


Xen: Tips and Tweaks




SWITCH

Serving Swiss Universities

Alessandra Scicchitano

alessandra.scicchitano@switch.ch

Let's first have a look at a real Xen server

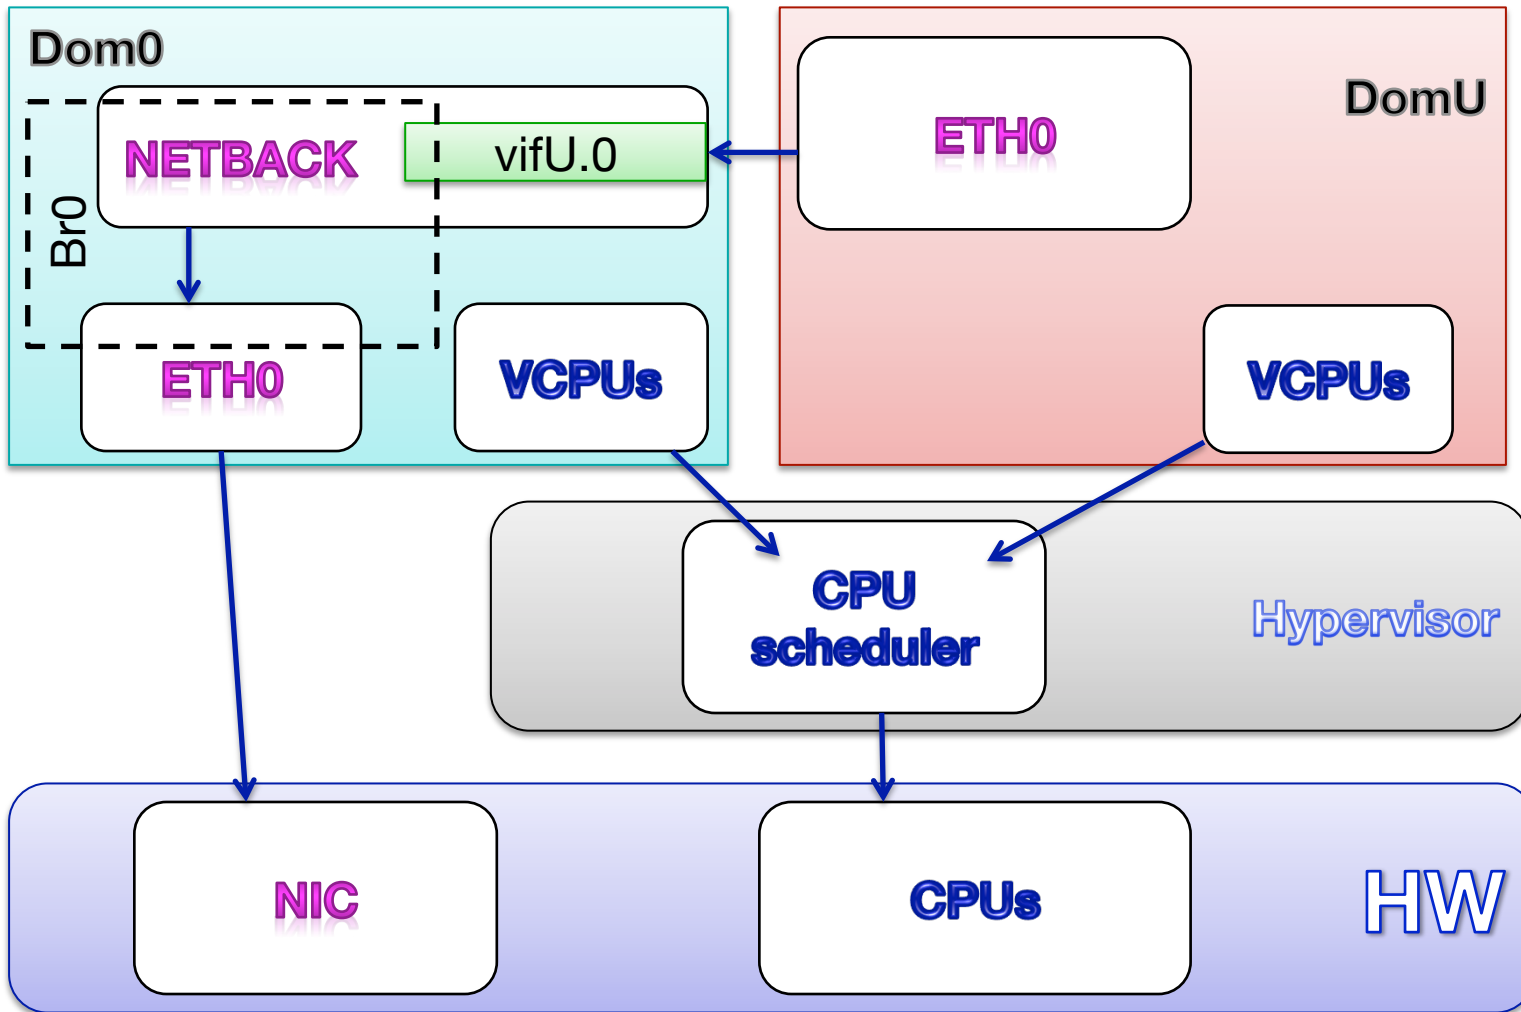


Diagnostic tools

- Performance tab
- Top and Xentop in DOM0
- Top in DOMU

Network testing Tools

- Iperf/nuttcp
- Traffic should be sent from bare metal to VM (on another host) and vice versa.
- Repeat the test more than once



Looking at the architecture:

- CPU
- I/O
- CPU & I/O

Possible cause

- Too many interrupts to handle
- Netback process takes away too much CPU
- No clear reason

CPU Bottleneck

- Most common problem: VCPU capacity
- Xentop + V on dom0: seconds spent in each VM's VCPU
- Assign a VM to a specific CPU

Be careful!!!!

VCPU measurement tools
inside VMs don't give reliable
results!!

Some adjustments

- Enable irqbalance
- Better distribute the load on the different VIFs and netbacks
- Consider assign a dedicated CPU to Dom0
- Tune TCP settings for both Dom0 and DomUs

Suggestion on TCP settings

Bandwidth Delay Product (BDP) = Route Trip Time (RTT) * Theoretical Bandwidth Limit

```
sysctl -w net.core.rmem_max= 'BDP'  
sysctl -w net.core.wmem_max= 'BDP'  
sysctl -w net.ipv4.tcp_rmem="x y 'BDP'"  
sysctl -w net.ipv4.tcp_wmem="x y 'BDP'"  
  
net.ipv4.tcp_congestion_control=cubic
```

Some more...

- Increase vif length

```
ifconfig vif<X>.<Y> txqueuelen 1024
```

- Enable jumbo frame