

LHCOPN-LHCONE meeting at RAL, Abingdon (UK) summary report

6th-7th March 2018 – v1.0
edoardo.martelli@cern.ch



Venue

**Hosted by RAL at STFC Cosener's House
Abingdon - UK**



Participants

- 31 participants
- 23 Institutes
- 3 Collaborations
- 6 Research Networks



LHCOPN update

- LHCOPN traffic steady over the year, peaked in October
- CH-CERN:
 - selected Juniper as data centre network equipment supplier
 - second network hub in production but not completed
 - plan for second datacentre on hold
- DE-KIT and NL-T1 will soon have 100G links to CERN
- FR-CCIN2P3 will get two additional 10G to CERN
- Tier1s routing policies: most T1s still using LHCOPN for T1-T1 traffic

<https://indico.cern.ch/event/681168/contributions/2844402/attachments/1611684/2559457/LHCOPNE-20180306-Abingdon-LHCOPN-update.pdf>

IT-INFN-CNAF update


Still recovering after the flood of November 2017

One power line and three chillers fully operational

All LHC data storage back in production

Additional computing resources at remote datacentre of Bari RECAS. Extension to CINECA Bologna almost ready

Planning to move datacentre in new hi-tech area being built in Bologna also for ECMWF



https://indico.cern.ch/event/681168/contributions/2916323/attachments/1611691/2559474/INFN-T1-20180306-LHC_OPN.pdf

BelleII allowed to use LHCOPN



BelleII traffic was allowed over LHCOPN after discussion at the previous meeting

On 19th Dec 2017, the WLCG Management Board agreed with the decision, as long as BelleII traffic doesn't impact WLCG data transfers

The LHCOPN AUP has been documented here

<https://twiki.cern.ch/twiki/bin/view/LHCOPN/LHCopnAUP>

<https://indico.cern.ch/event/681168/contributions/2848744/attachments/1611685/2559459/LHCOPNE-20180306-Abingdon-LHCOPN-for-BelleII.pdf>

ProtoDUNE, LHCOPN, QoS

FNAL has asked permission to use LHCOPN or LHCONE to transfer protoDUNE data from CERN

Expected (O)10Gbps from CERN at peak intervals. Not a concern for existing CERN-FNAL connectivity (100Gbps)

Discussed possibility to re-implement QoS to prioritize WLCG Tier0-Tier1 traffic over any other traffic on LHCOPN links.

FNAL and CERN will evaluate need and feasibility



ProtoDUNE outer vessel

Use of LHCOPN/ONE by other collaborations

Discussion followed whether allowing any collaboration to use LHCOPN and LHCONE. No consensus reached, but these points were noted:

- Different private networks for the different collaborations would be easy to create in backbone networks. They would allow a clearer separation of traffic for better statistics, security, billing
- End-sites serving multiple collaborations may have difficulties in separating the traffic to use different private networks. If funding agencies agree, sharing the same infrastructure would be preferable

LHCONE update

L3VPN status

- IHEP T2 in China now connected to LHCONE thanks to CSTNet, CERNet and TEIN
- MIT's LHCONE connection moved from Internet2 to Esnet
- UK-T1-RAL will soon connect

https://indico.cern.ch/event/681168/contributions/2853806/attachments/1611809/2559689/2018-03-06_ECapone_LHCONE.pdf

L3VPN operations

- ESnet tool to visualize BGP communities no longer on-line. The community expressed the interest to have it back and possibly expand it's visibility to all Network Providers

<https://indico.cern.ch/event/681168/contributions/2848473/attachments/1611718/2559653/LHCONE-OperationsUpdate-Abingdon.pdf>

Odd traffic inside LHCONe

DE-KIT and ESnet have run an investigation on unexpected packets coming from LHCONe connections


- mostly packets with private (RFC1918) source addresses
- packets from connected sites but with undeclared source addresses
- packets from not connected sites (one case)

Proposed action plan:

- keep looking for such traffic to eliminate it
- edge filtering (uRPF, antispoofing)
- end sites education

<https://indico.cern.ch/event/681168/contributions/2848474/attachments/1611723/2559528/LHCONe-Filter-Policy-Practice.pdf>

perfSONAR update



LHCONE MaDDash: improved situation

New MaDDash v2.0: changed color coding, several improvements

Upcoming new perfSONAR version 4.1:

- drops SLC6 support
- pScheduler now requires TCP port 443 to be opened
- support for TWAMP (two way measurement)
- support for Docker

Necessary to upgrade LHCOPN/ONE instances to v4.1

OSG/WLCg now storing perfSONAR measurements in Elasticsearch

OSG working on alarming for network issues

LHCOPN measures now visible on CERN [monit-grafana](#)

https://indico.cern.ch/event/681168/contributions/2855524/attachments/1611738/2559815/LHCONE_perfSONAR_update-RAL-2018-v1.2.pdf


BelleII update



- BelleII sites which are also WLCG Tier1s can use LHCOPN: US-T1-BNL, KR-KISTI, IT-INFN-CNAF, DE-KIT
- Throughput tests between KEK and BNL achieved 1.5GB/s sustained
- Estimation of future traffic patterns after 1yr of analyses: >92% of user traffic stays local on the sites; only <8% is sent to other sites
- Implemented MaDDash for BelleII sites

<https://indico.cern.ch/event/681168/contributions/2858913/attachments/1611831/2559808/BELLEII-UPDATE-LHCONE-2018-05-march-v4.0.pdf>

The DUNE collaboration



Fermilab presented the DUNE experiment:

- 2018: prototyping with ProtoDUNE at CERN
- 2019-2025 construction
- 2026 First neutrino beams from Fermilab to Sanford

Expected 30PB/yr of raw data (similar to LHC Run2).

Primary storage at Fermilab, data replication at CERN

Due to large overlap between LHC and DUNE sites, use of LHCONE may be envisaged

https://indico.cern.ch/event/681168/contributions/2867987/attachments/1611851/2559796/DUNE_Intro_and_NOvA_Update.pdf

GEANT Network evolution

- Awarded new contract for low cost transmission devices with coherent DWDM optics (DCI). They will be deployed alongside existing Infinera transmission boxes
- Planning to use cheaper routers in smaller POPs. Candidate router: Juniper MX-240
- Evaluating extension of fibres footprint
- Preparing to tender for next transmission equipment
- Testing software for whitebox switches



https://indico.cern.ch/event/681168/contributions/2913557/attachments/1611820/2559982/2018-03-06_ECapone_Geant_Network_Evolution.pdf

IPv6 deployment in WLCG



- WLCG needs to be ready for any opportunistic IPv6 only CPU resources
- Tier-1s should be providing production storage accessible over IPv6 by April 2018 (8/14 are currently able)
- Tier-2s by the end of 2018 (~20% done)

ATLAS now running production jobs on IPv6-only worker nodes

CERN EOS instances are now all dual-stack

Approximately 12% of FTS transfers go over IPv6 (which appear more efficient/reliable)

111 of 271 perfSONAR hosts now reporting IPv6-enabled

OPN, USCMS, USATLAS, UK and Belle II perfSONAR meshes now dual-stack

Some interest in examining performance of IPv6 compared to IPv4

<https://indico.cern.ch/event/681168/contributions/2876165/attachments/1611424/2558924/IPv6-update-20180305-v3.pdf>

Community White Paper



The WLCG and the HEP Software Foundation Community White Paper has been released

A roadmap for the HEP Software and Computing R&D for the 2020s: 13 topical sections summarizing R&D in a variety of technical areas for HEP Software and Computing

CWP identifies the main areas needing investment for the future HEP Software Upgrade. Must now turn them into concrete actions

https://indico.cern.ch/event/681168/contributions/2865947/attachments/1612178/2560398/WLCG_and_HEP-Computing-After-CWP.pdf


MTU size recommendation

Communication issues are still traced back to MTU mismatch

Jumbo frames must be configured on routers interconnection links in LHCOPN and LHCONE. It is recommended to configure at least 9000B for the IPv4 and IPv6 MTU

ICMP packets necessary to the Path MTU Discovery protocol must be allowed in firewalls

No consensus on the MTU to be recommended on the segments connecting servers and clients. A working group will produce these recommendations



LHCONE R&D update

AutoGOLE

- MEICAN has been chosen for the dynamic provisioning of circuits on a multi domain network
- Actively supported by RNP (Brazilian NREN)
- Collaboration with GNA to grow MEICAN adoption in Global Exchange Points

Data Transfers Nodes (DTNs)

- demo at SC17
- test with DTNs at iCAIR, PacificWave, UvA, CERN. Circuit over AutoGOLE infrastructure
- Test to compare new MDTMFTP and GridFTP run between CERN and Fermilab

https://indico.cern.ch/event/681168/contributions/2903390/attachments/1612487/2561034/LHCONE_-_RAL_-_March_2018_-_Gerben_van_Malenstein.pdf

https://indico.cern.ch/event/681168/contributions/2844417/attachments/1612567/2561206/International_WAN_DTN_Experients_V2.pdf

LHCONE R&D update

Towards a production service

- interest in developing a DTN service; all necessary parts are available
- if not for LHC , it may be useful to other experiments
- Upcoming new WLCG “Data lake” model may be the opportunity to use DTN service as a front-end
- challenges: complexity must be hidden to end users; global identity federation for authentication and authorization

Conclusions

Summary

- LHCOPN: preparing upgrades for last part of Run2 and Run3
- BelleII allowed to use LHCOPN. ProtoDUNE may use it as well
- LHCONE: China has joined. Edge filtering may be added to get rid of unwanted packets
- Monitoring: perfSONAR v4.1 and MaDDash v2.0 coming with new features. A route-server for LHCONE routes may be added
- IPv6 adoption expanding thanks to WLCG/HEPiX follow up activity. More dual-stack storage available
- R&D: improvements in circuit provisioning and DTN services

Actions

- Evaluate implementation of QoS to prioritize T0-T1 traffic on LHCOPN links (EdoardoM, PhilDM)
- If CERN and FNAL will use LHCOPN for ProtoDUNE, inform WLCG MB (EdoardoM)
- Follow-up on chase of illegal packets in LHCONE (BrunoH, MikeOC)
- Recommendation for MTU size (ShawnMK and others)
- Evaluate use of route-server to display available routes and tagged prefixes (MikeOC, MagnusB, EdoardoM)

Next Meeting

Proposed dates: last week of October 2018, most likely Tuesday 29th and Wednesday 30th 2018

Proposed host: Fermilab in Batavia Illinois (US)

Date and exact location will be confirmed

Links

LHCOPN and LHCONE meeting agenda:

<https://indico.cern.ch/event/681168/>

Questions?

edoardo.martelli@cern.ch