

Network Requirements and Computing Model R&D for the HL-LHC Era

Tuesday, 11 August 2020 10:00 (8 minutes)

Summary: Recent estimates of network capacity requirements for the HL-LHC era indicate that these cannot be met through technology evolution and price/performance improvements alone within a constant budget. An in-depth consideration of HL-LHC Computing Model is thus needed, and an R&D program to formulate, design and prototyping of the new Model is recommended. This program could take advantage of current development projects that provide the capability to set up, allocate and end-to-end network paths with bandwidth guarantees, and to coordinate the use of network resources with computing and storage resources.

2020 Update in the Outlook for Network Requirements:

In January, at the 43rd LHCOPN/LHCONE meeting at CERN, the LHC experiments expressed the need for Terabit/sec links by the start of HL-LHC operations in 2027-28, preceded by the usual Computing and Storage (and Network) challenges starting during LHC Run3 (2021-4). This was reinforced by the requirements presented by the DOMA project which “foresees requiring 1 Tbps links by HL-LHC (ballpark) to support WLCG needs. This is for the network backbones and larger sites...”

The quoted network capacity requirements are an order of magnitude greater than what is available now through the present national and transoceanic networks based on 100GE links. As discussed at the LHCONE meeting, in the GNA-G Leadership group meeting that followed, and in the HEPiX Techwatch technology tracking group, these requirements cannot be accommodated solely through the exploitation of technology evolution within a constant budget. As a result, the further development of managed end-to-end services for the LHC and other science programs and the associated plans presented in this note, could be of pivotal importance. Work in this direction should also be guided by DOMA statements that “caching/latency hiding will be important. DOMA is exploring XCache as a mechanism, which provides latency hiding and support for diskless sites (with regional data lakes). Production of AODs (using RAW) will be a network driver, especially regionally. Effectively the ‘site’ is expanded to encompass a ‘region’.”

- It was agreed in subsequent discussions that the HEPiX Technology Watch WG and/or the Global Network Advancement (GNA-G) leadership group that was formed in the fall of 2019, can help define how much of it can be satisfied through technology evolution by 2027, and by 2024 in the preparatory phase.
- The rest will involve a change in paradigm including the end-to-end services involving sites and networks, and orchestration, as is being developed in projects such as SENSE, SANDIE and NOTED (described below). Ongoing discussions should continue to conceptualize and define what the new class(es) of service required entail.
- An important part of this is the persistent testbed being deployed by SENSE in collaboration with AutoGOLE and other collaborating projects. This is proceeding starting with the current SENSE testbed sites, plus extensions to UCSD, CERN, Starlight in Chicago, and a few other sites in the US and overseas.

Primary author: Prof. NEWMAN, Harvey (Caltech)

Presenter: Prof. NEWMAN, Harvey (Caltech)

Session Classification: Community Feedback