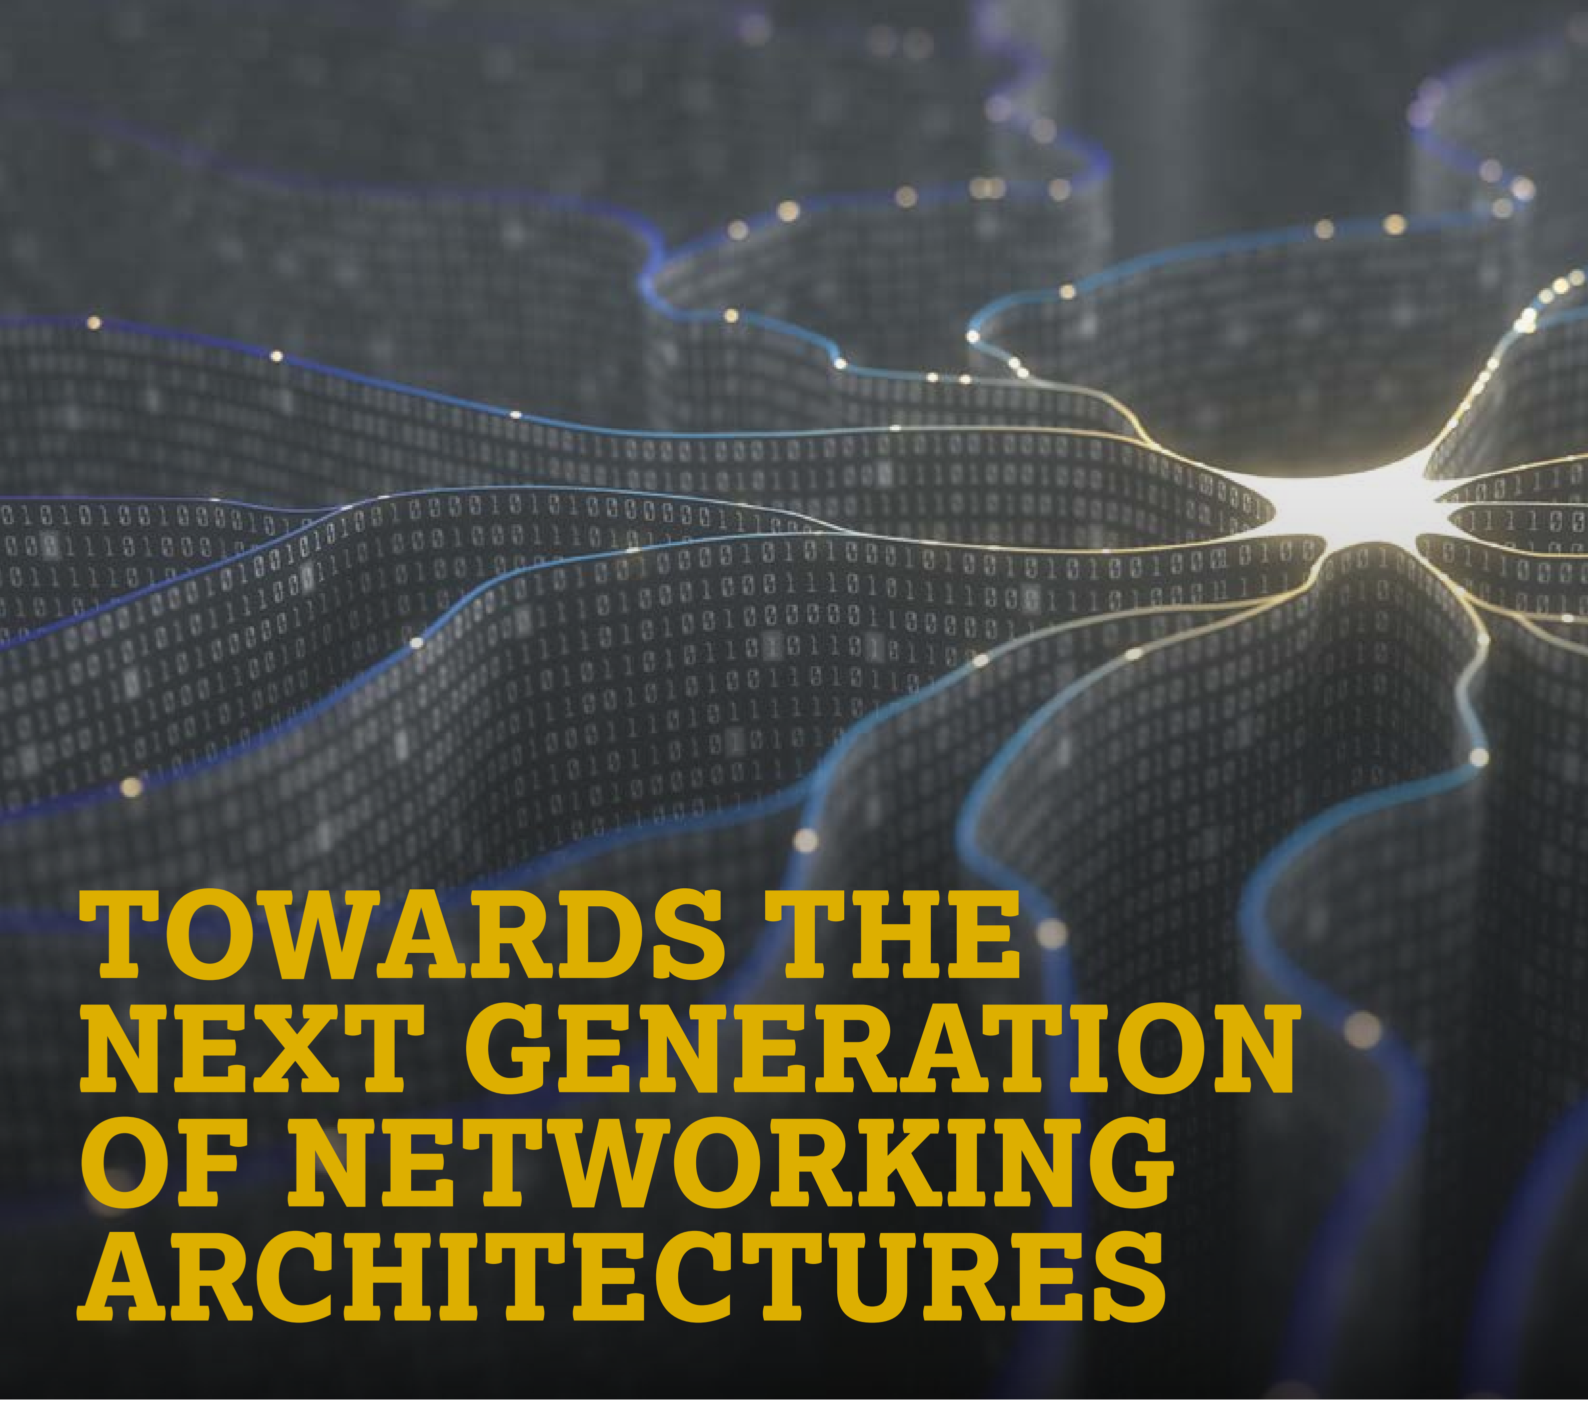

TOWARDS THE NEXT GENERATION OF NETWORKING ARCHITECTURES

The evolution of networking architectures has been a long path towards increased speed, capacity and functionality with different techniques used to support the demands of users and applications. GÉANT is working with the leading suppliers and other operators to help guide the direction of the next generation of networking services to future-proof our network and provide the services that NRENs and users need into the next decade.

The evolution of advanced networking

Until recently equipment operators have been looking to converge their platforms into unified communications equipment able to handle multiple network layers together in the same device. The key driver for convergence was a unified network management system which was used to manage layer 2/3 service and network elements. The operators benefit from convergence as it reduces the number of boxes (network elements) and simplified the management along with reducing the CAPEX/OPEX.

To support the move from Synchronous Digital Hierachy (SDH) based architectures to the more flexible MPLS structures, GÉANT converged its layer 2 and 3 equipment in 2011 on the basis that both layers could support similar functionality and a single box would make management easier and reduce OPEX.

However this trend towards equipment convergence seems to have stalled. The problem with convergence is the network operators end up with a monolithic block of hardware and software. Operators become dependent on a single vendor's innovation curve for new features in both hardware and software. In many cases the vendor's Operating System software has become so large that innovation cycles have become excessively slow. This restricts the ability of operators to offer their users new, innovative services or to leverage emerging techniques such as Alien Waves to reduce costs and pass savings onto their users.

Software Defined Networking and the benefits of disaggregation

Software Defined Networking (SDN) has reversed this convergence trend and in essence enabled providers to disaggregate not only the software from the hardware, but from vendors. This separation of hardware and software enables them to innovate independently and further disaggregation (modularisation) of functions within software and hardware allows operators to buy the only the modules that meet their requirements.

Using this approach the operators can build the network using best of breed hardware or software component instead of relying on a single vendor to deliver new features in both hardware and software. A router, switch or optical network can be built using the best modules/blocks available to deliver at scale and using the best technology available.

Efficient scaling is one of the major benefits of disaggregation in optical layer. A 'building-block' approach to the transport layer allows for low initial spend on an open line system and enables operators to grow their capacity incrementally as traffic increases. It also allows operators to exploit third party coherent optics or external transponders. This means when a new interface type or feature is needed it can be bought separately and deployed to run on top of an existing open line system.

No gain without pain

Disaggregation enables innovation but also presents some challenges especially in optical networks. The DWDM optical systems are 'analogue' with associated complexities in planning, configuration and control. An intelligent software management system for optimisation and control is needed to simplify deployment of disaggregated packet-optical solutions.

Applying the concept of disaggregation in a network means that network operators have to take responsibility to glue together different components of hardware and software. It's not just a technical change it's a cultural shift which breaks long-standing networking policy and would complicate support. For example, who do you call when there is a problem? The hardware company or the software company? It also requires an entirely new skill set in the operations team, and strong abilities in the areas of software architecture and coding are required.

Industry initiatives – GÉANT and Industry working together

There are many initiatives by web-scale companies to facilitate the development of network hardware and software in an open and collaborative environment and GÉANT is actively involved in two of the groups:

- OpenConfig.
- Telecom Infra Project (TIP).

The OpenConfig is a working group of network operators sharing the goal of moving networks towards more dynamic and programmable infrastructure by adopting SDN principles.

Telecom Infrastructure Project

The main objective of the Telecom Infra Project (TIP) is to develop an open optical transport system and disaggregate hardware and software, sustainably reduce TCO (total cost of ownership) of the transport network equipment and create open source management framework which allows for rapid adoption of new technologies. TIP has led the development of the Voyager 1U white-box switch with the objective of accelerating innovation while reducing costs.

What does this mean for GÉANT?

The aim for GÉANT is to be able to deliver a flexible, cost-effective and sustainable network architecture to support the next generation of networking requirements and support our partners and users. By working alongside vendors and network operators, GÉANT is well placed to influence and benefit from the next generation of SDN systems and infrastructure.

To find out more about the Telecom Infrastructure Project visit **https://telecominfraproject.com/**

CONNECT talks to the GÉANT Operations Team to find out more about the network, how efficiently it is run and to share their plans on how to handle the exponential growth in traffic.

How much traffic does the GÉANT network carry every day?

During 2016, GÉANT received an average of 3.9 Petabytes (that's 3.9 million Gigabytes) of data per day. That adds up to a total of more than 1.4 Exabytes of data accepted and transmitted by the GÉANT network in 2016.

What sort of traffic is this, and who are the big data producers?

GÉANT receives traffic on its two main networks: the Lambda network delivered by the Infinera DWDM system and the GÉANT MPLS/IP network delivered on Juniper MXs.

GÉANT Lambdas provide customers with 10Gbps and 100Gbps point-to-point Ethernet services and this accounts for about 40% of the total GÉANT traffic. Of this, CERN has the lion's share, amounting to approximately 80%, with other major users such as PRACE (see page 16), LOFAR (Low Frequency Array), GTS (the GÉANT Testbed Service) and others with much lower volumes.

The GÉANT MPLS/IP network provides all other services, such as GÉANT R&E IP, LHCONE, Internet access, GÉANT Plus, MDVPN and BoD, and receives the remaining 60% of the total traffic. 89% of the traffic received by this network is absorbed by three main services: GÉANT R&E IP 41%, LHCONE 23%, Internet access 25%, whilst the remaining 11% is shared among all other services. Also here CERN, with its high-energy physics community, is still the biggest data producer.

Of the NRENs, the largest users are DFN (Germany) and Jisc (UK), each accounting for over 15% of the total amount of data produced by NRENs, followed by GARR (Italy) and RENATER (France), then RedIris (Spain), with a share of 4% whilst all other NRENs' share of the traffic is lower than 3%. Data for consumption may show a different distribution, but the largest NRENs remain the main users.

What makes the network so special?

Firstly, GÉANT is a virtually lossless network, meaning that any data received by GÉANT is transmitted to its destination without dropping any packets. In order to achieve this, GÉANT actively manages capacity to accommodate for bursts and periodic changes in the traffic volume. Transmission of traffic at line rate should be possible at all times with minimal or no buffering in order to maximise throughput of host data flows. This way data is transferred unchanged and protocols running on hosts can be tuned to work on the assumption of transiting a virtually lossless environment. GÉANT also allows for elephant flows (single flows of very high rate use by HPCs) by keeping its core links running on the highest speed interfaces available on the market.

Secondly, GÉANT is a high-speed network and links are procured to ensure that delay between GÉANT PoPs in major European cities is as small as possible. This also ensures the best response time between any two entry points to the GÉANT network.

Thirdly, GÉANT strives to connect all countries with enough capacity to allow for high-speed networking. This enables the community to bridge the digital divide with the provision of adequate connectivity in countries where such resources are scarce and the market is closed.

How has the amount of traffic grown over the past few years?

From 2015 to 2016 traffic on the IP/MPLS network grew year-on-year by 64%, meaning that our traffic volume on this network is doubling every 15 months.

The growth has been driven by R&E, where the traffic volume has grown by more than 70% (LHCONE and GÉANT R&E IP), while in contrast, commercial Internet access traffic has grown by just 30%. Interestingly, this GÉANT commercial Internet traffic growth is in line with the Internet traffic growth seen by commercial providers.

The R&E community is therefore faced with a major challenge: its traffic growth rate is over twice as large as the commercial Internet's growth rate.

The growth rate for the Lambda services amounted to 12% approximately, however growth in this type of network is mainly related to new service uptake rather than to traffic growth within existing services.

What implications does this have?

Should R&E traffic growth continue at this rate, in 10 years' time GÉANT would receive 140 times the amount of traffic it receives today. By comparison, within the same time frame, a commercial provider with a growth rate of 30% would see its traffic increase 14 times. The main implication of this is that classic/standard solutions cannot work for GÉANT in the longer-term and GÉANT must look at innovative ways of delivering its network.

What is GÉANT doing to plan for this?

Classic standard solutions are based on pieces of equipment with very broad sets of features, such as a Juniper MX router. Devices like this are able to provide all services and in the past were used to aggregate services delivery into a single high feature-set box. The downside of this is that the cost-per-bit on those devices is the cost-per-bit required by the data-stream that is the most complex to deliver.

Any other data-stream whose requirements, in terms of flow handling, are much more limited, still transits through the same expensive ASICS and is handled by the same complex software.

This results in the payment of premium to transit traffic that could be delivered by a much simpler and less expensive network.

We are analysing traffic, gaining a deeper understanding of the various flows transiting our network and of the requirements from the network point of view for each flow. GÉANT is striving to understand which minimum set of features is required by the equipment that needs to service each group of flows with common requirements. We will subsequently look at groups of flows with minimum common requirements, whose aggregate traffic is large enough to justify disaggregation and find the most cost-effective way of delivering such traffic. All this will then come together in a future network architecture where the cost-per-bit is as optimised as possible and only the small amount of traffic requiring complex handling by expensive ASICS is transiting the devices providing the required capabilities.

The result is a modular layered architecture where best-in-breed for each module and layer can be selected with minimised impact on other modules/layers. A guiding principle in this architecture is ensuring that flows are handled at the lowest possible layer, which are the layers closest to the physical media.

Are there technologies that GÉANT is investigating/deploying to meet future needs?

Faced with such a high network growth-rate, GÉANT needs to look for solutions investigated and deployed by organisations with similar extreme growth levels: the datacentre world. In the past few years datacentres/cloud service providers have needed to develop their own solutions in order to be able to deal with their own growth. Primary examples include: Google, Microsoft, Amazon and Facebook. These content providers had to start optimising their networks; but as early adopters, they needed to drive the industry to produce the hardware and software they need.

These solutions have initially been tailored to the datacentre arena and consequently, in many cases, result in solutions that are of little use to ISP networks such as GÉANT. Now, content providers are expanding beyond the datacentre, into the WAN, and have started generating major disruption in the WAN market whilst driving the industry to quickly provide solutions. GÉANT is looking at this phenomenon very closely and preparing to deploy the right combination of technologies as they become available.

In line with this, GÉANT has been following trends such as open line systems, alien waves, packet optical integration, coherent optical networking, open-source hardware projects (TIP), merchant silicon evolution, SDN and white boxes.

In particular, a big trend in the industry is the move toward centralisation of control plane, multilayer orchestration and programmability: forward-looking decision making based on a more centralised and holistic view evaluating the use of a richer set of inputs. GÉANT plans to follow this approach to enable its partners to interact more directly with network resources. NRENs and other e-infrastructures will have a view and understanding of resources utilisation, they will be able to reserve capacity whilst influencing the network behaviour to fit their specific needs. Bi-directional communication between software running the network and software utilising the network as a resource should improve, and in the longer term this will improve service and customer experience as well as network utilisation and costs.

Acronyms table

LHCONE - LHC specific L3VPN network

GÉANT Plus - point to point Ethernet over MPLS service

MDVPN - Multi Domain VPN, Carrier of Carriers service

BoD - Bandwidth on Demand, dynamic point to point multidomain Ethernet service

HPC - High speed computing

SDN - Software defined networking

TIP - Telecom Infra Project, an opensource hardware initiative sponsored by Facebook, among others