Smart and Secure Cyber Infrastructure for Research and Education.

Cees de Laat
University of Amsterdam

on the Occasion of the 25th Anniversary of the Czech Republic's Connection to the Internet

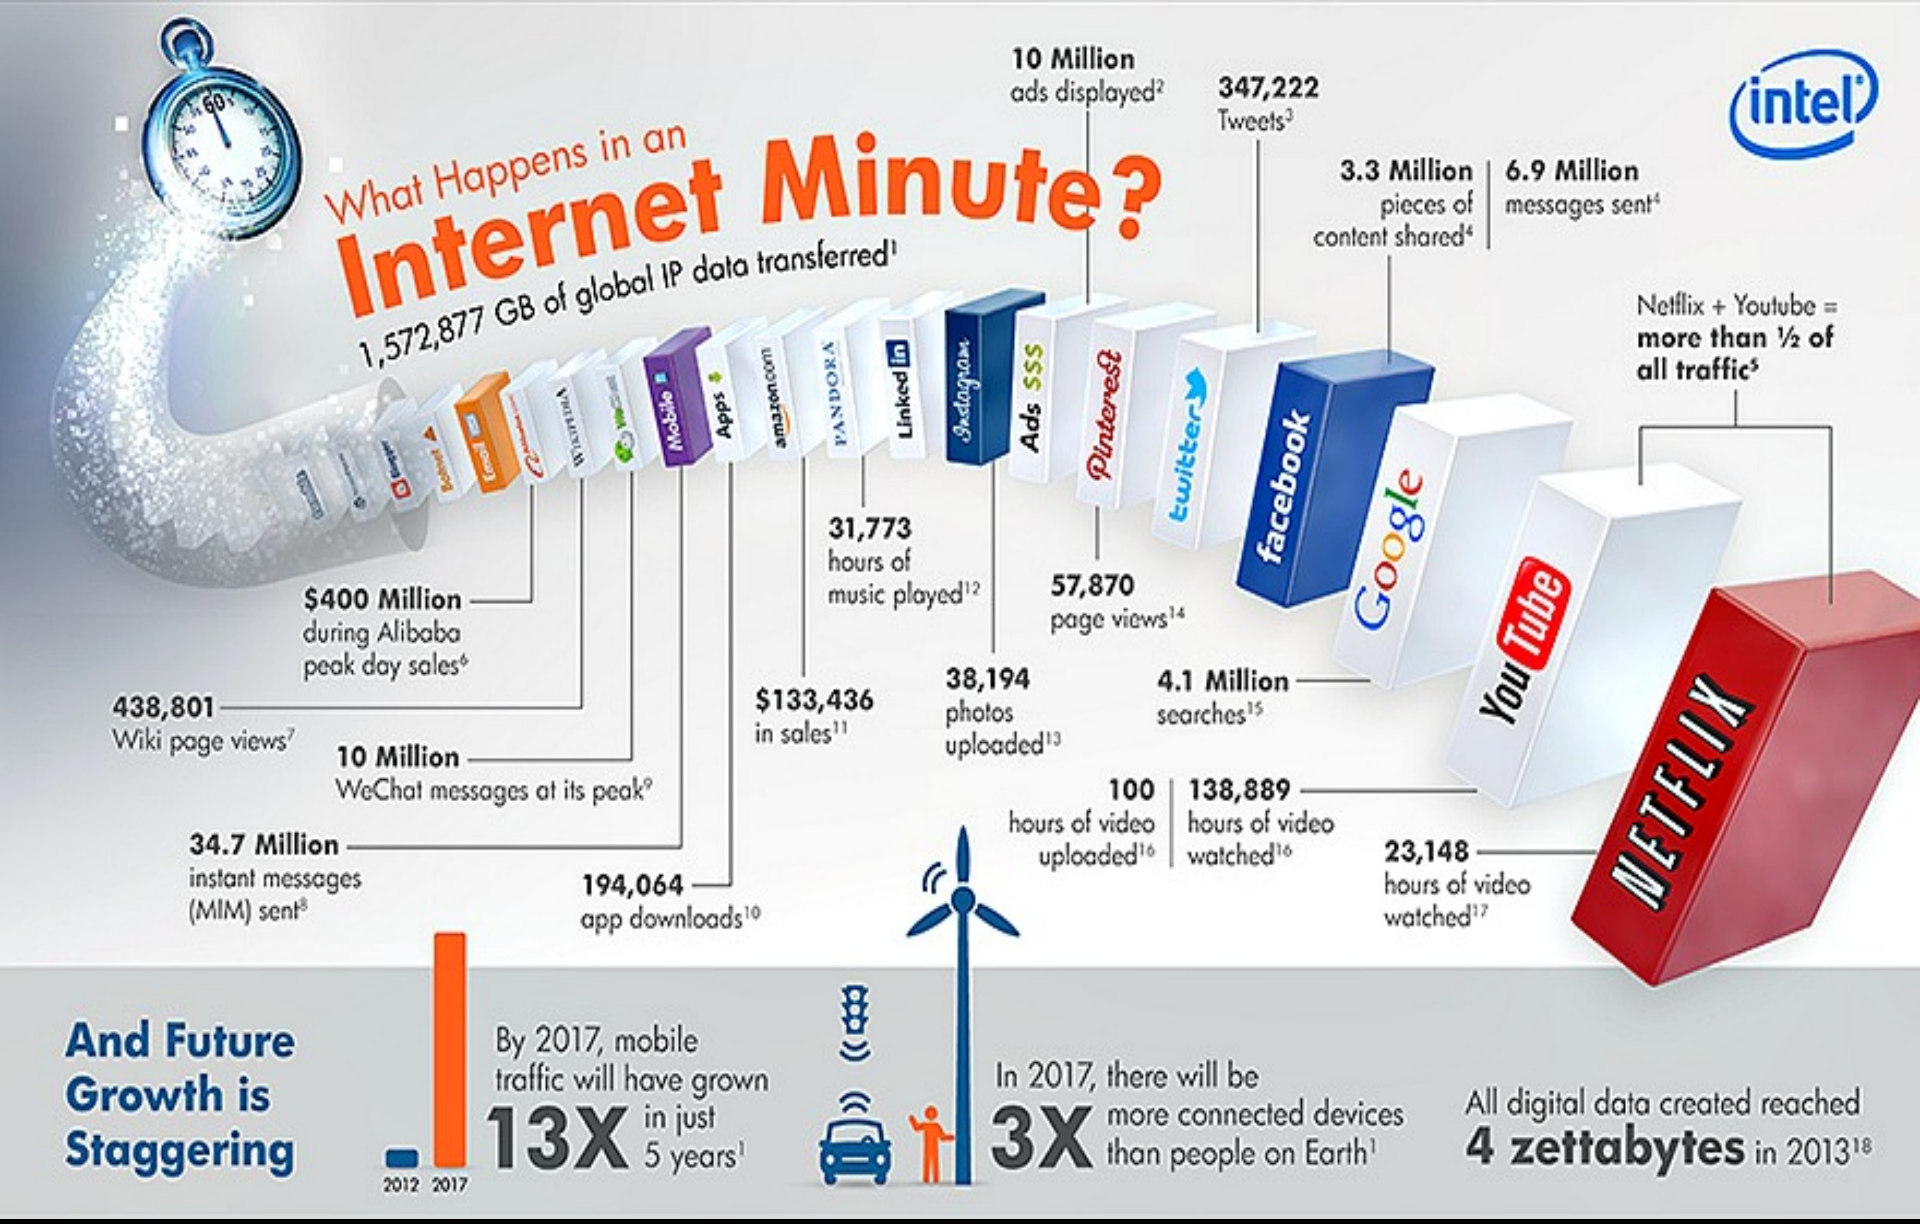

1,572,877 GByte/minute =
(8*1,572,877*10^9/60 bit/s)/(10*10^12 bit/s per fiber) =
21 fibers with each about 100 * 100 Gb/s channels

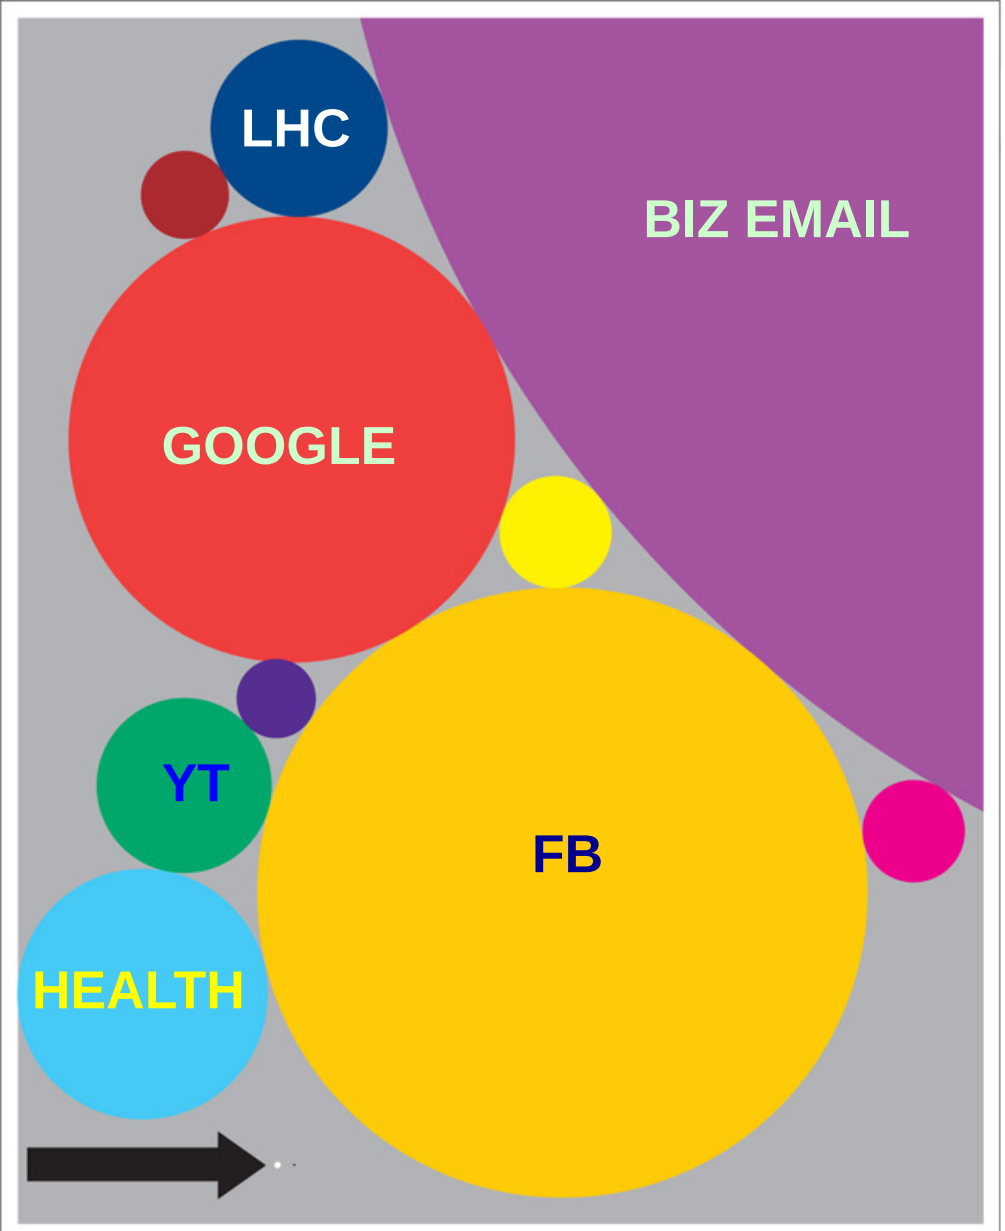

Size of data sets in terabytes

Business email sent per year	2,986,100	National Climactic Data Center database	6,144
Content uploaded to Facebook each year	182,500	Library of Congress' digital collection	5,120
Google's search index	97,656	US Census Bureau data	3,789
Kaiser Permanente's digital health records	30,720	Nasdaq stock market database	3,072
Large Hadron Collider's annual data output	15,360	Tweets sent in 2012	19
Videos uploaded to YouTube per year	15,000	Contents of every print issue of WIRED	1.26

There is always a bigger fish

SC16 Salt Lake City, hpc Utah matters.

3,2 Tbit/s

SCinet

amsix amsterdam internet exchange

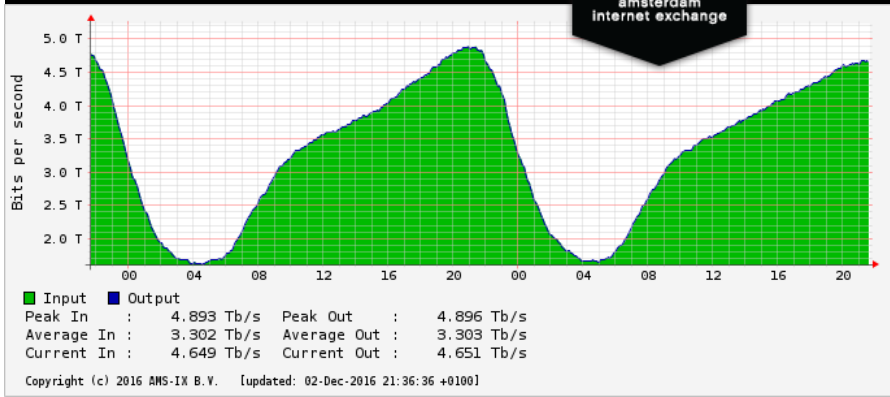

Moore's and Kryders Law

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.

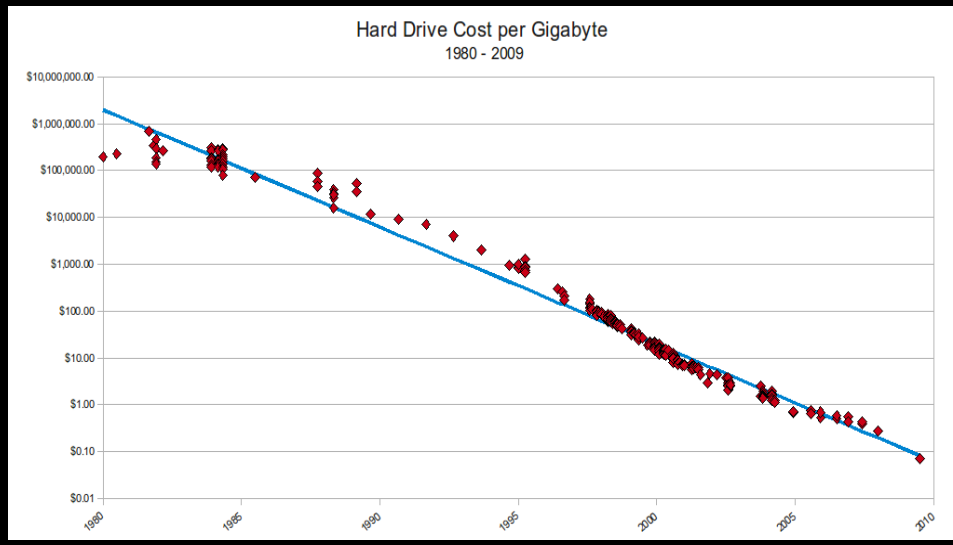

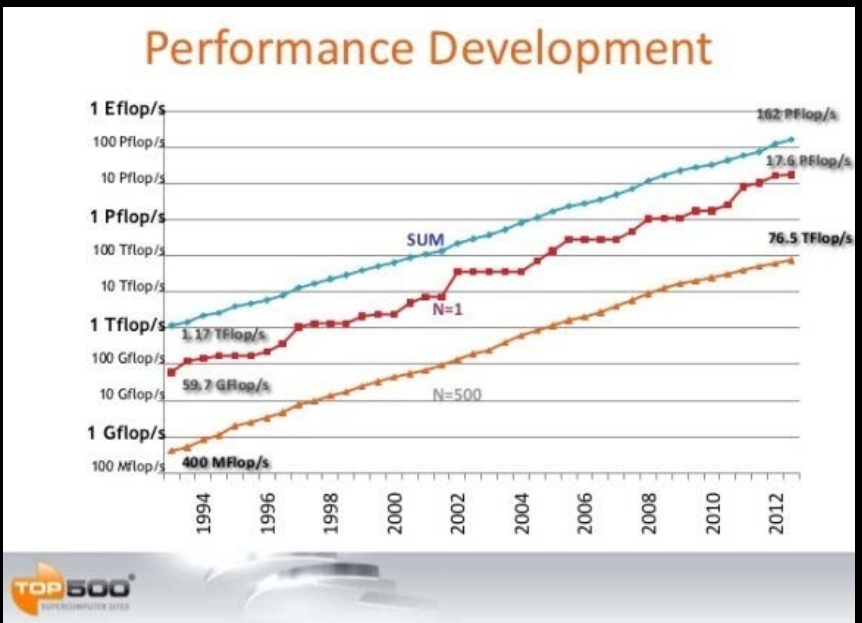

We will soon reach the limits of what is currently feasible and controllable.

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

Trends in Networking

... more data!

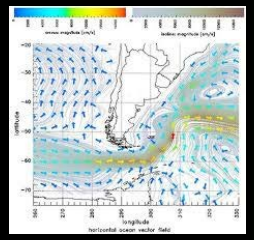

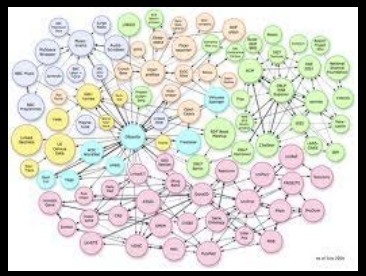

... more realtime!

... more users!

Trends in Networking

... more data!

Speed

Volume

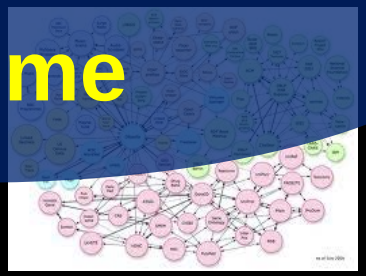

Deterministic

Real-time

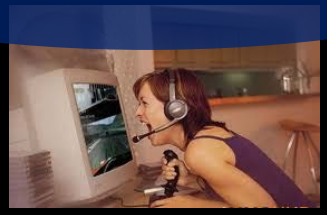

Scalable

Secure

... more users!

Why NREN's?

- Capacity in commercial world much bigger now!
- We don't have scientific water, electricity!
- However, we need an NREN if and only if:
  1. Need feature not delivered on market
     - E.g. dark fiber for SKA, new protocols, SDN/SDX
  2. Business model does not fit our needs
     - E.g. not willing to deliver dark fiber, only services we don't need
  3. Law, privacy, policy or security reasons
     - E.g. data privacy shield, medical data, safe email
- NREN's & Supercenters have opportunity to work on integrated end to end services!

Yesterday's Media Transport Method on the KL601 AMS-LAX-SAN!

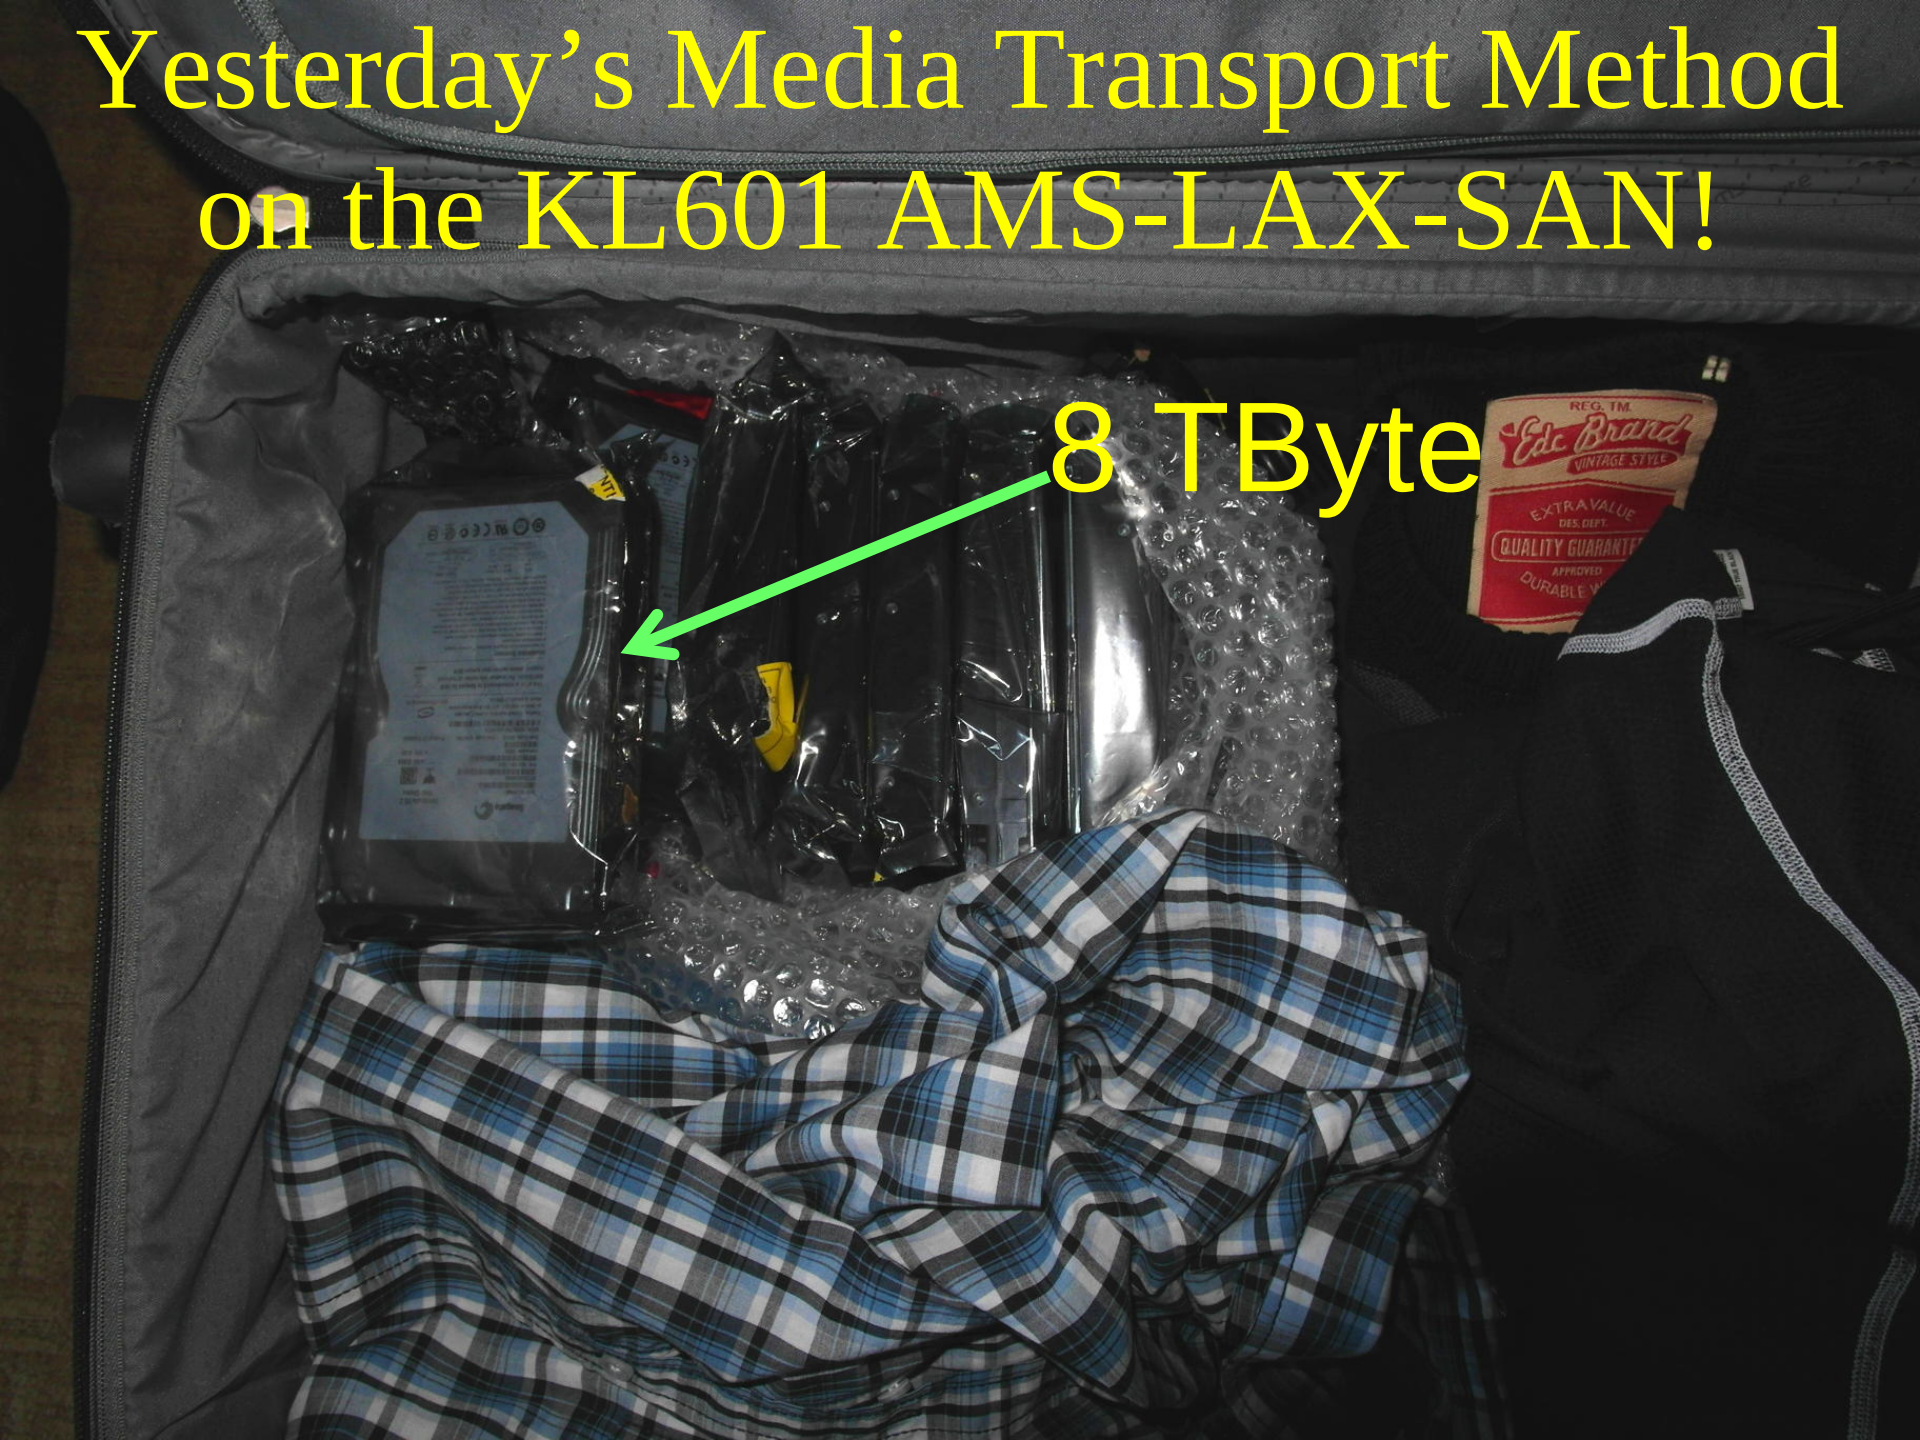

The GLIF – LightPaths around the World

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.

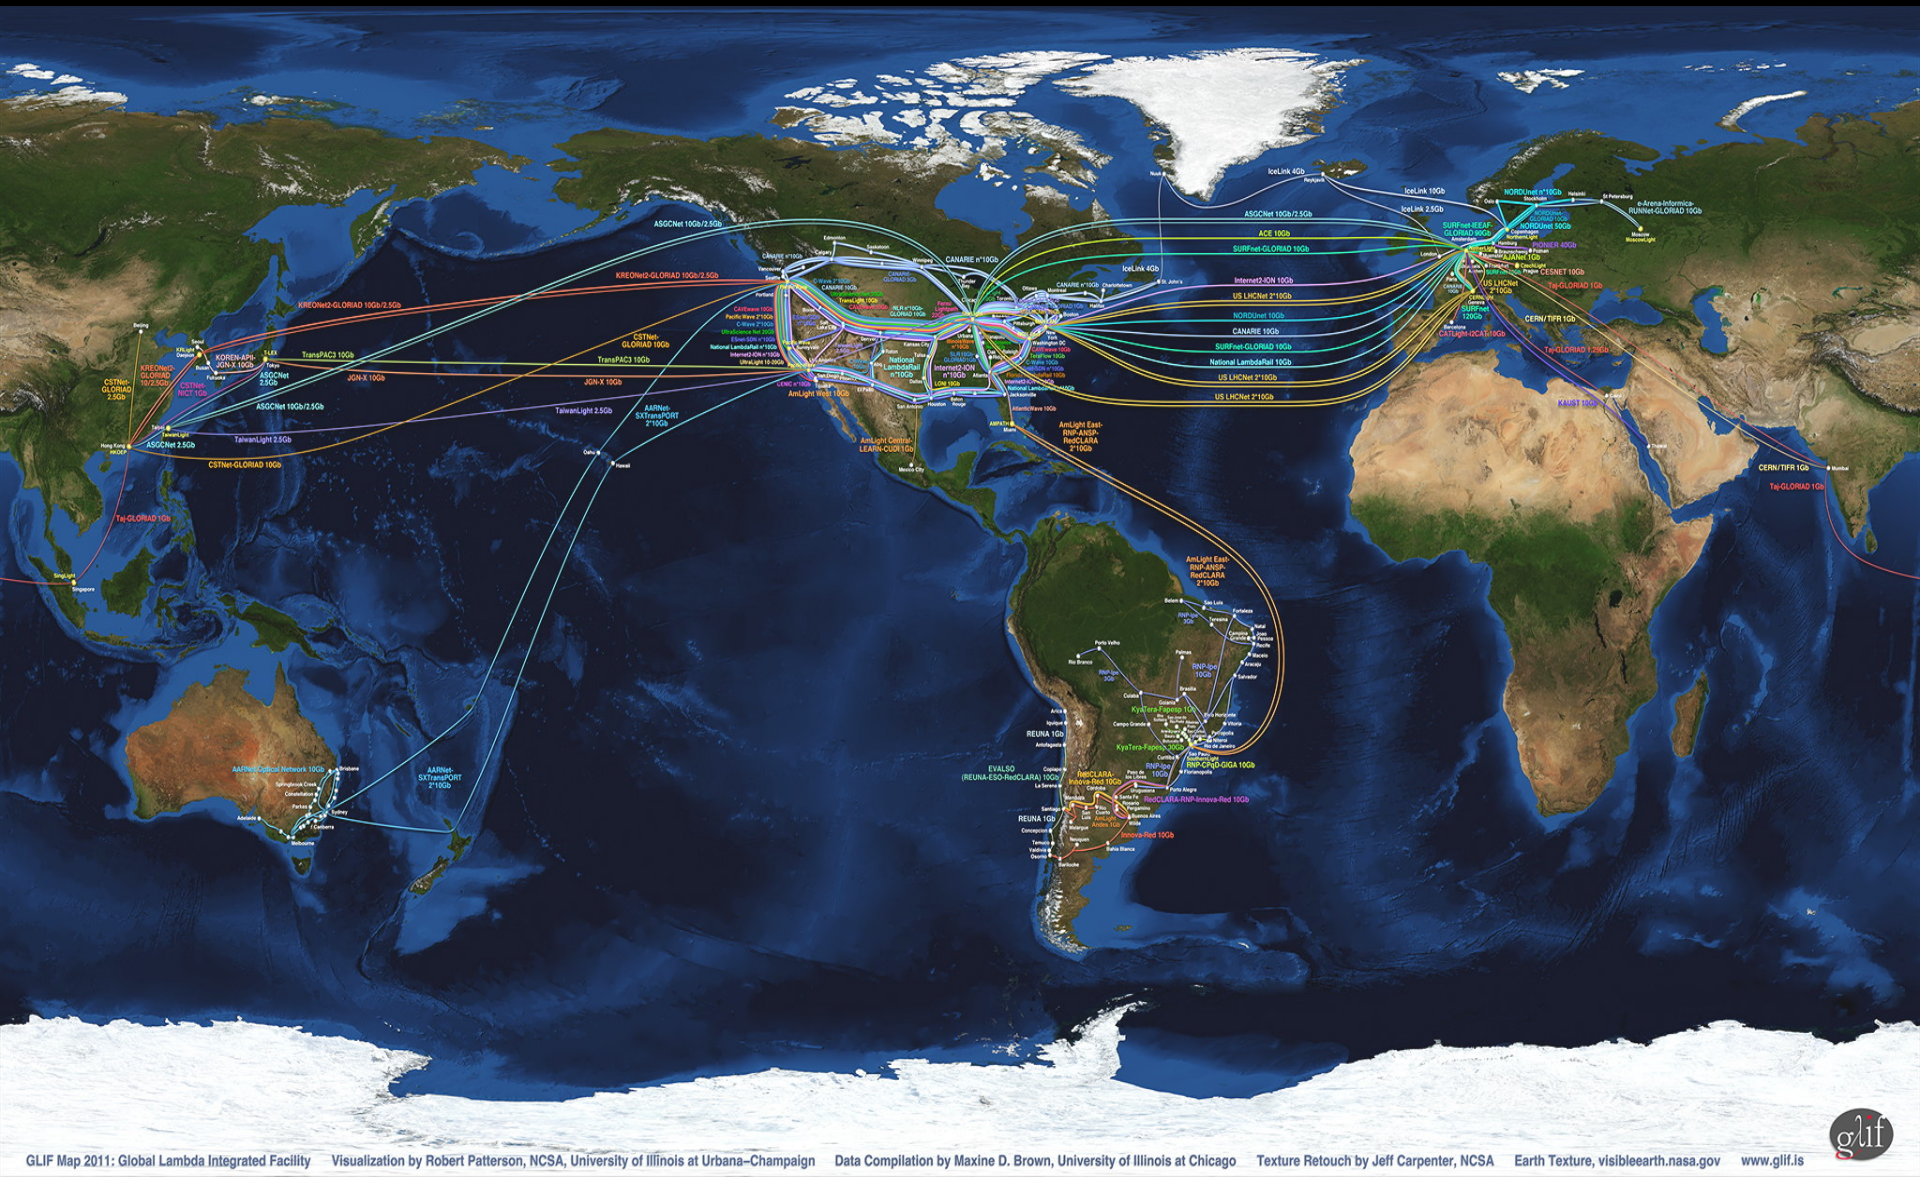

Amsterdam is a major hub in The GLIF

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.

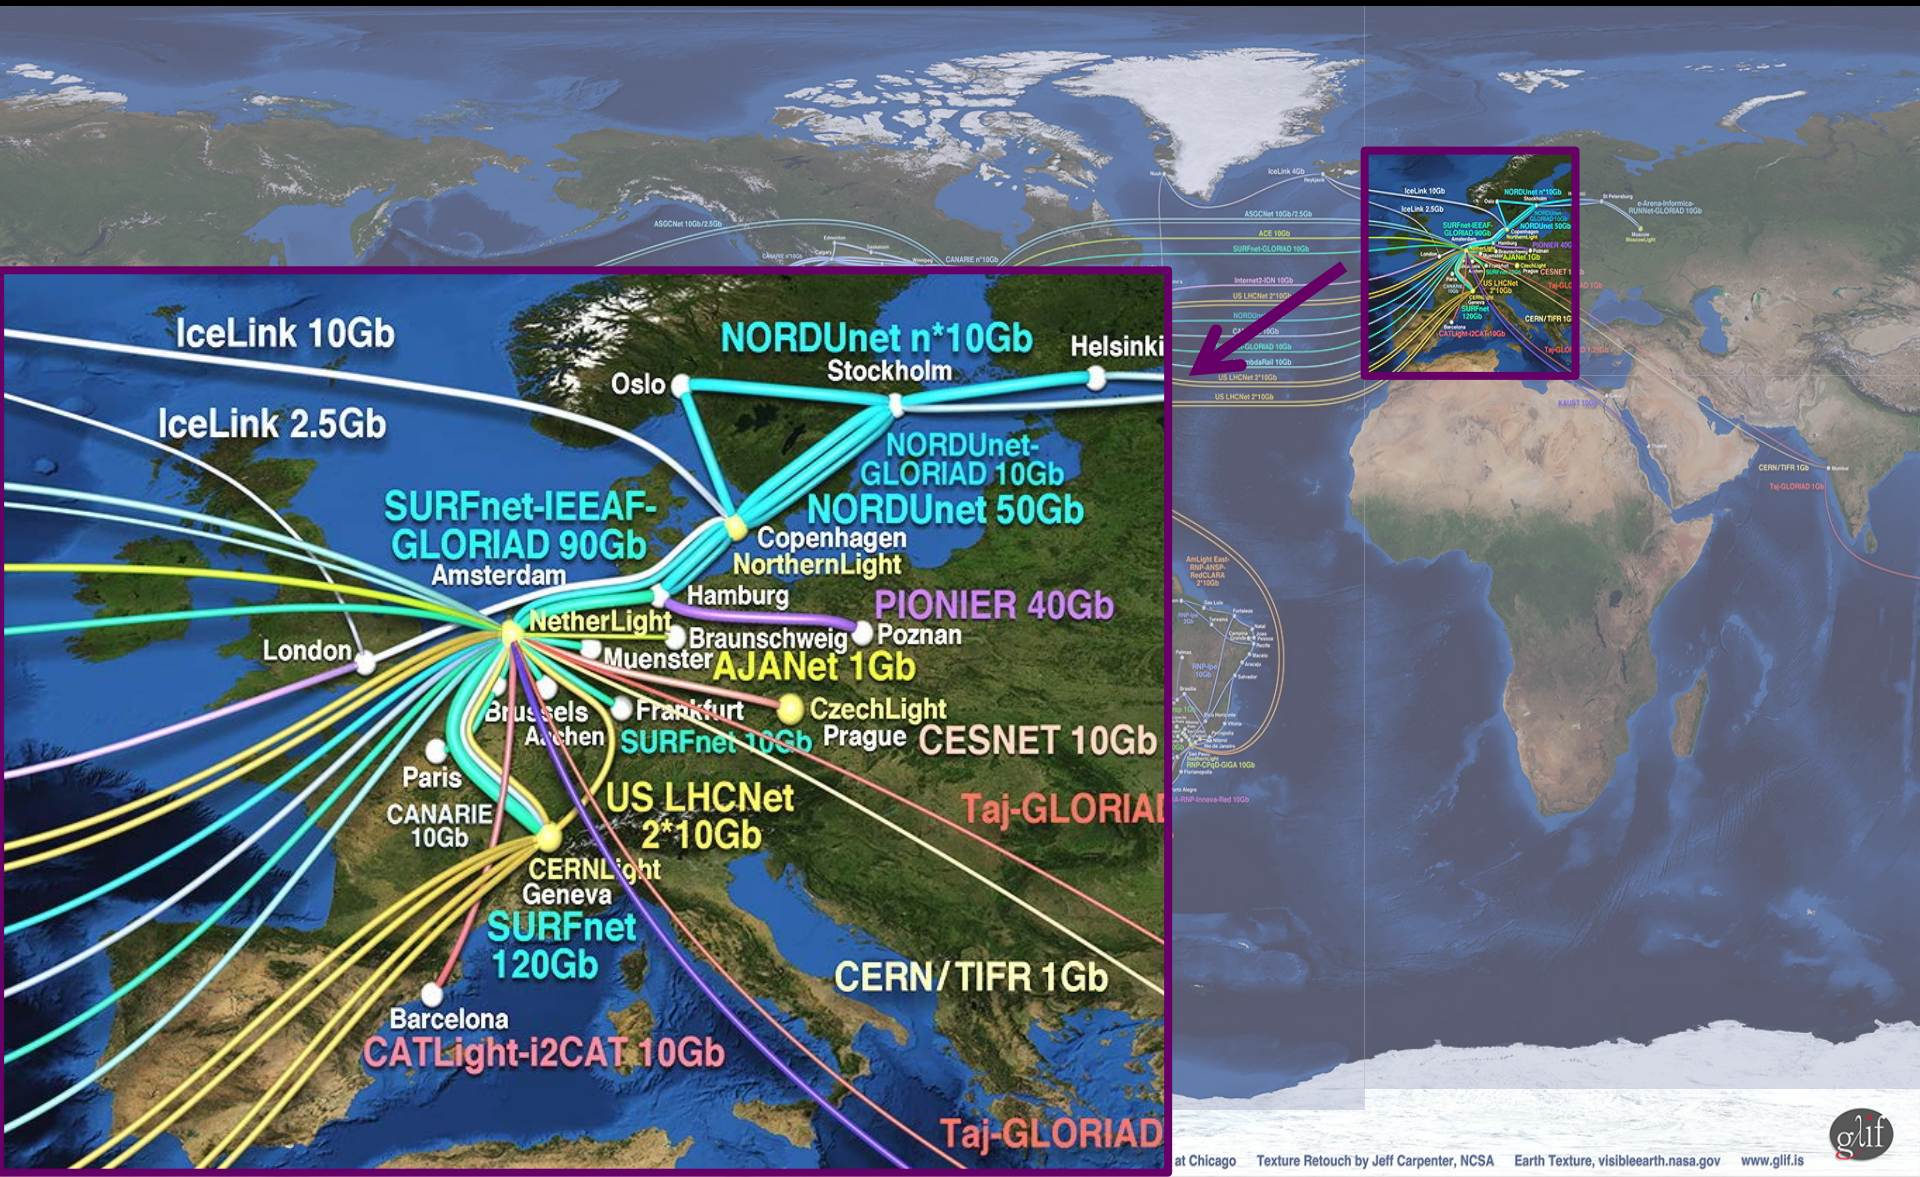

Alien light From idea to realisation!

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

NCF

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.

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

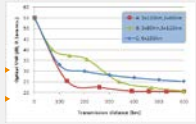

	FOM
A	5504
B	5504
C	1897

telindus

Alien light From idea to realisation!

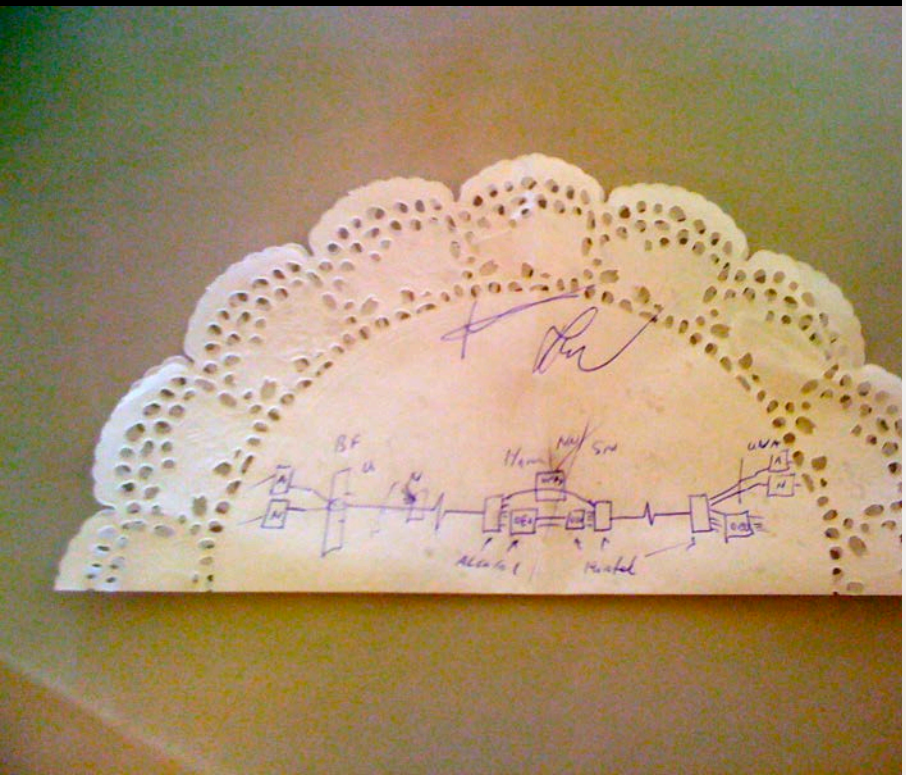

40Gb/s alien wavelength transmission via a multi-vendor 10Gb/s DWDM infrastructure

NCF

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.

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

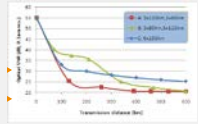

	FOM
A	5504
B	5504
C	1897

telindus

ExoGeni @ OpenLab - UvA

Installed and up June 3th 2013

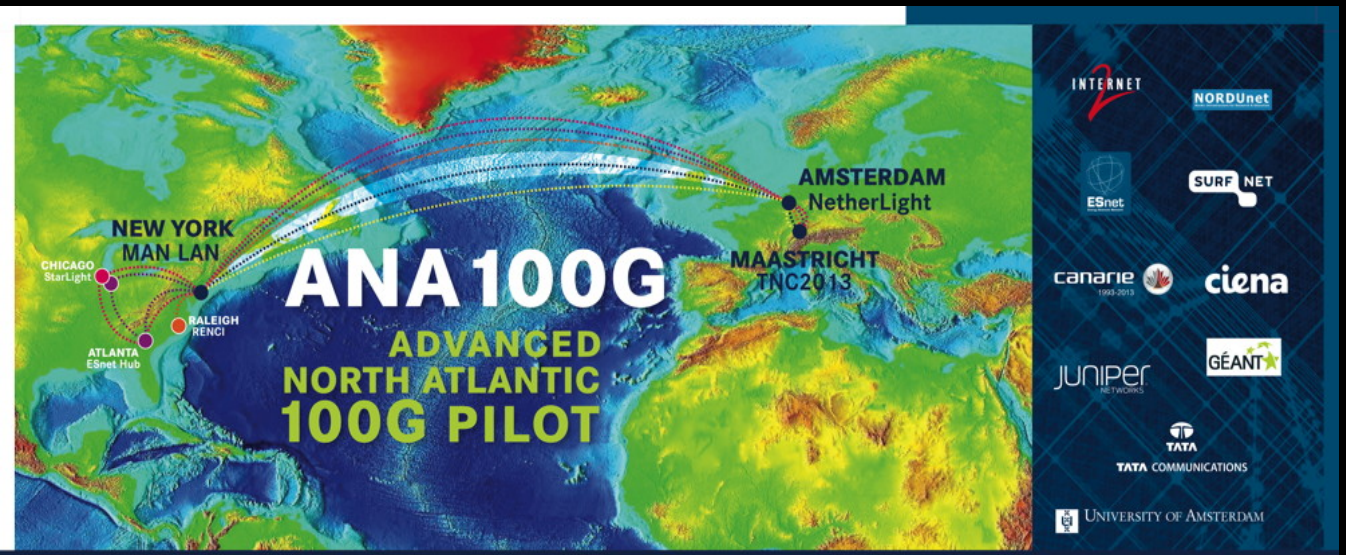

TNC2013 DEMOS JUNE, 2013

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)+ 2x 10GE (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 ALZS 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

Ciena's CENI topology

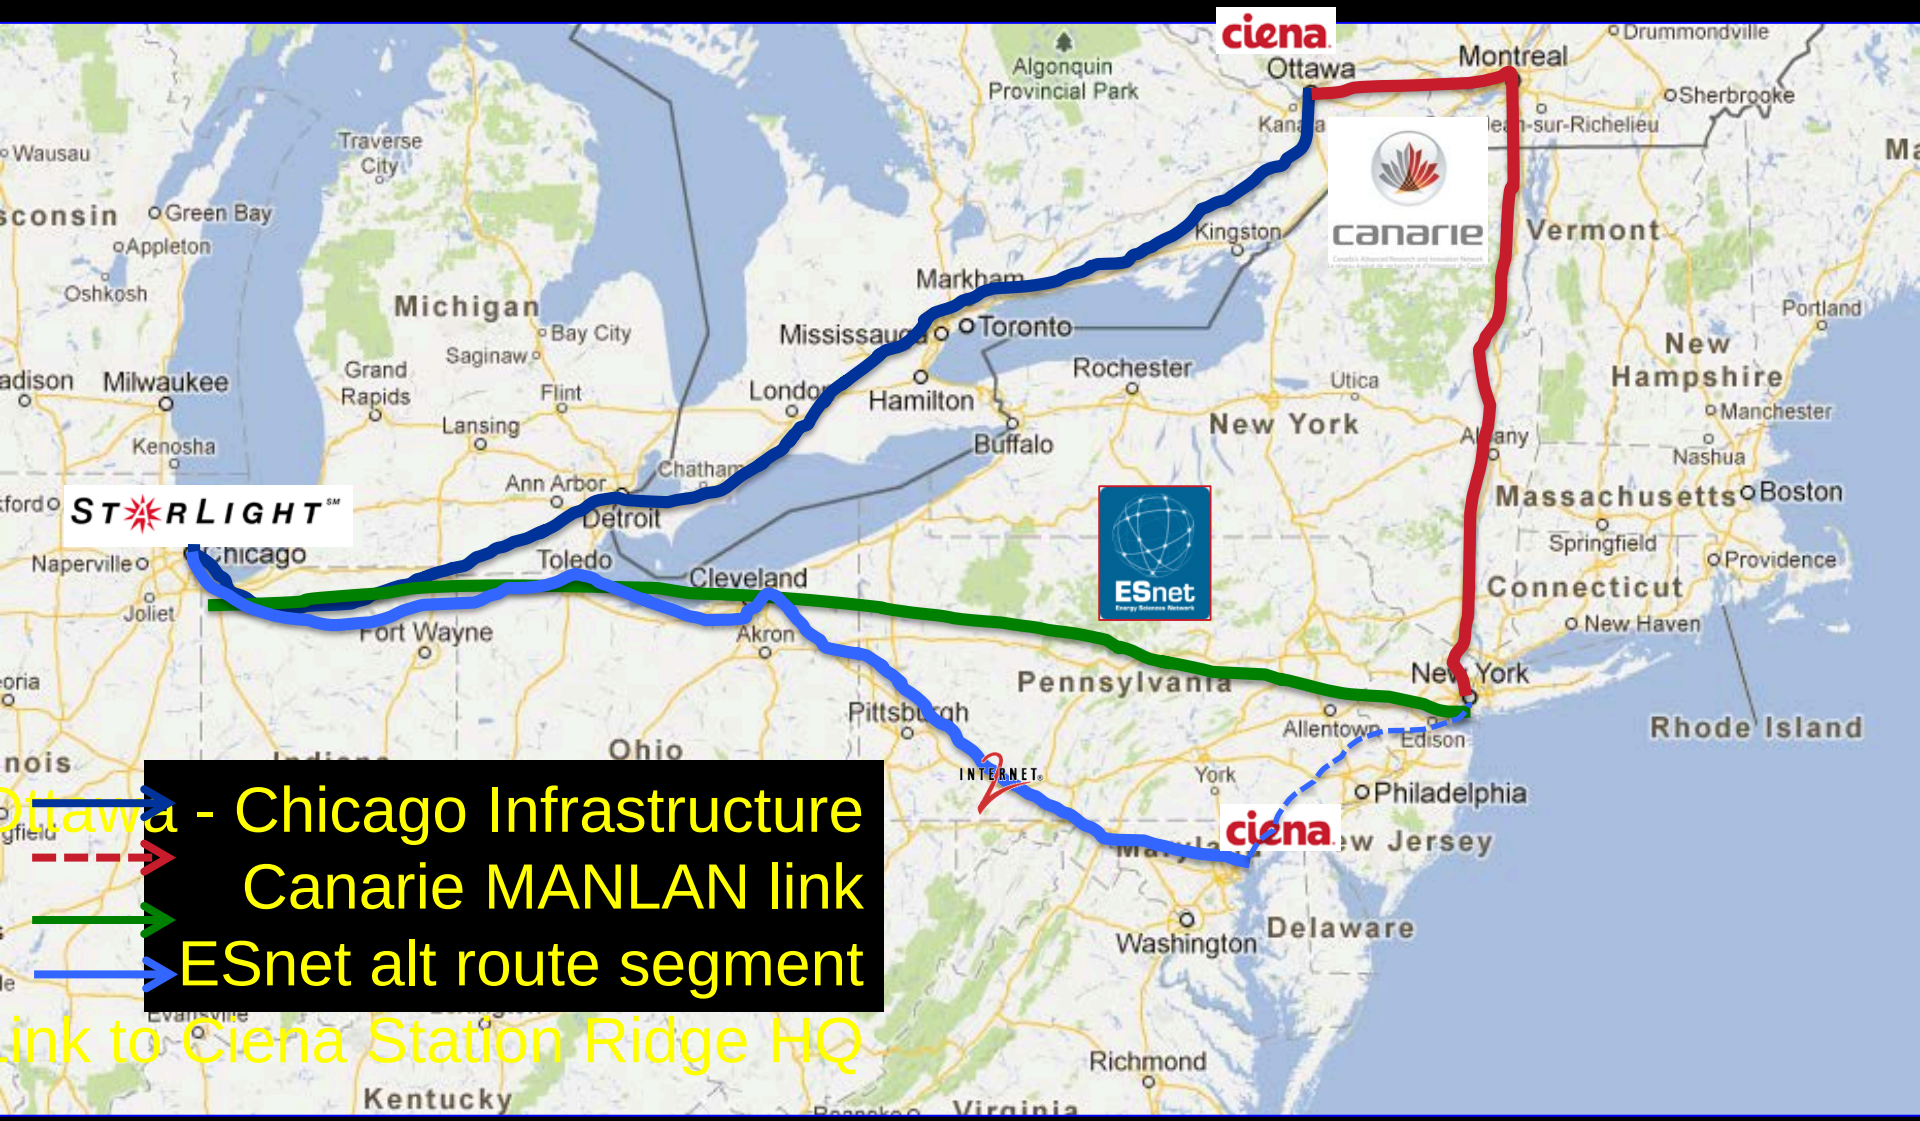

ciena

CENI, International extension to University of Amsterdam

Research Triangle Project. Operation Spring of 2015

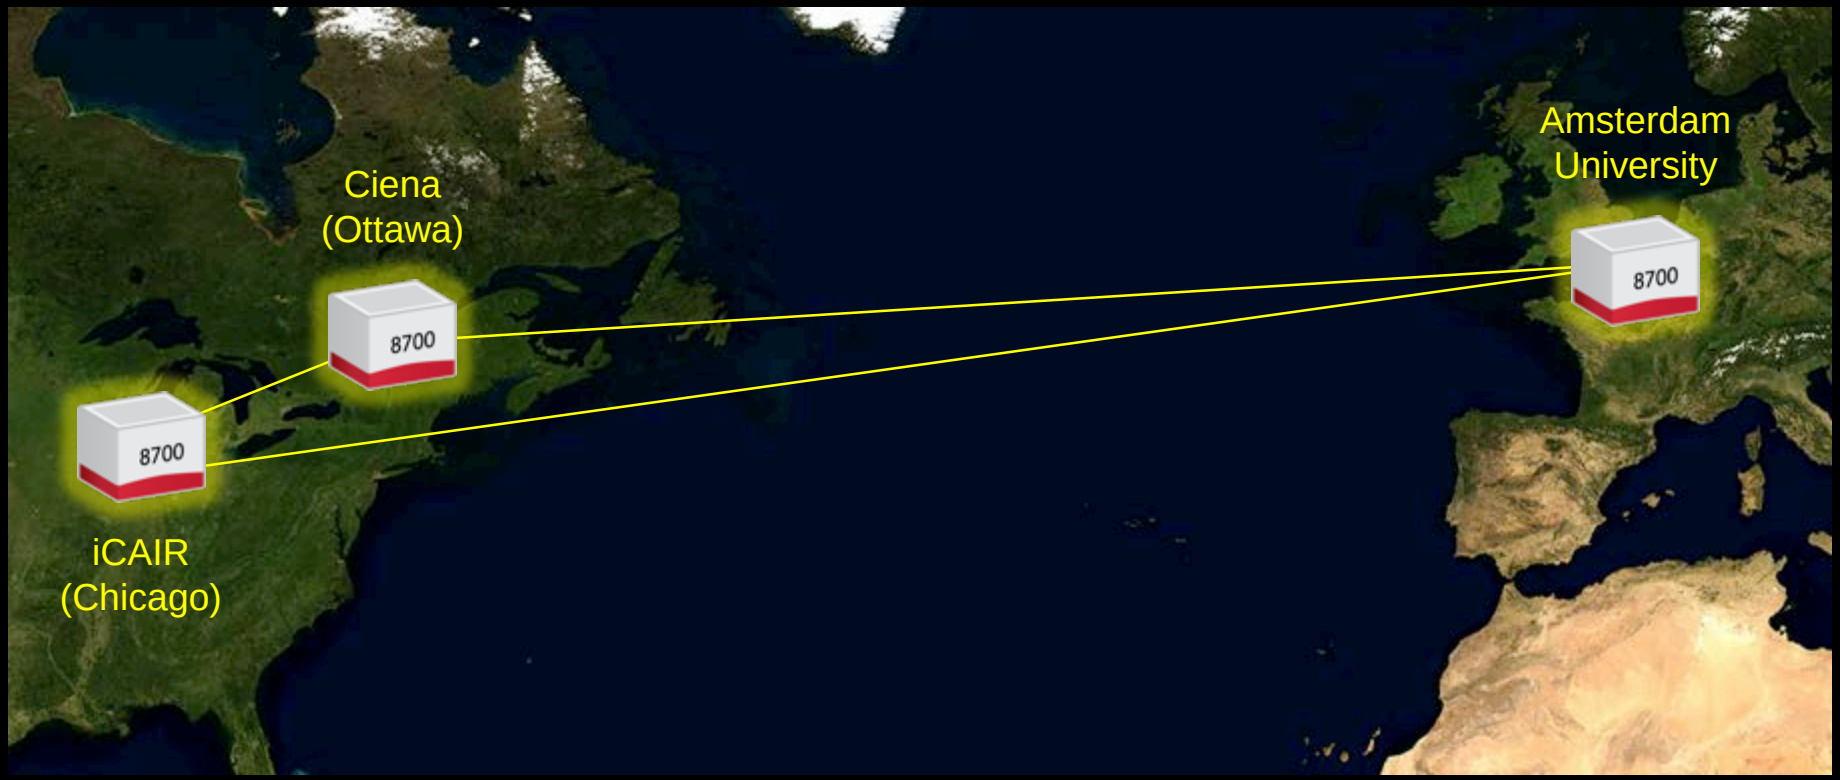

National Science Foundations ExoGENI racks, installed at UvA (Amsterdam), Northwestern University (Chicago) and Ciena's labs (Ottawa), are connected via a high performance 100G research network and trans-Atlantic network facilities using the Ciena 8700 Packetwave platform. This equipment configuration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

ciena

Ambition to put capabilities into fieldlab

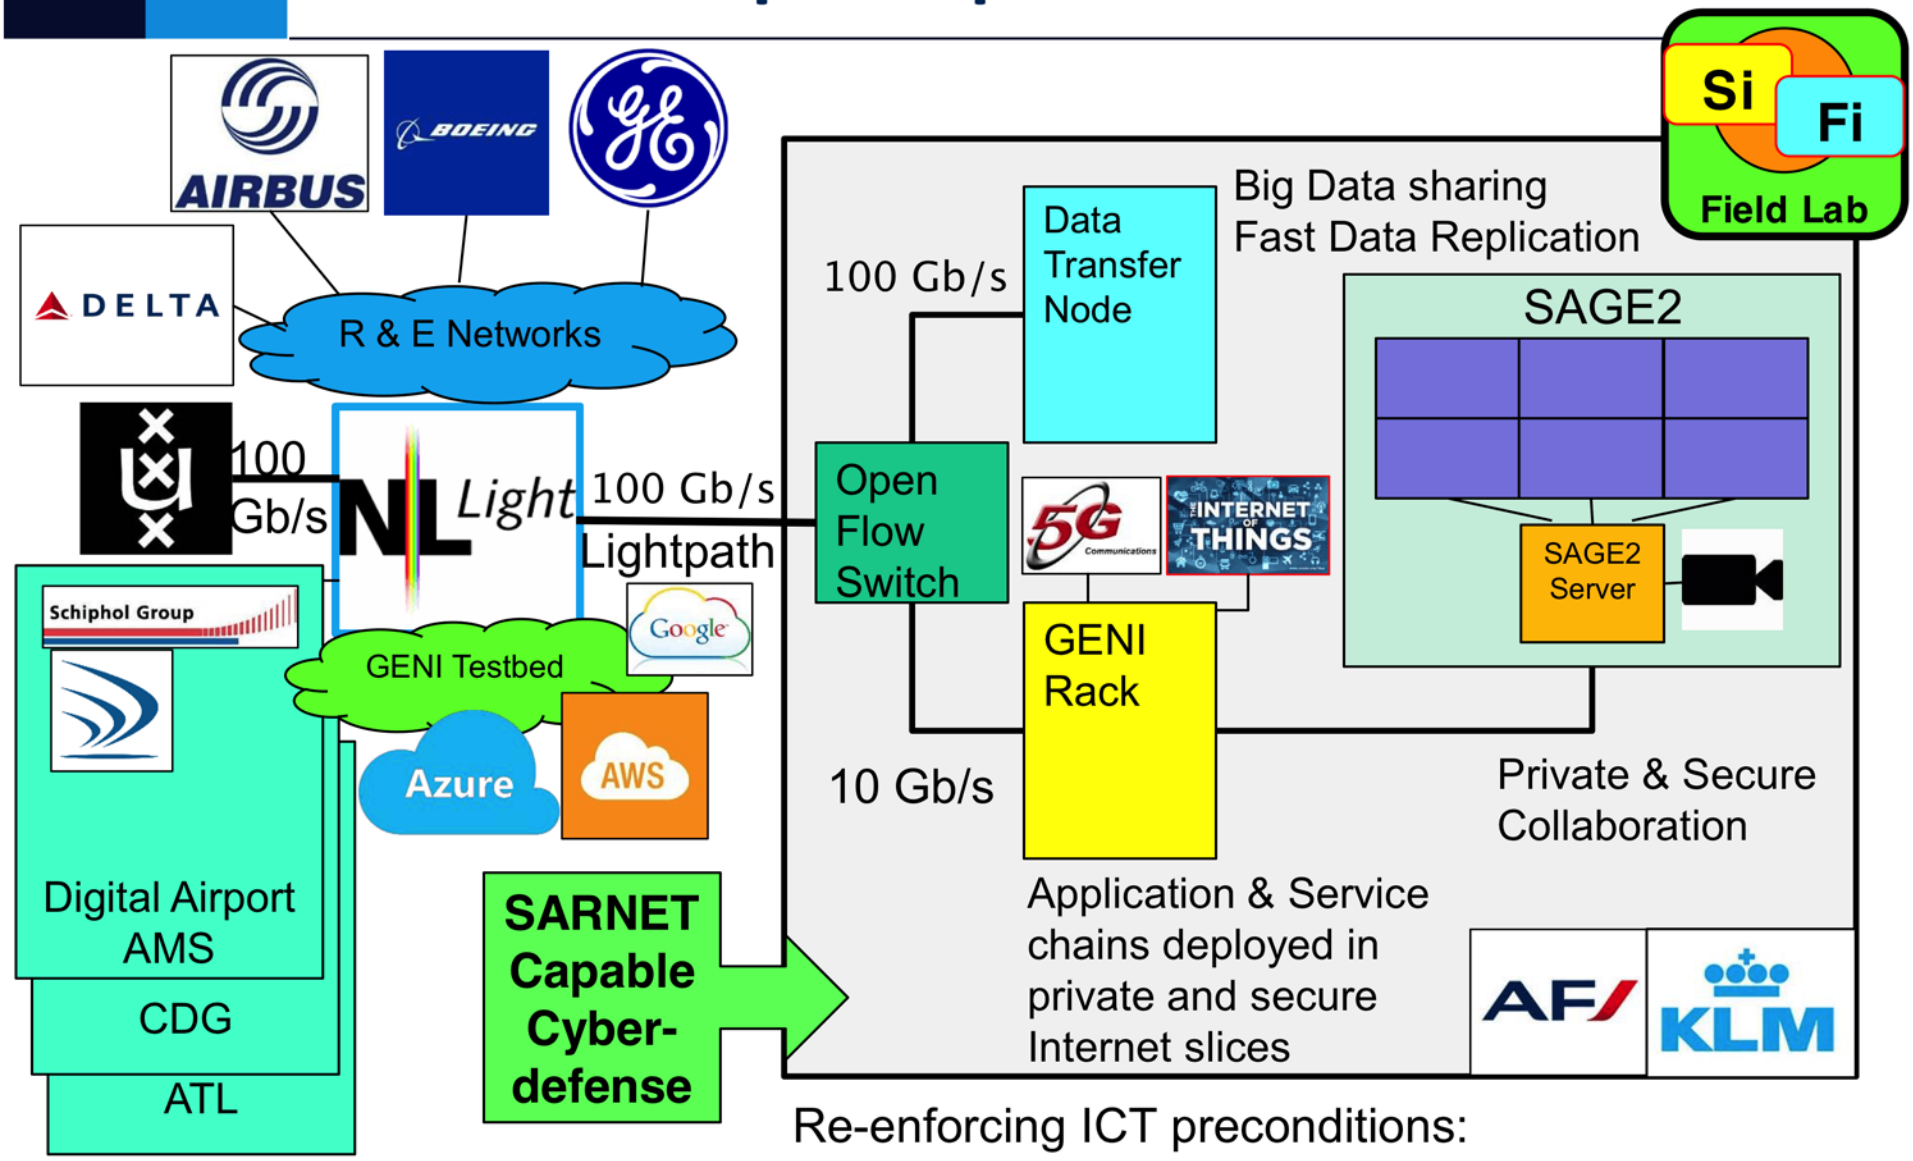

Re-enforcing ICT preconditions:
Each envisaged site has similar elements

ciena SURFNET AIRFRANCE KLM

Participation in Pacific Research Platform

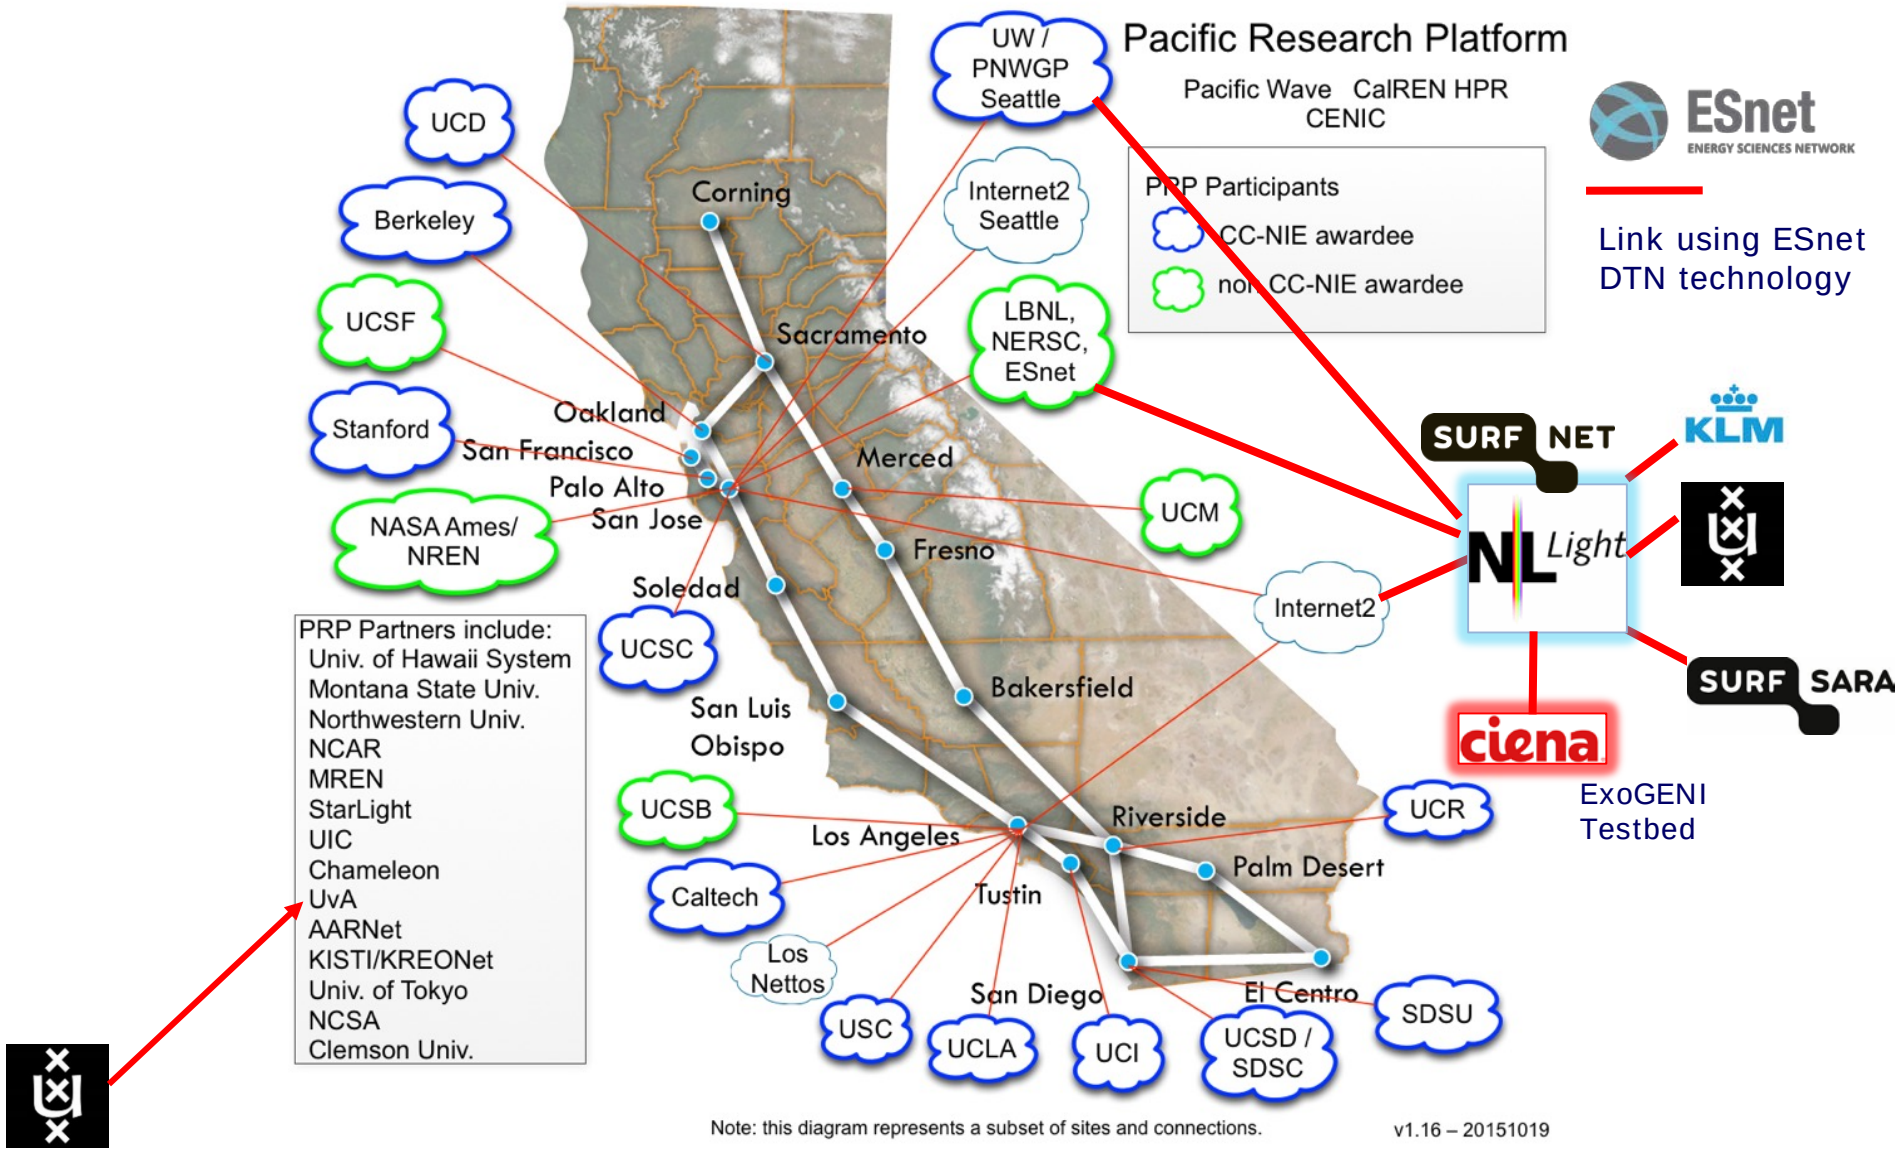

“Learning by Doing”
Early CineGrid Projects

CineGrid @ iGrid 2005

CineGrid @ AES 2006

CineGrid @ Holland Festival 2007

CineGrid @ GLIF 2007

4K interactive digital cinema color grading

realtime 4K uncompressed streaming over IP

CinePOST@Prague ⟷ Calit2@San Diego

Directing Remote Live Shoot of Virtual Set Acting with Live Compositing in the Cloud

Live action camera, actors, green screen at NFTA (Amsterdam #1)
Virtual set compositing at SARA (Amsterdam #2)
Remote viewing and direction at UCSD/Calit2 Vroom (San Diego)

Trucks of Tapes

WSJ Nov 30, 2016

THE WALL STREET JOURNAL.

Amazon Uses Trucks to Drive Data Faster

PERSONAL TECHNOLOGY

The Cable-Cutting Dream Is Kind ...

Altice Plans Fiber Upgrade That Could Leave Rivals in the Dust

Netflix Now Lets You Download, But Many Top Shows Are Off Limits

TECH

Amazon Uses Trucks to Drive Data Faster

Cloud-computing unit, Amazon Web Services, unveils new offerings at annual conference in Las Vegas

Amazon unveiled the 'Snowmobile' service on Wednesday in Las Vegas. *PHOTO: AMAZON WEB SERVICE*

The tractor-trailer hauls a massive storage device, dubbed Snowmobile, in the form of a 45-foot shipping container that holds 100 petabytes of data. A petabyte is about 1 million gigabytes.

1 fiber does about 16 Tbit/s
= 2 Tbyte/s
⇒ 500000 s/ExaByte
⇒ One week/ExaByte

The company, however, isn't promising lightning speed. Ten Snowmobiles would reduce the time it takes to move an exabyte from on-premises storage to Amazon's cloud to a little less than six months, from about 26 years using a high-speed internet connection, by the company's calculations.

By JAY GREENE By LAURA STEVENS
Updated Nov. 30, 2016 7:19 p.m. ET

LAS VEGAS—In Amazon Web Services, Amazon.com Inc. has built one of the most powerful computing networks in the world, on pace to post more than $12 billion in revenue this year.

But the retail giant on Wednesday proposed a surprising way to move data from large corporate customers' data centers to its public cloud-computing operation: by truck.

Networks can move massive amounts of data only so fast. Trucks, it turns out, can move it faster.

1 fiber does
100 Petabytes
in one day,
if you can fill it!

Science-DMZ

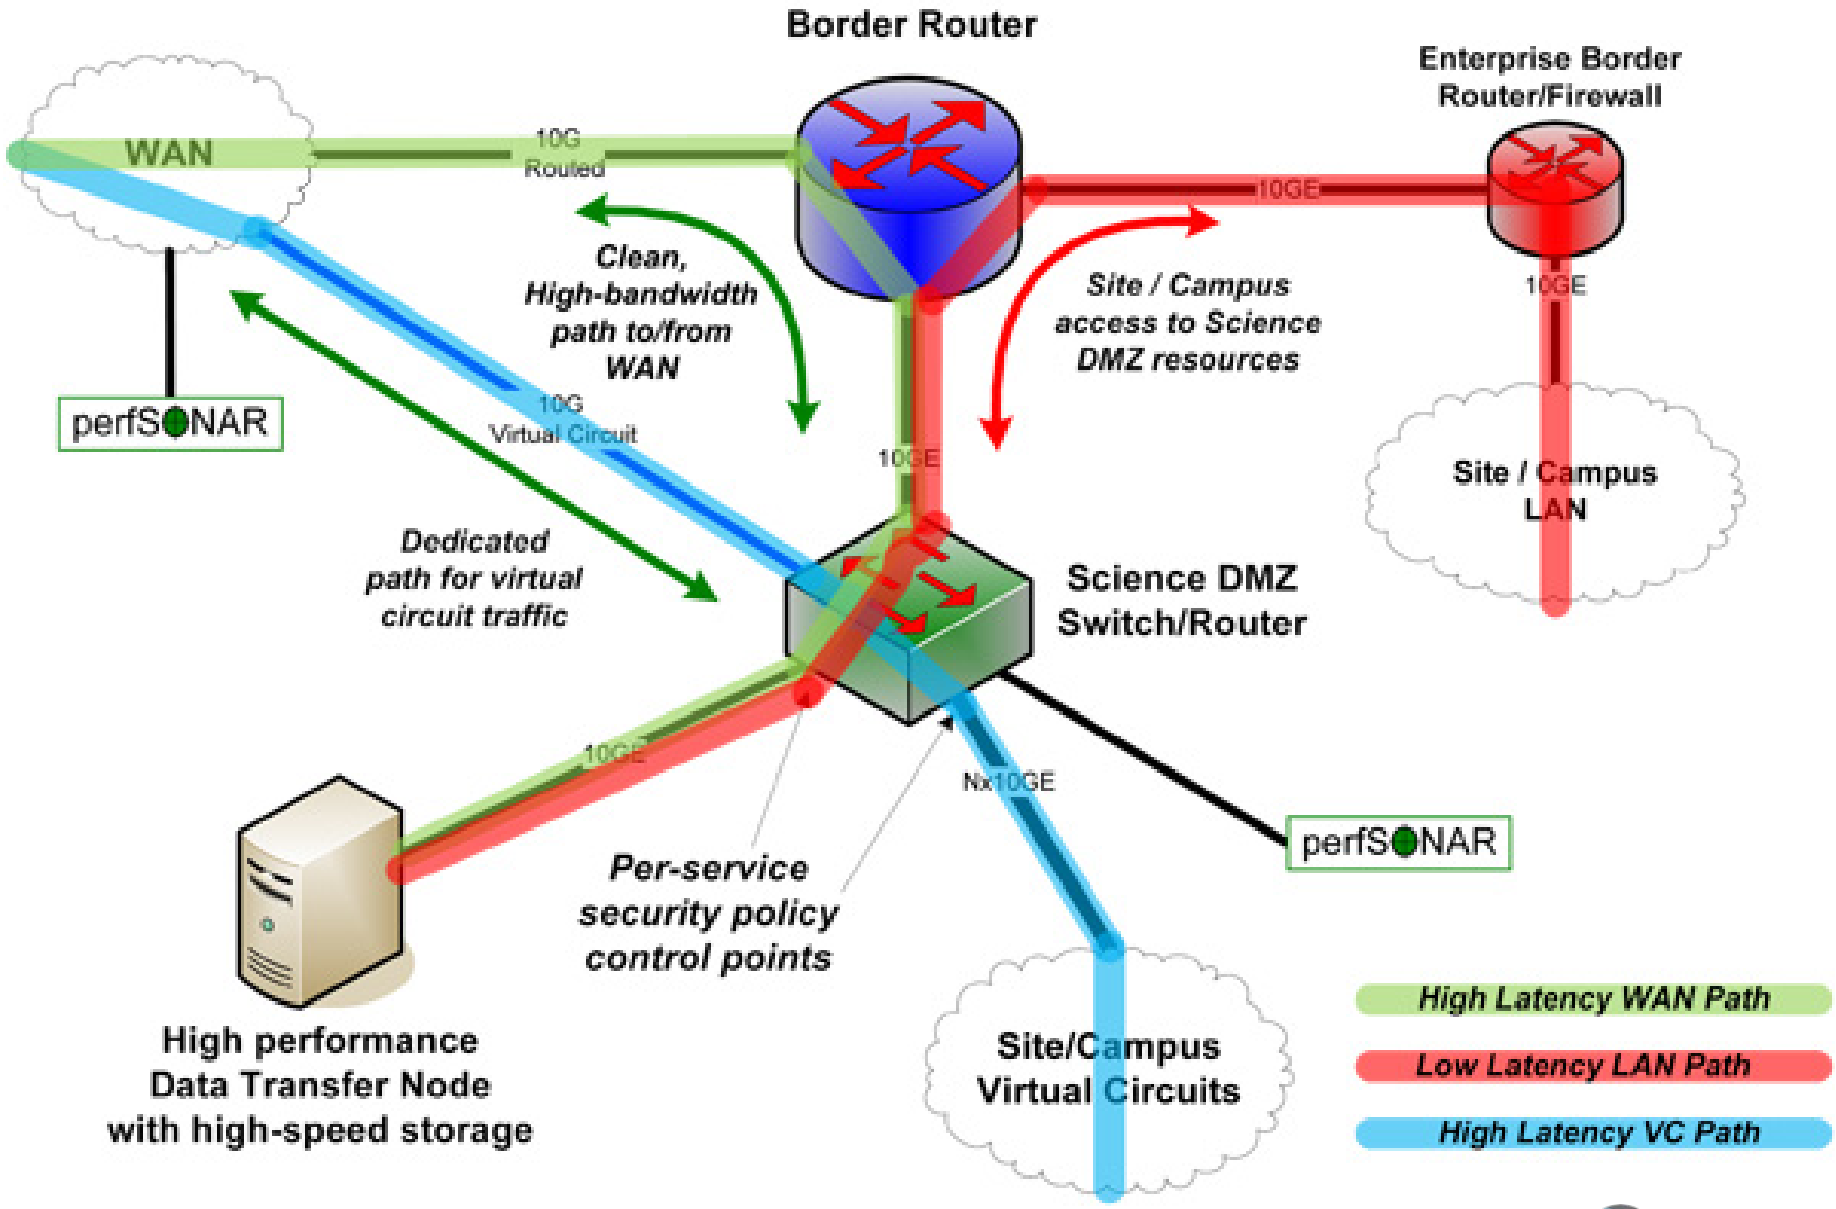

CENIC

ESnet

Networks of ScienceDMZ's & SDX's

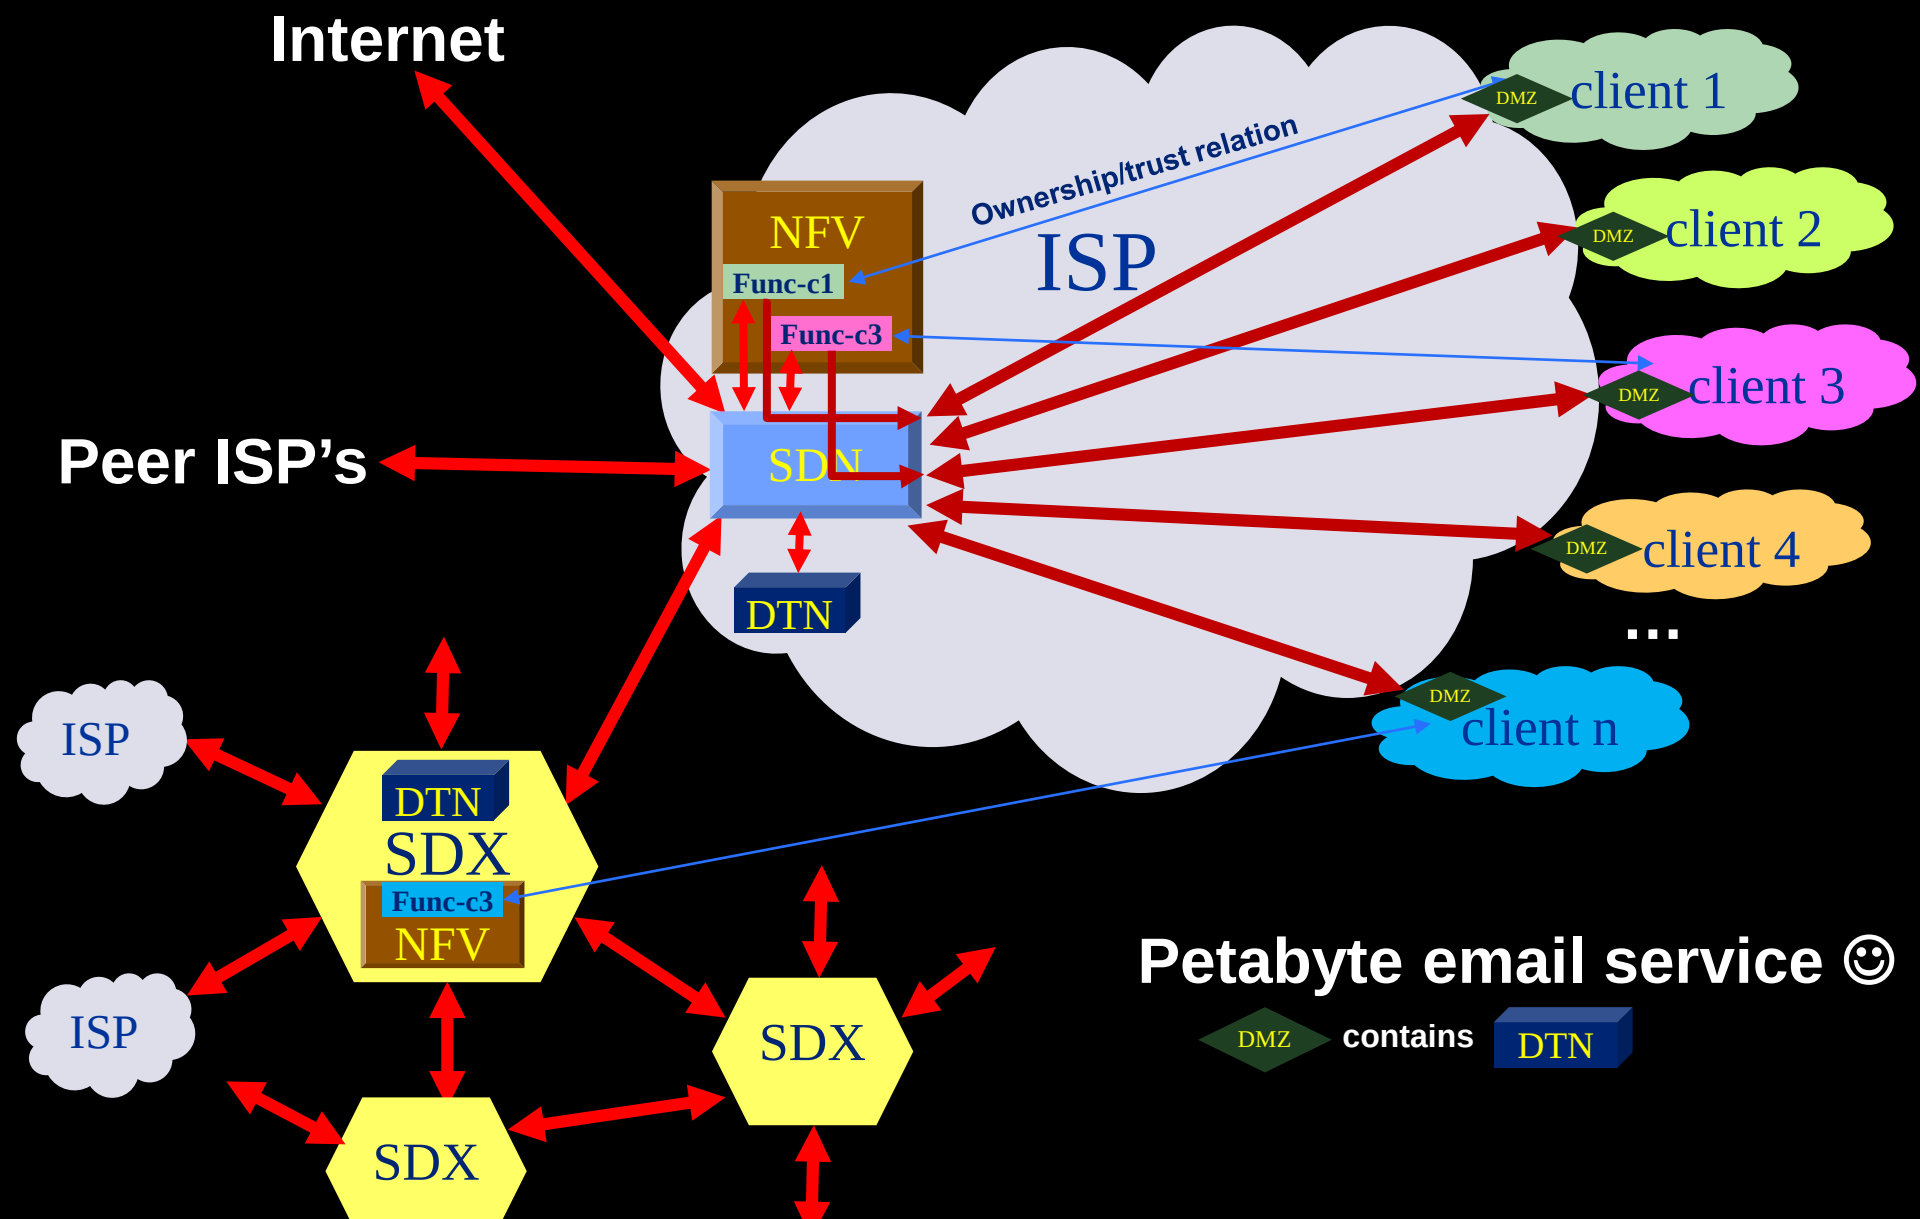

Basic operating system loop

Secure Policy Enforced Data Processing

- Bringing data and processing software from competing organisations together for common goal
- Docker with encryption, policy engine, certs/keys, blockchain and secure networking
- Data Docker (virtual encryped hard drive)
- Compute Docker (protected application, signed algorithms)
- Visualization Docker (to visualize output)

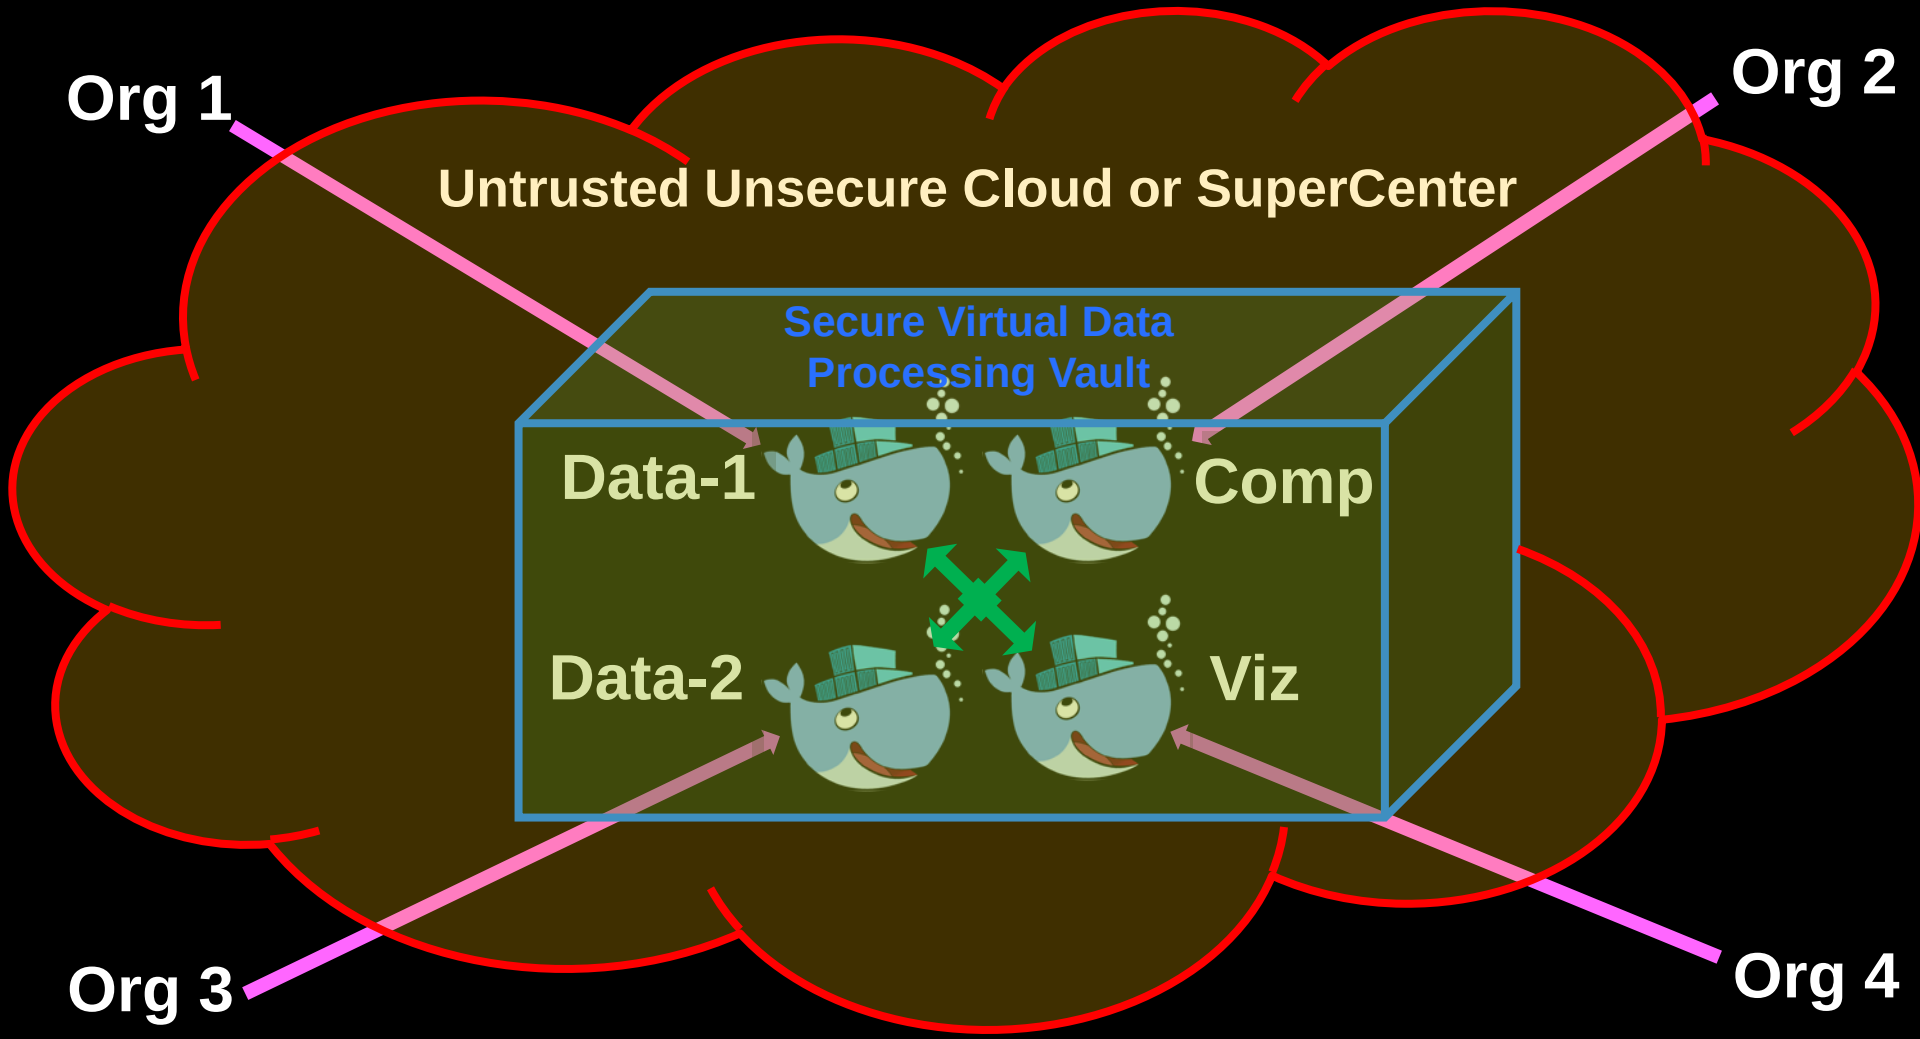

I want to

"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!

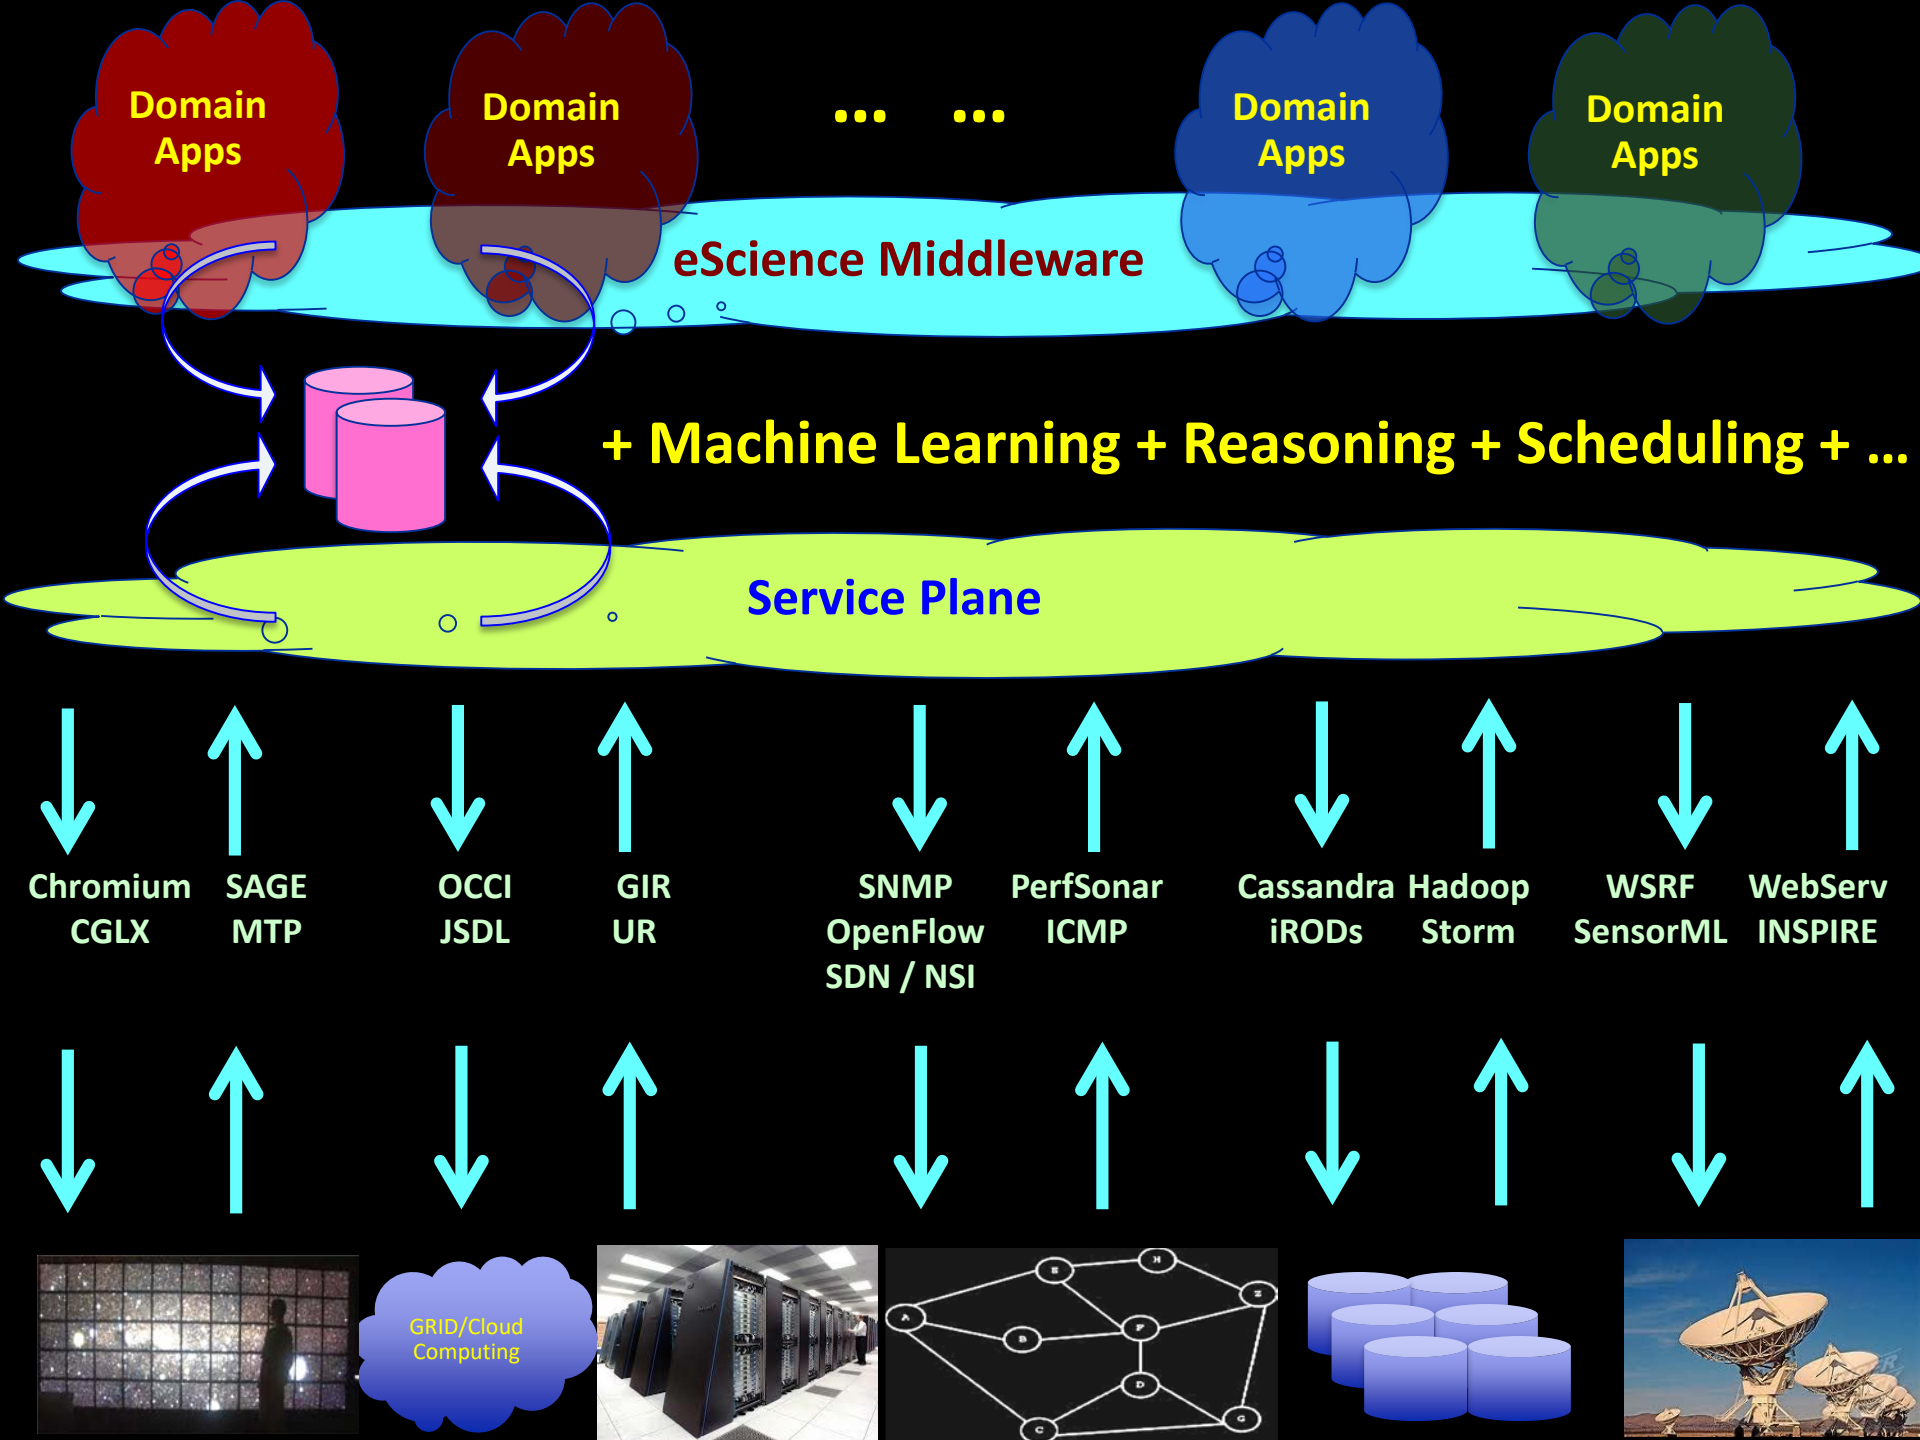
Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
+ Machine Learning + Reasoning + Scheduling + …
Service Plane
Chromium CGLX
SAGE MTP
OCCI JSDL
GIR UR
SNMP OpenFlow SDN / NSI
PerfSonar ICMP
Cassandra iRODs
Hadoop Storm
WSRF SensorML
WebServ INSPIRE
GRID/Cloud Computing

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

Doing Science

ICT to enable Science

Scientists live here!

Science App Store?

Wis dom

Knowledge

Information

Data

a.o. from ESFRI's

e-IRG

Workflows Schedulers

OWL

XML, RDF, rSpec, SNMP, Java based, etc.

MAGIC DATA CARPET

curation - description - trust - security - policy – integrity


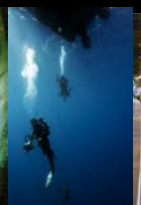




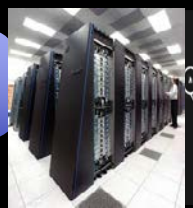
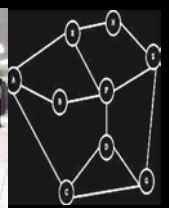
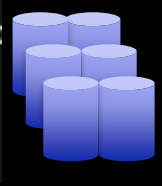

Machine Learning & Artificial Intelligence

- This is the era of Big Data & Machine Learning
- What if we apply ML&AI on all configuration, netflow and IDS data
- ➔ SmartNetworks
  - DoE workshop december 8-9 2016
  - http://www.orau.gov/smarthp2016/
- Challenges
  - How to get the user experience feed back to the ISP?
  - Privacy, integrity and security issues
  - Do we know what the AI actually learns?
    - Harness the complexity, e.g. in IOT.
    - It might actually come up with its own version of a packet travel ban! ;-)

TimeLine

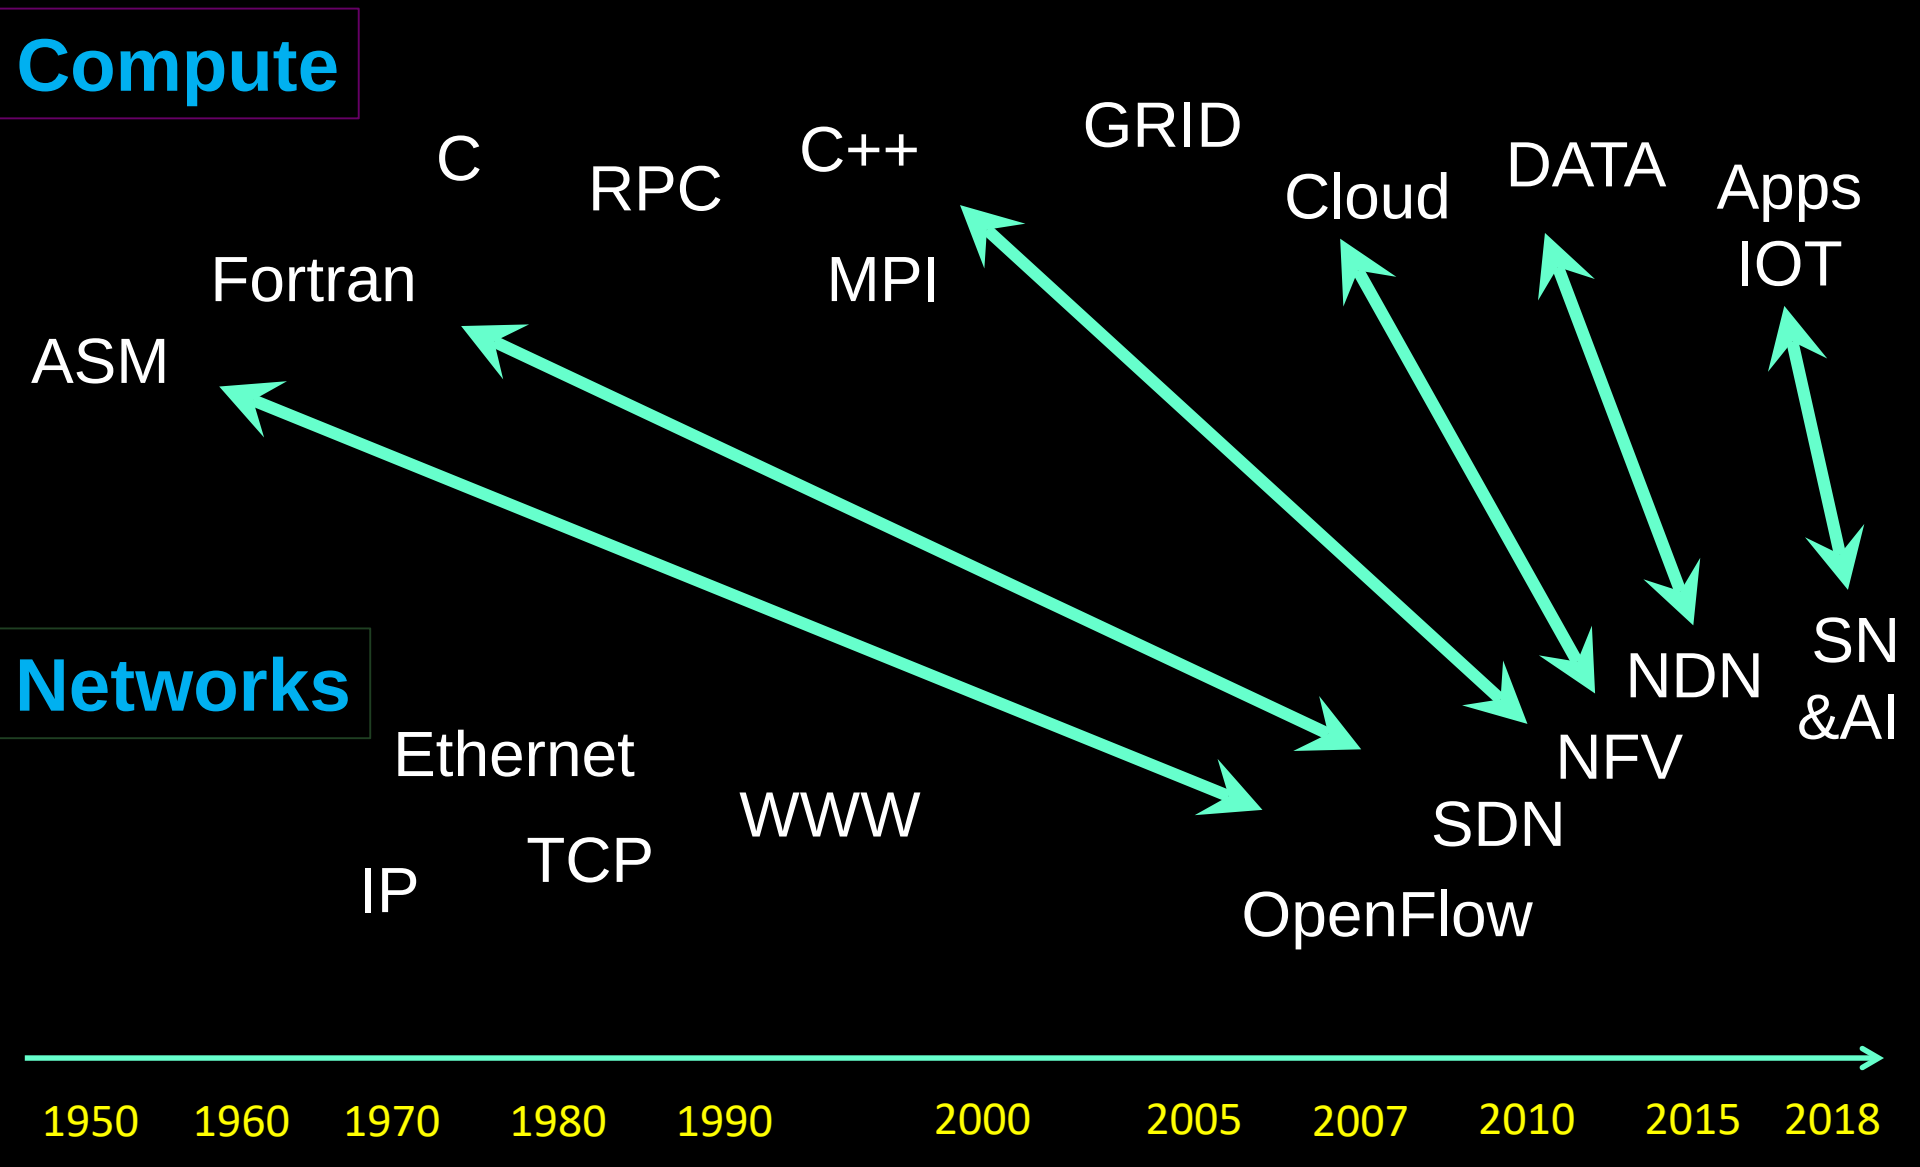
Compute
C
RPC
C++
GRID
Cloud
DATA
Apps
IOT
Fortran
MPI
ASM
Networks
Ethernet
WWW
TCP
IP
SDN
NFV
NDN
SN
&AI
OpenFlow
1950
1960
1970
1980
1990
2000
2005
2007
2010
2015
2018

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, Unix, c, SmallTalk, DECnet, TCP/IP, c++, Internet, WWW, Semantic Web, Photonic networks, Google, grid, cloud, BIG Data, Twitter, SDN, AI

to:

DDOS attacks destroying Banks and Bitcoins.

Congratulations CESNET!