The Internet Still Works. Amazing!

Cees de Laat
System & Network Engineering
University of Amsterdam

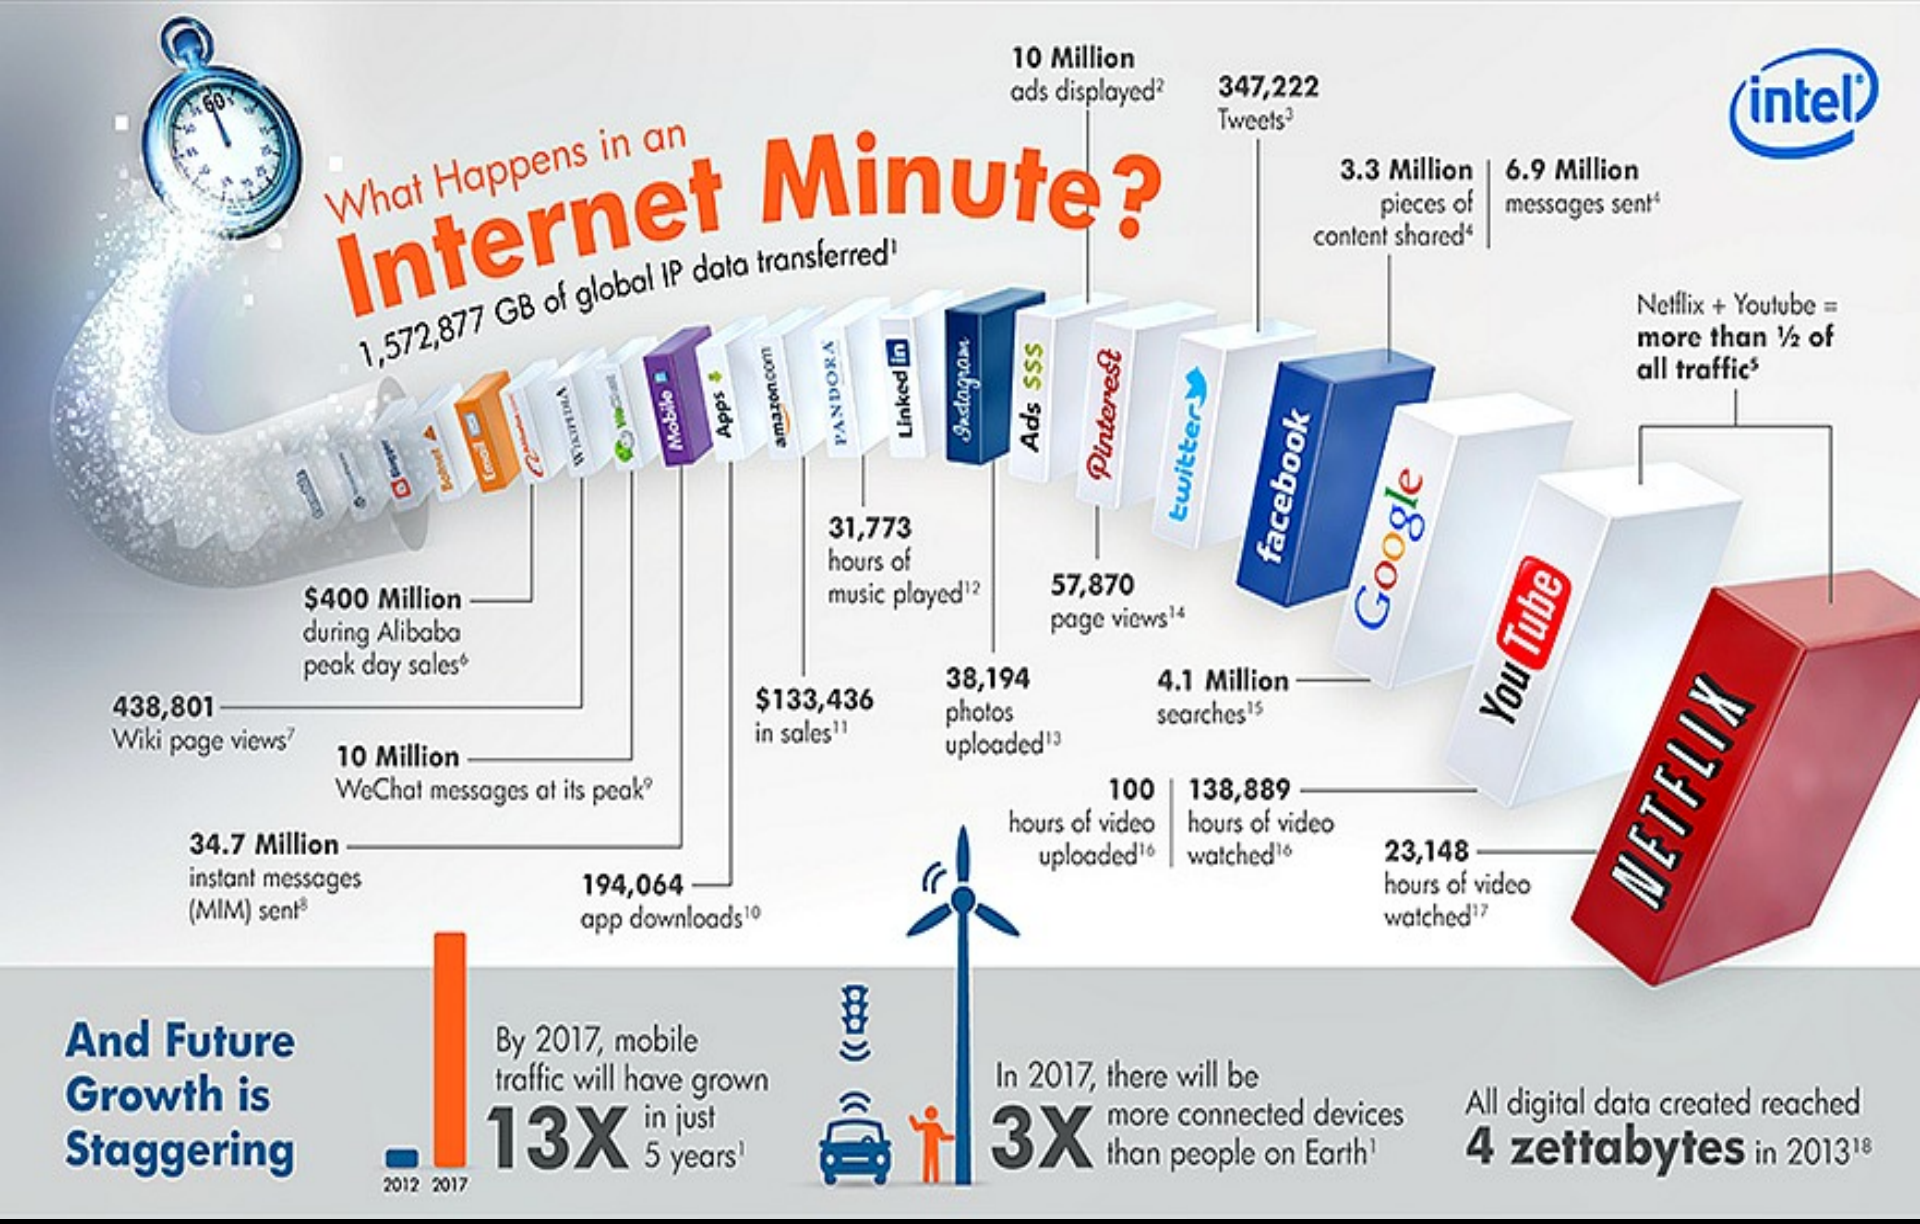

1,572,877 GByte/minute =
(8*1,572,877*10^9/60 bit/s)/(10*10^12 bit/s per fiber) =
21 fibers with each about 100 * 100 Gb/s channels

ATLAS detector @ CERN Geneve

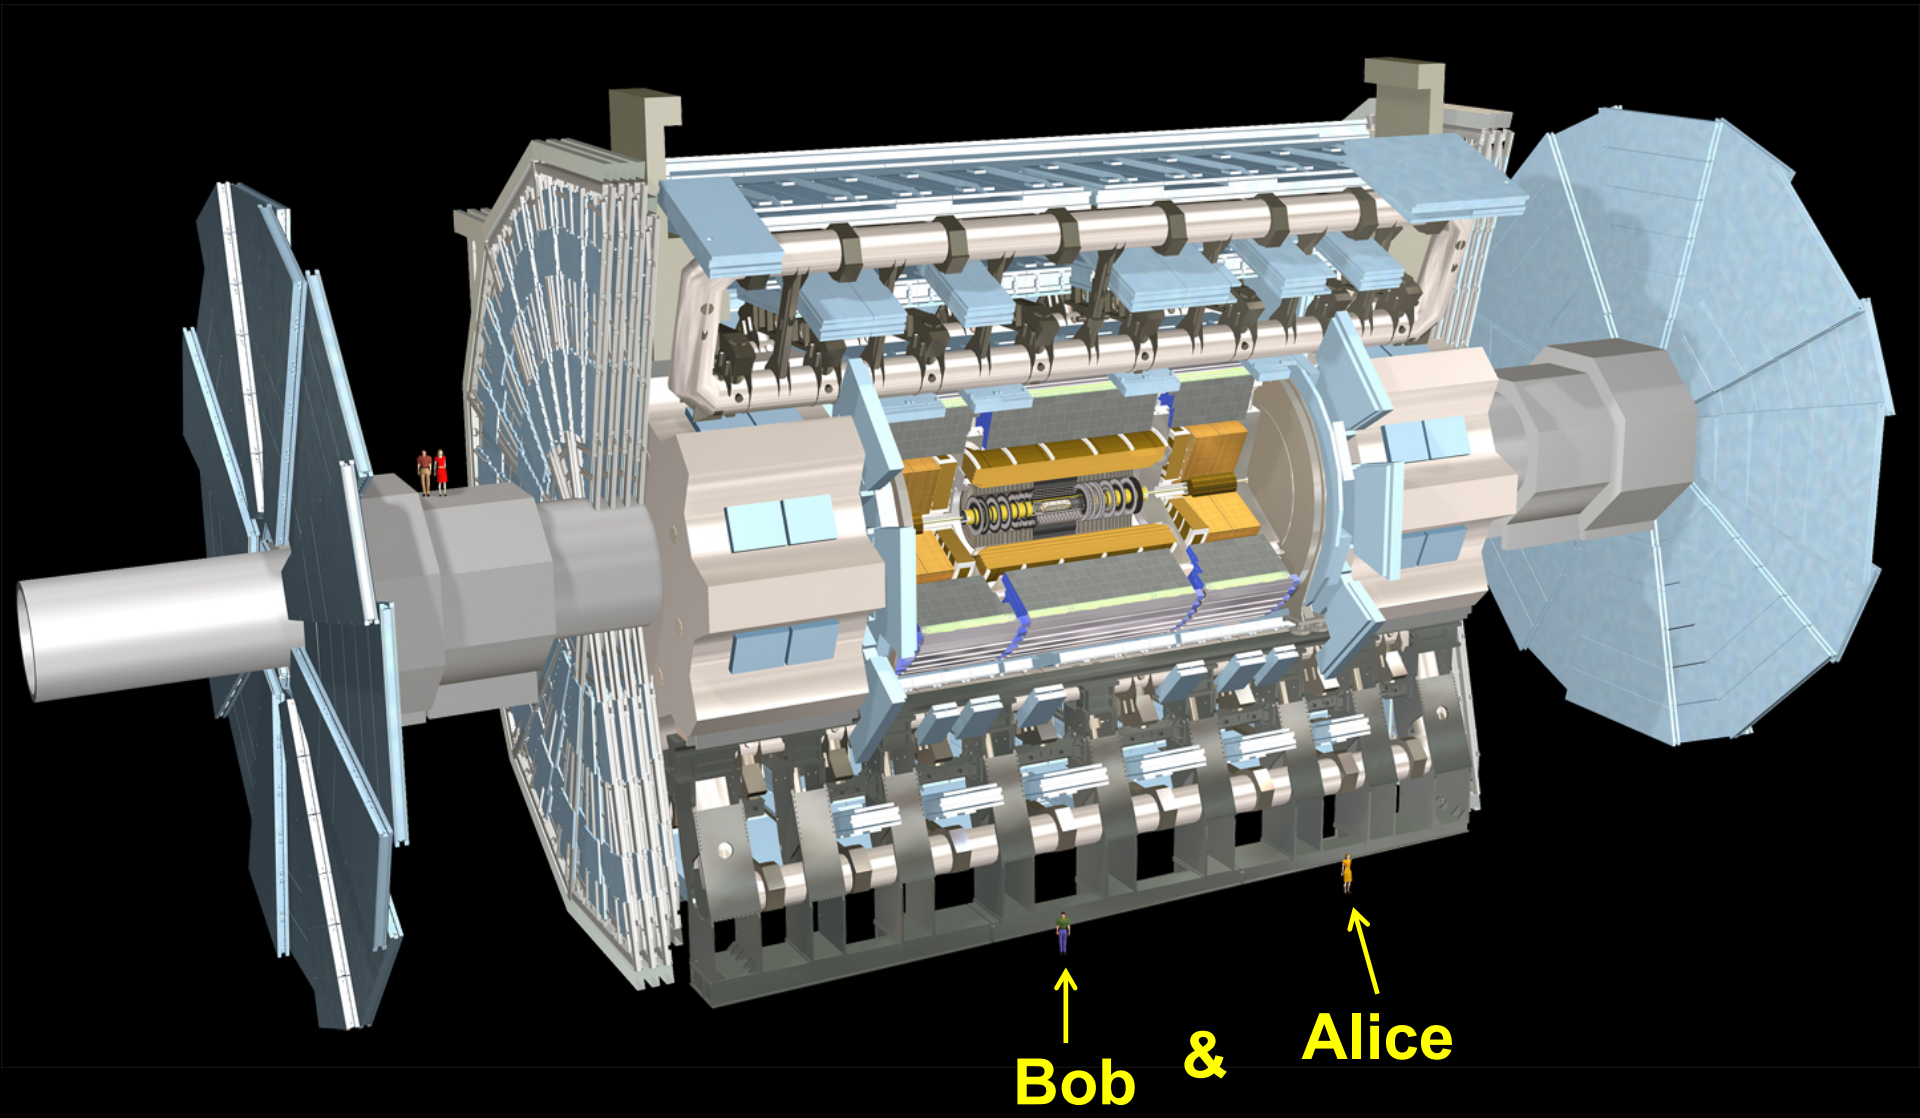

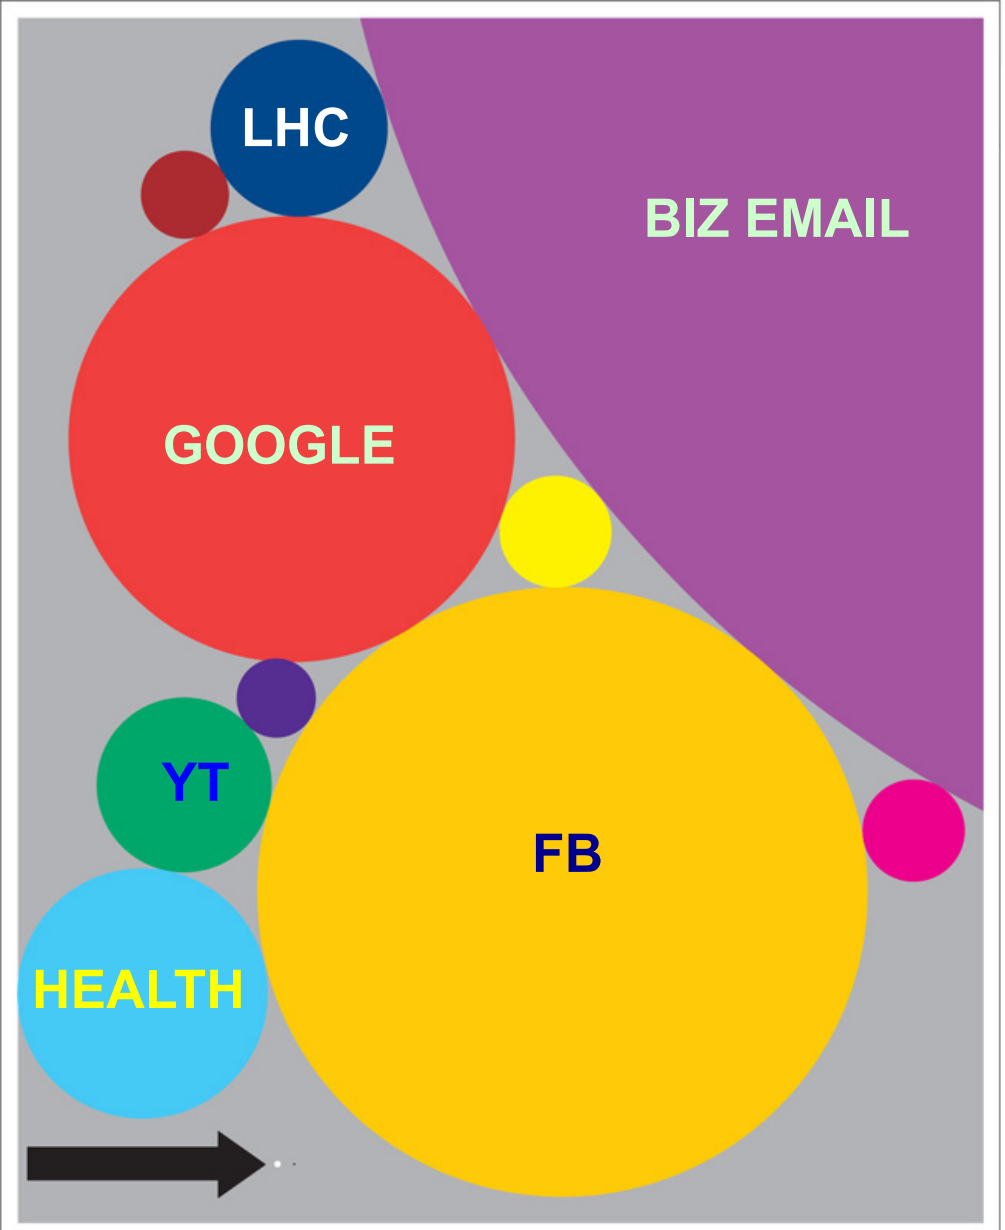

Size of data sets in terabytes

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

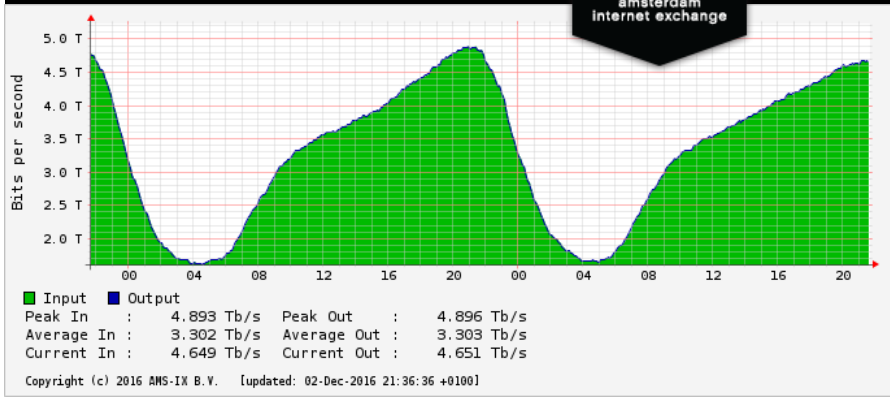

The Internet

- Developed between 1960's and 1980's
- Started out as series of experiments between a few hosts
- Lots of very very rough consensuses
- Spread exploded in the 1990's because of browsers
- Since the 1990's it is only patching because of installed base
- It seems to work, so it is used for ever more purposes
- About 4 Billion (Miljard) people online
- Now already essential for daily life
- HOWEVER:

HOWEVER:

- Developed between 1960's and 1980's
  - With technical limitations of those days
- Started out as series of experiments between a few hosts
  - Everybody trusted each other, buddies
- Lots of very very rough consensuses
  - Not invented here, political power, commercial pressure
- Spread exploded in the 1990's because of browsers
  - DNS and WWW made it usable for non scientists
- Since the 1990's it is only patching because of installed base
  - IPv6 std in mid to end 90's, scrambling with NAT, loc/ident
- It seems to work, so it is used for ever more purposes
  - Entertainment, newspapers, IOT, phone calls, business
- About 4 Billion (Miljard) people online
  - Ran out off addresses
- Now already essential for daily life
  - Talking with government, industry, services providers
- So?

Trends in Networking

... more data!

... more realtime!

... more users!

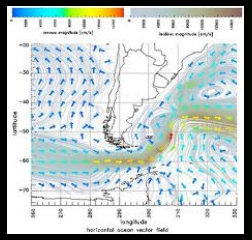

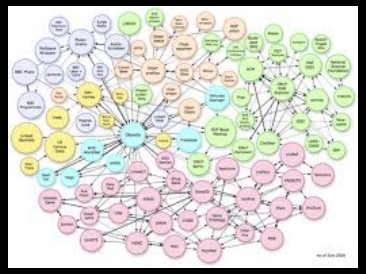

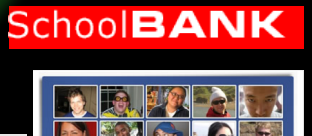

Trends in Networking

... more data!

Speed

Volume

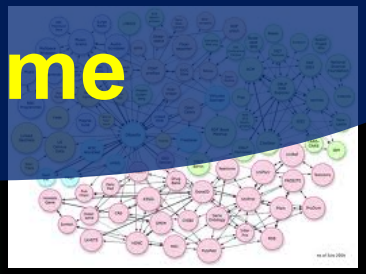

Deterministic

Real-time

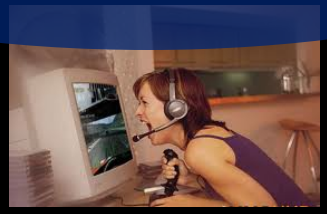

Scalable

Secure

... more users!

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

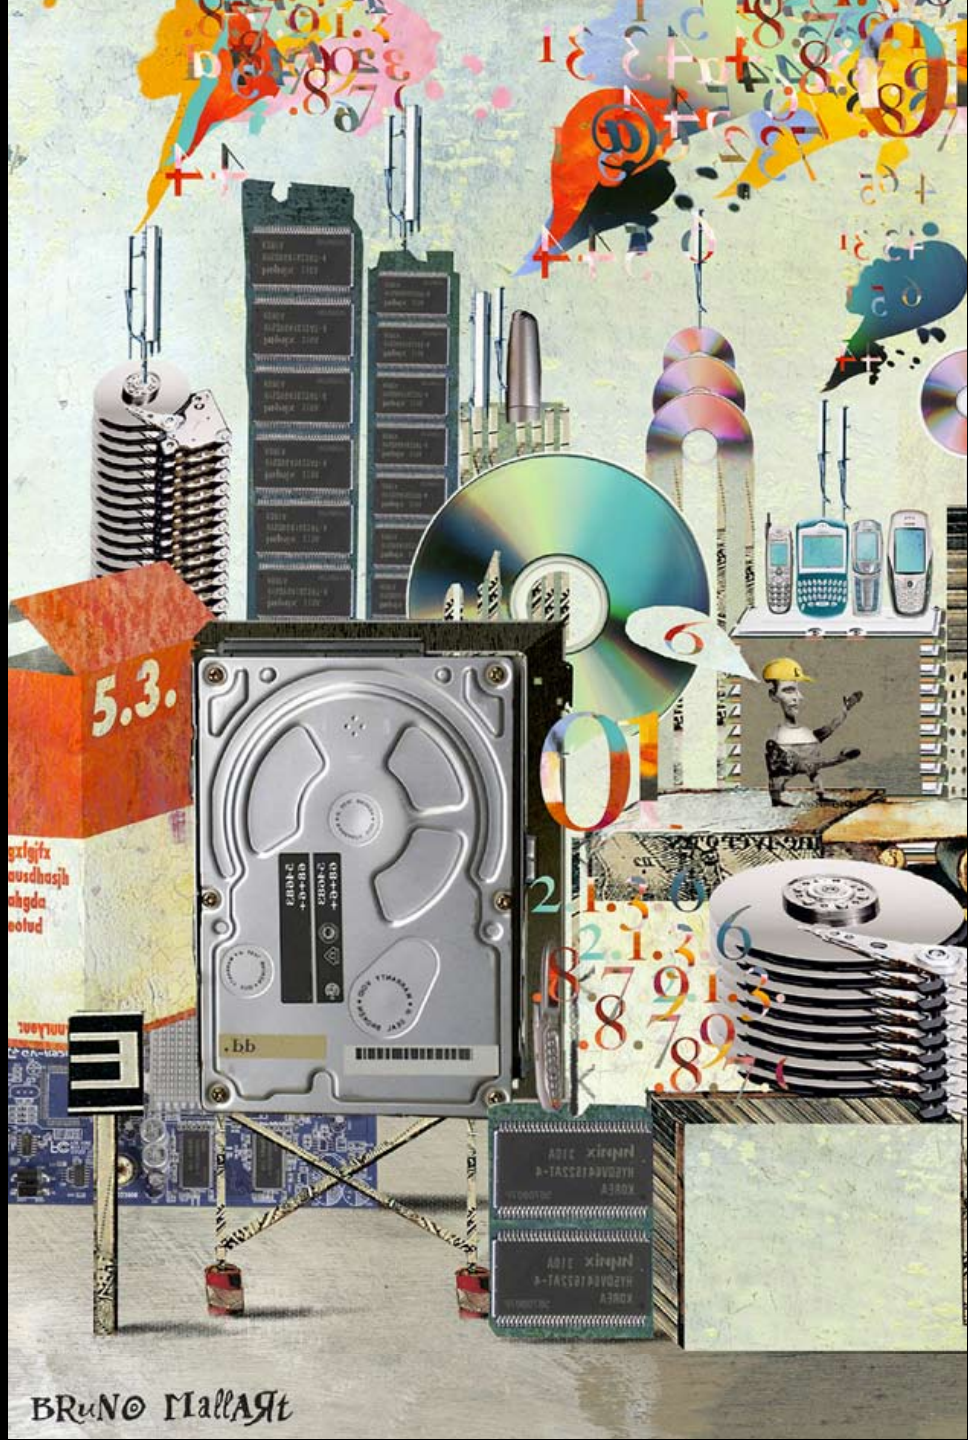

Reduction of Complexity by Integration

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?

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

Supers & Cloud

- Science computing dwarfed by Cloud
- Sweet point between general computing (cloud) and Mission computing
- In 5 to 10 years science computers may be hard to defend
- Cloud providers
  - Economy of scale
  - 24 * 7 operations
  - Big buying power -> define what the market delivers
  - Logistics
  - but no knowledge on Science algorithms
  - ➔ Software as a Service!
  - ➔ Learn to map algorithms to cloud!

So who has the world's largest data center?
We've seen a lot of huge data centers in our travels, and have identified 10 that we believe are the largest found anywhere. These data fortresses range between 400,000 and 1.1 million square feet.

http://www.datacenterknowledge.com/special-report-the-worlds-largest-data-centers/worlds-largest-data-center-350-e-cermak/

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.

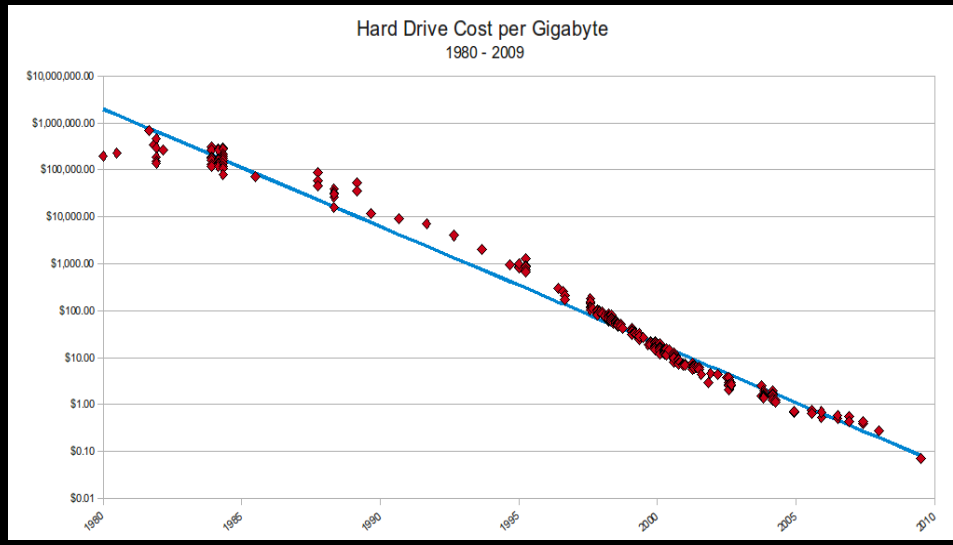

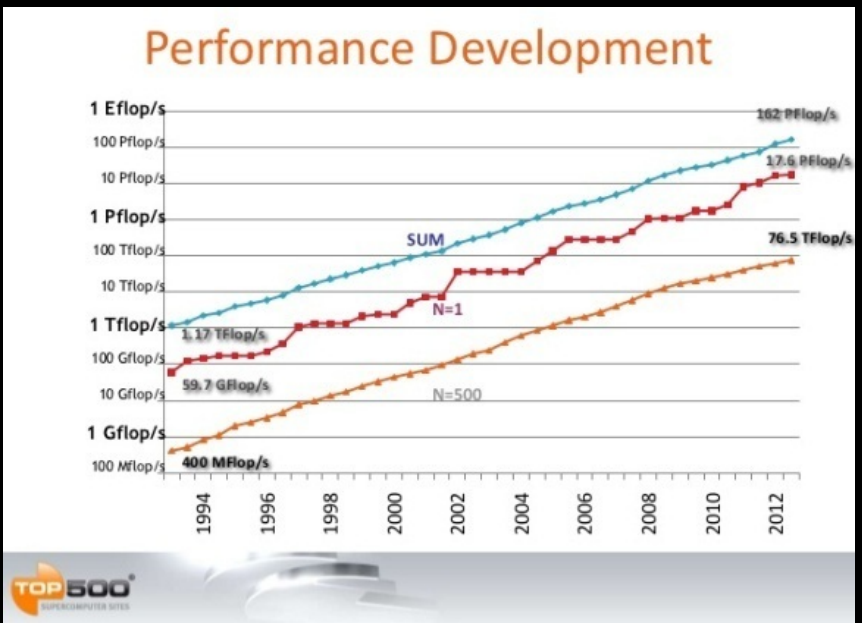

We will soon reach the limits of what is currently feasible and controllable.

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

Trucks of Tapes

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

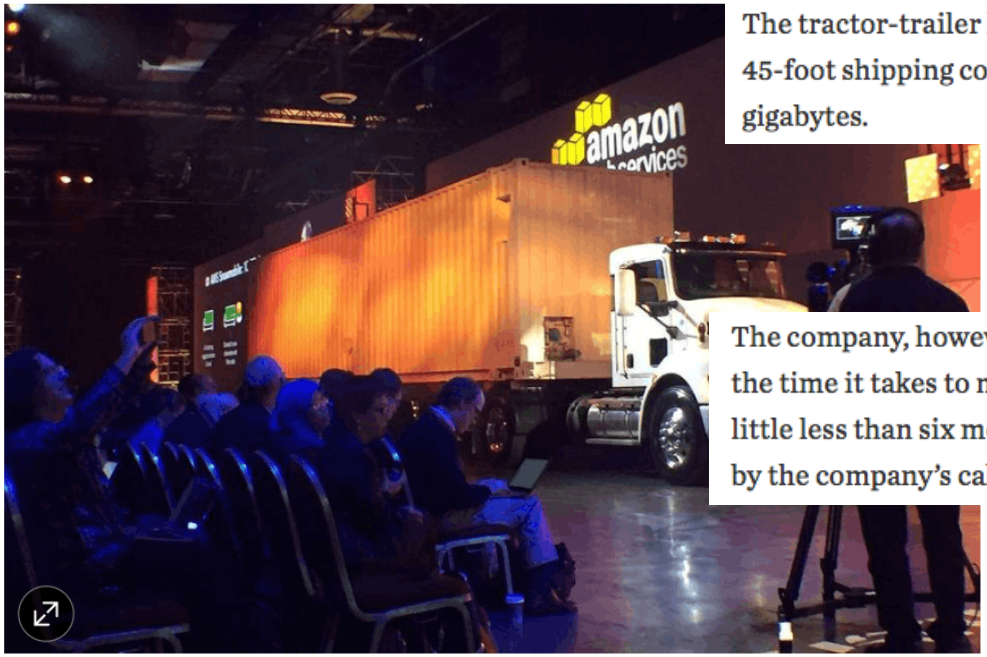

Amazon unveiled the 'Snowmobile' service on Wednesday in Las Vegas. *PHOTO: AMAZON WEB SERVICES*

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.

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

4 COMMENTS

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.

Most Popular

1. U.S. to Fo... Least $10... Student ... Coming ...
2. Opinion: Trump's ... Pick Sca...
3. Trump's ... His Busi... Draws Q...
4. Creator ... Mac Dies
5. Trump's ... Choice S... Absolute...

1 fiber does about 16 Tbit/s
= 2 Tbyte/s
⇒ 500000 s/ExaByte
⇒ One week/ExaByte
Or stick Joe and Harvey in a RV for 2 months.

Fading Trust in Internet

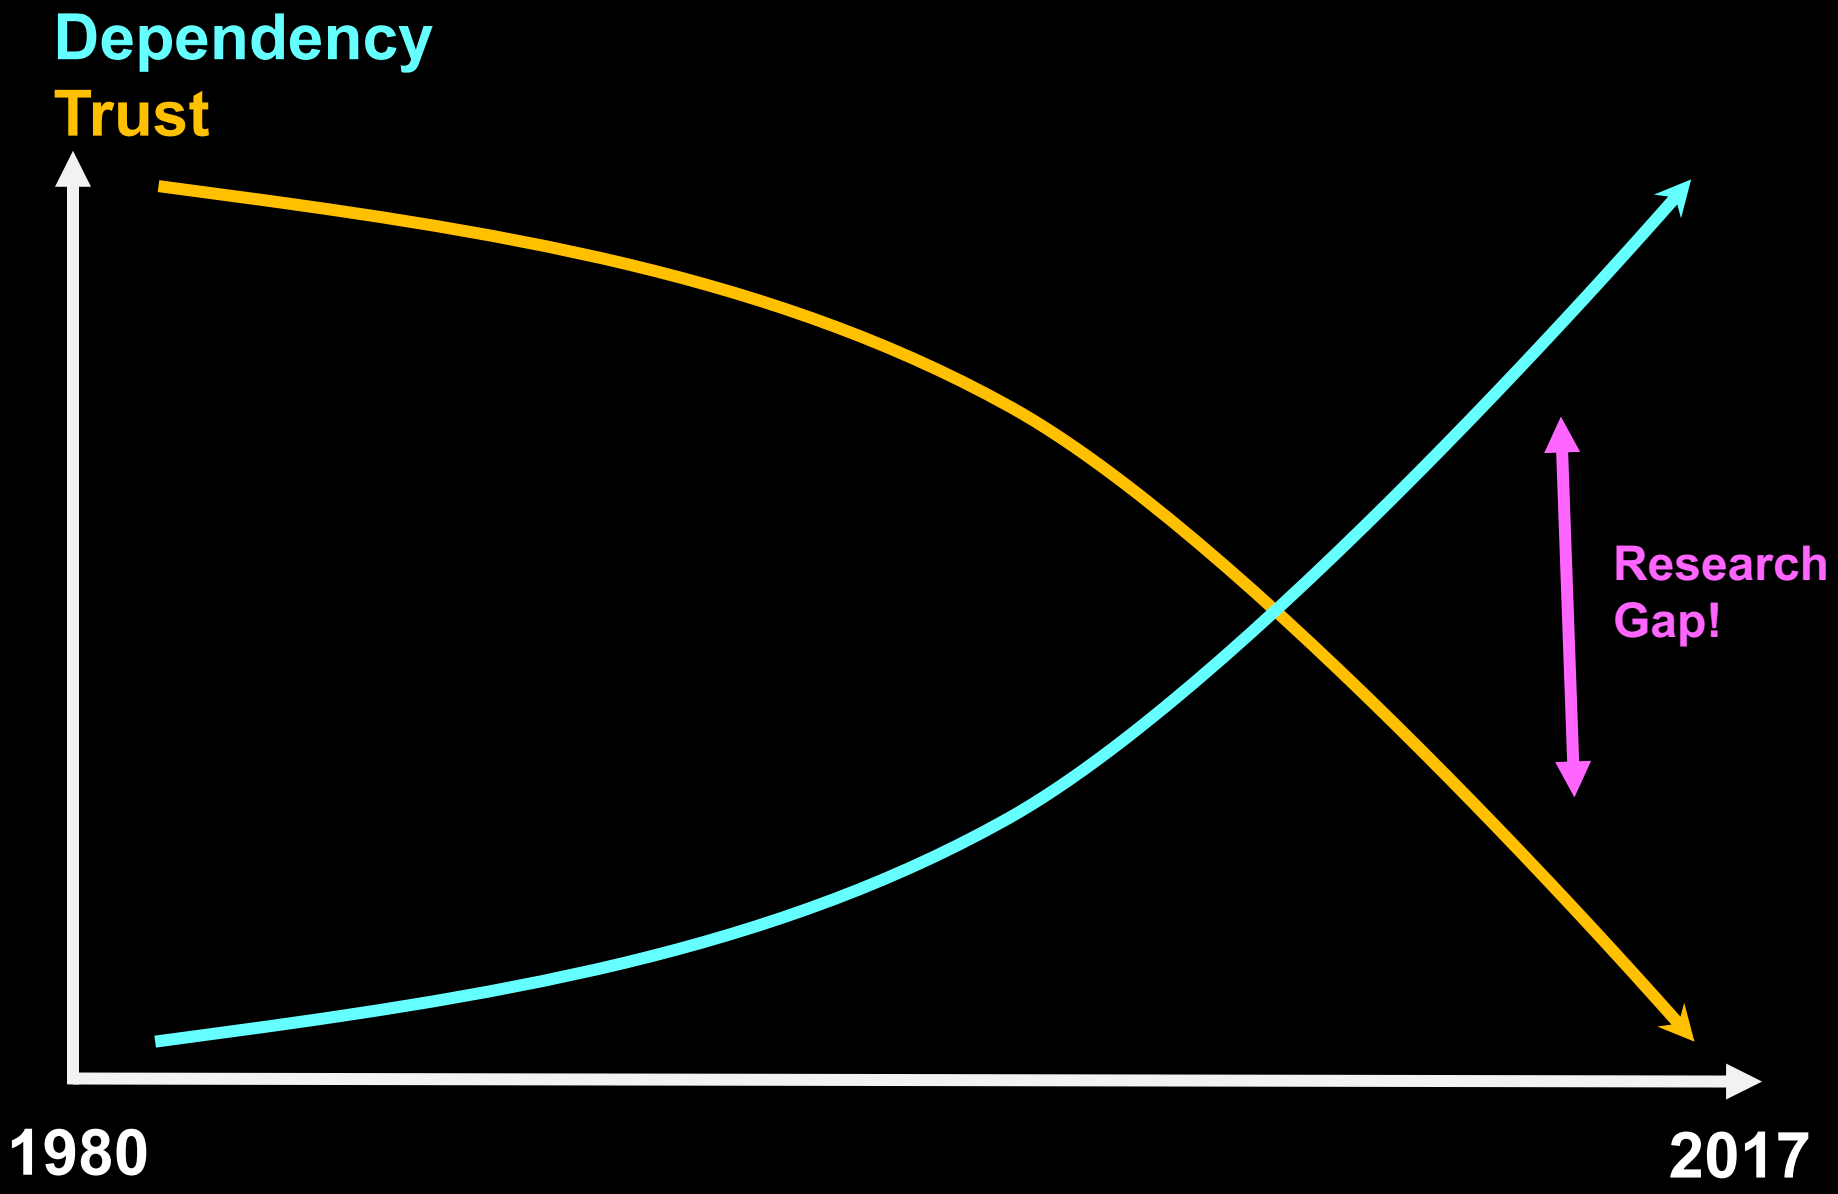

Technical developments

- ~1999: Dark Fiber, TELCO's out of the way
  - Multi Layer services, open infrastructure, unbundling
- ~2002: High Speed Networking
  - 100's of Gbit/s transport protocols
- ~2004: AAA and GRID
  - Cross domain resource reservation
- ~2008: Software defined Networking
  - Separation of control and forwarding plane
- ~2010: Clouds
  - Infra as a service, slices
- ~2012: Overlay networks
  - Fed4Fire, GENI, PlanetLab
- 2015: Internet of Things
  - Smart Cities, personal Health,

Back to Internet

2006: Raj Jain, Washington univ in Saint Louis MO

Internet 3.0: Ten Problems with Current Internet Architecture and Solutions for the Next Generation

http://www.cse.wustl.edu/~jain/papers/ftp/gina.pdf

IV. TOP TEN FEATURES REQUIRED IN THE NEXT GENERATION INTERNET ARCHITECTURE

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

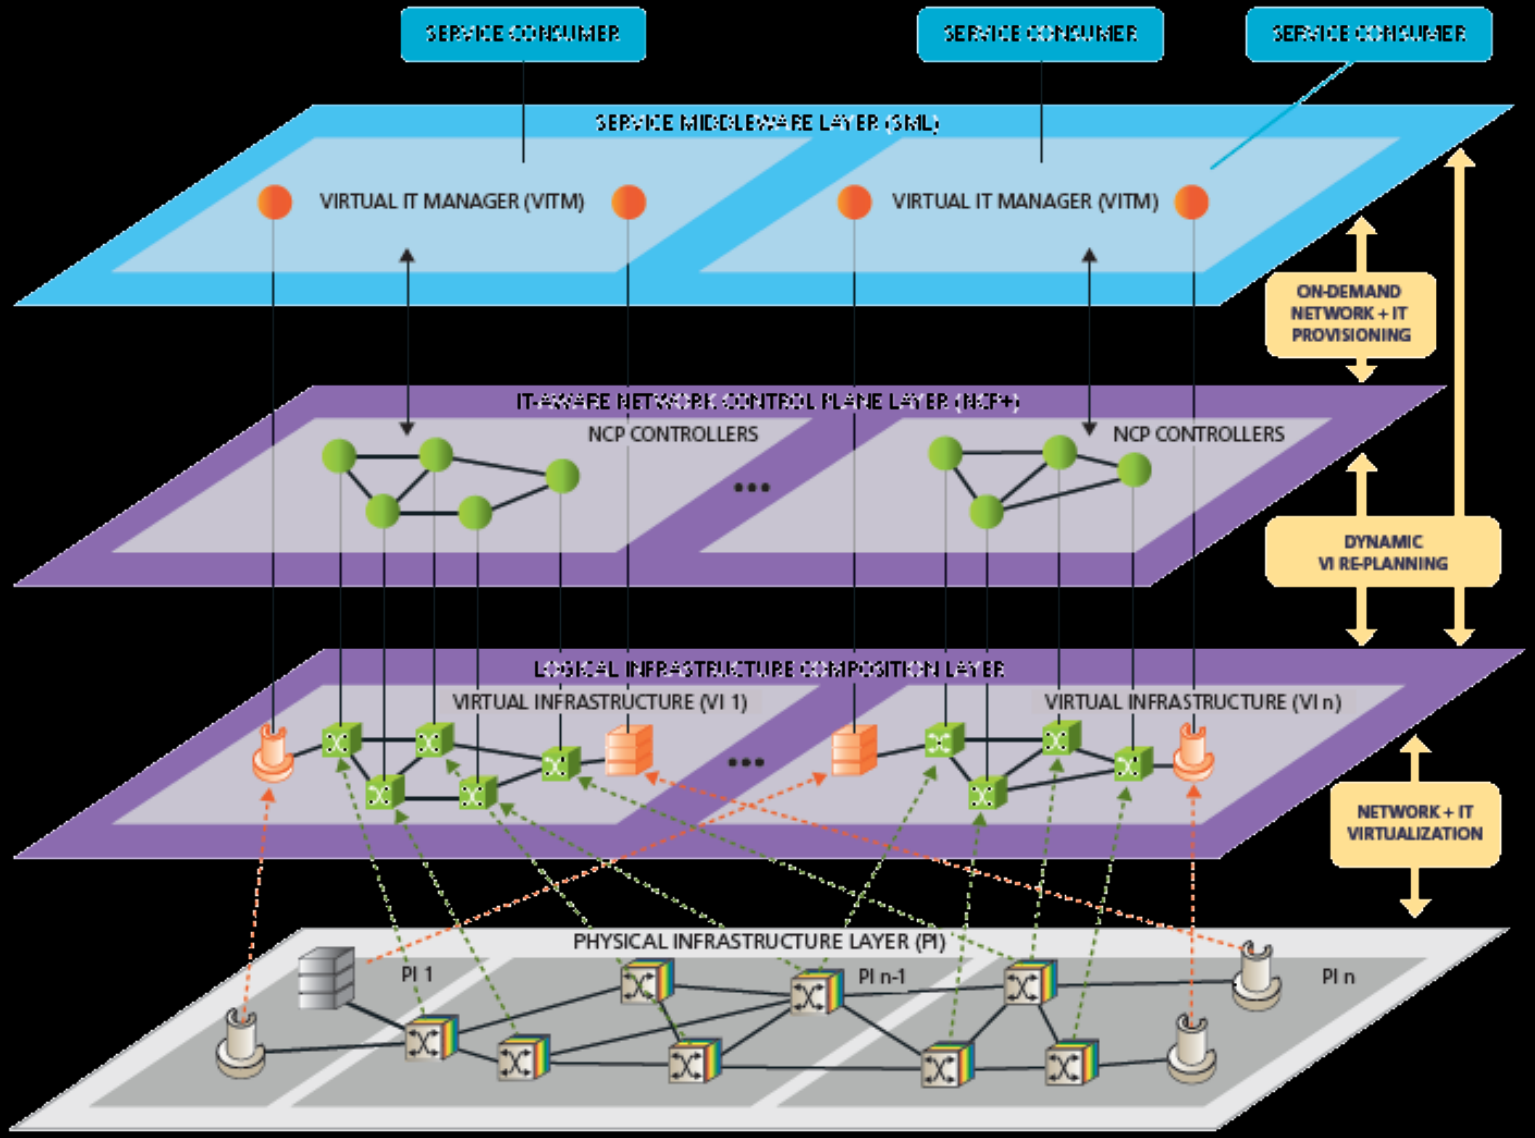
SERVICE CONSUMER
SERVICE CONSUMER
SERVICE CONSUMER
VIRTUAL IT MANAGER (VITM)
VIRTUAL IT MANAGER (VITM)
ON-DEMAND NETWORK + IT PROVISIONING
NCP CONTROLLERS
NCP CONTROLLERS
DYNAMIC VI RE-PLANNING
VIRTUAL INFRASTRUCTURE (VI 1)
VIRTUAL INFRASTRUCTURE (VI n)
NETWORK + IT VIRTUALIZATION
PHYSICAL INFRASTRUCTURE LAYER (PI)
PI 1
PI n-1
PI n

Envisioned role of the SPG: define slice archetypes?

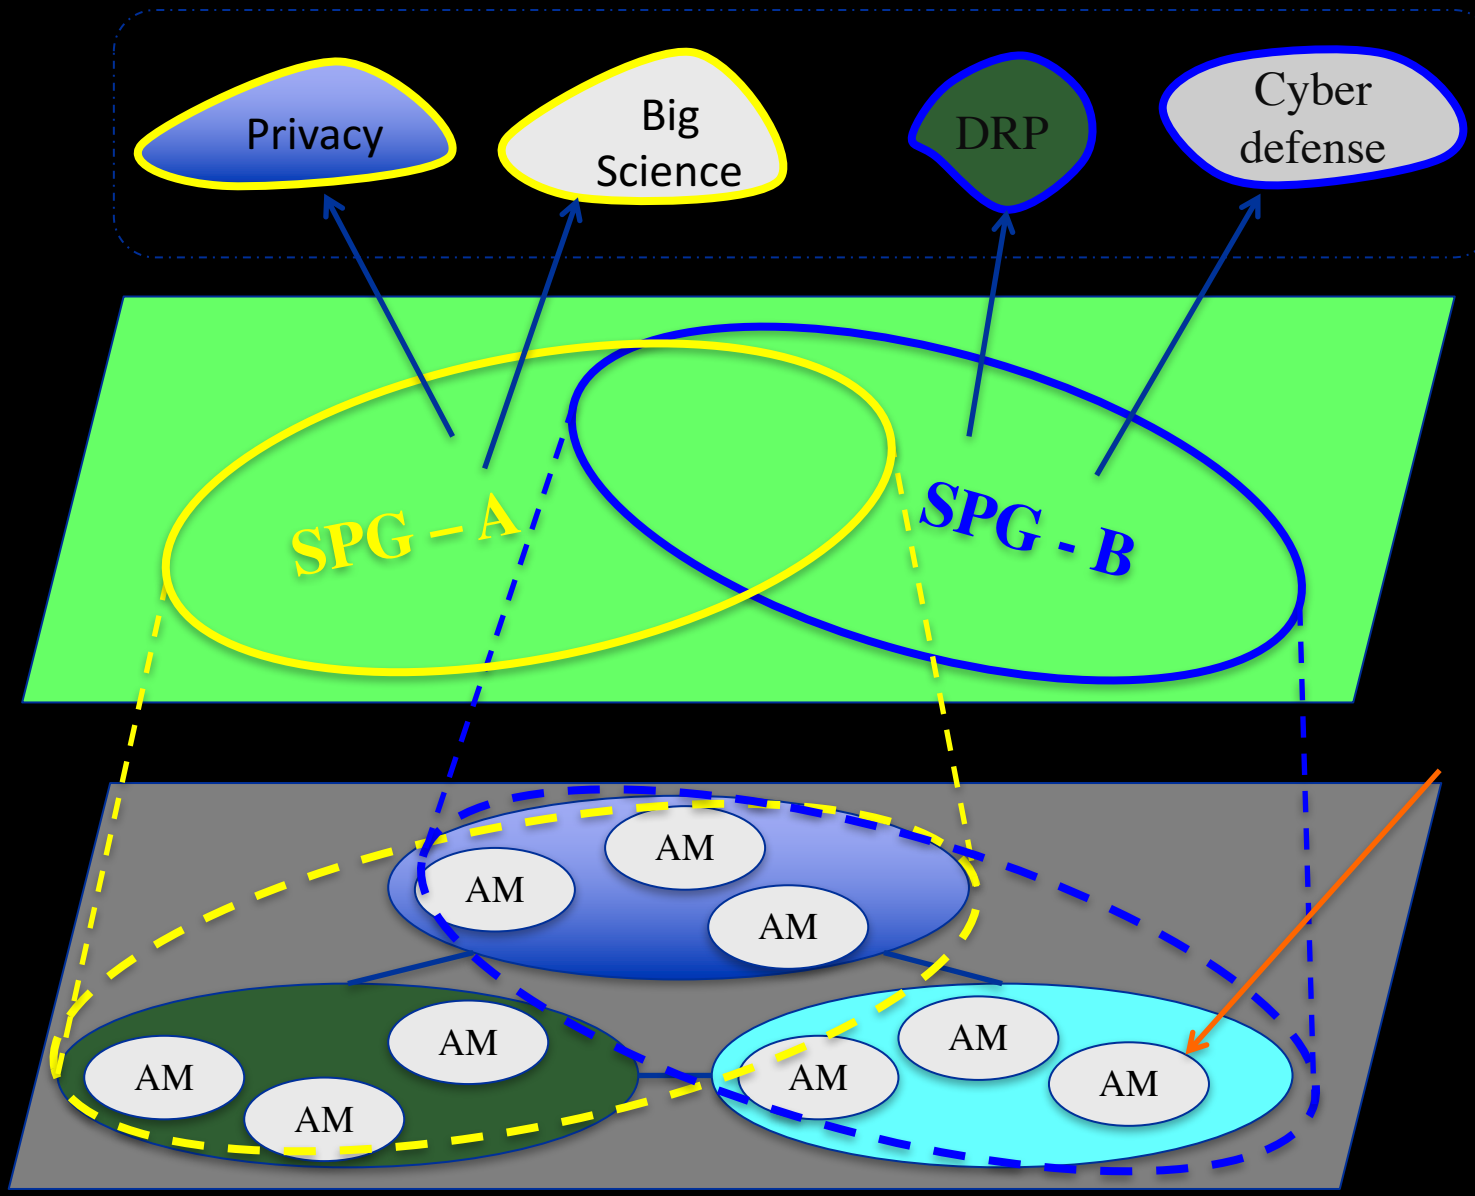

ExoGeni

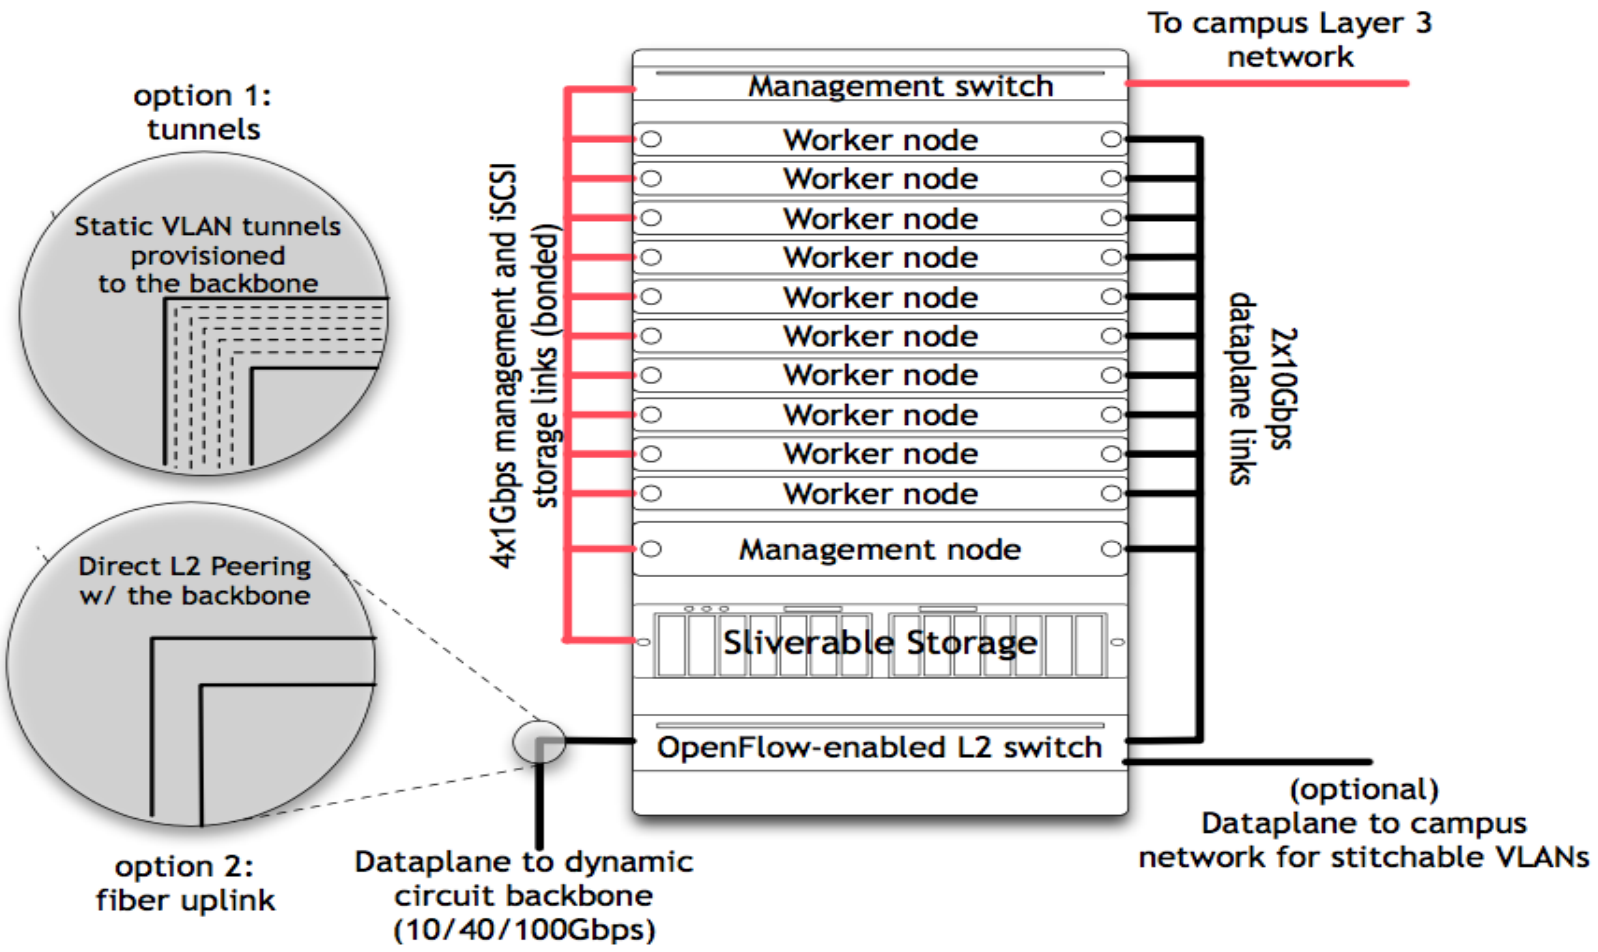

Fig. 1.2. Structure of an ExoGENI site rack for the initial deployment. Each rack has low-bandwidth IP connectivity for management and a high-bandwidth hybrid OpenFlow switch for the slice dataplanes. The site ORCA server controls L2 dataplane connections among local nodes and external circuits.

Fed4Fire

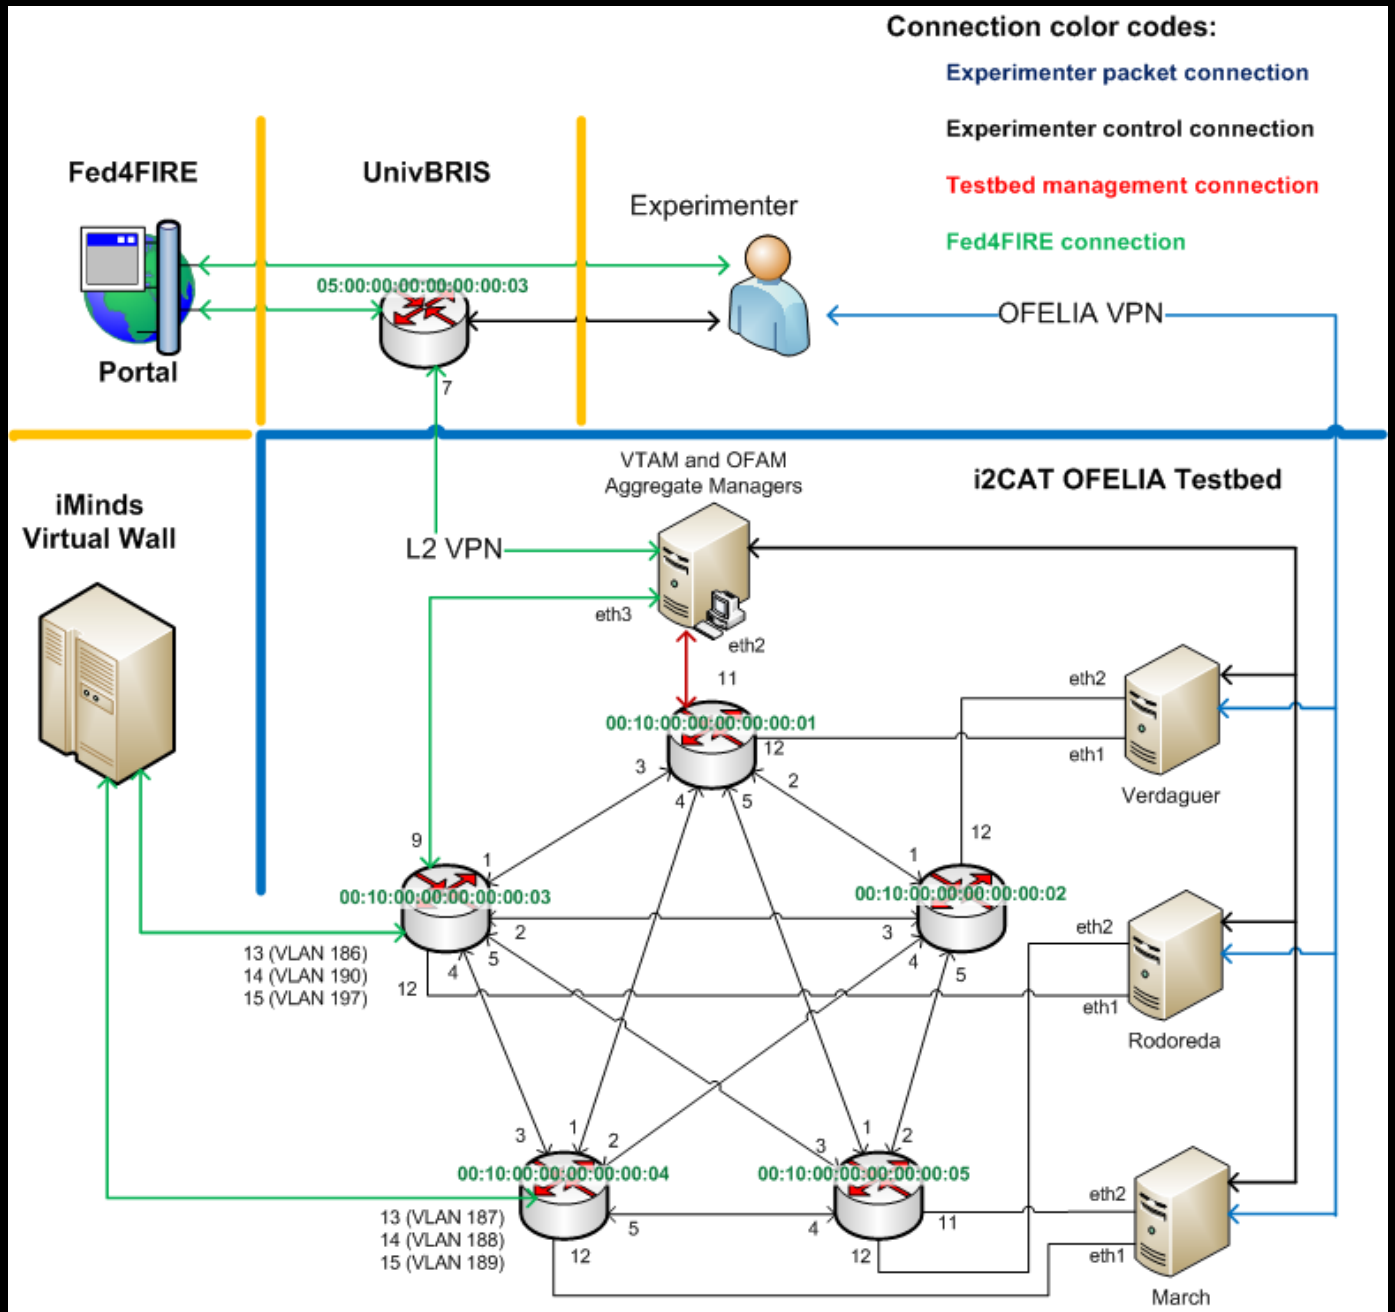
Connection color codes:
Experimenter packet connection
Experimenter control connection
Testbed management connection
Fed4FIRE connection
Fed4FIRE
UnivBRIS
Experimenter
Portal
05:00:00:00:00:00:00:03
7
OFELIA VPN
VTAM and OFAM
Aggregate Managers
i2CAT OFELIA Testbed
iMinds
Virtual Wall
L2 VPN
eth3
eth2
11
00:10:00:00:00:00:00:01
12
3
4
5
2
eth2
eth1
Verdaguer
9
1
00:10:00:00:00:00:00:03
2
5
4
12
13 (VLAN 186)
14 (VLAN 190)
15 (VLAN 197)
1
00:10:00:00:00:00:00:02
3
4
5
12
eth2
eth1
Rodoreda
3
1
2
00:10:00:00:00:00:00:04
5
12
13 (VLAN 187)
14 (VLAN 188)
15 (VLAN 189)
3
1
2
00:10:00:00:00:00:00:05
4
11
12
eth2
eth1
March

Fed4Fire

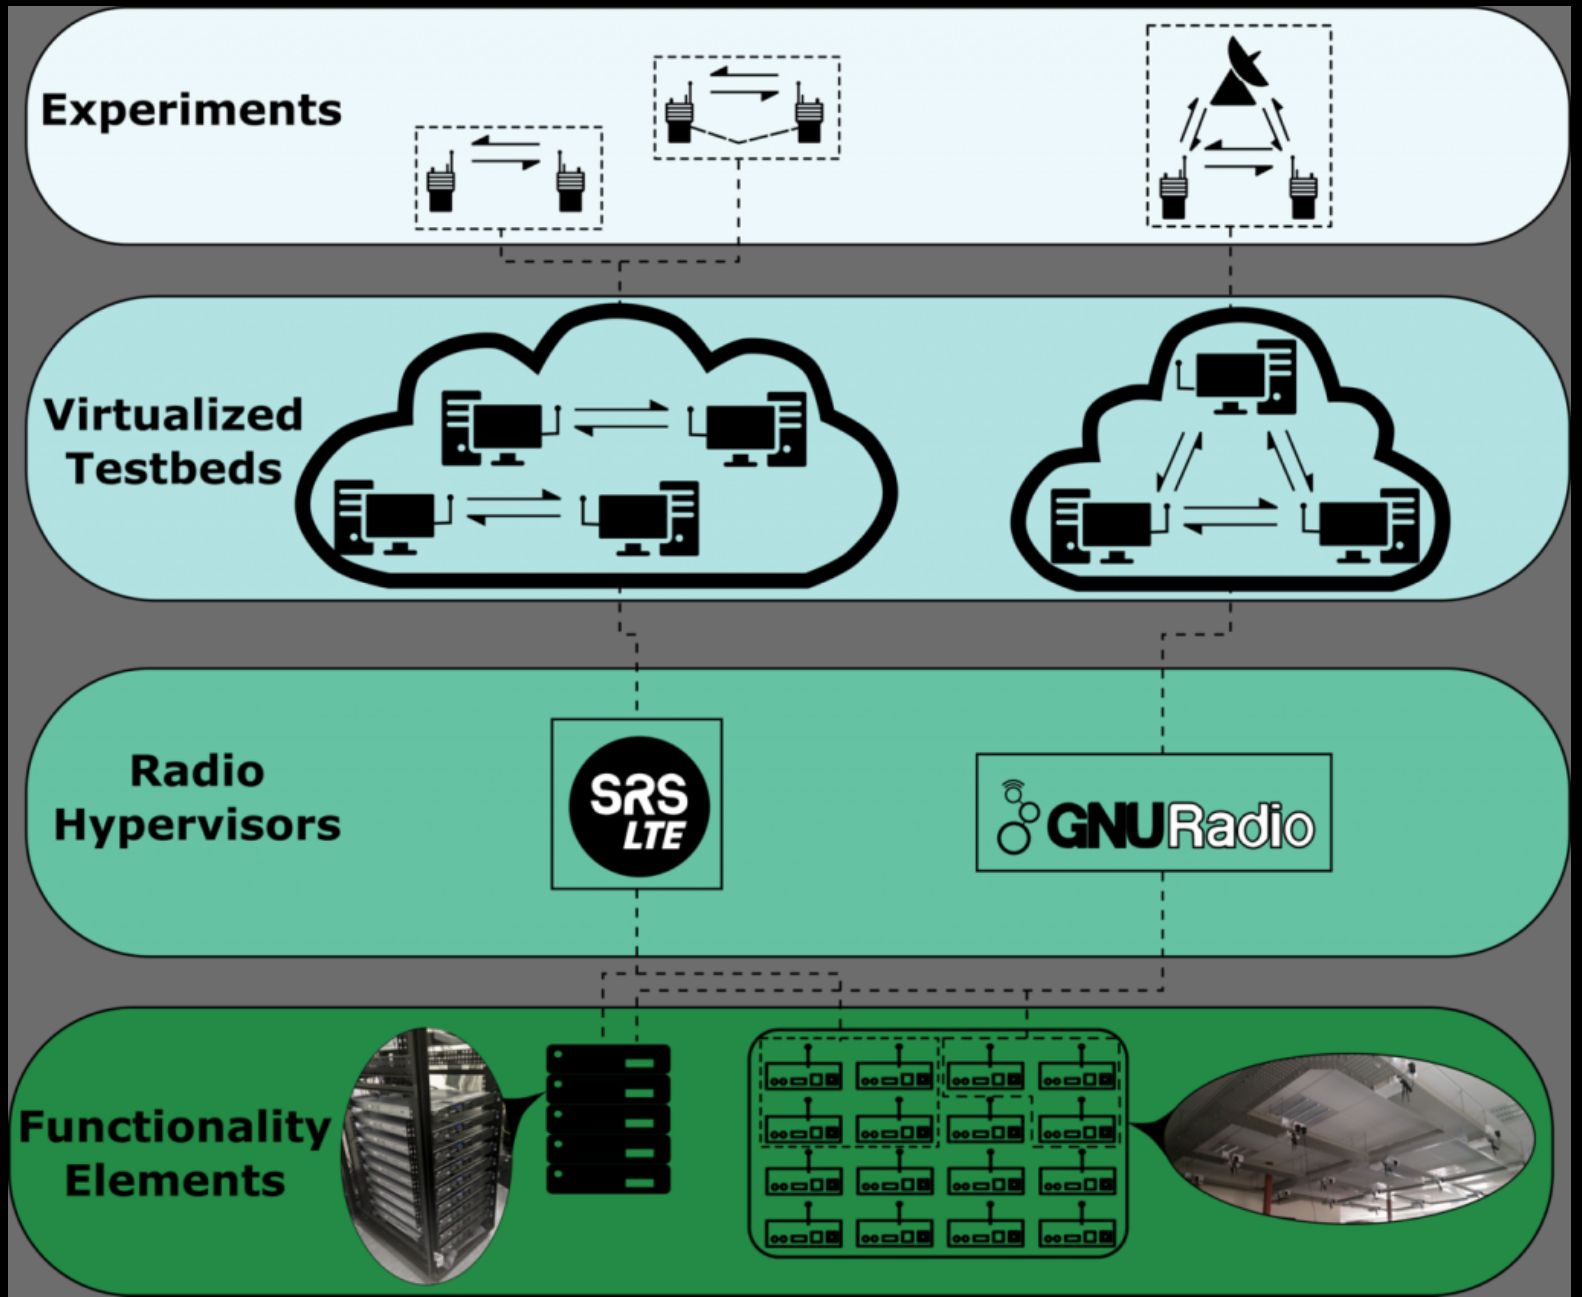
Experiments
Virtualized Testbeds
Radio Hypervisors
SRS LTE
GNURadio
Functionality Elements

Fed4Fire

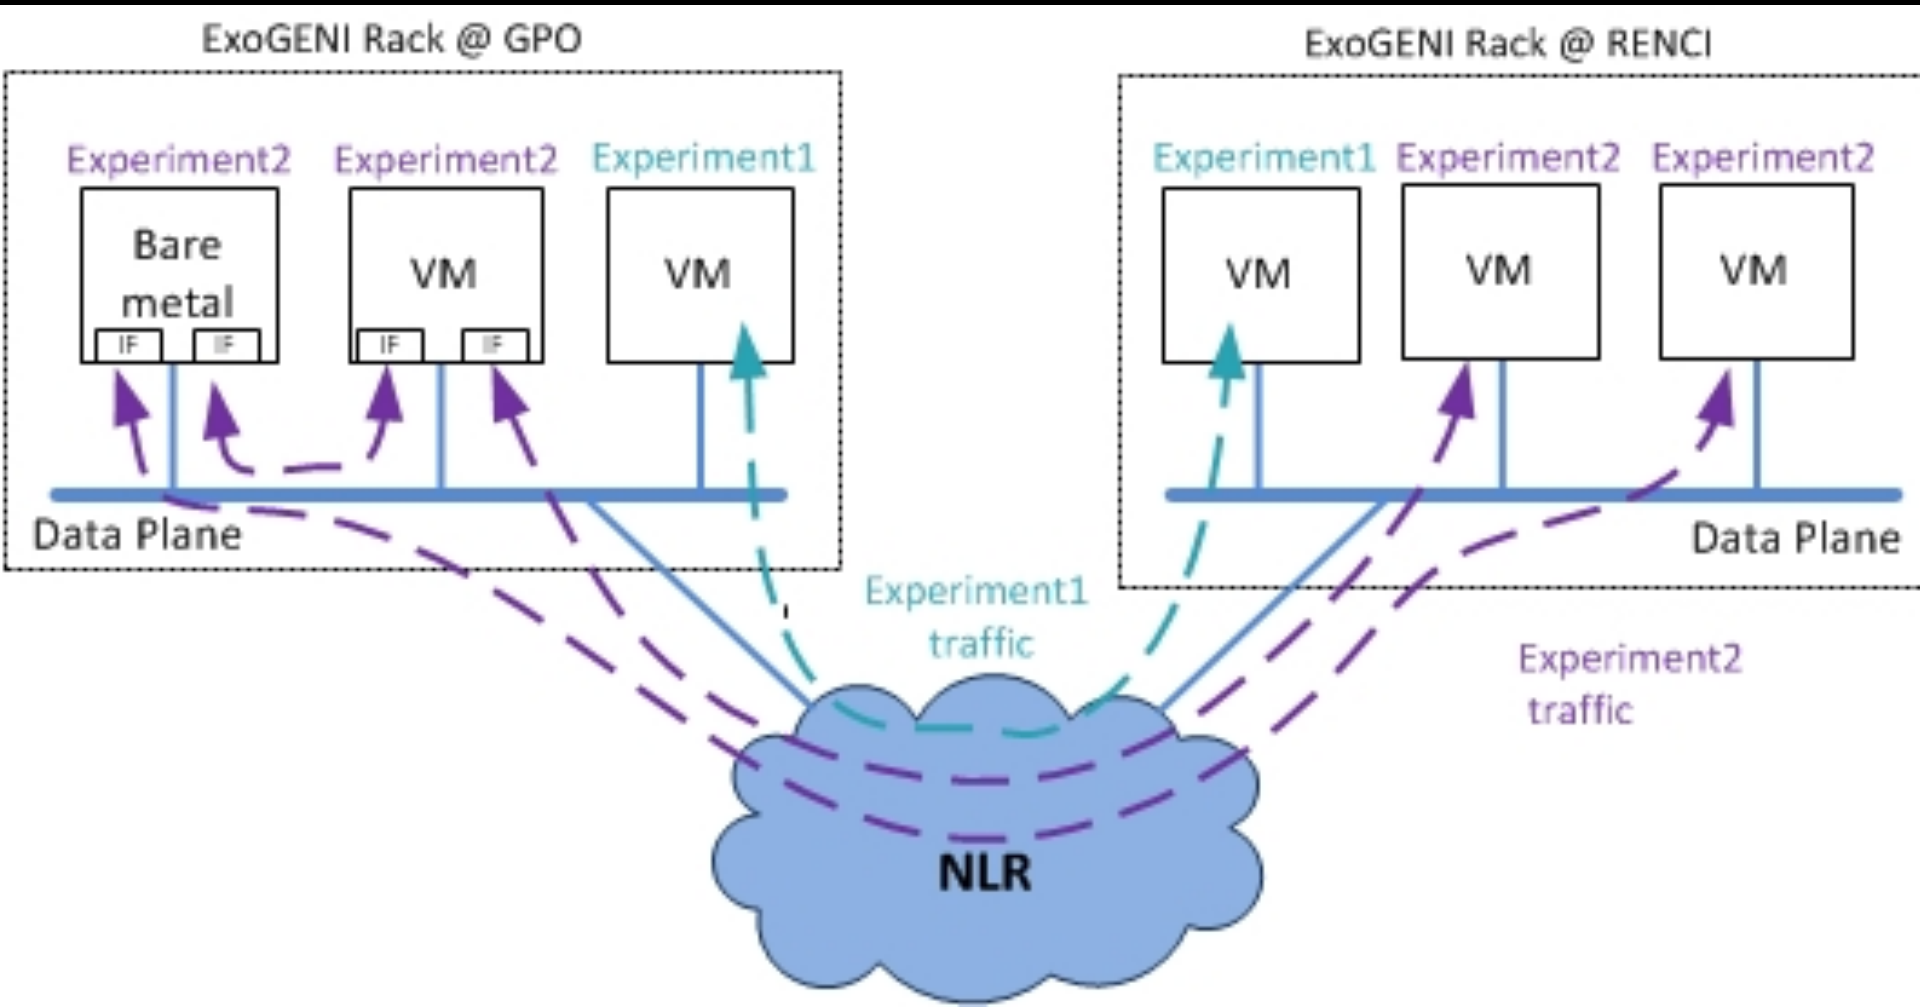
ExoGENI Rack @ GPO
Experiment2
Experiment2
Experiment1
Bare metal
VM
VM
IF
IF
IF
IF
Data Plane
ExoGENI Rack @ RENCI
Experiment1
Experiment2
Experiment2
VM
VM
VM
Data Plane
Experiment1 traffic
Experiment2 traffic
NLR

Some developments: RINA

- Recursive InterNetwork Architecture (RINA)
  - Recursive InterNetwork Architecture (RINA) is a computer network architecture that unifies distributed computing and telecommunications.
  - RINA's fundamental principle is that computer networking is just Inter-Process Communication or IPC. RINA reconstructs the overall structure of the Internet, forming a model that comprises a single repeating layer, the DIF (Distributed IPC Facility), which is the minimal set of components required to allow distributed IPC between application processes.
  - RINA inherently supports
    - mobility, multi-homing and Quality of Service without the need for extra mechanisms,
    - provides a secure and programmable environment,
    - motivates for a more competitive marketplace,
    - and allows for a seamless adoption.

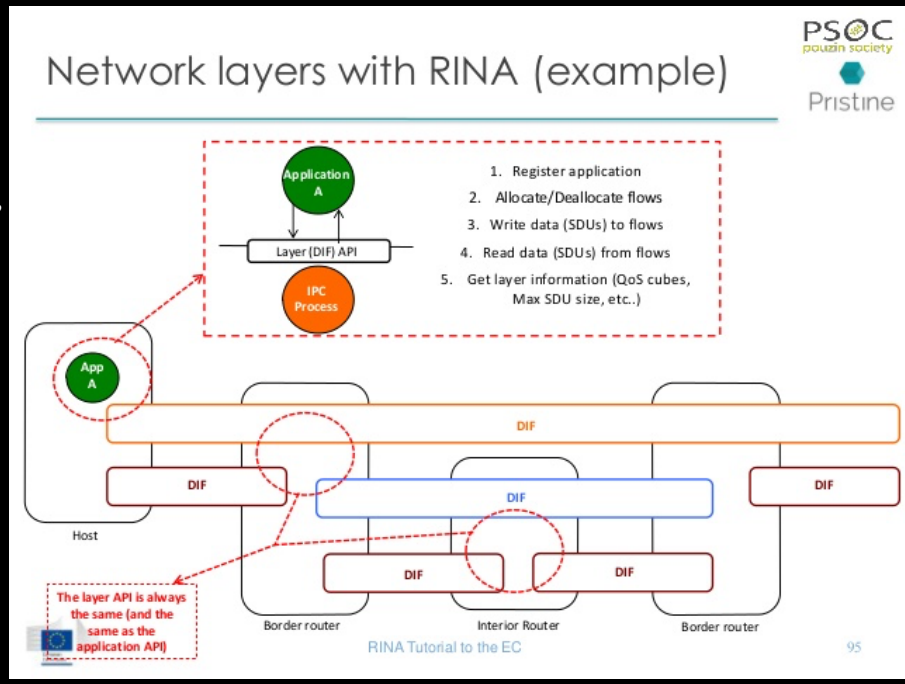

Learned from Scinet & INDIS

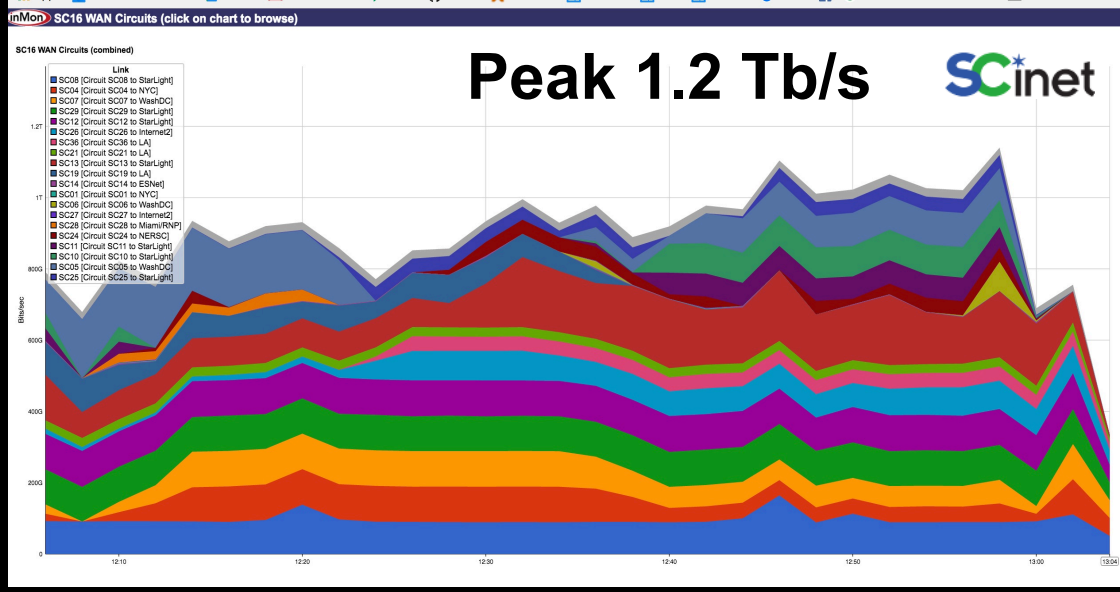

- 2013 - 2016
  - SDN
  - Security
  - Traffic management, policing, control
  - Hybrid – optical ring - approach to reach Tb/s
- 2017 - 2020
  - NFV
  - SDX
  - DTN @ core ➔ petabyte email network
  - Data abstractions (e.g. NDN)

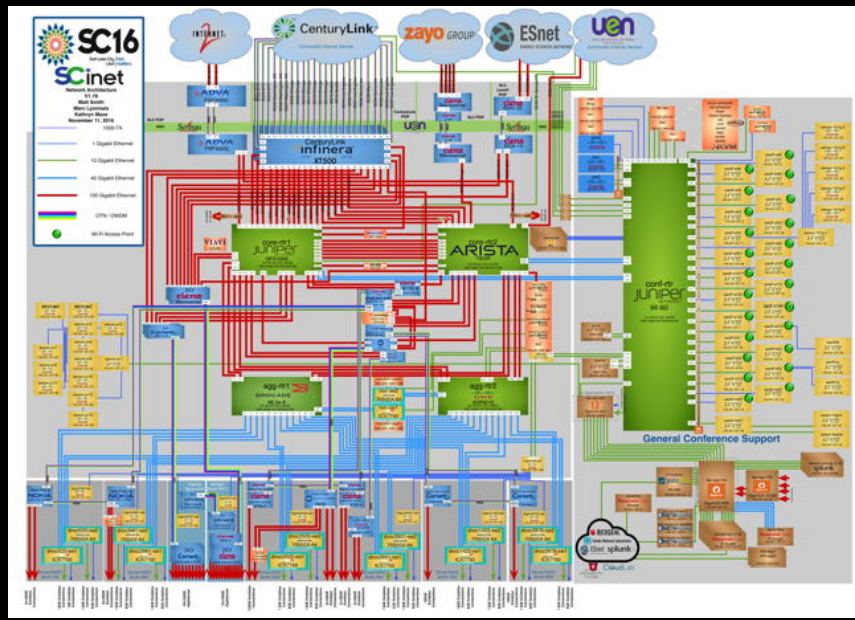

Some developments: GENI

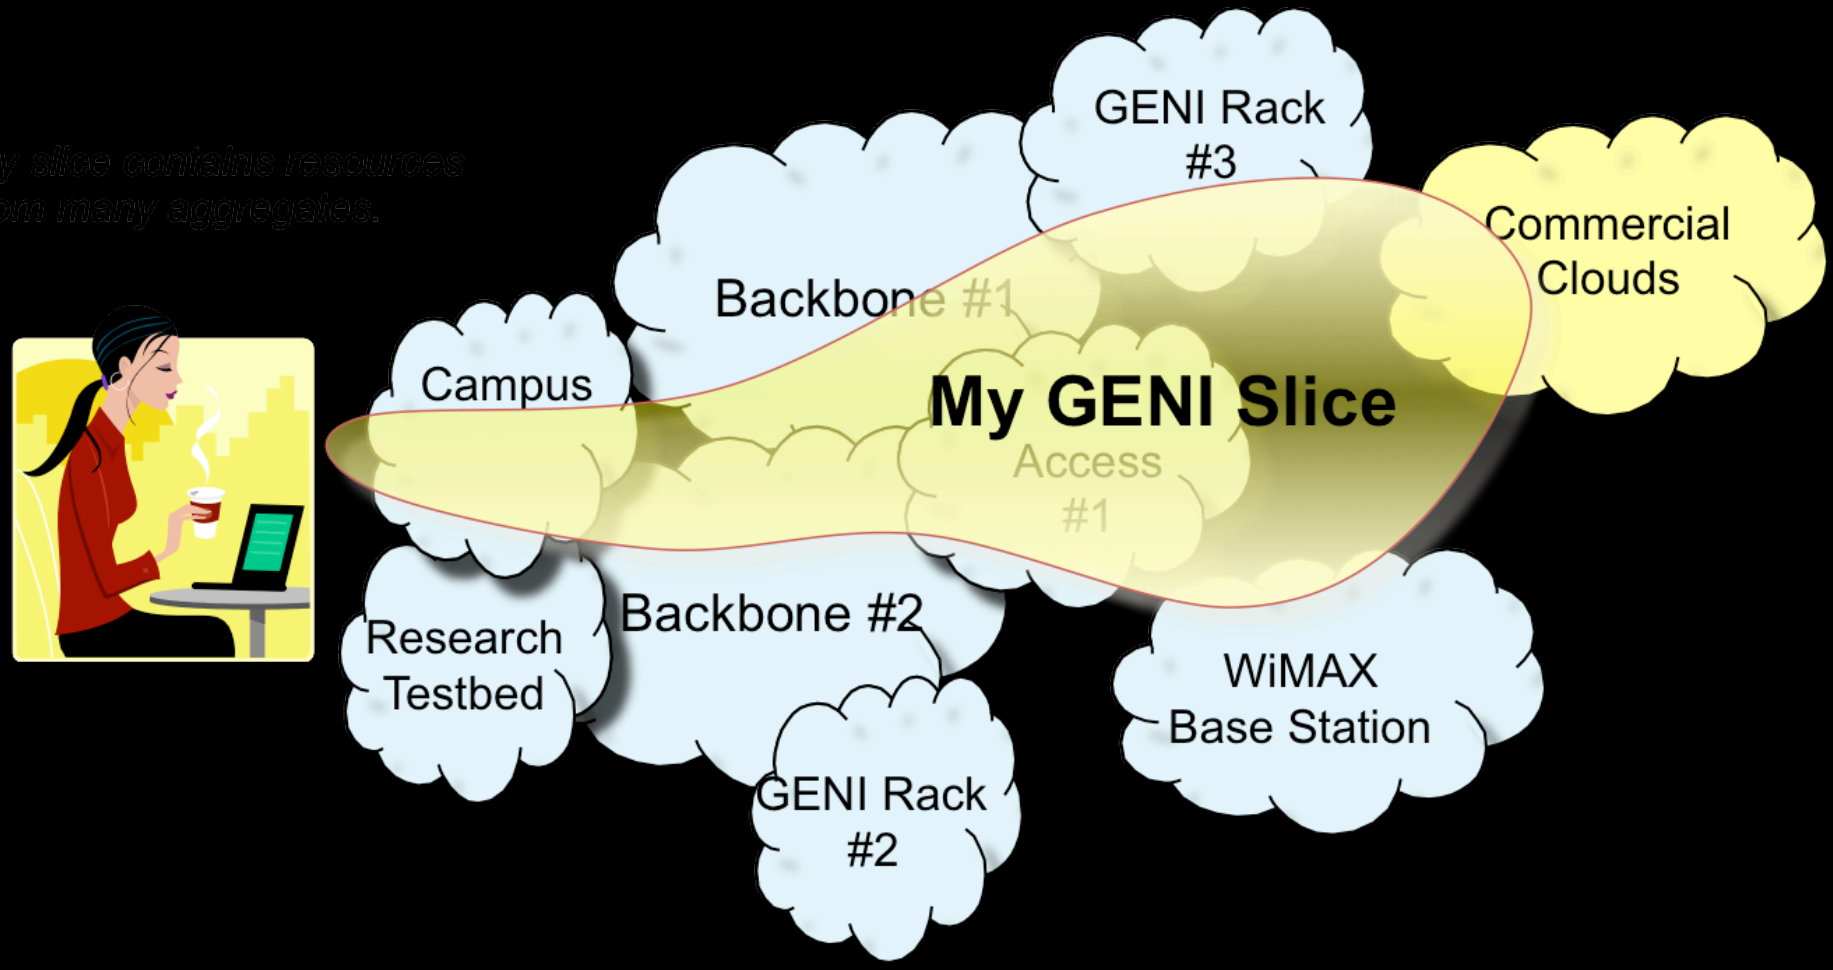

SKA: Depending on analysis load & physics mode they want to investigate to use SDN in real time to direct bursts of data to different compute resources and do load balancing.

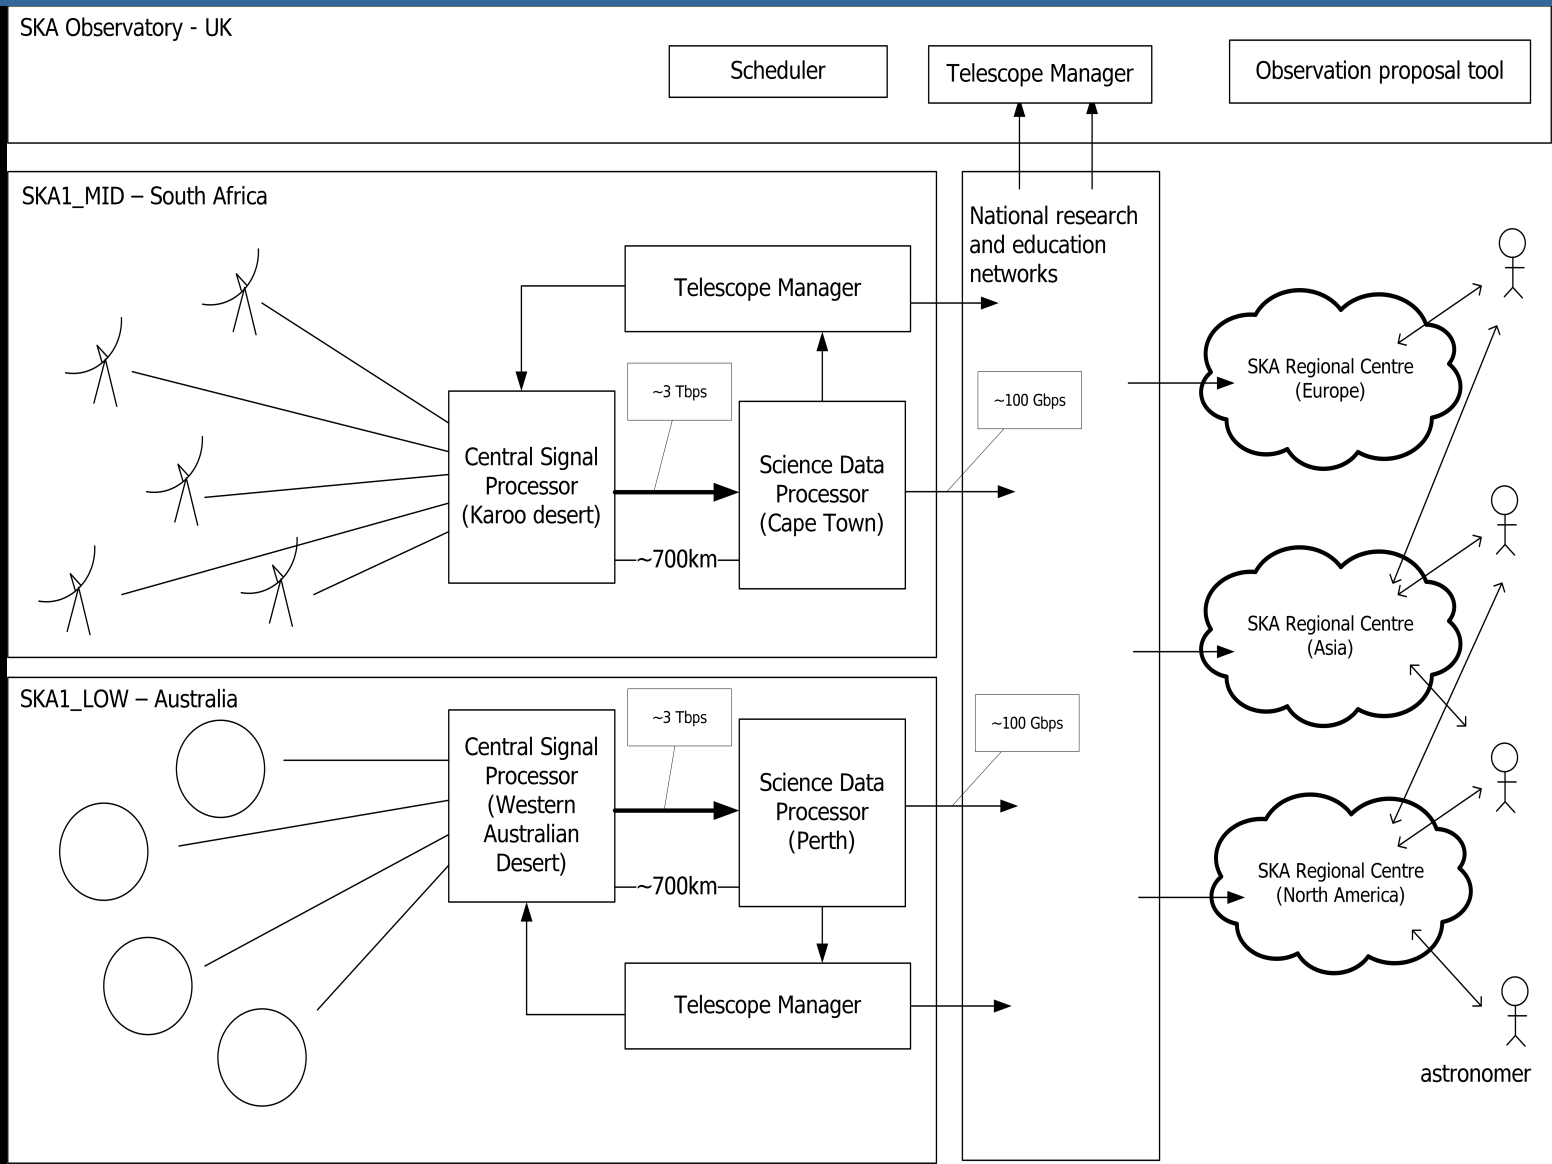

NFV & Security upstream

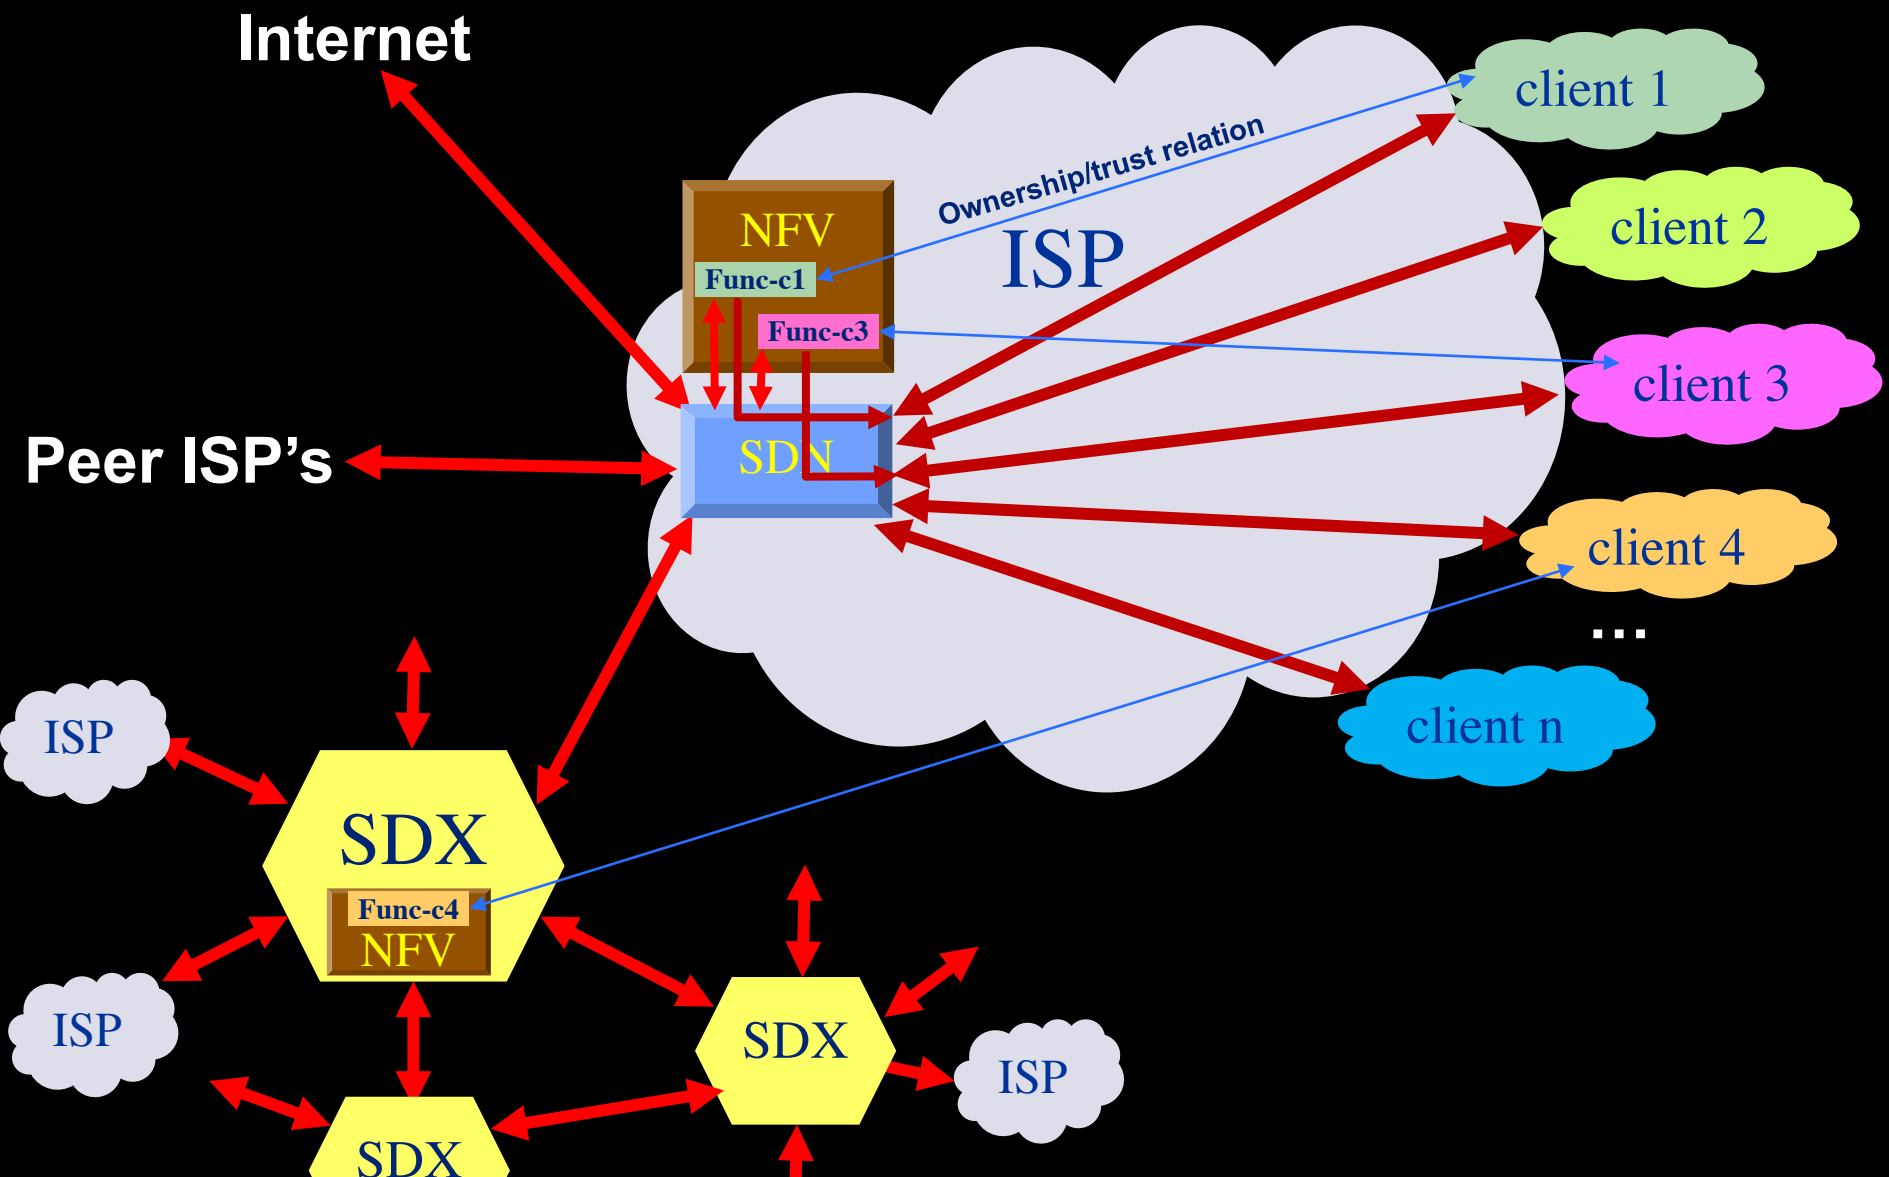

Networks of ScienceDMZ's & SDX's

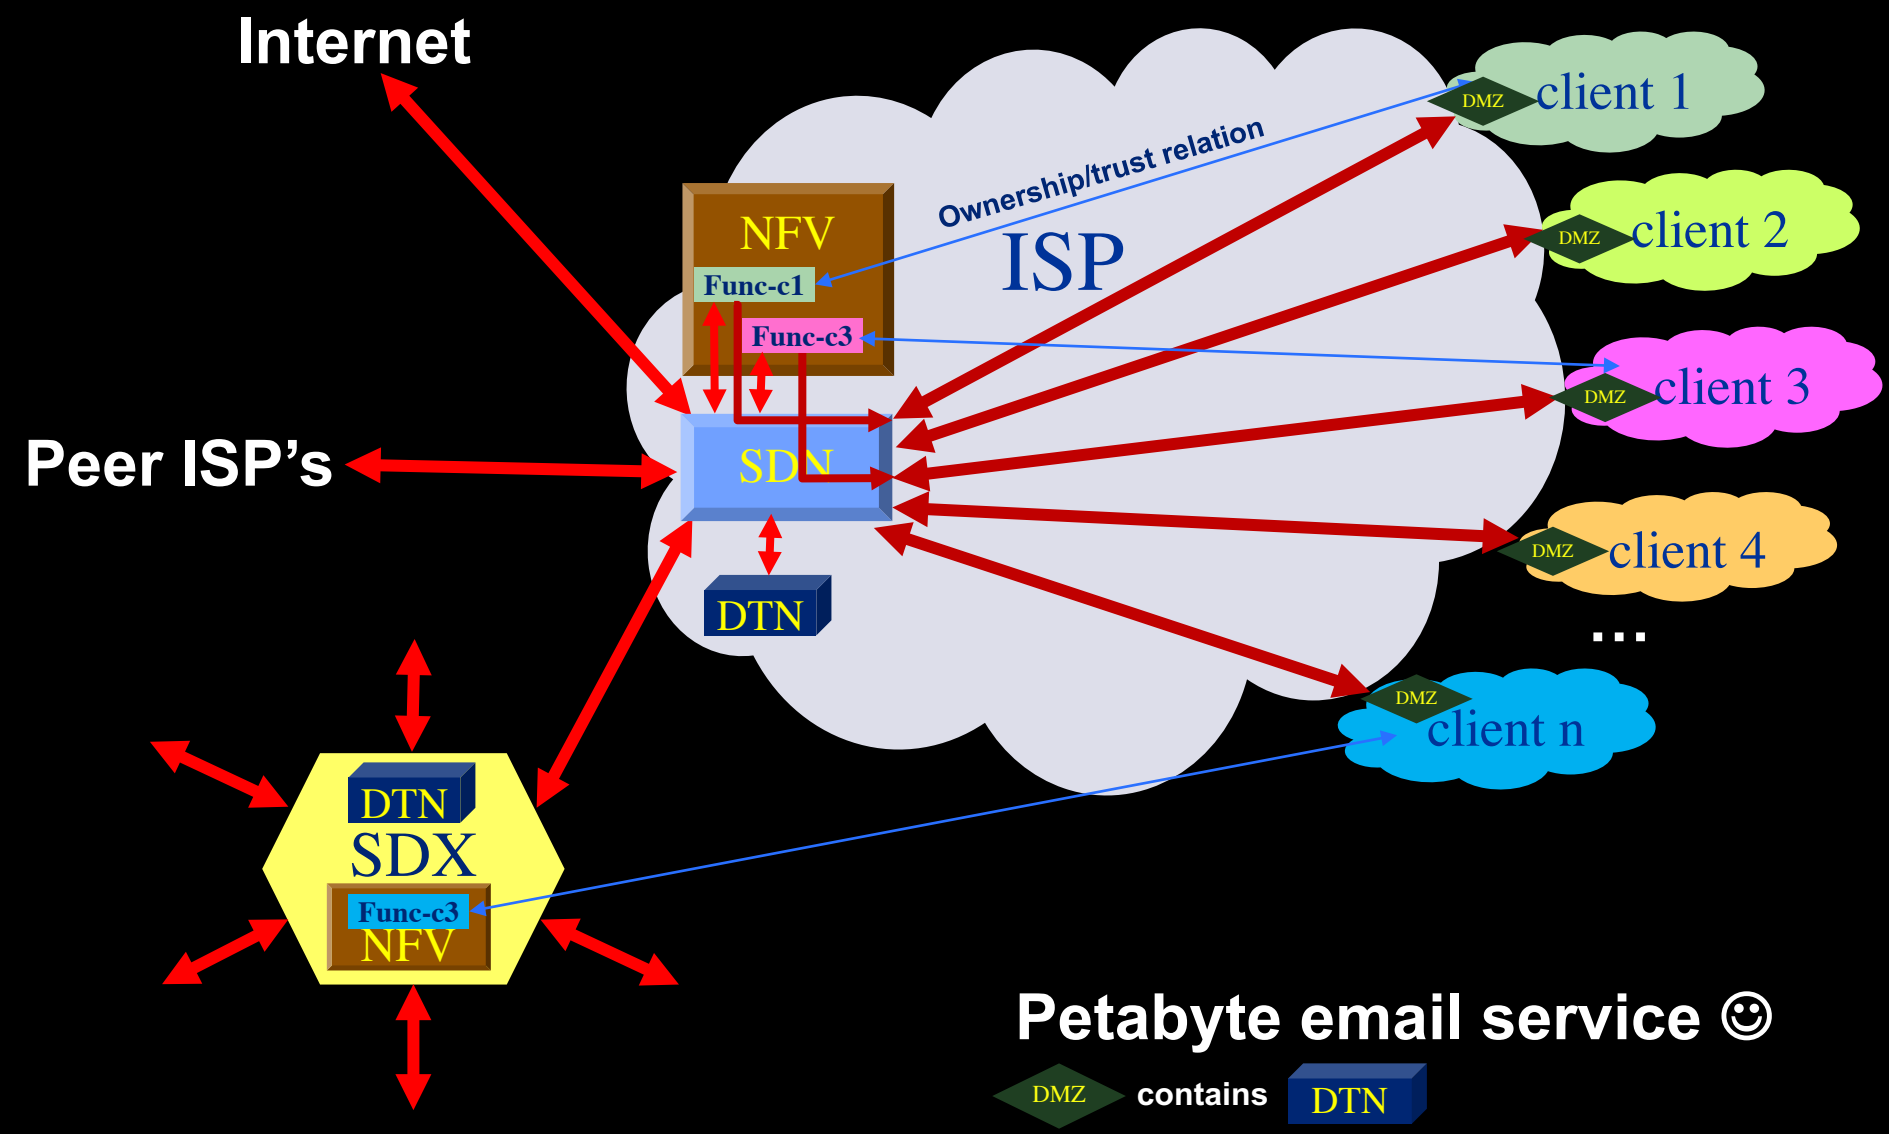

Cyber security program

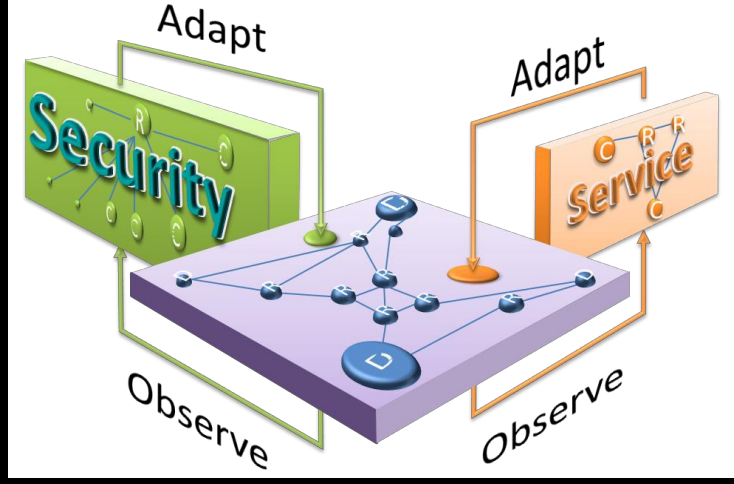

Research goal is to obtain the knowledge to create ICT systems that:

- 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

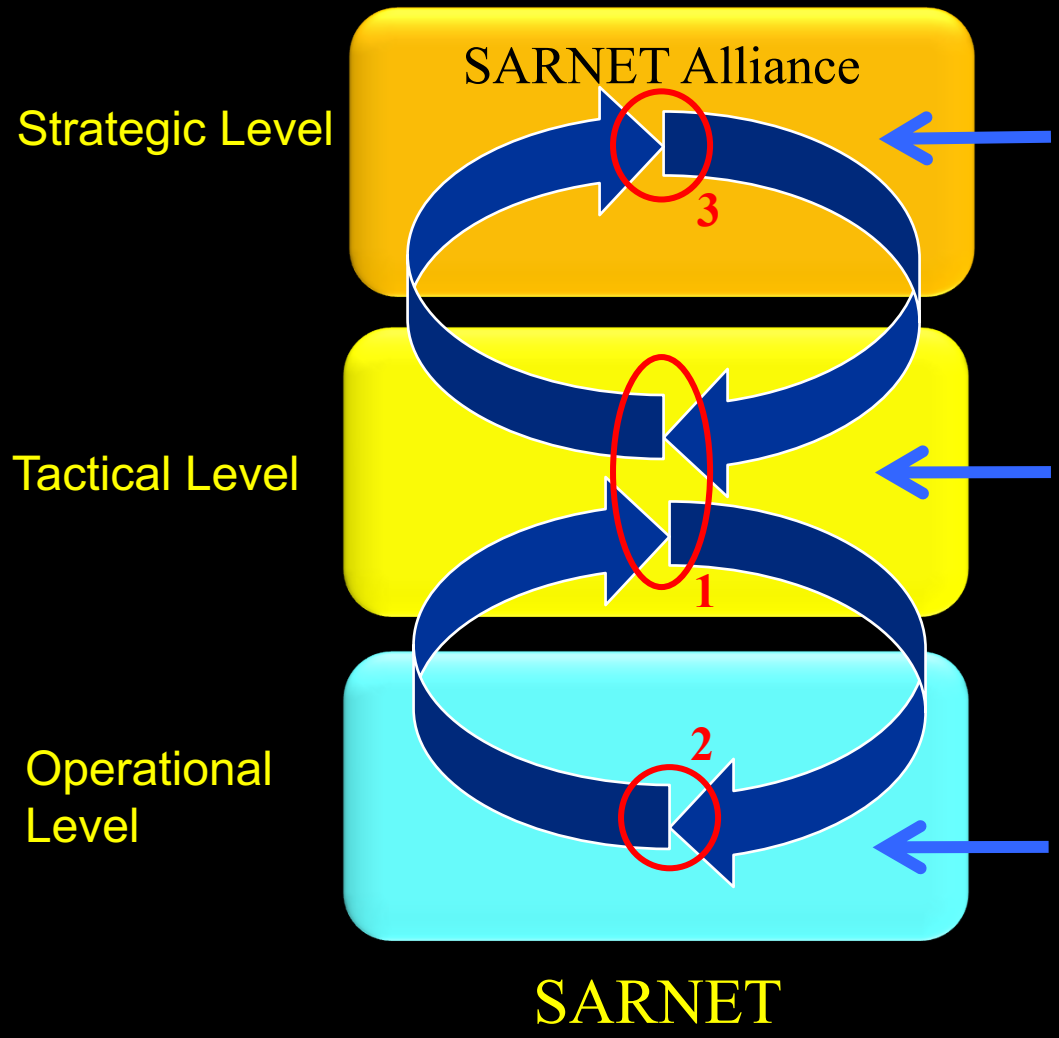

Ameneh Deljoo (PhD):
Why create SARNET Alliances?
Model (3) autonomous SARNET behaviors to identify risk and benefits for SARNET stakeholders

Gleb Polevoy (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).

Ambition to put capabilities into fieldlab

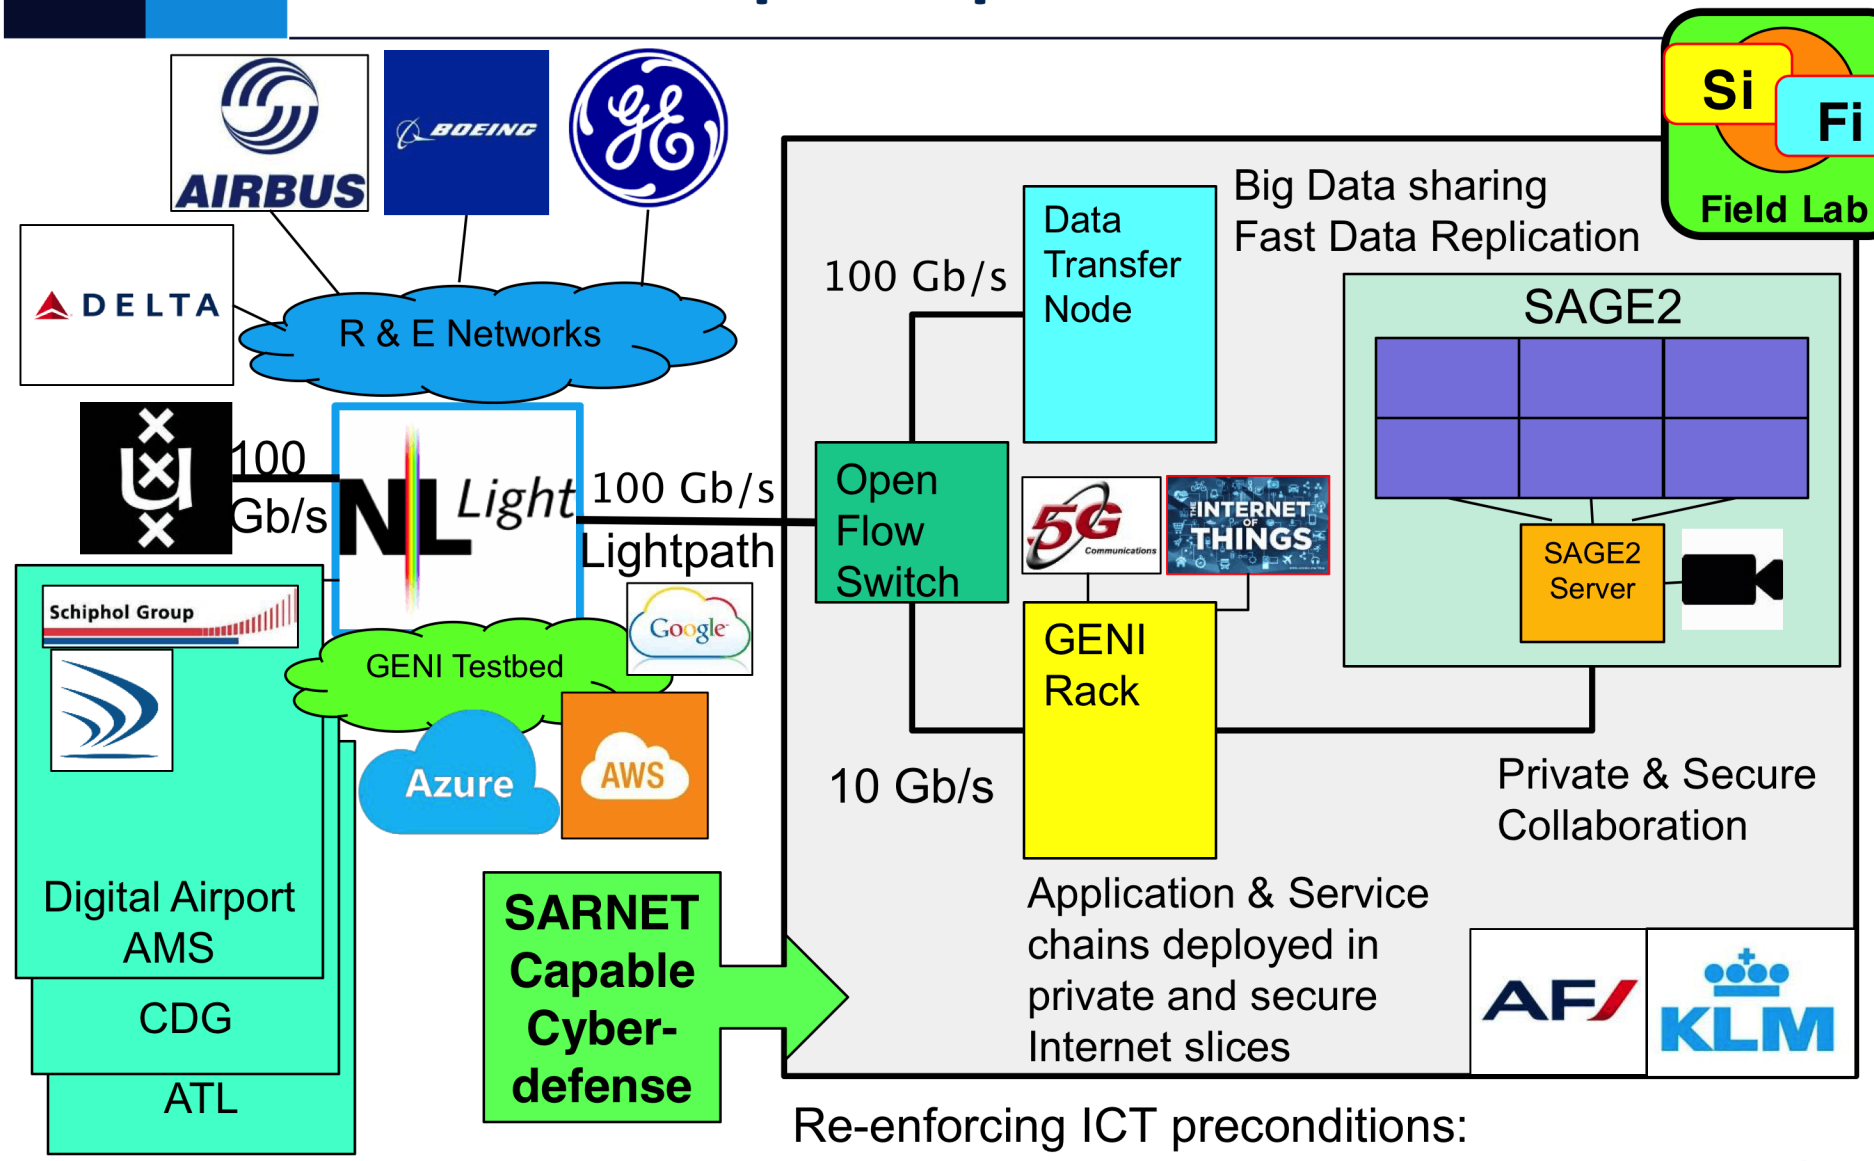

Re-enforcing ICT preconditions:
Each envisaged site has similar elements

AIRFRANCE KLM

Basic operating system loop

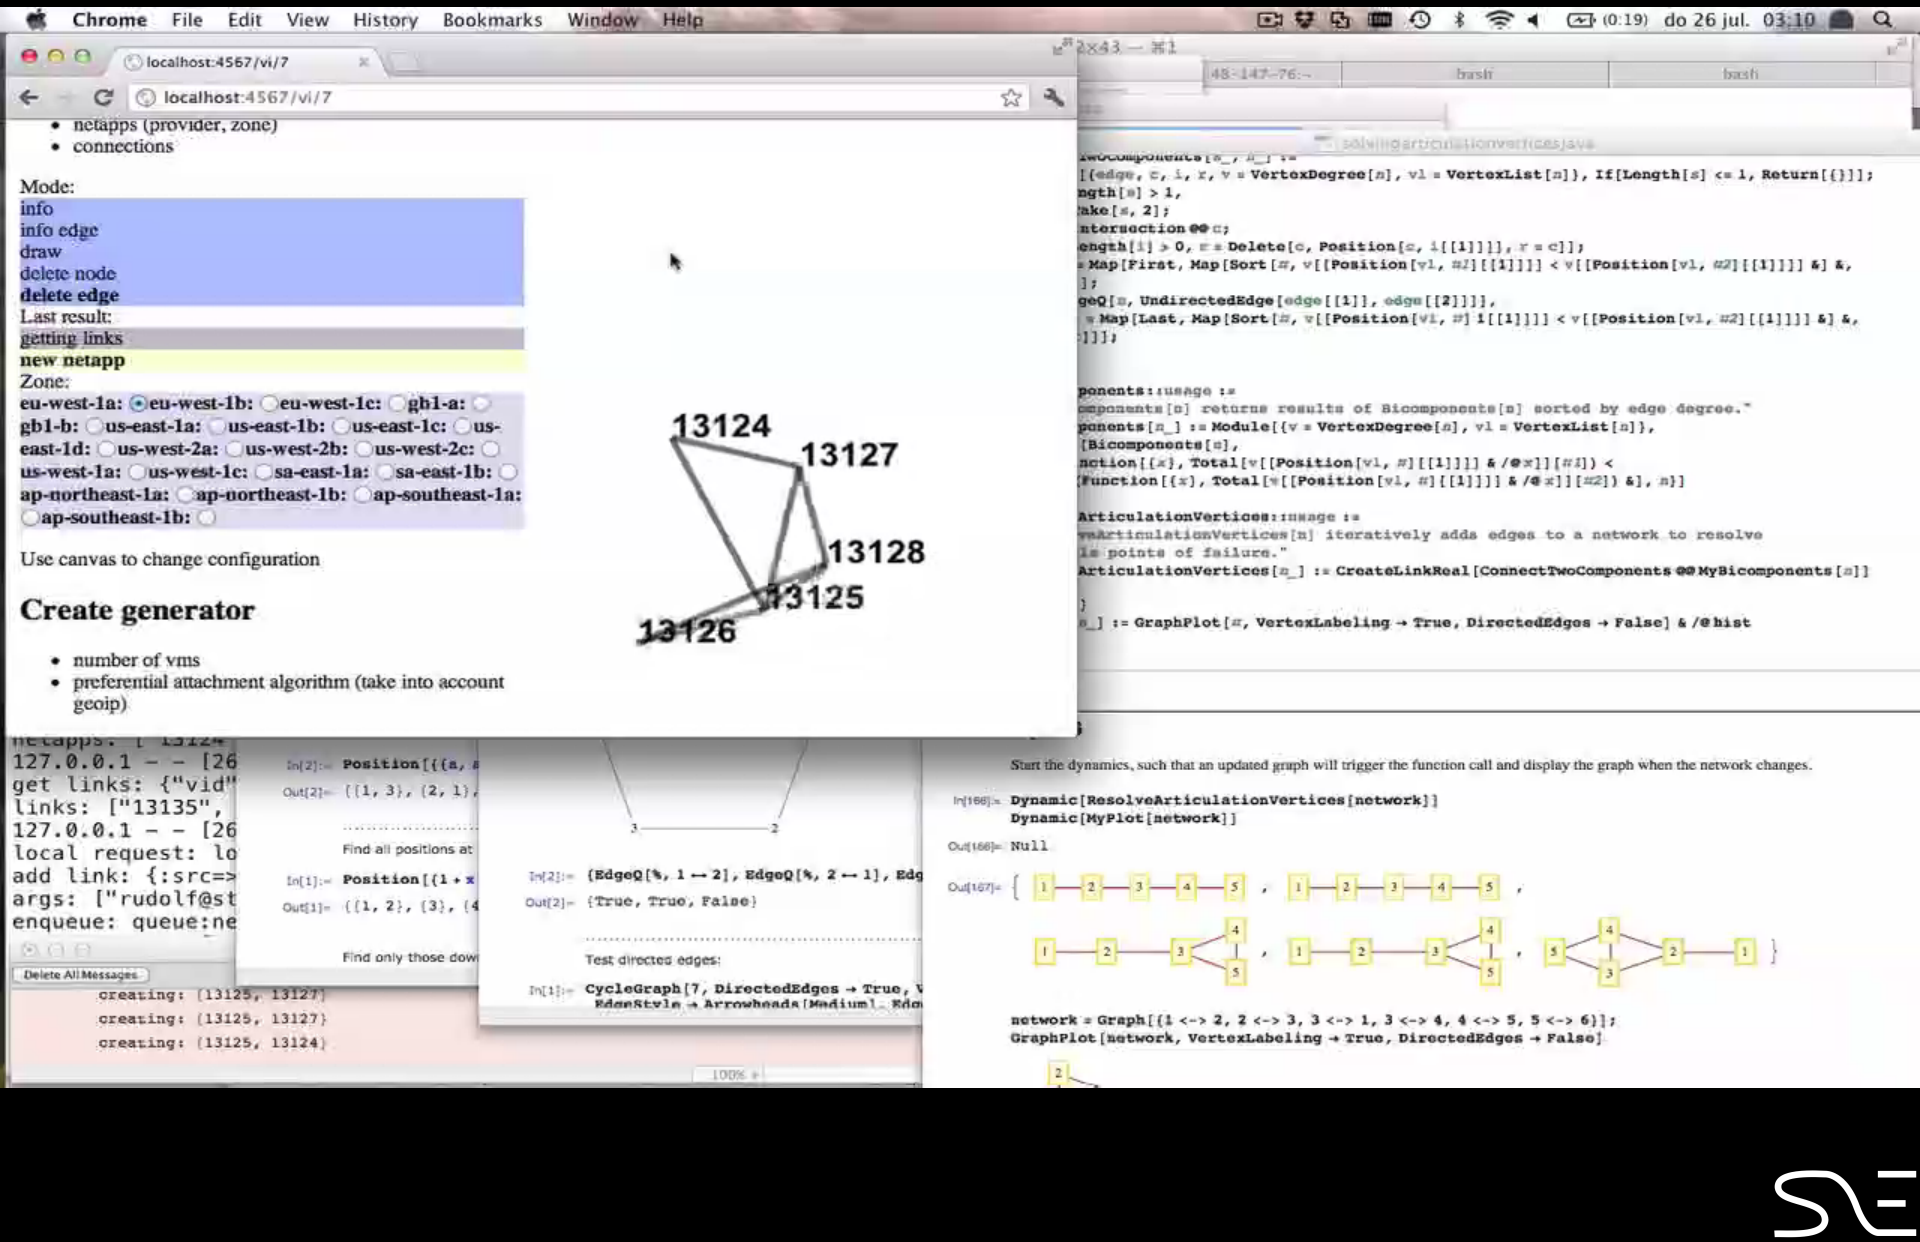

Status SARNET Operational Level

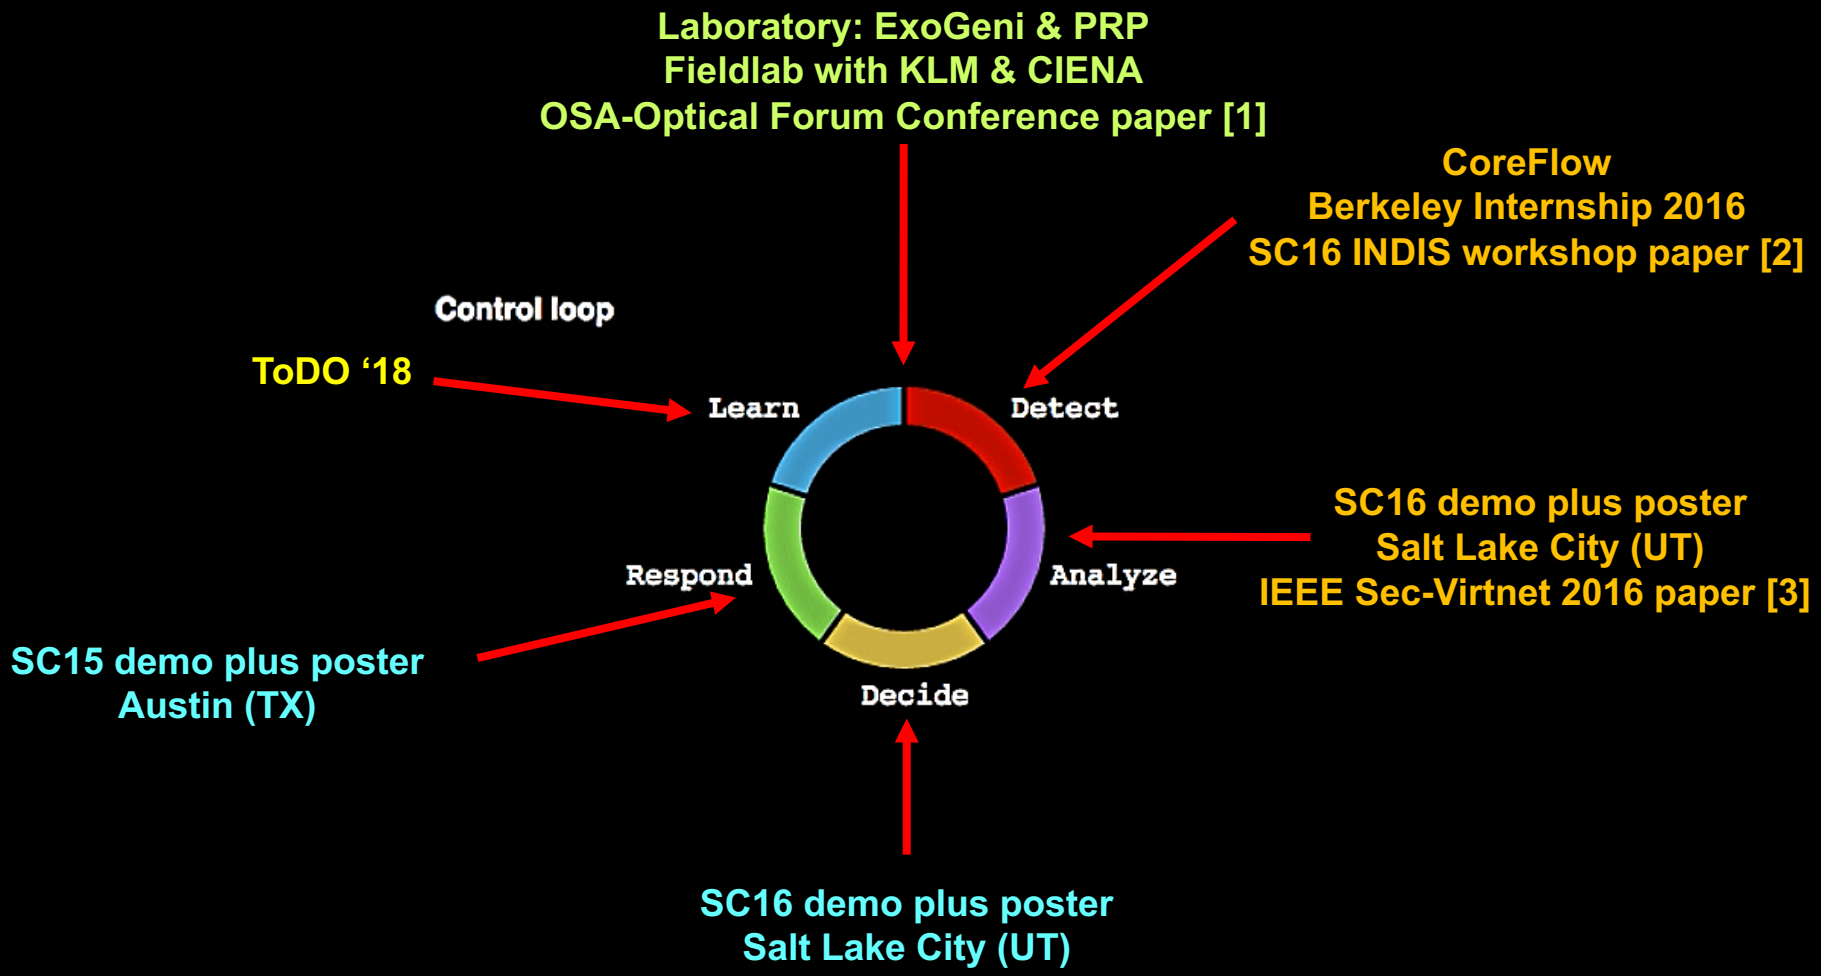

1. Paper: R. Koning, A. Deljoo, S. Trajanovski, B. de Graaff, P. Grosso, L. Gommans, T. van Engers, F. Fransen, R. Meijer, R. Wilson, and C. de Laat, "Enabling E-Science Applications with Dynamic Optical Networks: Secure Autonomous Response Networks ", OSA Optical Fiber Communication Conference and Exposition, 19-23 March 2017, Los Angeles, California.
2. Paper: Ralph Koning, Nick Buraglio, Cees de Laat, Paola Grosso, "CoreFlow: Enriching Bro security events using network traffic monitoring data", SC16 Salt Lake City, INDIS workshop, Nov 13, 2016.
3. Paper: Ralph Koning, Ben de Graaff, Cees de Laat, Robert Meijer, Paola Grosso, "Analysis of Software Defined Networking defences against Distributed Denial of Service attacks", The IEEE International Workshop on Security in Virtualized Networks (Sec-VirtNet 2016) at the 2nd IEEE International Conference on Network Softwarization (NetSoft 2016), Seoul Korea, June 10, 2016.

SC16 DEMO SARNET Operational Level

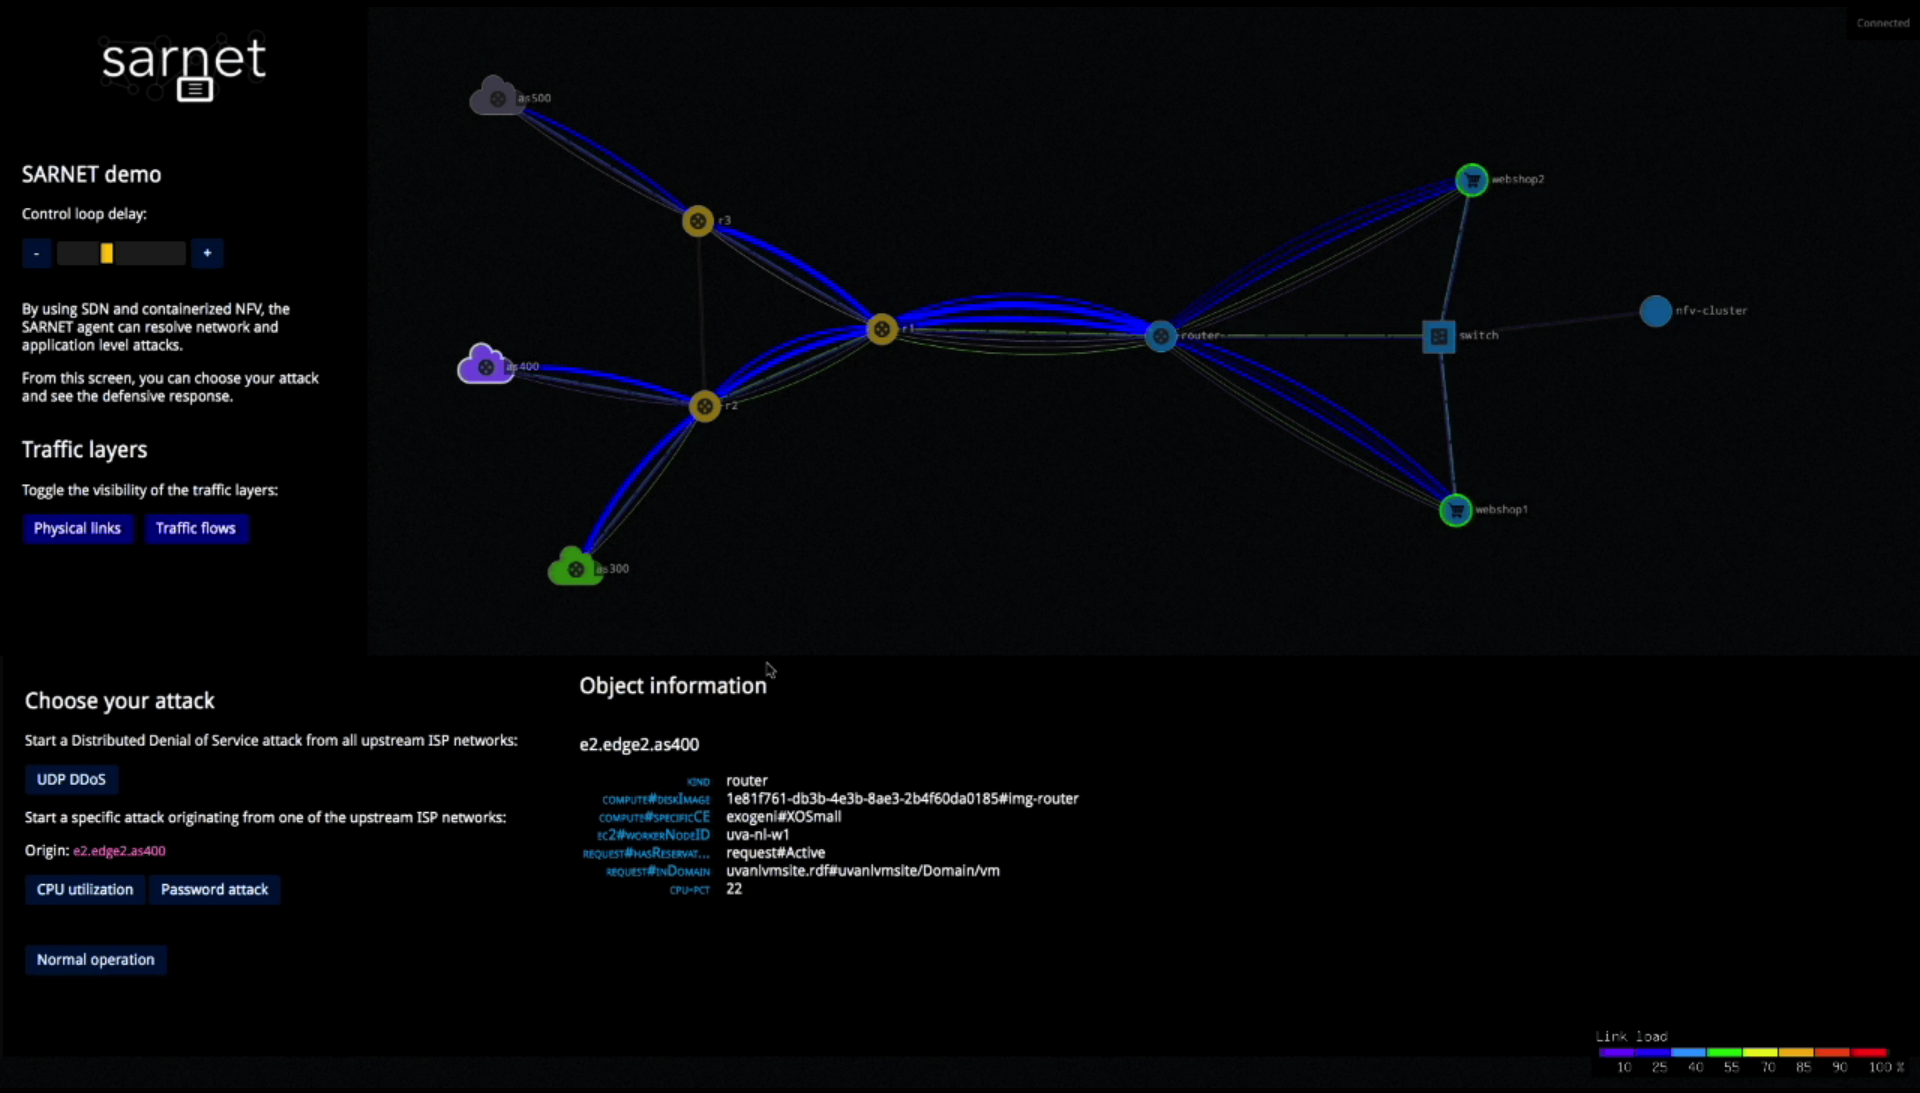

SC16 DEMO SARNET Operational Level

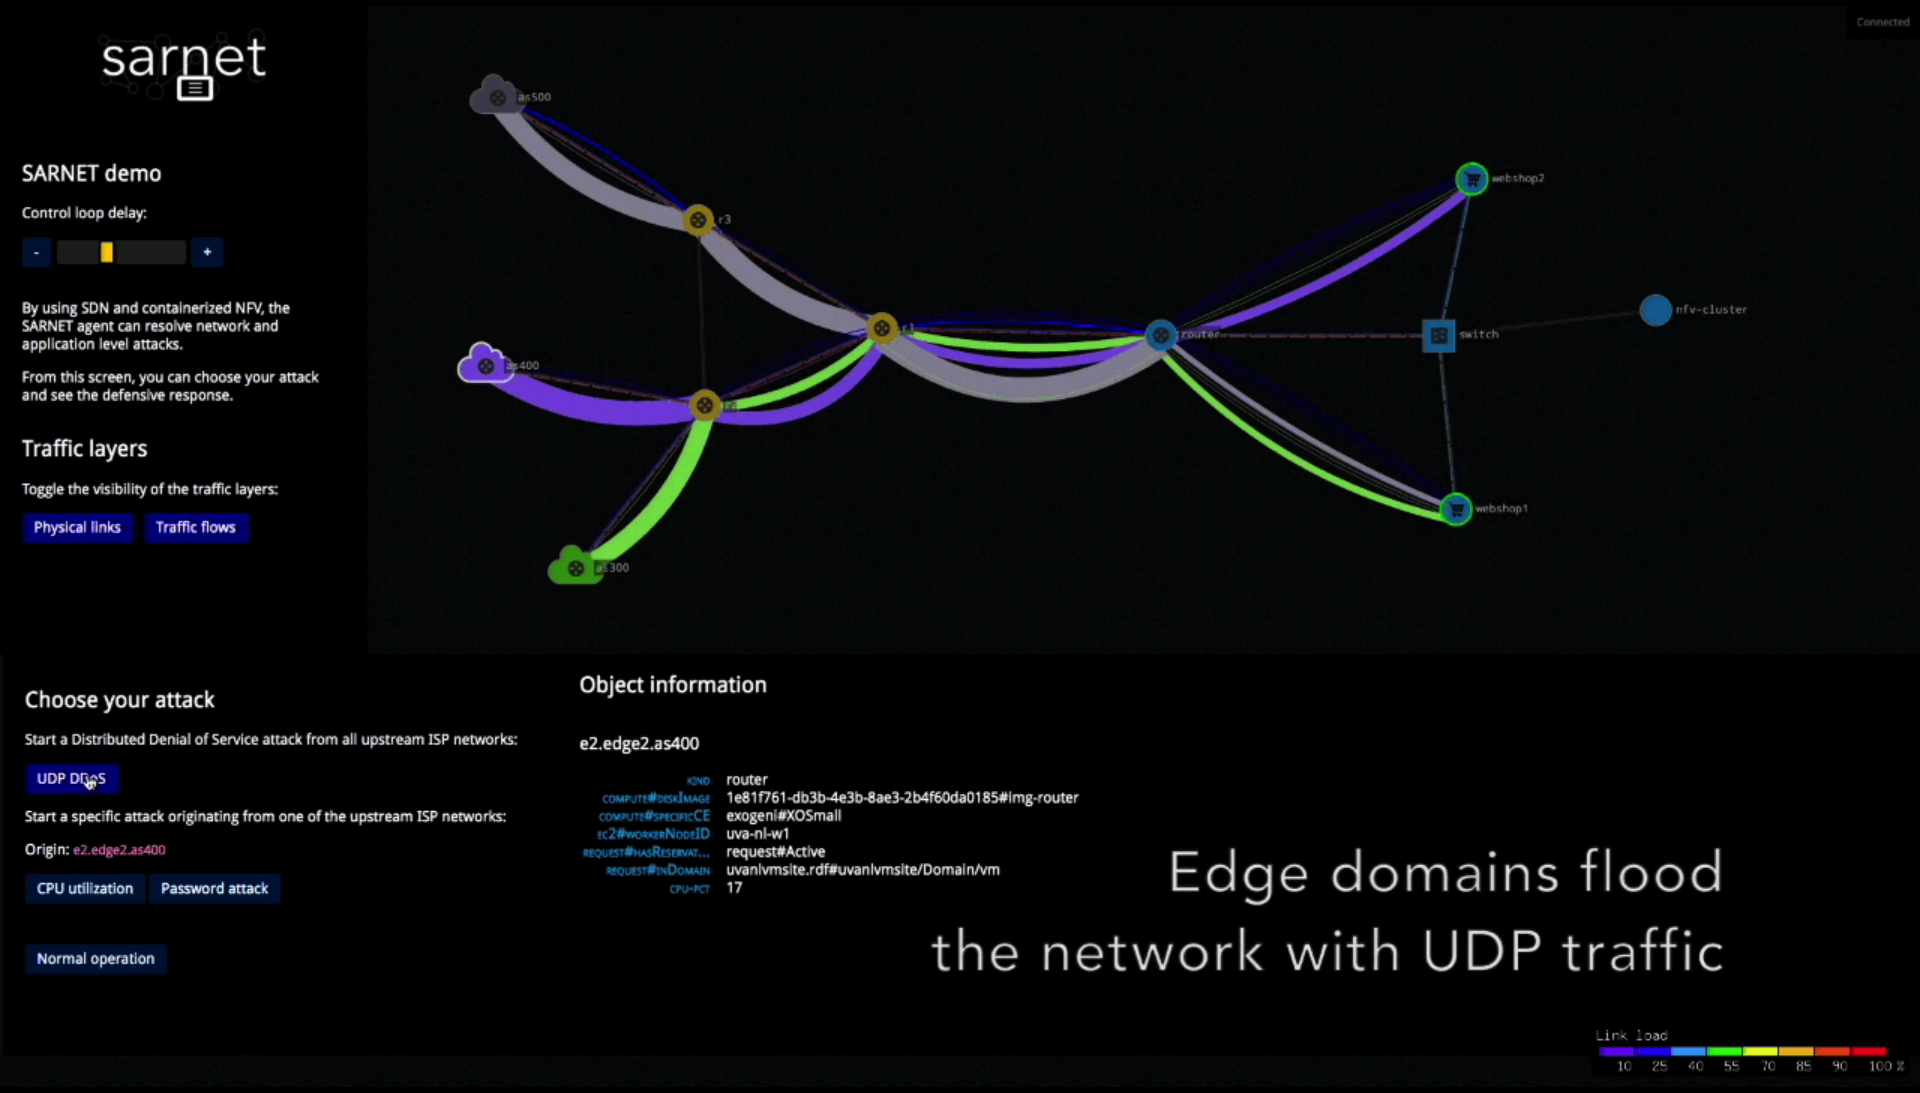

SC16 DEMO SARNET Operational Level

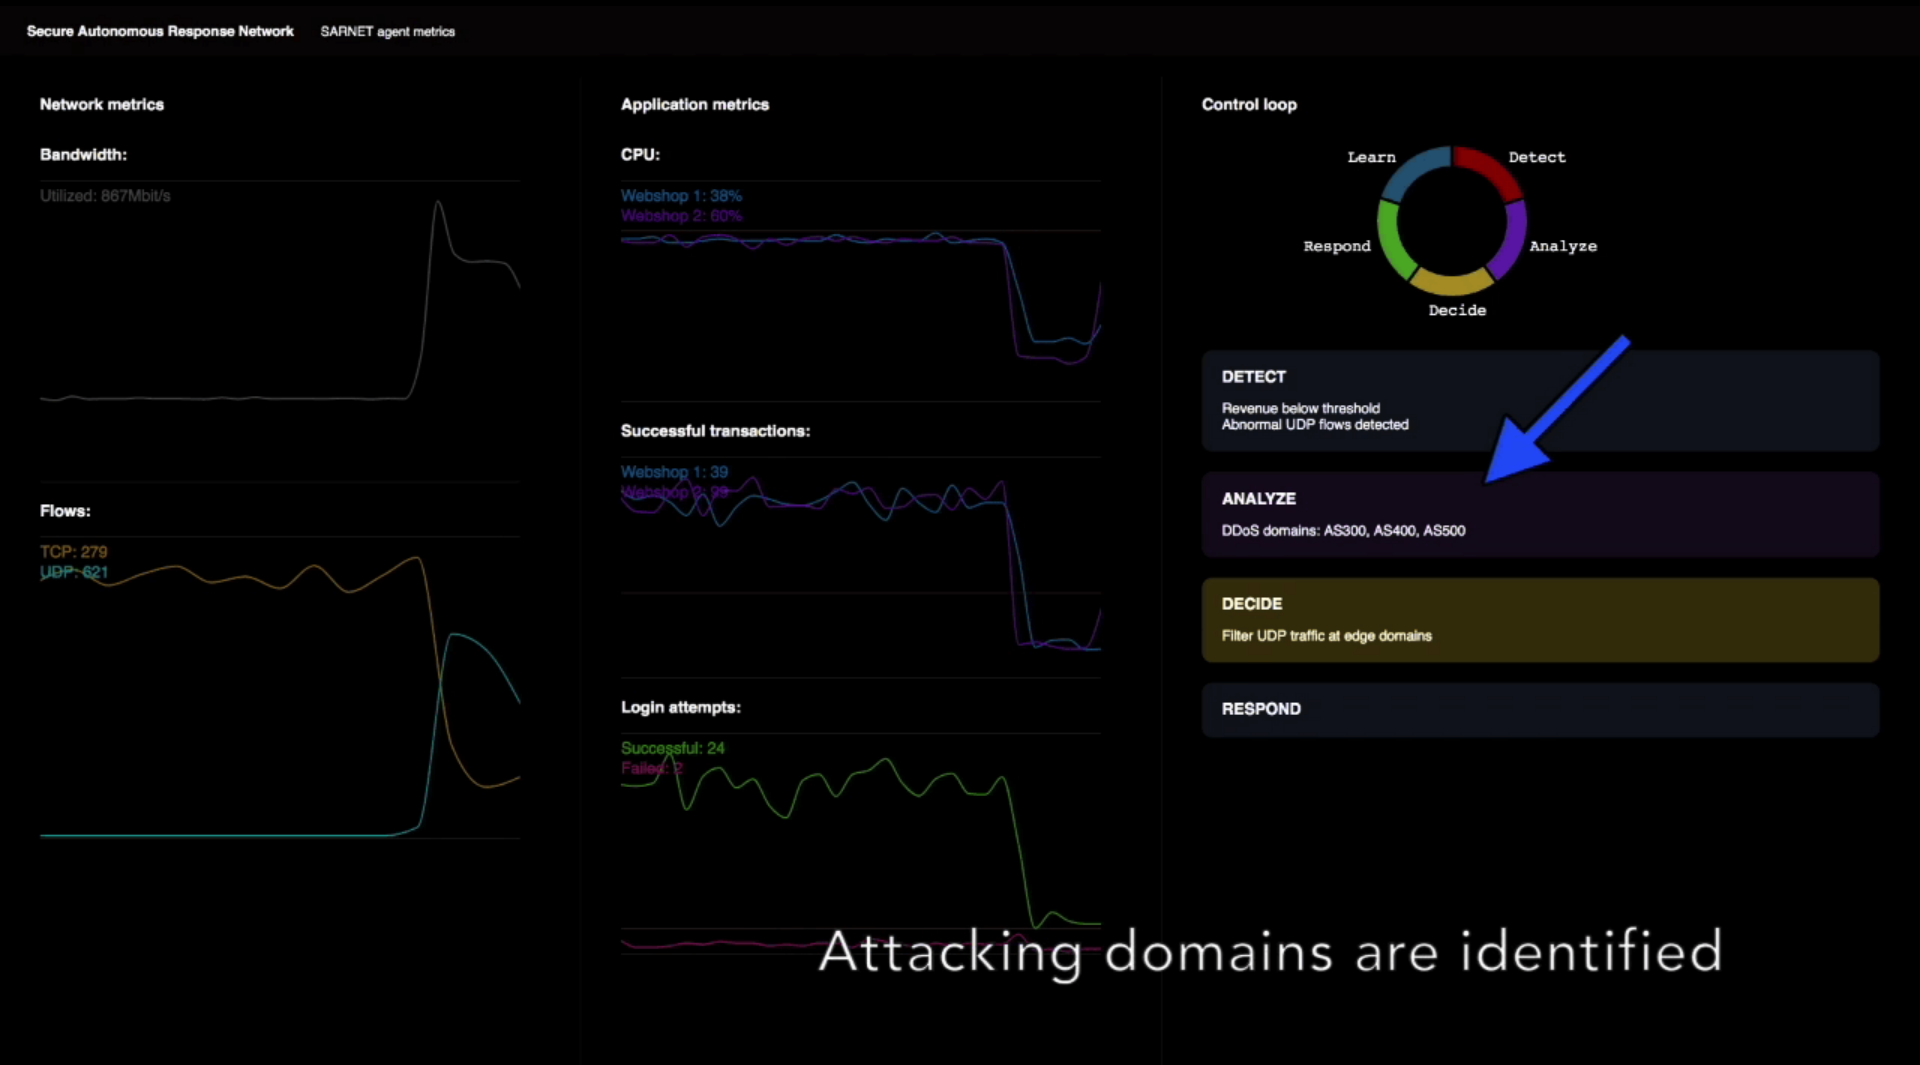

Attacking domains are identified

SC16 DEMO SARNET Operational Level

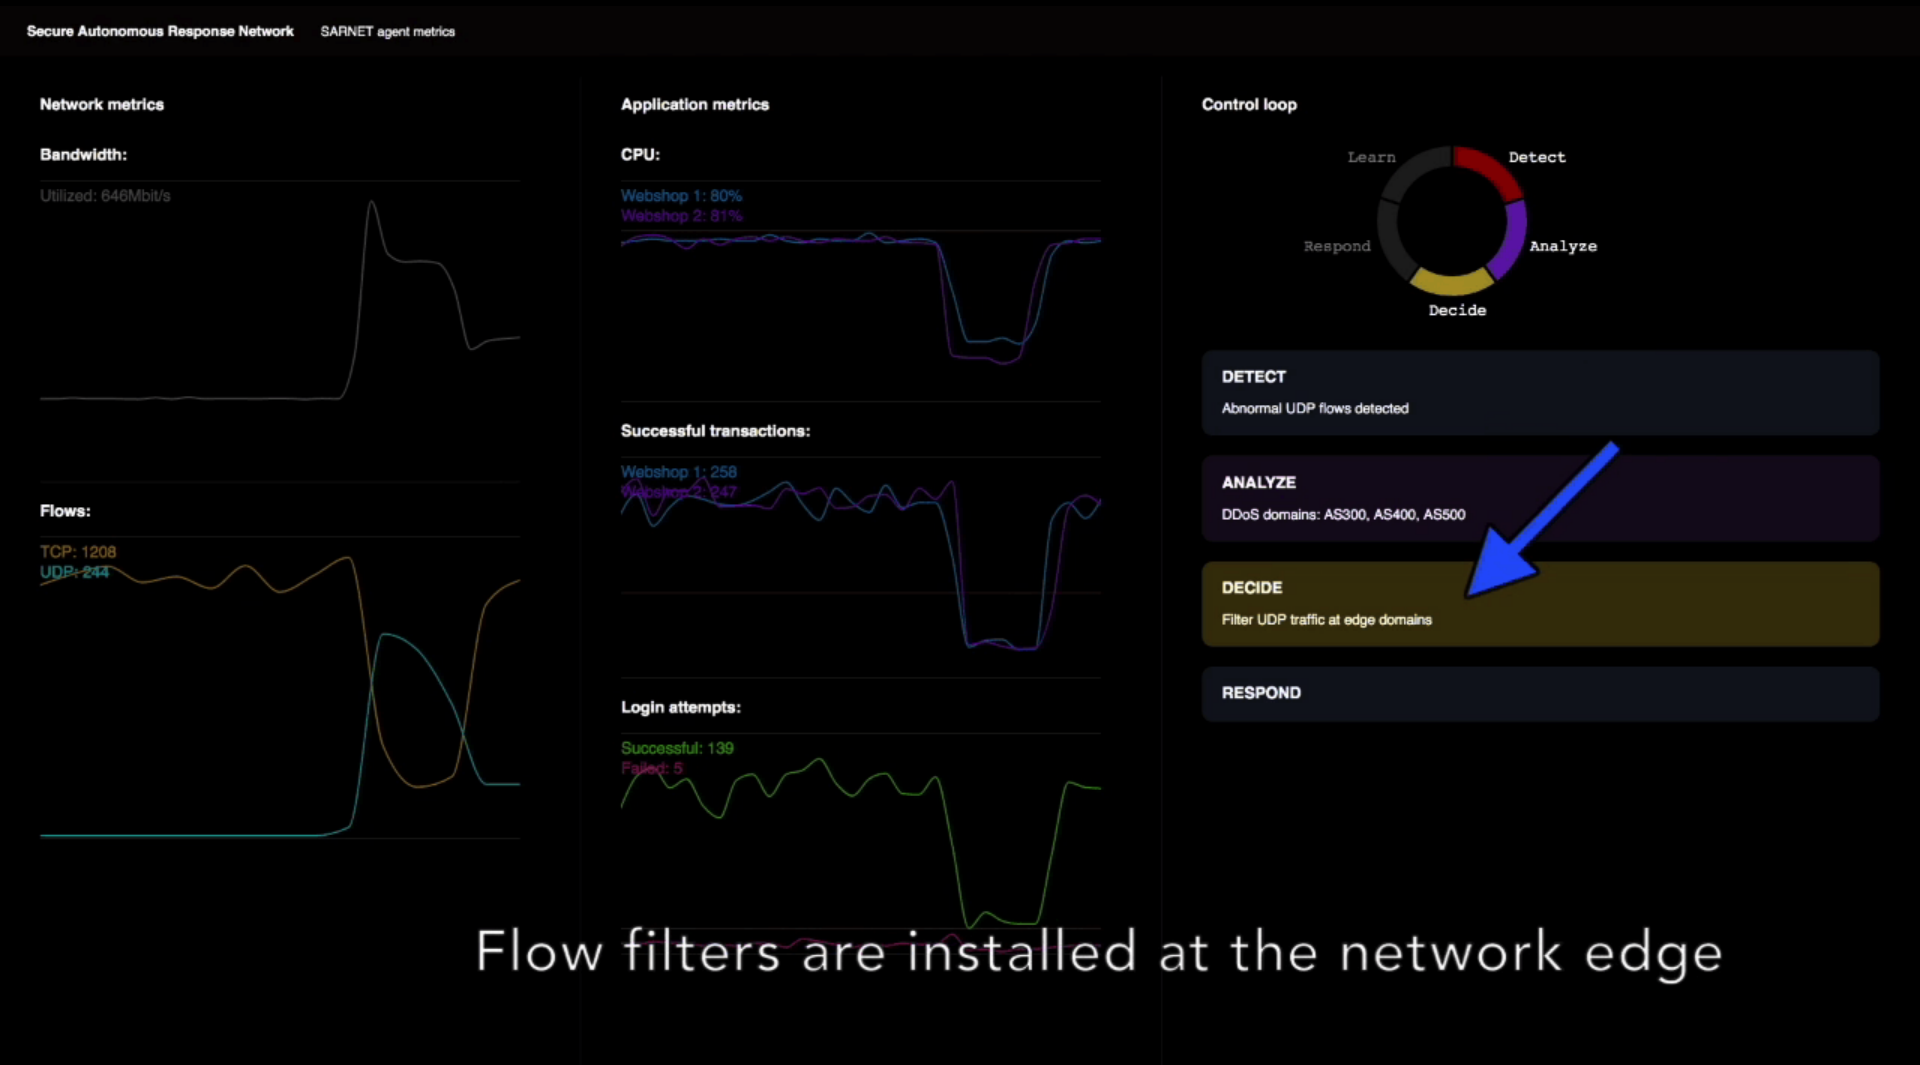

Flow filters are installed at the network edge

SC16 DEMO SARNET Operational Level

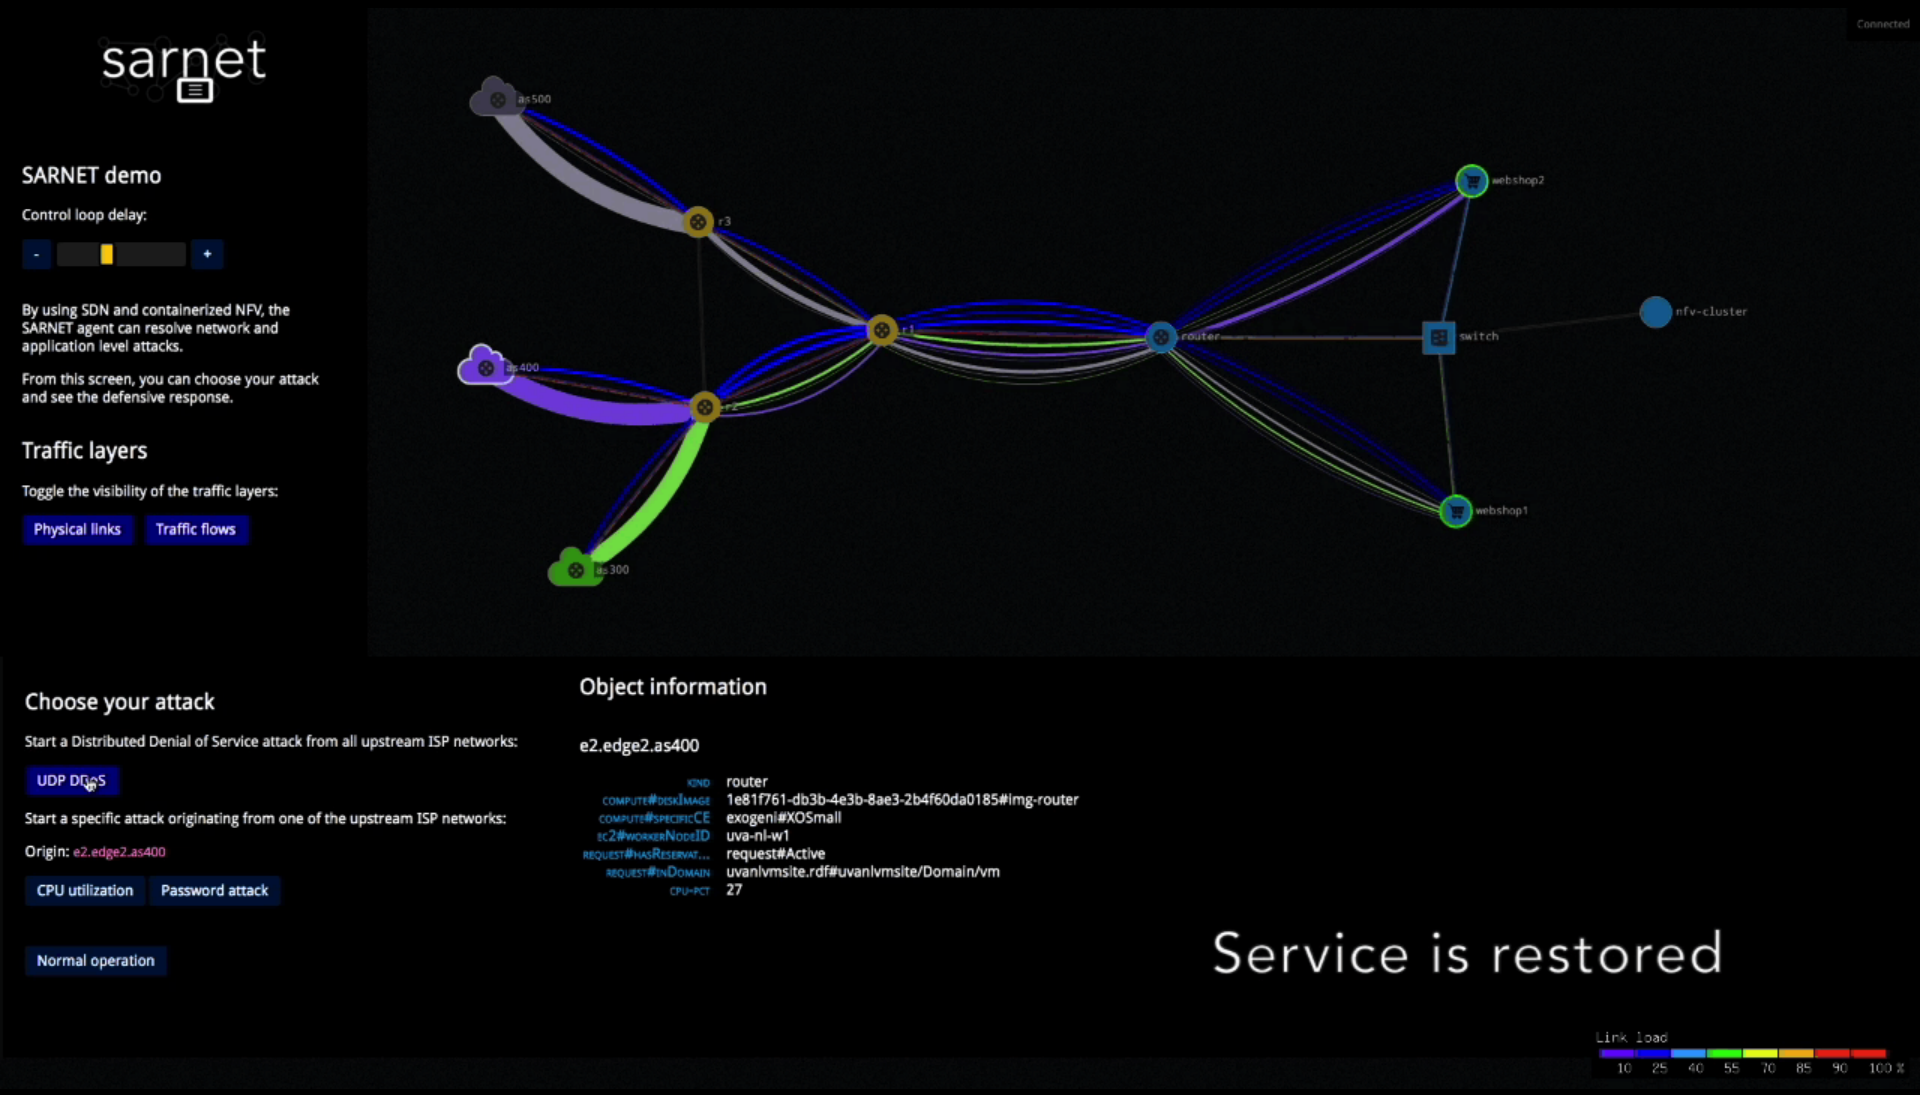

Control loop

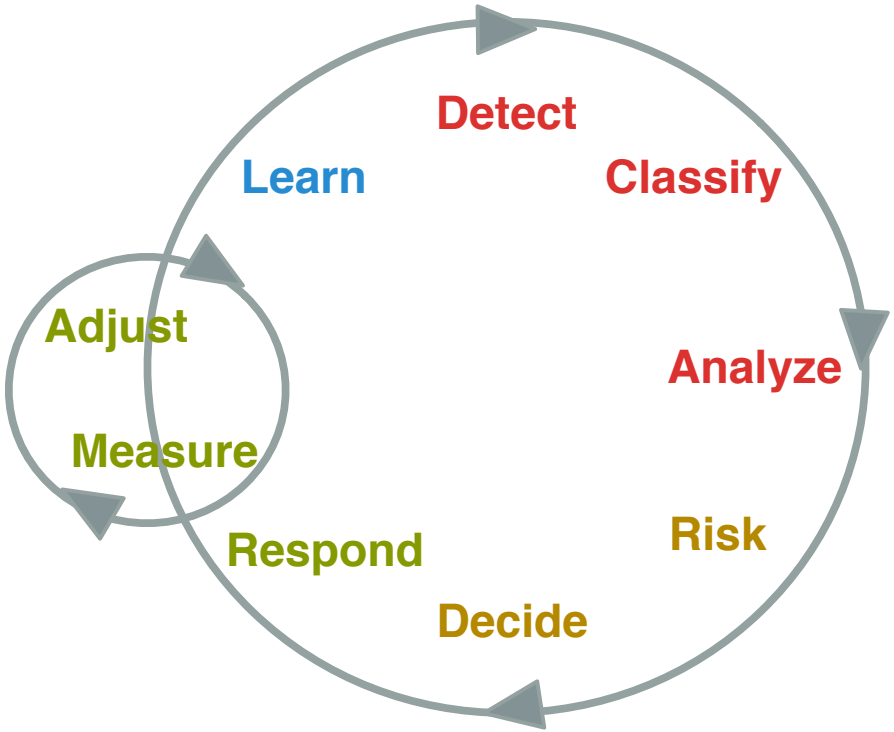

Detection phase: Detect, Classify, Analyze

Decision phase: Risk, Decide

Response phase: Respond, Adjust, Measure

Learn phase: Learn (with input form other phases)

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

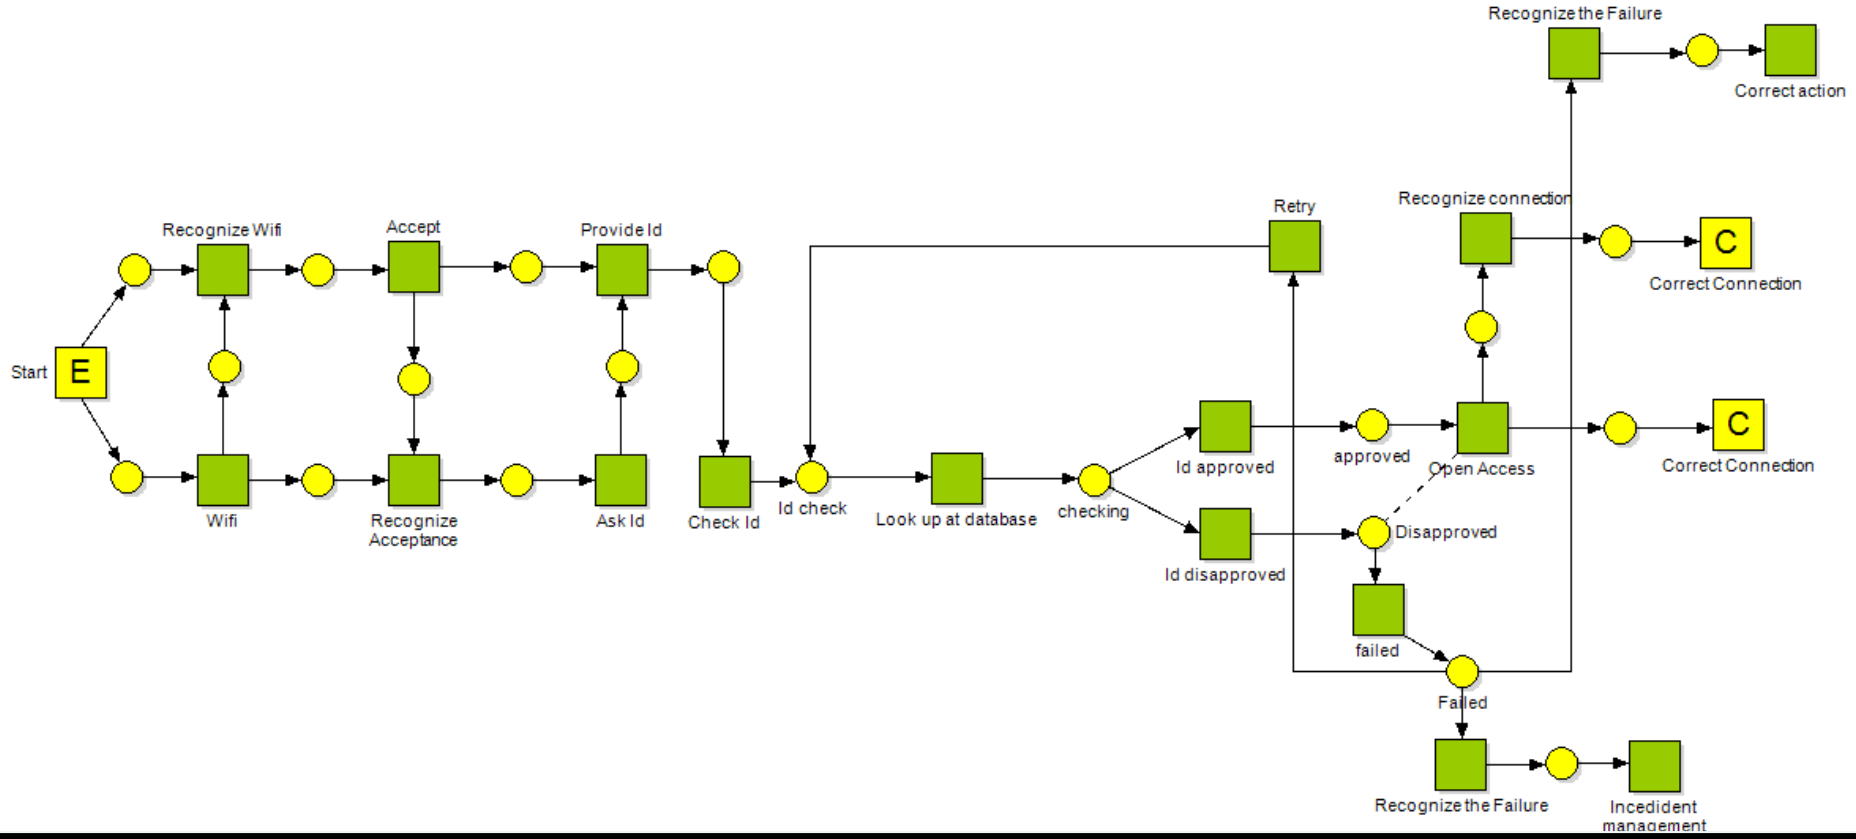

Petri net of EduRoam Case
(first step)

Describing Intentions, Motivations and Actions

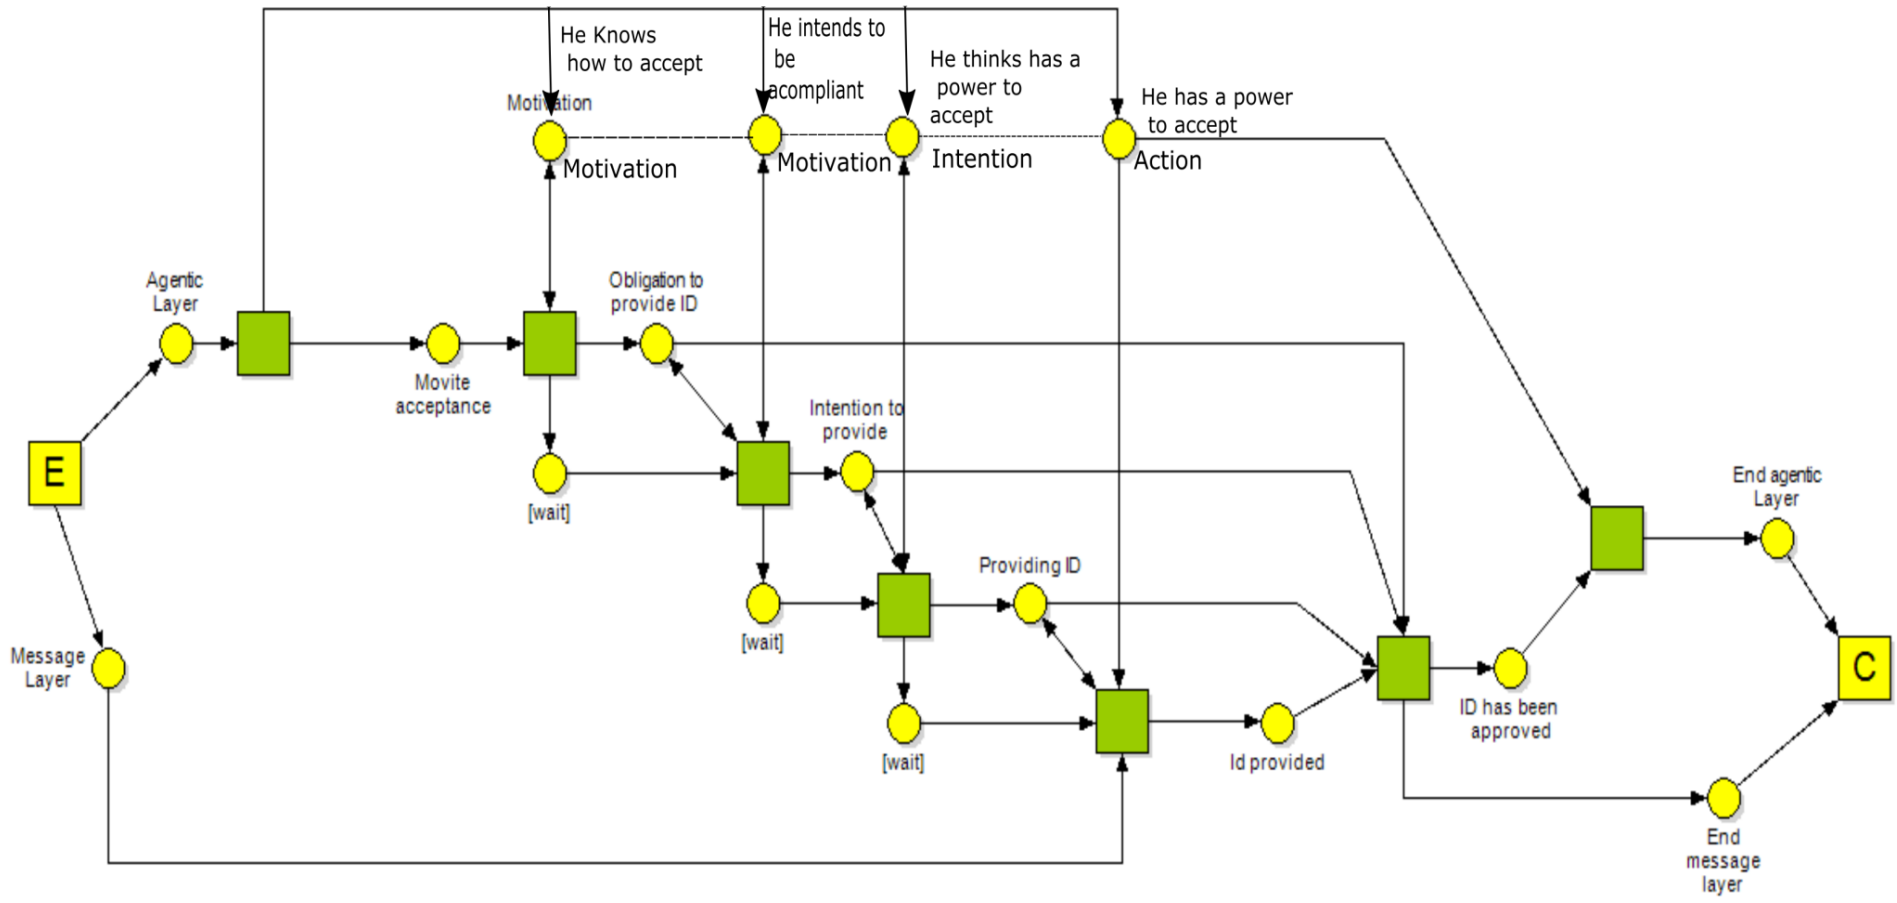

Petri net of EduRoam Case

The VM Turntable Demonstrator

iGrid2005
SC2005

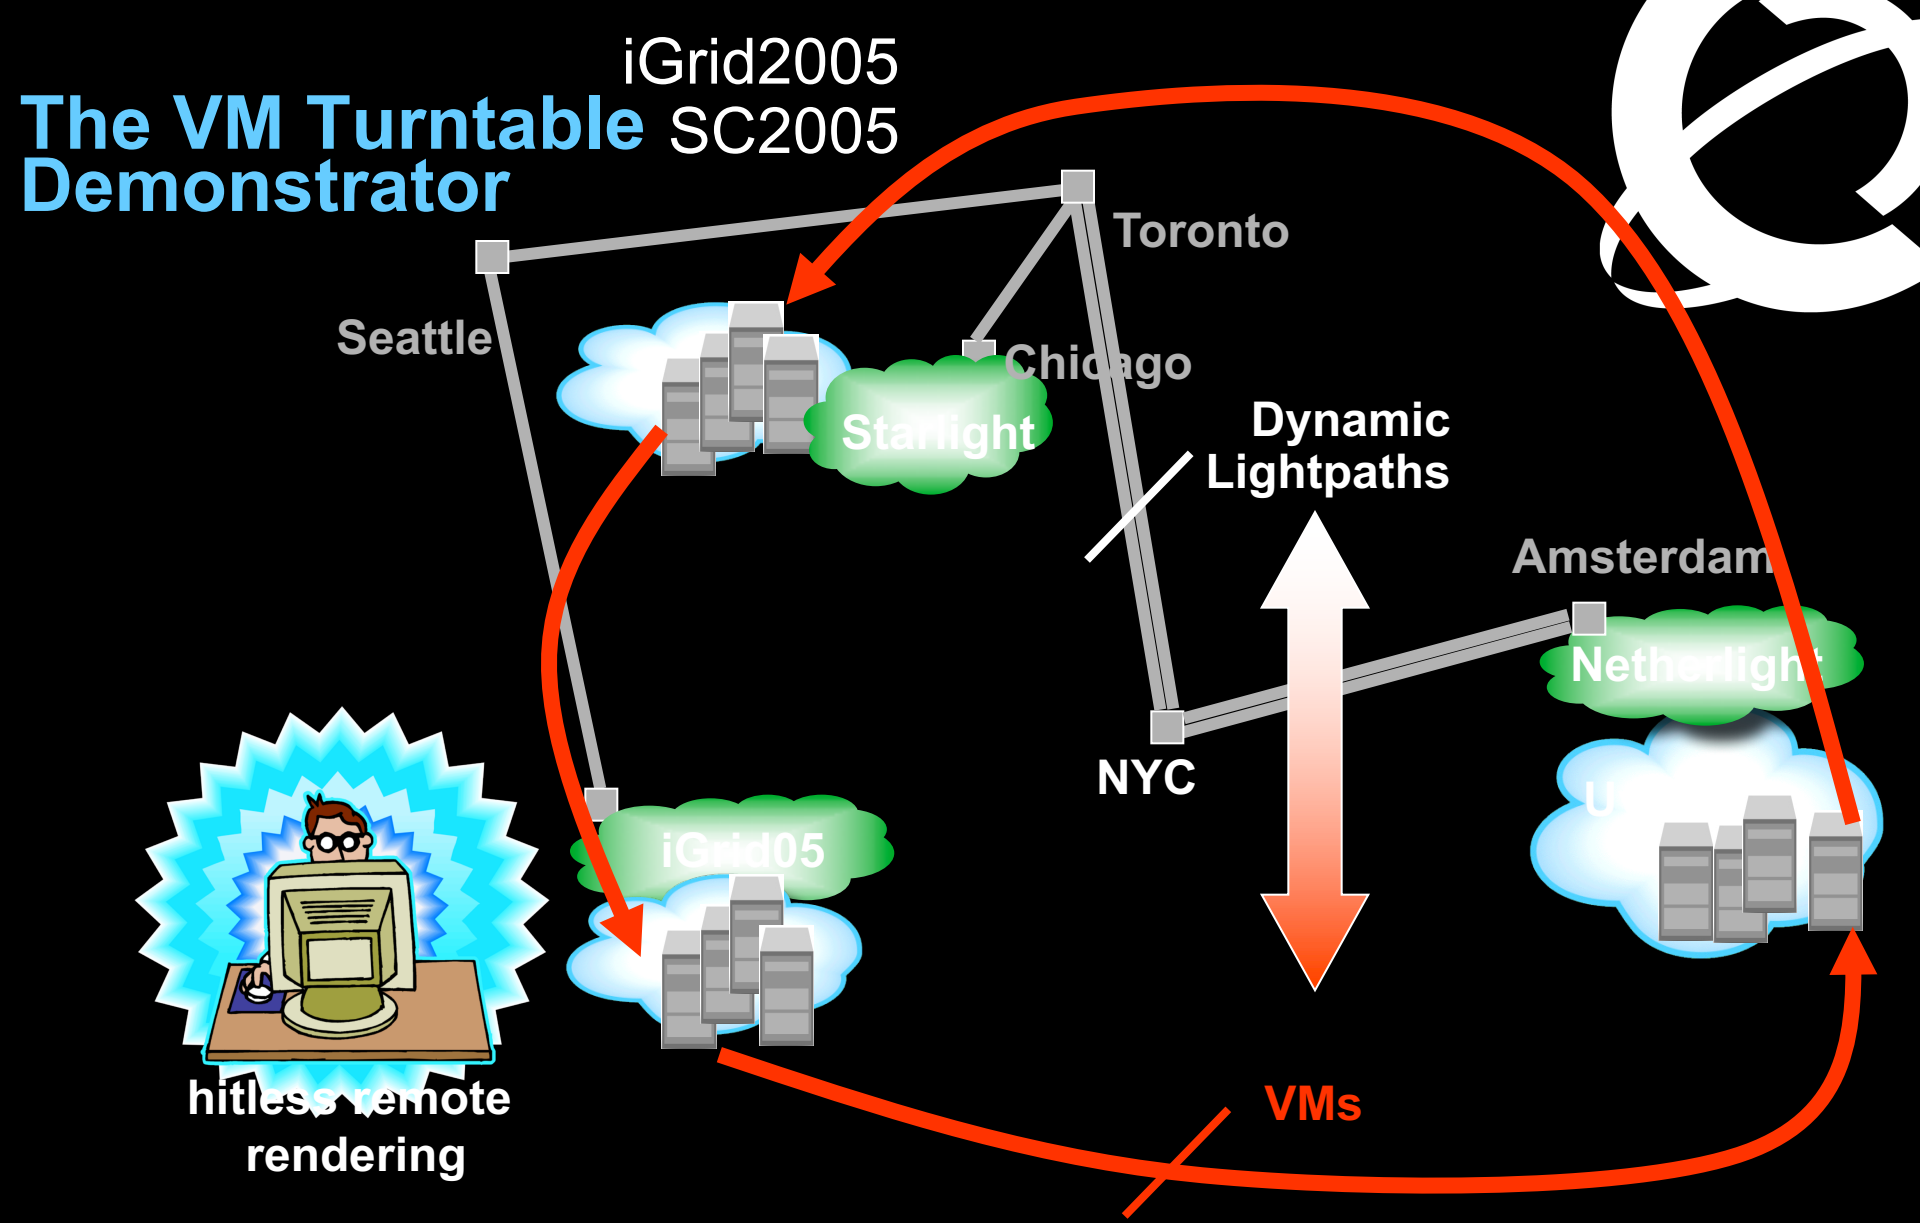

The VMs that are live-migrated run an iterative search-refine-search workflow against data stored in different databases at the various locations. A user in San Diego gets hitless rendering of search progress as VMs spin around

Experiment outcomes
Note, this was in 2005

We have demonstrated seamless, live migration of VMs over MAN/WAN

For this, we have realized a network service that

- **Exhibits predictable behavior; tracks endpoints**
- **Flex bandwidth upon request by credited applications**
- **Doesn't require peak provisioning of network resources**

Pipelining bounds the downtime in spite of high RTTs

- **San Diego – Amsterdam, 1GE, RTT = 200 msec, downtime <= 1 sec**
- **Back to back, 1GE, RTT = 0.2-0.5 msec, downtime = ~0.2 sec***
- **Clark et al. NSDI 05 paper. Different workloads*

VM + Lightpaths across MAN/WAN are deemed a powerful and general alternative to RPC, GRAM approaches

We believe it's a representative instance of active cpu+data+net orchestration

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)

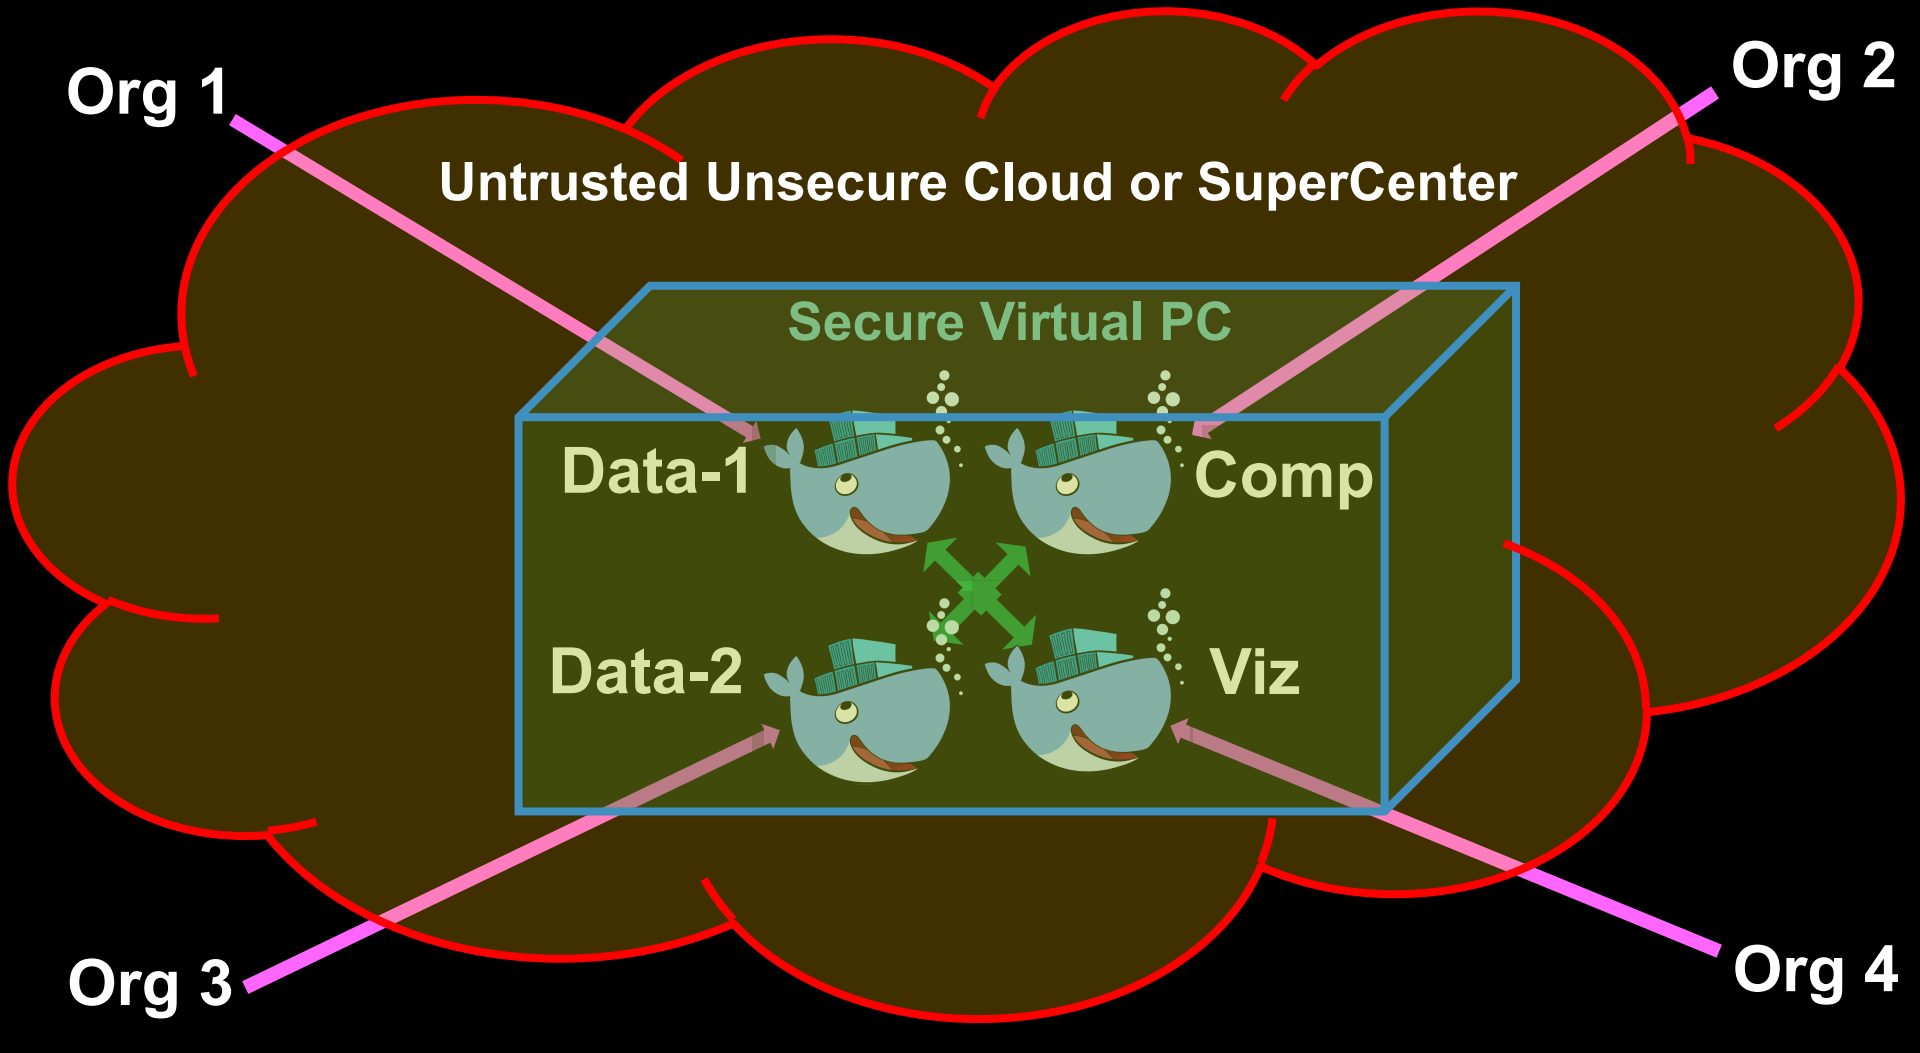

I want to

"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!

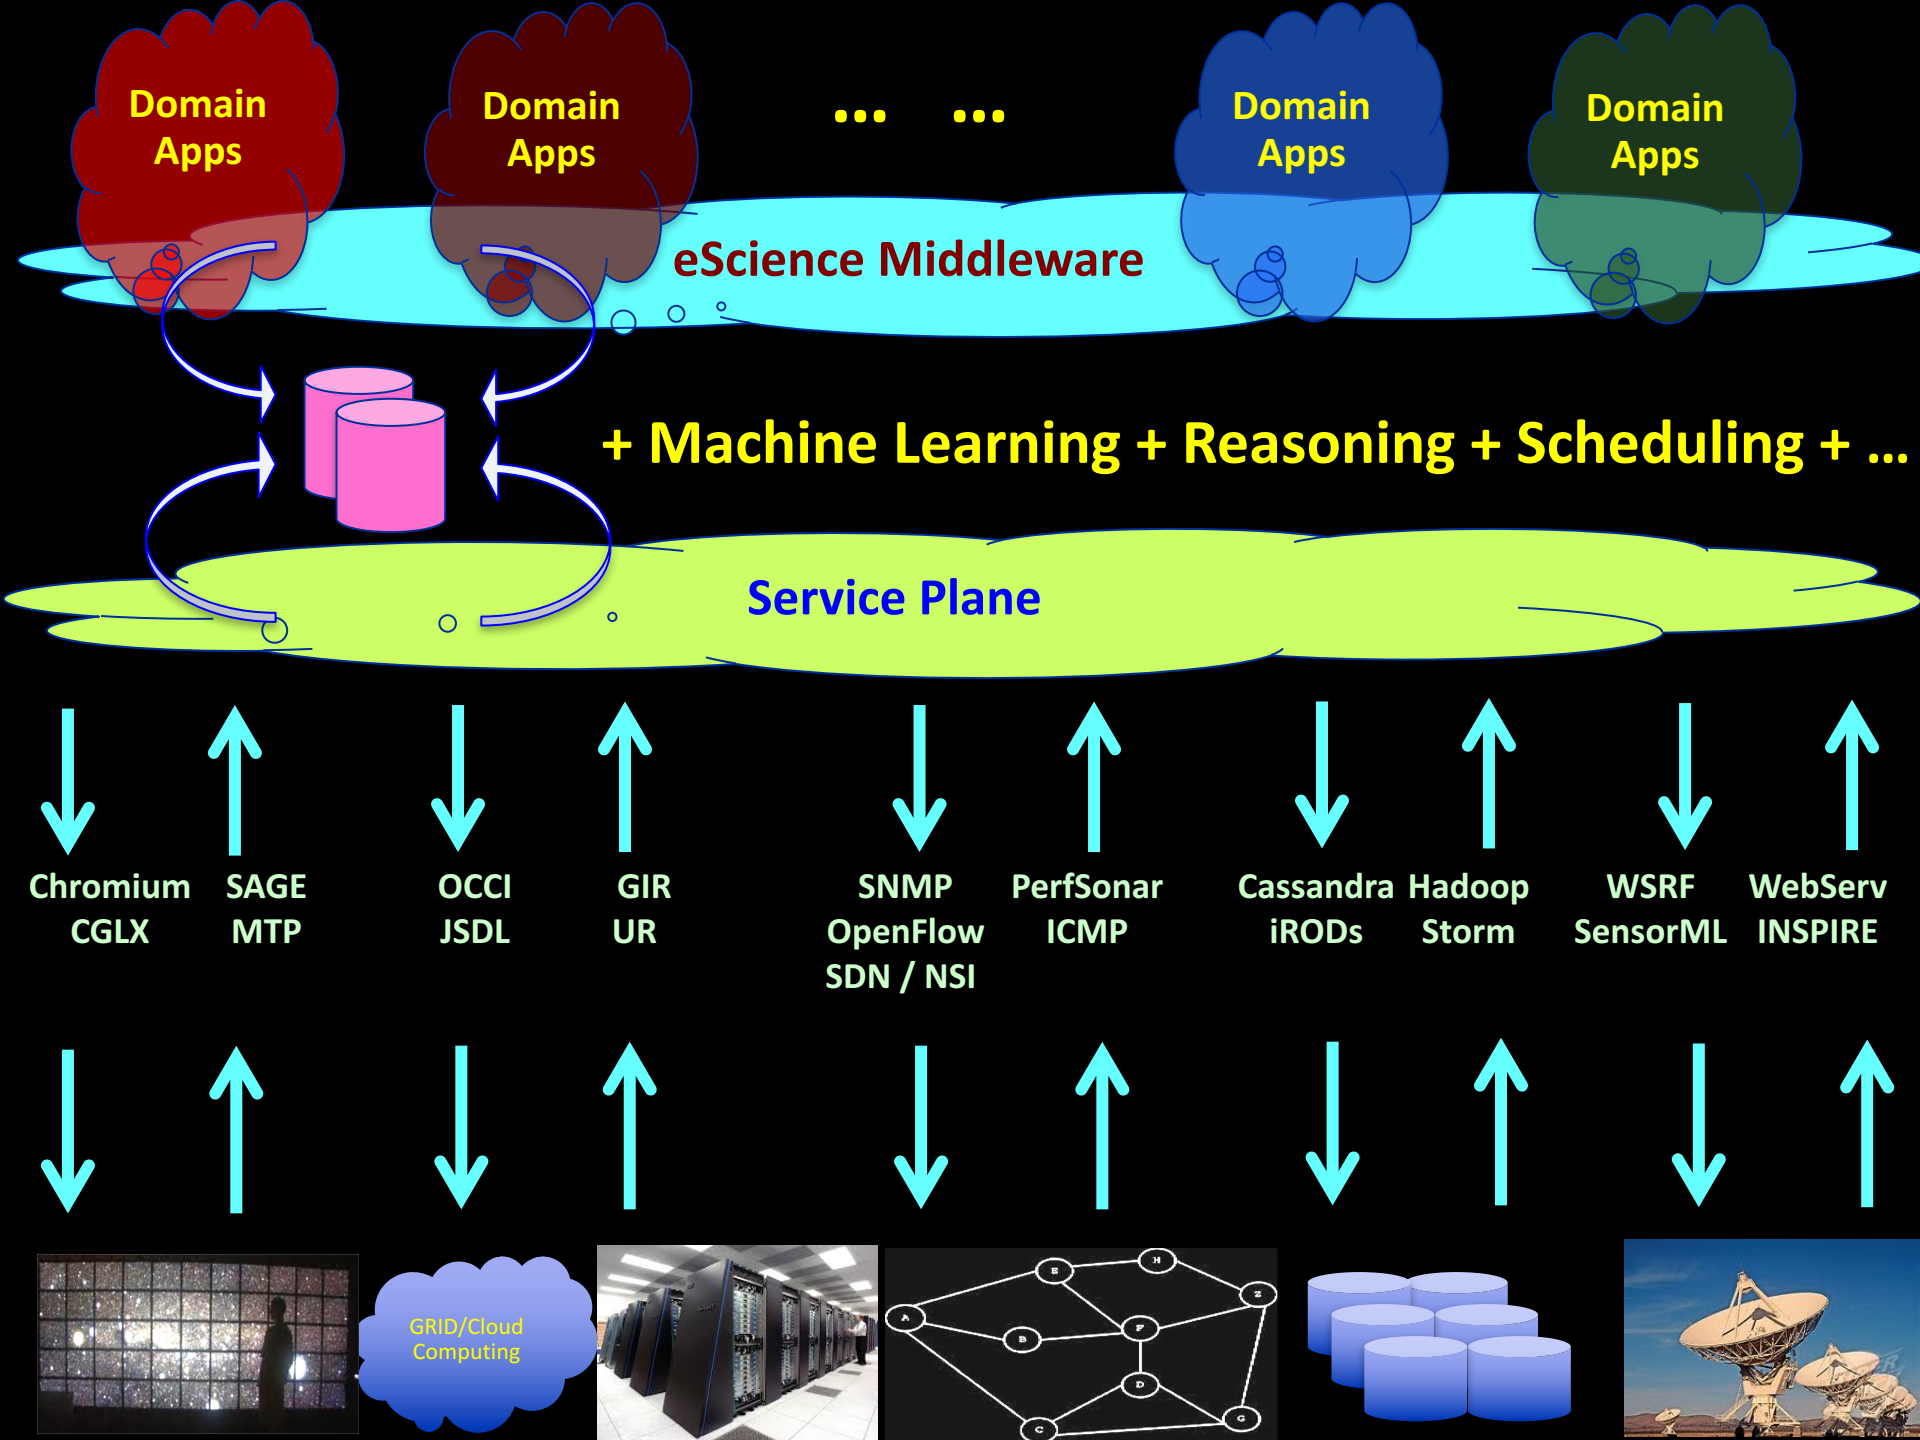
Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
+ Machine Learning + Reasoning + Scheduling + …
Service Plane
Chromium CGLX
SAGE MTP
OCCI JSDL
GIR UR
SNMP OpenFlow SDN / NSI
PerfSonar ICMP
Cassandra iRODs
Hadoop Storm
WSRF SensorML
WebServ INSPIRE
GRID/Cloud Computing

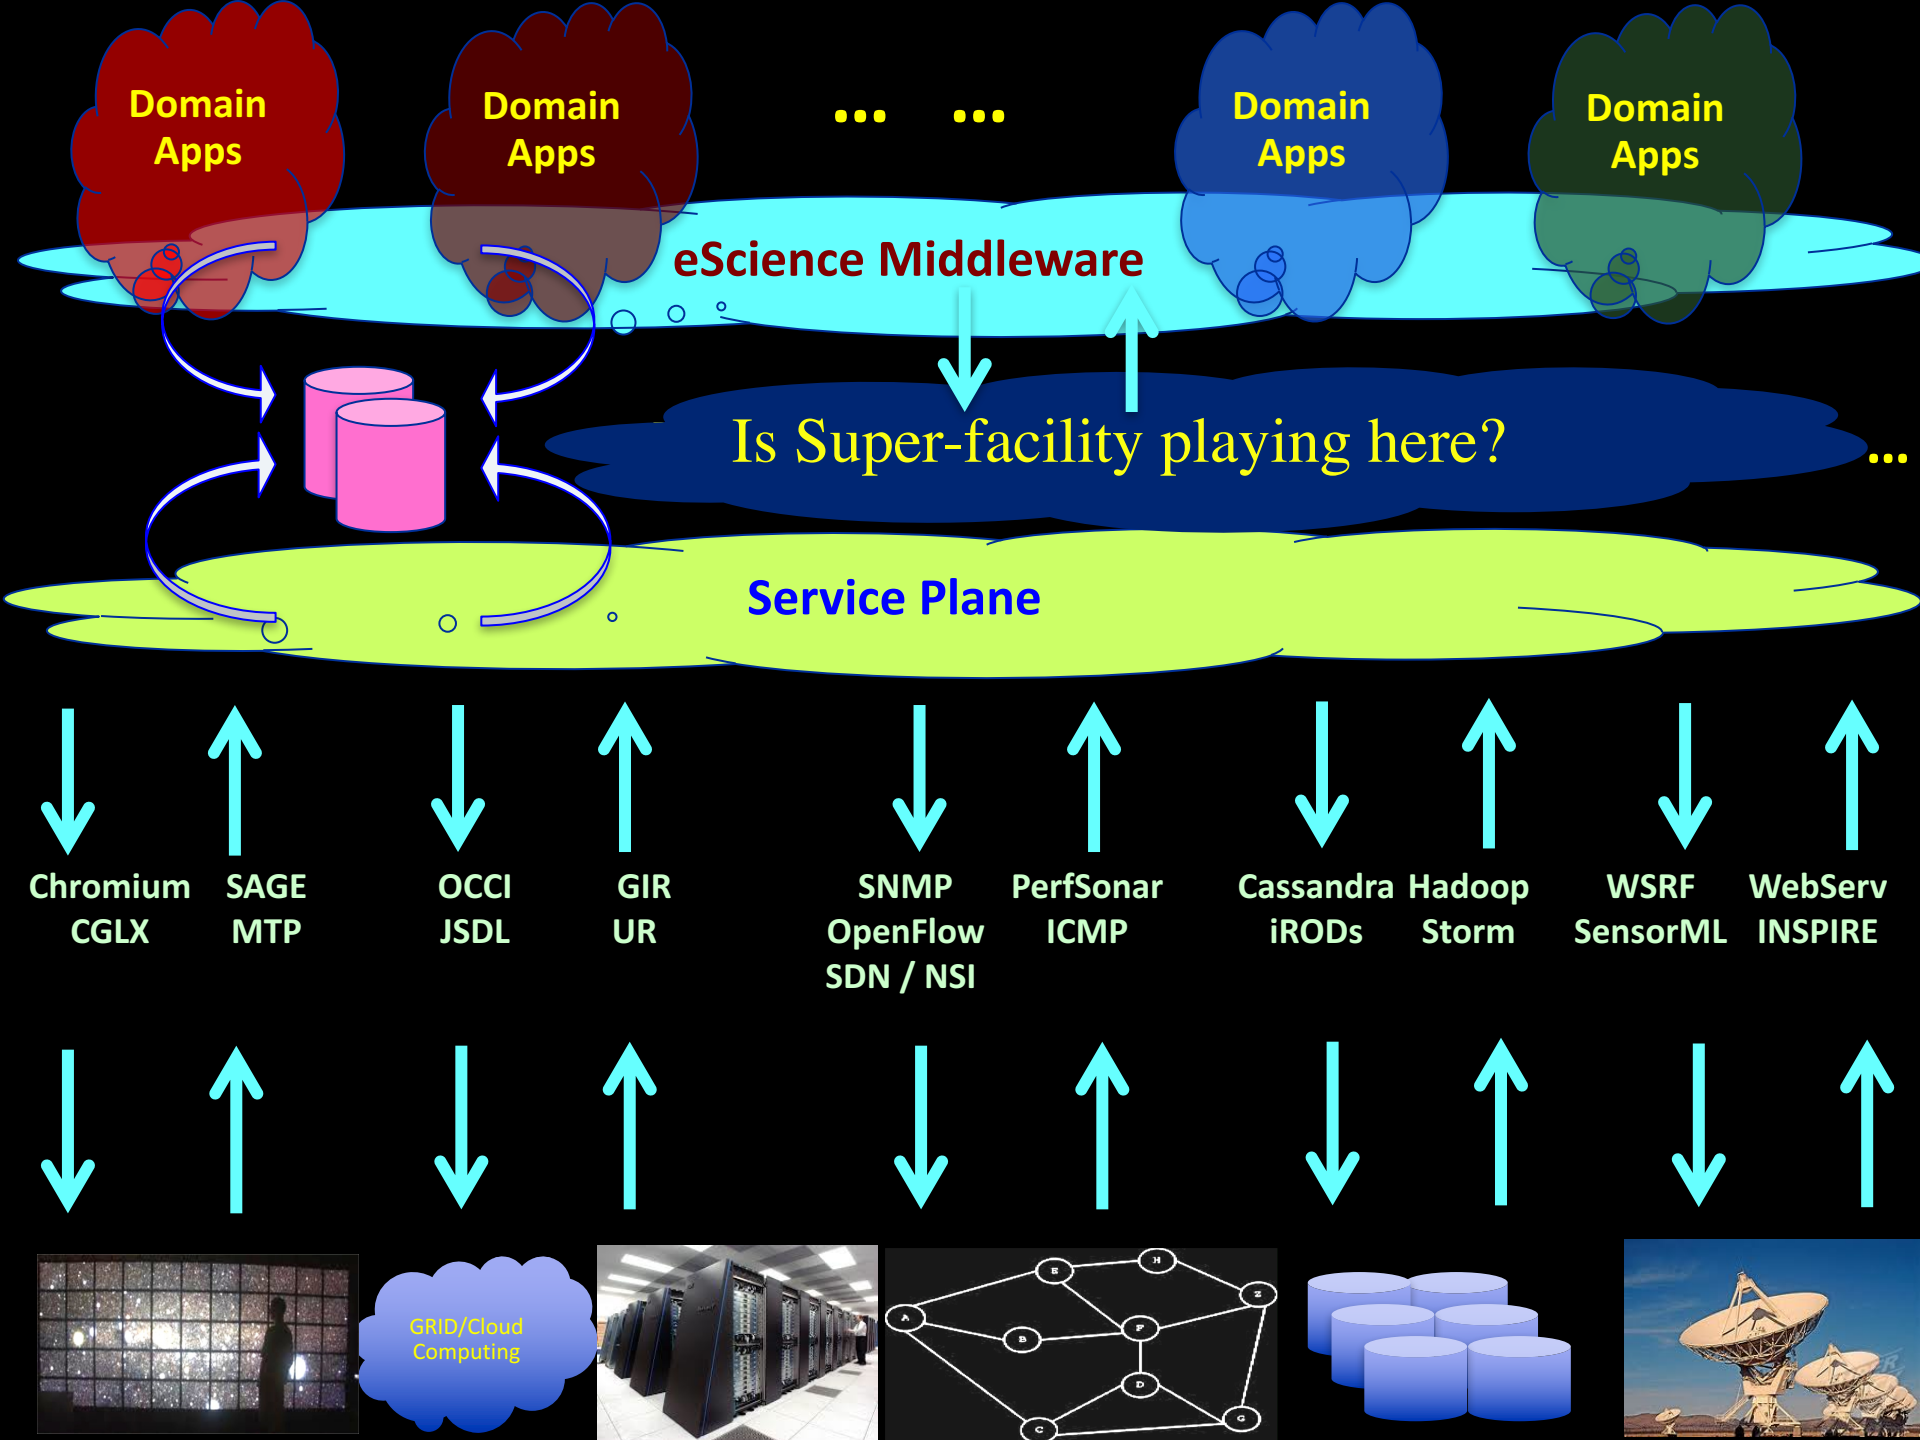
Domain Apps
Domain Apps
... ...
Domain Apps
Domain Apps
eScience Middleware
Is Super-facility playing here?
...
Service Plane
Chromium CGLX
SAGE MTP
OCCI JSDL
GIR UR
SNMP OpenFlow SDN / NSI
PerfSonar ICMP
Cassandra iRODs
Hadoop Storm
WSRF SensorML
WebServ INSPIRE
GRID/Cloud Computing

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

e-IRG

Workflows Schedulers

MAGIC DATA CARPET

curation - description - trust - security - policy – integrity

Information

OWL

Data

a.o. from ESFRI's

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



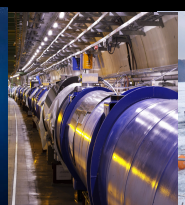


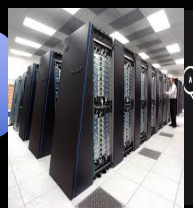
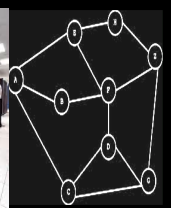


TimeLine

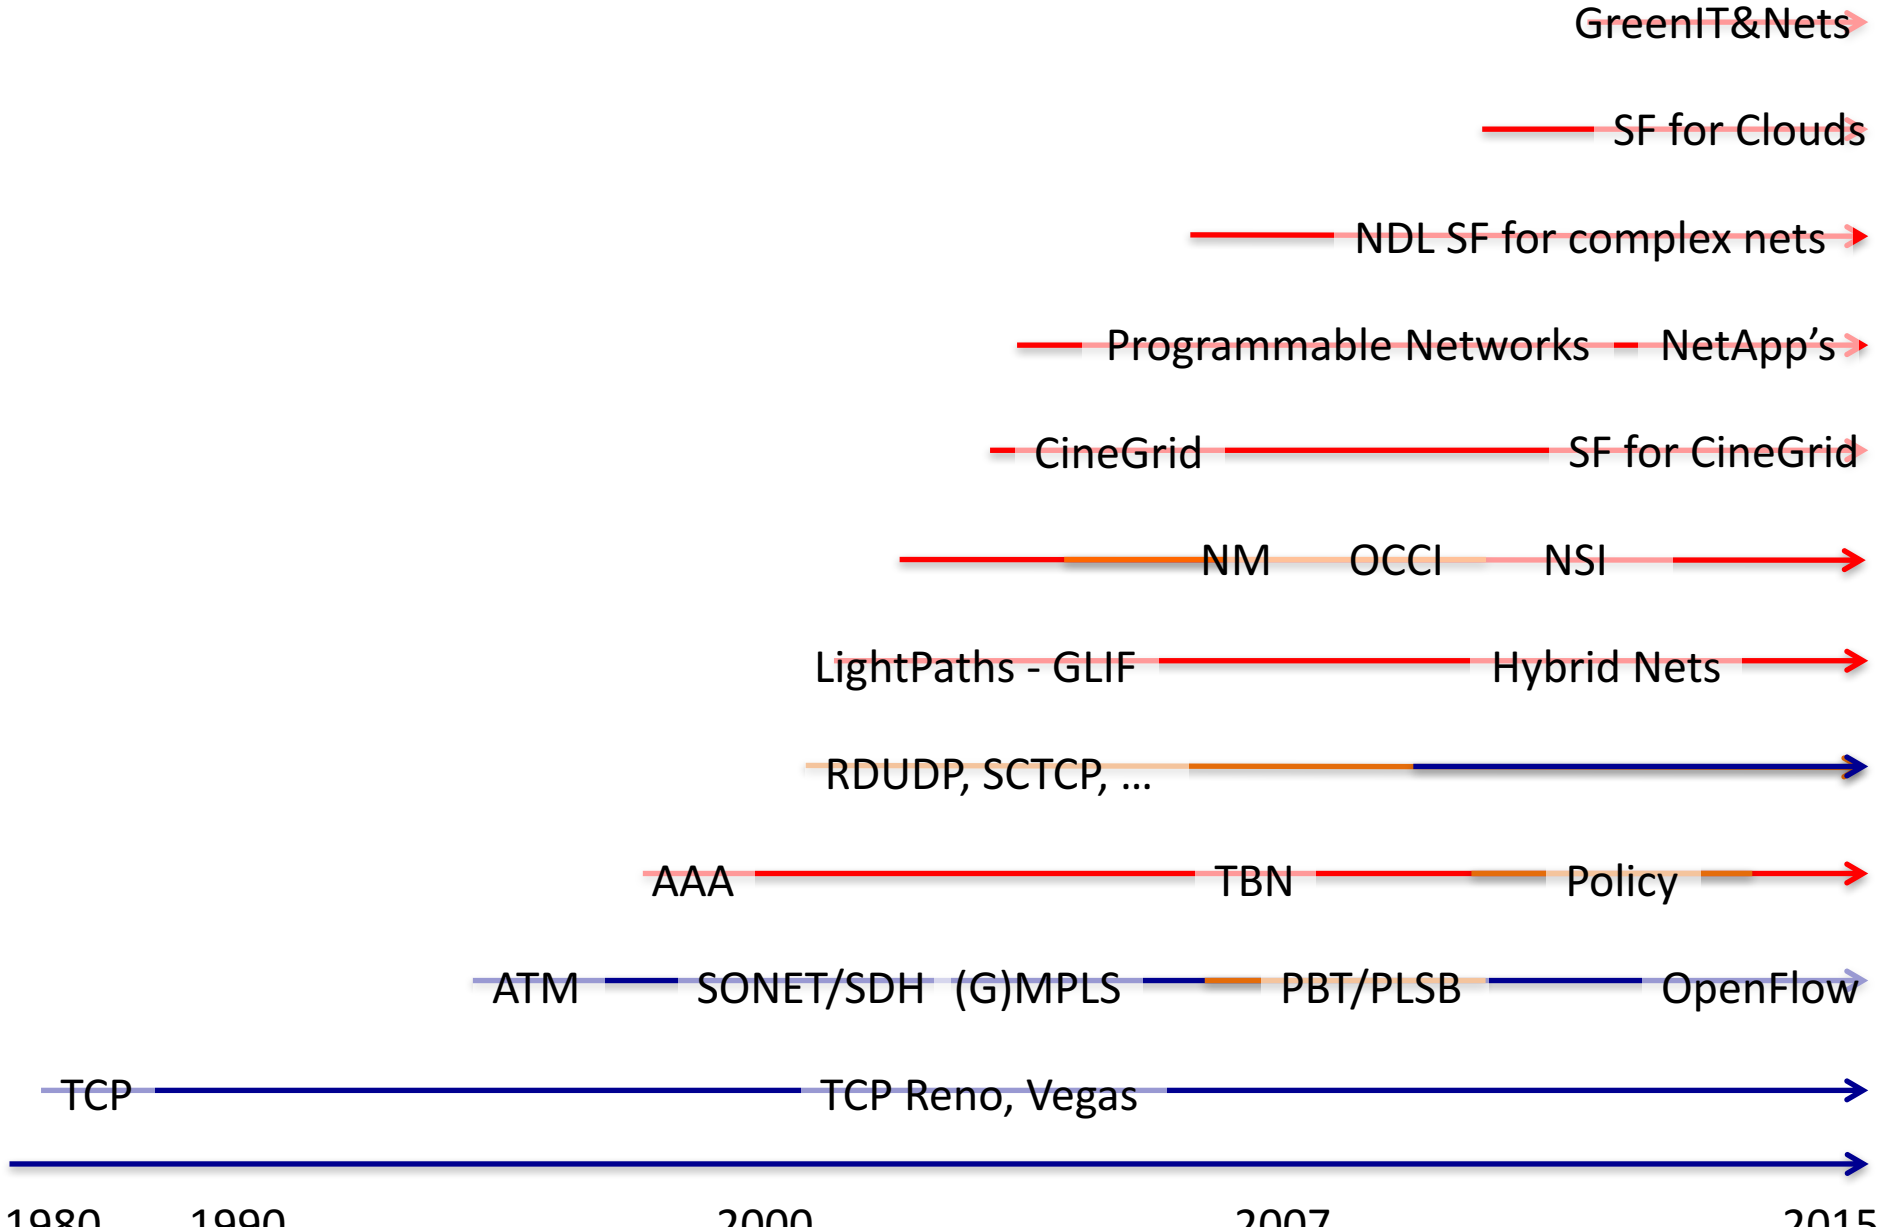
GreenIT&Nets
SF for Clouds
NDL SF for complex nets
Programmable Networks
NetApp's
CineGrid
SF for CineGrid
NM
OCCI
NSI
LightPaths - GLIF
Hybrid Nets
RDUDP, SCTCP, …
AAA
TBN
Policy
ATM
SONET/SDH
(G)MPLS
PBT/PLSB
OpenFlow
TCP
TCP Reno, Vegas
1980
1990
2000
2007
2015

TimeLine

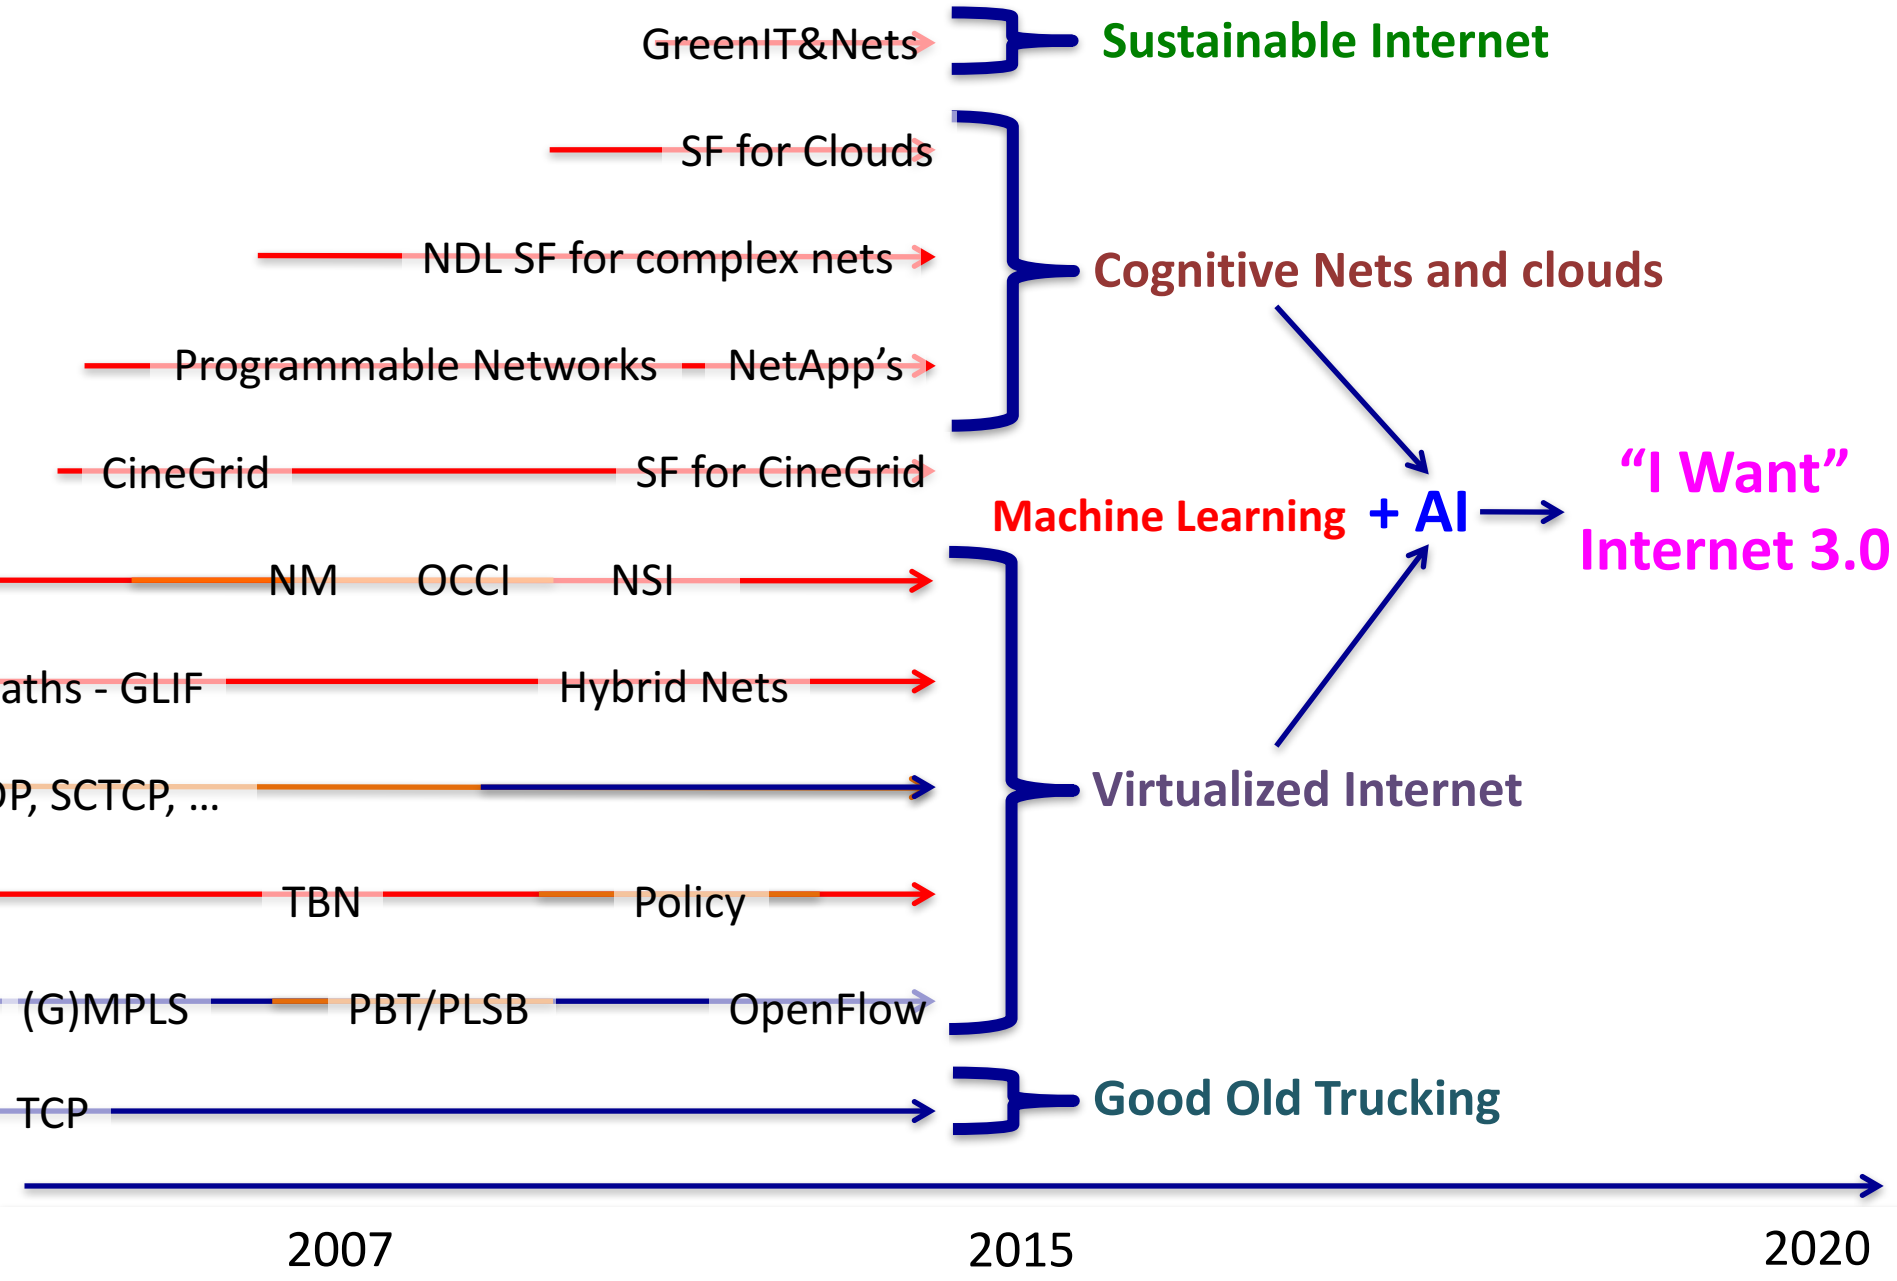
GreenIT&Nets
Sustainable Internet
SF for Clouds
NDL SF for complex nets
Programmable Networks
NetApp's
Cognitive Nets and clouds
CineGrid
SF for CineGrid
Machine Learning
+ AI
"I Want"
Internet 3.0
NM
OCCI
NSI
aths - GLIF
Hybrid Nets
DP, SCTCP, …
Virtualized Internet
TBN
Policy
(G)MPLS
PBT/PLSB
OpenFlow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

TimeLine

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

I retire

2020

2040

TimeLine

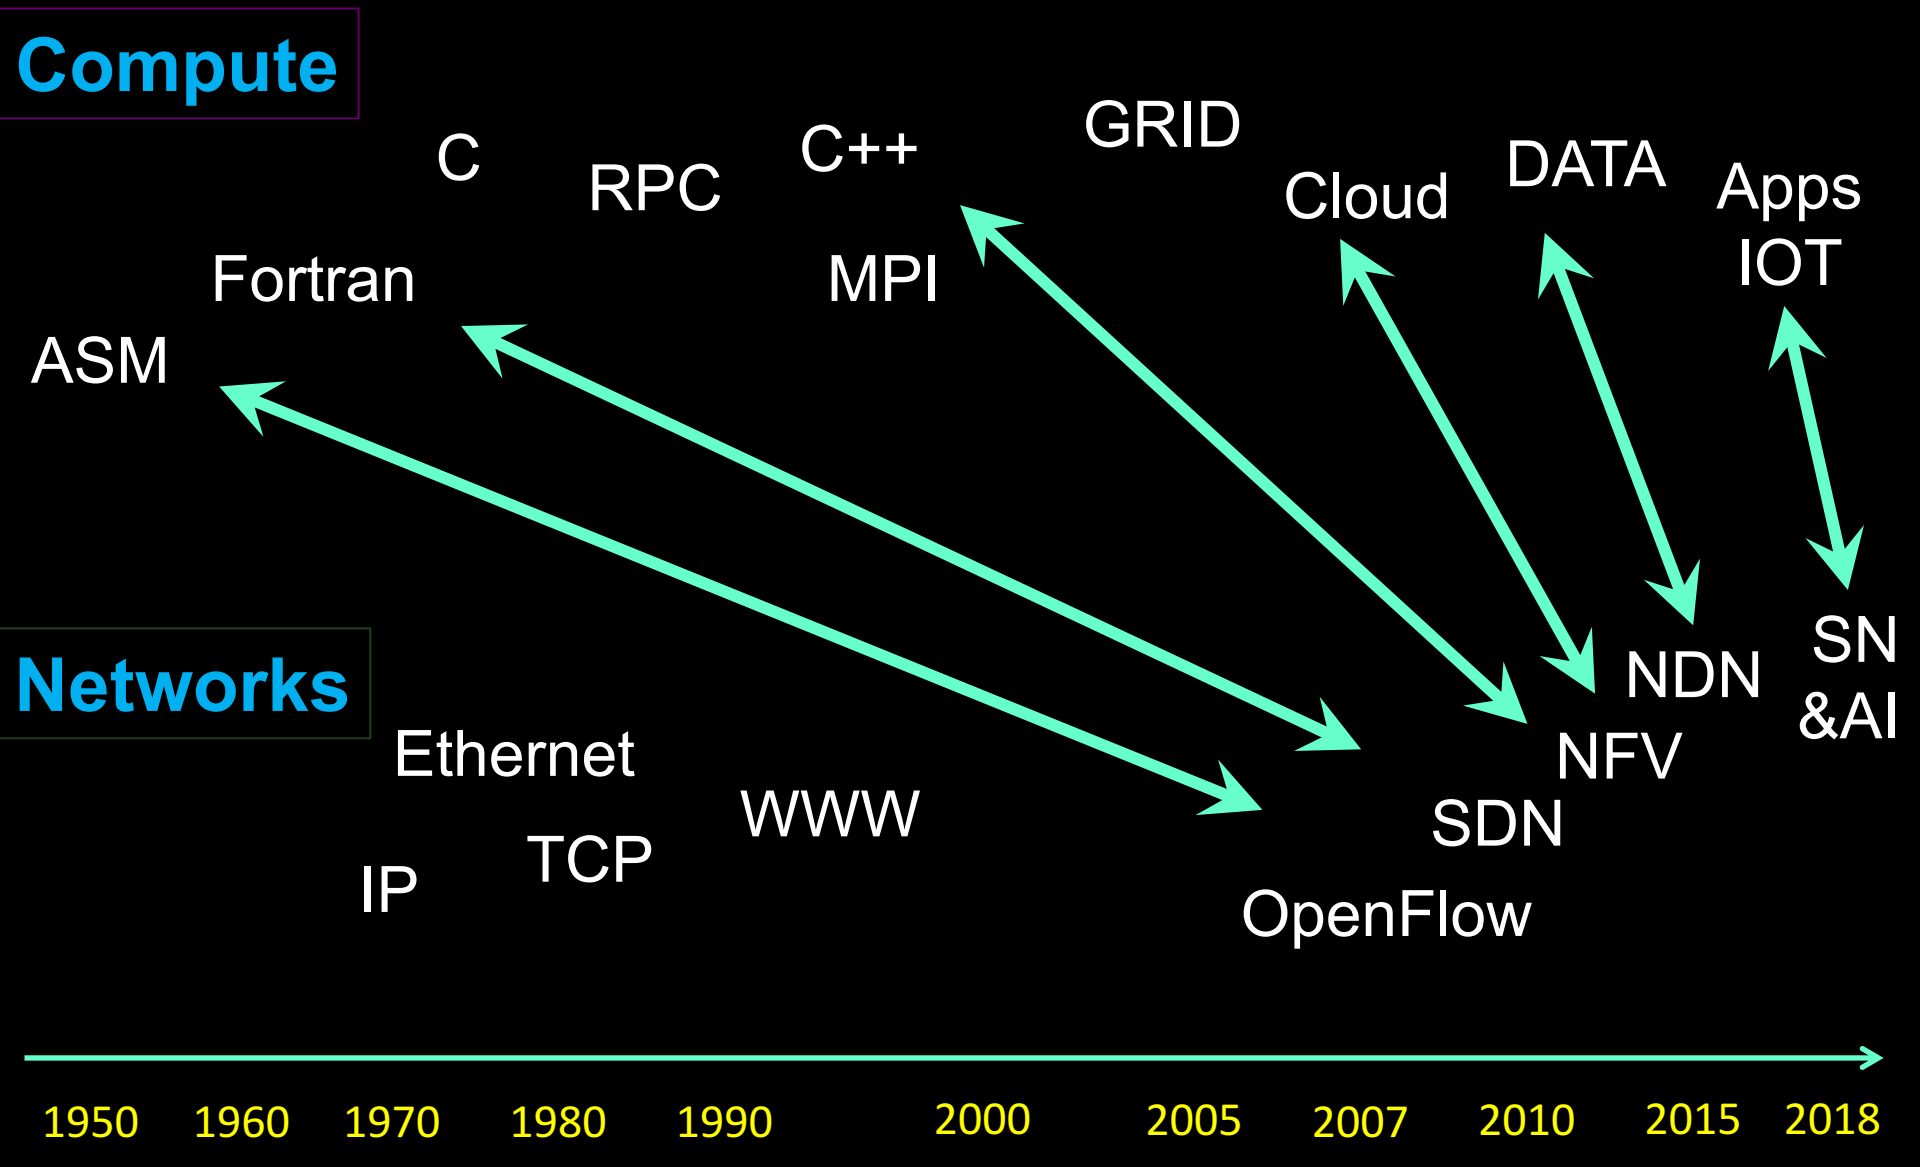
Compute
C
RPC
C++
GRID
Cloud
DATA
Apps
IOT
Fortran
MPI
ASM
Networks
Ethernet
WWW
IP
TCP
SDN
NFV
NDN
SN
&AI
OpenFlow
1950
1960
1970
1980
1990
2000
2005
2007
2010
2015
2018

What has AI to do with the Dutch National Science quiz 2013?

- Q13: For an illness that 1 out of 1000 people suffer, a 99% accurate test is developed. You are tested with that method and found bearer of the illness. What is the probability that you really have the specific illness?
- Choose: [A: 99%, B: 50%, C: 9%]
- Answer C: because you are in the set of true and false positives!
- Suppose the accuracy of PRISM, Tempora, Xkeyscore, etc. is 99% and 1 out of 100000 of the subjects are indeed terrorists
- False positives among 100k … ~1000 !
- Send in the drones: http://www.businessinsider.com/nsa-cia-drone-program-2013-10?international=true&r=US&IR=T

Areas of research

- Each domain its own AI on networks.
  - Multiple AI's fighting on my behalf?
- A-B-C slide
  - Where makes what AI sense?
- Many layers of complexity and abstraction.
  - Can AI help to understand and debug?
  - Can it explicitly understand? Reveal a model?
- Probabilities are badly understood in AI
  - How to deal with false positives?
  - Ethical issues?
  - Trust issues?
  - Intention issues?

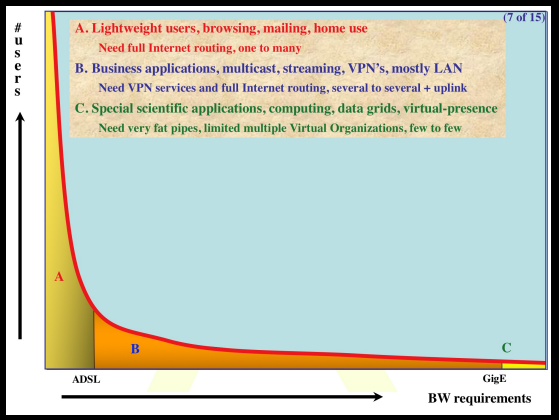

Critical notes

- We created complexity
- Huge number of actors (devices)
- Millions of lines of codes
- We have shrinking trust in the Internet
- Let’s throw in another hunderd-thousend lines of code! Good luck…
- Complexity encapsulation
- Do we have enough information for RL - ML?
- Do we understand what the Machine needs to learn?

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, Data^3, App, AI

to:

DDOS attacks destroying Banks and Bitcoins.

Conclusion:

Need for Safe, Smart, Resilient, Sustainable Infrastructure.