

LA-UR-22-23012

Approved for public release; distribution is unlimited.

Title: MPI Sessions – Investigate ability of UCX to support extended CIDs

Author(s): Pritchard, Howard Porter Jr.

Intended for: Report

Issued: 2022-03-31



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA00001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

MPI Sessions –

Investigate ability of UCX to support extended CIDs

ECP WBS	2.3.1.17 / OMPI-X
Epic/Story	STPR17-86
PI	David Bernholdt, ORNL
Milestone Lead	Howard Pritchard, LANL
Members	ORNL, LANL, LLNL, SNL, UTK

Scope and objectives

- ECP OMPI-X develops a production grade MPI library
- Investigate ability of Open MPI's UCX point to point messaging layer (PML) to support Sessions (extended CIDs)

Better support for multi-component applications

- Different components of an application can specify different levels of thread safety
- Specify different operational characteristics of MPI (e.g. spin polling vs blocking) on a per component basis
- Backward compatibility with apps using MPI_Init

Impact

- **New Way to Use MPI:** better isolation of application components use of MPI (thread support, polling models, more scalable resource requirements specificity etc.)
- **New Application spaces:** only components of application that need to use MPI need do so, potentially expanding range of applications which can use MPI
- Better support for more dynamic process models

Project accomplishment

- Developed a PMIx-based approach to handling extended CIDs, greatly simplifying the modifications required to make a PML usable with MPI Sessions
- Working with OpenPMIx developers to optimize the PMIx-based approach

Deliverables Memo describing results of Sessions/UCX investigation – [STPR17-86](#)

Modified UCX PML prototype on Github: https://github.com/hpritcha/ompi/tree/topic/ucx_excid_support

¹ LA-UR-22-XXYYYY

