



## **The Wire, Node, and Processor: Messaging Bottlenecks are Everywhere**

Ian Karlin

Lawrence Livermore National Laboratory

Karlin1@llnl.gov

**Abstract** – please provide a short abstract of your presentation (e.g., a paragraph or two).

With the deployment of the Sierra supercomputer new communication bottlenecks are showing up in HPC applications at Lawrence Livermore National Lab. Sierra only has 25% more injection bandwidth per node as the machine its replacing Sequoia, and only 2x the bandwidth of our commodity clusters despite having 10x or more the per node compute capability. In this talk, we present early experiences of applications running on Sierra and what is limiting their messaging. We find that some applications are limited by interconnect bandwidth and impacted by our 2:1 tapered network. Many other applications remain latency limited. In addition, many applications are being limited by message packing and other on-node challenges. Throughout the talk we will discuss our findings and how and what network improvements could lead to better application performance.

**Bio** – please provide your bio.

Ian Karlin is a computer scientist at Lawrence Livermore National Laboratory. His work focuses on application and computer system performance with a current emphasis on how best to balance systems for overall application throughput. He is a deputy for the CSSE program at LLNL focusing on applications and benchmarking. He also is the co-lead for the LLNL CORAL-2 benchmark team and the LLNL Technical representative for the Intel PathForward contract. Past achievements include best paper awards at Cluster 2015, IPDPS 2013 and the LLNL Deputy Director of Science and Technology Excellence in Publication Award for the best Math/Computer Science paper in 2015 and 2018. He was the deputy leader of the Institutional Center of Excellence project, which ran from 2015-2018 and aimed at porting LLNL applications to the unclassified version of the Sierra system