

## Open Infrastructure - Task #12343

### Evaluate cloud-hypervisor (Q12024)

01/07/2024 04:23 PM - Nico Schottelius

<b>Status:</b>	In Progress	<b>Start date:</b>	01/07/2024
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Nico Schottelius	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>			
<b>PM Check date:</b>			
<b>Description</b>			
<ul style="list-style-type: none"><li>• Might be a lightweight option for running in k8s</li><li>• Requires quite some work around it</li><li>• No management tooling</li><li>• Storage<ul style="list-style-type: none"><li>◦ For ceph probably need to use RBD mapped, rook supports that</li><li>◦ For creating thin provisioning, probably need to create a wrapper/controller</li></ul></li><li>• live migration<ul style="list-style-type: none"><li>◦ seems to be supported</li></ul></li><li>• Networking unclear<ul style="list-style-type: none"><li>◦ macvtap support, see below</li></ul></li></ul>			
<pre>"tap=&lt;if_name&gt;,ip=&lt;ip_addr&gt;,mask=&lt;net_mask&gt;,mac=&lt;mac_addr&gt;,fd=&lt;fd1:fd2...&gt;,iommu=on off,num_queues=&lt;number_of_queues&gt;,queue_size=&lt;size_of_each_queue&gt;,id=&lt;device_id&gt;,vhost_user=&lt;vhost_user_enable&gt;,socket=&lt;vhost_user_socket_path&gt;,vhost_mode=client server,bw_size=&lt;bytes&gt;,bw_one_time_burst=&lt;bytes&gt;,bw_refill_time=&lt;ms&gt;,ops_size=&lt;io_ops&gt;,ops_one_time_burst=&lt;io_ops&gt;,ops_refill_time=&lt;ms&gt;"</pre>			
<b>live migration</b>			
<ul style="list-style-type: none"><li>• examples use same host migration</li></ul>			
<pre>% cloud-hypervisor \   --kernel ./hypervisor-fw \   --disk path=focal-server-cloudimg-amd64.raw \   --cpus boot=2 \   --memory size=1024M \   --net "tap=,mac=,ip=,mask=" --api-socket=/tmp/apil</pre>			
<pre>[17:42] nb3:~% cloud-hypervisor --api-socket=/tmp/apil2</pre>			
<pre># receive VM ch-remote --api-socket=/tmp/apil2 receive-migration unix:/tmp/sock</pre>			
<pre># send VM - fails % ch-remote --api-socket=/tmp/apil1 send-migration --local unix:/tmp/sock Error running command: Server responded with an error: InternalServerError: ApiError(VmSendMigration(MigrateSend(Local migration requires shared memory or hugepages enabled)))</pre>			
<b>Sketch for running VMs with cloud-hypervisor</b>			
<ul style="list-style-type: none"><li>• Manage networking outside<ul style="list-style-type: none"><li>◦ pod running potentially in hostnetwork</li><li>◦ creating bridge depending on which customer it is</li><li>◦ Potentially running IPAM on a per customer basis</li><li>◦ Could potentially utilise netbox as a backend, but needs to be written</li></ul></li><li>• Console access<ul style="list-style-type: none"><li>◦ read only via pod</li><li>◦ serial forwarding unclear</li></ul></li><li>• Disk management<ul style="list-style-type: none"><li>◦ Thin provisioning / templates needs to be built</li></ul></li></ul>			

## History

---

**#1 - 01/07/2024 04:23 PM - Nico Schottelius**

- Description updated

**#2 - 01/07/2024 04:27 PM - Nico Schottelius**

### running simple VM

```
cloud-hypervisor \  
  --kernel ./hypervisor-fw \  
  --disk path=focal-server-cloudimg-amd64.raw \  
  --cpus boot=2 \  
  --memory size=1024M \  
  --net "tap=,mac=,ip=,mask="
```

- Comes up with a login after a few seconds

**#3 - 01/07/2024 04:28 PM - Nico Schottelius**

- Description updated

**#4 - 01/07/2024 04:33 PM - Nico Schottelius**

- Description updated

**#5 - 01/07/2024 04:36 PM - Nico Schottelius**

- Description updated

**#6 - 01/07/2024 04:40 PM - Nico Schottelius**

- Description updated

**#7 - 01/07/2024 04:41 PM - Nico Schottelius**

- Description updated

**#8 - 01/07/2024 04:52 PM - Nico Schottelius**

- Description updated

**#9 - 01/07/2024 05:33 PM - Nico Schottelius**

- Project changed from 45 to Open Infrastructure