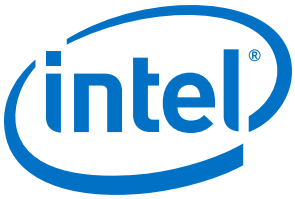

# Intel® FPGA SDK for OpenCL™ Pro Edition

## Version 19.3 Release Notes

Updated for Intel® Quartus® Prime Design Suite: **19.3**



RN-OCL004 | 2019.09.30
Latest document on the web: PDF | HTML

# Contents

# 1. Intel® FPGA SDK for OpenCL™ Pro Edition Version 19.3 Release Notes

The *Intel® FPGA SDK for OpenCL™ Pro Edition Release Notes* provides late-breaking information about the Intel FPGA Software Development Kit (SDK) for OpenCL[1][2] Pro Edition and the Intel FPGA Runtime Environment (RTE) for OpenCL Pro Edition Version 19.3.

## 1.1. New Features and Enhancements

The Intel FPGA SDK for OpenCL Pro Edition and the Intel FPGA RTE for OpenCL Pro Edition include the following new features:

- Added the Schedule Viewer (alpha) to the high-level design report. The Schedule Viewer displays a static view of the scheduled cycle and latency of kernels, blocks, clusters, and instructions in your design.
- In the high-level design report, enhanced the Graph Viewer to include scheduled cycle and latency information for nodes in the block and cluster views.
- Added *CL_ALLOW_GLOBAL_MEM_AT_NULL_ADDRESS_INTELFPGA* environment variable to enable 100% of the device's memory to be used (including the reserved memory).
- Added support for out-of-order command execution in the OpenCL host runtime command.
- Deprecated `numreadports` and `numwriteports` kernel memory attributes.
- Added `max_replicates(N)` and `simple_dual_port_memory` kernel memory attributes.
- Added `max_interleaving` pragma to limit the number of interleaved invocations of an inner loop that can be executed simultaneously.
- Added `fp contract` and `fp reassoc` pragmas to influence the intermediate rounding and conversions of floating-point operations and the ordering of arithmetic operations in your kernel at finer granularity than the Intel FPGA SDK for OpenCL Offline Compiler command options.
- Added `-ffp-reassoc aoc` command option to direct the Intel FPGA SDK for OpenCL Offline Compiler to relax the order of arithmetic floating-point operations using a balanced tree hardware implementation.

(1) OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission of the Khronos Group™.

(2) The Intel FPGA SDK for OpenCL is based on a published Khronos Specification, and has passed the Khronos Conformance Testing Process. Current conformance status can be found at www.khronos.org/conformance.

- Added the `-ffp-contract=fast` `aoc` command option to direct the Intel FPGA SDK for OpenCL Offline Compiler to remove intermediary floating-point rounding operations and conversions whenever possible, and to carry additional bits to maintain precision.
- Added new built-ins for controlling the styles of load-store units (LSUs) used for accessing the global memory.
- Added `fpga_crossgen` and `fpga_libtool` tools for creating libraries. These tools accept RTL, OpenCL and HLS C++ source code as inputs for `aoc` and `hls` targets. In a future release, the old method of library creation will be deprecated.

## 1.2. Operating System Support

Information about OS support for the Intel FPGA SDK for OpenCL is available on the Operating System Support page of the Intel FPGA website.

**Related Information**

Operating System Support

## 1.3. Changes to Software Behavior

Items listed in the following table represent cases in which the behaviors of the current release of the Intel FPGA SDK for OpenCL and the Intel FPGA RTE for OpenCL differ from the previous version.

| Description | Required Actions |
|---|---|
| By default, the OpenCL runtime reserves a small portion of the memory at the start of the device's memory to facilitate `NULL`-pointer checking. This ensures that valid global memory pointers always have a non-zero value and allows a zero value to be interpreted as a `NULL` pointer. | To enable 100% of the device's memory to be used, including the reserved memory, set the `CL_ALLOW_GLOBAL_MEM_AT_NULL_ADDRESS_INTEL` FPGA environment variable to 1.<br>If the `CL_ALLOW_GLOBAL_MEM_AT_NULL_ADDRESS_INTEL` FPGA environment variable is set to 1, a valid global memory pointer can be 0 when observed by the kernel. In other words, you cannot reliably check inside the kernel that a `NULL` global memory pointer is invalid. |
| (**Linux-specific**)The emulator in Intel FPGA SDK for OpenCL Pro Edition version 19.3 is built with GCC 6.3.0 as part of the offline compiler. When executing the host program for an emulated OpenCL device, the version of `libstdc++.so` must be at least that of GCC 6.3.0. In other words, the *LD_LIBRARY_PATH* environment variable must ensure that the correct version of `libstdc++.so` is found.<br>If the correct version of `libstdc++.so` is not found, the call to `clGetPlatformIDs` function fails to load the FPGA emulator platform and returns `CL_PLATFORM_NOT_FOUND_KHR (error code -1001)`. Depending on which version of `libstdc++.so` is found, the call to `clGetPlatformIDs` may succeed, but a later call to the `clCreateContext` function may fail with `CL_DEVICE_NOT_AVAILABLE (error code -2)`. | If `LD_LIBRARY_PATH` does not point to a sufficiently new `libstdc++.so`, use the following syntax to invoke the host program:<br>`env LD_LIBRARY_PATH=<path to sufficiently new libstdc++.so>:$LD_LIBRARY_PATH <host> [host arguments]` |
| Streaming and semi-streaming LSUs are no longer inferred by the compiler. | — |

*continued...*

| Description | Required Actions |
|---|---|
| The OpenCL headers provided by the Intel FPGA SDK for OpenCL have been updated to the latest version distributed by Khronos. | Refer to the OpenCL 2.0 Headers in the *Intel FPGA SDK for OpenCL Pro Edition: Programming Guide* for more information. |
| The `cl.hpp` header is now deprecated. | Update your OpenCL host program code to use the `cl2.hpp` header that supports all versions of OpenCL. |
| The OpenCL host runtime now returns `CL_PROFILING_INFO_NOT_AVAILABLE` when you call `clGetEventProfilingInfo` function on a `cl_event` associated with a `cl_command_queue` that was not created with the `CL_QUEUE_PROFILING_ENABLE` property. Previous versions of the Intel FPGA SDK for OpenCL did not enforce this requirement of the OpenCL Specification version 1.2. | Update your OpenCL host program to create the `cl_command_queue` with the `CL_PROFILING_INFO_ENABLE` property if you want to enable profiling. |
| The formatting of `printf` command output of OpenCL vector types has been modified to be conformant with the OpenCL Specification version 1.2. | — |
| Mandatory requirement for all host programs that use Intel FPGA-specific APIs and enumerations. | If your host programs use Intel FPGA-specific APIs and enumerations, such as `clReadPipeIntelFPGA`, `clGetProfileInfoIntelFPGA`, and other APIs that end with `IntelFPGA`, you must explicitly include the `CL/cl_ext_intelfpga.h` header file in your source code. |
| The host pipe feature requires the use of OpenCL 2.0 features. | Enable support for OpenCL 2.0 APIs in all of your host programs that use host pipes in their source code. Refer to the OpenCL 2.0 Headers in the *Intel FPGA SDK for OpenCL Pro Edition: Programming Guide* for more information. |
| In previous versions of Intel FPGA SDK for OpenCL, the offline compiler automatically disabled the local memory replication for all memory systems if it estimated an M20K utilization of over 100%. The offline compiler no longer disables the local memory replication automatically. | To restore the behavior of disabling the local memory replication automatically, use the `-no-local-mem-replication` option in the `aoc` command.<br>Instead, consider using memory attributes to limit memory replication to maintain performance. |
| Automatic loop unrolling is now disabled. | If you want to enable it, use LLVM options `-unroll-threshold` and `-unroll-partial-threshold` in the `aoc` command.<br>Recommended settings:<br>• For optimization levels `-O0`, `-O1`, and `-O2`, set `-unroll-threshold=150` and `-unroll-partial-threshold=150`.<br>• For optimization level `-O3`, set `-unroll-threshold=300` and `-unroll-partial-threshold=150`. |
| Support for passing pipes or channels by reference is removed. | Update your code to pass pipes or channels by value. |
| The `restrict` keyword is renamed to `__restrict`. | Update your code to use the `__restrict` keyword instead of the `restrict` keyword. |

Items listed in the following table represent cases in which the behaviors of the current release of the Intel FPGA SDK for OpenCL Custom Platform Toolkit and Reference Platforms differ from the previous version.

| Description | Workaround |
|---|---|
| N/A | N/A |

## Related Information

OpenCL 2.0 Headers

## 1.4. Known Issues and Workarounds

This section provides information about known issues that affect the Intel FPGA SDK for OpenCL and the Intel FPGA RTE for OpenCL Version 19.3.

| Description | Workaround |
|---|---|
| Block latency is removed from the $f_{MAX}$ II report and system view of the Graph Viewer. | Use the Schedule Viewer to identify the block latency. |
| Certain high-level design reports fail to launch correctly in the Microsoft Internet Explorer* browser. | Intel recommends using one of the following browsers to view the reports:<br>• Google Chrome*<br>• Microsoft Edge*<br>• Mozilla Firefox* |
| When compiling an OpenCL kernel containing calls to library functions containing HLS tasks, incremental compile may trigger recompilation for unaffected kernels. | No known workaround. However, this is not a functional bug. It may result in a more conservative incremental compile. |
| In the fast emulator, the memory used for modeling channel depth scales differently than under the legacy emulator. This means that some designs (that compile for the legacy emulator) fail to compile for the fast emulator and display the following warning and error message:<br>`*Internal compiler error* Unable to allocate section memory!` | If possible, scale down the design or limit the channel depth. If everything else fails, use the legacy emulator. |
| The OpenCL emulator does not return `CL_INVALID_GLOBAL_OFFSET` when a kernel with the `uses_global_work_offset(0)` kernel attribute set is enqueued with a non-zero or non-`NULL` `global_work_offset` argument. | Set the `global_work_offset` argument to `NULL` or zero in all work group dimensions so that this issue does not change the kernel output between the emulator and hardware. |
| OpenCL kernels with names longer than 61 characters might fail in the Intel Quartus® Prime Pro Edition compiler with an error similar to the following error:<br>`Error (16045): Instance "...\|<long_kernel_name>_cra_slave_inst" instantiates undefined entity "<long_kernel_name>_function_cra_slave" File: <filename> Line: <linenumber>` | Reduce the size of the OpenCL kernel name. |
| OpenCL kernel pipes cannot be passed as arguments in some cases. The symptom is the runtime will receive a `CL_INVALID_BUFFER_SIZE (-61)` error when you enqueue your kernel. | Modify your design to use channels instead of pipes. |
| The emulator runtime results in an assertion error if a kernel is enqueued 16,000 times. | Do not enqueue a kernel more than 16,000 times. |
| When alternatively using sub-buffers and their parent buffers, changes written to one might not be reflected in the other. | Unmapping and mapping a buffer forces the sub-buffers and their parent buffers to be synced. Unmapping and mapping a buffer between buffer uses should prevent this issue. |
| In the OpenCL runtime, making more than one OpenCL context in a multithreaded environment might cause a segmentation fault. | — |

This section provides information about known issues that affect the current release of the Intel FPGA SDK for OpenCL Custom Platform Toolkit and Reference Platforms. These issues might also affect Custom Platforms you create for use with the Intel FPGA SDK for OpenCL.

| Description | Workaround |
|---|---|
| Race conditions can occur between enqueue and dequeue buffer operations and host pipe operations. These conditions can results in incorrect data being read or written. | Manually make sure that no enqueue or dequeue buffer operations occur in parallel with host pipe API calls.<br>A way to ensure that buffer operations do not occur in parallel with host pipe API calls is to do buffer operations as blocking calls before the first host pipe operation and after you are certain that the last host pipe operation has been completed (for example, all the data is read back). |
| For Windows, when the host application queries the number of devices, calls to `clGetDeviceIDs` return 128 devices regardless of the actual number of devices present.<br>*Note:* You can find the actual available devices at the beginning of the device list returned by `clGetDeviceIDs`.<br>This issue affects the Intel Arria® 10 GX FPGA Development Kit Reference Platform and the Intel Stratix® 10 GX FPGA Development Kit Reference Platform. | Perform one of the following workarounds:<br>• Rewrite the host application to limit the query for `clGetDeviceIDs` to the actual number of devices.<br>• Rewrite the host application to use `clGetDeviceInfo` to query which devices are available. Calling `clGetDeviceInfo` with the `CL_DEVICE_AVAILABLE` flag correctly reports that extraneous devices are unavailable.<br>• Rewrite the host application to only call `clCreateContext` with the actual number of devices. Calling `clCreateContext` with extraneous devices will fail with the error `CL_DEVICE_NOT_AVAILABLE`.<br>• Set the environment variable *`CL_OVERRIDE_NUM_DEVICES_INTELFPGA`* to the correct number of devices. Doing so fixes the erroneous behavior of `clGetDeviceIDs`. |

### Latest Known Intel FPGA SDK for OpenCL Software Issues

For additional known issue information for the current Intel FPGA SDK for OpenCL version and for previous versions, refer to the Knowledge Base web page.

### Related Information

Knowledge Base

## 1.5. Software Issues Resolved

The following issues were corrected or otherwise resolved in the Intel FPGA SDK for OpenCL and the Intel FPGA RTE for OpenCL Version 19.3.

**Table 1. Issues Resolved in the Intel FPGA SDK for OpenCL and the Intel FPGA RTE for OpenCL Version 19.3**

| Customer Service Request Numbers |
|---|
| 00449824 |

## 1.6. Software Patches Included in this Release

The following software patches are included in this release:

**Table 2. Software Patches Included in the Intel FPGA SDK for OpenCL**

| Software Version | Patch | Customer Service Request Number |
|---|---|---|
| Intel FPGA SDK for OpenCL Version 19.3 | 0.16cl | 1409450435 |
| | 0.12cl | 1409322299 |

## 1.7. Intel FPGA SDK for OpenCL Pro Edition Release Notes Archives

If the table does not list a software version, the user guide for the previous software version applies.

| Intel Quartus Prime Version | User Guide |
|---|---|
| 19.2 | Intel FPGA SDK for OpenCL Pro Edition Version 19.2 Release Notes |
| 19.1 | Intel FPGA SDK for OpenCL Pro Edition Version 19.1 Release Notes |
| 18.1 | Intel FPGA SDK for OpenCL Pro Edition Version 18.1 Release Notes |
| 18.0 | Intel FPGA SDK for OpenCL Pro Edition Release Notes |
| 17.1 | Intel FPGA SDK for OpenCL Release Notes |
| 17.0 | Intel FPGA SDK for OpenCL Release Notes |
| 16.1 | Intel FPGA SDK for OpenCL Release Notes |
| 16.0 | Altera SDK for OpenCL Version 16.0 Release Notes |
| 15.1 | Altera SDK for OpenCL Version 15.1 Release Notes |
| 15.0 | Altera SDK for OpenCL Version 15.0 Release Notes |

## 1.8. Document Revision History of the Intel FPGA SDK for OpenCL Pro Edition Release Notes

| Document Version | Intel Quartus Prime Version | Changes |
|---|---|---|
| 2019.09.30 | 19.3 | Initial release. |