

The GÉANT network: addressing current and future needs of the HEP community

Vincenzo Capone, Mian Usman

GÉANT – Cambridge (UK)

vincenzo.capone@geant.org, mian.usman@geant.org

Abstract. The GÉANT infrastructure is the backbone that serves the scientific communities in Europe for their data movement needs and their access to international research and education networks. Using the extensive fibre footprint and infrastructure in Europe the GÉANT network delivers a portfolio of services aimed to best fit the specific needs of the users, including Authentication and Authorization Infrastructure, end-to-end performance monitoring, advanced network services (dynamic circuits, L2-L3VPN, MD-VPN). This talk will outline the factors that help the GÉANT network to respond to the needs of the High Energy Physics community, both in Europe and worldwide.

The Pan-European network provides the connectivity between 40 European national research and education networks. In addition, GÉANT also connects the European NRENs to the R&E networks in other world region and has reach to over 110 NREN worldwide, making GÉANT the best connected Research and Education network, with its multiple intercontinental links to different continents e.g. North and South America, Africa and Asia-Pacific.

The High Energy Physics computational needs have always had (and will keep having) a leading role among the scientific user groups of the GÉANT network: the LHCONE overlay network has been built, in collaboration with the other big world REN, specifically to address the peculiar needs of the LHC data movement. Recently, as a result of a series of coordinated efforts, the LHCONE network has been expanded to the Asia-Pacific area, and is going to include some of the main regional R&E network in the area.

The LHC community is not the only one that is actively using a distributed computing model (hence the need for a high-performance network); new communities are arising, as BELLE II. GÉANT is deeply involved also with the BELLE II Experiment, to provide full support to their distributed computing model, along with a perfSONAR-based network monitoring system. GÉANT has also coordinated the setup of the network infrastructure to perform the BELLE II Trans-Atlantic Data Challenge, and has been active on helping the BELLE II community to sort out their end-to-end performance issues.

In this talk we will provide information about the current GÉANT network architecture and of the international connectivity, along with the upcoming upgrades and the planned and foreseeable improvements. We will also describe the implementation of the solutions provided to support the LHC and BELLE II experiments.



1. About GÉANT

The GÉANT network [1] is the fast, reliable pan-European communications infrastructure backbone serving Europe's research and education community and essential to the success of many research projects in Europe. The network and associated program of activities is co-funded by the European Commission as part of Horizon 2020, in the form of a 7-years Framework Program Agreement, under the name of GN4, which started on May 2015. Matching funding is provided by the European National Research and Education Network (NREN) project partners connected to the network. The GN4 project partners are 39 European NRENs, GEANT Ltd. (the former DANTE) and GÉANT Association (the former TERENA). As the most advanced research network in the world, GÉANT has developed progressively over the last 20 years to ensure that European researchers are at the forefront of international and global collaboration.

The GÉANT network is a world leader in terms of capacity, reach and services offered, and is helping to drive the development of new networking technologies and services. GÉANT is currently based on a new 500Gbps dark fibre backbone. At present, the network connects 40 European countries, reaching 50 million users in Europe in over 10,000 institutions. Over 1,000 terabytes of data is transferred via the GÉANT IP backbone every day. As a result, this backbone has become not just an infrastructure for e-science but a positive demonstration of European integration. With connection to 65 countries outside of Europe, bringing the total number of users reached to 100 million, GÉANT is at the heart of a truly global network.

Responsibility for the network is logically allocated: NREN organisations operate and manage their respective national network, while GÉANT maintains and operates the GÉANT network on a pan-European basis. GÉANT is able to offer more flexible and bespoke options than commercial operators because of its unique remit to deliver real value and benefit to society by enabling research communities to transform the way they collaborate, and due to its not-for-profit status.

GÉANT is not only network connectivity, anyway, but also offers a portfolio of added-value innovative services to the R&E community, relating in such areas as trust and identity, security and certification, mobility and access, test-bed facilities, cloud services and media and real-time communications

Some of the services that are here worth mentioning are:

- eduROAM [2] (federated Wi-Fi access)
- eduGAIN [3] (trust and identity federation)
- perfSONAR [4] (network monitoring tool)
- TCS [5] (security certification authority)

Something that recently has been created to better serve the scientific community is the One-stop-shop concept. Given the structure of the GÉANT+NRENs network, for the international collaborations that spread among several countries this would mean for every single site in every country to engage the community only by means of the local NREN. With the One-stop-shop function, it's possible for the community to have a single point of contact in GÉANT, which will take care of liaising with all the involved NRENs and will also help in providing solutions with a holistic view of the community, with a better understanding of their collective needs.

2. GÉANT backbone

During 2012 the GÉANT network went through a deep re-engineering phase. The result has been the migration to a new transport and routing/switching platform, completed in 2013. The new network is based on the latest technology available, both on the optical and on the Ethernet/IP layers, providing optical transmission of up to 500G SuperChannels (on the part of the network implemented on dark

In particular, the MX series offers support for point-to-point as well as multipoint-to-multipoint, native support for 100GE, with a backplane throughput of 120Gb/slot. The routing capabilities are also on par with other market-leaders' products, with the additions of virtualization and SDN capabilities.

3. International connectivity

One of the main objectives of GÉANT is connecting the European NREN community with the other partner networks around the world. This is achieved in different ways: directly peering with the partner network, peering via an Internet Exchange, or in some cases taking an active part to the development of




Figure 2 Global map of international connectivity

new regional networks. A remarkable example of the first case is the peering with ESNNet, on whose behalf GÉANT has realized the intra-European infrastructure of the ESNNet extension to Europe [8].

The latter case mainly happens in less developed world regions (e.g. Middle-East, Mediterranean-African countries, Asia-Pacific, Africa), or in areas where the cost of connectivity is particularly high (e.g. China), requiring additional help in terms of funding from the European Commission. In these cases, new projects have been created and funded, managed by GÉANT on behalf of the regional R&E communities. GÉANT also fosters and promotes the creation of the required skills and infrastructure, so to handover the management of the newly created network to professionals from the community that the network serves. Winning examples of this strategy are the TEIN network, connecting the community of the Asia-Pacific countries, the UbuntuNet Alliance, connecting the east-African countries, and the RedCLARA Project, that connects the Latin American NRENs to the GÉANT network. In some cases, like for the CAREN and EUMEDCONNECT project, has not been possible so far to achieve this goal, so GÉANT is still in charge of operating and managing these networks.

Here follows a series of tables showing the connections to the different world regions, with details on the type of connection and the capacity.

3.1. North America

NREN	Type of connection	Capacity
Internet2, Canarie	Exchange point	100G GÉANT link Paris-New York 30G from the ACE collaboration
Internet2, Canarie	Exchange point	200G from ANA-200G
ESNet	Direct peering	340G (LON, AMS, GVA)

3.2. Asia

NREN	Type of connection	Capacity
SINET4	Exchange point	20G via North America
SINET5 (exp. 1H2016)	Direct peering	20G direct + 10G via N.A.
TEIN4	Direct peering	2.5G Madrid-Mumbai
TEIN4 (exp. 3Q2015)	Direct peering	10G from Singapore
CERNET/TEIN4	Direct peering	10G to ORIENT+
CARNET	Direct peering	622Mb + 155Mb

3.3. Africa

NREN	Type of connection	Capacity
UbuntuNet Alliance	Direct peering	20G via AfricaConnect
EumedCONNECT3	Direct peering	2 x 622Mb

3.4. Latin America

NREN	Type of connection	Capacity
RedCLARA	Direct peering	5G




Figure 3 The LHCOPN network (in red the links provided by GÉANT)

4. The support to High Energy Physics

4.1. LHC

The LHC computing, that avails of the World LHC Computing Grid (WLCG) organization, is currently the most network-intensive science activity in the R&E world. To better satisfy the needs of this community, two special network services have been designed and built over the years: LHCOPN and LHCONE.

The LHCOPN is an Optical Private network, a closed network infrastructure that was originally intended to serve T0-T1 traffic, and was later broadened to include T1-T1 traffic. Currently, 13 Tier1s are connected directly to CERN over LHCOPN and there are also several connections between large T1s. LHCOPN is built as a collection of privately procured circuits, by each Tier1, to connect to CERN. Some Tier1s have a double connection to CERN, while some have connections to other Tier1s as well.

The rationale behind the design of the LHCOPN was the MONARC [9] model, the strict hierarchical data distribution model, based on regional aggregations of T1 and T2s, that constituted the basis for the original LHC computing model. That model was gradually abandoned over time, leading to a more chaotic, full-mesh transfer pattern between the WLCG sites, allowing any T2/T3 to exchange data with any T1/T2 [10]. The answer to this change, in terms of networking, was the creation of the LHCONE (LHC Open Network Environment).

The current LHCONE network [11], which was implemented in early 2012, is based on interconnected VRF (Virtual Routing and Forwarding) instances, instantiated on the backbone networks of the participating regional and national RENs (European NRENs, GÉANT, ESNNet, and Internet2). These VRFs are directly connected to each other or via several Open Exchange points (ManLan, WIX, StarLight, CERNLight, and NetherLight) where these networks establish peerings with each other. Currently, over 50 Tier2s and 11 Tier1s are connected to LHCONE, over 15 different network operators, in Europe, North America and Asia. This infrastructure is on the verge of being extended also to new world regions. The work is rapidly proceeding to connect Latin America, via the RedCLARA link between GÉANT and the NRENs on the continent, starting with RNP (the Brazilian NREN) and a few sites in Brazil. There is already interest from Argentinian and Mexican WLCG sites to connect, as soon as the infrastructure is in place, which is expected to happen before summer.

A similar work is going on in Asia-Pacific, where GÉANT is liaising with TEIN to establish a LHCONE connection to the area: the first candidate country should be Pakistan, where two important WLCG sites are present, and PERN, which is the country's NREN.

4.2. BELLE II

The liaison between GÉANT and the BELLE II community has started at a very early stage of the experiment, since the first contacts date back to 2013. The first concrete steps were to help the community setup a perfSONAR-MDM based monitoring infrastructure. In a period of time stretching from January to May 2014, all major sites had a perfSONAR measurement point in place, also thanks to the GÉANT support to the local IT people at the sites; a User Interface was also installed at Pacific Northwest National Laboratory.

The BELLE II experiment was also involved in the users' testing campaign over the ANA-100G circuit: this was the result of an agreement between GÉANT, ESNNet, Internet2, CANARIE, NORDUNet and SURFNet, to procure and operate the first 100G transatlantic link. A number of science projects were then involved to test high-bandwidth transfers over this link, and BELLE II was one of them.

The joint work of BELLE II computing people and GÉANT users support led to the planning and implementation of the Trans-Atlantic Data Challenge, which was coordinated by GÉANT, with the combined efforts of GÉANT, ESNNet, GARR (Italian NREN) and DFN (German NREN). An international testbed infrastructure was set up, with the purpose to measure the achievable data rates during the simulation of the movement of 1 days' worth of raw data production (approximately 25




Figure 4 Intra-USA and Trans-Atlantic topology

TBytes) from PNNL to three European sites: Italy’s INFN-Napoli and INFN-CNAF, and Germany Karlsruhe Institute of Technology (KIT). The activity was spread across two timeframes, 12th-15th of May 2014 and 10th-20th of June, and has been very helpful to the BELLE II computing team to properly measure network throughput without bottlenecks on the path, as well as to address some last-mile network configuration issues.

The configuration for the Data Challenge was coordinated by GÉANT, and realized by engineers from the involved NRENs (GÉANT, ESNet, GARR and DFN) and the end sites (PNNL, KIT, CNAF,




Figure 5 The intra-European topology

INFN-Napoli), and as can be seen in Figure 4 and 5, the final topology spanned across the US, the Atlantic Ocean and Europe, up to Germany and Italy.

To start with, a single VLAN ID was agreed on by all parties, so to ease the configurations and troubleshooting, avoiding the use of VLAN ID swapping. Starting from the source of the traffic, the Pacific Northwest National Laboratory, the following actions were put in place:

- a 10G link was established to the ESNNet PoP over a dedicated connection
- a 10G best-effort Lambda Switched Path was created over the ESNNet MPLS domain
- dedicated VRFs were created on the ESNNet, GEANT and GARR backbones for this data challenge
- the ESNNet VRF was then connected to the GÉANT VRF via the two exchange point of MANLAN, in New York, and NetherLight, in Amsterdam, using the single VLAN ID over the ANA-100G link between New York and Amsterdam
- the GEANT VRF was connected to the GARR VRF to connect to the two end-sites in Italy
- DFN provided a L2 p2p link from their GEANT PoP in Frankfurt to KIT, still using the same VLAN ID.

The result of all these efforts was a full 100G capacity between ESNNet, GEANT, GARR and DFN.

5. Conclusions

The recently established GÉANT international user support service (One-stop-shop) has proved itself a valuable service to users, providing a single point of contact/coordination for the users to interact with the NREN community. The BELLE II Data Challenge case, especially, is the visible outcome of such activity.

The High Energy Physics community has been and will remain a top-priority user group for GÉANT, currently representing almost half of the total traffic flowing on the backbone. LHCONE, in particular, stands out as one of the most successful examples of multi-domain collaboration between the NRENs, proved by the growing number of physics collaborations interested in joining it.

Also, the constant improvement of the extra-EU connectivity will provide a tangible benefit to the HEP community overall.

Lastly, the recent deep renovation and improvement of the network backbone has made the GÉANT network already capable to meet tomorrow's user needs.

References

- [1] <http://www.geant.org>
- [2] <https://www.eduroam.org/>
- [3] <http://services.geant.net/edugain/Pages/Home.aspx>
- [4] <http://services.geant.net/perfsonar/Pages/Home.aspx>
- [5] <https://www.terena.org/activities/tcs/>
- [6] <http://www.infinera.com/products/DTN-X.html>
- [7] <http://www.juniper.net/us/en/products-services/routing/mx-series/>
- [8] <http://www.es.net/news-and-publications/esnet-news/2014/esnet-extends-100g-connectivity-across-atlantic/>
- [9] M. Campanella, L. Perini, "The analysis model and the optimization of geographical distribution of computing resources: a strong connection" - <http://www.mi.infn.it/~cmp/rd55/rd55-1-98.html>
- [10] <http://lhcone.web.cern.ch/node/19>
- [11] <https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome>