

SARA's LHCOPN Routing Challenges



pieter@sara.nl – www.sara.nl

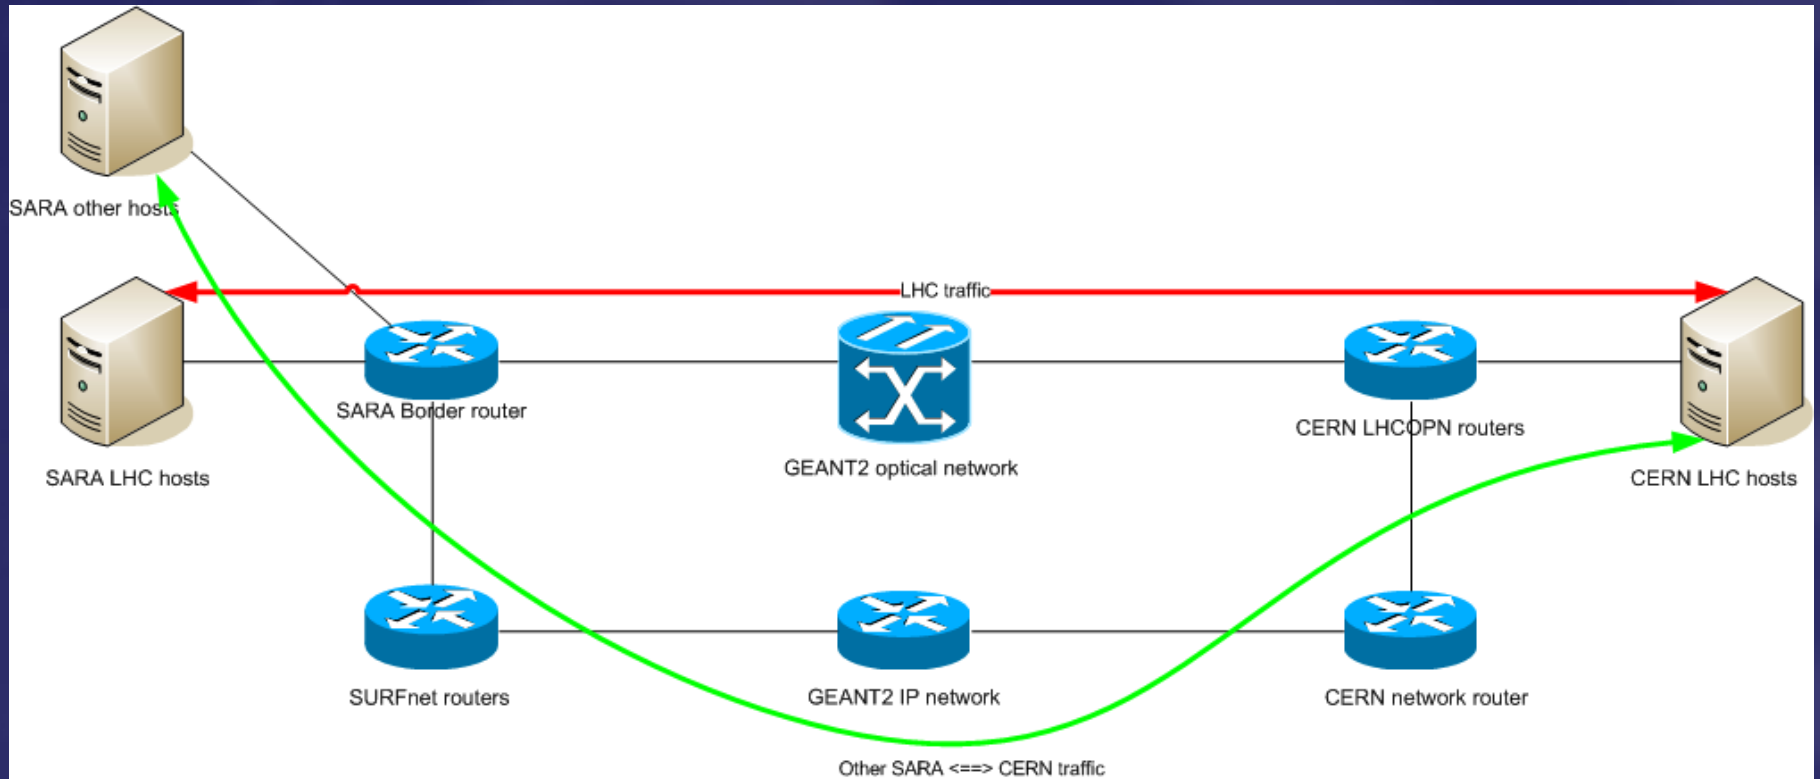
- <http://www.this-page-intentionally-left-blank.org/>

Our internal network

- Highly concentrated with only a few routers
- Router connected to the LHCOPN is not dedicated for the OPN, serves a lot of other stuff
- No policy based routing, QoS or other evil in our network today.

Implementing the OPN (enabling BGP)

- BGP session to CERN up, advertise our prefix, filters in place
- Should be fine
- But instead our entire Tier-1 goes down, very efficiently



So what's the problem

- It appears our grid services servers reside somewhere else in the network and are not allowed over the OPN (part of a compute cluster with other users) We'll fix this
- Other users in our network start to complain they can't reach the CERN LHC hosts. Apparently are they used for other purposes as well, so we need some way to do diverse routing...

- Solve it in the application, multiple NIC's in PC's static routes there, no source based routing needed, ACL's will fix all, but grid certificates don't like this approach
- PBR, but then we can't install any LHCOPN related prefix into our routing tables and have to statically route traffic somewhere with a PBR policy. Not to redundant and far worse then what BGP can achieve
- VRF, but crap our 6509/sup720A doesn't support that, need PFC3B(XL) and route leaking between VRF's and global table can only be statically (in between VRF's can be dynamically)

The Solution

- So VRF was the most likely
 - Hardware, through in another switch
 - Software, do VRF-lite, but needed to buy additional hardware
- Went for the software solution, bought a new PFC

Current situation

- We have a BGP session with CERN, but don't advertise nor accept anything
- Waiting for new PFC to arrive and get installed
- Instead we're pushing everything via our general internet connection via SURFnet over GEANT
- CBF with FZK is up, but FZK only want's to see traffic from our LHC hosts as this link is connected behind there firewall. No BGP yet

