

Intel Omni-Path Fabric Performance Tuning

With Intel MPI 2017, Update 3 or newer, the default for `I_MPI_FABRICS` changed from `shm:tmi` to `shm:ofi`. **`I_MPI_FABRICS=shm:ofi`** implies the OpenFabrics Interface (OFI) library will be used for inter-node communications, and the following settings have been added to the Intel MPI 2017 and 2018 environments via modulefile:

```
export I_MPI_FABRICS=shm:ofi
export I_MPI_OFI_PROVIDER=psm2
```

With the Intel Omni-Path Software 10.5 release (the current installed version on Marconi), the OFI library will be included in the Basic and IFS packages.

This change of default could provide an unexpected change in performance from the previous Intel Omni-Path Software and Intel MPI releases. The change to OFI over PSM2 could provide better or worse performance than the previous default, the TMI library over PSM2 that is provided with the Intel MPI package.

You can always switch to the TMI library by setting the following variables in your submission script, after loading the `intelmpi` module:

```
export I_MPI_FABRICS=shm:tmi
export I_MPI_TMI_PROVIDER=psm2
```

Please refer to the official [Intel Omni-Path Fabric Performance Tuning Guide](#) for additional details and optimization tips.