

OpenNebula - Feature #3028

support shared vgpu, eg: nvidia GRID

07/10/2014 07:49 AM - ☐ ☐ ☐

Status:	Closed	Start date:	07/10/2014
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Core & System	Estimated time:	0.00 hour
Target version:	Release 4.14	Pull request:	
Resolution:	fixed		
Description			
<p>Nvidia has released GRID, it can be shared by more than one vm. XenDesktop and WindowsServer has supported it.</p> <p>Can opennebula integrate this feature?</p>			

Associated revisions

Revision 6b8a16fc - 08/17/2015 09:14 PM - Ruben S. Montero

feature #3028: Adds PCI_DEVICES to HostShare

Revision 70433a6d - 08/20/2015 04:20 PM - Javi Fontan

feature #3028: kvm probe to get PCI devices

Revision 00427a09 - 08/20/2015 09:22 PM - Ruben S. Montero

feature #3028: Fix PCI set bugs

Revision 67ec62e0 - 08/20/2015 11:08 PM - Ruben S. Montero

feature #3028: Add PCI capacity when deploying a VM

Revision 17f2380e - 08/21/2015 08:46 AM - Ruben S. Montero

feature #3028: Fix PCI device assignment

Revision c4c64164 - 08/21/2015 10:00 AM - Ruben S. Montero

feature #3028: Delete pci devices from host when removing the VM. Always include VMID in PCI elements

Revision 6d8ea1a3 - 08/21/2015 11:46 AM - Ruben S. Montero

feature #3028: Print HostShare and HostXML into streams. Scheduler now parses PCI device lists

Revision 371a96c5 - 08/21/2015 11:48 AM - Javi Fontan

feature #3028: show pci devices in onehost show

Revision 2b7a2e99 - 08/21/2015 02:11 PM - Javi Fontan

feature #3028: add pci passthrough to deployment file

Revision b1baa06f - 08/21/2015 03:25 PM - Ruben S. Montero

feature #3028: Scheduler check and set PCI device requests

Revision 17c33241 - 08/22/2015 03:00 PM - Ruben S. Montero

feature #3028: Add comments to PCI format. Init pci list in get_* methods

Revision a2e7ccff - 08/25/2015 08:43 AM - Carlos Martín

Feature #3028: Add PCI_DEVICES to onedb migrator

Revision 18946ff3 - 08/25/2015 11:17 AM - Carlos Martín

Feature #3028: Add PCI dev checks to onedb fsck

Revision 415e0b6c - 08/26/2015 09:40 AM - Javi Fontan

feature #3028: check all needed pci params for deployment file

Revision a0a9504f - 08/26/2015 09:43 AM - Carlos Martín

Feature #3028: Add PCI info to im_dummy driver

Revision 991b6719 - 08/26/2015 09:43 AM - Carlos Martín

Feature #3028: Fix sched log formatting

Revision e348918e - 08/26/2015 10:34 AM - Carlos Martín

Feature #3028: Check pci devices in Host::test_capacity

Revision a34b2e27 - 08/27/2015 10:10 AM - Carlos Martín

Feature #3028: Add pci devices to sunstone

Revision 6c708296 - 08/27/2015 11:03 AM - Carlos Martín

Feature #3028: Set maxlength for pci inputs

Revision 6757fa66 - 08/27/2015 02:20 PM - Ruben S. Montero

feature #3028: Remove devices not shown in monitor from host. Recover constness of test method. Get rid of unneeded methods

Revision f2eb5ede - 08/27/2015 02:24 PM - Ruben S. Montero

feature #3028: Keep used PCI devices when updating monitoring

Revision 5c7e6291 - 08/27/2015 02:35 PM - Javi Fontan

feature #3028: tidy up pci devices output

Revision 6dd908ba - 08/27/2015 02:55 PM - Carlos Martín

Feature #3028: Fix pci table for 1 element, remove table columns

Revision 8e8da326 - 09/01/2015 07:42 PM - Ruben S. Montero

feature #3028: Make PCI restricted attribute for VMs

Revision 057063c2 - 09/02/2015 09:21 AM - Javi Fontan

feature #3028: bug assigning pci devices to VMs

Revision fe479a30 - 09/02/2015 09:27 AM - Carlos Martín

Feature #3028: Update host xsd with pci_devices

Revision 8178f90f - 09/04/2015 11:22 AM - Carlos Martín

Feature #3028: Ignore cpu,mem when capacity is not enforced

Revision 4125ef4c - 09/04/2015 01:12 PM - Ruben S. Montero

feature #3028: Update HostShare::get_pci_value check for empty and wrong values

Revision 19f3e879 - 09/07/2015 09:58 AM - Ruben S. Montero

feature #3028: Fix get_pci_value logic

Revision 24639575 - 09/08/2015 09:06 AM - Carlos Martín

Feature #3028: Another fix for get_pci_value

History

#1 - 07/10/2014 09:27 AM - Ruben S. Montero

- *Tracker changed from Feature to Request*

This depends on the underlying hypervisor, moving it to request to find out how this can be configured.

#2 - 07/16/2014 05:01 AM -

Now a days, physical servers are coming up with graphic cards that have multiple GPUs, VMs running on cloud can leverage the high computation power of GPU to meet high graphics processing requirements like Auto-CAD, Photoshop etc. Also, there are cards in market which supports sharing of a GPU cards among multiple VMs by creating vGPUs for each VM, e.g. NVIDIA has introduced vGPU capable cards GRID K1 and K2 which allow multiple vGPUs on a single physical GPU.

With vGPU technology, the graphics commands of each virtual machine are passed directly to the underlying dedicated GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced and shared across multiple VMs.

Xenserver has added the support for NVIDIA GRID K1 and GRID K2 cards. It allows the VMs on XenServer hosts to use the GPU cards in following ways:

GPU-passthrough: It allows the hypervisor to assign the entire PGPU to a VM, this is useful for a power users.

VGPU: It allows the VM to share a PGPU device with other VMs, this is useful for tier 2 users.

#3 - 07/16/2014 05:01 AM -

<https://cwiki.apache.org/confluence/display/CLOUDSTACK/GPU+and+vGPU+support+for+CloudStack+Guest+VMs>

#4 - 09/04/2015 11:57 AM - Carlos Martin

- *Tracker changed from Request to Feature*

- *Category set to Core & System*

- *Status changed from Pending to Closed*

- *Target version set to Release 4.14*

- *Resolution set to fixed*