

SC24 Network Research Exhibition: Demonstration Abstract

Multi Domain experiments using ESnet SENSE on the National Research Platform / PacWave / FABRIC

GNA-G AutoGOLE/SENSE WG

Abstract

The San Diego Supercomputer Center (SDSC) and the Prototype National Research Platform (PNRP) NSF Award 2112167 Category II are providing twelve Xilinx Alveo U55C FPGAs in six servers for researchers at the SC24 Network Research Exhibition (NRE). These devices will be utilized in ESnet SmartNIC P4 Segment Routing over IPv6 (SRv6) using Unstructured Segment Identifier (uSID) experiments, connecting from the FABRIC testbed StarLight site (FABRIC-STAR) to various sites. Each U55C FPGA features dual 100 Gigabit Ethernet interfaces, integrated into the SDSC Automated Global Optical Lightpath Exchange (AutoGOLE) Science Environment for Network and Computer Engineering (SENSE) network topology.

The PacWave NSF Award 2029306 (IRNC Core Improvement: Accelerating Scientific Discovery & Increasing Access - Enhancing & Extending the Pacific Wave Exchange Fabric) has supplied seven 1U DC-powered measurement and monitoring servers, known as Interactive Global Research Observatory Knowledgebase (IGROK) nodes. Each node is equipped with a Bluefield 2 Data Processing Unit (DPU), also with dual 100 Gigabit Ethernet interfaces, integrated into the PacWave AutoGOLE SENSE network topology.

Goals

- Highlight the exceptional value of the 400G Arista switch hosted at PacWave Los Angeles.
- Demonstrate ESnet SmartNIC SRv6 uSID capabilities, showcasing the use of FABRIC sites at StarLight and the University of California, San Diego (UCSD), possibly using Xilinx SN1000 SmartNIC in the mobile FABRIC-CINE rack within SCinet during SC24.
- Utilize Bluefield 2 DPUs at seven PacWave peering sites to participate in GP4L using freeRouter (freeRtr) and additional Vector Packet Processing (VPP) experiments running on the second 100G interface:
 - SDSC, PacWave-LA, PacWave-Sunnyvale, PacWave-Seattle, StarLight, NYSERNet, UoG (Guam)
- Demonstrate Bluefield 2 Smart Network Access Point (SNAP) NVMe over Fabrics (NVMe/OF) offloading regional storage devices.
- Use ESnet High Touch to remotely analyze a mirrored 400G port channel on the SDSC AS9516 400G SONiC Edgecore P4 switch.

- Local FPGA Packet Capture (pcap) streaming of the mirrored 400G port channel on the SDSC AS9516 400G SONiC Edgecore P4 switch.
- Show the sFlow-based measuring and monitoring stack available on the NRP Nautilus platform.
- Illustrate how SENSE and FABRIC can provision Layer 2 (L2) paths in the datacenter and over global Research & Education (R&E) networks using campus-facing FABRIC 400G facility ports.

Resources

- P4 Implementation of uSID SRv6 compiled and loaded onto 12 Xilinx Alveo U55Cs in 6 servers at SDSC using the ESnet SmartNIC Framework.
- Seven Bluefield 2 DPUs, distributed at PacWave peering sites, each with dual 100 Gigabit Ethernet interfaces.
- Xilinx Alveo U280 FPGAs deployed at FABRIC testbed sites..
- Xilinx SN1000 SmartNIC installed in the FABRIC-CIEN mobile site rack at SC24 SCinet.

Involved Parties

- John Graham, jjgraham@ucsd.edu
- Mohammad Firas Sada, mfsada@ucsd.edu
- Frank Würthwein, fkw@ucsd.edu
- Tom DeFanti, UCSD, tdefanti@ucsd.edu
- Dmitry Mishin, UCSD, dmishin@ucsd.edu
- Tom Hutton, SDSC, hutton@sdsc.edu
- Joe Mambretti, iCAIR, j-mambretti@northwestern.edu
- Jim Chen, iCAIR, jim-chen@northwestern.edu
- Fei Yeh, iCAIR, fyeh@northwestern.edu
- Tom Lehman, tlehman@es.net
- Xi Yang, xiyang@es.net
- Justas Balcas, jbaldas@es.net
- Marcos Felipe Schwarz, marcos.schwarz@mp.br
- Christopher Bruton, cbruton@cenic.org
- Neil McKee, neil.mckee@inmon.com

CENIC - IRNC Core Improvement:

https://www.nsf.gov/awardsearch/showAward?AWD_ID=2029306

PNRP - Category II:

https://www.nsf.gov/awardsearch/showAward?AWD_ID=2112167