

Multicore Software Development Kit for High-Performance Computing

Release Notes

Applies to Release: 03.00.01

Publication Date:

Nov 14, 2014

Document License

This work is licensed under the Creative Commons Attribution-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Copyright (C) 2012-2014 Texas Instruments Incorporated - <http://www.ti.com>

Contents

1	Overview	1
2	Licensing.....	1
3	Documentation.....	1
4	Platform and Device Support	1
5	Technical Support and Product Updates.....	1
6	Releases.....	3
6.1	Components/Versions packaged in this release	3
6.2	Release Features	4

MCSDK-HPC Release 3.0.1

1 Overview

This document is the Release Notes for Release 3.0.x of the Multicore Software Development Kit for High-Performance Computing (MCSDK-HPC). MCSDK-HPC is built as an add-on on top of the foundational Multicore Software Development Kit (MCSDK). MCSDK-HPC, along with MCSDK, provides the complete environment to develop HPC applications on a cluster of TI Keystone II devices.

2 Licensing

Please refer to the software manifest which outlines the licensing status for all packages included in this release.

3 Documentation

- **README:** Provides information on setting up the environment to compile the provided components and build out-of-box demos.
- **MCSDK-HPC Software Manifest:** Provides license information on software included in the MCSDK-HPC release. This document can be found in the release at [INSTALL-DIR]/mcsdk-hpc_<ver>/docs.

4 Platform and Device Support

The device and platforms supported with this release include:

Platform	Supported Devices	Supported EVM
[K2H]	66AK2H12	XTCIEVMK2X

5 Technical Support and Product Updates

For technical discussions and issues, please visit:

- [KeyStone Multicore forum](#)
- [BIOS Embedded Software forum](#)
- [Linux Embedded Software forum](#)
- [Code Composer Studio forum](#)

- [TI C/C++ Compiler forum](#)
- [Embedded Processors wiki](#)

For local support in China, please visit

- [China Support forum](#)

6 Releases

6.1 Components/Versions packaged in this release

Component	Version
framework_components	03.31.00.02
openmp_dsp	02.01.16.02
ti-openmpi	01.00.00.14
xdais	07.24.00.04
fftlib	02.00.01.03
ti-opencl	00.15.00
blas	03.11.00.03
ti-openmpacc	01.01.01
gdbc6x	1.0.0
gdbserver-c6x	1.0.0
cmem, cmem-mod-dkms	4.00.03.14
debugss-mod-dkms	1.0.1
dsptop	1.1.5
gdb-proxy-mod-dkms	1.0.1
mpm	2.0.1.6
mpm-transport	1.0.5.2
temperature-mod-dkms	1.0.1
ti-llvm-3.3	3.3
xdc-headers	3.25.05.94
ti-cgt-c6x	8.0.0
uio-module-drv-dkms	01.00.00.00
edma3-ld	02.11.11.15

6.2 Release Features

6.2.1 DEV.MCSDK_HPC.03.00.01.05

Changes from previous release:

- OpenMPI
 - Use PDSP based packet forwarding, instead of SRIO hardware IP based packet forwarding
 - This brings in significant infrastructure under MPI
 1. Navigator/Queue Management Sub System
 2. Resource Manager: RM server daemon started as part of Upstart Job
 3. PDSP firmware
 4. SRIO type-11 wrapper for interfacing with SRIO BTL in MPI
- MPM-transport
 - Robustness in Hyperlink SERDES start-up sequence
 - Hyperlink @ 6.25 Gbps per lane support
 - Eliminate /dev/mem accesses to restore non-root user access for OpenMPI
- OpenCL
 - Fix an MPAX allocation issue where subbuffers >= 1Gb could not be used
 - Fix issue compiling opencl kernels with barriers
 - Fix a dsptop issue by adding a ulm termination call
 - Bug fix to work with PyOpenCL (Python OpenCL)
- MPM
 - Fix incorrect assumptions on installation order of uio_module_drv and the dsps. This was causing OpenCL to hang on certain reboots
- Update udev rule with permission 0660 and group keystone-hpc
- Additional 256 KB is reserved in MSMC for PDSP based SRIO routing
- Bug Fix in EDMAMgr (FC) in handling large strides in 3D transfers
- OpenMP accelerator: Install examples to EVM file system
- Debian packaging updates

6.2.2 DEV.MCSDK_HPC.03.00.01.04

Changes from previous release:

- MPM-transport
 - Reliable SERDES start-up sequence for Hyperlink
 - Implement timeout for Hyperlink port open, so that Hyperlink can be used only when 2 nodes are powered ON
- OpenMPI
 - Use CPU initiated transfers, instead of EDMA, for short messages over Hyperlink. This is for performance improvements
 - Extend support for 60 nodes interconnected via SRIO
 - Enhancement to do SRIO enumeration as part of MPI start-up
 - Bump up priority on NREAD requests (so they don't get timed out competing with packet forwarding data traffic)
 - Topology/Routing changes to avoid loops in a cluster (because, there'll be a LSU lock up with an SRIO loop). **NOTE:** After HPC installation is done, the file `/etc/cluster/srio_topology.bin` needs to be regenerated, as per the instructions [here](#)
 - Packet forwarding fixes to exclude a node's own id, and to have no packet forwarding entries for end point nodes

- Temporary workaround for LSU errors: Software throttling of SRIO traffic. Will be removed in next drop
- EDMAMgr (FC) update to handle large strides in 3D transfers
- Debian packaging updates
- **NOTE:** Only Root users can run MPI (MPM Transport users /dev/mem to poke hyperlink PSC registers). This restriction will be removed in future release

6.2.3 DEV.MCSDK_HPC.03.00.01.03

Changes from previous release:

- Addressed Intel MPI Benchmarking Test issues with Hyperlink
 - Random hangs in 30min to 3 hours with 2 nodes
 - Fails in all_gatherv w/ 128K buffer size with 4 nodes
- Hyperlink SERDES initialization moved from HPC upstart job to OpenMPI
- OpenMPI: Fix for Cartridge ID = 45 with BTL SRIO
- DSP Compiler (ti-cgt-c6x) upgraded to updated to 8.0.0 GA
- MPM Transport: Hyperlink SERDES initialization sequence changes and boost workaround
- DSPTOP: Documentation updates and changes to eliminate dsptop_sync file extension and source permissions
- OpenMP Accelerator:
 - Addition of host versions of DSP functions for cache and dynamic memory Management and a dspheap example
 - Bug Fix for memory leak in implementation for target update construct
- OpenCL:
 - Performance improvements of Improve kernels with barriers and those which uses 'reqd_work_group_size' attribute.
 - Bug fixes for clFlush() semantics and deadlock in event callback functions.
- FFTLIB: Fix for multicore overheads with OpenMP

6.2.4 DEV.MCSDK_HPC.03.00.01.02

Changes from previous release:

- New components included in MCSDK-HPC.
 - dsptop: This provides functionality similar to Linux top utility, i.e. visibility into usage data for TI DSP + ARM KeyStone II system-on-chips such as the 66AK2Hxx family. For more details, please see online [wiki documentation](#)
 - gdb6x: This allows developers to utilize the standard features of GDB (GNU Debugger) to gain visibility to and debug the C66x DSP cores in the heterogeneous DSP + ARM KeyStone II system-on-chips such as the 66AK2Hxx family. For more details, please see online [wiki documentation](#)
 - Supporting Kernel and user space modules for dsptop and gdb6x (debugss-mod, temperature-mod, gdbproxy-mod, gdbserver)
- OpenCL: Integration with dsptop, gdbc6x components
- OpenCL: Dangling pointer fix on the sgemm example
- OpenCL: Various bug fixes exposed by conformance testing
- OpenMP accelerator: GA release with OpenMP 4.0 target constructs
- Debian packaging updates in preparation for Partner Migration
- DSP Compiler (ti-cgt-c6x) upgraded to Beta-4 release.
- OpenMPI: Fixed the segmentation fault during SRIO routing simulation
- MPM: Updates for compatibility with Linux kernel 3.13
- MPM Transport: Serdes sequence updates for Hyperlink

- Framework Components: Updates to add support for 3D-edma transfers

6.2.5 DEV.MCSDK_HPC.03.00.01.01

Changes from previous release:

- Publicly available/ freely distributable C66x DSP Code Generation Tools (CGT) are provided with HPC installation (Users no longer need to download CGT separately)
- MPI: 6.25 Gbps over Hyperlink (needs new DTS, and updated MPM, MPM-transport components); Known Issues from 3.0.1.0 with OpenMPI+OpenCL and OpenMPI+OpenMPacc are addressed.
- OpenCL: Added opencl/openmp examples: sgemm, dgemm, fftlib_offload
- OpenCL: Separate LLVM into a separate package
- FFTLIB C66x Library updates to support OpenCL dispatch
- Several building blocks to support DSPTOP and SocTune are now available with MCSDK-HPC. Integration of debug features into OpenCL and collateral updates will happen as part of next release

6.2.6 DEV.MCSDK_HPC.03.00.01.00

Changes from previous release:

OpenCL

- Support for all clauses of target construct in OpenMPacc (SDSCM00050116)
- Added host memory allocation functions `__malloc_ddr()`, `__free_ddr()`, `__malloc_msmc()`, and `__free_msmc()`. These functions have the same interface as standard malloc/free but manages memory from host+dsp shareable memory. These can be used within the host application
- Added dsp memory allocation functions `__malloc_ddr()`, `__free_ddr()`, `__malloc_msmc()`, `__free_msmc()`, `__malloc_l2()`, and `__free_l2()`. These functions are available from dsp kernel code to manage memory from the dsp heaps. The OpenCL example 'dspheap' has been added to illustrate how the dsp heap memory can be allocated and managed between kernel invocations.
- Significantly improved OpenCL performance on cache coherency operations for buffer sizes of 4MB and above.
- The `async_workgroup_copy` and `async_workgroup_strided_copy` functions are now implemented to use EDMA channels when available. (SDSCM00050115)
- The number/size of Kernel arguments are now fully compliant with the OpenCL 1.1 spec. Up to 1024 bytes of arguments are supported and vector types can be passed as standalone arguments.
- On the dsp, the standard malloc() function will now allocate from an 8M pool of DDR.
- The OpenCL package version number can now be obtained by either invoking 'clocl --version' or making a standard OpenCL platform query. The OpenCL 'platform' example illustrates how to make an OpenCL platform query.
- The opencl runtime provides built in functions that would allow a user to explicitly change the L2 cache size. Since this call would occur in the DSP code, the user would need to ensure that the app did not allocate I2 SRAM that would conflict with their I2 cache size choice. (SDSCM00050114)
- Setting the environment variable `TI_OCL_CGT_INSTALL` is not required if the C6000 compiler runtime library and header files are located in the ppa package default install path. `TI_OCL_CGT_INSTALL` can still be used to specify an alternate search path.
- The environment variable used to enable OpenCL kernel debugging has been changed from `TI_OCL_DEBUG_KERNEL` to `TI_OCL_DEBUG`.

MPI

- Latency and Bandwidth Improvements with SRIO (SDOCM00108189)
 - Latency: 4uS, BW: up to 240MB/s (@5gbps)
 1. SRIO user space module for RMA

- 2. CMEM cached allocated message buffers
- Latency and Bandwidth Improvements with Hyperlink
 - Latency: 5.5uS, BW, up to 420MB/s (@4x3.125gbps)
 1. CMEM cached allocated message buffers
 2. EDMA linked transfers for messages >4Kbytes
 3. Reduced EDMA setup time w/ preallocated EDMA resources
- Sudo privilege was needed for using MPI over SRIO. Now, it is non-root users can also use MPI over SRIO (SDOCM00107864)
- MPI using SRIO was broken for cartridges enumerated above 7 (c7) (SDOCM00107865)
- Reliable SRIO mports don't come up (sometimes) when all the nodes are not booted at once on Slayton (SDOCM00107669)
- MPI latency measurement triggering memory leakage (BTL SRIO & Hyperlink) (SDOCM00107945)
- **KNOWN ISSUE:**
OPENMPI runtime by itself works but when used with OpenCL or OpenMPacc doesn't work. However, there is a work around for examples involving Hyperlink interchip connection. The work around is as below.
 - Vecadd_mpi example (OpenMPI+OpenMPacc), should use additional mca parameter as temporary workaround:


```
/opt/ti-openmpi/bin/mpirun -x LD_LIBRARY_PATH -x TI_OCL_INSTALL -x TI_OCL_CGT_INSTALL -x TI_OMPACC_INSTALL --mca btl_hlink_use_edma 0 --mca btl self,hlink -np 2 -host c7n1,c7n2 vecadd_mpi
```
 - Multimode_batch_fftdemo example (OpenMPI+OpenCL),should use additional mca parameter as temporary workaround:


```
/opt/ti-openmpi/bin/mpirun --mca btl_hlink_use_edma 0 --mca btl self,hlink -np 2 -x LD_LIBRARY_PATH -x PATH -x OPAL_PREFIX -host c7n1,c7n2 ../host/demo_multinode_batch_fft ../testfiles/fftw_testfile.inp ../testfiles/fftOut.bin
```

These issues will be addressed in upcoming release.

FFTLIB

- Provide pre-built Library, along with Makefiles for Library and Unit-tests, in addition to CCS projects (SDOCM00107667)

OpenMP-DSP

- Added dynamic heap management APIs on the DSP

Debian Security Violations

- Change CMEM auto-insertion to an upstart job, instead of modifying /etc/modules (modifying /etc/modules is a debian security violation) (SDOCM00107905)

6.2.7 DEV.MCSDK_HPC.03.00.00.19

Changes from previous release:

- Requires MCSDK 3.0.4.18 (GA). Incompatible with previous MCSDK releases.
- All bundled components promoted to GA status (except for OpenMPacc)
- OpenMPI over SRIO with packet forwarding for switch-less topologies
- Default Hyperlink bandwidth increased from 3.125 to 6.25 Gbps
- Support to natively compile all out of box examples on EVM
- All out of box examples (opencl, openmp, openmpi, hpc) have been consolidated and installed to /usr/share/ti/examples

- mpi_examples present as separate package in previous release have been removed
- Dropbear patch has been removed
- Additional documentation and collateral on wiki pages
- HPC additions to the file system are delivered as IPKs and via apt-get

6.2.8 DEV.MCSDK_HPC.03.00.00.18

- First public release (Beta)
- Release bundles the following components:
 - OpenMP DSP runtime.
 - OpenCL
 - OpenMP Accelerator Model
 - OpenMPI w/ SRIO and hyperlink support
 - FFTLIB
 - BLAS
 - Framework Components with EdmaMgr support
 - Examples exercising combinations of each of the components.
- More details available at the MCSDK HPC Getting Started Guide (http://processors.wiki.ti.com/index.php/MCSDK_HPC_3.x_Getting_Started_Guide)