



Open MPI: 10^{15} Flops Can't Be Wrong

Open Source High Performance Computing




Jeff Squyres

Open MPI Architect

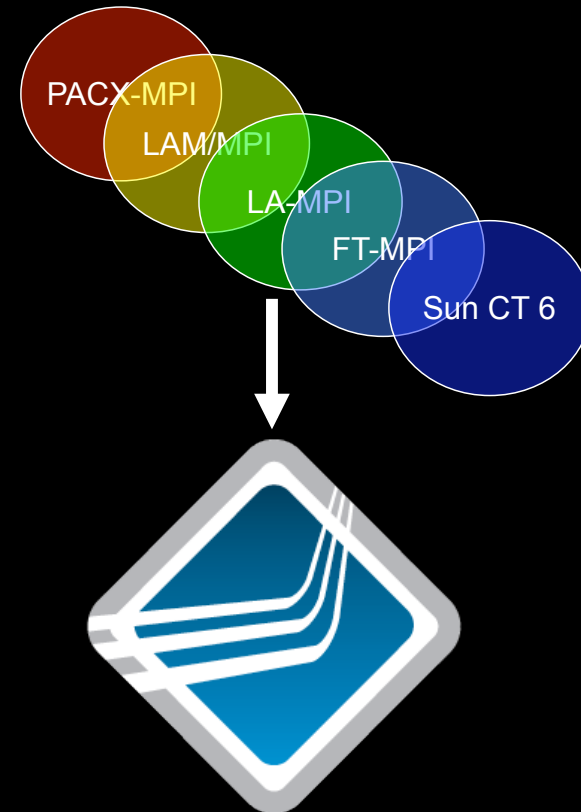
Petaflop!!

- Los Alamos Road Runner
- #1 on Nov. 2008 Top500
1.1 petaflops
- **Powered by Open MPI**
Significant community achievement



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-1 and MPI-2



OPEN MPI

15 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?
 - Who would want “Cisco MPI”?
 - Combined community resources
- Meshes with Cisco values
 - Standards-based
 - Open architectures
 - Consensus driven
 - Collaborate to innovate



Cisco votes “yes”
for community MPI

Cisco's Community Role

- Active development
 - Design, code
- [Very] Extensive testing
 - 300-500k regression tests/night
 - Data fed back to community
- Logistics support
 - Face-to-face engineering meetings
- Member, MPI Forum

The screenshot shows the Open MPI Test Reporter interface. The browser address bar displays the URL: <http://www.open-mpi.org/mtt/index.php?limit=&wrap=&trial=&yaxi>. The main content is a table with the following columns: #, ▲Org▼, ▲Platform name▼, ▲Hardware▼, ▲OS▼, ▲MPI name▼, ▲MPI version▼, and a 'Test run' section with sub-columns: ▲Pass▼, ▲Fail▼, ▲Skip▼, ▲Timed▼, and ▲Perf▼.

#	▲Org▼	▲Platform name▼	▲Hardware▼	▲OS▼	▲MPI name▼	▲MPI version▼	Test run				
							▲Pass▼	▲Fail▼	▲Skip▼	▲Timed▼	▲Perf▼
1	abssoft	Fortran_10.2_32_Suse9.3	i386	Linux	ompi-nightly-v1.2	1.2.9a0r19779	24	0	0	0	0
2	cisco	svbu-mpi	x86_64	Linux	ompi-nightly-trunk	1.4a1r19852	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	4	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	2	6	2	0
10	iu	IU_Odin	x86_64	Linux	ompi-nightly-v1.3	1.3b2r19861	6423	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	hurl-ct-v20z-10	x86_64	Linux	ompi-nightly-v1.2	1.2.9a0r19779	3200	8	248	0	56
15	sun	hurl-ct-v20z-12	x86_64	Linux	clustertools-8.1	1.3r19845-ct8.1-b04b-r17	4	0	78	750	4
16	sun	hurl-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

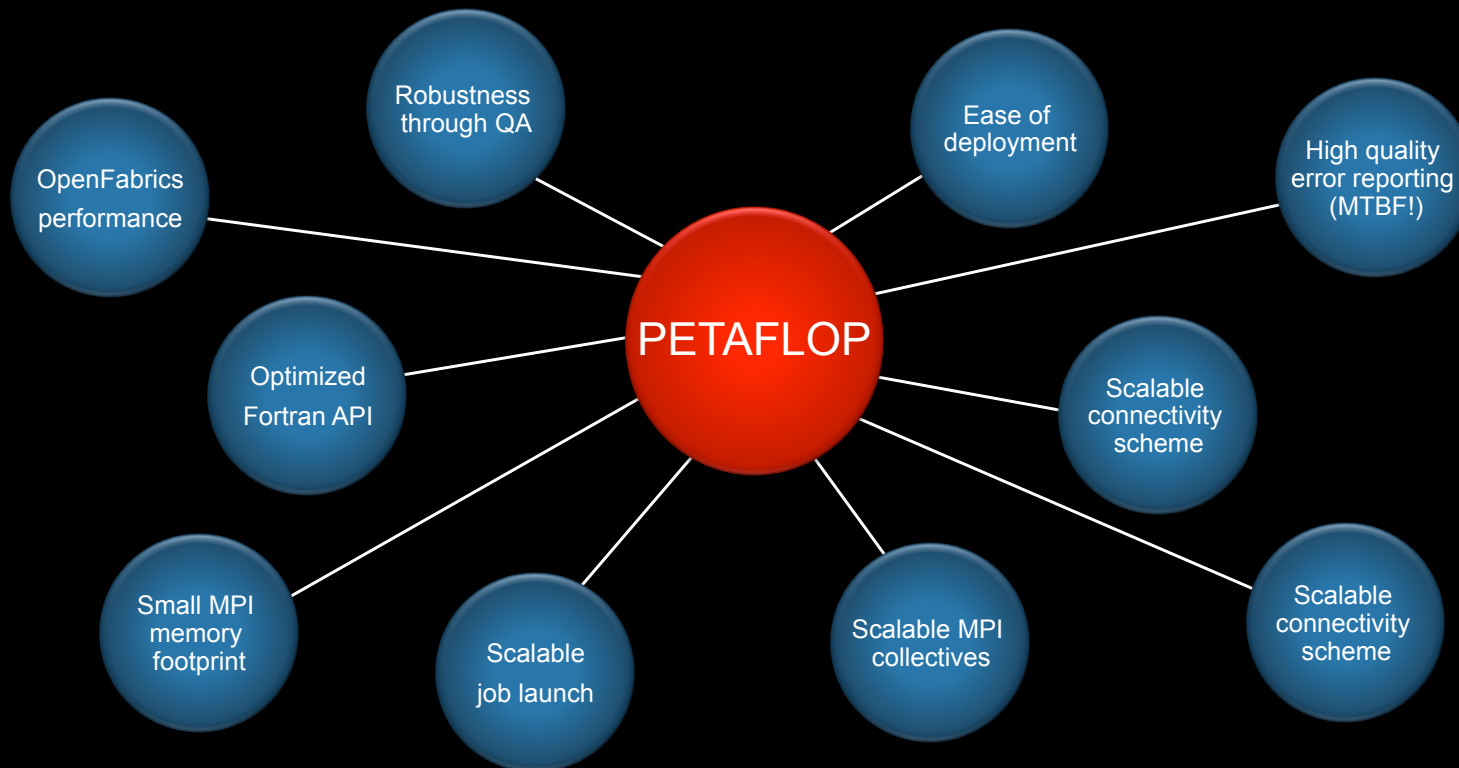
Technical

- Promote standards
 - Ethernet-based technologies
 - Commodity clusters
- Integrate with tools
 - Make parallel programming [a little] easier
- Understand and accelerate applications
 - RAB and DAL two (pseudo-HPC) examples

Non technical

- Promote community
 - Conferences, tradeshow
 - Contribute on open mailing lists
- Partner with academics and researchers
 - Foster cutting-edge research
- Perform “community service”
 - Example: Fortran API maintenance

...But How Does That Equal a Petaflop?



“Open source is decided by those who show up.”

Cisco is there. Come join us.

welcome to
the human network.

