

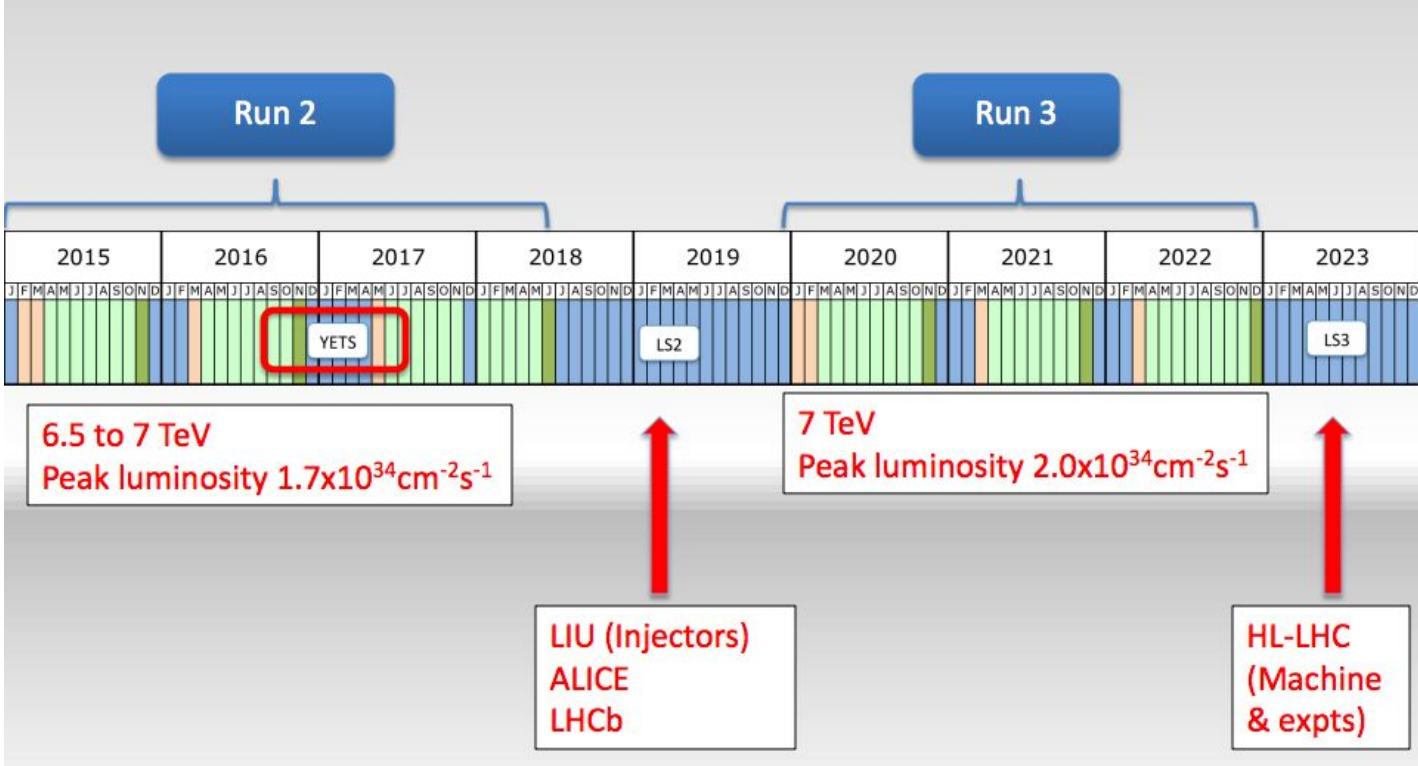
LHC Network Evolution

Marian Babik with input from
Shawn McKee and Edoardo Martelli
LHCOPN/LHCONE workshop, Helsinki

Introduction

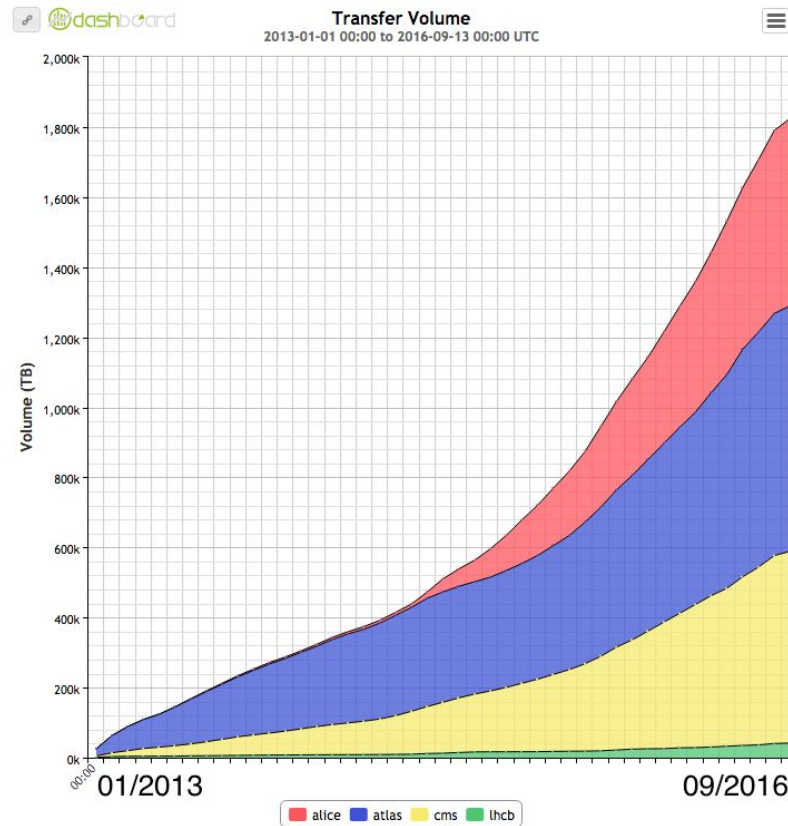
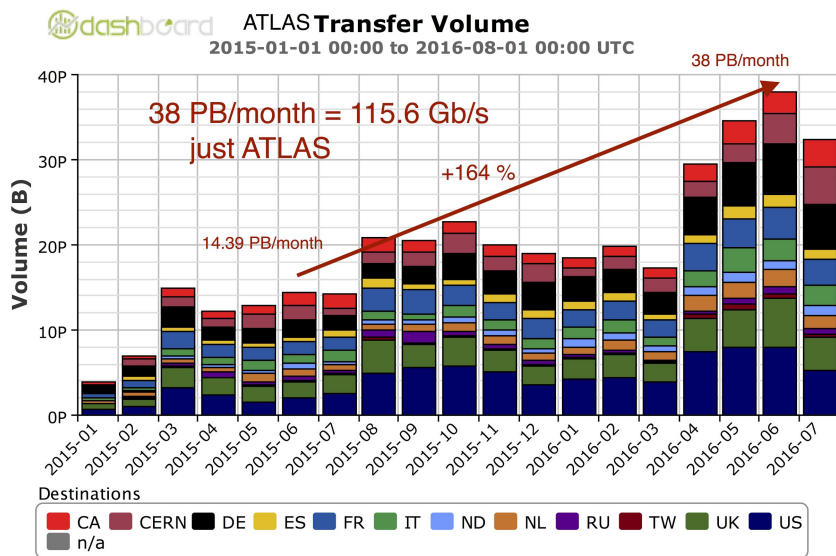
- High Energy Physics (HEP) has significantly benefited from strong relationship with Research and Education (R&E) network providers
 - Thanks to LHCOPN/LHCONE community and NREN contributions, experiments enjoy almost “infinite” capacity at relatively low (or no-direct) cost
 - NRENs were able to continually expand their capacities to overprovision the networks relative to the experiments needs and use
- Right now, experiments have no reasons to believe this situation will change in the long-term
 - Aim of this talk is to stimulate discussion on the topic of LHC network evolution (5 years)
 - Discuss emerging trends, their potential impact and how to better engage experiments in the process of evolving our networks

LHC schedule




Traffic growth

Experiments have been transferring exponentially increasing amount of data since startup. This trend is likely to continue as it's driven by increasing data volumes, more capable infrastructure and excellent networks.



100 Gbit LAN

- Historically WAN capacity has not always had a stable relationship compared to data-centre
 - In recent history WAN technologies grew rapidly and for a while outpaced LAN or even local computing bus capacities
 - Today 100Gbps WAN links are the typical high-performance network speed, but LANs are also getting in the same range
 - List price for 100Gbit dual port card is ~ \$1000, but significant discounts can be found (as low as \$400), list price for 16 port 100Gbit switch is \$9000
- Today it is easy to over-subscribe WAN links (in terms of \$ of local hardware at many sites)
- Will WAN be able to keep up ?
 - So far, not a problem, by 2020 800 Gbps waves likely available
 - But detailed planning of the hardware upgrades (at sites) will be needed



R&E Networking

- R&E network providers has been working closely with HEP community for a long time
 - HEP has been representative of the future data intensive science domains
 - Often serving as testbed environment for early prototypes
- Big data analytics requiring high throughput is no longer limited to just HEP domain
 - SKA (Square Kilometer Array) plans to operate at data volumes 200xLHC scale
 - Besides Astronomy there are MANY science domains anticipating data scales beyond LHC: Health, Bioinformatics, Engineering, etc.
- **What if N more HEP-scale science domains start competing for the same network resources ?**

Network Operations

- Deployment of perfSONARs at all WLCG sites makes it possible to see and debug end-to-end network problems
- A group focusing on helping sites and experiments with network issues using perfSONAR was formed recently (WLCG Network Throughput)
 - Reports of non-performing links are actually quite common (almost on a weekly basis)
 - Most of the end-to-end issues are due to faulty switches or mis-configurations at sites
 - Few cases also due to link saturation (recently in LHCOPN) or issues at NRENs
- Recent network analytics of LHCOPN/LHCONE perfSONAR data also point out some very interesting facts:
 - Packet loss greater than 2% for a period of 3 hours on almost 5% of all LHCONE links
- It is becoming increasingly important to focus on site-based network operations

Software Defined Networks (SDN)

- SDN is a set of technologies offering solutions for some of the future challenges
- Many different point-to-point efforts and successes reported within LHCOPN/LHCONE
- Still major challenge remains how to bring this to the production-level and how to better engage experiments
 - Potential benefits/impact of deploying production level SDNs unclear to experiments
 - How will SDN impact experiments data management systems and operations ?

SD-WAN

- Large Network as a Service providers include several well established CSPs such as Amazon, Rackspace, AT&T, Telefonica, etc.
- Recently more niche NaaS providers have appeared offering SD-WAN solutions
 - Aryaka, Cloudgenix, Pertino, VeloCloud, etc.
 - Their offering is currently limited and not suitable for high throughput, but evolving fast
- SD-WAN market is estimated to grow to \$6 billion in 2020 (sdxcentral)
- Will low cost WAN become available in a similar manner we are now buying cloud compute and storage services ?

Containers

- Recently there has been a strong interest in the container-based systems such as Docker
 - They offer a way to deploy and run distributed applications
 - Containers are lightweight - many of them can run on a single VM or physical host with shared OS
 - Greater portability since application is written to container interface not OS
- Obviously networking is a major limitation to containerization
 - Network virtualization, network programmability and separation between data and control plane are essential
 - Tools such as Flocker or Rancher can be used to create virtual overlay networks to connect containers across hosts and over larger networks (data centers, WAN)
- Containers have great potential to become disruptive in accelerating SDN and merging local and wide area networks

What's next ?

- WLCG workshop 8-9 Oct in San Francisco with dedicated network session
 - Covering IPv6, LHCONE/LHCOPN and LHC network evolution
- Input of LHCOPN/LHCONE community needed on questions that will likely be raised:
 - What new network technologies will become available (production-level) within the next 5 years and what will be their benefits/impact for the experiments ?
 - Should experiments plan to contribute effort to networking ? If so, what will be the motivation and what areas/projects would be important to cover ?
 - What other domains (HEP or non-HEP) with similar scale as LHC will enter production within the next 5-10 years ?
 - Do we need to change the way we plan network deployment/upgrades and capacity ?
- pre-GDB focused on network is planned January 10th
 - Good opportunity to talk F2F with the experiments

Backup

LHC schedule

