



Open MPI: 10¹⁵ Flops Can't Be Wrong

Open Source High
Performance Computing



Jeff Squyres
Open MPI Architect

Cisco SC'08 Booth Presentation

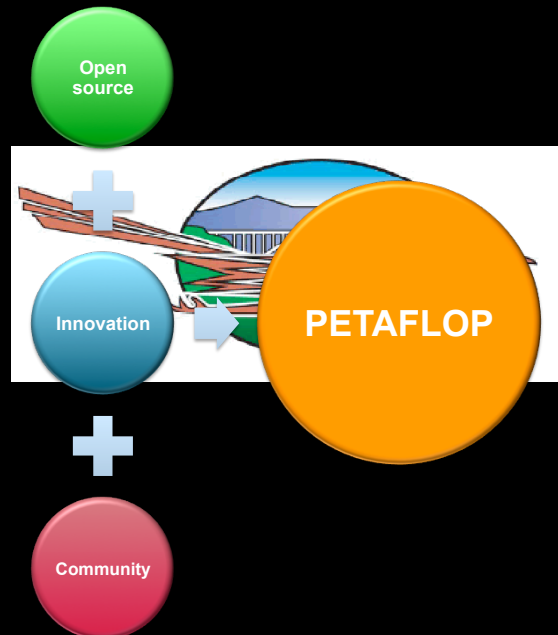
© 2008 Cisco Systems, Inc. All rights reserved.

Cisco Public

1

Petaflop!!

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



Cisco SC'08 Booth Presentation


© 2008 Cisco Systems, Inc. All rights reserved.

Cisco Public

2

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



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

#	Org	Platform name	Hardware	OS	MPI name	MPI version	Test Run				
							Adm	App	AppV	AppV	
1	ibm	PowerPC	ppc64	Linux	openmpi-1.2	1.2.9a6c19739	248	0	0	0	
2	cisco	vxba-mpi	x86_64	Linux	openmpi-1.2	1.4a119874	2452	0	0	0	
3	cisco	vxba-mpi	x86_64	Linux	openmpi-1.2	1.4a119874	83656	196	199	3212	
4	cisco	vxba-mpi	x86_64	Linux	openmpi-1.3	1.3b2c19861	224785	181	278	2284	
5	iu	BigRed	ppc64	Linux	openmpi-1.2	1.4a119874	2562	14	18	4	
6	iu	BigRed	ppc64	Linux	openmpi-1.2	1.2.9a6c19739	2592	20	18	0	
7	iu	BigRed	ppc64	Linux	openmpi-1.3	1.3b2c19861	2564	14	18	2	
8	iu	Odin	x86_64	Linux	openmpi-1.2	1.4a119874	8737	21	12	10	
9	iu	Odin	x86_64	Linux	openmpi-1.2	1.2.9a6c19739	1315	2	6	2	
10	iu	Odin	x86_64	Linux	openmpi-1.3	1.3b2c19861	6542	21	12	0	
11	iu	Sif	x86_64	Linux	openmpi-1.2	1.4a119874	4572	19	12	5	
12	iu	Sif	x86_64	Linux	openmpi-1.3	1.3b2c19861	4714	25	12	25	
13	indiana	mlnx-mpi	x86_64	Linux	openmpi-1.2	1.3b2c19861	3310	2	0	12	
14	sun	cs-x20c-10	x86_64	Linux	openmpi-1.2	1.2.9a6c19739	3200	8	248	0	
15	sun	cs-x20c-12	x86_64	Linux	openmpi-1.1	1.2b19445-c8b1304brc13	4	0	78	26	
16	sun	cs-x20c-2	SPARC	SunOS	openmpi-1.2	1.2.9a6c19739	2576	1980	228	2	
Totals							354428	2597	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, tradeshows
 - Contribute on open mailing lists
- Partner with academics and researchers
 - Foster cutting-edge research
- Perform "community service"
 - Example: Fortran API maintenance

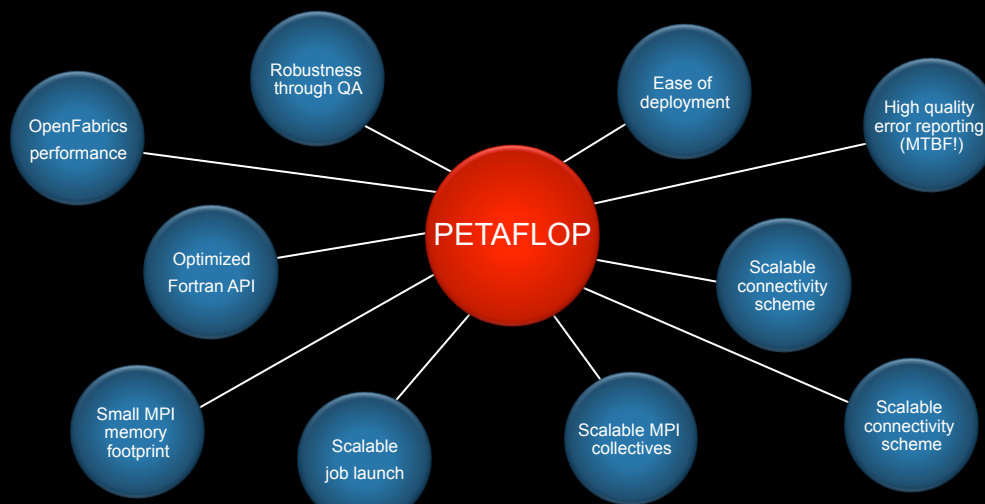
Cisco SC'08 Booth Presentation

© 2008 Cisco Systems, Inc. All rights reserved.

Cisco Public

9

...But How Does That Equal a Petaflop?



Cisco SC'08 Booth Presentation

© 2008 Cisco Systems, Inc. All rights reserved.

Cisco Public

10

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

Cisco is there. Come join us.

welcome to
the human network.

