v1.0

Research Networking Technical Working Group

For Users, Providers and Researchers

Joint LHCONE/LHCOPN, HEPiX and DOMA activity

Introduction

This document is the charter for a proposed Networking Technical Working group to develop practical networking capabilities in three specific areas that are deemed to be important by both WLCG experiments and National Research and Education Networks (NRENs). The group is intended to bring together experts from the experiments, NRENs, software developers and network researchers to explore, prototype and document possible solutions for at least three specific areas identified at that January 2020 LHCONE/LHCOPN meeting.

Overview and Motivation

High Energy Physics (HEP) experiments, as well as many similar data intensive, global science domains, rely on networks as one of the critical components of their infrastructure both within the participating laboratories and sites as well as globally to interconnect those sites, data centres and experiments instrumentation. Recent work by the HEPiX Network Function Virtualization working group has created a Phase I report[1] about activities underway both within HEP and more broadly within industry. The report surveyed ongoing activities relevant to HEP from the datacenter to the WAN and suggested possible areas of interest that HEP may want to explore with the various NRENs and network researchers. The report was presented and discussed at the January 13-14th 2020 LHCONE/LHCOPN meeting at CERN and resulted in a recommendation to pursue (at least) three areas of network development considered useful to at least one or more WLCG and HEP experiments.

A primary concern for WLCG and its constituent collaborations is whether it will have the networking capabilities required to most effectively exploit LHC data for the lifetime of the LHC.

Technical Working Areas

Our high level goal is to identify and refine the relevant areas of work that will make practical differences for users of high performance global networks. For each work area we should clearly define the proposed scope, deliverables and success criteria. This charter provides the overview and suggests initial areas that we have achieved some level of consensus about but the overall working group may choose to add, change or remove activities. We initially have three areas proposed which we cover in the next three sections.

[1] The report is available on Google (living document) and at:
https://zenodo.org/record/3565563#.Xepa7ZNKhTZ

Network Packet Marking

Analyzing the HEP traffic flows in detail is critical for understanding how the various complex systems developed by the LHC experiments are actually using the network. The potential work here is to identify how we might label our traffic to indicate which experiment and task it is a part of. This is especially important for sites which support many experiments simultaneously where any worker node or storage system may quickly change between different users. With a standardized way of marking traffic, any NREN or end-site could quickly provide detailed visibility into HEP traffic to and from their site.

There are **two components** for the technical work. One would encompass how to mark traffic at the network level, defining a standard set of markings and providing the tools to the experiments to make it easy for them to participate. We note that the increasing use of VMs and containers might make marking traffic easier where those technologies are in use. The second part of the work would involve documentation and prototyping of how to consume the markings. We will need the ability to make available various views, summaries and dashboards showing the details of the marked traffic, provided by the NRENs, end-sites or others involved with the packet flows. We will need to understand required hardware and/or software that would support making the packet marking visible and useful.

Network Shaping

It remains a challenge for HEP storage endpoints to utilize the network efficiently and fully. An area of potential interest to the experiments is traffic shaping. **Without traffic shaping**, network packets are emitted by the network interface in bursts corresponding to the wire speed of the interface. This short train of back-to-back packets (a microburst) can overflow queues along the path, causing packet-loss. The impact on TCP, especially for high-bandwidth transfers on long network paths can be significant. **If instead, flows are shaped**, spacing out the packets on the wire to better match the end-to-end usable throughput, the result is much smoother flows able to operate at higher bandwidth. A significant extra benefit is that these smooth flows are much friendlier to other users of the network by not bursting and causing buffer overflows which can impact all users of that link.

The challenge for the working group is how to enable each source for traffic shaping. The group will need to explore tools and options that can easily be deployed, managed and monitored for the research domains adopting traffic shaping. Initial work should focus on verifying the impact of traffic shaping to better understand the benefit we can derive for large scale WAN data flows.

Network Orchestration

Within our data centers, technologies like OpenStack and Kubernetes are being leveraged to create very dynamic infrastructures to meet a range of needs. Critical for these technologies is a level of automation for the required networking using both software defined networking and network function virtualization. As we look toward high luminosity LHC, the LHC experiments are trying to find tools, technologies and improved workflows that may help bridge the anticipated gap between the resources we can afford and what will actually be required to extract new physics from massive data we expect to produce. **The ways inwhich we organize our computing and storage resources will need to evolve**. Architectures like Data Lakes, federated or distributed Kubernetes and multi-site resource orchestration will certainly benefit from (or require) some level of WAN network orchestration to be effective. To support this type of resource organization evolution, we need to begin to prototype and understand what services and interactions are required between the network and experiment frameworks. We would suggest a sequence of limited scope proof-of-principle activities in this area would be beneficial for all our stakeholders.

Timeline

The working group will need to finalize any timelines and associated deliverables. It was hoped that at least the packet marking activity could produce working examples by the end of calendar year 2020.

Working Group Operational Details

Membership

Proposal: Membership is open to those interested and participating as part of either global-scale research domains, network related research projects, Universities, national laboratories and regional or national research and education networks.

Membership is currently defined by the CERN e-group (see Tools below).

Tools

- Mailing list: CERN e-group mailing list has been created. Self-signup at http://cern.ch/simba3/SelfSubscription.aspx?groupName=net-wg
- Shared Documents: What technology should we use to share documents?

    - Google drive has our working group directory at https://drive.google.com/drive/folders/1_ELBx8jD3mLGRmmdhGMmIwYjactMwN7J?usp=sharing (Default is 'view' access, just request 'edit' access if you need it)
- Wiki/Web page: Initial Wiki page setup at https://twiki.cern.ch/twiki/bin/view/LCG/WLCGNetworking
    - Discourse: https://discourse.web.cern.ch/ We need someone to setup a skeleton area for our use-case so we can test this.
- Other tools:
    - CERN Mattermost (https://mattermost.web.cern.ch/) vs SLACK
    - Calendar
    - Code managed by Github or ? We need someone to setup a skeleton Github area for use, configuring for our use so we can evaluate it.

Reporting

We assume the members of the working group will be making regular reports at LHCONE/LHCOPN, DOMA and HEPiX meetings about our progress. Additional special presentations should be targeted for various other relevant meetings to be determined by the working group or the specific sub-groups.

Meetings

Bi-weekly, every 3 weeks, monthly, other? The meeting cadence will need to be determined by the initial working group. It should be noted that we expect sub-groups, focused on specific technical areas, will have their own meetings and schedule.

Deliverables

What will the group (or sub-groups) deliver? Code, Best Practices, Documentation, Libraries? This is another area the working group (and associated sub-groups) will need to address as they define the scope, deliverables, timeline and metrics for success.

APPENDIX

The report of the HL-LHC computing review in May 2020 was recently released and, for WLCG, it noted the following about the network:

"*Network: The availability of a strong and reliable network as a backbone for the data lake model to work is of great importance. Close collaboration with LHC partners and with NRENs is required to ensure that the evolution of the networks is sufficient on the international, trans-atlantic, and national level. The network has to be shared with other experiments and it is crucial to continue to work on having a common monitoring framework (marking the traffic by*

experiments or activities). We further encourage simulation of the new model to understand how well it will accommodate the experiment's production and analysis workflows."