



Open MPI: Open Source + Innovation + Community = Petaflop

Jeff Squyres

Open MPI Architect

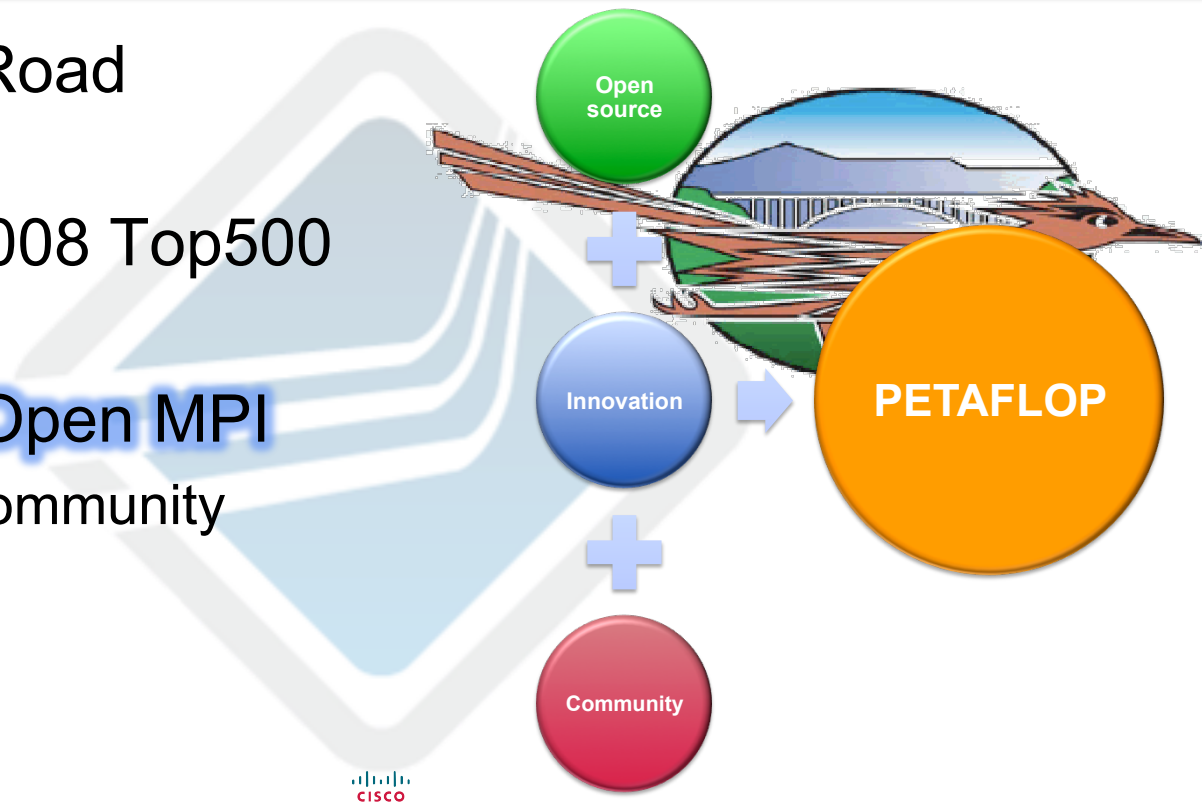


Aside: MPI-2.1 and MPI-3.0

- MPI-2.1 is complete
 - Merged MPI-1 and MPI-2 documents (yay!)
 - \$22 printed books (586 pages!), HLRS booth #1353
- The MPI Forum needs your help!
 - What do you want to see in MPI-3.0?
 - What do you not want to see in MPI-3.0?

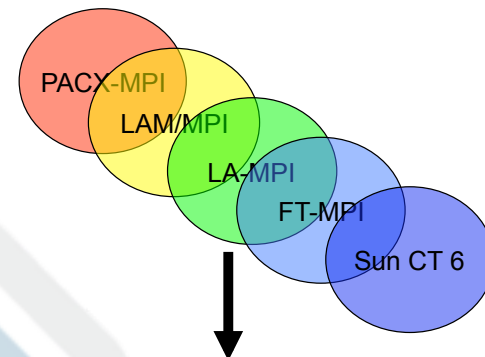
Petaflop!!

- 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

Why Does Open MPI Exist?


- Maximize all MPI expertise, including:
 - Research / academia
 - Vendors
 - Customers, enterprise
- Utilize years of MPI research and experience
- The sum is greater than the parts

15 Members, 9 Contributors, 2 Partners



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!



Alexander Graham Bell, 1877

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

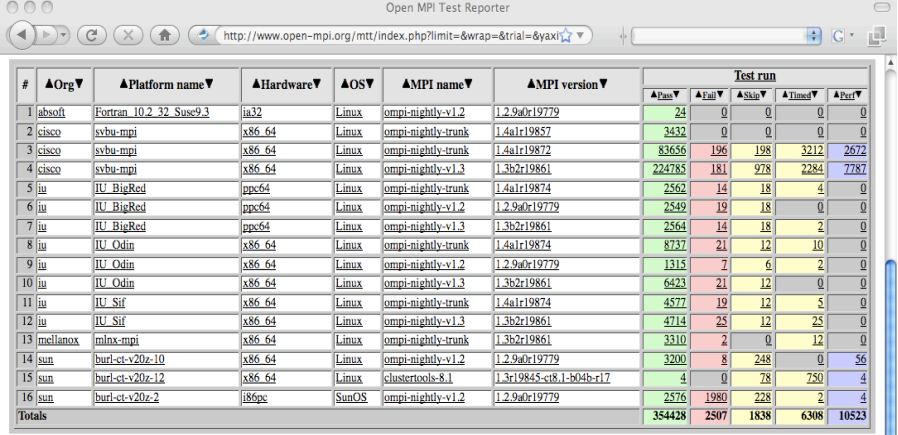
How Does It Work?

- It's all about the collaboration
 - Communicate, communicate, communicate
 - Developers scattered across three continents
- Our greatest development tool:



Give To Get

- Nightly community regression testing
 - 300-500k tests per night
 - Web-based analysis tools
- Strive for consensus
 - But realize it isn't always possible (or necessary)
- Perform “community service”
 - Example: Fortran API maintenance

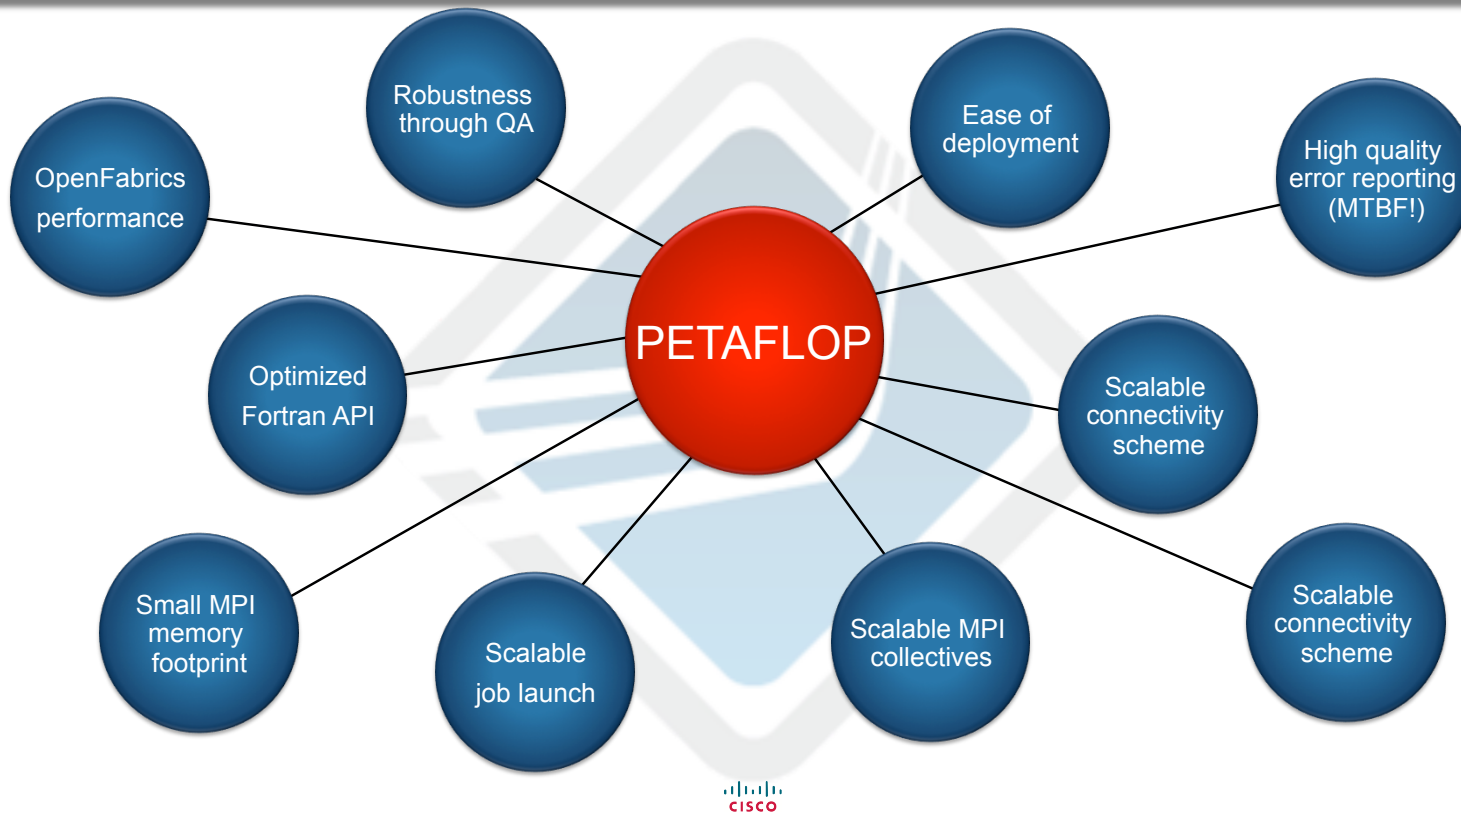


Open MPI Test Reporter

<http://www.open-mpi.org/mtt/index.php?limit=&wrap=&trial=&yaxi>

#	▲Org▼	▲Platform name▼	▲Hardware▼	▲OS▼	▲MPI name▼	▲MPI version▼	Test run				
							▲Pass▼	▲Fail▼	▲Skip▼	▲Time▼	▲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	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	7	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	burl-et-v20z-10	x86_64	Linux	ompi-nightly-v1.2	1.2.9a0r19779	3200	8	248	0	56
15	sun	burl-et-v20z-12	x86_64	Linux	clustertools-8.1	1.3r19845-et8.1-b04b-r17	4	0	78	750	4
16	sun	burl-et-v20z-2	ib6pc	SunOS	ompi-nightly-v1.2	1.2.9a0r19779	2576	1980	228	2	4
Totals							354428	2507	1838	6308	10523

...But How Does That Equal a Petaflop?



Community Makes It Possible



Come Join Us!

“Open source is decided by those who show up”

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

welcome to
the human network.

