

NL-T1's routing update

still having fun/challenges with those fancy lightpaths/OPN's

Pieter de Boer (pieter@sara.nl)

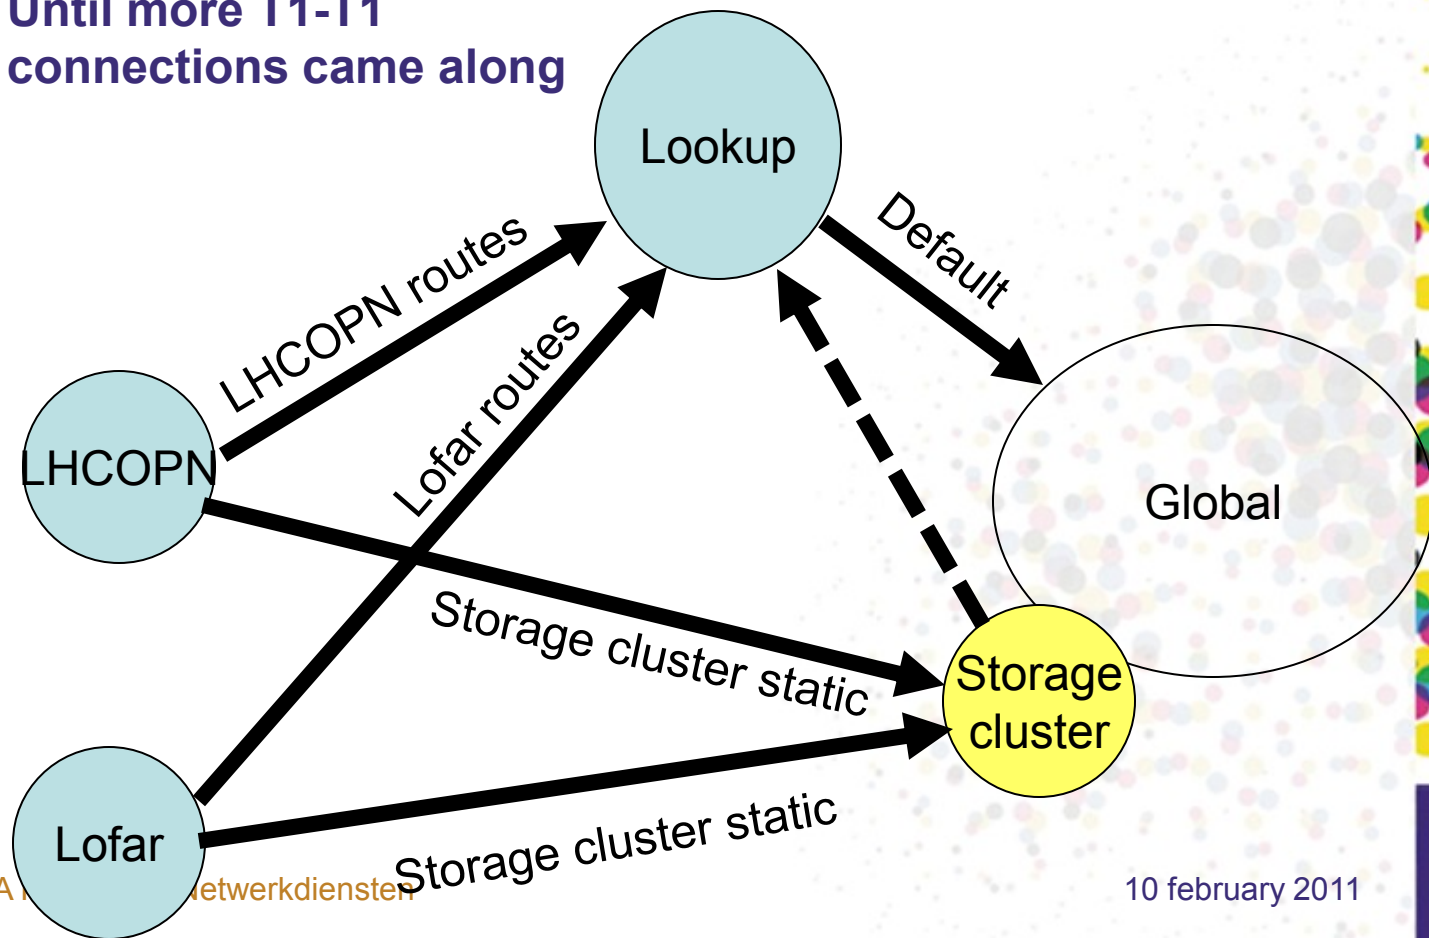
Senior Netwerkspecialist

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

- Hanno left SARA as of 31/12/2010
- I'll be Hanno's successor in relation to LHCOPN matters

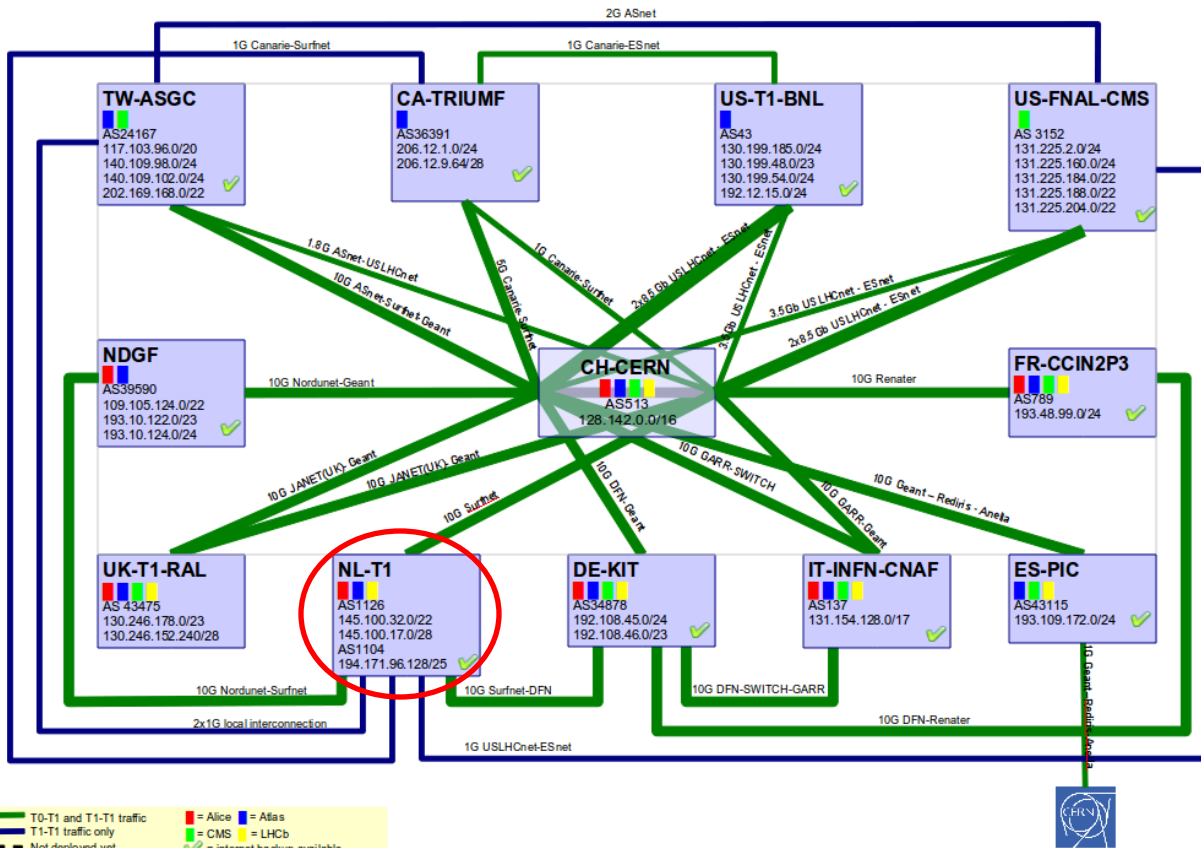
Last update @LHCOPN meeting april 2009 Utrecht

- ▶ Just migrated to VRF's on Juniper MX (feb 2009)
- ▶ Everything was (almost) perfect
- ▶ Until more T1-T1 connections came along



Current connections

LHCOPN



— T0-T1 and T1-T1 traffic
— T1-T1 traffic only
— Not deployed yet
(thick) >=10Gbps
(thin) <10Gbps
■ = Alice ■ = Afaa
■ = CMS ■ = LHCb
✓ = internet backup available
 p2p prefix: 192.16.166.0/24
 edoardo.martelli@cern.ch 20110208

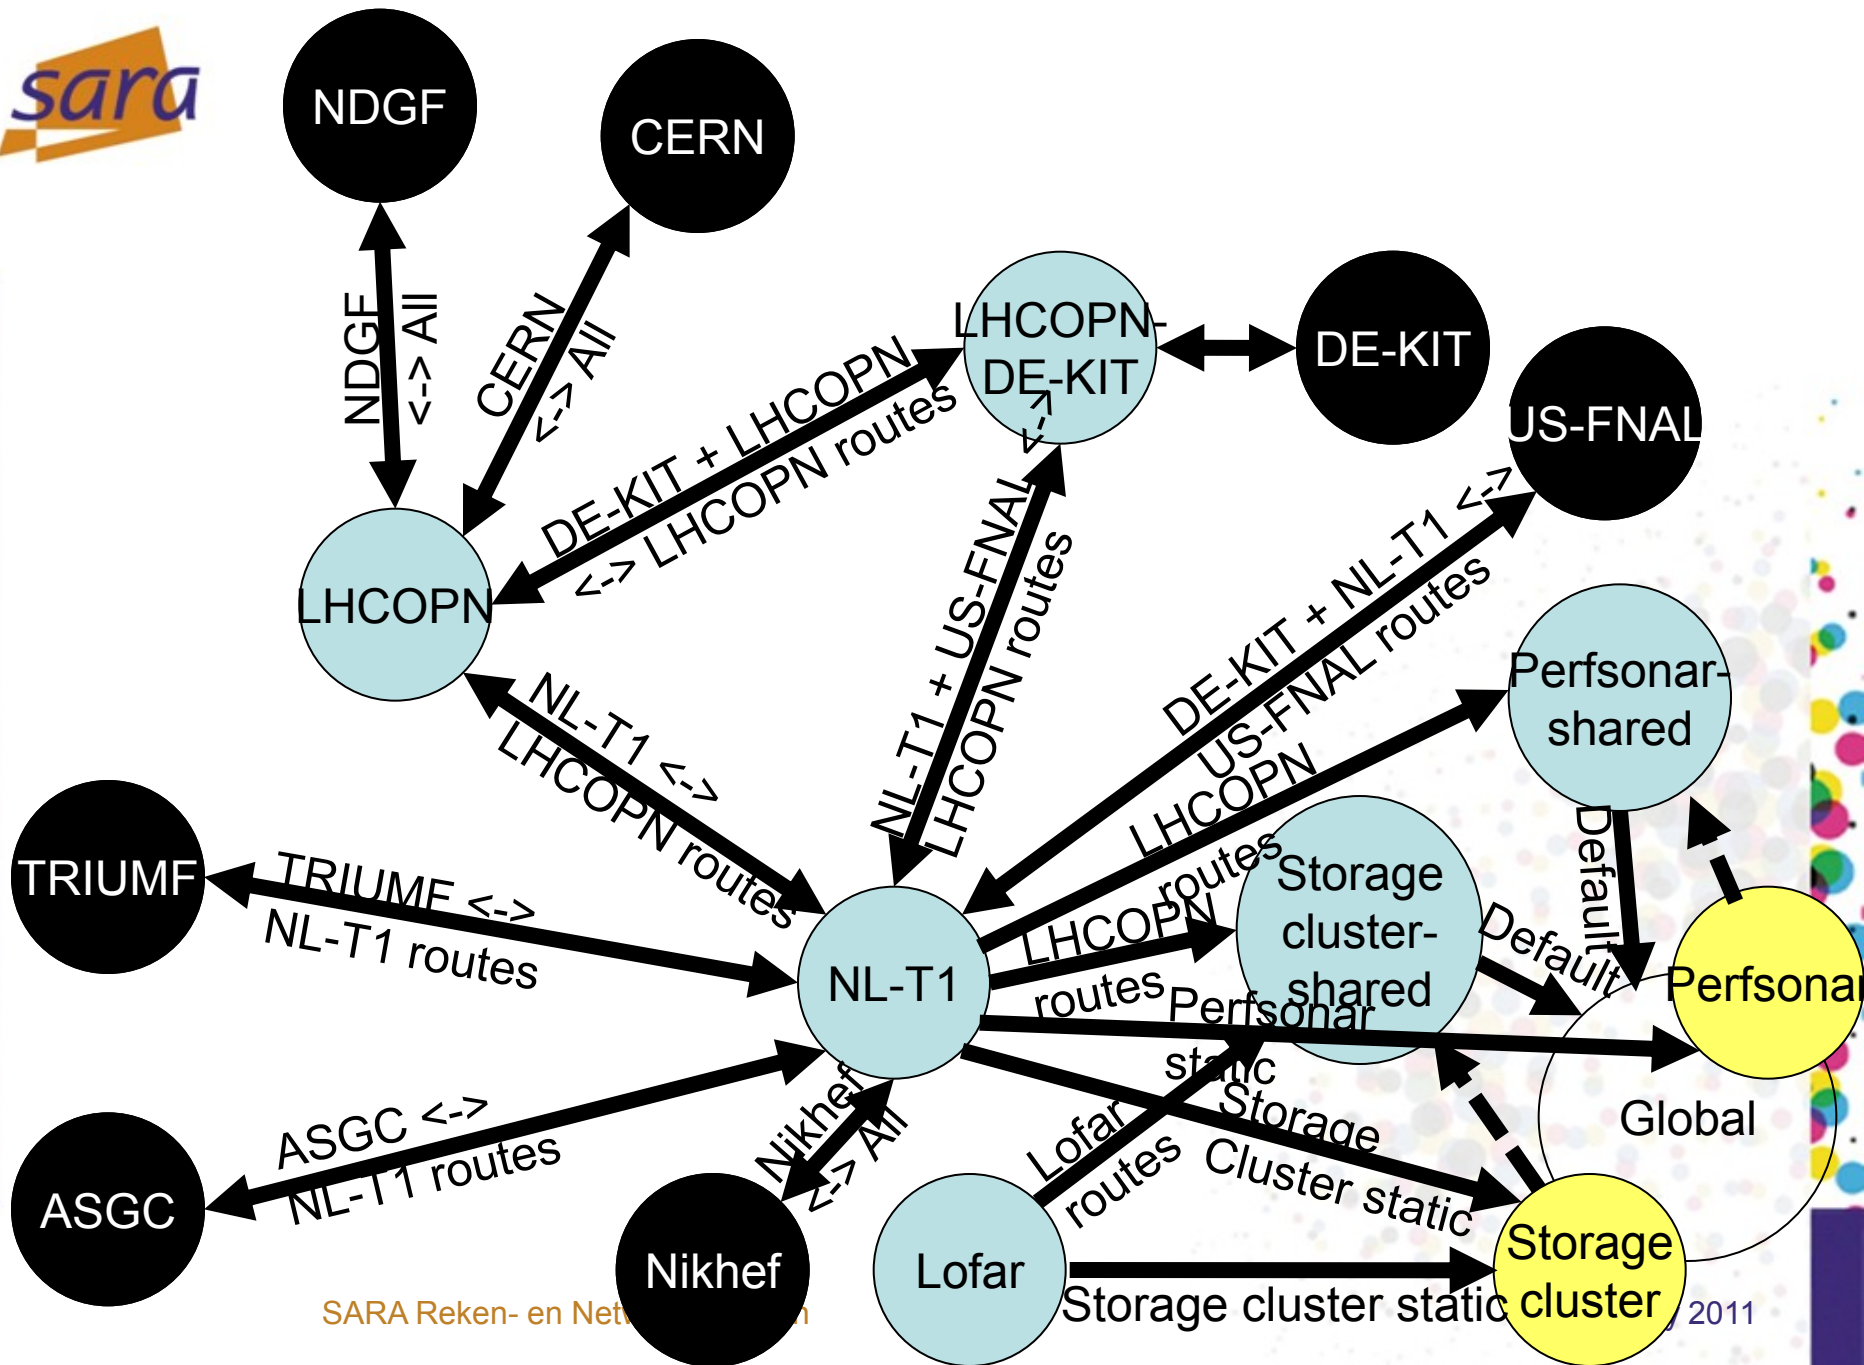
- Normal
 - CERN
 - NDGF
- Limited
 - TW-ASGC
 - CA-TRIUMF
 - US-FNAL-CMS
- Special
 - DE-KIT
 - NL-T1
 - Nikhef

The issue

- **Combination of limited view T1-T1 peerings limits ability to provide decent backup to other T1 sites (NDGF/DE-KIT)**
- **Outdated routing policy (actually still the one we used on our 6509)**
- **And in some cases T1 sites seem to have a policy to only allow traffic from their own site over links to CERN**
- **Let to complaints from NDGF about our limited route advertisement.**

New generation of our VRF design

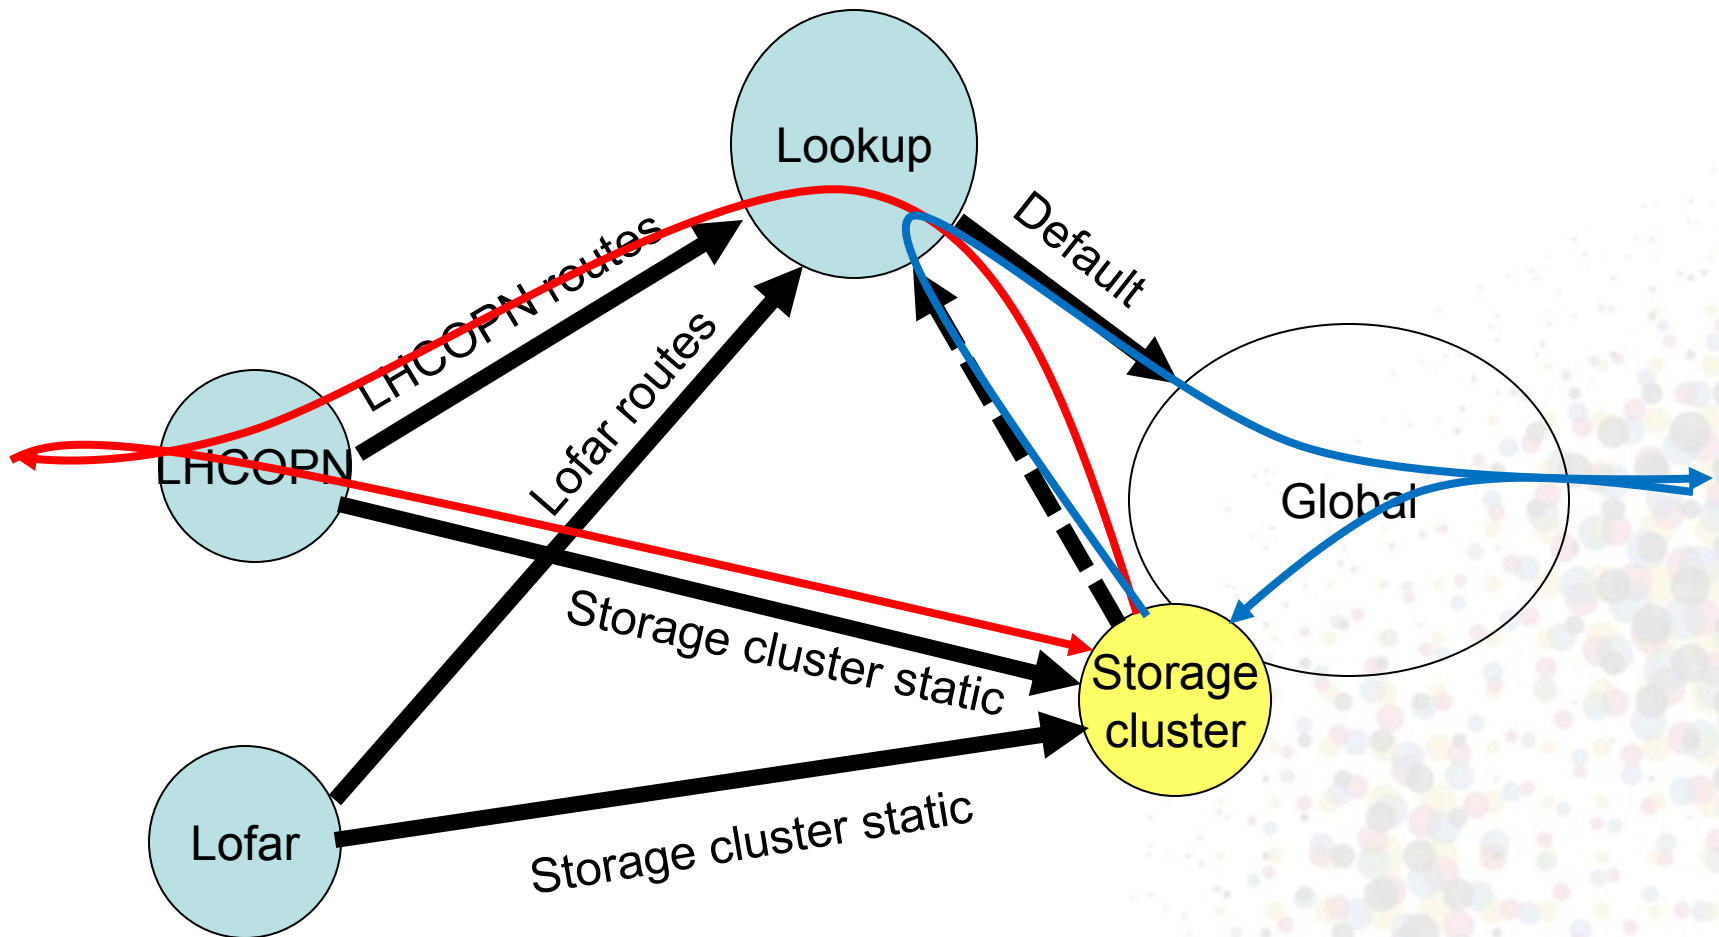
- ▶ **Seperate VRF's for different group**
 - ▶ Normal T1 peers
 - ▶ Limited T1 peers
 - ▶ Special T1 peers / NL-T1 peers
- ▶ **More route leaking**



Results

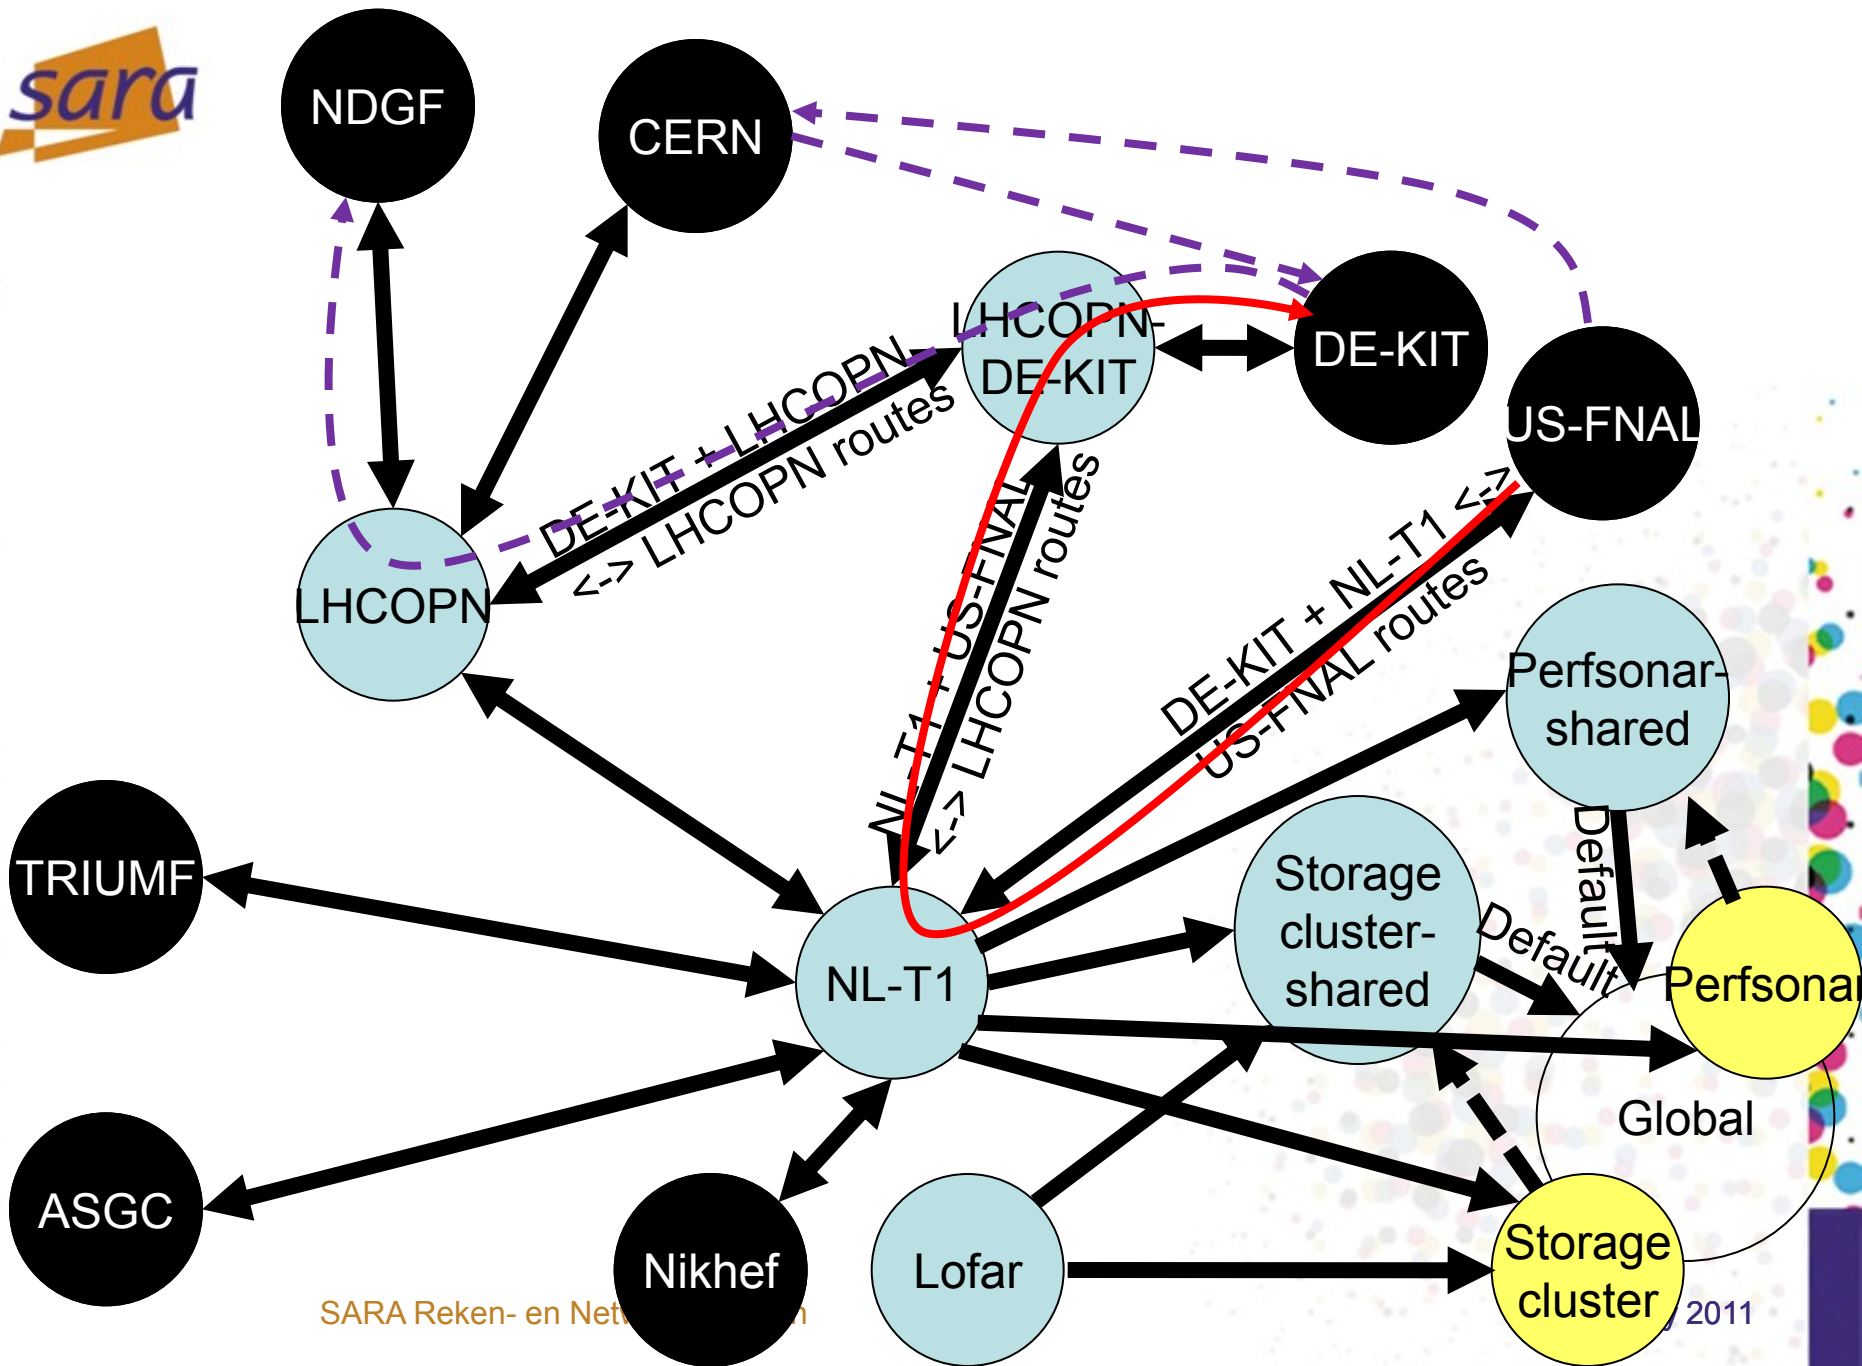
- **Implemented january 18th**
- **No complaints sofar**
- **Performs really well and allows for further growth**
- **Still able to guarantee no unintended route leaking due to seperation of in and out for hosts that have access to multiple VRF's. So routes taken into a lookup VRF have no way back**

Lookup VRF



Results continued

- ▶ **Still one corner case we can't fix**
 - ▶ **If both NL-T1 and NDGF lose there link towards CERN.**
 - ▶ **NL-T1 gets full backup of DE-KIT and has other links to TW-ASGC, CA-TRIUMF, US-FNAL-CMS**
 - ▶ **But NDGF will not get a route to US-FNAL-CMS**



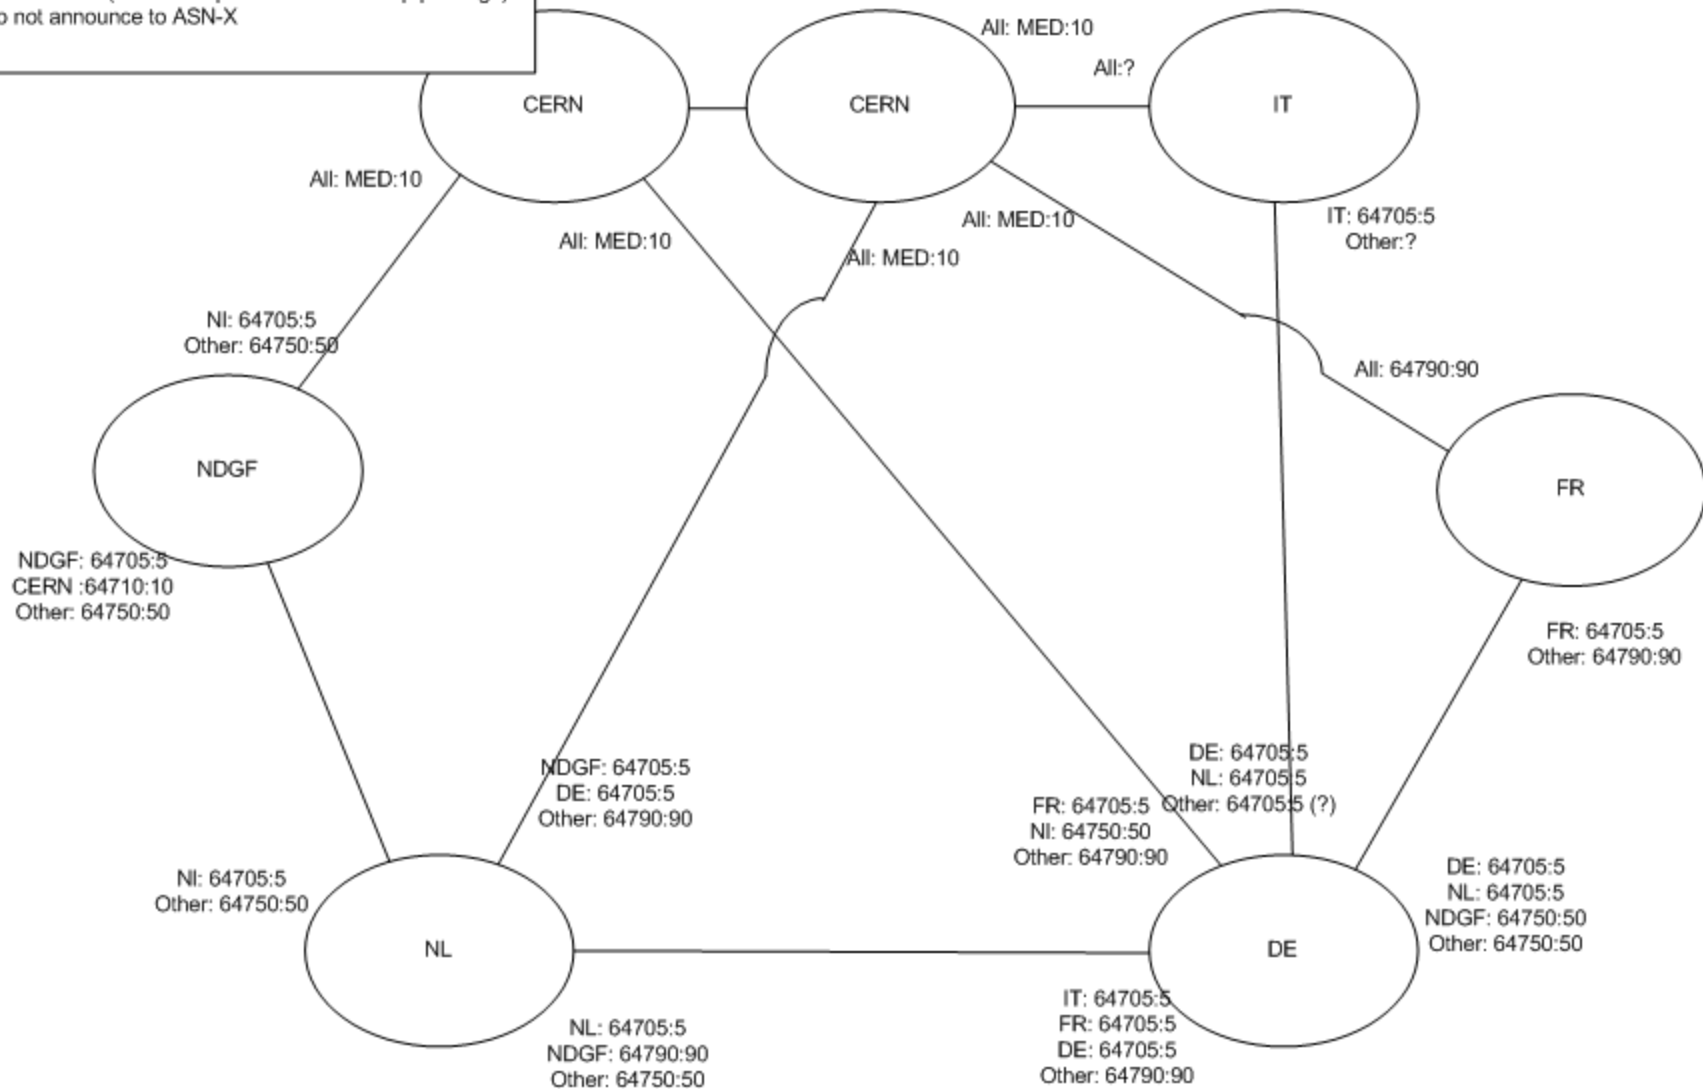
Traffic flapping

- **GGUS ticket 62381**
- **My co-worker Bas Kreukniet works on it**
- **Idea is to resolve this by smartly applying metric's in coordination with DE-KIT, FR-IN2P3 and other, use always compare MED's to make them a tiebreaker**
- **Hopefully implemented the next few weeks.**



Communities are advertised out, metric will be added at the remote site before importing the routes to the routingtable.

localASN:64705 Set MED=5 (best)
 localASN:64710 Set MED=10 (default for primary peerings from CERN).
 localASN:64750 Set MED=50
 localASN:64790 Set MED=90 (default for prefixes from backup peerings)
 localASN:ASN-X Do not announce to ASN-X



Design implications:

White paper / GEANT 3 JRA2 Task 2 (last mile studies)

- We are working on a white paper explaining our setup
- See <http://nrq.sara.nl>
- We've contributed this work to GEANT3 JRA2 TASK2 (last mile studies)
- We are happy to share information with you



Questions/coments



Or send me an e-mail pieter@sara.nl
or find me in a break