SC19 Network Research Exhibition NRE-022

Toward Unified Resource Discovery and Programming in Multi-Domain Networks

Submitted on behalf of the team by: Harvey Newman, Caltech, newman@hep.caltech.edu, Qiao Xiang and Jensen Zhang, Yale University, {qiao.xiang, jingxuang.zhang}@yale.edu

Abstract

The Yale, IBM, ESNet and Caltech team will demonstrate a novel, unified multi-domain resource discovery and programming system for data-intensive collaborative sciences. Specifically, this system provides three key components: (1) a fine-grained, accurate, highly-efficient multi-domain multi-resource discovery framework (a substantial extension of the team's SC'18 Mercator paper), (2) a strong machine learning component to provide accurate performance prediction for dynamic, reactive science workflows, and (3) a high-level resource programming and composition framework. This demonstration will include: (1) efficient discovery of multiple available resources in a multi-domain wide-area collaborative science network connecting Los Angeles and Denver, (2) real-time, accurate performance prediction for dynamic, reactive science workflows in this wide-area network, and (3) high-level resource programming and composition in science network with automatic resource orchestration update.

Goals

1. Demonstrate how the proposed framework can discover fine-grained, global multi-resource information across networks while preserving the privacy of different networks [1-3];
2. Demonstrate how machine learning techniques can be utilized to provide accurate performance prediction for dynamic, reactive science workflows [4];
3. Demonstrate how a high-level resource programming language simplifies the resource orchestration in science networks [5].

Resources

This demo is composed of three domains. In particular, we will use 2-4 data transfer nodes (DTNs) and 2 switches in the Caltech booth at SC19 exhibit floor to form one network. This network will be connected to the Caltech SDN testbed located at Pasadena, California via a 100 Gbps WAN circuit, provided by SCinet, CenturyLink and CENIC Los Angeles. In the SDN testbed, several switches and DTNs will be used to form two other domains.

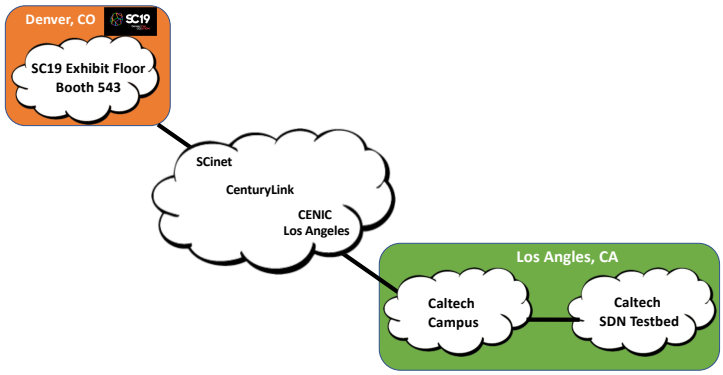

Involved Parties

- Qiao Xiang, Yale University, qiao.xiang@cs.yale.edu
- Jensen Zhang, Yale University, jingxuan.zhang@yale.edu
- Harvey Newman, California Institute of Technology, newman@hep.caltech.edu
- Y. Richard Yang, Yale University, yry@cs.yale.edu
- Franck Le, IBM T. J. Watson Research Center, fle@us.ibm.com
- Chin Guok, Lawrence Berkeley National Laboratory, chin@es.net
- John MacAuley, Lawrence Berkeley National Laboratory, macauley@es.net

Publications

[1] Xiang, Qiao, Jingxuan Jensen Zhang, Xin Tony Wang, Yang Jace Liu, Chin Guok, Franck Le, John MacAuley, Harvey Newman, and Y. Richard Yang. "Toward Fine-Grained, Privacy-Preserving, Efficient Multi-Domain Network Resource Discovery." *IEEE Journal on Selected Areas in Communications* 37, no. 8 (2019): 1924-1940.

[2] Xiang, Qiao, X. Tony Wang, J. Jensen Zhang, Harvey Newman, Y. Richard Yang, and Y. Jace Liu. "Unicorn: Unified resource orchestration for multi-domain, geo-distributed data analytics." *Future Generation Computer Systems* 93 (2019): 188-197.

[3] Xiang, Qiao, J. Jensen Zhang, X. Tony Wang, Y. Jace Liu, Chin Guok, Franck Le, John MacAuley, Harvey Newman, and Y. Richard Yang. "Fine-grained, multi-domain network resource abstraction as a fundamental primitive to enable high-performance, collaborative data sciences." In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 58-70. IEEE, 2018.

[4] Gao, Kai, Jingxuan Zhang, Y. Richard Yang, and Jun Bi. "Prophet: Fast Accurate Model-Based Throughput Prediction for Reactive Flow in DC Networks." In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pp. 720-728. IEEE, 2018.

[5] Gao, Kai, Taishi Nojima, and Y. Richard Yang. "T rident: toward a unified SDN programming framework with automatic updates." In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pp. 386-401. ACM, 2018.