

---

# **os-traits Documentation**

*Release 3.4.1.dev1*

**OpenStack Foundation**

**Apr 03, 2025**



# CONTENTS

- 1 Installation Guide** **3**
- 1.1 Installation . . . . . 3
- 2 Usage Guide** **5**
- 2.1 Usage . . . . . 5
- 3 Contributor Guide** **7**
- 3.1 Contributing . . . . . 7
- 4 Reference** **9**
- 4.1 Reference . . . . . 9
- 4.2 Available Traits . . . . . 10



*os-traits* is a library containing standardized trait strings.

Traits are strings that represent a feature of some resource provider. This library contains the catalog of constants that have been standardized in the OpenStack community to refer to a particular hardware, virtualization, storage, network, or device trait.



## INSTALLATION GUIDE

### 1.1 Installation

At the command line:

```
$ pip install os-traits
```

Or, if you have virtualenvwrapper installed:

```
$ mkvirtualenv os-traits  
$ pip install os-traits
```





## USAGE GUIDE

### 2.1 Usage

*os-traits* is primarily composed of a set of constants that may be referenced by simply importing the `os_traits` module and referencing one of the modules traits constants:

```
$ python3
Python 3.6.8 (default, Oct 7 2019, 12:59:55)
[GCC 8.3.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import os_traits as ot
>>> print(ot.HW_CPU_X86_SSE42)
HW_CPU_X86_SSE42
```

You can get a list of the `os_traits` symbols by simply doing `dir(os_traits)`.

Want to see the trait strings for a subset of traits? There's a method for that too:

```
>>> import pprint
>>> pprint.pprint(ot.get_traits(prefix='HW_CPU_X86_'))
['HW_CPU_X86_FMA3',
 'HW_CPU_X86_AVX',
 'HW_CPU_X86_MMX',
 'HW_CPU_X86_MPX',
 'HW_CPU_X86_CLMUL',
 'HW_CPU_X86_AVX512VL',
 'HW_CPU_X86_AVX512CD',
 'HW_CPU_X86_BMI',
 'HW_CPU_X86_AVX512DQ',
 'HW_CPU_X86_SSE3',
 'HW_CPU_X86_ABM',
 'HW_CPU_X86_SSE4A',
 'HW_CPU_X86_AESNI',
 'HW_CPU_X86_F16C',
 'HW_CPU_X86_VMX',
 'HW_CPU_X86_SVM',
 'HW_CPU_X86_TSX',
 'HW_CPU_X86_AVX512PF',
 'HW_CPU_X86_SSE41',
 'HW_CPU_X86_ASF',
```

(continues on next page)

(continued from previous page)

```
'HW_CPU_X86_SGX',  
'HW_CPU_X86_SSE',  
'HW_CPU_X86_SSSE3',  
'HW_CPU_X86_SHA',  
'HW_CPU_X86_TBM',  
'HW_CPU_X86_SSE42',  
'HW_CPU_X86_3DNOW',  
'HW_CPU_X86_BMI2',  
'HW_CPU_X86_AVX512BW',  
'HW_CPU_X86_XOP',  
'HW_CPU_X86_AVX2',  
'HW_CPU_X86_AVX512F',  
'HW_CPU_X86_SSE2',  
'HW_CPU_X86_FMA4',  
'HW_CPU_X86_AVX512ER']
```

## CONTRIBUTOR GUIDE

### 3.1 Contributing

#### 3.1.1 Trait lifecycle policy

It is the policy of this project that once registered, traits should never be removed, even those which will never be used by code (e.g. as a result of pivots in design or changes to the namespaces).

The general principle behind this policy is simply that an extensible-only enumeration is easier to manage than one that can be shrunk. One particular example concerns the need for the placement service to keep its database in sync with the strings in os-traits. Whenever a placement service sees a new version of os-traits it syncs up its database with the strings that are in the package, creating a row in the traits table, with an id that becomes a foreign key in other tables. If traits could be removed, then extra clean-up code might be needed in several places to handle this, and this would be particularly error-prone when execution of that code would need to be correctly orchestrated across multiple projects.

#### 3.1.2 Generic instructions for contributing

If you would like to contribute to the development of OpenStack, you must follow the steps in this page:

<https://docs.openstack.org/infra/manual/developers.html>

If you already have a good understanding of how the system works and your OpenStack accounts are set up, you can skip to the development workflow section of this documentation to learn how changes to OpenStack should be submitted for review via the Gerrit tool:

<https://docs.openstack.org/infra/manual/developers.html#development-workflow>

Pull requests submitted through GitHub will be ignored.

Bugs should be filed on StoryBoard, not GitHub:

<https://storyboard.openstack.org#!/project/openstack/os-traits>

The code is hosted at:

<https://opendev.org/openstack/os-traits>



## 4.1 Reference

- *CUDA*
- *AMD SEV*

### 4.1.1 CUDA

Applications that need to perform massively parallel operations, like processing large arrays, may use the CUDA framework to accelerate their processing on graphics processing units (GPUs). The CUDA framework has two complementary pieces to it.

There are a set of GPU instruction set extensions that are implemented by various graphics cards. These instruction set extensions are known as the CUDA Compute Capability.

The second part of the framework is an SDK that allows developers to take advantage of the hardware's instruction set extensions of a particular version (a specific CUDA Compute Capability version, that is).

An application will link with a version of the CUDA SDK, and the version of the CUDA SDK controls which CUDA Compute Capability versions the application will be able to work with.

The `os_traits.hw.gpu.cuda` module contains traits for both the CUDA compute capability version as well as the CUDA SDK version. For example, `os_traits.hw.gpu.cuda.COMPUTE_CAPABILITY_V3_2` and `os_traits.hw.gpu.cuda.SDK_V6_5`.

The `os_traits.hw.gpu.cuda` module contains a utility function called `compute_capabilities_supported()` that accepts a trait indicating the CUDA SDK version and returns a `set()` containing the matching CUDA compute capability traits that that version of the CUDA SDK knows how to utilize.

Here is an example of listing the CUDA compute capability version traits that the CUDA SDK 8.0 is capable of working with:

```
>>> from os_traits.hw.gpu import cuda
>>> import pprint
>>>
>>> sdk8_caps = cuda.compute_capabilities_supported(cuda.SDK_V8_0)
>>> pprint.pprint(sdk8_caps)
set(['HW_GPU_CUDA_COMPUTE_CAPABILITY_V2_0',
     'HW_GPU_CUDA_COMPUTE_CAPABILITY_V2_1',
```

(continues on next page)

(continued from previous page)

```
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V3_0',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V3_2',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V3_5',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V3_7',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V5_0',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V5_2',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V5_3',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V6_0',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V6_1',  
'HW_GPU_CUDA_COMPUTE_CAPABILITY_V6_2']])
```

For more information on CUDA, see the [Wikipedia article](#).

### 4.1.2 AMD SEV

While data is typically encrypted today when stored on disk, it is stored in DRAM in the clear. This can leave the data vulnerable to snooping by unauthorized administrators or software, or by hardware probing. New non-volatile memory technology (NVDIMM) exacerbates this problem since an NVDIMM chip can be physically removed from a system with the data intact, similar to a hard drive. Without encryption any stored information such as sensitive data, passwords, or secret keys can be easily compromised.

AMDs SEV (Secure Encrypted Virtualization) is a VM protection technology which transparently encrypts the memory of each VM with a unique key. It can also calculate a signature of the memory contents, which can be sent to the VMs owner as an attestation that the memory was encrypted correctly by the firmware. SEV is particularly applicable to cloud computing since it can reduce the amount of trust VMs need to place in the hypervisor and administrator of their host system.

The `os_traits.hw.cpu.amd.SEV` trait is reserved in order to indicate that a compute host contains support for SEV not only on-CPU, but also in all other layers of the hypervisor stack required in order to take advantage of this feature.

## 4.2 Available Traits

Below is a list of all traits currently available.

- COMPUTE\_DEVICE\_TAGGING
- COMPUTE\_NODE
- COMPUTE\_TRUSTED\_CERTS
- COMPUTE\_SAME\_HOST\_COLD\_MIGRATE
- COMPUTE\_RESCUE\_BFV
- COMPUTE\_ACCELERATORS
- COMPUTE\_SOCKET\_PCI\_NUMA\_AFFINITY
- COMPUTE\_REMOTE\_MANAGED\_PORTS
- COMPUTE\_MEM\_BACKING\_FILE
- COMPUTE\_MANAGED\_PCI\_DEVICE
- COMPUTE\_ADDRESS\_SPACE\_PASSTHROUGH

- COMPUTE\_ADDRESS\_SPACE\_EMULATED
- COMPUTE\_SHARE\_LOCAL\_FS
- COMPUTE\_ARCH\_AARCH64
- COMPUTE\_ARCH\_PPC64LE
- COMPUTE\_ARCH\_MIPSEL
- COMPUTE\_ARCH\_S390X
- COMPUTE\_ARCH\_RISCV64
- COMPUTE\_ARCH\_X86\_64
- COMPUTE\_CONFIG\_DRIVE\_REGENERATION
- COMPUTE\_EPHEMERAL\_ENCRYPTION
- COMPUTE\_EPHEMERAL\_ENCRYPTION\_PLAIN
- COMPUTE\_EPHEMERAL\_ENCRYPTION\_LUKS
- COMPUTE\_EPHEMERAL\_ENCRYPTION\_LUKSV2
- COMPUTE\_FIRMWARE\_BIOS
- COMPUTE\_FIRMWARE\_UEFI
- COMPUTE\_GRAPHICS\_MODEL\_BOCHS
- COMPUTE\_GRAPHICS\_MODEL\_CIRRUS
- COMPUTE\_GRAPHICS\_MODEL\_GOP
- COMPUTE\_GRAPHICS\_MODEL\_NONE
- COMPUTE\_GRAPHICS\_MODEL\_QXL
- COMPUTE\_GRAPHICS\_MODEL\_VGA
- COMPUTE\_GRAPHICS\_MODEL\_VIRTIO
- COMPUTE\_GRAPHICS\_MODEL\_VMVGA
- COMPUTE\_GRAPHICS\_MODEL\_XEN
- COMPUTE\_IMAGE\_TYPE\_AKI
- COMPUTE\_IMAGE\_TYPE\_AMI
- COMPUTE\_IMAGE\_TYPE\_ARI
- COMPUTE\_IMAGE\_TYPE\_ISO
- COMPUTE\_IMAGE\_TYPE\_QCOW2
- COMPUTE\_IMAGE\_TYPE\_RAW
- COMPUTE\_IMAGE\_TYPE\_VDI
- COMPUTE\_IMAGE\_TYPE\_VHD
- COMPUTE\_IMAGE\_TYPE\_VHDX
- COMPUTE\_IMAGE\_TYPE\_VMDK

- COMPUTE\_IMAGE\_TYPE\_PLOOP
- COMPUTE\_MIGRATE\_AUTO\_CONVERGE
- COMPUTE\_MIGRATE\_POST\_COPY
- COMPUTE\_NET\_ATTACH\_INTERFACE
- COMPUTE\_NET\_ATTACH\_INTERFACE\_WITH\_TAG
- COMPUTE\_NET\_VIRTIO\_PACKED
- COMPUTE\_NET\_VIF\_MODEL\_E1000
- COMPUTE\_NET\_VIF\_MODEL\_E1000E
- COMPUTE\_NET\_VIF\_MODEL\_LAN9118
- COMPUTE\_NET\_VIF\_MODEL\_NETFRONT
- COMPUTE\_NET\_VIF\_MODEL\_NE2K\_PCI
- COMPUTE\_NET\_VIF\_MODEL\_PCNET
- COMPUTE\_NET\_VIF\_MODEL\_RTL8139
- COMPUTE\_NET\_VIF\_MODEL\_SPAPR\_VLAN
- COMPUTE\_NET\_VIF\_MODEL\_SRIOV
- COMPUTE\_NET\_VIF\_MODEL\_VIRTIO
- COMPUTE\_NET\_VIF\_MODEL\_VMXNET
- COMPUTE\_NET\_VIF\_MODEL\_VMXNET3
- COMPUTE\_NET\_VIF\_MODEL\_IGB
- COMPUTE\_SECURITY\_TPM\_1\_2
- COMPUTE\_SECURITY\_TPM\_2\_0
- COMPUTE\_SECURITY\_TPM\_TIS
- COMPUTE\_SECURITY\_TPM\_CRB
- COMPUTE\_SECURITY\_TPM\_SECRET\_SECURITY\_USER
- COMPUTE\_SECURITY\_TPM\_SECRET\_SECURITY\_HOST
- COMPUTE\_SECURITY\_TPM\_SECRET\_SECURITY\_DEPLOYMENT
- COMPUTE\_SECURITY\_UEFI\_SECURE\_BOOT
- COMPUTE\_SECURITY\_STATELESS\_FIRMWARE
- COMPUTE\_SOUND\_MODEL\_SB16
- COMPUTE\_SOUND\_MODEL\_ES1370
- COMPUTE\_SOUND\_MODEL\_PCSPK
- COMPUTE\_SOUND\_MODEL\_AC97
- COMPUTE\_SOUND\_MODEL\_ICH6
- COMPUTE\_SOUND\_MODEL\_ICH9



- COMPUTE\_SOUND\_MODEL\_USB
- COMPUTE\_SOUND\_MODEL\_VIRTIO
- COMPUTE\_STATUS\_DISABLED
- COMPUTE\_STORAGE\_BUS\_FDC
- COMPUTE\_STORAGE\_BUS\_IDE
- COMPUTE\_STORAGE\_BUS\_LXC
- COMPUTE\_STORAGE\_BUS\_SATA
- COMPUTE\_STORAGE\_BUS\_SCSI
- COMPUTE\_STORAGE\_BUS\_USB
- COMPUTE\_STORAGE\_BUS\_VIRTIO
- COMPUTE\_STORAGE\_BUS\_UML
- COMPUTE\_STORAGE\_BUS\_XEN
- COMPUTE\_STORAGE\_VIRTIO\_FS
- COMPUTE\_USB\_MODEL\_QEMU\_XHCI
- COMPUTE\_USB\_MODEL\_NEC\_XHCI
- COMPUTE\_VIOMMU\_MODEL\_INTEL
- COMPUTE\_VIOMMU\_MODEL\_SMMUV3
- COMPUTE\_VIOMMU\_MODEL\_VIRTIO
- COMPUTE\_VIOMMU\_MODEL\_AUTO
- COMPUTE\_VOLUME\_ATTACH
- COMPUTE\_VOLUME\_ATTACH\_WITH\_TAG
- COMPUTE\_VOLUME\_EXTEND
- COMPUTE\_VOLUME\_MULTI\_ATTACH
- HW\_ARCH\_ALPHA
- HW\_ARCH\_ARMV6
- HW\_ARCH\_ARMV7
- HW\_ARCH\_ARMV7B
- HW\_ARCH\_AARCH64
- HW\_ARCH\_CRIS
- HW\_ARCH\_I686
- HW\_ARCH\_IA64
- HW\_ARCH\_LM32
- HW\_ARCH\_M68K
- HW\_ARCH\_MICROBLAZE

- HW\_ARCH\_MICROBLAZEEL
- HW\_ARCH\_MIPS
- HW\_ARCH\_MIPSEL
- HW\_ARCH\_MIPS64
- HW\_ARCH\_MIPS64EL
- HW\_ARCH\_OPENRISC
- HW\_ARCH\_PARISC
- HW\_ARCH\_PARISC64
- HW\_ARCH\_PPC
- HW\_ARCH\_PPCLE
- HW\_ARCH\_PPC64
- HW\_ARCH\_PPC64LE
- HW\_ARCH\_PPCEMB
- HW\_ARCH\_S390
- HW\_ARCH\_S390X
- HW\_ARCH\_SH4
- HW\_ARCH\_SH4EB
- HW\_ARCH\_SPARC
- HW\_ARCH\_SPARC64
- HW\_ARCH\_UNICORE32
- HW\_ARCH\_X86\_64
- HW\_ARCH\_XTENSA
- HW\_ARCH\_XTENSAEB
- HW\_CPU\_HYPERTHREADING
- HW\_CPU\_AARCH64\_FP
- HW\_CPU\_AARCH64\_ASIMD
- HW\_CPU\_AARCH64\_EVTSTRM
- HW\_CPU\_AARCH64\_AES
- HW\_CPU\_AARCH64\_PMULL
- HW\_CPU\_AARCH64\_SHA1
- HW\_CPU\_AARCH64\_SHA2
- HW\_CPU\_AARCH64\_CRC32
- HW\_CPU\_AARCH64\_FPHP
- HW\_CPU\_AARCH64\_ASIMDHP

- HW\_CPU\_AARCH64\_ASIMDRDM
- HW\_CPU\_AARCH64\_ATOMICS
- HW\_CPU\_AARCH64\_JSCVT
- HW\_CPU\_AARCH64\_FCMA
- HW\_CPU\_AARCH64\_LRCPC
- HW\_CPU\_AARCH64\_DCPPOP
- HW\_CPU\_AARCH64\_SHA3
- HW\_CPU\_AARCH64\_SM3
- HW\_CPU\_AARCH64\_SM4
- HW\_CPU\_AARCH64\_ASIMDDP
- HW\_CPU\_AARCH64\_SHA512
- HW\_CPU\_AARCH64\_SVE
- HW\_CPU\_AARCH64\_CPUID
- HW\_CPU\_AMD\_SEV
- HW\_CPU\_PPC64LE\_POWER8
- HW\_CPU\_PPC64LE\_POWER9
- HW\_CPU\_X86\_AVX
- HW\_CPU\_X86\_AVX2
- HW\_CPU\_X86\_CLMUL
- HW\_CPU\_X86\_FMA3
- HW\_CPU\_X86\_FMA4
- HW\_CPU\_X86\_F16C
- HW\_CPU\_X86\_MMX
- HW\_CPU\_X86\_SSE
- HW\_CPU\_X86\_SSE2
- HW\_CPU\_X86\_SSE3
- HW\_CPU\_X86\_SSSE3
- HW\_CPU\_X86\_SSE41
- HW\_CPU\_X86\_SSE42
- HW\_CPU\_X86\_SSE4A
- HW\_CPU\_X86\_XOP
- HW\_CPU\_X86\_3DNOW
- HW\_CPU\_X86\_AVX512F
- HW\_CPU\_X86\_AVX512CD

- HW\_CPU\_X86\_AVX512PF
- HW\_CPU\_X86\_AVX512ER
- HW\_CPU\_X86\_AVX512VL
- HW\_CPU\_X86\_AVX512BW
- HW\_CPU\_X86\_AVX512DQ
- HW\_CPU\_X86\_AVX512VNNI
- HW\_CPU\_X86\_AVX512VBMI
- HW\_CPU\_X86\_AVX512IFMA
- HW\_CPU\_X86\_AVX512VBMI2
- HW\_CPU\_X86\_AVX512BITALG
- HW\_CPU\_X86\_AVX512VAES
- HW\_CPU\_X86\_AVX512GFNI
- HW\_CPU\_X86\_AVX512VPCLMULQDQ
- HW\_CPU\_X86\_AVX512VPOPCNTDQ
- HW\_CPU\_X86\_ABM
- HW\_CPU\_X86\_BMI
- HW\_CPU\_X86\_BMI2
- HW\_CPU\_X86\_TBM
- HW\_CPU\_X86\_AESNI
- HW\_CPU\_X86\_SHA
- HW\_CPU\_X86\_MPX
- HW\_CPU\_X86\_SGX
- HW\_CPU\_X86\_TSX
- HW\_CPU\_X86\_ASF
- HW\_CPU\_X86\_VMX
- HW\_CPU\_X86\_SVM
- HW\_CPU\_X86\_PDPE1GB
- HW\_CPU\_X86\_STIBP
- HW\_CPU\_X86\_AMD\_SEV
- HW\_CPU\_X86\_AMD\_SEV\_ES
- HW\_CPU\_X86\_AMD\_SVM
- HW\_CPU\_X86\_AMD\_IBPB
- HW\_CPU\_X86\_AMD\_NO\_SSB
- HW\_CPU\_X86\_AMD\_SSBD

- HW\_CPU\_X86\_AMD\_VIRT\_SSBD
- HW\_CPU\_X86\_INTEL\_MD\_CLEAR
- HW\_CPU\_X86\_INTEL\_PCID
- HW\_CPU\_X86\_INTEL\_SPEC\_CTRL
- HW\_CPU\_X86\_INTEL\_SSBD
- HW\_CPU\_X86\_INTEL\_VMX
- HW\_GPU\_API\_DIRECTX\_V10
- HW\_GPU\_API\_DIRECTX\_V11
- HW\_GPU\_API\_DIRECTX\_V12
- HW\_GPU\_API\_DIRECT2D
- HW\_GPU\_API\_DIRECT3D\_V6\_0
- HW\_GPU\_API\_DIRECT3D\_V7\_0
- HW\_GPU\_API\_DIRECT3D\_V8\_0
- HW\_GPU\_API\_DIRECT3D\_V8\_1
- HW\_GPU\_API\_DIRECT3D\_V9\_0
- HW\_GPU\_API\_DIRECT3D\_V9\_0B
- HW\_GPU\_API\_DIRECT3D\_V9\_0C
- HW\_GPU\_API\_DIRECT3D\_V9\_0L
- HW\_GPU\_API\_DIRECT3D\_V10\_0
- HW\_GPU\_API\_DIRECT3D\_V10\_1
- HW\_GPU\_API\_DIRECT3D\_V11\_0
- HW\_GPU\_API\_DIRECT3D\_V11\_1
- HW\_GPU\_API\_DIRECT3D\_V11\_2
- HW\_GPU\_API\_DIRECT3D\_V11\_3
- HW\_GPU\_API\_DIRECT3D\_V12\_0
- HW\_GPU\_API\_VULKAN
- HW\_GPU\_API\_DXVA
- HW\_GPU\_API\_OPENCL\_V1\_0
- HW\_GPU\_API\_OPENCL\_V1\_1
- HW\_GPU\_API\_OPENCL\_V1\_2
- HW\_GPU\_API\_OPENCL\_V2\_0
- HW\_GPU\_API\_OPENCL\_V2\_1
- HW\_GPU\_API\_OPENCL\_V2\_2
- HW\_GPU\_API\_OPENGL\_V1\_1

- HW\_GPU\_API\_OPENGL\_V1\_2
- HW\_GPU\_API\_OPENGL\_V1\_3
- HW\_GPU\_API\_OPENGL\_V1\_4
- HW\_GPU\_API\_OPENGL\_V1\_5
- HW\_GPU\_API\_OPENGL\_V2\_0
- HW\_GPU\_API\_OPENGL\_V2\_1
- HW\_GPU\_API\_OPENGL\_V3\_0
- HW\_GPU\_API\_OPENGL\_V3\_1
- HW\_GPU\_API\_OPENGL\_V3\_2
- HW\_GPU\_API\_OPENGL\_V3\_3
- HW\_GPU\_API\_OPENGL\_V4\_0
- HW\_GPU\_API\_OPENGL\_V4\_1
- HW\_GPU\_API\_OPENGL\_V4\_2
- HW\_GPU\_API\_OPENGL\_V4\_3
- HW\_GPU\_API\_OPENGL\_V4\_4
- HW\_GPU\_API\_OPENGL\_V4\_5
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V1\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V1\_1
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V1\_2
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V1\_3
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V2\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V2\_1
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V3\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V3\_2
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V3\_5
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V3\_7
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V5\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V5\_2
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V5\_3
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V6\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V6\_1
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V6\_2
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V7\_0
- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V7\_1

- HW\_GPU\_CUDA\_COMPUTE\_CAPABILITY\_V7\_2
- HW\_GPU\_CUDA\_SDK\_V6\_5
- HW\_GPU\_CUDA\_SDK\_V7\_5
- HW\_GPU\_CUDA\_SDK\_V8\_0
- HW\_GPU\_CUDA\_SDK\_V9\_0
- HW\_GPU\_CUDA\_SDK\_V9\_1
- HW\_GPU\_CUDA\_SDK\_V9\_2
- HW\_GPU\_CUDA\_SDK\_V10\_0
- HW\_GPU\_MAX\_DISPLAY\_HEADS\_1
- HW\_GPU\_MAX\_DISPLAY\_HEADS\_2
- HW\_GPU\_MAX\_DISPLAY\_HEADS\_4
- HW\_GPU\_MAX\_DISPLAY\_HEADS\_6
- HW\_GPU\_MAX\_DISPLAY\_HEADS\_8
- HW\_GPU\_RESOLUTION\_W320H240
- HW\_GPU\_RESOLUTION\_W640H480
- HW\_GPU\_RESOLUTION\_W800H600
- HW\_GPU\_RESOLUTION\_W1024H600
- HW\_GPU\_RESOLUTION\_W1024H768
- HW\_GPU\_RESOLUTION\_W1152H864
- HW\_GPU\_RESOLUTION\_W1280H720
- HW\_GPU\_RESOLUTION\_W1280H768
- HW\_GPU\_RESOLUTION\_W1280H800
- HW\_GPU\_RESOLUTION\_W1280H1024
- HW\_GPU\_RESOLUTION\_W1360H768
- HW\_GPU\_RESOLUTION\_W1366H768
- HW\_GPU\_RESOLUTION\_W1440H900
- HW\_GPU\_RESOLUTION\_W1600H900
- HW\_GPU\_RESOLUTION\_W1600H1200
- HW\_GPU\_RESOLUTION\_W1680H1050
- HW\_GPU\_RESOLUTION\_W1920H1080
- HW\_GPU\_RESOLUTION\_W1920H1200
- HW\_GPU\_RESOLUTION\_W2560H1440
- HW\_GPU\_RESOLUTION\_W2560H1600
- HW\_GPU\_RESOLUTION\_W3840H2160

- HW\_GPU\_RESOLUTION\_W7680H4320
- HW\_NIC\_SRIOV
- HW\_NIC\_MULTIQUEUE
- HW\_NIC\_VMDQ
- HW\_NIC\_PROGRAMMABLE\_PIPELINE
- HW\_NIC\_ACCEL\_SSL
- HW\_NIC\_ACCEL\_IPSEC
- HW\_NIC\_ACCEL\_TLS
- HW\_NIC\_ACCEL\_DIFFIEH
- HW\_NIC\_ACCEL\_RSA
- HW\_NIC\_ACCEL\_ECC
- HW\_NIC\_ACCEL\_LZS
- HW\_NIC\_ACCEL\_DEFLATE
- HW\_NIC\_DCB\_PFC
- HW\_NIC\_DCB\_ETS
- HW\_NIC\_DCB\_QCN
- HW\_NIC\_OFFLOAD\_TSO
- HW\_NIC\_OFFLOAD\_GRO
- HW\_NIC\_OFFLOAD\_GSO
- HW\_NIC\_OFFLOAD\_UFO
- HW\_NIC\_OFFLOAD\_LRO
- HW\_NIC\_OFFLOAD\_LSO
- HW\_NIC\_OFFLOAD\_TCS
- HW\_NIC\_OFFLOAD\_UCS
- HW\_NIC\_OFFLOAD\_SCS
- HW\_NIC\_OFFLOAD\_L2CRC
- HW\_NIC\_OFFLOAD\_FDF
- HW\_NIC\_OFFLOAD\_RXVLAN
- HW\_NIC\_OFFLOAD\_TXVLAN
- HW\_NIC\_OFFLOAD\_VXLAN
- HW\_NIC\_OFFLOAD\_GRE
- HW\_NIC\_OFFLOAD\_GENEVE
- HW\_NIC\_OFFLOAD\_TXUDP
- HW\_NIC\_OFFLOAD\_QINQ



- HW\_NIC\_OFFLOAD\_RDMA
- HW\_NIC\_OFFLOAD\_RXHASH
- HW\_NIC\_OFFLOAD\_RX
- HW\_NIC\_OFFLOAD\_TX
- HW\_NIC\_OFFLOAD\_SG
- HW\_NIC\_OFFLOAD\_SWITCHDEV
- HW\_NIC\_SRIOV\_QOS\_TX
- HW\_NIC\_SRIOV\_QOS\_RX
- HW\_NIC\_SRIOV\_MULTIQUEUE
- HW\_NIC\_SRIOV\_TRUSTED
- HW\_NUMA\_ROOT
- HW\_PCI\_LIVE\_MIGRATABLE
- HW\_PCI\_ONE\_TIME\_USE
- MISC\_SHARES\_VIA\_AGGREGATE
- OWNER\_CYBORG
- OWNER\_NOVA
- STORAGE\_DISK\_HDD
- STORAGE\_DISK\_SSD