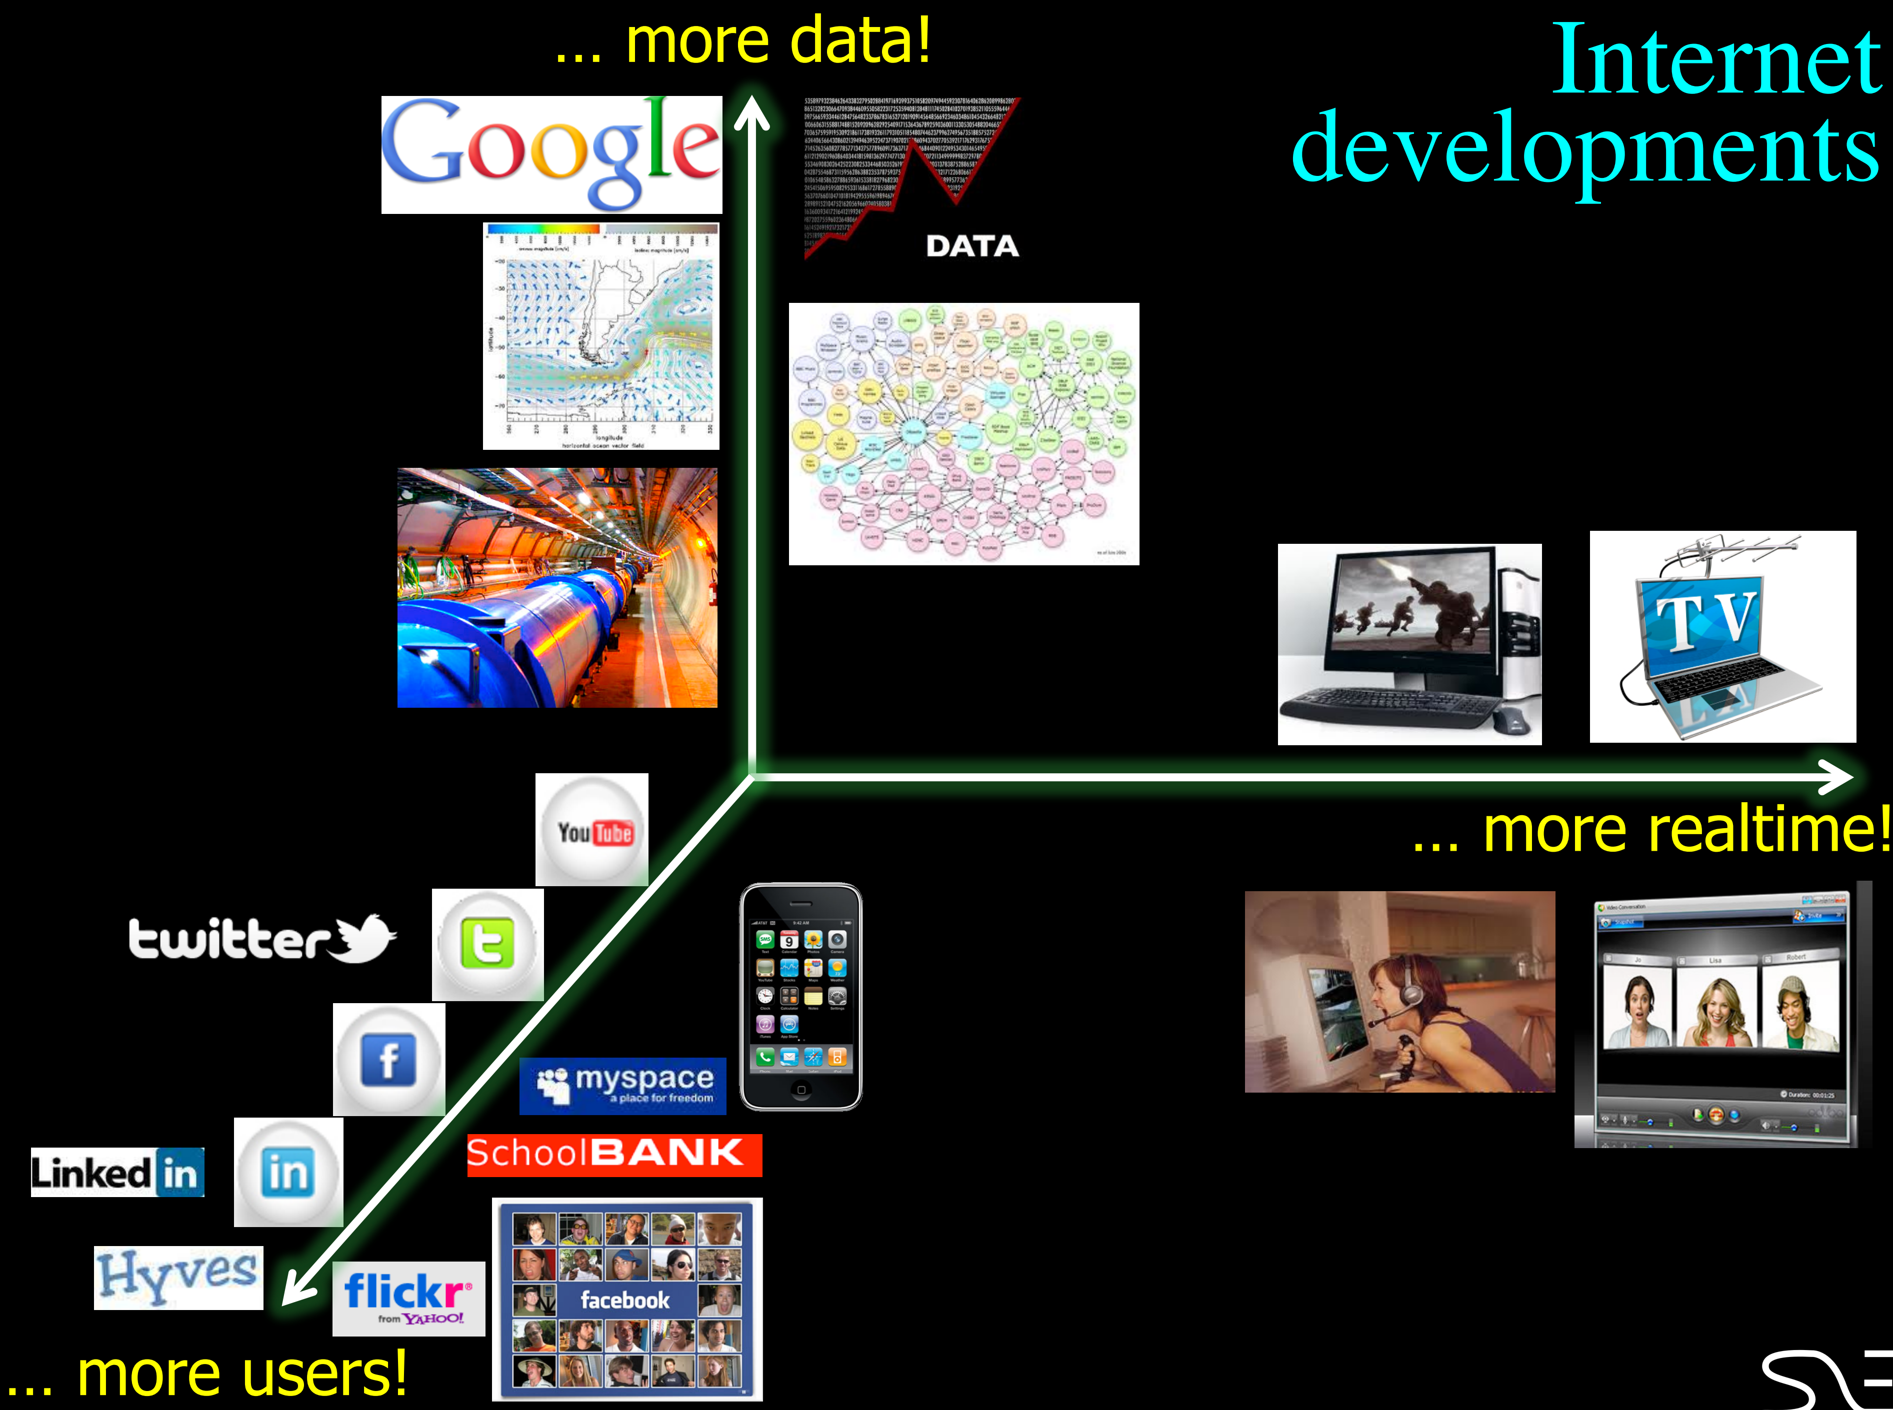

Internet developments

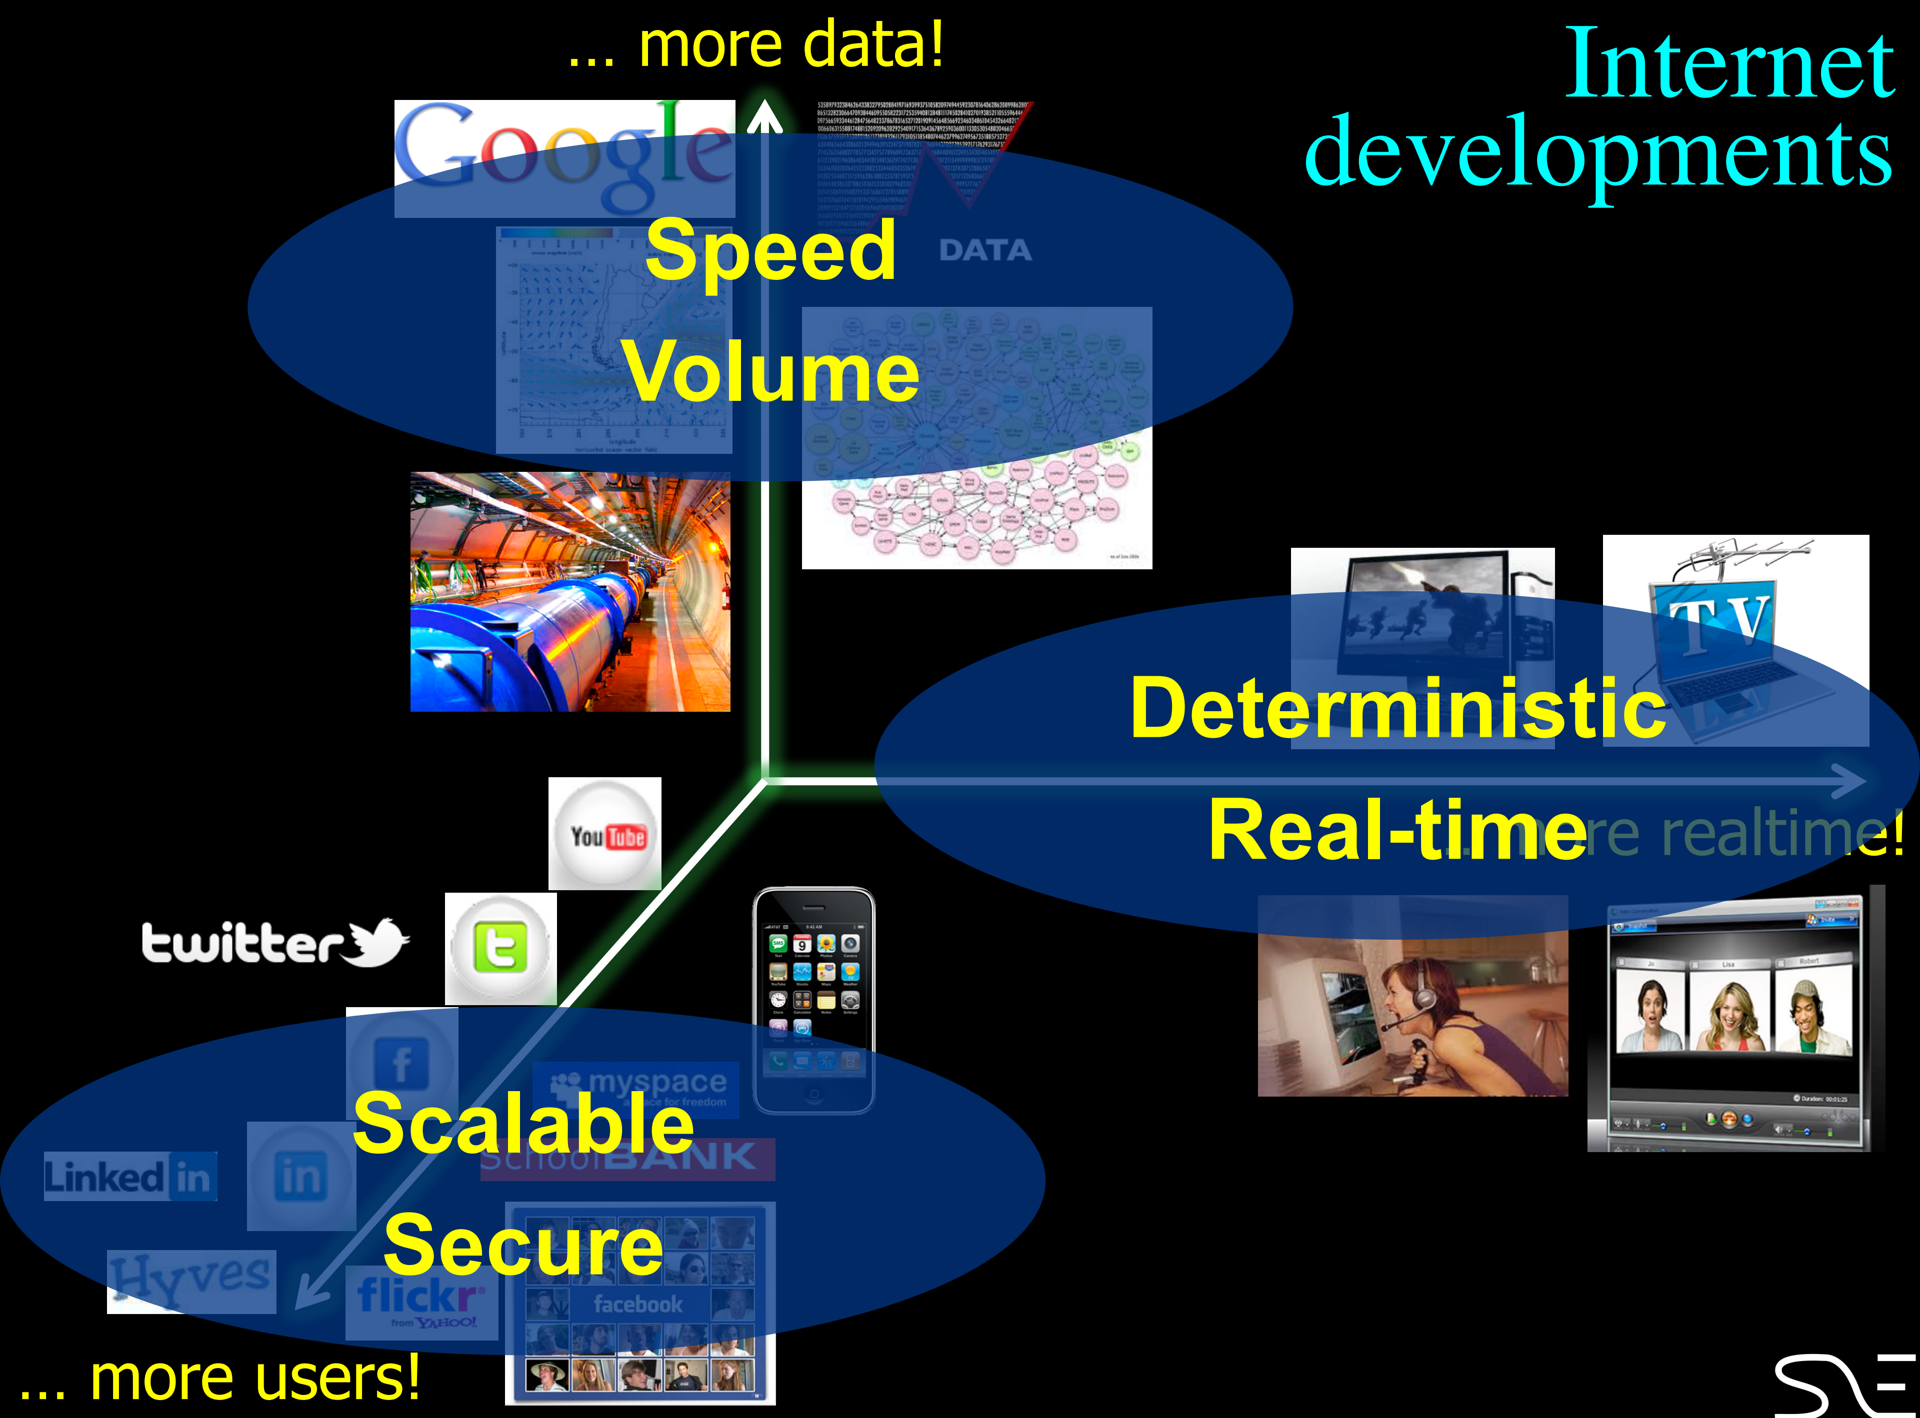

Multiple colors / Fiber

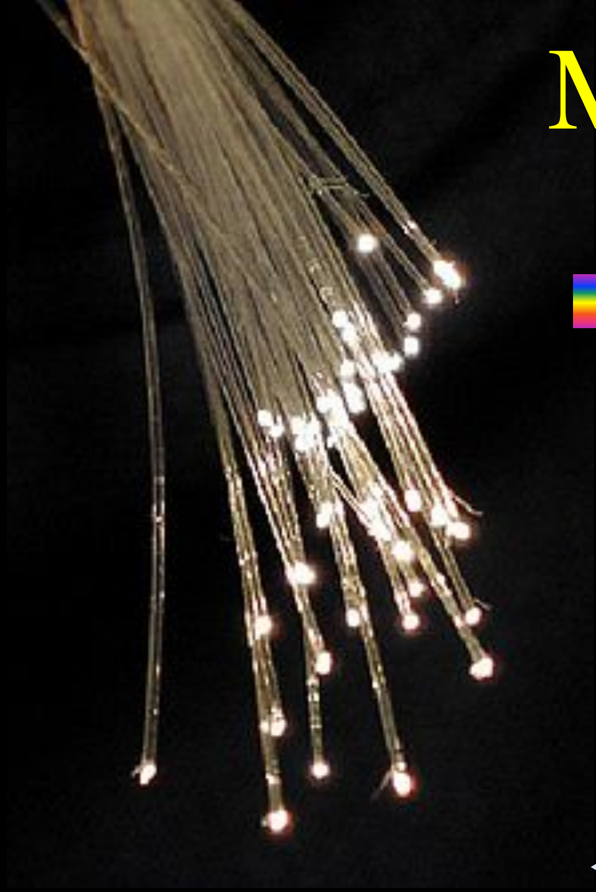

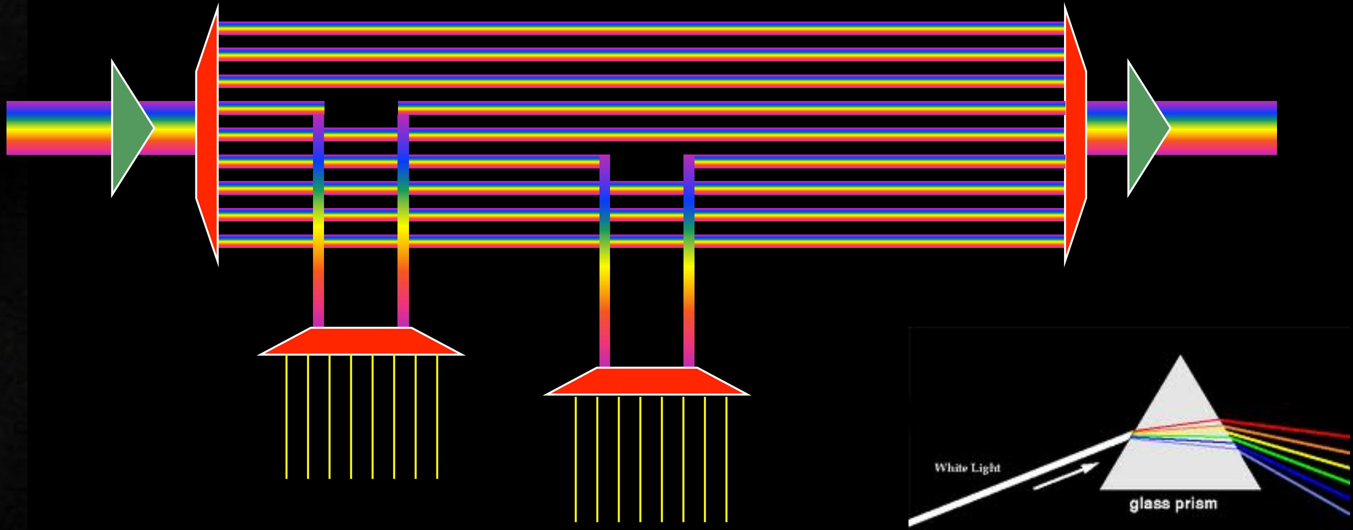

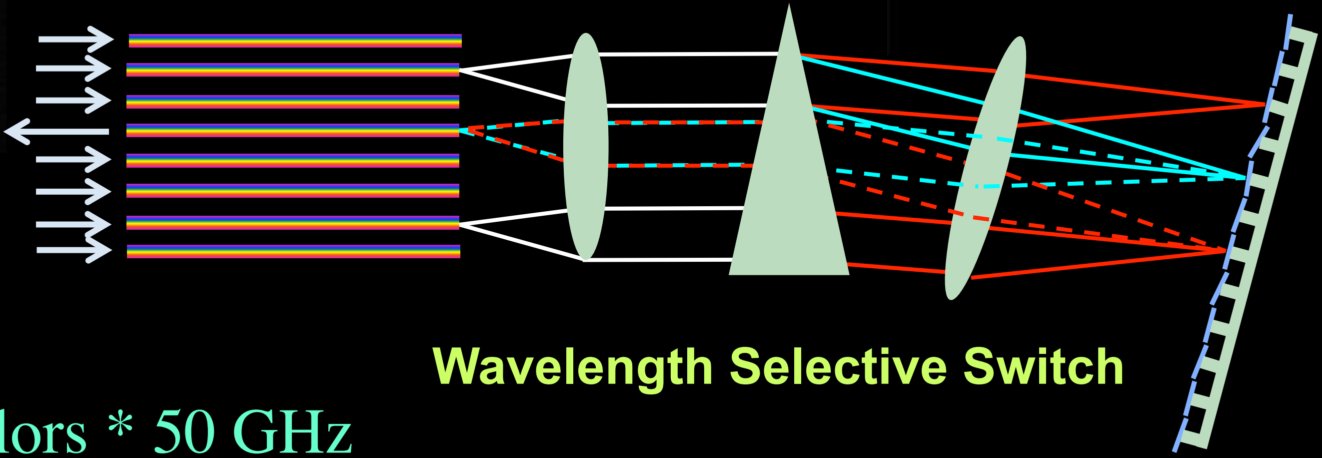

Wavelength Selective Switch

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!

Wireless Networks

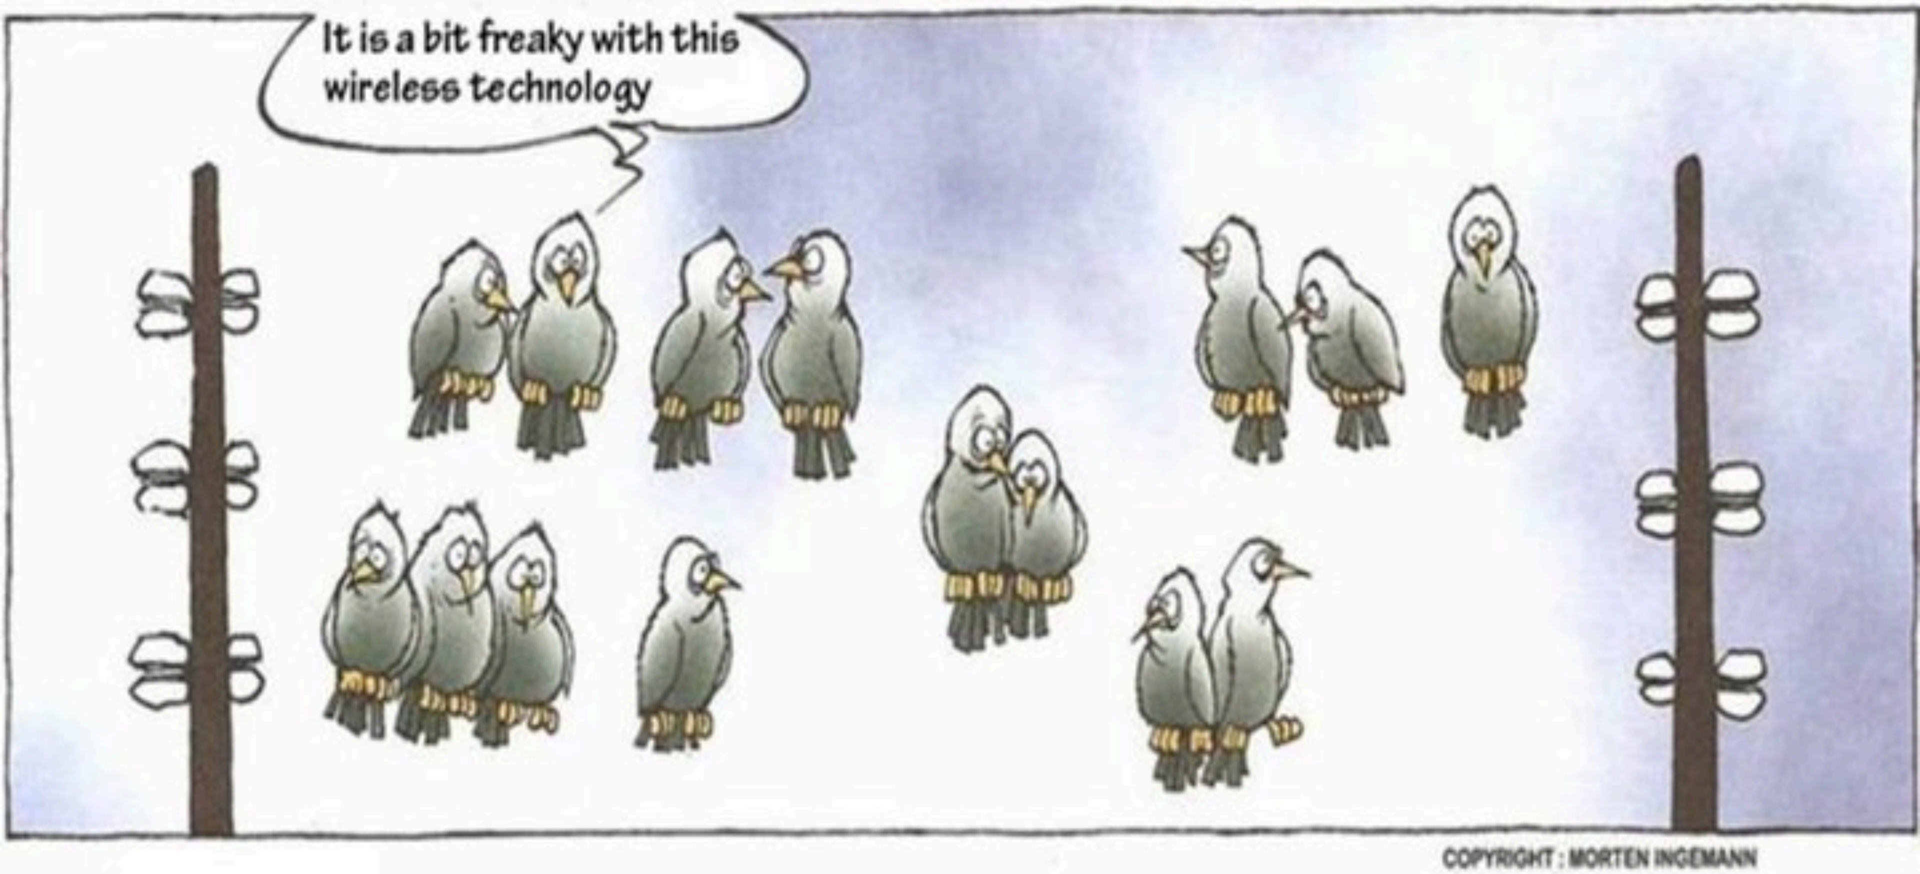

protocol LAN due to the easy comparison and convenience in the digital home. While consumer PC products has just started to migrate to a much higher bandwidth of 802.11n wireless LAN now working on next-generation standard definition is already in progress.

GPU cards are distruptive!

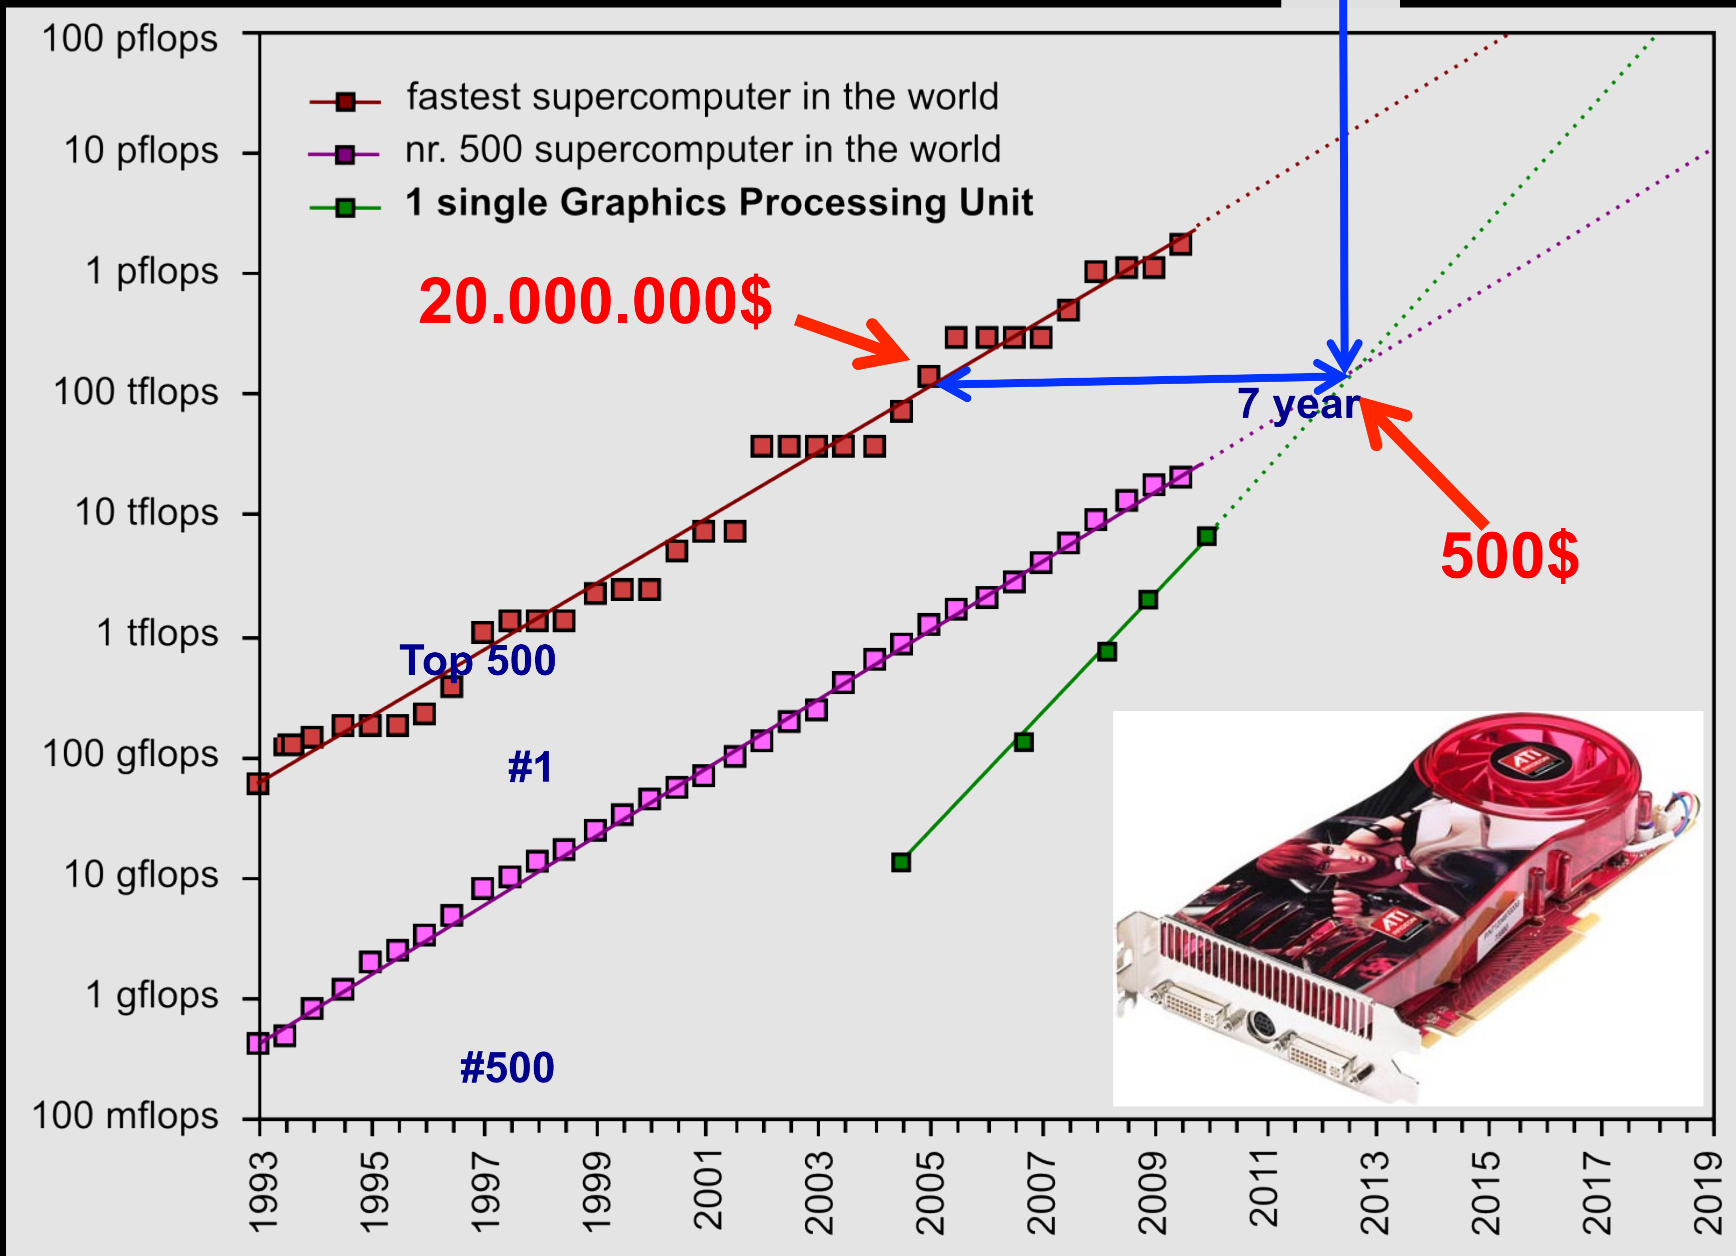

Reliable and Safe!

This omnipresence of IT makes us not only strong but also vulnerable.

- A virus, a hacker, or a system failure can instantly send digital shockwaves around the world.

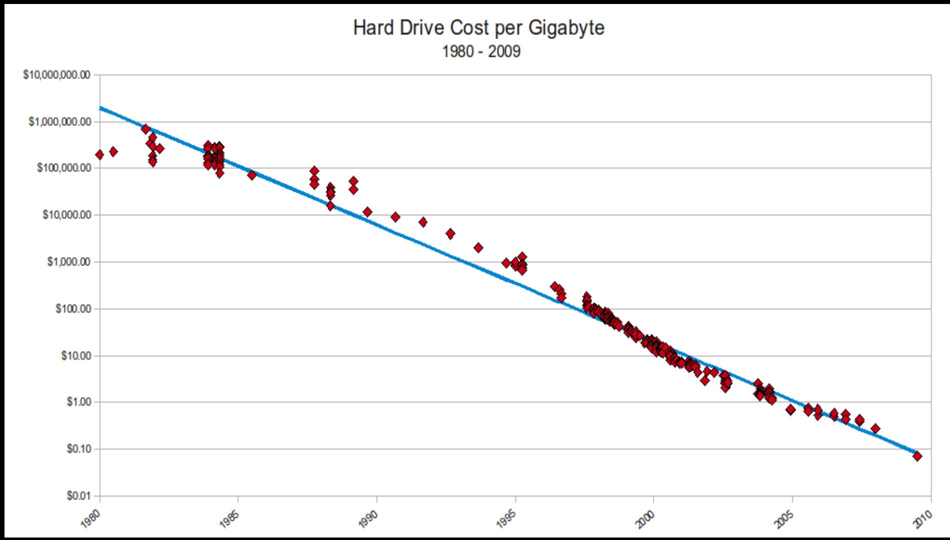

The hardware and software that allow all our systems to operate is becoming bigger and more complex all the time, and the capacity of networks and data storage is increasing by leaps and bounds.

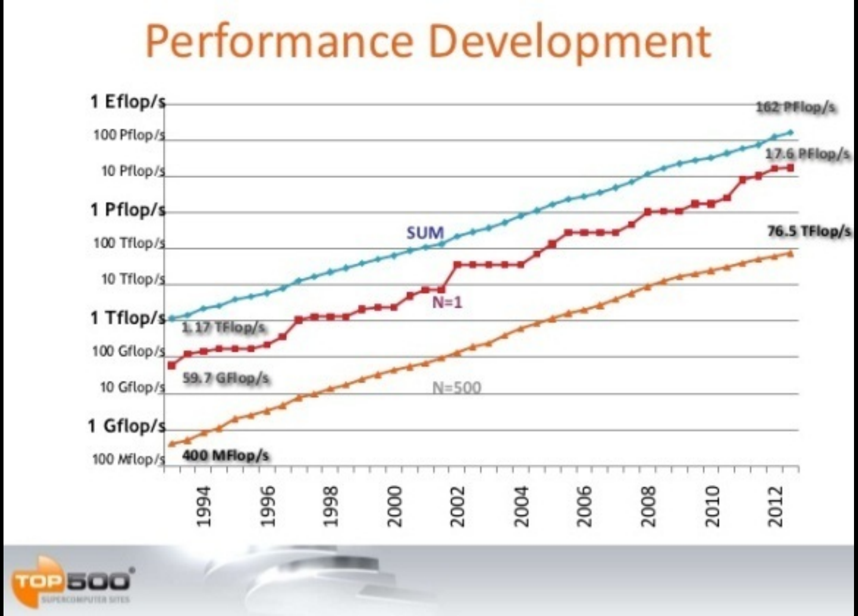

We will soon reach the limits of what is currently feasible and controllable.

https://www.knaw.nl/shared/resources/actueel/publicaties/pdf/20111029.pdf

The GLIF – LightPaths around the World

F Dijkstra, J van der Ham, P Grosso, C de Laat, “A path finding implementation for multi-layer networks”, Future Generation Computer Systems 25 (2), 142-146.

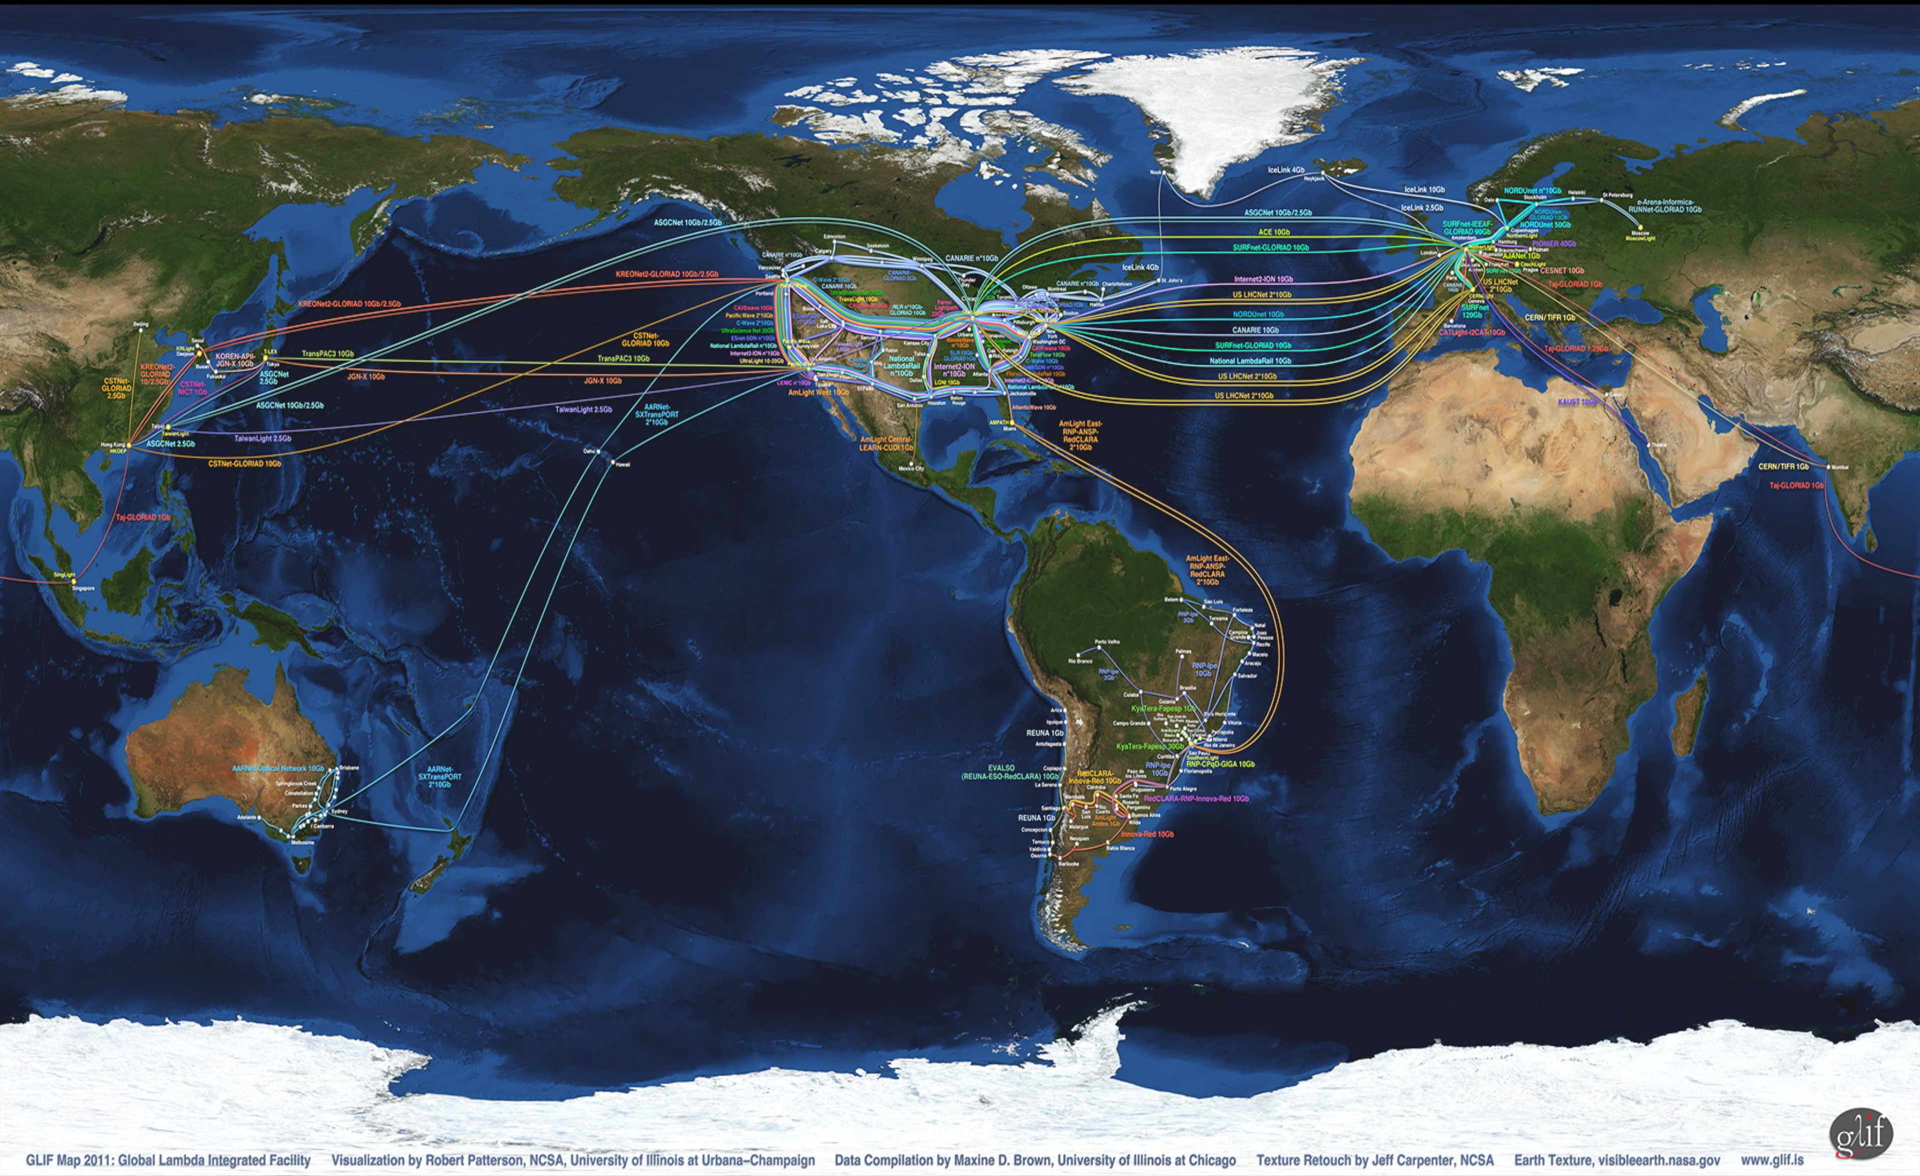

Amsterdam is a major hub in The GLIF

F Dijkstra, J van der Ham, P Grosso, C de Laat, "A path finding implementation for multi-layer networks", Future Generation Computer Systems 25 (2), 142-146.

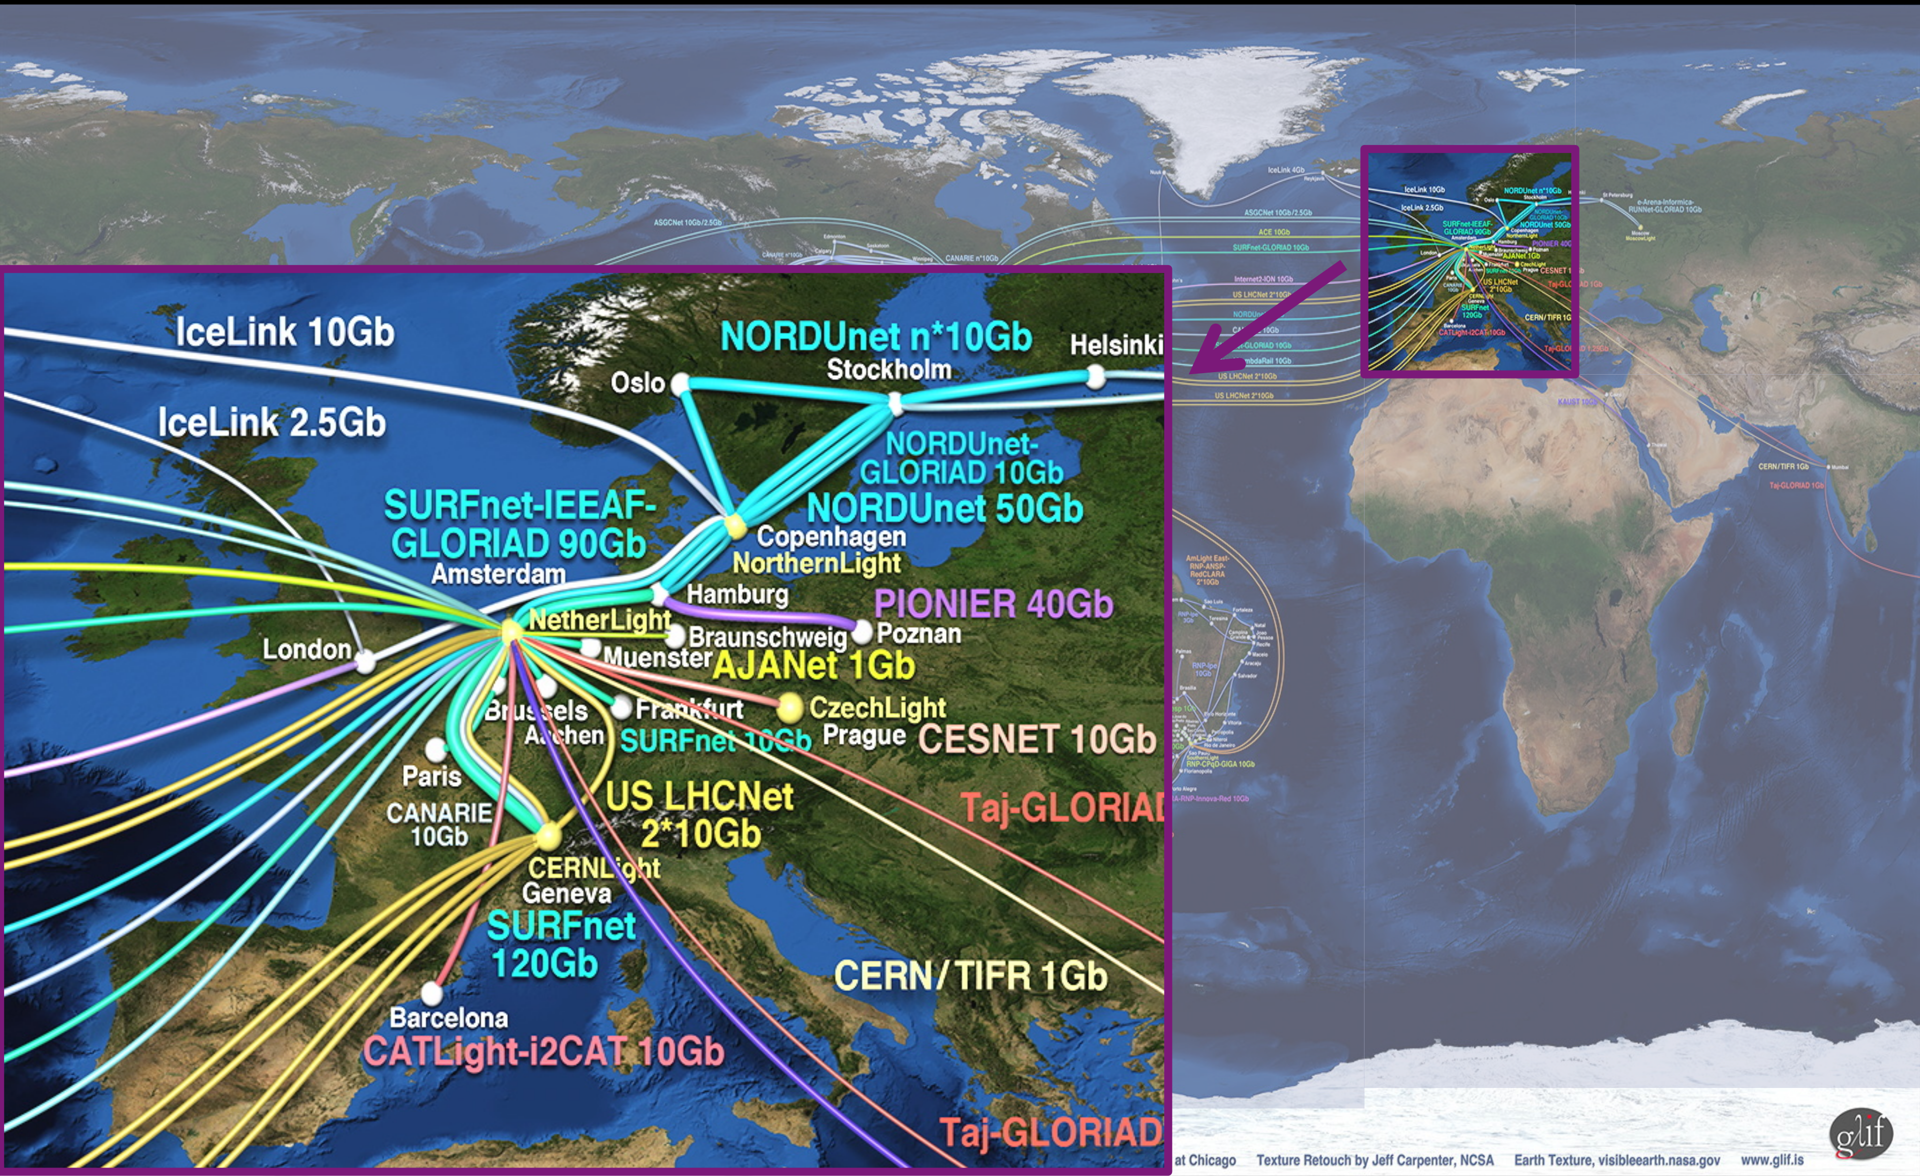

ExoGeni @ OpenLab - UvA

Installed and up June 3th 2013

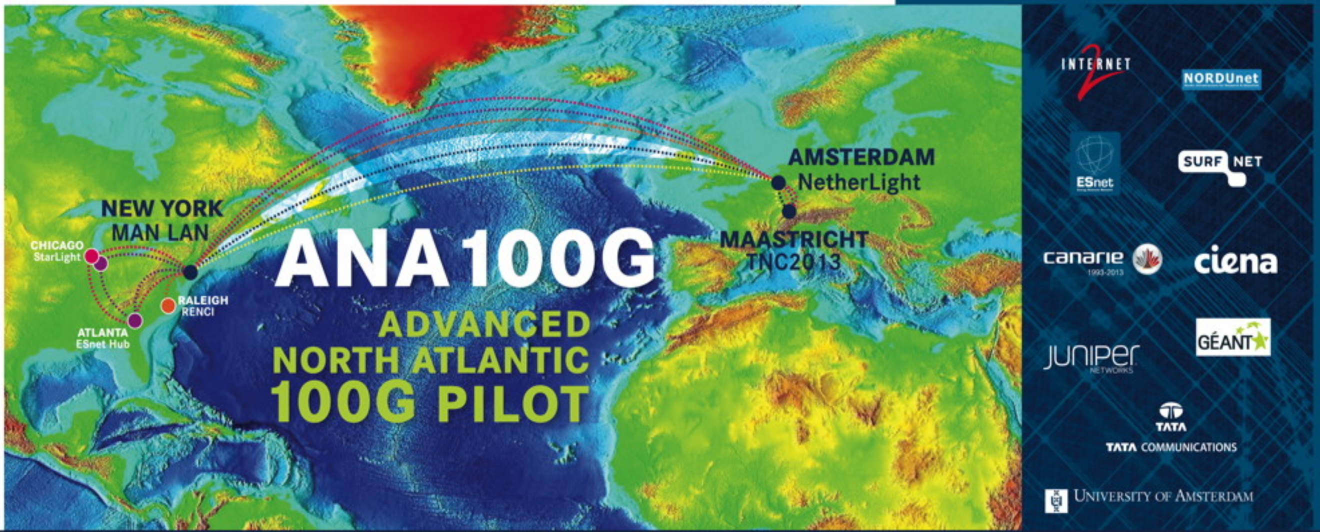

TNC2013 DEMOS JUNE, 2013

DEMO	TITLE	OWNER	AFFILIATION	E-MAIL	A-SIDE	Z-SIDE	PORTS(S) MAN LAN	PORTS(S) TNC2013	DETAILS
1	Big data transfers with multipathing, OpenFlow and MPTCP	Ronald van der Pol	SURFnet	ronald.vanderpol@surfnet.nl	TNC/MECC, Maastricht NL	Chicago, IL	Existing 100G link between internet2 and ESnet	2x40GE (Juniper)+ 2x10GE (OME6500)	In this demonstration we show how multipathing, OpenFlow and Multipath TCP (MPTCP) can help in large file transfers between data centres (Maastricht and Chicago). An OpenFlow application provisions multiple paths between the servers and MPTCP will be used on the servers to simultaneously send traffic across all those paths. This demo uses 2x40G on the transatlantic 100G link. ESnet provides 2x40G between MAN LAN and StarLight, ACE and USLHCnet provide additional 10GEs.
2	Visualize 100G traffic	Inder Monga	ESnet	imonga@es.net					Using an SNMP feed from the Juniper switch at TNC2013,and/or Brocade AL2S node in MANLAN, this demo would visualize the total traffic on the link, of all demos aggregated. The network diagram will show the transatlantic topology and some of the demo topologies.
3	How many modern servers can fill a 100Gbps Transatlantic Circuit?	Inder Monga	ESnet	imonga@es.net	Chicago, Ill	TNC showfloor	1x 100GE	8x 10GE	In this demonstration, we show that with the proper tuning and tool, only 2 hosts on each continent can generate almost 80Gbps of traffic. Each server has 4 10G NICS connected to a 40G virtual circuit, and has iperf3 running to generate traffic. ESnet's new "iperf3" throughput measurement tool, still in 'beta', combines the best features from other tools such as iperf, nuttcp, and netperf. See: https://my.es.net/demos/tnc2013/
4	First European ExoGENI at Work	Jeroen van der Ham	UvA	vdham@uva.nl	RENCI, NC	UvA, Amsterdam, NL	1x 10GE	1x 10GE	The ExoGENI racks at RENCI and UvA will be interconnected over a 10G pipe and be on continuously, showing GENI connectivity between Amsterdam and the rest of the GENI nodes in the USA.
5	Up and down North Atlantic @ 100G	Michael Enrico	DANTE	michael.enrico@dante.net	TNC showfloor	TNC showfloor	1x 100GE	1x 100GE	The DANTE 100GE test set will be placed at the TNC2013 showfloor and connected to the Juniper at 100G. When this demo is running a loop @ MAN LAN's Brocade switch will ensure that the traffic sent to MAN LAN returns to the showfloor. On display is the throughput and RTT (to show the traffic traveled the Atlantic twice)

Connected via the new 100 Gb/s transatlantic To US-GENI

Alien light
From idea to realisation!

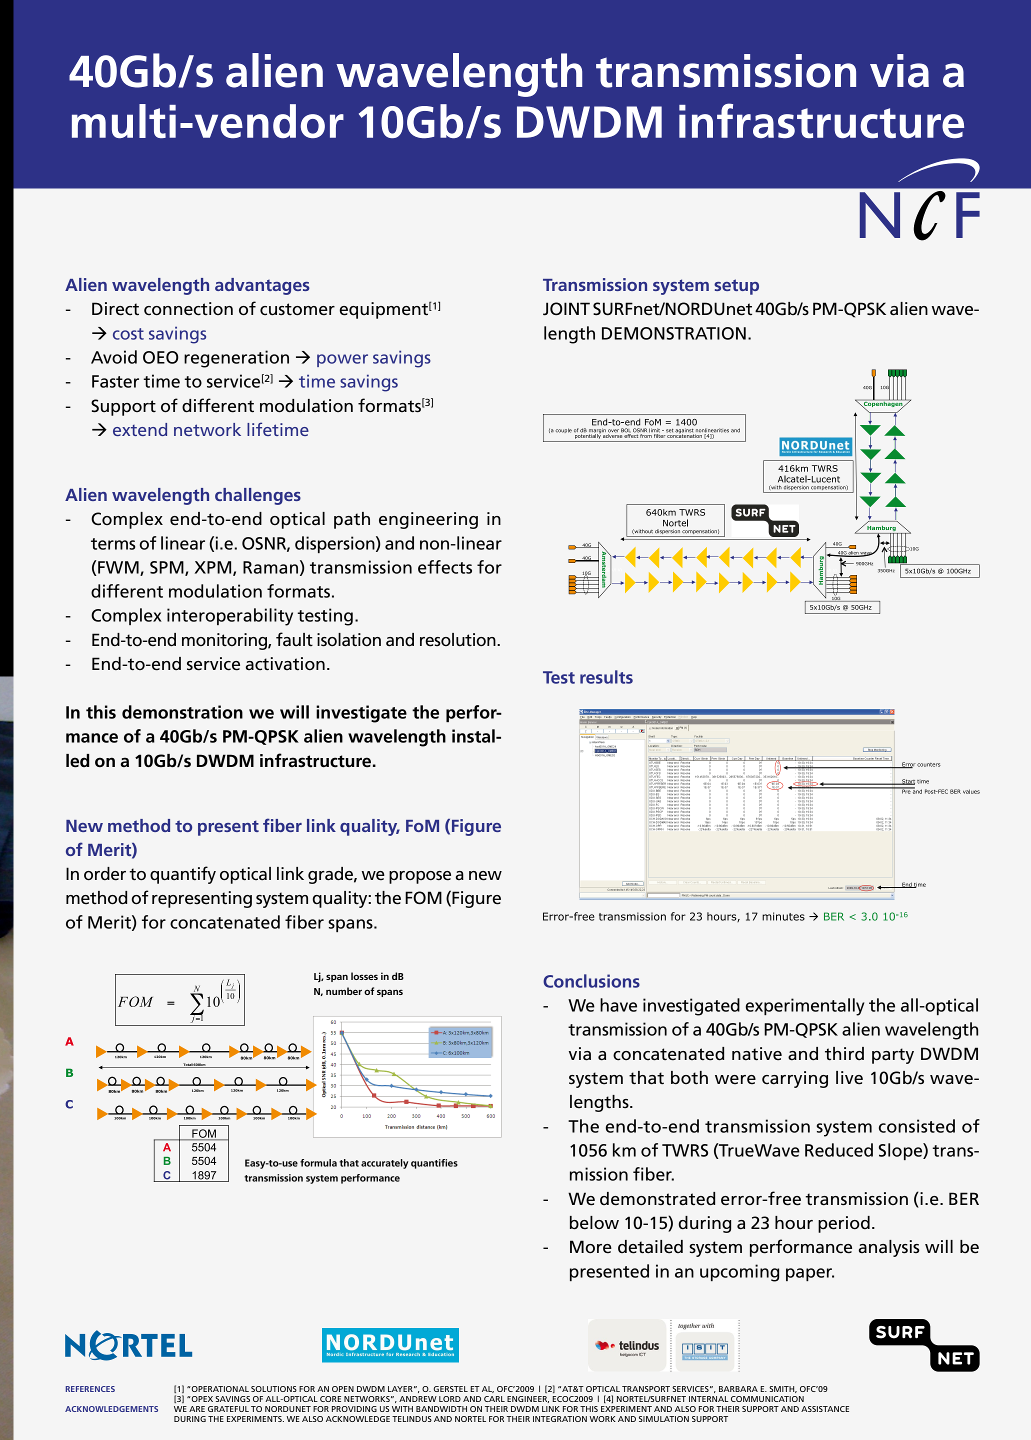

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

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!

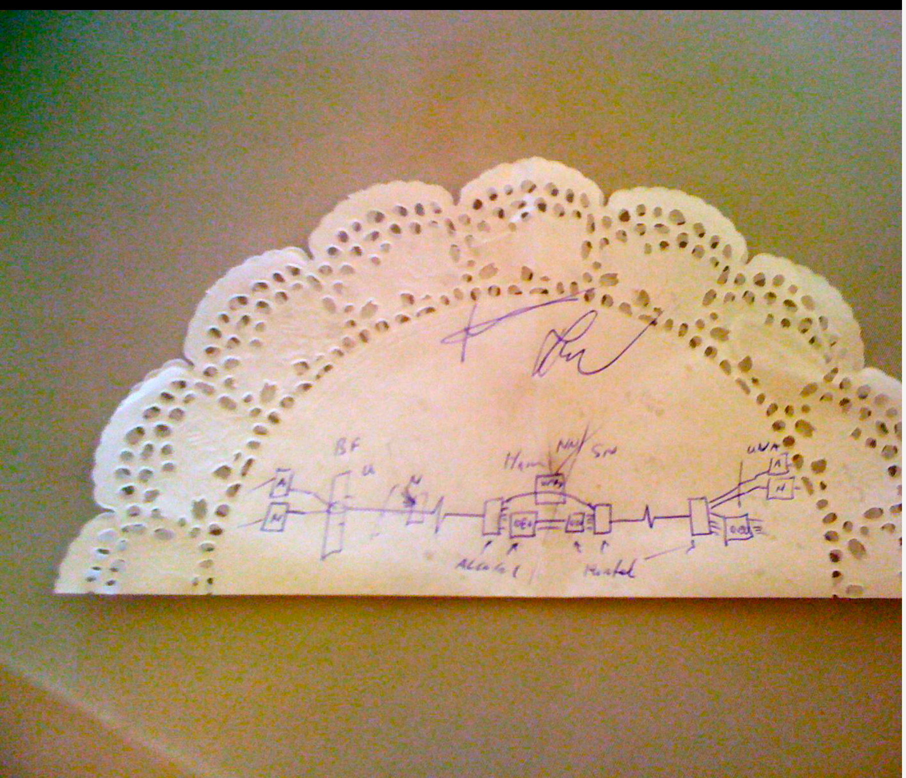

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

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

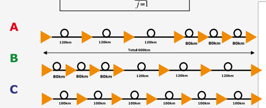

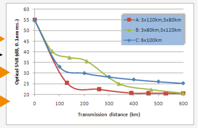

	FOM
A	5504
B	5504
C	1897

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

Transmission system setup

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

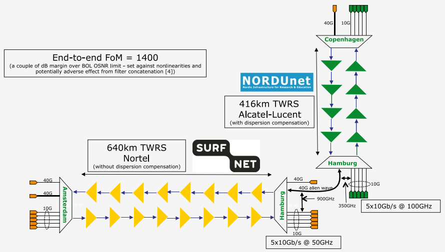

Test results

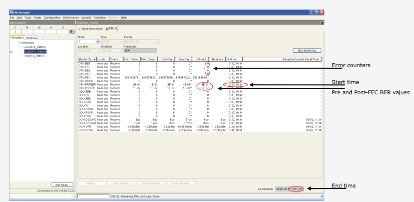

Error-free transmission for 23 hours, 17 minutes → BER < $3.0\ 10^{-16}$

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 Nordic Infrastructure for Research & Education telindus belgacom ICT together with ISIT THE STORAGE COMPANY 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

ClearStream @ TNC2011

Setup codename: FlightCees

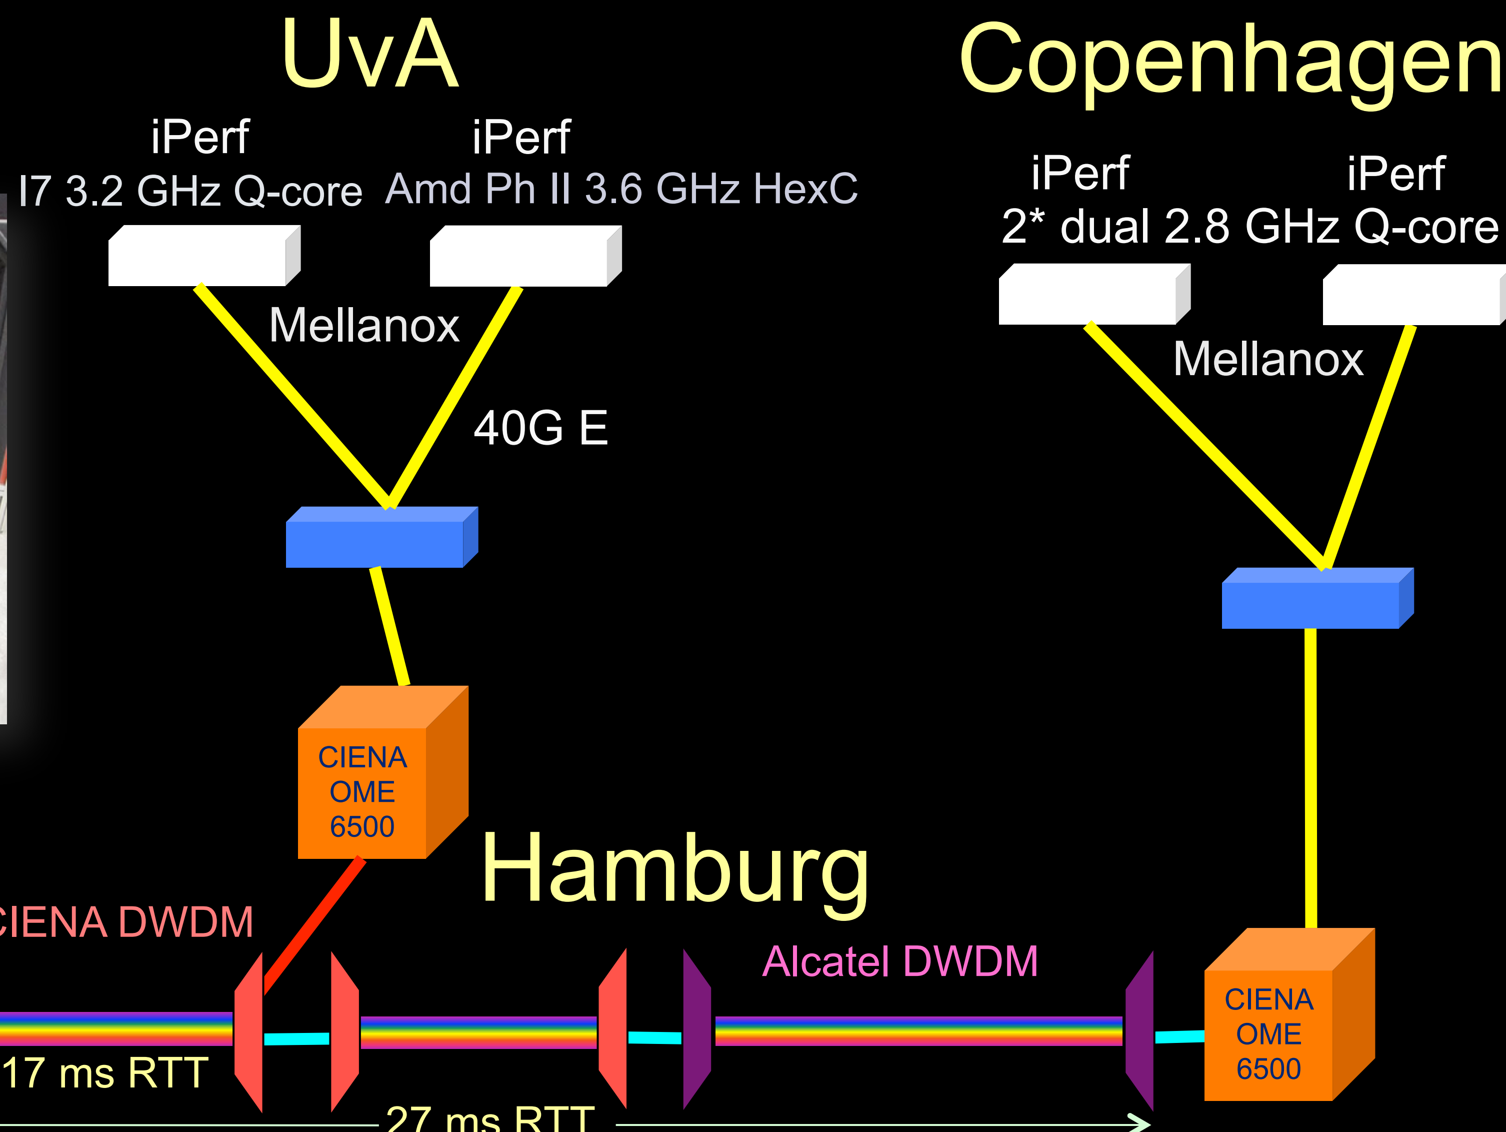

Amsterdam – Geneva (CERN) – Copenhagen – 4400 km (2700 km alien light)

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

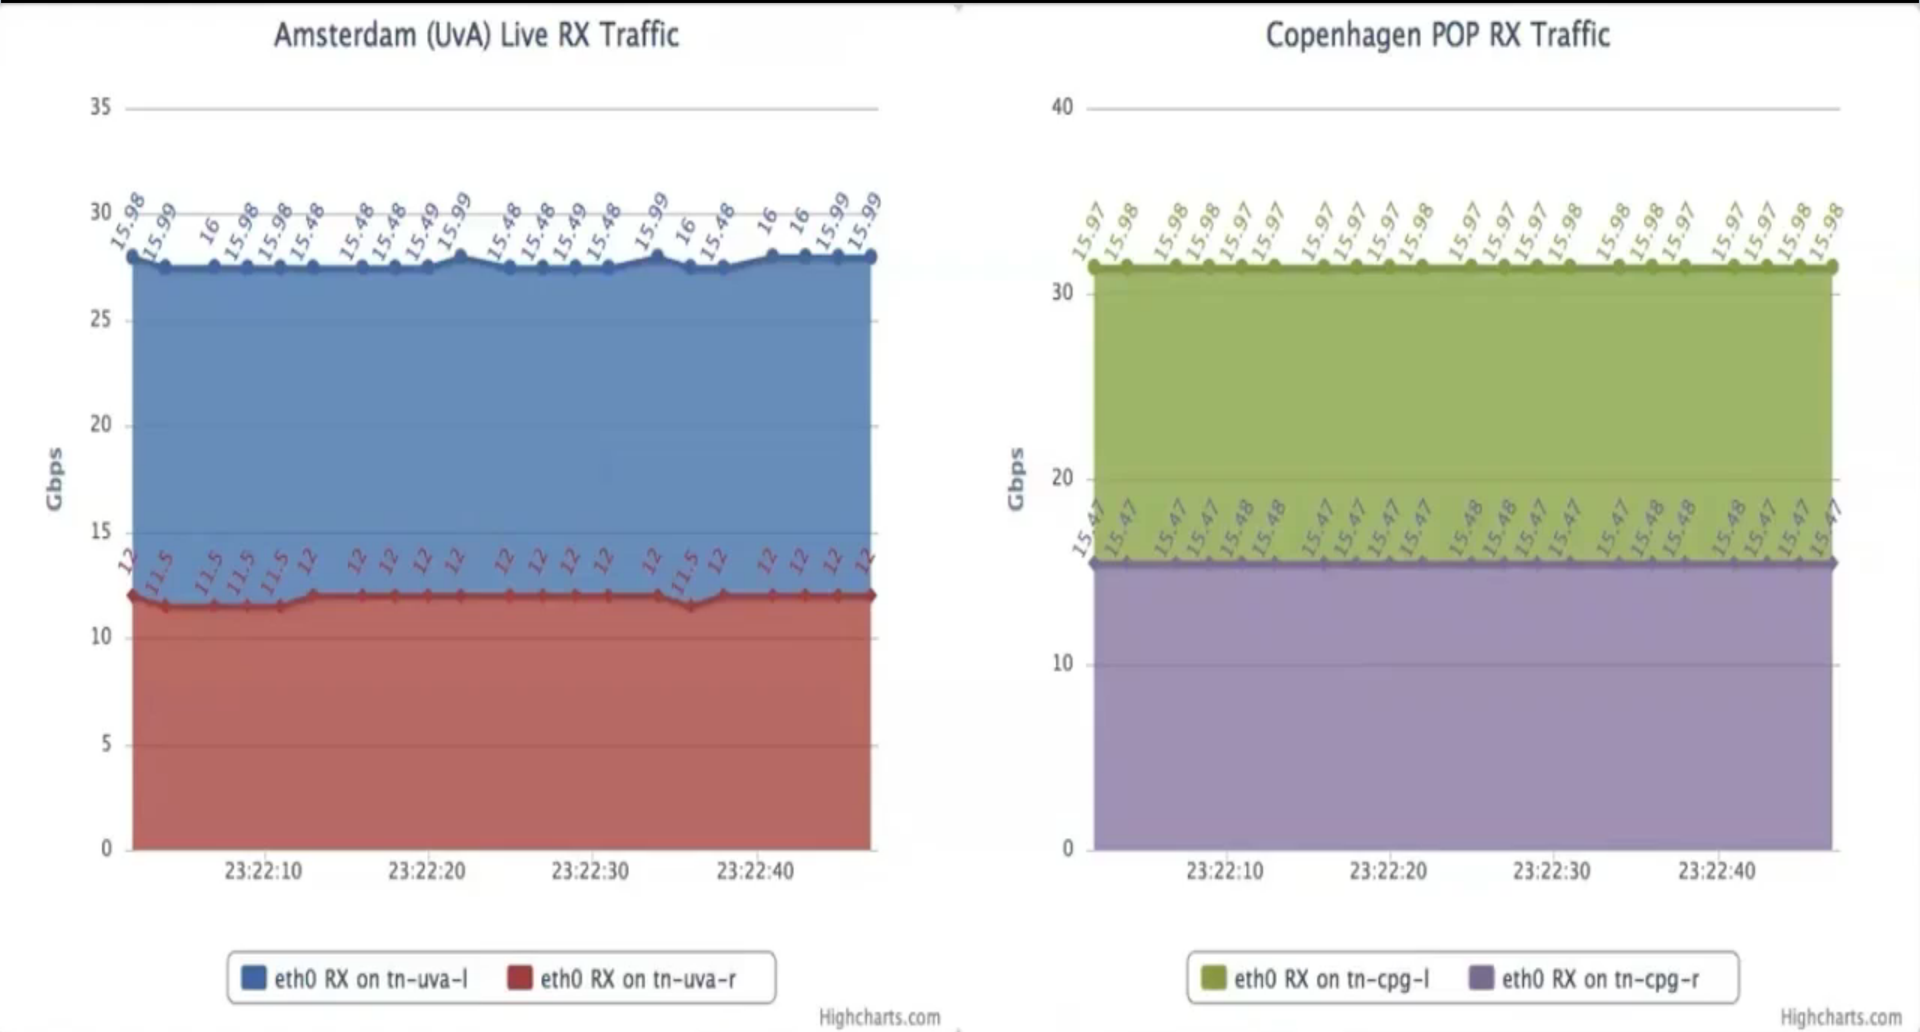

27.99 Gbps to Amsterdam <-> 31.45 Gbps to Copenhagen

Total Throughput 59.44 Gbps RTT 44.010 ms

Server Architecture

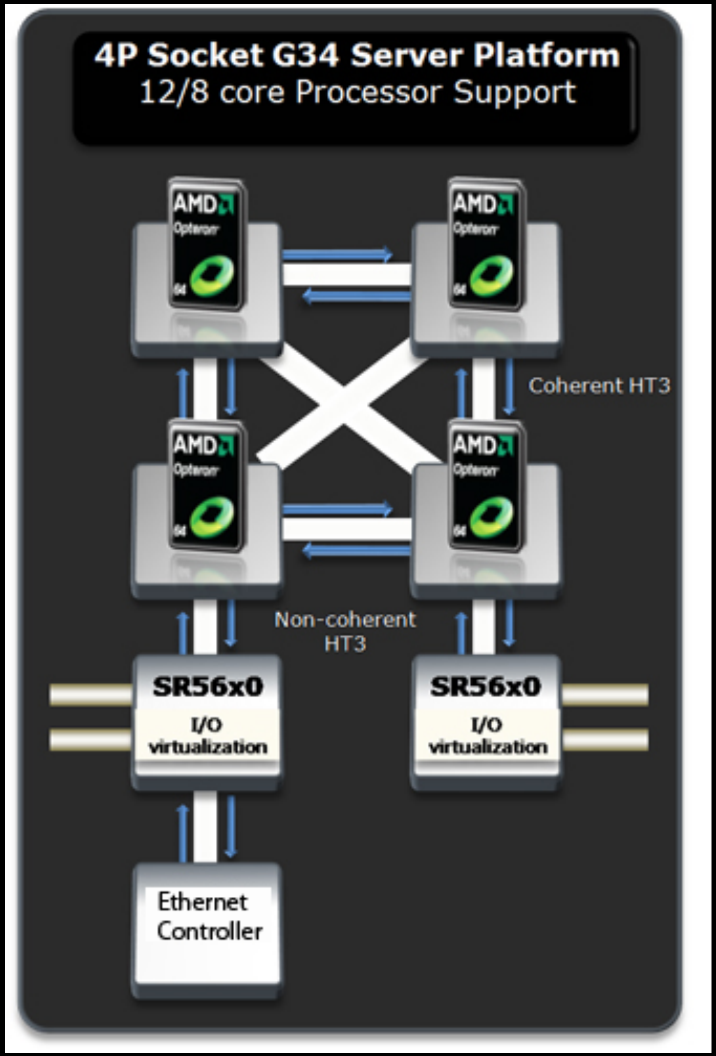

DELL R815
4 x AMD Opteron 6100

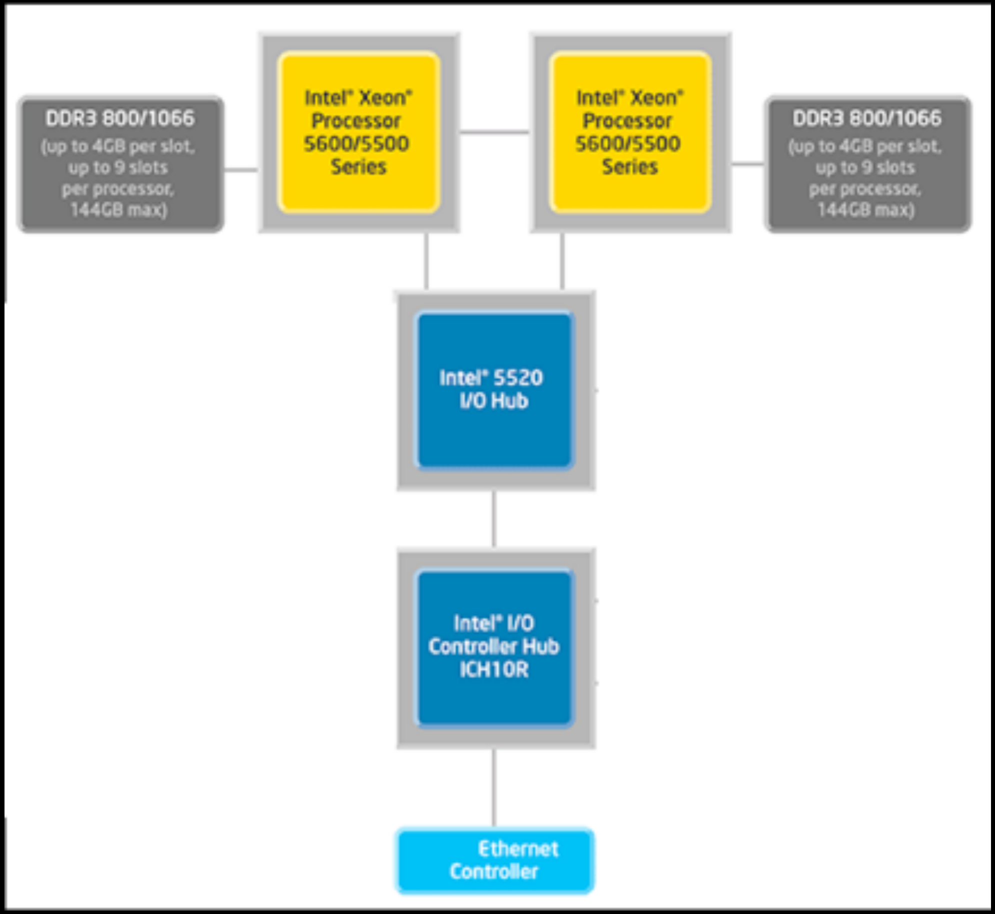

Supermicro X8DTT-HIBQF
2 x Intel Xeon

CPU Topology benchmark

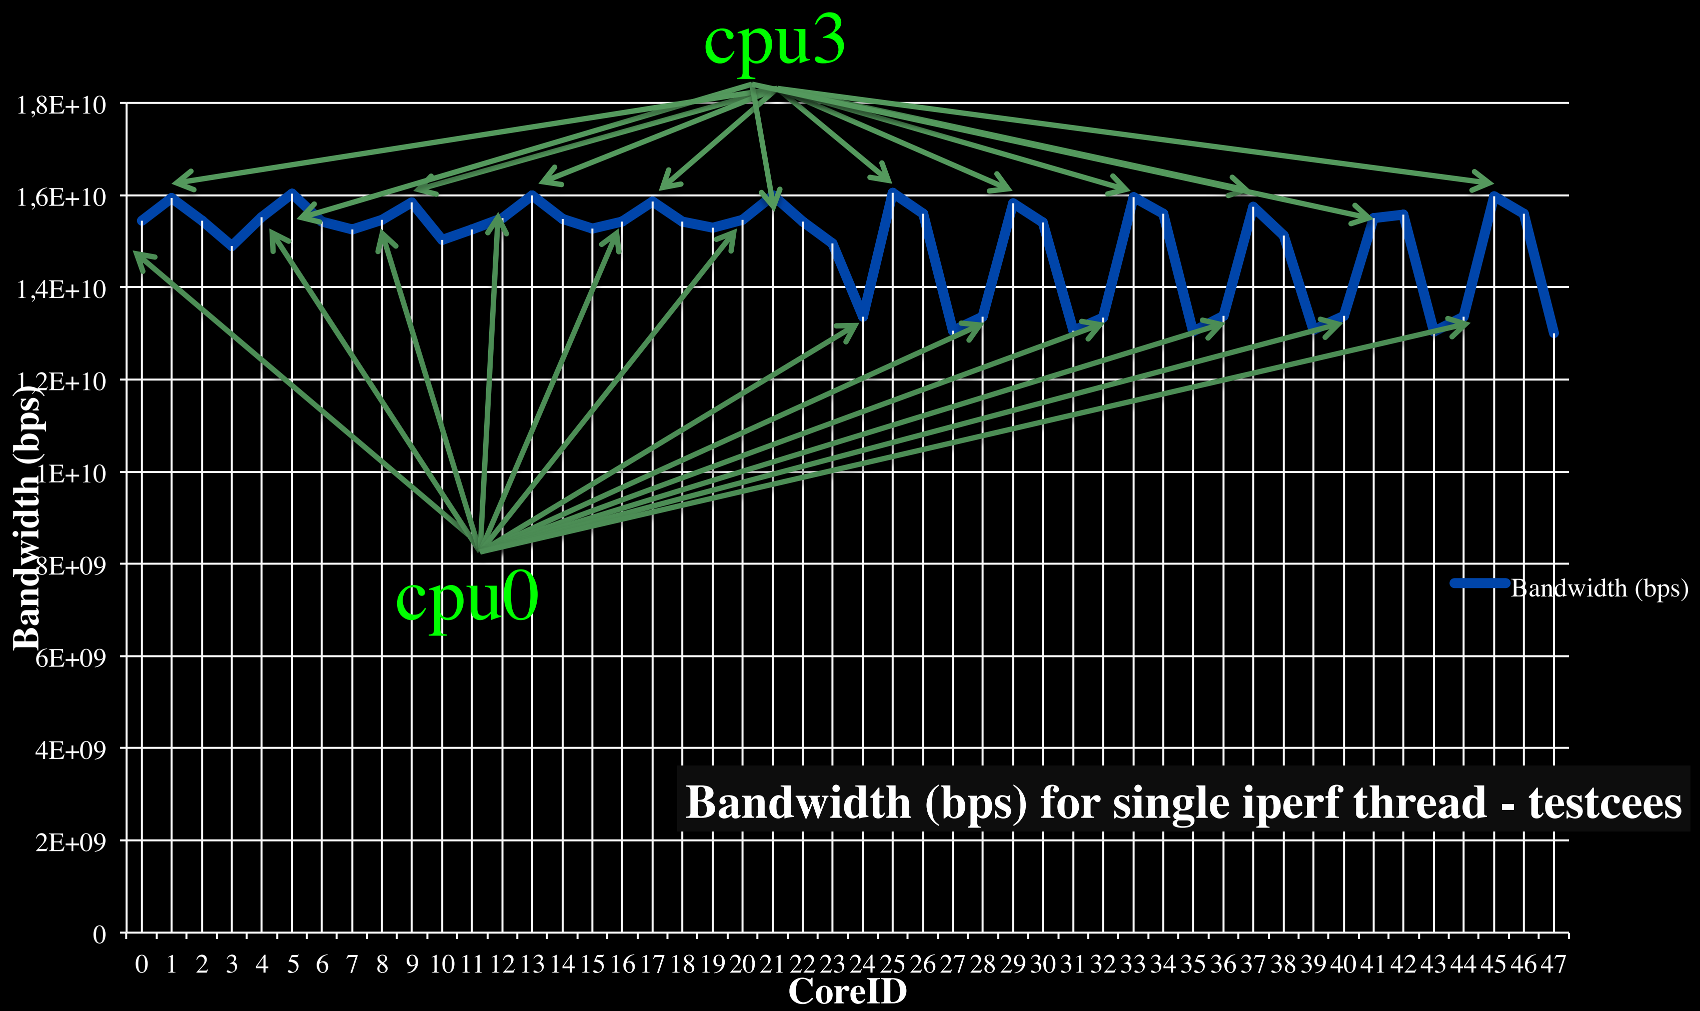

We used numactl to bind iperf to cores

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- *Capacity*
  - *Bandwidth on demand, QoS, architectures, photonics, performance*
- ***Capability***
  - ***Programmability, virtualization, complexity, semantics, workflows***
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- *Sustainability*
  - *Greening infrastructure, awareness*
- *Resilience*
  - *Systems under attack, failures, disasters*

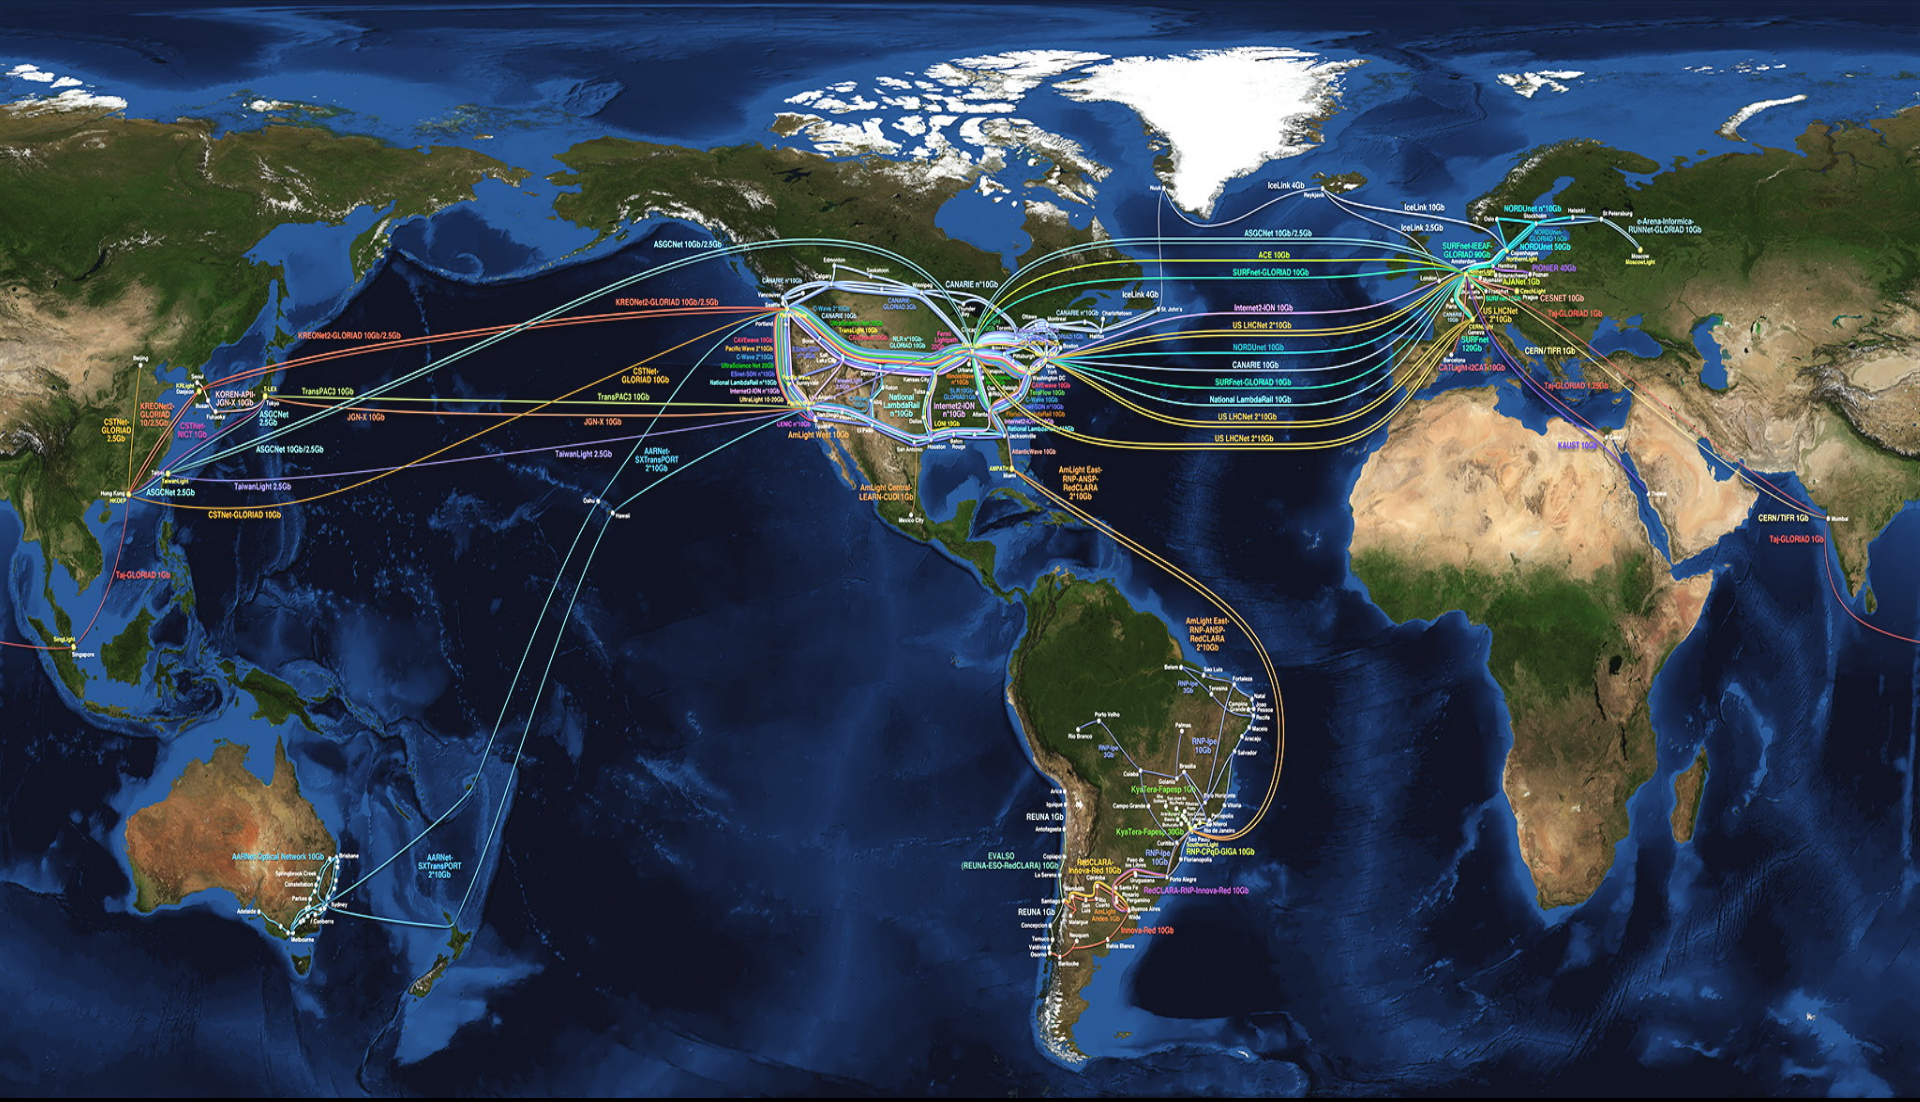

We investigate: TomTom for complex networks!

California Institute for Telecommunications and Information Technology (Calit2) Wide Area Network

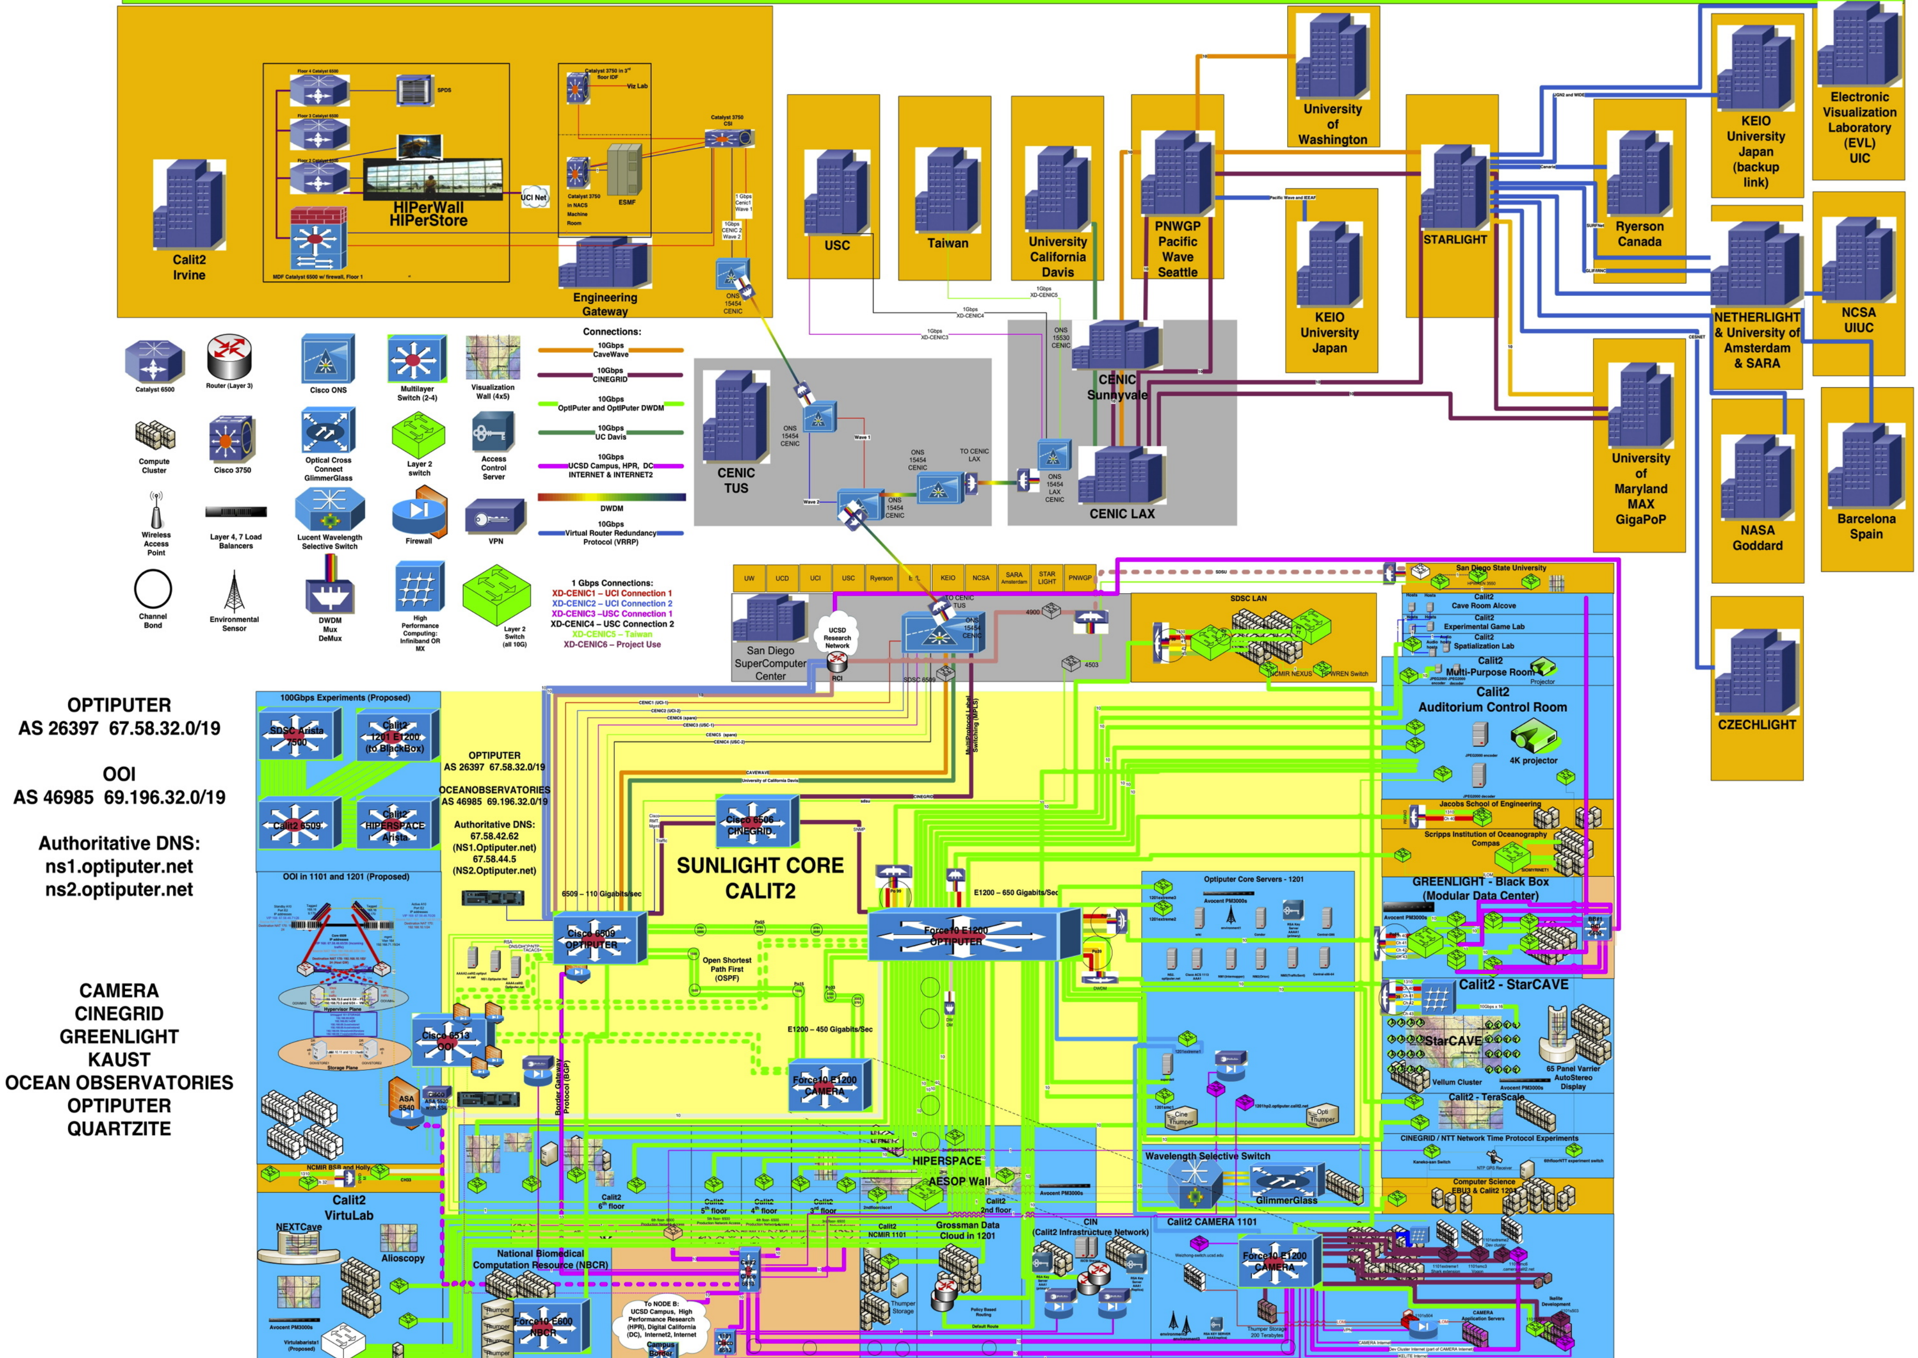

LinkedIN for Infrastructure

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

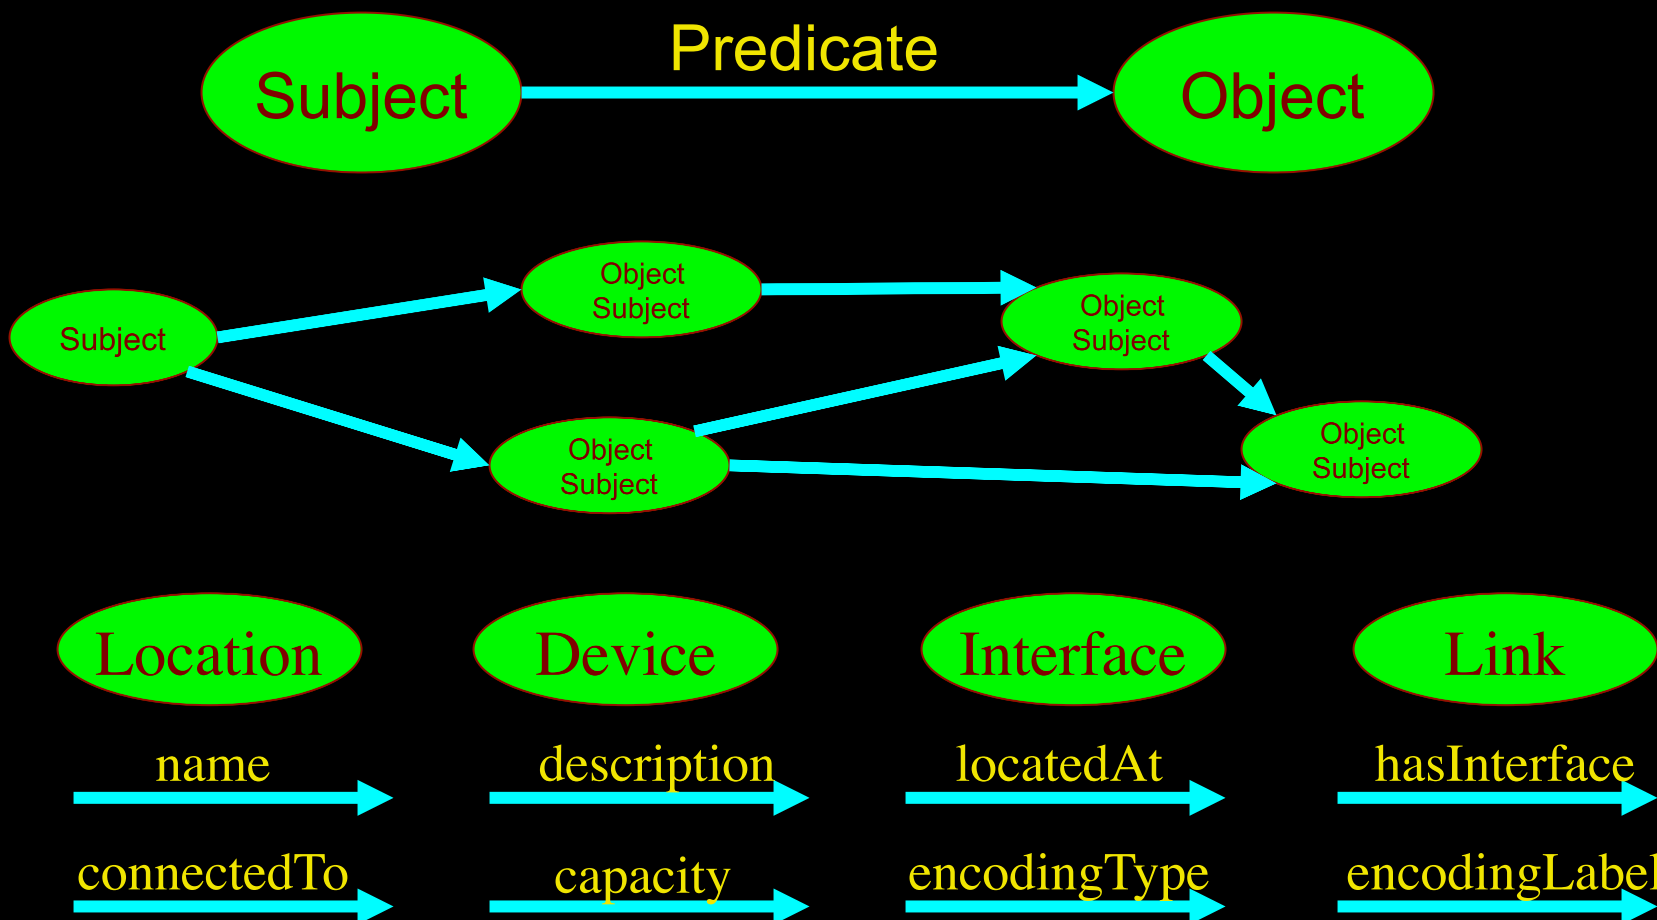

Network Description Language

Choice of RDF instead of XML syntax

Grounded modeling based on G0805 description:

Article: F. Dijkstra, B. Andree, K. Koymans, J. van der Ham, P. Grosso, C. de Laat, *"A Multi-Layer Network Model Based on ITU-T G.805"*

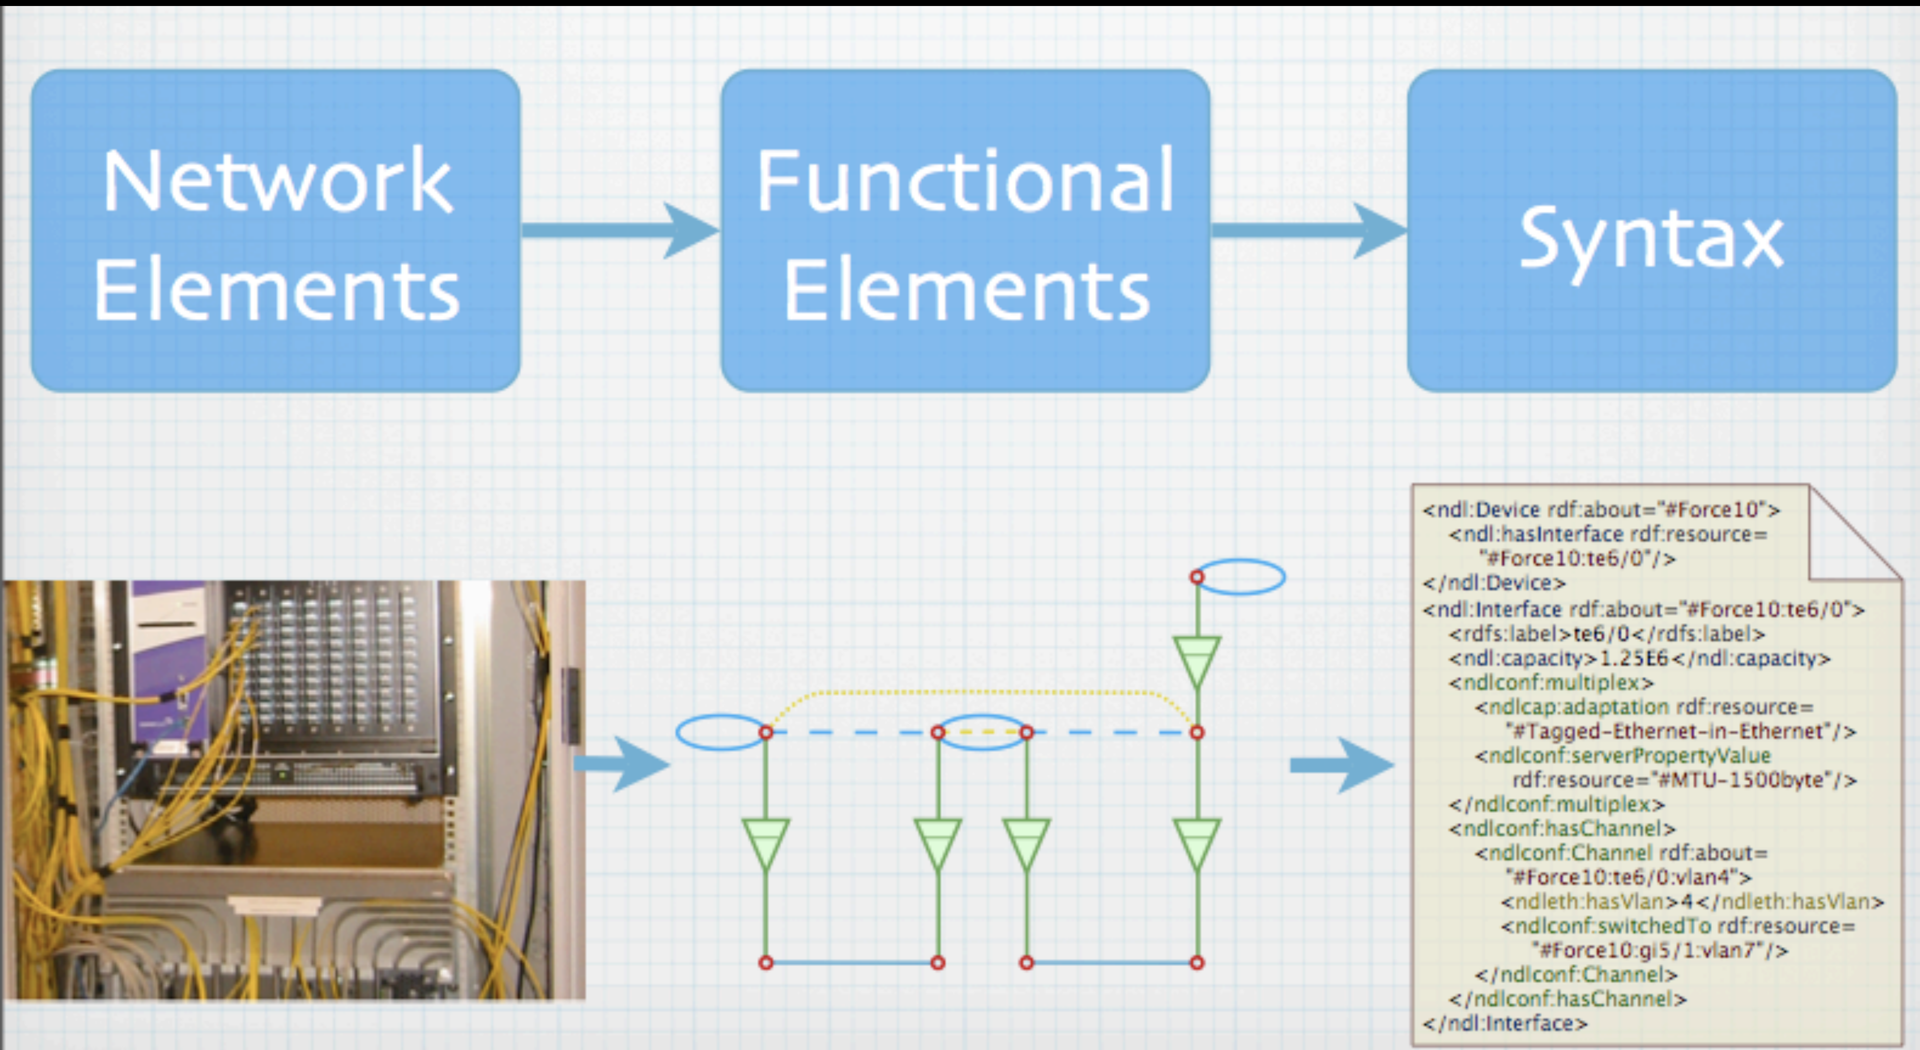

Multi-layer descriptions in NDL

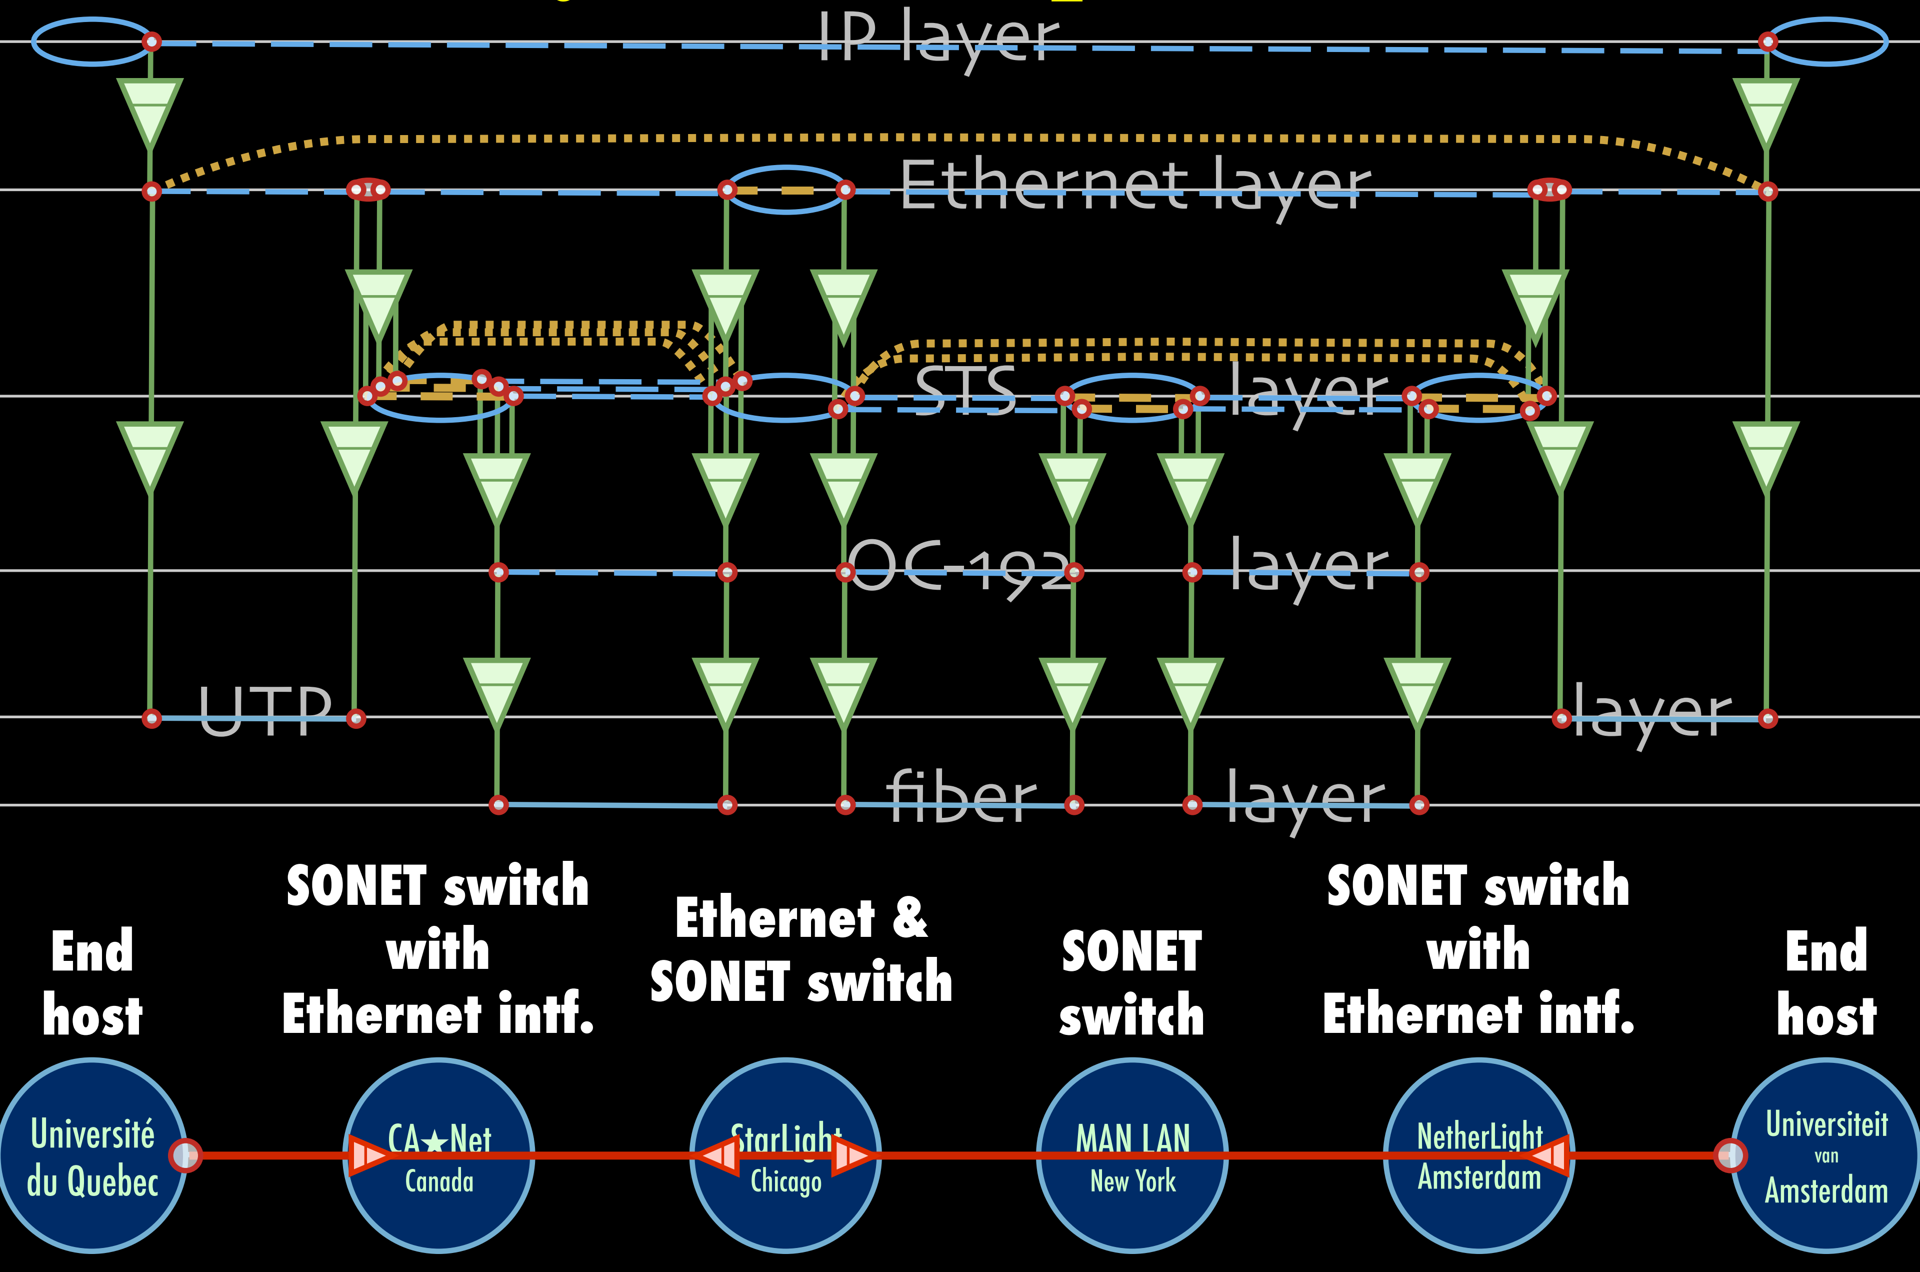

NML
OGF spec
iNDL

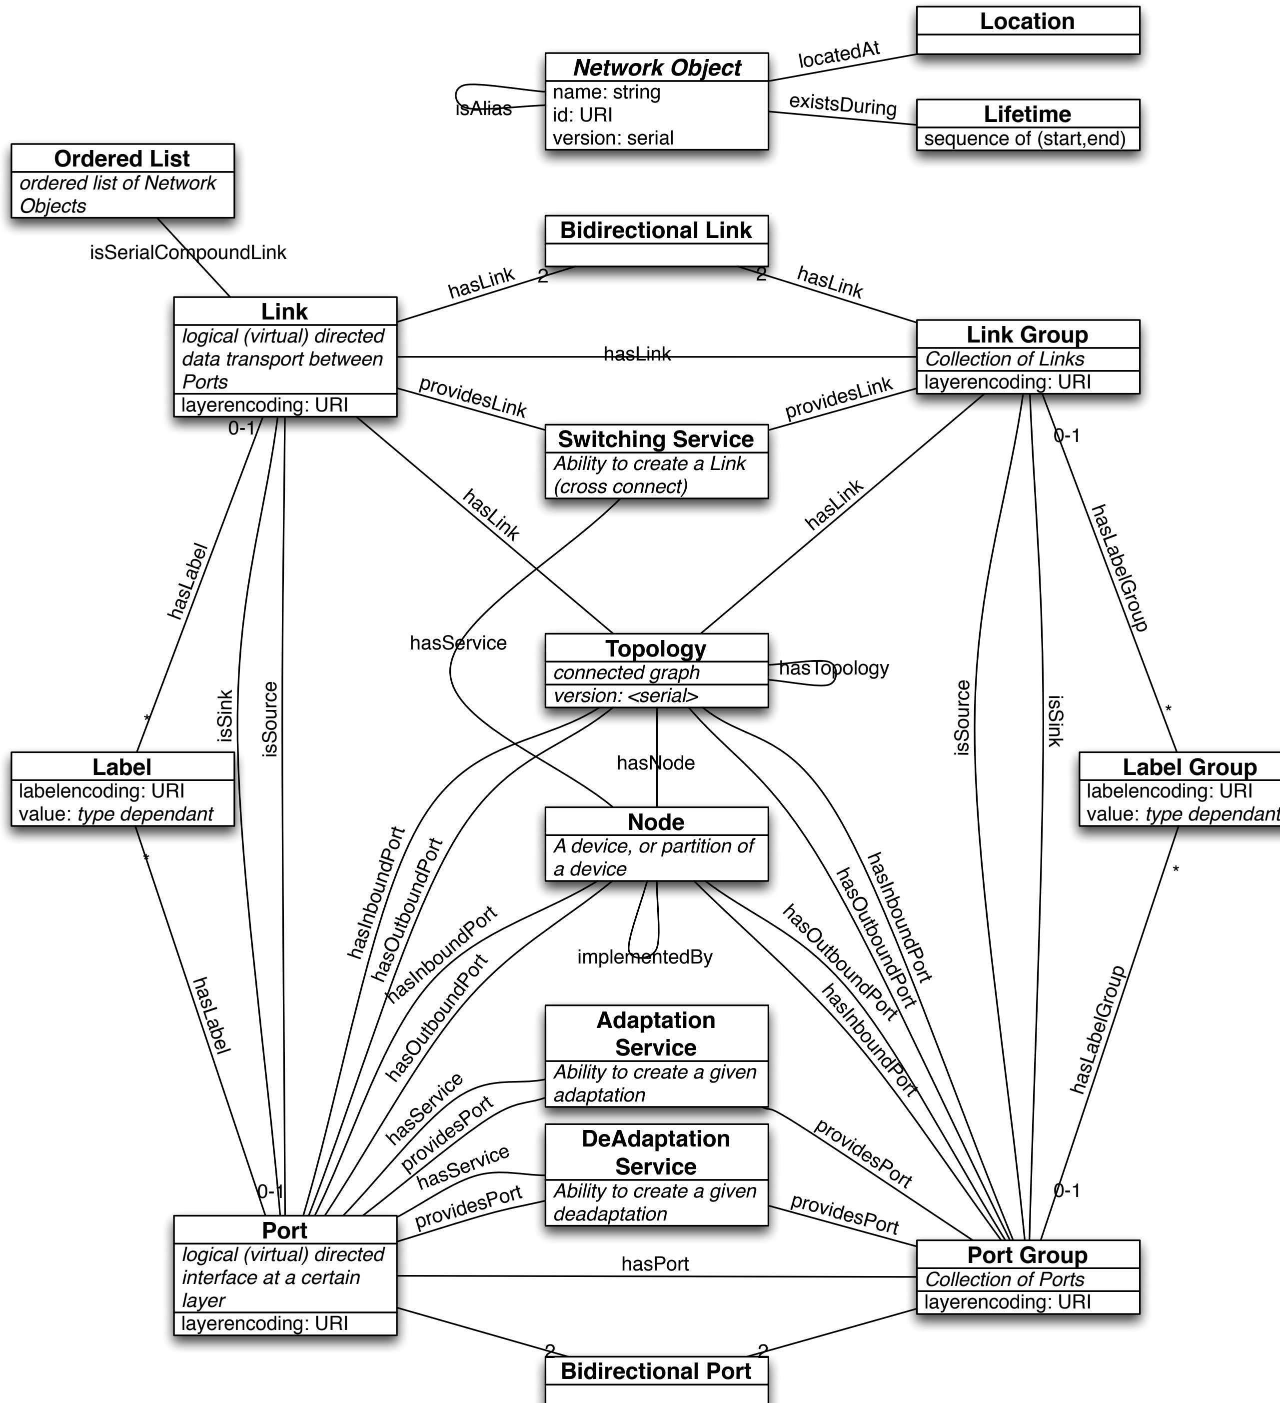

Network Topology Description

Network topology research supporting automatic network provisioning

- Inter-domain networks
- Multiple technologies
- Based on incomplete information
- Possibly linked to other resources

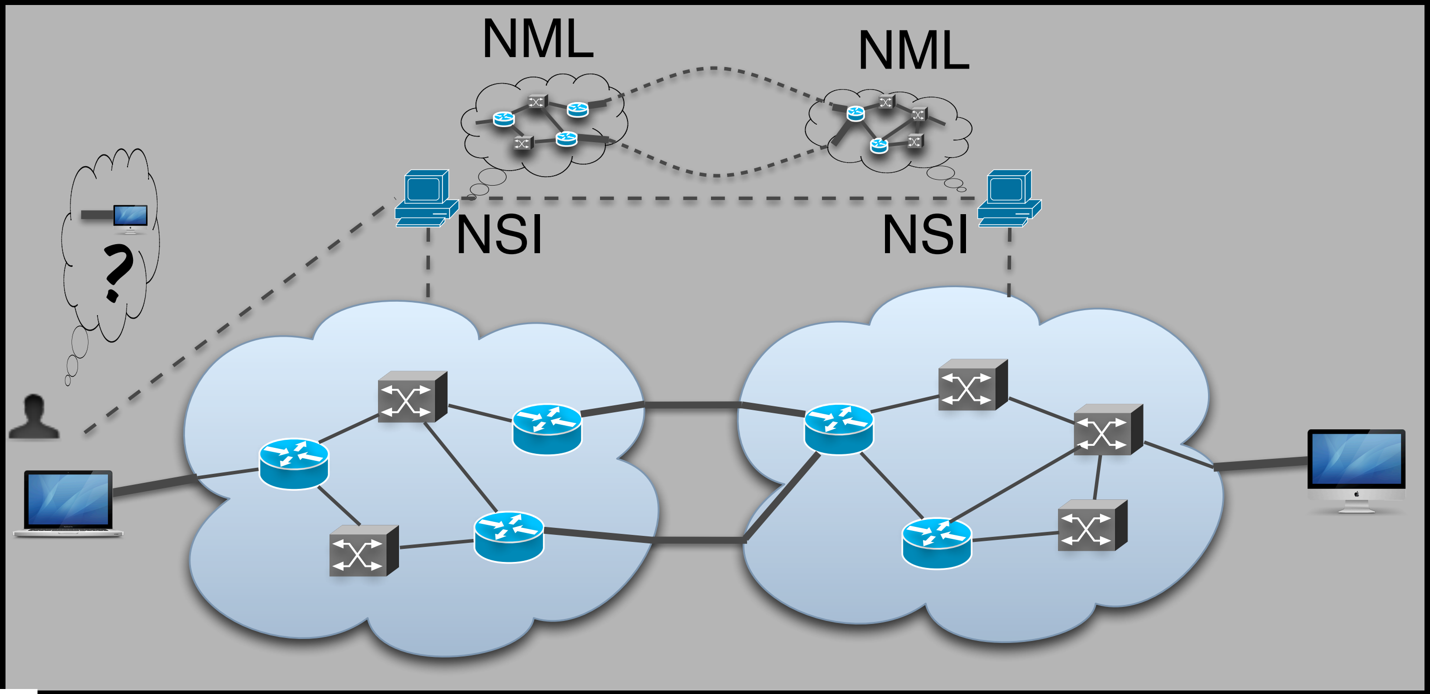

http://redmine.ogf.org/projects/nml-wg
http://redmine.ogf.org/projects/nsi-wg

http://sne.science.uva.nl/ndl

GLIF 2013 in NML

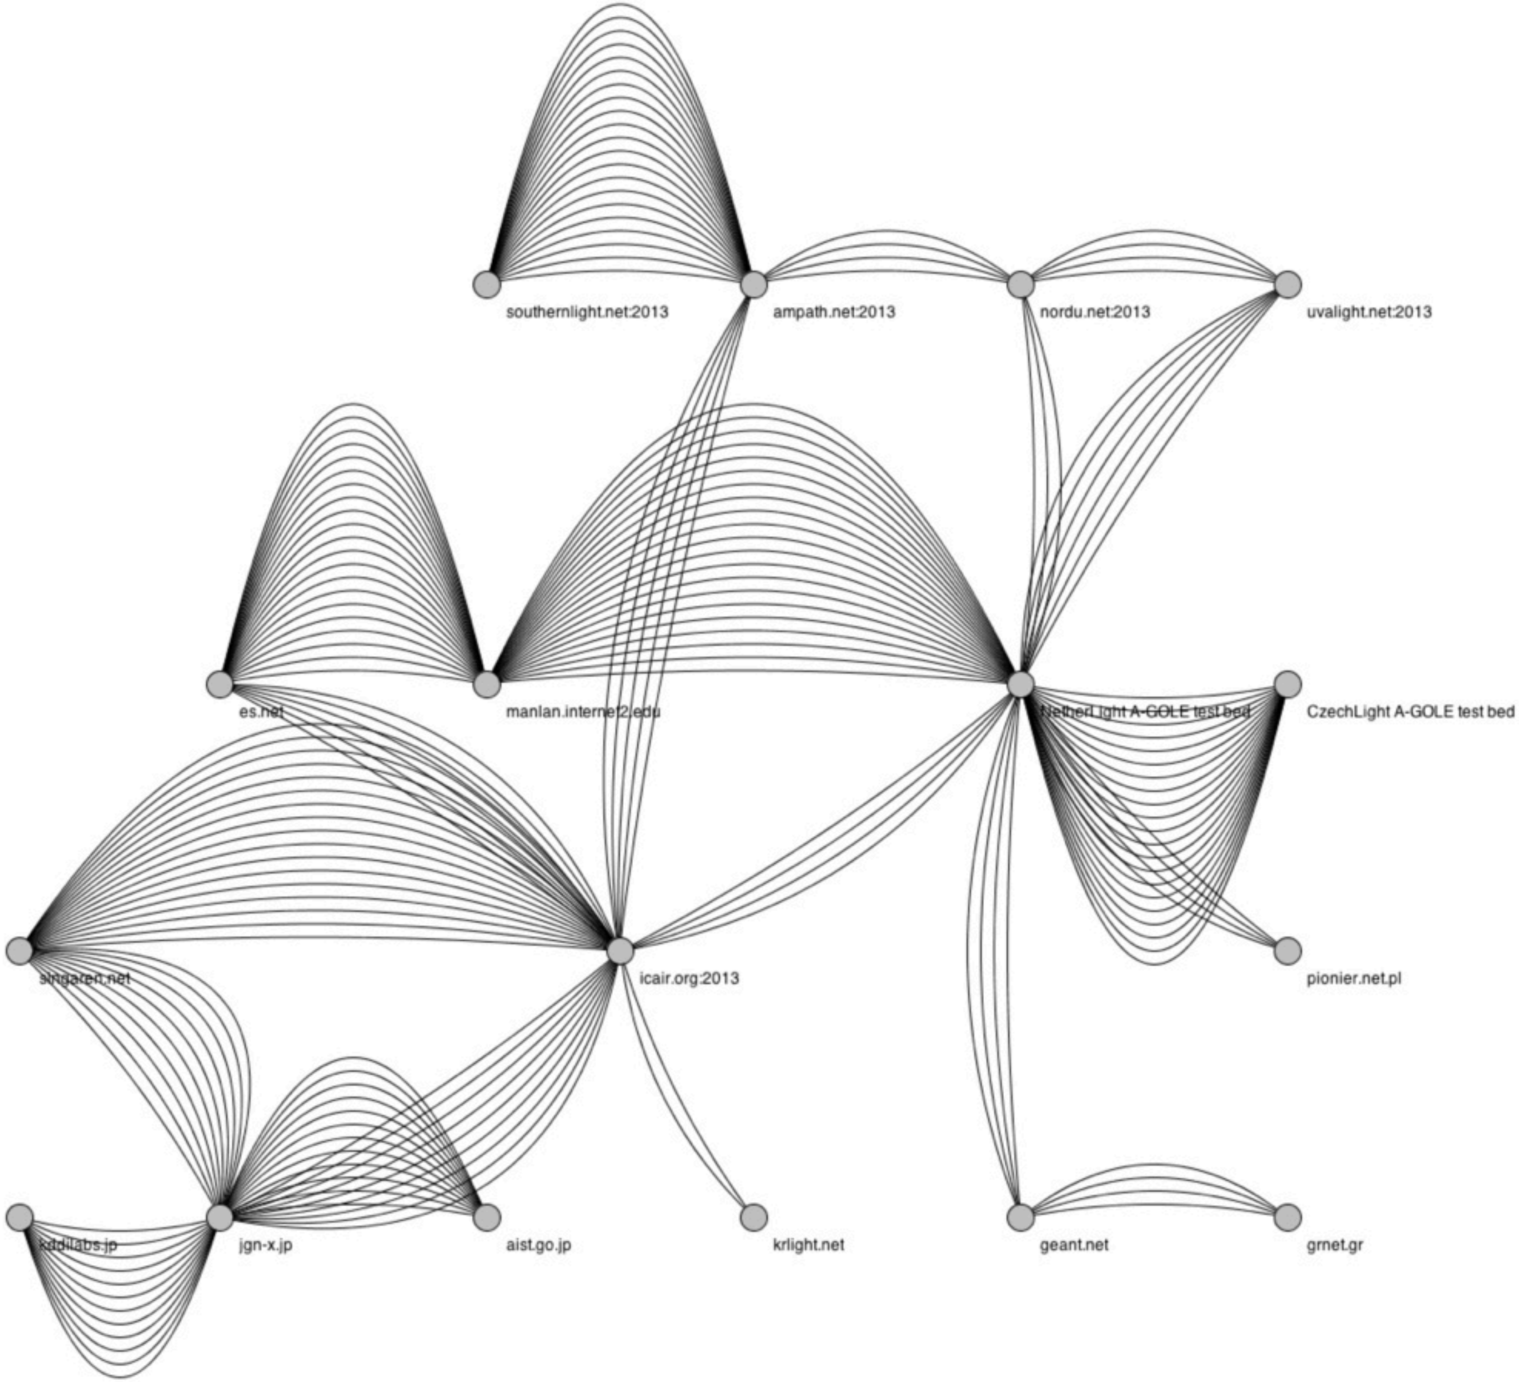

CdL

Applications and Networks become aware of each other!

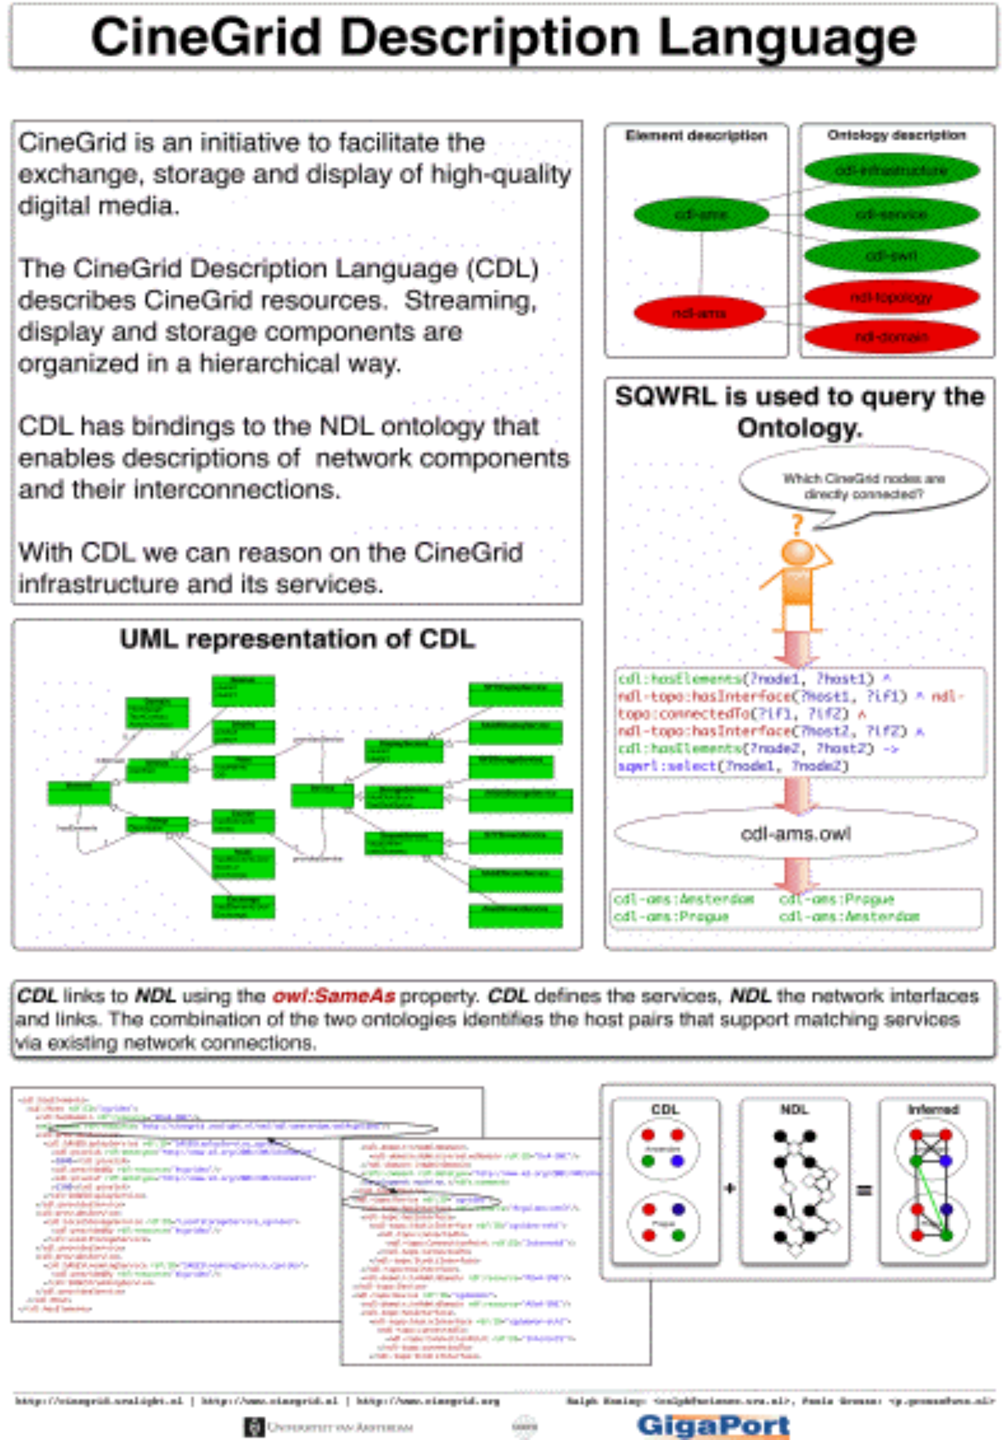

Information Modeling

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

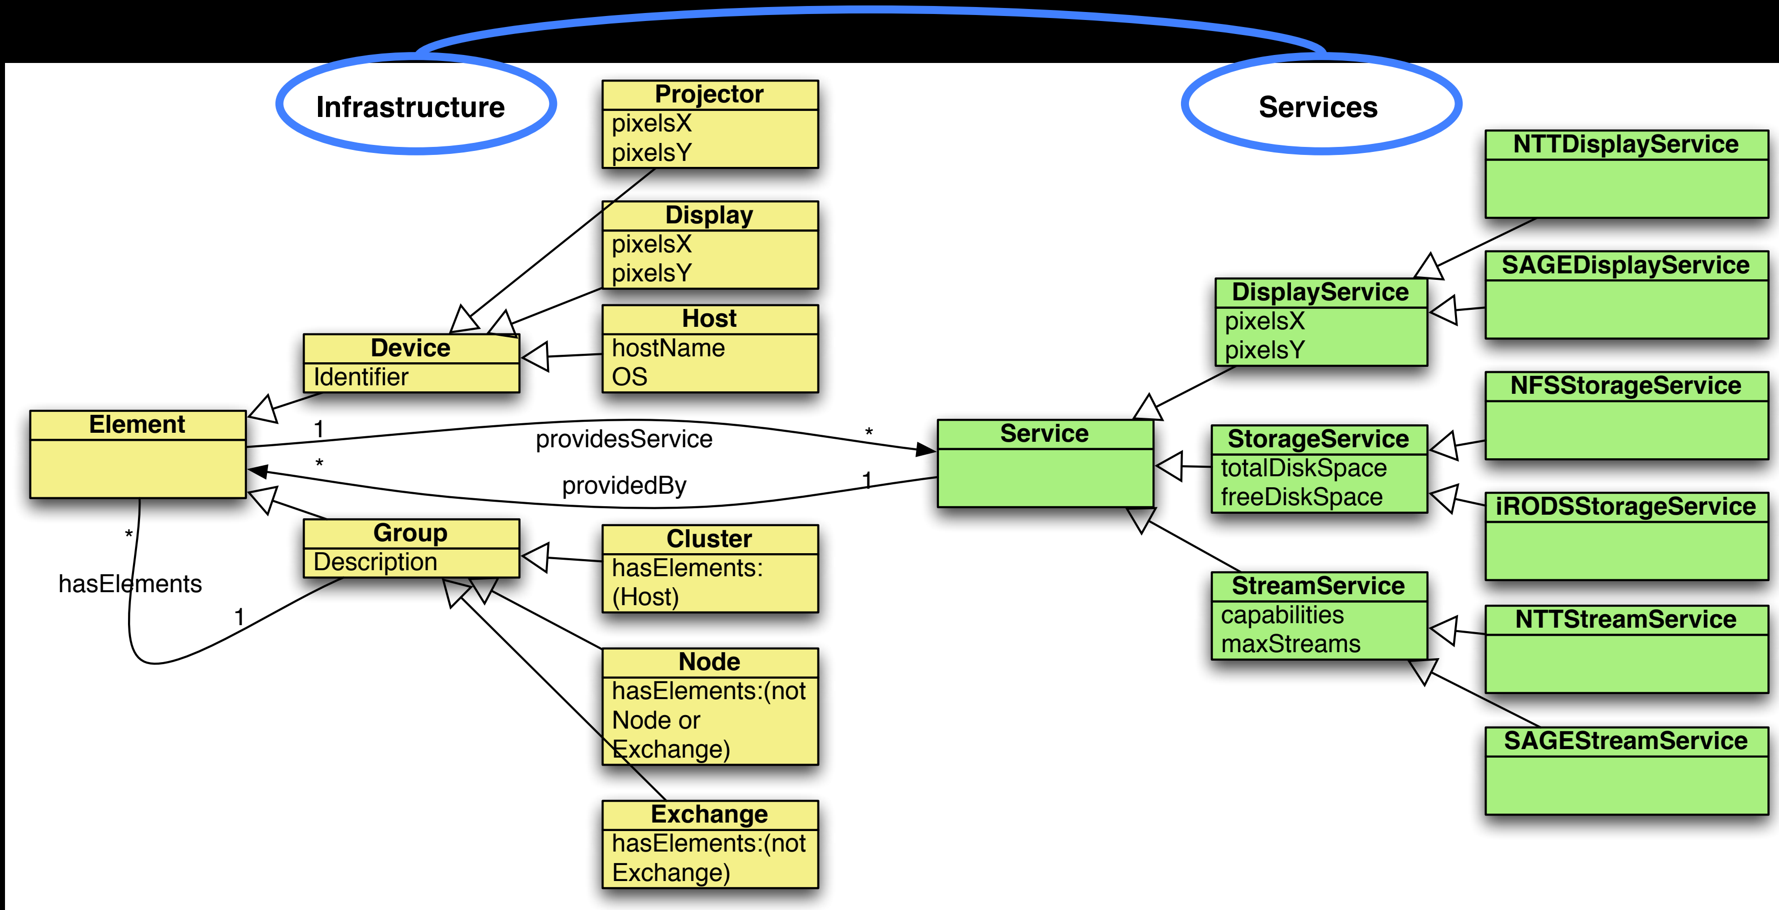

J. van der Ham, F. Dijkstra, P. Grosso, R. van der Pol, A. Toonk, C. de Laat
A distributed topology information system for optical networks based on the semantic web, Elsevier Journal on Optical Switching and Networking, Volume 5, Issues 2-3, June 2008, Pages 85-93

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

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- *Capacity*
  - *Bandwidth on demand, QoS, architectures, photonics, performance*
- *Capability*
  - *Programmability, virtualization, complexity, semantics, workflows*
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- ***Sustainability***
  - ***Greening infrastructure, awareness***
- *Resilience*
  - *Systems under attack, failures, disasters*

ECO-Scheduling

Green scheduling

Network infrastructures

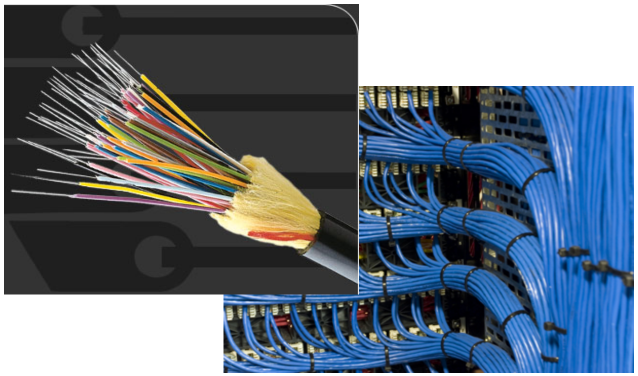

Green energy sources

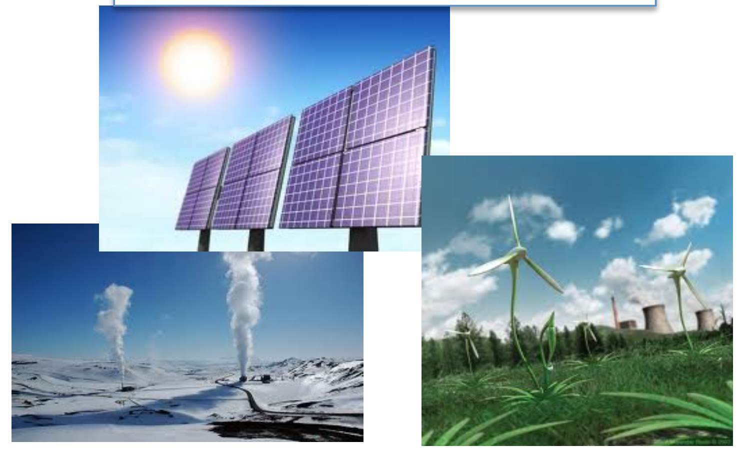

CO_2 footprint;
Energy needed and lost

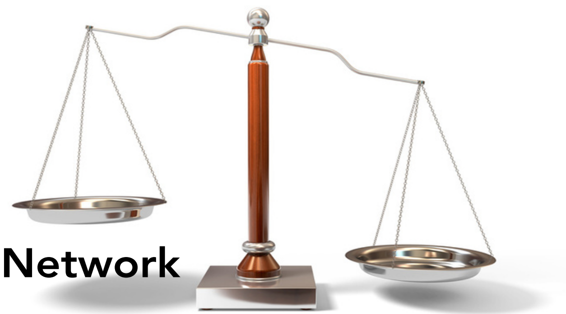

Network

Energy transport

CO_2 footprint;
Energy needed and lost

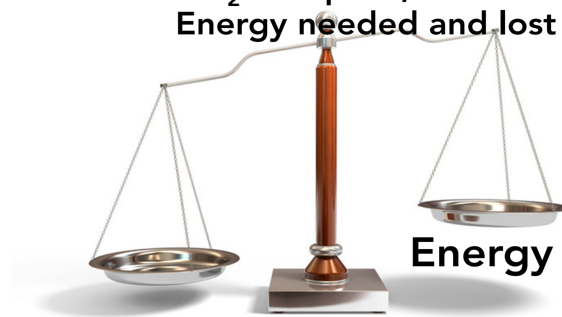

Energy transport

Network

Bits to energy

Energy to bits

Energy Efficient Ethernet (802.3az)

Power savings techniques in hardware can be leveraged in architecturing communication patterns in data centra

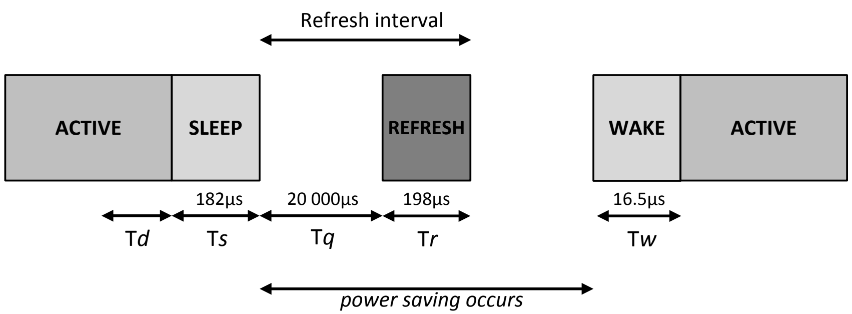

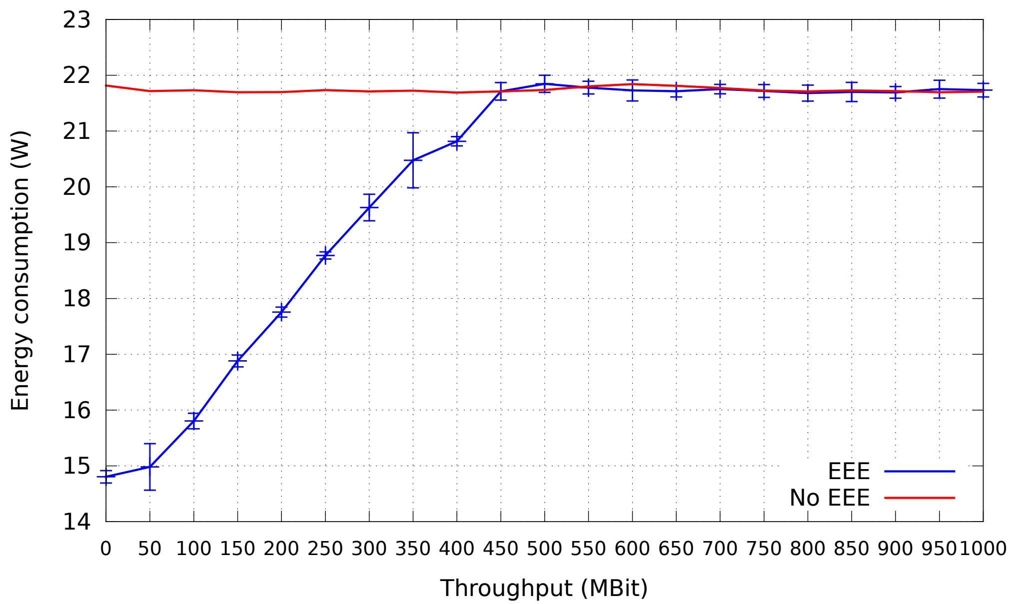

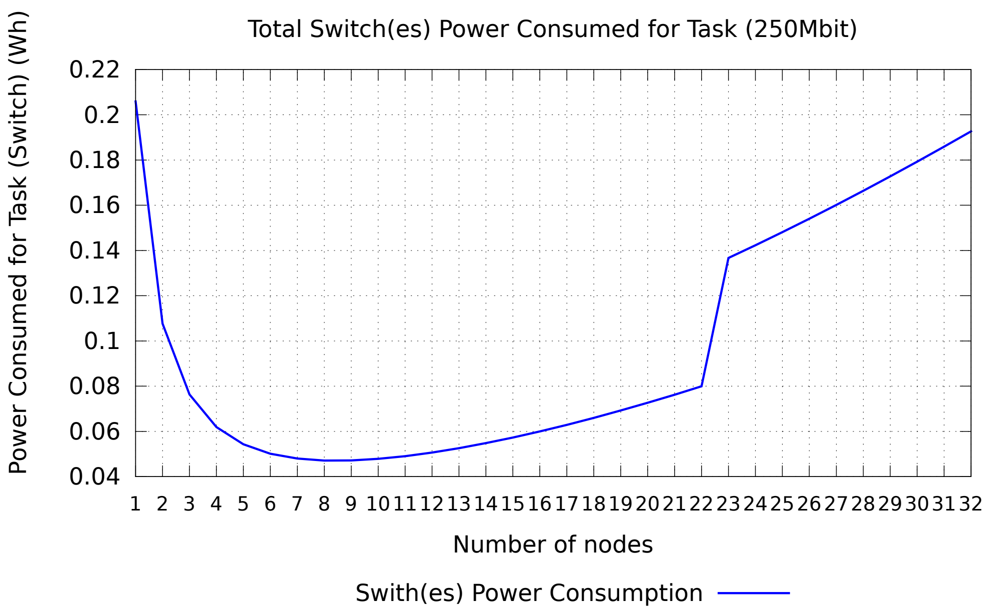

D. Pavlov and J. Soert and P. Grosso and Z. Zhao and K. van der Veldt and H. Zhu and C. de Laat
Towards energy efficient data intensive computing using IEEE 802.3az
In: DISCS 2012 workshop - Nov 2012

Energy saving in clouds

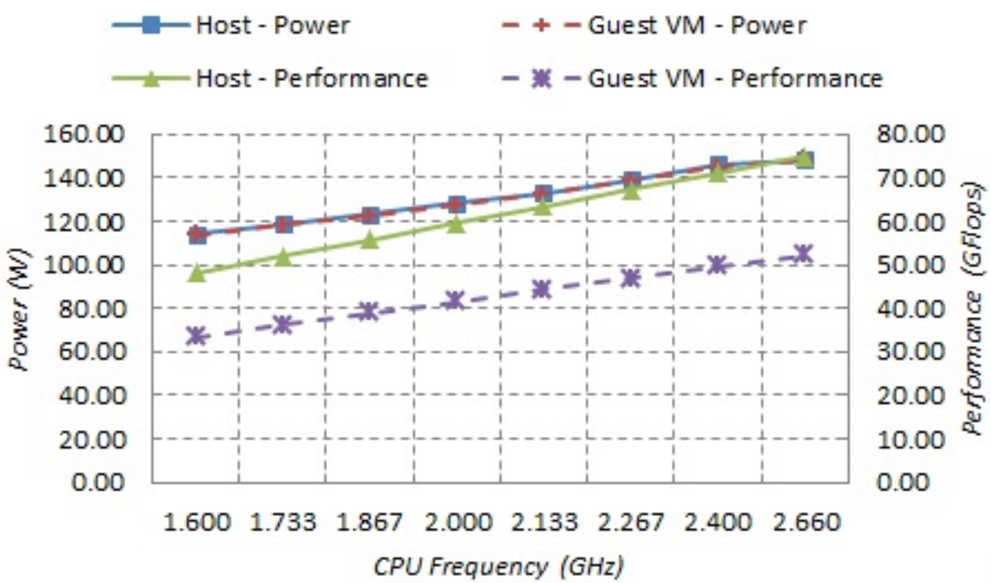

Quantifying the energy performance of VMs is the first step toward energy-aware job scheduling.

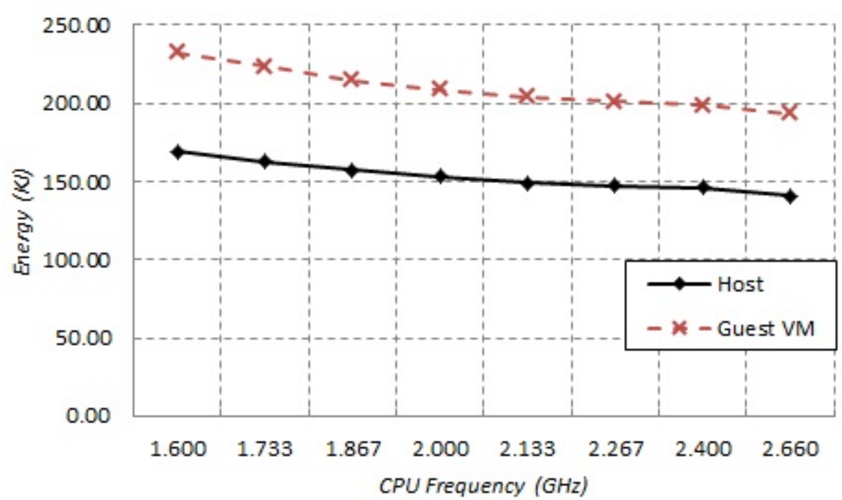

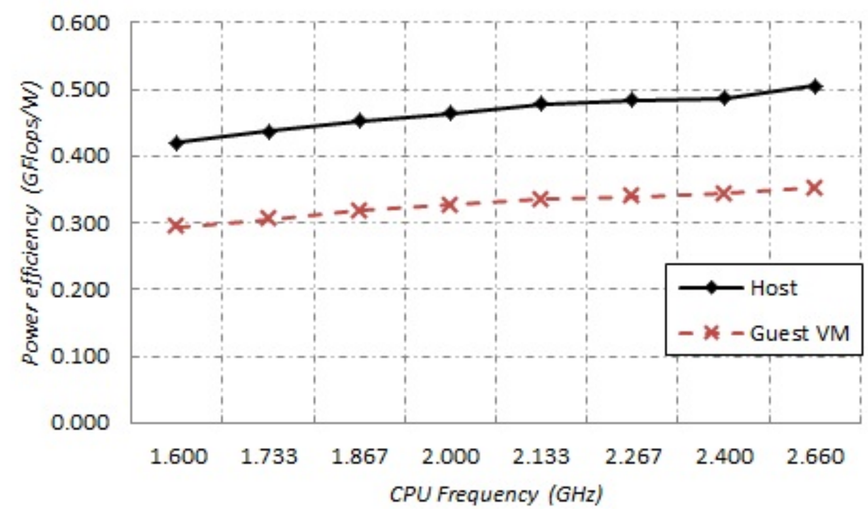

Q. Chen, P. Grosso, K. van der Veldt, C. de Laat, R. Hofman and H.Bal.
Profiling energy consumption of VMs for green cloud computing
In: International Conference on Cloud and Green Computing (CGC2011), Sydney December 2011

Networks and CO_2

- Take a network (ESnet, working on using SURFnet data)
- Define the traffic model running on it
- Use the energy monitoring information and energy costs data
- Compare path selection strategies : shortest, cheapest and greenest

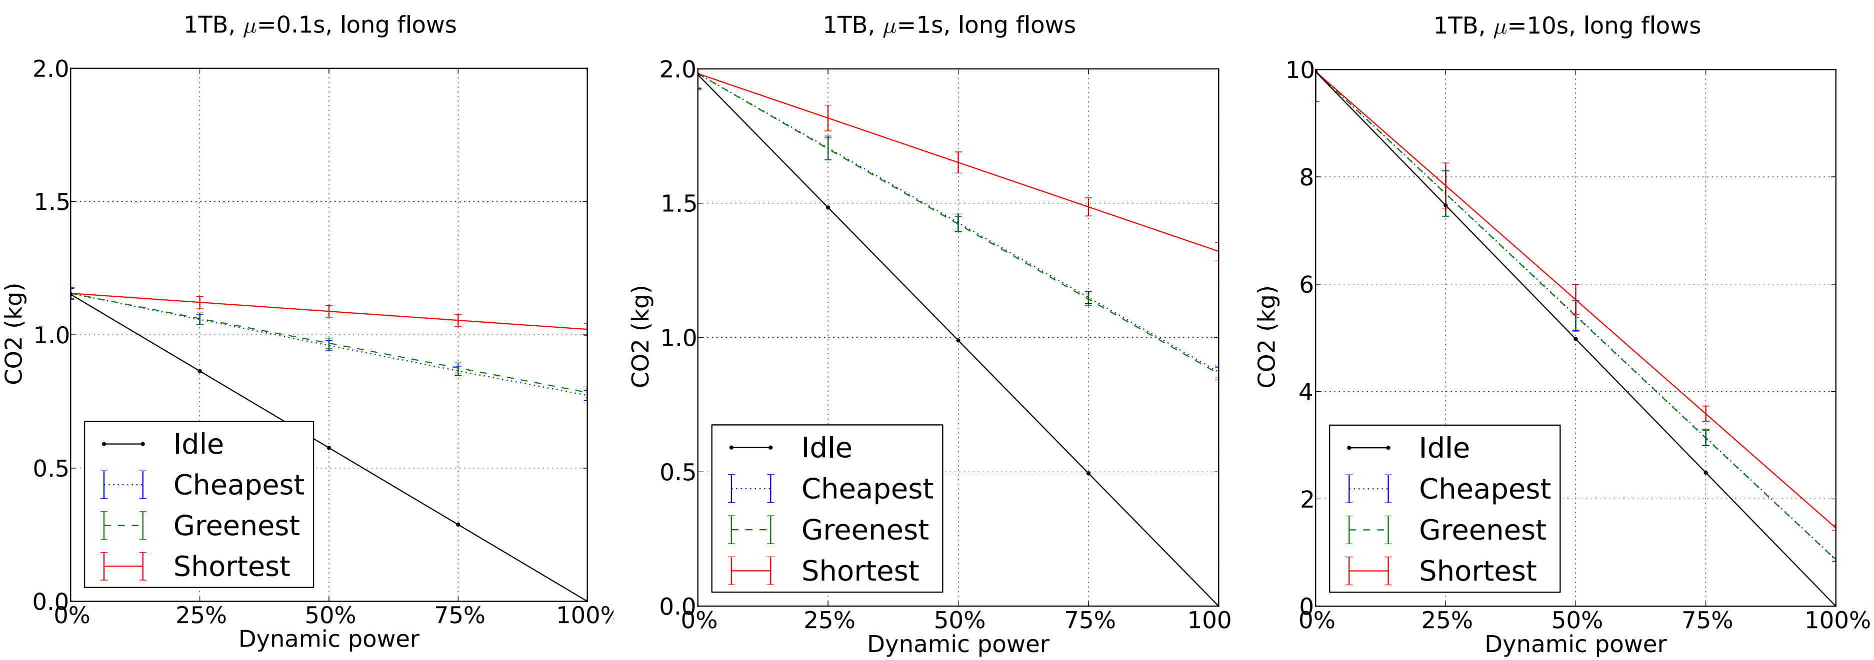

“A motivation for carbon aware path provisioning for NRENs” (submitted to eEnergy2014)

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- *Capacity*
  - *Bandwidth on demand, QoS, architectures, photonics, performance*
- ***Capability***
  - ***Programmability, virtualization, complexity, semantics, workflows***
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- *Sustainability*
  - *Greening infrastructure, awareness*
- ***Resilience***
  - ***Systems under attack, failures, disasters***

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 -> OpenFlow & SDN
  - -> Virtualized Programmable Networks

User Programmable Virtualized Networks allows the results of decades of computer science to handle the complexities of application specific networking.

- The network is virtualized as a collection of resources
- UPVNs enable network resources to be programmed as part of the application
- Mathematica, a powerful mathematical software system, can interact with real networks using UPVNs

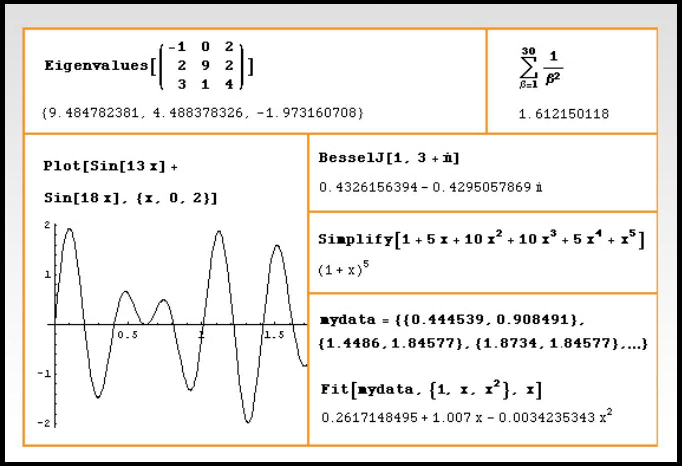

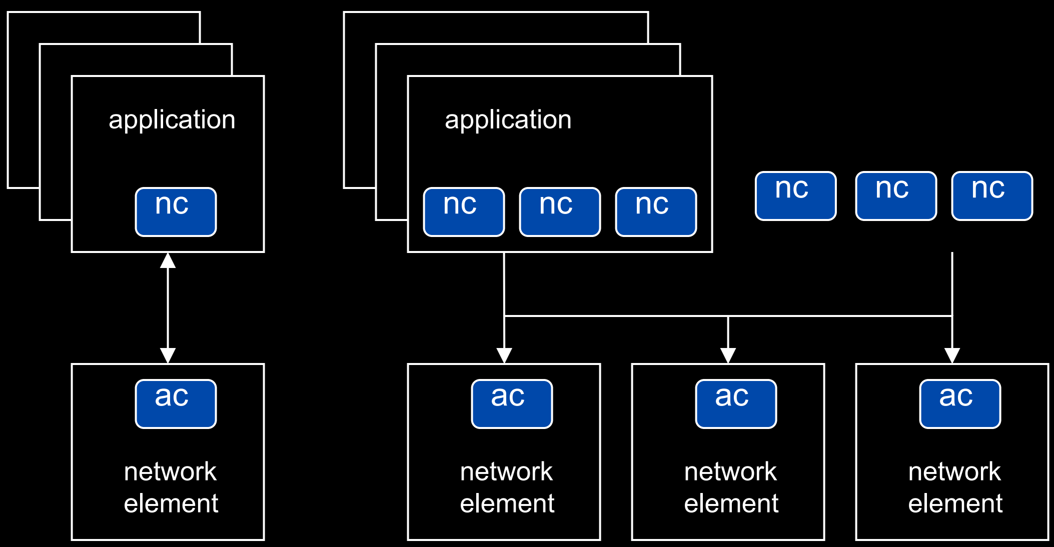

Mathematica enables advanced graph queries, visualizations and real-time network manipulations on UPVNs

Topology matters can be dealt with algorithmically
Results can be persisted using a transaction service built in UPVN

Initialization and BFS discovery of NEs

```
Needs["WebServices`"]
<<DiscreteMath`Combinatorica`
<<DiscreteMath`GraphPlot`
InitNetworkTopologyService["edge.ict.tno.nl"]

Available methods:

{DiscoverNetworkElements,GetLinkBandwidth,GetAllIpLinks,Remote,
NetworkTokenTransaction}

Global`upvnverbose = True;

AbsoluteTiming[nes = BFSDiscover["139.63.145.94"];][[1]]

AbsoluteTiming[result = BFSDiscoverLinks["139.63.145.94", nes];][[1]]

Getting neigbours of: 139.63.145.94
Internal links: {192.168.0.1, 139.63.145.94}
(...)
Getting neigbours of:192.168.2.3

Internal links: {192.168.2.3}
```

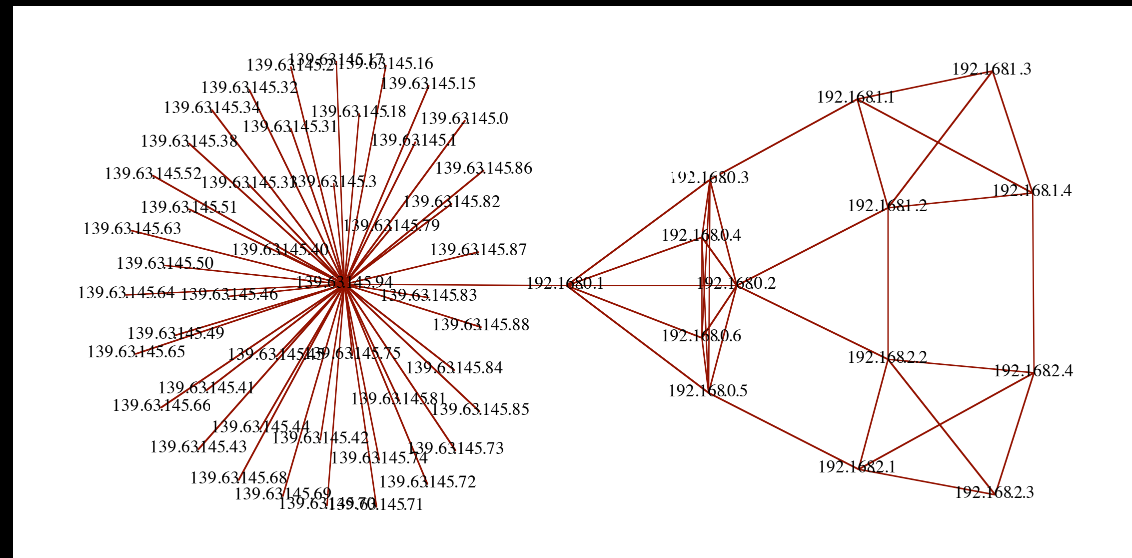

Transaction on shortest path with tokens

```
nodePath = ConvertIndicesToNodes[
        ShortestPath[
            g,
            Node2Index[nids,"192.168.3.4"],
            Node2Index[nids,"139.63.77.49"]],
        nids];
Print["Path: ", nodePath];
If[NetworkTokenTransaction[nodePath, "green"]==True,
    Print["Committed"], Print["Transaction failed"]];

Path:
{192.168.3.4,192.168.3.1,139.63.77.30,139.63.77.49}

Committed
```

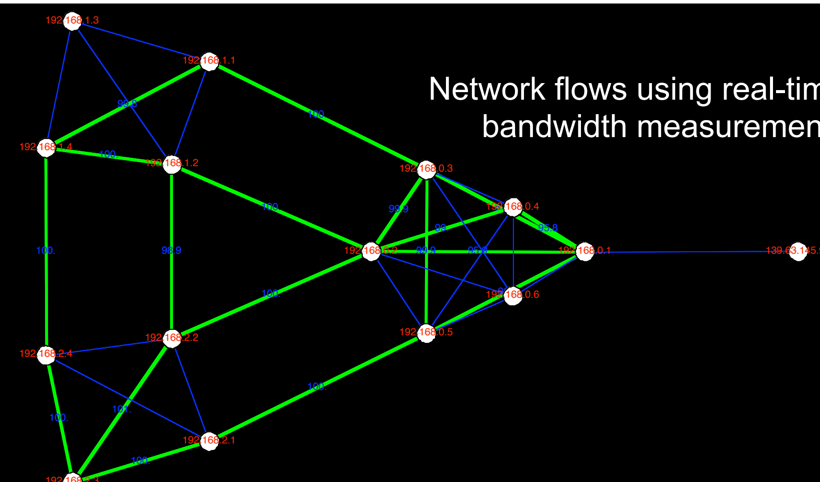

Network flows using real-time bandwidth measurements

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.

StarPlane

CDN on Demand in the cloud

Interactive programmable networks

Basic operating system loop

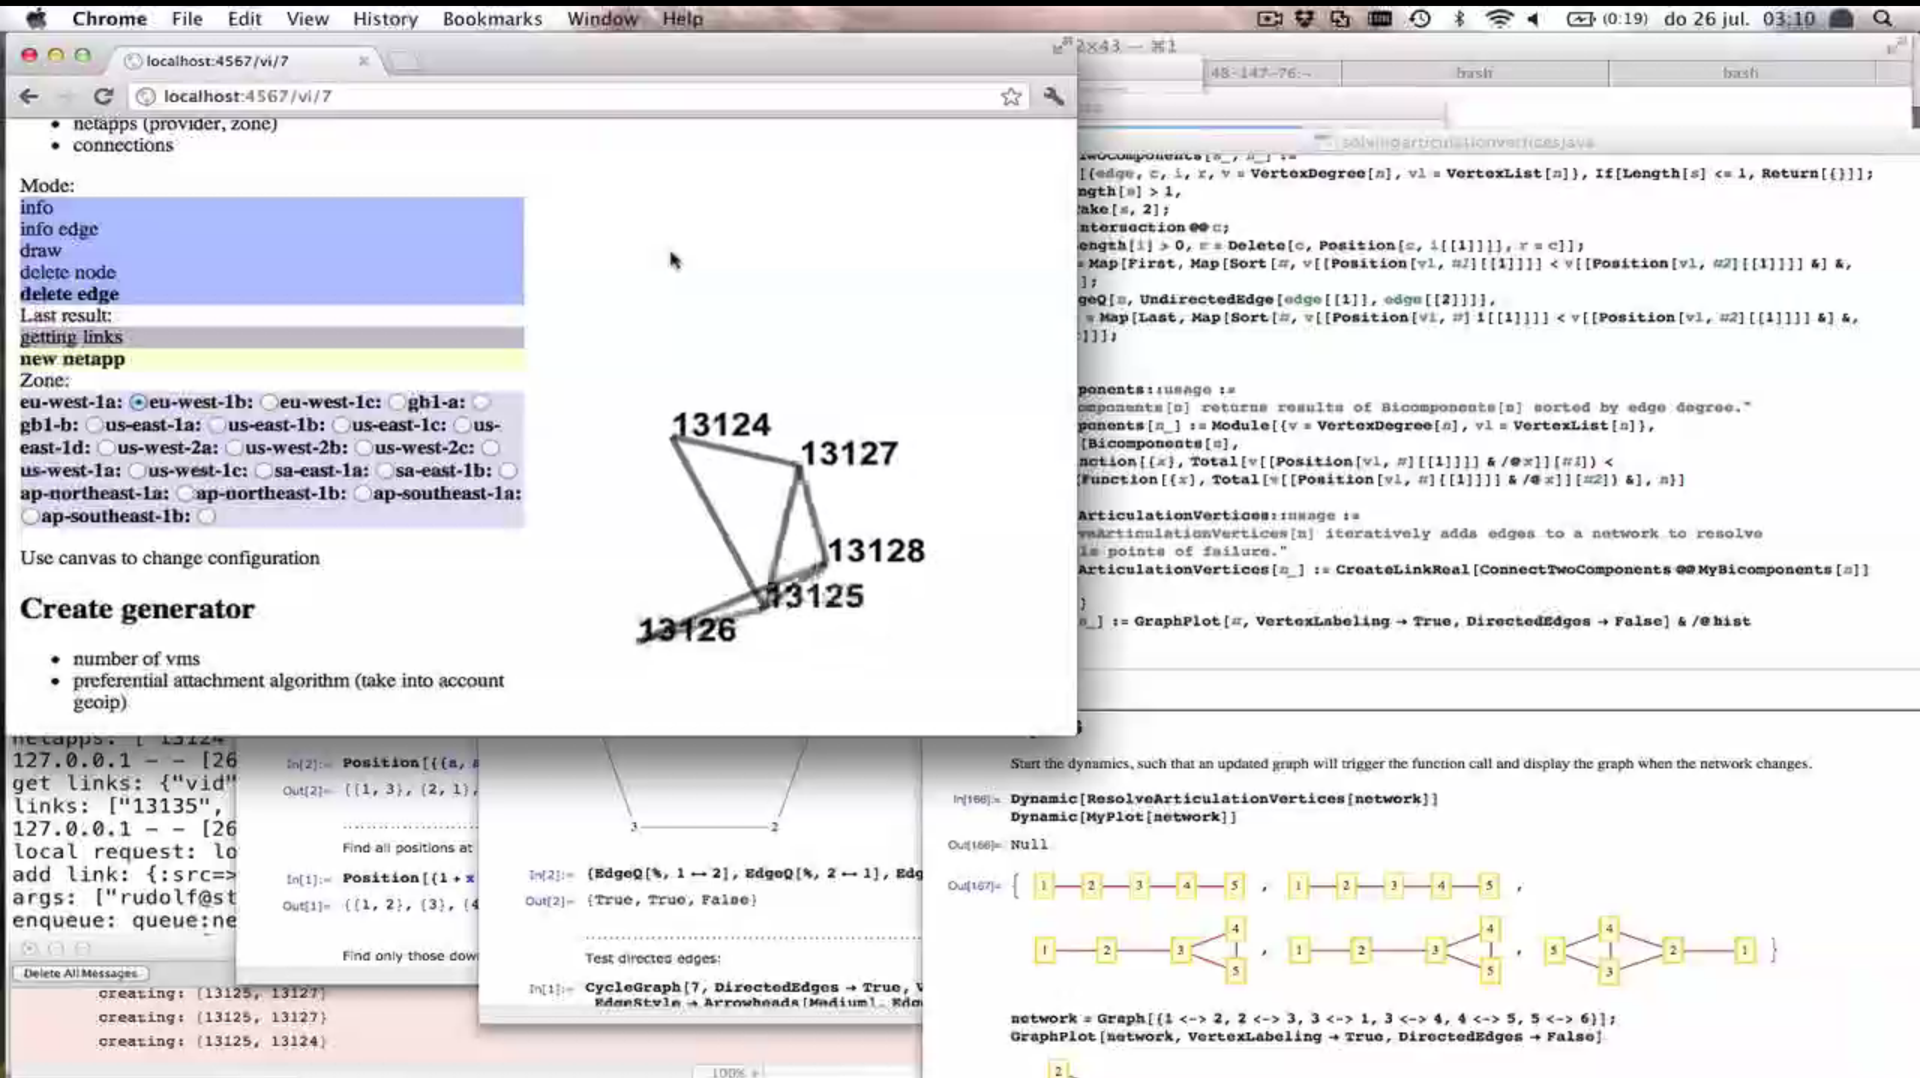

Mission

Can we create smart and safe data processing infrastructures that can be tailored to diverse application needs?

- *Capacity*
  - *Bandwidth on demand, QoS, architectures, photonics, performance*
- *Capability*
  - *Programmability, virtualization, complexity, semantics, workflows*
- *Security*
  - *Anonymity, integrity of data in distributed data processing*
- *Sustainability*
  - *Greening infrastructure, awareness*
- *Resilience*
  - *Systems under attack, failures, disasters*

SMART

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!

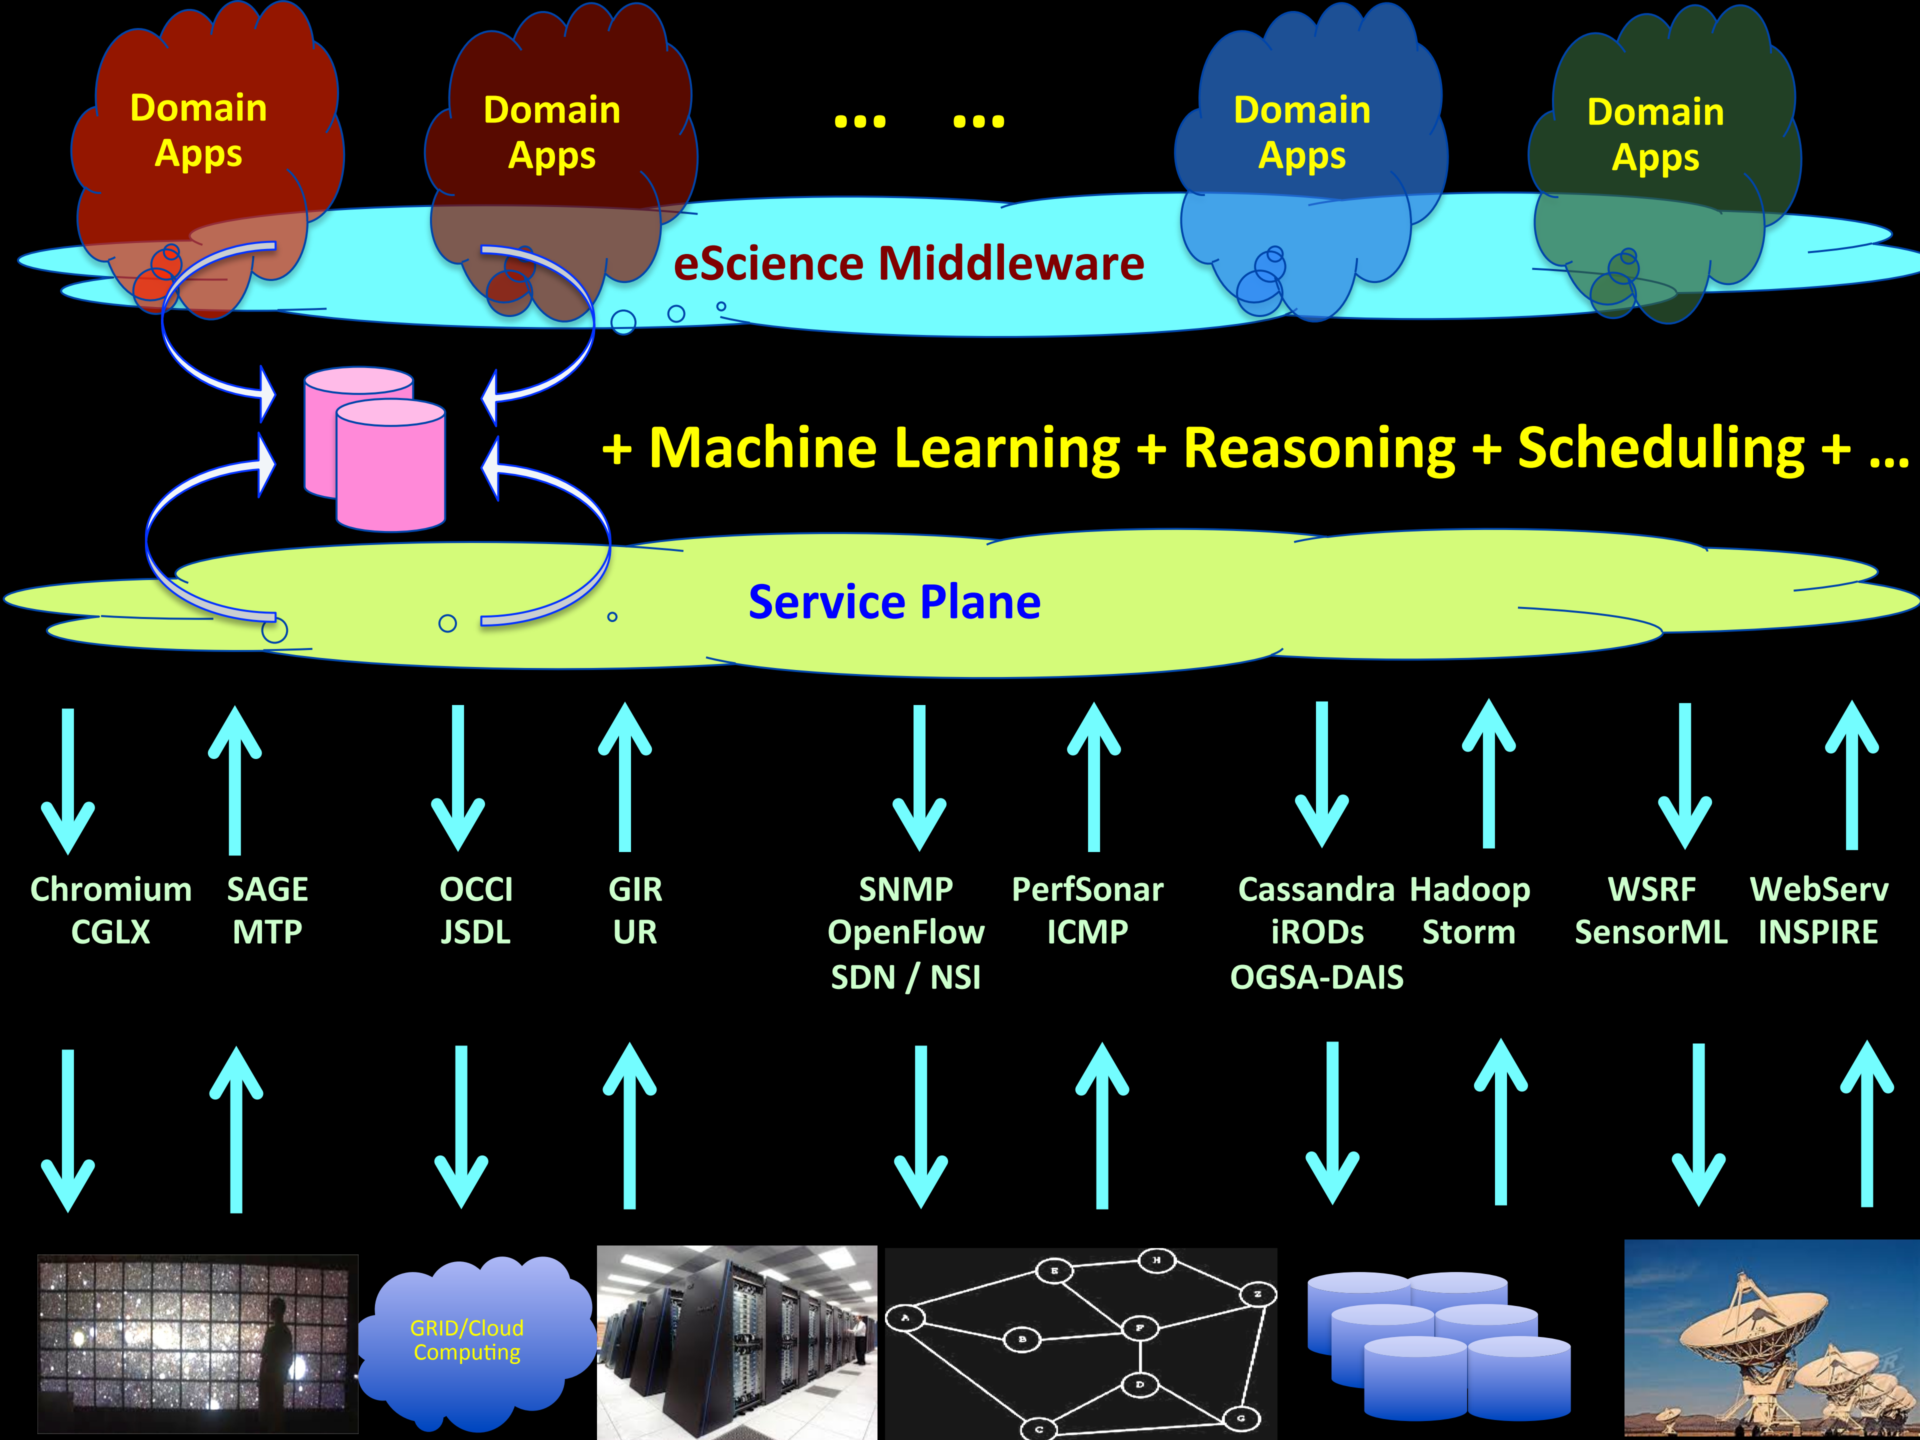

Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
+ Machine Learning + Reasoning + Scheduling + ...
Service Plane
Chromium CGLX
SAGE MTP
OCCI JSDL
GIR UR
SNMP OpenFlow SDN / NSI
PerfSonar ICMP
Cassandra iRODs OGSA-DAIS
Hadoop Storm
WSRF SensorML
WebServ INSPIRE
GRID/Cloud Computing

Layers

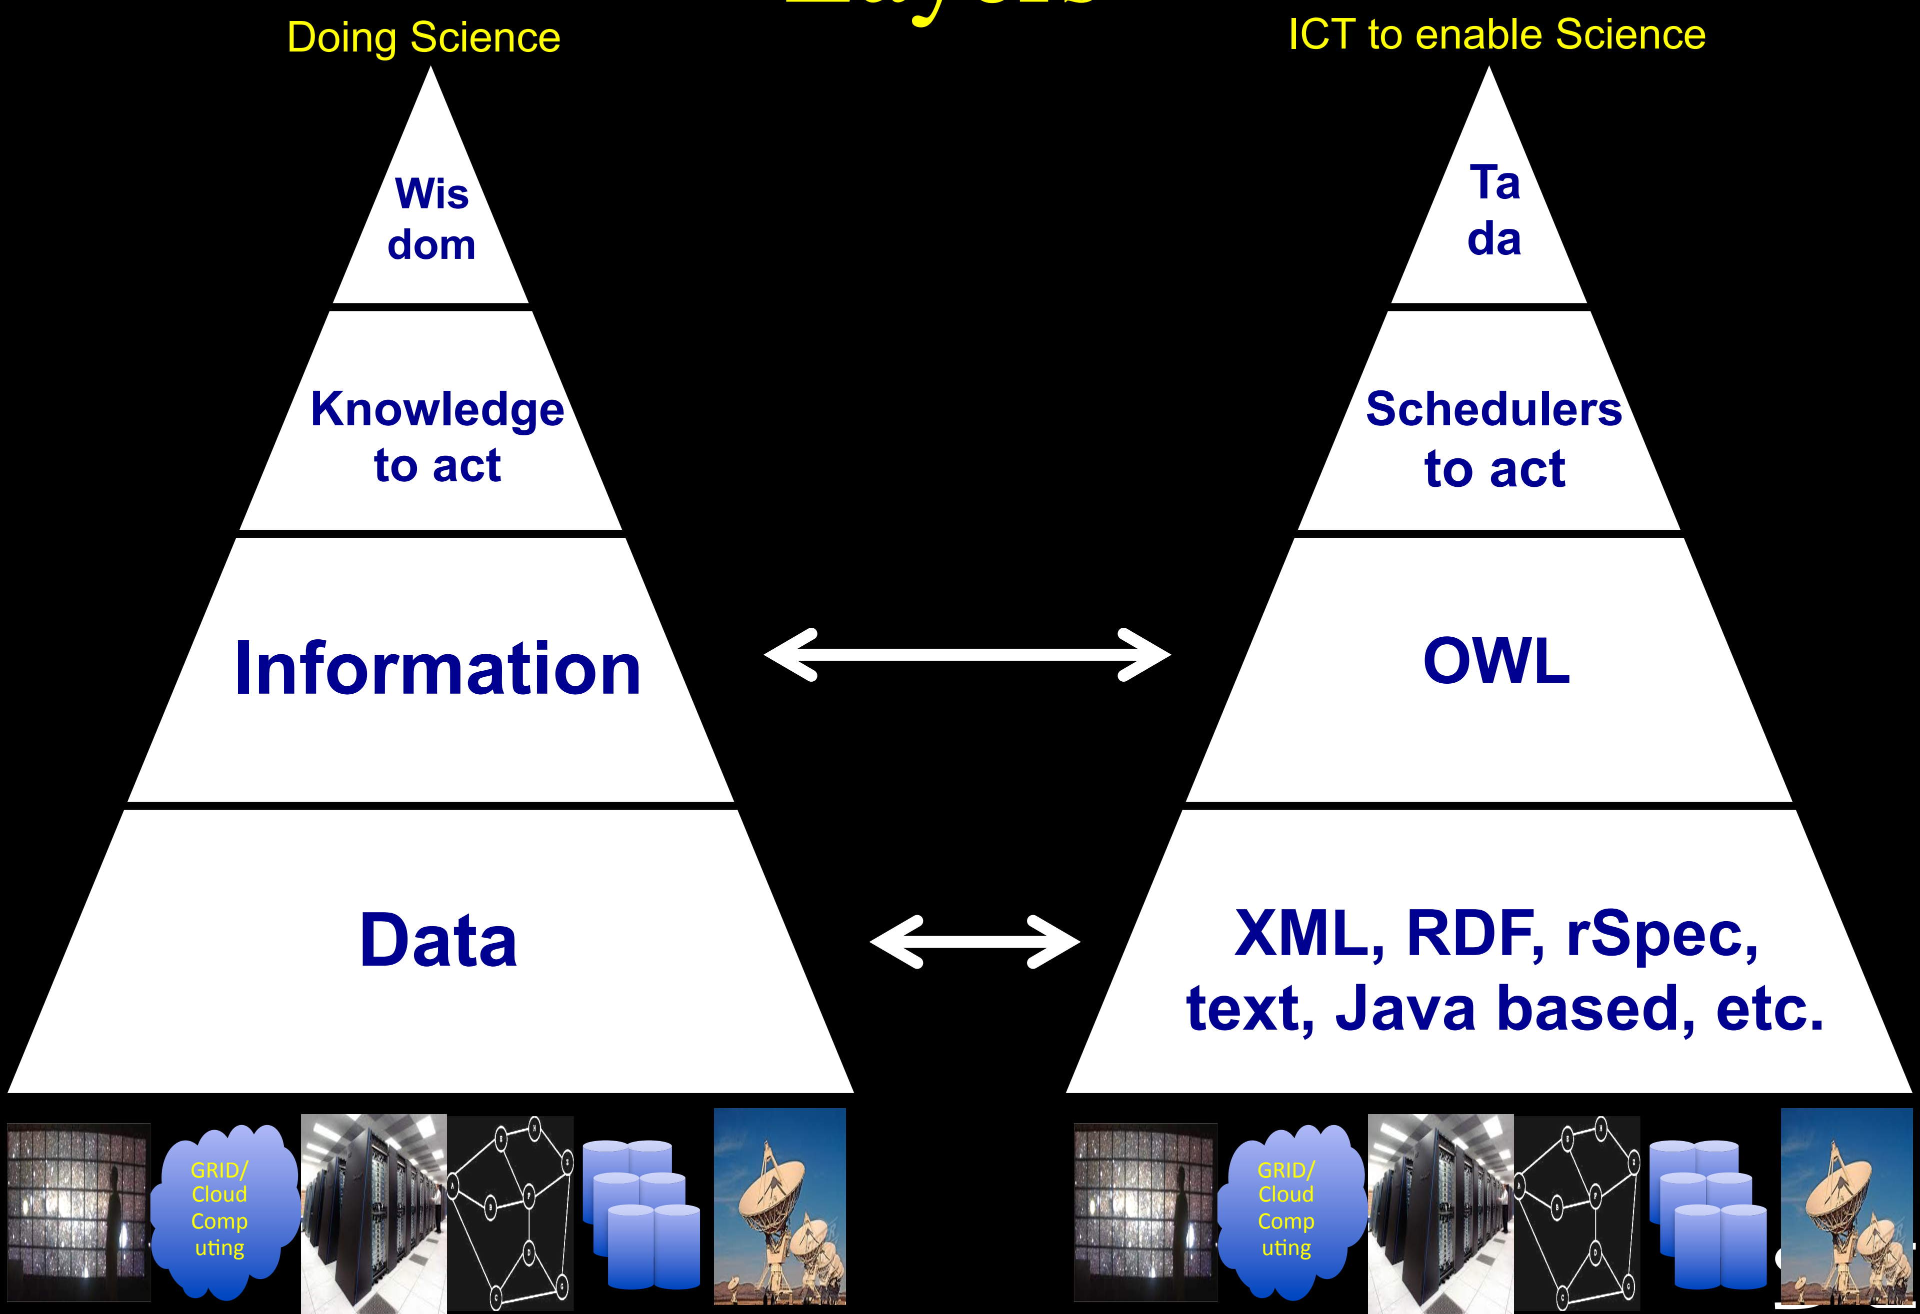
Doing Science
Wis dom
Knowledge to act
Information
Data
GRID/ Cloud Comp uting
ICT to enable Science
Ta da
Schedulers to act
OWL
XML, RDF, rSpec, text, Java based, etc.
GRID/ Cloud Comp uting

The Big Data Challenge

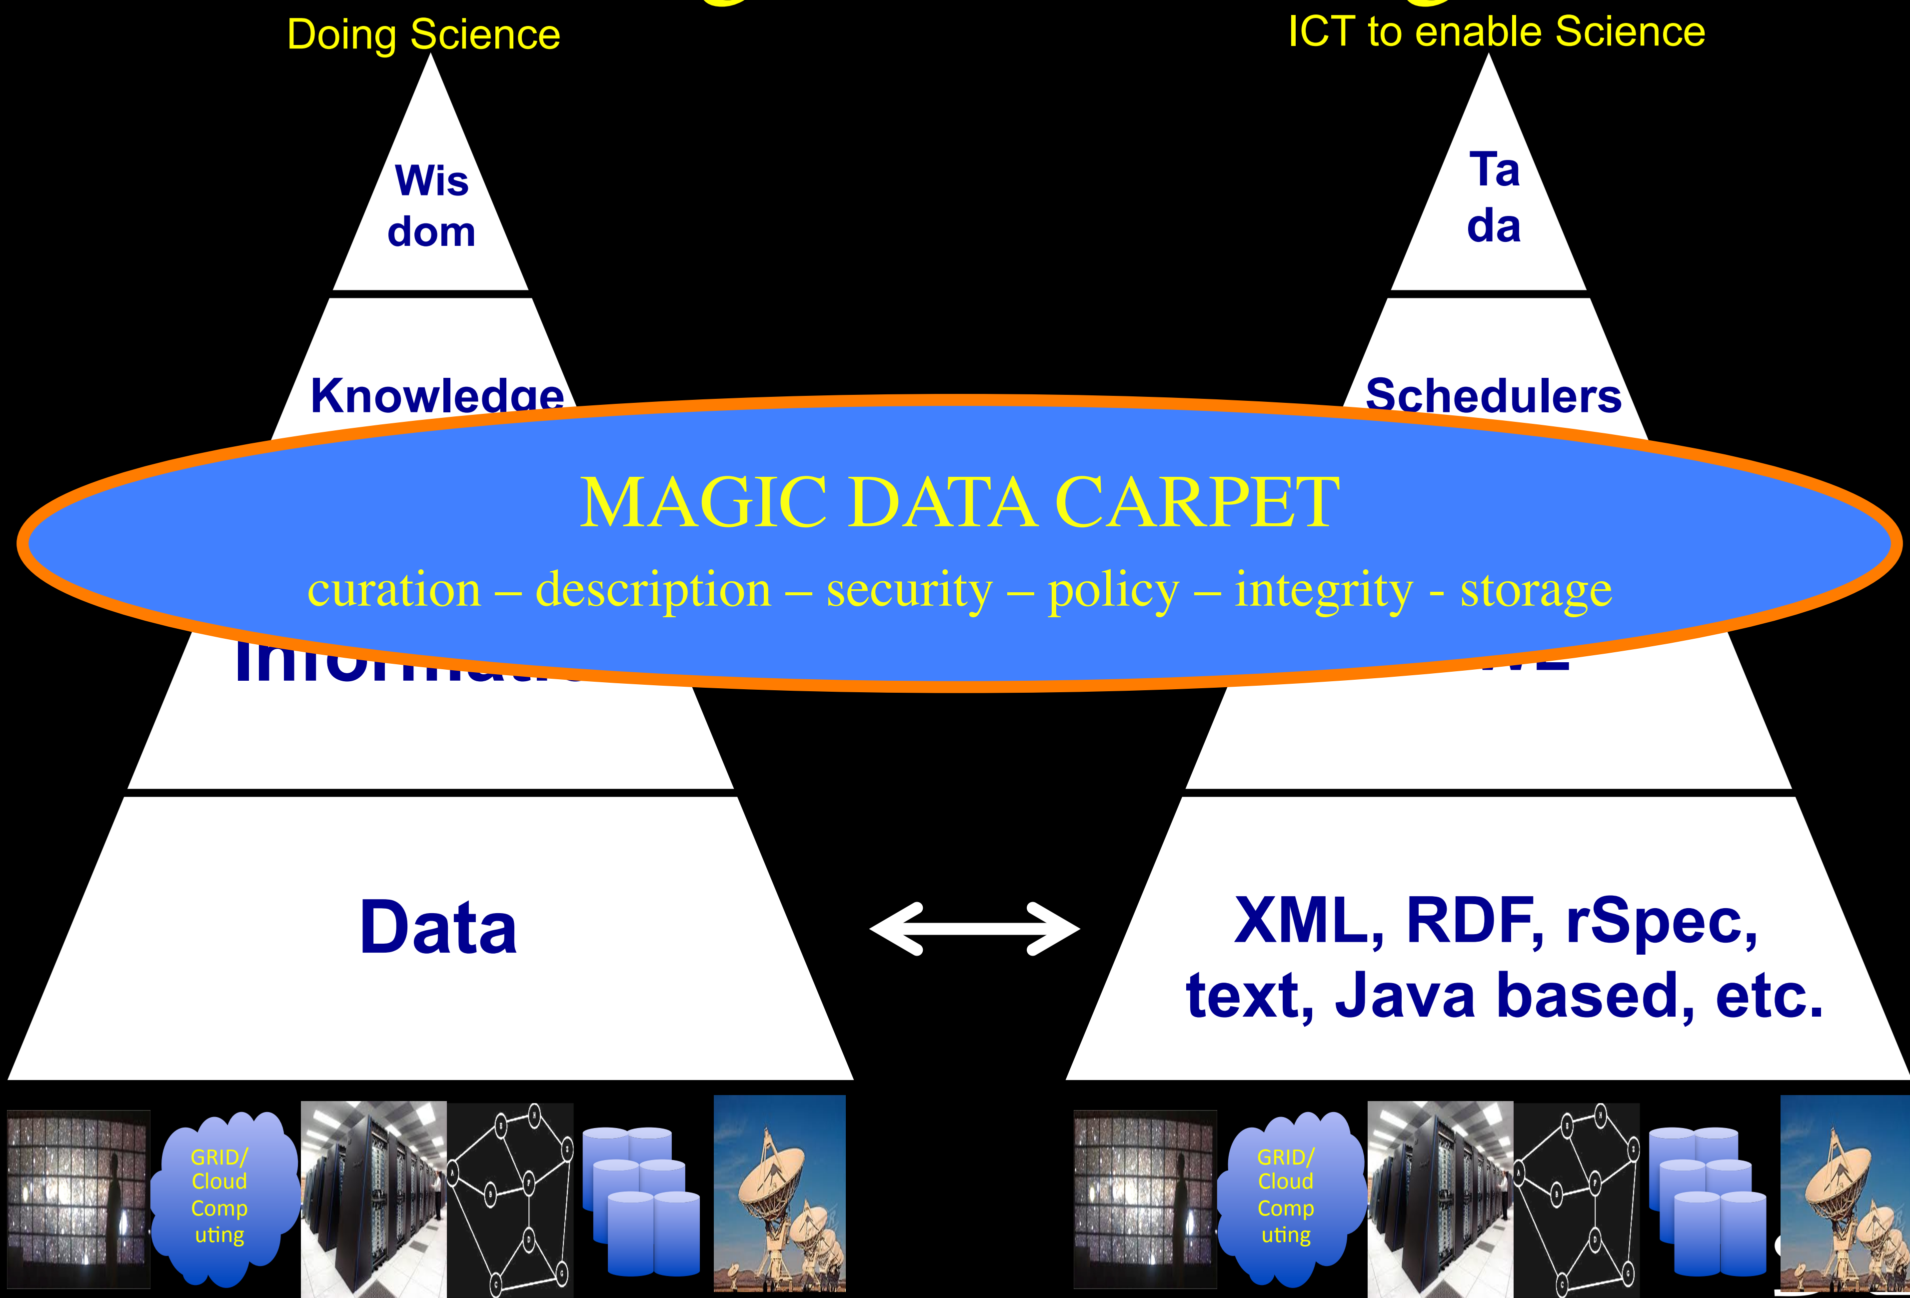

TimeLine

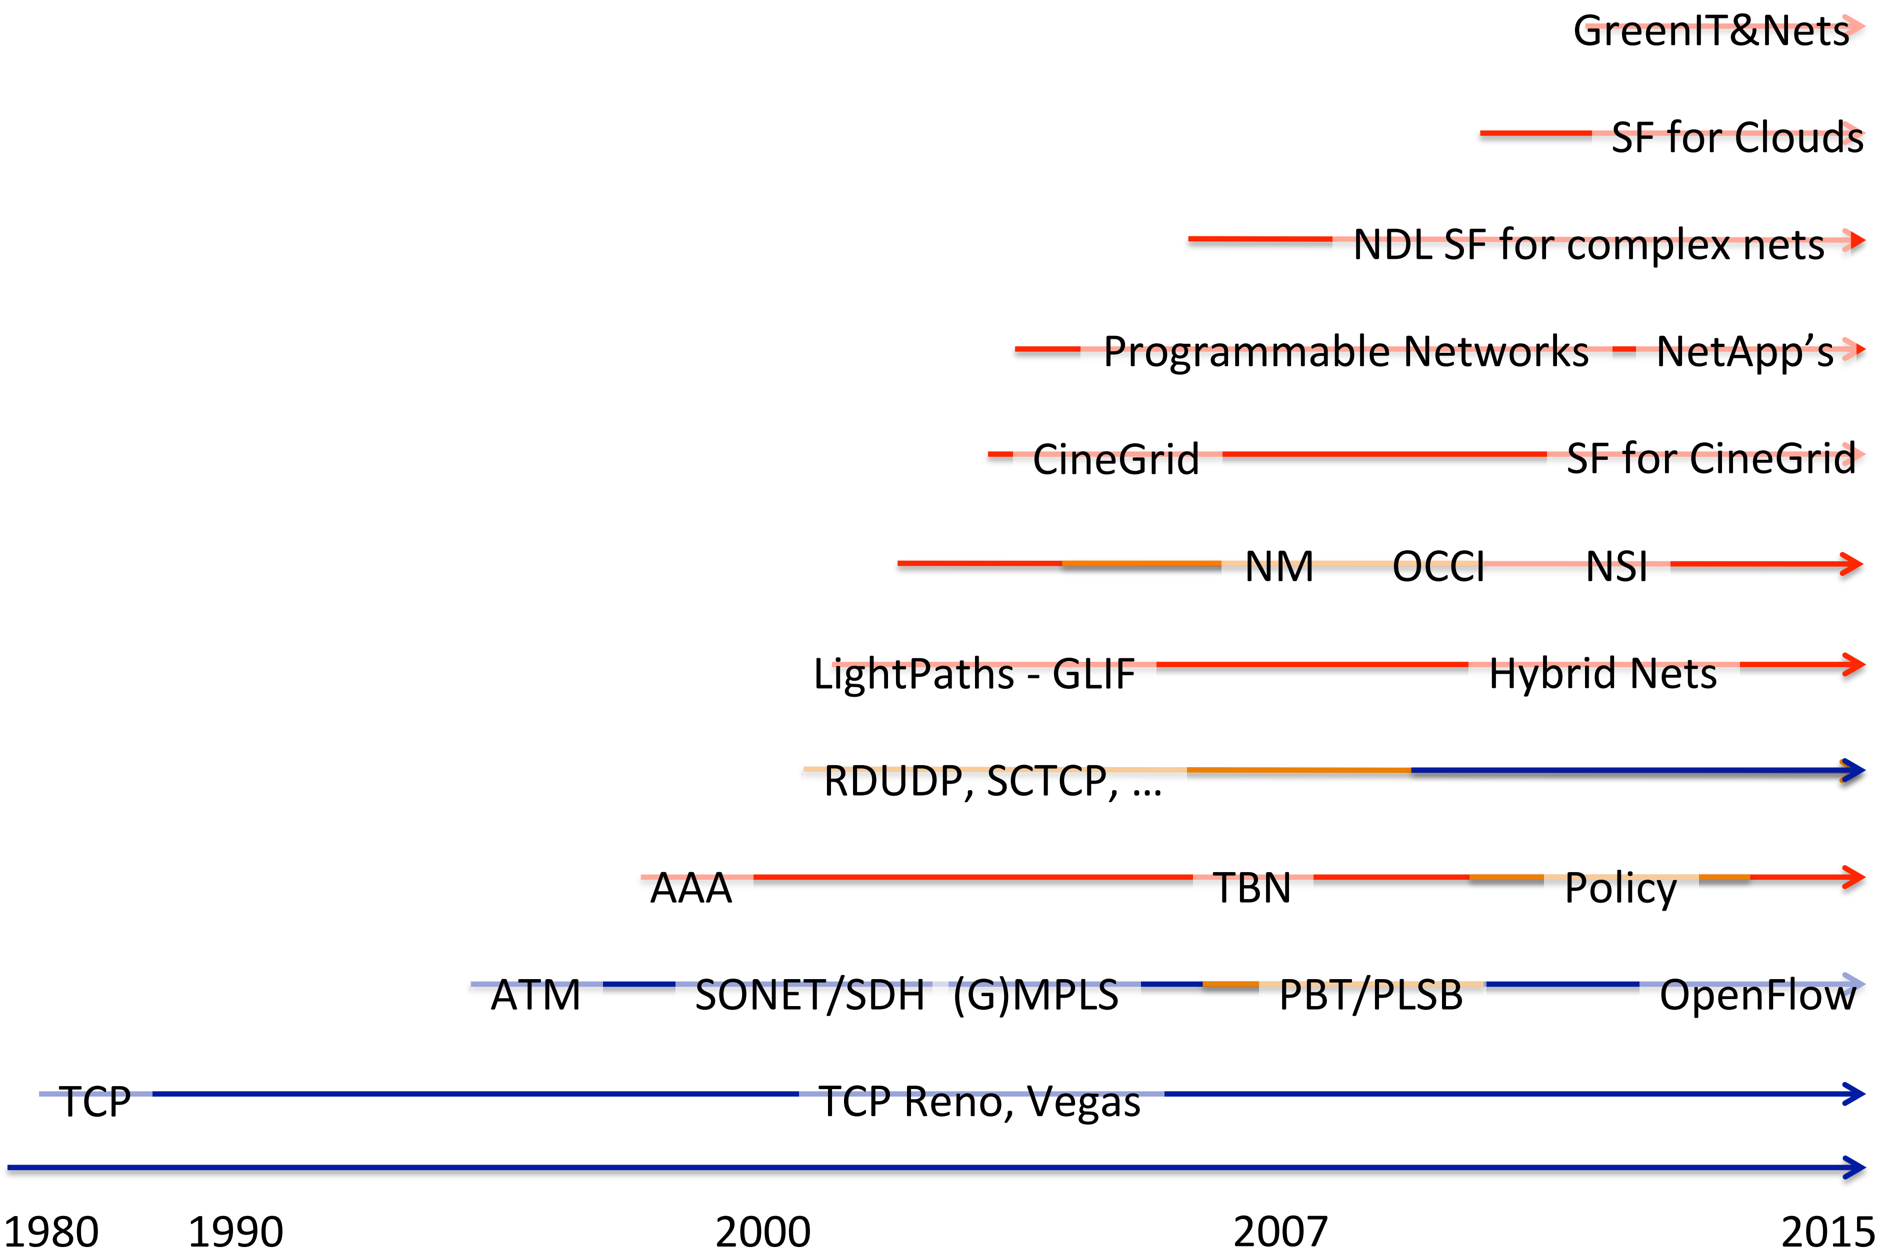
GreenIT&Nets
SF for Clouds
NDL SF for complex nets
Programmable Networks
NetApp's
CineGrid
SF for CineGrid
NM
OCCI
NSI
LightPaths - GLIF
Hybrid Nets
RDUDP, SCTCP, …
AAA
TBN
Policy
ATM
SONET/SDH
(G)MPLS
PBT/PLSB
OpenFlow
TCP
TCP Reno, Vegas
1980
1990
2000
2007
2015

TimeLine

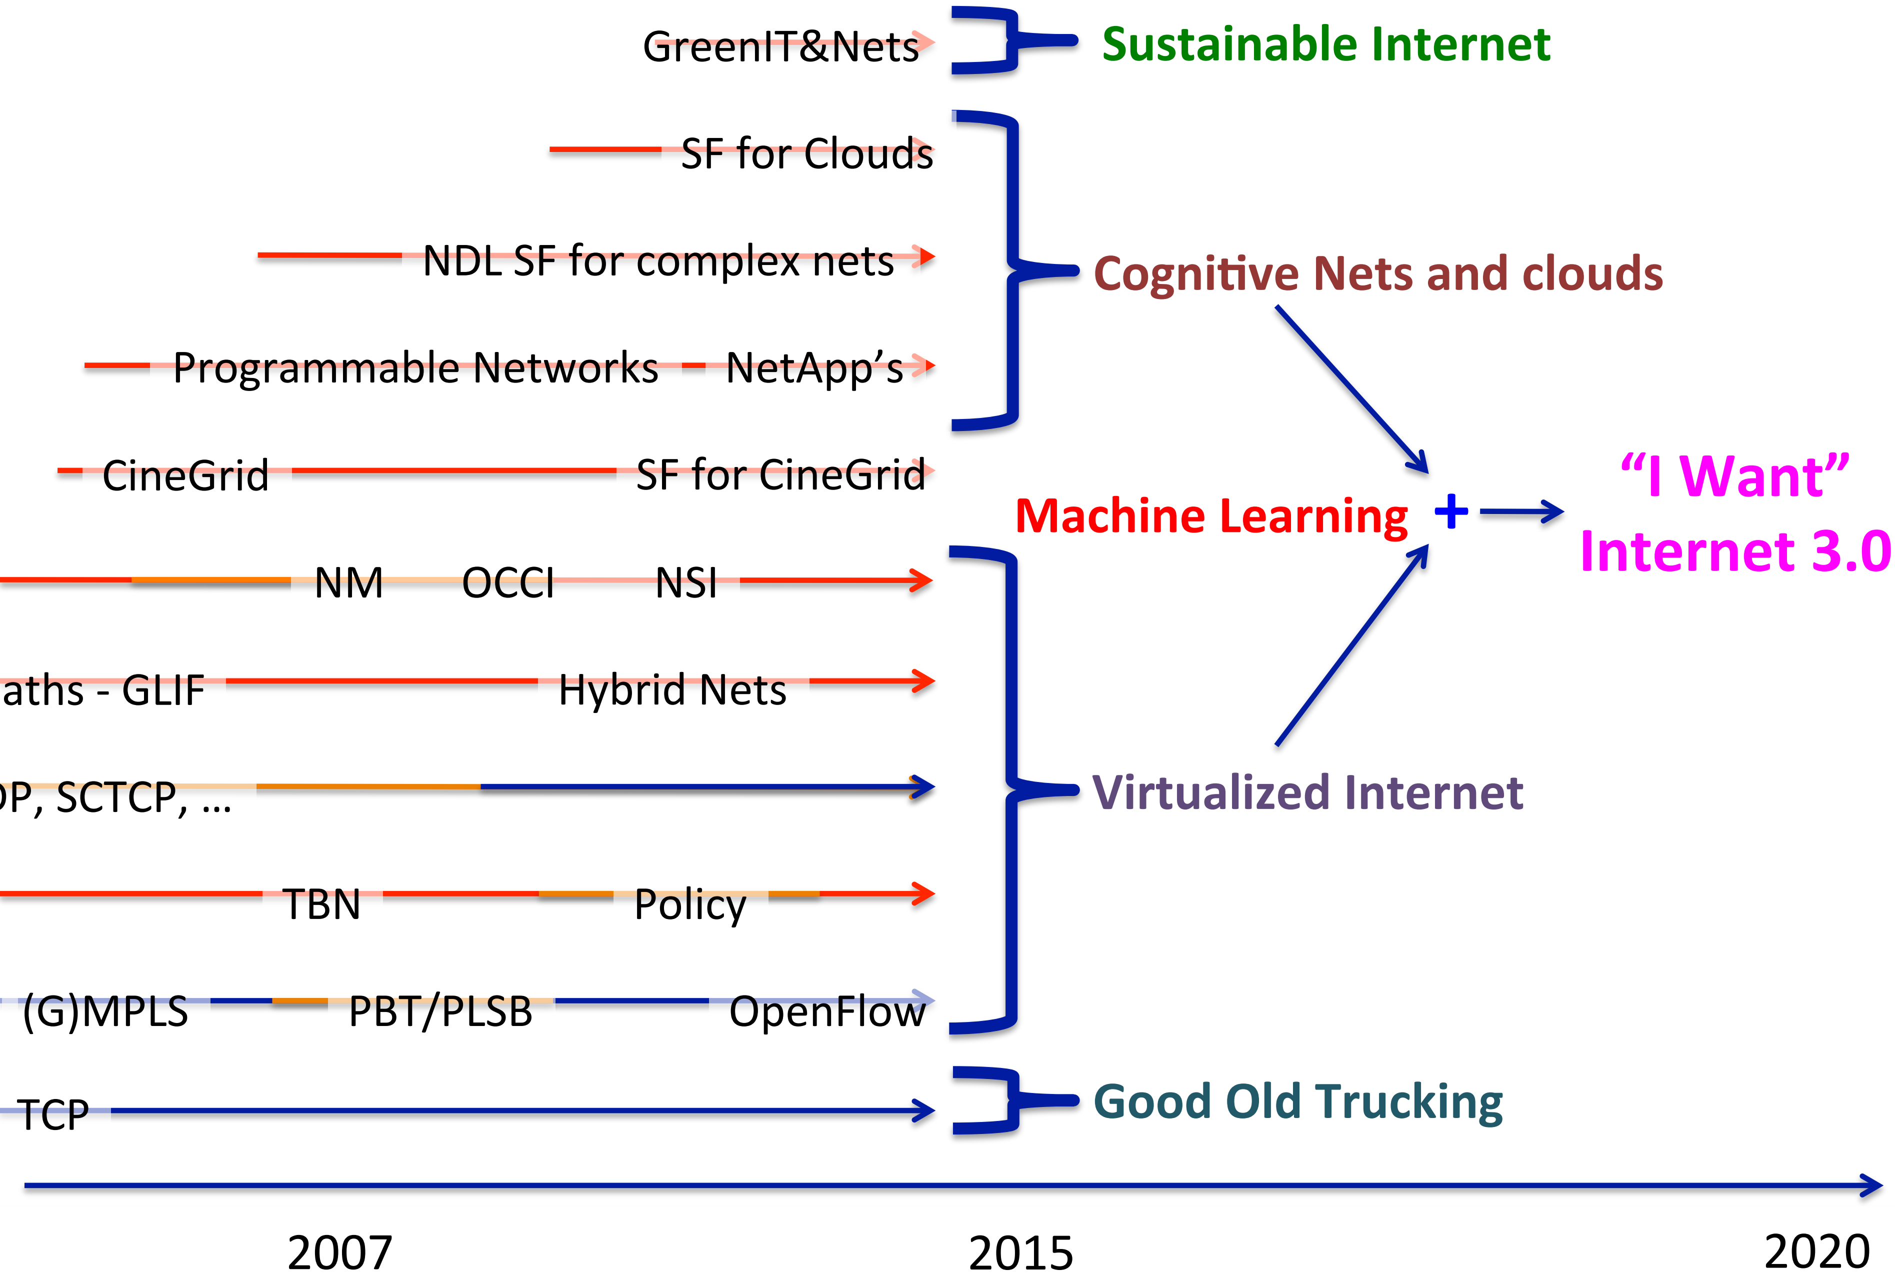

GreenIT&Nets
Sustainable Internet
SF for Clouds
NDL SF for complex nets
Cognitive Nets and clouds
Programmable Networks
NetApp's
CineGrid
SF for CineGrid
Machine Learning
+
"I Want"
Internet 3.0
NM
OCCI
NSI
aths - GLIF
Hybrid Nets
P, SCTCP, …
Virtualized Internet
TBN
Policy
(G)MPLS
PBT/PLSB
OpenFlow
TCP
Good Old Trucking
2007
2015
2020

TimeLine

- Sustainable Internet

- Cognitive Nets and clouds

lachine Learning + → "I Want" Internet 3.0

- Virtualized Internet

- Good Old Trucking

I retire

2020

2040

TimeLine

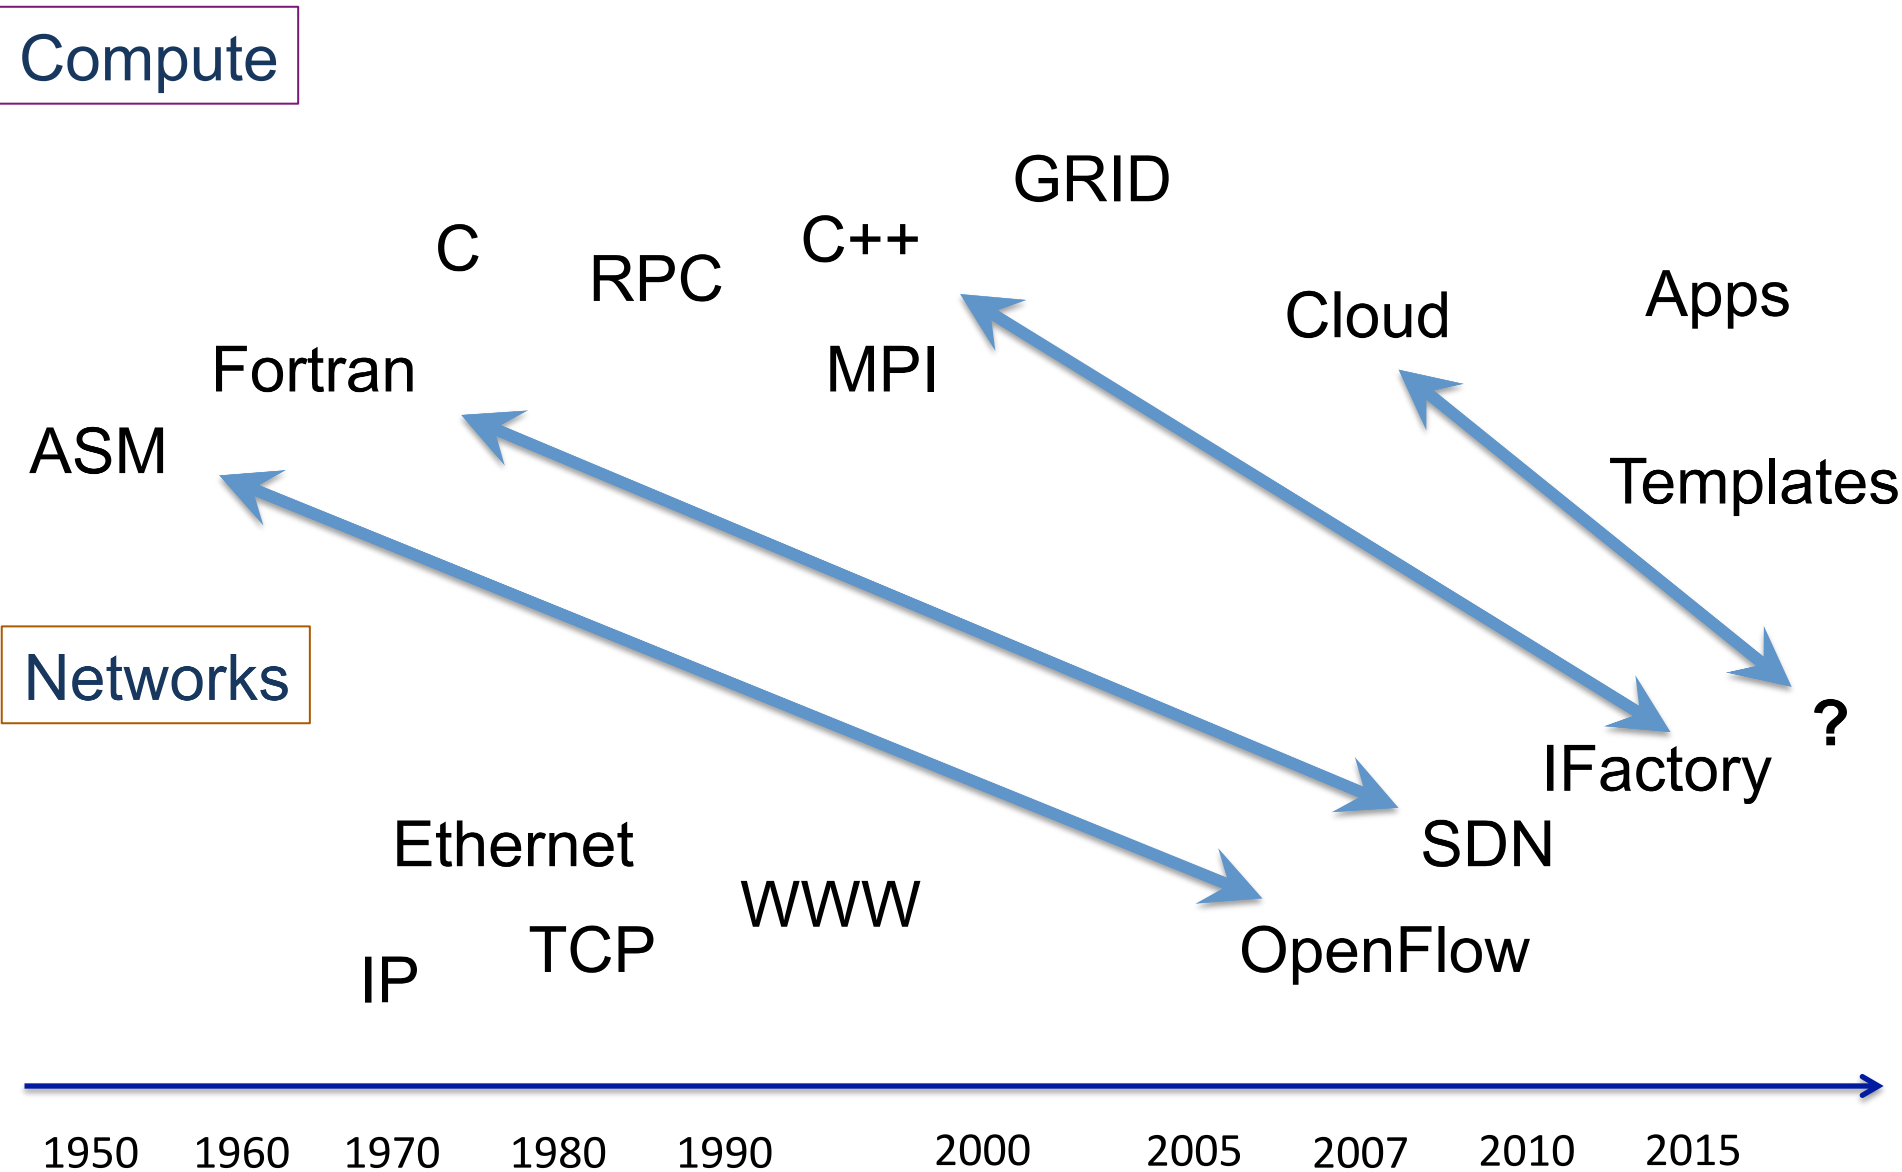

Compute
GRID
C
RPC
C++
Cloud
Apps
Fortran
MPI
ASM
Templates
Networks
?
IFactory
SDN
Ethernet
WWW
TCP
OpenFlow
IP
1950
1960
1970
1980
1990
2000
2005
2007
2010
2015

Mission

RECENT GRANTS

Can we create smart and safe data processing … infras… …uctur… that can be tailored to diverse … …atio… needs?

- *C…*
  - *Ba… …widt… …demand, QoS, a… …hit… …ures, …pho… …ics, p… …formance*
- *…apa… …lity*
  - *Prog… …ammability, virtualization, complexity, semantics, workflows*
- *Security*
  - *Anonymity, integrity of d… …in distr… …ed d… …a proc… …sing*
- *S… …abi… …y*
  - *Greenin… …ructur… …ess…*
- *Resili… …ce*
  - *…stems under attack, failures, disasters*

SARNET: Security Autonomous Response with programmable NETworks

Cees de Laat

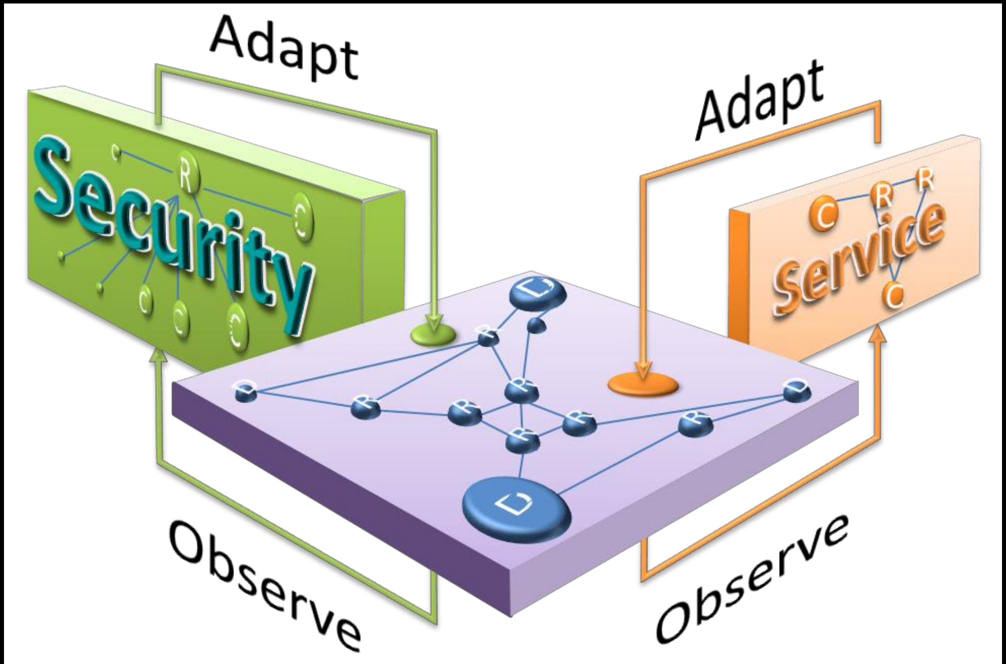

SARNET

Cyber Security program

PI: CdL

Co-Pi's: RM, LG, RW

- 400 + 285 + 300 kEuro:
- 3 PhD's
- Programming and
- Engineering manpower

- Network virtualizations and SDN
- Reasoning
- Risk evaluation
- Trust groups
- Execute response & adaptation

delaat.net/sarnet

Cyber security program

- Research goal:
  - obtain the knowledge to create ICT systems that model their state (situation), discover by observations and reasoning if and how an attack is developing and calculate the associated risks. In addition, our goal is to have the knowledge to calculate the effect of counter measures on states and their risks, choose and execute one. In short, a we research the concept of networked computer infrastructures exhibiting SAR:

Security Autonomous Response.

Line of research

- 1997: Need for authorization framework for combination of resources across domains
- 1998: AAA-ARCHitecture research in IRTF
- 2000: RFC 2903-2906, 3334
- 2005: open versus not so open exchanges
- 2006: start of trust research (also in rfc 2904)
- 2012: I2-spring session presenting line of research
- 2014: PhD defense of research plus publication
- 2015: Here we are.

Different modes of control

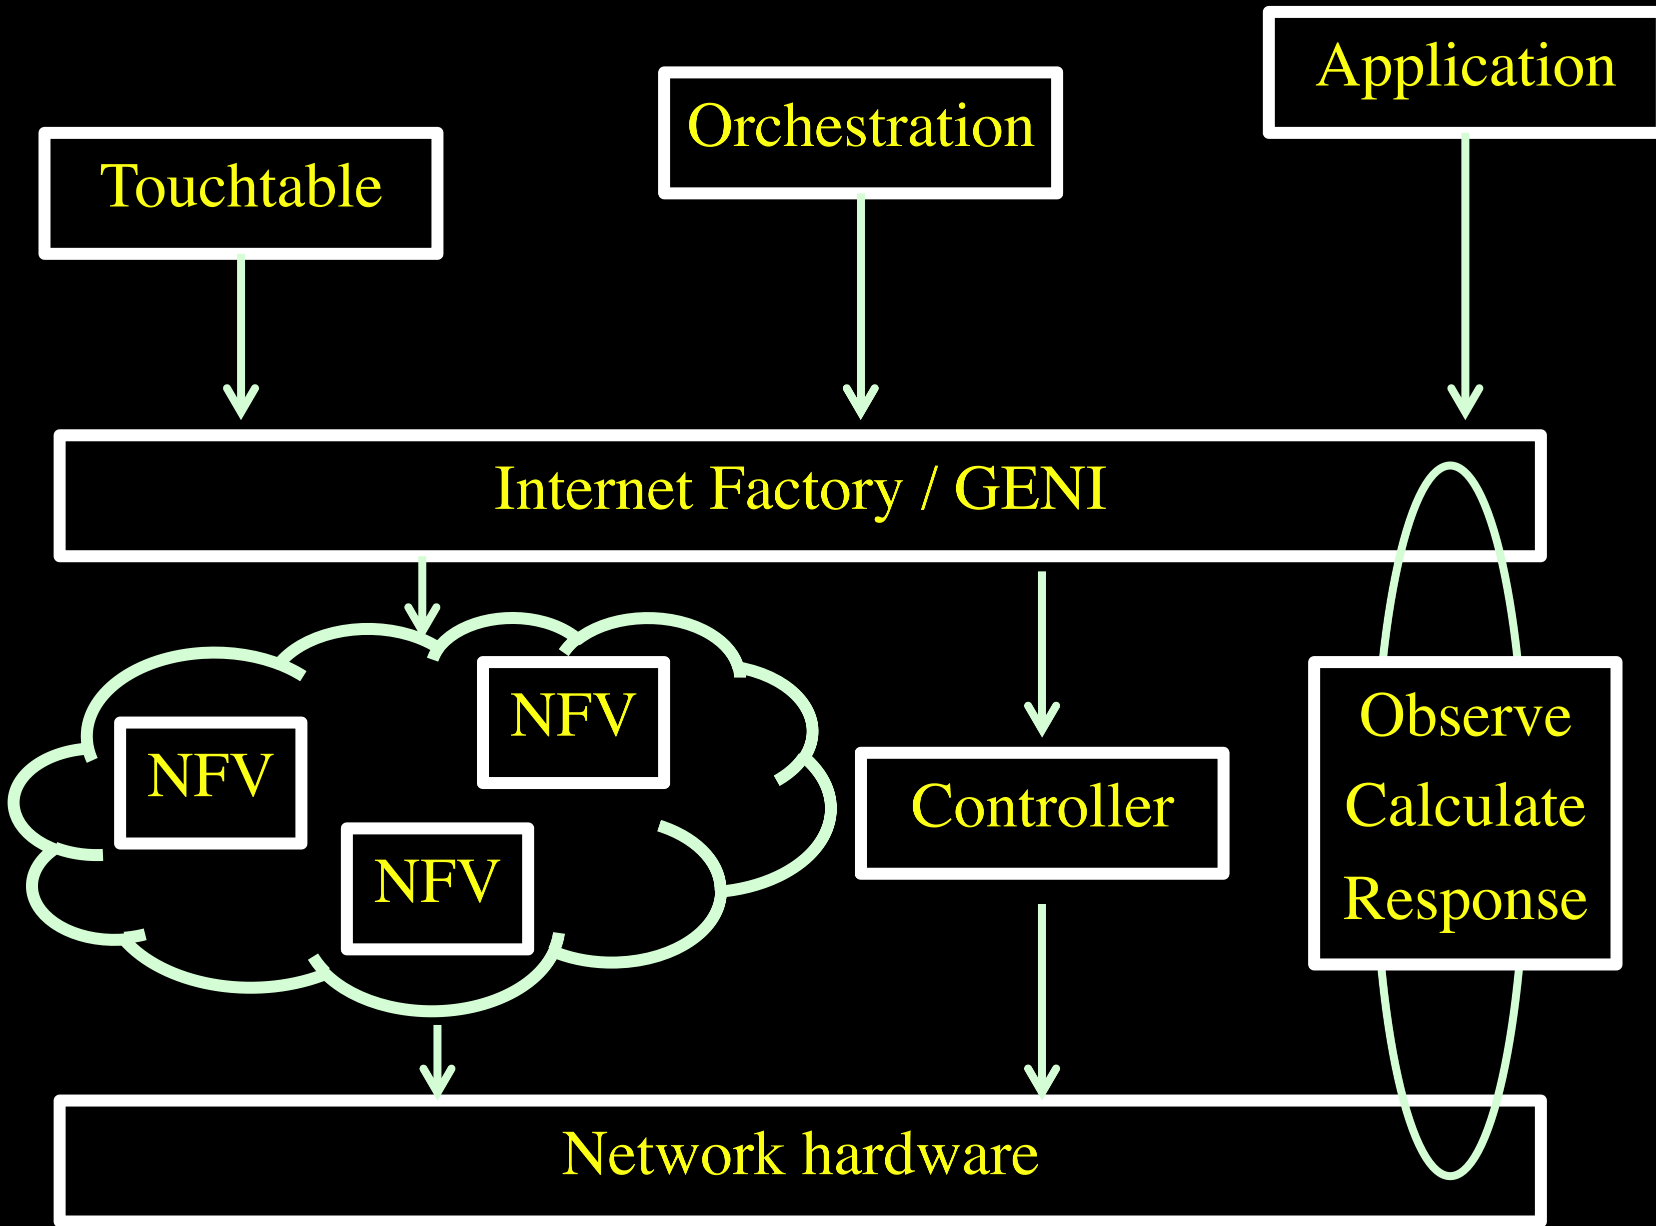

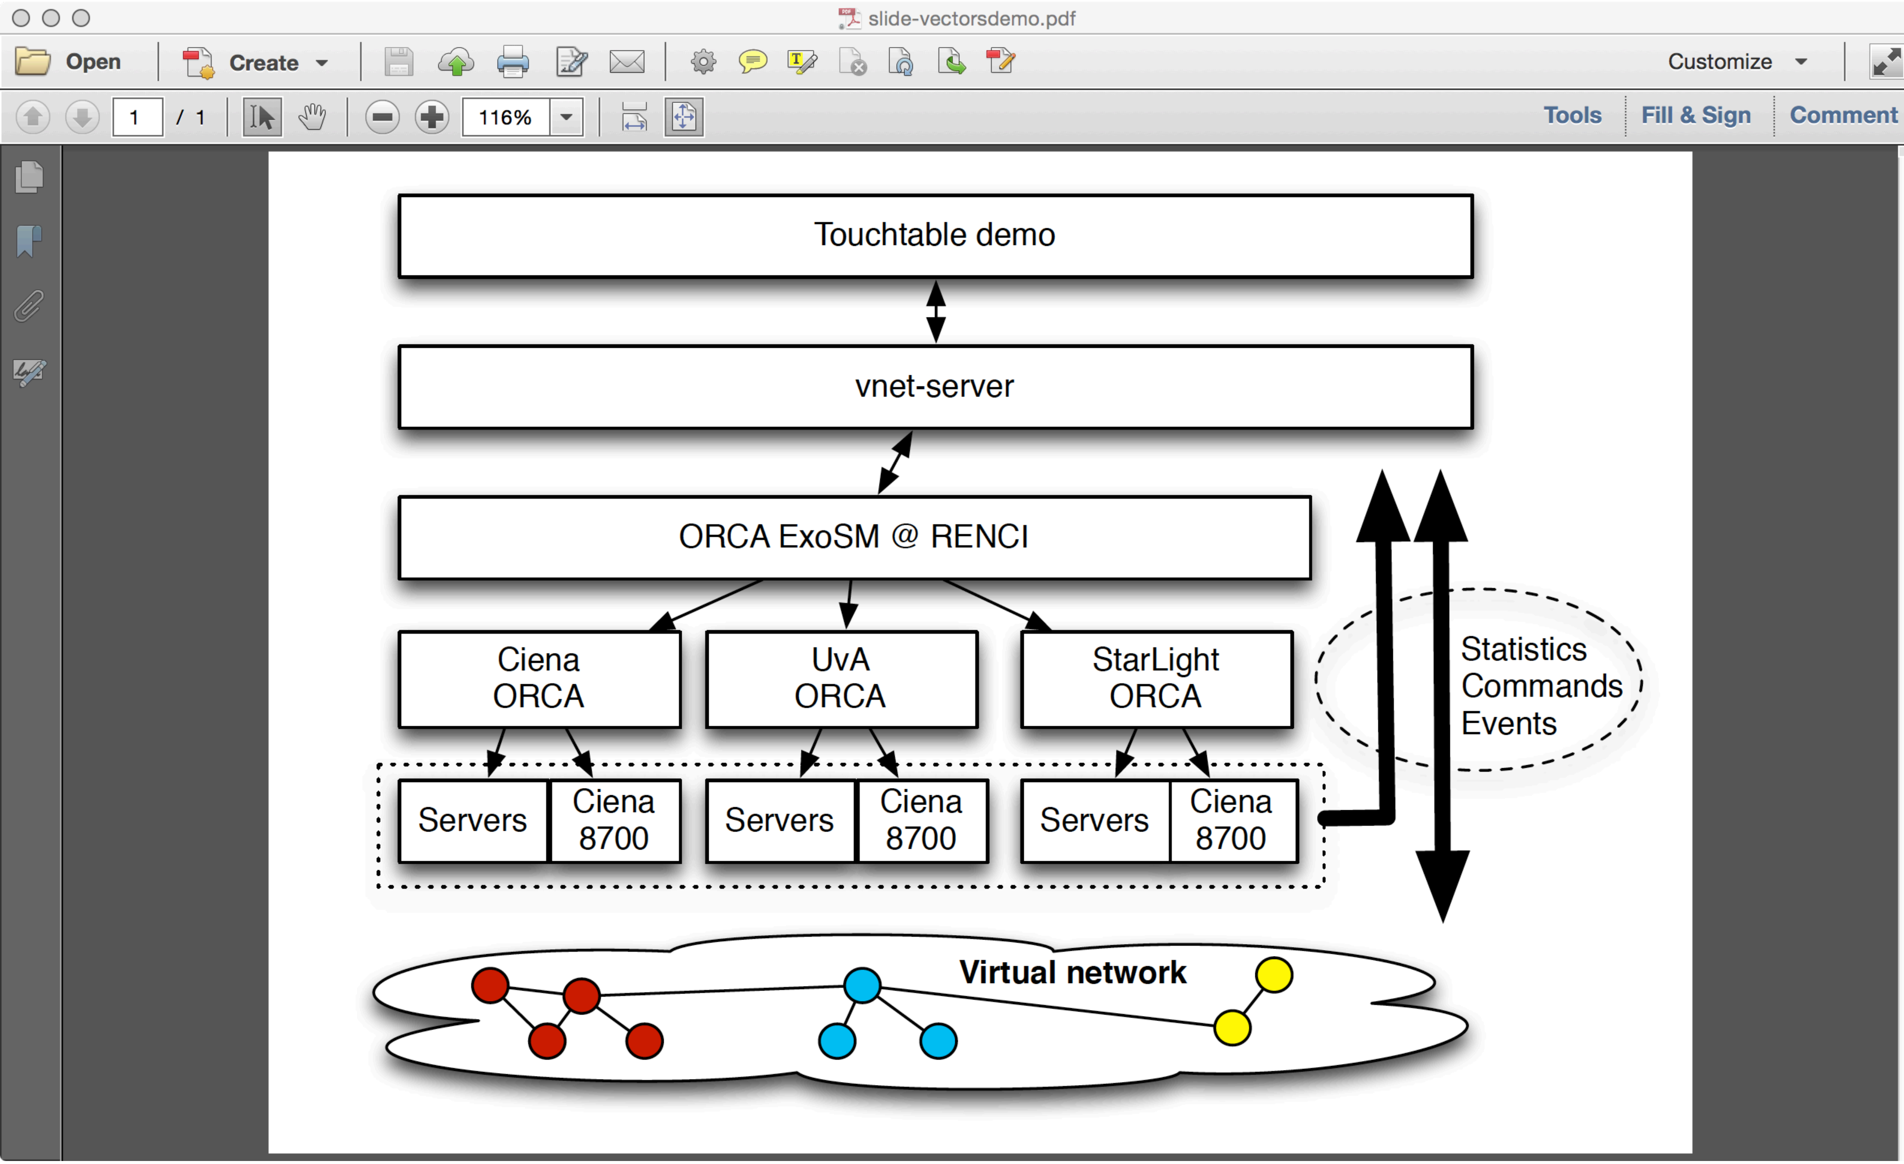
slide-vectorsdemo.pdf
Touchtable demo
vnet-server
ORCA ExoSM @ RENCI
Ciena ORCA
UvA ORCA
StarLight ORCA
Statistics Commands Events
Servers
Ciena 8700
Servers
Ciena 8700
Servers
Ciena 8700
Virtual network


Cloud

Program

16h30	Cees de Laat	Opening and Introduction
16h35	Leon Gommans	Creating a SARNET alliance
16h55	Rodney Wilson	SARNET testbed
17h05	Cees de Laat	Panel introduction
17h06	Chip Elliot	
17h10	Ken Klingenstein	
17h14	Inder Monga	
17h18	Jerry Sobieski	
17h22	panel	discussion
17h29	Next steps	
17h30	end	

More Info

- http://delaat.net/sarnet
- http://delaat.net/pubs/2014-7-3.pdf
- Contact us:
  - delaat@uva.nl
  - l.gommans@uva.nl
  - rwilson@ciena.com
  - Robert.meijer@tno.nl
  - T.M.vanEngers@uva.nl

Hot from the Press Grants!

- EDISON – 554 kEuro
  - Education for Data Intensive Science to Open New science frontiers.
  - Yuri Demchencko
- VRE4EIC – 457 kEuro
  - A Europe-wide Interoperable Virtual Research Environment to Empower Multidisciplinary Research Communities and Accelerate Innovation and Collaboration
  - Zhiming Zhao

ENVRI+ Challenges in system level environmental sciences

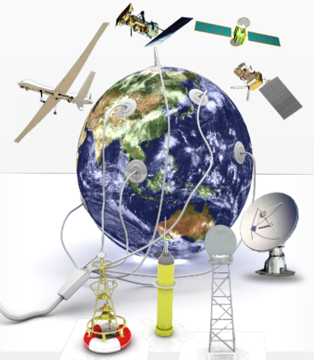

Quality challenge: Quality control **for large quantity** and **nearly real time** sensor data

Collaboration challenge in documenting **data processing workflows** and **sharing among communities.**

Identification/ Citation/ challenge
Identifying and **citing** data objects, and in **publishing** data

Execution challenge in **executing** applications on **distributed computing infrastructures**

Access challenge: **accessing** data from **different sources**

Processing challenge in **combining** different data **processing models**


EISCAT 3D
• Upper space observing


EMSO
• Deep sea observing

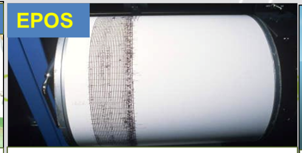
EPOS
• Volcano, earthquake observing

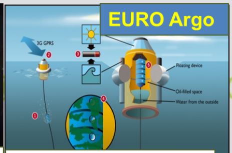
EURO Argo
• Open sea observing

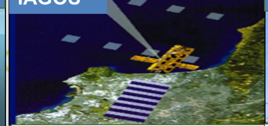
IAGOS
• Aircraft for global observing


ICOS
• Greenhouse observing


SIOS
• Svalbard arctic Earth observing

……

The *Data for Science* theme

Objectives:

1) **Analysis and design interoperable** ICT solutions for prioritized problems current RIs are facing
2) **Develop** the solutions together with the RIs
3) **Validate** the solutions and **deploy** them on e-Infrastructures for future operation

UvA: 840 kEuro

Theme 1: **Technical innovation** WP 1,2,3,4 (5MEuro)	Theme2: **data for science** WP 5,6,7,8,9 (5M Euro)	Theme 3: **Access:** WP 10, 11 (1M)	Theme 4: **Social aspects** WP 12, 13, 14 (1M)	Theme 5: **Knowledge transfer** WP 15, 16 (1M)	Theme 6: **Dissemination** WP 17, 18 (1M)	Theme 7: **management** WP19 (1M)

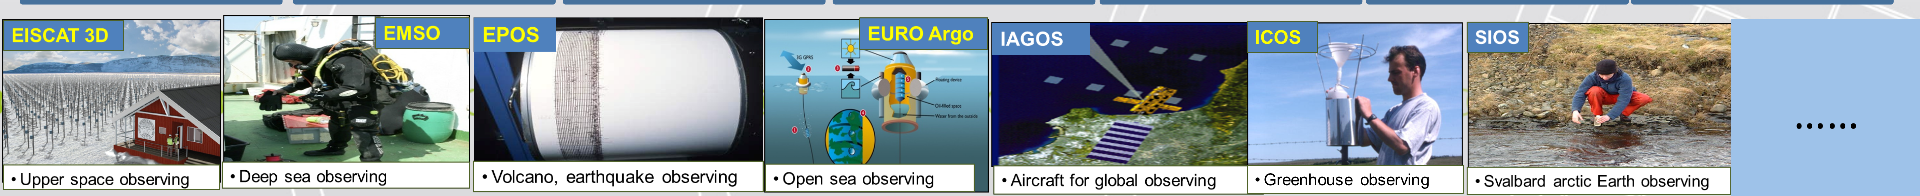

Principles for the solutions

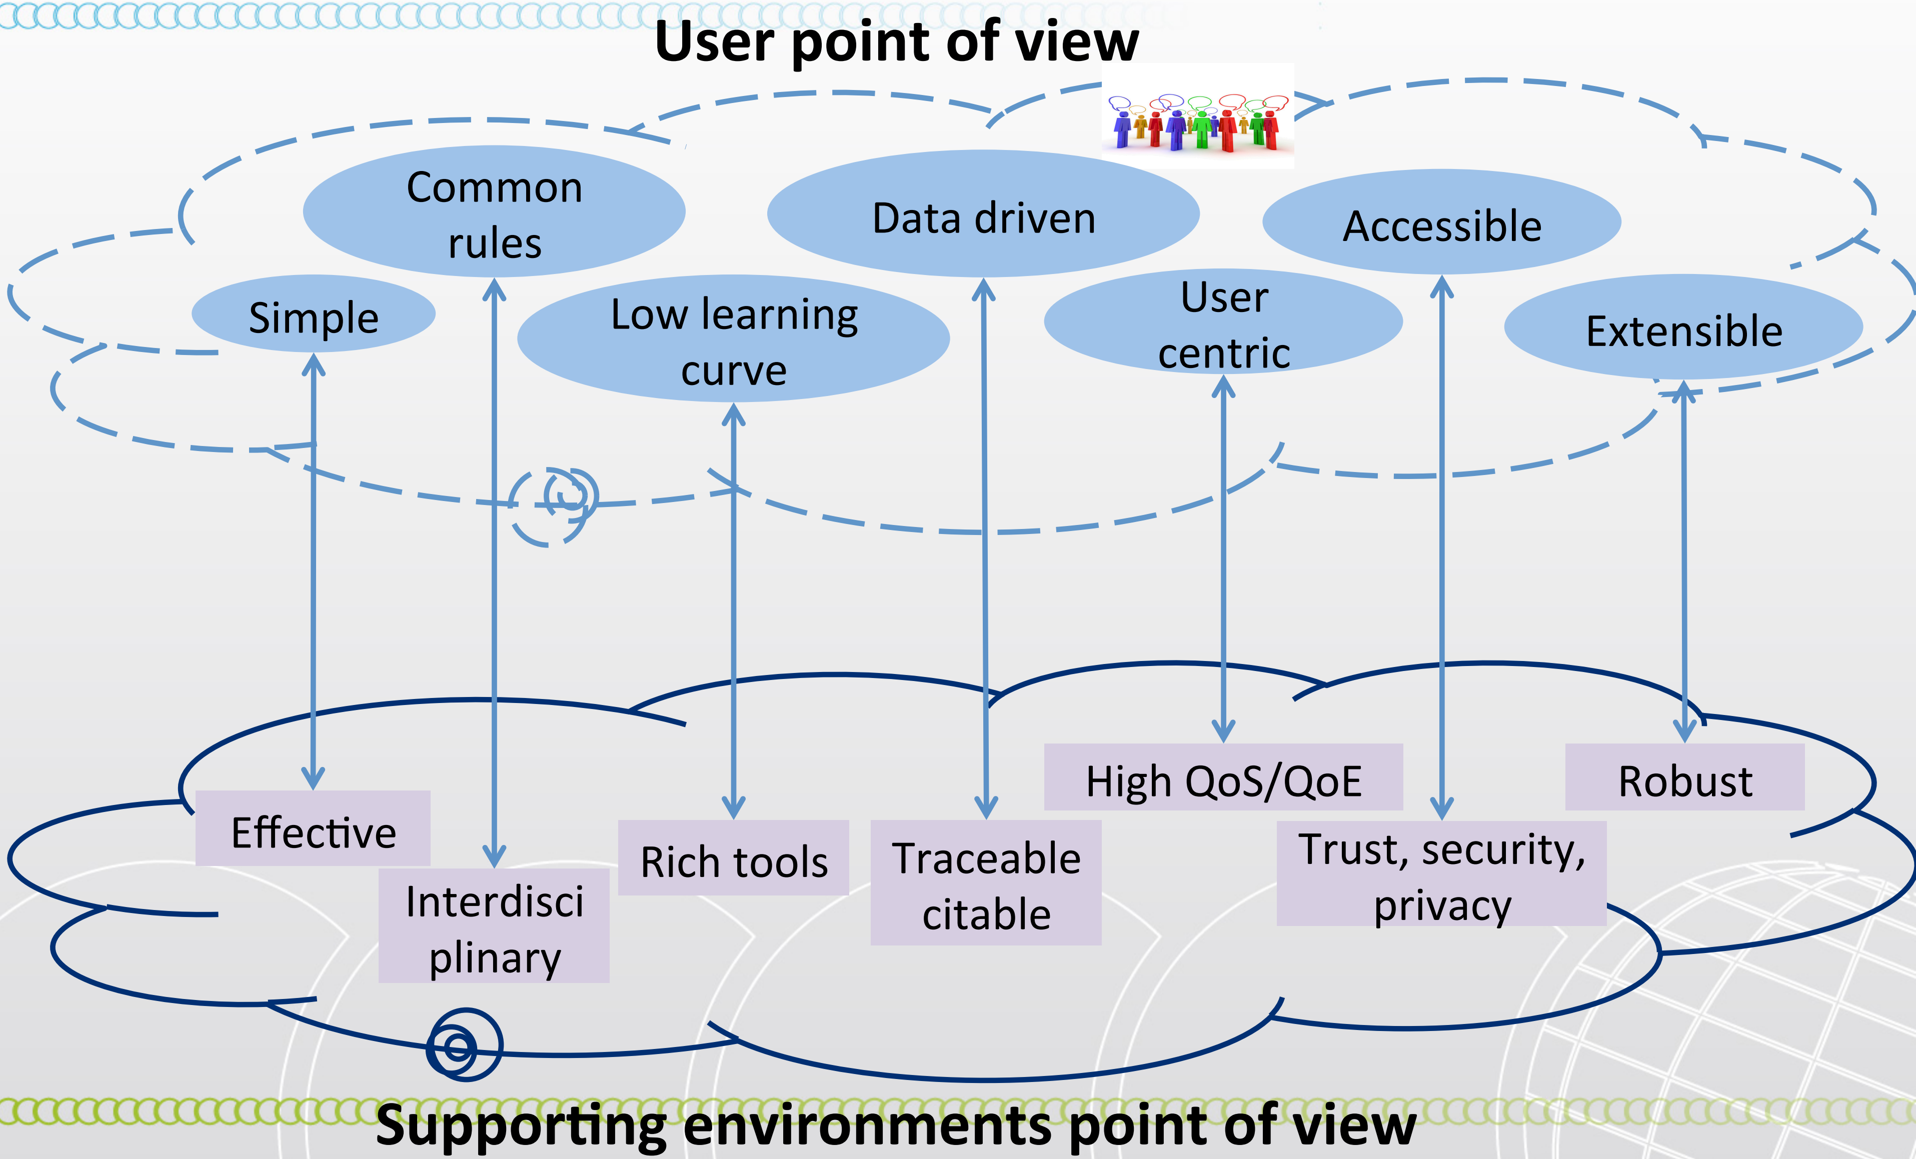

Kick-off meeting, 10-Feb-2015, Amsterdam, the Netherlands

EU H2020-ICT-2014

SWITCH

Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud applications

PI - Dr. Zhiming Zhao
PI - Prof. Cees de Laat
System and network engineering
University of Amsterdam

3 MEuro – 780 kEuro

Universiteit van Amsterdam · Wellness Telecom · Cardiff University · Univerza v Ljubljani · Beia · MOG

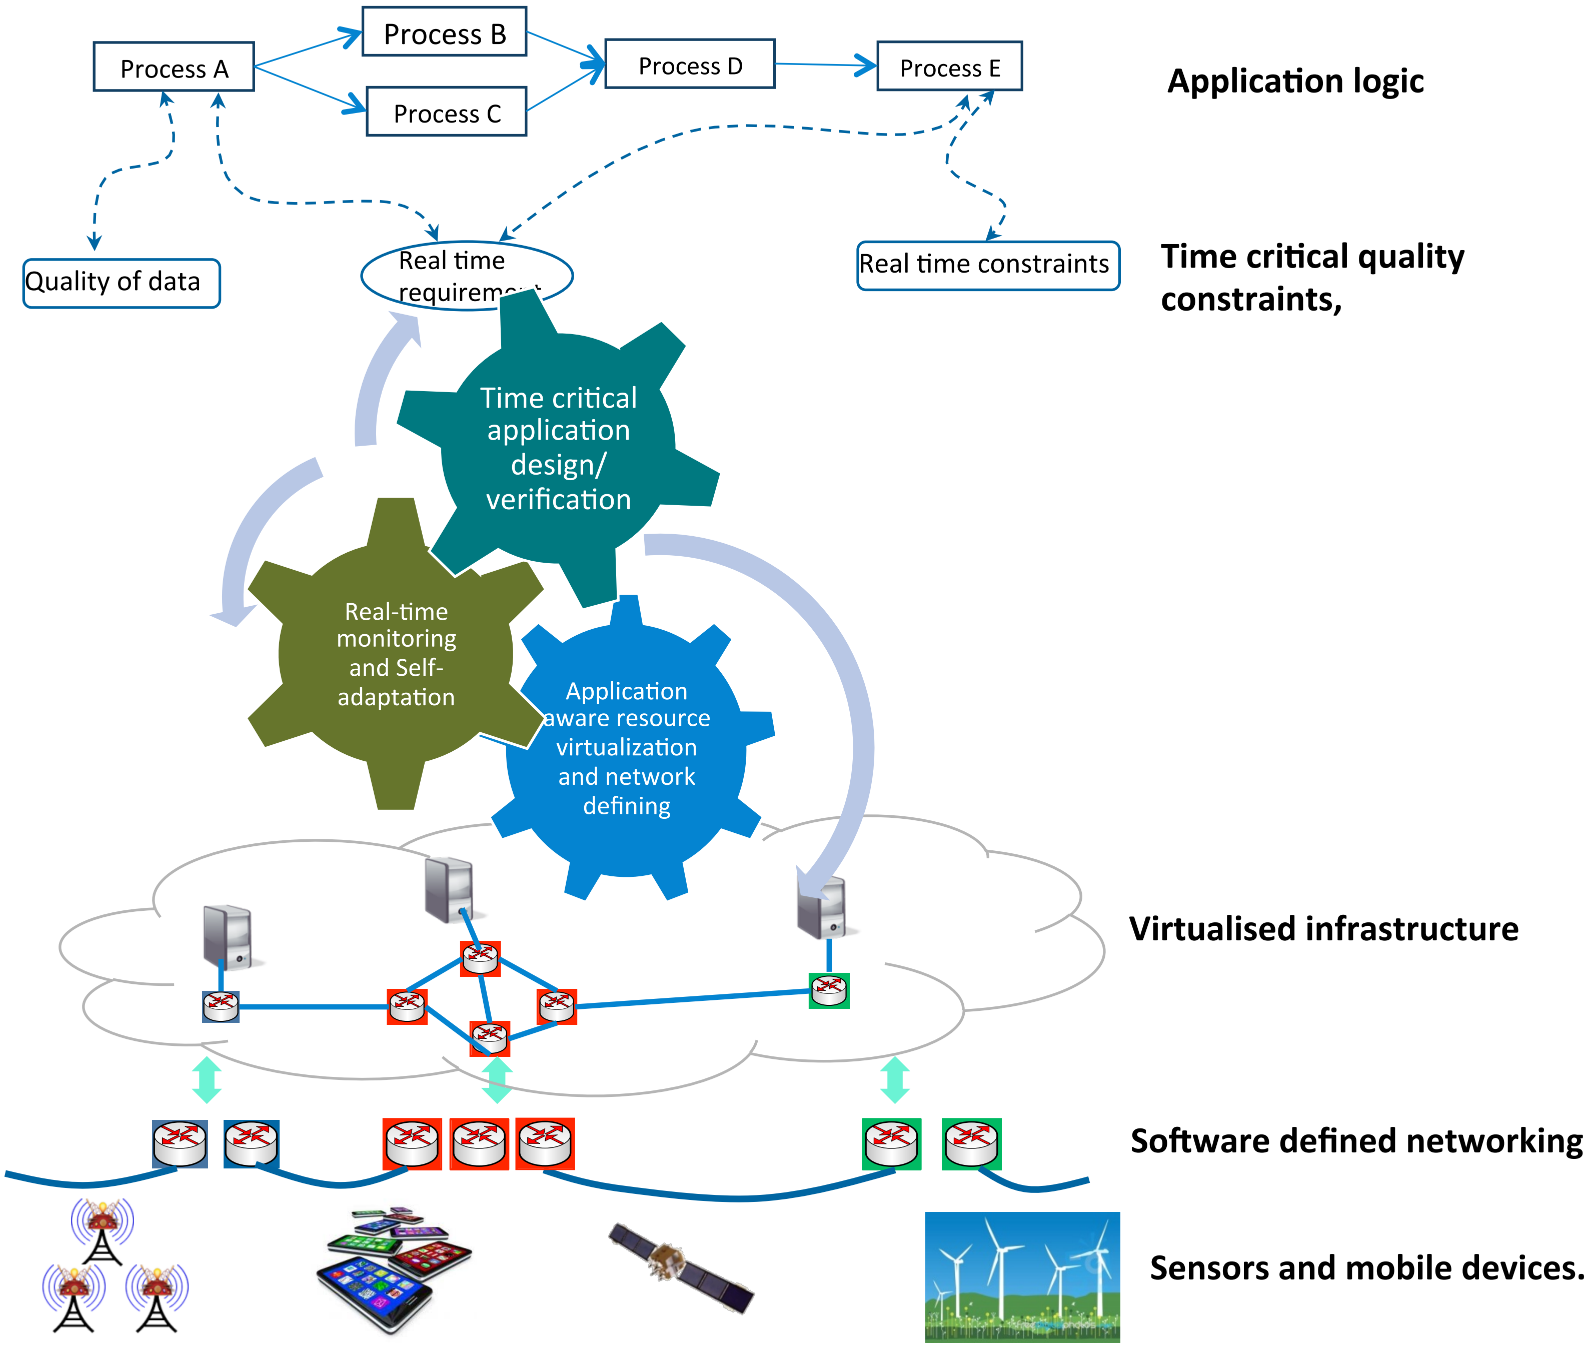

Process A
Process B
Process C
Process D
Process E
Application logic
Quality of data
Real time
Real time constraints
Time critical quality constraints,
Time critical application design/ verification
Real-time monitoring and Self-adaptation
Application aware resource virtualization and network defining
Virtualised infrastructure
Software defined networking
Sensors and mobile devices.

Other Grants

- CYCLONE – Intercloud – 450 kEuro
  - Yuri Demchencko
- Information complexity – 500 kEuro
  - Pieter Adriaans & Frank van Harmelen
- iDapt – DSL extra functional behavior 500 kEuro
  - Andy Pimentel, Rafael Poss
- Extra - Software Defined Computing ☺ - 500 kEuro
  - Ana Varbanescu

The constant factor in our field is Change!

The 50 years it took Physicists to find one particle, the Higgs,
we came from:

Assembler, Fortran, COBOL, VM, RSX11, Unix, c, Pascal, SmallTalk, DECnet, VMS, TCP/IP, c++, Internet, WWW, ATM, Semantic Web, Photonic networks, Google, Grid, Phyton, FaceBook, Twitter, Cloud, SDN, Data^3, App's

to:

DDOS attacks destroying Banks and BitCoins!

Conclusion:

Need for Safe, Smart, Resilient Sustainable Infrastructure.

Questions?

Arie Taal
Paola Grosso Ana Oprescu
Cees de Laat
Marc Makkes Ralph Koning
Bas Terwijn
Leon Gommans Fahimeh Alizadeh
Pieter Adriaans Cosmin Dumitru Karst Koymans
Yuri Demchenko Rob Meijer Karel van der Veldt
Reggie Cushing
Rudolf Strijkers Miroslav Zivkovic
Jan Sipke van der Veen
Naod Duga Jebessa Spiros Koulouzis Hao Zhu Sander Klous
Jaap van Ginkel Guido van 't Noordende
Jeroen van der Ham
Mikolaj Baranowski Steven de Rooij
Ngo Tong Canh Souley Madougou Paul Klint
Adianto Wibisono
Magiel Bruntink
Zhiming Zhao
Anna Varbanescu Marijke Kaat
Niels Sijm Hans Dijkman Gerben de Vries
Adam Belloum Arno Bakker Marian Bubak
Daniel Romao Erik-Jan Bos
Peter Bloem

http://delaat.net

http://sne.science.uva.nl

http://www.os3.nl/

http://sne.science.uva.nl/openlab/

http://pire.opensciencedatacloud.org

http://staff.science.uva.nl/~delaat/pire/

https://rd-alliance.org

http://envri.eu

SNE
Data Science Research Center Amsterdam
Supported by: COMMIT/ ENVRI