



# Open MPI OpenFabrics Update

Jeff Squyres



# Current Membership

- 14 members, 9 contributors, 1 partner
  - 4 US DOE labs
  - 8 universities
  - 10 vendors
  - 1 individual



# Current Status

- Stable release series: 1.2
  - Current community release: v1.2.4
  - OFED includes: v1.2.4
    - Includes ConnectX configuration fix
  - Currently has a libsysfs dependency issue (!)
  - Possibly release v1.2.5 with libsysfs fix...?

# Upcoming v1.3 Series

- BSRQ (presented at Sonoma)
  - Multiple receive queues with different size buffers (I.e., send to Q with closest size)
  - More efficient use of registered memory
- ConnectX XRC support
  - Pasha/Mellanox just submitted patch
  - Being reviewed by Cisco, Voltaire, ORNL
- LSF support
  - In addition to Torque, SLURM, ...etc.

# Upcoming v1.3 Series

- [Partial] Overhaul of progress engine
  - Decrease work per polling cycle
  - Allow possibility of blocking
- RDMA CM
  - iWarp support
  - Take advantage of VLs
- Exploring multi-core / NUMA / NUNA possibilities
  - Processor / memory / HCA / NIC affinity

# Upcoming v1.3 Series

- PERUSE interface
  - Exporting internal MPI statistics to tools
  - Fragment information / counters
  - Internal timing information
  - ...etc.
- Tighter integration with parallel tools
  - DDT parallel debugger “understands” opaque MPI handles
  - VampirTrace integration (tracefile / post-mortem analysis)

# Upcoming v1.3 Series

- MPI\_THREAD\_MULTIPLE support
  - Unlikely
  - Not enough demand to prioritize resources
- Dropped VAPI support
- Estimate CY 1H08
  - Difficult to estimate for all the organizations



**Come Join Us!**

<http://www.open-mpi.org/>