



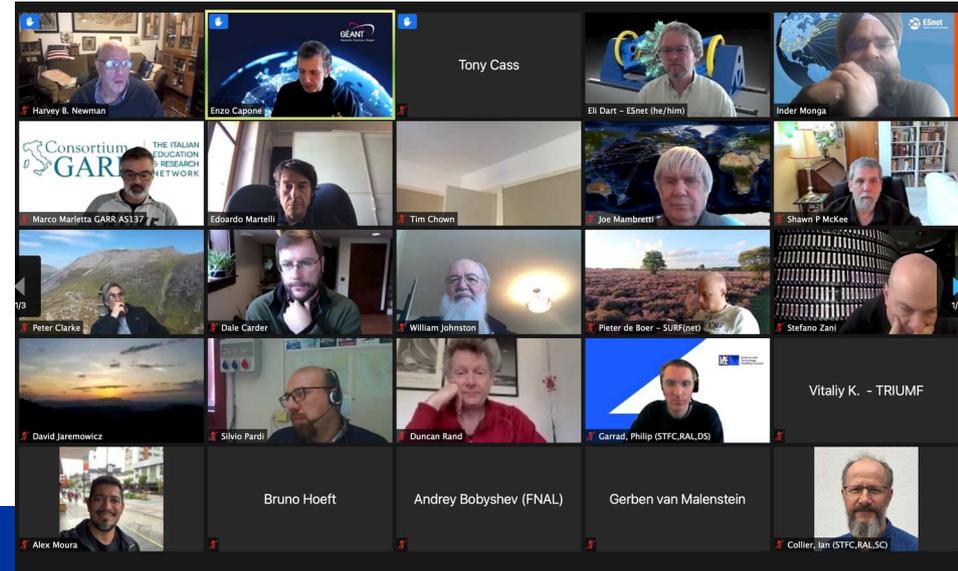
LHCOPN-LHCONE meeting #46 summary notes

GDB 12 May 2021

edoardo.martelli@cern.ch

Venue

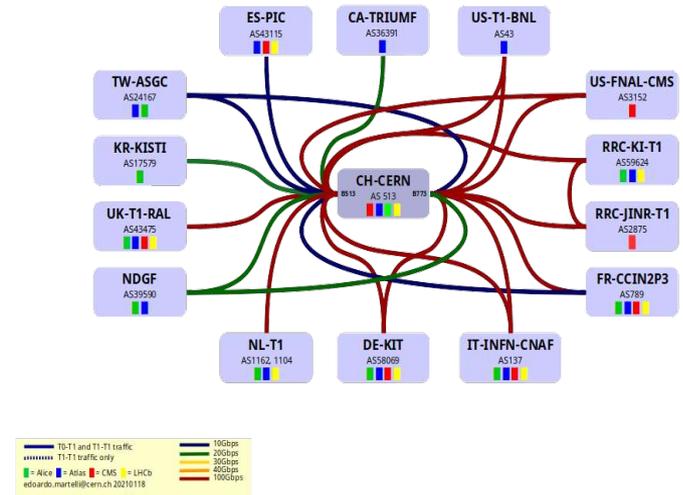
- On 23-24 of March 2021
 - On video conference only, for the third time
 - Two sessions of 3 hours (and a bit more) in two days
 - 87 subscribed
 - 66 participants on the first day, 57 on the second
- Agenda at <https://indico.cern.ch/e/LHCOPNE46>



LHCOPN update

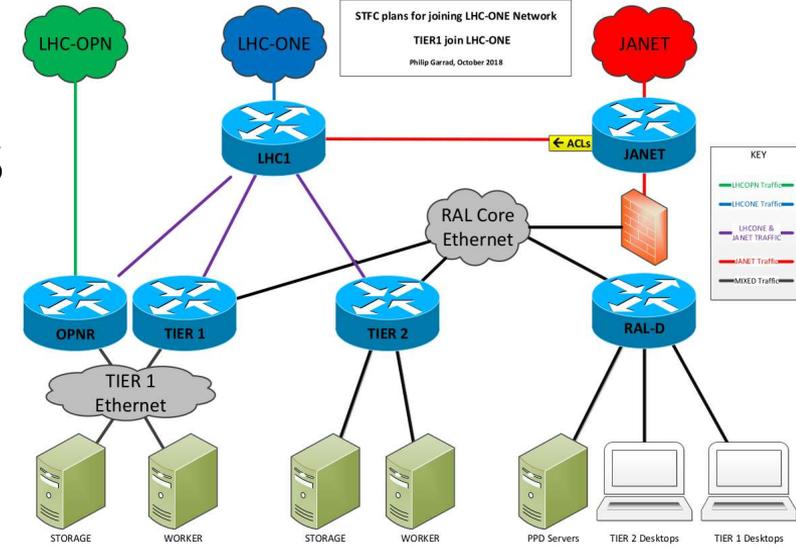
- More 100G upgrades: RAL, DE-KIT (second)
- Upcoming: PIC 100G, IN2P3 second 100G
- FNAL and BNL will soon connect to ESnet6 at 400Gbps
- Total of 1.3Tbps from the Tier0 to the Tier1s
- Traffic stats: moved 269PB in one year. -6% compared to previous year

LHCOPN



UK-RAL update

- Campus network upgrade: new firewall, 200Gbps off-site connectivity
- New data-centre network: leaf/spine design, based on Mellanox switches and Cumulus OS
- RAL Tier1 will connect to LHCONE in May/June 2021
- New network fully deployed by Summer 2021
- Will add second 100G for LHCOPN during Run3

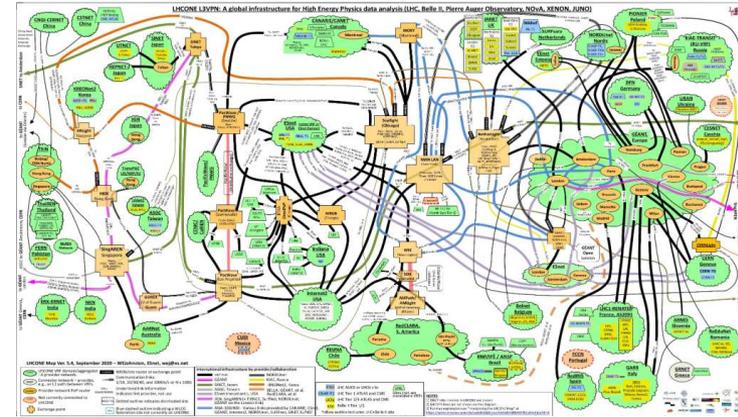


LHCONE update

Topology: no new network providers, no new sites

Traffic volumes:

- general decrease of 20% in all provider networks
- decrease between EU and North America
- increase between Asia and North America, probably due to BelleII data taking
- constant between EU-Asia and EU-South America
- campus networks have seen a lower decrease (6-10%)



Transatlantic Network Capacity

- Network capacity planning analyses made by ESnet.
- ESnet is planning the capacity of its network using historical data and requirements and forecasts from the largest user communities
- LHC Run 3 will continue to take advantage of ESnet capacity, without needing significant additions initially. Upgrades will happen during Run3, as utilization grows
- ESnet European network will also grow, to match the transatlantic capacity
- Transatlantic costs will reduce overtime. No economical worries to meet Run3 needs
- ESnet needs user input to best plan future investments for HL-LHC

European Demand and Capacity Forecasts (updated March 2021)



ESnet
ENERGY SCIENCES NETWORK

<https://indico.cern.ch/event/983436/contributions/4231750/attachments/2213488/3746723/dart-monga-lhcone-transatlantic-2021-03.pdf>

BelleII update

- So far collected 2PB of raw data, stored at KEK and BNL. 1.2PB to be collected in 2021
- Long shutdown will start July 2022 for machine upgrade
- LHCONE: 30% of sites connected, but they generate 80% of the BelleII traffic
- All data flow are managed by Rucio. The adoption of Rucio has brought many improvements
- New data distribution schema: no longer only in BNL, but shared among all the raw data centres.
- Because of this change, a new round of data challenges is planned. Also to simulate the expected increase of capacity (45TB/day). Pilot tests starting now with KIT and CNAF
- BelleII will contribute to the packet marking activity.
- Tier0 is deploying IPv6, but no prefix advertised yet



Database for LHCONE prefixes



- Network information related to LHCOPN, LHCONE and monitoring agreed to be added to CRIC
- A WHOIS route-set RS-LHCONE will be automatically updated from CRIC
- The route-set will be used by NREN and sites to build security routing filters

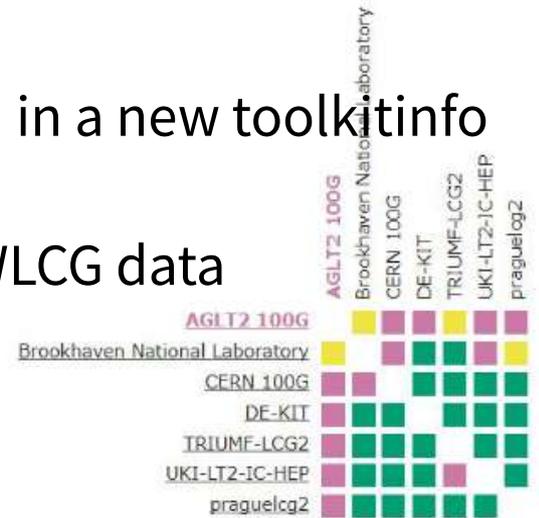
Next steps:

- implement records
- populate database, from LHCONE twiki page and LHCONE routing tables
- agree with WLCG operations on how to keep records up to date

LHCONE monitoring update

- perfSONAR latest version is 4.3.4. 4.3 was a major upgrade with the move to Python3
- RNP Brazil has joined the perfSONAR collaboration
- 207 production endpoints
- Improvement in the 100G mesh
- Using new Kibana dashboard. All resources are organized in a new toolkit info web page
- Additional monitoring will be needed for the upcoming WLCG data challenges

perfSONAR



<https://indico.cern.ch/event/983436/contributions/4232157/attachments/2213899/3748225/LHCONE%20Monitoring%20Update%20Spring%202021.pdf>

JUNO update



JUNO is a neutrino detector in China. It detects the neutrino generated by two nuclear plants nearby

- Status: civil work started, caves almost ready
- Data taken foreseen for late 2021 has been postponed due to pandemic
- 78 institutes participate in JUNO
- Estimated to produce 2PB/year
- Data stored at: iHEP (Tier-0, CN) JINR (RU), MSU (RU), IN2P3 (FR), CNAF (IT)
- Data will be moved first to CNAF, then from there to the other data centres
- Rucio will be adopted
- Request to join LHCONE completed



<https://indico.cern.ch/event/983436/contributions/4276086/attachments/2213944/3747717/Andronico%20JUNO%20update.pdf>

LHCONE AUP review

LHCONE Acceptable Use Policy being reviewed:

- full text review, more readable now
- security part shortened and pointer to relevant WLCG groups included
- more specific to WLCG

Agreed draft available here:

<https://docs.google.com/document/d/1BUjk51LZ4ivYzvAGmxEL2obxihRVTfx6JVykT0YHVkU/edit>

New version just published in the LHCONE twiki:

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



WLCG Network and Data challenges

- The computing model of the experiments at HL-LHC will be different from what they have today
- ATLAS and CMS will produce 350PB/year/experiment. To be exported in real time to Tier1s
- This would require 4.8Tbps from CERN to the Tier1s, of which 1.25Tbps over the Atlantic
- Two scenarios considered. The minimalistic one is the minimum capacity expected. The Flexible scenario would allow the experiments to exercise better the system and try to improve the efficient use of the resources.
- Larger Tier1s are supposed to get connected to CERN and to their Tier2s at 1Tbps (1Tbps in from Tier0, 1Tbps out to Tier2s)
- Based on these targets, a plan for data challenges is being proposed for the years preceding Run4. First challenges will start at the end of 2021
- Data challenges will use the production infrastructure and will co-exist with production activities.
- Data challenges are being discussed in the DOMA-TPC sub-wg



WLCG
Worldwide LHC Computing Grid

<https://indico.cern.ch/event/983436/contributions/4226012/attachments/2213578/3746895/LHCOPN-LHCONE-Mar2021.pdf>

Data-centre network architecture

Agreed to organize workshop on data-centre network architectures

Pre-GDB of the 7-8 of June 2021

<https://indico.cern.ch/event/1028690/>



NOTED achievements

NOTED is a framework that can detect large FTS data transfers and trigger network optimizations to speed up their executions

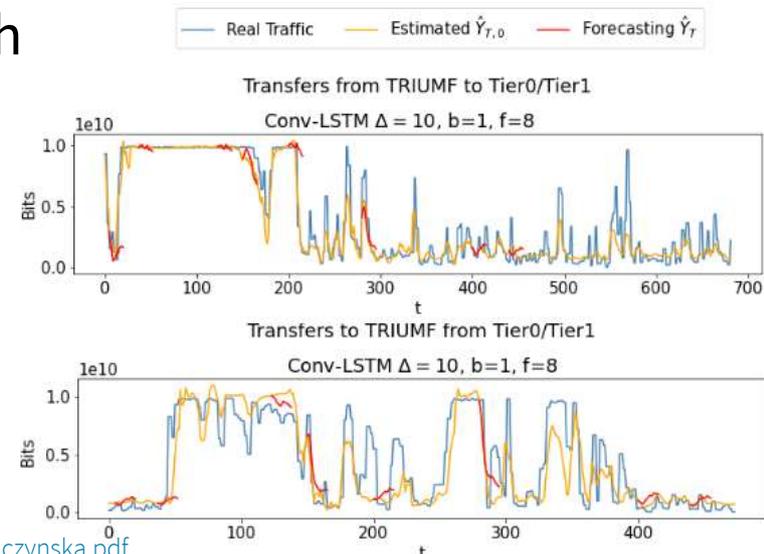
Two successful tests recently:

- CERN-PIC with LHCOPN-LHCONE load balancing
- CERN-TRIUMF with activation of larger bandwidth circuit

Developed new algorithm using machine learning to predict volume and duration of the transfer

Two papers submitted to vCHEP 2021

https://indico.cern.ch/event/983436/contributions/4226042/attachments/2214496/3749482/NOTED_achievements_Waczynska.pdf



Research Network Technology WG - update

Working on network research projects, with contributions not only from WLCG and NREN, but also from outside, like RFC editors, Linux kernel developers

Packet marking activity: considered many options: multiple addresses, IPv6 headers, MPLS. IPv6 flowlabel seems to be the more promising

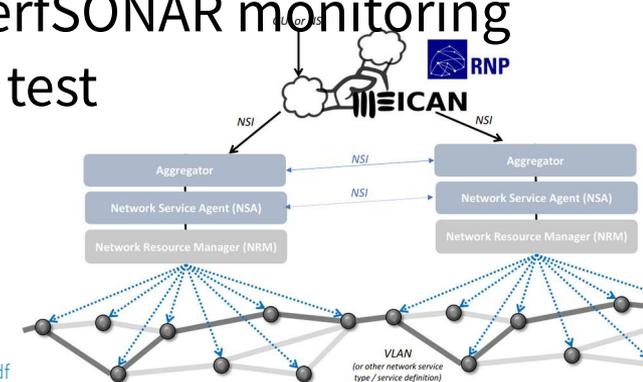
- Proposed a packet marking schema for IPv6 flowlabel field
- Testing already started using iperf3 and other tools
- Targetting implementation in perfSONAR and XRootD

BitPattern	ScienceDomain	Application	Hdr Bit 12	Hdr Bit 13	Hdr Bit 14	Hdr Bit 15	Hdr Bit 16	Hdr Bit 17	Hdr Bit 18	Hdr Bit 23	Hdr Bit 24	Hdr Bit 29	Hdr Bit 30	Hdr Bit 31
xx100000000x000001xx	ATLAS	perfSONAR	x	x	1	0	0	0	0	x	0	1	x	x
xx010000000x000001xx	CMS	perfSONAR	x	x	0	1	0	0	0	x	0	1	x	x
xx110000000x000001xx	LHCb	perfSONAR	x	x	1	1	0	0	0	x	0	1	x	x
xx001000000x000001xx	ALICE	perfSONAR	x	x	0	0	1	0	0	x	0	1	x	x
xx101000000x000001xx	BelleII	perfSONAR	x	x	1	0	1	0	0	x	0	1	x	x
xx011000000x000001xx	SKA	perfSONAR	x	x	0	1	1	0	0	x	0	1	x	x
xx111000000x000001xx	LSSST	perfSONAR	x	x	1	1	1	0	0	x	0	1	x	x
xx000100000x000001xx	DUNE	perfSONAR	x	x	0	0	0	1	0	x	0	1	x	x

https://indico.cern.ch/event/983436/contributions/4231754/attachments/2214549/3748864/Research%20Networking%20Technical%20Working%20Group%20Update%20%28March%202021%20LHCOPN_LHCONE%29.pdf

AutoGOLE and DTNs activity update

- AutoGOLE infrastructure is growing. On-going effort to connect Guam, Hawaii, Singapore, Kaust, South Africa
- The SURFnet AutoGOLE connects Moxy (CA), ESnet and GEANT in Amsterdam, CERN
- Circuit provisioning demonstration done at SC20
- AutoGOLE could be enhanced to provide more advanced service beyond layer2, like multiple VRF, circuits for DTN transfer, perfSONAR monitoring
- Provisioned dynamic circuit CERN-TRIUMF for NOTED test



<https://indico.cern.ch/event/983436/contributions/4226052/attachments/2215117/3749870/Gerben%20van%20Malenstein%20-%20LHCONE46%20AutoGOLE%20SENSE.pdf>

GNA-G DIS Working Group update

Data Intensive Sciences WG: working on challenges presented by high requirements of HL-LHC and other sciences

A special focus of the group is to address the growing demand for

- Network-integrated workflows
- Comprehensive cross-institution data management
- Automation
- Federated infrastructures encompassing networking, compute, and storage

Many projects and testbeds:

- AutoGOLE and Sense
- RARE: working on a router process and P4
- caches for WLCG

ROBIN project update

Comparing Rucio/FTS vs Rucio/SENSE/DTNs

Used DTNs at Starlight and CERN

The tests results claim SENSE/DTNs outperforms FTS between 2 and 30 times for a single stream transfer over a long RTT link.

Suggested to repeat the test using FTS with HTTPS-TPC

Conclusions

Summary

- Traffic decreased during LS2, but upgrades are on-going. Getting ready for Run3
- LHCONE AUP reviewed, new version expected soon
- Network information being added to CRIC
- BelleII in full data taking, JUNO delayed by pandemic
- LHCONE community will contribute and support WLCG Data Challenges
- Several on-going R&D activities to prepare for Run4: packet marking, NOTED, AutoGOLE, SENSE, DTNs

Actions for next meeting

- Complete CRIC and route-set implementation
- Organize workshop on data-centre network architectures [scheduled]
- Publish new AUP [done]

Next meeting

11-12 of October 2021 at CERN (in person, if possible)

Agenda will be published here:

<https://indico.cern.ch/e/LHCOPNE47>

References

Meeting agenda and presentations:
<https://indico.cern.ch/e/lhcopne46>

Minutes:
<https://indico.cern.ch/event/983436/note/>

Questions?

edoardo.martelli@cern.ch

