SC22 Network Research Exhibition: Demonstration Preliminary Abstract

Packet Marking for Networked Scientific Workflows

Marian Babik, Shawn McKee, Dale W. Carder, Tristan Sullivan, Fatema Bannat Wala, Eli Dart, Mariam Kiran, Edoardo Martelli, Pavlo Svirin, Tim Chown, Marcos Schwarz, Florent Parent, Thomas Tam, Bruno Hoeft, Alexandr Zaytsev, Karl Newell, Joe Mambretti, Jim Chen, Fei Yeh, Eric Brown, Mario Lassnig, Michael Lambert, Tony Cass, Carmen Misa, Dale Carder, James Deaton, Andrew Hanushevsky

Abstract

Managing large scale scientific workflows over networks is becoming increasingly complex, especially as multiple science projects share the same foundation resources simultaneously yet are governed by multiple divergent variables: requirements, constraints, configurations, technologies etc. A key method to address this issue is employ techniques that provide high fidelity visibility into exactly how science flows utilize network resources end-to-end. This demonstration will showcase one such method, Scientific network tags (scitags), an initiative that is promoting identification of the science domains and their high-level activities at the network level. This open system initiative provides open source technologies to help R&E networks understand resource utilization while providing information to scientific communities on the behavior of their workflows network flows.

Overview

Given the increasing complexity of scientific workflows over shared networks and the number of parameters they are subjected to, enhanced methods of visibility into flows is required. Such visibility enables efficient workflows via optimization of resource utilization, and avoidance of impairments such as channel saturation and/or congestion, contention, latency, and packet loss. To enhance this visibility, this scitags initiative is undertaking the development of multiple building blocks to achieve both enhanced visibility into network flows and enhanced efficiency based on the information provided by that enhanced visibility. Objectives include a) developing an overall architecture with a defined standard set of markings, b) providing methods for marking network traffic that would be easy to implement, e.g., reliance on VMs and containers, c) creating techniques for reading/monitoring/validating/analyzing those markings, d) determining appropriate responses/actions based on those analysis, and e) communication channels that provide information of these methods to wider communities, including recommendations on specific hardware, software, protocols, and configurations.

The scitags initiative will organize several SC22 MRE demonstrations.

One demonstration, led by Tristan Sullivan, will showcase Packet Marking with the Extended Berkeley Packet Filter (eBPF). eBPF is a feature in the Linux kernel that allows code to be injected into the kernel without the need to write a kernel module. This code is executed whenever a specified kernel hook occurs. The capability of eBPF to mark packets at high rate

will be demonstrated, with an aim to showing that the impact on transfer performance is negligible

A second demonstration, led by Andrew Hanushevsky and Marian Babik, will showcase packet marking using Linux advanced socket interface with XRootD. It will demonstrate how the native socket facilities of the recent Linux kernel can be used to set/update IPv6 header flow label field and how it works at high throughput disk-to-disk transfers. The aim will be to demonstrate that it has negligible impact on the transfer performance.

Another demonstration, led by Edoardo Martelli and Carmen Misa Moreira, will showcase the accounting of flow label tagged packets using a P4 programmable switch and the FreeRTR routing stack developed by the RARE GEANT project. The switch will count the packets tagged with specific flow label values set by the other demonstration of this exhibition. Counters will be exported using the telemetry capabilities of FreeRTR.

Goals

The SC22 goals of this NRE experiment/demonstration will be to showcase the capabilities of the scitags architecture and methods for optimizing data intensive science.

Resources

This NRE demonstration will be conducted using resources of the collaborating domains, including: global high-bandwidth paths among DTNs at or connected to international Open Exchange Points and the SC22 venue, especially the paths developed for that venue by the SCinet WAN group. Participating Open Exchange Points include CERNLight, the Montreal Open Exchange (MOXY), the StarLight International/National Communications Exchange Facility, and NetherLight,

Involved Parties

The scitag initiative is an open international collaboration that originated as a WLCG Networking Technical Working group developing practical networking capabilities for HEP research using National Research and Education Networks (NRENs), including the LHCOPN and LHCONE. Participating in these demonstrations will be multiple individuals from many organizations.

Marian Babik (CERN) marian.babik@cern.ch
Shawn McKee (Michigan), smckee@umich.edu
Dale W. Carder (LBNL/ESnet), dwcarder@es.net
Fatema Bannat Wala (LBNL/ESNet),
Eli Dart (LBNL/ESnet), dart@es.net
Mariam Kiran (ESnet) mkiran@es.net
Edoardo Martelli (CERN), edoardo@cern.ch
Pavlo Svirin (CERN/UTA),
Tim Chown (Jisc), tim.chown@jisc.ac.uk
Tristan Sullivan (University of Victoria) tssulliv@uvic.ca
Florent Parent (DTI, University of Lavel)) florent.parent@dti.ulaval.ca

Thomas Tam (CANATRIE) Thomas.tam@canarie.ca
Marcos Schwarz (RNP), marcos.schwarz@rnp.br
Alexandr Zaytsev (BNL),
Joe Mambretti (StarLight, iCAIR NU, MREN) j-mambretti@northwestern.edu
Jim Chen (StarLight, iCAIR NU, MREN) jim-chen@northwestern.edu
Fei Yeh (StarLight, iCAIR NU, MREN) fyeh@northwestern.edu
Tony Cass (CERN) tony.cass@cern.ch
Eric Brown (Virginia Tech),
Mario Lassnig (CERN), mario.lassnig@cern.ch
Michael Lambert (PSC), lambert@psc.edu
Andrew Hanushevsky (SLAC), abh@slac.stanford.edu
Carmen Misa (CERN) carmen.misa@cern.ch
Garhan Attebury (UNL) garhan.attebury@unl.edu
Bruno Hoeft (SCC)DE-KIT, bruno.hoeft@kit.edu
James Deaton (Great Plains) james@greatplains.net