

Cray Support of the MPICH ABI Compatibility Initiative

Steve Oyanagi
2/17/2015

Introduction

The goal of the MPICH ABI Compatibility Initiative is to provide and maintain ABI (Application Binary Interface) compatibility amongst the MPI implementations provided by the initiative members. Cray is a member of the MPICH ABI Compatibility Initiative along with Intel, IBM, and Argonne National Laboratory's MPICH development group (ANL MPICH). Cray, Intel, and IBM's MPI implementations are all derived from the ANL MPICH implementation.

ABI compatibility allows dynamically linked applications built with one ABI-compatible MPI to use a different ABI-compatible MPI at run time. The MPICH ABI Compatibility Initiative defines initiative compliance to be ABI compatibility by conforming to a specific ABI plus additional changes that allow ABI compatibility to be easily used. These additional changes include requiring all implementations have the same dynamic shared library names and version numbers. Further details on initiative compliance can be found on ANL's ABI Compatibility Initiative WIKI page¹.

ANL MPICH 3.1.1 was the first ANL MPICH release to achieve full compliance with the MPICH ABI Compatibility Initiative. The other initiative member MPI implementations have become compliant with the initiative by transitioning their MPI implementation to be based on ANL MPICH 3.1.1 or later.

Cray MPICH has transitioned to be based on ANL MPICH 3.1.2 but is not fully compliant with the MPICH ABI Compatibility Initiative. Cray supports the execution of applications built with an MPICH ABI Compatibility Initiative MPI implementation on a Cray. Building applications on Cray systems to be run with other Initiative compliant MPI implementations is not supported.

It should be noted that the IBM MPI implementation included in the MPICH ABI Compatibility Initiative is not Platform MPI. At this time there is no indication that Platform MPI will become part of the MPICH ABI Compatibility Initiative. The IBM MPI implementation included in the initiative currently supports POWER. It is unclear if this version of IBM MPI supports the x86 architecture.

¹ http://wiki.mpich.org/mpich/index.php/ABI_Compatibility_Initiative

Current Status of Cray with the MPICH ABI Compatibility Initiative

Cray MPICH ABI Compatibility Initiative support for running applications built with other initiative MPI implementations on Cray systems has many dependencies and restrictions that the users should be aware of. The following section discusses these dependencies and limitation.

Cray MPICH ABI compatibility has multiple dependencies. Obviously the build target CPU architecture must match the architecture that the application is run on. Be aware that within architectures newer chip families often have new instructions that will not run on older chip families. The build target OS must be compatible with the OS that the application will run on. The compiler used to build the application may require libraries that would have to be present when the application runs.

Cray MPICH ABI compatibility only supports applications built with the Intel and GNU compilers. Please consult the appropriate Cray Programming Environments Release Announcement² to determine what GNU and Intel compiler versions are compatible with a particular MPT release. An application using ABI compatibility to run with Cray MPICH must be built with a GNU or Intel compiler at a version compatible with that Cray MPICH/MPT release. Previous GNU or Intel compiler versions may be partially or completely compatible with the listed compiler version for an MPT release. Please see the appropriate GNU or Intel documentation for that information.

Cray MPICH ABI compatibility depends on OS compatibility between the target OS the application was built for and the Cray CLE version that the application is to be run on. CLE is a based on the SUSE SLES Linux distribution. Please consult the appropriate CLE Software Release Overview³ to determine what version of SLES the target CLE release is based on, and the appropriate SLES documentation to determine what other OS distributions and releases a particular SLES release is compatible with.

Cray MPICH ABI Compatibility is intended for use on Cray XE/XC systems in Extreme Scaling Mode (ESM). This feature is not intended for Cray Compatibility Mode (CCM) and should not be used there.

Full compliance with the MPICH ABI Compatibility Initiative requires dynamic shared library name and version number changes. These changes would cause downward ABI incompatibilities with previous MPT versions and existing Cray PE software built with previous MPT releases. Cray has chosen to maintain its current library names and version numbers to avoid incompatibilities with existing Cray PE software and to allow users to switch back to older MPT versions. Cray's current dynamic shared library names and version numbers are supported by the existing "cray-mpich" module. Support

² Programming Environment Release Announcements can be found on the Cray documentation web site <http://docs.cray.com/>

³ Cray Linux Environment (CLE) Software Release Overview documents can be found on the Cray documentation web site <http://docs.cray.com/>

for ABI compatibility is provided by the additional module "cray-mpich-abi". The "cray-mpich-abi" module should be loaded instead of the existing "cray-mpich" module to use ABI compatibility.

MPICH ABI Compatibility Support in MPT 7.1.3 and beyond

To use ABI Compatibility the application must conform to the following requirements:

1. The application must be built with an Intel or GNU compiler that is compatible with the Cray MPT release the application will be linked to when the application is run.
2. The application must be built with an MPI implementation that is ABI compatible with ANL MPICH 3.1.1.
3. The application must be linked dynamically to the MPI libraries.

To run the application using Cray MPICH on a Cray XE or XC system using MPT 7.1.3 or later:

1. Make sure that either the Intel or GNU Programming Environments (PrgEnv-intel or PrgEnv-gnu) is loaded
2. module swap cray-mpich cray-mpich-abi
3. Prepend the environment variable CRAY_LD_LIBRARY_PATH to the environment variable LD_LIBRARY_PATH.
tcsh/csh: setenv LD_LIBRARY_PATH
 \${CRAY_LD_LIBRARY_PATH}:\$LD_LIBRARY_PATH
bash/sh/ksh: export LD_LIBRARY_PATH=
 \${CRAY_LD_LIBRARY_PATH}:\$LD_LIBRARY_PATH
4. Launch the application using aprun.

Intel MPI Compatibility Details

Intel released their final MPI 5.0 late in 2014, allowing us to complete our ABI Compatibility support for Intel MPI 5.0. Beginning with MPT 7.1.3 Cray MPICH supports ABI compatibility with Intel MPI 5.0 and ANL MPICH 3.1.1 and newer releases. At this time Cray supports only the execution of applications built with Intel MPI 5.0 and Intel or GNU compilers on Cray XE/XC/XK systems in ESM using MPT 7.1.3.

Compatibility With Previous Intel MPI Versions

Cray MPICH ABI compatibility with previous Intel MPI releases depends on Intel MPI 5.0's ABI compatibility with those previous versions. Intel MPI 5.0 is mostly upwardly compatible with several previous Intel MPI releases. Many applications built with older versions of Intel MPI should run with Cray MPICH. Users should consult the release

notes for Intel MPI 5.0 to understand where previous Intel MPI versions are incompatible with Intel MPI 5.0, and thus also be incompatible with the Cray MPICH. Please note that Intel does specify incompatibilities between Intel MPI 5.0 and previous versions of Intel MPI. For convenience Cray has provided support in the “cray-mpich-abi” module to allow users to attempt to run Intel MPI 4.1, 4.0, and 3.2 applications with Cray MPICH.

Intel Xeon Phi Support

Instructions for Intel Xeon Phi autonomous mode are the same as above. Be sure to include the appropriate options for running an Intel Xeon Phi application in autonomous mode on the aprun command. Please note that only Intel MPI 4.1 and later support Intel Xeon Phi.

MPICH ABI Compatibility Support in MPT Releases Before MPT 7.1.3

Cray MPICH became ABI compatible with MPICH 3.1.1 and other ABI Compatibility Initiative MPI implementations beginning with MPT 7.0.0 released in June of 2014. MPT releases prior to MPT 7.1.3 provided ABI compatibility support only for applications built with the Intel compiler and Intel MPI 4.1 or 4.0. Please note that the Cray MPICH included in these earlier MPT releases is ABI compatible with Intel MPI 5.0 and not Intel MPI 4.1 or 4.0. Users should consult the Intel MPI 5.0 release notes to understand where earlier Intel MPI releases are not ABI compatible with Intel MPI 5.0, and thus incompatible with Cray MPICH. For convenience Cray provided support in the “cray-mpich-abi” module to allow users to attempt to run Intel MPI 4.1 and 4.0 applications with Cray MPICH.

Running Applications Built with the Intel compiler and Intel MPI Using Cray MPICH on Cray XC/XE Systems Including Intel Xeon Phi™ Autonomous Mode

To use ABI Compatibility the application must conform to the following requirements:

1. The application must be built with Intel compiler version 14.0 or a compatible Intel compiler version. Please see the Intel compiler documentation for what Intel compiler versions are compatible with Intel 14.0.
2. The application must be built with Intel MPI 4.0 or 4.1. For Intel Xeon Phi only Intel MPI 4.1 is supported. Intel MPI 4.0 does not support Intel Xeon Phi.
3. The application must be dynamically linked to the Intel MPI libraries.

To run the application using Cray MPICH on a Cray XE or XC system using MPT 7.0.3 through MPT 7.1.2:

1. Make sure that the Intel Programming Environment (PrgEnv-intel) is loaded
2. module swap cray-mpich cray-mpich-abi

3. Prepend the environment variable `CRAY_LD_LIBRARY_PATH` to the environment variable `LD_LIBRARY_PATH`.
tcsh/csh: `setenv LD_LIBRARY_PATH
 ${CRAY_LD_LIBRARY_PATH}:${LD_LIBRARY_PATH}`
bash/sh/ksh: `export LD_LIBRARY_PATH=
 ${CRAY_LD_LIBRARY_PATH}:${LD_LIBRARY_PATH}`
4. Launch the application using `aprun`.

To run the application using Cray MPICH on a Cray XE or XC system using MPT 7.0.0, MPT 7.0.1, or MPT 7.0.2:

1. Make sure that the Intel Programming Environment (PrgEnv-intel) is loaded.
2. Create a directory in a file system accessible from the compute nodes and change directory into it.
3. Create the following symbolic links in the directory:
 - `ln -s /opt/cray/mpt/7.0.x/gni/mpich2-intel/140/lib/libmpich_intel.so.3
libmpi.so.4`
 - `ln -s libmpi.so.4 libmpi_mt.so.4`
 - `ln -s /opt/cray/mpt/7.0.x/gni/mpich2-intel/140/lib/libmpichcxx_intel.so.3
libgc4.so.4`
 - `ln -s /opt/cray/mpt/7.0.0/gni/mpich2-intel/140/lib/libmpichf90_intel.so.3
libigf.so.4`
4. Prepend the directory containing the symbolic links to `LD_LIBRARY_PATH`
5. Launch the application using `aprun`.

Instructions for Intel Xeon Phi autonomous are the same as above. Be sure to use the appropriate options for running an Intel Xeon Phi application in autonomous mode on the `aprun` command.

Future Plans

Cray is investigating adding support to the `cray-mpich-abi` module for building applications on Cray systems that could be run with other ABI Compatibility Initiative MPI implementations.