

Number of platforms

2

Platform Name
Platform Vendor
Platform Version
Platform Profile

Intel(R) OpenCL
Intel(R) Corporation



OpenCL 1.2 LINUX
FULL_PROFILE
cl_khr_icd cl_khr_global_int32_base_atomics
cl_khr_global_int32_extended_atomics
cl_khr_local_int32_base_atomics
cl_khr_local_int32_extended_atomics
cl_khr_byte_addressable_store
cl_khr_depth_images cl_khr_3d_image_writes
cl_intel_exec_by_local_thread cl_khr_spir
cl_khr_fp64

Platform Extensions
Platform Extensions function suffix

INTEL

Platform Name
Platform Vendor
Platform Version
Platform Profile

NVIDIA CUDA
NVIDIA Corporation
OpenCL 1.2 CUDA 8.0.0



FULL_PROFILE
cl_khr_global_int32_base_atomics
cl_khr_global_int32_extended_atomics
cl_khr_local_int32_base_atomics
cl_khr_local_int32_extended_atomics cl_khr_fp64
cl_khr_byte_addressable_store cl_khr_icd
cl_khr_gl_sharing cl_nv_compiler_options
cl_nv_device_attribute_query cl_nv_pragma_unroll
cl_nv_copy_opts

Platform Extensions
Platform Extensions function suffix

NV

Platform Name	Intel(R) OpenCL
Number of devices	1
<u>Device Name</u>	<u>Intel(R) Xeon(R) CPU E5-2670 v2 @ 2.50GHz</u>
Device Vendor	Intel(R) Corporation
Device Vendor ID	0x8086
Device Version	OpenCL 1.2 (Build 25)
Driver Version	1.2.0.25
Device OpenCL C Version	OpenCL C 1.2
<u>Device Type</u>	<u>CPU</u>
Device Available	Yes
Device Profile	FULL_PROFILE
Max compute units	40
Max clock frequency	2500MHz
Device Partition	(core)
Max number of sub-devices	40
Supported partition types	by counts
Max work item dimensions	3
Max work item sizes	8192x8192x8192
Max work group size	8192
Compiler Available	Yes
Linker Available	Yes
Preferred work group size multiple	128
Preferred / native vector sizes	
char	1 / 16
short	1 / 8
int	1 / 4
long	1 / 2
half	0 / 0 (n/a)
float	1 / 8
double	1 / 4 (cl_khr_fp64)
Half-precision Floating-point support	(n/a)
Single-precision Floating-point support	(core)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	No
Round to infinity	No
IEEE754-2008 fused multiply-add	No
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	No
Double-precision Floating-point support	(cl_khr_fp64)
Denormals	Yes



Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	No
Address bits	64
Global memory size	134987472896 (125.7GiB)
Error Correction support	No
Max memory allocation	33746868224 (31.43GiB)
Unified memory for Host and Device	Yes
Minimum alignment for any data type	128 bytes
Alignment of base address	1024 bits (128 bytes)
Global Memory cache type	Read/Write
Global Memory cache size	262144
Global Memory cache line	64 bytes
Image support	Yes
Max number of samplers per kernel	480
Max size for 1D images from buffer	2109179264 pixels
Max 1D or 2D image array size	2048 images
Max 2D image size	16384x16384 pixels
Max 3D image size	2048x2048x2048 pixels
Max number of read image args	480
Max number of write image args	480
Local memory type	Global
Local memory size	32768 (32KiB)
Max constant buffer size	131072 (128KiB)
Max number of constant args	480
Max size of kernel argument	3840 (3.75KiB)
Queue properties	
Out-of-order execution	Yes
Profiling	Yes
Local thread execution (Intel)	Yes
Prefer user sync for interop	No
Profiling timer resolution	1ns
Execution capabilities	
Run OpenCL kernels	Yes
Run native kernels	Yes
SPIR versions	1.2
printf() buffer size	1048576 (1024KiB)
Built-in kernels	
Device Extensions	cl_khr_fp64

Platform Name	NVIDIA CUDA
Number of devices	3
<u>Device Name</u>	<u>Tesla K40m</u>
Device Vendor	NVIDIA Corporation
Device Vendor ID	0x10de
Device Version	OpenCL 1.2 CUDA
Driver Version	375.26
Device OpenCL C Version	OpenCL C 1.2
<u>Device Type</u>	<u>GPU</u>
Device Available	Yes
Device Profile	FULL_PROFILE
Device Topology (NV)	PCI-E
Max compute units	15
Max clock frequency	745MHz
Compute Capability (NV)	3.5
Device Partition	(core)
Max number of sub-devices	1
Supported partition types	None
Max work item dimensions	3
Max work item sizes	1024x1024x64
Max work group size	1024
Compiler Available	Yes
Linker Available	Yes
Preferred work group size multiple	32
Warp size (NV)	32
Preferred / native vector sizes	
char	1 / 1
short	1 / 1
int	1 / 1
long	1 / 1
half	0 / 0 (n/a)
float	1 / 1
double	1 / 1 (cl_khr_fp64)
Half-precision Floating-point support	(n/a)
Single-precision Floating-point support	(core)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	Yes



Double-precision Floating-point support	(cl_khr_fp64)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	No
Address bits	64
Global memory size	11995578368 (11.17GiB)
Error Correction support	Yes
Max memory allocation	2998894592 (2.793GiB)
Unified memory for Host and Device	No
Integrated memory (NV)	No
Minimum alignment for any data type	128 bytes
Alignment of base address	4096 bits (512 bytes)
Global Memory cache type	Read/Write
Global Memory cache size	245760
Global Memory cache line	128 bytes
Image support	Yes
Max number of samplers per kernel	32
Max size for 1D images from buffer	134217728 pixels
Max 1D or 2D image array size	2048 images
Max 2D image size	16384x16384 pixels
Max 3D image size	4096x4096x4096 pixels
Max number of read image args	256
Max number of write image args	16
Local memory type	Local
Local memory size	49152 (48KiB)
Registers per block (NV)	65536
Max constant buffer size	65536 (64KiB)
Max number of constant args	9
Max size of kernel argument	4352 (4.25KiB)
Queue properties	
Out-of-order execution	Yes
Profiling	Yes
Prefer user sync for interop	No
Profiling timer resolution	1000ns
Execution capabilities	
Run OpenCL kernels	Yes
Run native kernels	No

Kernel execution timeout (NV)	No
Concurrent copy and kernel execution (NV)	Yes
Number of async copy engines	2
printf() buffer size	1048576 (1024KiB)
Built-in kernels	cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_fp64 cl_khr_byte_addressable_store cl_khr_icd cl_khr_gl_sharing cl_nv_compiler_options cl_nv_device_attribute_query cl_nv_pragma_unroll cl_nv_copy_opts
Device Extensions	

Device Name

Device Vendor

Device Vendor ID

Device Version

Driver Version

Device OpenCL C Version

Device Type

Device Available

Device Profile

Device Topology (NV)

Max compute units

Max clock frequency

Compute Capability (NV)

Device Partition

Max number of sub-devices

Supported partition types

Max work item dimensions

Max work item sizes

Max work group size

Compiler Available

Linker Available

Preferred work group size multiple

Warp size (NV)

Preferred / native vector sizes

char

short

Tesla K40m

NVIDIA Corporation

0x10de

OpenCL 1.2 CUDA

375.26

OpenCL C 1.2

GPU

Yes

FULL_PROFILE

PCI-E

15

745MHz

3.5

(core)

1

None

3

1024x1024x64

1024

Yes

Yes

32

32

1 / 1

1 / 1



int	1 / 1
long	1 / 1
half	0 / 0 (n/a)
float	1 / 1
double	1 / 1 (cl_khr_fp64)
Half-precision Floating-point support	(n/a)
Single-precision Floating-point support	(core)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	Yes
Double-precision Floating-point support	(cl_khr_fp64)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	No
Address bits	64
Global memory size	11995578368 (11.17GiB)
Error Correction support	Yes
Max memory allocation	2998894592 (2.793GiB)
Unified memory for Host and Device	No
Integrated memory (NV)	No
Minimum alignment for any data type	128 bytes
Alignment of base address	4096 bits (512 bytes)
Global Memory cache type	Read/Write
Global Memory cache size	245760
Global Memory cache line	128 bytes
Image support	Yes
Max number of samplers per kernel	32
Max size for 1D images from buffer	134217728 pixels
Max 1D or 2D image array size	2048 images
Max 2D image size	16384x16384 pixels
Max 3D image size	4096x4096x4096 pixels
Max number of read image args	256

Max number of write image args	16
Local memory type	Local
Local memory size	49152 (48KiB)
Registers per block (NV)	65536
Max constant buffer size	65536 (64KiB)
Max number of constant args	9
Max size of kernel argument	4352 (4.25KiB)
Queue properties	
Out-of-order execution	Yes
Profiling	Yes
Prefer user sync for interop	No
Profiling timer resolution	1000ns
Execution capabilities	
Run OpenCL kernels	Yes
Run native kernels	No
Kernel execution timeout (NV)	No
Concurrent copy and kernel execution (NV)	Yes
Number of async copy engines	2
printf() buffer size	1048576 (1024KiB)
Built-in kernels	

Device Name	Tesla K40m
Device Vendor	NVIDIA Corporation
Device Vendor ID	0x10de
Device Version	OpenCL 1.2 CUDA
Driver Version	375.26
Device OpenCL C Version	OpenCL C 1.2
<u>Device Type</u>	<u>GPU</u>
Device Available	Yes
Device Profile	FULL_PROFILE
Device Topology (NV)	PCI-E
Max compute units	15
Max clock frequency	745MHz
Compute Capability (NV)	3.5
Device Partition	(core)
Max number of sub-devices	1
Supported partition types	None
Max work item dimensions	3
Max work item sizes	1024x1024x64
Max work group size	1024
Compiler Available	Yes
Linker Available	Yes
Preferred work group size multiple	32
Warp size (NV)	32
Preferred / native vector sizes	
char	1 / 1
short	1 / 1
int	1 / 1
long	1 / 1
half	0 / 0 (n/a)
float	1 / 1
double	1 / 1 (cl_khr_fp64)
Half-precision Floating-point support	(n/a)
Single-precision Floating-point support	(core)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	Yes



Double-precision Floating-point support	(cl_khr_fp64)
Denormals	Yes
Infinity and NaNs	Yes
Round to nearest	Yes
Round to zero	Yes
Round to infinity	Yes
IEEE754-2008 fused multiply-add	Yes
Support is emulated in software	No
Correctly-rounded divide and sqrt operations	No
Address bits	64
Global memory size	11995578368 (11.17GiB)
Error Correction support	Yes
Max memory allocation	2998894592 (2.793GiB)
Unified memory for Host and Device	No
Integrated memory (NV)	No
Minimum alignment for any data type	128 bytes
Alignment of base address	4096 bits (512 bytes)
Global Memory cache type	Read/Write
Global Memory cache size	245760
Global Memory cache line	128 bytes
Image support	Yes
Max number of samplers per kernel	32
Max size for 1D images from buffer	134217728 pixels
Max 1D or 2D image array size	2048 images
Max 2D image size	16384x16384 pixels
Max 3D image size	4096x4096x4096 pixels
Max number of read image args	256
Max number of write image args	16
Local memory type	Local
Local memory size	49152 (48KiB)
Registers per block (NV)	65536
Max constant buffer size	65536 (64KiB)
Max number of constant args	9
Max size of kernel argument	4352 (4.25KiB)
Queue properties	
Out-of-order execution	Yes
Profiling	Yes
Prefer user sync for interop	No
Profiling timer resolution	1000ns
Execution capabilities	
Run OpenCL kernels	Yes
Run native kernels	No

Kernel execution timeout (NV)	No
Concurrent copy and kernel execution (NV)	Yes
Number of async copy engines	2
printf() buffer size	1048576 (1024KiB)
Built-in kernels	<ul style="list-style-type: none"> cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_fp64 cl_khr_byte_addressable_store cl_khr_icd cl_khr_gl_sharing cl_nv_compiler_options cl_nv_device_attribute_query cl_nv_pragma_unroll cl_nv_copy_opts
Device Extensions	
NULL platform behavior	