Smart Cyber Infrastructure for Big Data Processing

Cees de Laat

followed by mini talks from Paola Grosso & Zhiming Zhao

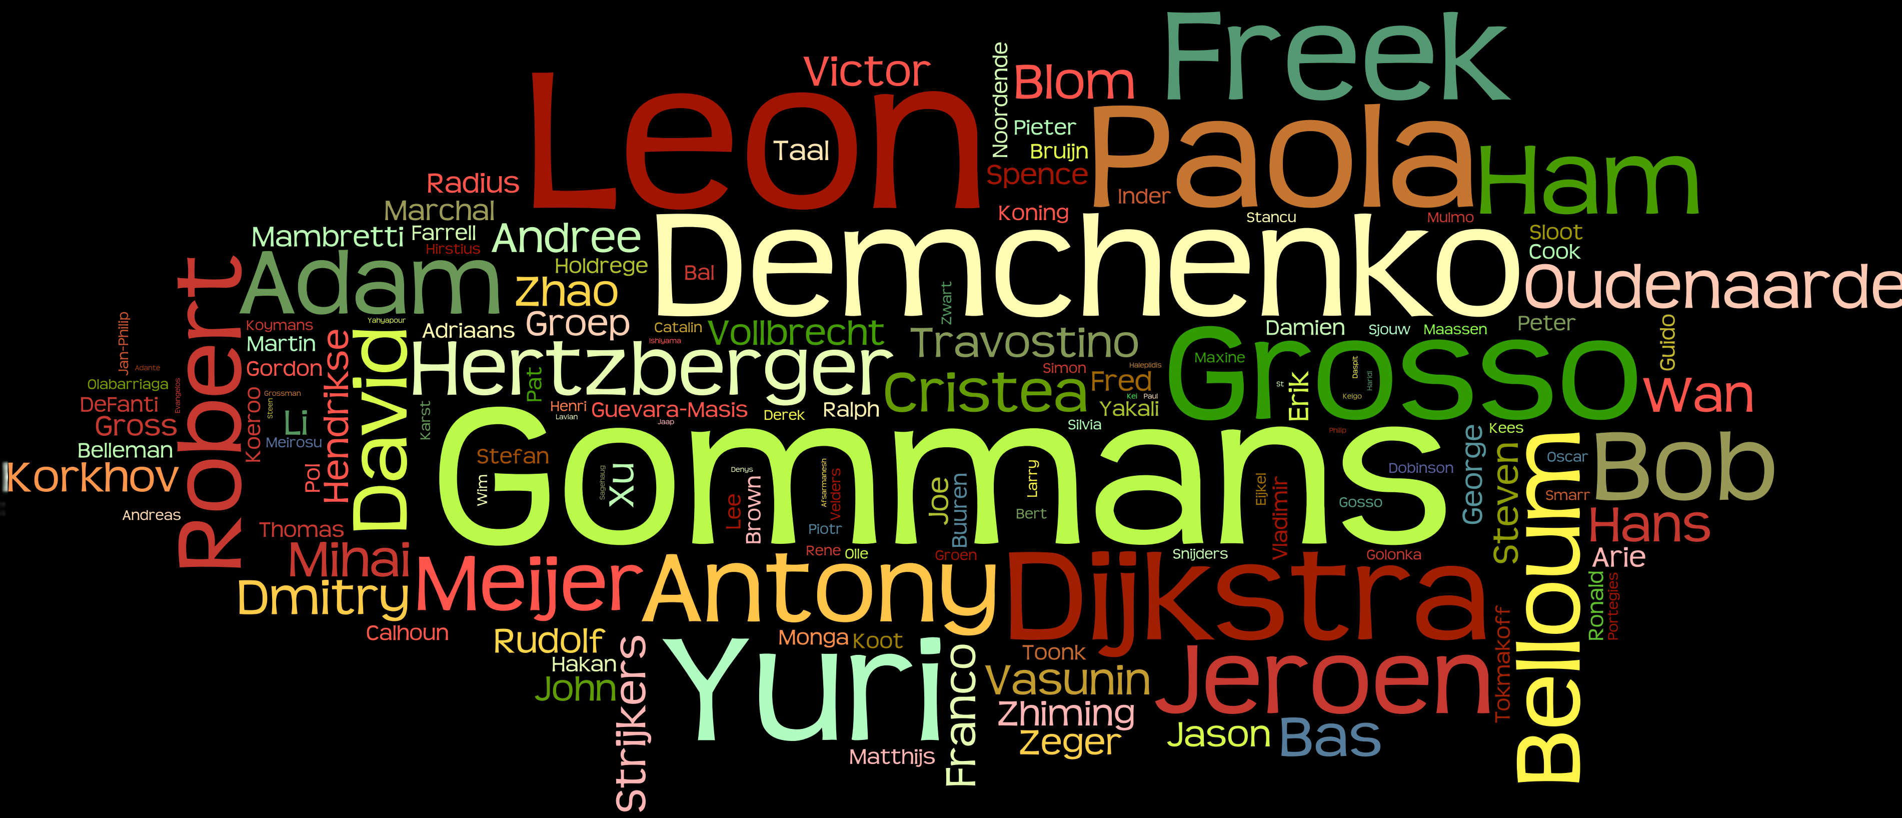

Science Faculty @ UvA

Informatics Institute

- AMLAB: Machine Learning (Prof. dr. M. Welling)
- FCN: Federated Collaborative Networks (Prof. dr. H. Afsarmanesh)
- ILPS: Information and Language Processing Systems (Prof. dr. M. de Rijke)
- ISIS: Intelligent Sensory Information Systems (Prof. dr. ir. A.W.M. Smeulders)
- CSL: Computational Science Laboratory (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)

SNE - 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*
  - *Policy, Trust, Anonymity, Privacy, Integrity*
- *Sustainability*
  - *Greening infrastructure, Awareness*
- *Resilience*
  - *Failures, Disasters, Systems under attack*

SNE - Staffing

Group leader: *prof. C. de Laat*

Deputy group leaders: *dr. Andy Pimentel, dr. Paola Grosso*

- 1 full prof (CdL)
- 2 part time professors
- 2 endowed professors
- *2 senior researchers*
- 1 associate prof
- 4 assistant professors
- ~12 postdoc's
- *About 15 phd students*
- ~10 guests

- *Yearly turnover ~ 3,5 MEuro*

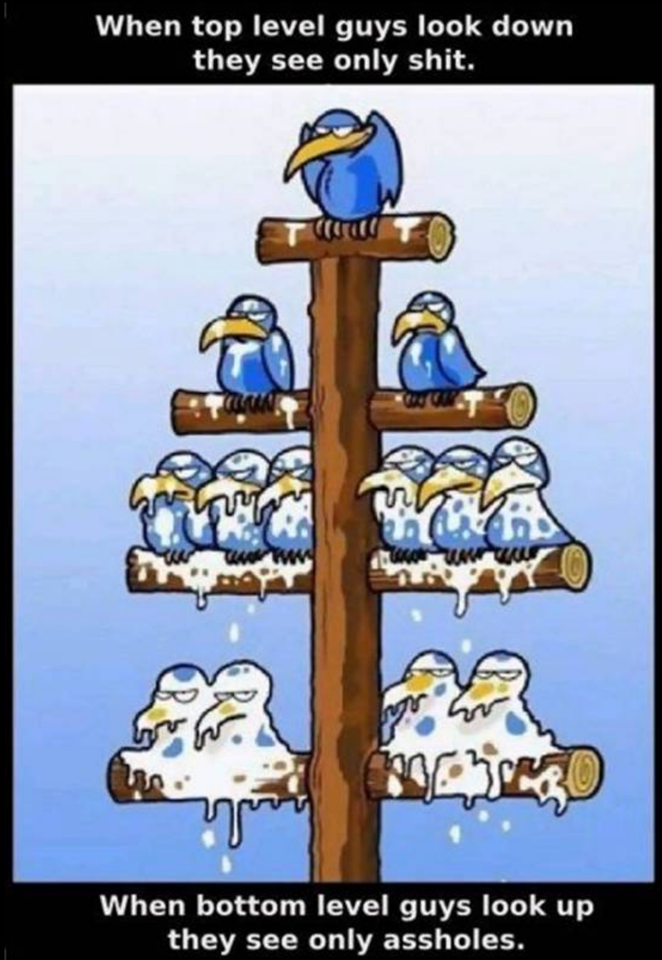

SNE - Mission

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*
  - *Policy, Trust, Anonymity, Privacy, Integrity*
- *Sustainability*
  - *Greening infrastructure, Awareness*
- *Resilience*
  - *Failures, Disasters, Systems under attack*

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**

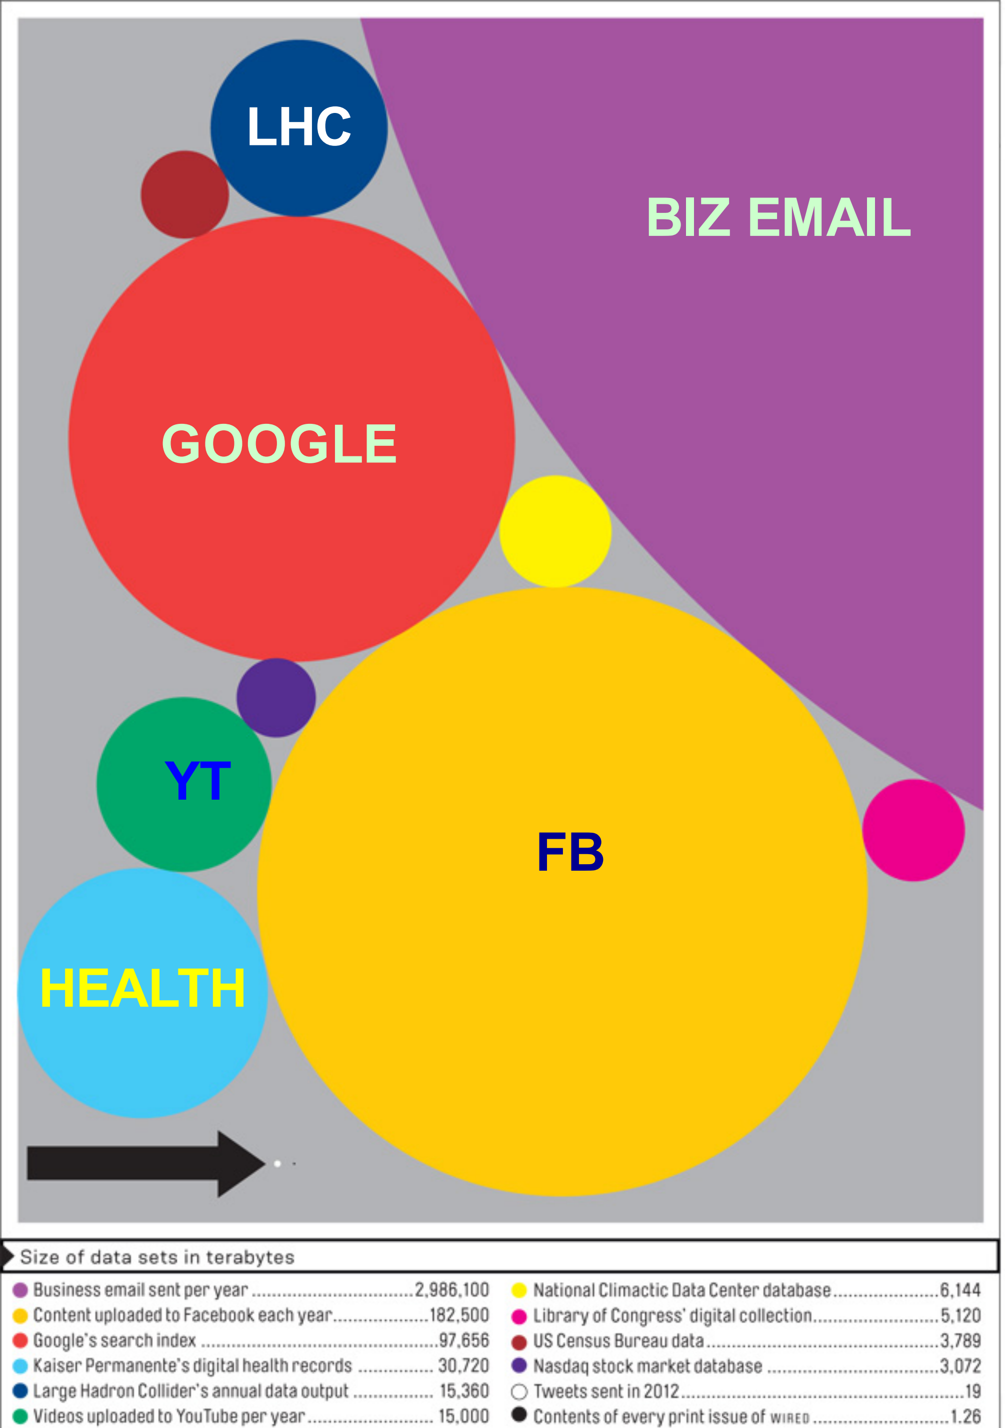

There is always a bigger fish

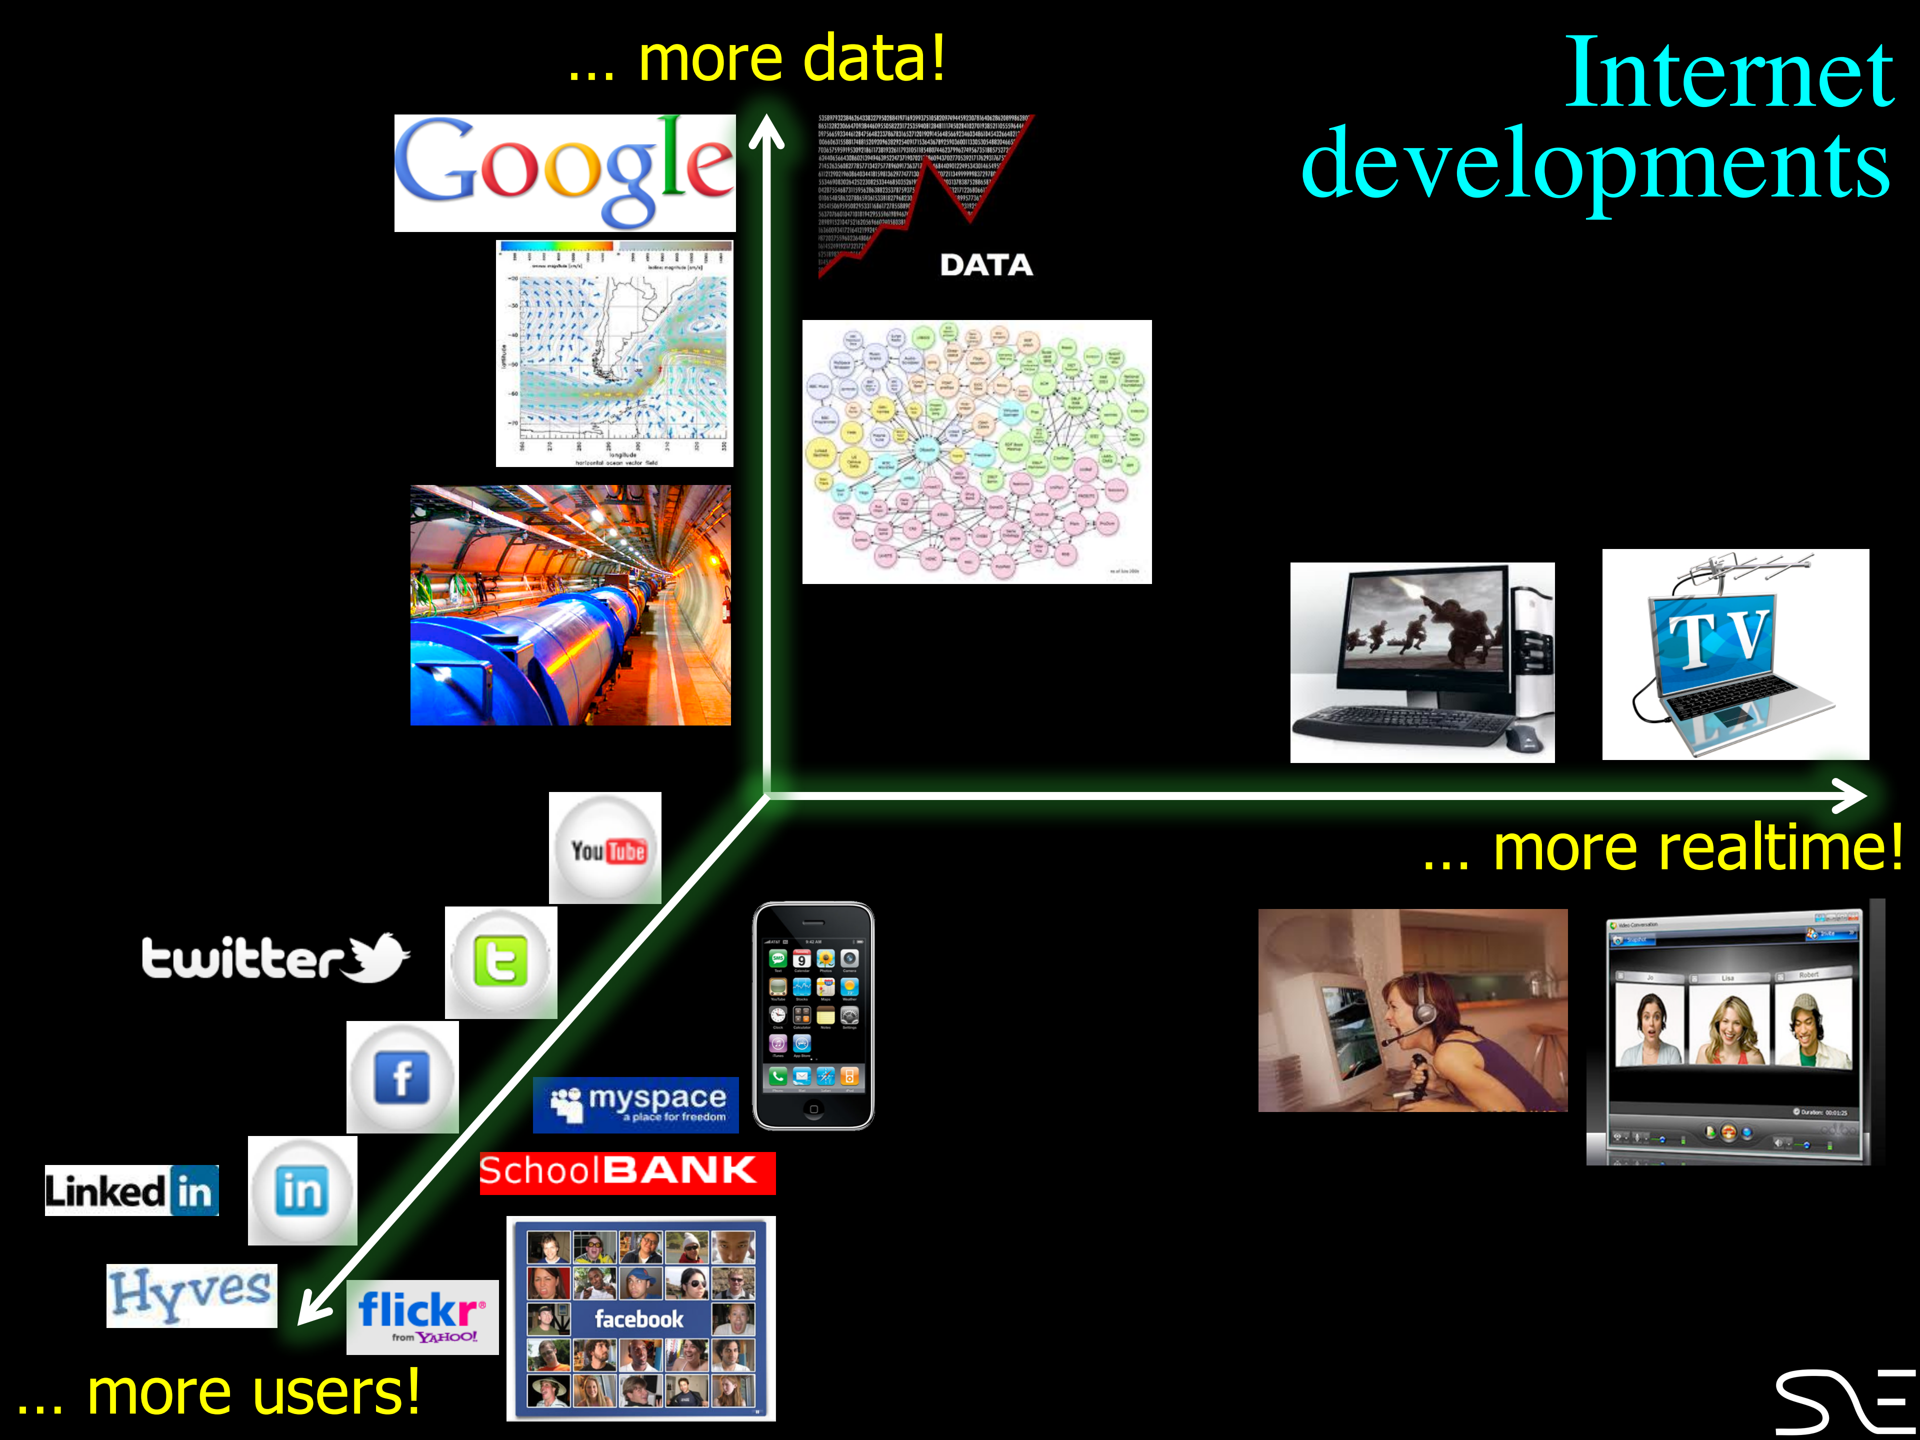
Internet developments
... more data!
Google
DATA
TV
... more realtime!
You Tube
twitter
myspace
a place for freedom
SchoolBANK
Linked in
Hyves
flickr
from YAHOO!
facebook
... more users!

Internet developments

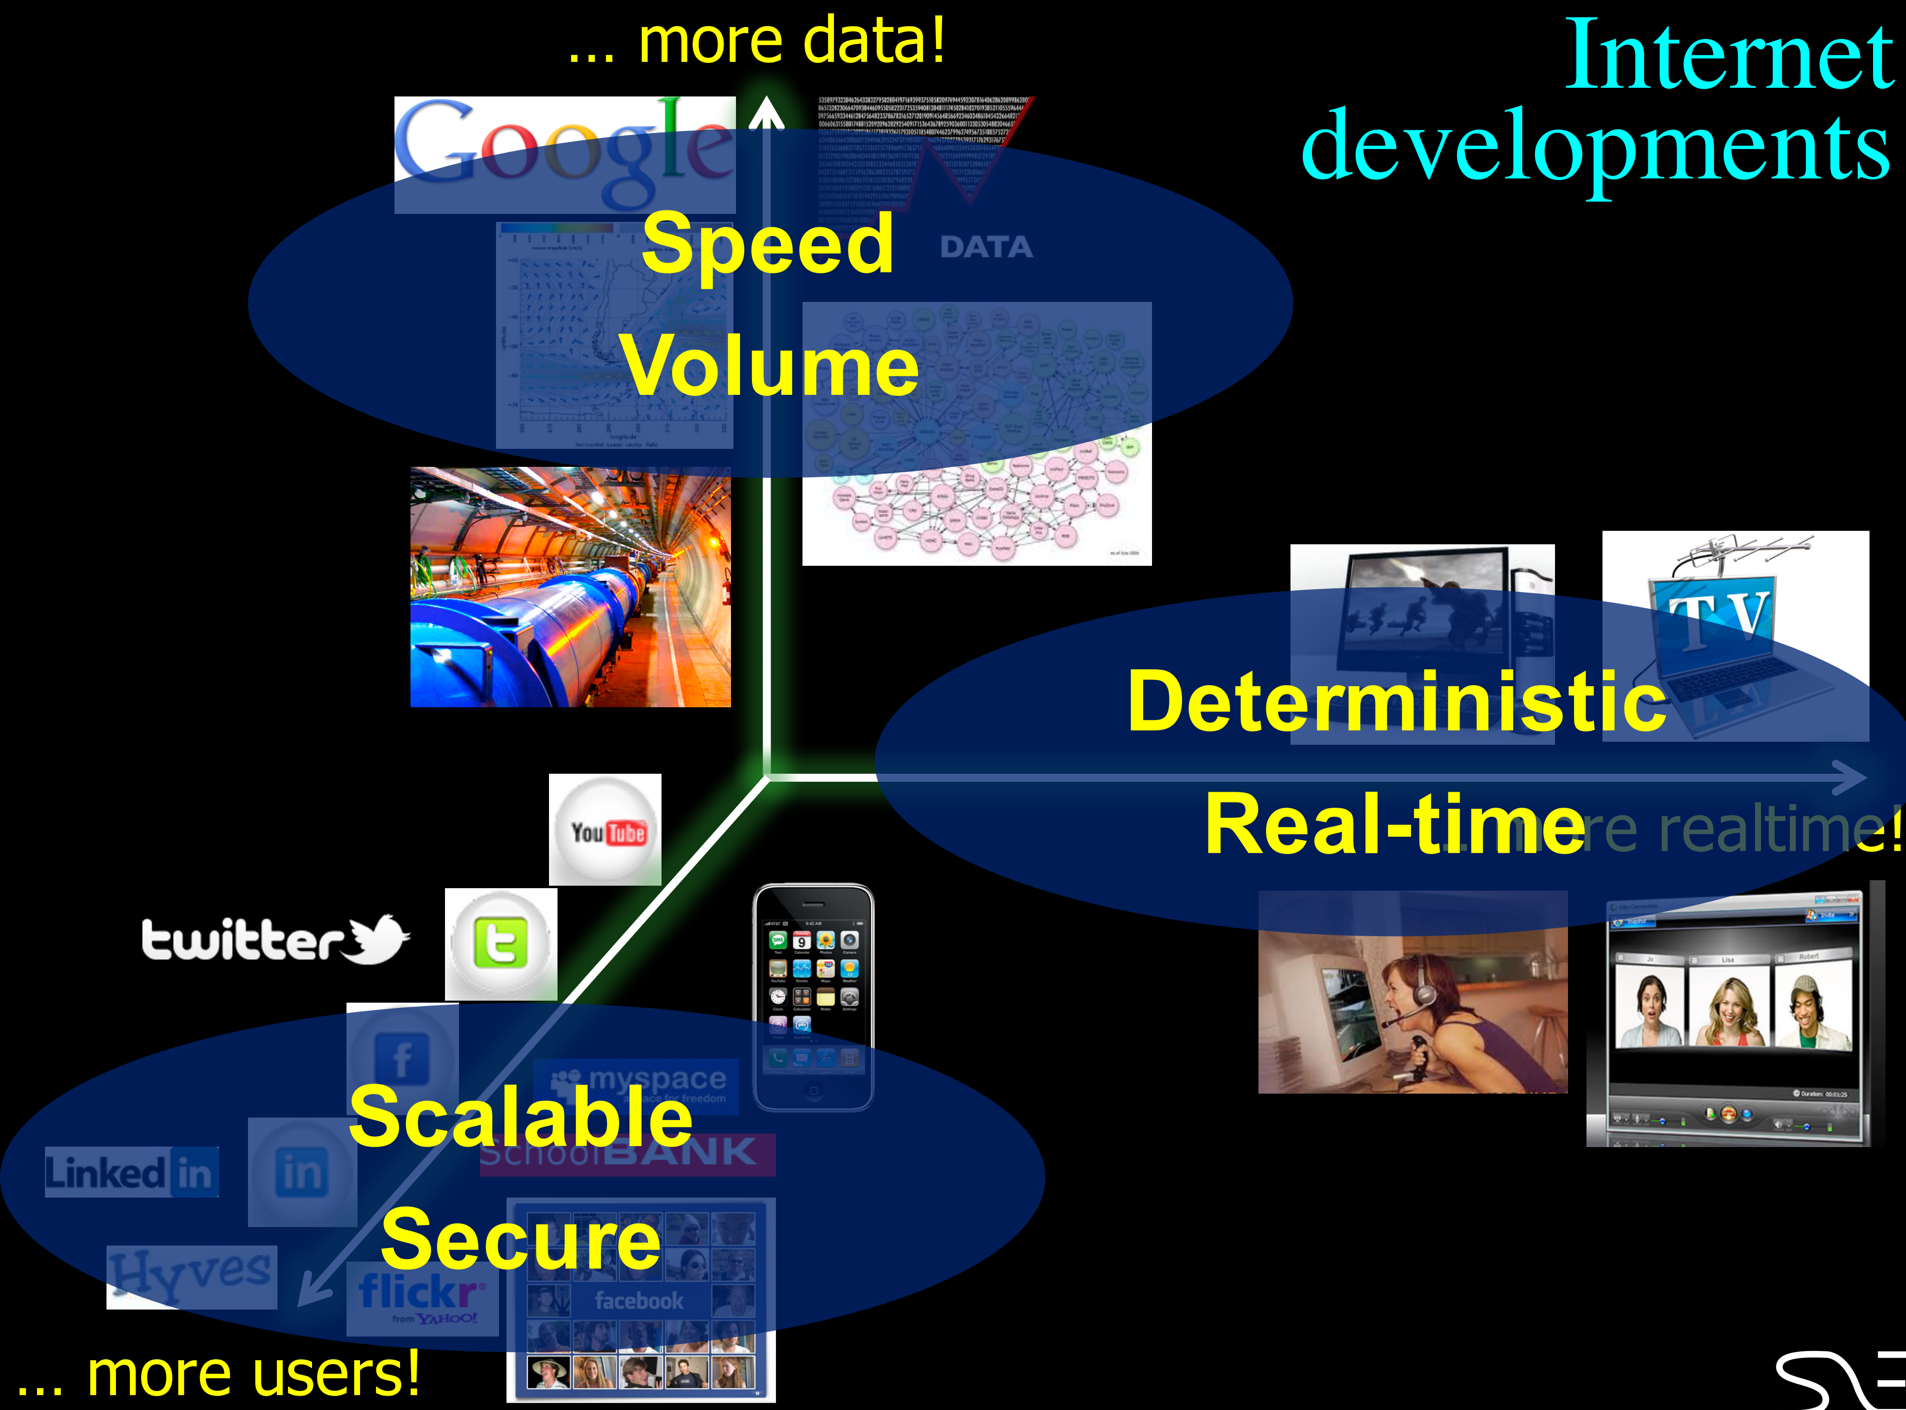

Multiple colors / Fiber

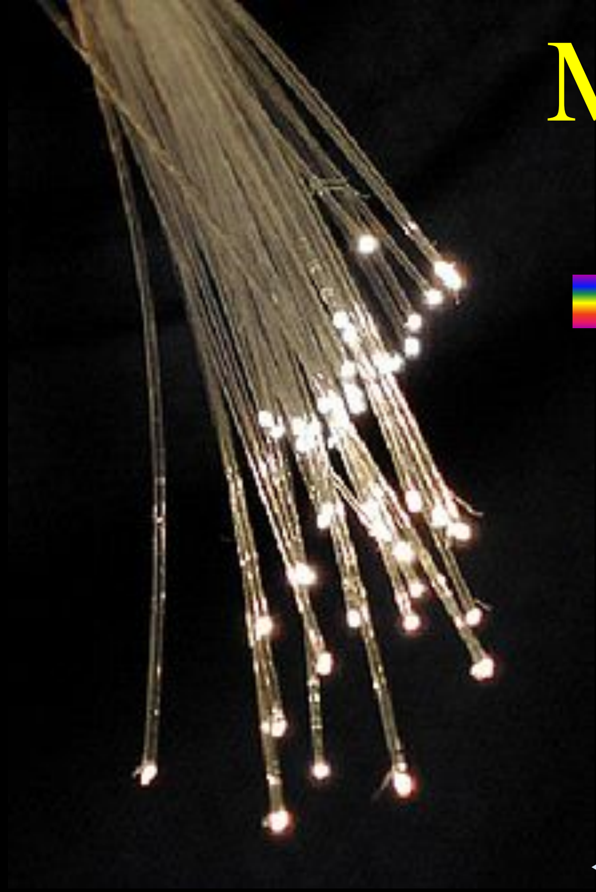

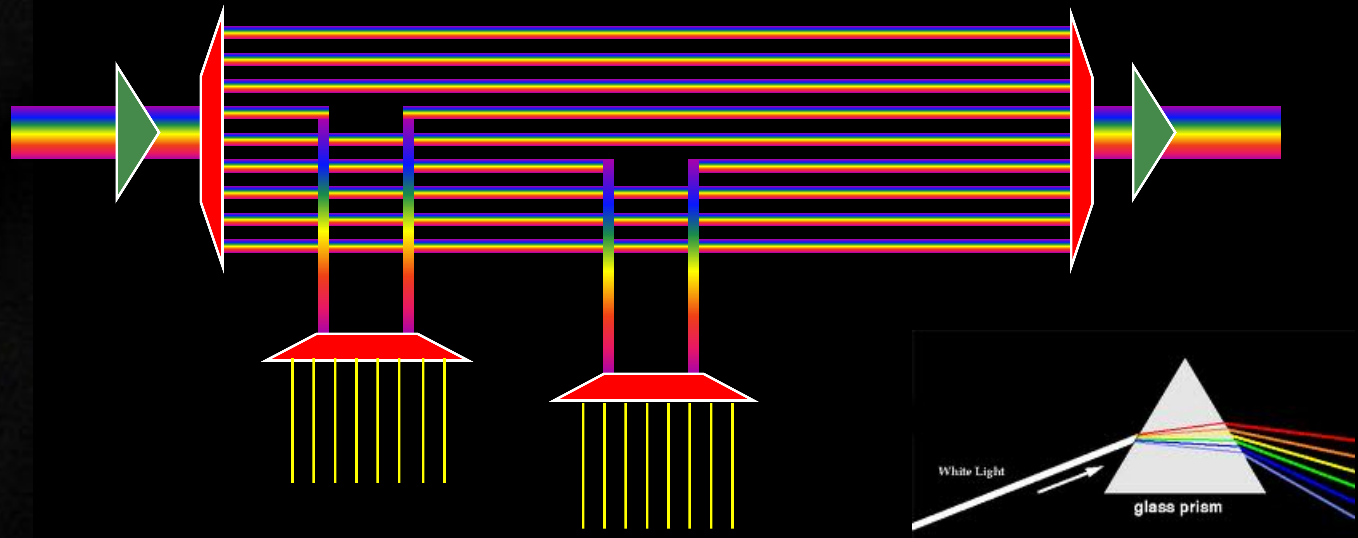

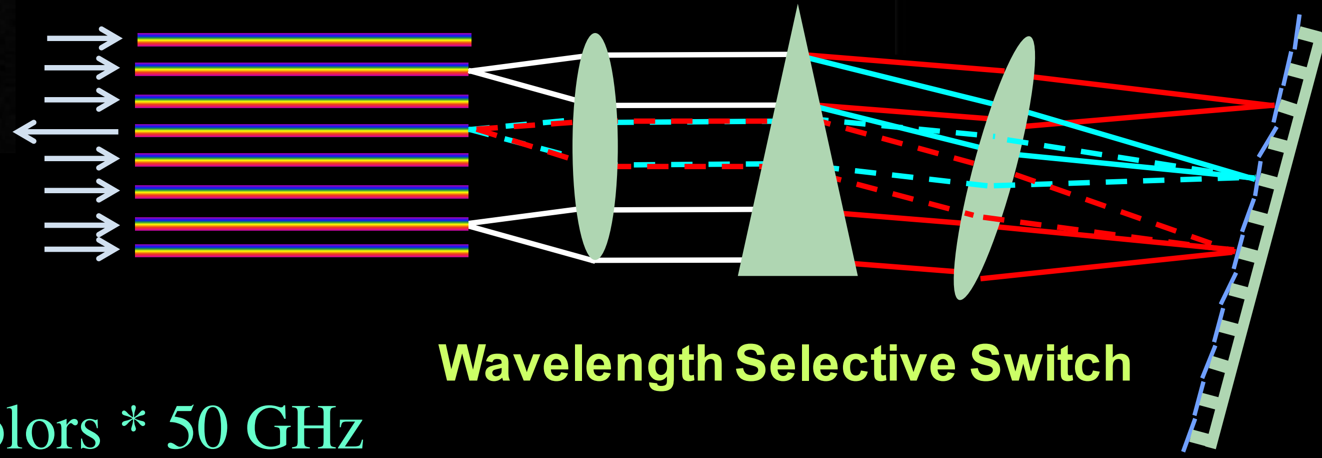

Wavelength Selective Switch

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!

GPU cards are distruptive!

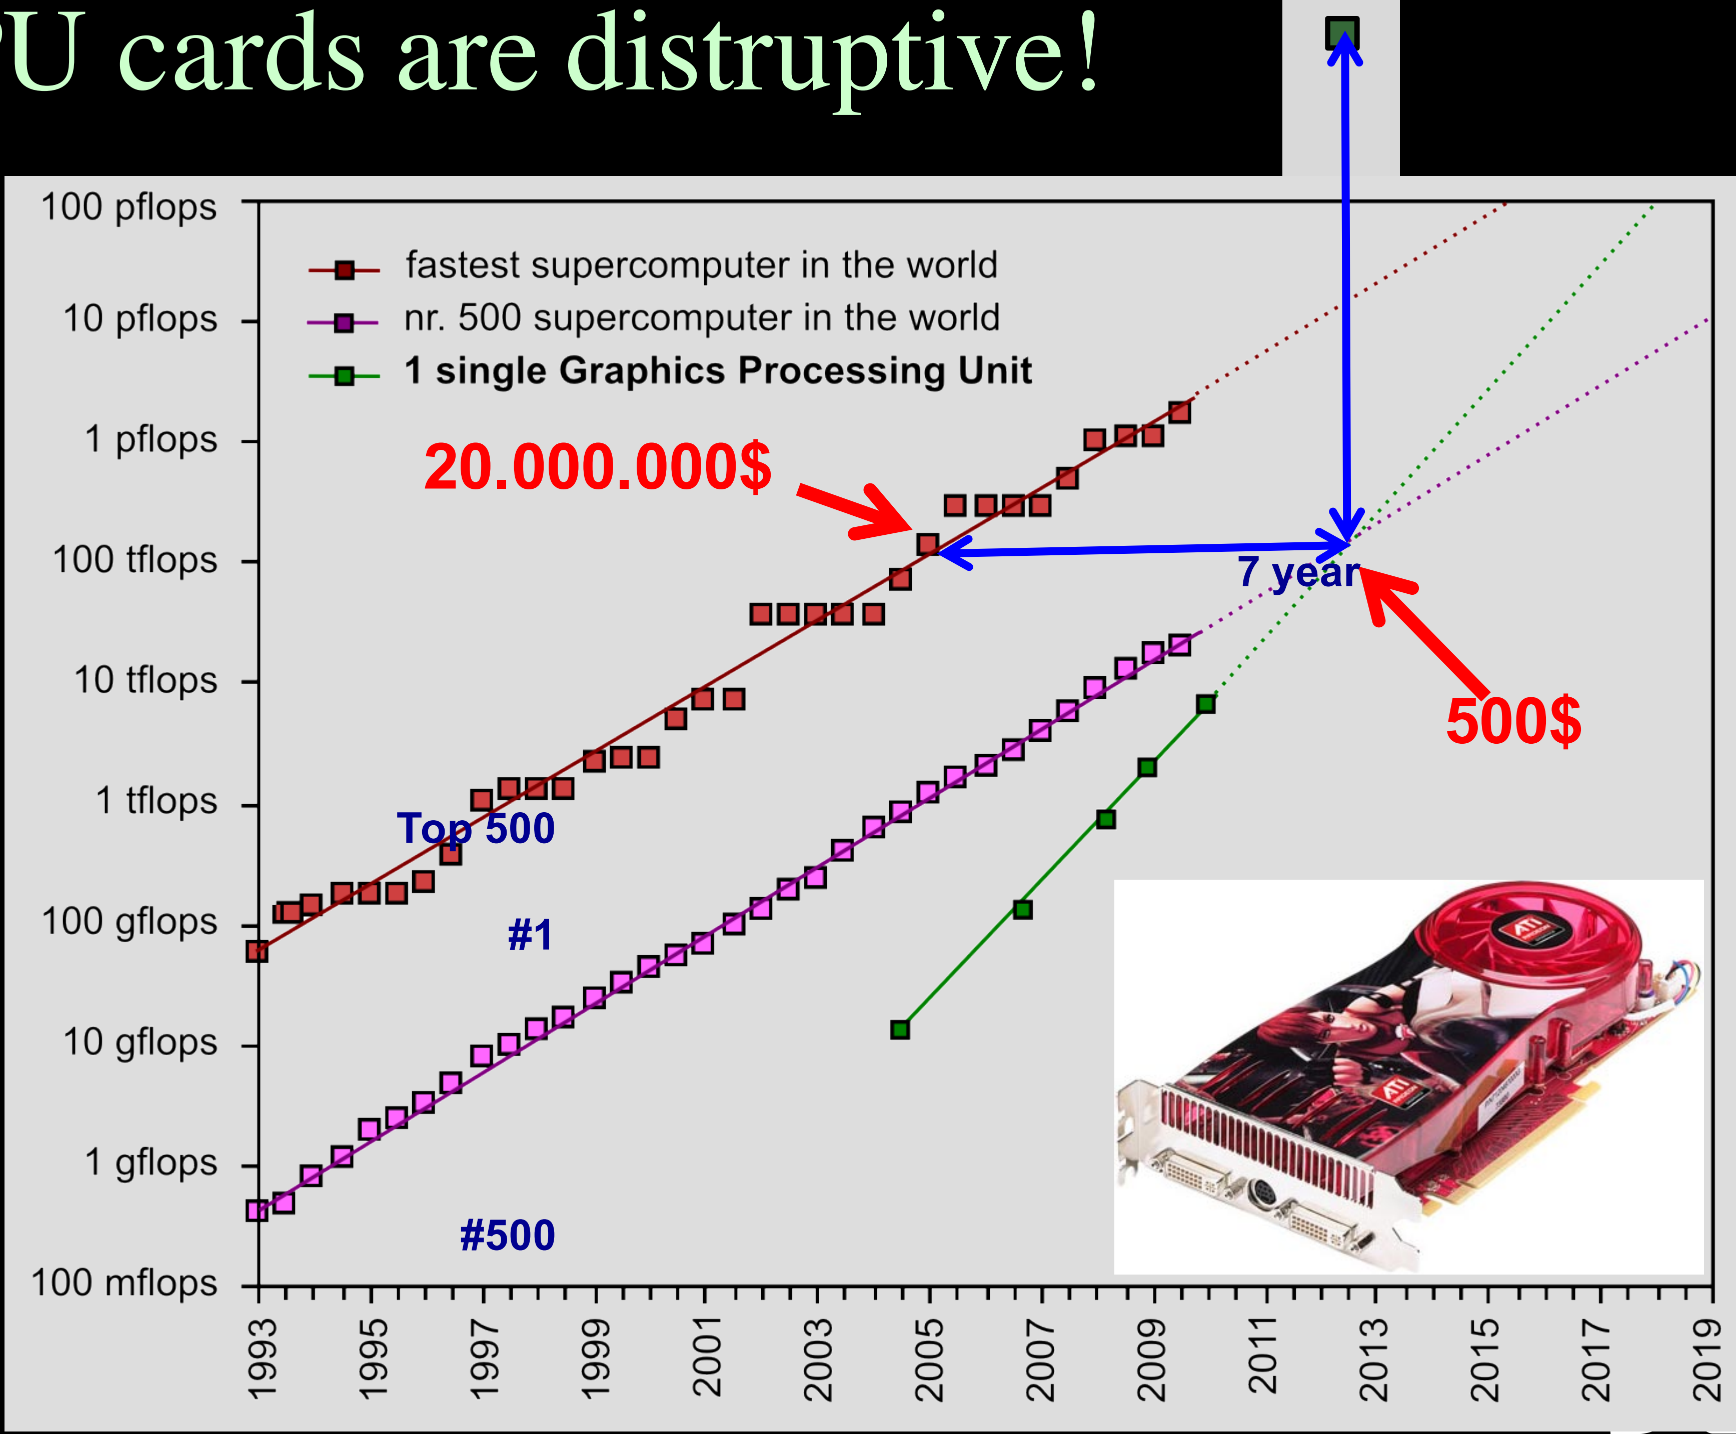

Reliable and Safe!

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.

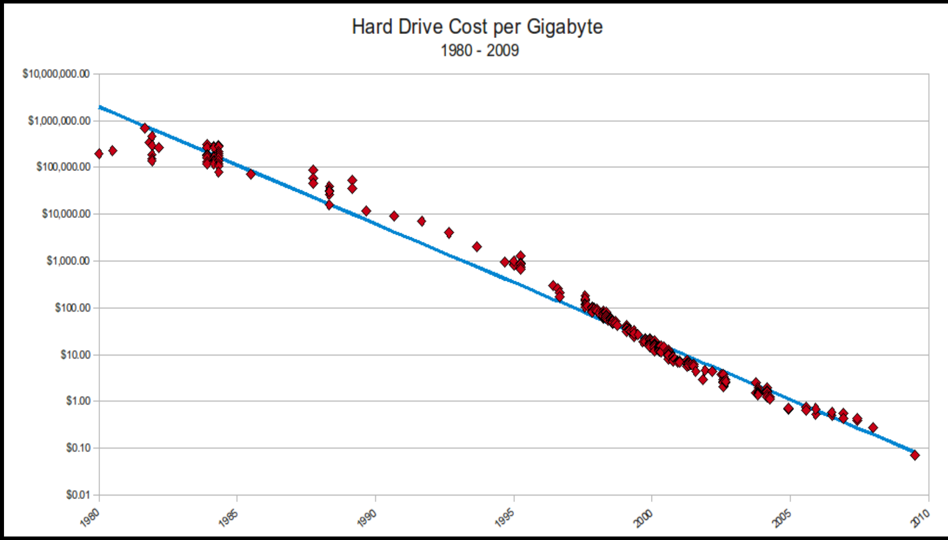

The hardware and software that allow all our systems to operate is becoming bigger and more complex all the time, and the capacity of networks and data storage is increasing by leaps and bounds.

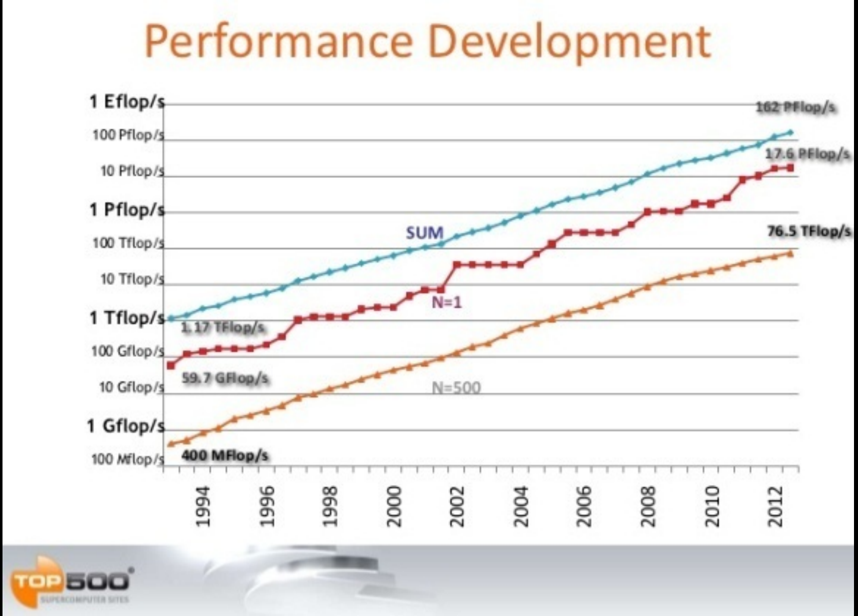

We will soon reach the limits of what is currently feasible and controllable.

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.

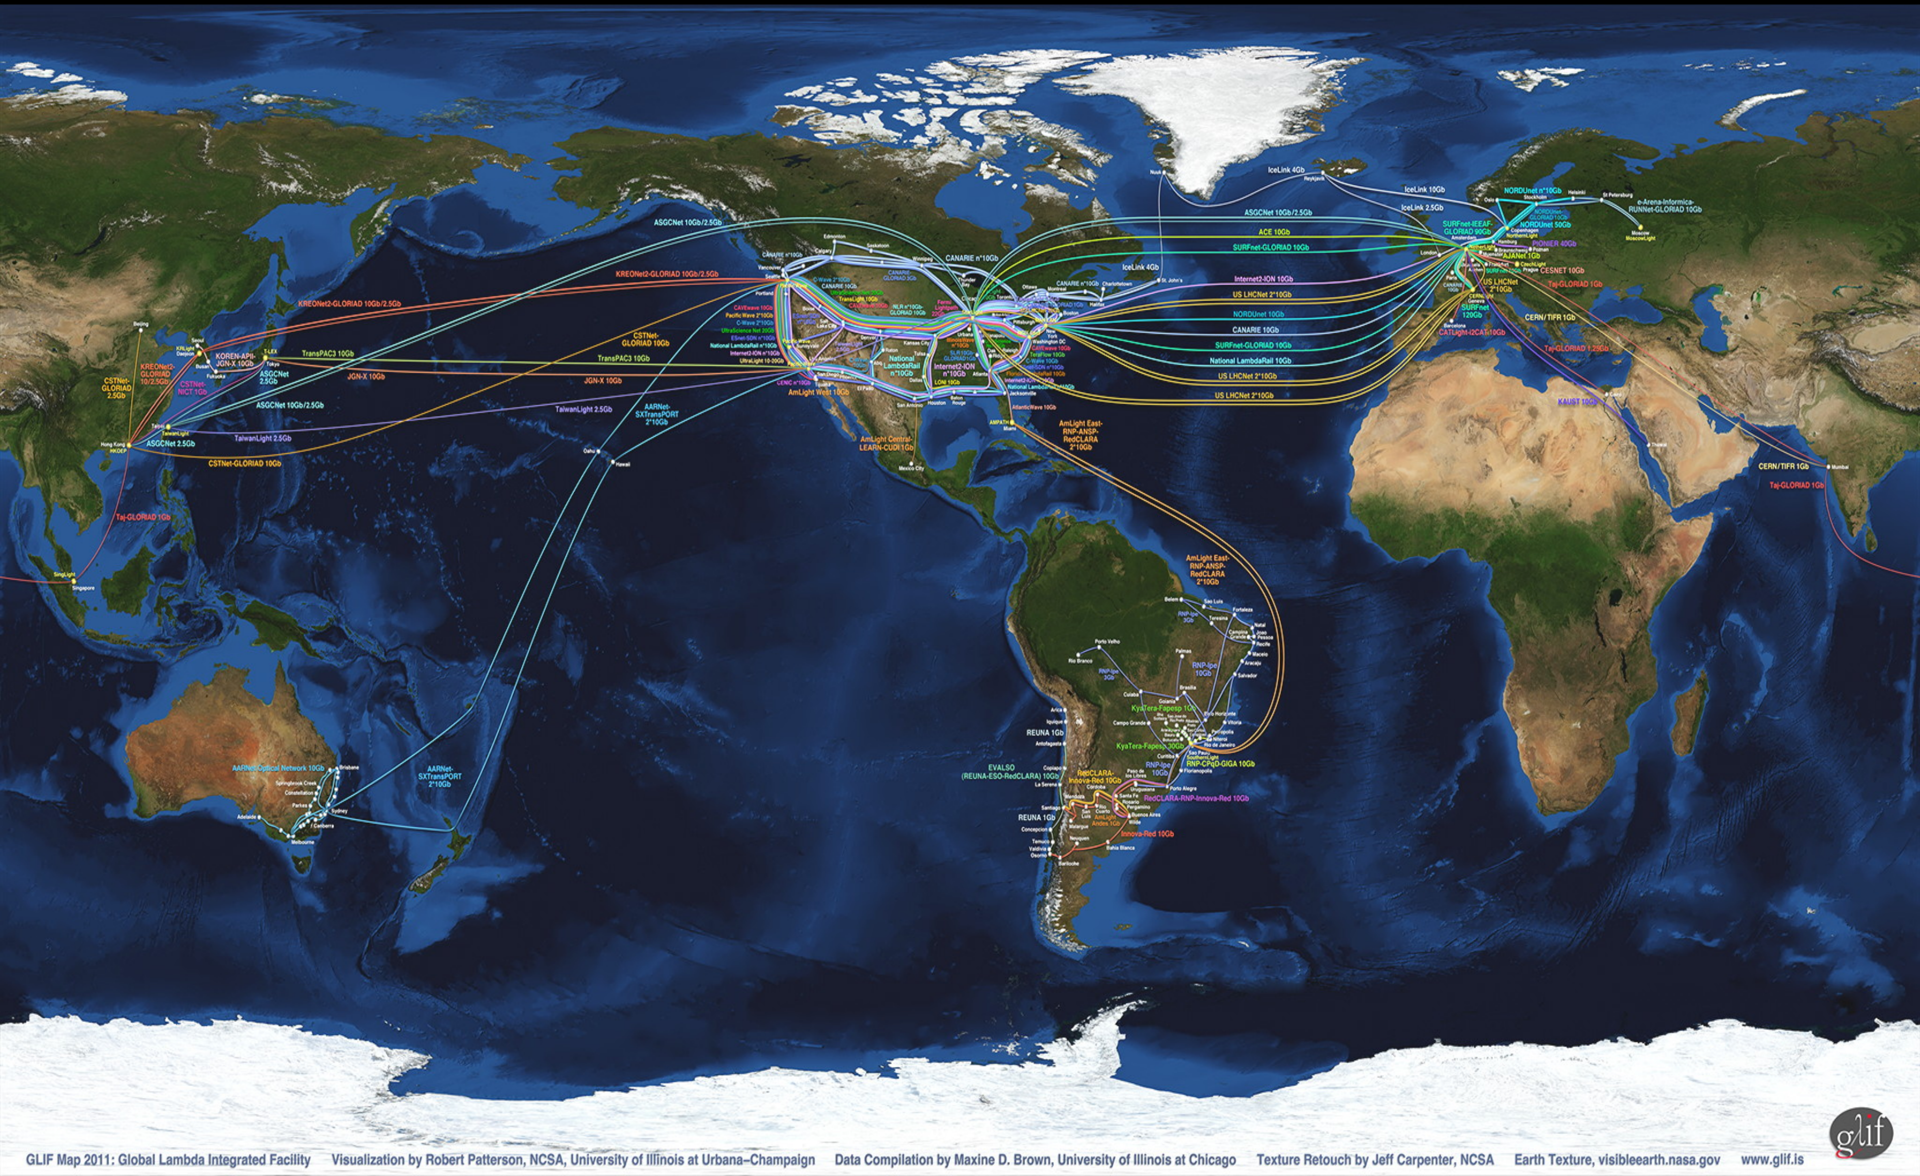

Amsterdam is a major hub in The GLIF

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.

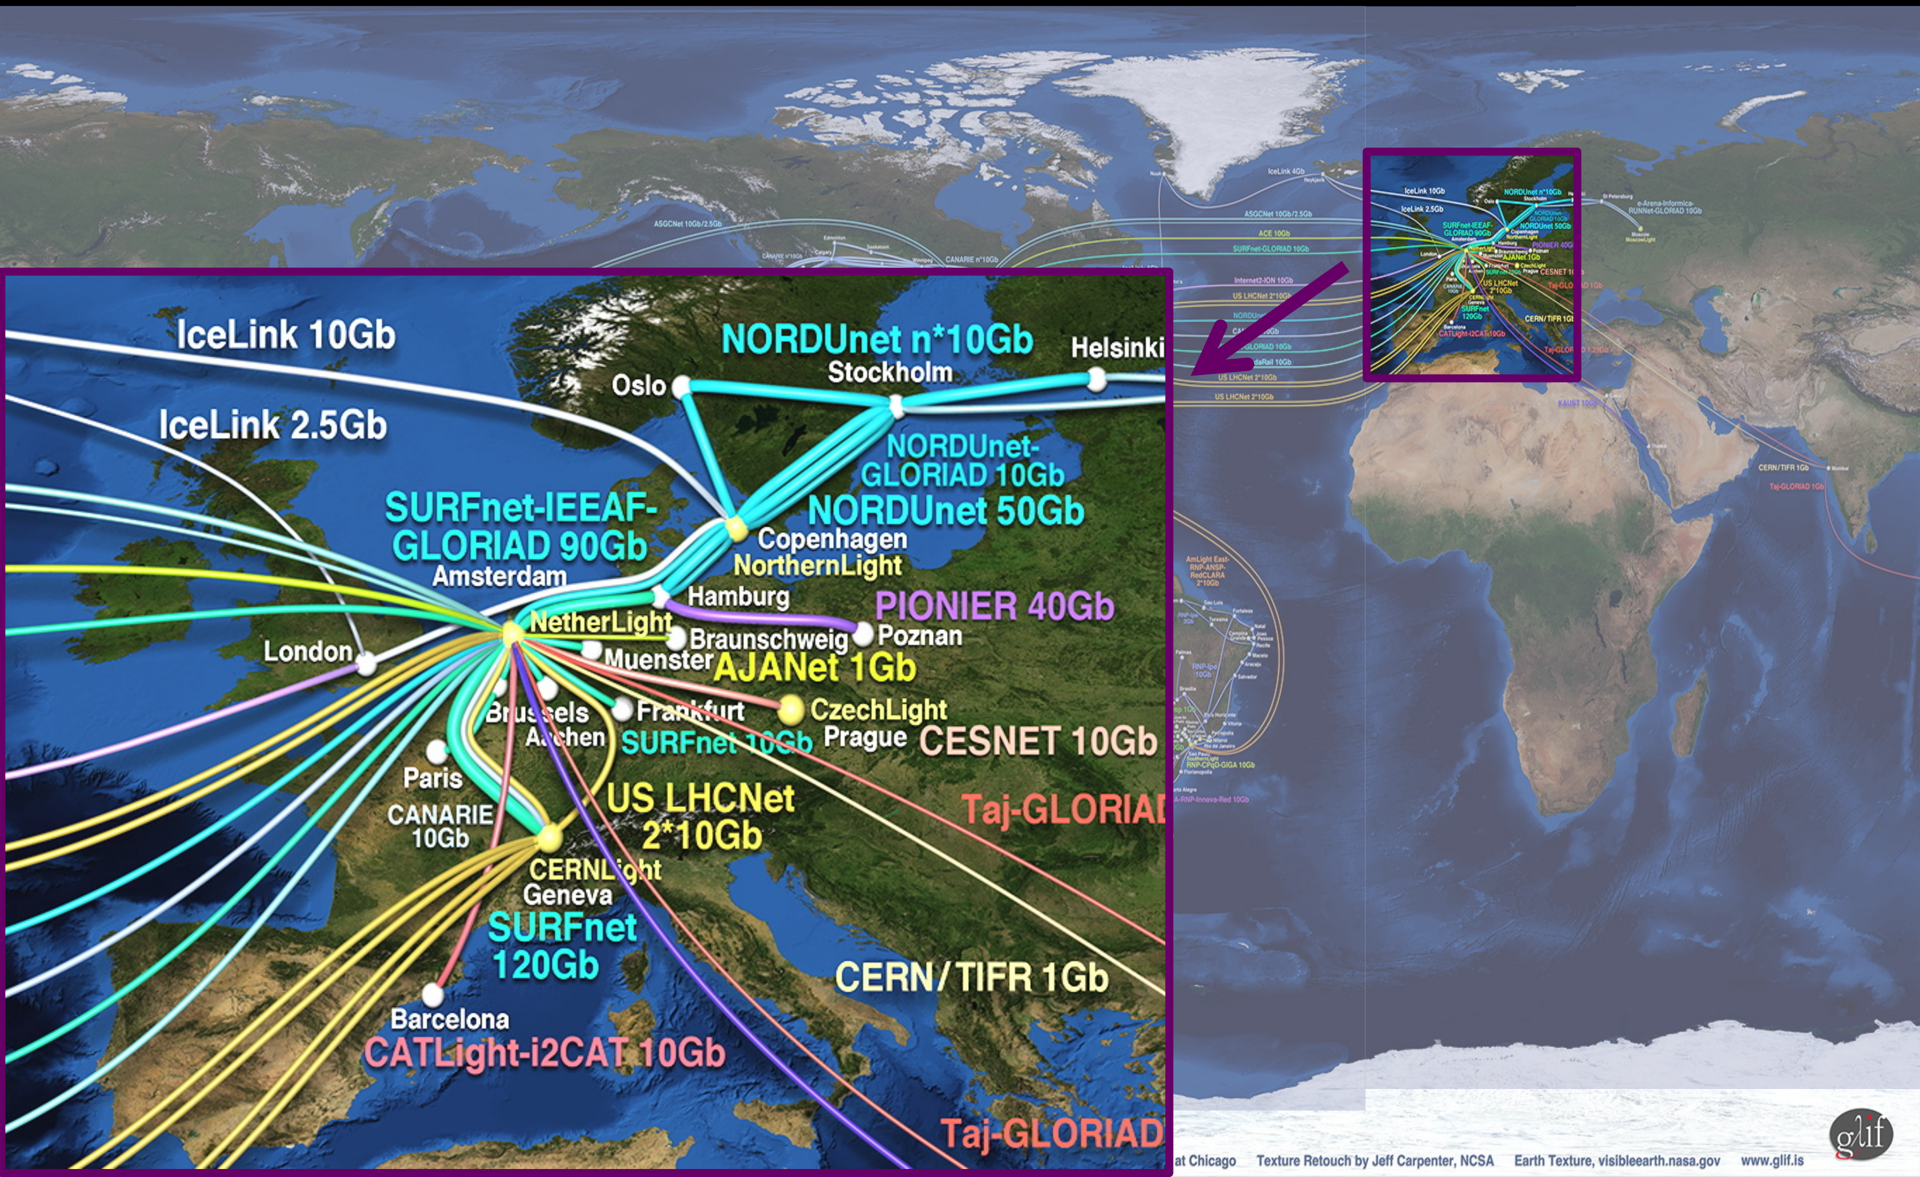

ExoGeni @ OpenLab - UvA

Installed and up June 3th 2013

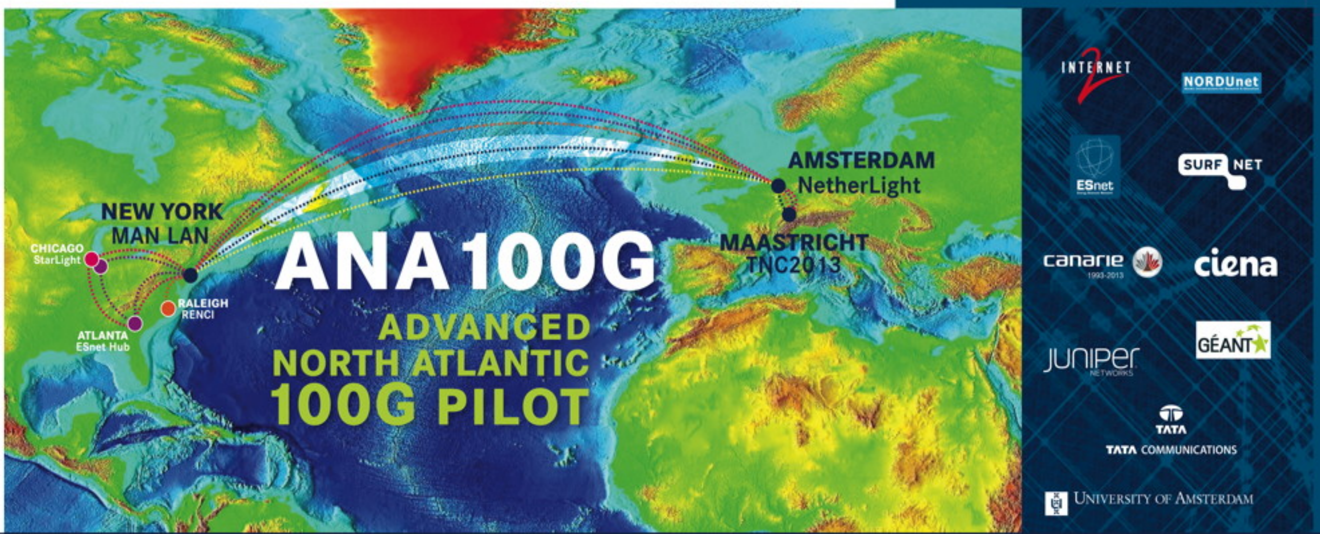

TNC2013 DEMOS JUNE, 2013

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

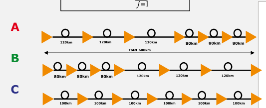

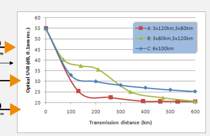

	FOM
A	5504
B	5504
C	1897

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

Transmission system setup

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

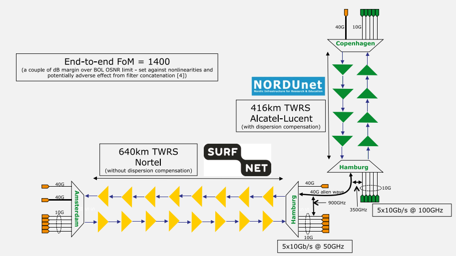

Test results

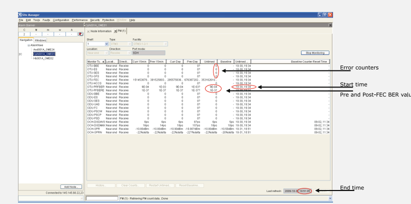

Error-free transmission for 23 hours, 17 minutes → BER < 3.0 10^{-16}

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 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

Alien light
From idea to realisation!

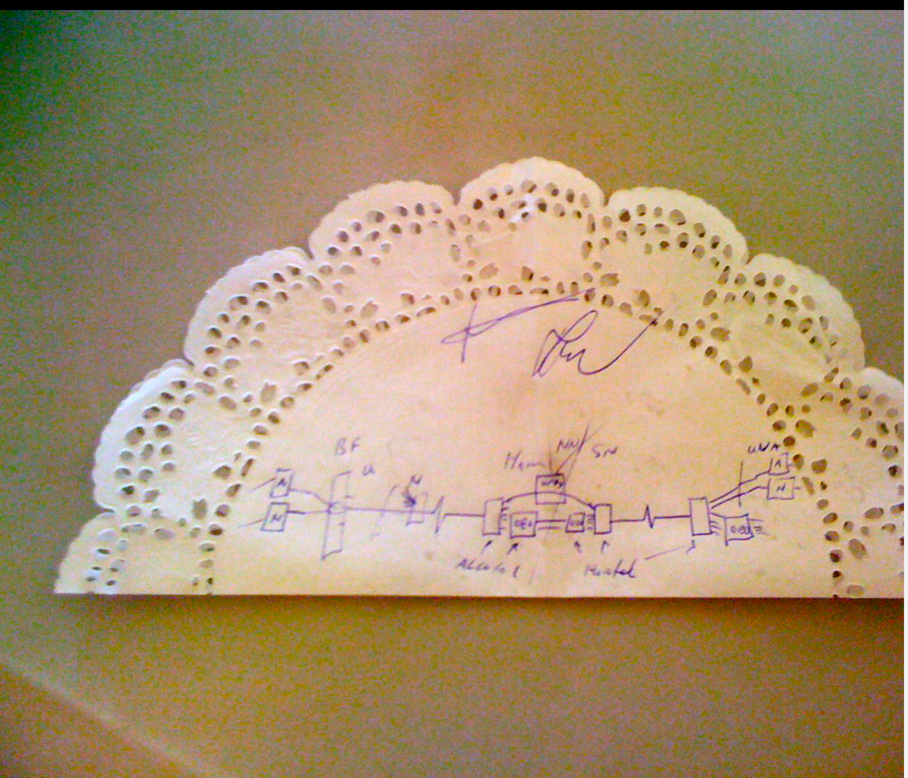

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

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

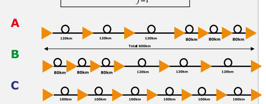

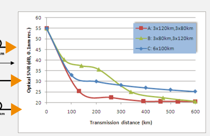

	FOM
A	5504
B	5504
C	1897

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

Transmission system setup

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

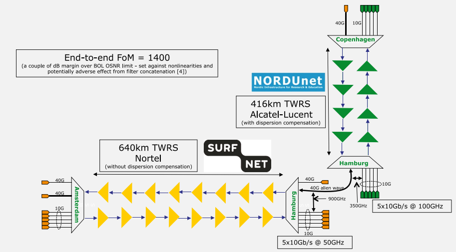

Test results

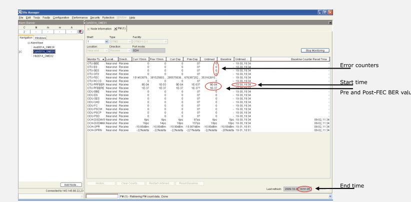

Error-free transmission for 23 hours, 17 minutes → BER < 3.0 10^{-16}

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

CPU Topology benchmark

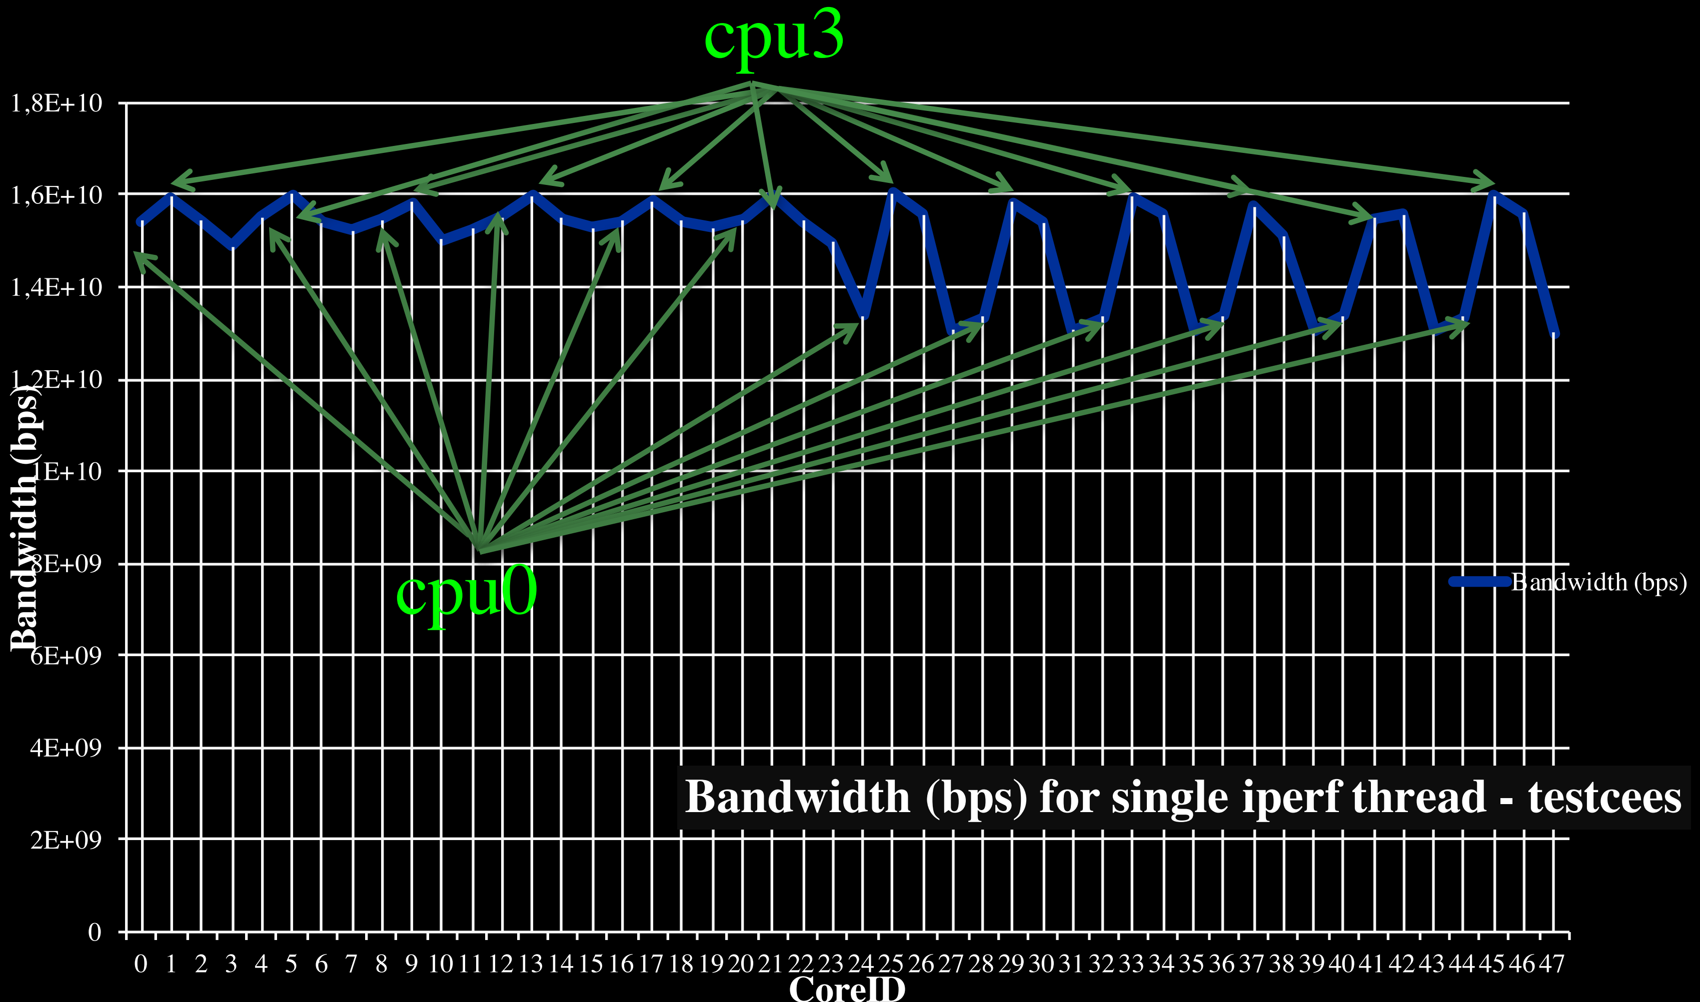

We used numactl to bind iperf to cores

SNE - 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***
  - ***Policy, Trust, Anonymity, Privacy, Integrity***
- *Sustainability*
  - *Greening infrastructure, Awareness*
- ***Resilience***
  - ***Failures, Disasters, Systems under attack***

SARNET: Security Autonomous Response with programmable NETworks

Cees de Laat
Leon Gommans, Rodney Wilson, Rob Meijer
Tom van Engers, Marc Lyonais, Paola Grosso, Frans Franken,
Ameneh Deljoo, Ralph Koning, Ben de Graaff, Stojan Trajanovski

UNIVERSITY OF AMSTERDAM

TNO

ciena

AIRFRANCE KLM

NWO
Netherlands Organisation for Scientific Research

COMMIT/

Cyber security program

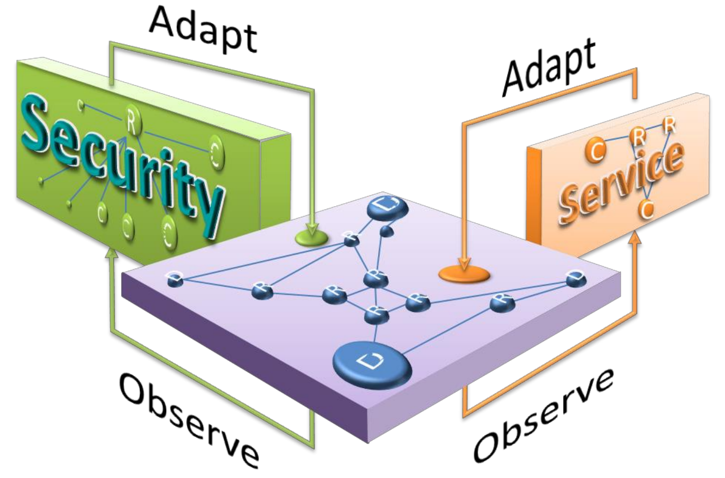

Research goal is to obtain the knowledge to create ICT systems that:

- model their state (situation)
- discover by observations and reasoning if and how an attack is developing and calculate the associated risks
- have the knowledge to calculate the effect of counter measures on states and their risks
- choose and execute one.

In short, we research the concept of networked computer infrastructures exhibiting SAR: Security Autonomous Response.

Context & Goal

Security Autonomous Response NETwork Research

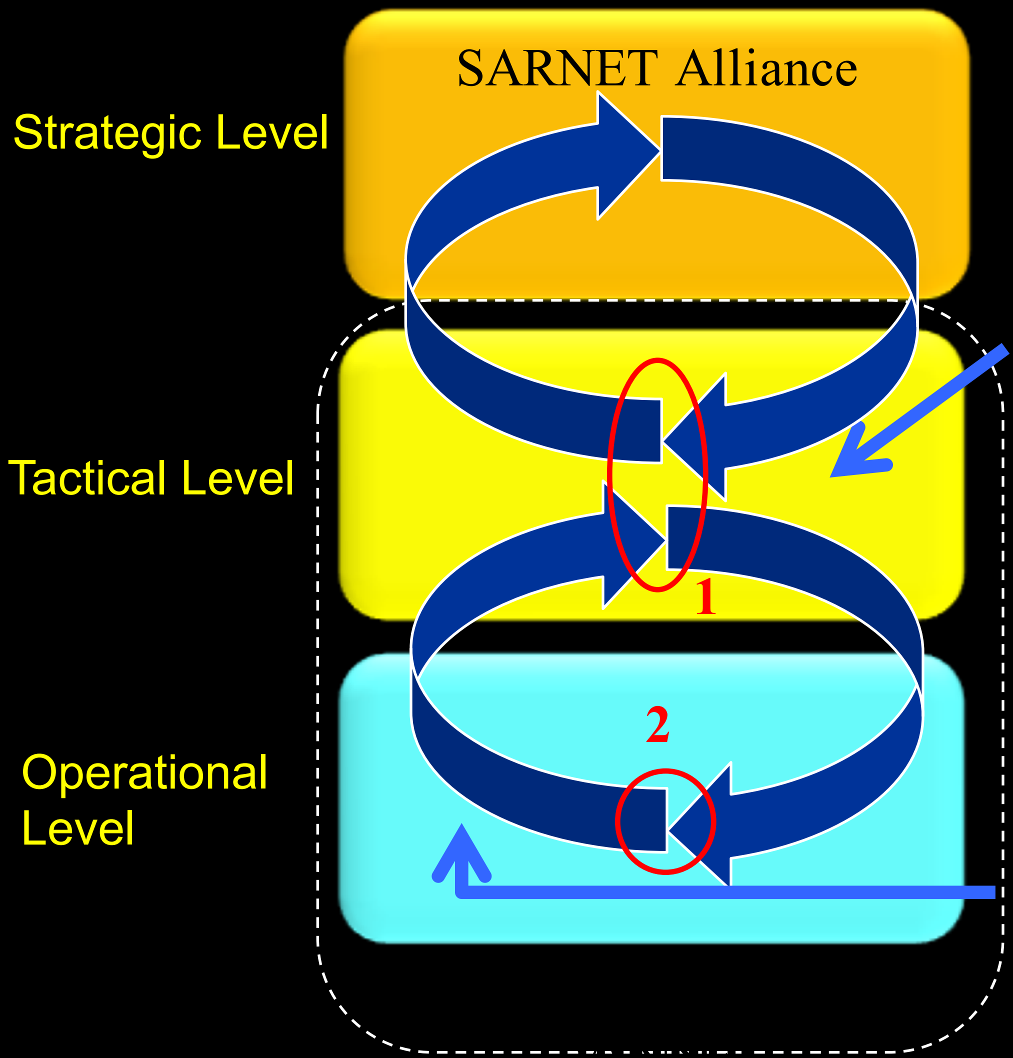

Ameneh Deljoo (PhD):
Why create SARNET Alliances?
Model autonomous SARNET behaviors to identify risk and benefits for SARNET stakeholders

Stojan Trajanovski (PD):
Determine best defense scenario against cyberattacks deploying SARNET functions (1) based on security state and KPI information (2).

Ralph Koning (PhD)
Ben de Graaff (SP):
1. Design functionalities needed to operate a SARNET using SDN/NFV
2: deliver security state and KPI information (e.g cost)

Ciena’s CENI topology

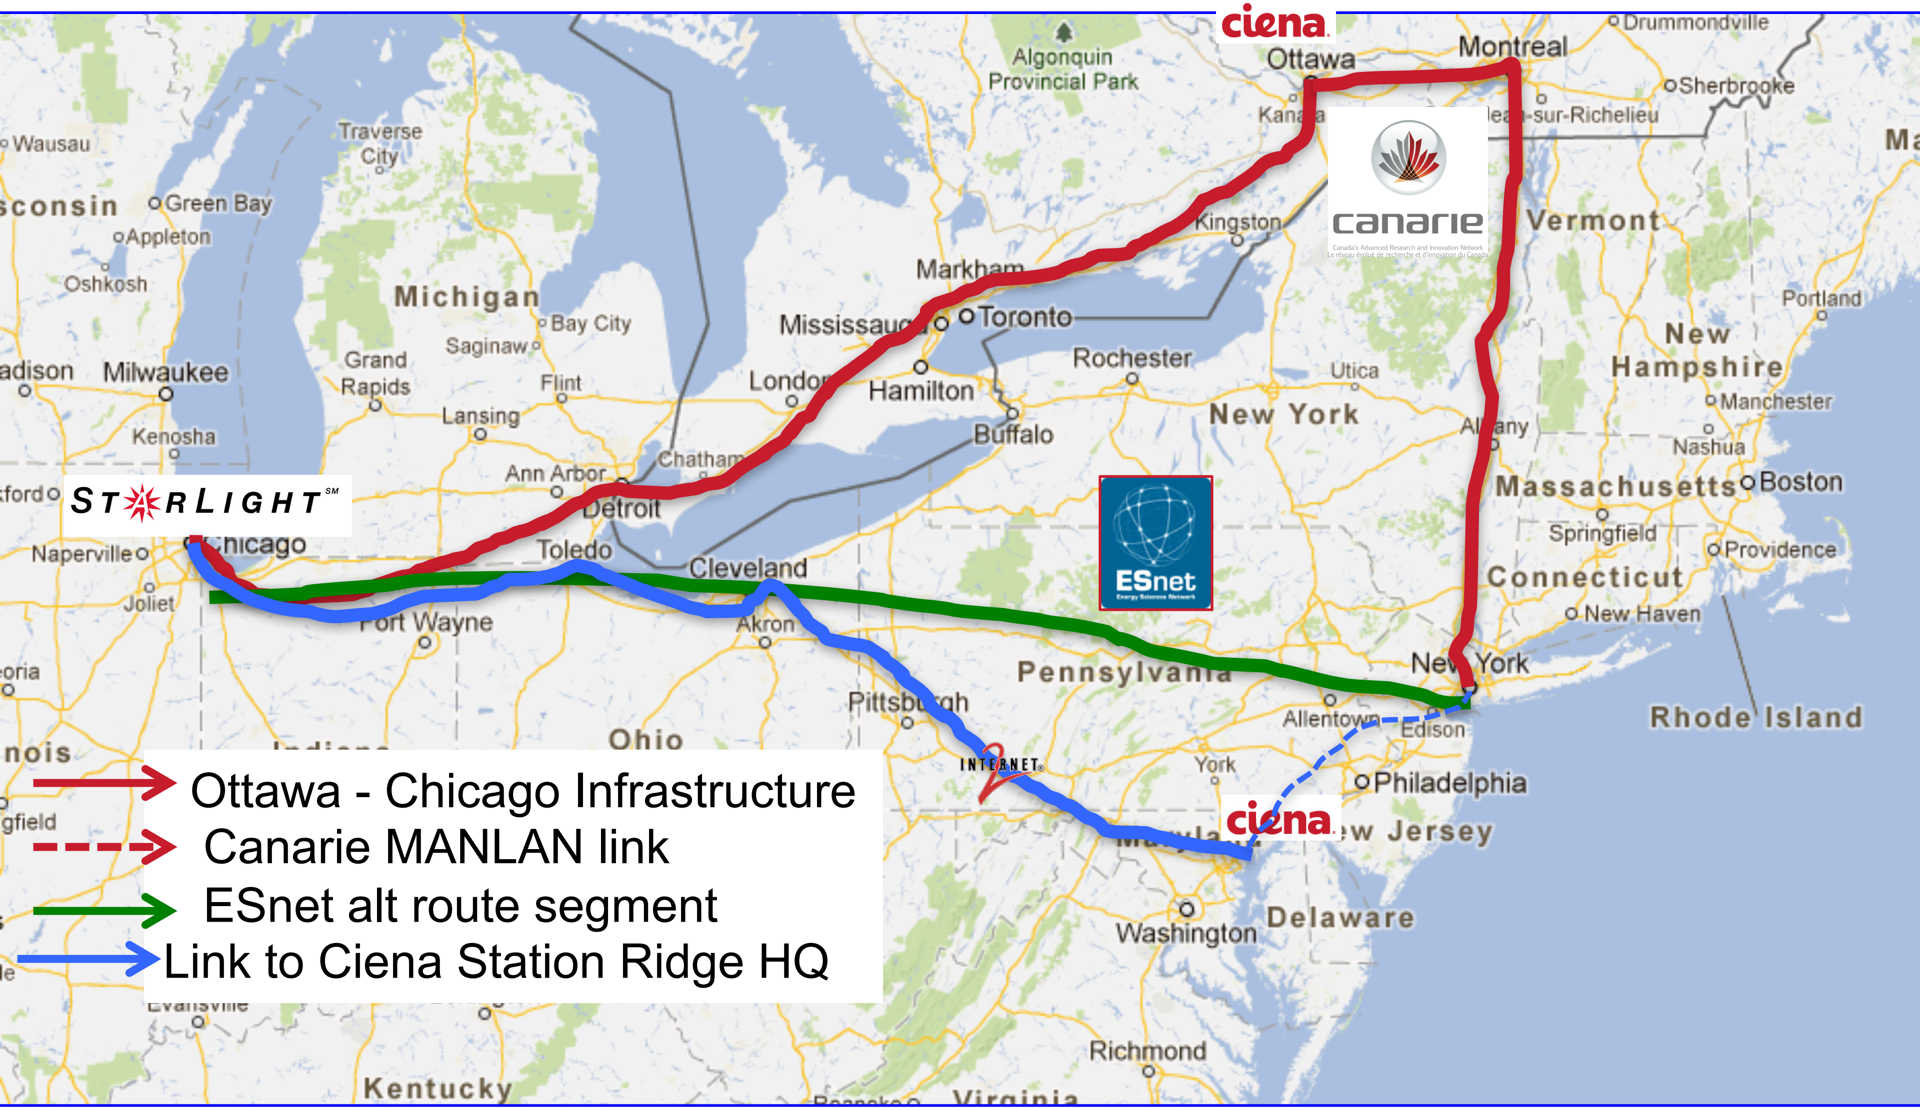

CENI, International extension to University of Amsterdam

Research Triangle Project.   Operation Spring of 2015

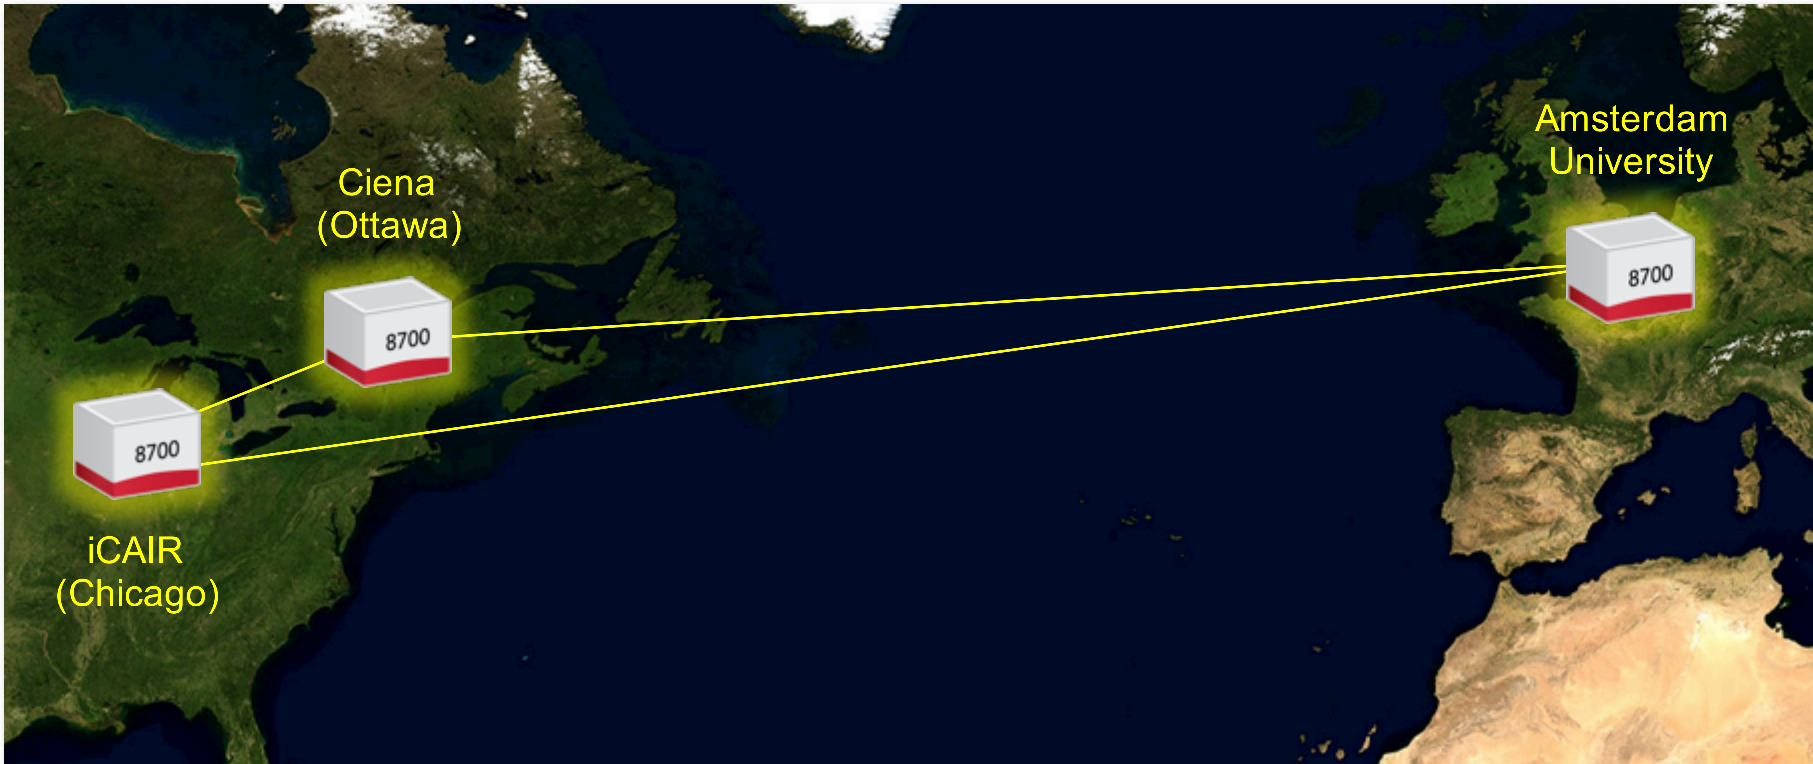

National Science Foundations ExoGENI racks, installed at UvA (Amsterdam), Northwestern University (Chicago) and Ciena's labs (Ottawa), are connected via a high performance 100G research network and trans-Atlantic network facilities using the Ciena 8700 Packetwave platform. This equipment configuration is used to create a computational and storage test bed used in collaborative demonstrations.

Position of demo @ SC15

Objective

- To get a better understanding for cyber attack complexity by visually defend a network suffering from basic volumetric attacks.
- To find a way to visualize future research in automated response.

Demo highlights

- Pre-programmed attack scenarios that are able to show defense functions.
- Virtual sales + income from web services
- Defense cost

DDoS Defence functions.

- Filtering
- Blocking
- Resource Scaling

Demo

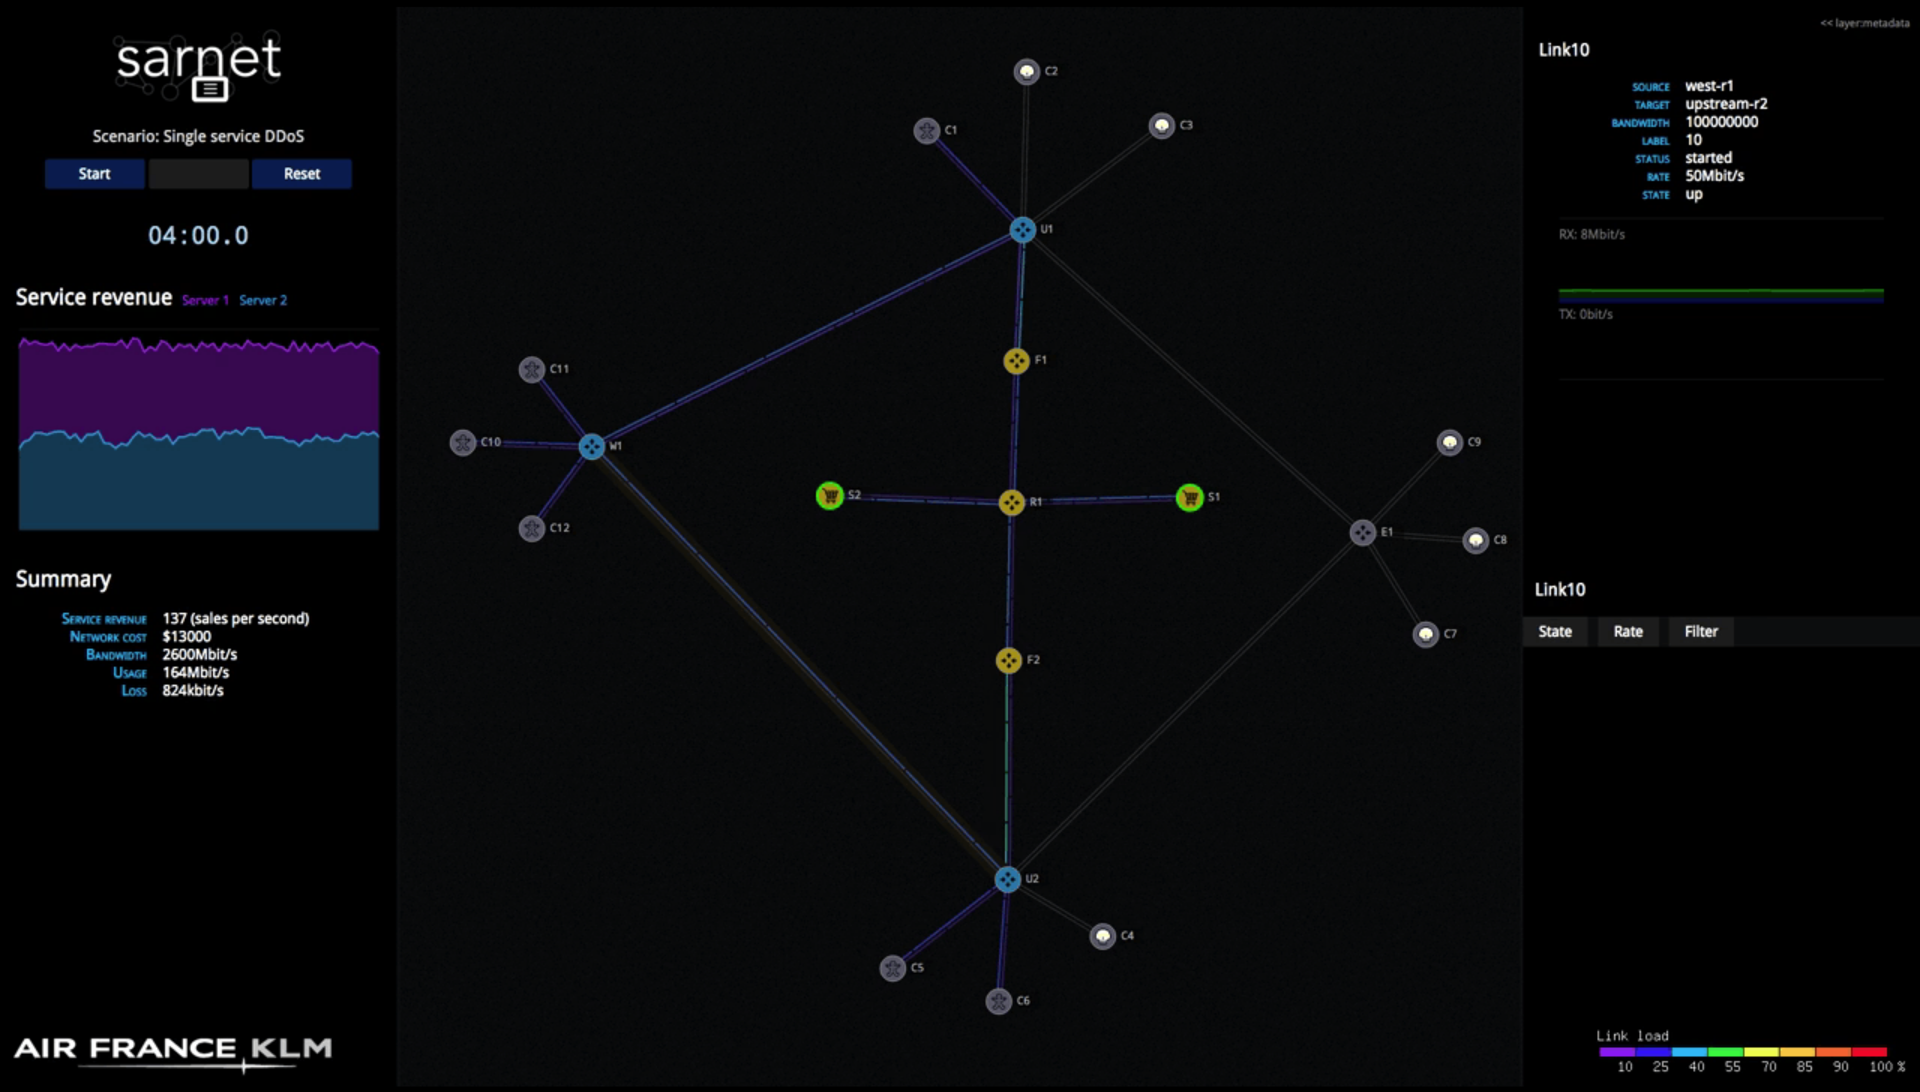
sarnet
Scenario: Single service DDoS
Start
Reset
04:00.0
Service revenue
Server 1
Server 2
Summary
Service revenue 137 (sales per second)
Network cost $13000
Bandwidth 2600Mbit/s
Usage 164Mbit/s
Loss 824kbit/s
AIR FRANCE KLM
C2
C1
C3
U1
F1
C11
C10
W1
S2
R1
S1
C9
E1
C8
C12
C7
F2
U2
C4
C5
C6
<< layer:metadata
Link10
source west-r1
target upstream-r2
bandwidth 100000000
label 10
status started
rate 50Mbit/s
state up
RX: 8Mbit/s
TX: 0bit/s
Link10
State
Rate
Filter
Link load
10 25 40 55 70 85 90 100 %

Service Provider Group framework

A Service Provider Group (SPG) is an organisation structure providing a defined service only available if its members collaborate.

Examples:

Envisioned role of the SPG: define slice archetypes?

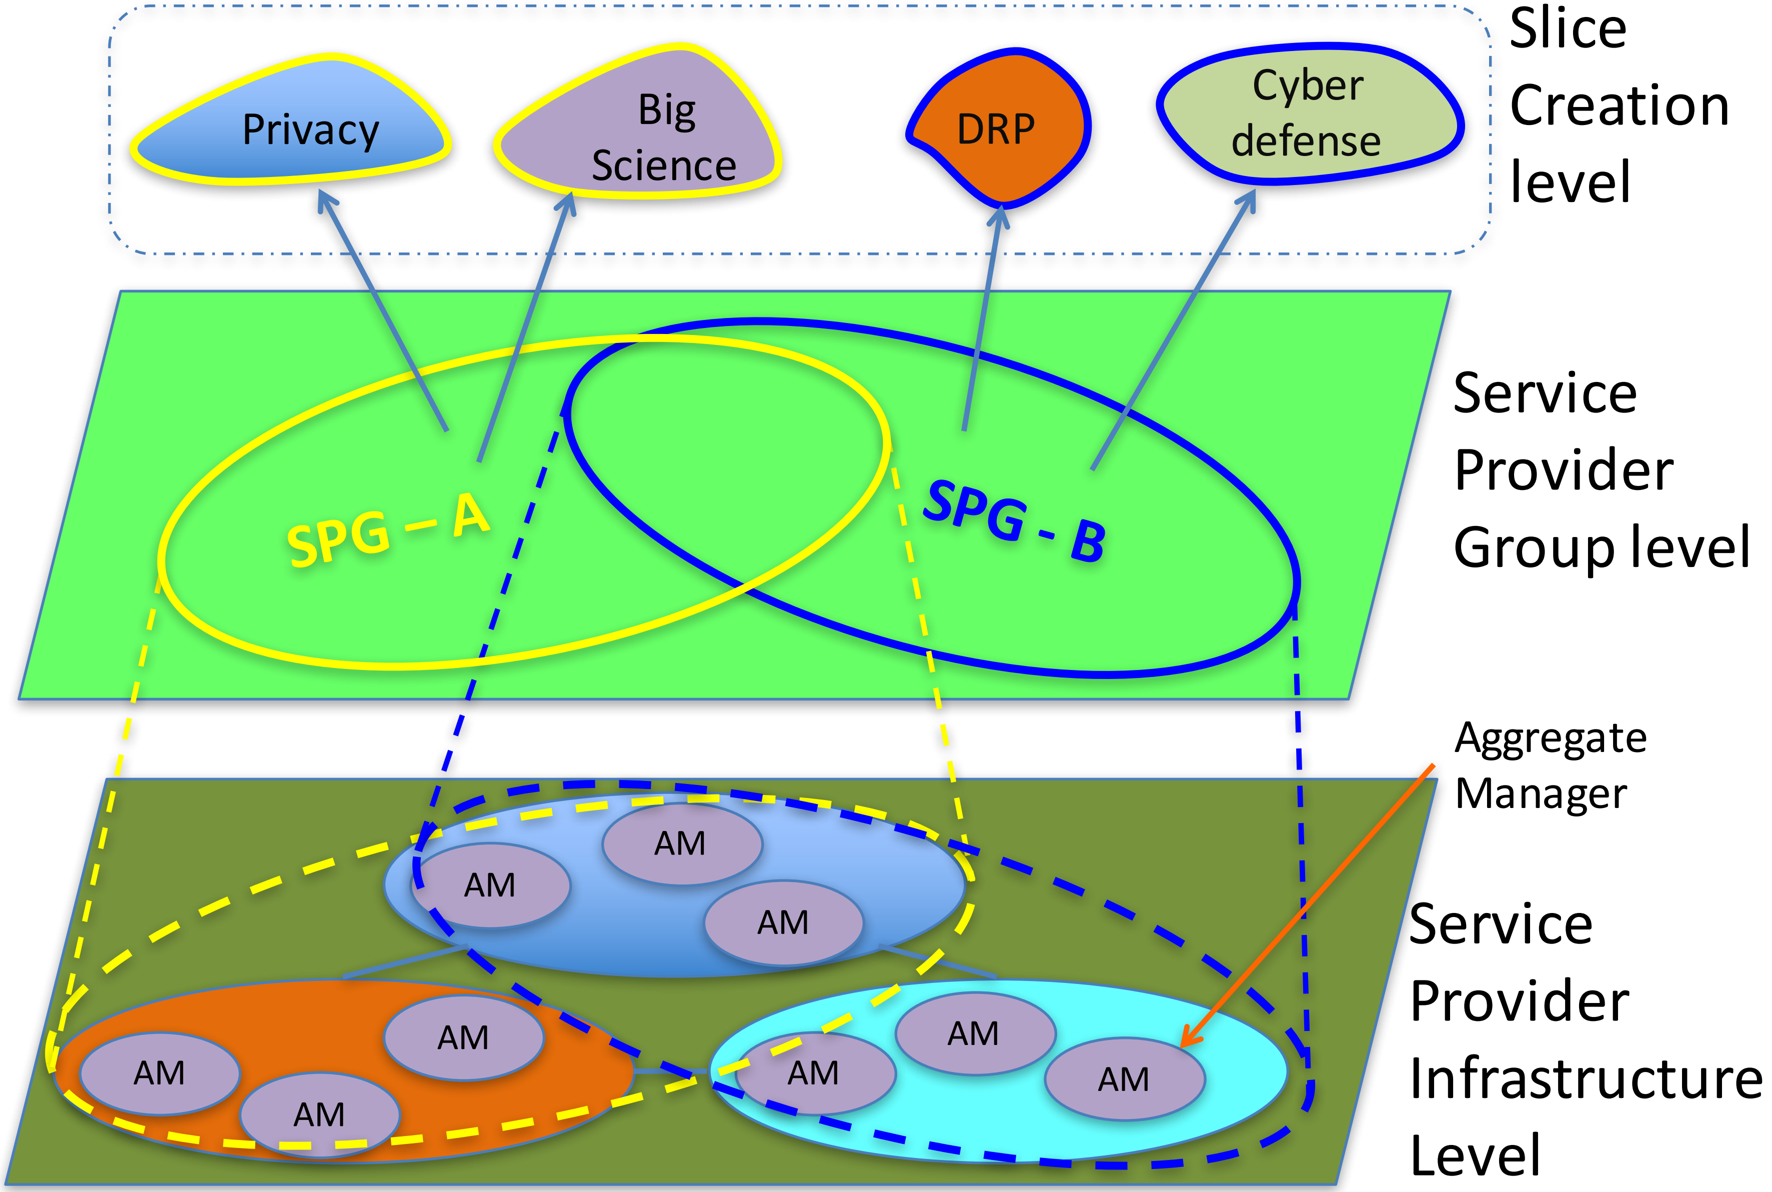

Agent Based Modelling Framework

	Main component
Signal layer	Message / Act
Action layer	Action / Activity
Intentional layer	Intention
Motivational layer	Motive

In our model, we refer to four layers of components:

- the signal layer— describes acts, side-effects and failures showing outcomes of actions in a topology.
- the action layer—actions: performances that bring a certain result,
- the intentional layer—intentions: commitments to actions, or to build up intentions,
- the motivational layer—motives: events triggering the creation of intentions.

Simplified Eduroam case at signalling layer

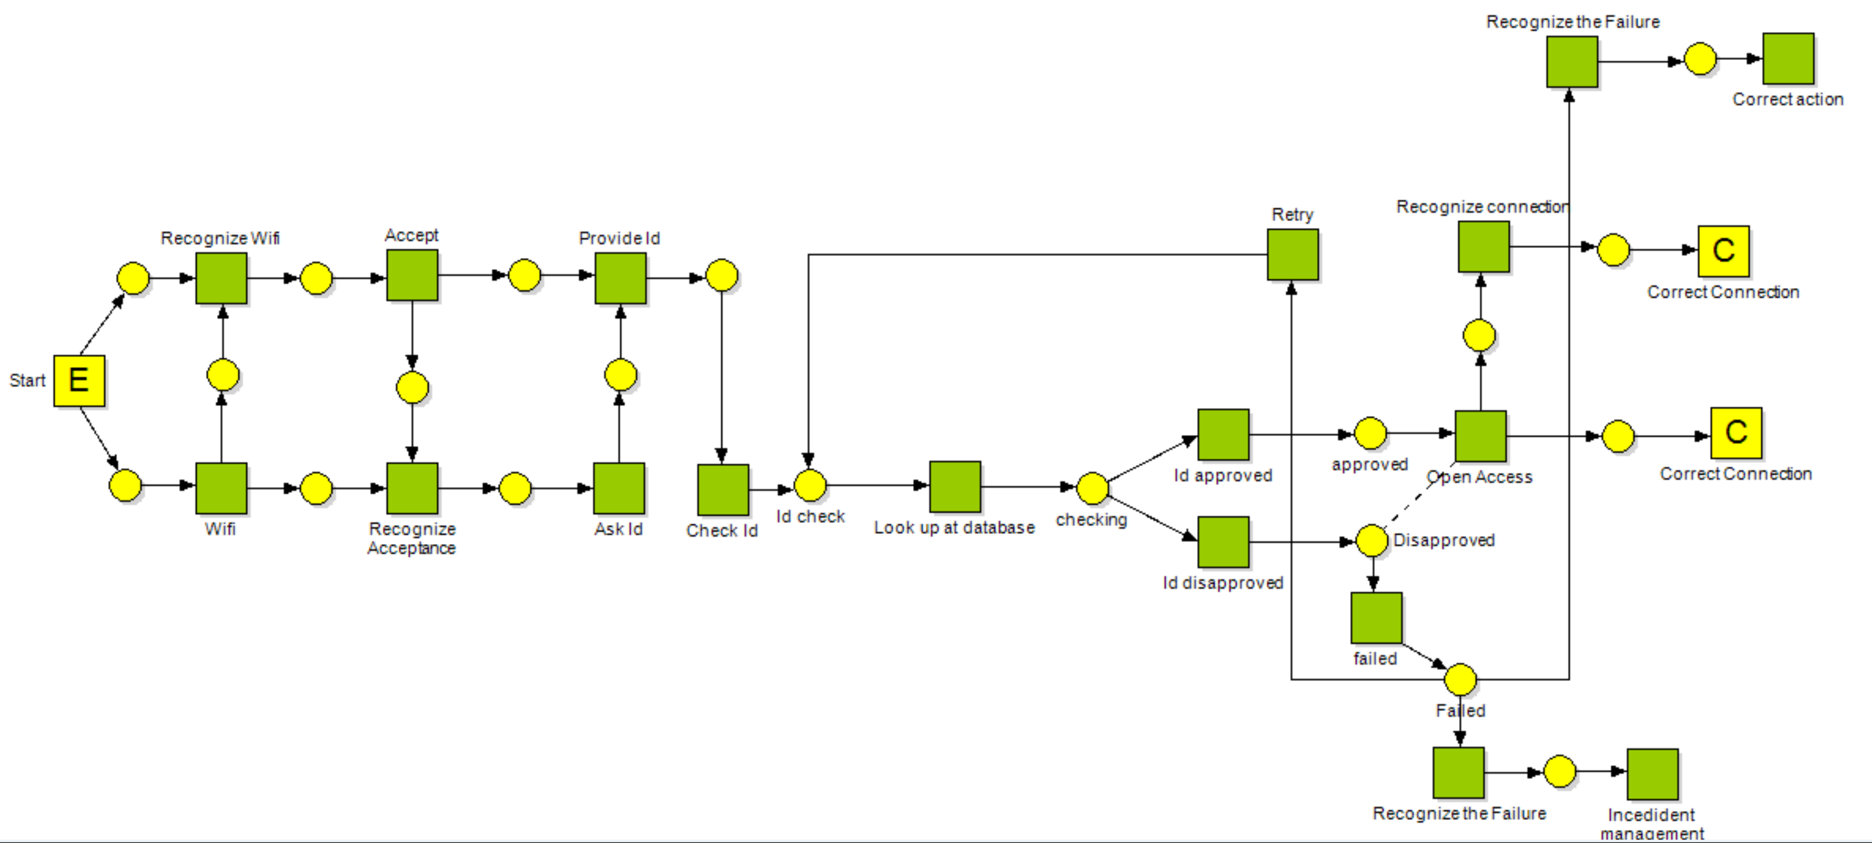

Petri net of EduRoam Case
(first step)

Describing Intentions, Motivations and Actions

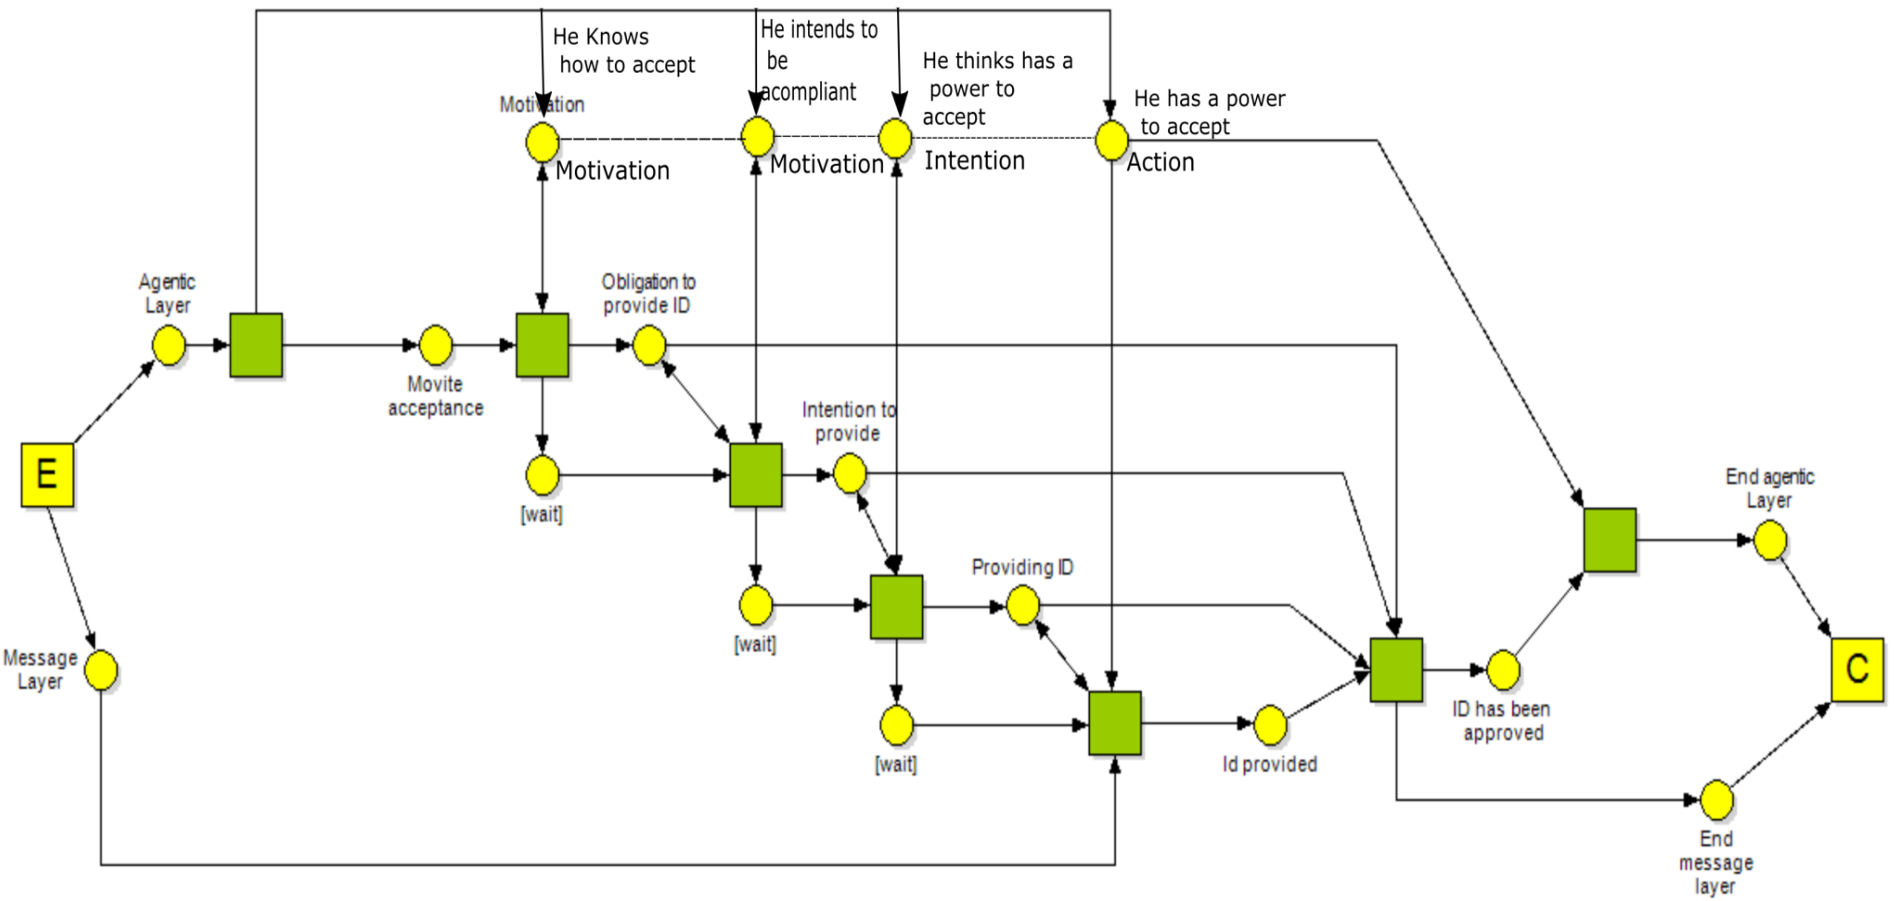

Petri net of EduRoam Case

SNE - 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*
  - *Policy, Trust, Anonymity, Privacy, Integrity*
- *Sustainability*
  - *Greening infrastructure, Awareness*
- *Resilience*
  - *Failures, Disasters, Systems under attack*

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!

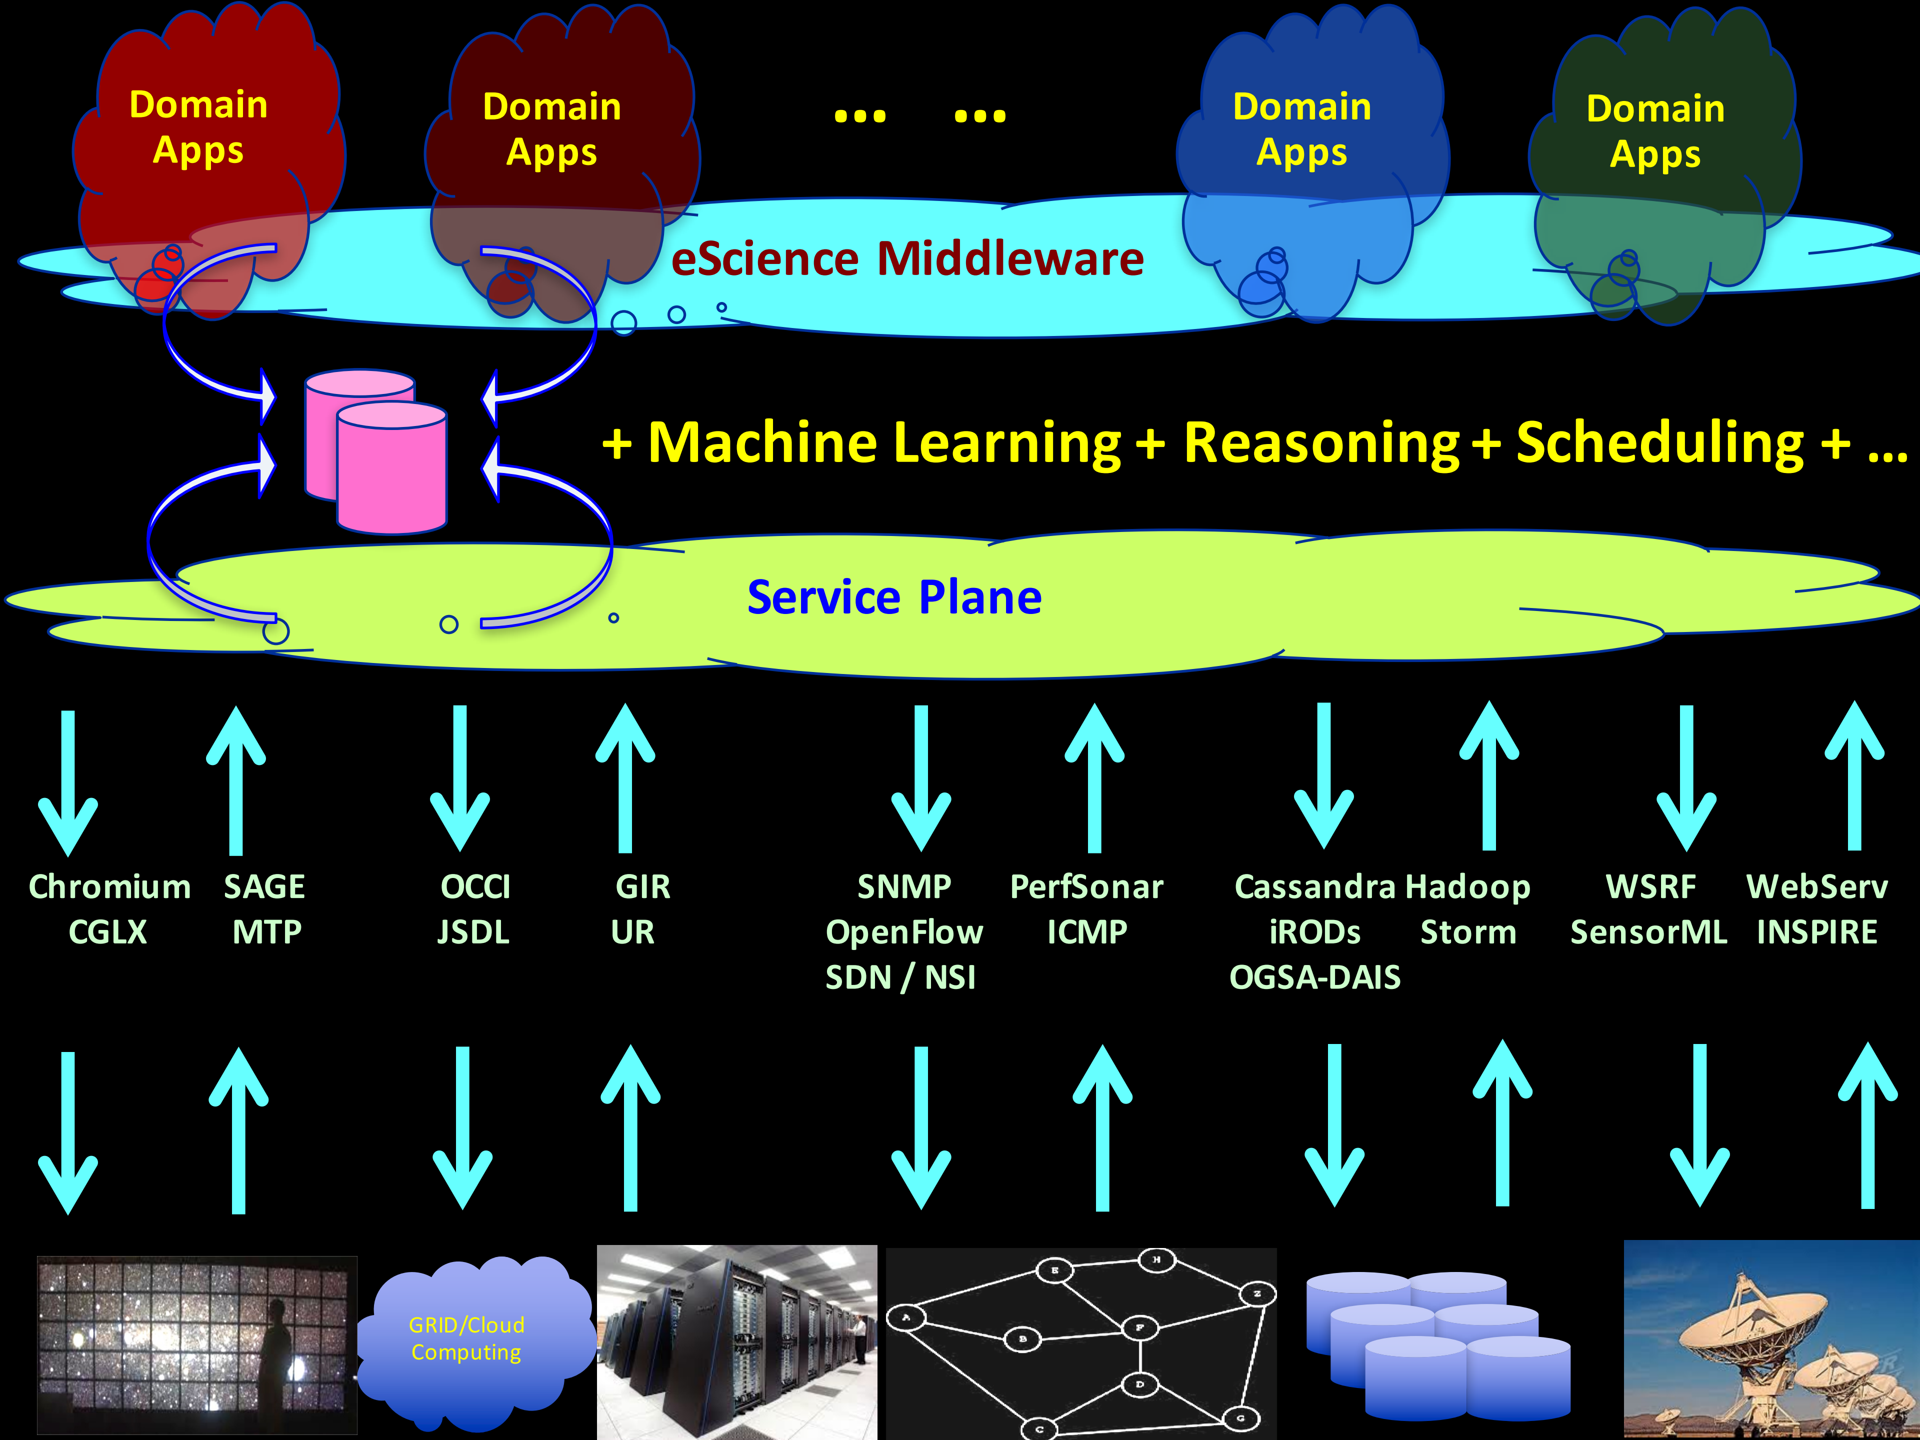

Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
+ Machine Learning + Reasoning + Scheduling + ...
Service Plane
Chromium CGLX
SAGE MTP
OCCI JSDL
GIR UR
SNMP OpenFlow SDN / NSI
PerfSonar ICMP
Cassandra iRODs OGSA-DAIS
Hadoop Storm
WSRF SensorML
WebServ INSPIRE
GRID/Cloud Computing

The Big Data Challenge

Doing Science

- Wis dom
- Knowledge to act
- Information
- Data
  a.o. from ESFRI's

ICT to enable Science

- e-IRG
- Workflows Schedulers to act
- OWL
- XML, RDF, rSpec, SNMP, Java based, etc.

The Big Data Challenge

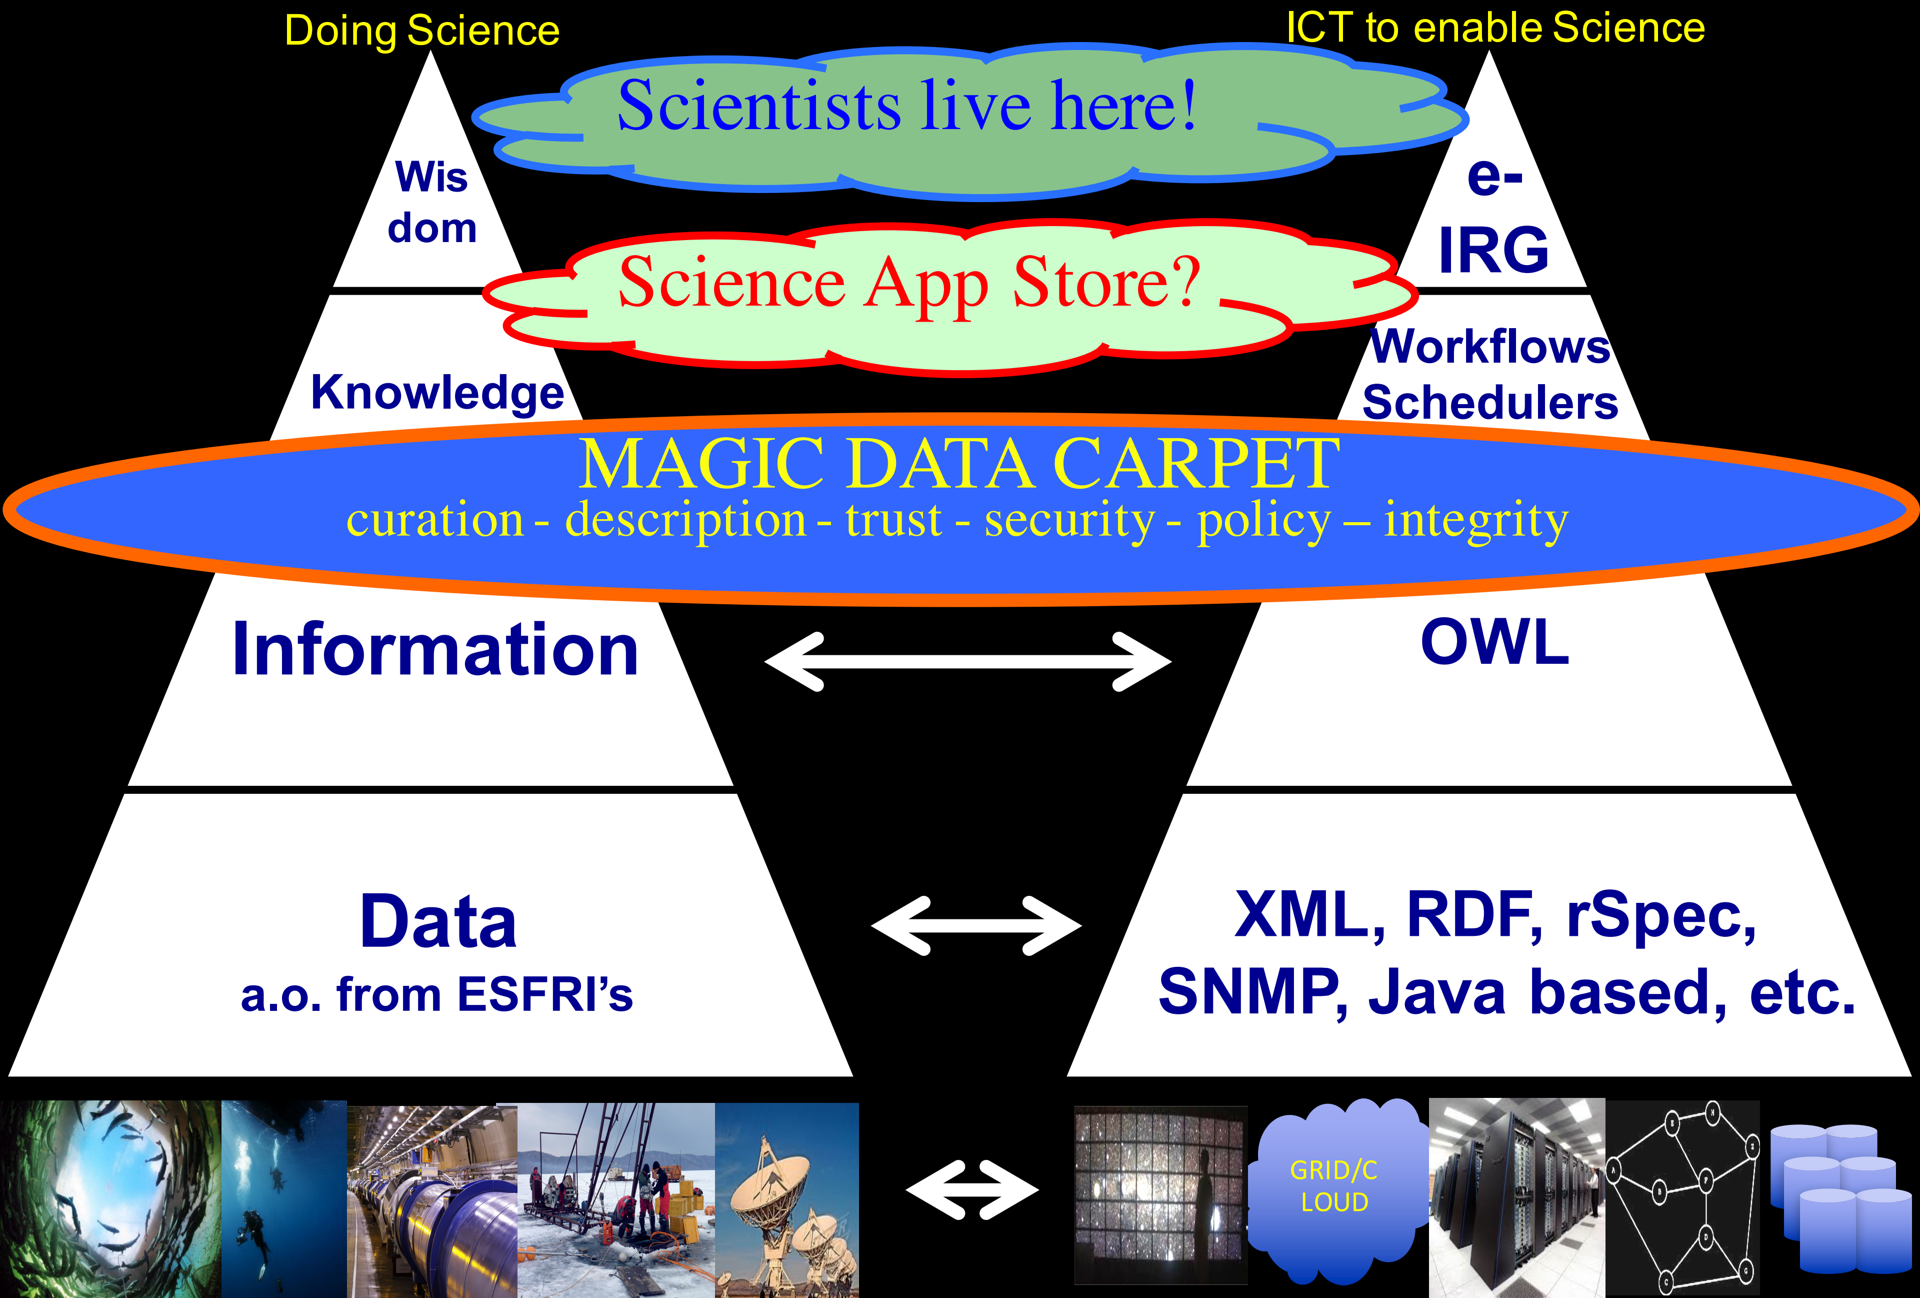

Mission

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

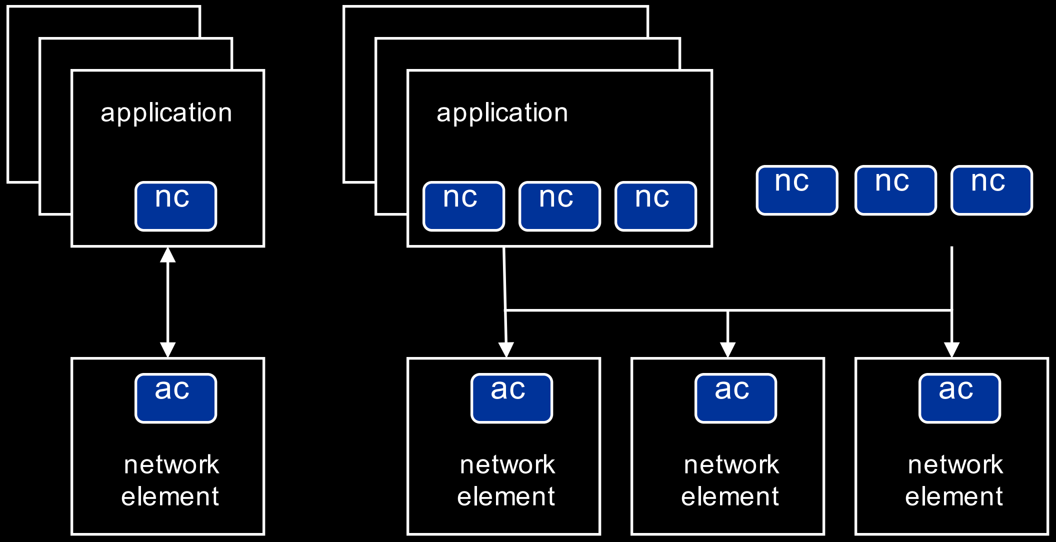

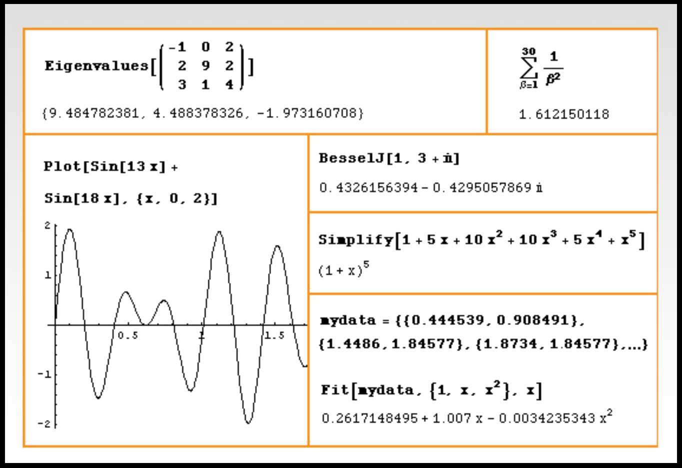

Mathematica enables advanced graph queries, visualizations and real-time network manipulations on UPVNs

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}
```

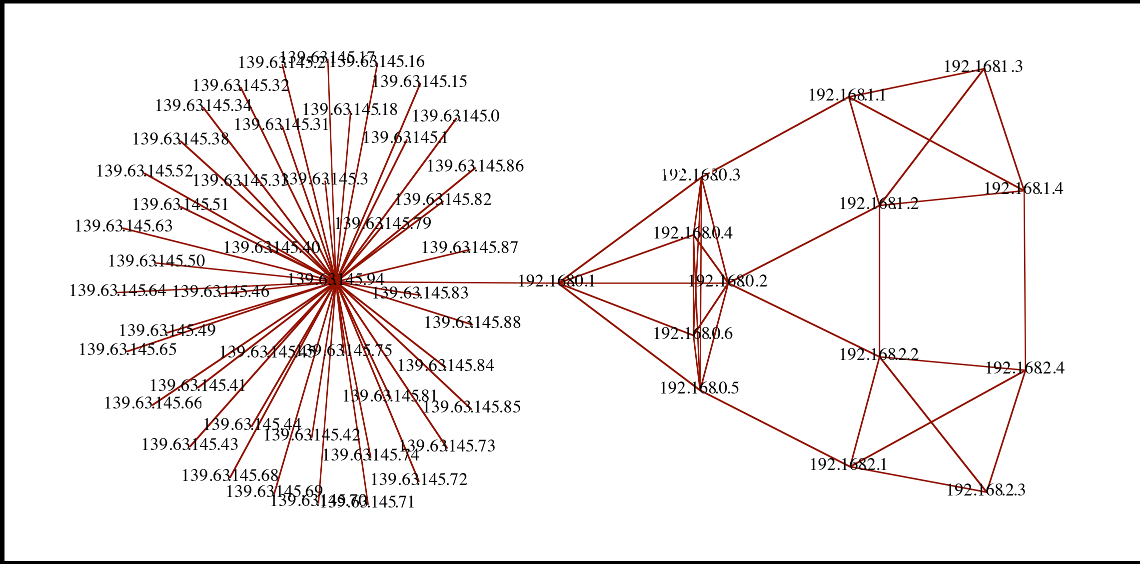

Transaction on shortest path with tokens

```
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
```

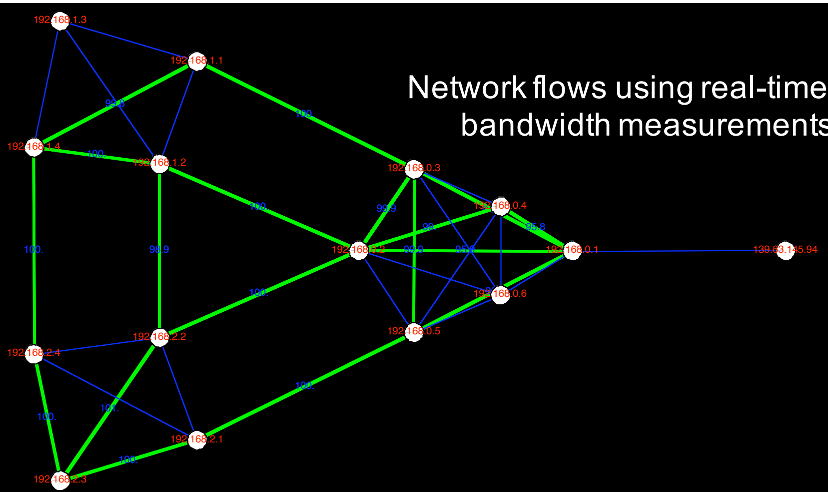

Network flows using real-time bandwidth measurements

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.

CDN on Demand in the cloud

Interactive programmable networks

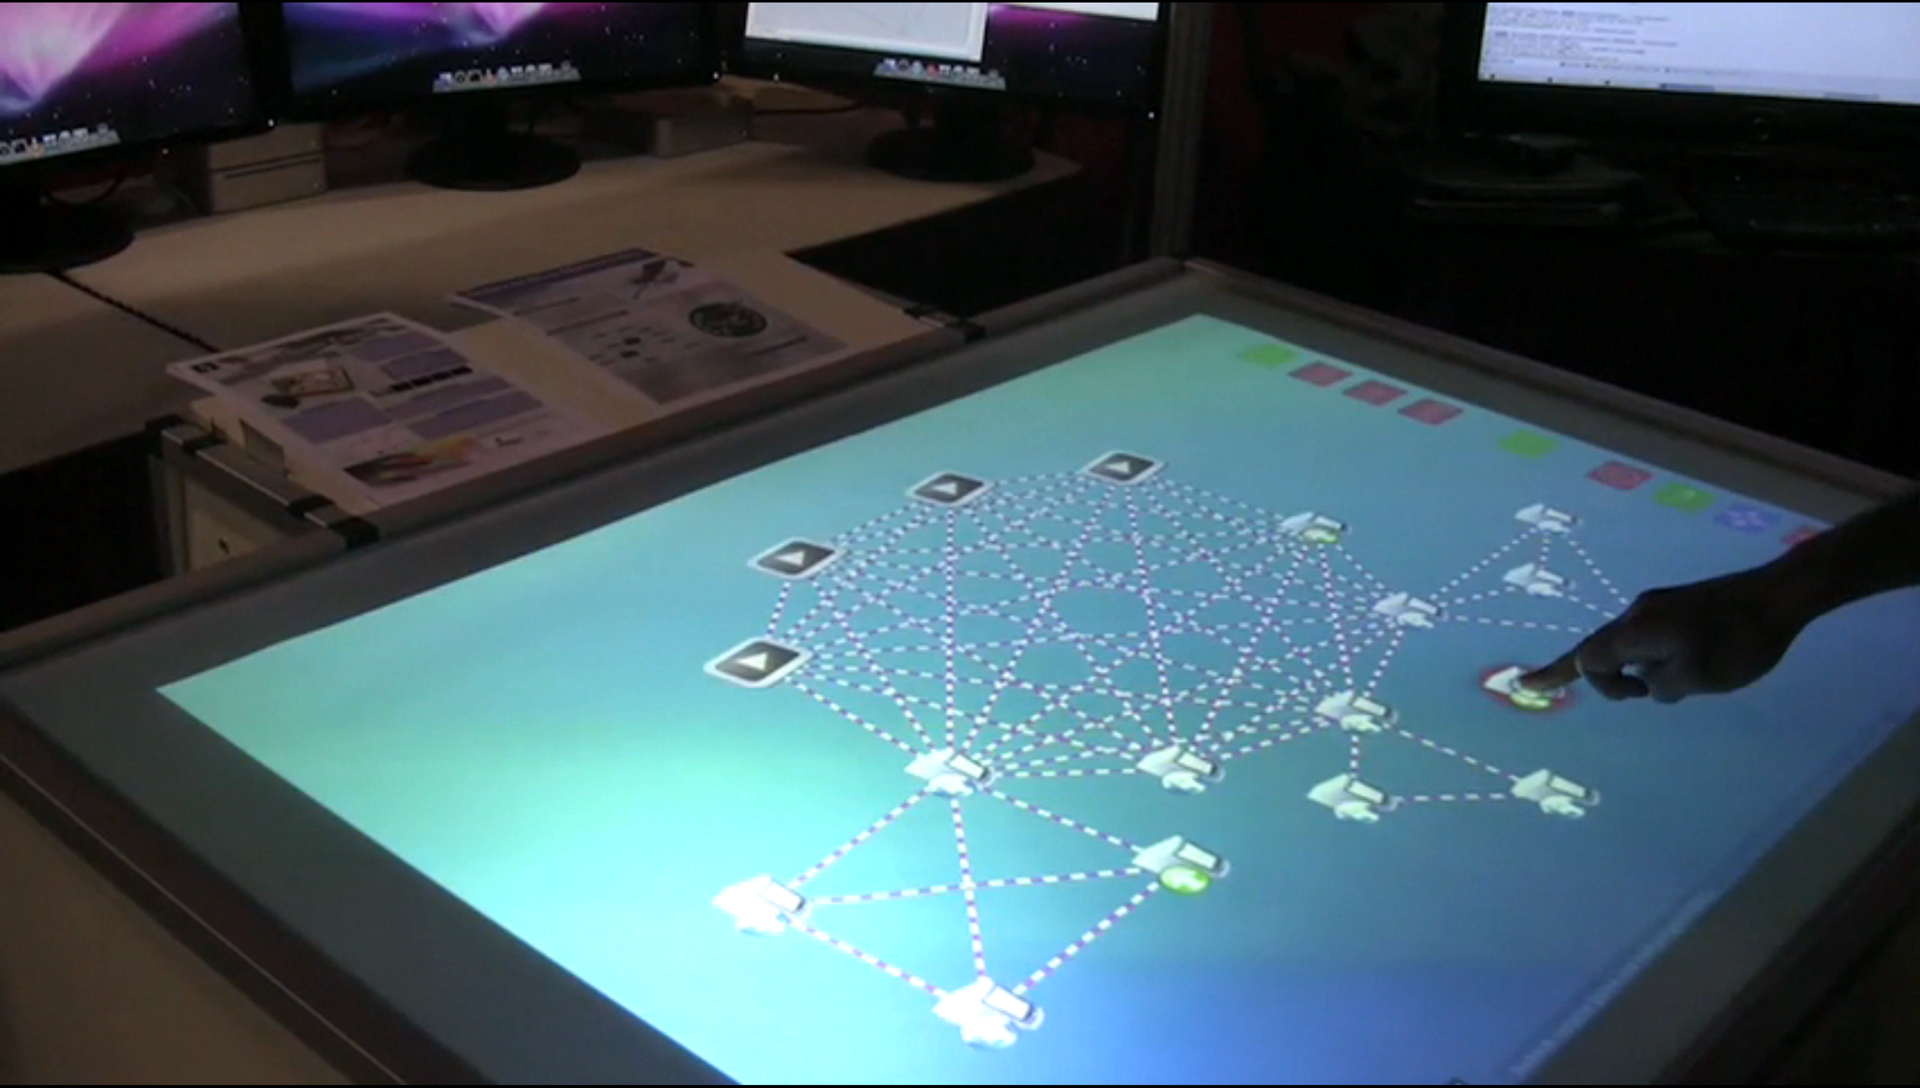

Basic operating system loop

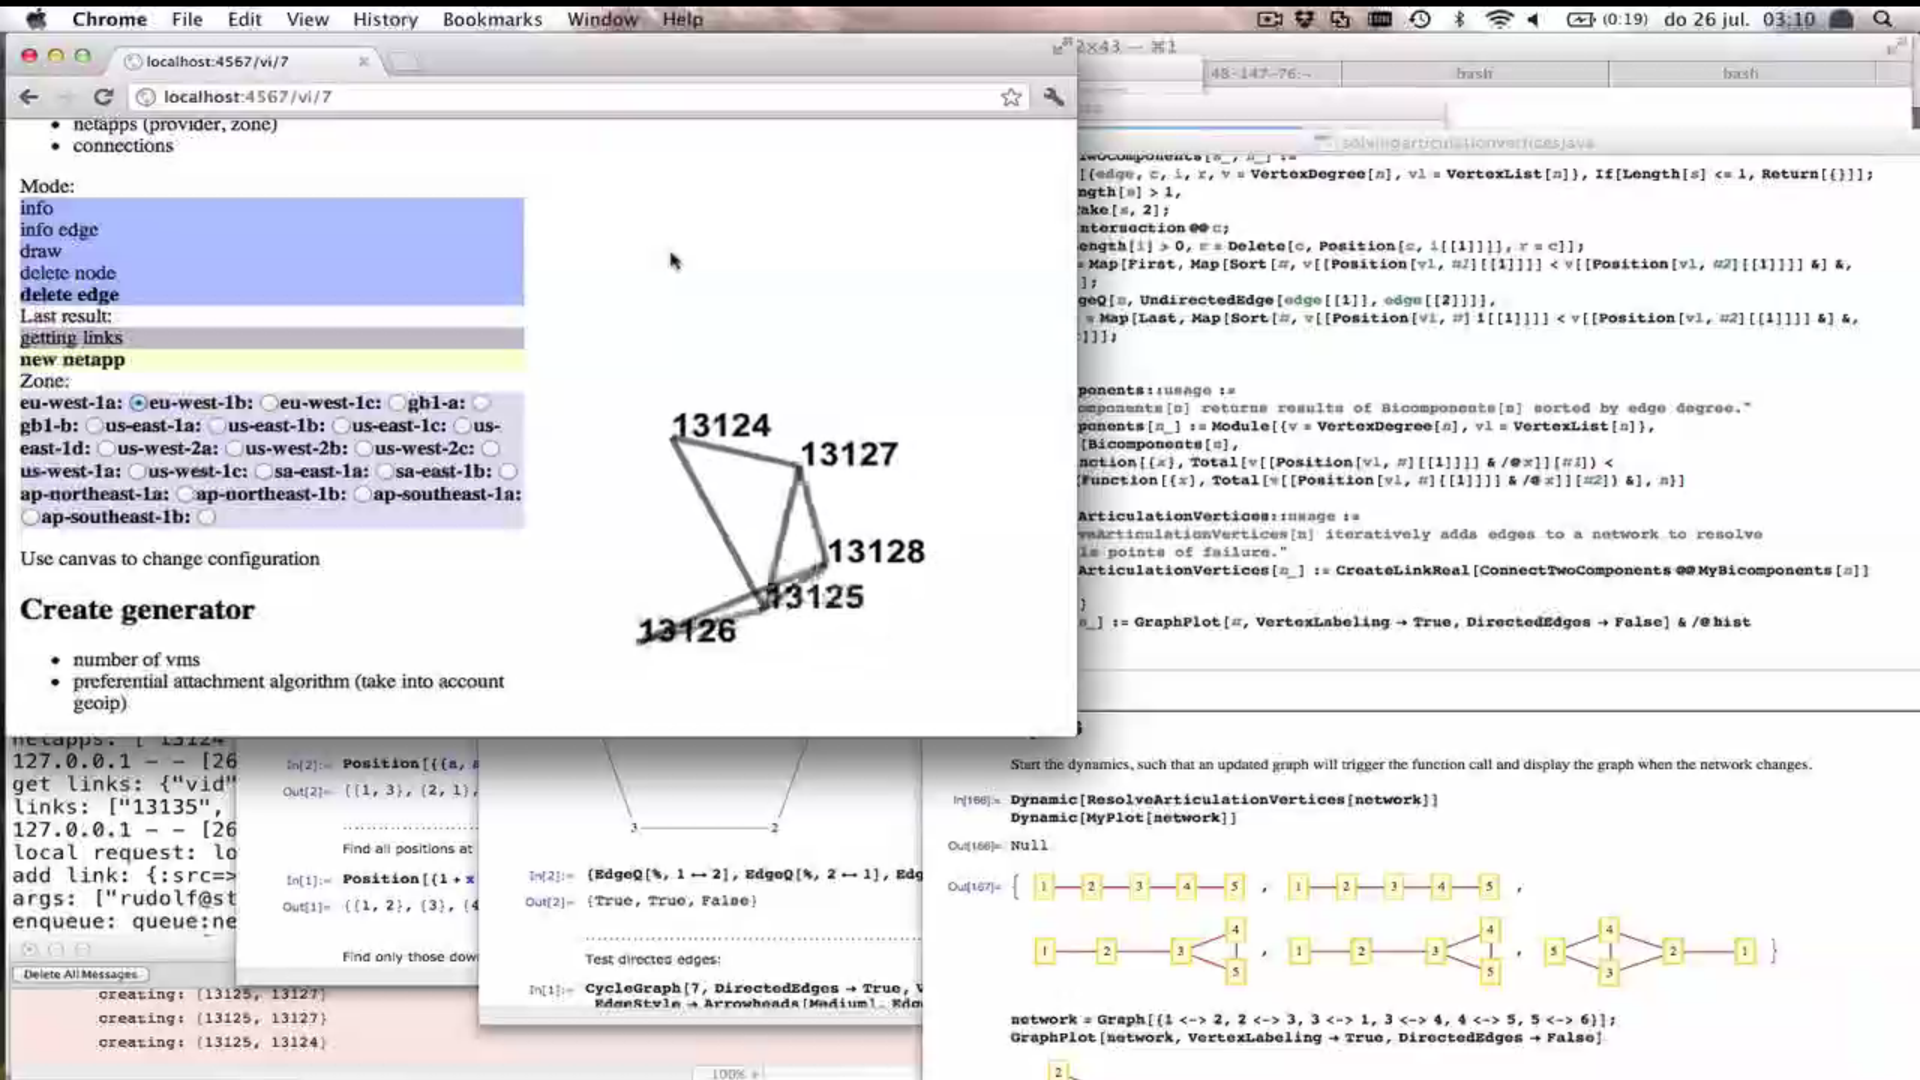

TimeLine

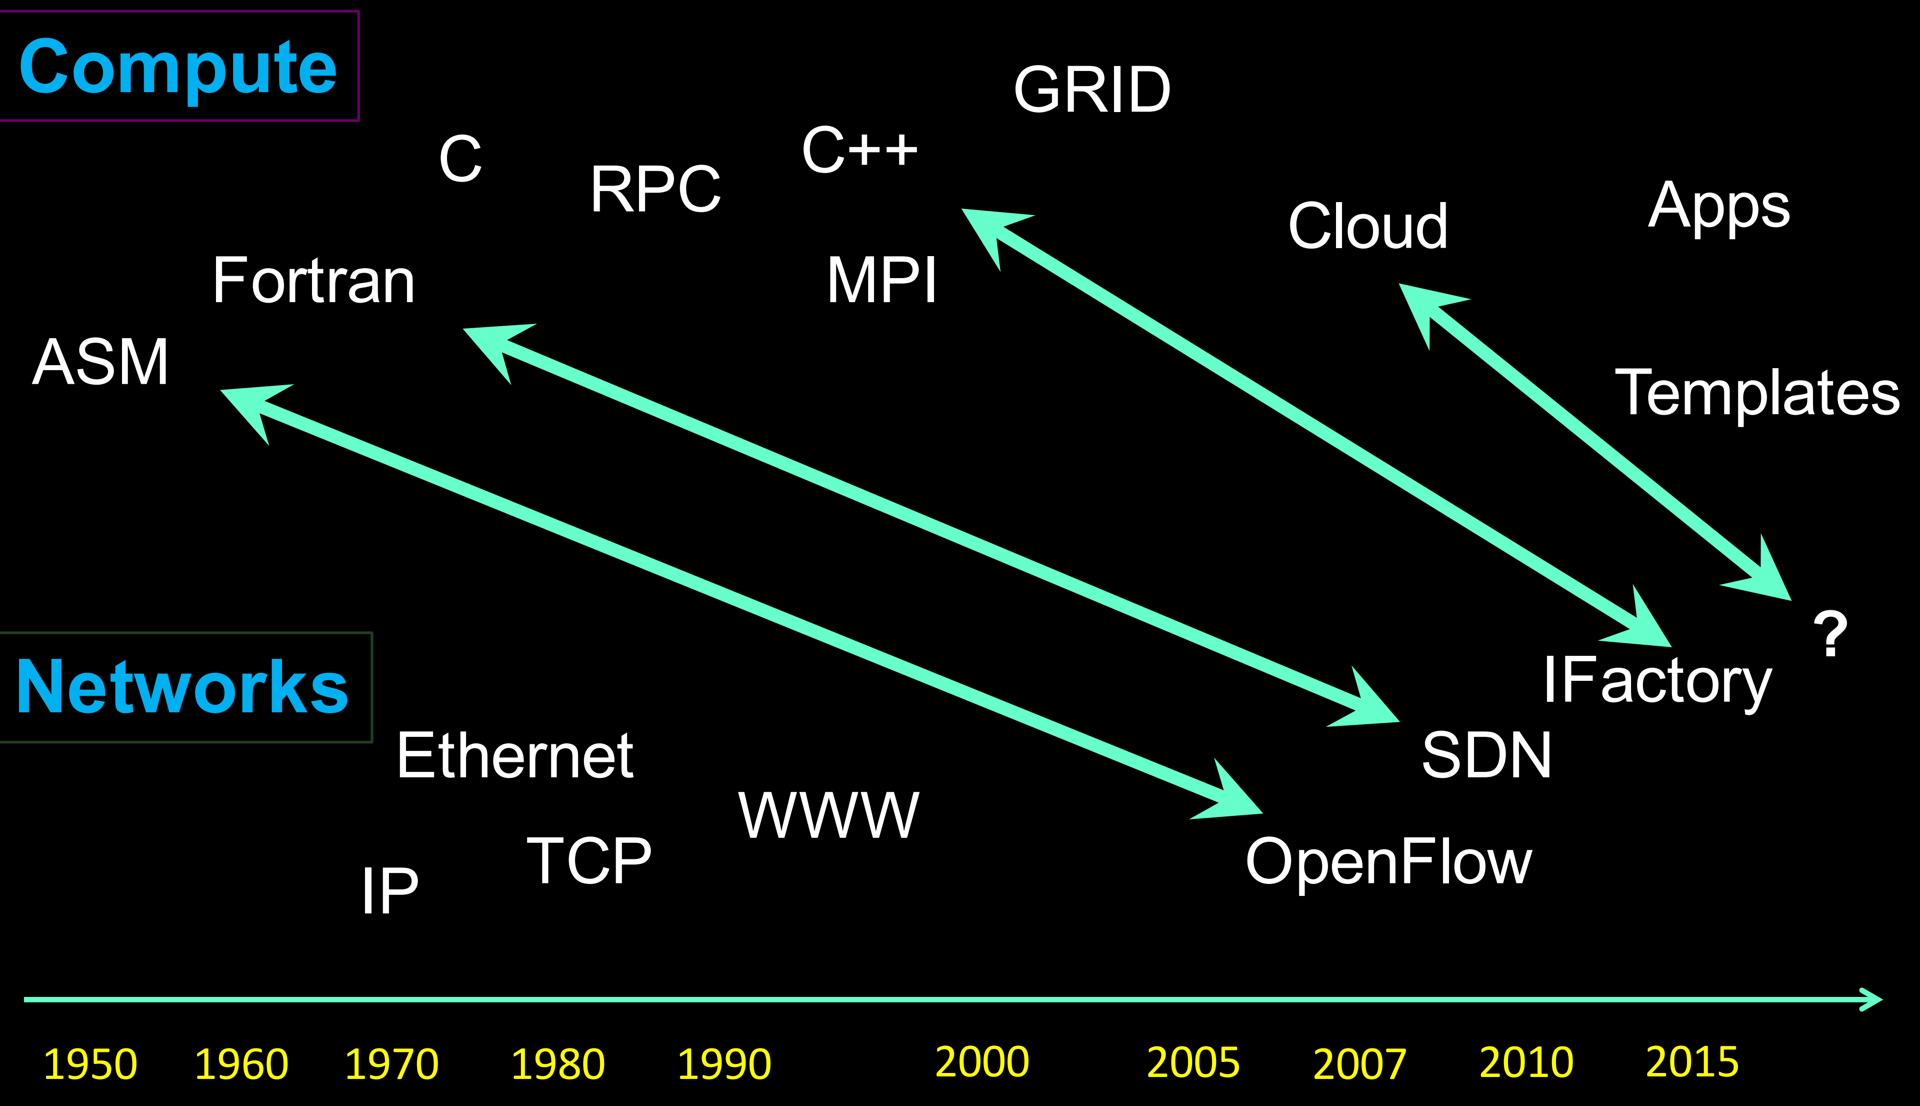
Compute
C
RPC
C++
GRID
Fortran
MPI
Cloud
Apps
ASM
Templates
?
IFactory
Networks
SDN
Ethernet
WWW
TCP
OpenFlow
IP
1950
1960
1970
1980
1990
2000
2005
2007
2010
2015

Questions?

http://delaat.net

http://delaat.net/sarnet

Leon Gommans, "Multi-Domain Authorization for e-Infrastructures", UvA, Dec 2014.

http://delaat.net/pubs/2014-t-3.pdf

Rudolf Strijkers, "Internet Factories", UvA, Nov 2014.

http://delaat.net/pubs/2014-t-2.pdf

Contact us:

delaat@uva.nl

p.grosso@uva.nl

z.zhao@uva.nl

SNE - 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*
  - *Policy, Trust, Anonymity, Privacy, Integrity*
- ***Sustainability***
  - ***Greening infrastructure, Awareness***
- *Resilience*
  - *Failures, Disasters, Systems under attack*

The smart network, the smart infrastructure

Dr. Paola Grosso
System and Network Engineering (SNE) research group
UvA
Email: p.grosso@uva.nl

Universiteit van Amsterdam

System and Network Engineering

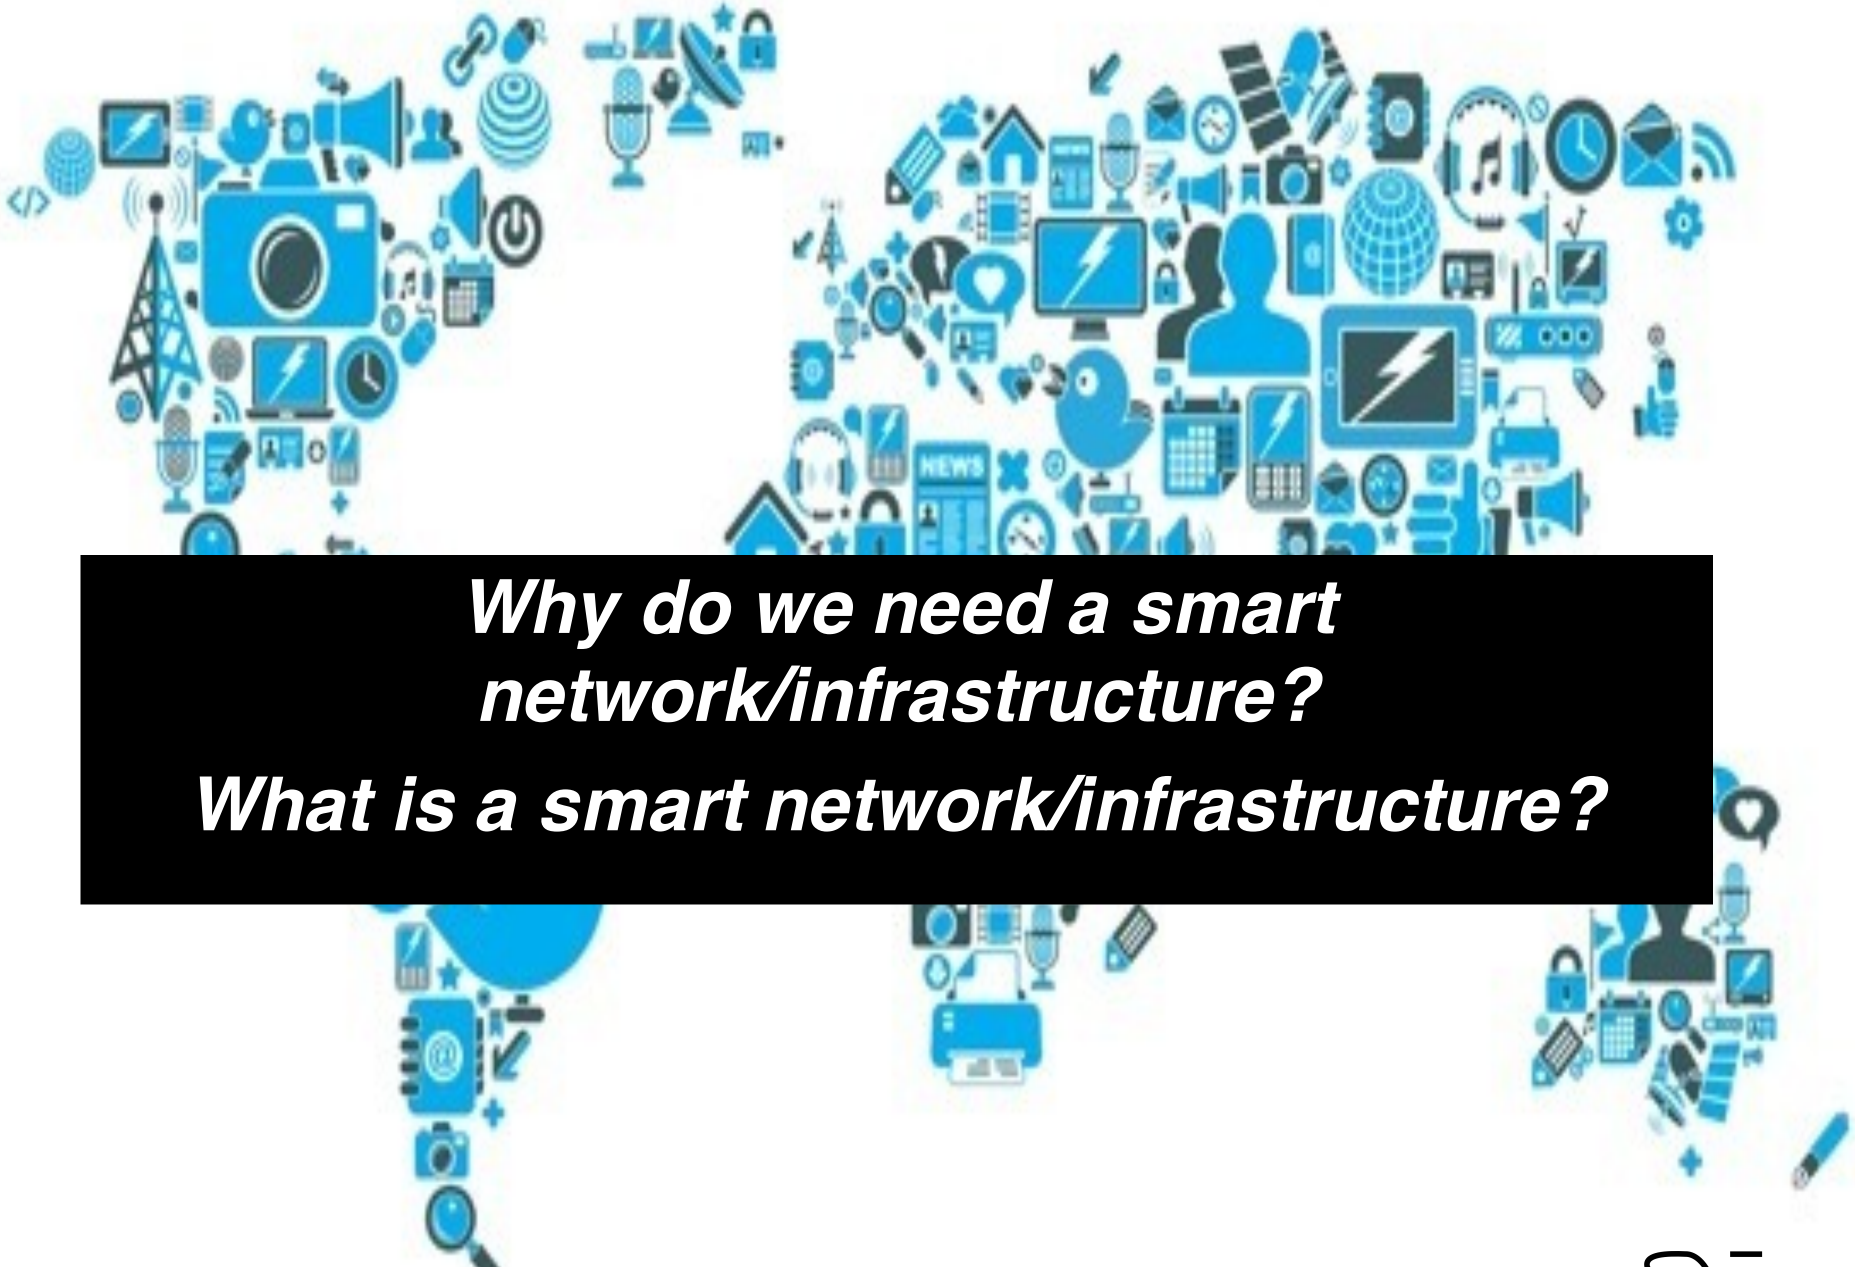

Why do we need a smart network/infrastructure?

What is a smart network/infrastructure?

We exploit **richer network and infrastructure descriptions** to *deliver federated network and clouds services.*

We leverage the **SDN paradigm** to *align network behavior closer and applications needs*.

We get:

- more energy-efficient,
- more secure,
- more adaptable networks

Ontology families

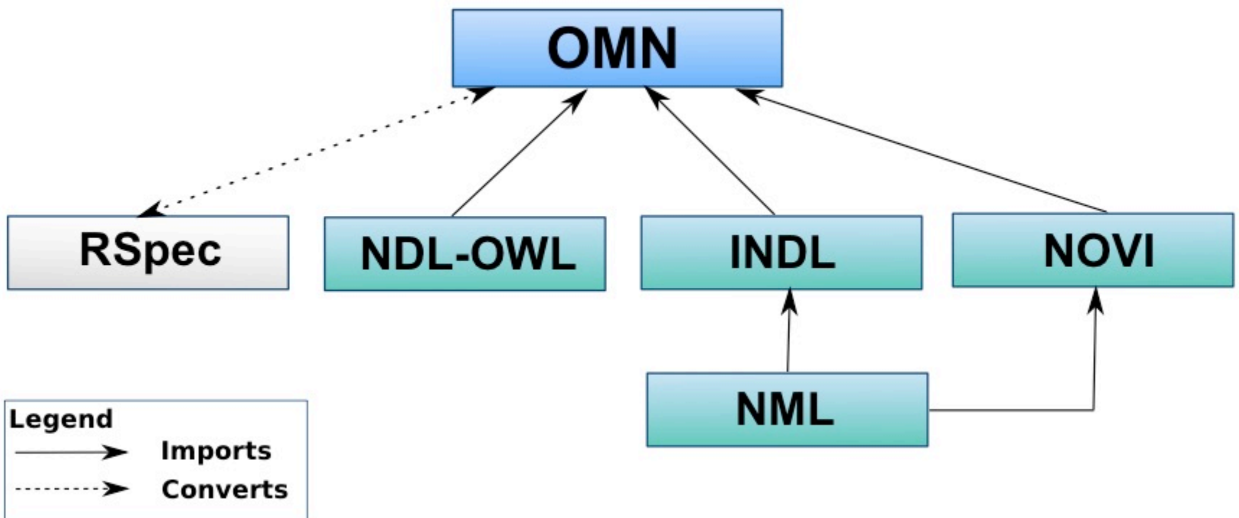

M. Ghijsen, , J. v/d Ham, P. Grosso, C. Dumitru, H. Zhu, Z. Zhao, C. De Laat
"A Semantic-Web Approach for Modeling Computing Infrastructures," Computers and Electrical Engineering, vol. 39 (8), pp. 2553–2565, 2013

OMN – Open Multinet

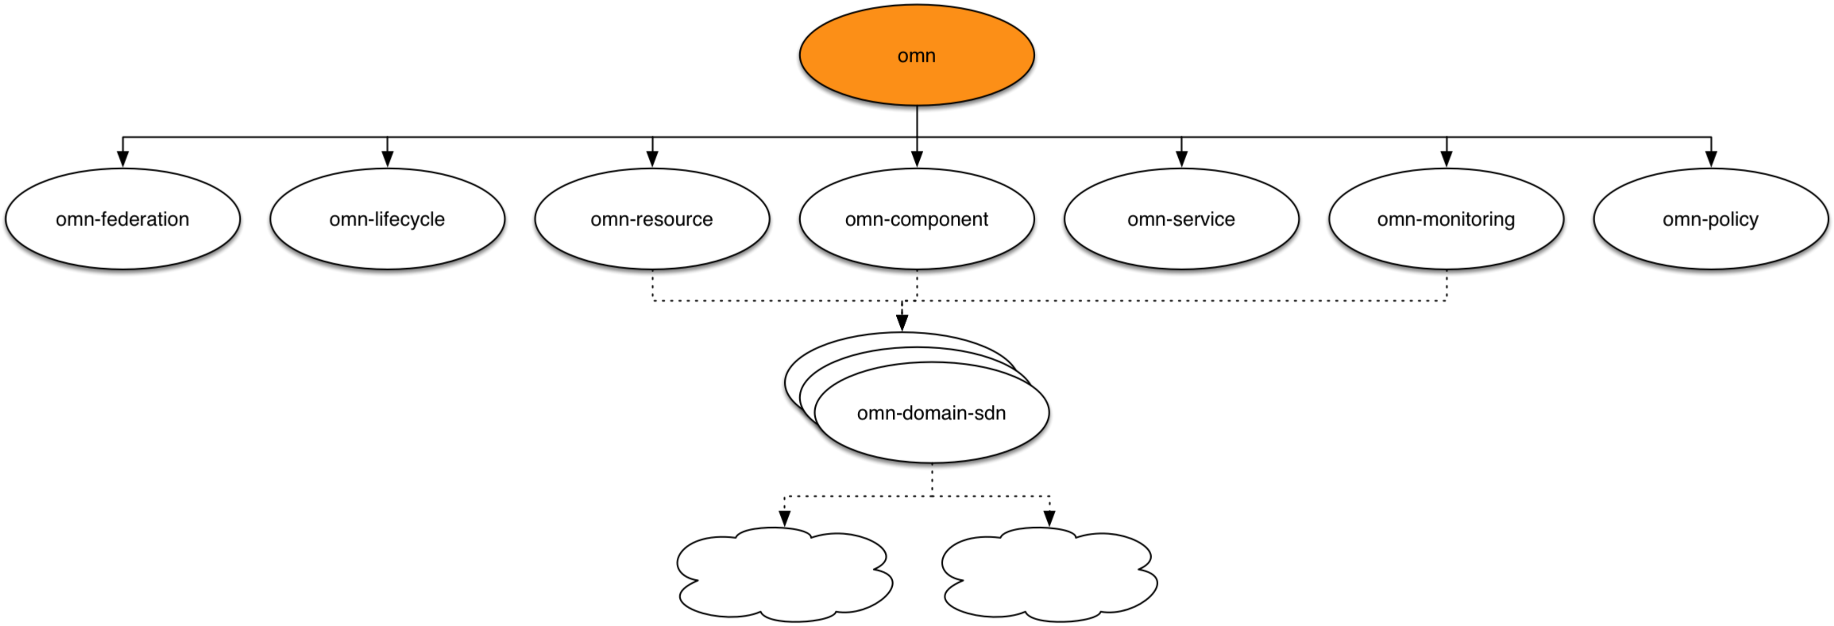

https://github.com/open-multinet
https://github.com/open-multinet/playground-rspecs-ontology/tree/master/omnlib/ontologies

A. Willner, C. Papagianni, M. Giatili, P. Grosso, M. Morsey, Al-Hazmi Y., I. Baldin
The Open-Multinet Upper Ontology - Towards the Semantic-based Management of Federated Infrastructures
The 10th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (TRIDENTCOM 2015)

Transitioning from OpenNaas to SemNaaS

Developed a semantic enabled Network-as-a-Service (NaaS) system

Applied NML2/OMN on that system.

M Morsey, H Zhu, I Canyameres, P Grosso
SemNaaS: Add Semantic Dimension to the Network as a Service
The Semantic Web: ESWC 2015 conference

UNIVERSITEIT VAN AMSTERDAM

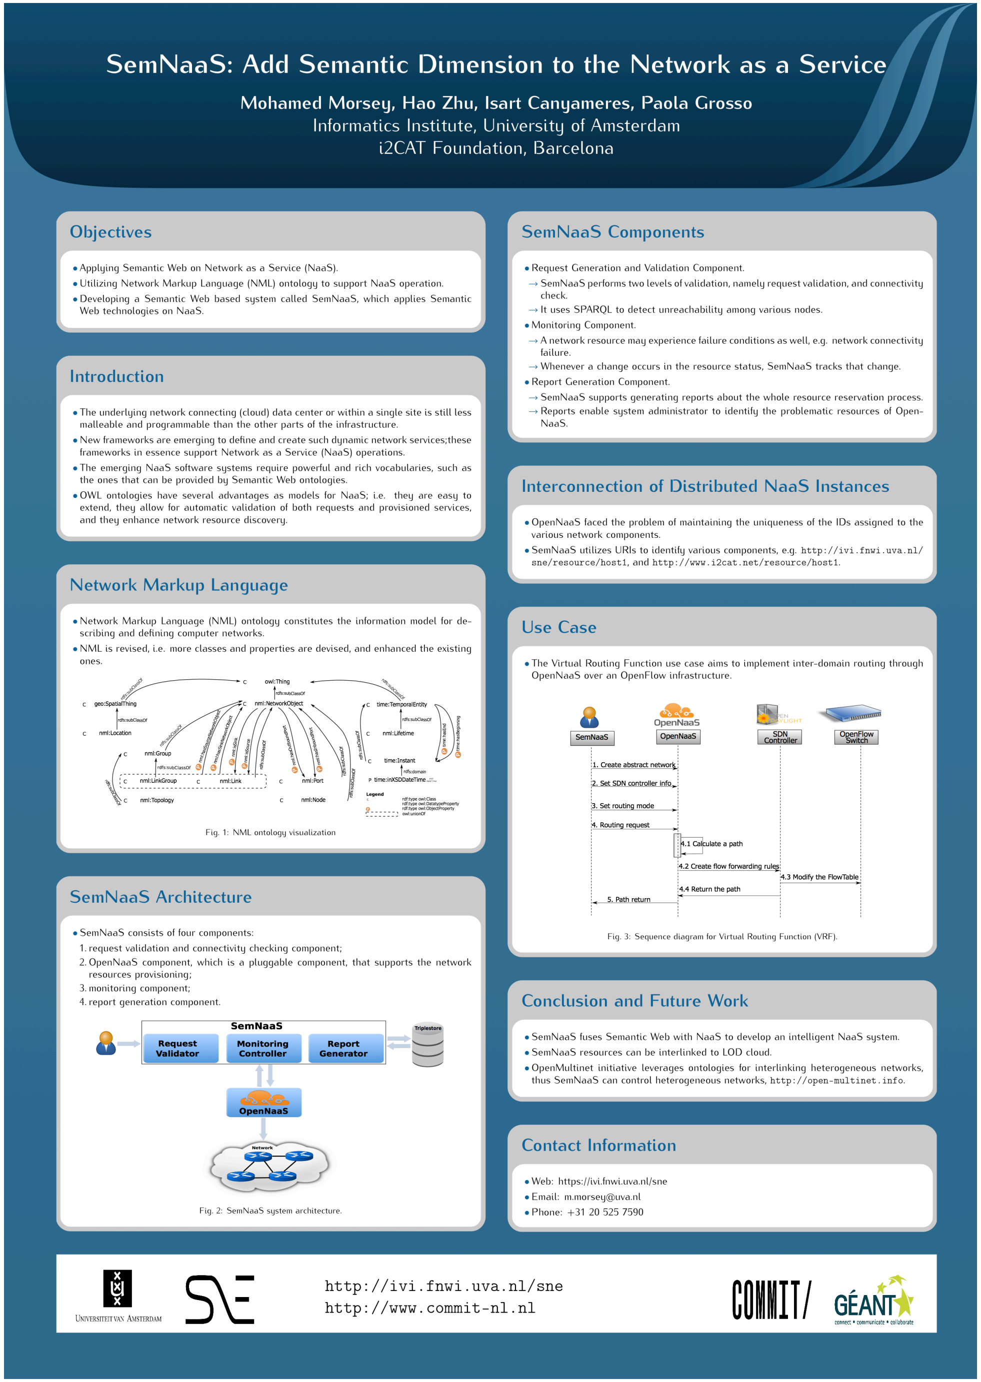

SWSDI 2016 Workshop

The First International Workshop on Semantic Web for Federated Software Defined Infrastructures (SWSDI2016)
Co-located with ESWC 2016, Anissaras, Crete, Greece. May, 2016

Contents [hide]

Workshop Objectives

The main objective of SWSDI 2016 is to study the applicability and maturity of Semantic Web based methodologies for modelling the newly emerging software-defined (computing and networking) infrastructures, particularly federated infrastructures and federated Clouds. Furthermore, SWSDI 2016 aims to identify how the Semantic Web surpasses other approaches, such as the exchange of simple XML or JSON serialized tree data structures.

Workshop organizers
Jorge Cardoso, University of Coimbra, Portugal
Paola Grosso, University of Amsterdam, The Netherlands
Mohamed Morsey, University of Amsterdam, The Netherlands
Alexander Willner, TU Berlin, Germany

Deadlines
Paper submission : Friday March 4th, 2016.
Notification of acceptance: Friday April 1st, 2016.

Emerging interest in SDN for energy efficiency

- Emerging studies improve on the energy consumption of servers by the VM migration.
- Some change the OpenFlow protocol to be energy-aware.
- All of them are implemented in intra-data center scale.
- All have a fixed initial traffic matrix.

Moghaddam, F. A., Lago, P., & Grosso, P.
Energy-Efficient Networking Solutions in Cloud-Based Environments: A Systematic Literature Review.
ACM Computing Surveys (CSUR), 47(4), 64. (2015)

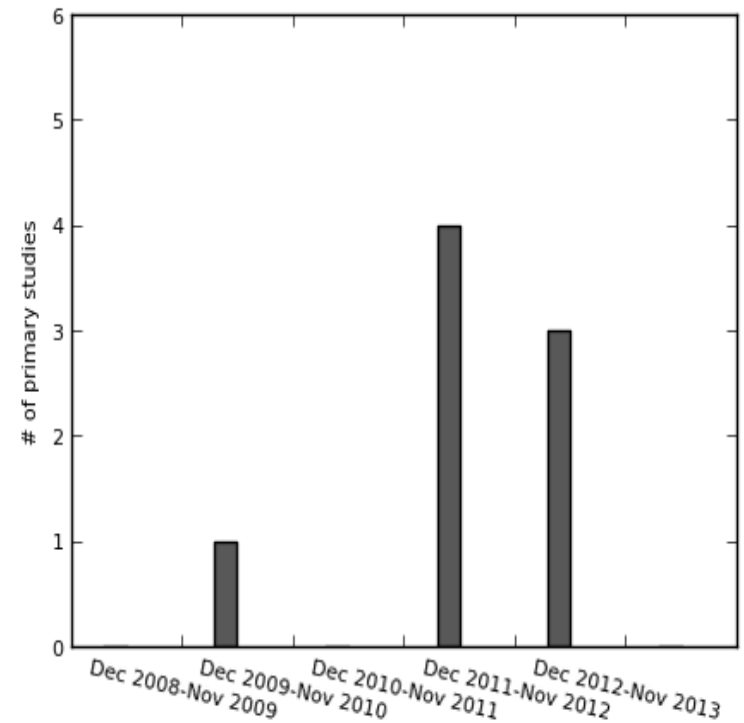

Yearly distribution of the OpenFlow technique adoption by decision frameworks from December 2008 to November 2013

Green routing with SDN

Make a routing decision to aggregate traffic over a subset of links and devices in over-provision networks and switch off unused network components

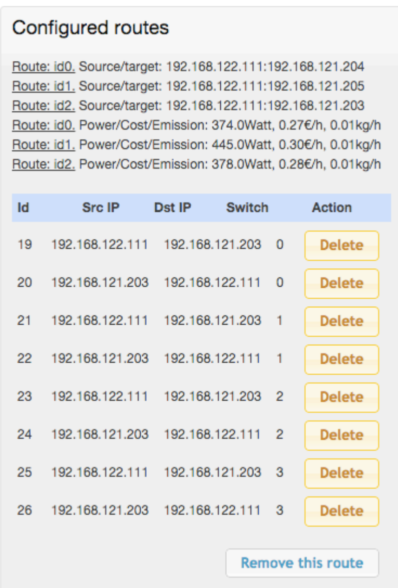

Configured routes

Route: id0. Source/target: 192.168.122.111:192.168.121.204
Route: id1. Source/target: 192.168.122.111:192.168.121.205
Route: id2. Source/target: 192.168.122.111:192.168.121.203
Route: id0. Power/Cost/Emission: 374.0Watt, 0.27€/h, 0.01kg/h
Route: id1. Power/Cost/Emission: 445.0Watt, 0.30€/h, 0.01kg/h
Route: id2. Power/Cost/Emission: 378.0Watt, 0.28€/h, 0.01kg/h

Id	Src IP	Dst IP	Switch	Action
19	192.168.122.111	192.168.121.203	0	Delete
20	192.168.121.203	192.168.122.111	0	Delete
21	192.168.122.111	192.168.121.203	1	Delete
22	192.168.121.203	192.168.122.111	1	Delete
23	192.168.122.111	192.168.121.203	2	Delete
24	192.168.121.203	192.168.122.111	2	Delete
25	192.168.122.111	192.168.121.203	3	Delete
26	192.168.121.203	192.168.122.111	3	Delete

Remove this route

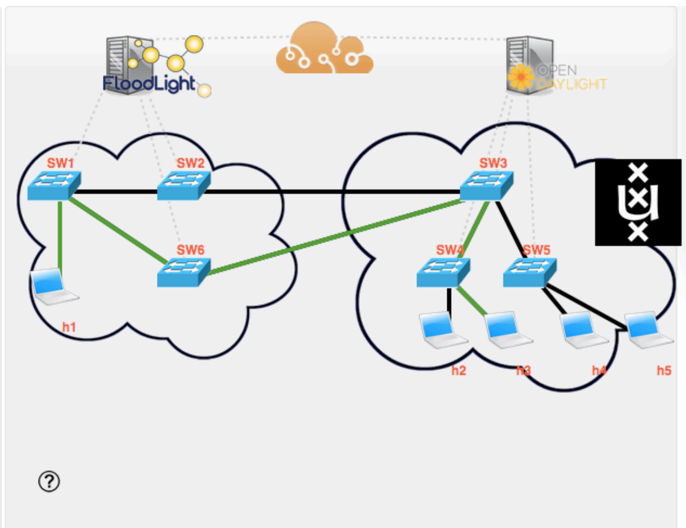

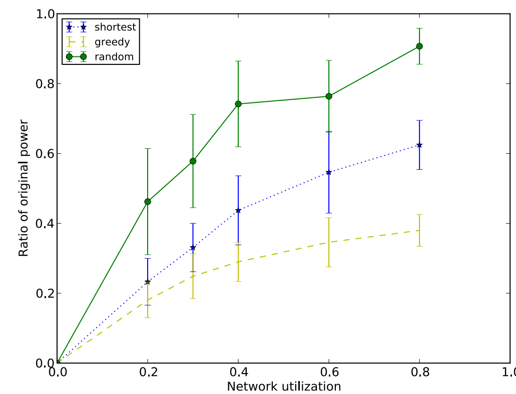

H. Zhu. X. Liao, C. de Laat and P. Grosso. (2015)
Joint flow routing-scheduling for energy efficient software defined data center networks
(to appear in Elsevier Journal of Network and Computer Applications)

Universiteit van Amsterdam

System and Network Engineering

How much savings?

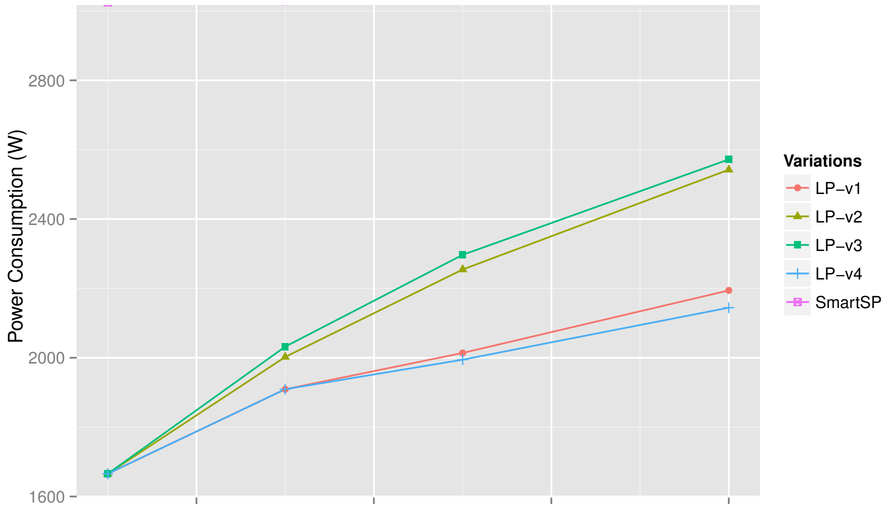

We adopt linear programming to determine how to program flows in the network.

We show that in FatTree networks, where switches can save up to 60% of power in sleeping mode, we can achieve 15% minimum improvement assuming a one-to-one traffic scenario.
Two of our algorithm variations privilege performance over power saving and still provide around 45% of the maximum achievable savings.

F. A.. Moghaddam and P. Grosso. (2016)
Linear Programming Approaches for Power Savings in Software-defined Networks
(Under review at the NetSoft conference 2016)

4th International Conference on ICT for Sustainability (ICT4S)

Aug 30 - Sep 2, 2016 - Amsterdam, The Netherlands

Home Calls & Dates Program Team ict4s.org Venue Search

General chair
Anwar Osseyran, SURFsara & University of Amsterdam, The Netherlands

Program chairs
Paola Grosso, University of Amsterdam, The Netherlands
Patricia Lago, VU University Amsterdam, The Netherlands

Deadlines
Paper submission : Monday April 11th, 2016.
Notification of acceptance: Tuesday May 31st, 2016.

Universiteit van Amsterdam

System and Network Engineering

SARNET

SARNET
Secure Autonomous Response Networks

Ralph Koning (UvA), Ameneh Deljoo (UvA), Robert Meijer (TNO), Leon Gommans (KLM),
Tom van Engers (UvA), Rodney Wilson (Ciena), Cees de Laat (UvA)

SARNET

SARNET, Secure Autonomous Response NETworks, is a project funded by the Dutch Research Foundation. The University of Amsterdam, TNO, KLM, and Ciena conduct research on **automated methods against attacks** on computer **network infrastructure**.

The research goal of SARNET is to obtain the knowledge to create ICT systems that

- **model** the system's state based on the emerging behaviour of its components,
- discover by observations and **reasoning** if and how an attack is developing and calculate the associated risks,
- have the **knowledge** to calculate the effect of countermeasures on states and their risks, and
- choose and **execute** the most effective **countermeasure**.

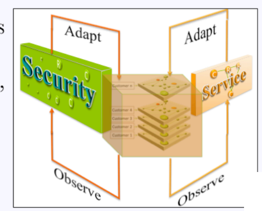

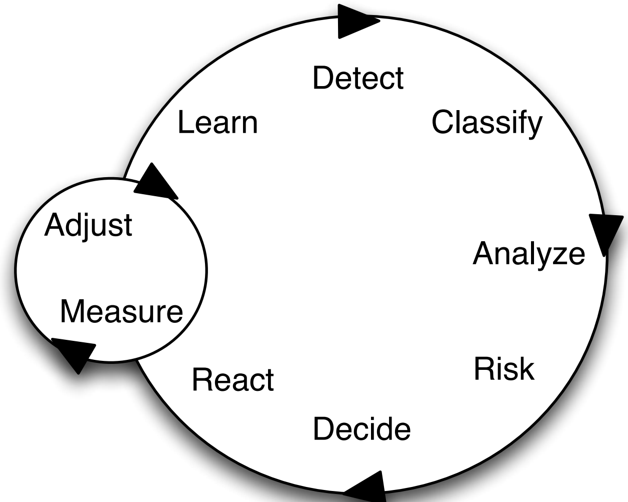

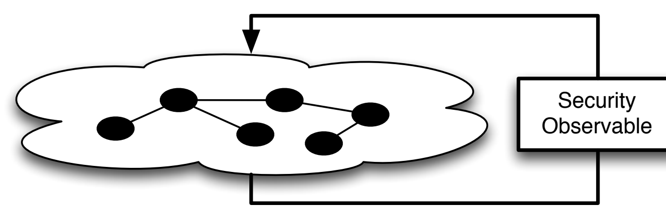

Universiteit van Amsterdam

System and Network Engineering

Container networking

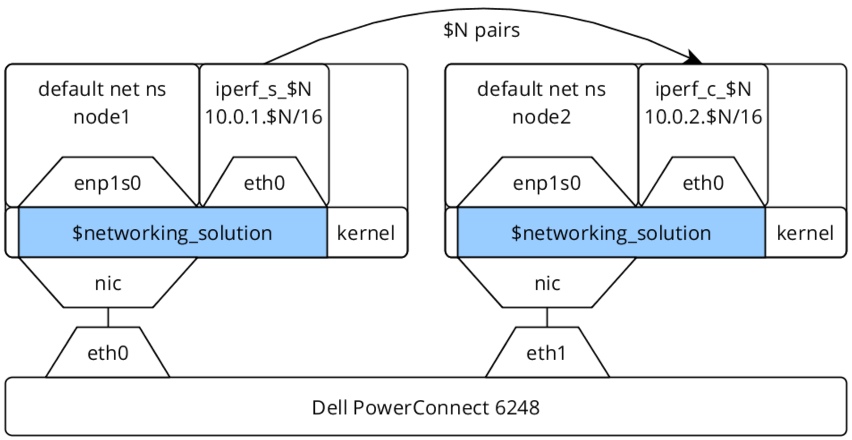

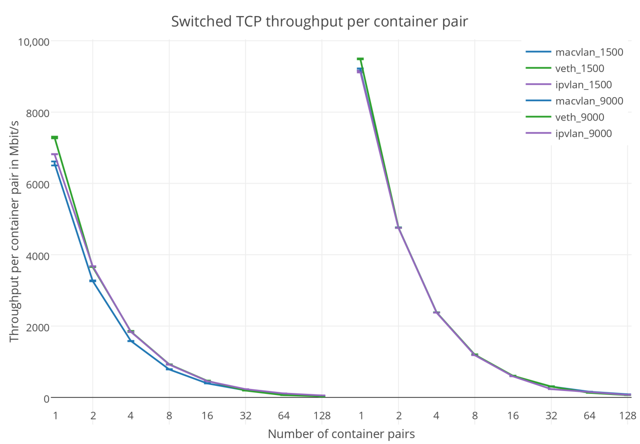

We benchmarked three kernel modules: **veth**, **macvlan** and **ipvlan**, to quantify their respective raw TCP and UDP performance and scalability.

Our results show that the macvlan kernel module outperforms all other solutions in raw performance. All kernel modules seem to provide sufficient scalability to be deployed effectively in multi- containers environments.

J. Claassen, R. Koning and P. Grosso. (2016)
Linux containers networking: performance and scalability of kernel modules
Accepted at NOMS 2016

Open research directions

- Can we create Semantic NaaS in federated environments?
- How can software services exploit SDN for energy efficiency of the applications?
- Are containers and (SDN) overlays the solution for secure networks?

Dr. Zhiming Zhao

Senior Researcher
System and Network Engineering
University of Amsterdam

EU H2020 SWITCH (**Scientific Coordinator**)
EU H2020 ENVRI^{PLUS} (**Theme Leader**)
EU H2020 VRE4EIC (**Task Leader**)

Email: z.zhao@uva.nl
Web: http://staff.fnwi.uva.nl/z.zhao/

Modeling, Developing and Controlling *Quality Critical Distributed Systems* on Programmable Infrastructures.

Environmental Research Infrastructures Providing Shared Solutions for Science and Society

Research topics

- **Programming, provisioning** and **controlling** models for time critical applications on programmable infrastructures
- **Interoperable** research infrastructures for system level of **big data sciences**
- **Virtual research environments** for large scale research communities

Programming, provisioning and **controlling** models for time critical applications on programmable infrastructures

Time critical applications

- Have **very high** business potential or social impacts, e.g.,
  - live event broadcasting,
  - disaster early warning, and
  - real-time business collaboration;
- have **very critical** quality requirements for services, e.g.,
  - video quality, system interaction, or data delivery;
- But are **very expensive in** implementation and operation.

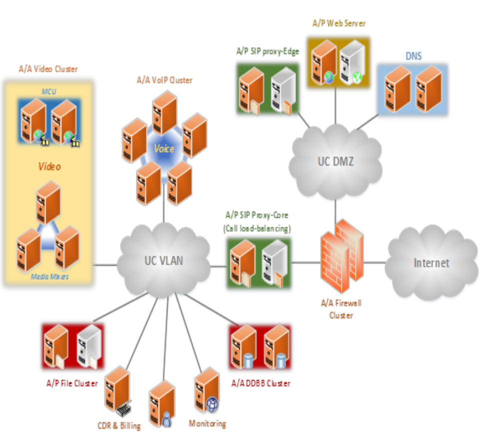

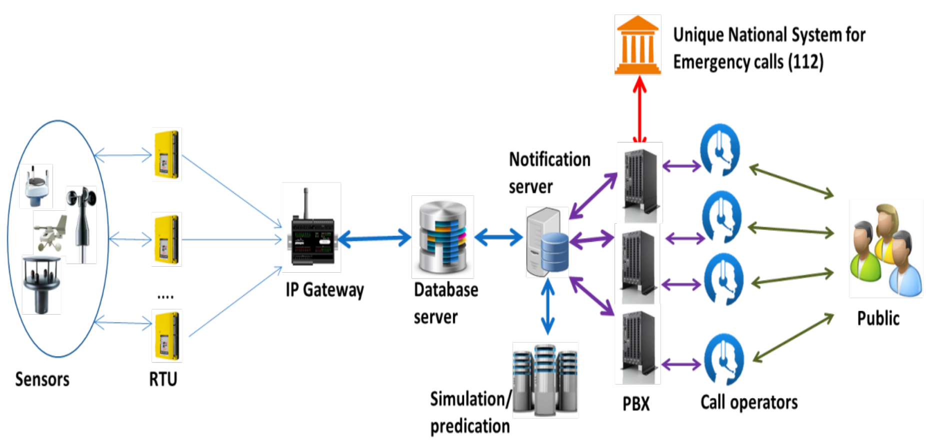

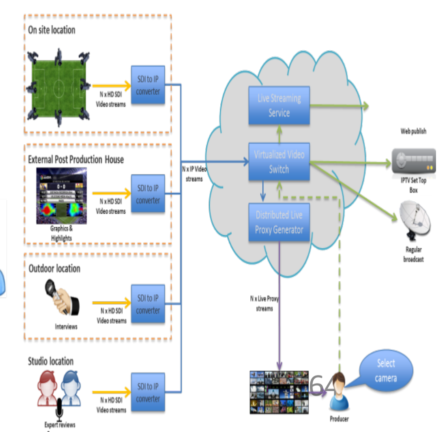

Challenges and difficulties

- Development challenges
  - QoS/QoE between different levels
  - Verification
  - Optimization
- Provisioning challenges
  - Infrastructure customization
  - Fast provisioning
- Operation challenges
  - Run-time monitoring adaptation
  - Autonomous/human-in-the-loop control

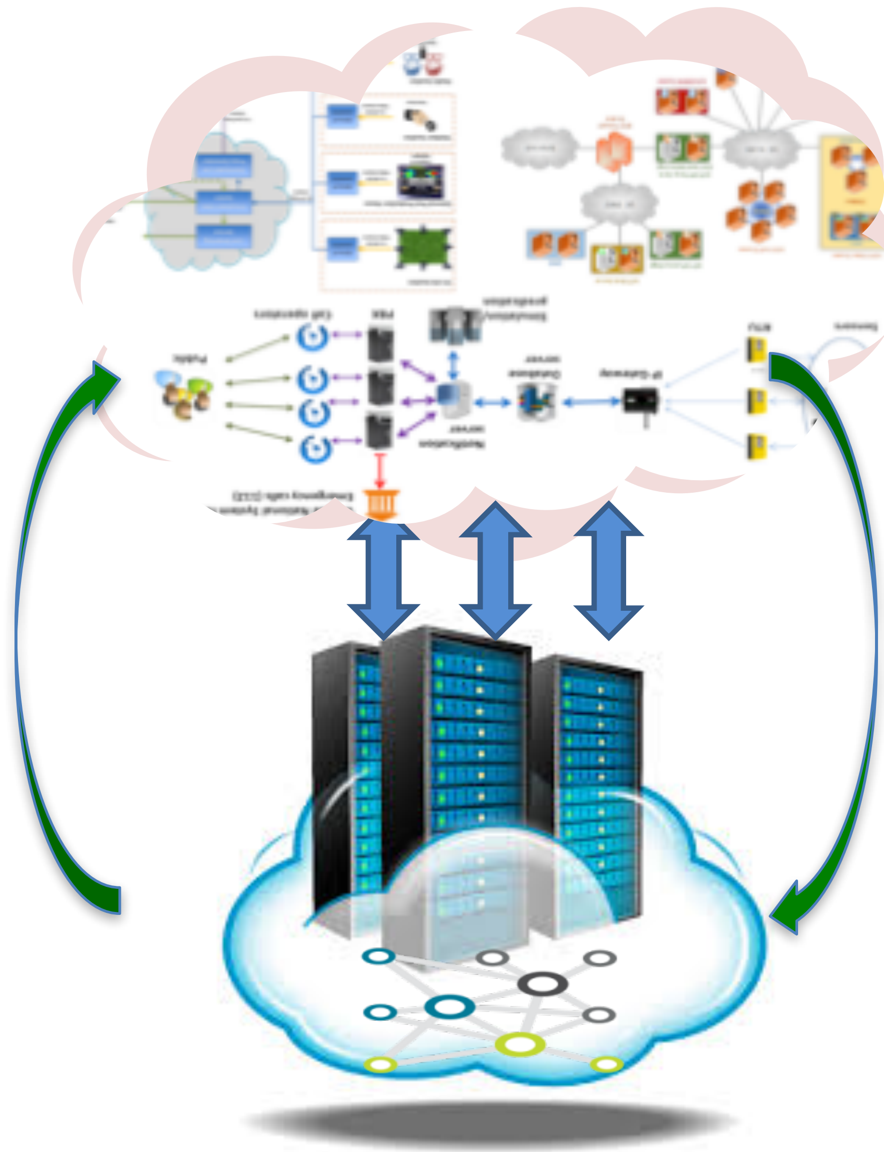

The SWITCH approach

SWITCH addresses the **entire life-cycle** of time-critical, self-adaptive Cloud applications by developing **new middleware and front-end tools** to enable users to **specify** their time-critical requirements for an application interactively using a direct manipulation user interface, **deploy** their applications and **adapt** the infrastructure to changing requirements **at runtime** either automatically (using the specified requirements) or by human intervention if desired.

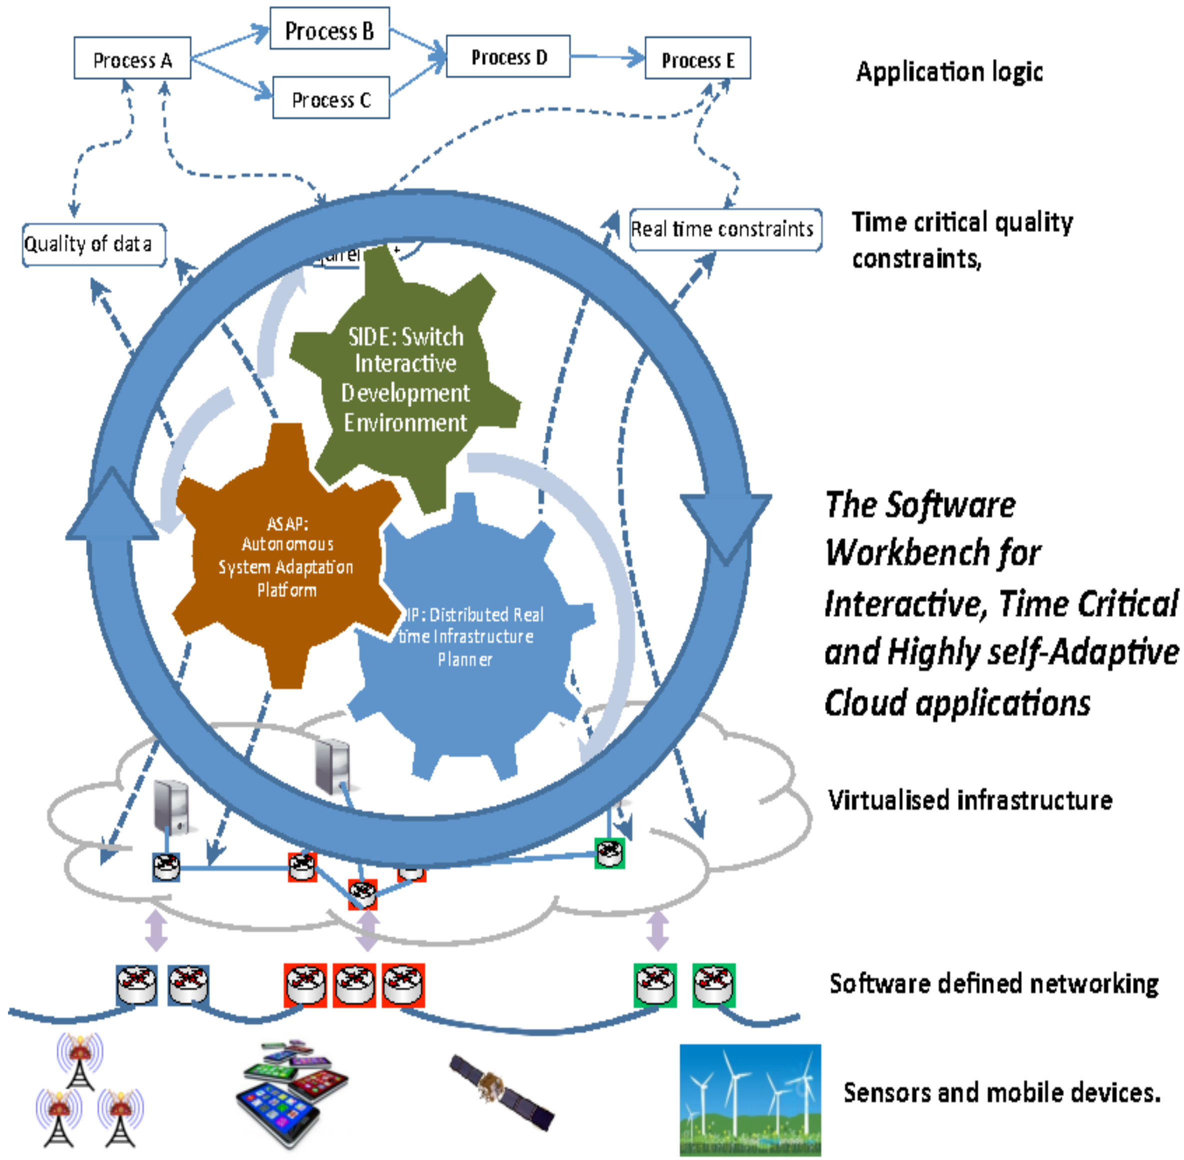

Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud applications

- EU H2020 ICT RIA project
- Funding 3M, 6 partners
- Coordinator: University of Amsterdam
- Scientific coordinator: **Zhiming Zhao**
- Duration: 3 years

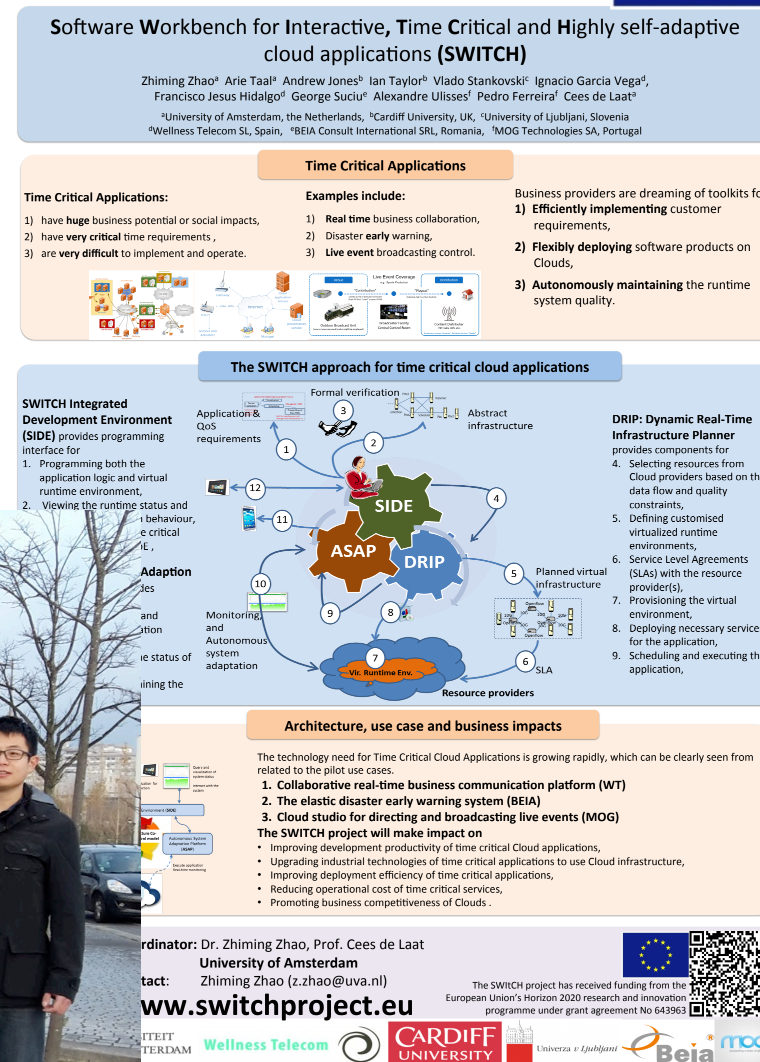

ENVRI plus

Environmental Research Infrastructures Providing Shared Solutions for Science and Society

Interoperable research infrastructures for system level of **big data sciences**

Motivation: societal challenges- system level of environmental sciences

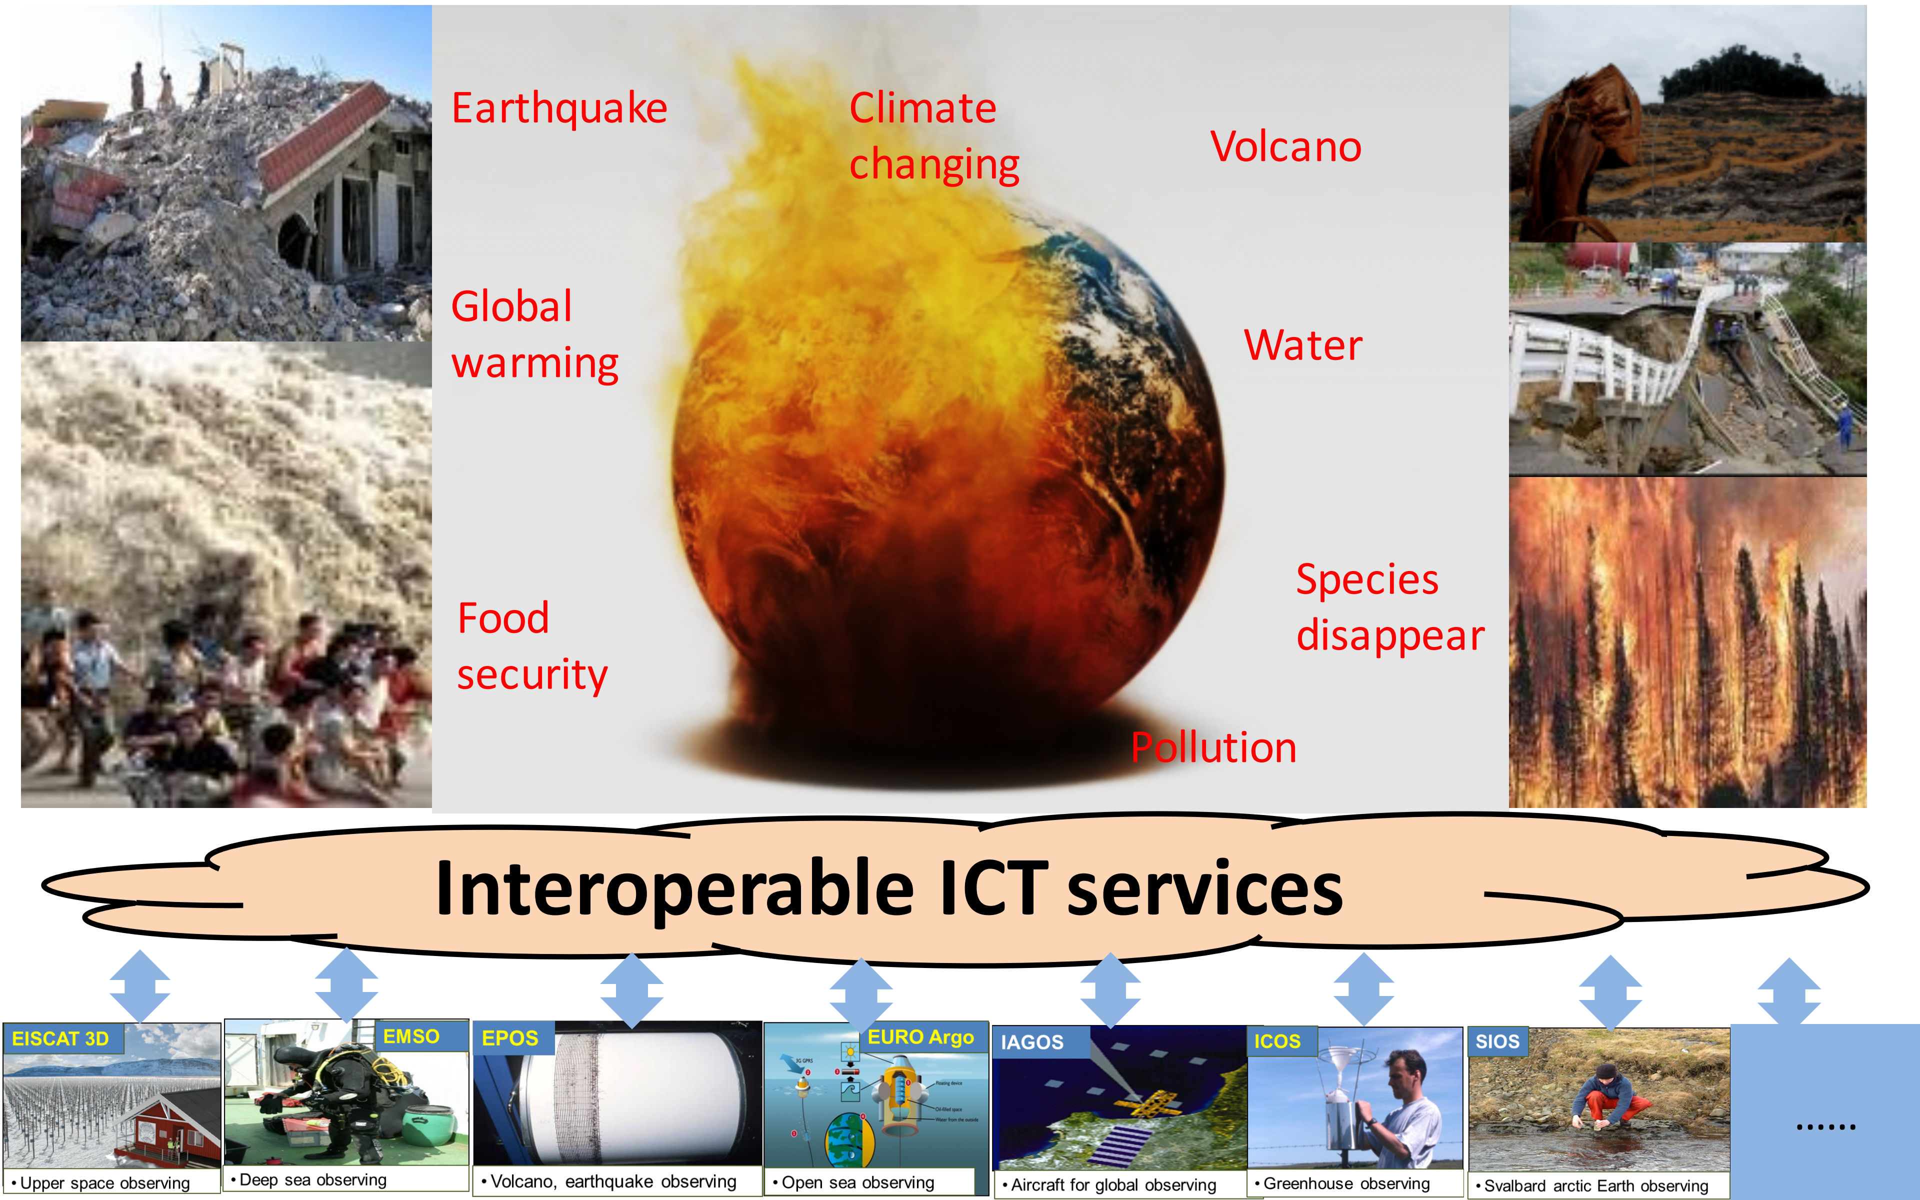

Interoperable infrastructures for environmental sciences

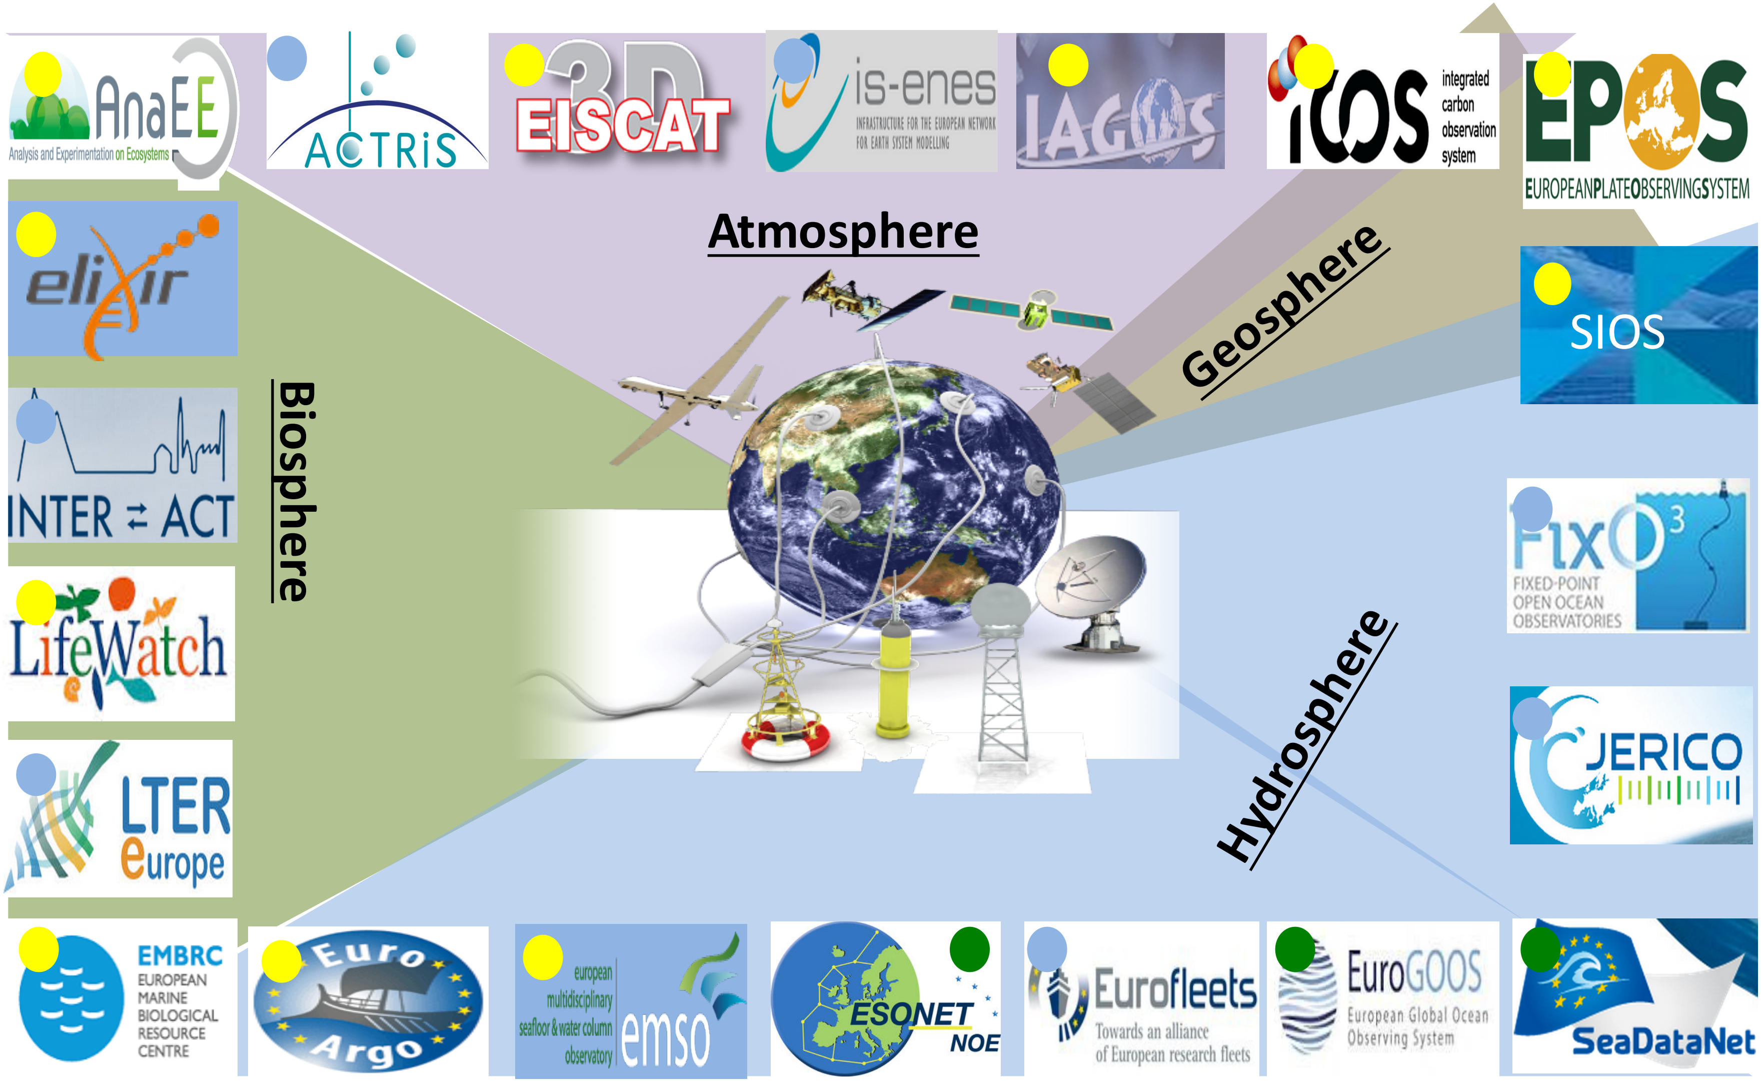

Research Infrastructures, I3, and ESFRIs in environmental Sciences

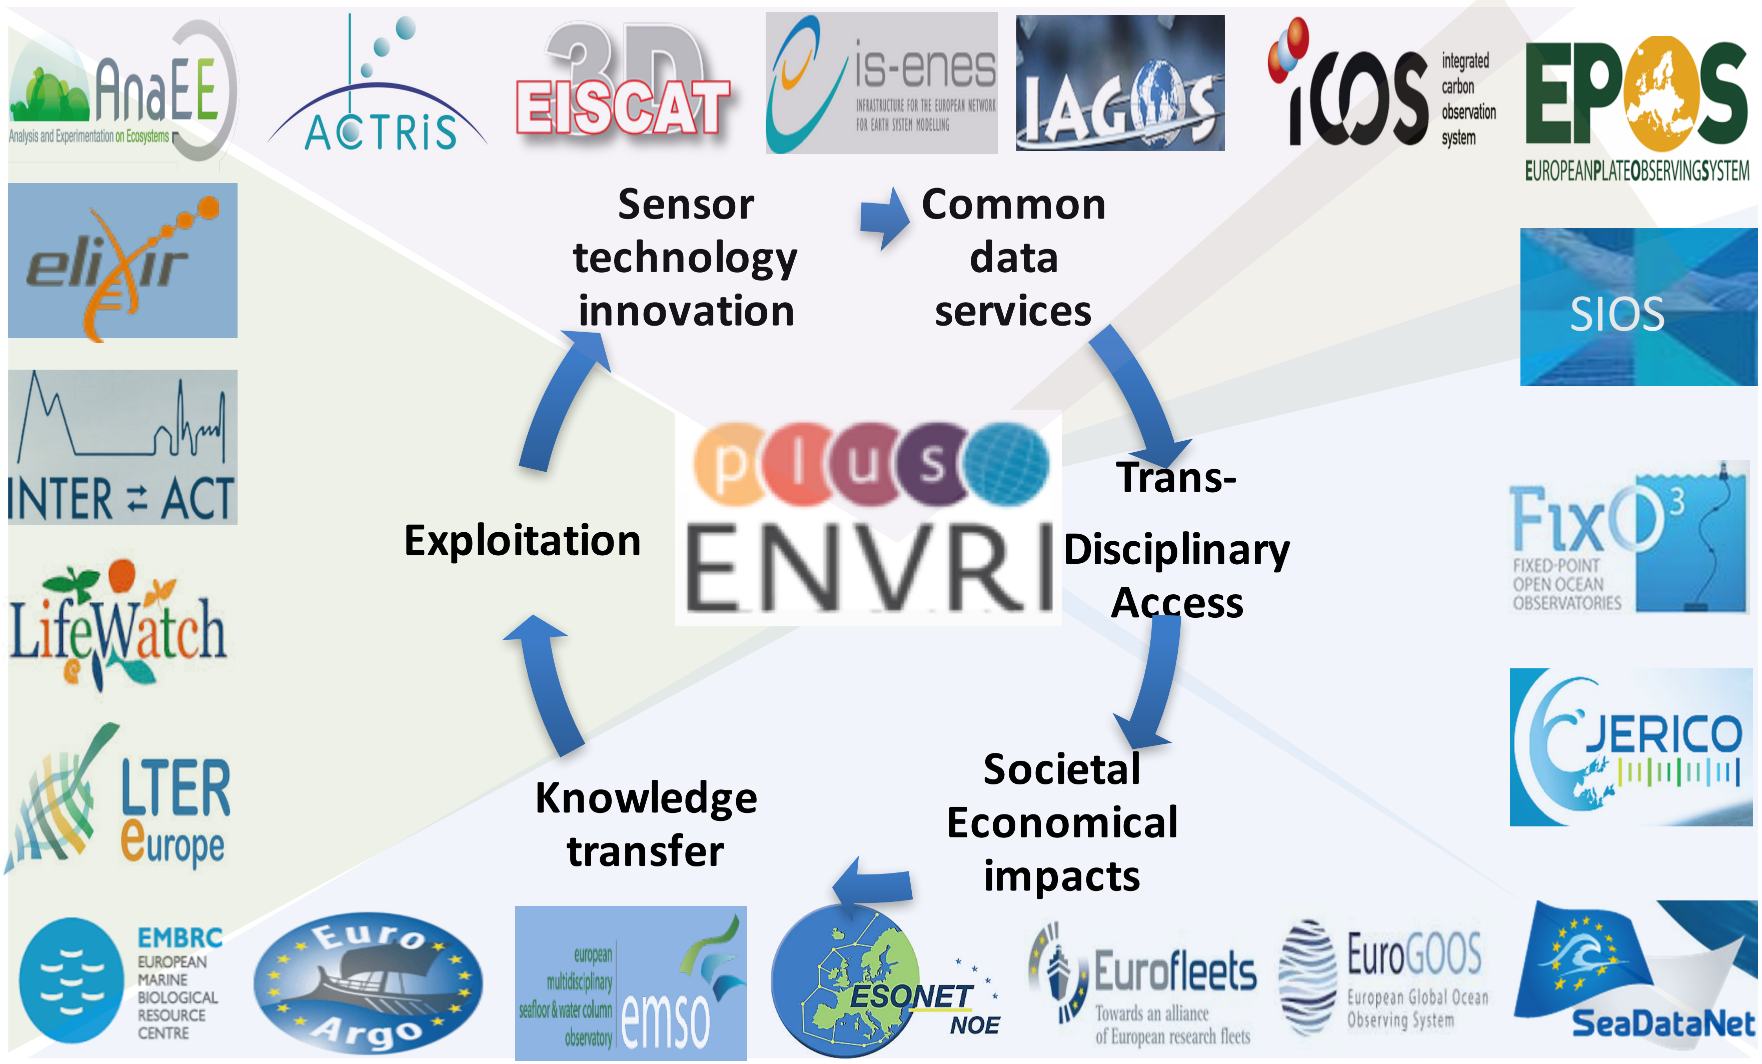

ENVRI plus

Environmental Research Infrastructures Providing Shared Solutions for Science and Society

Common data services

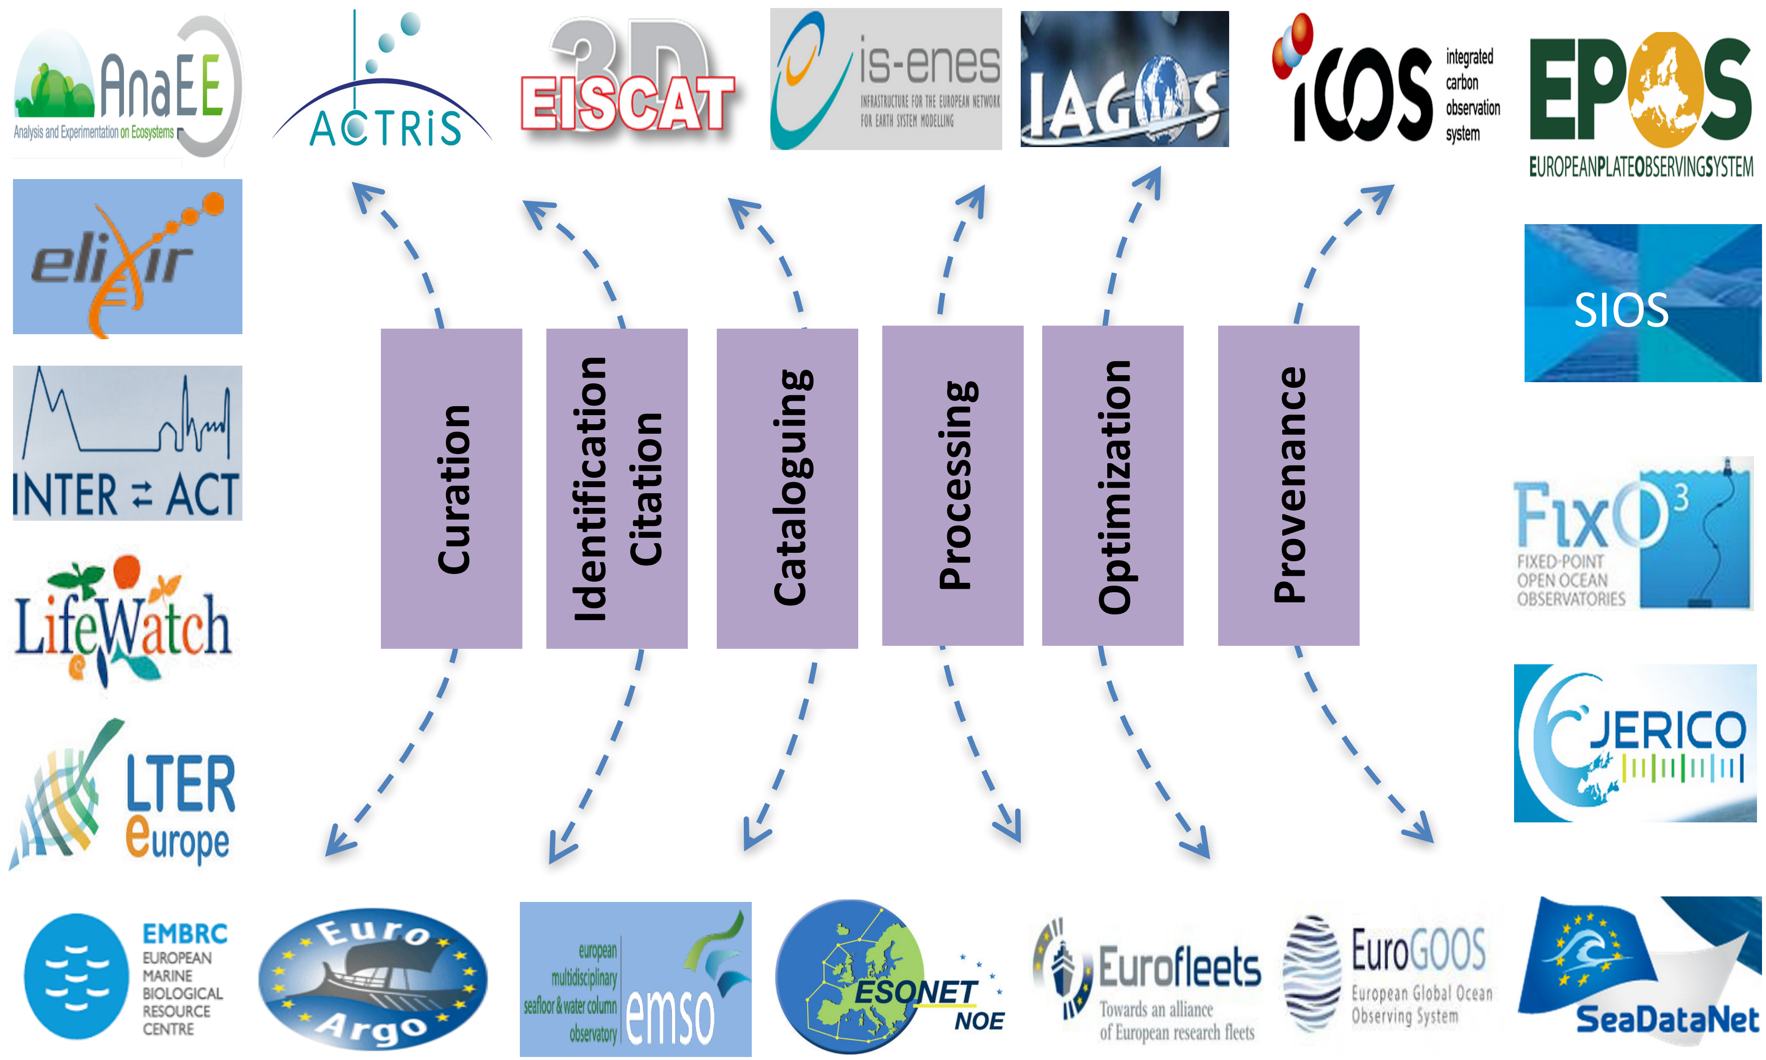

ENVRI plus

Environmental Research
Infrastructures Providing Shared
Solutions for Science and Society

Data for science theme in ENVRI^{PLUS}

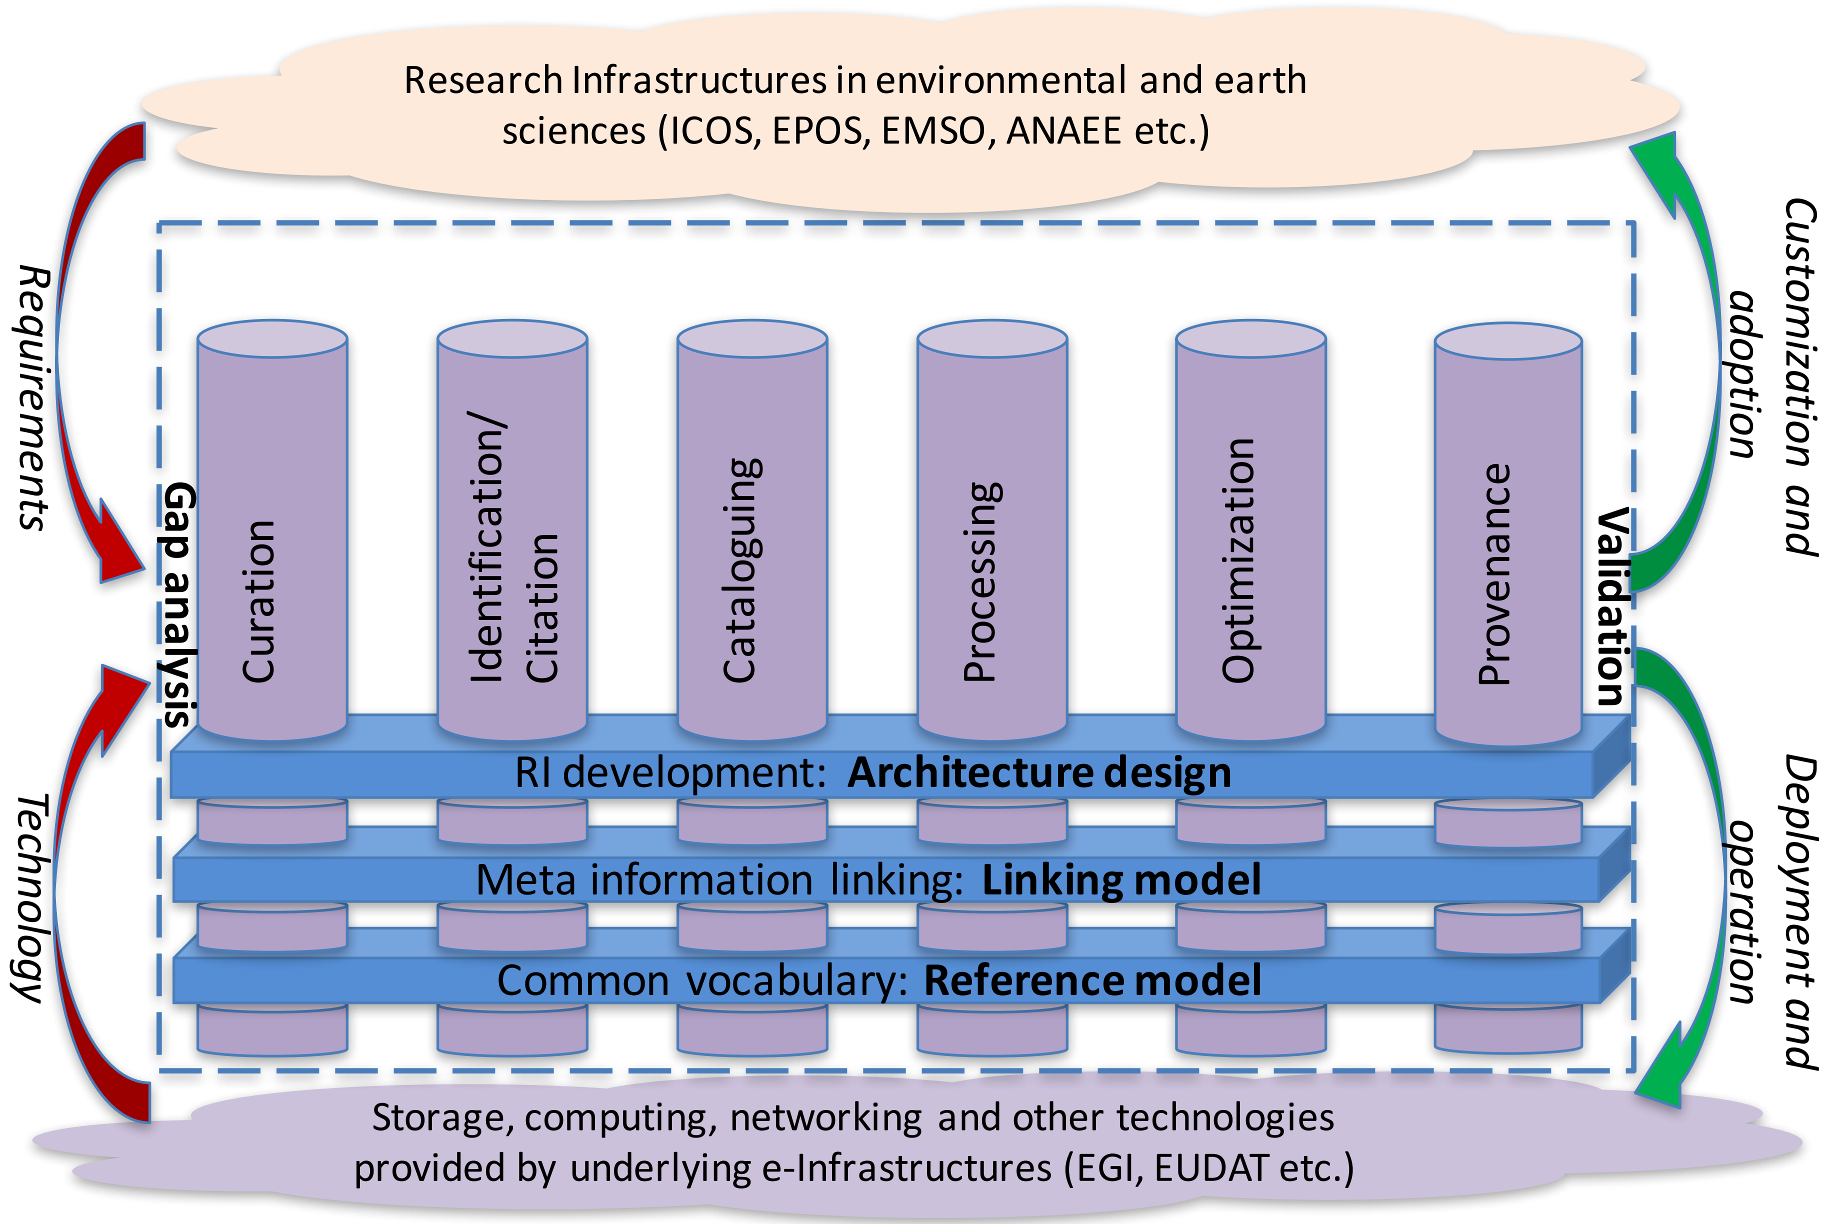

Use cases: Data processing from different RIs

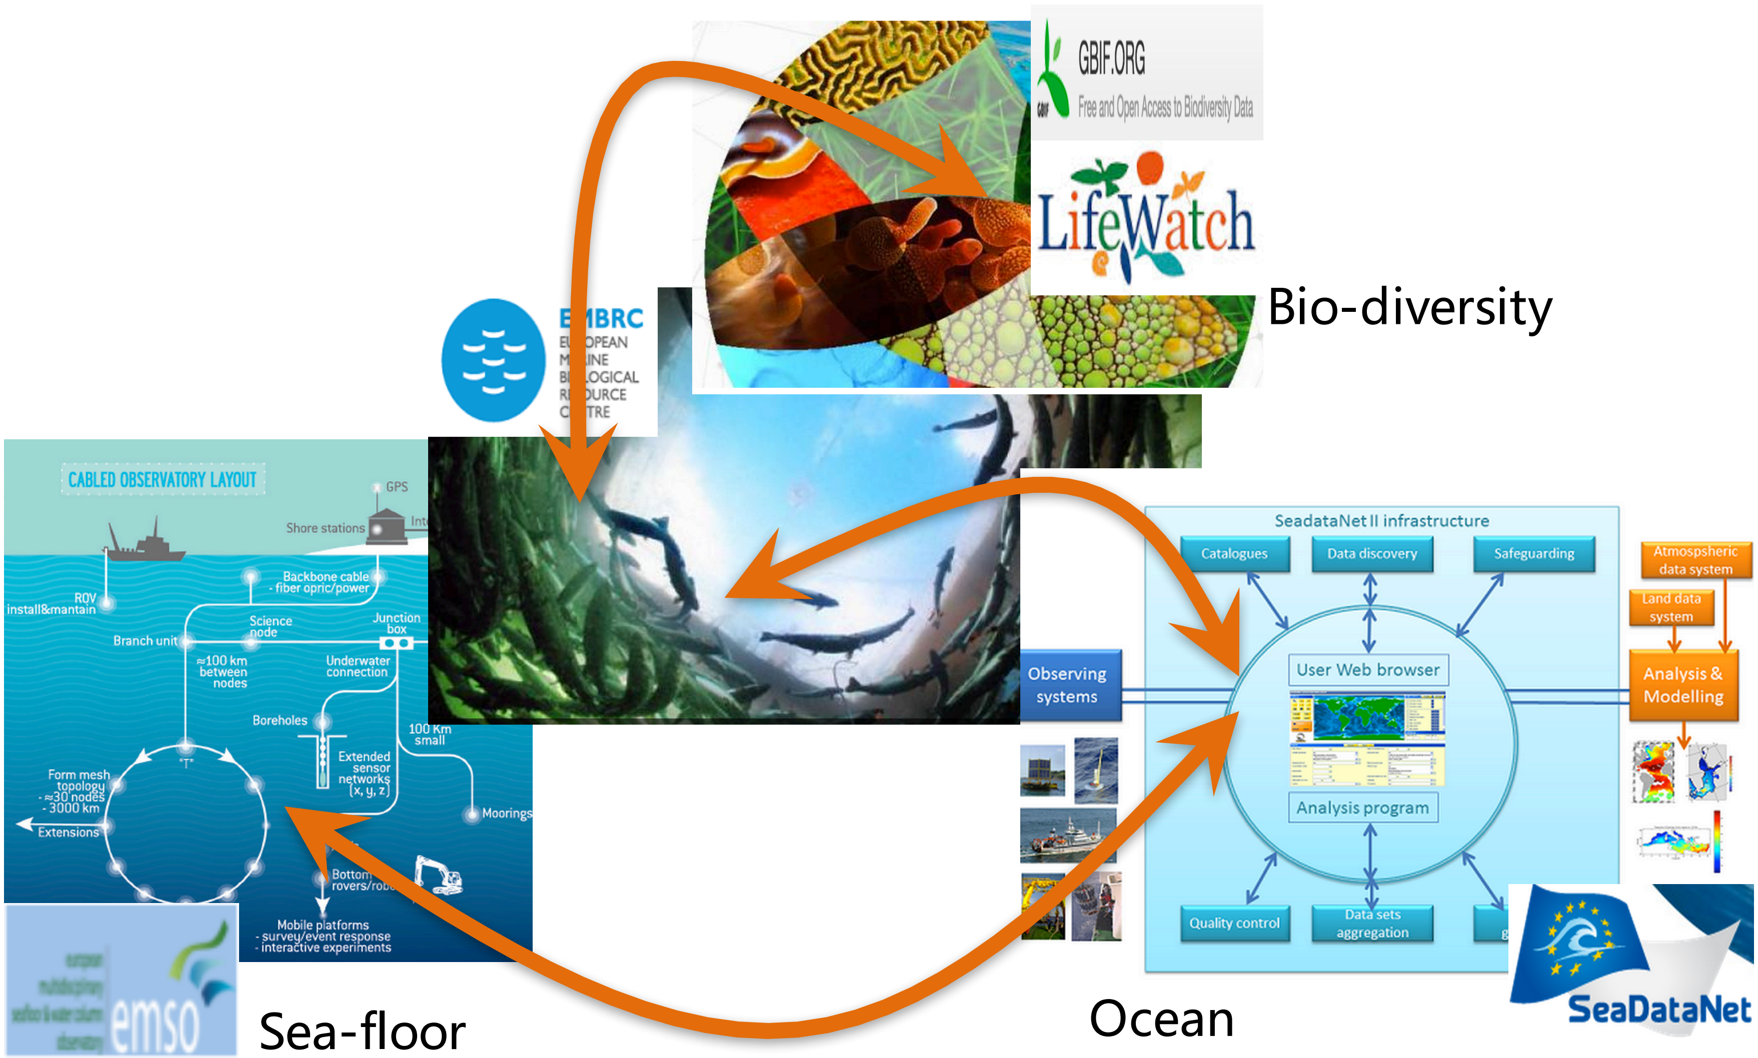

ENVRI plus

Environmental Research
Infrastructures Providing Shared
Solutions for Science and Society

ENVRI^{PLUS}

- ENVRI^{PLUS}: www.envriplus.eu, 15M Euro, 4 years
- The data for science theme: 5M Euro
- Partners: 37
- Theme leader: University of Amsterdam (**Zhiming Zhao**)

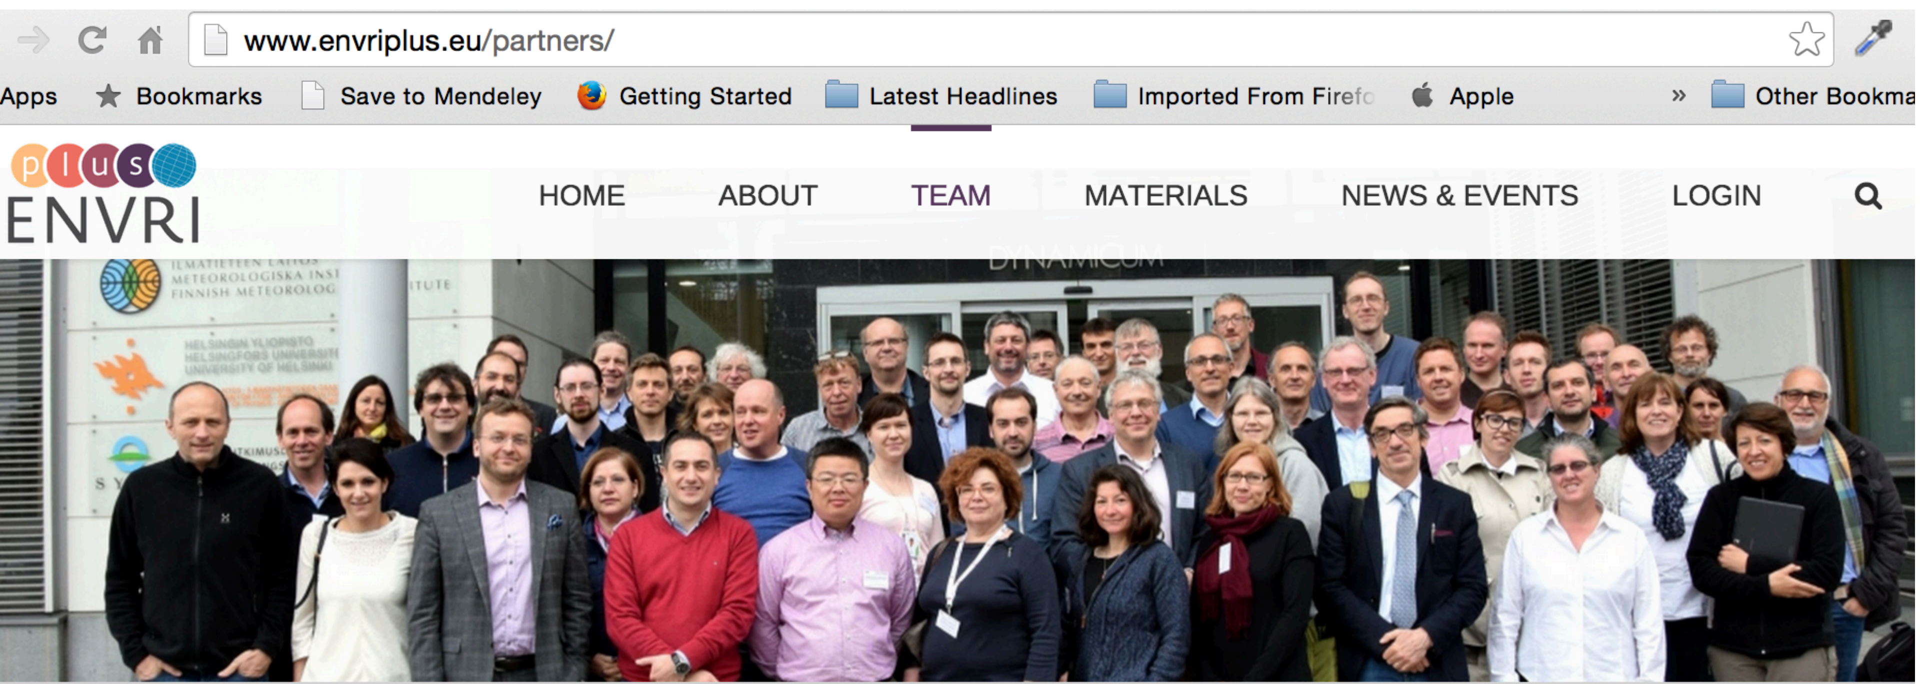

V R E 4 E I C

Virtual research environments for large scale research communities

VRE4EIC

- EU H2020: VRE4EIC project
- 4.75M Euro in total, UvA shares 0.47MEuro
- Virtual Research Environment (VRE)
- Our contribution: lead the task of research sustainability, exploitation of VRE development to the ENVRI^{PLUS} community.

User environments bridging communities and the infrastructures

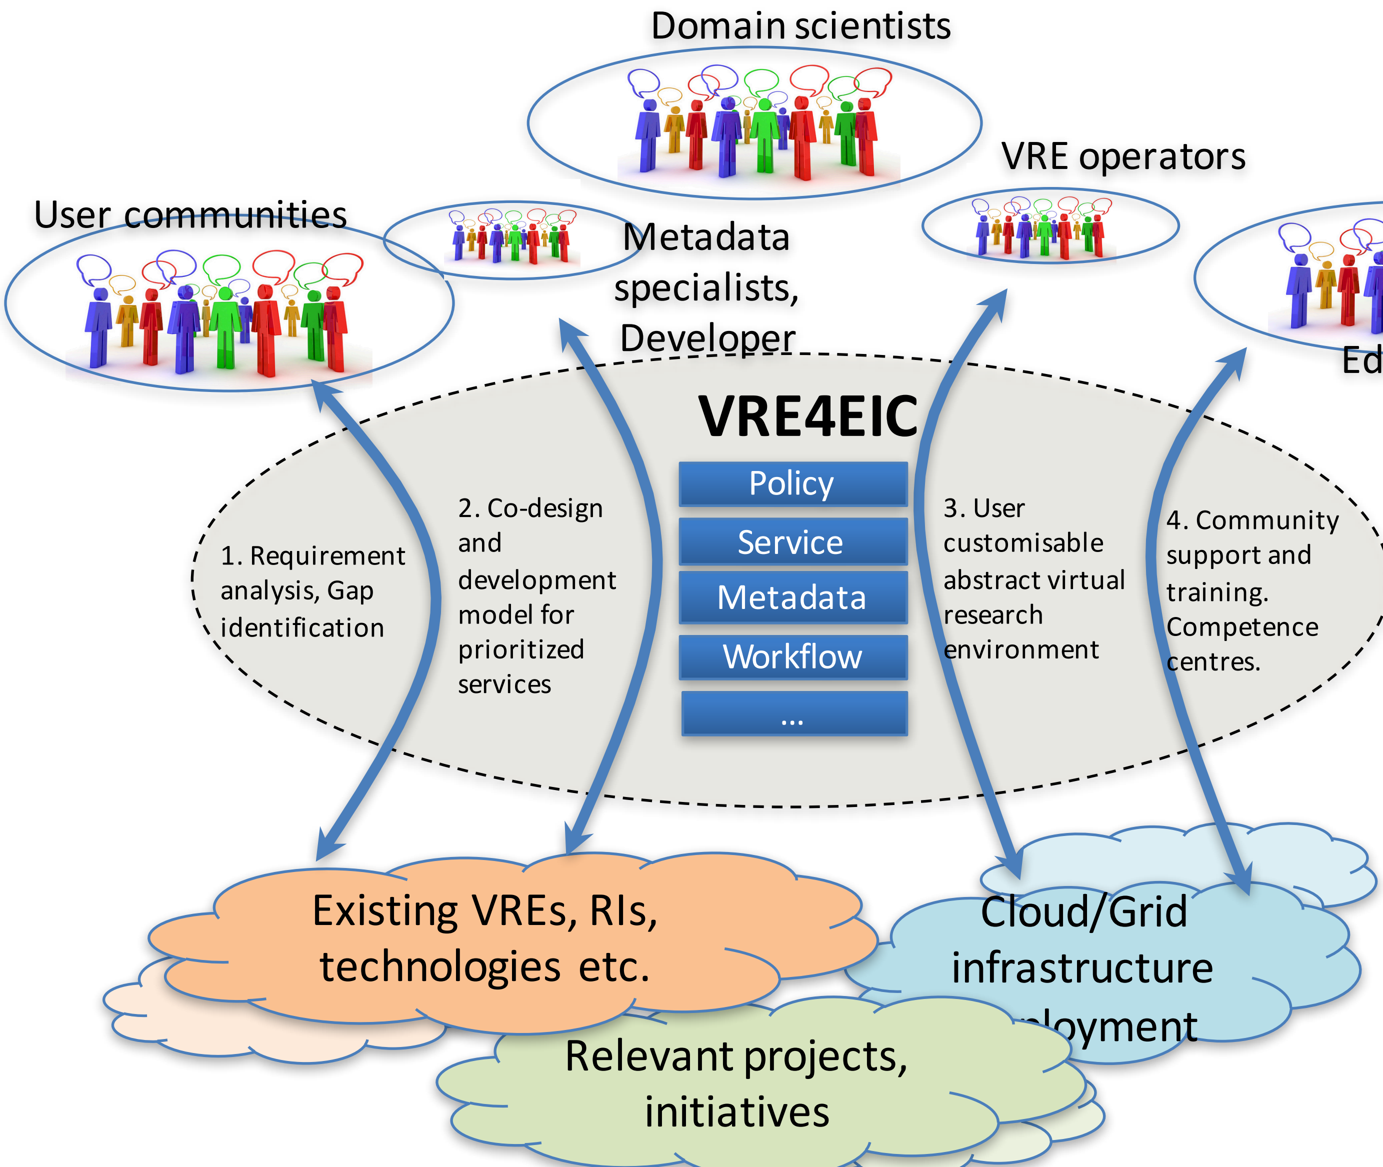

VRE to **E**mpower multi-disciplinary research communities and accelerate **I**nnovation and **C**ollaboration

Summary

- Topics
  - Programming models for time/quality critical cloud applications
    - Application-Infrastructure programming, SLA, Self-adaptability, etc.
  - System level optimization in big data infrastructures
    - Quality critical big data applications, infrastructure interoperability, etc.
  - User centered virtual research environments
    - Infrastructure interoperability

Summary

- References:
  - http://staff.fnwi.uva.nl/z.zhao/ Publications, projects, presentations
  - z.zhao@uva.nl
- Collaborations
  - Exchange of research experiences, ideas etc.
  - Joint research projects/proposals
    - EU-China
    - NWO- NSFC
  - Student projects

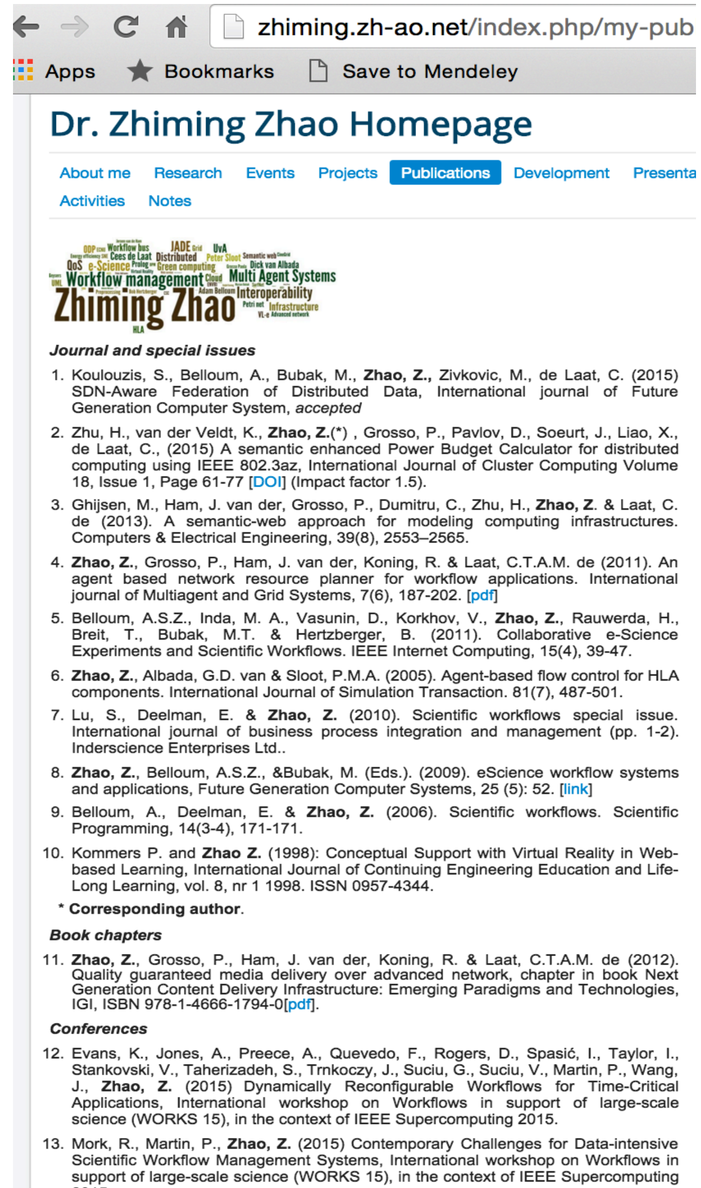

zhiming.zh-ao.net/index.php/my-pub

Apps Bookmarks Save to Mendeley

Dr. Zhiming Zhao Homepage

About me Research Events Projects Publications Development Presenta

Activities Notes

Journal and special issues

1. Koulouzis, S., Belloum, A., Bubak, M., **Zhao, Z.,** Zivkovic, M., de Laat, C. (2015) SDN-Aware Federation of Distributed Data, International journal of Future Generation Computer System, *accepted*
2. Zhu, H., van der Veldt, K., **Zhao, Z.**(*) , Grosso, P., Pavlov, D., Soeurt, J., Liao, X., de Laat, C., (2015) A semantic enhanced Power Budget Calculator for distributed computing using IEEE 802.3az, International Journal of Cluster Computing Volume 18, Issue 1, Page 61-77 [DOI] (Impact factor 1.5).
3. Ghijsen, M., Ham, J. van der, Grosso, P., Dumitru, C., Zhu, H., **Zhao, Z.** & Laat, C. de (2013). A semantic-web approach for modeling computing infrastructures. Computers & Electrical Engineering, 39(8), 2553–2565.
4. **Zhao, Z.**, Grosso, P., Ham, J. van der, Koning, R. & Laat, C.T.A.M. de (2011). An agent based network resource planner for workflow applications. International journal of Multiagent and Grid Systems, 7(6), 187-202. [pdf]
5. Belloum, A.S.Z., Inda, M. A., Vasunin, D., Korkhov, V., **Zhao, Z.,** Rauwerda, H., Breit, T., Bubak, M.T. & Hertzberger, B. (2011). Collaborative e-Science Experiments and Scientific Workflows. IEEE Internet Computing, 15(4), 39-47.
6. **Zhao, Z.,** Albada, G.D. van & Sloot, P.M.A. (2005). Agent-based flow control for HLA components. International Journal of Simulation Transaction. 81(7), 487-501.
7. Lu, S., Deelman, E. & **Zhao, Z.** (2010). Scientific workflows special issue. International journal of business process integration and management (pp. 1-2). Inderscience Enterprises Ltd..
8. **Zhao, Z.,** Belloum, A.S.Z., &Bubak, M. (Eds.). (2009). eScience workflow systems and applications, Future Generation Computer Systems, 25 (5): 52. [link]
9. Belloum, A., Deelman, E. & **Zhao, Z.** (2006). Scientific workflows. Scientific Programming, 14(3-4), 171-171.
10. Kommers P. and **Zhao Z.** (1998): Conceptual Support with Virtual Reality in Web-based Learning, International Journal of Continuing Engineering Education and Life-Long Learning, vol. 8, nr 1 1998. ISSN 0957-4344.

* **Corresponding author**.

Book chapters

11. **Zhao, Z.,** Grosso, P., Ham, J. van der, Koning, R. & Laat, C.T.A.M. de (2012). Quality guaranteed media delivery over advanced network, chapter in book Next Generation Content Delivery Infrastructure: Emerging Paradigms and Technologies, IGI, ISBN 978-1-4666-1794-0[pdf].

Conferences

12. Evans, K., Jones, A., Preece, A., Quevedo, F., Rogers, D., Spasić, I., Taylor, I., Stankovski, V., Taherizadeh, S., Trnkoczy, J., Suciu, G., Suciu, V., Martin, P., Wang, J., **Zhao, Z.** (2015) Dynamically Reconfigurable Workflows for Time-Critical Applications, International workshop on Workflows in support of large-scale science (WORKS 15), in the context of IEEE Supercomputing 2015.
13. Mork, R., Martin, P., **Zhao, Z.** (2015) Contemporary Challenges for Data-intensive Scientific Workflow Management Systems, International workshop on Workflows in support of large-scale science (WORKS 15), in the context of IEEE Supercomputing 2015