

Open MPI:

What we've done and where we're going

Andrew Lumsdaine
Open Systems Lab
Pervasive Technology Labs
Indiana University




1



What is Open MPI?

- ❑ Open source implementation of MPI-2
- ❑ Combined expertise from 4+ previous MPIs
- ❑ High performance & robust
- ❑ Works with most interconnects
- ❑ Modular Component Architecture
 - Combinatorial capabilities
 - Function pointers faster than shared library calls




2



Parentage of Open MPI

- ❑ LAM/MPI (Indiana U.)
- ❑ FT-MPI (U. of Tennessee)
- ❑ LA-MPI (Los Alamos, Sandia)
- ❑ PACX-MPI (HLRS, U. Stuttgart)



3



Motivation for a new MPI

- ❑ Each prior project had different strong points
 - Could not easily combine into one code base
- ❑ New concepts could not easily be accommodated in old code bases
- ❑ Easier to start over
 - Start with a blank sheet of paper
 - Many years of collective implementation experience



4



Open MPI was born in 2003

- Developers of FT-MPI, LA-MPI, LAM/MPI
 - Kept meeting at conferences in 2003
 - Culminated at SC 2003: Let's start over
 - Open MPI was born
- Started serious design and coding work
January 2004
 - All of MPI except one-sided operations
 - First release 1Q 2005



5



Open MPI Today

- Version 1.0 released in November 2005
- Version 1.1 released in June 2006
 - Added support for one-sided communications
 - Version 1.1.2 released in October 2006
 - Version 1.1.3 soon
- Version 1.2b1 (beta) released today
 - Library level matching: Myrinet/MX & Portals
 - First set of Tuned Collectives
 - Sun N1 Grid Engine
 - Data Reliability



6



Where we are going...

- Fault Tolerance Support
 - Checkpoint/Restart
 - Process Migration & Batch Scheduling
 - Data Reliability
- Multicore optimizations
 - Processor Affinity
 - Memory Affinity
 - Shared Memory Collectives



7



Where we are going...

- Improved Collectives
 - Topology Aware Collectives
 - Non-blocking Collectives
- Next generation one-sided primitives
- Tool support & Integration
 - Vampir NG
 - Enhanced TotalView support
- All Major Interconnects



8



Where we are going...

- ❑ Open MPI is an open extensible research platform
- ❑ Where we go will also be determined by you, the user community!



9



Open MPI Events @ IU Booth

- ❑ Today @ 2:00pm
iPod nano giveaway!
- ❑ Wednesday 4:00 - 4:50pm
Open MPI: A research platform
- ❑ Thursday 11:00 - 11:50am
Dealing with disaster: Fault Tolerance in Open MPI



10



Questions?



What you all came here for...

The iPod nano giveaway!

