

LA-UR-15-22242

Approved for public release; distribution is unlimited.

Title: Open MPI and Open Fabrics Software

Author(s): Pritchard, Howard Porter Jr.

Intended for: Open Fabrics User Group 2015, 2015-03-19 (Monterey, California, United States)
presentation

Issued: 2015-03-27

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.



OpenFabrics
Software
User Group
Workshop

Open MPI and Open Fabrics Software

Howard Pritchard (howardp@lanl.gov)

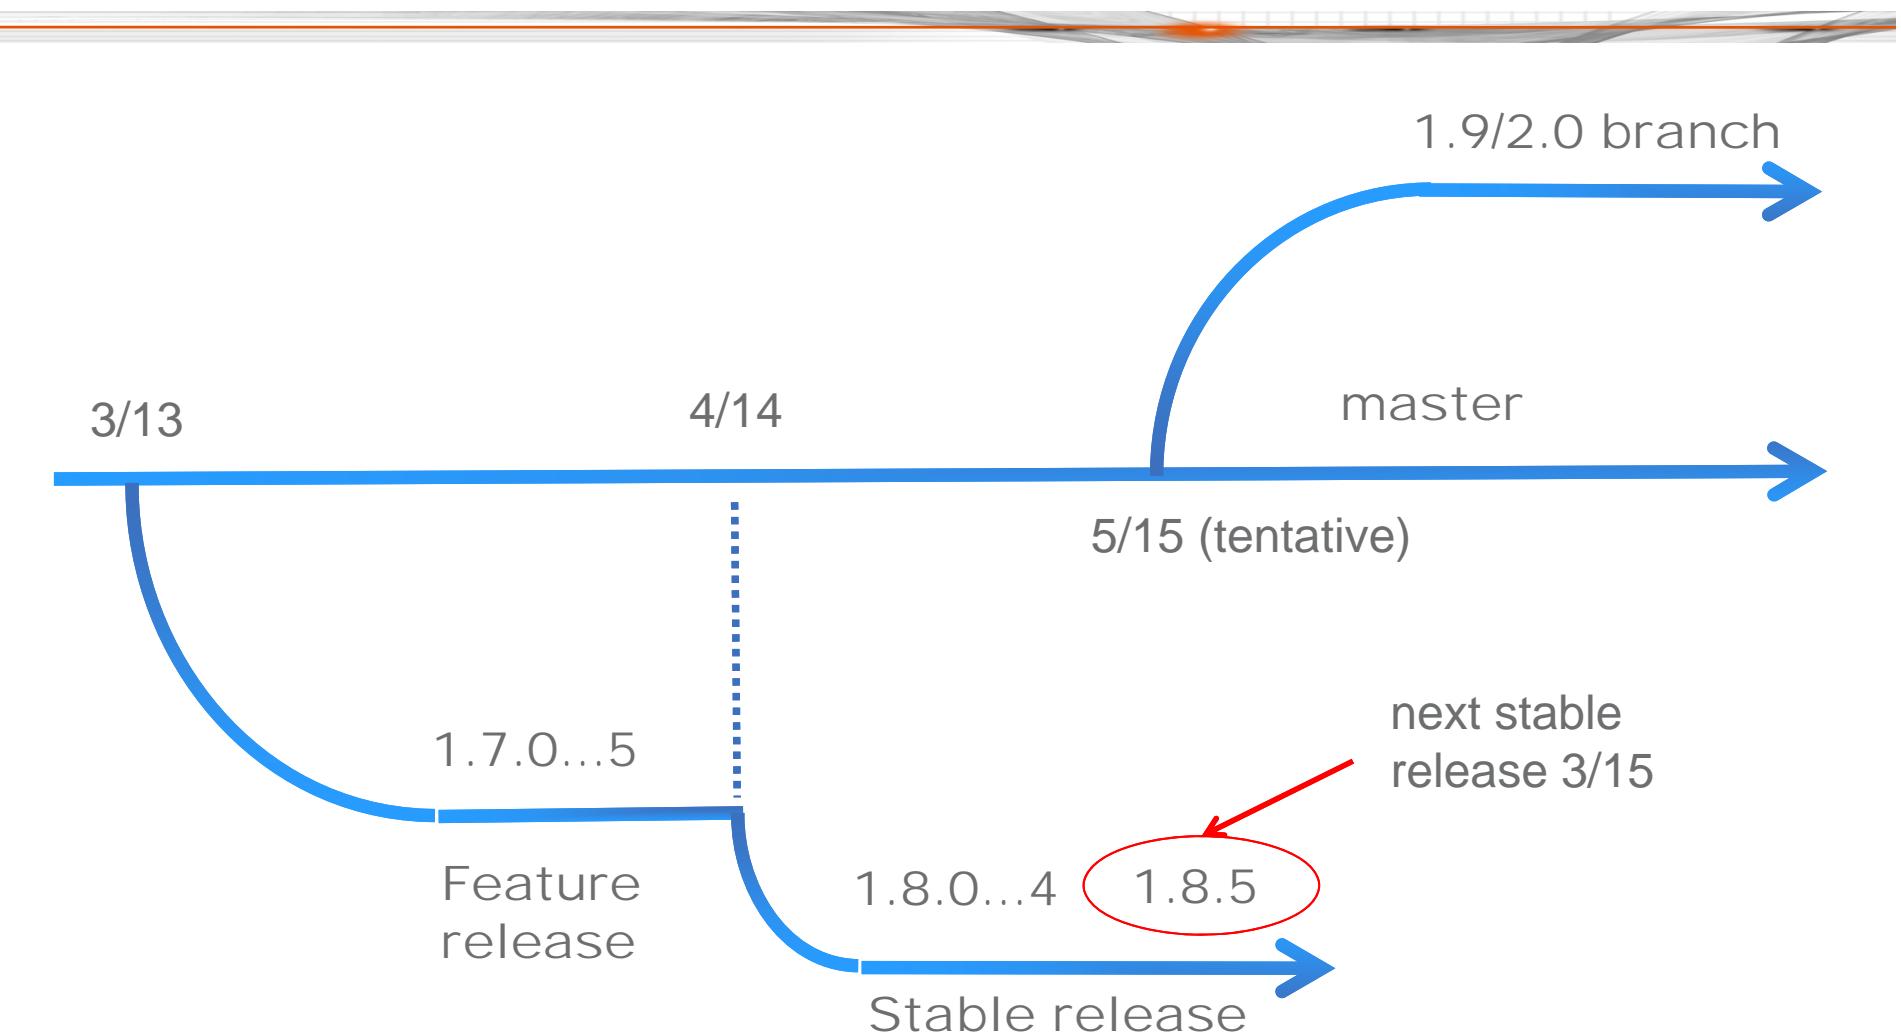
#OFSUserGroup

Outline



- Release timeline
- Github
- Upcoming Features
- Status of Open MPI and Open Fabrics software
- What's trending

Release Timeline



<https://github.com/open-mpi/mpi>



This screenshot shows the GitHub repository page for `open-mpi / mpi`. The repository has 22,339 commits, 1 branch, 1 release, and 59 contributors. The `master` branch is selected. A list of recent commits is shown, with the latest commit being `ccba8ce856` by `hjelmn` 9 hours ago. The commits are categorized by file type (e.g., config, contrib, examples, ompi, opal, orte, oshmem, test, .gitignore, .mailmap, AUTHORS) and include descriptions and timestamps. The GitHub interface includes a header with navigation links (Explore, Gist, Blog, Help), a user profile for `hpritcha`, and a sidebar with links to Code, Issues (213), Pull requests (18), Wiki, Pulse, and Graphs. There is also an `HTTPS` clone URL and download options for `Clone in Desktop` and `Download ZIP`.

Open MPI main development repository

22,339 commits 1 branch 1 release 59 contributors

branch: master ompi / +

Merge pull request #457 from hjelmn/mpit_fixes ...

hjelmn authored 9 hours ago latest commit ccba8ce856

File	Description	Time Ago
config	opal_check_pmi: protect un-prefixed shell variables	6 days ago
contrib	libldl: remove libldl from the tree	9 days ago
examples	Add missing files to tarball	a month ago
ompi	Merge pull request #457 from hjelmn/mpit_fixes	9 hours ago
opal	Merge pull request #457 from hjelmn/mpit_fixes	9 hours ago
orte	grpcomm: fixed brks and rcd algorithms - added enough space for masks...	11 hours ago
oshmem	(ortelompiloshmem)*info tools: convert to opal_dl interface	9 days ago
test	datatype: trivial spelling fix	19 days ago
.gitignore	Update orte-submit manpage, add the ompi-* versions of orte-dvm and o...	2 months ago
.mailmap	Update the emails from UTK folks.	a year ago
AUTHORS	AUTHORS: Updated with Git email addresses (vs. SVN IDs)	6 months ago

Code Issues (213) Pull requests (18) Wiki

Pulse Graphs

HTTPS clone URL <https://github.com/>

You can clone with [HTTPS](#), [SSH](#), or [Subversion](#).

Clone in Desktop Download ZIP

Github move

- Special thanks to David Goodell and Jeff Squyres for extra effort to make this happen
- Open git issue if you wish to report a bug
- Better yet, do a fork, fix the problem, then issue a pull request - @jsquyres or @hpritcha
 - Note a major feature may require signing project's collaborator agreement
- Check the [git wiki page](#) for info on working with Open MPI on github, etc.

Upcoming Features – 1.9/2.0 release stream



- MPI 3.0 errata and MPI 3.1 compliant (assuming 3.1 is passed at June '15 forum meeting)
 - Possible exception of ticket [#273](#)
- MPI_THREAD_MULTIPLE support for important BTLs
- Job startup improvements and reduced memory footprint connection management info
- Libfabric adoption (usNIC, OFI MTL)
- Mellanox specific improvements
- Streamlined build process for Cray systems
- Improved asynchronous progress for some BTLs
- ScoreP support
- One or more MPI-4 features (as MPIX_*)

Deprecated Features in 1.9/2.0

- Cray XT systems
 - May terminate support for Cray XE CLE 4.X
- MX BTL
- SCIF BTL (?)
- Coll ML component (?)



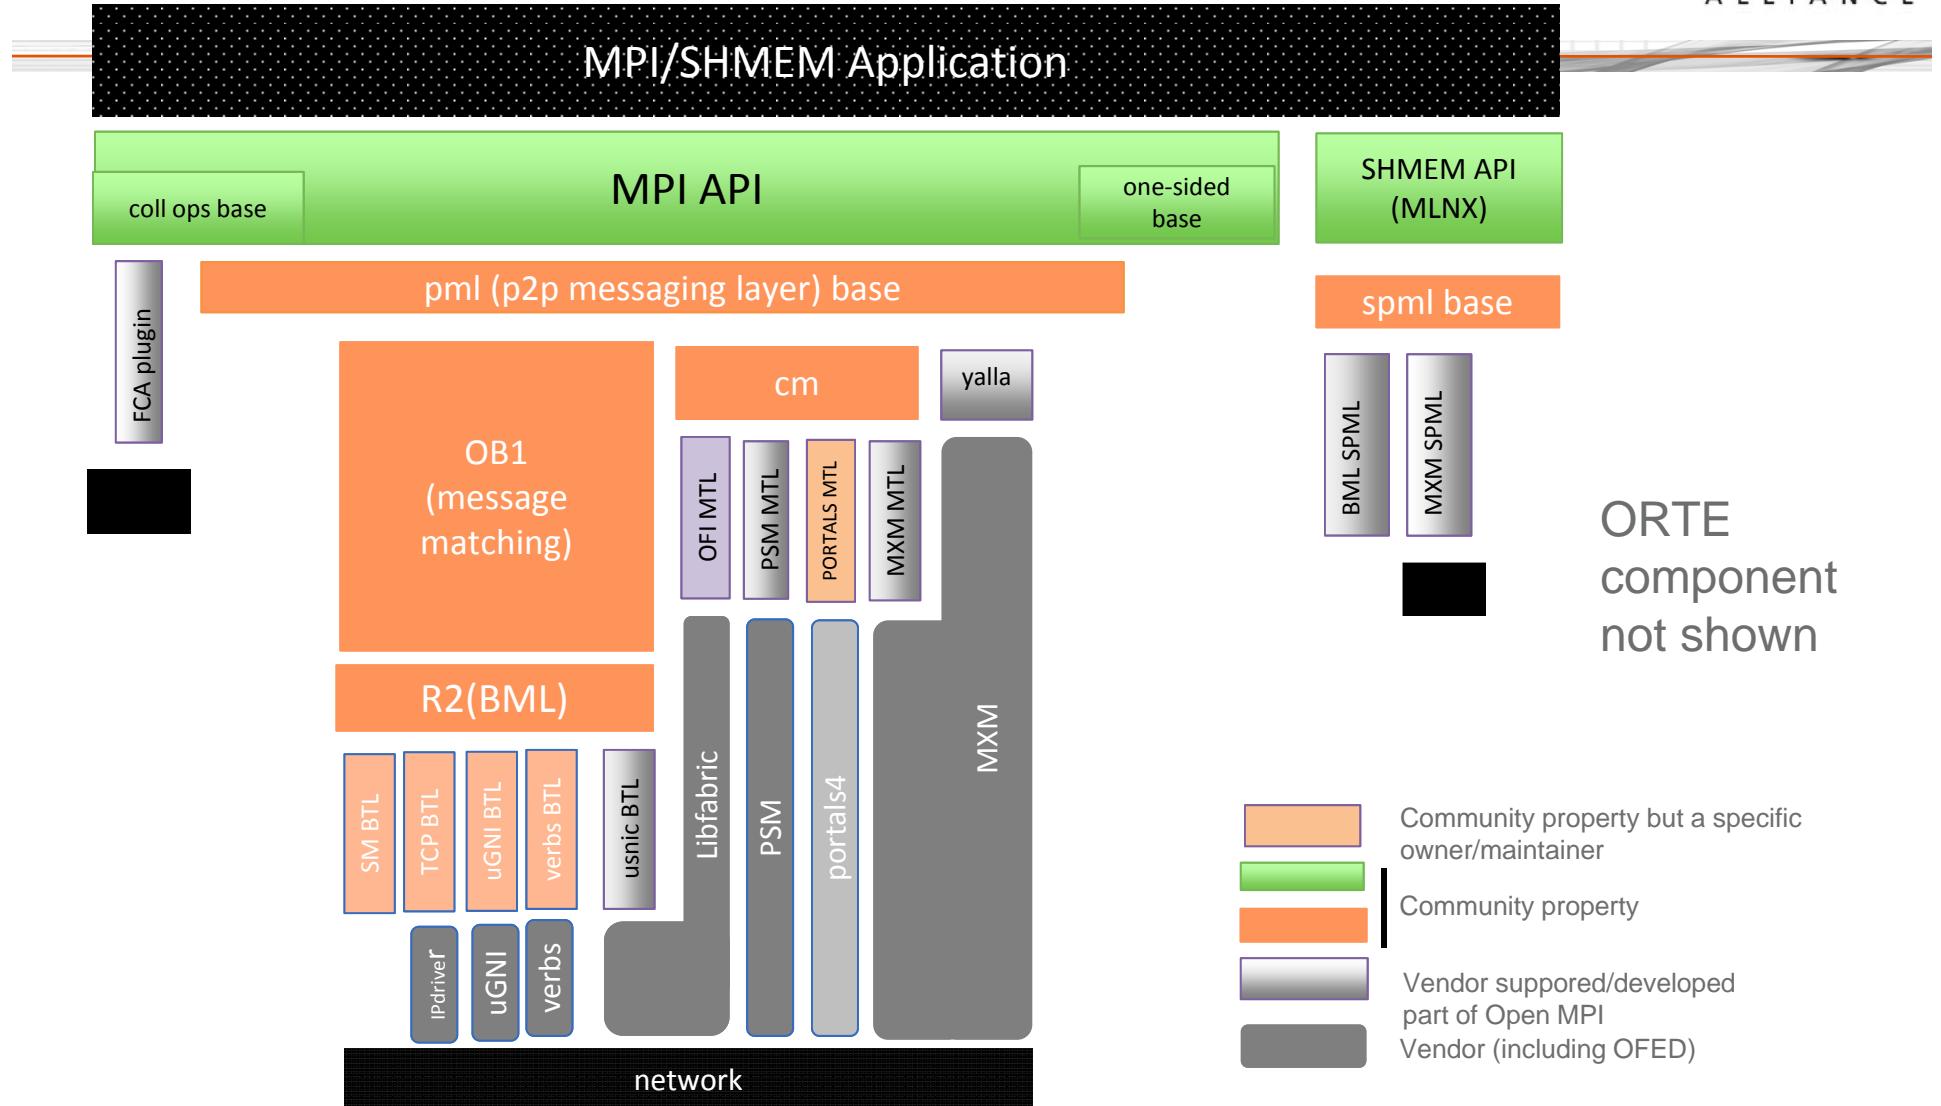
Open Fabrics and Open MPI

OFED Software use in Open MPI (master)



- Open IB BTL
- MXM, PSM, and libfabric MTLs
- Mellanox FCA collective accelerator package
- Out of band control network (used by mpirun)
 - Uses UD

Open MPI Project Code Structure



What's in/what's not

- For Mellanox HCA's use MXM/FCA
 - Well supported by the vendor
 - Actively being developed/enhanced
 - export MXM_TLS=rc,shm,self for best performance, for scalability ud,shm,self, latest HCA/firmware, dc,shm,self
- For Intel InfiniPath HCA's use PSM
 - OFI libfabric MTL is available for a check-out
- The openib BTL in Open MPI will be maintained, but no development is planned
- May implement a slimmed down openib BTL specifically for iWarp
- The openib BTL may be deprecated in next release stream (2.1/2.2)

Questions?



Thank You



OpenFabrics Software
User Group Workshop

#OFSUserGroup