Access to Hybrid Networks

End Site Challenges

Ronald van der Pol

SARA

rvdp@sara.nl

Overview

- Traditional versus hybrid interconnect
- Multihoming & routing
- L2 versus L3
- Addressing
- Dynamic lightpath challenges
- Performance challenges

Traditional Interconnect

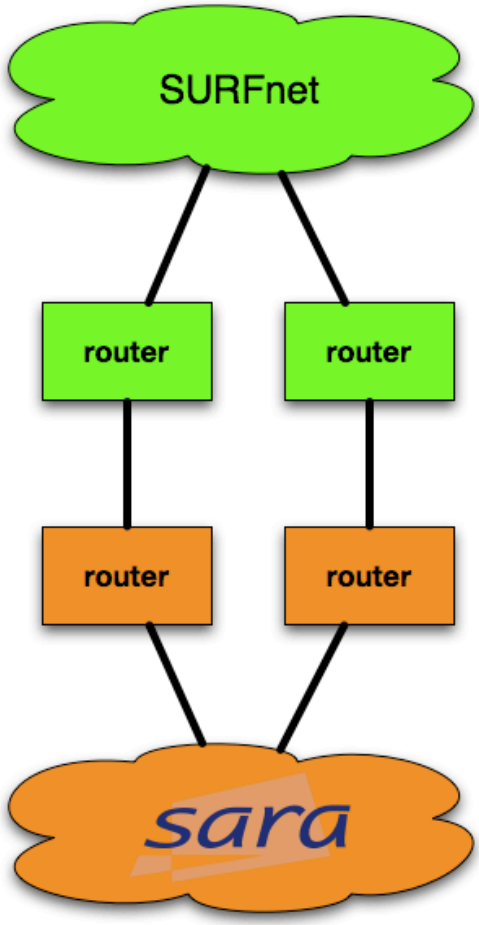

- Router connection
- Well understood
- BGP
- IP ACLs
- Admin separation

Hybrid Interconnect

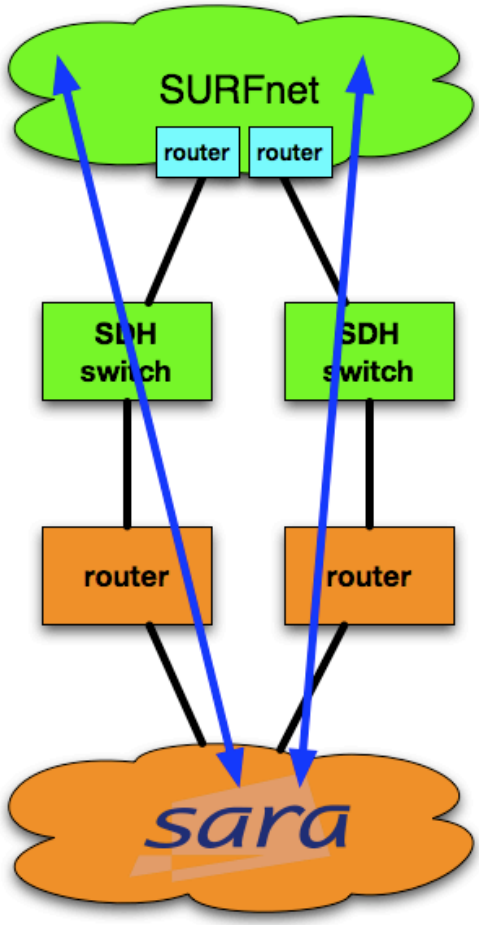

- L1/L2 connection
- Router by-pass?
- No BGP?
- No ACLs?
- Admin boundary?

SARA's Lightpath Connectivity

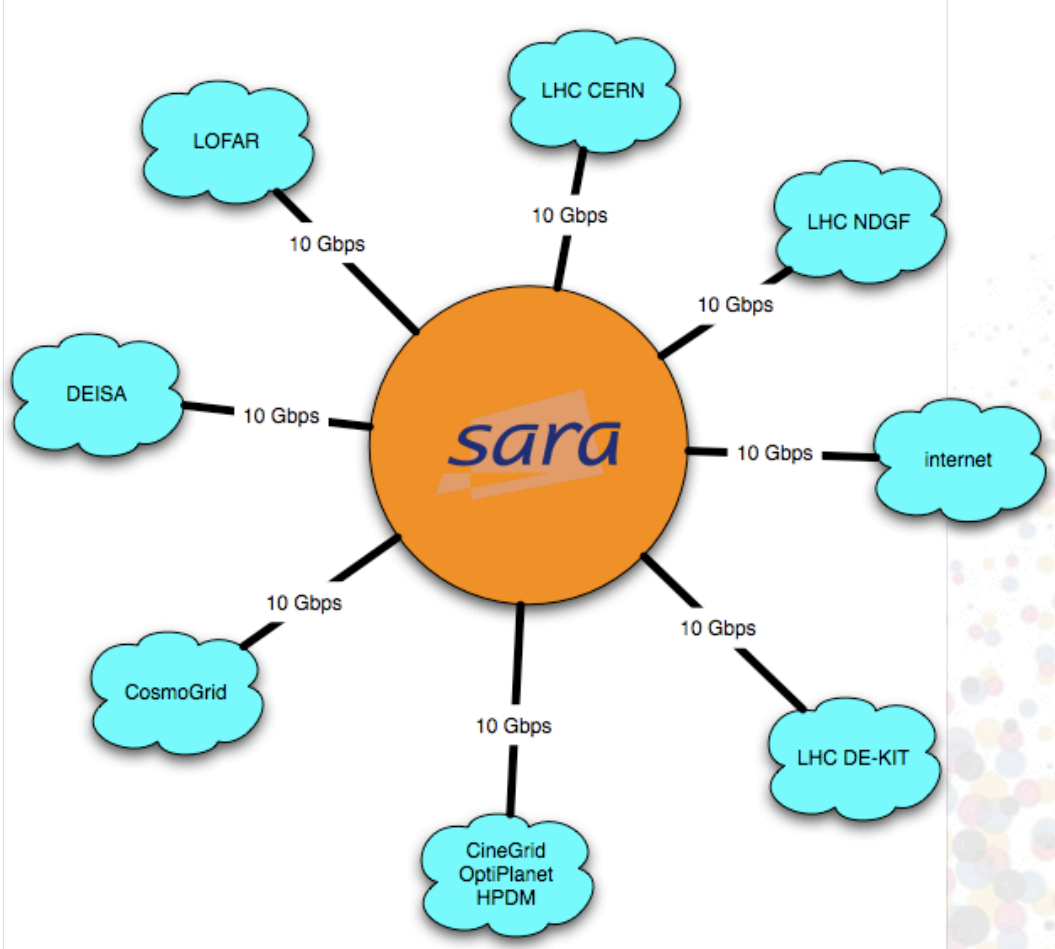

Routing Policies

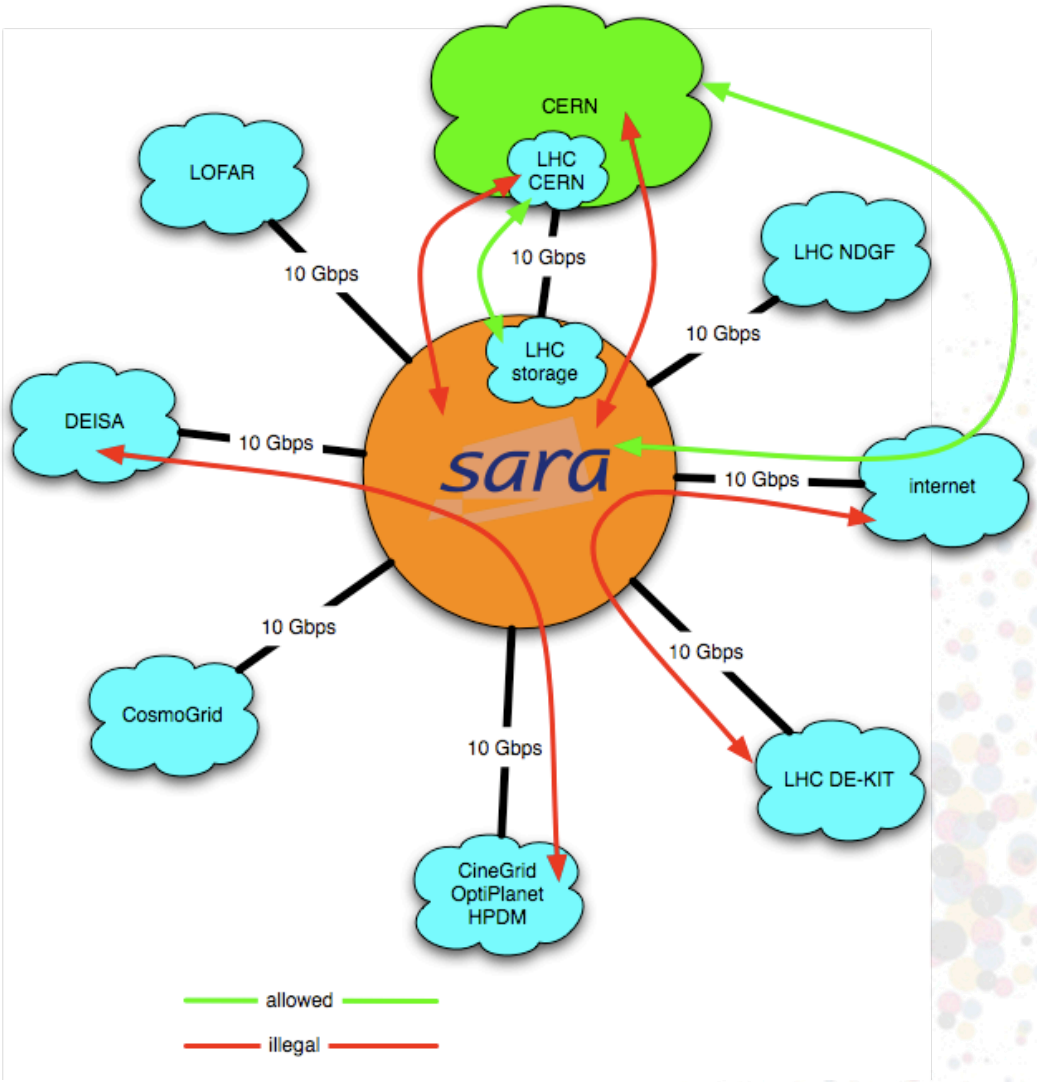

L2 Versus L3

- Lightpaths are L1/L2
  - Share L1/L2 between sites?
  - How to protect against misconfigurations?
    - Ethernet loops
    - Broadcast storms
  - How to do firewalling?
- Connection at L3 well understood
  - Clear separation of administrative domains
  - BGP for routing/traffic policy
  - Extensive experience with firewalls/ACLs
  - More expensive
- Complexity and cost is pushed to the edge
  - Much work to be done
  - Many challenges

Addressing

lightpath between two end nodes

- What addresses to use on lightpath endpoints?
- 2nd address + VLANs on node interface? Second interface?
- Manually configured? Autoconf?
- Use private or public addresses?
- Lightpath endpoints are in the same subnet
- Who provides public address range?
  - Site A?
  - Site B?
  - NREN?

Dynamic Lightpaths

- Server or site has multiple on-demand p2p connections
- What addresses to use?
- Put all sites and servers in the same subnet?
  - How does this scale?
- How to do routing?
  - Links come and go
  - Re-route traffic or drop traffic?
    - Complex routing policy
- How to do monitoring?
  - Is a link provisioned, but not working?
  - When to act on connectivity problems?
  - Which alarms are real and which are “normal”
- How to debug connectivity problems?
  - Local site? Remote site? NREN? Middleware?

Performance Challenge

- Performance gap between bandwidth at end sites and bandwidth that users need and core networks can provide
- NRENs are rolling out 40G and exploring 100G
- Server limitations:
  - 1-10 Gbps network I/O
    - Next step 40G Ethernet or 100G Ethernet?
    - When available at affordable cost?
  - ~ 100-200 MB/s I/O per disk (< 2 Gbps)
- User requirements are a lot higher
- Streaming and transferring data:
  - 27 Gbps is needed to drive 35 Mpixel TPD @ 30 fps
  - 10 TB file @ 10 Gbps takes more than 2 hours
  - 10 TB file @ 40 Gbps takes more than 33 minutes
  - 10 TB file @ 100 Gbps takes more than 13 minutes
  - 10 TB file @ 1 Tbps takes 80 seconds

SARA's 3x3 TPD (35 Mpixel)

- SC09
  - 10 Gbps
  - 10 fps
- SC10
  - 40 Gbps
  - 8K video
  - 30 fps

Thank You

Ronald van der Pol

SARA

rvdp@sara.nl