IRIS-HEP Project Proposal

Optimization of the Rucio-SENSE Data Movement Manager (DMM)

Fellowship Applicant: Camille Sicat

Mentors: Aashay Arora and Diego Davila

Project Duration: 12 weeks

Proposed Start Date: 24 June 2024

Project Summary

Problem Statement

We are approaching the exa-scale computing era for most large collaborative experiments such as the High Luminosity LHC [4]. With such large volumes of data being produced, we will require highly predictable and accountable network usage in order to meet the throughput demand. Software defined networking (SDN) has emerged as a strong contender for allowing end-to-end accountability in bandwidth utilization and allowing different experiments to manage the priorities of their data flows [3]. It provides the mechanisms to enable multi-domain orchestration for a variety of network and other cyberinfrastructure resources in a highly customized manner [3]. Therefore, to meet the demands posed by future experiments, we need to explore the integration of SDN with the existing data transfer software stack [2]. Using Rucio/FTS + XRootD as the control ground setup, we are trying to integrate it with SENSE, the SDN service developed by ESNet [4].

Solution

The Data Movement Manager (DMM) is the proposed interface that works between Rucio and SENSE [1]. From Rucio, DMM obtains transfer metadata such as the priority and the size of the transfer and uses it to make decisions about bandwidth allocation [2]. DMM then uses SENSE to construct peer-to-peer VLANs for each set of endpoints, adjusting data flows with relevant Quality of Service (QoS) to ensure each transfer gets the bandwidth promised to them according to their priority. In addition, it monitors the status and performance of all data flows. With this information, DMM can generate reports of underperforming flows, allowing it to identify the point of failure in a data flow chain.

Goals

For this project, the main goals are:

- Implementation of DMM's monitoring system to correlate FTS monit data with host level information to generate performance reports for transfers
- Optimization of DMM's interactions with Rucio
- Optimization of DMM's bandwidth decision algorithm

Proposed Timeline

Week	Tasks
1-2	• Project overview presentation • In-depth familiarization of Rucio, SENSE through documentation • Familiarization with DMM through documentation and example
3-4	• Update Rucio messaging interactions • Implement DMM monitoring • Familiarization with DMM bandwidth allocation algorithm through documentation and simulation
5-7	• Mid quarter project report on current progress, challenges faced, future goals • Optimization of DMM bandwidth allocation algorithm • Identify and address potential issues
8-10	• Optimization of DMM interaction with front-end • Modularization of DMM for further scalability, future simulation • Begin documentation
11-12	• Finalize documentation • Propose further tests/optimization • Final presentation on challenges faced, lessons learned

References

[1] F. Würthwein, J. Guiang, A. Arora, D. Davila, J. Graham, D. Mishin, T. Hutton, I. Sfiligoi, H. Newman, J. Balcas, T. Lehman, X. Yang, & C. Guok. (2022). Managed Network Services for Exascale Data Movement Across Large Global Scientific Collaborations. In 2022 4th Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP). IEEE.

[2] T. Lehman, X. Yang, C. Guok, F. Wuerthwein, I. Sfiligoi, J. Graham, A. Arora, D. Mishin, D. Davila, J. Guiang, T. Hutton, H. Newman, and J. Balcas, "Data transfer and network services management for domain science workflows," 2022. [Online]. Available: https: //arxiv.org/abs/ 2203.08280

[3] J. Zurawski, D. Brown, B. Carder, E. Colby, E. Dart, K. Miller et al., "2020 high energy physics network requirements review final report," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-2001398, Jun 2021. [Online]. Available: https://escholarship.org/uc/item/78j3c9v4

[4] I. Monga, C. Guok, J. MacAuley, A. Sim, H. Newman, J. Balcas, P. DeMar, L. Winkler, T. Lehman, and X. Yang, "Software- defined network for end-to-end networked science at the exascale," Future Generation Computer Systems, vol. 110, pp. 181–201, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X19305618