

## OpenNebula - Backlog #2869

### Add the ability to configure the CPU model exposed to KVM virtual machines.

05/06/2014 06:02 AM - Vladislav Gorbunov

|   |               |                        |                   |
|---|---------------|------------------------|-------------------|
| <b>Status:</b>  | Pending       | <b>Start date:</b>     | 05/06/2014        |
| <b>Priority:</b>  | Normal        | <b>Due date:</b>       |                   |
| <b>Assignee:</b>  |               | <b>% Done:</b>         | 0%                |
| <b>Category:</b>  | Core & System | <b>Estimated time:</b> | 0.00 hour         |
| <b>Target version:</b>  |               |                        |                   |
| <b>Description</b>  |               |                        |                   |
| <p>One of the most important missing features was the ability to configure the CPU model exposed to KVM virtual machines. There are a couple of reasons for wanting to specify the CPU model</p> <ul style="list-style-type: none"><li>- To maximise performance of virtual machines by exposing new host CPU features to the guest</li><li>- To ensure a consistent default CPU across all machines, removing reliance of variable QEMU defaults.</li></ul> <p>In libvirt, the CPU is specified by providing a base CPU model name (which is a shorthand for a set of feature flags), a set of additional feature flags, and the topology (sockets/cores/threads). The libvirt KVM driver provides a number of standard CPU model names (defined in /usr/share/libvirt/cpu_map.xml):</p> <p>"486", "pentium", "pentium2", "pentiumpro", "coreduo", "n270", "pentiumpro", "qemu32", "kvm32", "cpu64-rhel5", "cpu64-rhel5", "kvm64", "pentiumpro", "Conroe", "Penryn", "Nehalem", "Westmere", "pentiumpro", "cpu64-rhel5", "cpu64-rhel5", "Opteron_G1", "Opteron_G2", "Opteron_G3, "Opteron_G4"</p> <p>It is also possible to request the host CPU model in two ways</p> <ul style="list-style-type: none"><li>- "host-model" - this causes libvirt to identify the named CPU model which most closely matches the host from the above list, and then request additional CPU flags to complete the match. This should give close to maximum functionality/performance, which maintaining good reliability/compatibility if the guest is migrated to another host with slightly different host CPUs.</li><li>- "host-passthrough" - this causes libvirt to tell KVM to passthrough the host CPU with no modifications. The difference to host-model, instead of just matching feature flags, every last detail of the host CPU is matched. This gives absolutely best performance, and can be important to some apps which check low level CPU details. The guest can only be migrated to an exactly matching host CPU.</li></ul> <p>This path is add CPUMODEL support:<br/>OS = [ CPUMODEL="host" ] for host-passthrough<br/>or for example<br/>OS = [ CPUMODEL="core2duo" ] for custom CPU model. Full list of supported cpu shown by the command:<br/>kvm -cpu ?</p> |               |                        |                   |
| <b>Related issues:</b>  |               |                        |                   |
| Related to Backlog # 796: Include cpu tune & topology parameters for qemu/lib...  |               | <b>Pending</b>         | <b>09/06/2011</b> |
| Related to Bug # 3459: OpenNebula generates template with 12 sockets instead ...  |               | <b>Closed</b>          | <b>12/16/2014</b> |
| Duplicated by Backlog # 3026: support CPU model and topology  |               | <b>Closed</b>          | <b>07/10/2014</b> |

## History

### #1 - 05/06/2014 09:18 AM - Ruben S. Montero

- Category set to Core & System
- Status changed from Pending to New

Thanks for the patch! :) Scheduling this 4.8

### #2 - 06/24/2014 02:13 AM - Ruben S. Montero

- Assignee set to Javi Fontan

**#3 - 06/27/2014 10:02 AM - Ruben S. Montero**

- Tracker changed from Feature to Backlog
- Status changed from New to Pending
- Assignee deleted (Javi Fontan)
- Target version deleted (Release 4.8)

**#4 - 07/10/2014 09:22 AM - Ruben S. Montero**

- Duplicated by Backlog #3026: support CPU model and topology added

**#5 - 12/29/2014 04:17 PM - Ruben S. Montero**

- Related to Backlog #796: Include cpu tune & topology parameters for qemu/libvirt drivers added

**#6 - 12/29/2014 04:19 PM - Ruben S. Montero**

- Related to Bug #3459: OpenNebula generates template with 12 sockets instead of 12 cores added

**#7 - 07/19/2015 04:26 PM - Eugene W**

Is it working now in OpenNebula 4.12?

I tried to add

```
OS = [ boot = "hd", arch = "x86_64", CPUMODEL="host", machine = "q35" ]
```

into

```
/etc/one/vmm_exec_kvm.conf
```

But it is can not help (I tried make a new template).

**#8 - 07/19/2015 04:59 PM - Eugene W**

Have found solution for CPU host-passthrough.

If need to to expose the underlying physical CPU to the virtual machine.

To do this, you need to modify your VM's template to include the following raw attributes:

```
RAW=[ DATA="<cpu mode='host-passthrough'/>", TYPE="kvm" ]
```

You can add these attributes by following these steps:

1. Choose your VM template
2. Click on 'UPDATE'
3. Click on 'Other'
4. From the 'RAW data' dropdown menu, choose 'KVM'
5. In the DATA text box, insert:

```
"<cpu mode='host-passthrough'/>"
```

**Files**

---

cpumodel.diff

1.2 KB

05/06/2014

Vladislav Gorbunov