

Optimizing Virtual Machines

KVM and libvirt

Toshaan Bharvani - VanTosh bvba

<toshaan@vantosh.com>



FroSCon 2013



24 August 2013

\$ whoami

Toshaan Bharvani

- From Antwerp, Belgium
- Currently self-employed : VanTosh
<http://www.vantosh.com>
- Involved with Enterprise Linux : RHEL, CentOS, SLES, ...
- Likes to keep everything secure : SELinux, Web, ...
- Lives in a virtual world : KVM, Xen, LXC, ...
- Likes automation : Ansible
- Works on both hardware and software side
- Wants to take over the world
- Twitter : [@toshywoshy](#)
- Blog : <http://www.toshaan.com>

Table of contents

1 Introduction

2 CPU

3 Memory

4 IO

5 Network

6 Conclusion

Introduction

CPU

Memory

IO

Network



Conclusion

The End

1

Introduction

Architecture



#170752

Internals

Introduction

CPU

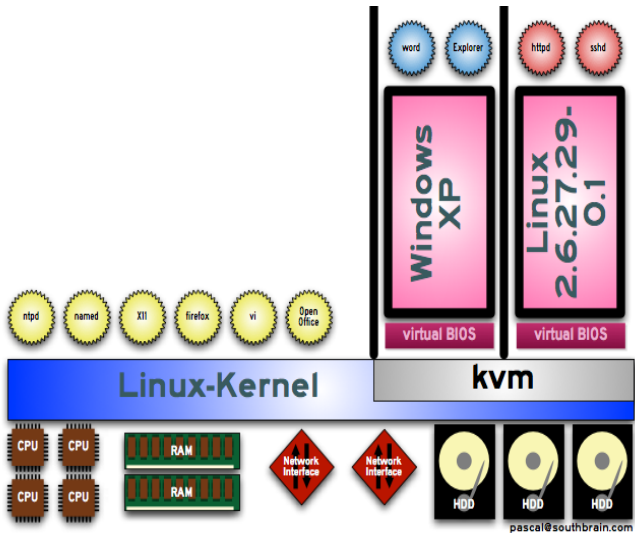
Memory

IO

Network

Conclusion

The End



pascal@southbrain.com

Introduction

CPU

Memory

IO

Network

Conclusion

The End

2

CPU

CPU optimization (1)

- Check hyperthreading and verify logical mapping
- Thin provision CPU

```
1 <vcpu placement='static' current="2">8</vcpu>
```

- Set the right CPU type

```
1 <cpu mode='custom' match='exact'>  
2   <model fallback='allow'>Penryn</model>  
3   <vendor>Intel</vendor>  
4   <topology sockets='4' cores='4' threads='2'>  
5 </cpu>  
6  
7 <cpu mode='host-passthrough'>  
8  
9 <cpu mode='host-model'>  
10   <model fallback='forbid'>  
11   <topology sockets='1' cores='2' threads='1'>  
12 </cpu>
```


CPU optimization (2)

Introduction

CPU

Memory

IO

Network

Conclusion

The End

● Pin CPU

```
1 <cputune>
2 <vcputune vcpu="0" cpuset="1-4,^2" />
3 <vcputune vcpu="1" cpuset="0,1" />
4 <vcputune vcpu="2" cpuset="2,3" />
5 <vcputune vcpu="3" cpuset="0,4" />
6 <emulatorpin cpuset="1-3" />
7 <shares>2048</shares>
8 <period>1000000</period>
9 <quota>-1</quota>
10 <emulator_period>1000000</emulator_period>
11 <emulator_quota>-1</emulator_quota>
12 </cputune>
```

● virt-install

```
1 virt-install \
2   --vcpus=8,maxvcpus=4,sockets=2,cores=2,threads=2
3   --cpuset=0,2,^3,5
4   --cpu core2duo,+x2apic,disable=vmx
5   --cpu host
```

Introduction

CPU

Memory

IO

Network

Conclusion

The End

3

Memory

Memory

- Overcommitting memory (aka ballooning)

```
1 <memory unit='KiB'>2097152</memory>  
2 <currentMemory unit='KiB'>524288</currentMemory>
```

- Backing Pages

```
1 <memoryBacking>  
2   <hugepages/>  
3   <nosharepages/>  
4   <locked/>  
5 </memoryBacking>
```

- Memory Tuning

```
1 <memtune>  
2   <hard_limit unit='G'>1</hard_limit>  
3   <soft_limit unit='M'>128</soft_limit>  
4   <swap_hard_limit unit='G'>2</swap_hard_limit>  
5   <min_guarantee unit='bytes'>67108864</min_guarantee>  
6 </memtune>
```

Introduction

CPU

Memory

IO

Network

Conclusion

The End

4

IO

IO (3)

- Use the VirtIO driver for native speed
- Use the writecache and threading

```
1 <disk type='file' device='disk'>
2   <driver name='qemu' type='qcow2' cache='writeback' io='
   threads' />
3   <source file='/virtual/builder6-boot.qcow2' />
4   <target dev='vda' bus='virtio' />
5 </disk>
6 <disk type='block' device='lun'>
7   <driver name='qemu' type='raw' />
8   <source dev='/dev/sda' />
9   <target dev='sda' bus='virtio' />
10  <address type='drive' controller='0' bus='0' target='3'
   unit='0' />
11 </disk>
```

IO (2)

● Block IO Priority

```
1 <blkio tune>
2   <weight>800</weight>
3   <device>
4     <path>/dev/sda</path>
5     <weight>1000</weight>
6   </device>
7   <device>
8     <path>/dev/sdb</path>
9     <weight>500</weight>
10  </device>
11 </blkio tune>
```

● virt-install

```
1 virt-install --disk \
2   path=/path/to/file ,
3   device=disk ,
4   size=11,
5   cache=writeback ,
6   format=qcow2 ,
7   io=threads ,
8   bus=virtio
```

Introduction

CPU

Memory

IO

Network

Conclusion



The End

5

Network

Network (1)

Topology : Bridges vs vSwitch



Network (2)

- Use the VirtIO driver for native speed (upto 10Gbps)

```
1 <devices>
2   <interface type='bridge'>
3     <mac address='52:54:00:dc:fa:29' />
4     <source bridge='br0' />
5     <model type='virtio' />
6     <address type='pci' domain='0x0000' bus='0x00' slot='0
       x03' function='0x0' />
7   </interface>
8   <interface type='network'>
9     <source network='default' portgroup='engineering' />
10    <model type='virtio' />
11    <target dev='vnet7' />
12    <mac address="00:11:22:33:44:55" />
13    <virtualport>
14      <parameters instanceid='09b11c53-8b5c-4eeb-8f00-
       d84eaa0aaa4f' />
15    </virtualport>
16  </interface>
17 </devices>
```

- virt-install

```
1 virt-install \
2   --network bridge=br0,model=virtio \
3   --network network=virbr0,model=virtio
```

6

Conclusion

Conclusion

Introduction

CPU

Memory

IO

Network

Conclusion

The End

- Use your hardware better
- Optimize for your application/usage
- Increase performance by :
 - CPU : fully optimized = +20%
 - Memory : fully optimized = +5%
 - IO : fully optimized = +15%
 - Networking : fully optimized = +10%

The End



Thank You



Toshaan Bharvani - VanTosh bvba <toshaan@vantosh.com>

<http://www.vantosh.com/publications>

Made with Beamer \LaTeX
a \TeX based Presentation program