Reference ID: 11224248227_Wuerthwein

Submission Date and Time: 12/15/2019 8:27:15 PM

Author Names & Affiliations

Submitting author: Frank Wuerthwein - UCSD, SDSC, OSG

Additional authors: None

Contact Email Address (for NSF use only): (hidden)

Research domain(s), discipline(s)/sub-discipline(s)

All of Open Science

Title of Response

NETFLIX for all of Open Science

Abstract

NETFLIX operates a Content Delivery Network (CDN) that we all know how to use. To good approximation, it provides streaming data access to searchable curated data from anywhere on the internet, at anytime, to any subscriber. To be useful for all of open science, the CDN moreover needs to be federated, i.e. anybody's data can be shared via the CDN from a storage origin they locally own.

Requirements on origins need to be minimal in order to maximally stimulate sharing. Access is mediated via caches in the network and at processing endpoints to keep performance requirements on the origins minimal. Access performance is determined by location and performance of the closest cache. Groups of users need to be able to share data securely among each other. These groups ought to be locally managed and defined. There are thus locally defined access policies that are globally enforced. The CDN is domain science agnostic, supporting any file format with the appropriate client. Metadata, and thus curation is as federated as the data. The CDN is implemented via a plug-in architecture to allow experimentation with new caching algorithms, new network protocols (e.g. NDN), etc. on subsets of the production infrastructure.

Question 1 (maximum 400 words) – Data-Intensive Research Question(s) and Challenge(s). Describe current or emerging data-intensive/data-driven S&E research challenge(s), providing context in terms of recent research activities and standing questions in the field. NSF is particularly interested in cross-disciplinary challenges that will drive requirements for cross-disciplinary and disciplinary-agnostic data-related CI.

Science is a team sport, and midscale as well as large scale research endeavors commonly require international teams. A large scale example is the LIGO Scientific Collaboration transitioning to IGWN, the international gravitational wave observatory network, via the collaboration of LIGO, Virgo, and KAGRA. Other big science examples include DUNE and the LHC experiments ATLAS and CMS. Midscale examples include XENON, MINERVA, VERITAS, GlueX, Nova, and many others. All of these benefit from global resource sharing, requiring global data access, sometimes, as in the example of IGWN, including both private and public data. As a second example, imagine multiple hospitals wanting to share their de-identified clinical data from a variety of sources. The quest for advances in health care make it desirable for biomedical informaticists to have access to the same data from as many hospitals as possible. At the same time, some of the hospitals may want to restrict sharing to only partner hospitals that share equally, in part because of concerns of the intellectual property potentially derivable from that data. There is thus a need for potentially fine grained, and rather complex access policy definitions that are locally defined by the owners of the data, and need to be globally enforced as part of the CDN. A third example is that of a single PI and her group producing data at a national facility (experimental or computational). The PI may want to share this data with her collaborators from multiple institutions, possibly as part of mutual data sharing agreements. The different collaborators may have computing resource access at a variety of different centers, and via different modalities ranging from owned, to rented, to allocated, to openly shared. The entire collaboration of PIs will want to access the data at all the computing resources they have access to. They thus want to federate compute and data resources. In addition, a PI may have allocations at a variety of national facilities funded by different federal agencies, subject to different data retention policies. Data thus needs to be brought home to their institutions for longer term archiving and use. Despite the movement of data during the data lifecycle, its access ought to be unchanged for the collaborators that enjoy access. Data origins should thus be independent of the global namespace the data lives in such that the former can change without affecting the latter.

Question 2 (maximum 600 words) – Data-Oriented CI Needed to Address the Research Question(s) and Challenge(s). Considering the end-to-end scientific data-to-discovery (workflow) challenges, describe any limitations or absence of existing data-related CI capabilities and services, and/or specific technical and capacity advancements needed in data-related and other CI (e.g., advanced computing, data services, software infrastructure, applications, networking, cybersecurity) that must be addressed to accomplish the research question(s) and challenge(s) identified in Question 1. If possible, please also consider the required end-to-end structural, functional and performance characteristics for such CI services and capabilities. For instance, how can they respond to high levels of data heterogeneity, data integration and interoperability? To what degree can/should they be cross-disciplinary and domain-agnostic? What is required to promote ease of data discovery, publishing and access and delivery?

Obviously, NETFLIX, while evocative, is a poor model for what is needed in a CDN for Open Science. A useful system will benefit from ideas previously implemented by NETFLIX, Napster, Cloudflare, and others. In addition to the already mentioned challenges of federation (anybody needs to be able to provide content and storage) and access policies (data owners need to be able to define who has access to the data they share), there are resource provisioning issues to address, both short and long term. Short term, resource provisioning in a CDN is about what data gets cached where in order to provide the desired quality of access while avoiding thrashing of caches. At what point is it more effective to move the input data as part of the scheduled workflow, and how would a workflow system benefit from a CDN either to hide access latency or reduce network traffic and/or performance requirements from the data origin. Long term it is about sustainability and growth of the infrastructure. How does the infrastructure grow the storage in the network and at endpoints to provide the desired performance? How does the CDN grow the variety and volume of the data content shared across it? How are responsibilities for hardware maintenance, service operations, and science support shared across multiple teams that are generally employed by different organizations in different countries? Given differences in scale it seems highly likely that the CDN may be operated as independent data federations that use largely the same software. E.g. the Exascale data volumes at the LHC experiments ATLAS and CMS almost certainly deserve their own individual CDNs, while smaller endeavors are probably better served sharing a CDN infrastructure. In addition, cache owners are likely to want to dynamically decide on policies of use for their caches, e.g. to protect quality of service for the primary domain they serve. Finally, such an infrastructure ought to be open and modular at multiple layers of its architecture. E.g. caching policies, network transfer protocols (including support/exploration of NDN), storage resource provisioning services might all be interchangeable via an appropriate plugin architecture in the core CDN, and in fact, different approaches may coexist within the same CDN. Additional namespace and metadata functionality, including but not limited to data curation, publication, and search could be built on top of the core CDN by projects independent of that core. In a sense, we are envisioning an open infrastructure software and services ecosystem that is built around basic notions of streaming compute access to any data anytime from anywhere with performance characteristics, or “quality of service” that is reasonably predictable.

Question 3 (maximum 300 words) – Other considerations. Please discuss any other relevant aspects, such as organization, processes, learning and workforce development, access and sustainability, that need to be addressed; or any other issues more generally that NSF should consider.

Within the OSG Consortium, we operate a production prototype CDN that implements some of the core functionalities mentioned above. We presently have 15 caches in production worldwide, including ones in Amsterdam (Netherlands), Cardiff (UK), Bologna (Italy), and Seoul (Korea). We are working within global cybersecurity bodies on policies that will facilitate operations at network POPs around the world by separating cleanly the responsibilities of storage provider from CDN service provider. Internet2 and ESNet as collaborators have started adding storage into their network, making it available to us to operate as part of our CDN. This CDN is presently used by several multi-institutional science collaborations, including IGWN, DUNE, Minerva, DES, and Nova, among others. In addition, ATLAS and CMS use much of the same infrastructure software for their own emerging CDNs. There is thus significant critical mass around an existing technology that is open, and modular, and can provide a proof of principle and starting point for experimentation and future evolution, while operating an at scale production system. We firmly are convinced that operating the prototype as a production system is essential to achieve global deployment and widespread adoption. A CDN like the one we envision will emerge gradually via evolution of a production worthy system, rather than as a funded project designed from the ground up. Sustainability and adoption is achieved via processes that are not tied to technology, and even radical evolution of CI software is possible when planned carefully over appropriate timescales. The vision outlined here requires multiple CI teams across multiple awards in parallel and series to contribute their respective expertise. We believe that the OSG Consortium as a coalition of the willing, governed by an unfunded council, provides an excellent ecosystem for the necessary collaborative evolution of federated infrastructure like the aforementioned CDN.

-- End Submission –