

The NREN community support to HEP science

SA3CC meeting 2025 - La Serena, CL 7th May 2025 edoardo.martelli@cern.ch



Introduction to CERN, LHC, WLCG, LHCOPN, LHCONE

How the NREN community supports HEP





CERN and LHC



LHC accelerator



LHC major experiments









Next: the HL-LHC project



The High-Luminosity Large Hadron Collider (HL-LHC) is an upgraded version of the LHC

It will operate at a higher luminosity, i.e. it will produce more collisions and data

The HL-LHC will enter service in 2030, **increasing the volume of data** produced by the experiments **by a factor of 10**





WLCG and Networks

LHC Computing Model



Tier 0 (1x) Data source Full data on Tape Data reconstruction



Tier 1s (15x)

Distributed 2nd copy on Tape Simulations, Data analyses



Tier 2s (~150x)

Data caches Data analyses

WLCG

The Worldwide LHC Computing Grid (WLCG) is a large, distributed computing and storage infrastructure and the software framework to exploit it



WLCG networks

Computer **Networks are an essential component of WLCG**; they connect all the computing resources distributed in more than 150 institutes around the world



WLCG sites





LHCOPN

Private network connecting Tier1s to the Tier0

Secure:

- Dedicated to LHC data transfers
- Only declared IP prefixes can exchange traffic
- Can connect directly to Science-DMZ, bypass perimeter firewalls

Technologies:

- L2 VPN
- BGP communities for traffic engineering







LHCOPN



Line speeds:	Experiments:	
20Gbps 100Gbps 200Gbps 400Gbps 800Gbps	= Alice = Atlas = CMS = LHCb	
	Last update: 20240823 edoardo.martelli@cern.ch	

CERN

https://twiki.cern.ch/twiki/bin/view/LHCOPN/OverallNetworkMaps

Numbers
17 sites for 15 Tier1s + 1 Tier0
14 countries in 3 continents
2.88 Tbps to the Tier0

20 years of LHCOPN

First LHCOPN meeting: <u>https://indico.cern.ch/event/420981/</u> 20-21 of January 2005 in Amsterdam

CERN



Grid Deployment Board - GDB

Tier 0/1 Network Meeting

Opening Session

January 20/21 2005

David Foster david.foster@cern.ch





LHCONE L3VPN service



Private network connecting Tier1s and Tier2s

Secure:

- Dedicated to LHC data transfers
- Only declared IP prefixes can exchange traffic
- Can connect directly to Science-DMZ, bypass perimeter firewalls

Technologies:

- Multi domain L3 VPN
- BGP communities for traffic engineering





Open to other HEP collaborations













LHCONE status



- VRFs: 32 national and international Research Networks
- Connected sites: ~110 in Europe, North and South America, Asia, Australia
- Trans-Atlantic connectivity provided by ESnet, GEANT, Internet2, RedCLARA, NORDUnet, CANARIE and SURF
- Trans-Pacific connectivity provided by KREOnet, SINET, TransPAC
- Interconnections at Open Exchange Points including NetherLight, StarLight, MANLAN, WIX, CERNlight, Hong Kong, Singapore and others

WLCG network requirements for HL-LHC

Each Major Tier1s:

1 Tbps to the Tier0 (LHCOPN) 1 Tbps to the Tier2s (aggregated, LHCONE)

Each Major Tier2s:

>400 Gbps (LHCONE)

WLCG and the NREN community are already working on the implementation of these requirements

NREN community support

LHCOPN+LHCONE community meetings

- A forum for site mangers, NRENs, experiments to discuss requirements and policies, plan upgrades, design new features
- Meeting two times per year
- Just held meeting #54, hosted by SKAO in UK





NRENs supporting HEP



The NREN community brings

- network know-how and savoir-fair
- collaborating attitude that has allowed to connect HEP centres in any region of the world
- new tools and innovative services, allowing the development of advanced data distribution and computing orchestration frameworks
- advanced technologies, robustness and always increasing capacity, which have driven the evolution of the HEP computing models from mostly-local to mostly-distributed



Computing model evolution





Original MONARCH model

Model evolution in the 2010s'

Fasterdata

Long lasting and authoritative repository for:

- tuning networks and servers for fast, long distance (RTT) data transfers
- Science DMZ architecture

- performance tests

	myESnet	ESnet	
HOST TUNING NET	NORK TUNING SCIENCE DMZ DATA TRANSFER NODES DATA TRANSFER TOOLS	PERFORMANCE TESTING NSF DOCS	
Home » Science DMZ			
Science DMZ	Science DMZ	Short Cuts	
Viotivation	A Scalable Network Design Pattern for Optimizing Science Data	<u>Science DMZ Overview</u> <u>video</u>	
Architecture	Transfers		
Network Components	The Science DMZ is a portion of the network, built at or near the	Data Transfer Node Info	
Performance	campus or laboratory's local network perimeter that is designed	<u>Science DMZ FAQ</u>	
Monitoring	optimized for high-performance scientific applications rather	<u>Set of short Science</u> <u>DMZ videos (at</u> learn psrc org)	
Security	than for general-purpose business systems or "enterprise"		
Practices and Policy	computing.	• Longer Video by the	
AQ	Developed by ESnet engineers, the Science DMZ model	EPOC project	
_earn More	addresses common network performance problems encountered at research institutions by creating an environment	A Comprehensive Tutorial on a Science DMZ	
Science DMZ Community Presentation	that is tailored to the needs of high performance science applications, including high-volume bulk data transfer, remote experiment control, and data visualization.		
	The Science DMZ is scalable, incrementally deployable, and	Cite the Science DMZ	
	Citation information for the <u>SC13 Science DMZ</u>		



perfSONAR



perfSONAR network monitoring platform

- Developed by the collaboration of Internet2, GEANT, ESnet, RNP, Indiana University, University of Michigan
- Toolkits installed at NRENs PoPs and WLCG sites
- Essential to monitor WLCG network performances and investigate issues





LHCONE traffic engineering

LHCONE BGP communities for traffic engineering

Community	Туре	Meaning	Notes
65001:XXXX	Operational	prepend 1x to ASxxxx	Mandatory
65002:XXXX	Operational	prepend 2x to ASxxxx	Mandatory
65003:XXXX	Operational	prepend 3x to ASxxxx	Mandatory
65010:XXXX	Operational	do not announce to ASxxxx	Mandatory

- Implemented by all the NRENs providing LHCONE
- Used by end-sites to improve their trans-continental routing



scitags.org

Network Flow and Packet Marking for Global Scientific Computing

Science Tags: marking of data packets and flows with Experiment and Application IDs for better network accounting

Two options being implemented:

- Tag in the IPv6 flowlabel field (proposed IETF draft: draft-cc-v6ops-wlcg-flow-label-marking)
- Tags (and more info) in UDP fireflies (UDP packets sent in parallel to each flow)

NRENs run fireflies collectors, supports deployment and testing





CNAF-CERN DCI over GEANT-GARR

The LHCOPN link of IT-INFN-CNAF Tier1 is implemented using shared spectrum over GEANT and GARR (Italian NREN) dark fibres

- 4x100Gbps links between CERN and CNAF used for DC24 and now in production
- Ready to be upgraded to 4x 400Gbps
- cost effective technique to get >1Tbps LHCOPN connections already today





IPv6 deployment

IPv6 needed to overcome IPv4 address space shortage

- IPv6 Traffic in LHCOPN: ~90% of the total. Proposal to disable IPv4 in LHCOPN by 2029
- Supported by all NRENs



R&D: NOTED SDN

NOTED is a framework that detects large data transfers in real-time and trigger network optimization actions to speed up the execution of those transfers

- Already used with production data transfers during SC22 and SC23





ESnet SENSE

Software Defined Networking (SDN) for End-to-End Networked Science at the Exascale





Data Challenges for HL-LHC

WLCG has planned for a series of data challenges to prepare for HL-LHC data taking

- Demonstrate readiness for the expected HL-LHC data rates with:

- Increasing volume/rates
- Increase complexity (e.g. additional technology)

2021: 10% of HL-LHC requirements (480Gbps minimal – 960Gbps flexible)
2024: 25% of HL-LHC requirements (1.2Tbps minimal – 2.4Tbps flexible)
2027: 50% of HL-LHC requirements (date and % to be confirmed)
2029: 100% of HL-LHC requirements (date and % to be confirmed)
2030: start of HL-LHC (Run4) (4.8Tbps minimal – 9.6Tbps flexible)



ESnet High Touch

High Touch collected LHCONE transfers during DC24.

It allowed deep understanding of network flows: short duration, no retransmissions





Global P4 Lab

Network of P4 switches provided by the GNA-G AutoGOLE / SENSE Persistent Multi-Resource Testbed and the GÉANT P4 Lab

Used to demonstrate a LHCONE MultiONE implementation using Flowlabel policy routing





FABRIC and FAB



International research infrastructure funded by NSF and built by ESnet that enables the computer science and networking community to develop and test new architectures



Time and Frequency distribution

CERN is prototyping a frequency distribution service for the LHC experiments, using the precise frequency provided by the REFIMEVE network (French National Metrology Institute) **delivered to CERN by RENATER**

https://home.cern/news/news/experiments/optical-fibre-link-make-cern-more-time-ever





Conclusions



(Intercontinental) connectivity of European research networks, August 1996

From the 90s...





...to the Future Circular Collider...

Purely technical schedule, assuming green light to preparation work **A 70 years scientific program**

8 years preparation	10 years tunnel and FCC-ee construction	15 years FCC-ee operation	11 years FCC-hh preparation and installation	25 years FCC-hh operation pp/PbPb/eh
2020-2028		2038-2053		2064-2090





An effective collaboration

Computing and networking are essential components of the LHC instrument

The NREN community has been supporting WLCG and HEP science over many years: an outstanding collaboration

Nobel prize in Advanced Collaborative Science!!



Photo credits: Nobel Media AB 2013. Photo: Alexander Mahmoud



Questions?

edoardo.martelli@cern.ch

