

---

## 1 General Questions

### 1. Do I need to use additional software with the SDK?

To run an OpenCL™ application, you must have an OpenCL™ runtime on your system. If your system includes a recent AMD discrete GPU, or an APU, you also should install the latest Catalyst™ drivers, which can be downloaded from AMD.com. Information on supported devices can be found at [developer.amd.com/appsdk](http://developer.amd.com/appsdk). If your system does not include a recent AMD discrete GPU, or APU, the SDK installs a CPU-only OpenCL™ run-time.

Also, we recommend using the debugging profiling and analysis tools contained in the [AMD CodeXL heterogeneous compute tools suite](#).

### 2. Which versions of the OpenCL™ standard does this SDK support?

AMD APP SDK 2.9.1 supports the development of applications using the OpenCL™ Specification v 1.2.

### 3. Will applications developed to execute on OpenCL™ 1.1 still operate in an OpenCL™ 1.2 environment?

OpenCL™ is designed to be backwards compatible. The OpenCL™ 1.2 run-time delivered with the AMD Catalyst drivers run any OpenCL™ 1.1-compliant application. However, an OpenCL™ 1.2-compliant application will not execute on an OpenCL™ 1.1 run-time if APIs only supported by OpenCL™ 1.2 are used.

### 4. Does AMD provide any additional OpenCL™ samples, other than those contained within the SDK?

The most recent versions of all of the samples contained within the SDK are also available for individual download from the [developer.amd.com/appsdk](http://developer.amd.com/appsdk) “Samples & Demos” page. This page also contains additional samples that either were too large to include in the SDK, or which have been developed since the most recent SDK release. Check this web page for new, updated, or large samples.

### 5. How often can I expect to get AMD APP SDK updates?

Developers can expect that the AMD APP SDK may be updated two to three times a year. Actual release intervals may vary depending on available new features and product updates. AMD is committed to providing developers with regular updates to allow them to take advantage of the latest developments in AMD APP technology.

**6. What is the difference between the CPU and GPU components of OpenCL™ that are bundled with the AMD APP SDK?**

The CPU component uses the compatible CPU cores in your system to accelerate your OpenCL™ compute kernels; the GPU component uses the compatible GPU cores in your system to accelerate your OpenCL™ compute kernels.

**7. What CPUs does the AMD APP SDK v2.9.1 with OpenCL™ 1.2 support work on?**

The CPU component of OpenCL™ bundled with the AMD APP SDK works with any x86 CPU with SSE3 or later, as well as SSE2.x or later. AMD CPUs have supported SSE3 (and later) since 2005. Some examples of AMD CPUs that support SSE3 (or later) are the AMD Athlon™ 64 (starting with the Venice/San Diego steppings), AMD Athlon™ 64 X2, AMD Athlon™ 64 FX (starting with San Diego stepping), AMD Opteron™ (starting with E4 stepping), AMD Sempron™ (starting with Palermo stepping), AMD Phenom™, AMD Turion™ 64, and AMD Turion™ 64 X2.

**8. What APUs and GPUs does the AMD APP SDK v2.9.1 with OpenCL™ 1.2 support work on?**

For the list of supported APUs and GPUs, see the AMD APP SDK v2.9.1 System Requirements list at: <http://developer.amd.com/appsdk>

**9. Can my OpenCL™ code run on GPUs from other vendors?**

At this time, AMD does not plan to have the AMD APP SDK support GPU products from other vendors; however, since OpenCL™ is an industry standard programming interface, programs written in OpenCL™ 1.2 can be recompiled and run with any OpenCL-compliant compiler and runtime.

**10. What version of MS Visual Studio is supported?**

The AMD APP SDK v2.9.1 with OpenCL™ 1.2 supports Microsoft® Visual Studio® 2010 Professional Edition, Microsoft® Visual Studio 2012, and Microsoft® Visual Studio 2013. The AMD APP SDK installer includes project files for Visual Studio 2010 and 2012. Visual Studio 2013 projects can be generated using CMake files.

**11. Is it possible to run multiple AMD APP applications (compute and graphics) concurrently?**

Multiple AMD APP applications can be run concurrently, as long as they do not access the same GPU at the same time. AMD APP applications that attempt to access the same GPU at the same time are automatically serialized by the runtime system.

**12. Which graphics driver is required for the current AMD APP SDK v2.9.1 with OpenCL™ 1.2 CPU support?**

For the minimum required graphics driver, see the AMD APP SDK v2.9.1 System Requirements list at: <http://developer.amd.com/appsdk>. In general, it is advised that you update your system to use the most recent graphics drivers that are available for it.

**13. How does OpenCL™ compare to other APIs and programming platforms for parallel computing, such as OpenMP and MPI? Which one should I use?**

OpenCL™ is designed to target parallelism within a single system and provide portability to multiple different types of devices (GPUs, multi-core CPUs, etc.). OpenMP targets multi-core CPUs and SMP systems. MPI is a message passing protocol most often used for communication between nodes; it is a popular parallel programming model for clusters of machines. Each programming model has its advantages. It is anticipated that developers mix APIs, for example programming a cluster of machines with GPUs with MPI and OpenCL.

**14. If I write my code on the CPU version, does it work on the GPU version, or do I have to make changes.**

Assuming the size limitations for CPUs is considered, the code works on both the CPU and GPU components. Performance tuning, however, is different for each.

**15. What is the precision of mathematical operations?**

See Chapter 7, “OpenCL Numerical Compliance,” of the OpenCL™ 1.2 Specification for exact mathematical operations precision requirements.

<http://developer.amd.com/support/KnowledgeBase/Lists/KnowledgeBase/DispForm.aspx?ID=88>

**16. Are byte-addressable stores supported?**

Byte-addressable stores are supported.

**17. Are long integers supported?**

Yes, 64-bit integers are supported.

**18. Are operations on vectors supported?**

Yes, operations on vectors are supported.

**19. Is swizzling supported?**

Yes, swizzling (the rearranging of elements in a vector) is supported.

**20. How does one verify if OpenCL™ has been installed correctly on the system?**

Run `clinfo` from the command-line interface. The `clinfo` tool shows the available OpenCL™ devices on a system.

**21. How do I know if I have installed the latest version of the AMD APP SDK**

On installation of the SDK, if an internet connection is available, the installer states whether or not a newer SDK is available in the file `VersionInfo.txt` in the directory `C:\Users\<username>\Documents\AMD APP\<APP SDK version>\. Alternatively, you can check for the latest available version of the AMD APP SDK at http://developer.amd.com/appsdk.`

**22. Does it require administrator permissions to install the AMD APP SDK?**

On Windows, one needs administrator permissions to install the AMD APP SDK. On Linux, however, starting with APP SDK 2.9.1, users with root as well as non-root permissions can install the AMD APP SDK.

**23. Can multiple versions of the AMD APP SDK co-exist on a system?**

On Windows, the AMD APP SDK installer supports co-existence with earlier versions of the SDK. However on Linux, currently co-existence is not supported. The User must uninstall any existing version(s) of APP SDK before installing a new version.

## 2 OpenCL™ Questions

**24. What is OpenCL™?**

OpenCL™ (Open Computing Language) is the first truly open and royalty-free programming standard for general-purpose computations on heterogeneous systems. OpenCL™ lets programmers preserve their expensive source code investment and easily target both multi-core CPUs and the latest GPUs, such as those from AMD.

Developed in an open standards committee with representatives from major industry vendors, OpenCL™ gives users a cross-vendor, non-proprietary solution for accelerating their applications on their CPU and GPU cores.

**25. How much does the AMD OpenCL™ development platform cost?**

AMD bundles support for OpenCL™ as part of its AMD APP SDK product offering. The AMD APP SDK is offered to developers and users free of charge.

**26. What operating systems does the AMD APP SDK v2.9.1 with OpenCL™ 1.2 support?**

AMD APP SDK v2.9.1 runs on 32-bit and 64-bit versions of Windows and Linux. For the exact list of supported operating systems, see the AMD APP SDK v2.9.1 System Requirements list at: <http://developer.amd.com/appsdk>

**27. Can I write an OpenCL™ application that works on both CPU and GPU?**

Applications that program to the core OpenCL™ 1.2 API and kernel language should be able to target both CPUs and GPUs. At runtime, the appropriate device (CPU or GPU) must be selected by the application.

**28. Does the AMD OpenCL™ compiler automatically vectorize for SSE on the CPU?**

The CPU component of OpenCL™ that is bundled with the AMD APP SDK takes advantage of SSE3 instructions on the CPU. It also takes advantage of the AVX instructions where supported. In addition to AVX, OpenCL™ math library functions also leverage XOP and FMA4 capabilities on CPUs that support them.

**29. Does the AMD APP SDK v2.9.1 with OpenCL™ 1.2 support work on multiple GPUs (ATI CrossFire)?**

OpenCL™ applications can explicitly invoke separate compute kernels on multiple compatible GPUs in a single system. The partitioning of the algorithm to multiple parallel compute kernels must be done by the developer. It is recommended that ATI CrossFire be turned off in most system configurations so that AMD APP applications can access all available GPUs in the system.

ATI CrossFire technology allows multiple AMD GPUs to work together on a single graphics-rendering task. This method does not apply to AMD APP computational tasks because it is not compatible with the compute model used for AMD APP applications.

**30. Can I ship pre-compiled OpenCL™ application binaries that work on either CPU or GPU?**

By using OpenCL™ runtime APIs, developers can write OpenCL™ applications that can detect the available compatible CPUs and GPUs in the system. This lets developers pre-compile applications into binaries that dynamically work on either CPUs or GPUs that execute on targeted devices. Including LLVM IR in the binary provides a means for the binary to support devices for which the application was not explicitly pre-compiled.

**31. Is the OpenCL™ double precision optional extension supported?**

The Khronos and AMD double precision extensions are supported on certain devices. Your application can use the OpenCL™ API to query if this functionality is supported on the device in use.

**32. Is it possible to write OpenCL™ code that scales transparently over multiple devices?**

For OpenCL™ programs that target only multi-core CPUs, scaling can be done transparently; however, scaling across multiple GPUs requires the developer to explicitly partition the algorithm into multiple compute kernels, as well as explicitly launch the compute kernels onto each compatible GPU.

**33. What should I do if I get wrong results on the Apple platform with AMD devices?**

Apple handles support for the Apple platform; please contact them.

**34. Is it possible to dynamically index into a vector?**

No, this is not possible because a vector is not an array, but a representation of a hardware register.

**35. What is the difference between `local int a[4]` and `int a[4]`?**

`local int a[4]` uses hardware local memory, which is a small, low-latency, high-bandwidth memory; `int a[4]` uses per-thread hardware scratch memory, which is located in uncached global memory.

**36. How come my program runs slower in OpenCL™ than in CUDA/Brook+/IL?**

When comparing performance, it is better to compare code optimized for our OpenCL™ platform against code optimized against another vendor's OpenCL™ platforms. By comparing the same toolchain on different vendors, you can find out which vendors hardware works the best for your problem set.

**37. Why do read-write images not exist in OpenCL?**

OpenCL™ has a memory consistency model that requires certain constraints (see the *OpenCL Specification* for more information). Since images are special functional hardware units, they are different for reading and writing. This is different from pointers, which for the most part use the same hardware units and can guarantee the consistency that OpenCL™ requires.

### 38. Does prefetching work on the GPU?

Prefetch is not needed on the GPU because the hardware has a built-in mechanism to hide latency when many work-groups are running. The LDS can be used as a software-controlled cache.

### 39. How do you determine the max number of concurrent work-groups?

The maximum number of concurrent work-groups is determined by resource usage. This includes number of registers, amount of LDS space, and number of threads per work group. There is no way to directly specify the number of registers used by a kernel. Each SIMD has a 64-wide register file, with each column consisting of 256 x 32 x 4 registers.

### 40. Is it possible to tell OpenCL™ not to use all CPU Cores?

Yes, use the `device_fission` extension. Also, OpenCL 1.2 introduces the sub-device concept. The CPU can be partitioned into sub-devices/cores and the program can choose which sub-device/core to use. The Device Fission feature of OpenCL is demonstrated in the `DeviceFission` APP SDK sample.

## 3 OpenCL™ Optimizations

### 41. How do I use constants in a kernel for best performance?

For performance using constants, highest to lowest performance is achieved using:

- Literal values
- Constant pointer with compile time constants indexing.
- Constant pointer with runtime constant indexing that is the same for all threads.
- Constant pointer with linear access indexing that is the same for all threads.
- Constant pointer with linear access indexing that is different between threads.
- Constant pointer with random access indexing.

### 42. Why are literal values the fastest way to use constants?

Up to 96 bits of literal values are embedded in the instruction; thus, in theory, there is no limit on the number of usable literals. In practice, the limit is 16K unique literals in a compilation unit.

### 43. Why does `a * b + c` not generate a mad instruction?

Depending on the hardware and the floating point precision, the compiler may not generate a mad instruction for the computation of `a * b + c` due to the floating-point precision requirements in the OpenCL™ specification. Here, developers who want to exploit the mad instruction for performance can replace that computation with the `mad()` built-in function in OpenCL.

### 44. What is the preferred work-group size?

The preferred work-group size on the AMD platform is a multiple of 64.

**45. What is more efficient, the ternary operator `?:` or the `select` function?**

The `select` function compiles to the single cycle instruction, `cmov_logical`; in most cases, `?:` also compiles to the same instruction. In some cases, when memory is in one of the operands, the `?:` operator is compiled down to an IF/ELSE block. An IF/ELSE block takes more than a single instruction to execute.

**46. What is the difference between 24-bit and 32-bit integer operations?**

24-bit operations are faster because they use floating point hardware and can execute on all compute units. Many 32-bit integer operations also run on all stream processors, but if both a 24-bit and a 32-bit version exist for the same instruction, the 32-bit instruction executes only one per cycle.

## 4 Hardware Information

**47. How are 8/16-bit operations handled in hardware?**

The 8/16-bit operations are emulated with 32-bit registers.

**48. Do 24-bit integers exist in hardware?**

No, there are 24-bit instructions, such as `MUL24/MAD24`, but the smallest integer in hardware registers is 32-bits.

**49. What are the benefits of using 8/16-bit types over 32-bit integers?**

8/16-bit types take less memory than a 32-bit integer type, increasing the amount of data you are able to load with a single instruction. The OpenCL™ compiler up-converts to 32-bits on load and down-converts to the correct type on store.

**50. What is the difference between a GPR and a shared register?**

Although they are physically equivalent, the difference is whether the register offset in the hardware is absolute to the register file or relative to the wavefront ID.

**51. How often are wavefronts created?**

Wavefronts are created by the hardware to execute as long as resources are available. If they are created but cannot execute immediately, they are put in a wait queue where they stay until currently running wavefronts are finished.

**52. What is the maximum number of wavefronts?**

The maximum number of wavefronts is determined by which resource limits the number of wavefronts that can be spawned. This can be the number of registers, amount of local memory, required stack size, or other factors. Developers are recommended to profile the OpenCL application using CodeXL to figure out the resources that limit the number of wavefronts and that thereby affect the performance.

**53. Why do I get blue or black screens when executing longer running kernels?**

The GPU is not a preemptable device. If you are running the GPU as your display device, ensure that a compute program does not use the GPU past a certain time limit set by

Windows. Exceeding the time limit causes the watchdog timer to trigger; this can result in undefined program results.

**54. What is the cost of a clause switch?**

In general, the latency of a clause switch is around 40 cycles. Note that this is relevant only for Evergreen and Northern Island devices.

**55. How can I hide clause switch latency?**

By executing multiple wavefronts in parallel. Note that this is relevant only for Evergreen and Northern Island devices.

**56. How can I reduce clause switches?**

Clause switches are almost directly related to source program control flow. By reducing source program control flow, clause switches can also be reduced. This is only relevant for Evergreen and Northern Islands devices.

**57. How does the hardware execute a loop on a wavefront?**

The loop only ends execution for a wavefront once every thread in the wavefront breaks out of the loop. Once a thread breaks out of the loop, all of its execution results are masked, but the execution still occurs.

**58. How does flow control work with wavefronts?**

There are no flow control units for each individual thread, so the whole wavefront must execute the branch if any thread in the wavefront executes the branch. If the condition is false, the results are not written to memory, but the execution still occurs.

**59. What is the constant buffer size on GPU hardware?**

64 kB.

**60. What happens with out-of-bound memory operations?**

Writes are dropped, and reads return a pre-defined value.

## **5 Build Environment for the APP SDK samples**

**61. Can I use the SDK with Microsoft® Visual Studio® 2012 Express?**

Due to limitations in Microsoft® Visual Studio® Express, it is only possible to use build files for individual samples. Microsoft® Visual Studio® Express does not support building of all of the samples at the same time. The project files that build all of the samples are only supported by full versions of Microsoft® Visual Studio® 2010 or 2012.

**62. Why is CMake preferred over typical make files?**

CMake is a cross-platform open-source build system that can be freely downloaded from <http://www.cmake.org/>. CMake thus provides users with added flexibility.



## 6 Bolt

### 63. What is Bolt?

Bolt is a C++ template library optimized for heterogeneous computing. Bolt is designed to provide high-performance library implementations for common algorithms such as scan, reduce, transform, and sort. The Bolt interface was modeled on the C++ Standard Template Library (STL). Developers familiar with the STL will recognize many of the Bolt APIs and customization techniques.

### 64. Where can I learn more about Bolt ?

You can learn more about Bolt on this page: <https://github.com/HSA-Libraries/Bolt>.

### 65. Why does Bolt need TBB libraries ?

BOLT supports parallelization using Intel Threading Building Blocks (TBB (Intel)). One can switch between OpenCL /C++ AMP and TBB (Intel) calls by changing the control structure.

### 66. Where should one report the issues found in Bolt ?

One can report the issues found in Bolt on this page: <https://github.com/HSA-Libraries/bolt/issues>.

### 67. What is the selection order of the devices in Bolt?

As of the Bolt 1.2 release, the default mode is "Automatic" which means it will go into the OpenCL/C++ AMP path first, then TBB (Intel), then SerialCpu with the order mentioned below. The control will go to the other paths only if the selected path is not found.

The selection order is as follows:

- i. Run on GPU if AMD GPU is found and AMD APP SDK is installed.
- ii. If AMD GPU is not found, run on Multicore CPU if Intel TBB is installed.
- iii. If neither AMD GPU nor Intel TBB is found, run the CPU serial path.

### 68. How to force the APP SDK Bolt samples to run on a specific code path GPU/Multicore TBB/Serial CPU?

APP SDK Bolt sample provides the "--device" option to run on GPU or Multi-core CPU(TBB) or Serial CPU. The usage is as follows:

For Bolt OpenCL samples:

```
<BoltSample>.exe --device [auto|OpenCL|SerialCpu|MultiCoreCpu]
```

For Bolt C++ AMP Samples:

```
<BoltSample>.exe --device [auto|GPU|SerialCpu|MultiCoreCpu]
```

**69. Is installing TBB necessary ?**

No. The TBB libraries are needed if one wants to make use of the multi-core CPU path supported in BOLT.

**70. Do all APP SDK Bolt samples support running all 3 Bolt code paths GPU/Multicore TBB/Serial CPU?**

Except for the introductory samples, BoltIntro and BoltAMPIntro, all the other APP SDK Bolt samples support running in all 3 code paths GPU, Multicore TBB, and Serial CPU.

These Bolt projects include additional build configuration (Debug\_TBB and Release\_TBB) for enabling TBB path and have precompiler flag, "ENABLE\_TBB" set. Binaries generated in this build configuration have "\_TBB " as suffix and can be used to run in Multicore TBB path in addition to GPU and Serial CPU path. An additional requirement for running the multicore path is to have the environment variable TBB\_ROOT, that points to install location of Intel TBB.

Samples compiled with regular Debug and Release configuration can only be used to run in GPU and Serial CPU modes.

**71. Do all the Bolt APIs support the Multicore TBB path?**

Yes. Beginning with the Bolt 1.1 release, all the Bolt APIs support the Multicore TBB path.

**72. Which GPU compute technologies are supported by Bolt?**

APP SDK 2.9.1 is aligned with Bolt 1.2 libraries. The GPU compute technologies supported by Bolt 1.2 are OpenCL and C++ AMP.

**73. Does Bolt require linking to a library?**

Yes, Bolt includes a small, static library that must be linked into a user program to properly resolve symbols. Since Bolt is a template library, almost all functionality is written in header files, but OpenCL™ follows an online compilation model which makes sense to include in a library and not in header files.

**74. Does Bolt depend on any other libraries?**

Yes, Bolt contains dependencies on Boost. Bolt also depends on Intel TBB libraries for execution on the **multicore cpu** path.

**75. What algorithms does Bolt currently support?**

For the complete list of functions supported by Bolt 1.2 and the supported code paths, see: <https://hsa-libraries.github.io/Bolt/html/supportedFunction.html>.

**76. What version of OpenCL™ does Bolt currently require?**

Bolt uses features available in OpenCL™ v1.2.

**77. Does Bolt require the AMD OpenCL™ runtime?**

Yes, Bolt requires the use of templates in kernel code. At the time of this writing, AMD is the only vendor to provide this support.

## 78. Which Catalyst package should I use with Bolt?

Generally speaking, downloading and installing the latest Catalyst package contains the most recent OpenCL™ runtime. As of the time of this writing, the recommended Catalyst package is 14.4.

## 79. Which license is Bolt licensed under?

The Apache License, Version 2.0.

## 80. When should I use `device_vector` vs regular host memory?

`bolt::cl::device_vector` is used to manage device-local memory and can deliver higher performance on discrete GPU system. However, the host memory interfaces eliminate the need to create and manage `device_vectors`. If memory is re-used across multiple Bolt calls or is referenced by other kernels, using `device_vector` delivers higher performance

## 81. How do I measure the performance of OpenCL™ Bolt library calls?

The first time that a Bolt library routine is called, the runtime calls the OpenCL™ compiler to compile the kernel. Each unique template instantiation reuses the same OpenCL™ compilation, so two Bolt calls with the same functors are compiled only one time. When measuring the performance of Bolt, it is important to exclude the first call (with the compilation overhead) from the timing measurement. Here are two ways to time Bolt function calls:

The first example pulls the call outside the timer loop.

```
std::vector<int> a;
bolt::cl::transform(a.begin, a.end, z.begin(), bolt::cl::negate<int>);
startTimer(...);
for (int i=0; i<100; i++) {
    bolt::cl::transform(a.begin, a.end, z.begin(), bolt::cl::negate<int>);
}
stopTimer(...);
```

A second alternative is to explicitly exclude the first compilation from the timing region:

```
// Explicitly start timer after first iteration:
std::vector<int> a,z;
for (int i=0; i<100; i++) {
    if (i==1) {
        startTimer(...)
    }
    bolt::cl::transform(a.begin(), a.end(), z.begin(), bolt::cl::negate<int>);
}
stopTimer(...);
```

## 82. The Bolt library APIs accept host data structures such as `std::vector()`. Are these copied to the GPU or kept in host memory?

For host data structures, Bolt creates an OpenCL™ buffer with the `CL_MEM_USE_HOST_PTR` flag and uses the OpenCL™ map and unmap APIs to manage the host and device access. On the AMD OpenCL™ implementation, this results in data being kept in host memory when the host memory is aligned on 256-byte boundary. In other cases, the runtime copies the data as needed to the GPU. See section 4.5 of the *AMD Accelerated Parallel Processing OpenCL Programming Guide* for more information on memory optimization. Higher-

performance may be obtained by aligning the host data structures to 256-byte boundaries or by using the `device_vector` class for finer control over the allocation properties.

**83. Bolt APIs also accept `device_vector` inputs. When should use `device_vectors` for my inputs?**

The `bolt::cl::device_vector` interface is a thin wrapper around the `cl::Buffer` class; it provides an interface similar to `std::vector`. Also, the `device_vector` optionally accepts a `flags` argument that is passed to the OpenCL™ buffer creation. The following two flags can be useful:

`CL_MEM_ALLOC_HOST_POINTER`: Forces memory to be allocated in host memory, even on discrete GPUs. GPUs access to the data is slower (over the PCIe bus), but the allocation avoids the overhead of copying data to, and from, the GPU. This can be useful when the data is used in only a single Bolt call before being processed again on the CPU.

`CL_MEM_USE_PERSISTENT_MEM_AMD`: Provides device-resident zero copy memory. Use this for data that is write-only by the CPU. See section 4.5 of the *AMD Accelerated Parallel Processing OpenCL Programming Guide* for more information.

## 7 C++ AMP

**84. What is C++ AMP?**

C++ Accelerated Massive Parallelism (C++ AMP) accelerates execution of C++ code by taking advantage of data-parallel hardware, such as the compute units on and APU. C++ AMP is an open specification that extends regular C++ through the use of the `restrict` keyword to denote portions of code to be executed on the accelerated device.

**85. What AMD devices support C++ AMP?**

Any AMD APU or GPU that supports Microsoft® DirectX 11 Feature Level 11.0 and higher.

**86. Can you run a C++ AMP program on a CPU?**

You can run a C++ AMP program on a CPU, but its usage is not recommended in a production environment. For more information, see:  
<http://blogs.msdn.com/b/nativeconcurrency/archive/2012/03/10/cpu-accelerator-in-c-amp.aspx>

**87. Where can I get support for C++ AMP?**

C++ AMP is supported by the C++ compiler in Microsoft® Visual Studio® version 2012 and higher.

**88. Is C++ AMP supported on Linux?**

At this time, C++ AMP is not supported on Linux; however, as C++ AMP is defined as an open standard, Microsoft has opened the door for implementations on operating systems other than Windows.

**89. Where can I learn more about C++ AMP?**

There is an MSDN page on C++ AMP: <http://msdn.microsoft.com/en-us/library/vstudio/hh265137.aspx>. Also, the book *C++ AMP: Accelerated Massive Parallelism with Microsoft Visual C++*, by Kate Gregory and Ade Miller, may be useful.

## 8 Aparapi

### 90. What is Aparapi?

Aparapi is an API for expressing data parallel algorithms in Java and a runtime component capable of converting Java bytecode to OpenCL™ for execution on the GPU. If Aparapi cannot execute on the GPU at runtime, it executes the developer's algorithm in a Java thread pool.

For appropriate workloads, this extends Java's 'Write Once Run Anywhere' to include GPU devices.”

### 91. What AMD devices support Aparapi?

Any AMD APU, CPU, or GPU that supports OpenCL™ 1.1 and higher will work. See the list of AMD APP SDK v2.9.1 System Requirements list at: <http://developer.amd.com/appsdk>.

### 92. Where can I get support for Aparapi?

See: <http://code.google.com/p/aparapi/>

### 93. Is Aparapi supported on Linux?

Yes.

### 94. Which JDK/JRE is recommended for Aparapi?

The Oracle JDK/JRE is highly recommended for performance and for stability reasons.

### 95. Where can I learn more about Aparapi?

You can find out more about Aparapi from: <http://code.google.com/p/aparapi/>.

## 9 OpenCV

### 96. Where to download OpenCV libraries from?

OpenCV can either be installed by downloading pre-builts or by building from sources from [opencv.org](http://opencv.org).

### 97. Which version of OpenCV does APP SDK 2.9.1 support?

APP SDK 2.9.1 OpenCV samples require OpenCV 2.4.9 to be installed as a prerequisite. The Windows OpenCV 2.4.9 installation includes prebuilt binaries with OpenCL bindings.

### 98. Does OpenCV-CL APIs support interoperability with raw OpenCL kernels?

Yes. The CartoonFilter sample in the AMD APP SDK demonstrates it.

### 99. I am getting the `opencv2/core/core.hpp: No such file or directory` error while building.

Verify that the `OPENCV_DIR` environment variable is set/exported to the installed OpenCV directory and that the `OCVCL_VER` environment variable is set to OpenCV version number(249). For details, see the OpenCV section in the *AMD APP SDK Getting Started Guide* document.

**100.I am not able to run the built examples. I am getting the “error while loading shared libraries: libopencv\_core.so.2.x: cannot open shared object file: No such file or directory” error.**

Export LD\_LIBRARY\_PATH (in Linux) or set PATH (in Windows) to the installed OpenCV `lib` directory.

**101.I am getting the error - Unspecified error (The function is not implemented. Rebuild the library with Windows, GTK+ 2.x or Carbon support. If you are on Ubuntu or Debian, install libgtk2.0-dev and pkg-config, then re-run cmake or configure script).**

The reason for this error could be that the above tools might not have been installed before building OpenCV. The tools are required for displaying the images (cv::imshow calls). Install the above packages and rebuild the OpenCV library.

**102.Do the OpenCV samples depend on any other library apart from the OpenCV binaries and OpenCL?**

The Gesture Recognition sample depends on an additional library called `OpenNI`. The `OpenNI` library is used to extract frames from the input video.

**103.Where must one download the OpenNI libraries from?**

The `OpenNI` libraries can be downloaded from: <http://structure.io/openni>.

**104.What version of the OpenNI library is preferred?**

The `OpenNI` library version 2.2.0.33 Beta is preferred.

**105.I am getting the "Cannot find OpenNI.h" or "Cannot find OpenNI2.lib" error.**

For 32-bit builds, set `OPENNI2_INCLUDE` and `OPENNI2_LIB`, and for 64-bit builds, set `OPENI2_INCLUDE64` and `OPENNI2_LIB64`, to the corresponding include and library locations.

**106.Where can I know more about the OpenNI APIs?**

You can know more about the `OpenNI` APIs here: <http://structure.io/openni>.

---

#### Contact

Advanced Micro Devices, Inc.  
One AMD Place  
P.O. Box 3453  
Sunnyvale, CA, 94088-3453  
Phone: +1.408.749.4000

#### For AMD Accelerated Parallel Processing:

URL: [developer.amd.com/appsdk](http://developer.amd.com/appsdk)  
Developing: [developer.amd.com/](http://developer.amd.com/)  
Forum: [developer.amd.com/opencilforum](http://developer.amd.com/opencilforum)



The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

#### Copyright and Trademarks

© 2014 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo, Radeon, FireStream, and combinations thereof are trademarks of Advanced Micro Devices, Inc. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.