# Open Source High Performance Computing With Open MPI


**Jeff Squyres**

**Cisco Systems**

  

## High Performance

- Two of the top 10
  - #2 Roadrunner (Los Alamos)
  - #10 Red Sky (Sandia)
- Collaborate to innovate
  - – Vendors
  - – Academic researchers
  - – Lab scientists

  

## Open MPI Is…

- Evolution of several prior MPI implementations
- Open source project and community
  - Production quality
  - Vendor-friendly
  - Research- and academic-friendly
- All of MPI-2.1
  - Working on MPI-2.2

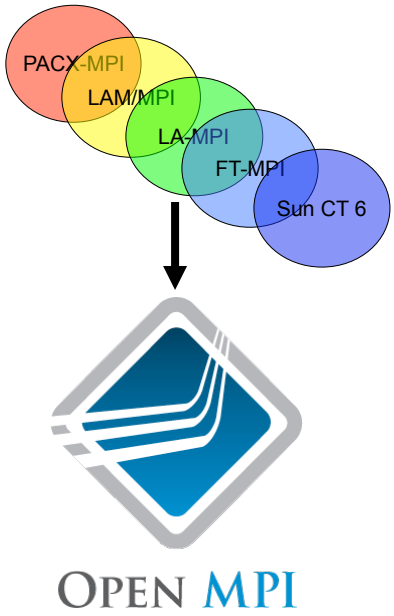




## 16 Members, 9 Contributors, 2 Partners

## Why Does Open MPI Exist?

- Maximize all MPI expertise
  - Research / academia
  - Vendors
  - Customers, enterprise
  - …elsewhere
- Capitalize on years of MPI research and implementation experience
- **The sum is greater than the parts**

 

"Great discoveries and improvements invariably involve the cooperation of many minds."

Alexander Graham Bell, 1877

 

## Cisco: Why Open MPI?

- It seems obvious to us!
  - Why re-invent the wheel?
  - Established, high quality MPI
  - Combined community resources
- Meshes with Cisco values
  - Standards-based
  - Open architectures
  - Consensus driven
  - Collaborate to innovate

Cisco votes “yes” for community MPI

 

## Why Open Source?

- Open source HPC is good for everyone
  - Room for research / new ideas
  - Open information transfer
  - Feed them back into production / commodity products
- Shorten the cycle from research to commodity
- Researchers have ideas; industry has production capability
  - There are smart people in both!

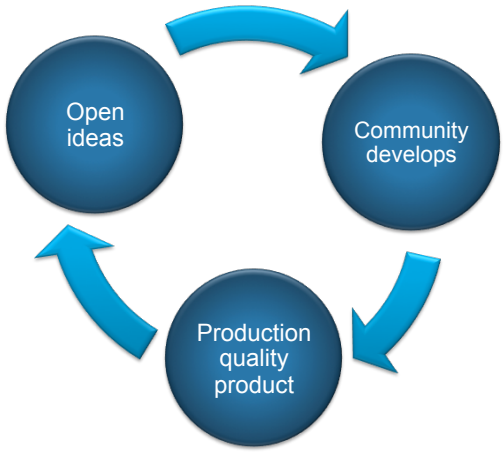


 

## Open MPI Sub Project: Hardware Locality

- Initially developed by INRIA
- Command line and C API
- Discover the topology of your machine
  - Processors: sockets, cores, threads
  - Memory: caches, main RAM
- *Incredibly important as we move towards manycore*

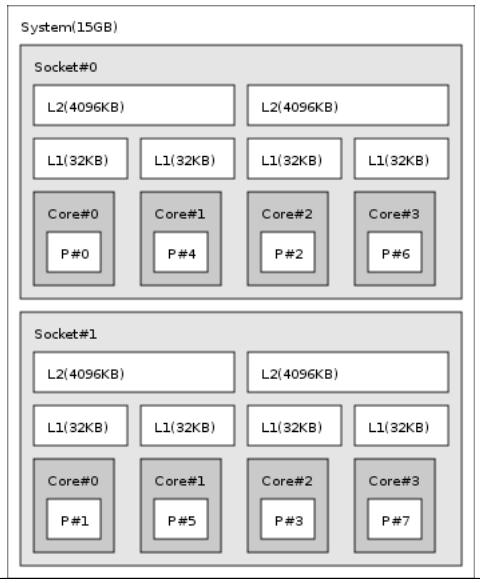






## Cisco's Open MPI Community Role

- Active development
  - Design, code
- [Very] Extensive testing
  - 100K's regression tests/night
  - Data fed back to community
- Logistics support
  - Collaboration, facilitation
  - Face-to-face engineering meetings
- Member, MPI Forum

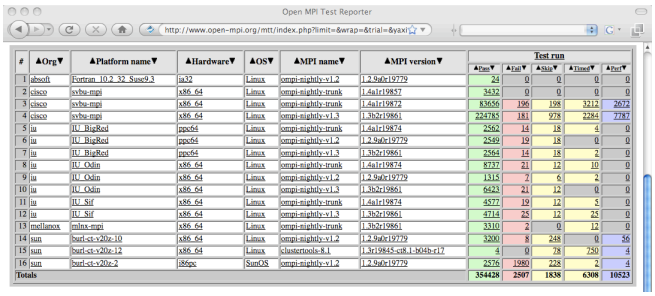


| # | Org | Platform name | Hardware | OS | MPI name | MPI version | Pass | Fail | Skip | Timed | Perf |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | absoft | Fortran_10.2_32_Suse9.3 | ia32 | Linux | ompi-nightly-v1.2 | 1.2.9a0r19779 | 24 | 0 | 0 | 0 | 0 |
| 2 | cisco | svbu-mpi | x86_64 | Linux | ompi-nightly-trunk | 1.4a1r19857 | 3432 | 0 | 0 | 0 | 0 |
| 3 | cisco | svbu-mpi | x86_64 | Linux | ompi-nightly-trunk | 1.4a1r19872 | 83656 | 196 | 198 | 3212 | 2672 |
| 4 | cisco | svbu-mpi | x86_64 | Linux | ompi-nightly-v1.3 | 1.3b2r19861 | 224785 | 181 | 978 | 2284 | 7787 |
| 5 | iu | IU_BigRed | ppc64 | Linux | ompi-nightly-trunk | 1.4a1r19874 | 2562 | 14 | 18 | 6 | 0 |
| 6 | iu | IU_BigRed | ppc64 | Linux | ompi-nightly-v1.2 | 1.2.9a0r19779 | 2549 | 19 | 18 | 0 | 0 |
| 7 | iu | IU_BigRed | ppc64 | Linux | ompi-nightly-v1.3 | 1.3b2r19861 | 2564 | 14 | 18 | 2 | 0 |
| 8 | iu | IU_Odin | x86_64 | Linux | ompi-nightly-trunk | 1.4a1r19874 | 8737 | 21 | 12 | 10 | 0 |
| 9 | iu | IU_Odin | x86_64 | Linux | ompi-nightly-v1.2 | 1.2.9a0r19779 | 1315 | 7 | 6 | 2 | 0 |
| 10 | iu | IU_Odin | x86_64 | Linux | ompi-nightly-v1.3 | 1.3b2r19861 | 6623 | 21 | 12 | 0 | 0 |
| 11 | iu | IU_Sif | x86_64 | Linux | ompi-nightly-trunk | 1.4a1r19874 | 4577 | 19 | 12 | 5 | 0 |
| 12 | iu | IU_Sif | x86_64 | Linux | ompi-nightly-v1.3 | 1.3b2r19861 | 4714 | 25 | 12 | 25 | 0 |
| 13 | mellanox | mlnx-mpi | x86_64 | Linux | ompi-nightly-trunk | 1.3b2r19861 | 3310 | 2 | 0 | 12 | 0 |
| 14 | sun | burl-ct-v20z-10 | x86_64 | Linux | ompi-nightly-v1.2 | 1.2.9a0r19779 | 3200 | 8 | 248 | 0 | 56 |
| 15 | sun | burl-ct-v20z-12 | x86_64 | Linux | clustertools-8.1 | 1.3r19845-cx8.1-b04b-r17 | 4 | 0 | 78 | 750 | 4 |
| 16 | sun | burl-ct-v20z-2 | i86pc | SunOS | ompi-nightly-v1.2 | 1.2.9a0r19779 | 2576 | 1980 | 228 | 2 | 4 |
| Totals | | | | | | | 354428 | 2507 | 1838 | 6308 | 10523 |

## Cisco's Open MPI Goals

| Technical | Non technical |
|---|---|
| ▪ Promote standards<br>– MPI Forum leadership<br>– Ethernet-based technologies<br>– Cisco Unified Computer Servers<br>– Commodity Clusters<br>▪ Integrate with tools<br>– Provide deep insight in to complex problems<br>– Make parallel programming [a little] easier | ▪ Promote community<br>– Conferences, tradeshows<br>– Contribute on open mailing lists<br>▪ Partner with academics and researchers<br>– Foster cutting-edge research<br>▪ Perform "community service"<br>– Example: Fortran API maintenance |




"Open source is decided by those who show up."

Cisco is there. Come join us.

welcome to
the human network.