

NT1 central report

NeIC NT1 Manager
Mattias Wadenstein
<maswan@ndgf.org>

2024-09-26
NDGF All Hands
Ljubljana, Slovenia

Overview

- Organisation
- Staff
- Services
- Hardware



Organisation

- NeIC, the e-Infrastructure development project organisation is moving out from NordForsk
 - New project office at CSC
 - Varying commitments from Nordic countries
 - Applying for funding to run more useful e-Infra service devel
- NT1, the Nordic Tier-1 operations activity stays at NordForsk during 2025
 - While investigating proper long-term funding and hosting solution
 - Less NeIC branding and oversight, but I don't think we need to cut ties
 - 2025 contracts should be circulating already
 - 2026 TBD



Staff

- Krishnaveni has moved on to a new job
- Darren will take on the dCache developer role as of 2025-01-01
 - But he can spend some time on it already, just not the full 0.5 FTE
 - He'll have both 0.5 FTE central ops and 0.5 FTE dCache devel
- We asked for a budget increase to do necessary development in ARC and dCache for HL-LHC but have been told that this money should “come from somewhere else”

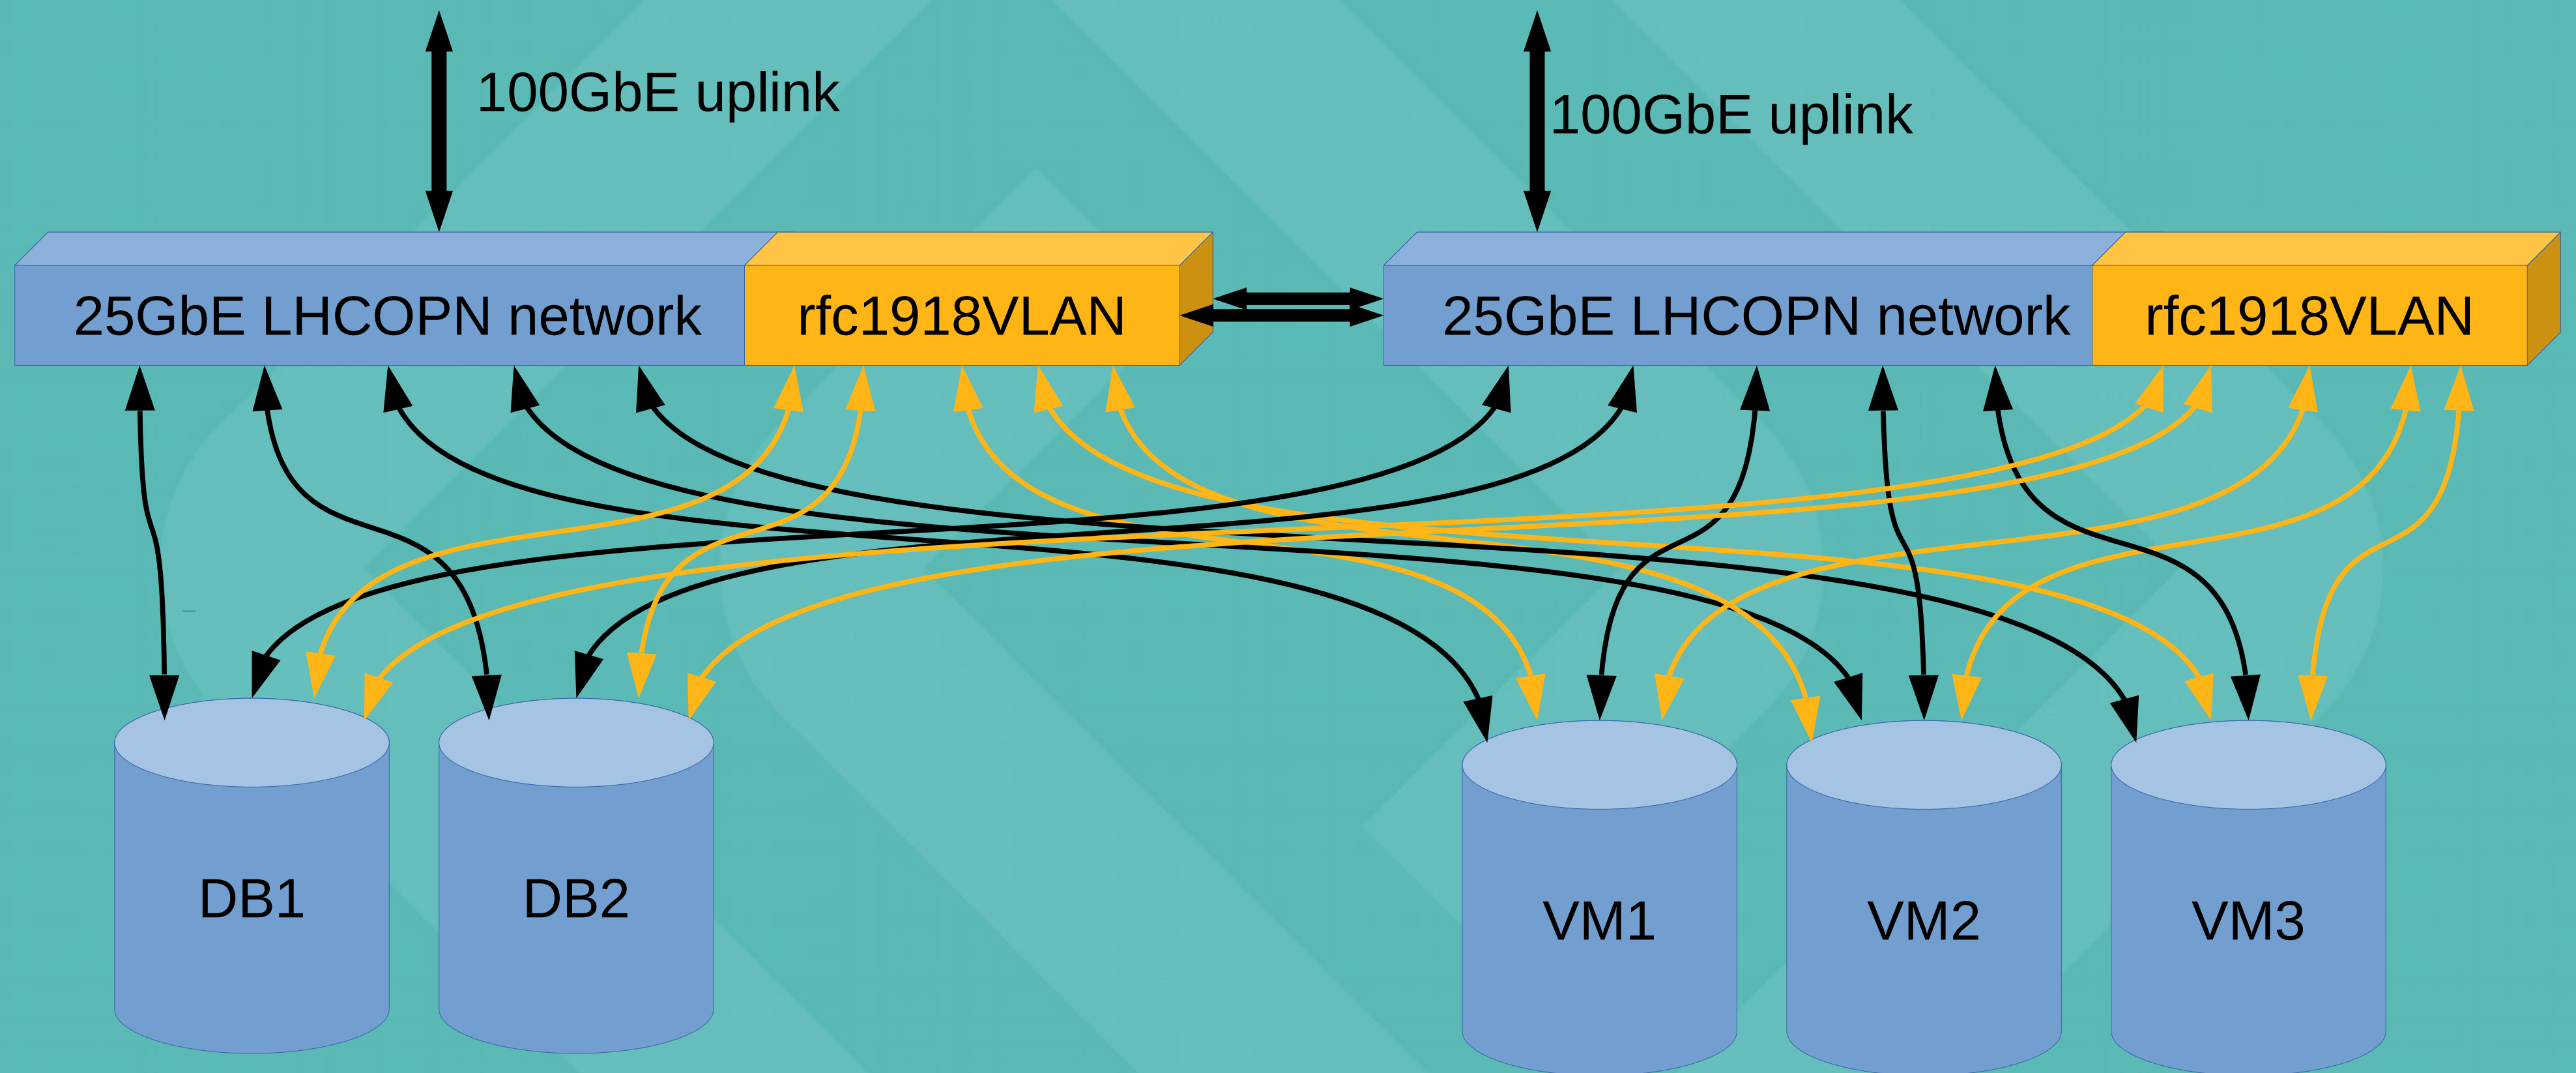


Services

- Trying to get rid of some old cruft
 - BDII
 - Ganglia
 - SRM
 - IGTF (but it'll be a few years still, due to slow token migration)
 - IPv4 (but it'll be a few years still, due to slow adoption for compute)
- New stuff
 - REST API for bulk operations as replacement for SRM
 - Bearer tokens for data and compute access
 - Grafana + prometheus/victoriametrics



Old headnodes



5 x Dell r640 with 2x5222 (2x4 core 3.8GHz) and 192GB RAM, 2TB SSD
4x25GbE for redundant (LACP) internal and external network
Network 2xS5212F-ON (12x25+3x100) with multi-chassi link aggregation



New headnodes

- Old stuff has worked fine in production
 - But it'll be really old next year and service will cost €++
 - Any resource we are short on?
- Any suggestions for what to improve?
 - Assuming roughly same budget
 - Might be some flexibility if we really need it, but I'll get challenged hard on any cost increase compared to budget, as well as not wasting money
 - At least N+1 redundancy so we can reboot them or manage workarounds (remotely) if a single piece of hardware dies
 - Scalability might be nice, but any major increase in needs of central services should come with a new project and money





Questions?