

Cray MPICH2 Past, Present and Future

**Larry Kaplan
Senior Principal Engineer
Cray Inc.**

Cray MPICH2 Past and Present

- Cray has worked closely with ANL MPICH2 developers since 2002
 - release of MPICH2 on Cray X2, Cray XT and now Cray XE hardware
- Used the ADI device interface on the Cray X2 and Cray XT
 - Scaling MPICH2 to over 220K cores on Jaguar at ORNL
- Worked closely with ANL MPICH2 developers in 2009 and 2010 to enable Nemesis on the new Cray XE systems
 - Have successfully scaled to over 150K cores on Hopper at NERSC
- We have our own standalone PMI library that interfaces with a Cray specific job launcher (aprun)
 - It has both PMI-1 and PMI-2 interfaces
- Our XE version of MPICH2 can recover from transient network faults
- MPT 5.3.4 (October 2011) is based off of MPICH2 1.3.1
 - Supports MPI-2.2 (with the exception of dynamic process management)

- Cray MPT 5.4 (November 2011) also will be based off of MPICH2 1.3.1 and has had some successful preliminary testing on using over 300K MPI ranks.

- Cray MPT 5.5 (April 2012) will incorporate a pre-release version of MPICH2 1.5 and MPT 5.6 (September 2012) will use the released MPICH2 1.5.

- Cray plans to continue to work closely with ANL MPICH2 developers
 - Work on enhancing our XE6 MPICH2 library
 - Work on support for our future interconnects

- Cray is an active participant in the MPI-3 Forum
 - We are very interested in making available on Cray platforms the initial MPICH2 implementations of the MPI-3 additions. in particular:
 - Non-blocking collectives
 - Fault tolerance
 - RMA