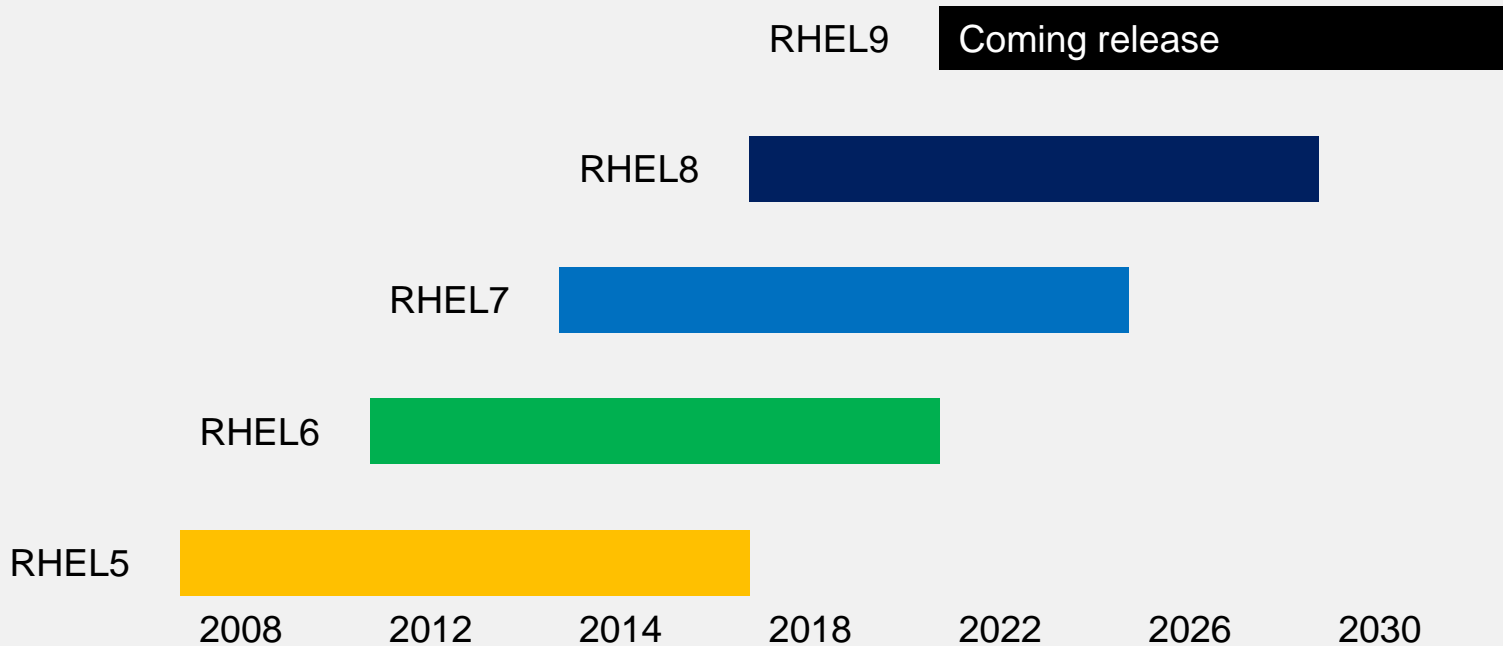


Integration MVAPICH2 into RedHat Enterprise Linux

Honggang Li
August 25, 2021

MVAPICH2 rpm packaging history



In 2009, RedHat packaged MVAPICH2 for RHEL5.3.

In later **major** releases of RHEL, MVAPICH2 always been added as RHEL specific package.

Using MVAPICH2

- Load environment module
 - **mpi/mvapich2-psm2-x86_64** *or* **mpi/mvapich2-x86_64**
- Compile MPI source code with mpicc/mpicxx/mpifort, and run
- Rpm packaging special require
 - Must install the binary executables under same directory of mpicc
 - https://bugzilla.redhat.com/show_bug.cgi?id=1958960
 - Workaround: explicitly requires mvapich2/mvapich2-psm/mvpaich2-psm2
 - Requires: mvapich2
 - See mpitests.spec (centos 9 stream) file for example
 - <https://gitlab.com/redhat/centos-stream/rpms/mpitests/-/blob/c9s/mpitests.spec>

Supported mainstream RDMA hardware

| | Verbs (InfiniBand, IWARP, RoCE) | | | | | | | | | PSM | PSM2 |
|-------|---------------------------------|-------|--------|---------|-------|-------|------|------|------|-----|------|
| | cxgb3 | cxgb4 | ocrdma | bnxt_re | ipath | mthca | mlx4 | mlx5 | qedr | qib | hfi1 |
| RHEL9 | | | | | | | | | | | |
| RHEL8 | | | | | | | | | | | |
| RHEL7 | | | | | | | | | | | |
| RHEL6 | | | | | | | | | | | |
| RHEL5 | | | | | | | | | | | |

1. Verbs based build also available for qib and hfi1/opa
2. Only one type of mlx4 hardware will be supported for rhel9
3. usnic, vma_pvrDMA, soft-rxe and soft-iwarp are not supported
4. psm3 code is available in upstream Linux kernel and libfabric, but not supported

Benchmarks

- OSU benchmarks
 - `osu_bw`, `osu_latency`, ...
- Intel MPI benchmarks
 - `MPI1`, `EXT`, `IO`, `MT`, `NBC`, `P2P`, `RMA`
- Only run test over **homogeneous** hardware
- Only test **GNU C/C++ compiler**, no Fortran
- Package built for **x64**, `aarch64`, `ppc64le`, but only tested `x64`

Upstream and downstream release schedule

| RHEL release schedule | | | | |
|-----------------------|------------------|---------------------|-----------|---------|
| Plan | Feature requests | | | |
| Devel | | Package/code rebase | Exception | Blocker |
| QE | | | | |



| release | Date |
|---------|------------|
| 2.3.6 | 05/11/2021 |
| 2.3.5 | 11/30/2020 |
| 2.3.4 | 06/01/2020 |

- Upstream release interval is about 5+ months
- Steady and predicable upstream release schedule is very important

Upstream SVN access

<https://scm.mvapich.cse.ohio-state.edu/svn/mpich/mvapich2/trunk/>

| Revision | Author | Commit Message | Time | File Count | Path |
|----------|----------|---------------------------|---------------------|------------|-------|
| 8626 | subramon | Update release number | 05/12/2021 03:55 AM | 1 | trunk |
| 8625 | subramon | MVAPICH2 2.3.6 GA Release | 05/12/2021 03:32 AM | 1345 | trunk |
| 8624 | subramon | MVAPICH2 v2.3.5 release | 12/01/2020 06:32 AM | 2975 | trunk |
| 8623 | subramon | MVAPICH2 v2.3.3 release | 01/11/2020 02:06 AM | 239 | trunk |
| 8622 | subramon | MVAPICH2 2.3.2 GA Release | 08/10/2019 03:53 AM | 437 | trunk |
| 8621 | subramon | Delete file | 03/02/2019 05:49 AM | 1 | trunk |
| 8620 | subramon | MVAPICH2 2.3.1GA Release | 03/02/2019 05:48 AM | 1268 | trunk |
| 8619 | subramon | MVAPICH2 2.3 release | 07/25/2018 11:51 PM | 1165 | trunk |

No access to MVAPICH2 commits history, it's impossible to narrow down regression issues with 'git bisect'. mpich and openmpi provides git repo access for public.

Problem for major RHEL release

- Because MVAPICH2 is unavailable in Fedora, MVAPICH2 always had been added as RHEL only/specific package for major release like RHEL-6.0, 7.0, 8.0 and 9.0.
 - Compilation security harden
 - SeLinux harden
 - Block CI testing
- Mpich and openmpi are inherited from corresponding Fedora releases.