


Performance and Virtual Machines

What happens when you lose your way...



SWITCH

Serving Swiss Universities

Alessandra Scicchitano

alessandra.scicchitano@switch.ch

The matter

“Virtualization” has become a very used and common word today.

Performance and Virtualization get along just fine(???)

Proactive Monitoring

- CPU Utilization
- Memory Utilization
- I/O latency
- Network

How much can we trust what we see? And how do we handle what we can't see?

Determining where the issue is,
may not be a trivial thing

Real case:
SWITCH DNS Server

SWITCH DNS Server

The most important performance indexes for a DNS Server are:

- a) Number of Queries per second that the server is able to answer without “dropping” (qps)
- b) And the time needed to answer a request

The numbers

Results BIND 9.8.1-P1

	bamus (1 vCPU)			bamus (2 vCPU)			niobe			
	-	+querylog	+querylog, +dsc	-	+querylog	+querylog, +dsc	-	+querylog	+querylog, +dsc	+dsc
manaro	~21'000 qps	~16'000 qps	~10'000 qps	~18'000 qps	~15'000 qps	~10'000 qps	~44'000 qps	~38'000 qps	~28'000 qps	~32'000 qps
asama	~18'500 qps	~15'500 qps	~10'000 qps	~16'000 qps	~14'000 qps	~10'000 qps	~32'000 qps	~31'000 qps	~28'000 qps	~28'000 qps
manaro + asama	~21'000 qps	~16'000 qps	~10'000 qps	~18'000 qps	~15'000 qps	~10'000 qps	~50'000 qps	~42'000 qps	~31'000 qps	~35'000 qps

Terrible qps on the VM versus very good qps on the real server

= The bottleneck is located somewhere between the Hardware Interface of the mother host and the BIND (DNS software) in the Virtual machine.

First Hypothesis

The mother host on which the VM is running has 2 CPUs Dual-Core and 20GB RAM. Not enough.

Suggestion:

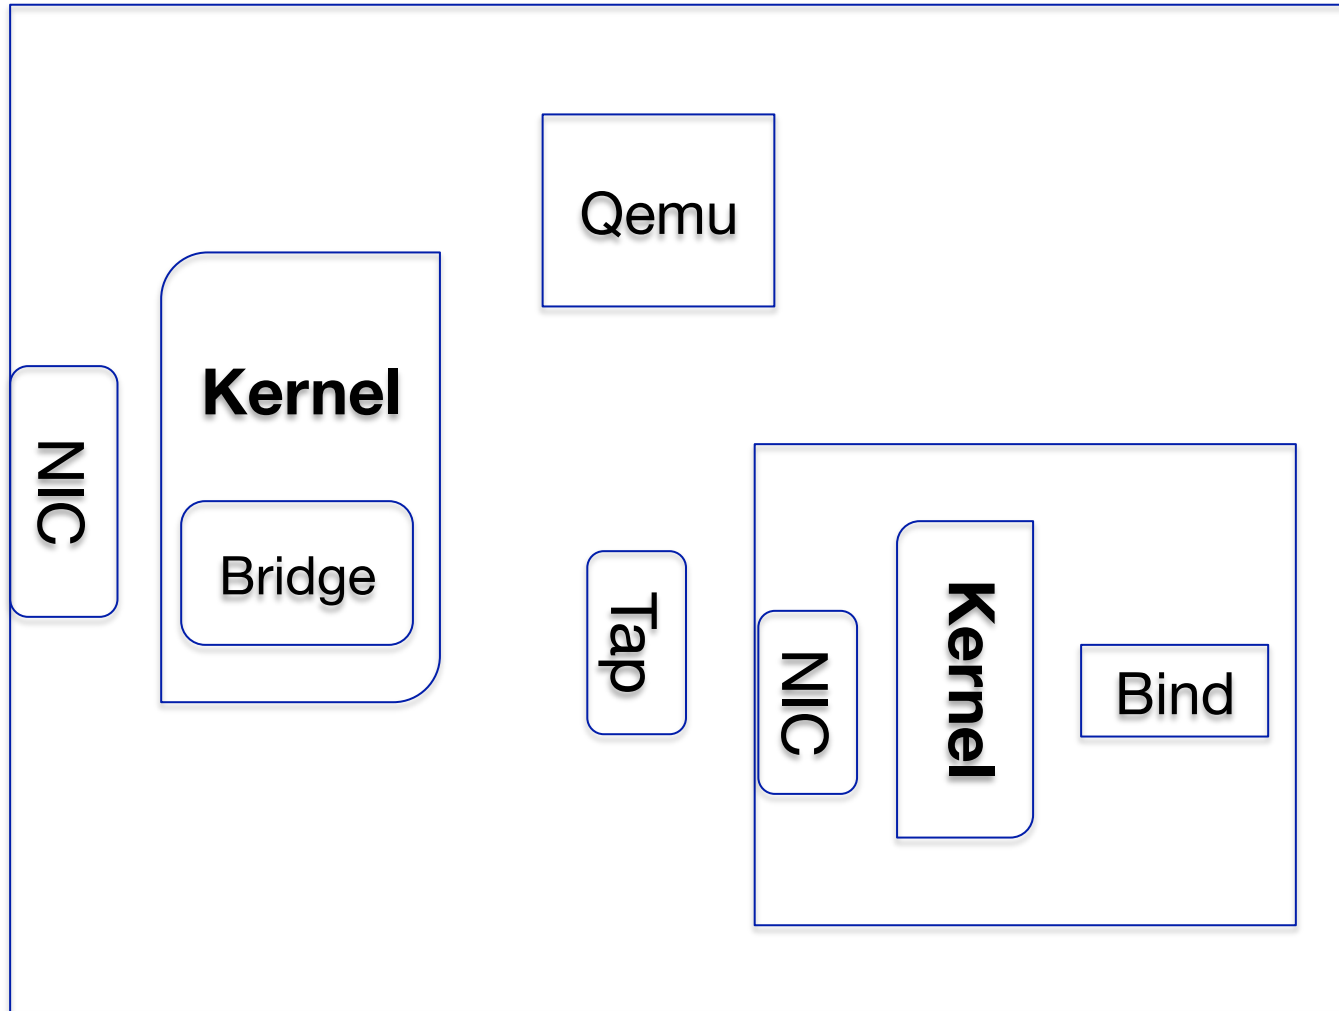
Migrating the VM on a machine with more cores.

Results:

Not as we hoped...

Running the same test on a machine with 24 logical cores showed almost no improvements...

Overview

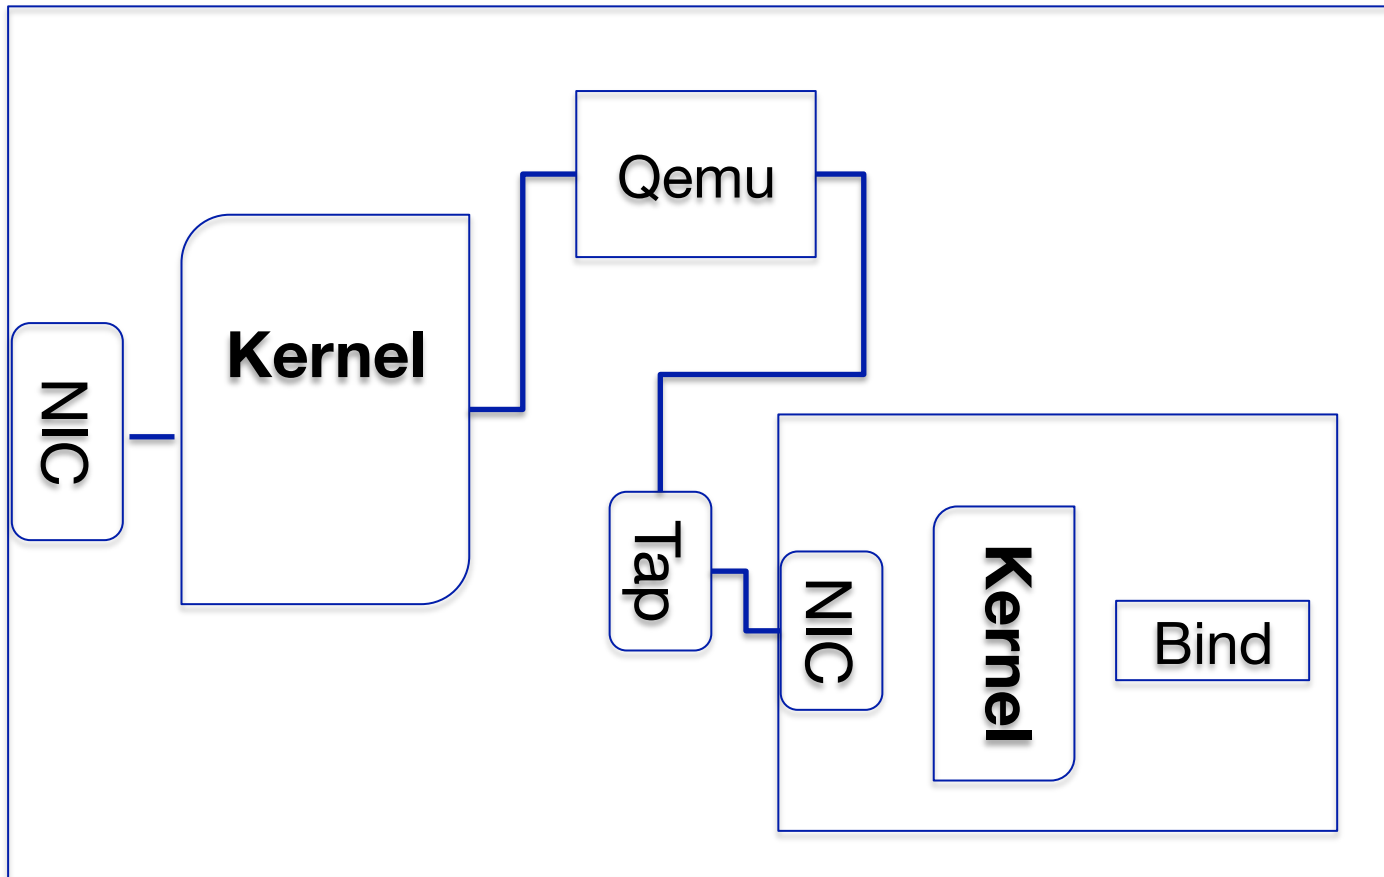


Client Hypervisor

Virtualization techniques can most easily be broken down into three groups:

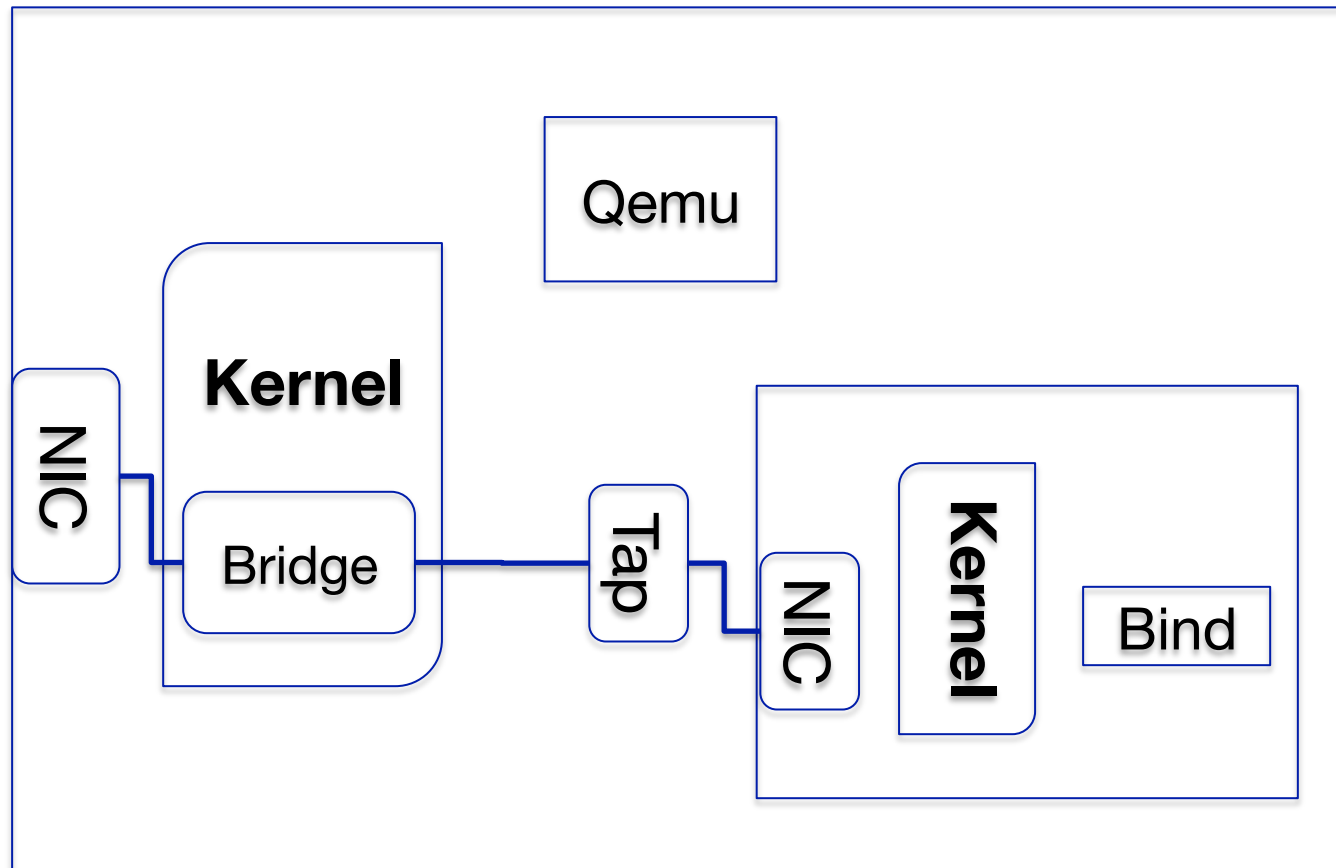
- Emulation
- Paravirtualization
- Hardware pass-through

Emulation



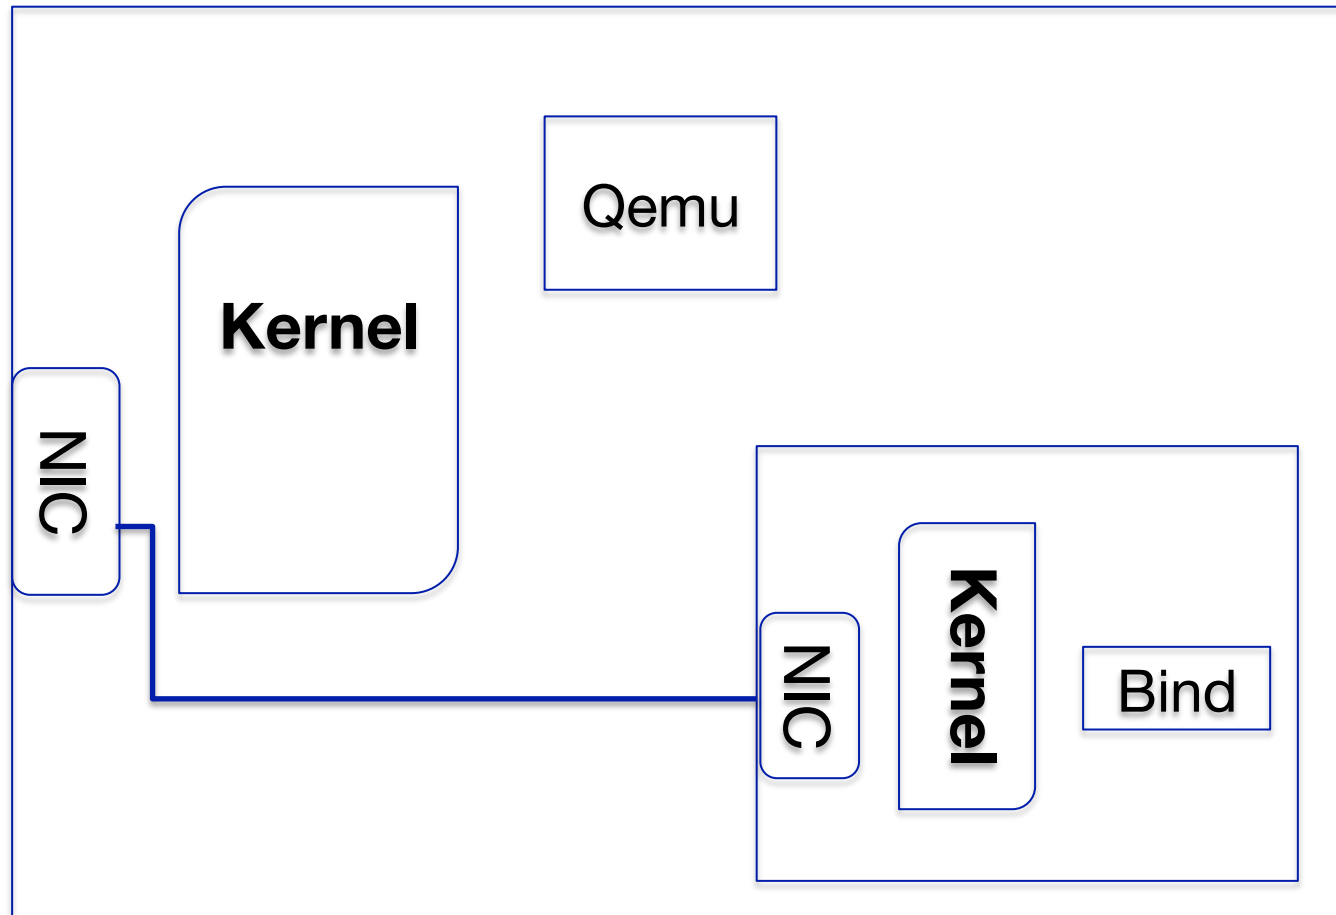
Emulation is where the hypervisor emulates a certain piece of hardware that it presents to the guest VM, regardless of what the actual physical hardware is.

Paravirtualization



Paravirtualization is a technique where the hypervisor exposes a modified version of the physical hardware interface to the guest VM.

Hardware pass-through



Hardware pass-through means that the guest VM has direct access to the physical hardware.

Second Hypothesis

The use of Qemu user-space costs too much...

Suggestion:

Paravirtualization. Use of Virtio-net paravirtual drivers and the linux vhost-net driver.

Results

	simbo (1 vCPU)			simbo (2 vCPU)			simbo (3 vCPU)			
	-	+querylog	+querylog, +dsc	-	+querylog	+querylog, +dsc	-	+querylog	+querylog, +dsc	+dsc
manaro	23'900 qps	16'200 qps	10'700 qps	45'800 qps	28'500 qps	19'500 qps	57'100 qps	36'400 qps	28'300 qps	46'200 qps
manaro + lopevi	27'000 qps	17'700 qps	11'600 qps	48'800 qps	29'500 qps	19'500 qps	73'000 qps	39'000 qps	30'000 qps	50'300 qps

Results

Excellent performance!!

Important: Both host and VM
were running Debian stable
(squeeze) + Backports for Kernel

Conclusion

- Troubleshooting an issue on the way from the mother host to the VM is not an easy thing.
- No monitoring tools available.
- What helps the most is knowledge and experience.