GLIF, the Global Lambda Integrated Facility

Measuring the use of the hybrid network

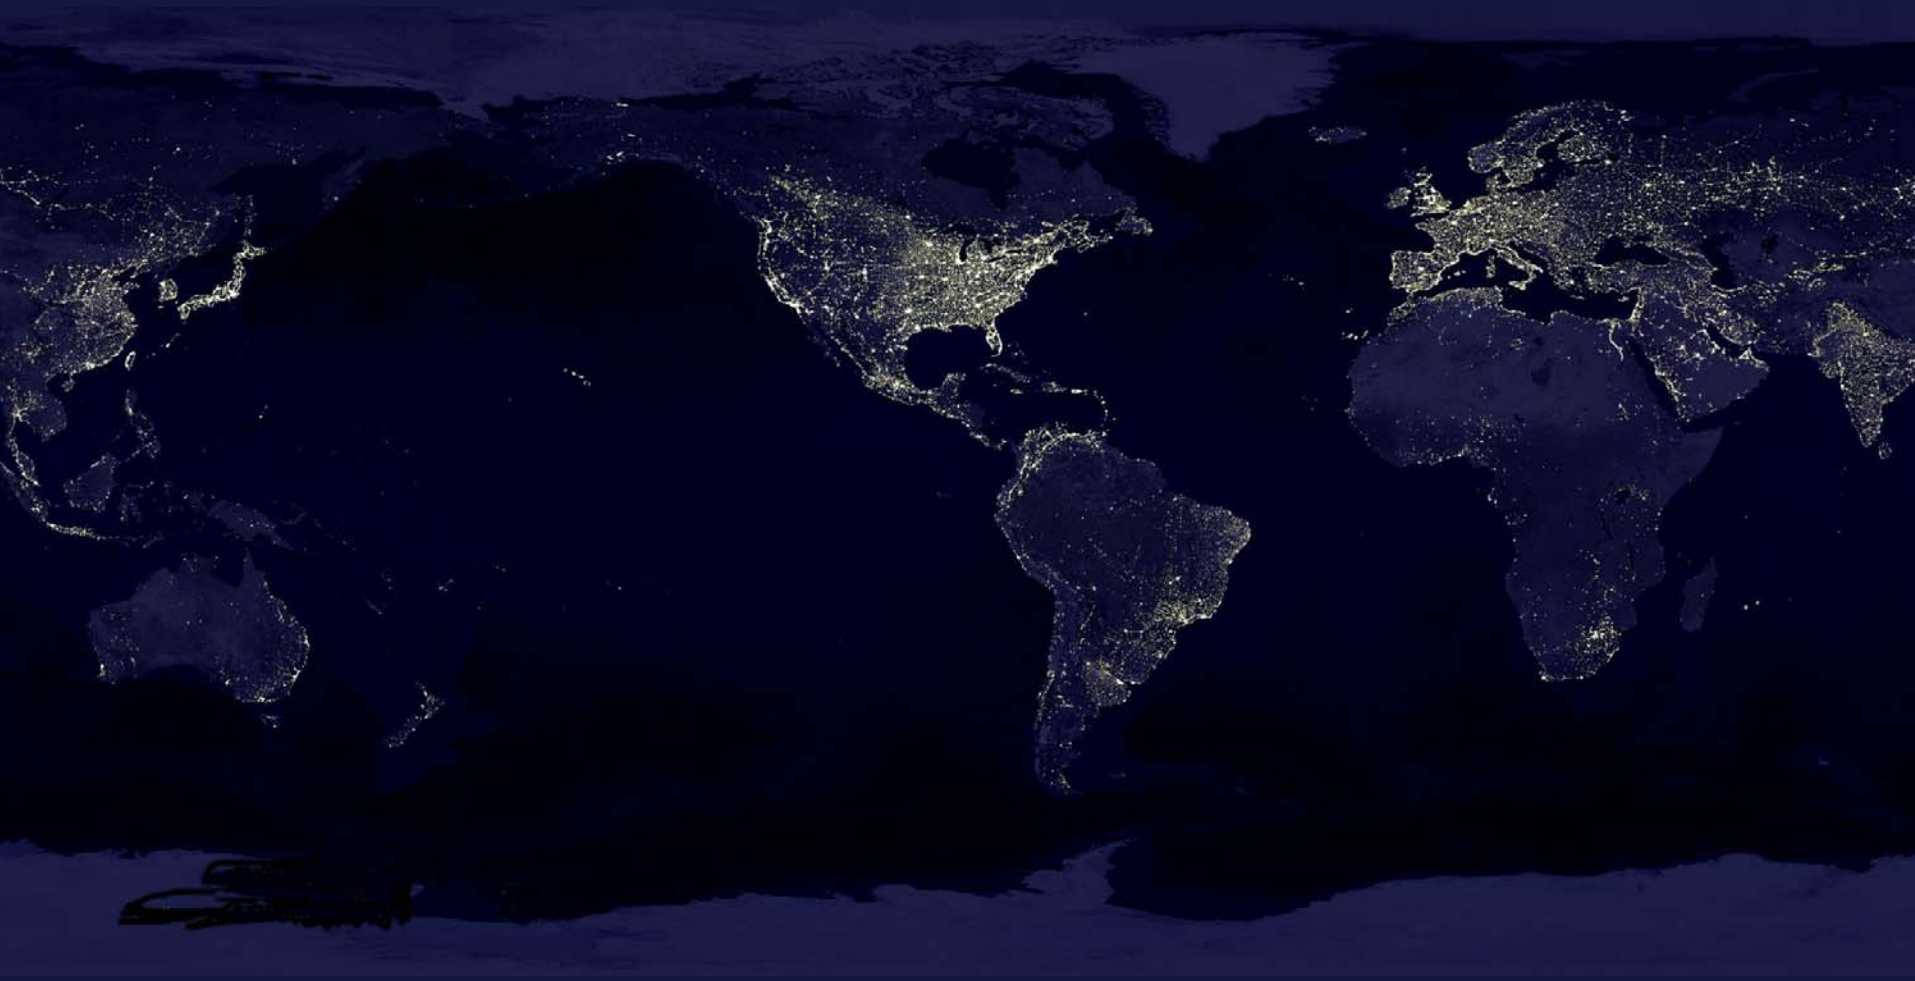

Erik-Jan Bos
Managing Director SURFnet

Minneapolis, MN, USA
February 14 & 15, 2007

The hybrid network enables a range of services

- Previous NREN generations offer only IP services
- The hybrid network offers services at different levels:
  - IP
  - Sub-lambda lightpaths (GFP-F / VCAT / LCAS)
  - Full lambda lightpaths
- Using modern network equipment and control plane, the NREN can make lightpaths more dynamic
  - Easy configuration for static lightpath services
  - User controlled lightpaths: scheduled or dynamic

SURFnet

glif

New metrics needed to measure usage in hybrid network

- Previous generations reported primarily on IP traffic
  - Metrics, measurement methods and tools are well known
  - Byte count provides sufficient information for many purposes
  - Graphs are easy to understand
- Lightpaths are different
  - Connection oriented
  - Static, scheduled, or dynamic

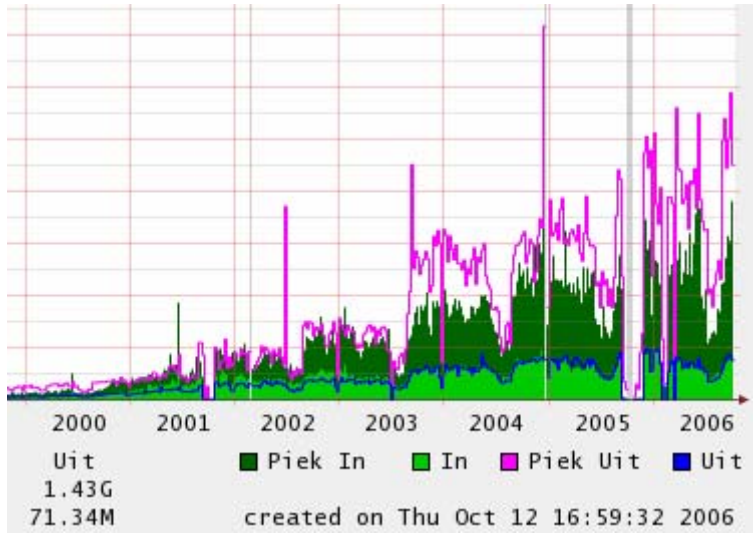

- How to report on services, and on the network as a whole?

SURFnet

glif

Start by mapping services to network capabilities

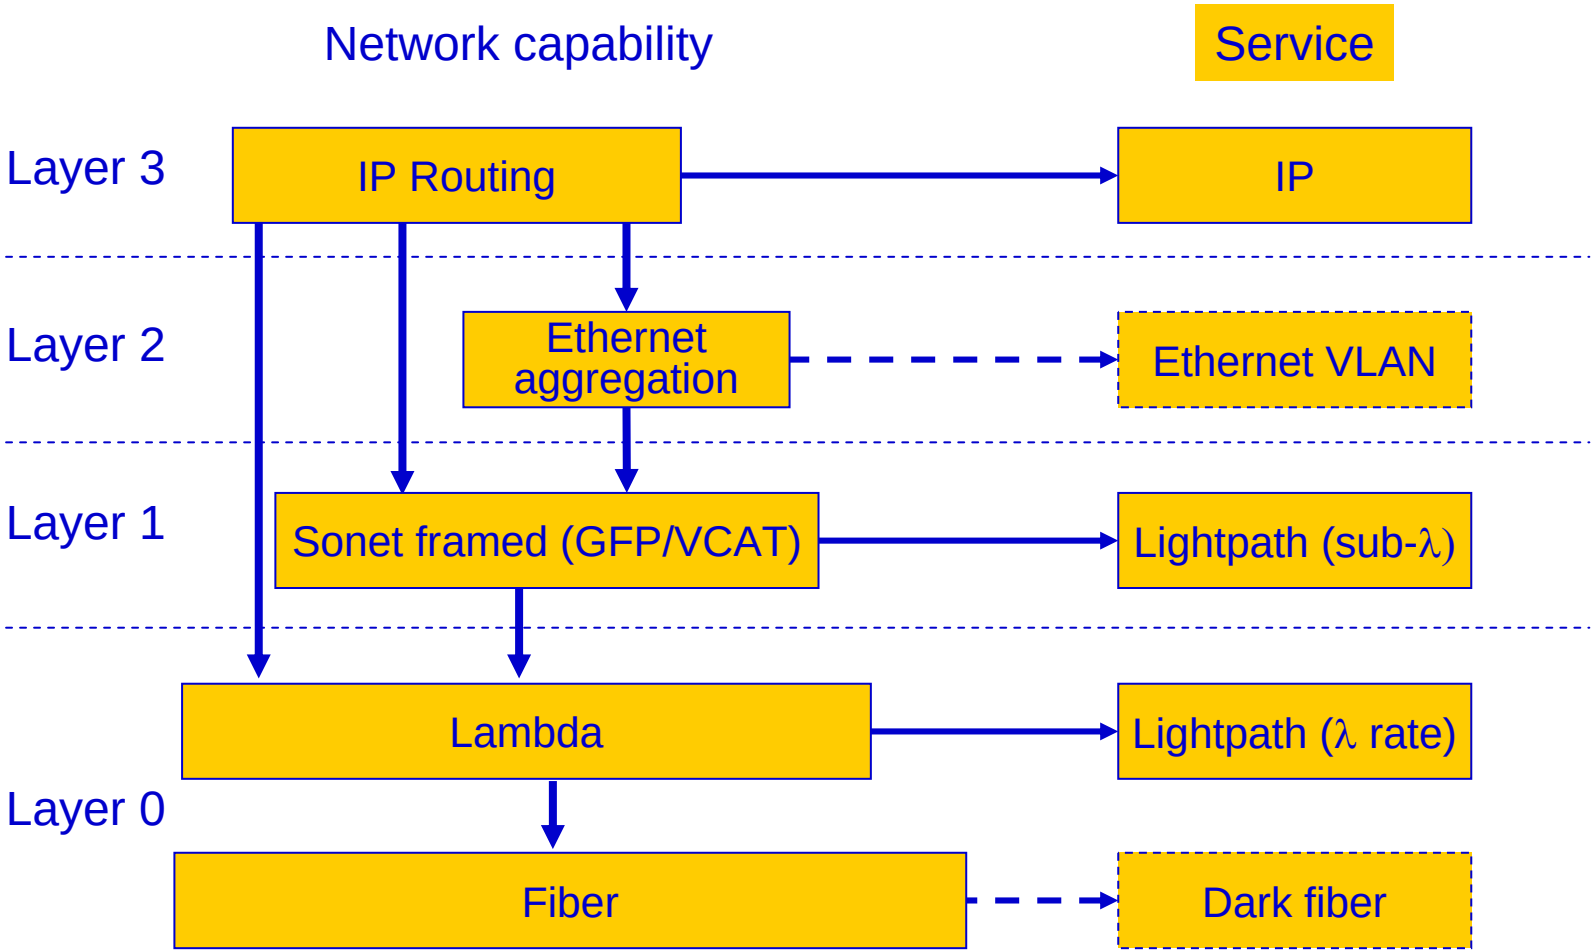

SURFnet

SURFnet service

Service not currently offered by SURFnet

Who is it for?

- General public and policy makers:
  - need to justify funding
- Connected organizations:
  - need to understand use of the service
  - need to justify costs
- Network operator:
  - need to allocate costs (at least internally)
  - need to plan network build-out
  - need to make routing decisions

SURFnet

glif

Different stakeholders, different levels of complexity

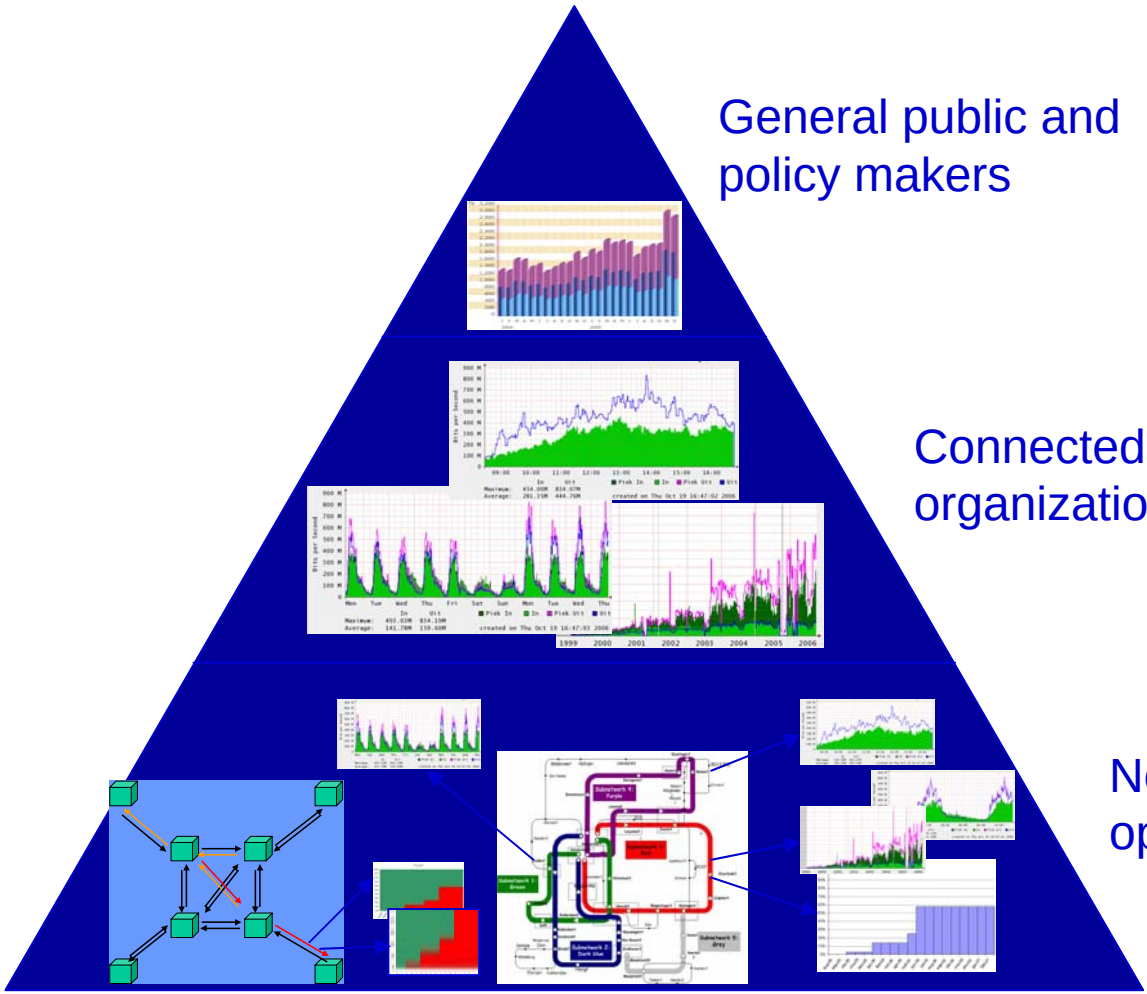

General public and policy makers

Connected organizations

Network operator

SURFnet

glif

General public & policy makers: justify spending money

- This group requires simple metrics: easy to explain, easy to understand

- Clear link to the value delivered through the network

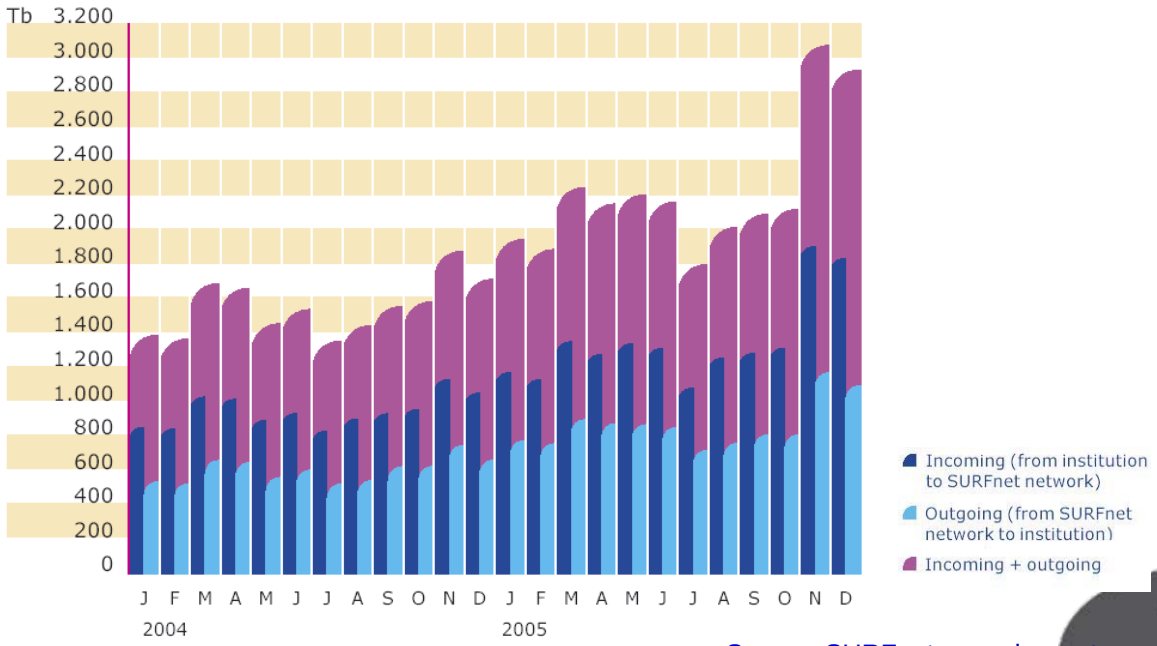

Source: SURFnet annual report

General public and policy makers: keep it simple

- **Simply count all the bytes transported by the network**

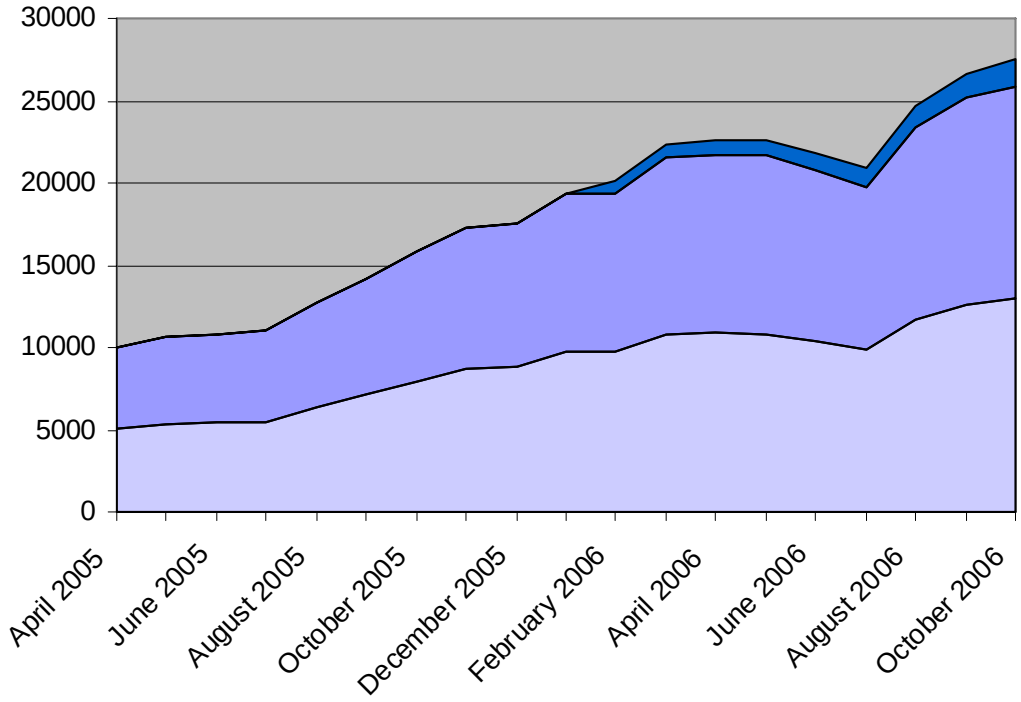

(hypothetical example)

SURFnet

glif

General public and policy makers: keep it simple

- …and report total capacity provided for all services

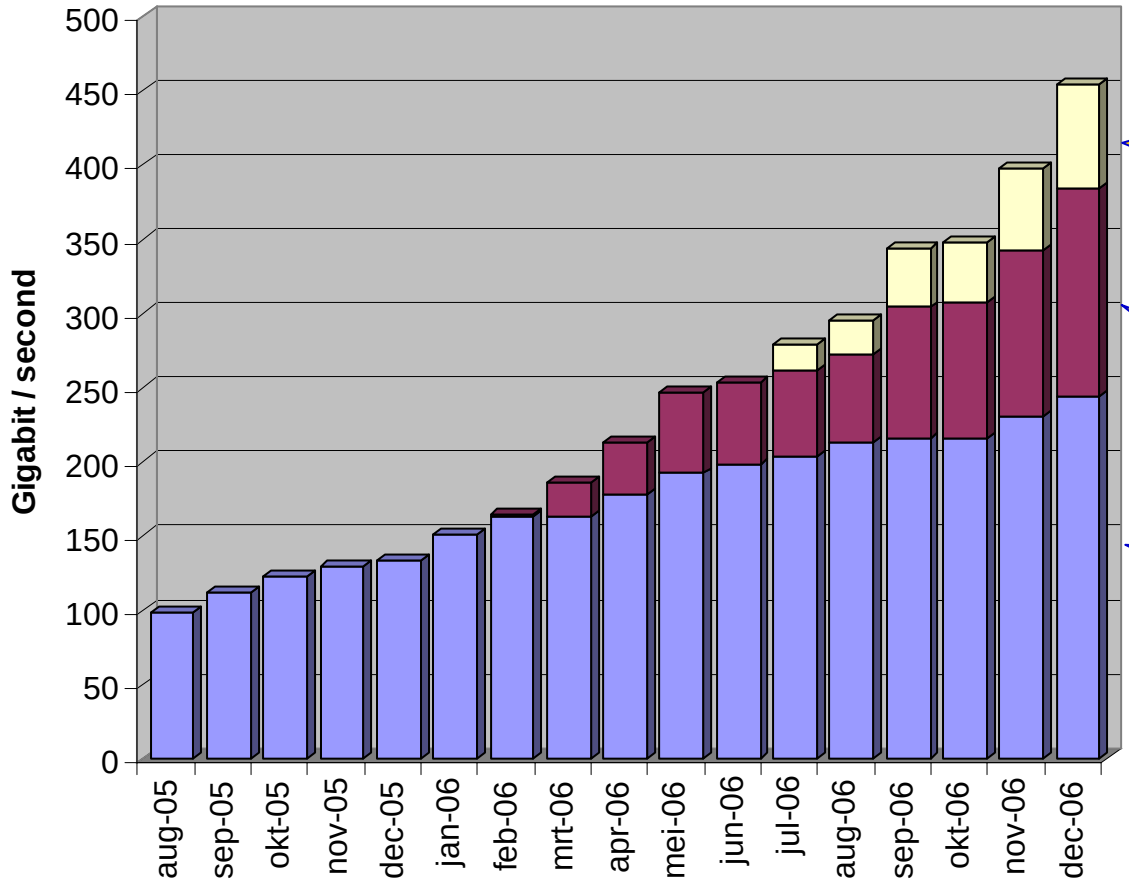

Dynamic lightpaths (bandwidth x % of time active)

Total bandwidth static lightpaths

Transport capacity for IP

(hypothetical example)

Connected organizations: understanding service delivered

- Want to understand price/performance
- Want to pinpoint capacity bottlenecks
- May want to charge internal user groups

- Therefore:
  - More detail, less aggregation
  - Easy access to information
  - Different timescales

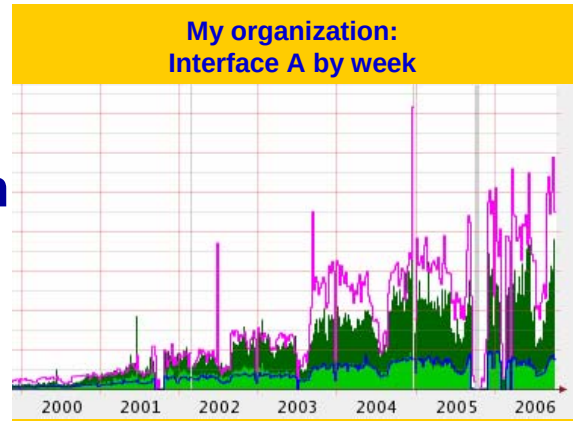

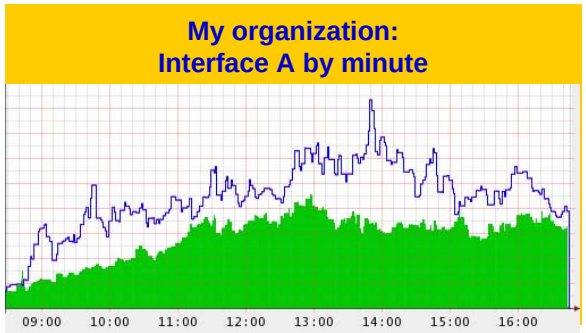

SURFnet

glif

Connected organizations: more detail, less aggregation

- Separate reporting per service delivered:
  - IP:
    - Packets and octets, peak and average, at different timescales
    - Packet loss
    - Flow monitoring above the IP layer?
  - Static lightpath:
    - Interface capacity
    - Transport bandwidth delivered to the interface
    - Octets transported?
  - Dynamic lightpath:
    - Interface capacity
    - Transport bandwidth requested and granted
    - Percentage of time active, number of requests, number denied
    - Octets transported?

SURFnet

glif

Network operator: accounting and charging

- Most NRENs have very simple pricing schemes
  - no need for detailed charging per customer
- However, they still need to make pricing decisions based on actual facts
- Therefore they need to understand:
  - the use of each service
  - the resulting claim on resources
  - and the cost of each resource

SURFnet

glif

Network operator: mapping to resource usage

- Report on the services delivered
  - Same metrics as provided to connected organizations
- Infrastructure
  - Static resources used by each service
    - e.g. IP connection using *n* VC-4's on a given link
  - Variable resources used by each service
    - e.g. dynamic lightpath using *n* VC-4's on average, *m* peak, *x* % of the time

SURFnet

glif

Network operator: mapping service to use of infrastructure

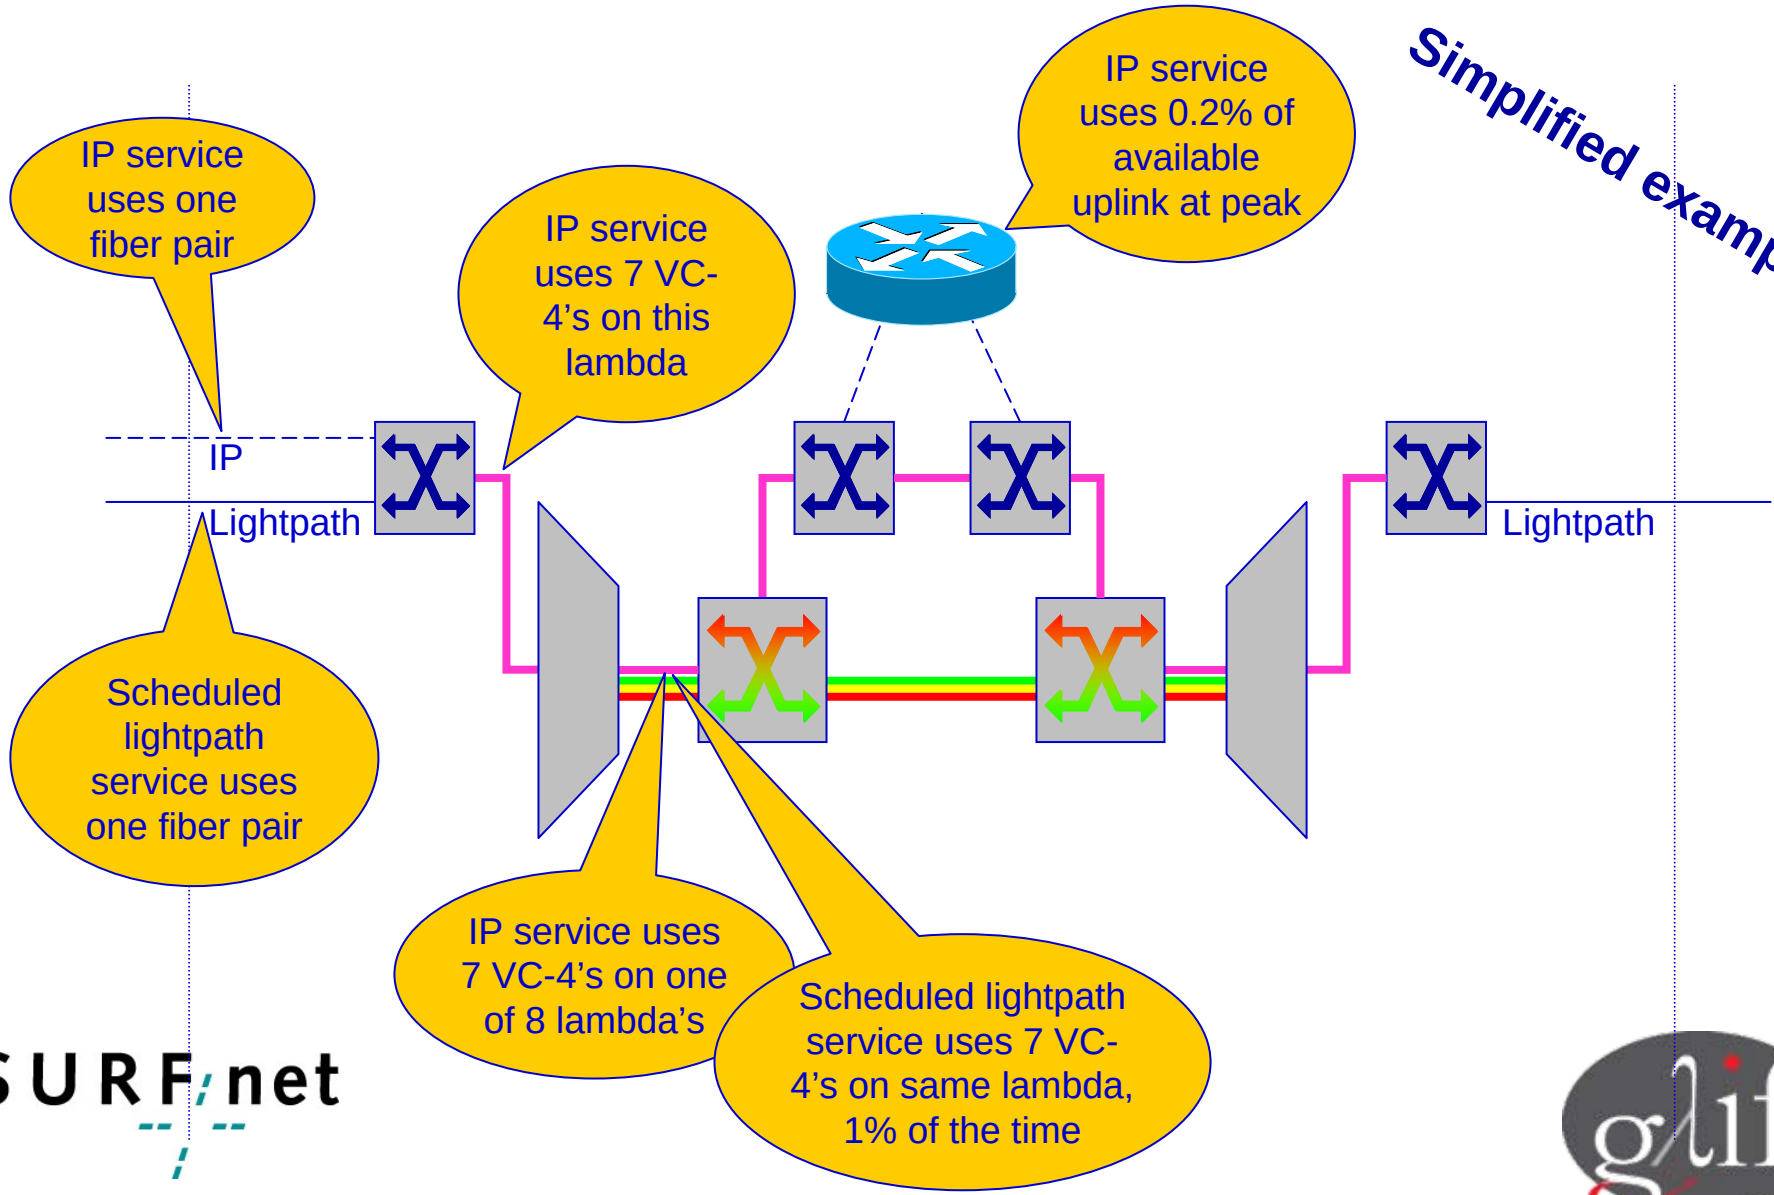

Network operator: network planning and design

- NREN needs to pinpoint bottlenecks…
  - preferably before they become bottlenecks!
- … at each layer in the network
  - but not in isolation: each layer is related to the others
- … at each device and each link
  - but not only at the local level: bottlenecks may be resolved by re-routing
- … at different timescales
  - busiest hour, busiest day...

SURFnet

glif

Network operator: local capacity monitoring at each layer

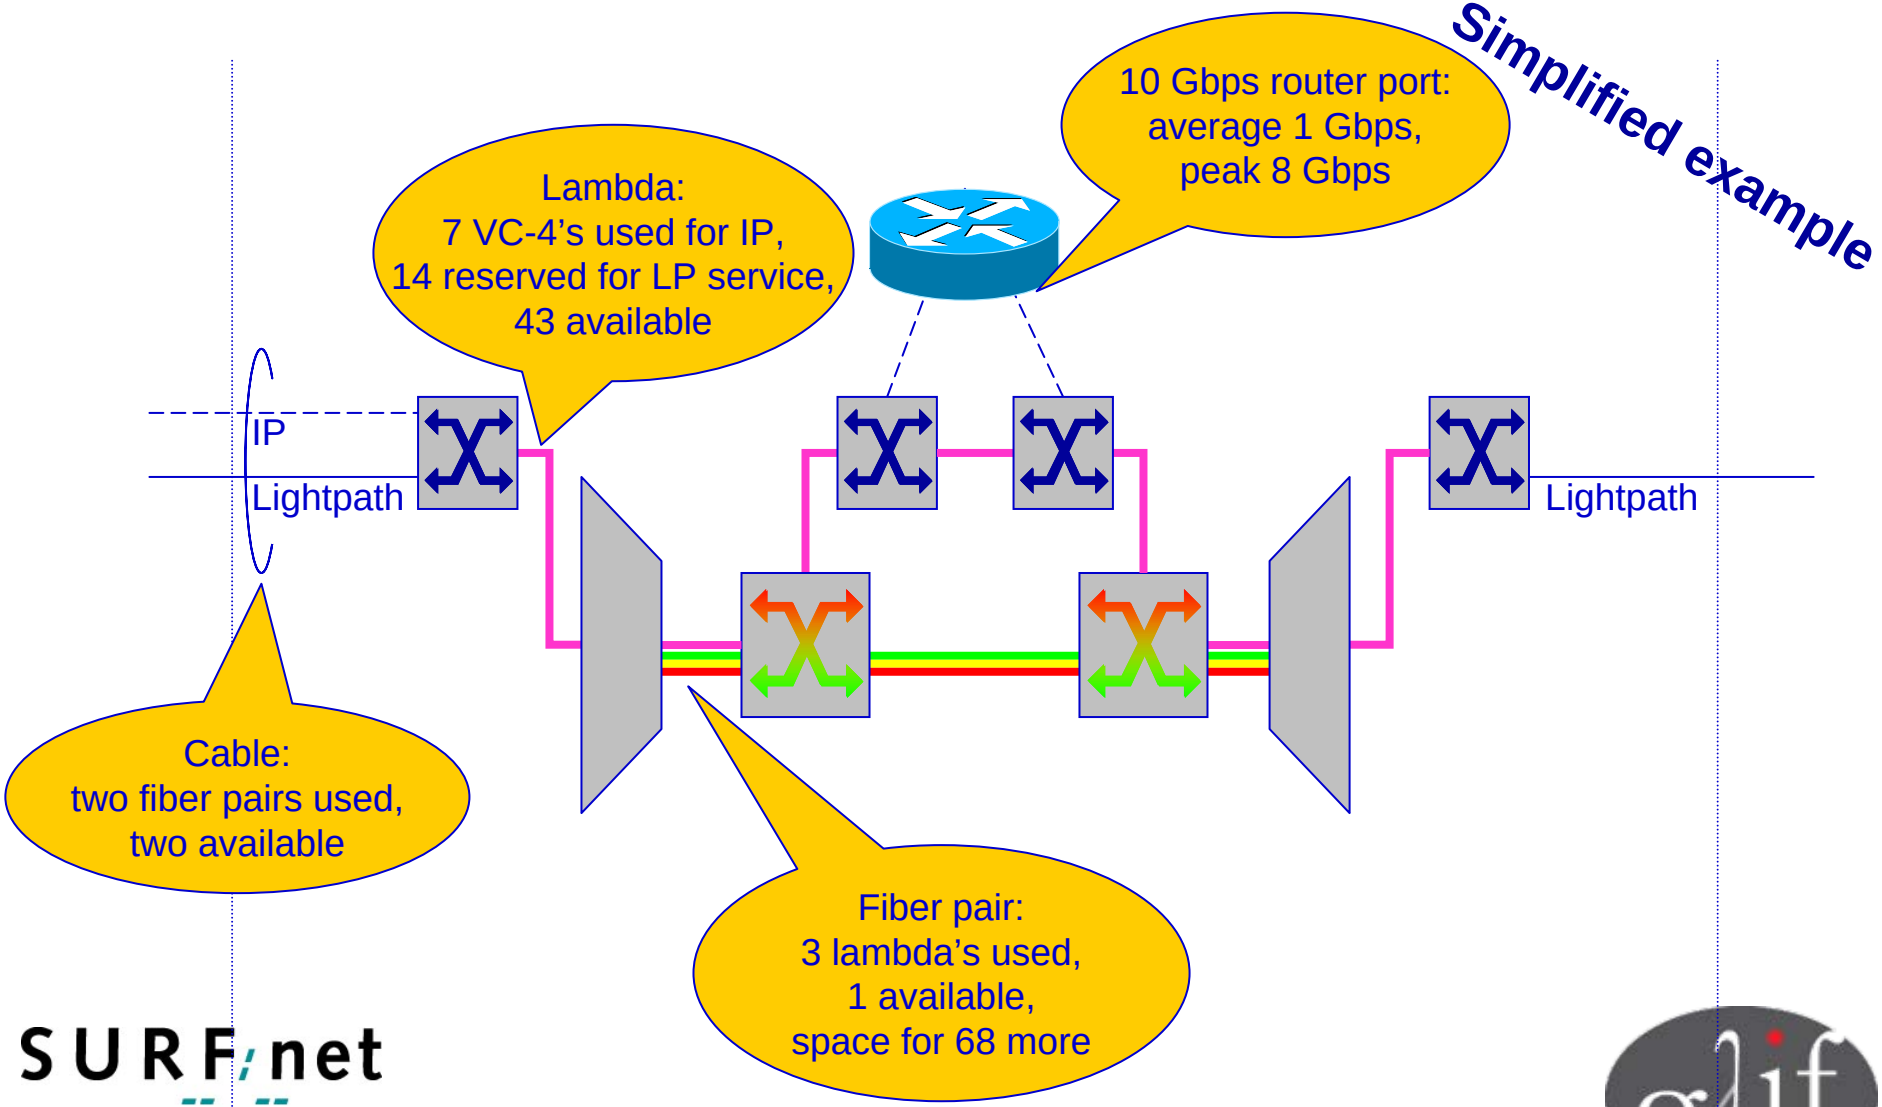

SURFnet

glif

From local to global: build up an end-to-end traffic matrix at each layer

IP Traffic (Mbps/sec) (peak day, peak hour)

Organization / Site	University A Site: A1	A2	University B Site:	Upstream	...
University A					
Site A1	-	0,			
Site A2	0,2	-			
University B					
Site B1	0,1	-			
Upstream	0,3	-			
...					

Lightpaths in use (VC-4's): static, dynamic (peak day, peak hour)

Organization / Site	University A Site: A1	A2	University B Site: B1	IP Routing	GEANT
University A					
Site A1	-	-	2 / 14	14	14
Site A2	-	-	-	2	
University B					
Site B1	2 / 7	-	-	7	7
IP Routing	14 / 7	2	7	-	-
GEANT	14		7		

SURFnet

gaif

Network operator: capacity planning for lightpath services

- IP is packet switched, but lightpaths are circuit switched - just like telephony
- The telephony world has a good capacity model: the Erlang model
- So can we use the Erlang model to predict capacity requirements in the lightpath environment?

SURFnet

glif

The Erlang model will not work for lightpaths: at least not in the near future

- The Erlang model assumes:
  - Stochastic distribution of arrival times
  - Stochastic distribution of duration
- Why does this work for telephony:
  - Very large number of potential callers, each contributing very little, with independent reasons for calling
    - All calling at random times ➔ Poisson distribution
    - Negative exponential distribution of duration
- Why does this not work for lightpaths
  - Small number of potential users at this time
  - Limited number of applications at this time

SURFnet

glif

Accurate forecasting of lightpath use will not be possible

- New service → no data to extrapolate from
- Small number of applications → large impact from use patterns of each application
- Small number of users → large impact from a single user decision
- Example: imagine an organization offering a remote backup service across scheduled lightpaths
  - What time does a scheduled backup start?
    Probably at 20:00 or 0:00
  - How long is the backup scheduled for?
    One hour, or two.
  - Fairly predictable pattern, but a change in backup strategies can disrupt the pattern

SURFnet

glif

Therefore: augment data from the network with knowledge

- Measure the use of resources in the network

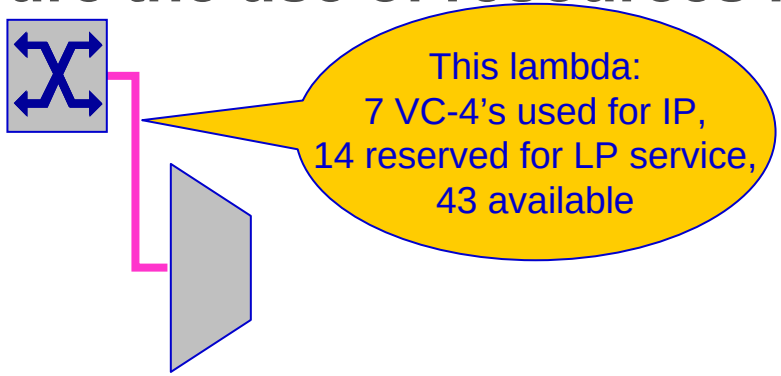

- Show historic data
  see first order tren

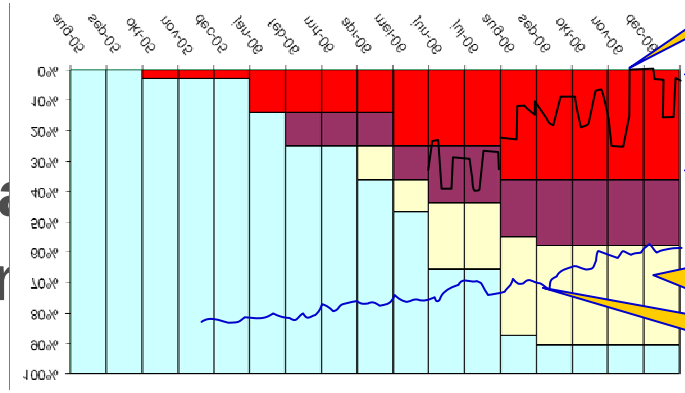

Actual use of dynamic lightpaths

Reserved for dynamic lightpaths

Static lightpaths

Used for IP

Actual IP traffic

SURFnet

glif

Questions?

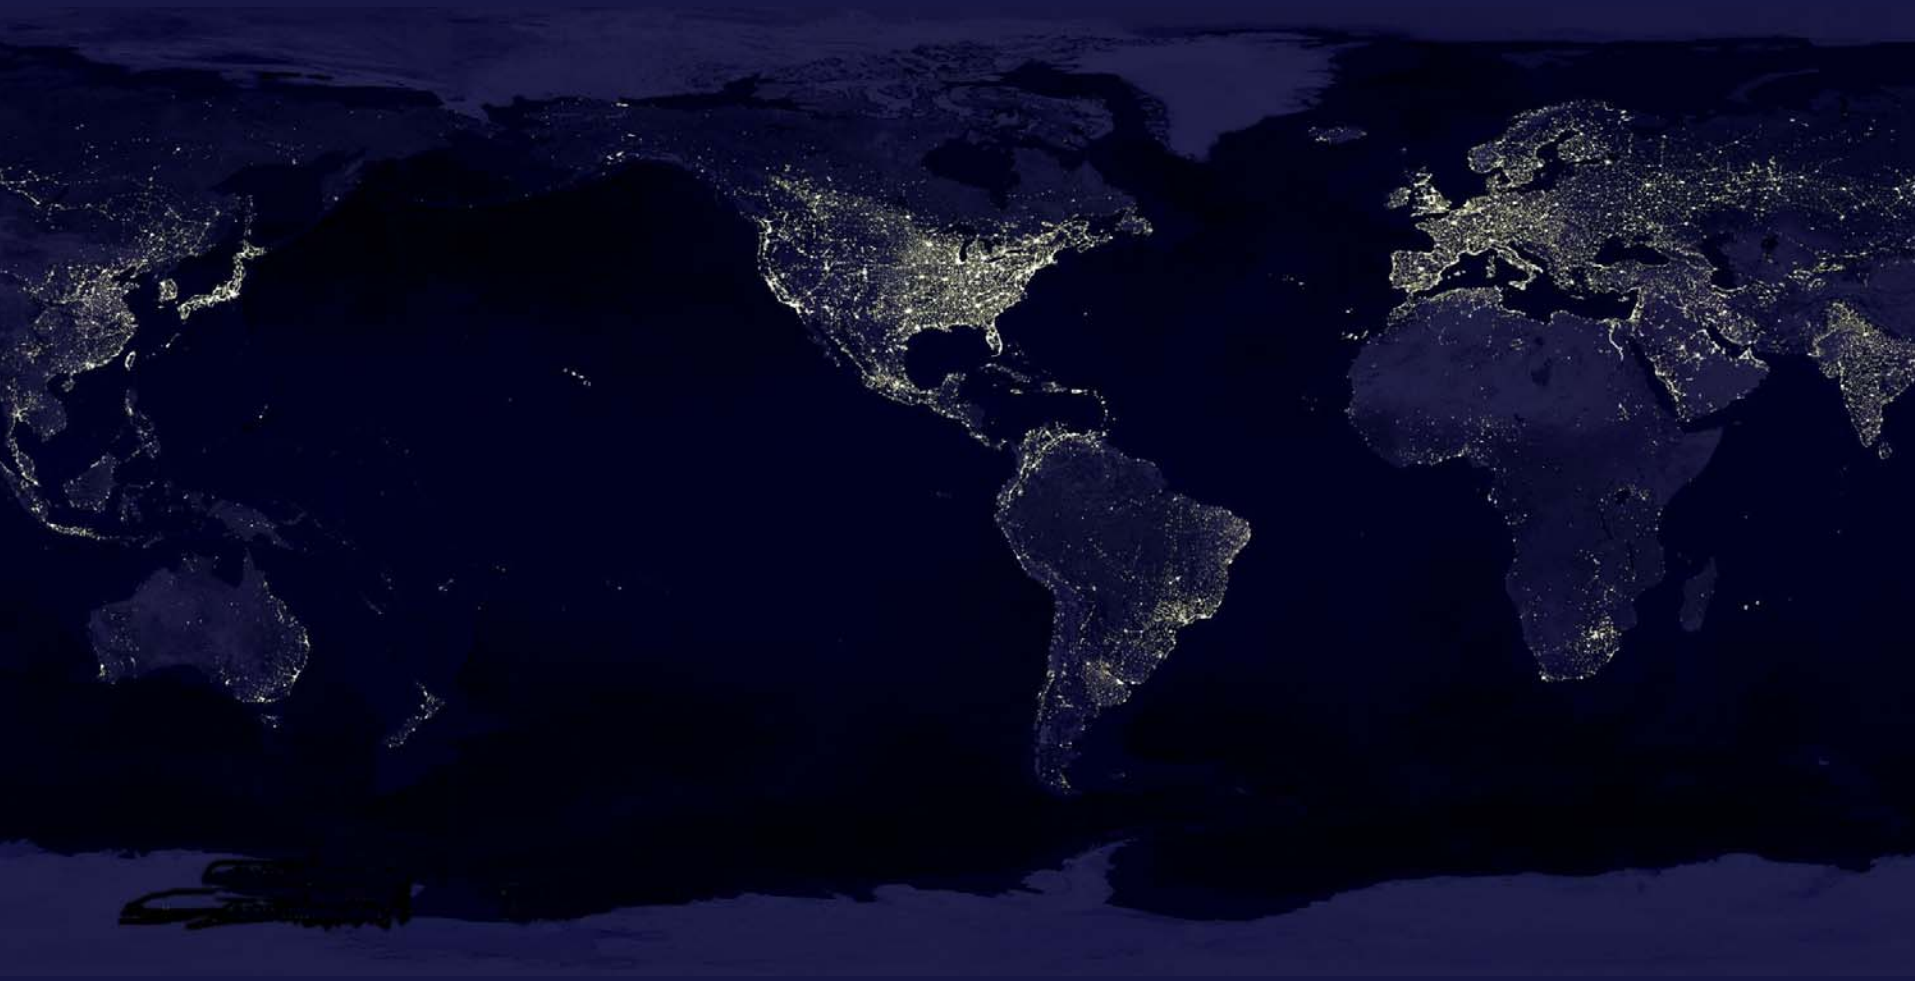

Potential metrics, layer by layer

- Layer 0: fibers and lambda's
  - kilometers of cables and fibers
  - fibers lit
  - lambda's available in WDM systems
  - lambda's in use and total bandwidth, by use type:
    - used for higher network layers
    - used as static lightpaths
    - reserved for dynamic lightpaths
  - use of dynamic lambda lightpaths:
    - bandwidth x hours of use
    - number of set-up requests
    - experienced blocking probability
  - octets transported?

SURFnet

glif

Potential metrics, layer by layer

- Layer 1: optical switching
  - number and size of lightpaths provided, by use type:
    - used for higher network layers
    - used as static lightpaths
    - reserved for dynamic lightpaths
  - use of dynamic sub-lambda lightpaths:
    - bandwidth x hours of use
    - number of set-up requests
    - experienced blocking probability
  - octets transported?

SURFnet

glif

Potential metrics, layer by layer

- Layer 2: Ethernet aggregation
  - number of interfaces
  - lightpath bandwidth
  - aggregation ratio
  - number of frames transported, dropped, errored
  - number of frames, by upper layer protocol
  - octets transported

Potential metrics, layer by layer

- Layer 3: IP
  - number of interfaces
  - transport capacity per interface (lightpath bandwidth)
  - packet loss
  - packets transported: IPv4/IPv6, unicast/multicast
  - octets transported: IPv4/IPv6, unicast/multicast
  - protocols above IP (netflow)