


Virtualization

Alessandra Scicchitano, SWITCH
EGI TF, Madrid
September 2013
alessandra.scicchitano@switch.ch

Virtualization is so popular...



...because it's a wonderful illusion.



Utilization vs Performance



Nothing is perfect, not even magic.

The illusion of having more resources doesn't mean that we have more. It only means that we use more what we have.

Sharing what we have means that sometimes we have to be patient and altruist.

Great performance != altruism

- ◆ CPU, memory, disk storage, and network access are shared among VMs on the same host.
- ◆ From a hardware point of view each VM is perfectly isolated.
- ◆ From a performance point of view...not so much.

Performance in a virtualized environment



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?

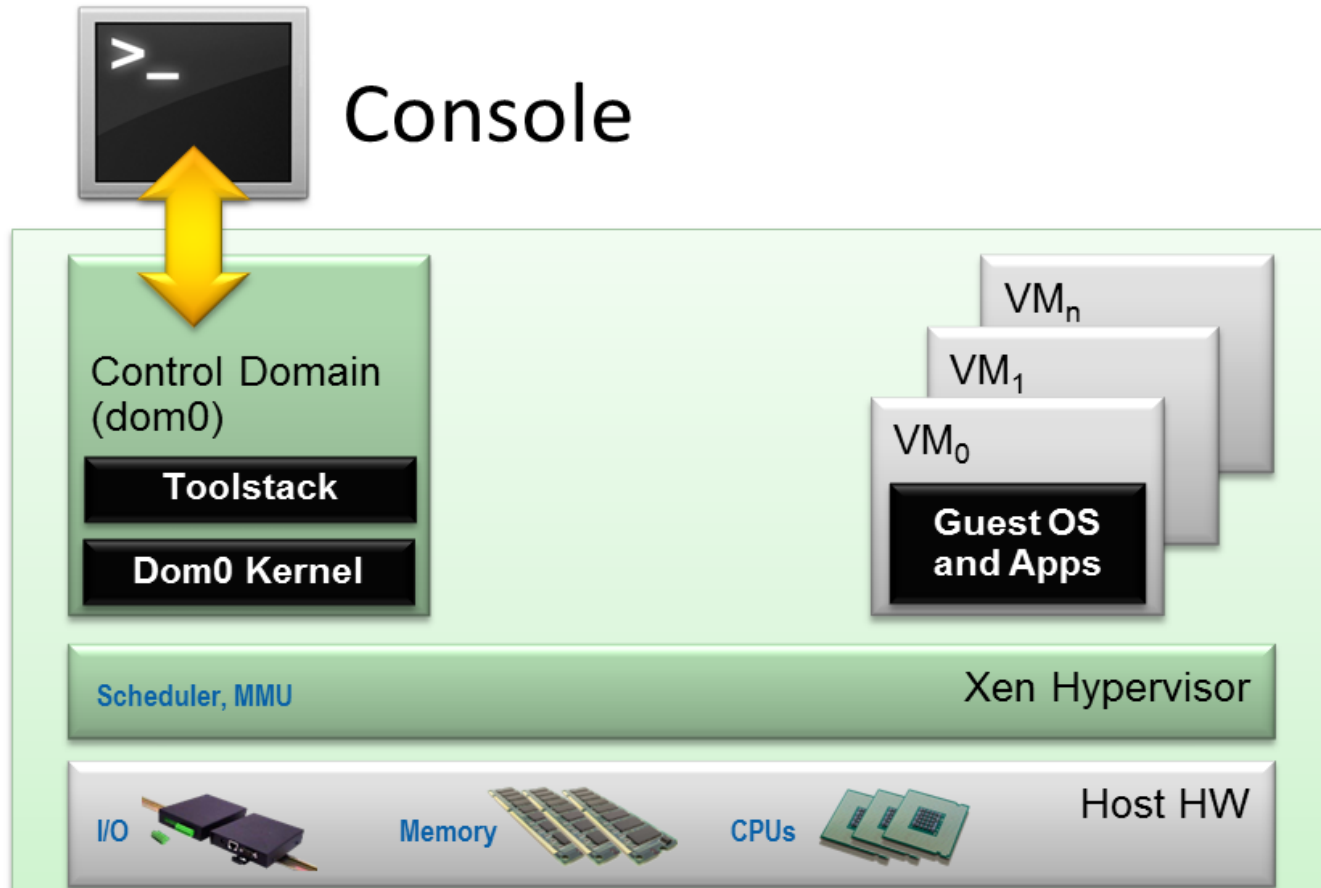
One step back.



- Resource sharing is managed by the hypervisor.
- The main task of the hypervisor is to handle all the incoming requests from the different VMs without creating interference (Is that even possible??)
- The hypervisor represents basically an extra layer in the system.

This is something to keep in mind

An example: XEN



Credit Scheduler:

- Based on *weight* and *cap* (optional) parameters
- VCPUs are labeled *under* if they have existing credits, *over* otherwise.
- Every 30ms the priorities (credits) are recalculated
- Features load balancing of Virtual CPUs across physical ones.

Everything goes fine until...



...the CPUs of the mother host are fully utilized.

When the utilization of the underlying system approaches 100%, VMs will have to wait their turns before getting a slice of CPU



Time keeping



The time keeping problem is very well known in virtualized environment (especially for VMWare)

Many steps have been done in the direction of a better synchronization of the machines (also during suspension)

However time keeping is still an issue when executing performance measurements.


When the physical memory is 100% used and more virtual memory is required, we can have either:

- Swapping: This is done indiscriminately and it leads to serious performance degradation

Or

- Ballooning: It improves resource utilization. It is important not to overcommit too much memory which leads back to swapping. (This might not be true in all virtual solution. Hyper-V for example doesn't support overcommitting)

Memory ballooning?



And?



Ok...everything very nice, but can we just go back to the point?

Yes.

How the hell do we measure performance in a VM?

Fair enough!!

Most of the numbers that we see on a monitor running within the VM can be trust, but that it will give us no idea of what happens on the mother host

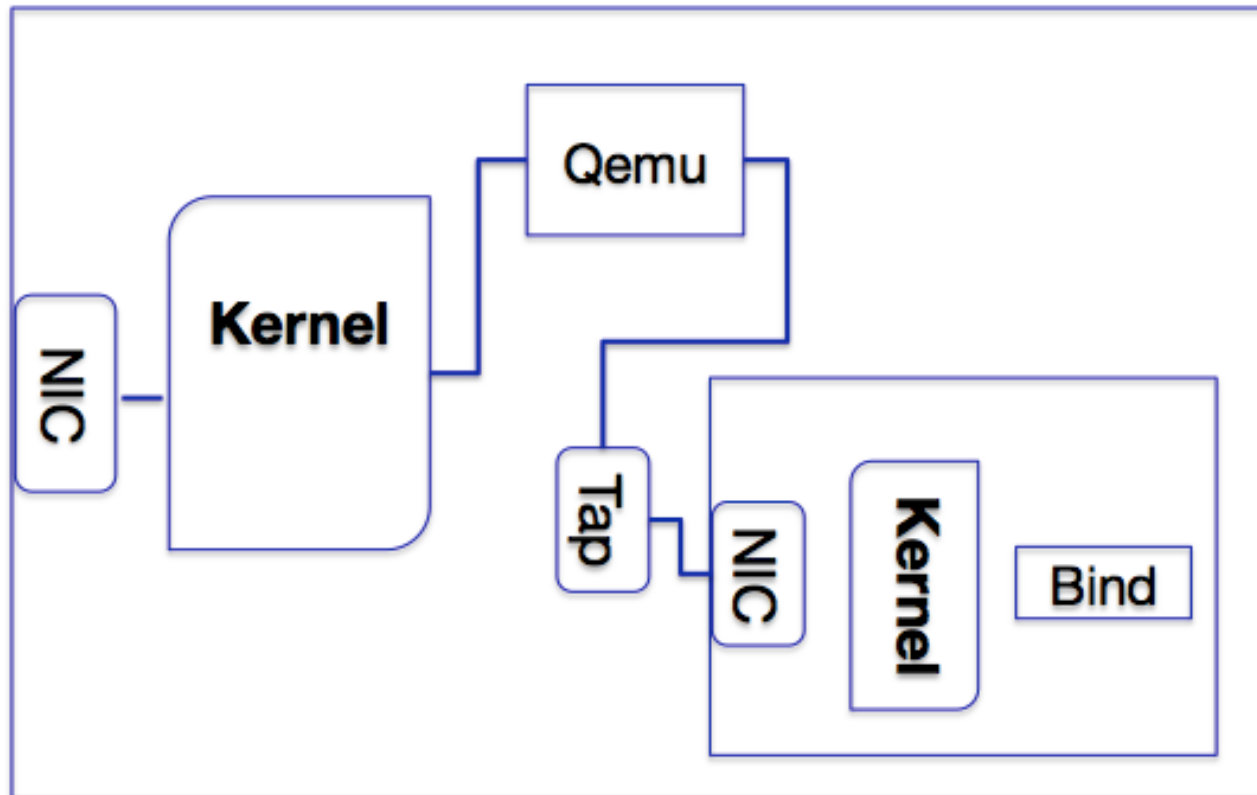
Something tricky to consider

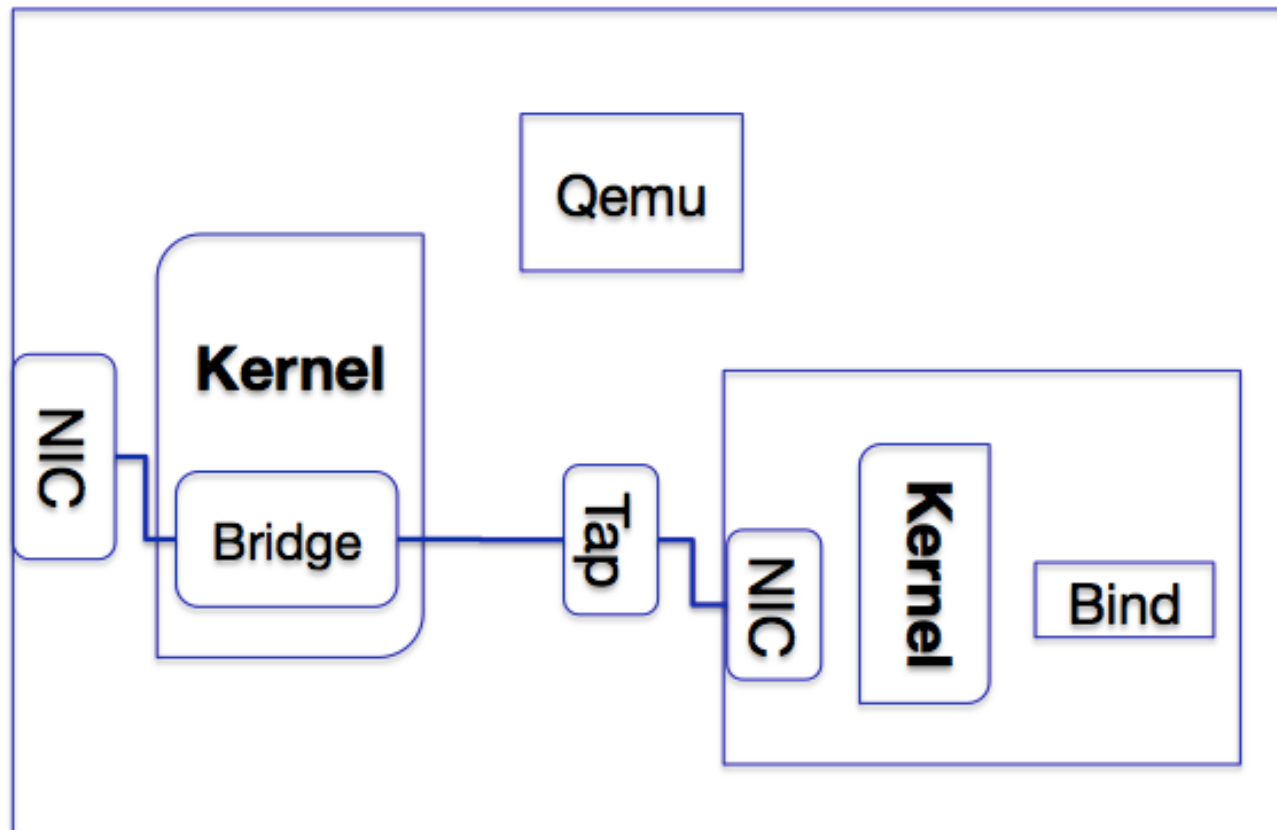


- Because of an application a VM needs more resources
- It is either stalled (has to wait), or it begins to “steal” resources from other VMs running on the same host.
- The other VMs will behave poorly, but they are not the problem.

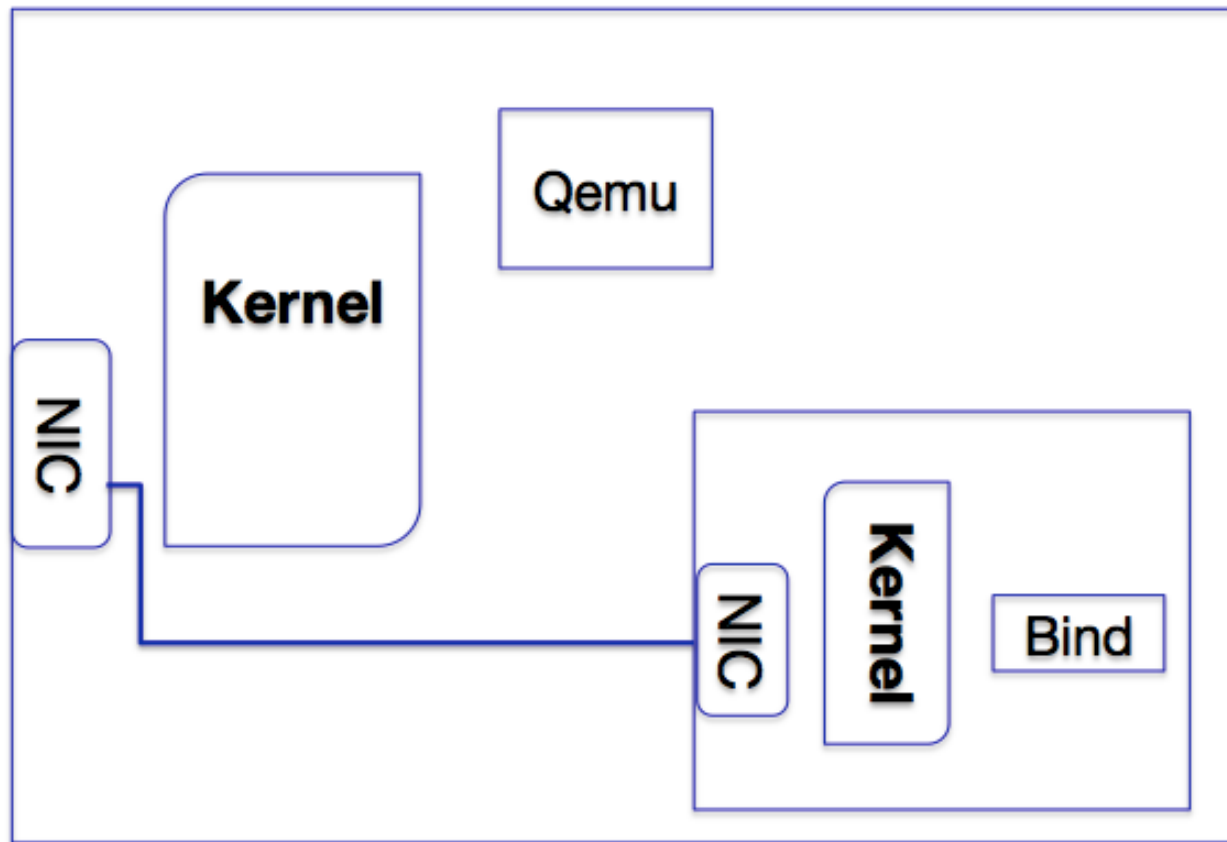
Virtualization techniques can most easily be broken down into three groups:

- Emulation
- Paravirtualization
- Hardware pass-through





Hardware pass-through



Paravirtualization is often the solution to boost performance