Why is optical networking interesting?

www.science.uva.nl/~delaat

Cees de Laat

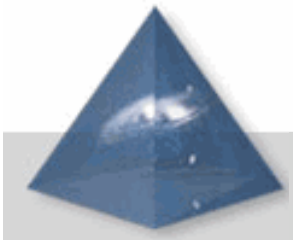

Faculty of Science

Why is optical networking interesting?

www.science.uva.nl/~delaat

Cees de Laat

EU

SURFnet

University of Amsterdam

SARA

NIKHEF

serena

What is this buzz about optical networking

- **Networks are already optical for ages**
- **Users won't see the light**
- **Almost all current projects are about SONET circuits and Ethernet (old wine in new bags?)**
- **Are we going back to the telecom world, do NRN's want to become telco's**
- **Does it scale**
- **Is it all about speed or is it integrated services**

Current technology + (re)definition

- Current (to me) available technology consists of SONET/SDH switches
- Changing very soon!
- DWDM+switching coming up
- Starlight uses for the time being VLAN's on Ethernet switches to connect [exactly] two ports
- So redefine a λ as:

"a λ is a pipe where you can inspect packets as they enter and when they exit, but principally not when in transit. In transit one only deals with the parameters of the pipe: number, color, bandwidth"

VLBI

ger term VLBI is easily capable of generating many Gb of data per
The sensitivity of the VLBI array scales w
(=data-rate) and there is a strong push to
. Rates of 8Gb/s or more are entirely feasibl
ider development. It is expected that parall
orrelator will remain the most efficient approa
s distributed processing may have an appli
lti-gigabit data streams will aggregate into la
or and the capacity of the final link to the da
tor.

Westerbork Synthesis Radio Telescope - Netherlands

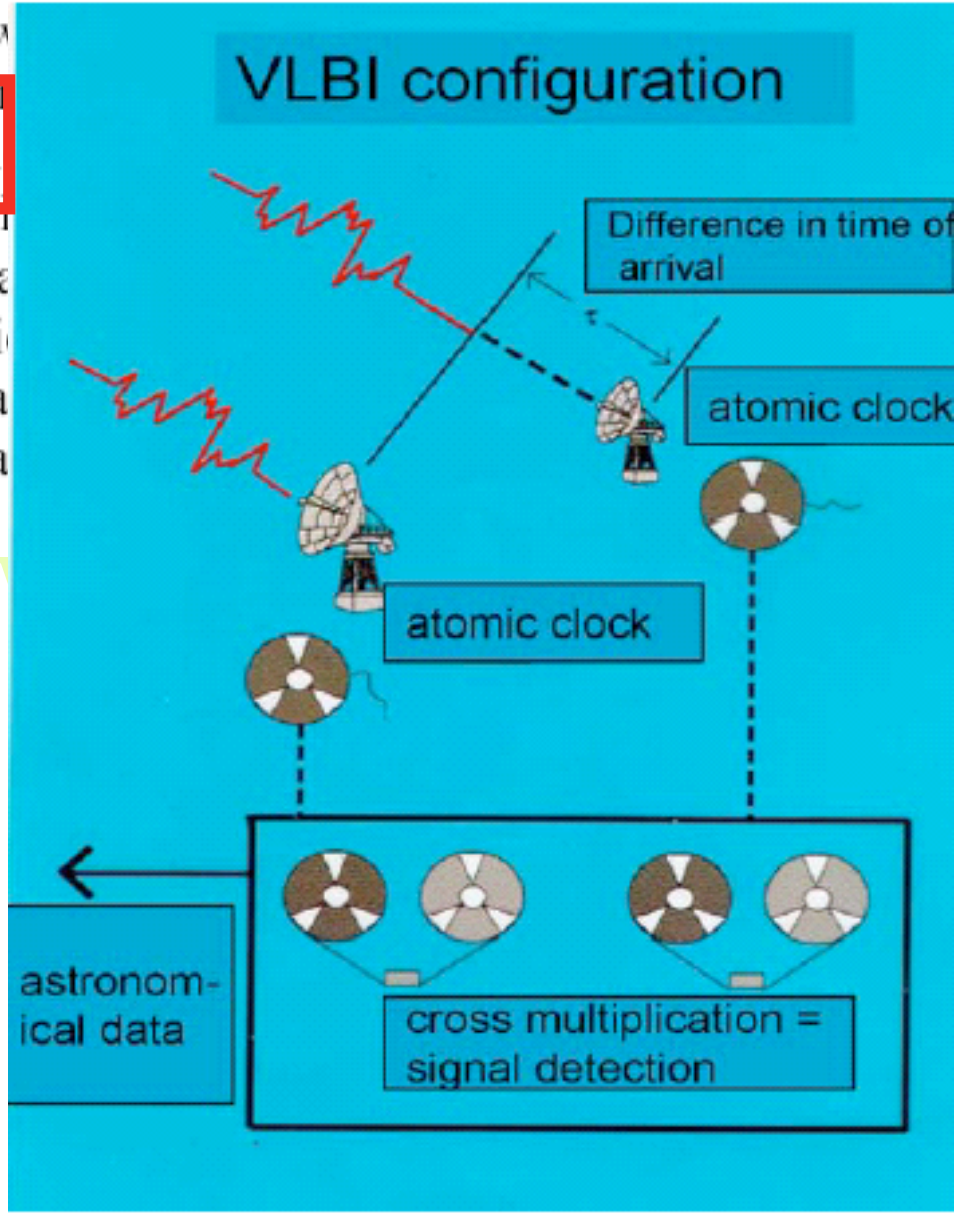

Know the user

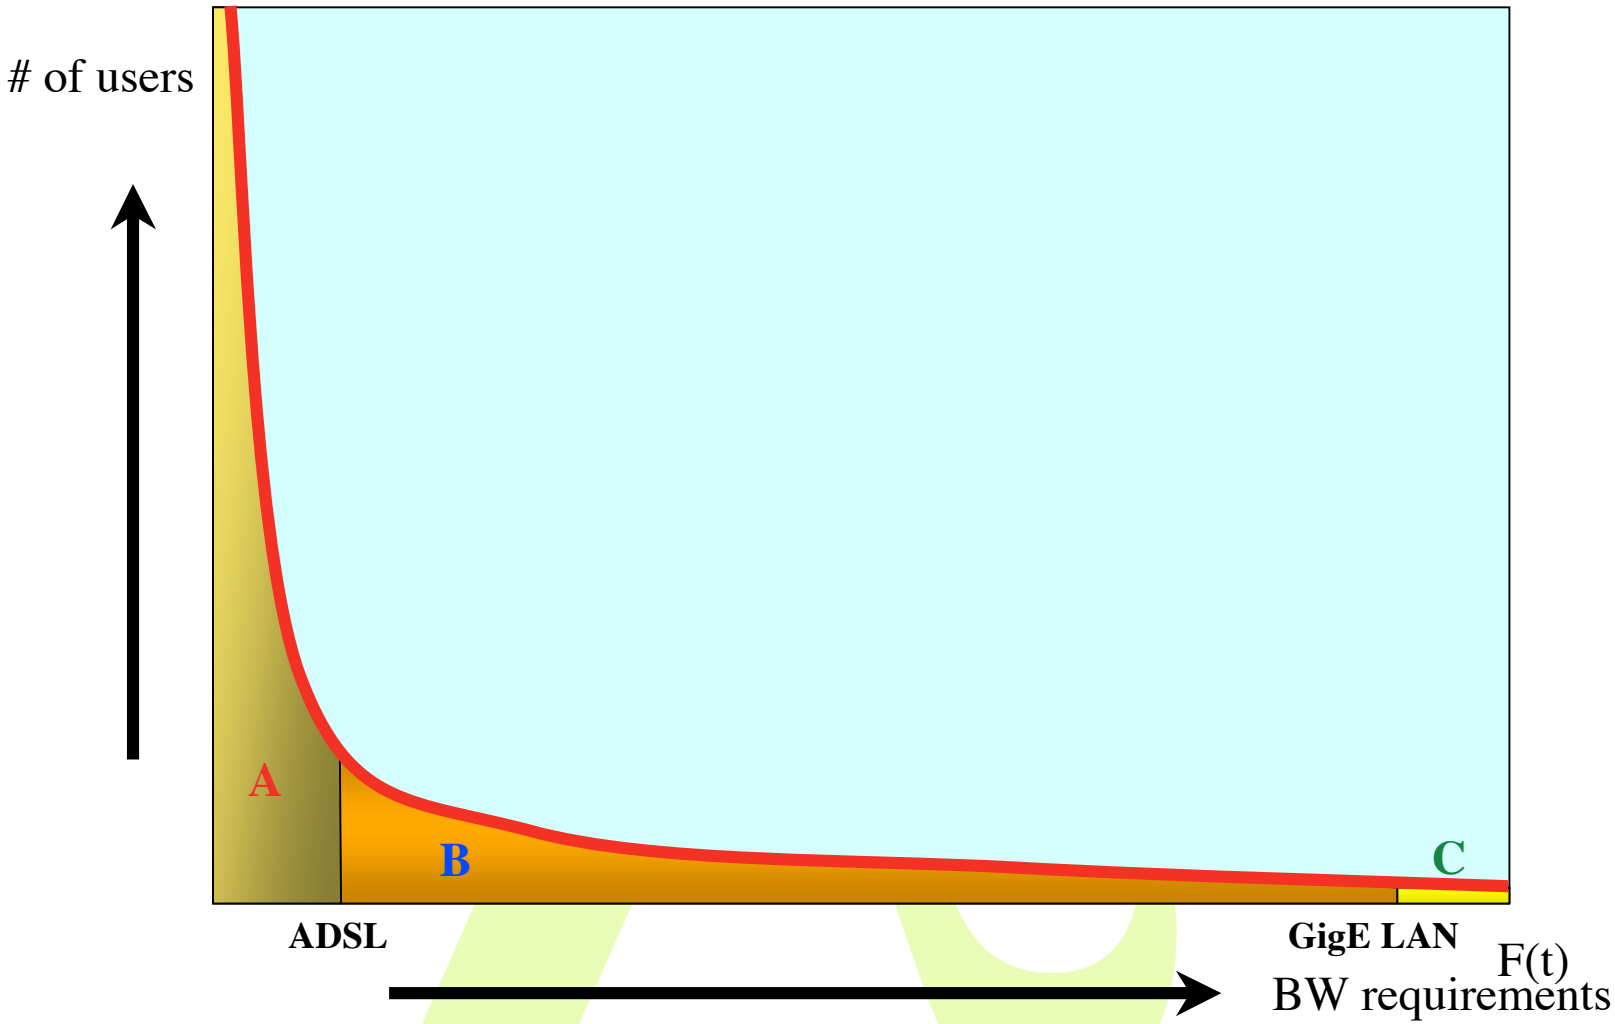

A -> Lightweight users, browsing, mailing, home use

B -> Business applications, multicast, streaming, VPN's, mostly LAN

C -> Special scientific applications, computing, data grids, virtual-presence

What the user

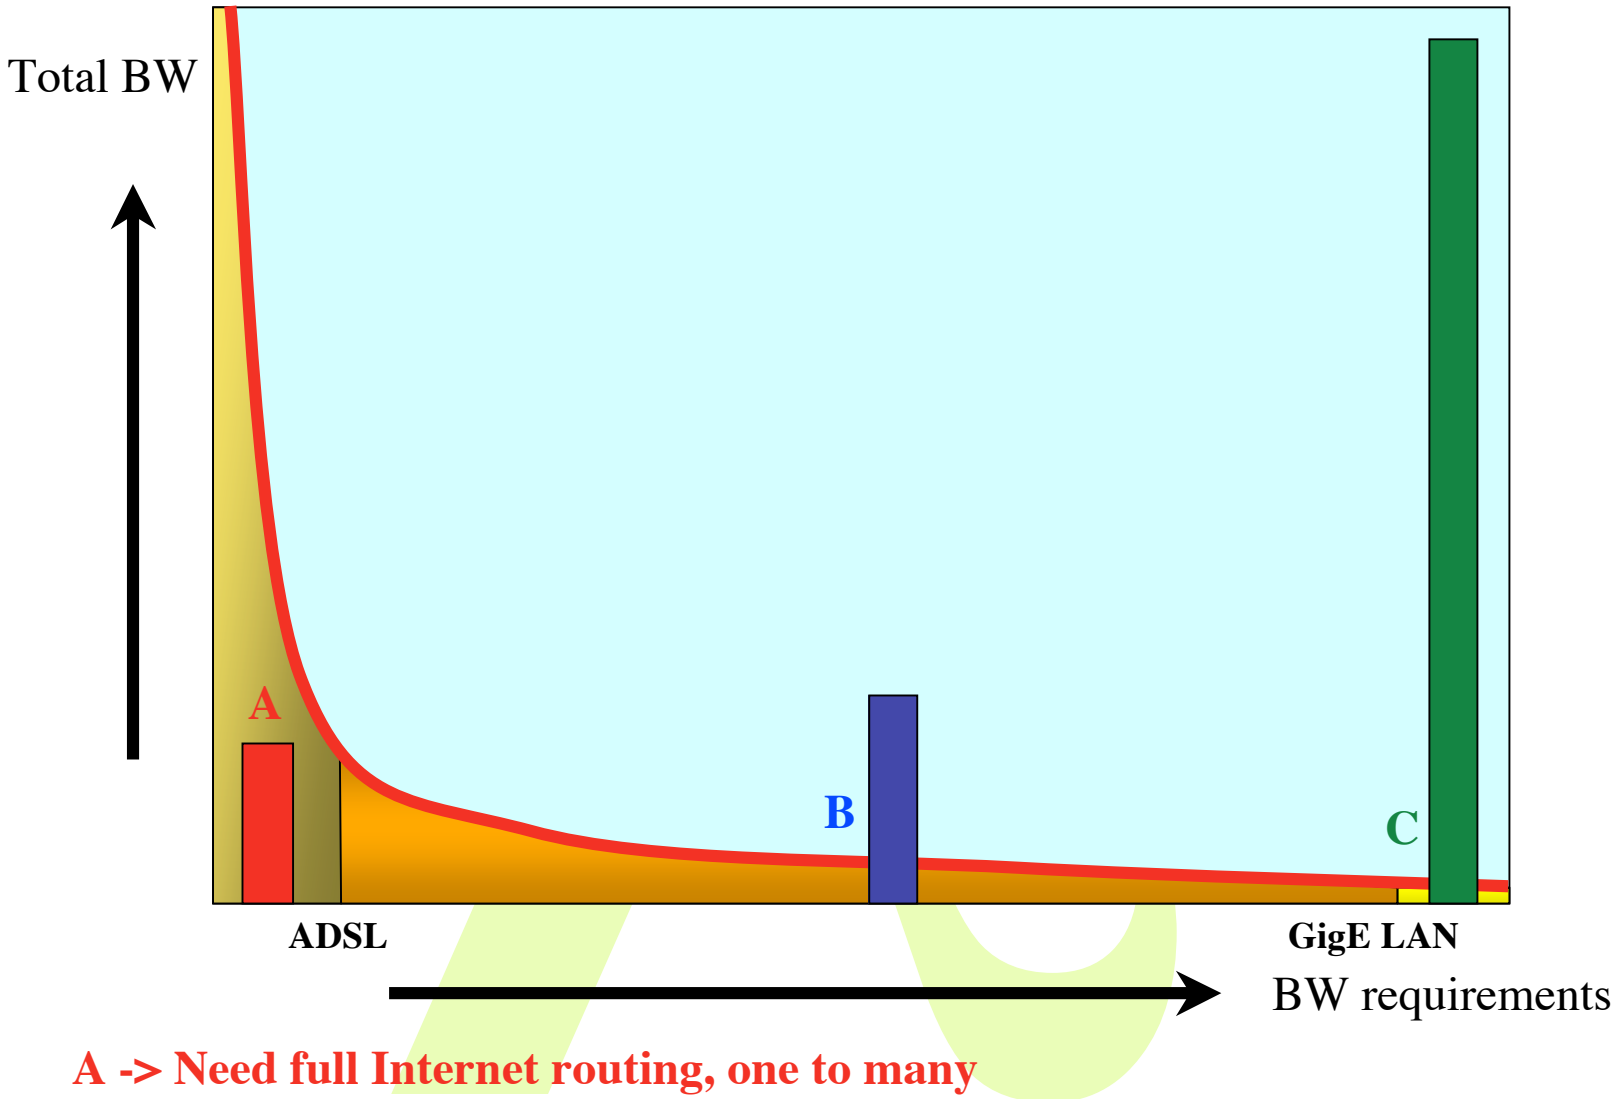

A -> Need full Internet routing, one to many

B -> Need VPN services on/and full Internet routing, several to several

C -> Need very fat pipes, limited multiple Virtual Organizations, few to few

So what are the facts

- Costs of fat pipes (fibers) are one/third of equipment to light them up
  - Is what Lambda salesmen tell me
- Costs of optical equipment 10% of switching 10 % of full routing equipment for same throughput
  - 100 Byte packet @ 10 Gb/s -> 80 ns to look up in 100 Mbyte routing table (light speed from me to you on the back row!)
- Big sciences need fat pipes
- Bottom line: create a hybrid architecture which serves all users in one consistent cost effective way

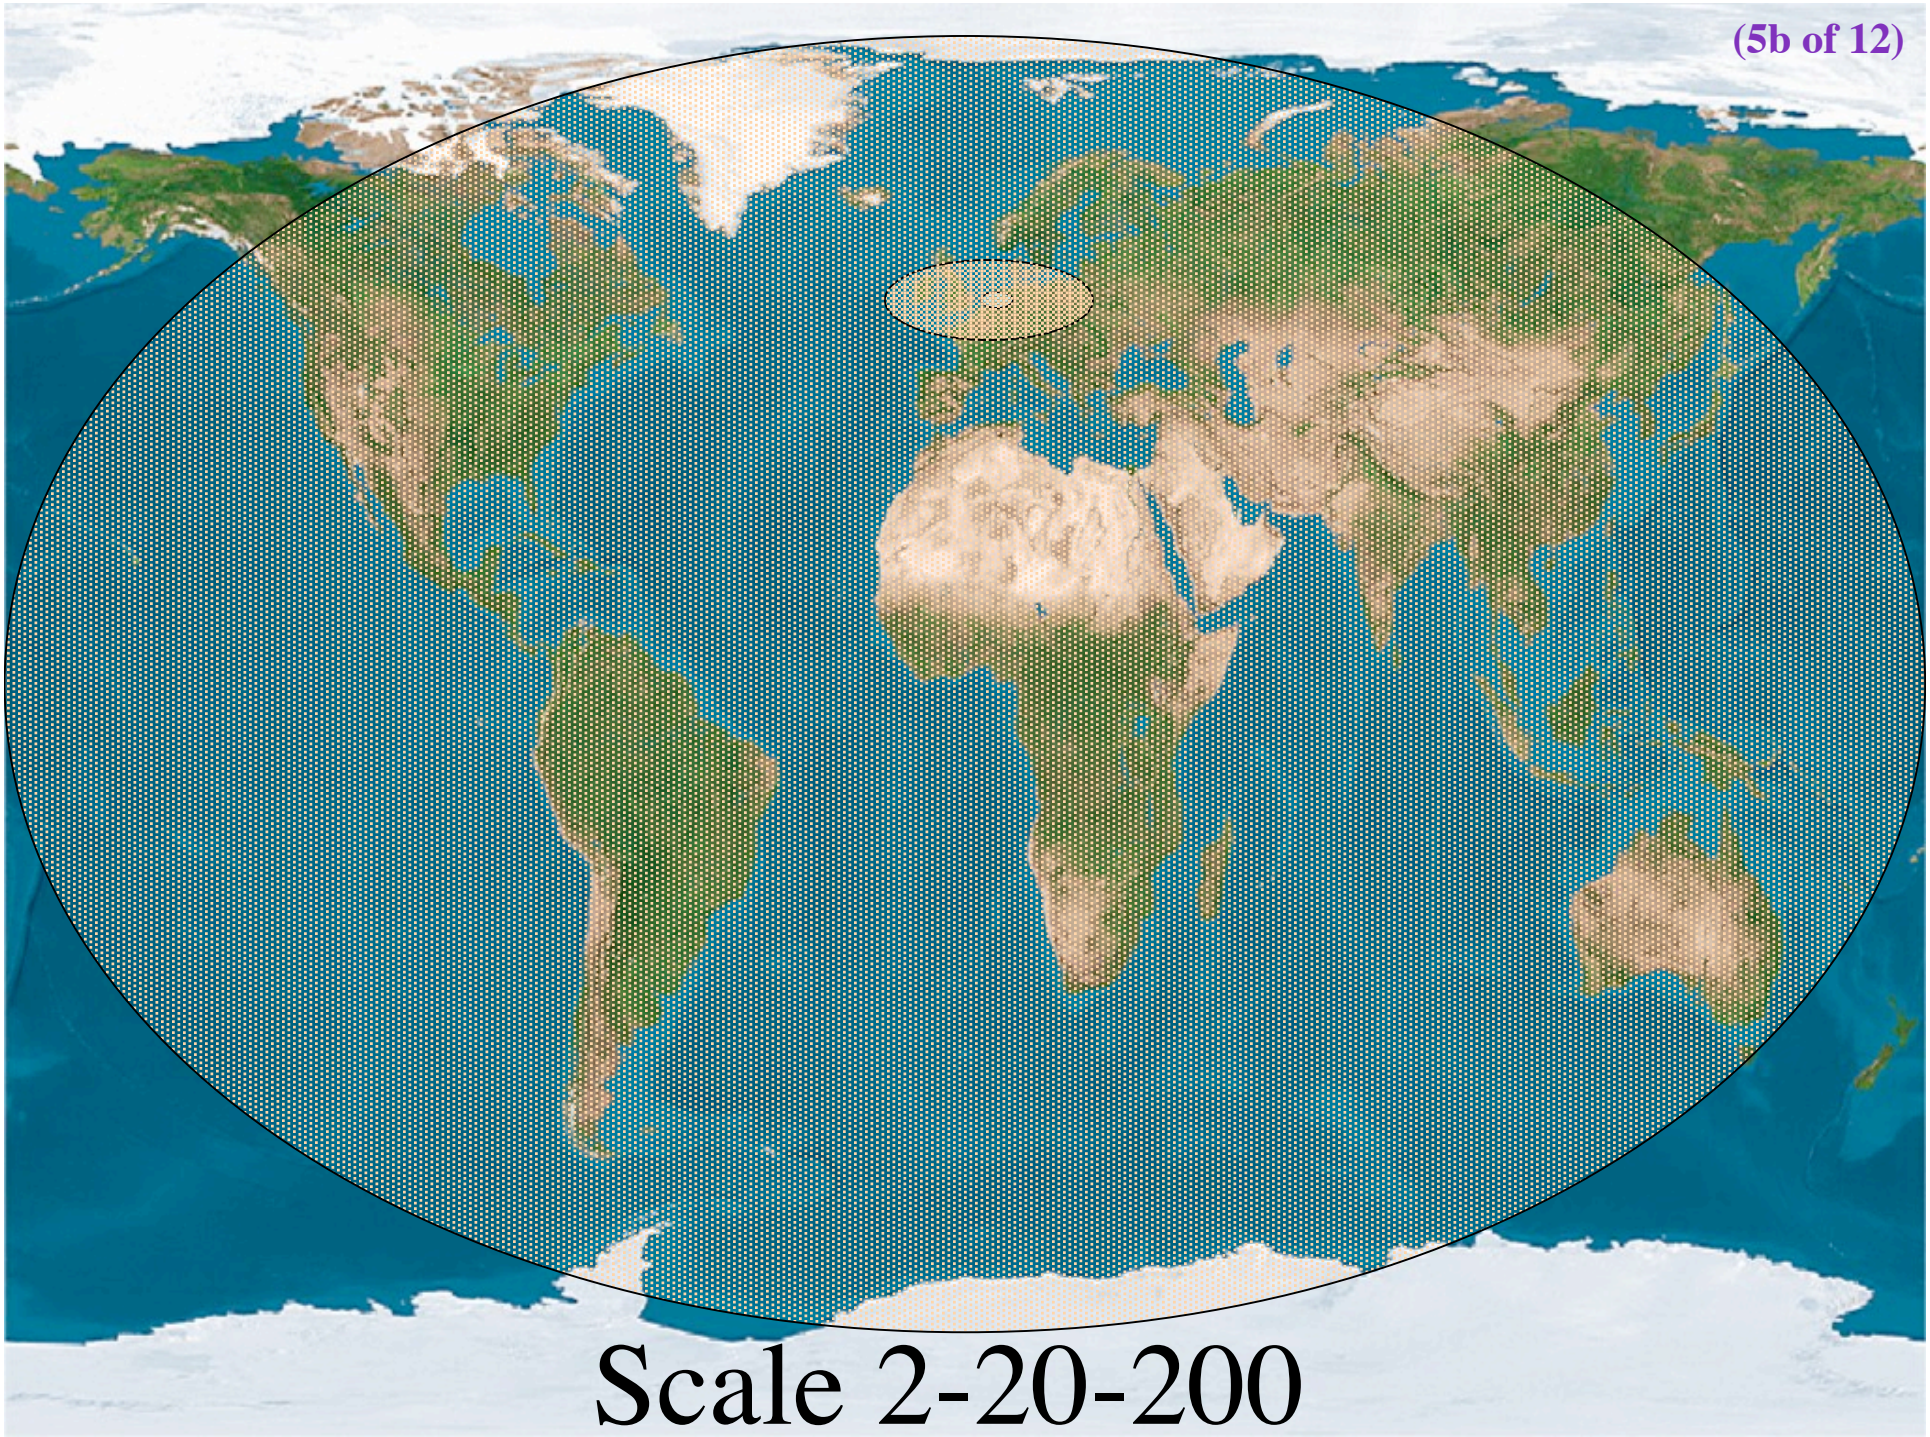

Scale 2-20-200

Services

	2 Metro	20 National/ regional	200 World
A	Switching/ routing	Routing	ROUTER$
B	VPN's, (G)MPLS	VPN's Routing	Routing
C #λ ~ 200/RTT	dark fiber Optical switching	Lambda switching	Sub- lambdas, ethernet- sdh

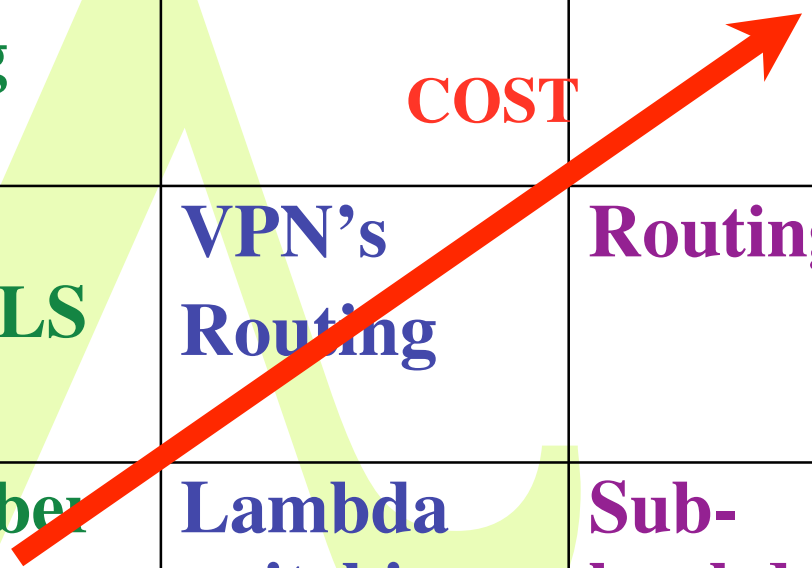

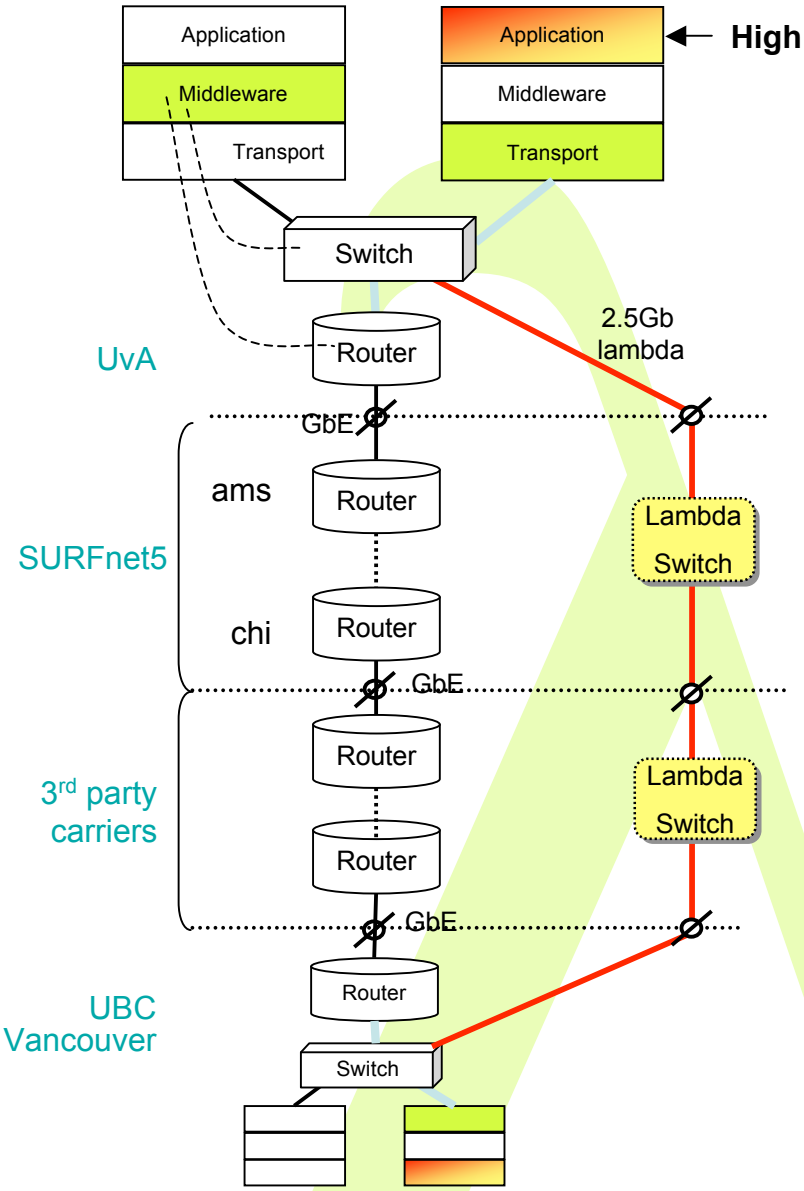

- lambda for high bandwidth applications
  - Bypass of production network
  - Middleware may request (optical) pipe
- RATIONALE:
  - Lower the cost of transport per packet

CA*net 4 Architecture

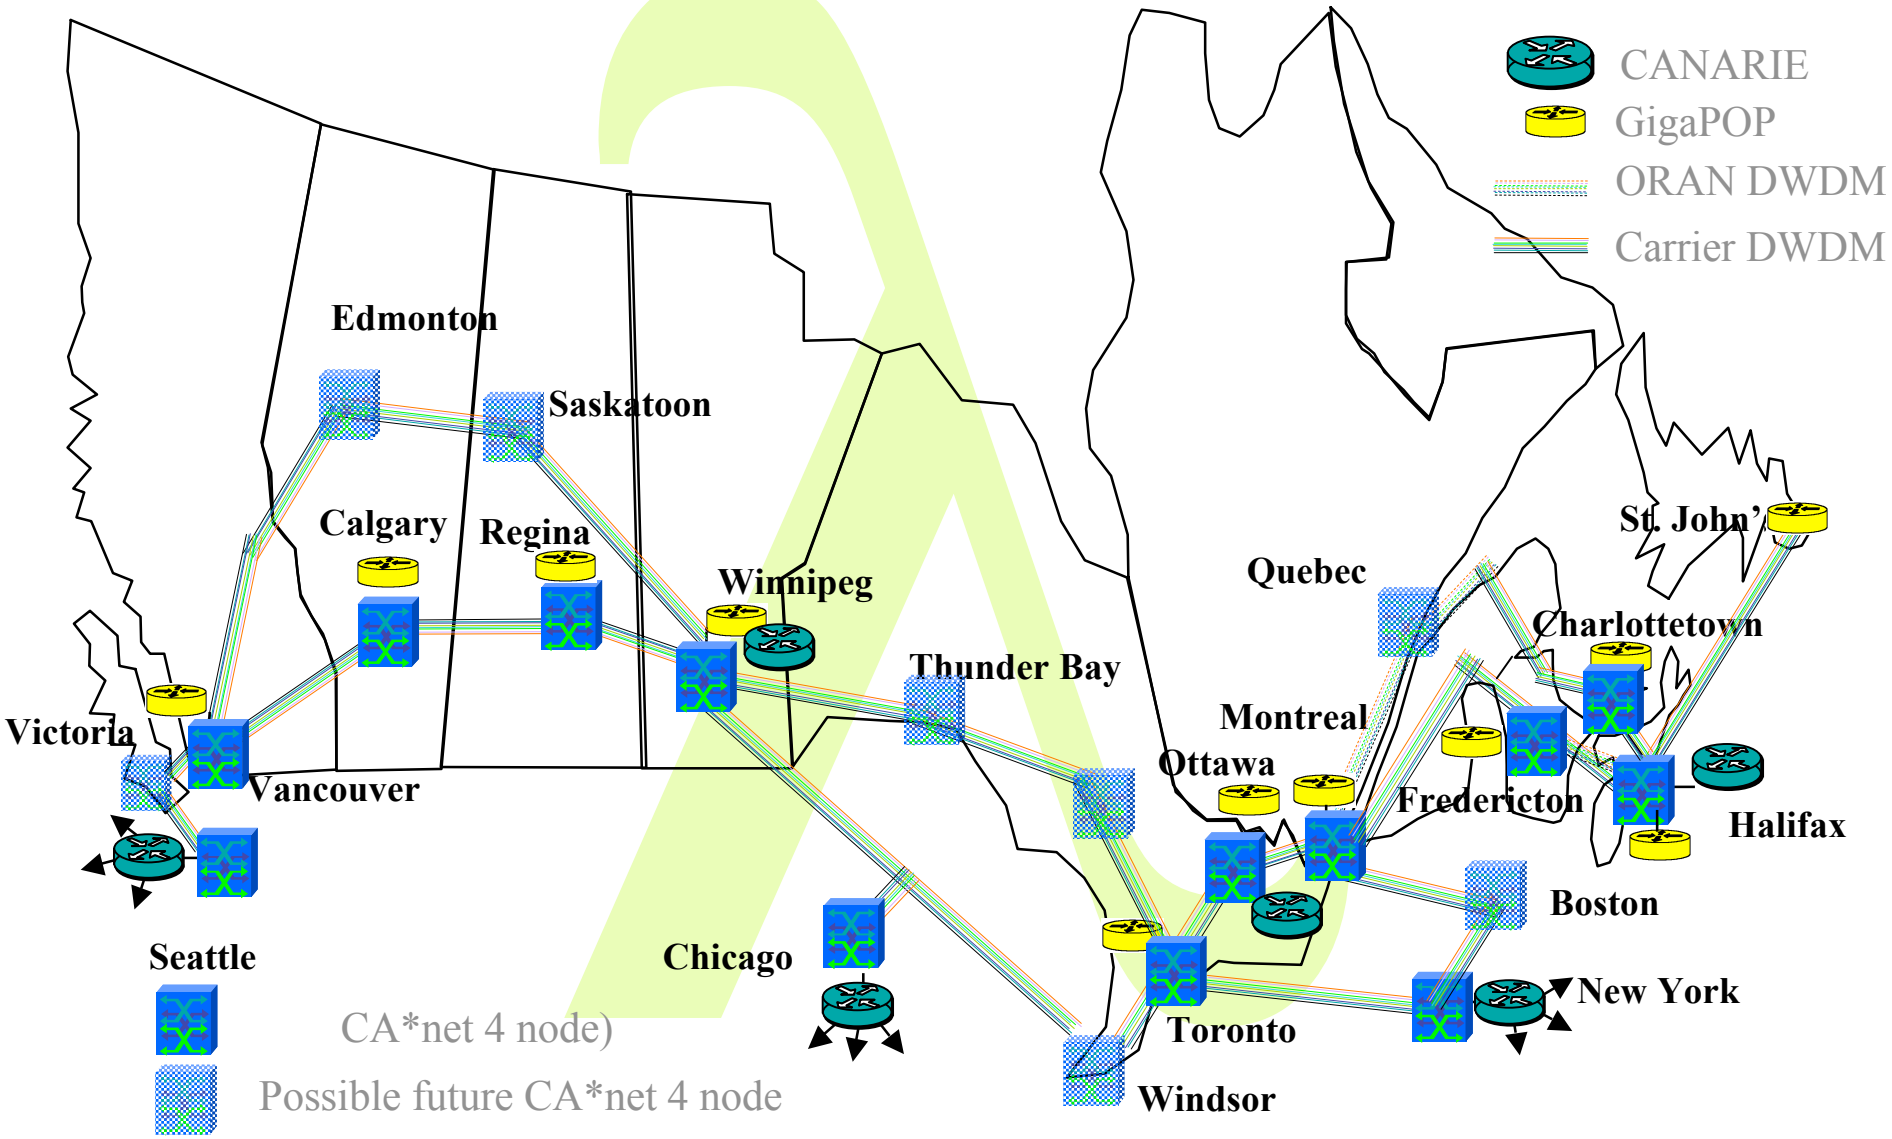

Architectures - L1 - L3

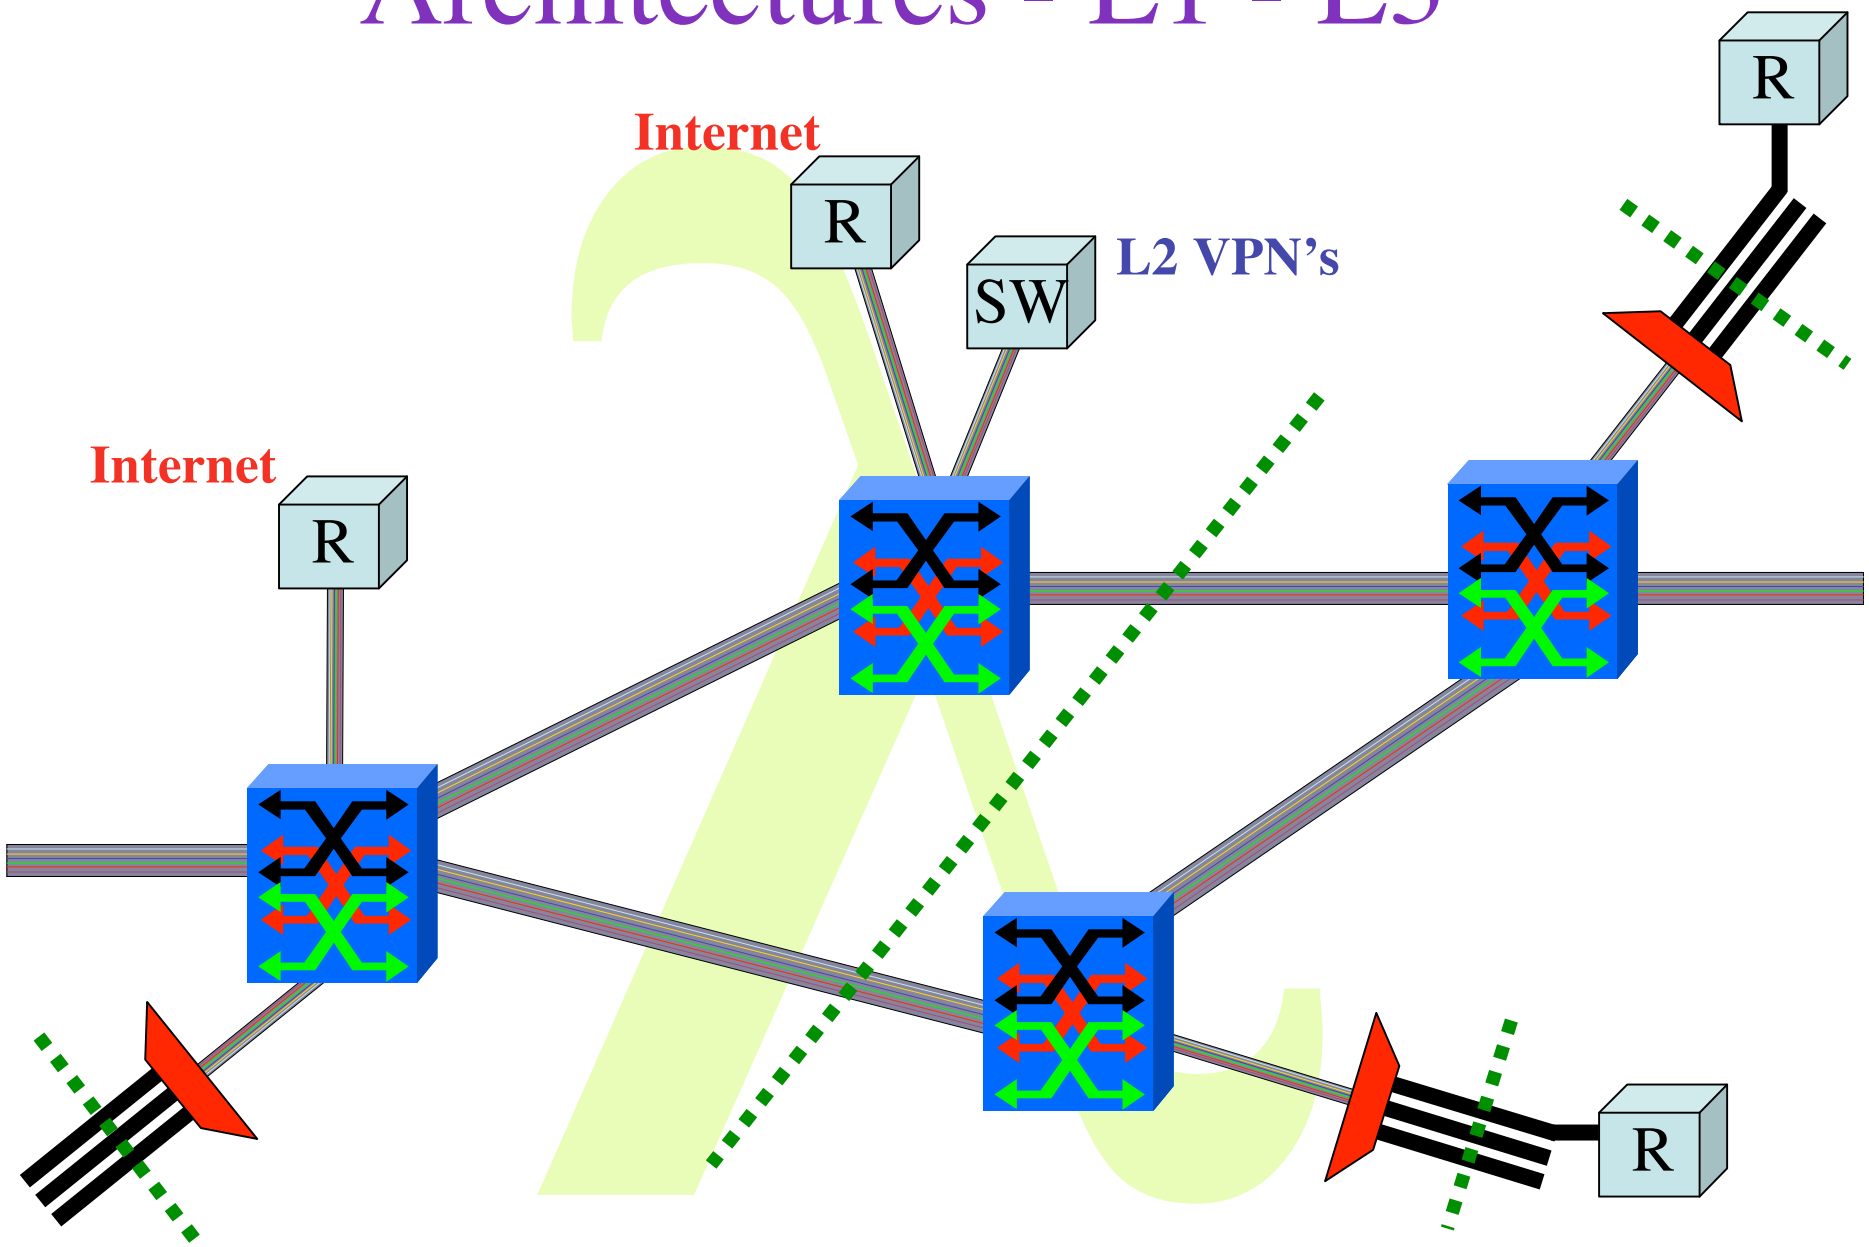

Bring plumbing to the users, not just create sinks in the middle of nowhere

Distributed L2

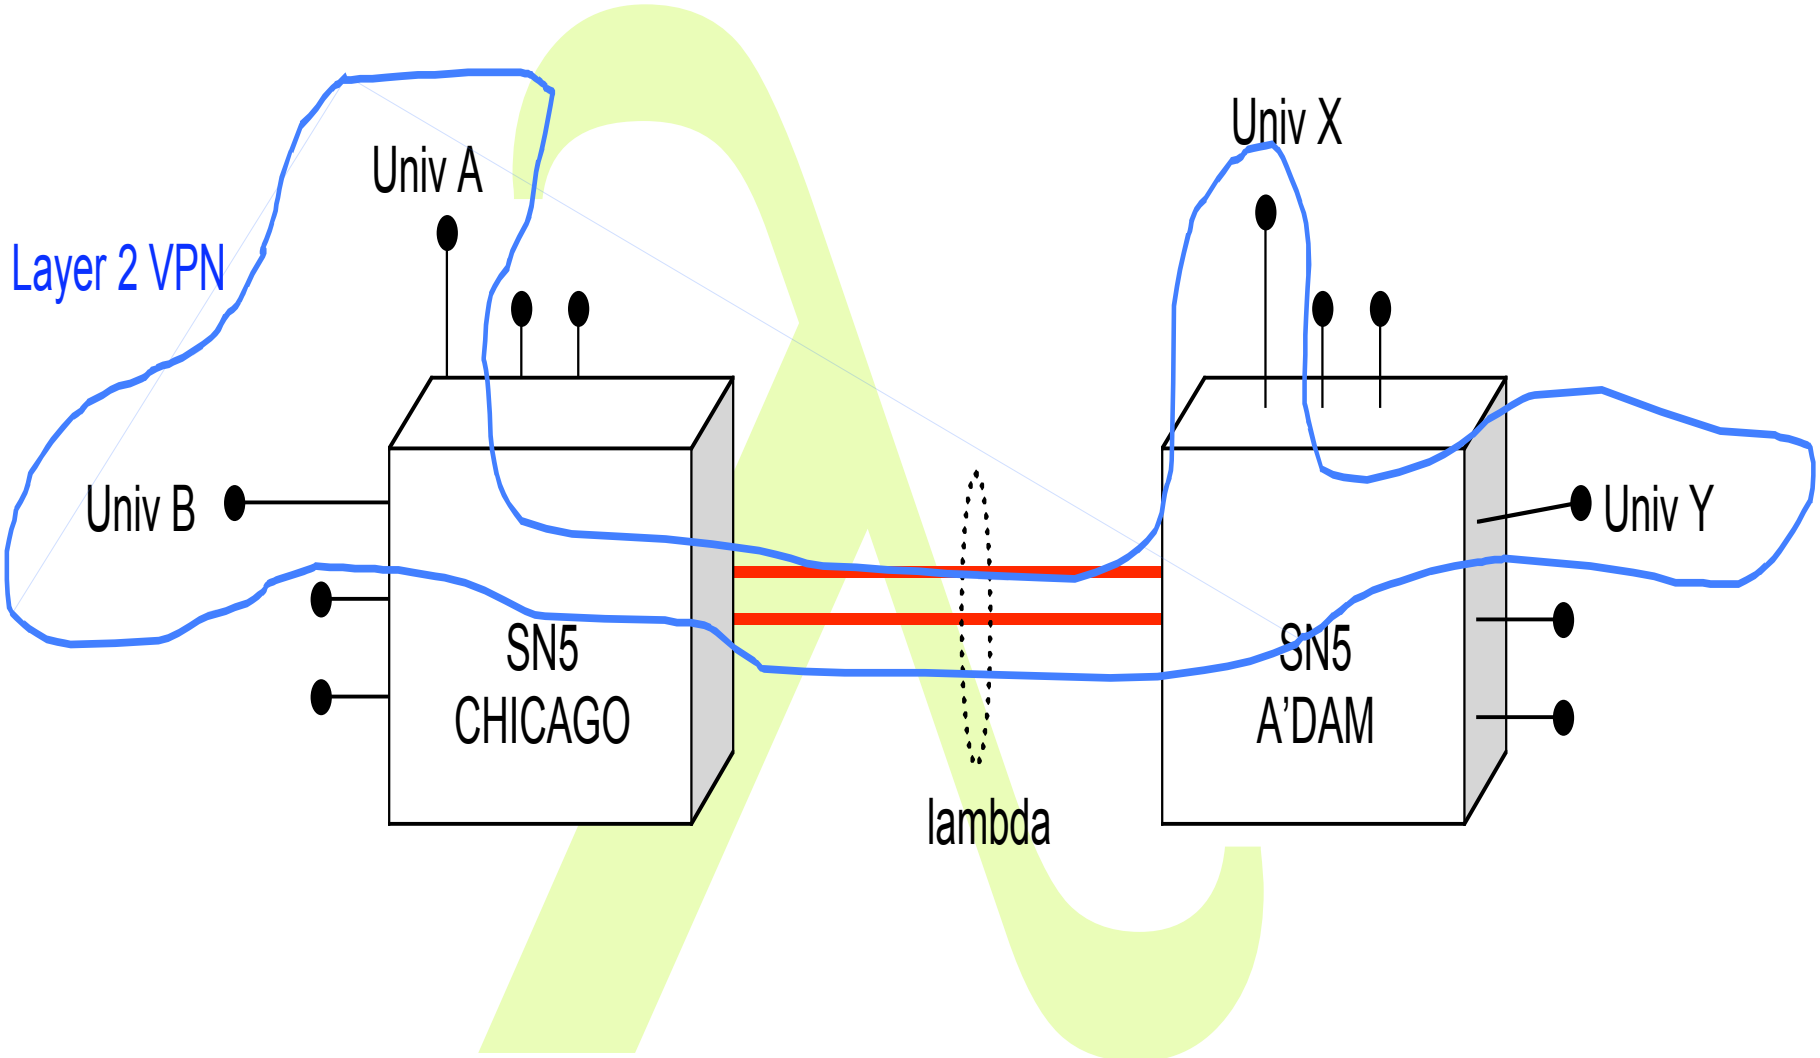

Transport in the corners

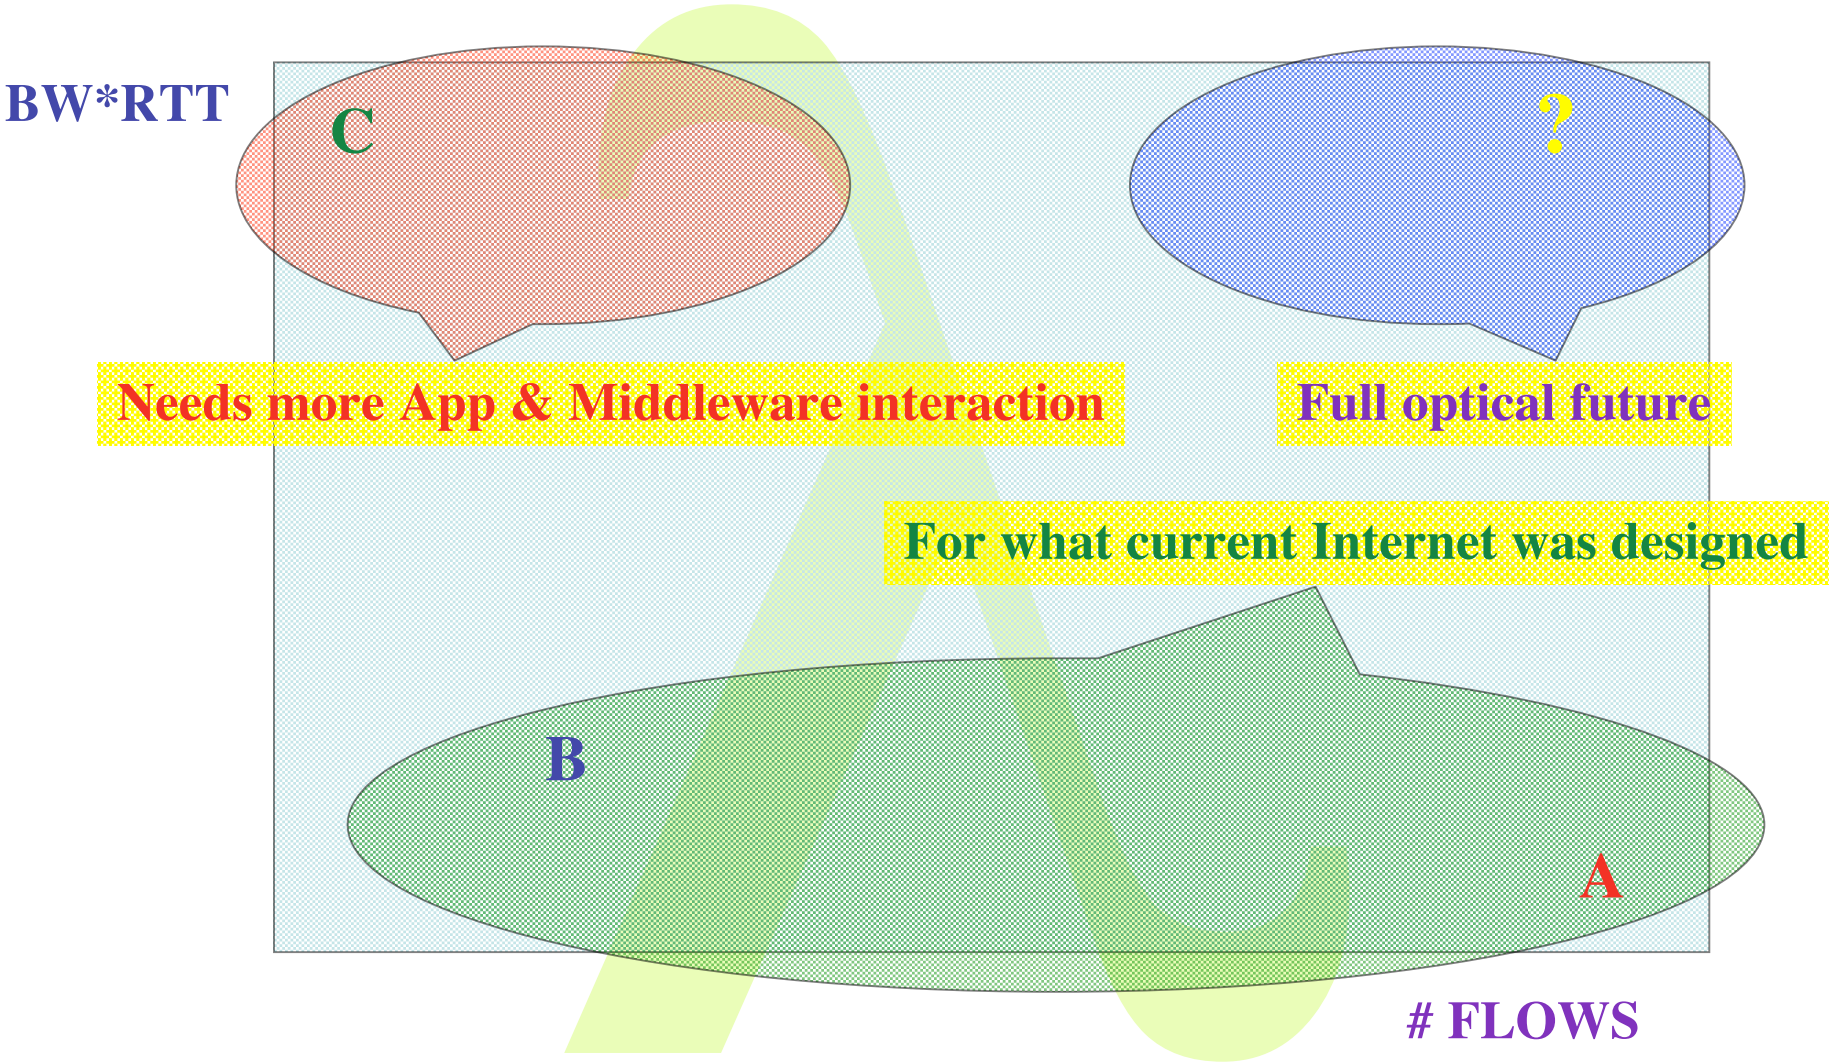

EVL => WCH

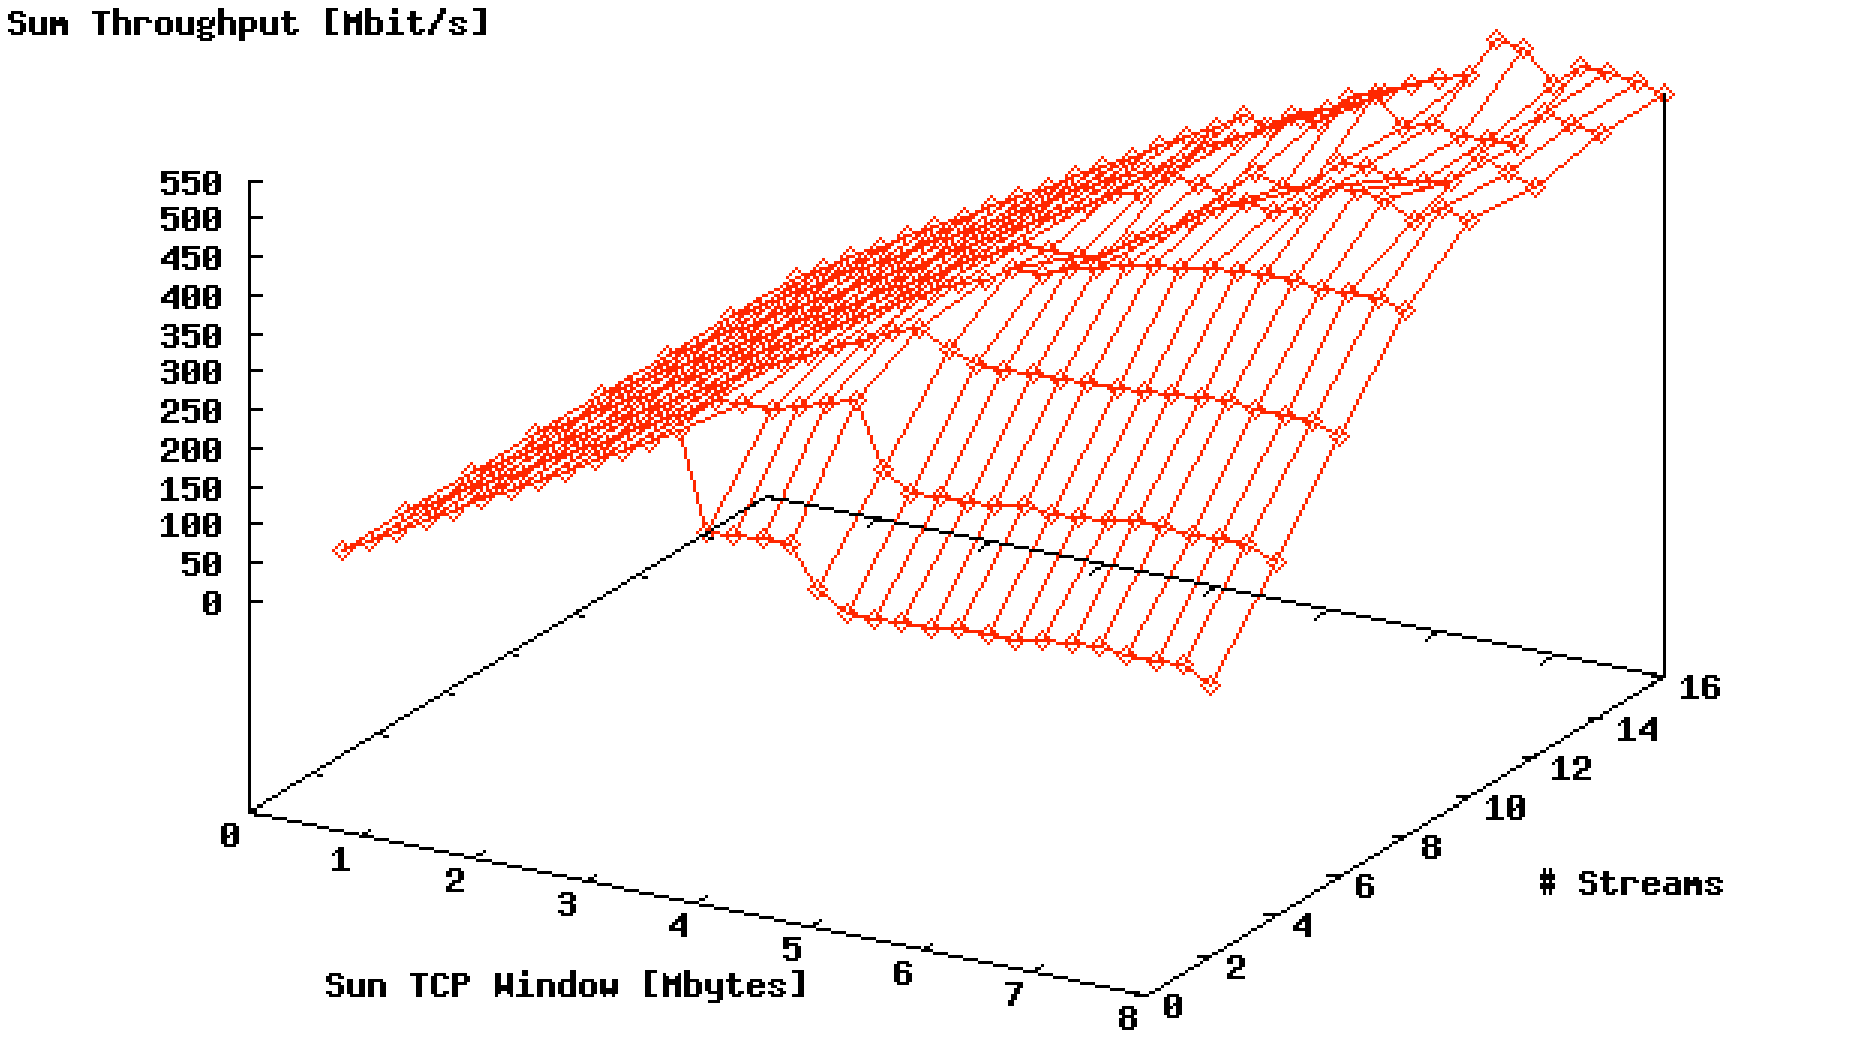
Sum Throughput [Mbit/s]
550
500
450
400
350
300
250
200
150
100
50
0
0
1
2
3
4
5
6
7
8
Sun TCP Window [Mbytes]
0
2
4
6
8
10
12
14
16
Streams

Layer - 2 requirements from 3/4

TCP is bursty due to sliding window protocol and slow start algorithm. So pick from menu:

- *Flow control*
- *Traffic Shaping*
- *RED (Random Early Discard)*
- *Self clocking in TCP*
- *Deep memory*

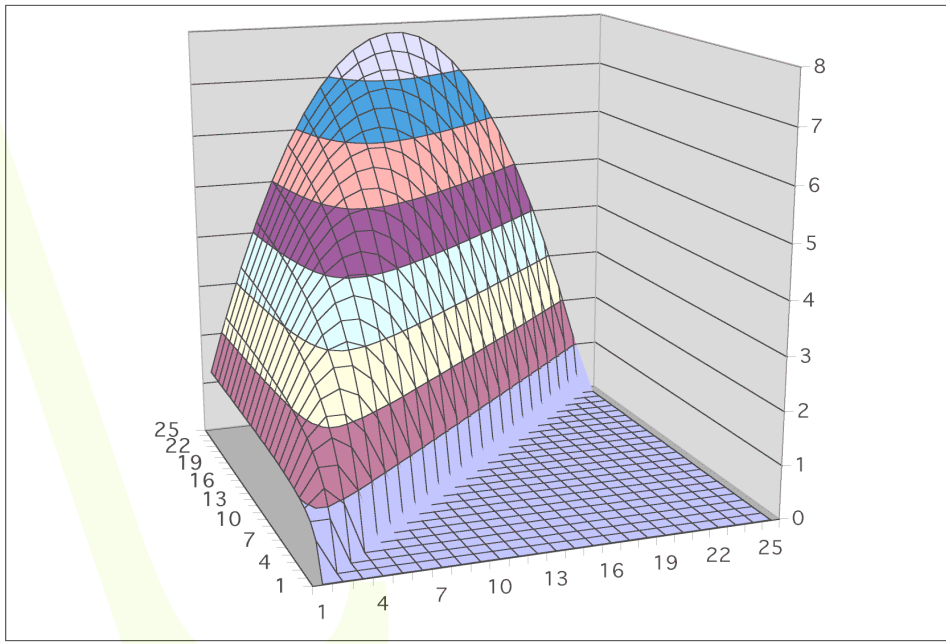

Window = BandWidth * RTT & BW == slow

$$\text{Memory-at-bottleneck} = \frac{\text{fast} - \text{slow}}{\text{fast}} * \text{slow} * \text{RTT}$$

Forbidden area, solutions for s when f = 1 Gb/s, M = 0.5 Mbyte AND NOT USING FLOWCONTROL

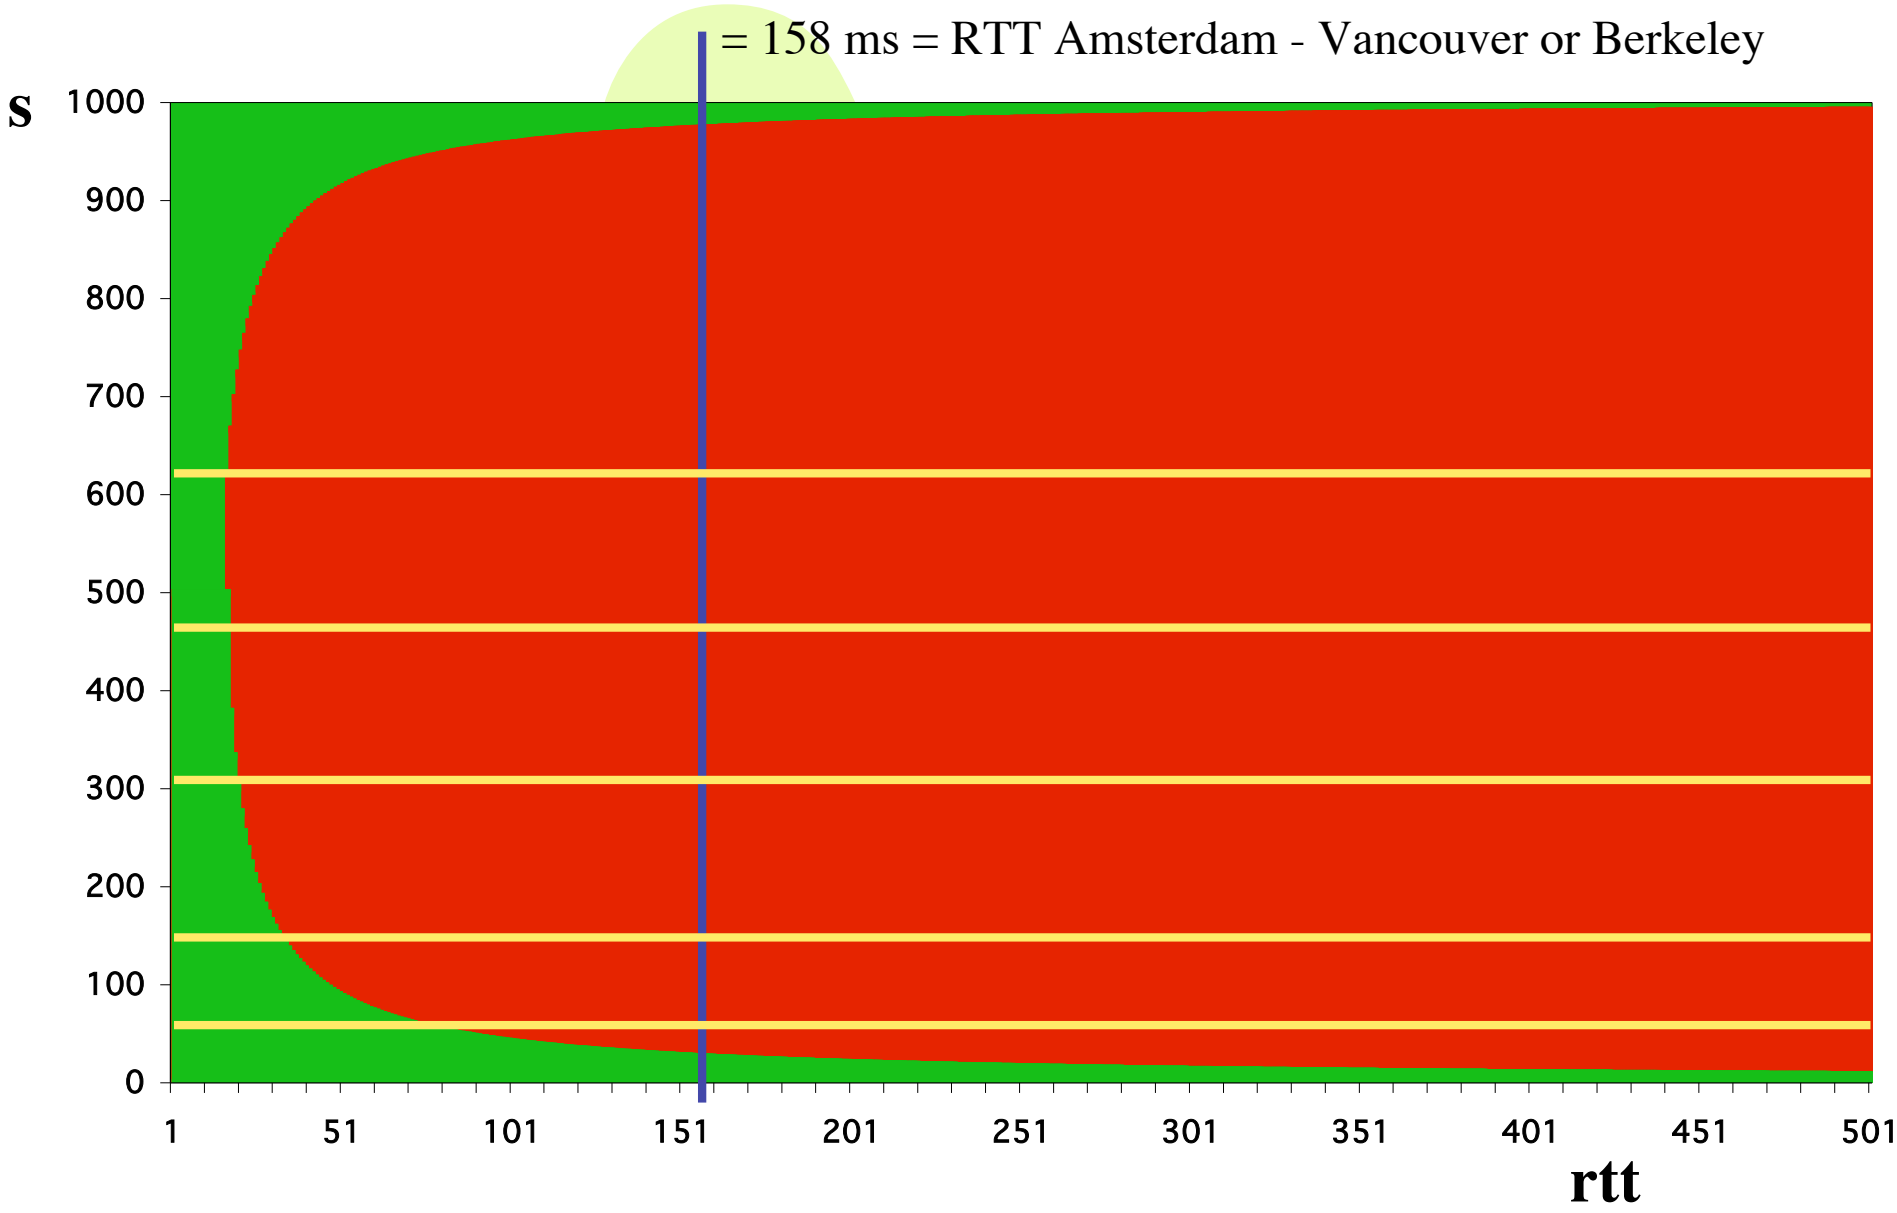

Daisy Chain control model of administrative domains

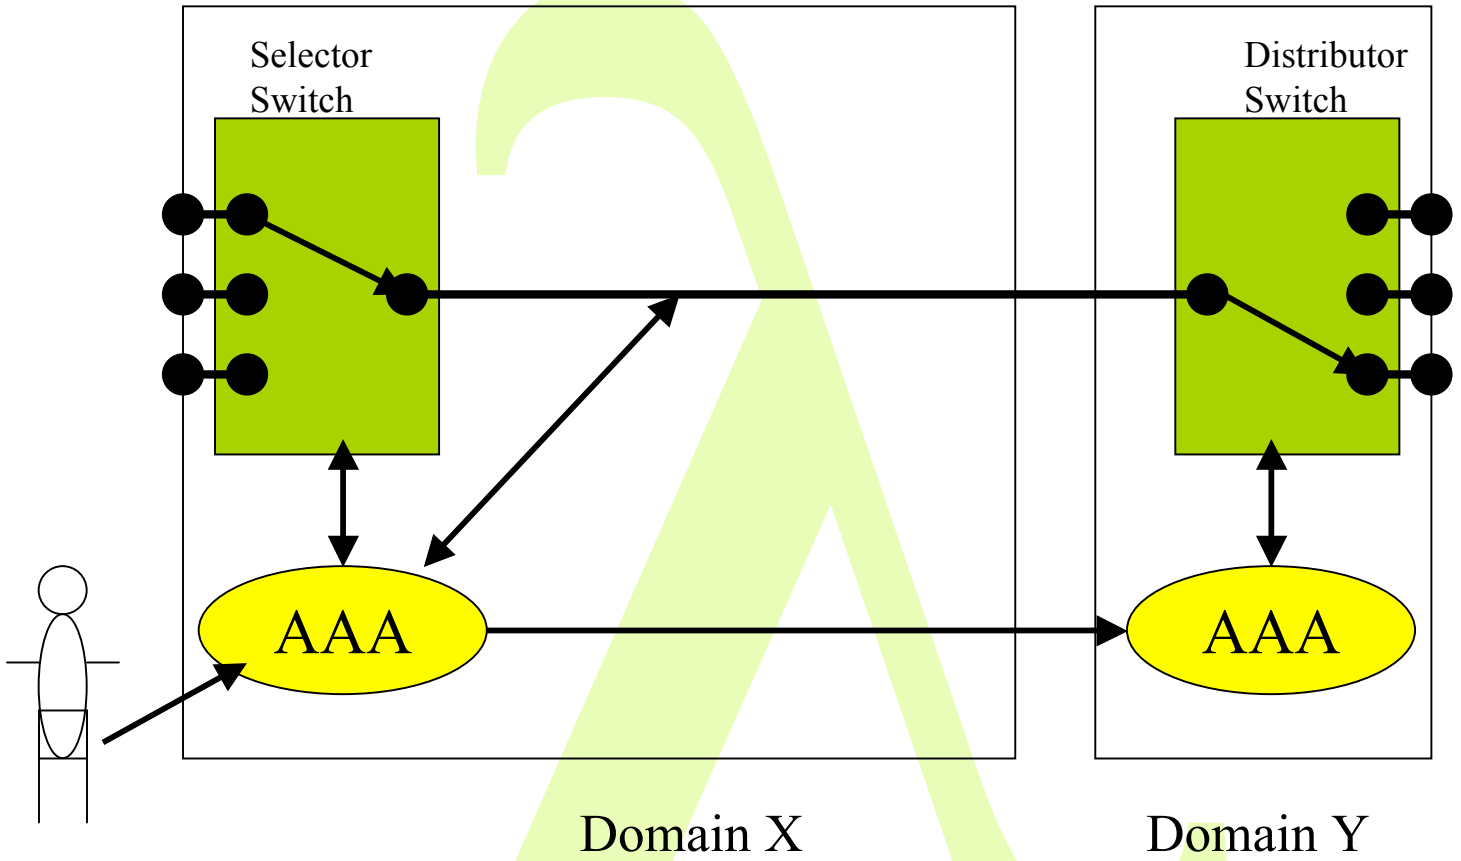

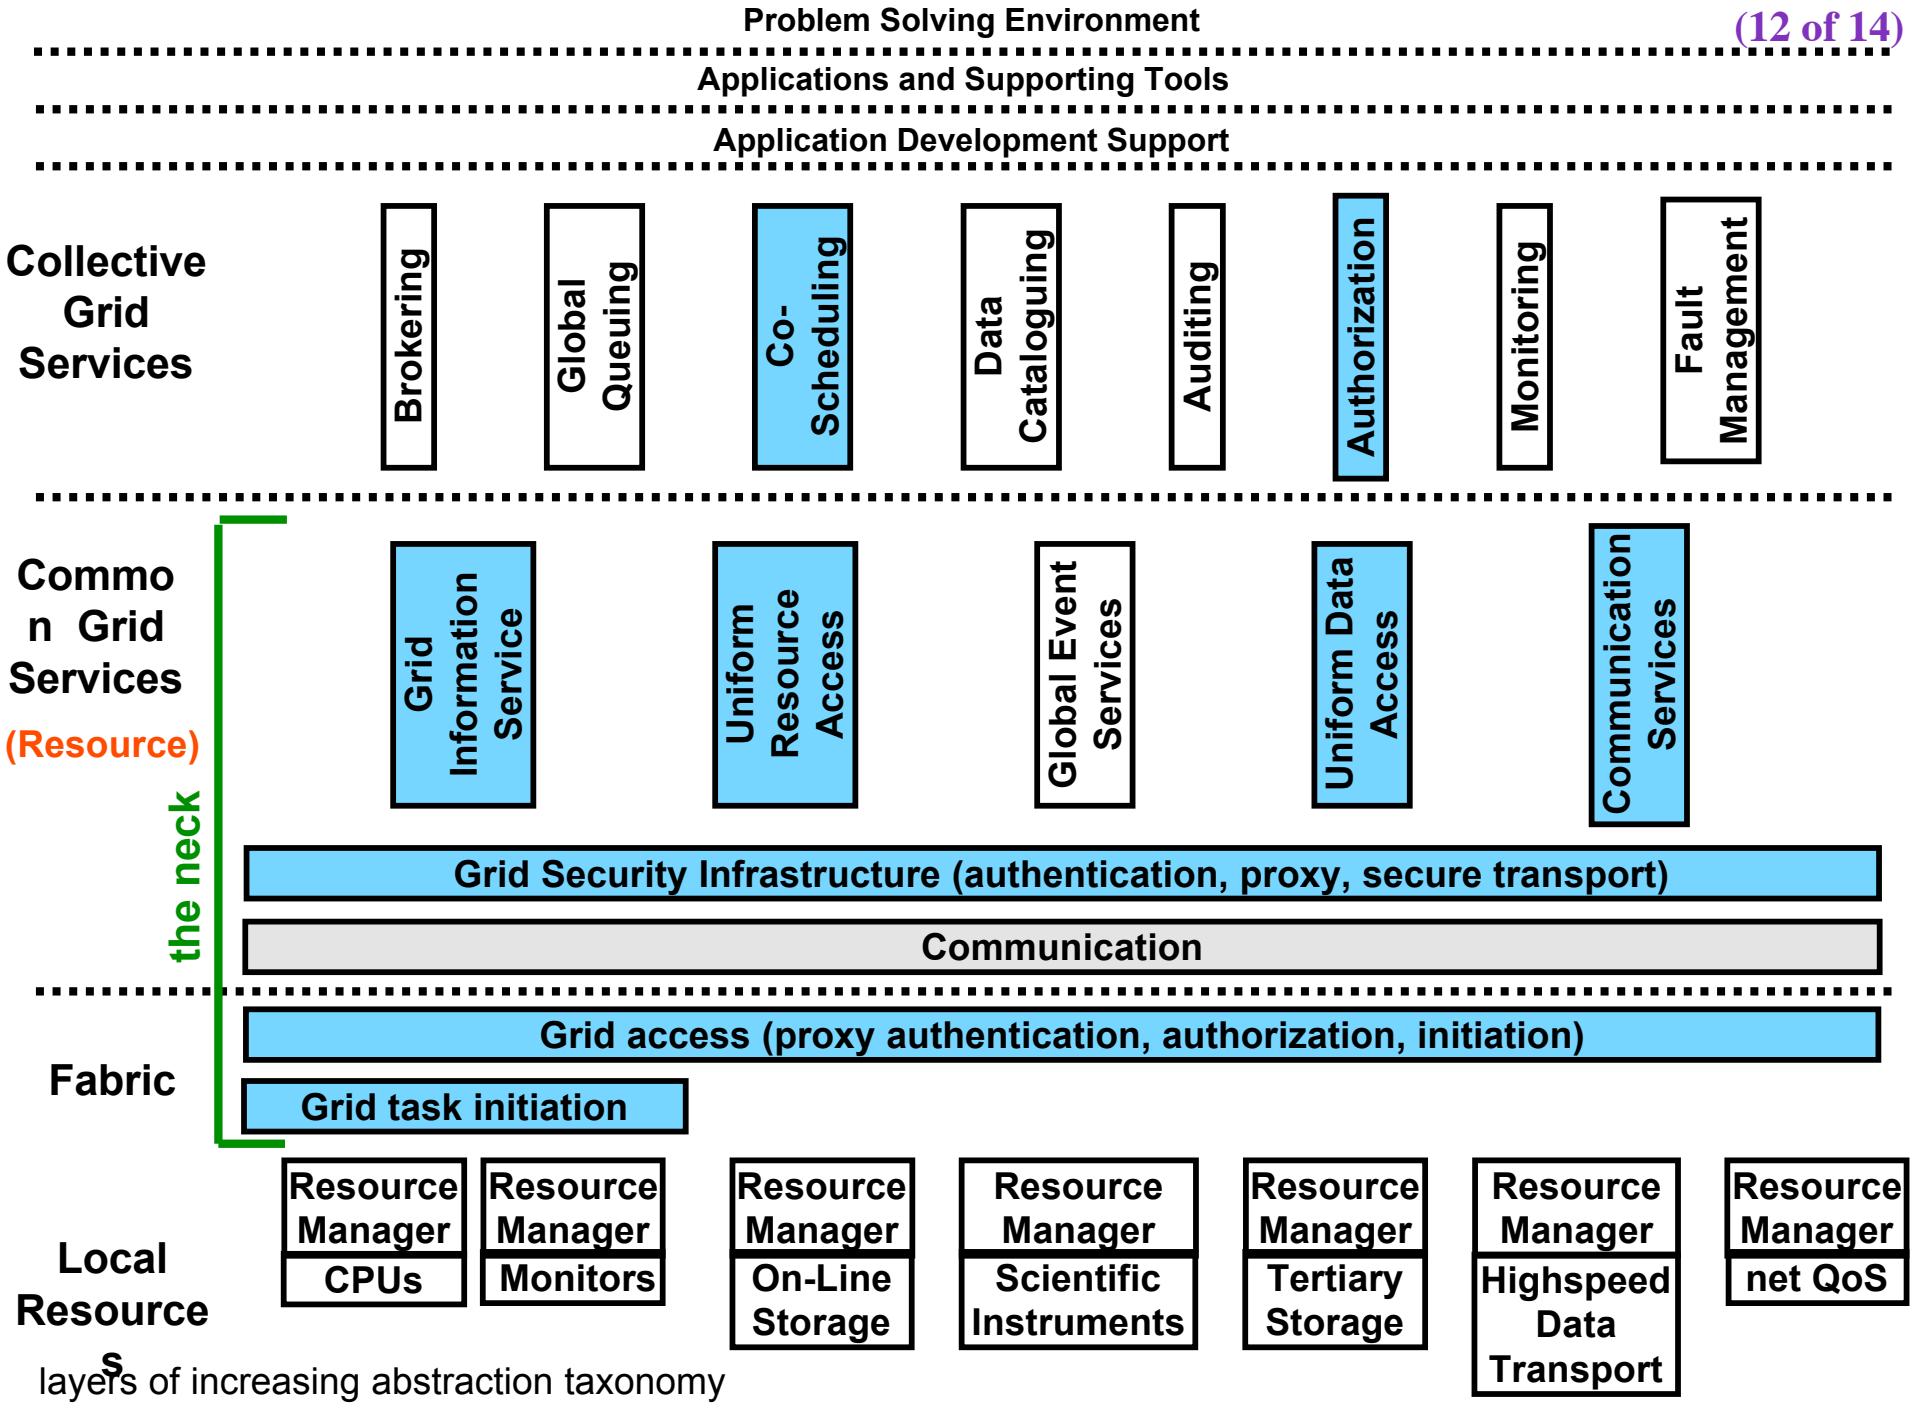

layers of increasing abstraction taxonomy

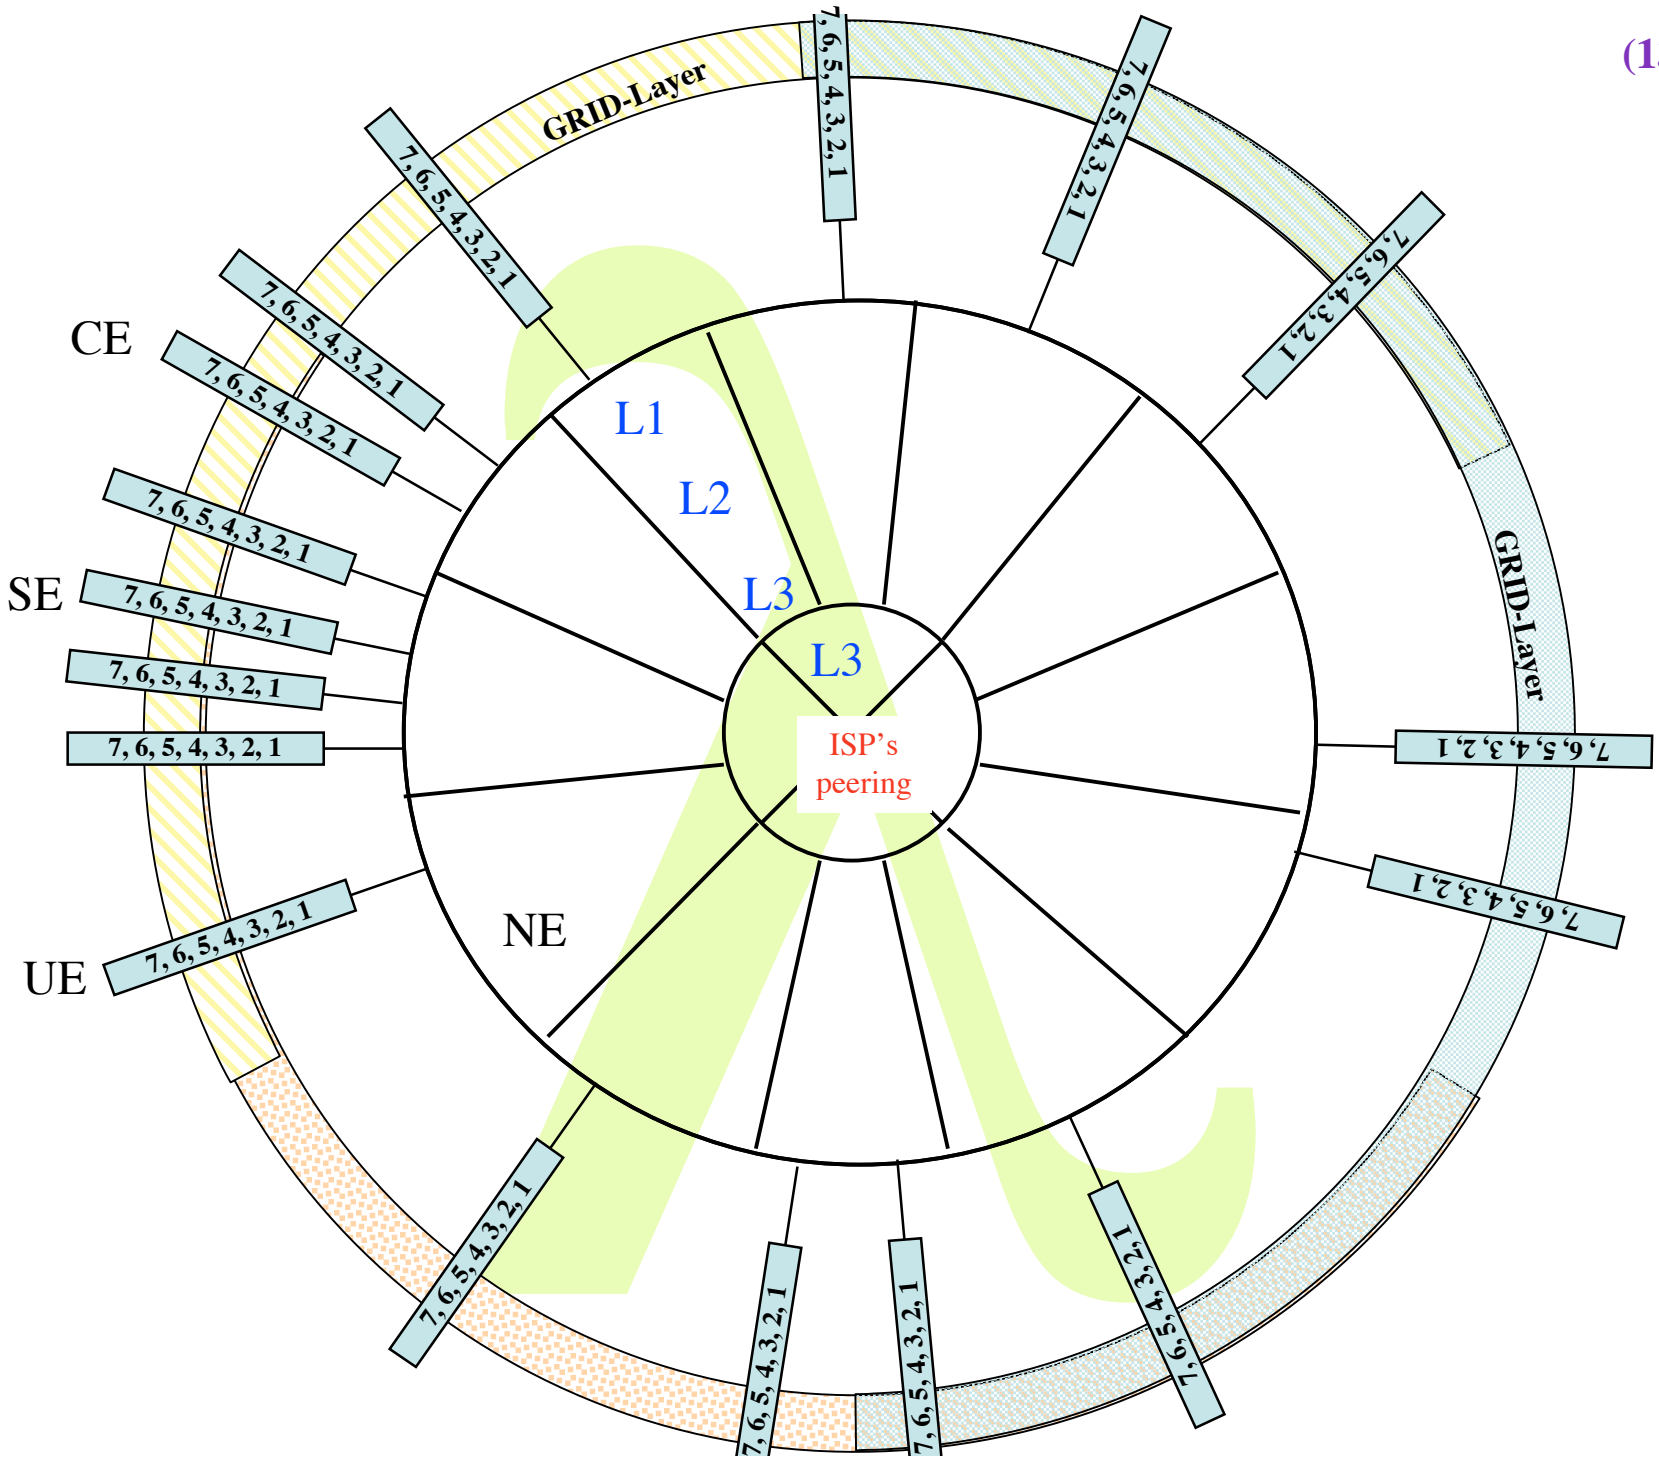
GRID-Layer
GRID-Layer
CE
SE
UE
NE
L1
L2
L3
L3
ISP's
peering
7, 6, 5, 4, 3, 2, 1

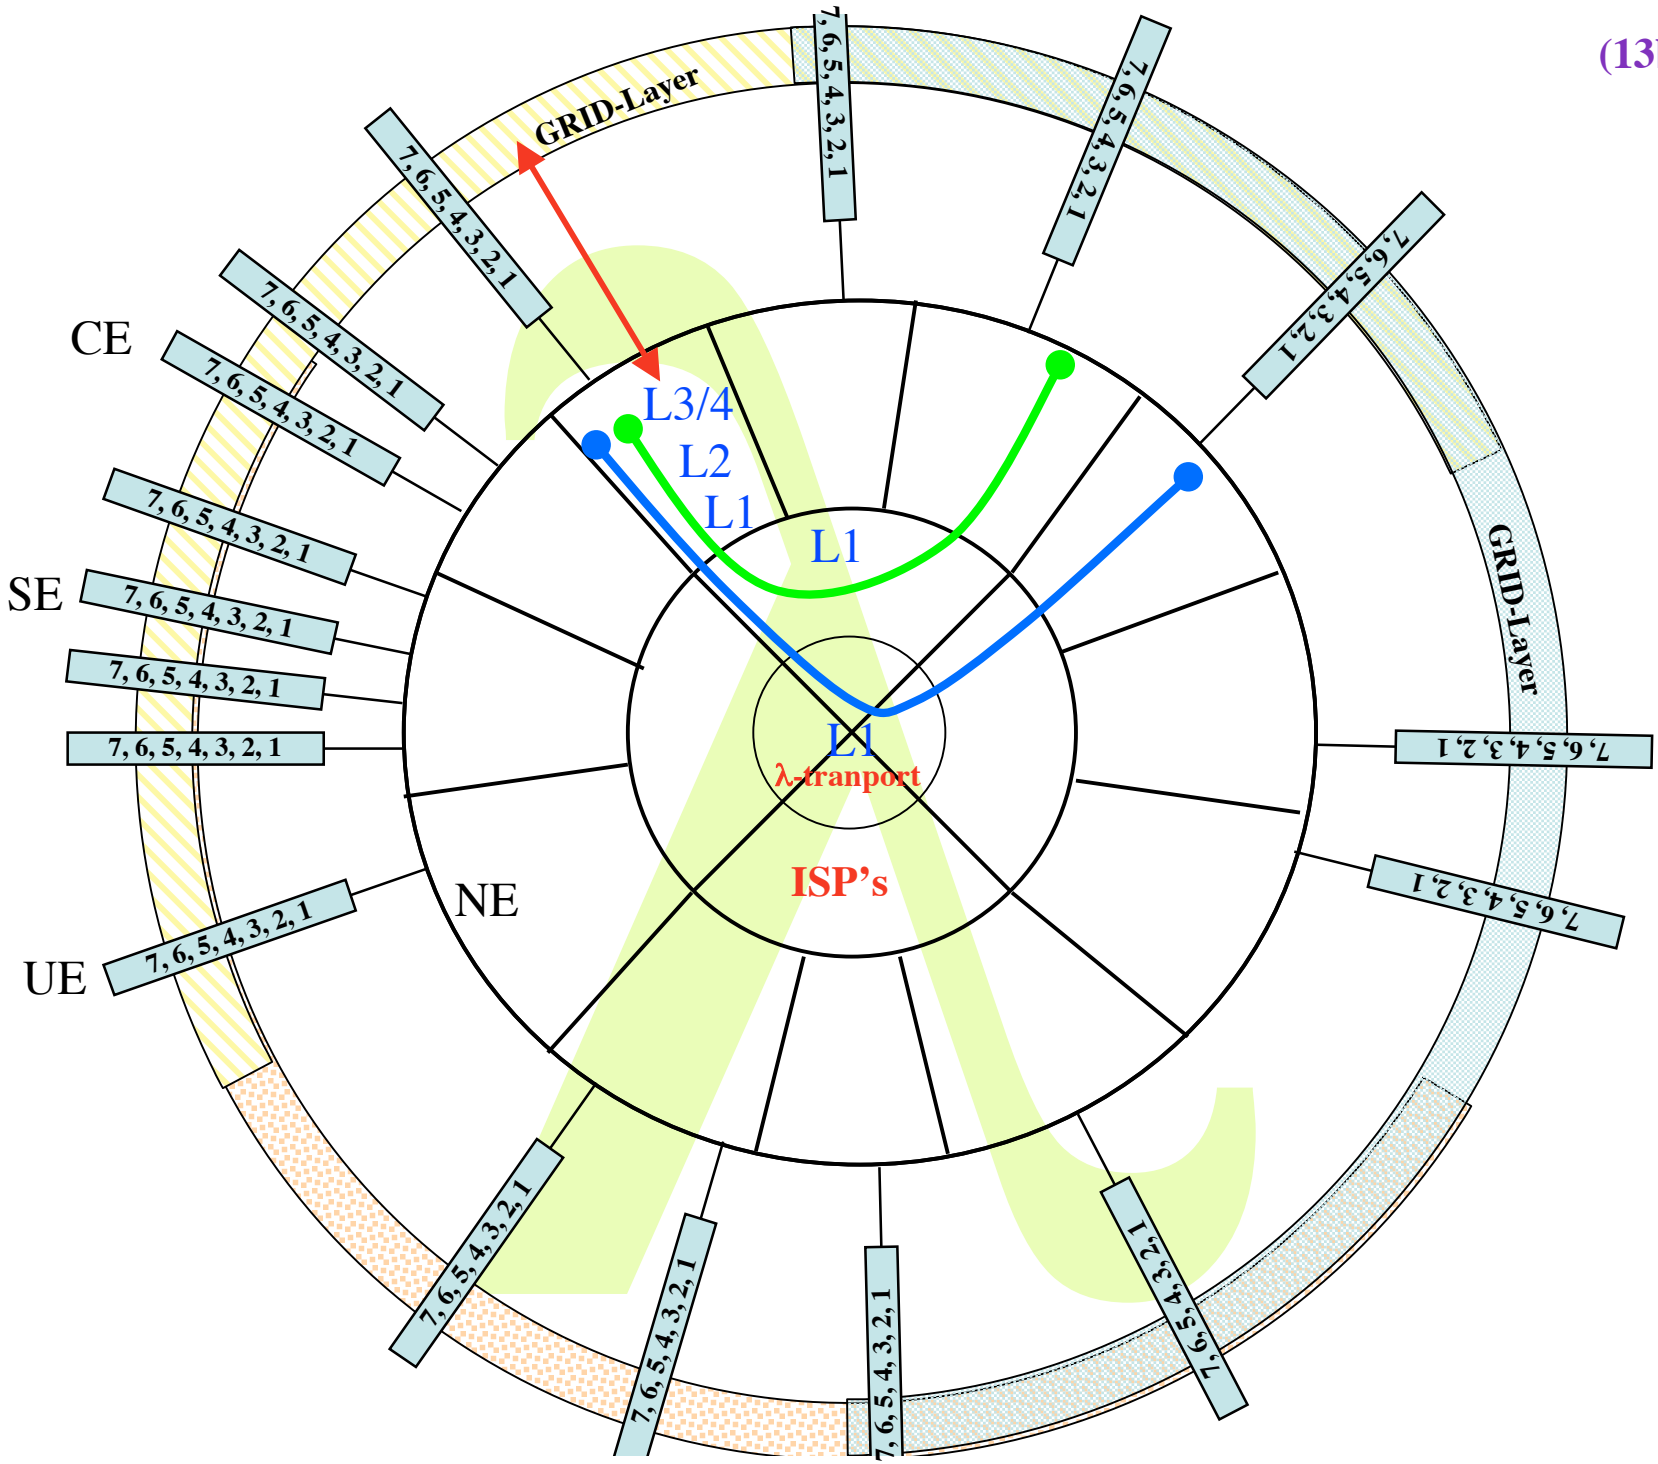
GRID-Layer
7, 6, 5, 4, 3, 2, 1
CE
SE
NE
UE
L3/4
L2
L1
L1
L1
λ-tranport
ISP's
GRID-Layer

Research needed

- Optical devices
- Internet Architecture
- Network Elements as Grid Resources
- Transport protocols get in other corners
- How dynamic must your optical underware be
- Don’t mix trucks and Ferrari’s

Revisiting the truck of tapes

Consider one fiber

- **Current technology allows 320 λ in one of the frequency bands**
- **Each λ has a bandwidth of 40 Gbit/s**
- **Transport: 320 * 40*10^9 / 8 = 1600 GByte/sec**

- **Take a 10 metric ton truck**
  - **One tape contains 50 Gbyte, weights 100 gr**
  - **Truck contains (10000 / 0.1) * 50 Gbyte = 5 PByte**

- **Truck / fiber = 5 PByte / 1600 GByte/sec = 3125 s ≈ one hour**
- **For distances further away than a truck drives in one hour (50 km) minus loading and handling 100000 tapes the fiber wins!!!**

The END

Thanks to

TERENA: David Williams

SURFnet: Kees Neggers

UIC&iCAIR: Tom DeFanti, Joel Mambretti

CANARIE: Bill St. Arnaud

This work is supported by:

SURFnet

EU-IST project DATATAG