# Embracing Programmable Data Planes for an Elastic Data Transfer Infrastructure


Joaquin Chung, Zhengchun Liu, Raj Kettimuthu, Ian T. Foster

Argonne National Laboratory
jchung@mcs.anl.gov,
zhengchun.liu@anl.gov,
kettimut@mcs.anl.gov, foster@mcs.anl.gov

Se-Young Yu, Jim Chen, Joe Mambretti

Northwestern University / StarLight
young.yu@northwestern.edu,
jim-chen@northwestern.edu,
j-mambretti@northwestern.edu

Eun Sung Jung

Hongik University
ejung@hongik.ac.kr


Data plane programmability has emerged as a response to the lack of flexibility in networking ASICs and the long product cycles that vendors take to introduce new protocols on their networking gear. Furthermore, it tries to bridge the gap between the SDN promise and OpenFlow implementations. Following the ASIC tradition, OpenFlow implementations have focused on defining matching protocol header fields in forwarding tables, and these matching protocols cannot be modified once the switch is manufactured. On the other hand, programmable data planes seek to allow network operators and programmers to define exactly how packets are processed in a reconfigurable switch chip (or in a virtual software switch). Such levels of programmability open up opportunities for offloading certain processing on the data to the network and obtaining more accurate state of the network.

An elastic data transfer infrastructure (DTI) [1] is an architecture that expands and shrinks data transfer node resources based on the demand (see Figure 1). A key element of the elastic DTI architecture is the statistics collector that feeds usage and performance information to a decision engine. Another key element of the elastic DTI architecture is the load balancer that distributes the load of incoming transfers among existing virtualized resources. State-of-the-art solutions rely on traditional network monitoring systems such as SNMP and sFlow to collect network state information. However, traditional network monitoring methods either use polling mechanism to query network devices, or use sampling when devices are allowed to push data in order to lower the communication overhead and save database storage space.

In-band Network Telemetry (INT) is a framework that allows the data plane to add telemetry metadata to each packet of a flow. Then the metadata is removed and sent to a collector/analyzer before the packet is forwarded to the final destination. In this demo we will present the impact of advanced network telemetry using programmable switches and the P4 (Programming Protocol-independent Packet Processors) language on the granularity of network monitoring measurements (see Figure 2). We will compare the detection gap between a programmable data plane approach and traditional methods such as sFlow.

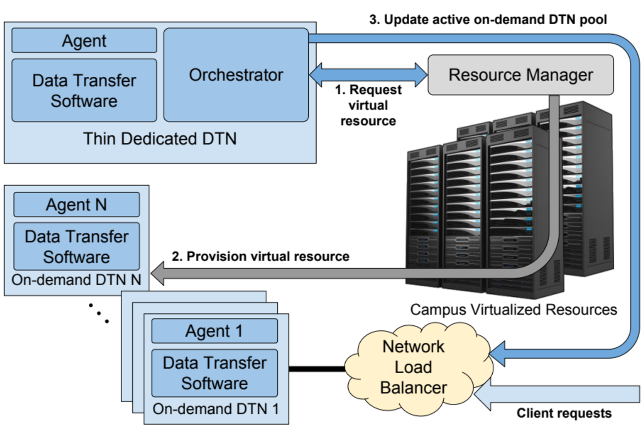


Figure 1. Elastic Data Transfer Infrastructure (DTI)

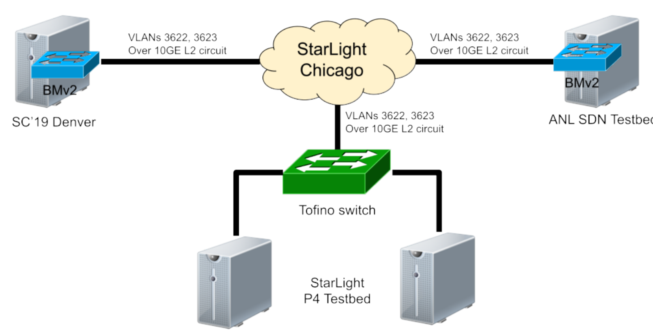


Figure 2. Demonstration setup


[1] Joaquin Chung, Zhengchun Liu, Rajkumar Kettimuthu, and Ian Foster. Toward an Elastic Data Transfer Infrastructure. In 15th International Conference on eScience (eScience) (eScience 2019). IEEE, 262–265. https://doi.org/10.1109/eScience.2019.00036