



Open MPI: Open Source + Innovation + Community = Petaflop

Jeff Squyres

Open MPI Architect

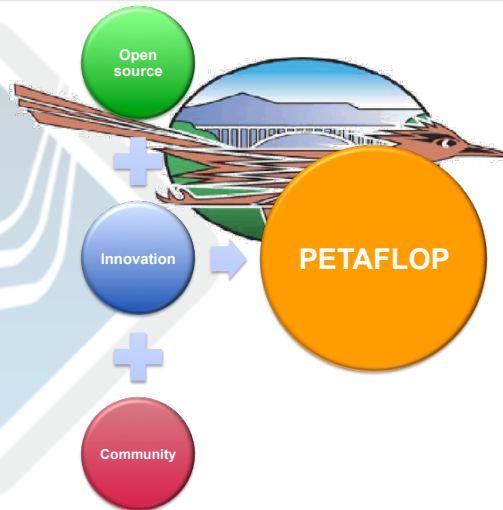


Aside: MPI-2.1 and MPI-3.0

- 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



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

 CISCO


15 Members, 9 Contributors, 2 Partners



 CISCO

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:

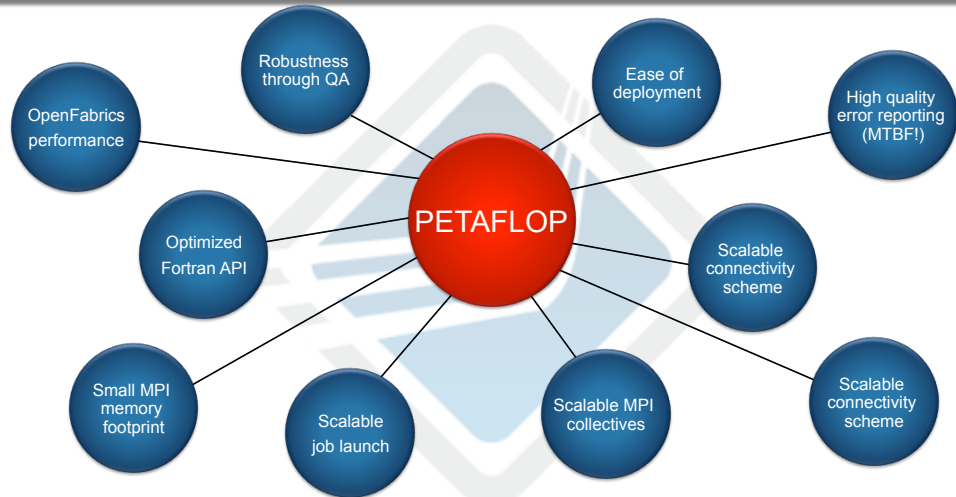


Give To Get

- 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

#	Org	Platform name	Hardware	AOS	MPI name	MPI version	Test run				
							Attn	Attn	Attn	Attn	
1	atsd	Fortran_102_32_Suse9.3	ia32	Linux	openmpi-1.2	1.2.96c19779	24	0	0	0	
2	cisco	infini-mpi	x86_64	Linux	openmpi-1.2-trunk	1.461119821	252	0	0	0	
3	cisco	infini-mpi	x86_64	Linux	openmpi-1.2-trunk	1.461119821	83056	196	199	2212	
4	cisco	infini-mpi	x86_64	Linux	openmpi-1.3	1.362119861	222785	181	978	2281	
5	in	ll_Suse9c	ppc64	Linux	openmpi-1.2-trunk	1.461119821	2561	14	15	4	
6	in	ll_Suse9c	ppc64	Linux	openmpi-1.2	1.2.96c19779	2561	14	15	4	
7	in	ll_Suse9c	ppc64	Linux	openmpi-1.3	1.362119861	2564	14	15	4	
8	in	ll_Osln	x86_64	Linux	openmpi-1.2-trunk	1.461119821	8737	21	11	10	
9	in	ll_Osln	x86_64	Linux	openmpi-1.2	1.2.96c19779	1315	2	6	2	
10	in	ll_Osln	x86_64	Linux	openmpi-1.3	1.362119861	6642	21	11	0	
11	in	ll_Sf	x86_64	Linux	openmpi-1.2-trunk	1.461119821	4371	10	11	4	
12	in	ll_Sf	x86_64	Linux	openmpi-1.3	1.362119861	4714	24	11	25	
13	netlinx	infini-mpi	x86_64	Linux	openmpi-1.2-trunk	1.362119861	3310	2	0	12	
14	in	infini-200c_10	x86_64	Linux	openmpi-1.2	1.2.96c19779	2300	8	288	0	
15	in	infini-200c_14	x86_64	Linux	openmpi-1.1	1.2019651c41.1305b111	4	0	78	267	
16	in	infini-200c_2	ppc64	Linux	openmpi-1.2	1.2.96c19779	3576	1980	228	2	
Totals							354428	2397	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.

