

GigaPort-SAC: Trends in Network Research and Engineering.

Cees de Laat

SURFnet

BSIK

EU

University of Amsterdam

System & Network Engineering @ UvA

- group has 5 sections, 24 people
 - Advanced Networking (GP, EU, TNO)
 - Authorization Concepts and Architectures (GP, EU, TI)
 - Security (SurfWorks)
 - Grid Middleware and Work Flow Management (VLE)
 - Sensor Grids - Intelligent networks (TNO)
- Home @ Science Park Amsterdam, co-located with:
 - NIKHEF (with SARA Tier-1 center)
 - SARA (SN6-NOC, LightHouse)
 - AMSIX
 - UvA Science faculty (Dutch e-Science program VL-e)



Optical networking research activities

Several areas of interest with a common vision:

- Network architectures and principles
- Network descriptions
- Path finding algorithms
- Fault detection
- Monitoring
- StarPlane

... research novel ways to describe, manage, monitor, and provision optical hybrid networks and lightpaths.

u
s
e
r
s

A. Lightweight users, browsing, mailing, home use

Need full Internet routing, one to many

B. Business/grid applications, multicast, streaming, VO's, mostly LAN

Need VPN services and full Internet routing, several to several + uplink

C. E-Science applications, distributed data processing, all sorts of grids

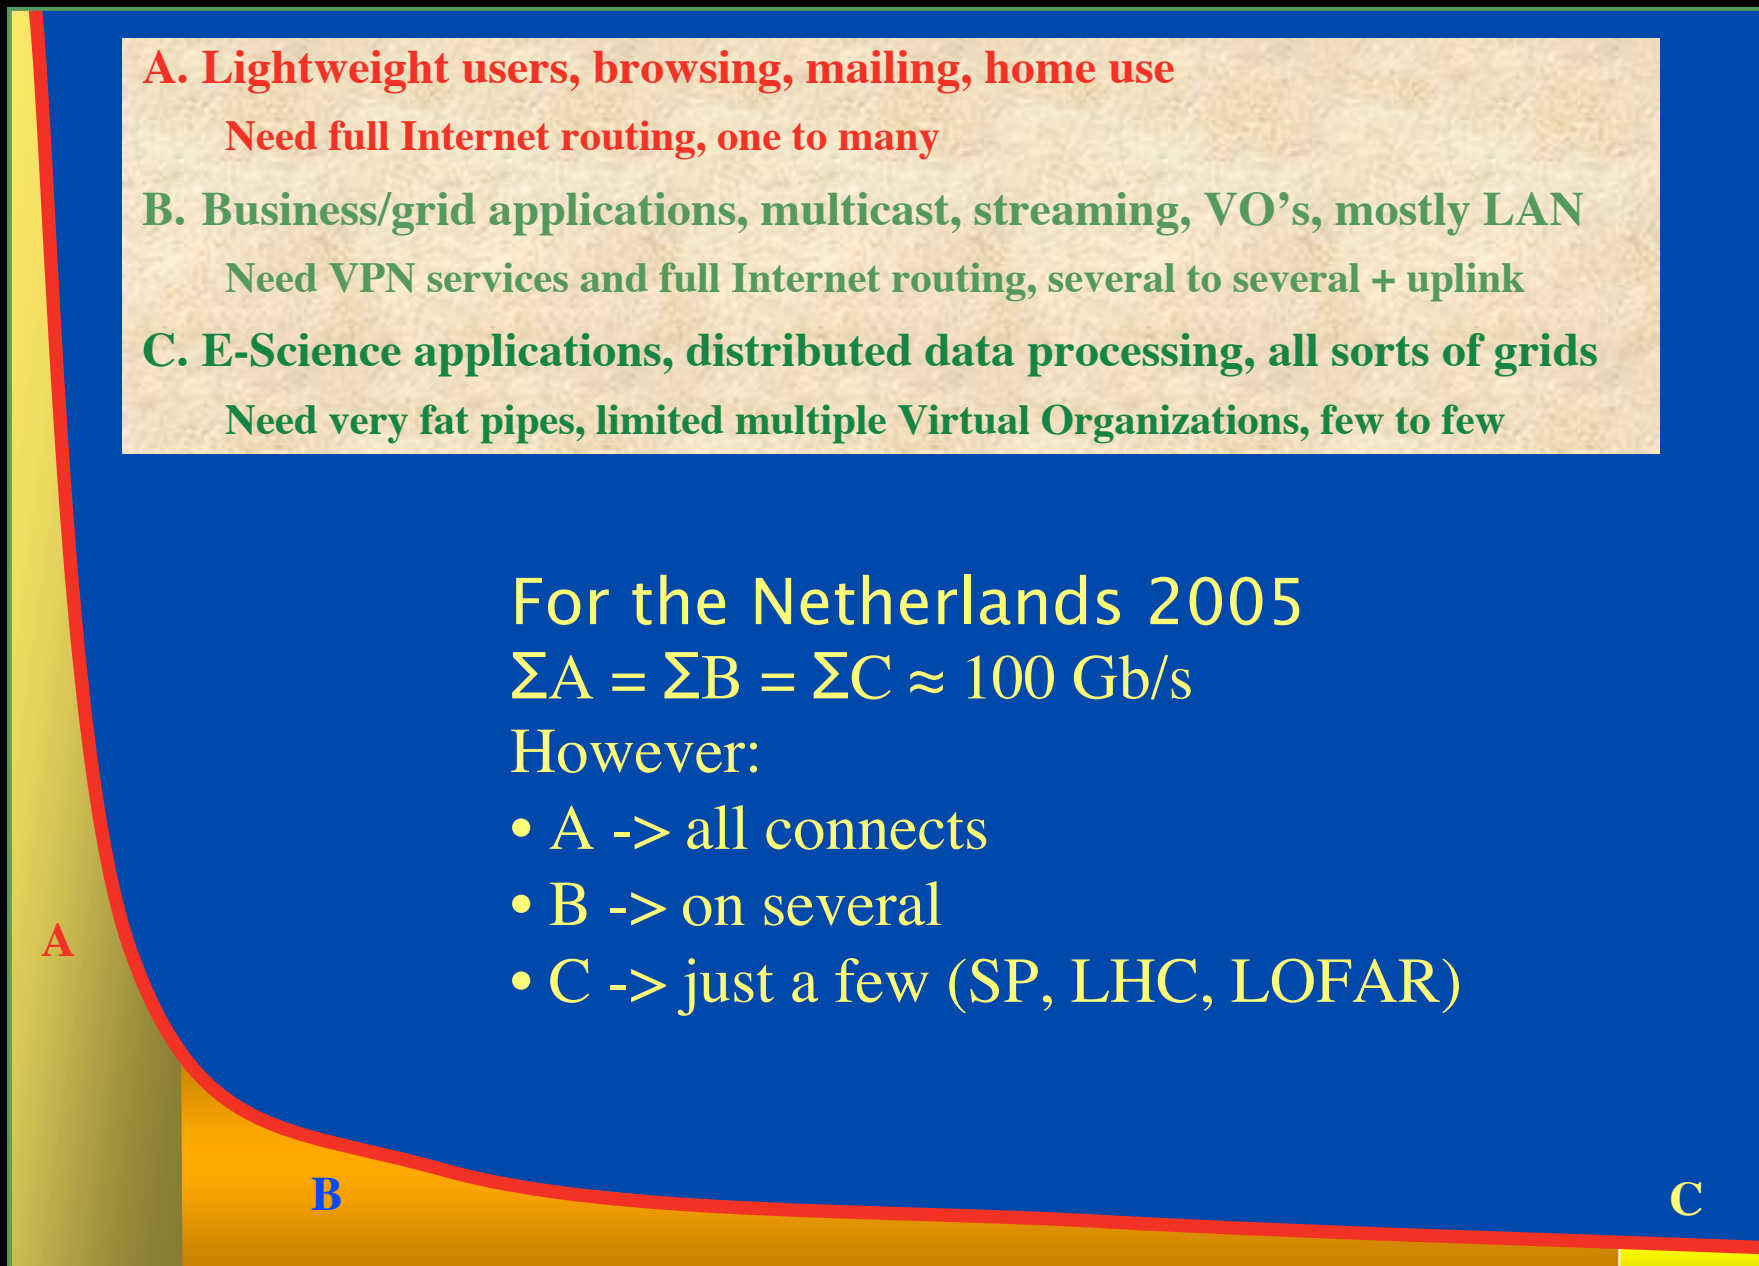
Need very fat pipes, limited multiple Virtual Organizations, few to few

For the Netherlands 2005

$$\Sigma A = \Sigma B = \Sigma C \approx 100 \text{ Gb/s}$$

However:

- A -> all connects
- B -> on several
- C -> just a few (SP, LHC, LOFAR)




ADSL (10 Mbit/s)

GigE

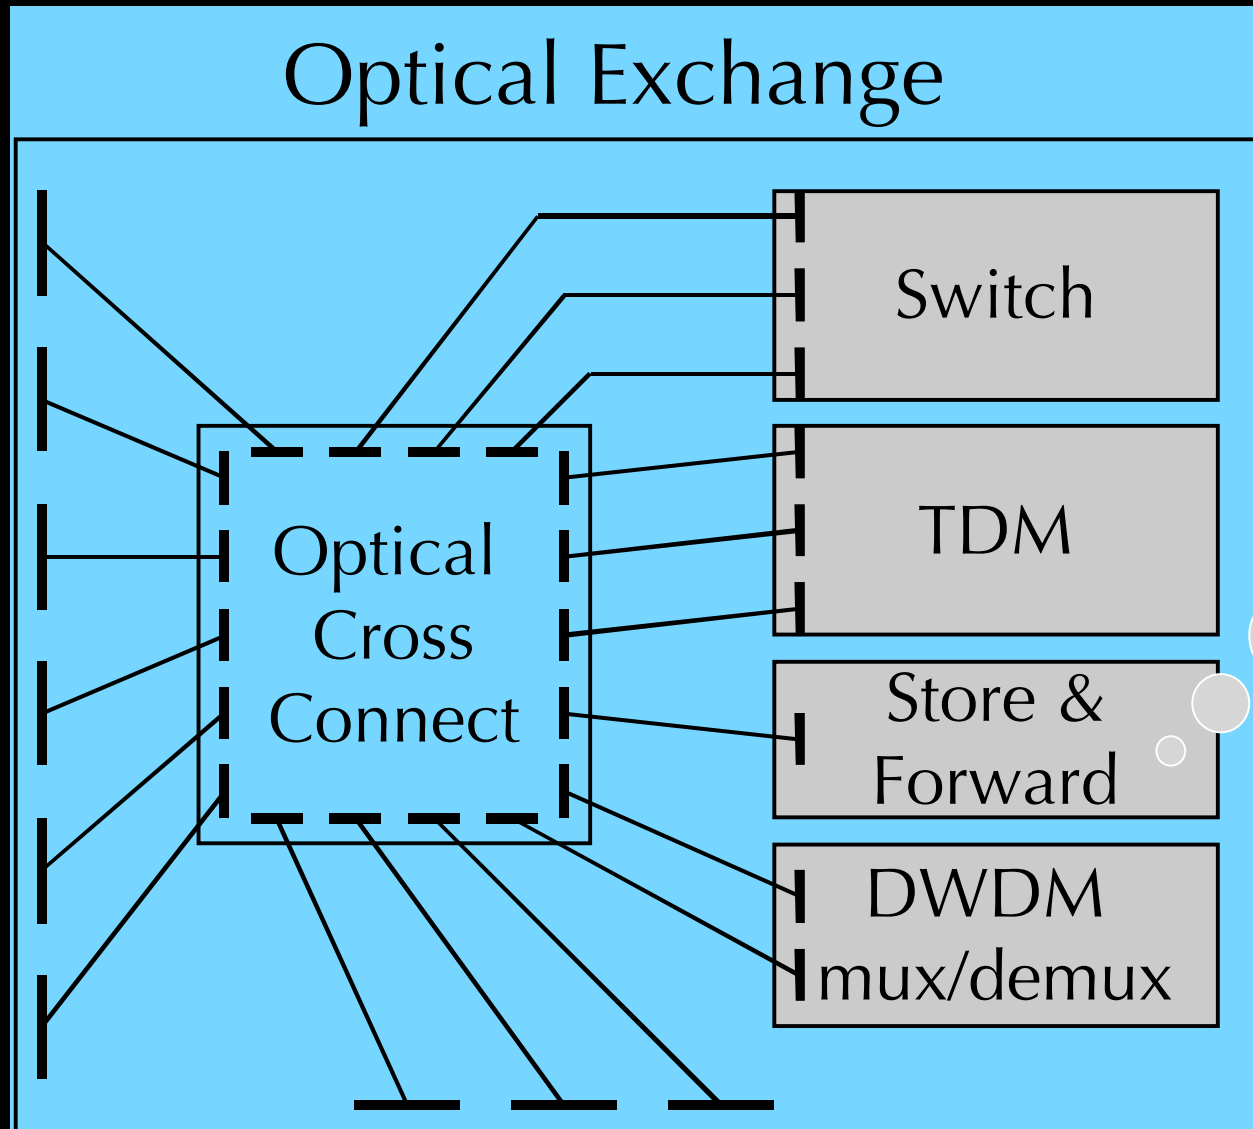
BW requirements



How low can you go?



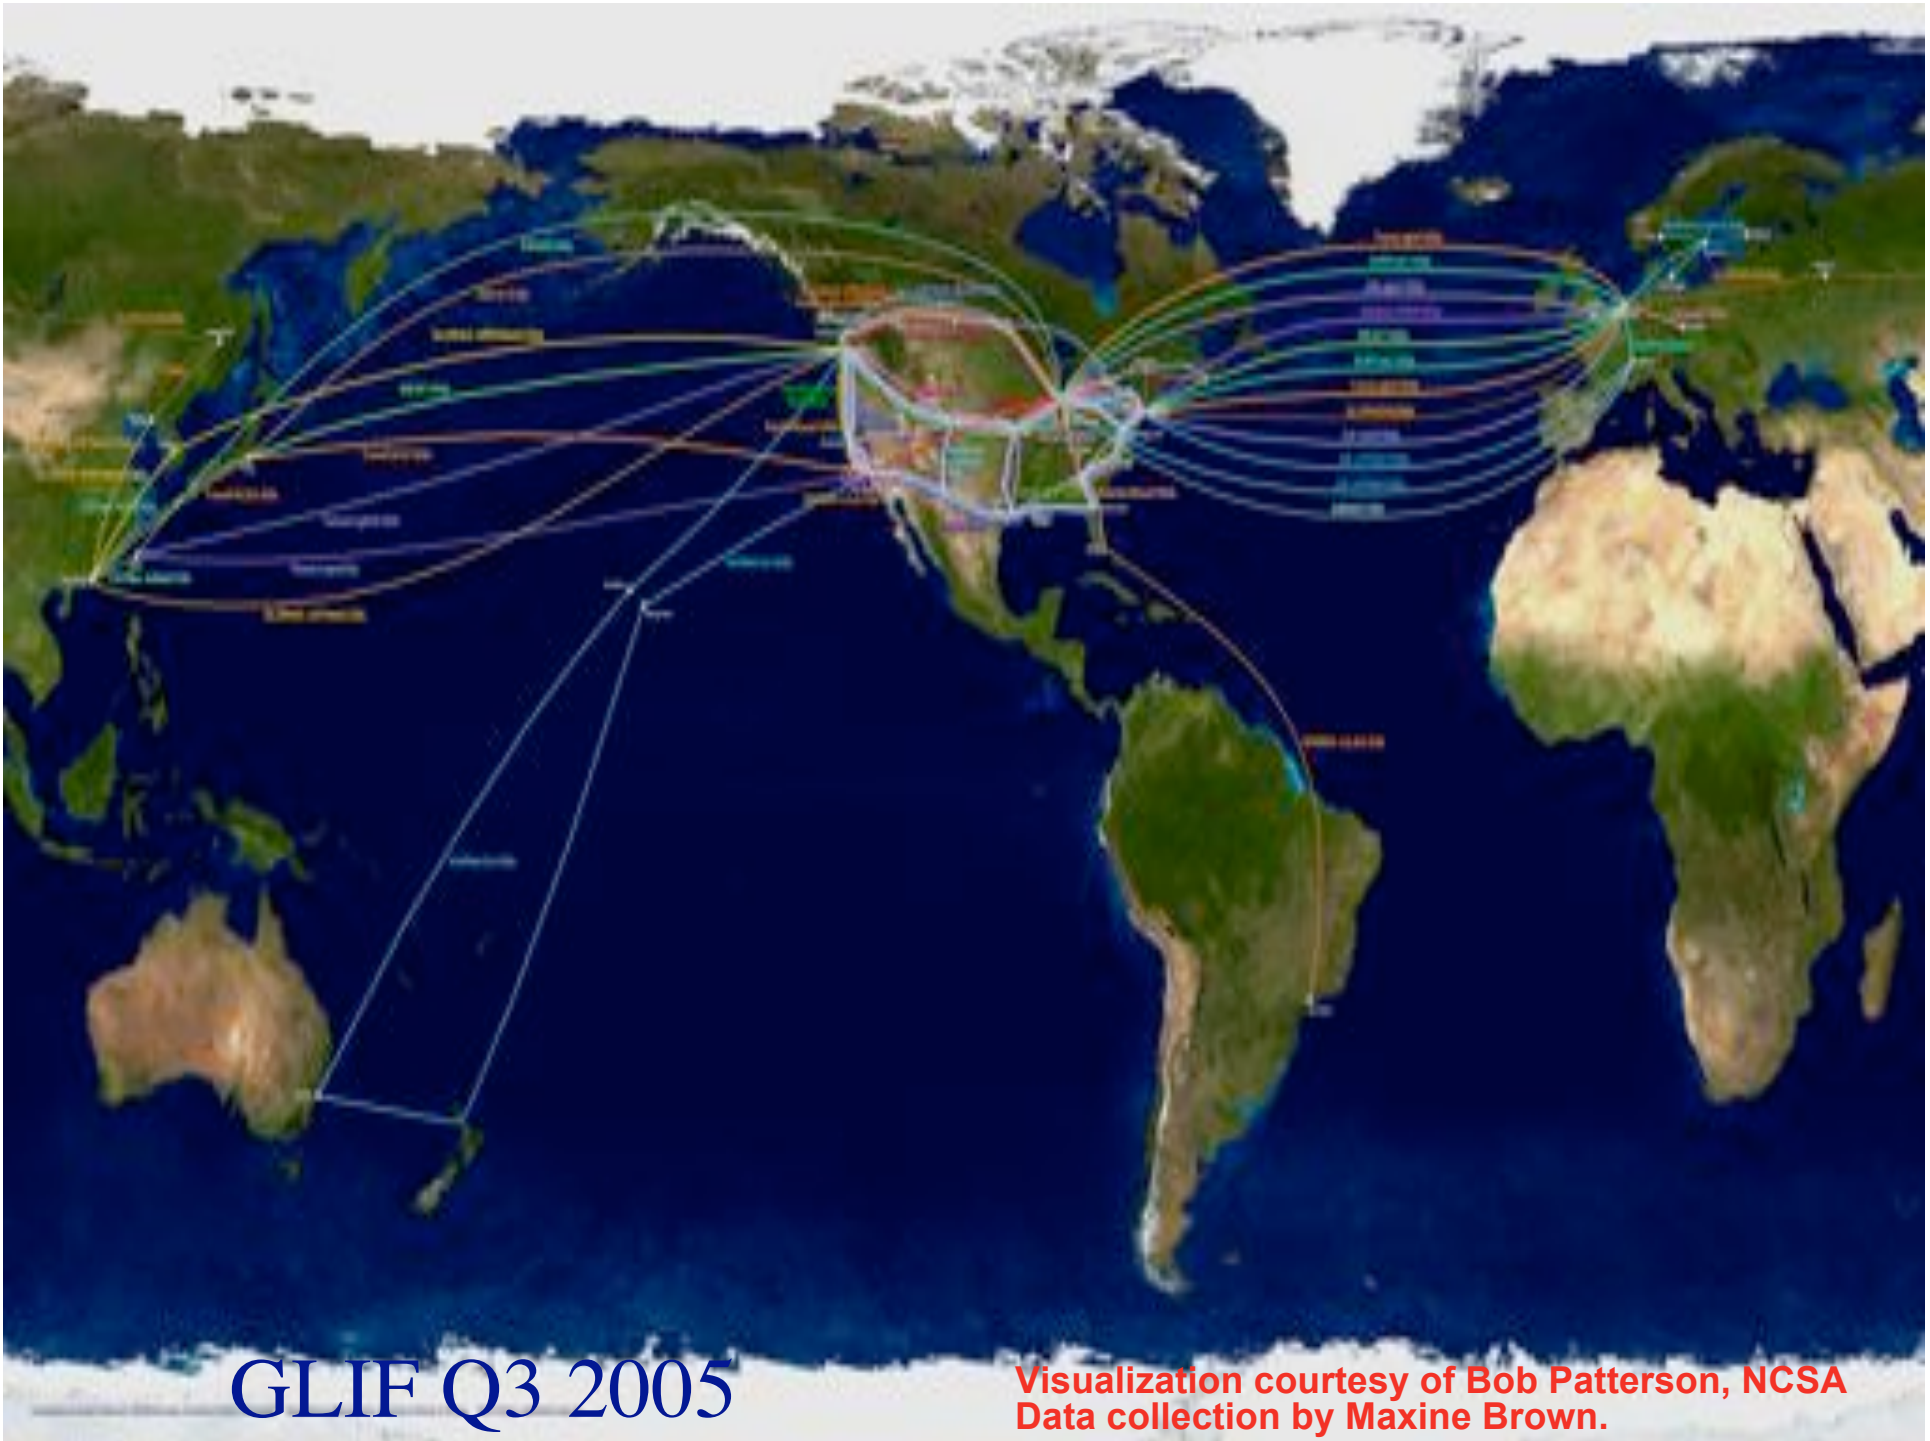
Optical Exchange as Black Box



TeraByte
Email
Service

Ref: gridnets paper by Freek Dijkstra, Cees de Laat






GLIF Q3 2005


Visualization courtesy of Bob Patterson, NCSA
Data collection by Maxine Brown.

Network Description Language

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



NDL Example



```
<ndl:Device rdf:about="#Rembrandt3">  
  <ndl:name>Rembrandt3</ndl:name>  
  <ndl:locatedAt rdf:resource="#Lighthouse"/>  
  <ndl:hasInterface rdf:resource="#Rembrandt3:eth0"/>  
</ndl:Device>  
<ndl:Interface rdf:about="#Rembrandt3:eth0">  
  <ndl:name>Rembrandt3:eth0</ndl:name>  
  <ndl:connectedTo rdf:resource="#Speculaas:port3"/>  
</ndl:Interface>
```

NDL Generator and Validator

NDL for the GLIF - NDL Validator

NDL - Network Description Language - is an ontology for description of (hybrid) networks, air provisioning. The GLIF collaboration makes use of NDL to describe each individual domain, maps.

This page will provide you with tools to validate an NDL file. We provide here two types of validation:

- Syntax validation
- Content validation

Syntax validation

We can validate that the NDL file you generated is written following the latest NDL schema. You will get back feedback on its validity.

Please paste your NDL file below:

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdflib="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:ndl="http://www.science.uva.nl/research/son/ndl#"
  xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#"
>
  <!-- Description of foo -->
  <ndl:Location rdf:about="#foo">
    <ndl:name>bar</ndl:name>
    <geo:lat>0</geo:lat>
    <geo:long>0</geo:long>
  </ndl:Location>

  <!-- Rem2 -->
  <ndl:Device rdf:about="#Rem2">
    <ndl:name>Rem2</ndl:name>
    <ndl:locatedAt rdf:resource="#foo"/>
    <ndl:hasInterface rdf:resource="#Rem2:eth0"/>
  </ndl:Device>
</rdf:RDF>
```

Submit

Content validation

Often NDL files reference information contained in other files managed by others. Such as for example when an interface on a local device connects to an interface to a remote device. The content validator performs a few basic checks to see that the information contained in cross-referencing NDL files is consistent.

Please enter the URL of the NDL file to be validated:

Submit

Step 1 - Location

Indicate the name and a short description of the network that is going to be described in NDL.

Name Description

Provide also the latitude and the longitude of this location: this will aid the visualization programs.

Both latitude and longitude should use floating point notation.

Latitude Longitude

Step 2 - Devices

Indicate the name of all the devices present in the network. If you need to describe more than 3 devices just "Add a Device"

Device

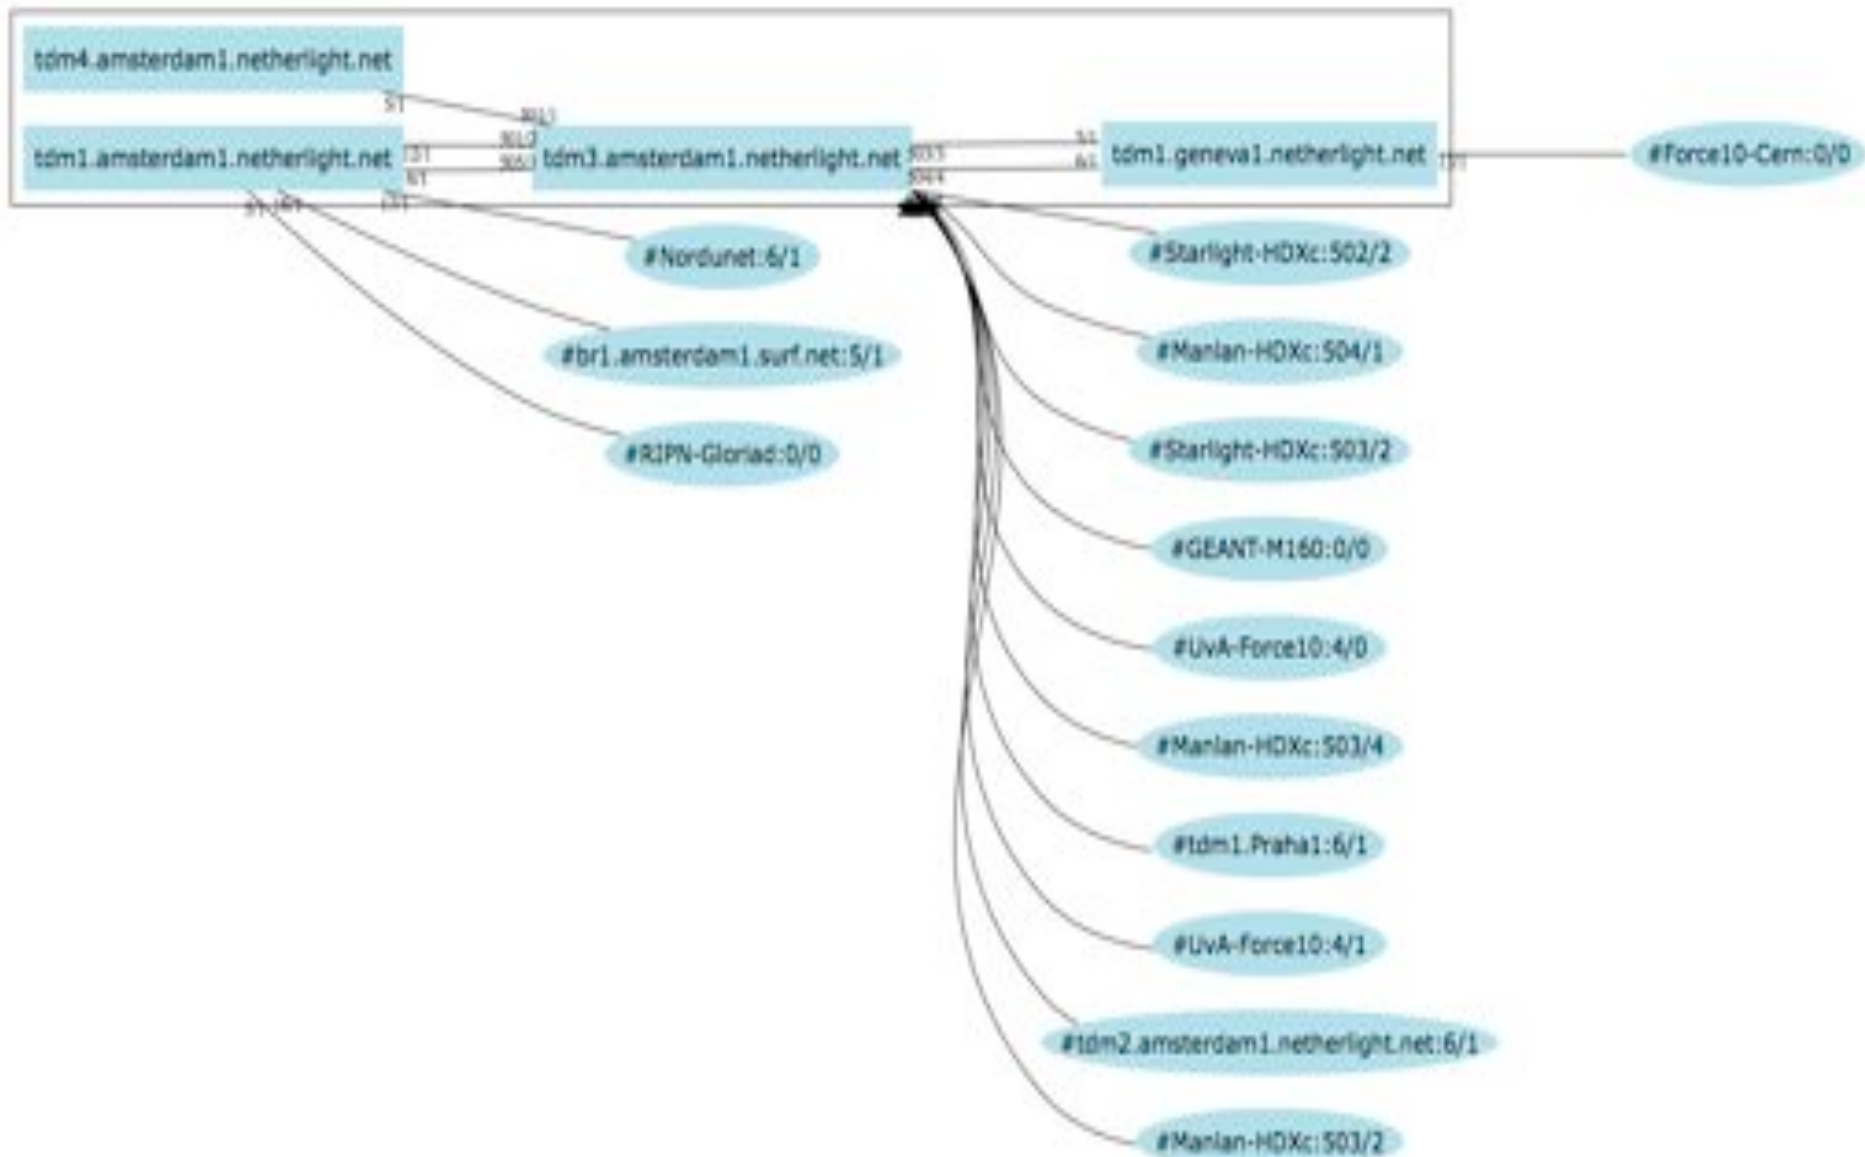
Device

Device

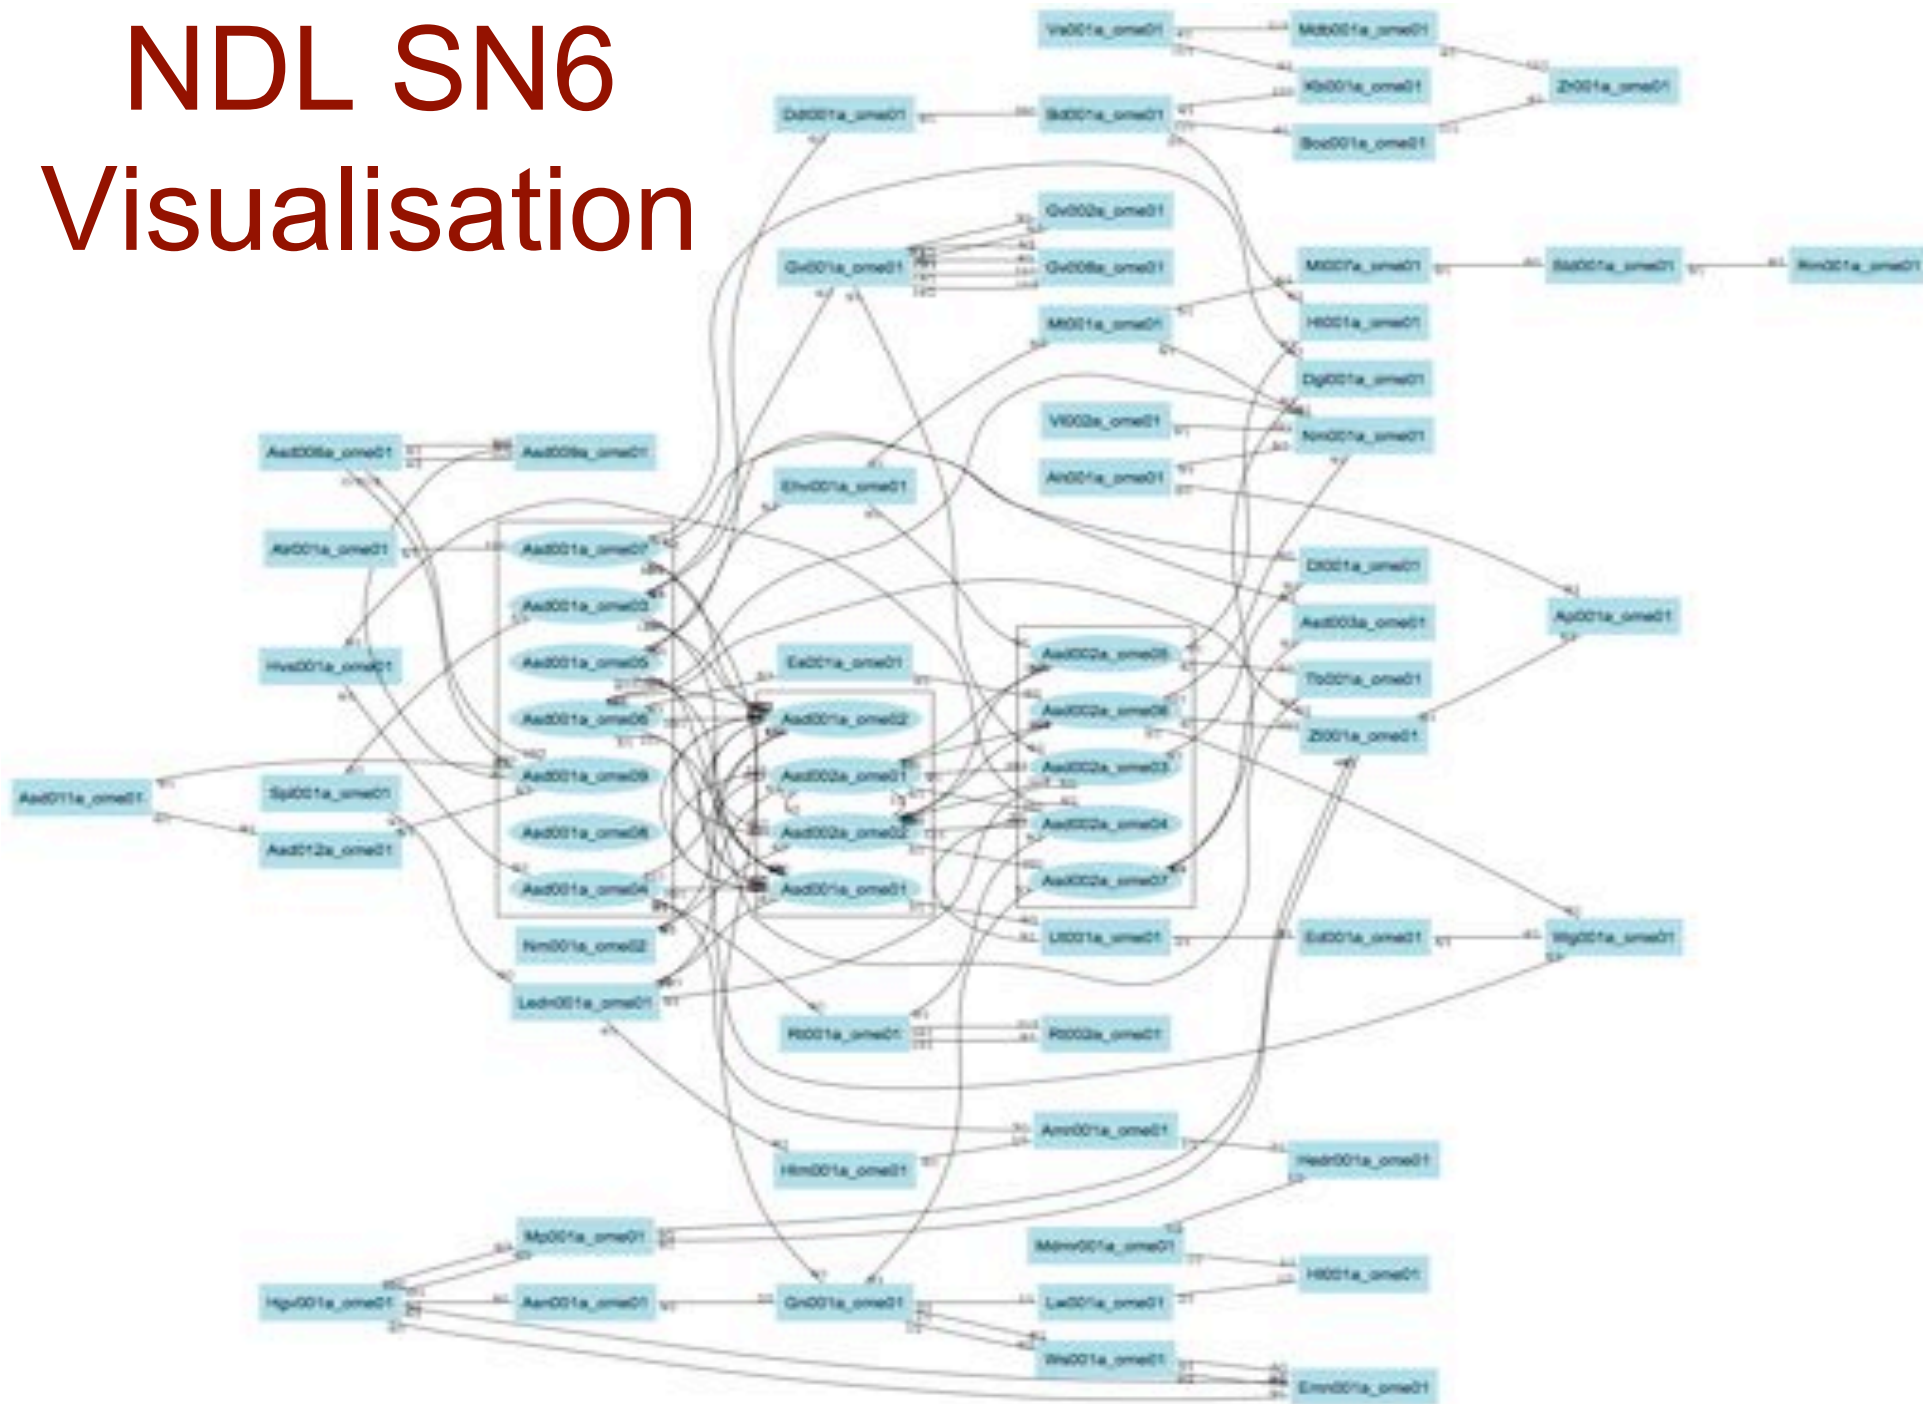
Add a Device

see <http://trafficlight.uva.netherlight.nl/NDL-demo/>

NDL Visualisation



NDL SN6 Visualisation



GOLE Descriptions

```
<rdfl:Interface rdf:about="#netherlight:if5">
  <rdfl:name>netherlight:if5</rdfl:name>
  <rdfl:connectedTo
rdf:resource="http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight
/starlight.rdf#starlight:if3"/>
</rdfl:Interface>

<!-- http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight/starlight
.rdf#starlight:if3 -->
<rdfl:Interface
rdf:about="http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight/st
arlight.rdf#starlight:if3">
  <rdfls:seeAlso
rdf:resource="http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight
/starlight.rdf"/>
</rdfl:Interface>

<!-- netherlight:if6 -->
<rdfl:Interface rdf:about="#netherlight:if6">
  <rdfl:name>netherlight:if6</rdfl:name>
  <rdfl:connectedTo
rdf:resource="http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight
/starlight.rdf#starlight:if4"/>
</rdfl:Interface>

<!-- http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight/starlight
.rdf#starlight:if4 -->
<rdfl:Interface
rdf:about="http://trafficlight.uva.netherlight.nl/JointDemo/GOLEs/starlight/st
arlight.rdf#starlight:if4">
```


Current status: NDL

NDL - **Network Description Language** - an RDF based model for hybrid network descriptions.

It leverages all the semantic web tools, to provide:

- parsing of the RDF files
- graphs and visualization of connections and lightpaths
- lightpath provisioning support at inter and intra domain level.


Latest developments were presented at the GLIF meeting in Sep. '06.



Google map and NDL...



...the GLIF connections described by NDL.

StarPlane DWDM backplane



DAS-3 Cluster Tender


http://www.clustervision.com/pr_das3_uk.html






QOS in a non destructive way!




- Destructive QOS:
 - have a link or λ
 - set part of it aside for a lucky few under higher priority
 - rest gets less service




- Constructive QOS:
 - have a λ
 - add other λ 's as needed on separate colors
 - move the lucky ones over there
 - rest gets also a bit happier!



Traffic engineering



GRID-Colocation problem space



What makes StarPlane fly?

- Wavelength Selective Switches
 - for the “low cost” photonics
- Sandbox by confining StarPlane to one band
 - for experimenting on a production network
- Optimization of the controls to turn on/off a Lambda
 - direct access to part of the controls at the NOC
- electronic Dynamically Compensating Optics (eDCO)
 - to compensate for changing lengths of the path
- traffic engineering
 - to create the OPN topologies needed by the applications
- Open Source GMPLS
 - to facilitate policy enabled cross domain signalling






No Change
Minimum Credit
Billing \$3
For questions, comments, or info
(408) 484-7665...
Office Hours: 9:00 AM -

SURFNET PREMIERE

HELP

net

Three Easy Steps :



Click the **START** button



Insert money...

\$0.25 per minute...

Example :

\$1 = 4 minutes

\$5 = 20 minutes

No change is provided!



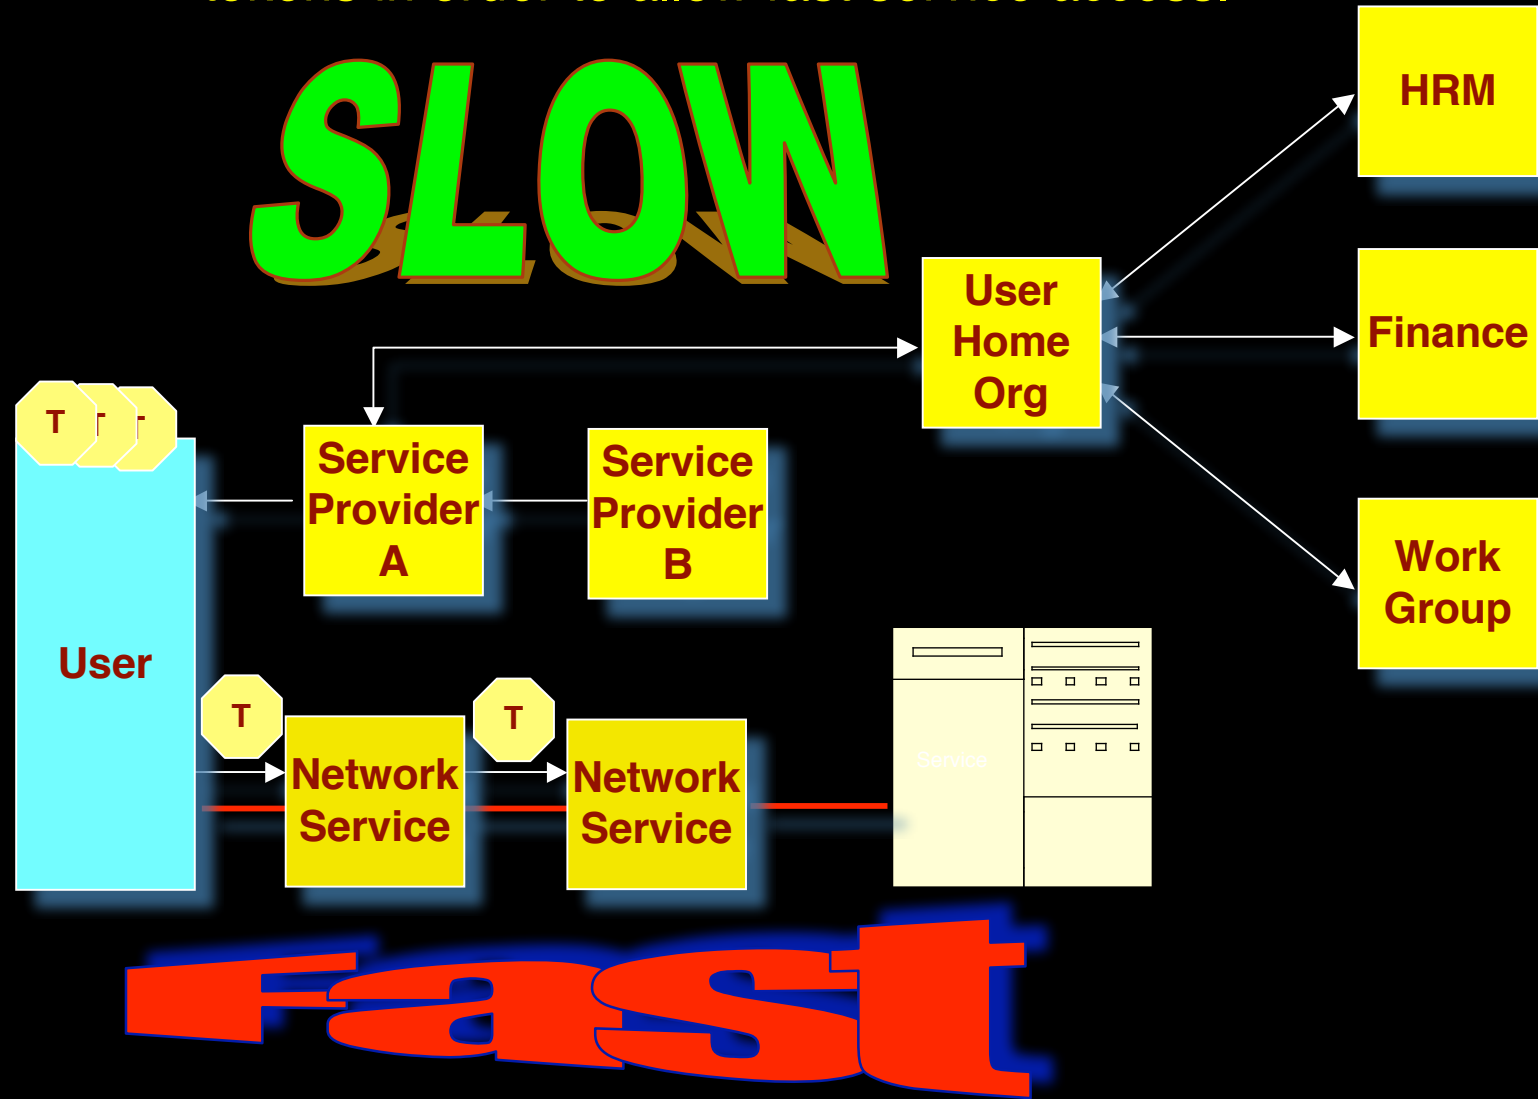
Surf the web!

Simple service access

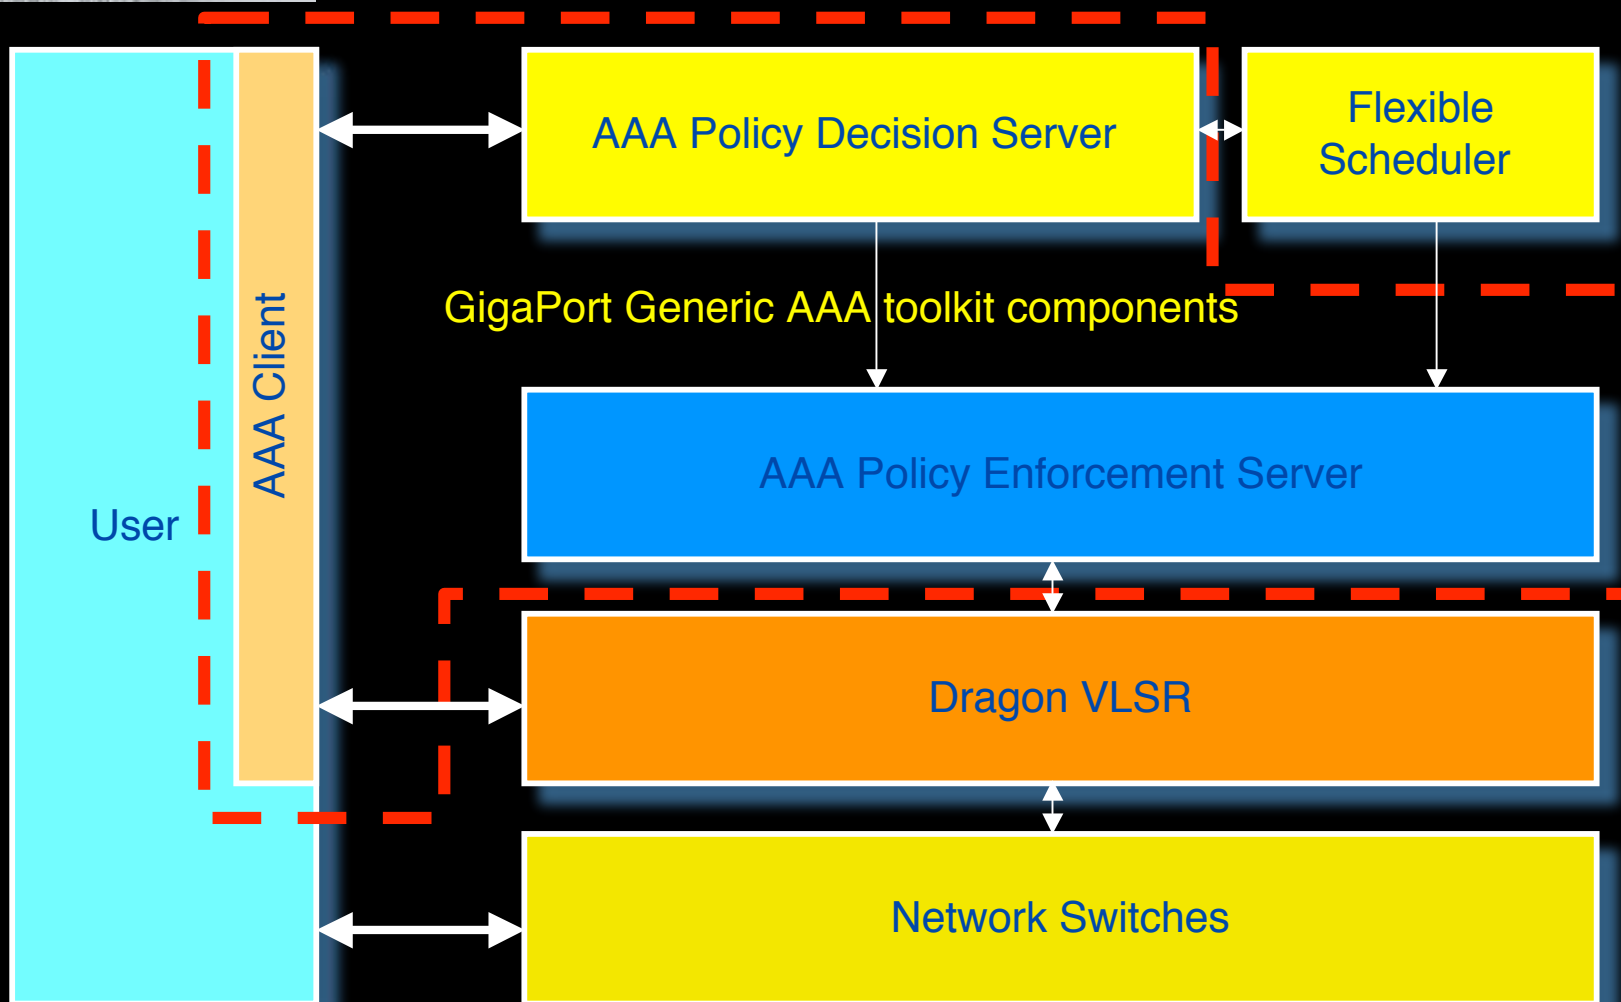


Pitlochry, Scotland - Summer 2005

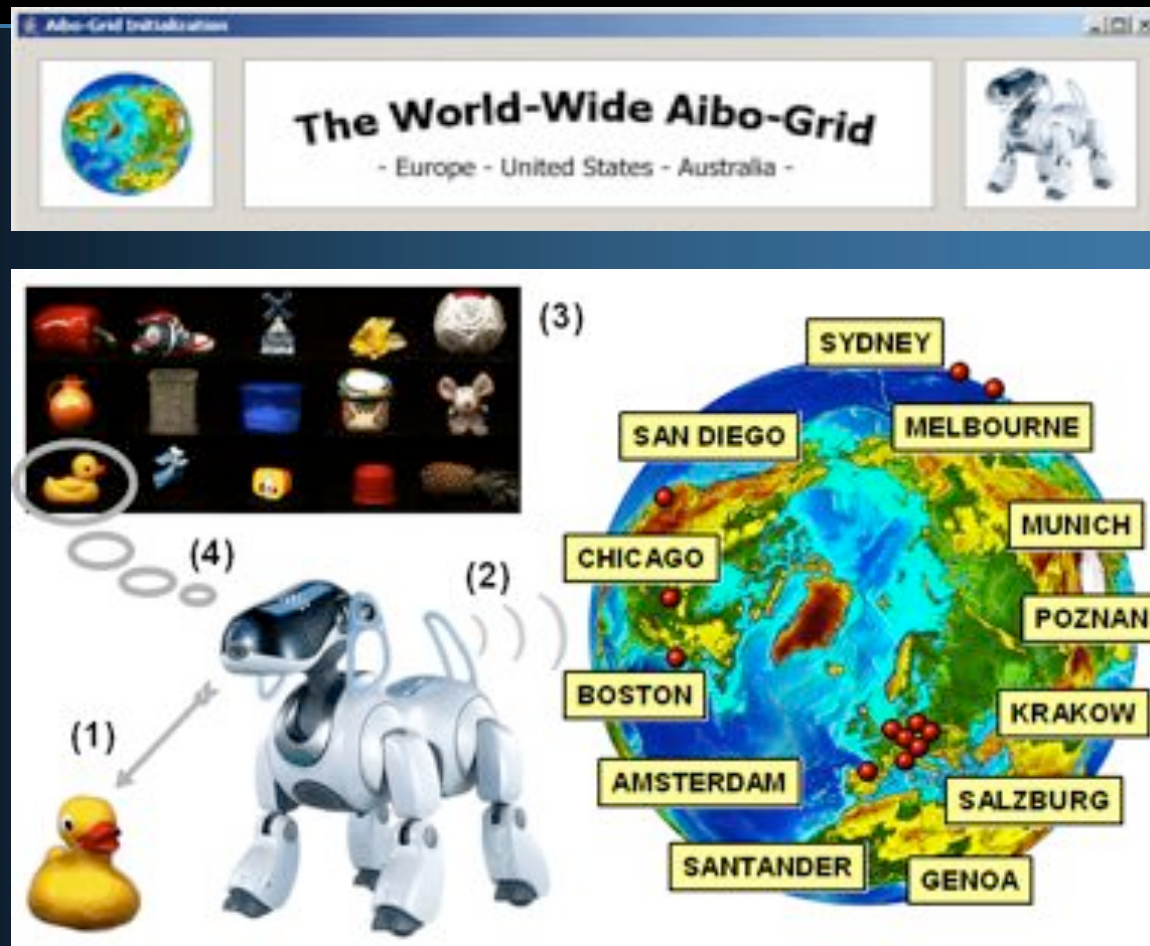
Use AAA concept to split (time consuming) service authorization process from service access using secure tokens in order to allow fast service access.



Components developed within GigaPort NG

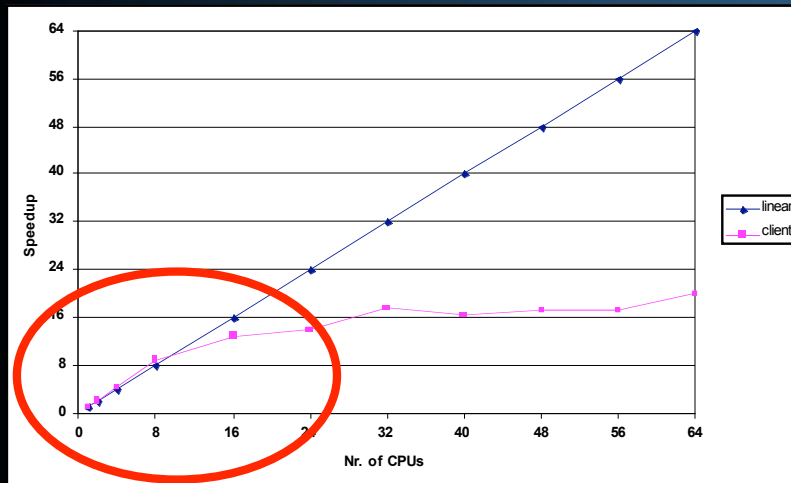


Example: Object Recognition

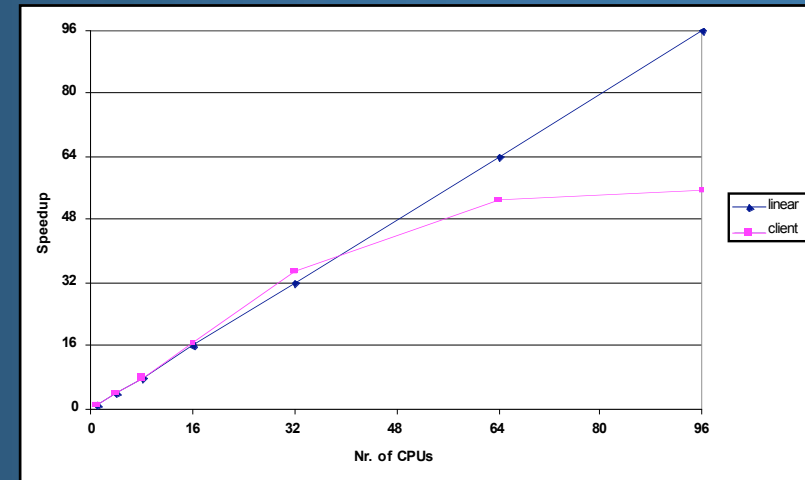


See also: <http://www.science.uva.nl/~fjseins/aibo.html>

Performance / Speedup on DAS-2



Single cluster, client side speedup



Four clusters, client side speedup

- Recognition on single machine: +/- 30 seconds
- Using multiple clusters: up to 10 frames per second
- Insightful: even 'distant' clusters can be used effectively for close to 'real-time' recognition

VL-e: Virtual Laboratory for e-Science project (2004-2008)

- 40 M€ Dutch project (20 M€ from government)
- 2 experimental environments:
 - Proof of Concept: applications research
 - Rapid Prototyping (using DAS): computer science
- Research on:
 - Applications (biodiversity, bioinformatics, food informatics, telescience, physics)
 - Computer science tools for visualization, workflow, ontologies, data management, PSEs, grid computing

vl·e



virtual laboratory for e-science

SC2004 “Dead Cat” demo

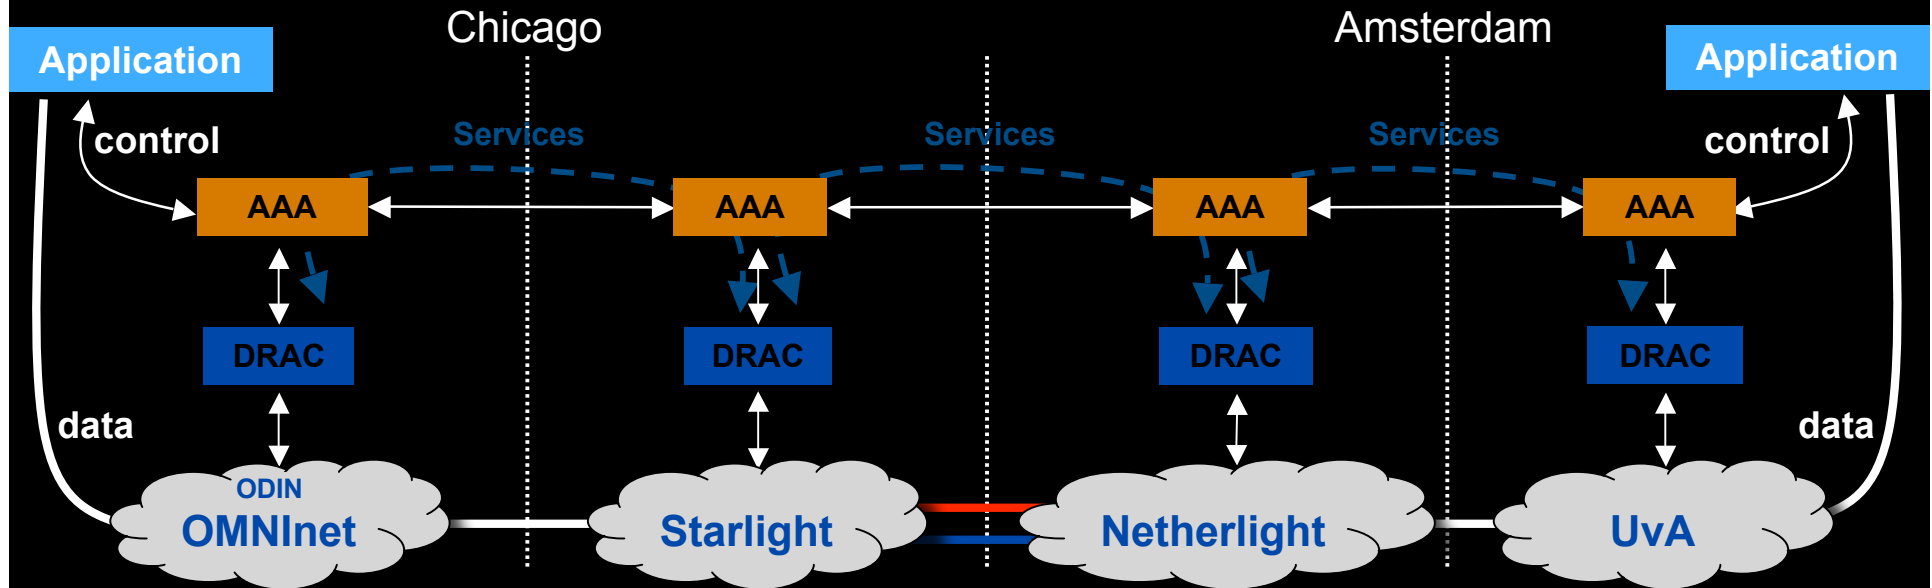
SC2004,
Pittsburgh,
Nov. 6 to 12, 2004

Produced by:
Michael Scarpa
Robert Belleman
Peter Sloat

Many thanks to:
AMC
SARA
GigaPort
UvA/AIR
Silicon Graphics, Inc.
Zoölogisch Museum



SC2004 Lambda Service Demonstrator




- finesse the control of bandwidth across multiple domains
- while exploiting scalability and intra-, inter-domain fault recovery
- thru layering of a novel SOA upon legacy control planes and NEs



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

Scientific Publications



- Elsevier FGCS journal iGrid2005 issue
 - 26 reviewed papers, 5 from GigaPort RON partners
 - many use GigaPort for their scientific work on and with the network and its principles
- Key publications (UvA):
 - Cees de Laat, Erik Radius, Steven Wallace, "The Rationale of the Current Optical Networking Initiatives", iGrid2002 special issue, Future Generation Computer Systems, volume 19 issue 6 (2003).
 - Tom DeFanti, Cees de Laat, Joe Mambretti, Kees Neggers, Bill St. Arnaud: "TransLight: a global-scale LambdaGrid for e-science", Communications of the ACM, Volume 46, Issue 11 (November 2003), Pages: 34 - 41.
 - Leon Gommans, Freek Dijkstra, Cees de Laat, Arie Taal, Alfred Wan, Bas van Oudenaarde, Tal Lavian, Inder Monga, Franco Travostino, "Applications Drive Secure Lightpath Creation across Heterogeneous Domains", IEEE Communications Magazine, vol. 44, no. 3, March 2006
 - Leon Gommans, Cees de Laat, Robert Meijer, "Token Based path authorization at Interconnection Points between Hybrid Networks and a Lambda Grid", IEEE GRIDNETS2005 proceedings, ISBN 0-7803-9277-9.
 - Freek Dijkstra, Cees de Laat, "Optical Exchanges", GRIDNETS conference proceedings, oct 2004, <http://www.broadnets.org/2004/workshop-papers/Gridnets/DijkstraF.pdf>.
 - **pubs about principles, protocols, methods, architectures, experiments & proofs of concept**
- About 30 publications (UvA) in journals and conf records.
 - see <http://www.science.uva.nl/~delaat/pubs.html>
- About 15 talks/year, many invited.
 - see <http://www.science.uva.nl/~delaat/talks.html>



Future transport research

- Scalability of optical/photonic networks
 - do we look for scale in fatter pipes (40 - 100 Gb) carved up with SONET?
 - do we want to go for 1000 moderate lambda's in one fiber?
 - go for a band-less system to provide more granularity?
 - what architectures can make the incremental costs for provisioning extra lambdas lower. (destructive versus constructive QoS)?
- Persistent technology
 - is SONET going to stay and is that fine or not?
 - is (G)MPLS going to stay and how do we think about that?
 - There was a time SONET was not to be mentioned and MPLS idem!
- Subsecond switching schemes
 - control plane issues!
 - physics issues
 - does it make sense to have a nanosecond switching device if it takes seconds to talk to?
 - does the remark: "the longer the circuit, the slower it can be switched" make sense?
 - argument: need anyway rtt's for signaling and traffic establishment!
- Dardy's argument, are we on the wrong technological path?
 - his proposal -> goto infiniband (RDMA)!
- Van Jacobson's proposal for transport
 - rework stacks into streams in user space!



Future AAA research

- Generic AAA toolkit version 2 development aiming at:
 - Creation of AAA scenario's using workflow management tools (BPEL).
 - Integration of Identity Management environments / AAls
 - Integration of Advance Flexible Schedulers (USC - Nextgrid)
 - Integration of GMPLS based network resources using tokens (DRAGON).
 - Integration of Network Description Language to describe authorization topology.
- Integration Trusted Network Computing architecture and AAA concepts to create a User Controlled Secure Environment in collaboration with Globus & EGEE.
- Development of concepts around User Programmable Virtualized Networks and integration of AAA token concepts.



Questions ?



Major contributions from Leon Gommans, Paola Grosso