

From
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!”

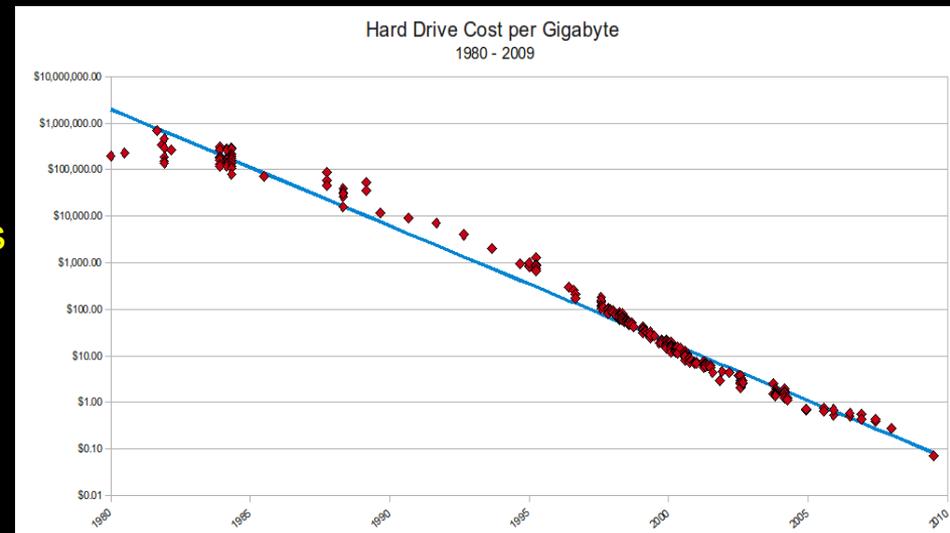


Reliable and Safe!

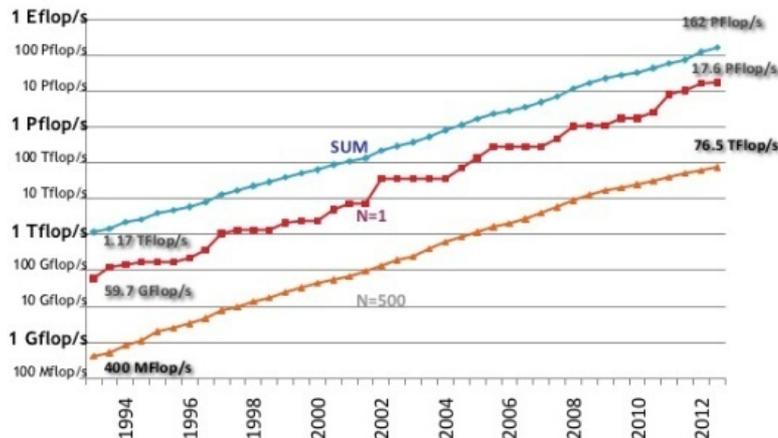
This omnipresence of IT makes us not only strong but also vulnerable.

- 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.



Performance Development



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

Mission

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

Mission

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

- *Capacity*
- *Capability*
- *Security*
- *Sustainability*
- *Resilience*

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*
 - *Authorization, Anonymity, integrity of data in distributed data processing*
- *Sustainability*
 - *Greening infrastructure, awareness*
- *Resilience*
 - *Systems under attack, failures, disasters*

Reduction of Complexity by Integration

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

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

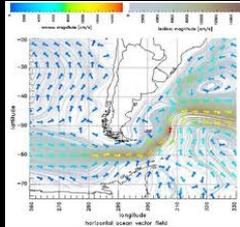
SNE addresses a.o. the **highlighted** questions!



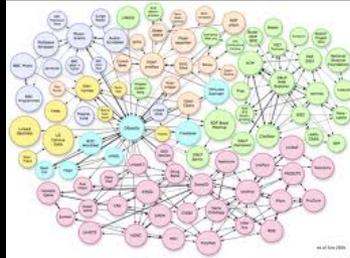
... more data!

Internet developments

Google



DATA



... more realtime!



twitter



myspace
a place for freedom



SchoolBANK



Linked in

Hyves

flickr
from YAHOO!



... more users!

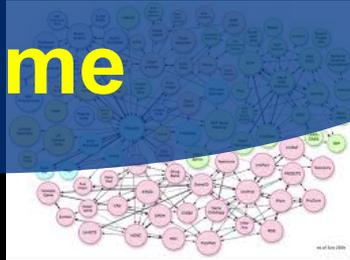
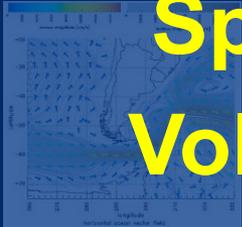
... more data!

Internet developments

Google

Speed
Volume

DATA



Deterministic

Real-time



twitter



Scalable

Secure

Linked in



myspace
SchoolBANK

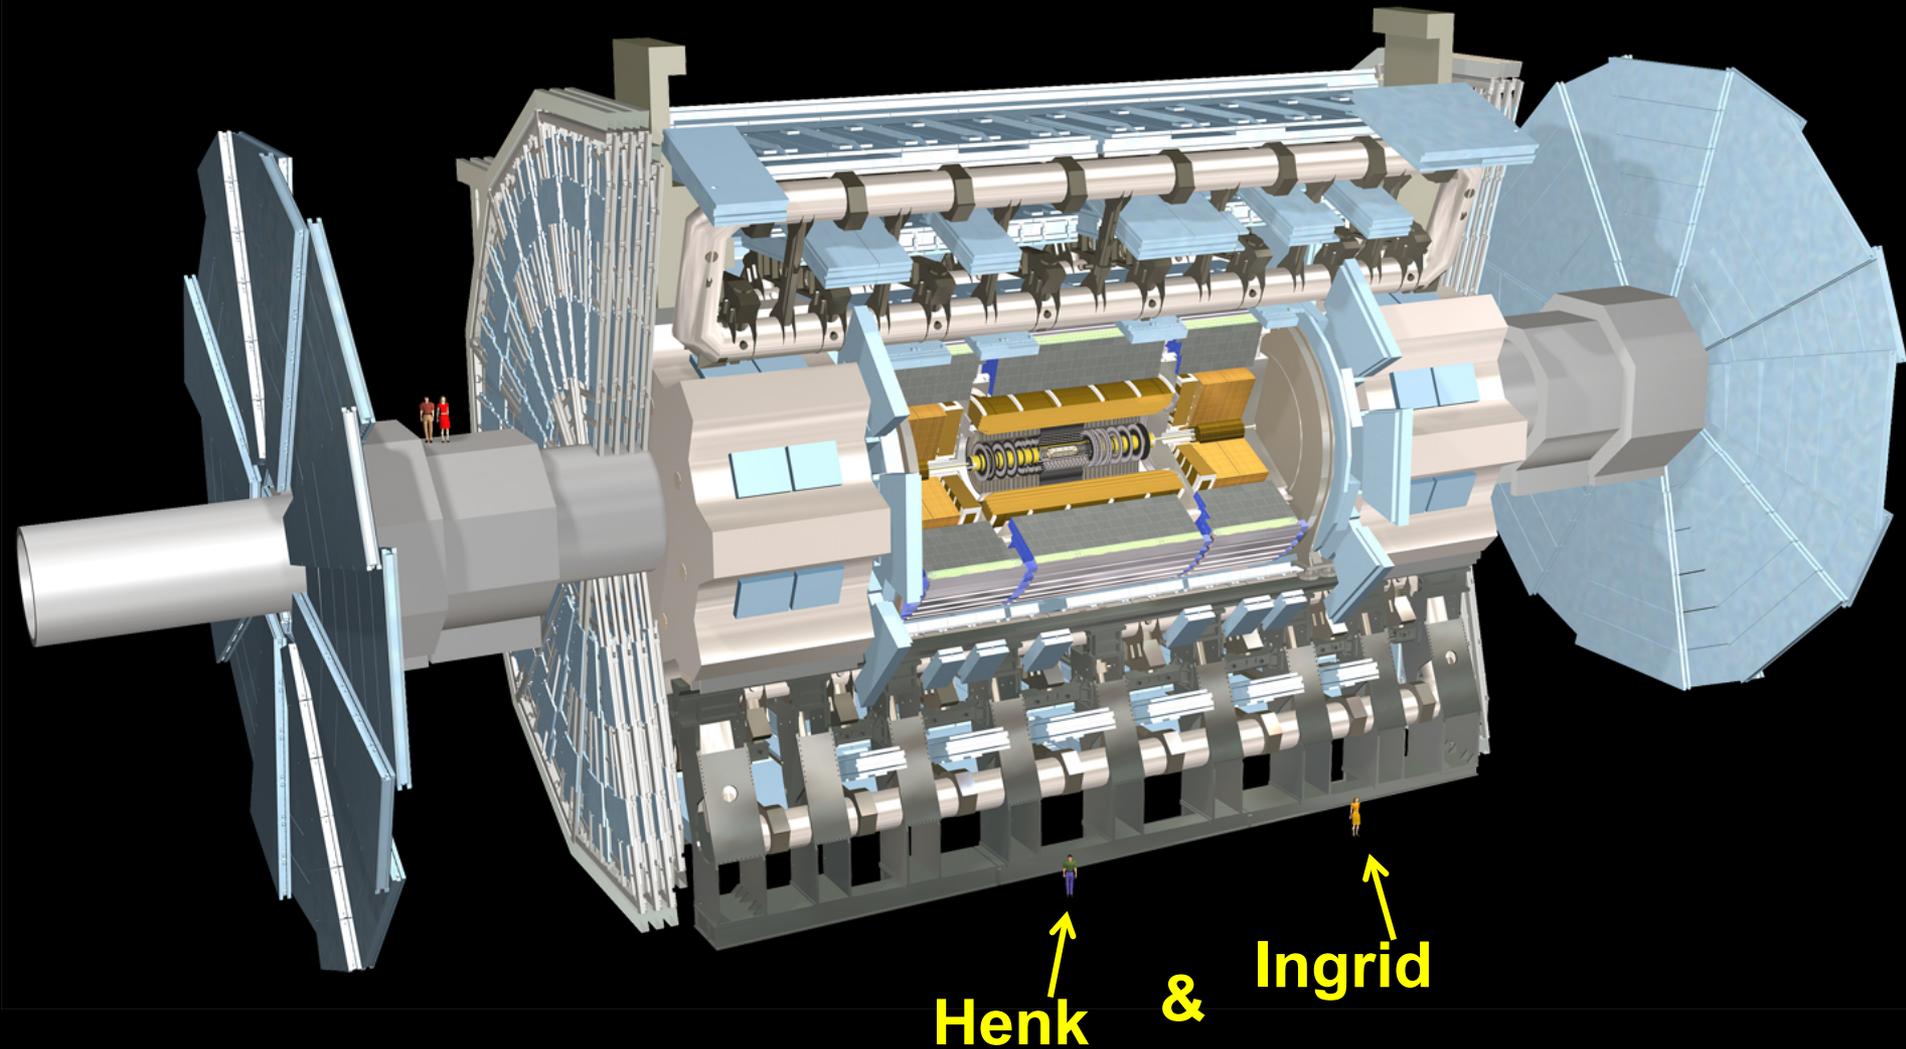
Hyves

flickr



... more users!

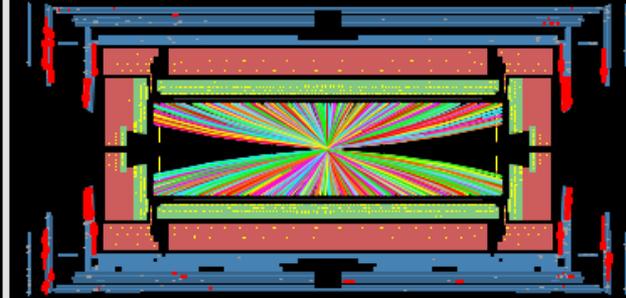
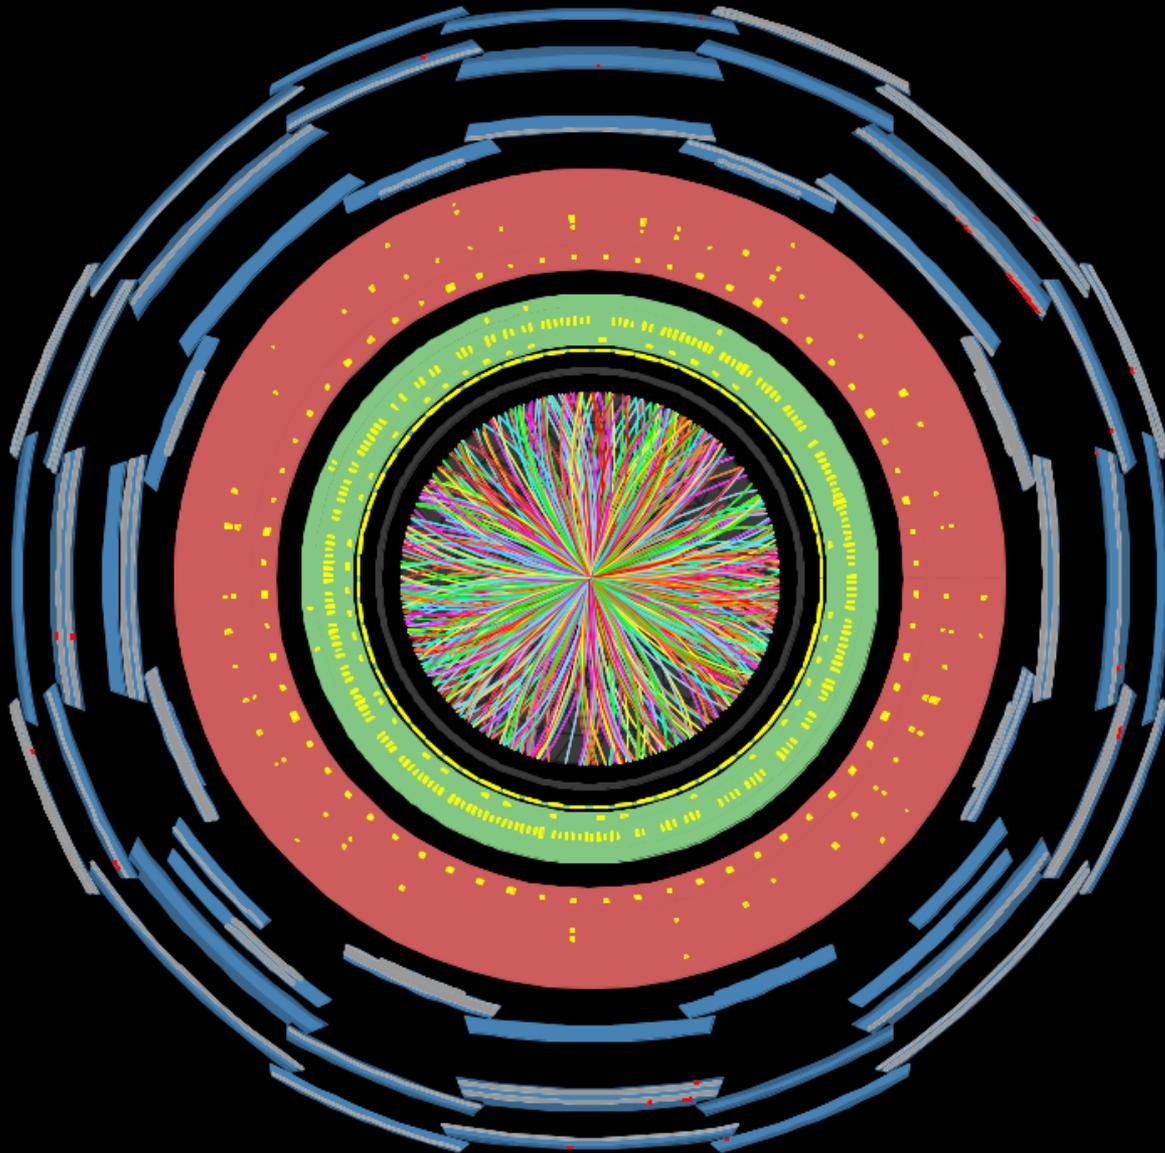
ATLAS detector @ CERN Geneve



ATLAS detector @ CERN Geneve



Een gebeurtenis



ATLAS
EXPERIMENT

Run Number: 170482, Event Number: 3936308

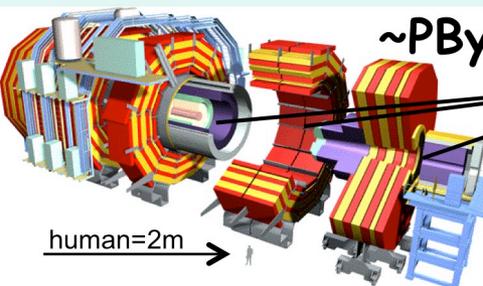
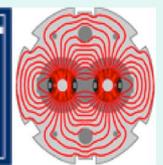
Date: 2010-12-06 17:21:31 CET

Snapshot of a heavy ion collision
directly from the ATLAS experiment



LHC Data Grid Hierarchy

CMS as example, Atlas is similar



human=2m →

CMS detector: 15m X 15m X 22m
12,500 tons, \$700M.

Online System

~PByte/sec

Tier 0 + 1

~100 MBytes/sec

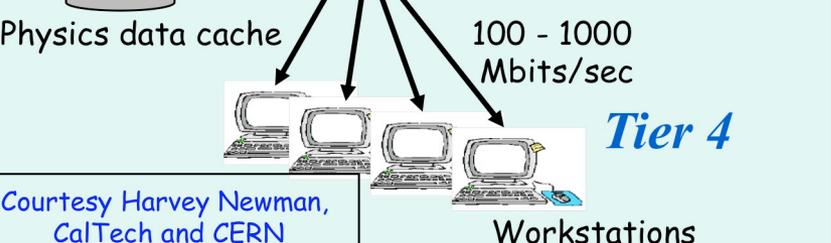
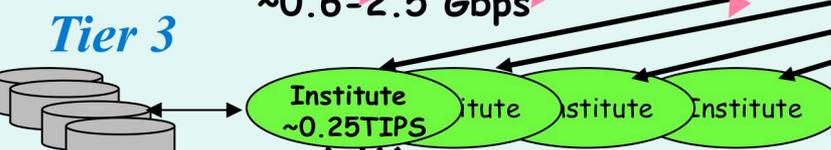
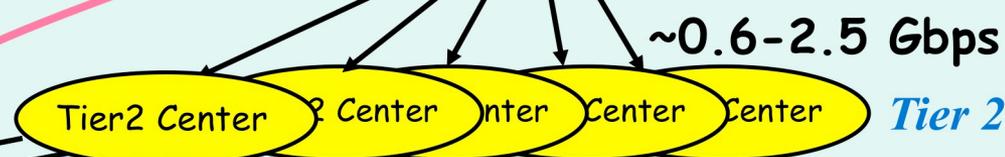
100000 flops/byte

10 Pflops/s

event simulation

event reconstruction

~2.5 Gbits/sec



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

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

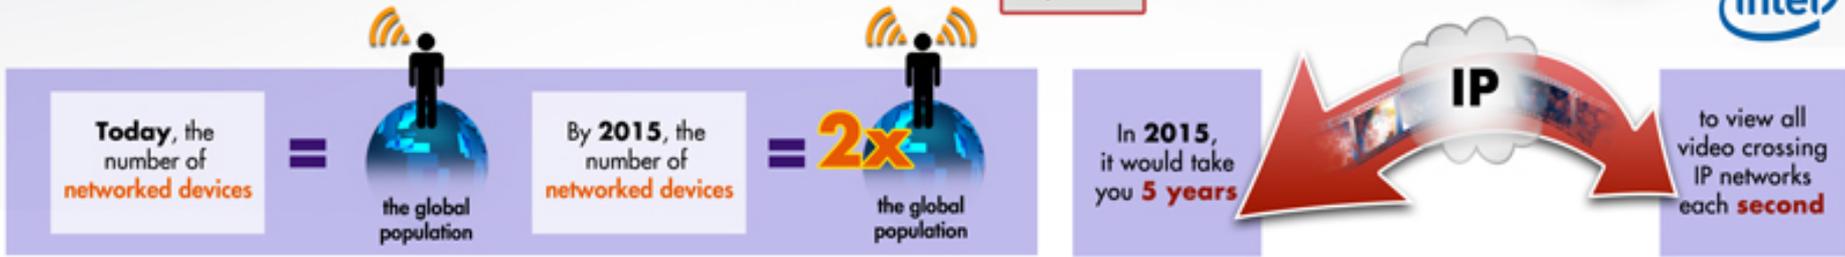
2000 physicists in 31 countries are involved in this 20-year experiment in which DOE is a major player.

Courtesy Harvey Newman, CalTech and CERN

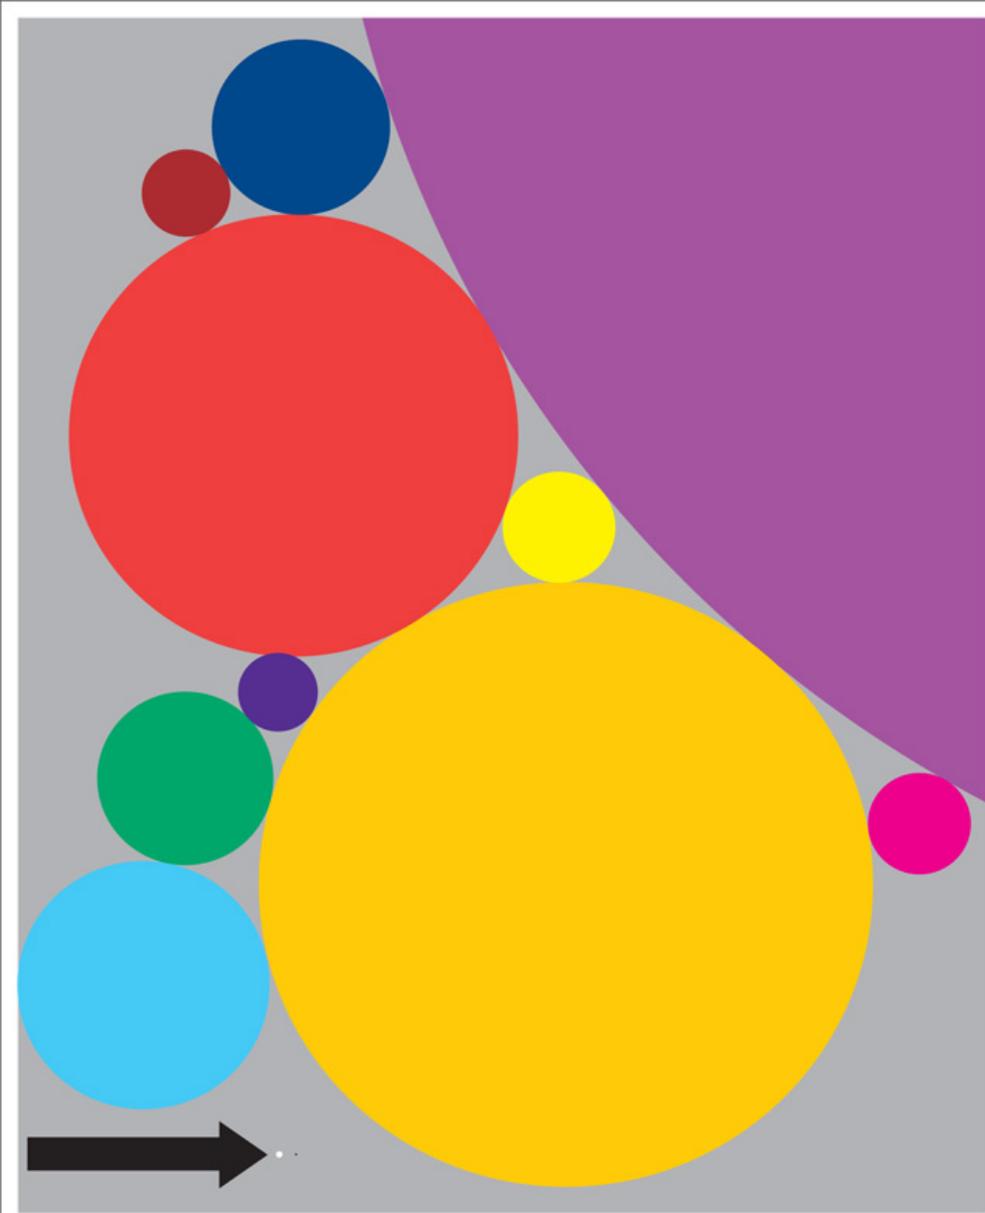
What Happens in an Internet Minute?



And Future Growth is Staggering



There
is
always
a
bigger
fish

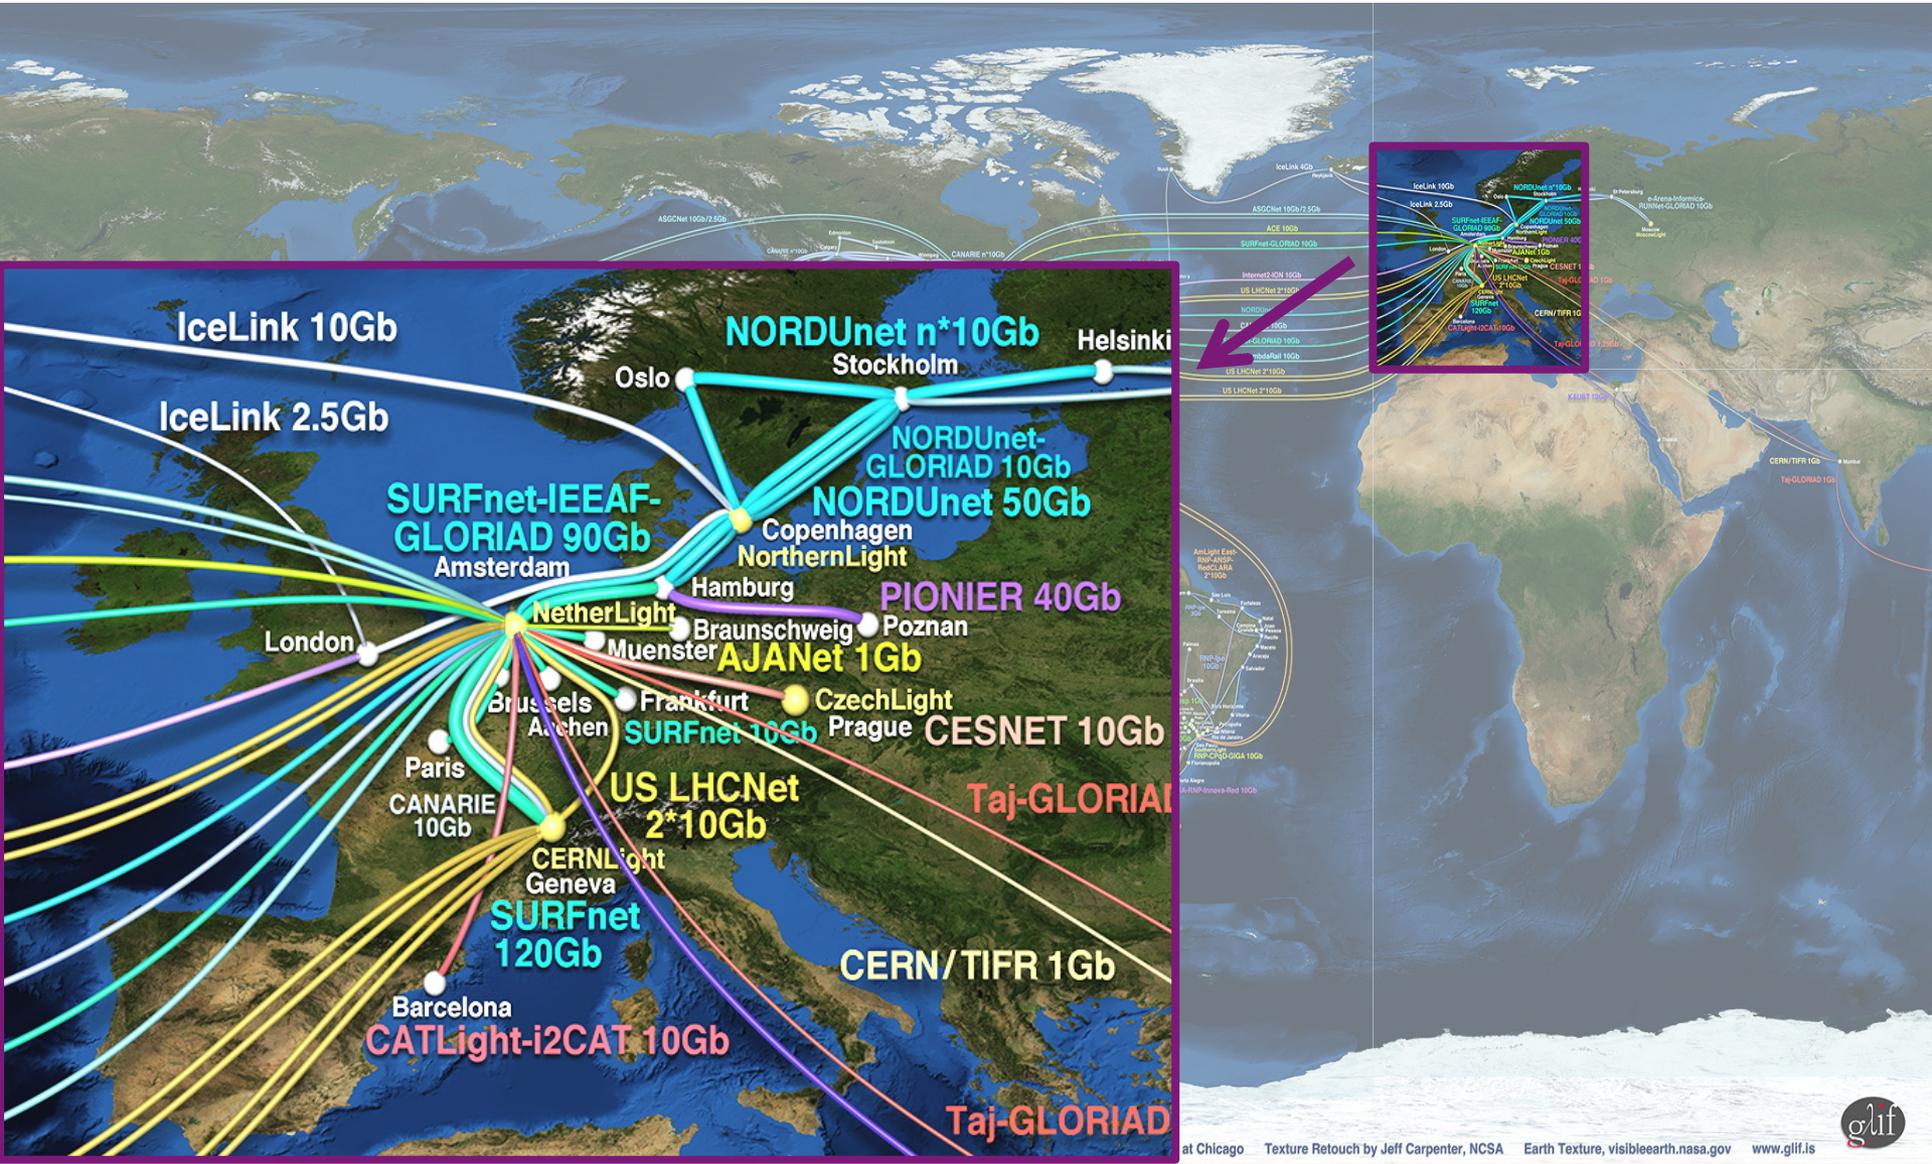


Size of data sets in terabytes

| | | | |
|--|-----------|---|-------|
| 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 |

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.



VIZUALIZATION

DataExploration

RemoteControl

TV

Medical

CineGrid



Gaming

Conference

Workflow

Clouds



Distributed

EventProcessing

GRID&CLOUD

Management

Mining

Web2.0



Meta

DATACENTER

Backup

Media

Visualisation

Security

NetherLight

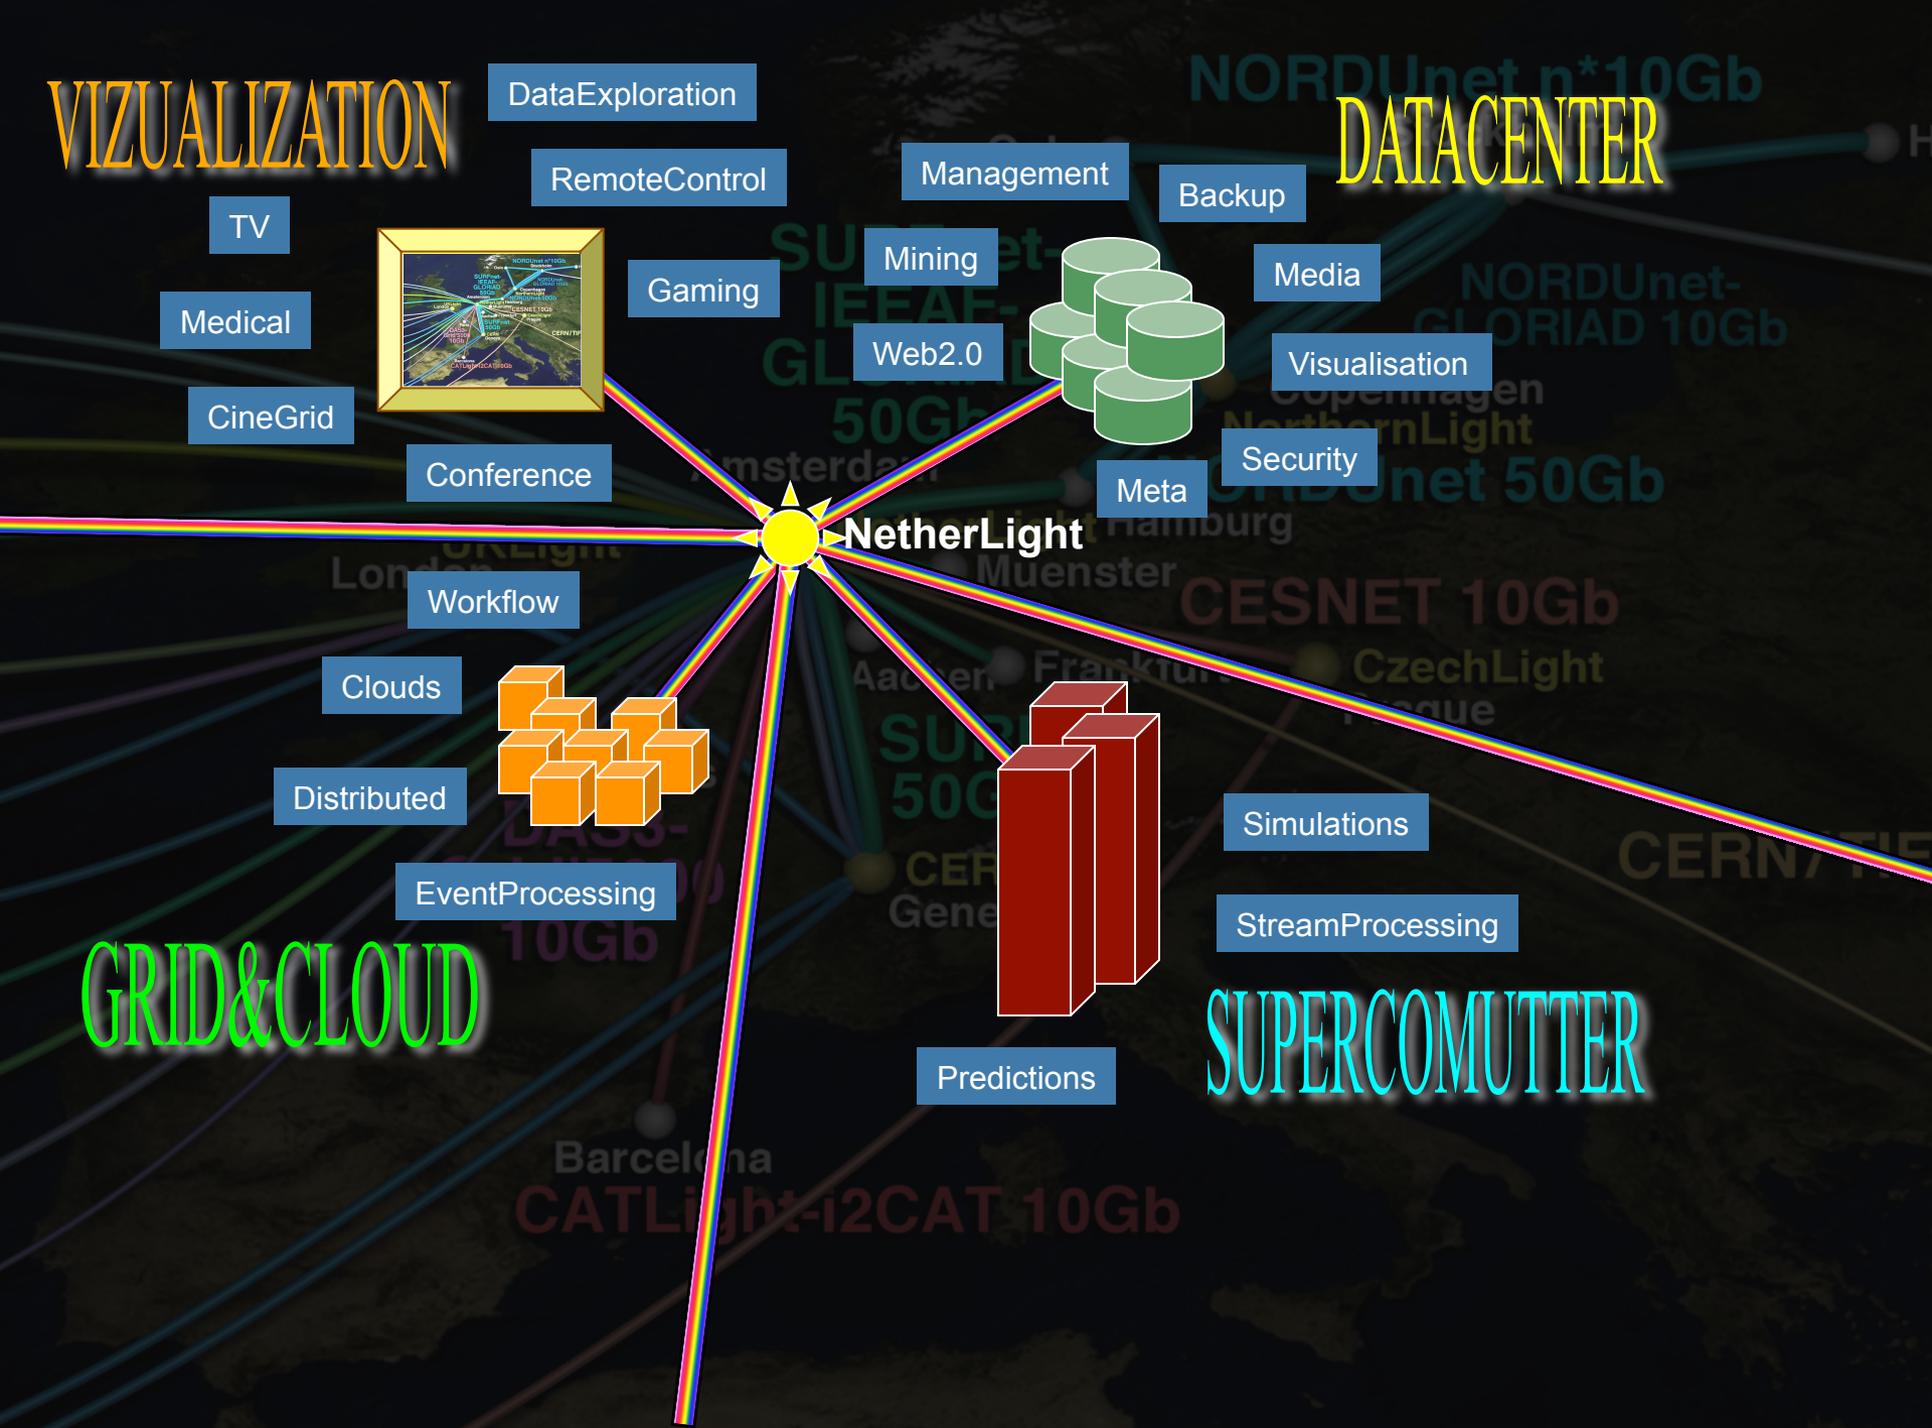
Simulations

StreamProcessing

SUPERCOMUTTER

Predictions

Barcelona
CATLight-i2CAT 10Gb





I want to



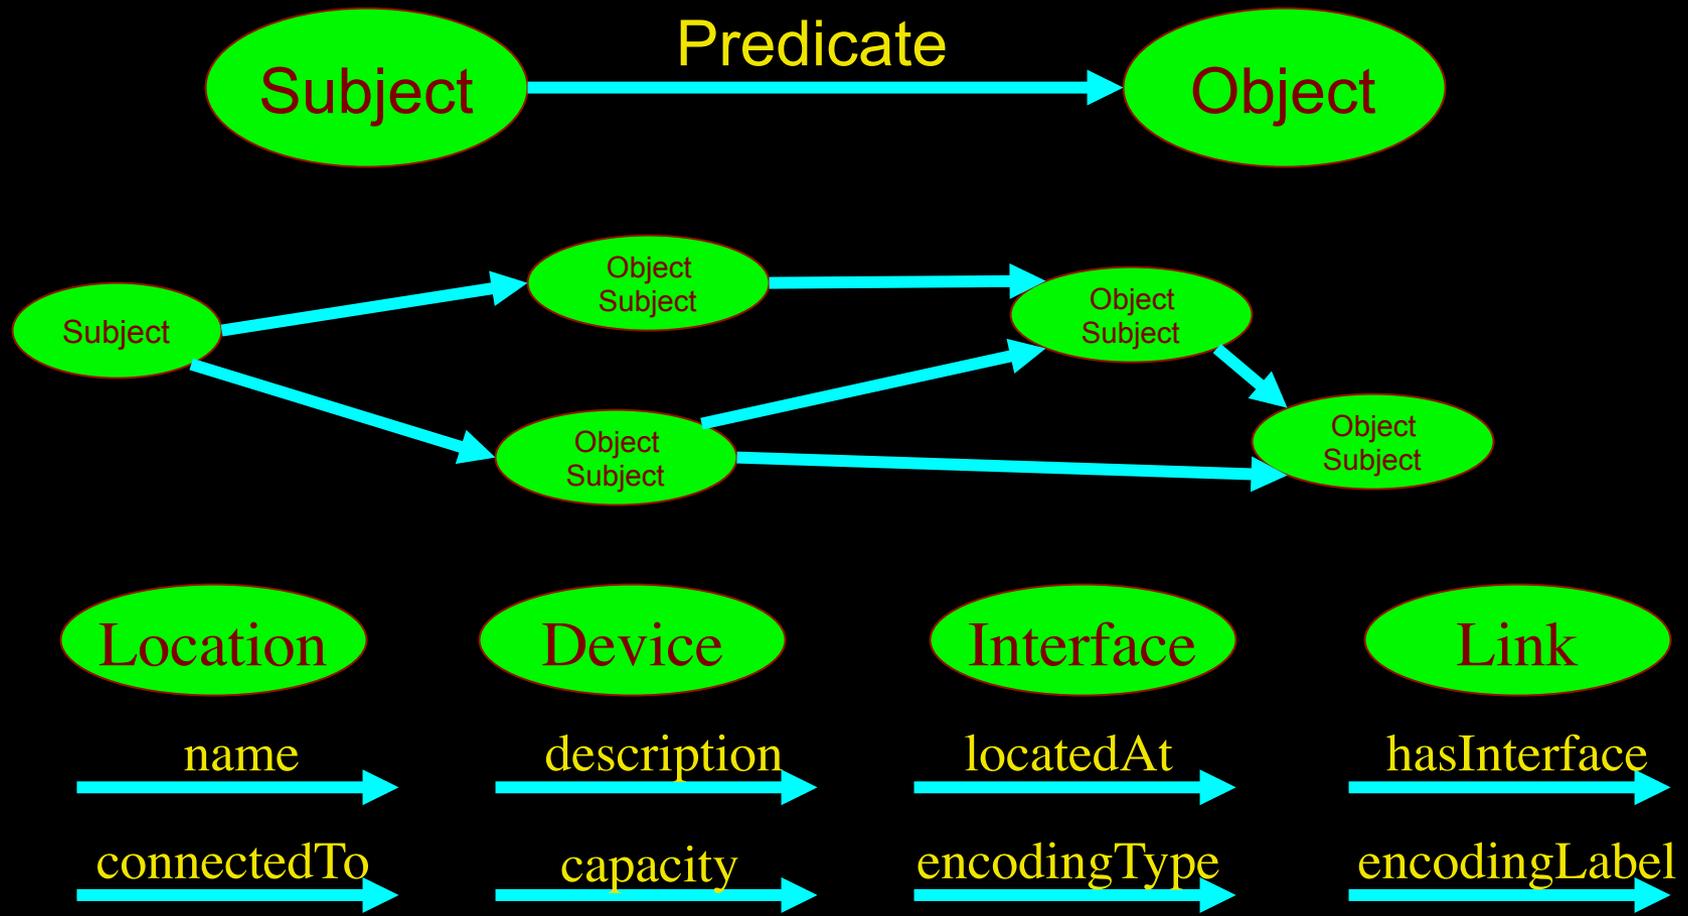
“Show Big Bug Bunny in 4K on my Tiled Display using green Infrastructure”

- 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.

LinkedIn for Infrastructure



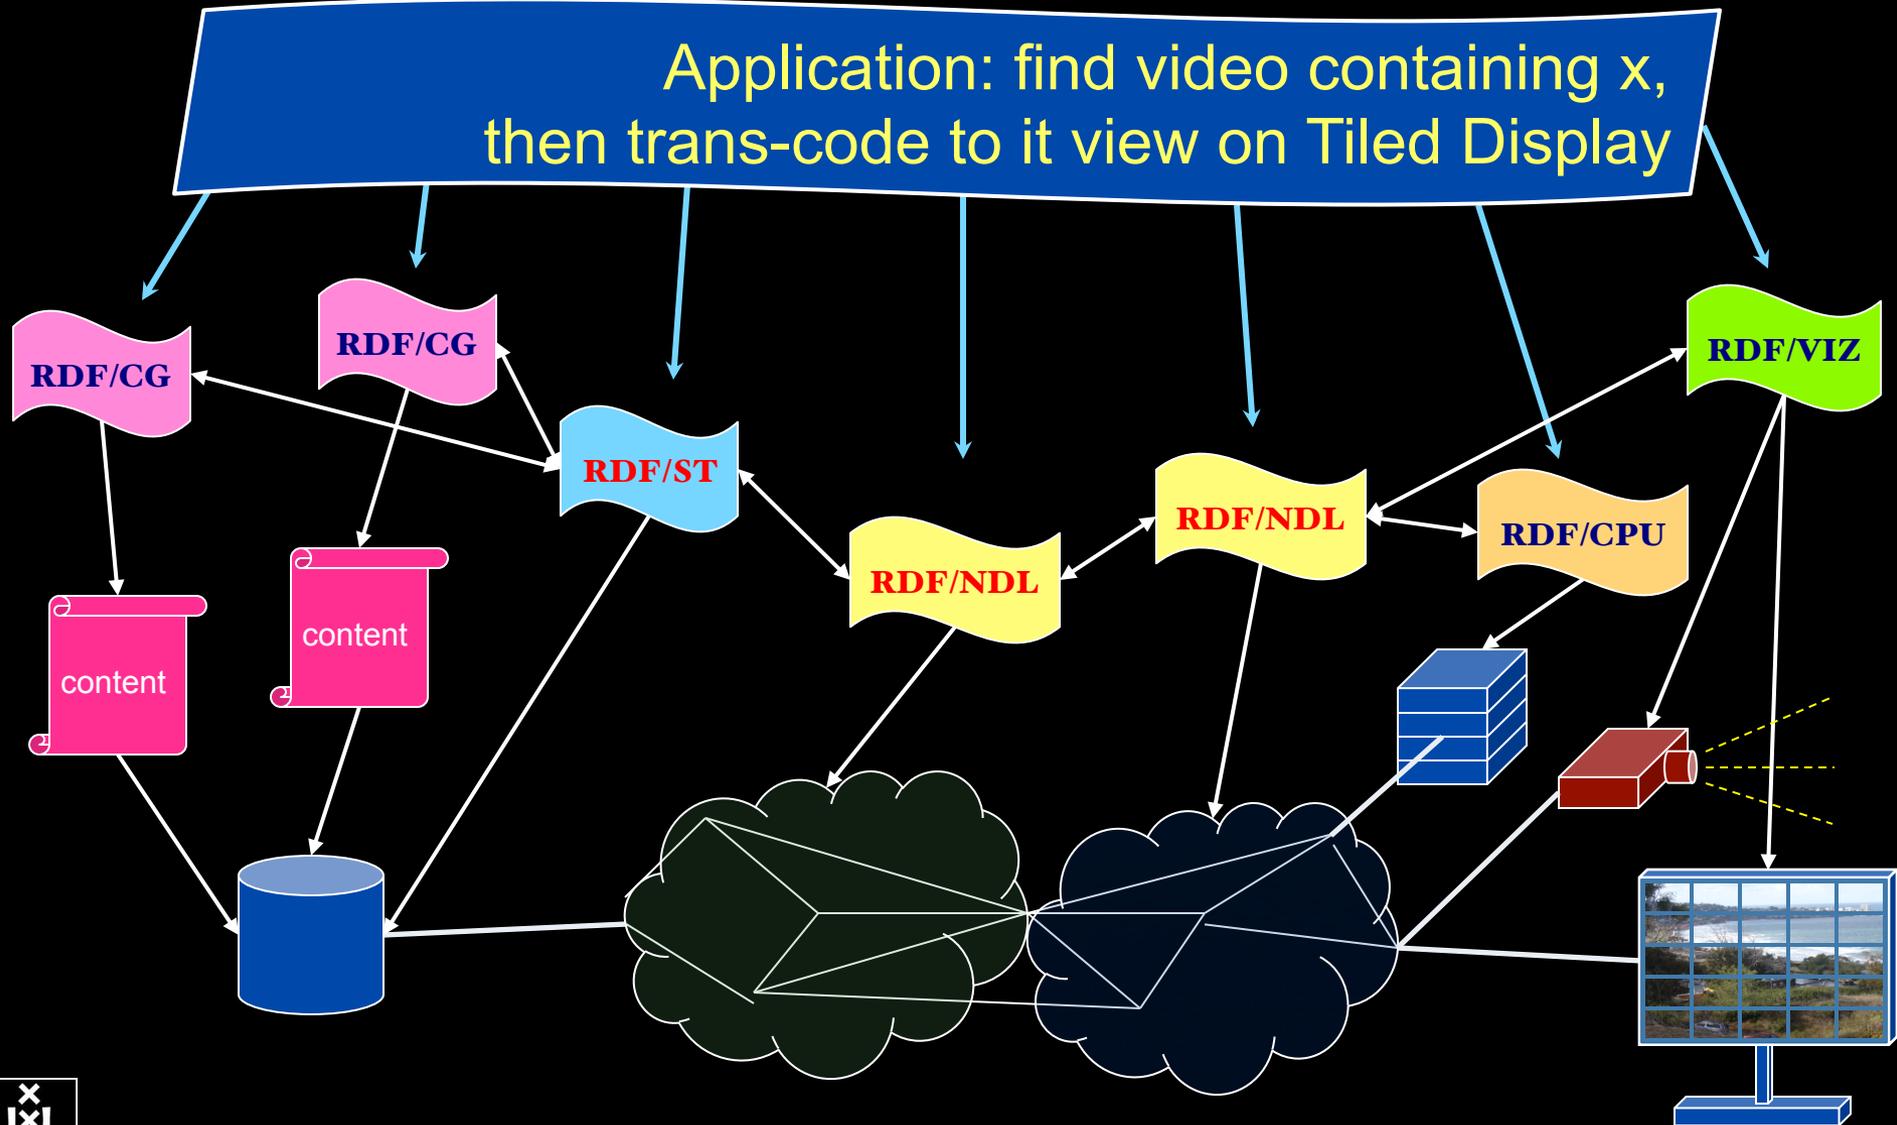
- From semantic Web / Resource Description Framework.
- The RDF uses XML as an interchange syntax.
- Data is described by triplets (Friend of a Friend):



RDF describing Infrastructure

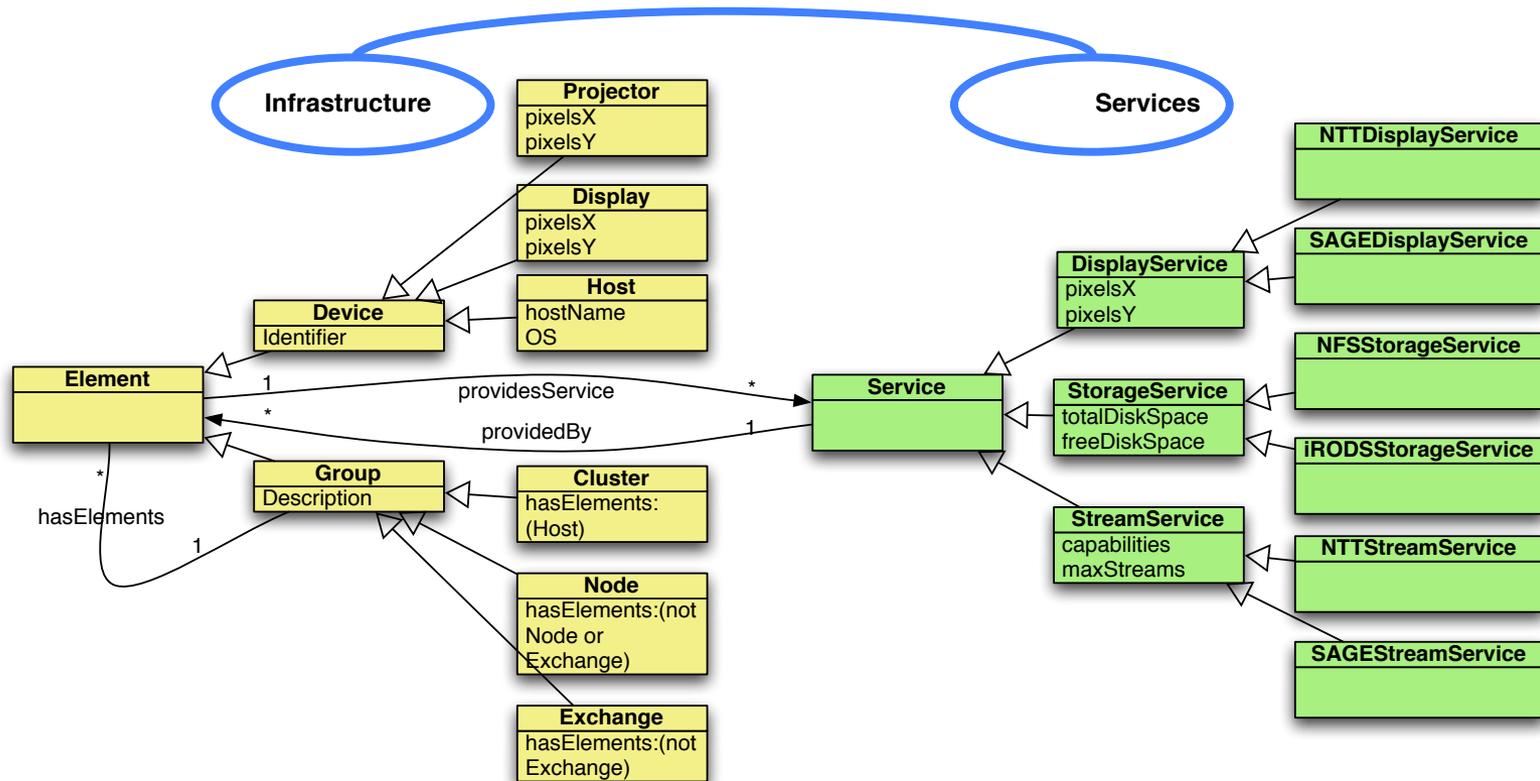
“I want”

Application: find video containing x,
then trans-code to it view on Tiled Display



Information Modeling

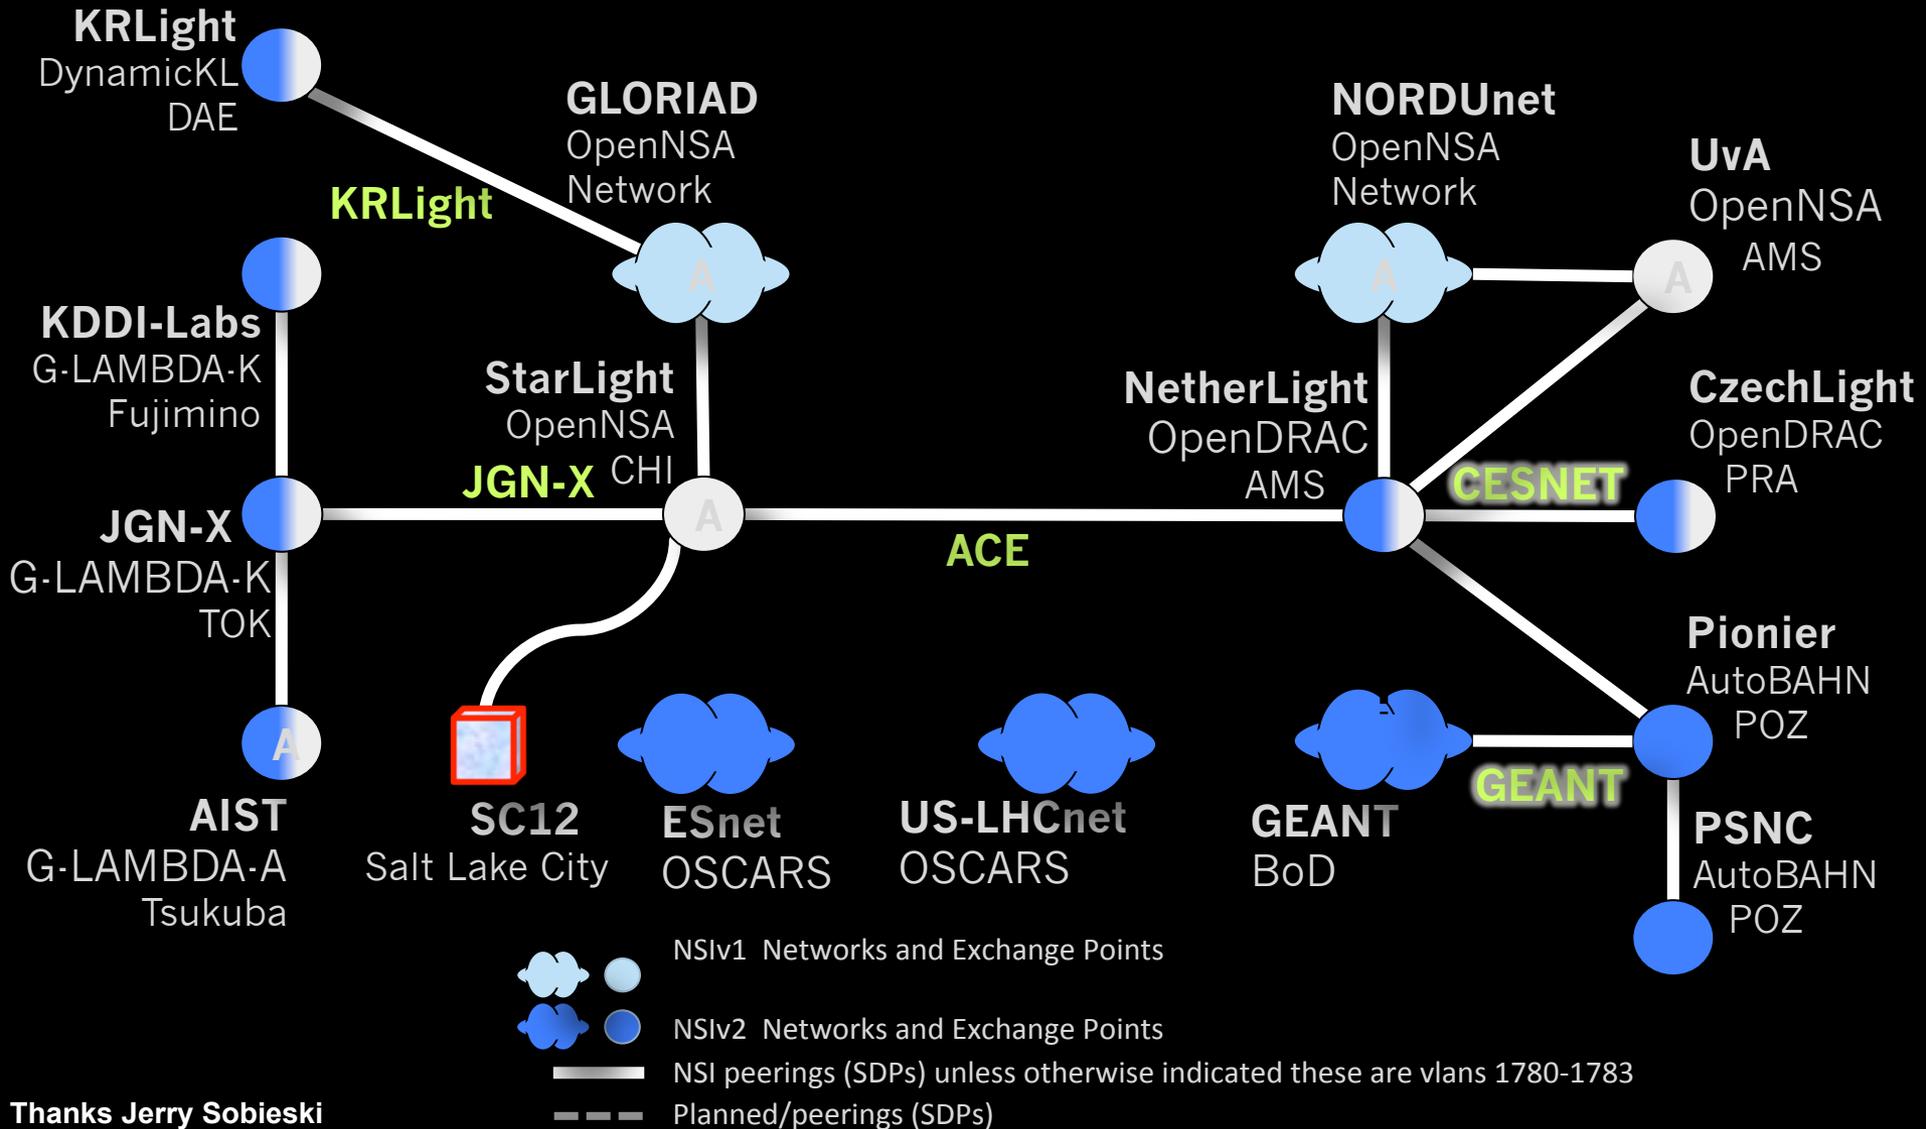
Define a common information model for *infrastructures* and *services*.
Base it on Semantic Web.



Automated GOLE + NSI

Joint NSI v1+v2 Beta Test Fabric Nov 2012

Ethernet Transport Service

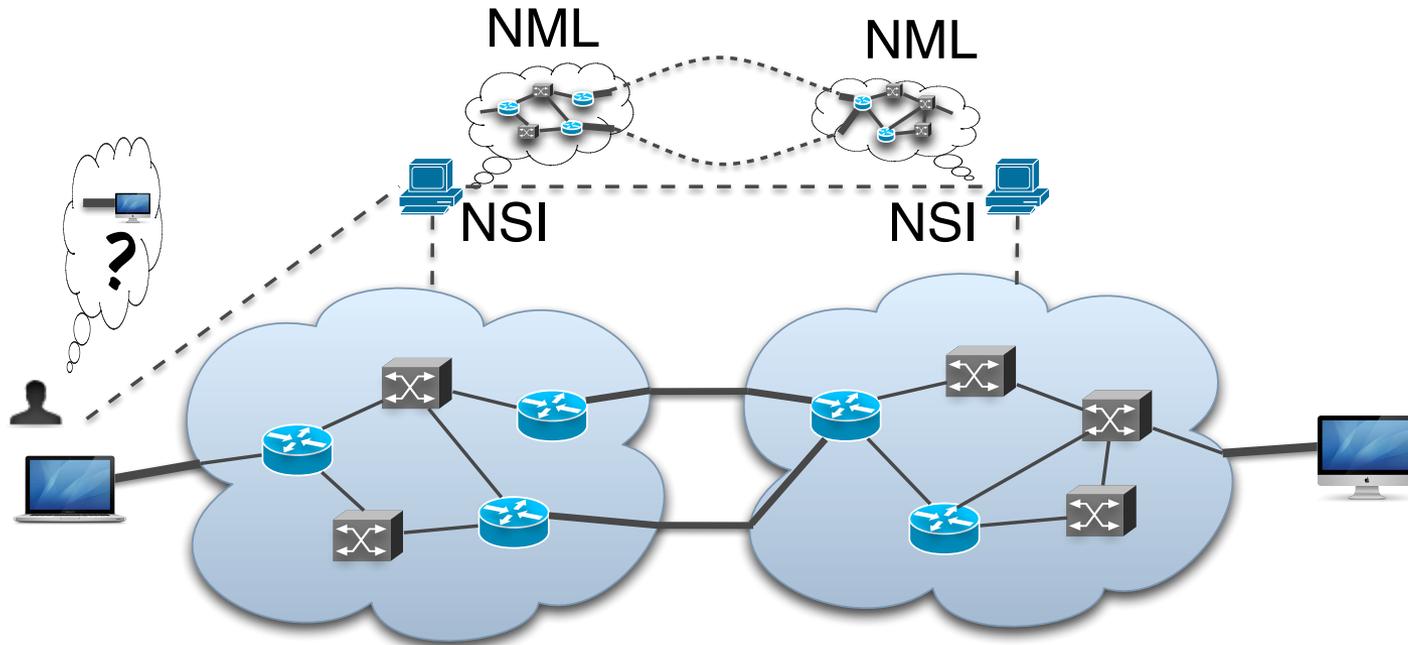


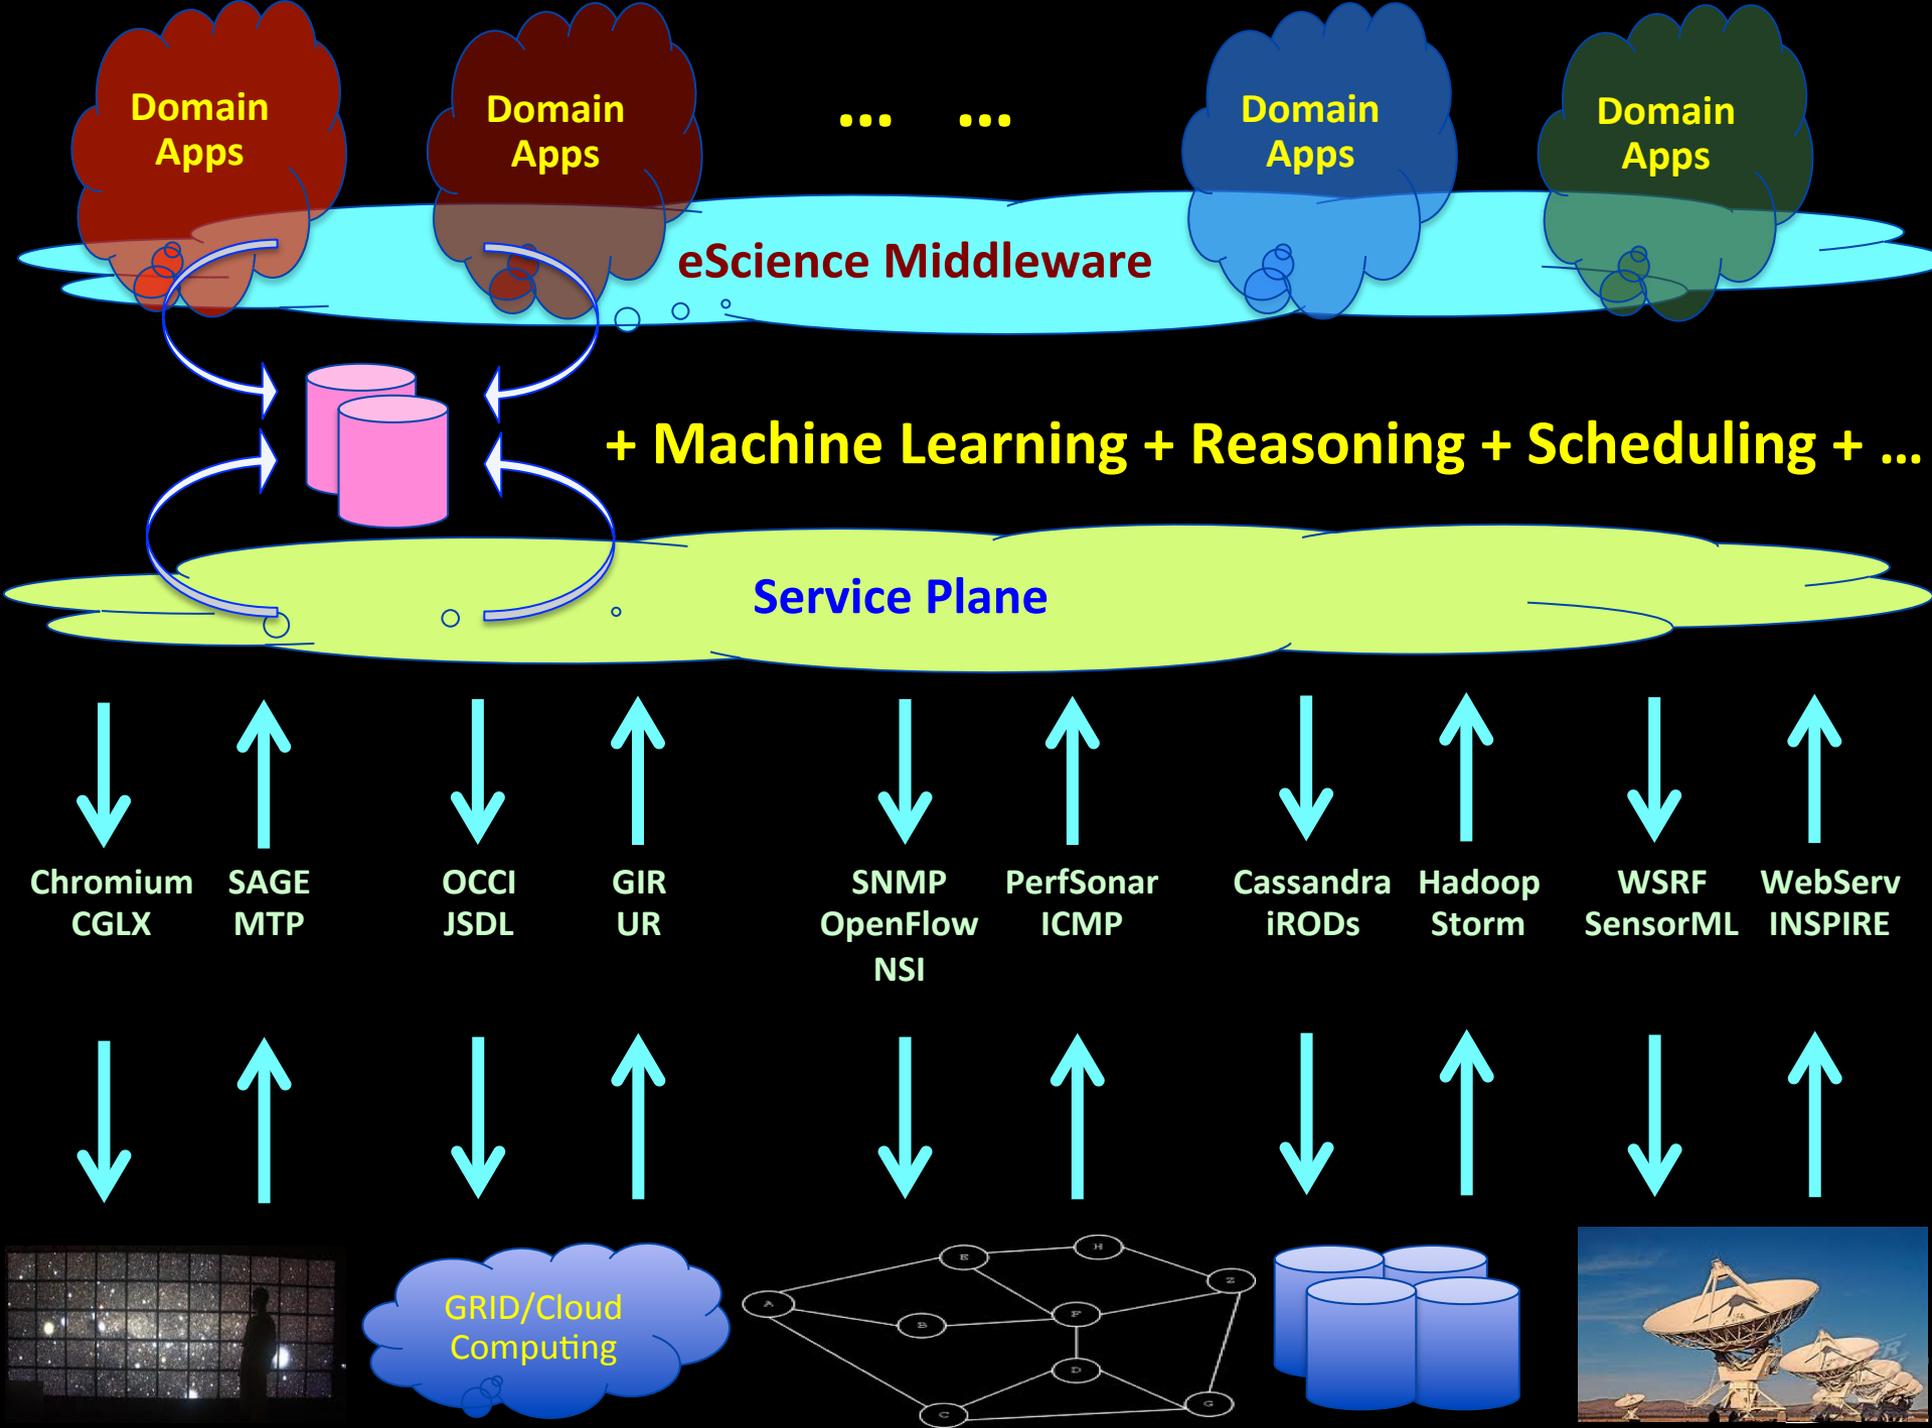
Thanks Jerry Sobieski

Network Topology Description

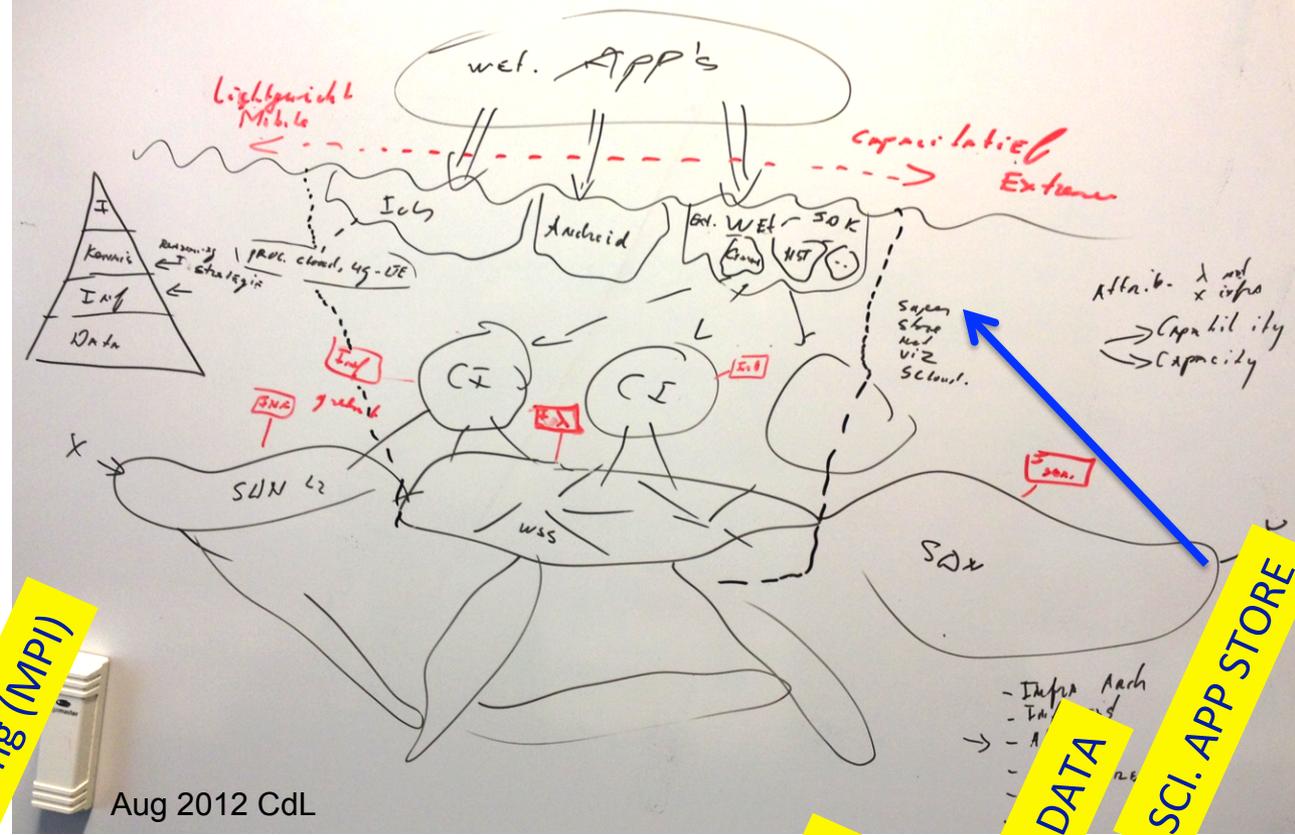
Network topology research supporting automatic network provisioning

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





TimeLine



Aug 2012 CdL

Remote Procedure Call

Distributed Computing (MPI)

GRID

CLOUD

BIG DATA

SCI. APP STORE

1980 1990 2000 2005 2013

TimeLine

-  we started this
-  we strongly participated
-  we use

 GreenIT&Nets

 SF for Clouds

 NDL SF for complex nets

 Programmable Networks  NetApp's

 CineGrid  SF for CineGrid

 NM  OCCI  NSI

 LightPaths - GLIF  Hybrid Nets

 RDUDP, SCTCP, ... 

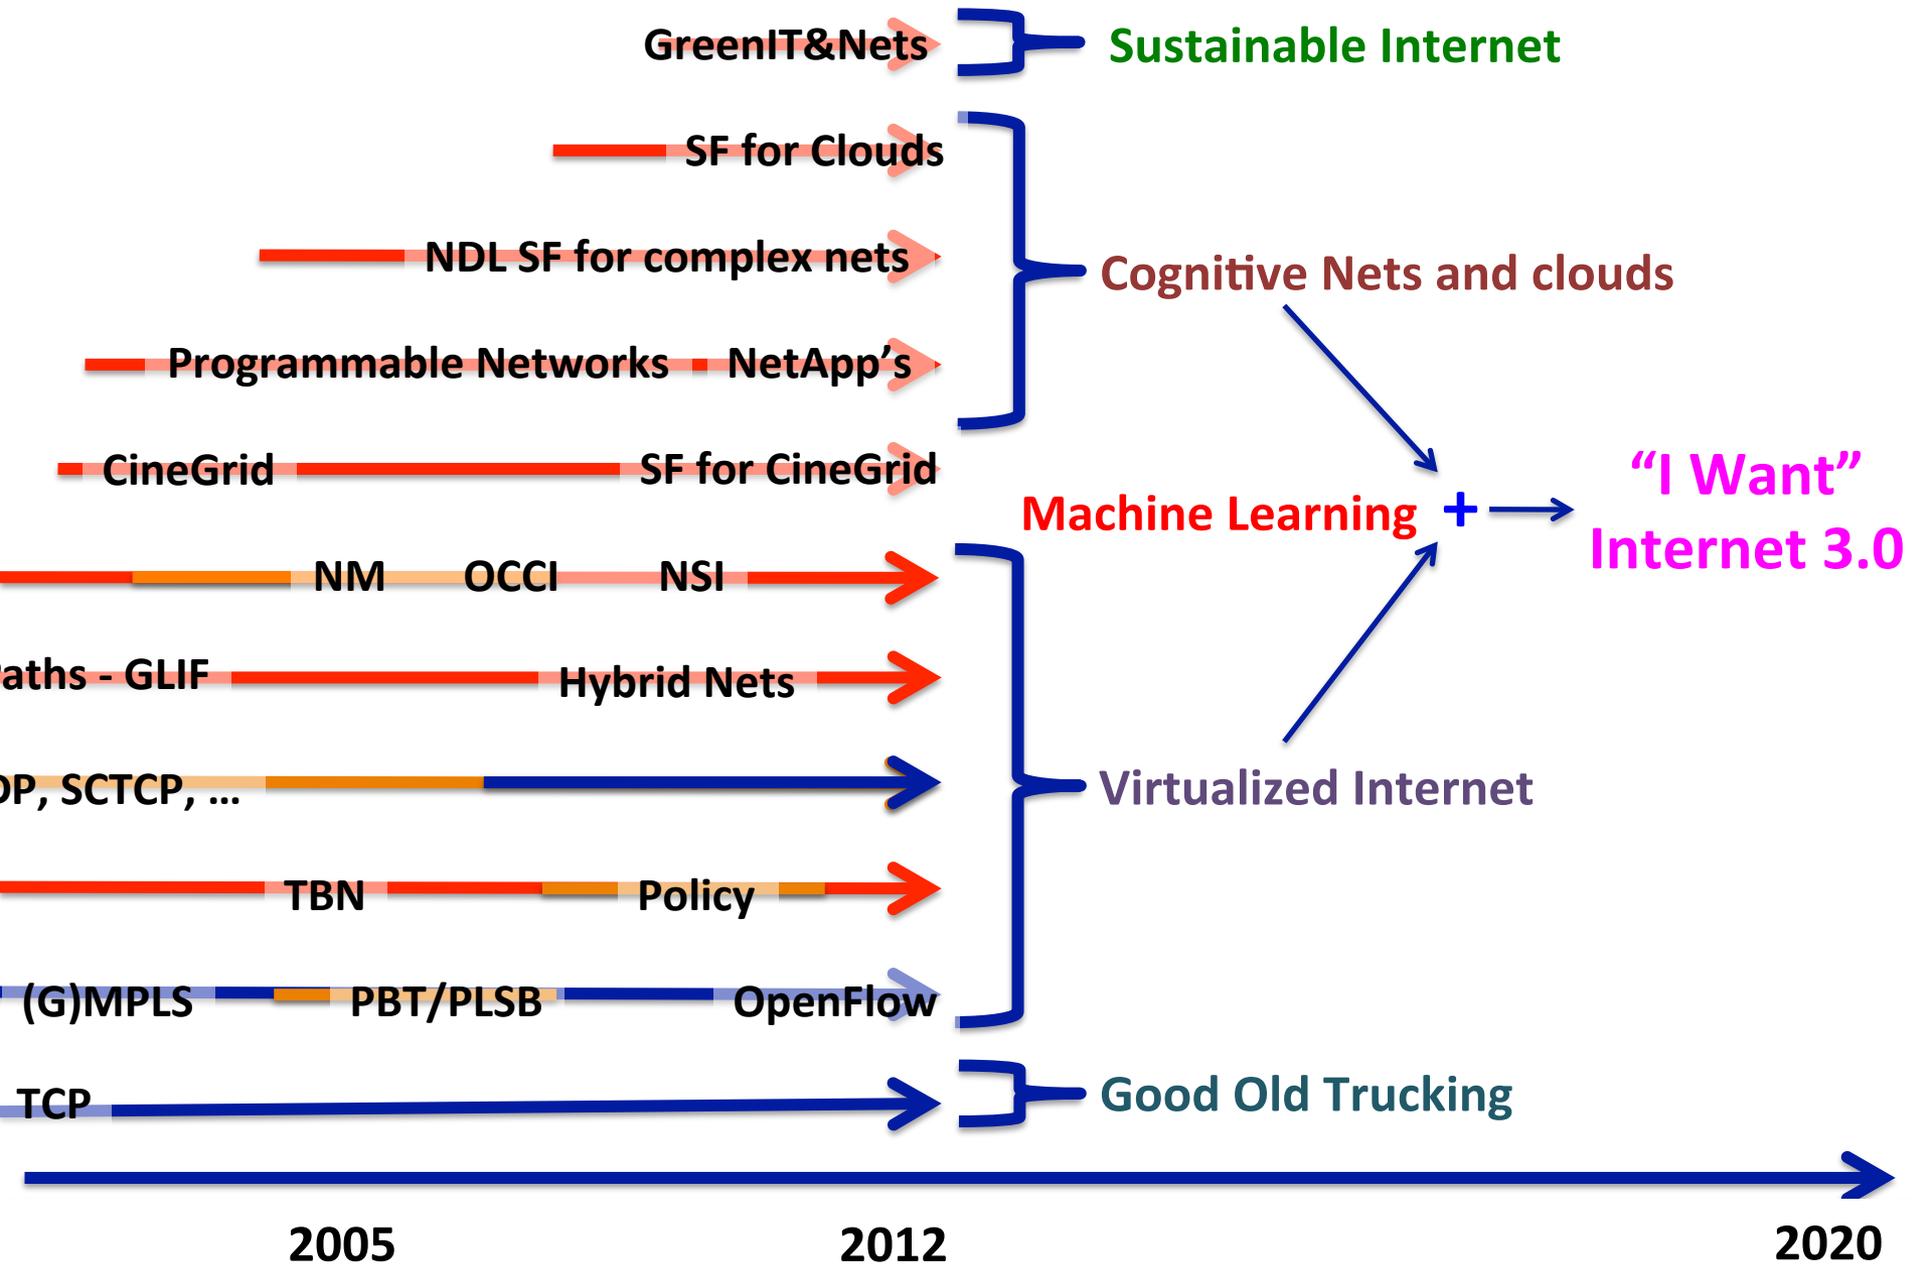
 AAA  TBN  Policy

 ATM  SONET/SDH  (G)MPLS  PBT/PLSB  OpenFlow

 TCP  TCP Reno, Vegas

1980 1990 2000 2005 2012

TimeLine



TimeLine

• Sustainable Internet

• Cognitive Nets and clouds

• Machine Learning +

• Virtualized Internet

• Good Old Trucking

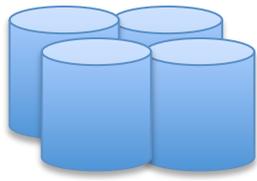
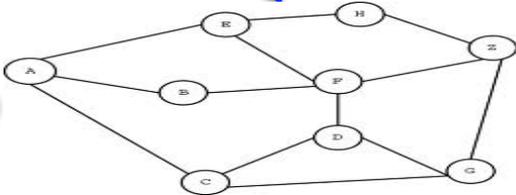
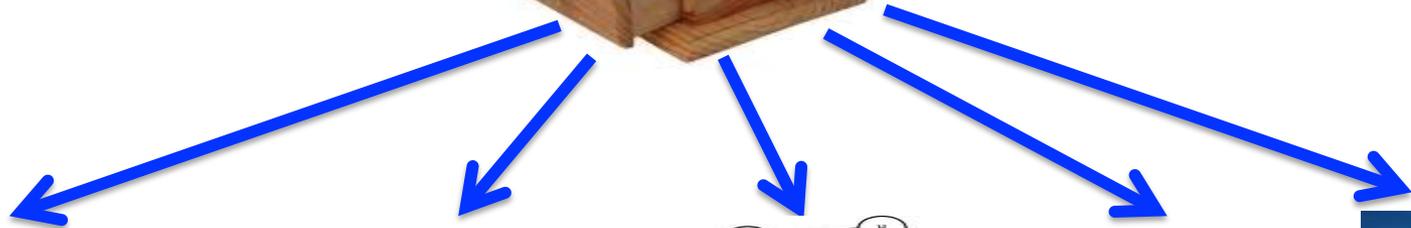
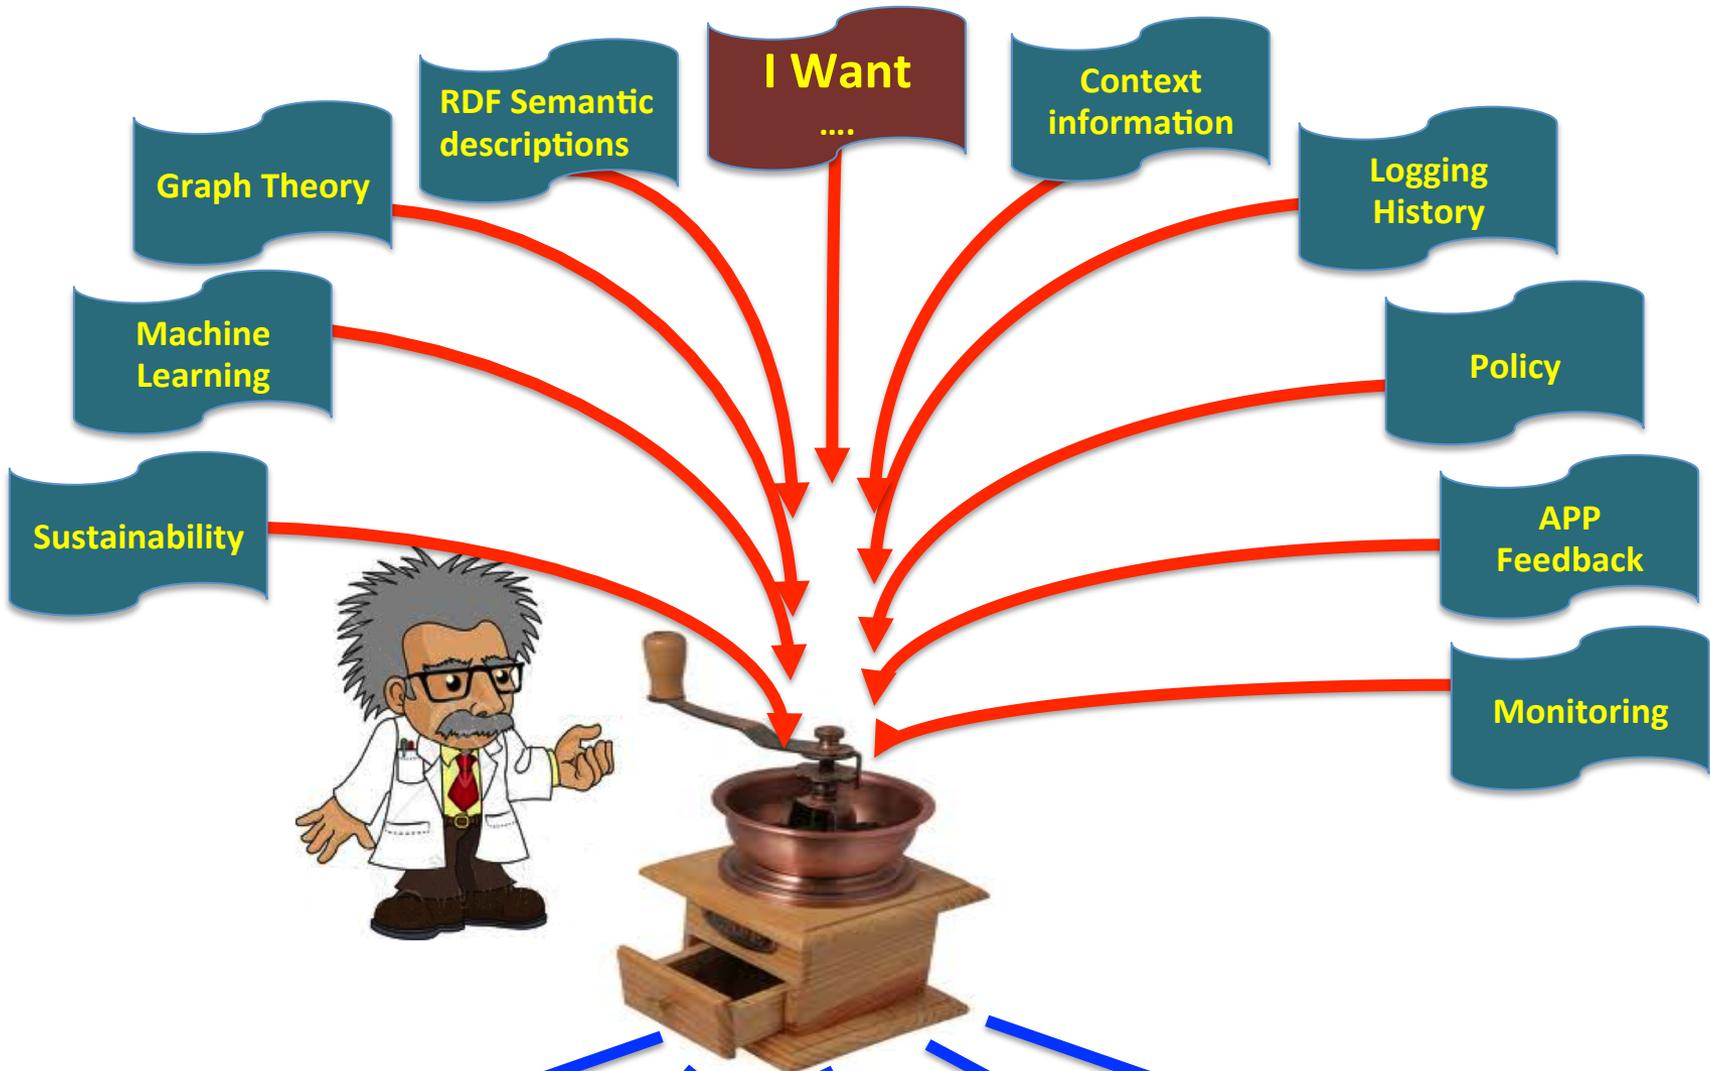
“I Want”
Internet 3.0



I
retire

2020

2040



Conclusion

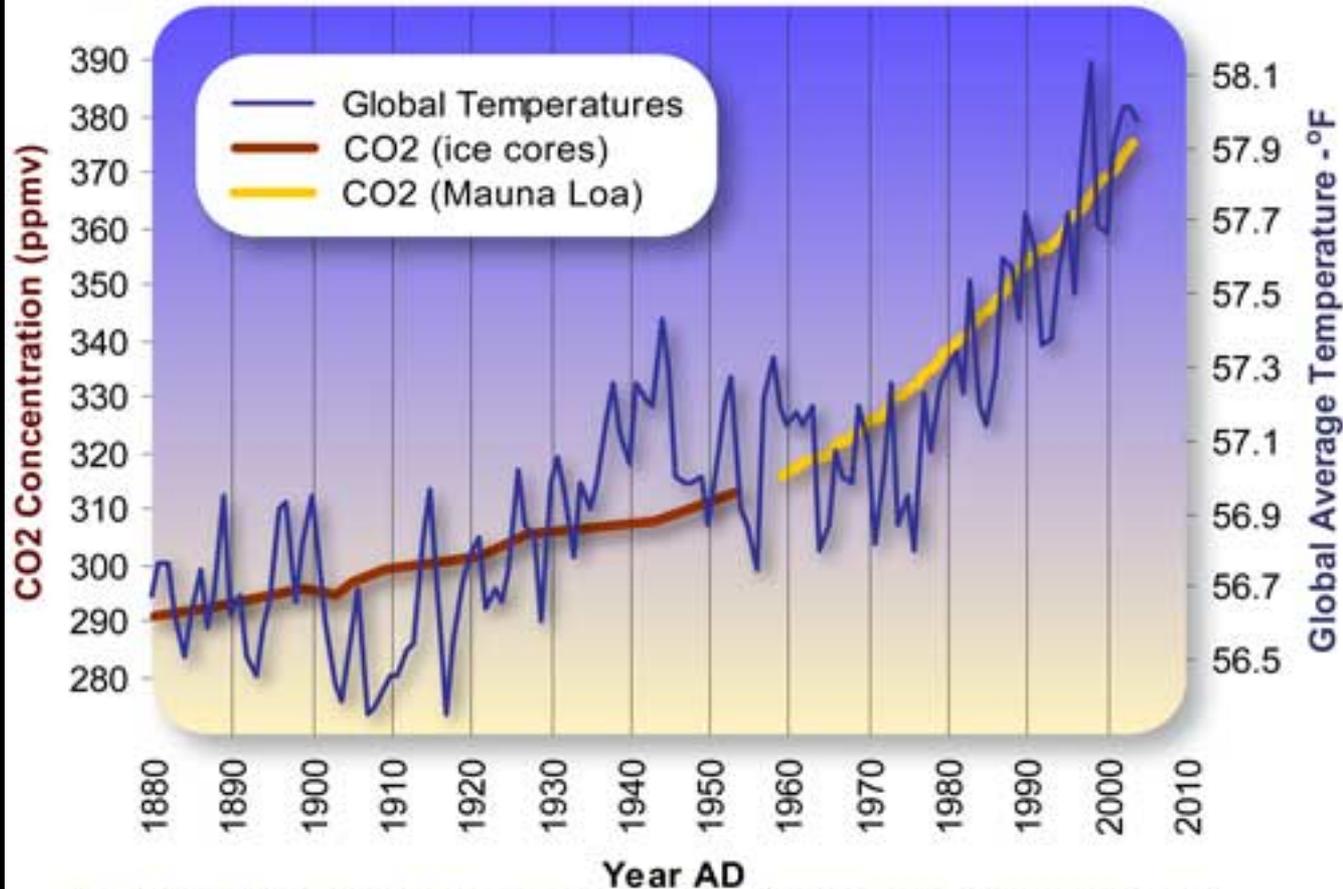
I want a MiS system!

Catchphrase first used in "Encounter At Farpoint" (28 September 1987) by Gene Roddenberry, and thereafter used in many episodes and films, instructing a crew member to execute an order.



Need for GreenIT

Global Average Temperature and Carbon Dioxide Concentrations, 1880 - 2004



Data Source Temperature: ftp://ftp.ncdc.noaa.gov/pub/data/anomalies/annual_land_and_ocean.ts

Data Source CO2 (Siple Ice Cores): <http://cdiac.esd.ornl.gov/ftp/trends/co2/siple2.013>

Data Source CO2 (Mauna Loa): <http://cdiac.esd.ornl.gov/ftp/trends/co2/maunaloa.co2>

Graphic Design: Michael Ernst, The Woods Hole Research Center



Greening the Processing System

Positive proof of global warming.





Turn Green Tech into Greenbacks

IT Certifications for Jobs That Make a Difference



Uptime Institute Accredited Tier Designer

The Uptime Institute has long been a proponent for green data center design and implementation. Its certification course on data center design embeds green principles into the curriculum.



SUSTAINABILITY
Your Career



ECO-Scheduling



Why?



Because we can!



The constant factor in our field is Change!

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

“Fortran goto”, Unix, c, SmallTalk, DECnet, TCP/IP, c++,
Internet, WWW, Semantic Web, Photonic networks, Google,
grid, cloud, Data³, App

to:

DDOS attacks destroying Banks and Bitcoins.

Conclusion:

Need for Safe, Smart, Resilient Sustainable Infrastructure.

Questions?

