

What makes an exchange open?



Cees de Laat
Freek Dijkstra
Leon Gommans
Bas van Oudenaarde

University of Amsterdam

www.glif.is

u
s
e
r
s

A. Lightweight users, browsing, mailing, home use


Need full Internet routing, one to many

B. Business applications, multicast, streaming, VPN's, mostly LAN


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

C. Scientific applications, distributed data processing, all sorts of grids

Need very fat pipes, limited multiple Virtual Organizations, few to few, p2p




Transport of flows




Towards Hybrid Networking!

- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
 - 10G routerblade -> 75-300 k\$, 10G switch port -> 5-10 k\$, MEMS port -> 0.5-1.5 k\$
 - DWDM lasers for long reach expensive, 10-50 k\$
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way ==> map A -> L3 , B -> L2 , C -> L1
- Give each packet in the network the service it needs, but no more !

L1 \approx 0.5-1.5 k\$/port




L2 \approx 5-10 k\$/port



L3 \approx 75+ k\$/port




How low can you go?





LIGHTHOUSE



QuickTime™ and a TIFF (Uncompressed) decompressor are needed to see this picture.

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.

UVA's

64*64

Optical Switch

@ LightHouse

**Costs 1/100th of
a similar
throughput router
or**

**1/10th of
a similar
throughput
Ethernet switch
but has only
specific services!**



Common Photonic Layer (CPL) in SURFnet6

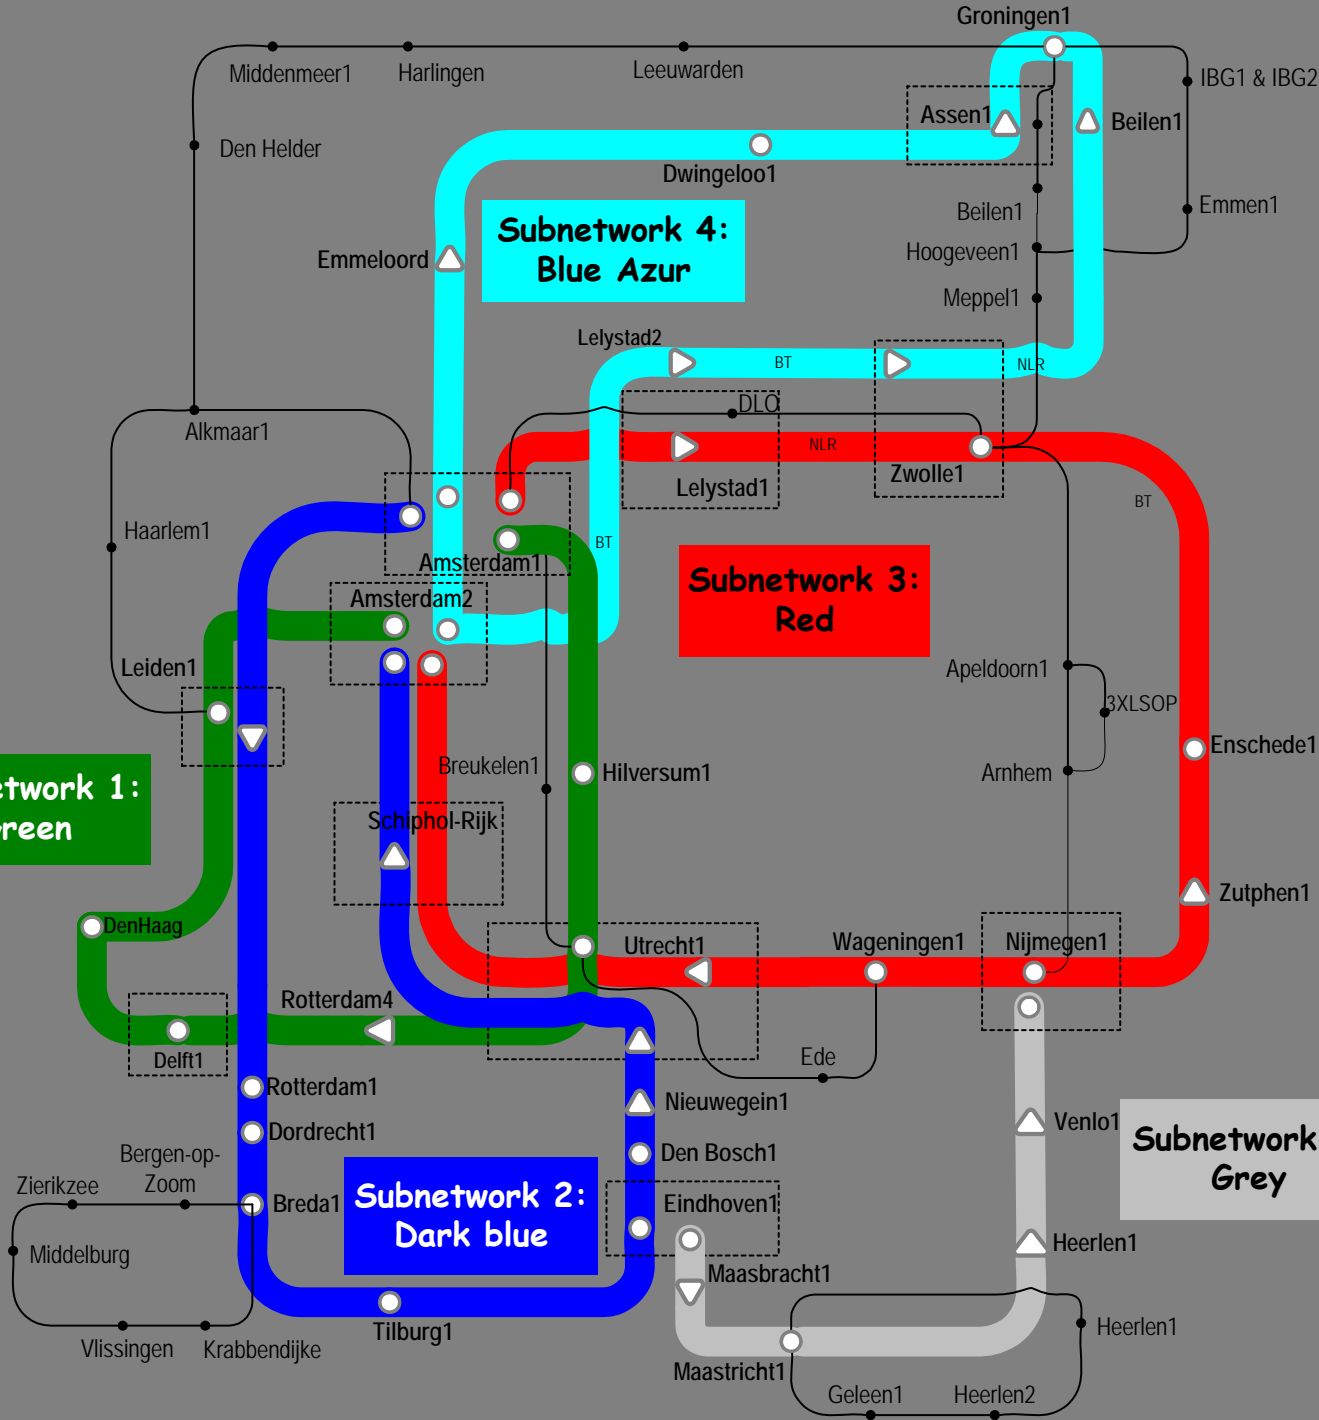
**Subnetwork 1:
Green**

**Subnetwork 2:
Dark blue**

**Subnetwork 4:
Blue Azur**

**Subnetwork 3:
Red**

**Subnetwork 5:
Grey**

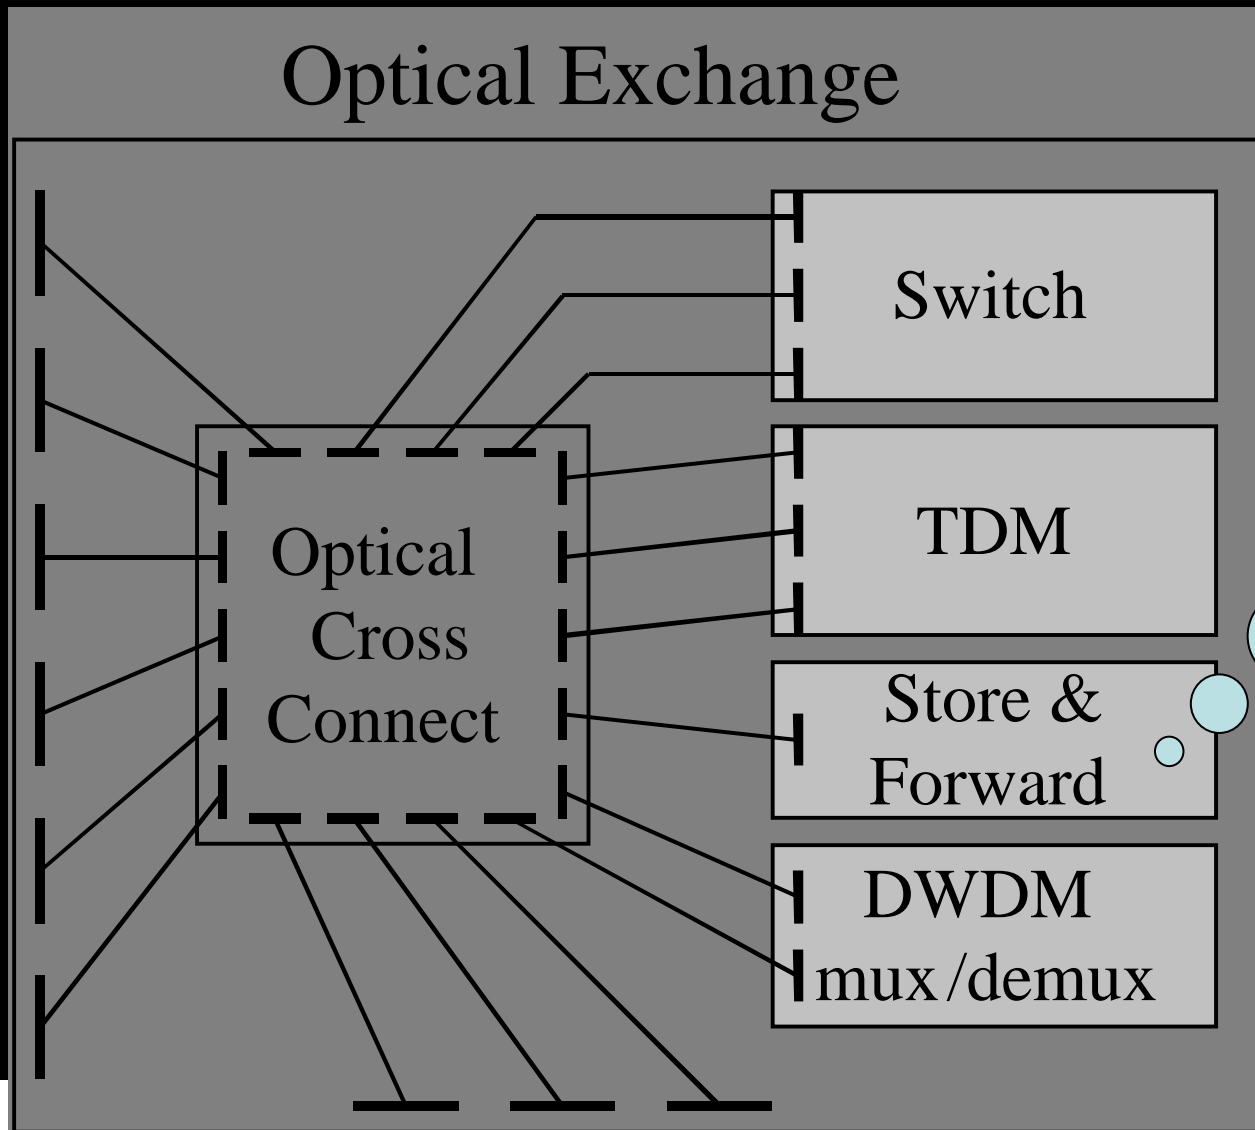




Services

<div>SCALE</div> <div>CLASS</div>	<div>2</div> <div>Metro</div>	<div>20</div> <div>National/ regional</div>	<div>200</div> <div>World</div>
A	<div>Switching/ routing</div>	<div>Routing</div>	<div>ROUTERS\$</div>
B	<div>Switches + E-WANPHY VPN's</div>	<div>Switches + E-WANPHY (G)MPLS</div>	<div>ROUTERS\$</div>
C	<div>dark fiber</div> <div>DWDM</div> <div>MEMS switch</div>	<div>DWDM, TDM / SONET</div> <div>Lambda switching</div>	<div>Lambdas,</div> <div>VLAN's</div> <div>SONET</div> <div>Ethernet</div>

Optical Exchange as Black Box



TeraByte
Email
Service

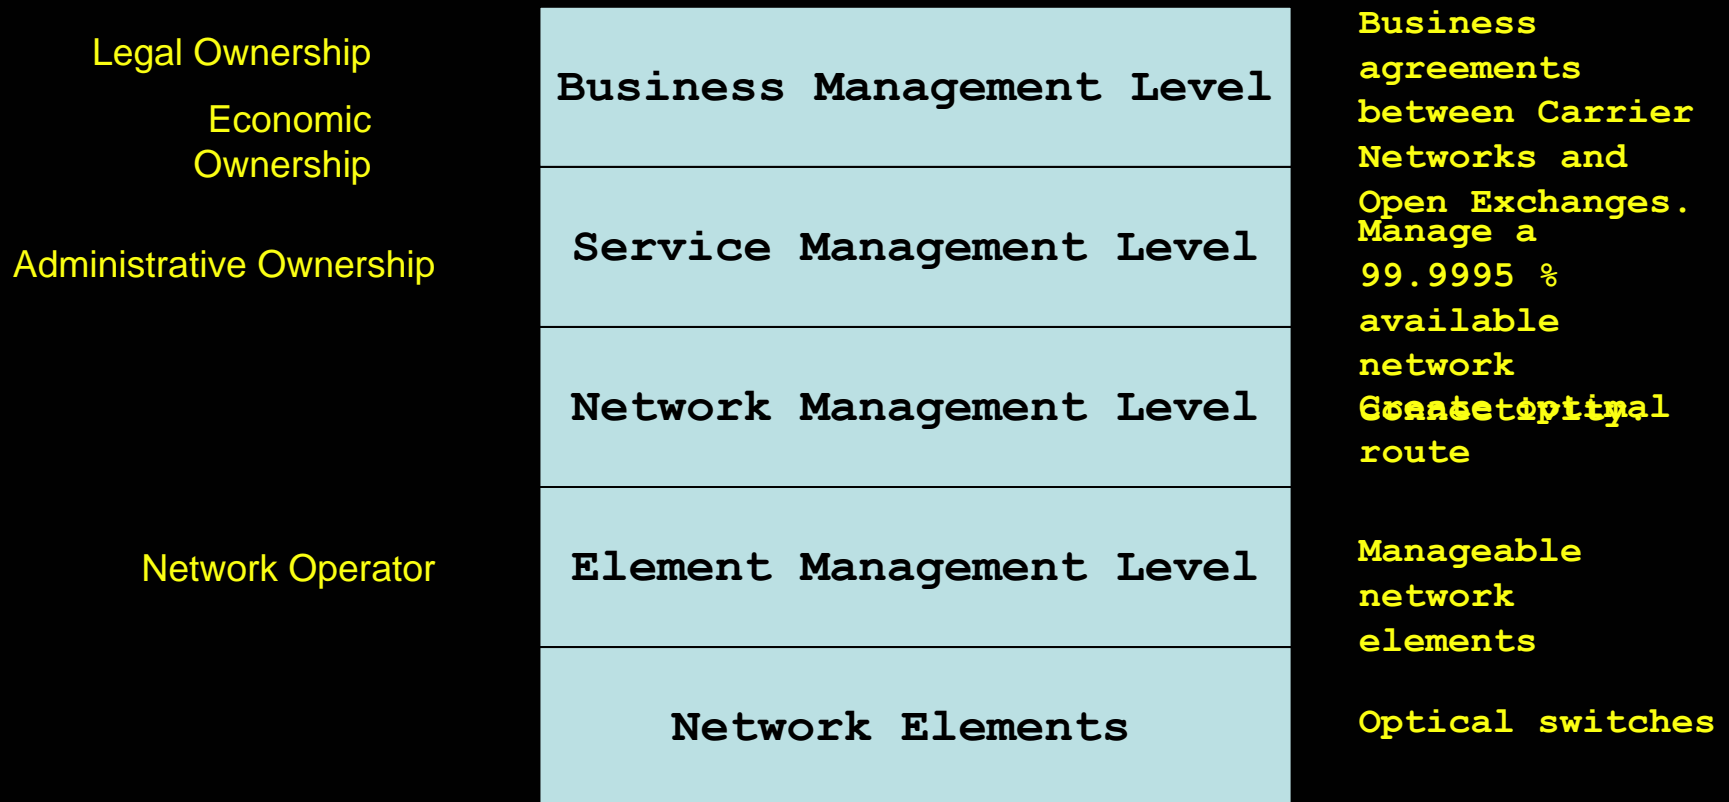
Service Matrix

To From	WDM (multiple λ)	Single λ , any bitstream	SONET/ SDH	1 Gb/s Ethernet	LAN PHY Ethernet	WAN PHY Ethernet	VLAN tagged Ethernet	IP over Ethernet
WDM (multiple λ)	cross-connect multicast, regenerate, multicast	WDM demux	WDM demux*	WDM demux *	WDM demux *	WDM demux *	WDM demux *	WDM demux *
Single λ , any bitstream	WDM mux	cross-connect multicast, regenerate, multicast	N/A *	N/A *	N/A *	N/A *	N/A *	N/A *
SONET/SDH	WDM mux	N/A *	SONET switch, +	TDM demux *	TDM demux ⁶	SONET switch	TDM demux *	TDM demux *
1 Gb/s Ethernet	WDM mux	N/A *	TDM mux	aggregate, Ethernet conversion +	aggregate, eth. convert	aggregate, Ethernet conversion	aggregate, VLAN encap	L3 entry *
LAN PHY Ethernet	WDM mux	N/A*	TDM mux ⁶	aggregate, Ethernet conversion	aggregate, Ethernet conversion +	Ethernet conversion	aggregate, VLAN encap	L3 entry *
WAN PHY Ethernet	WDM mux	N/A *	SONET switch	aggregate, Ethernet conversion	Ethernet conversion	aggregate, Ethernet conversion +	aggregate, VLAN encap	L3 entry *
VLAN tagged Ethernet	WDM mux	N/A *	TDM mux	aggregate, VLAN decap	aggregate, VLAN decap	aggregate, VLAN decap	Aggregate, VLAN decap & encap +	N/A
IP over Ethernet	WDM mux	N/A *	TDM mux	L3 exit *	L3 exit *	L3 exit *	N/A	Store & forward, L3 entry/exit+

Ownership of resources

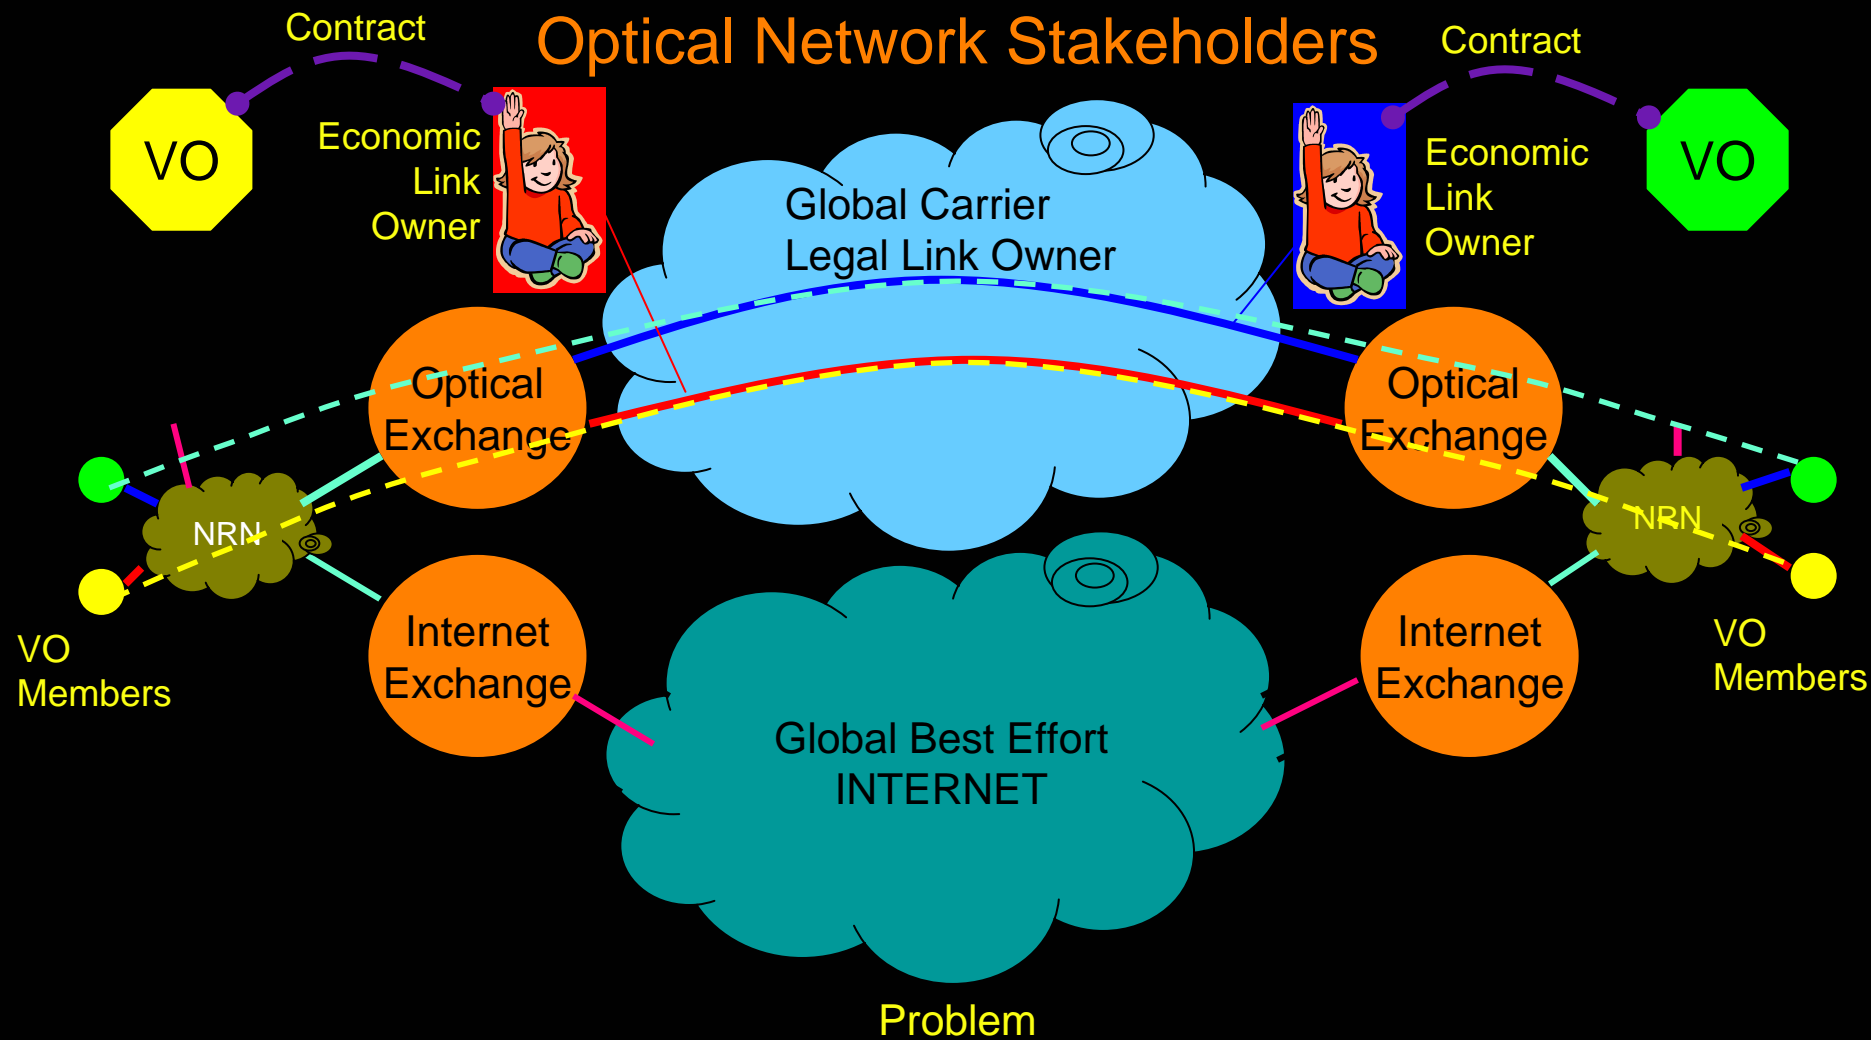
- **Legal Owner:**
 - Organization that legally owns a resource.
 - A legal owner may sell the right to economically use the resource.
- **Economic Owner:**
 - Acquires economic resource usage right a from legal resource owner.
 - A contract details terms by which a resource may be used.
 - Economic owners may outsource resource management to an Administrative Owner by means of a service level agreement.
- **Administrative Owner:**
 - Technically implements the terms of a service level agreement
 - Signals requests to other AO's and handles responses.
 - Collects accounting information.
- **Relationship between owners:**
 - Legal, economic and administrative owners may or may not be independent organizations.
 - Economic owners may acquire resources from different legal owners.
 - Administrative owners may serve different economic owners.
 - Economic owners may establish contracts with other economic owners to create more elaborate services. Technical details are delegated and implemented by Administrative Owners.

ISO Telecommunications Management Networks (TMN) reference model



TMN is based on the OSI management framework and uses an object-oriented approach, with managed information in network resources modeled as attributes in managed objects. TMN is defined in ITU-T M.3000 series recommendations

Optical Network Stakeholders



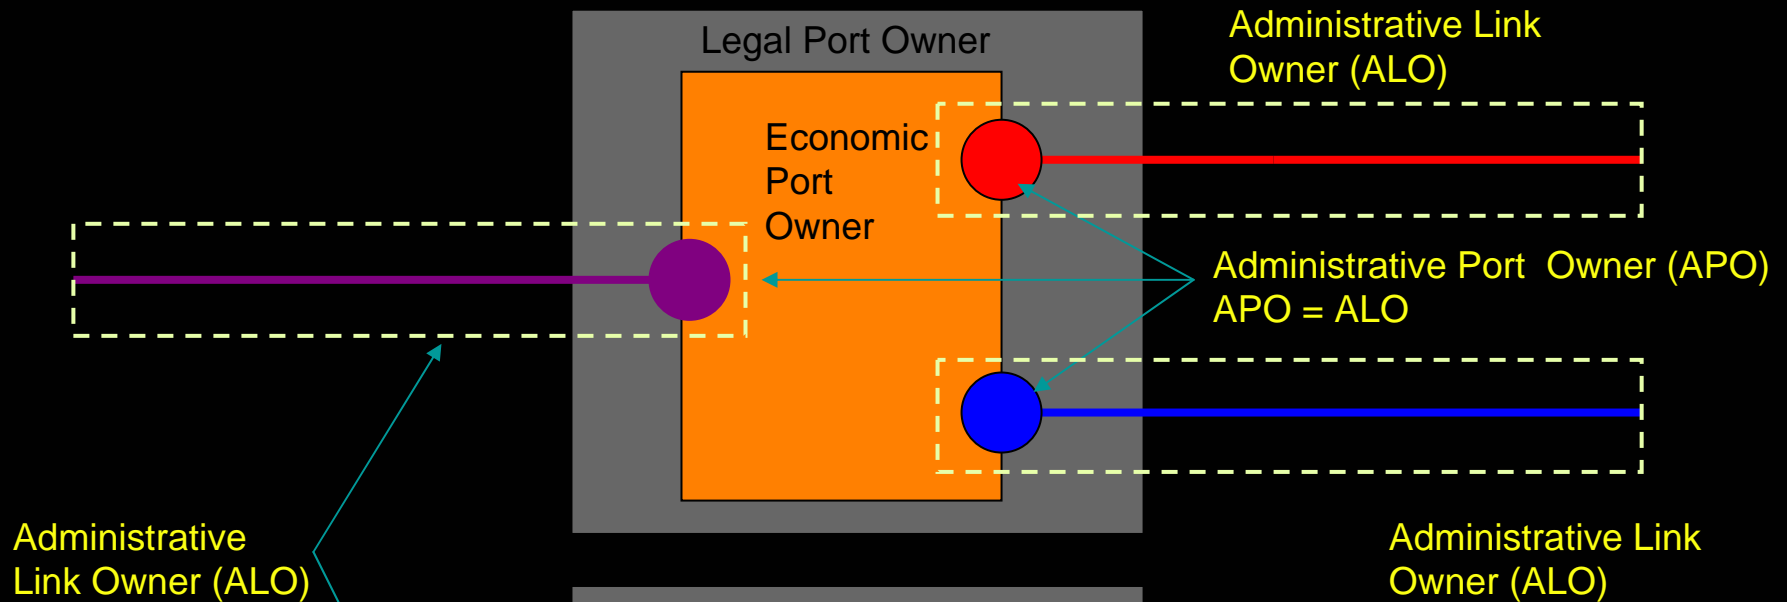
In order to enable a dynamic, cost effective VO business operation, Economic Link Owners Red and Blue need to create and have the ability to implement link usage contracts with VO's leading to the creation of Optical Private Network (OPN) between VO members.

Role definitions

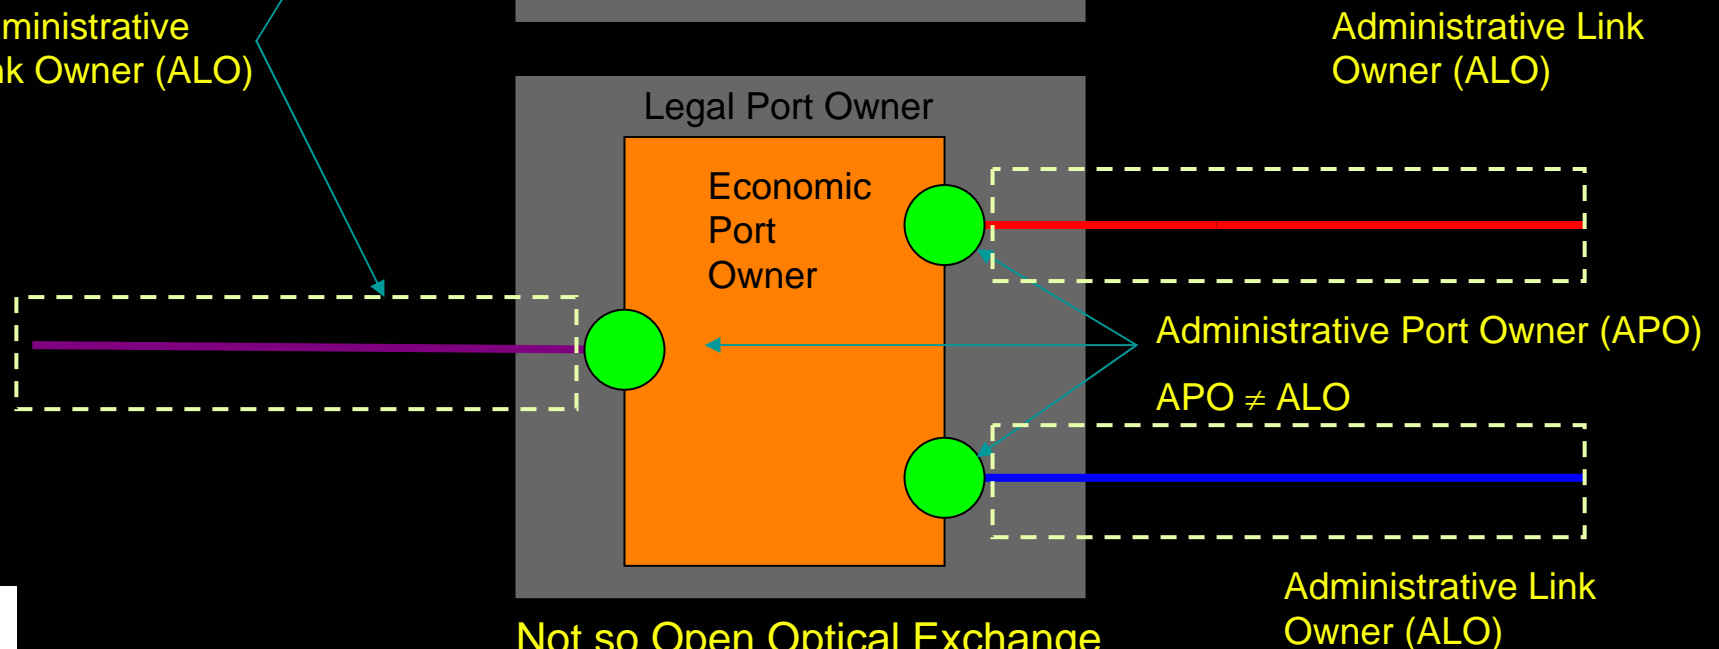
- Legal Link Owner (LLO): Sells the right to use a link to an ELO's
- Economic Link Owner (ELO): Acquires the right to use a link and creates agreements with Economic VO's about the usage of its links. ELO's will terminate a link at an optical exchange based on a contract with an EPO.
- Administrative Link Owner (ALO): Translates the ELO defined business rules governing link access to technical rules that are subsequently pushed to the APO for enforcement (optical link fibers have no electronic control).
- Legal Port Owner (LPO): Owns optical switch-ports. Usage rights are sold to EPO's. Multiple LPO's may be present within an Optical Exchange.
- Economic Port Owner (EPO): Acquires the usage right from one or more LPO's for one or more ports on the Optical Exchange. EPO's establishes contracts to allow peering with own or other EPO ports on behalf of ELO's.
- Administrative Port Owner (APO): an entity that accepts peering policies from ALO's. Peering policies are based on the agreements between ELO and a VO. Creates connections with own ports or other ports from different APO's based on requests with credentials from VO's members or its proxy .

Optical Exchange Stakeholders

Open Optical Exchange



Not so Open Optical Exchange



Discussion points

- Nobody wanted to be a closed “whatever”
- Naming is important
- Housing can be open or closed
- The impact of AUP’s
- Is a policy free exchange an open exchange
 - Always say yes
 - Content ignorant

Research and development issues on OEX's

- Economic & business models
- Hierarchical virtualization of exchange components
- Resource descriptions
 - Rdf
 - Central database
 - dns
 - web services, etc...
- Service plane
- Control plane
- Management plane
- AAA model applications
- Pick the right F to standardize

On the iGRID2005 DEMO front:

SARA & collaborators just before this session scored
a very significant record :

> 19 Gbit/s over two 10 Gbit/s Lambda's to run the
55 tile panel feeding data from Amsterdam!

Ref. Paul Wielinga

More shameless dutch pr on the iGRID2005 DEMO front:

NORTEL/UvA/NWU & collaborators showed
complete virtual system and running application
migration from Amsterdam to San Diego using
dynamical Lambdas with subsecond application
downtime

Novel token based Lambda on demand

Automatic addressing setup worked in the last demo!
iGRID is a very inspiring learning event!

The “Dead Cat” demo

Produced by:

Michael Scarpa
Robert Belleman
Peter Slood
Cees de Laat

Many thanks to:

AMC
SARA
GigaPort
UvA/AIR
Silicon Graphics
Zoölogisch Museum

iGRID2005 publication opportunity

"Future Generation Computer Systems (FGCS): The International Journal of Grid Computing: Theory, Methods and Applications" will publish a SPECIAL iGRID ISSUE in Spring/Summer 2006.

Guest editors: Larry Smarr, Tom DeFanti, Maxine Brown, Cees de Laat

We can accept around 20-25 papers, Papers will be reviewed

- * Maximum paper length is limited to 8 pages
- * Limit of 1 paper per demonstration.
- * Describe your iGrid experiences, results and performance measurements.
- * **DEADLINE for submission is ONE MONTH AFTER iGRID -> Oct 31.**

Submission must be via the FGCS website. For author guides and submission information, see <<http://ees.elsevier.com/fgcs/>>.

Contact: Cees de Laat: delaat@science.uva.nl

(need reviewers :-)

Questions ?

More info:

<http://www.science.uva.nl/~delaat>

delaat@uva.nl

Credits:

- Leon Gommans, Bas Oudenaarde, Freek Dijkstra, Bert Andree, Jeroen van der Ham, Karst Koymans, & team
- Paola Grosso, Hans Blom
- SURFnet / GigaPort, Kees Neggers, Erik-Jan Bos, et al!
- NORTEL: Franco Travostino, Kim Roberts
- SARA: Anwar Osseryan, Paul Wielinga, Pieter de Boer, Ronald van der Pol
- Joe Mambretti, Bill stArnaud, GLIF community
- Tom & Maxine & Larry, laurin, OptIPuter, OnVector team !!!!

