LA-UR-21-31331
Approved for public release; distribution is unlimited.

Title:	MPI Sessions WG Update
Author(s):	Pritchard, Howard Porter Jr.
Intended for:	The International Conference for High Performance Computing, Networking, Storage, and Analysis, 2021-11-14 (St. Louis, Michigan, United States)
Issued:	2021-11-15

MPI Sessions WG Update

Howard Pritchard
Los Alamos National Laboratory

LA-UR-21-XXXXX

MPI Sessions – current state

- In the MPI 4.0 standard
- Consider this as first step for Sessions

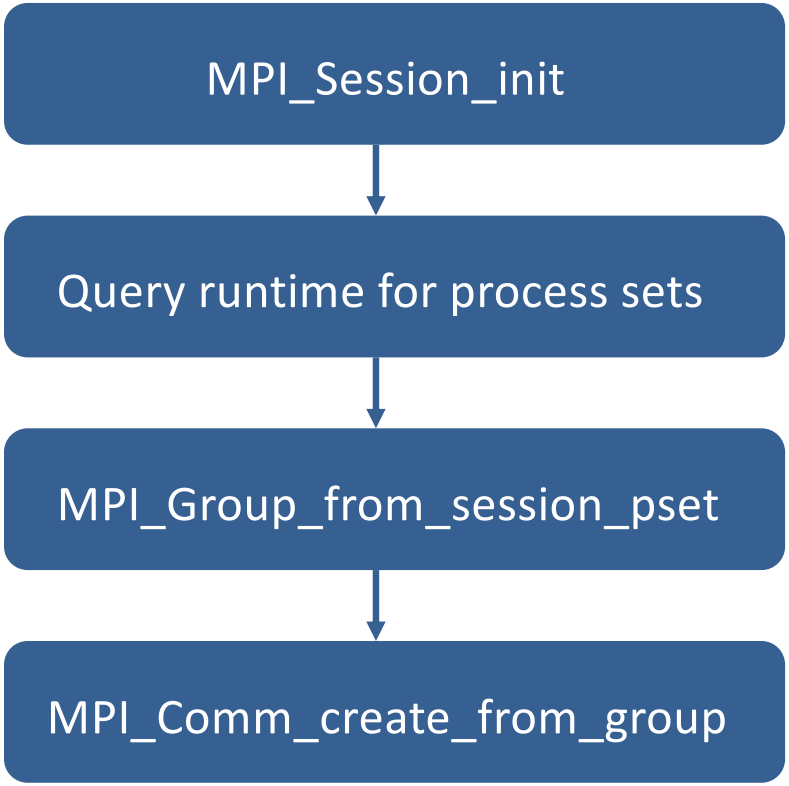

MPI Sessions – current state in MPI implementations

- Available in MPICH 4.0 release stream
- Prototype based on Open MPI is available at
  https://github.com/hpc/ompi/tree/sessions_pr
  (this branch is subject to rebasing!)
- Set of simple tests are available at
  https://github.com/open-mpi/ompi-tests-public

MPI WG Sessions – current activities

- For MPI 4.1 standard - clarifications of Sessions related text

- For MPI 5.0 - investigating requirements for more dynamic use of Sessions:
  - Presentation of available process sets in a manner more suitable for dynamic environments
  - Mechanisms for runtime notifying application of resource changes
  - Mechanisms for application to notify runtime about changing resource requirements
  - Adding/removing MPI processes (beyond MPI_Comm_spawn)
  - Working with FT WG to develop FT approaches that leverage Sessions functionality