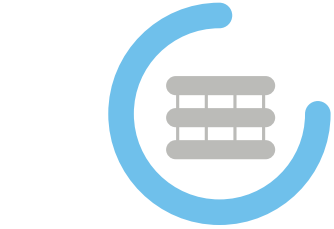

Meeting of the Technical Steering Committee (TSC) Board

Wednesday, November 03rd, 2021
11:00am ET

Antitrust Policy Notice

- Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.

- Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at http://www.linuxfoundation.org/antitrust-policy. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrove of the firm of Gesmer Updegrove LLP, which provides legal counsel to the Linux Foundation.

THE LINUX FOUNDATION

Agenda/Updates

- Announcements, upcoming talks and deadlines
  - SC'21 BoF – accepted
    - Wednesday, 17 November 12:15pm - 1:15pm CST
    - online (zoom) session
  - RedHat mini theater session
    - Wednesday, 17 November 10:15am - 10:50am CST

- hwloc updates
- Priority for SC release?

hwloc update

- Recall from earlier this summer the issue we encountered with hwloc changing .so numbering (first in Stream, then in CentOS 8.4)
  - broke slurm and openpbs builds
  - although we provide a user-oriented version of hwloc with modules, we were using system provided versions of hwloc for the resource managers to that daemons would fire up properly (ie, can't run "module load hwloc")
  - at the time, we talked about a desire to update our approach
    - tweak ohpc-build of hwloc so that it is installed into a fixed path (makes it more of admin package, could not support multiply installed versions)
    - update slurm and openpbs builds to use the ohpc variant of hwloc

```
+ yum -y --installroot=/opt/ohpc/admin/images/centos8.3 install ohpc-slurm-client
Last metadata expiration check: 0:00:21 ago on Tue 15 Jun 2021 06:45:36 PM CDT.
Error:
 Problem: conflicting requests
  - package ohpc-slurm-client-2.0-47.1.ohpc.2.0.x86_64 requires slurm-slurmd-ohpc, but none of the providers can be installed
  - package ohpc-slurm-client-2.1-5.1.ohpc.2.1.x86_64 requires slurm-slurmd-ohpc, but none of the providers can be installed
  - package ohpc-slurm-client-2.0-47.1.ohpc.2.0.aarch64 does not have a compatible architecture
  - nothing provides libhwloc.so.5()(64bit) needed by slurm-slurmd-ohpc-20.02.5-13.1.ohpc.2.0.x86_64
  - package ohpc-slurm-client-2.1-5.1.ohpc.2.1.aarch64 does not have a compatible architecture
  - nothing provides libhwloc.so.5()(64bit) needed by slurm-slurmd-ohpc-20.11.7-3.1.ohpc.2.2.x86_64
(try to add '--skip-broken' to skip uninstallable packages)
```

issues caused by CentOS8.4

- thought I would try to knock this out real quick for 2.4 release
  - gotchas encountered when verifying that hwloc-libs dependencies from OS were indeed removed

hwloc update (cont.)

- Both SLURM and OpenPBS provide configure options to point to alternate hwloc install
  - SLURM could build against our ohpc version with no problem
  - OpenPBS had issues
    - their "--with-hwloc" option for an alternate path is hardcoded to look for a static library that is not named libhwloc (libhwloc_embedded.a ??)
    - had to patch the `with_hwloc.m4` macro file to check for normal libhwloc.so instead
  - So, updating the builds initially was not too bad
- Problem was that the builds still added rpm dependency on libhwloc.so provided by OS

hwloc update (cont.)

- Both SLURM and OpenPBS provide configure options to point to alternate hwloc install
  - SLURM built against our ohpc version with no problem
  - OpenPBS had issues:
    - their "--with-hwloc" option for an alternate path is hardcoded to look for a static library that is not named libhwloc (libhwloc_embedded.a ??)
    - had to patch the `with_hwloc.m4` macro file to check for normal libhwloc.so instead
  - So, updating the builds initially was not too bad

- Problem was that the resulting builds still added rpm dependency on libhwloc.so provided by OS
  - fix here is more subtle and complicated
  - interaction with the (ohpc) coloring we introduced in 2.x using RPM plugin infrastructure to isolate ohpc-provided libraries from similar OS names

```
# rpm -q --provides hwloc-ohpc
hwloc-ohpc = 2.5.0-3.1.ohpc.2.4
hwloc-ohpc(x86-64) = 2.5.0-3.1.ohpc.2.4
libhwloc.so.15()(64bit)(ohpc)
```

hwloc update (cont.)

- Gotcha #1 is that we were only applying our dependency magic on ELF binaries that were installed into an /opt/ohpc path
  - SLURM and OpenPBS builds do not do that
  - ohpc RPM plugin was not exercised and we have to explicitly identify paths where we want to enable our plugin
  - have updated the plugin config path (owned by ohpc-buildroot) as follows:

```
%__ohpc_provides        /usr/lib/rpm/ohpc-find-provides
%__ohpc_requires        /usr/lib/rpm/ohpc-find-requires %{buildroot} %{OHPC_HOME}

%__ohpc_path            ^(%{OHPC_HOME})|(/usr/sbin/slurm.*)|(/opt/pbs/sbin/pbs.*)$
%__elf_exclude_path     ^(%{OHPC_HOME})|(/usr/sbin/slurm.*)|(/opt/pbs/sbin/pbs.*)$

%__ohpc_magic           ^ELF (32|64)-bit.*$
%__ohpc_flags           magic_and_path
```

  - seemed like the only choice unless we want to use are ohpc analyzer on all files

hwloc update (cont.)

- Changing the config path resolved dependency issues for openpbs
- Gotha #2
  - canonical SLURM build uses their own old-style dependency generator (so no plugin support)
  - had to disable these bits in our slurm build

```
cat > find-requires.sh <<'EOF'
exec %{__find_requires} "$@" | egrep -v '^libpmix.so|libevent|libnvidia-ml'
EOF
chmod +x find-requires.sh
global _use_internal_dependency_generator 0
global __find_requires %{_builddir}/%{buildsubdir}/find-requires.sh
```

hwloc update (cont.)

- If we can live with those changes, resulting dependencies now look good and rely on hwloc-ohpc

```
# rpm -q openpbs-server-ohpc --requires | grep hwloc
hwloc-ohpc
libhwloc.so.15()(64bit)(ohpc)
```

OpenPBS

```
# rpm -q slurm-slurmd-ohpc --requires | grep hwloc
libhwloc.so.15()(64bit)(ohpc)
```

SLURM

- SLURM tests running Rocky8.4 run fine, but CI does show an issue with OpenPBS
  - pbs_mom not resolving hwloc
  - note: slurm build automatically rpaths the external hwloc location
  - likely need another mod here for openpbs build to also rpath the external hwloc location
    - or, could add to ld.so.conf? (

```
# ldd /opt/pbs/sbin/pbs_mom
linux-vdso.so.1 (0x00007fff4c91c000)
libhwloc.so.15 => not found
...

# rpm -q hwloc-ohpc
hwloc-ohpc-2.5.0-3.1.ohpc.2.4.x86_64
```

Priorities for SC'21 efforts

- Obviously we have a BoF, presume the primary topic of interest is our treatment going forward with CentOS8 going away
- Unlikely to be able to get to everything we had hoped for a v2.4 and v1.3.10 release - need to prioritize or hold off
  - note that we do have people using Rocky already
  - but 2.3 only has 1 Rocky recipe (x86/WW/slurm)
- Top-level decision?
  1. have at least 1 release for SC'21 (and prioritize what we can do)
  2. hold off and get more done for later this year or Q1 2022
- If (1), input on items to prioritize for v2.4?
  - (high) finalize hwloc/rms changes
  - (high) convert all CentOS8 recipes to Rocky8 and test (this would enable aarch64)
  - (high) Lustre client on Rocky8 (have a build of this now, still no support for Leap 15.2)
  - (high|medium) Leap 15.3
  - (medium) Updated vendor compiler support (Arm and/or Intel oneAPI)
  - (medium) more component version updates
  - (low) updated PMIx support with RMS(s)
  - (medium|low) Warewulf4....