

Uses of virtualisation

Ewan Mac Mahon
Oxford

Uses of virtualisation

- Service nodes (so far, just grid nodes)
 - Basic server consolidation argument
 - Grid nodes in particular benefit from isolating services from each other
- Laptops
 - Need to run SL4, but also support modern laptop hardware, especially WiFi and power management.
- Cluster worker nodes
 - Better resource usage between local/grid systems
 - Better job isolation

Grid service nodes

- Most grid service nodes don't need much computational power.
- They do, however, benefit from keeping things as separate as possible (e.g. Not running the site BDII on a CE)
- Service nodes are (mostly) critical, so need to be on decent hardware (e.g. Dual PSU, RAID disks, UPS)

Current grid service nodes

t2ce02

t2ce03 - retired

t2ce04

t2bdii01 - retired(ish)

t2bdii02

t2mon02

t2torque

t2manage

t2wlcgnagios

- The key advantages have been:
 - Use of better hardware. If we'd used real machines for all this they'd have had to be old ones, or have cost a lot.
 - New machines with no cost.
 - New machines with no planning.
 - Archiving of old systems.
 - (Some) redundancy with manual migration of VMs.
- Loading
 - We're using two hosts, both 8core 16Gb systems, both are well under used with load averages of < 1

VMware server – pros and cons

- Our original choice of VMware server was mainly motivated by:
 - Pointy-clicky GUI goodness
 - It runs on SL4 (obvious potential impact to its stability though)
 - It doesn't cost anything.
- Downsides:
 - It's missing some feature found in higher-end (pay-for) VMware products, and also some Free alternatives.
 - Slightly complicates kernel updates as modules need to be rebuilt.
- Alternatives
 - Xen (yuck), KVM (not ready at the time)

Conclusions

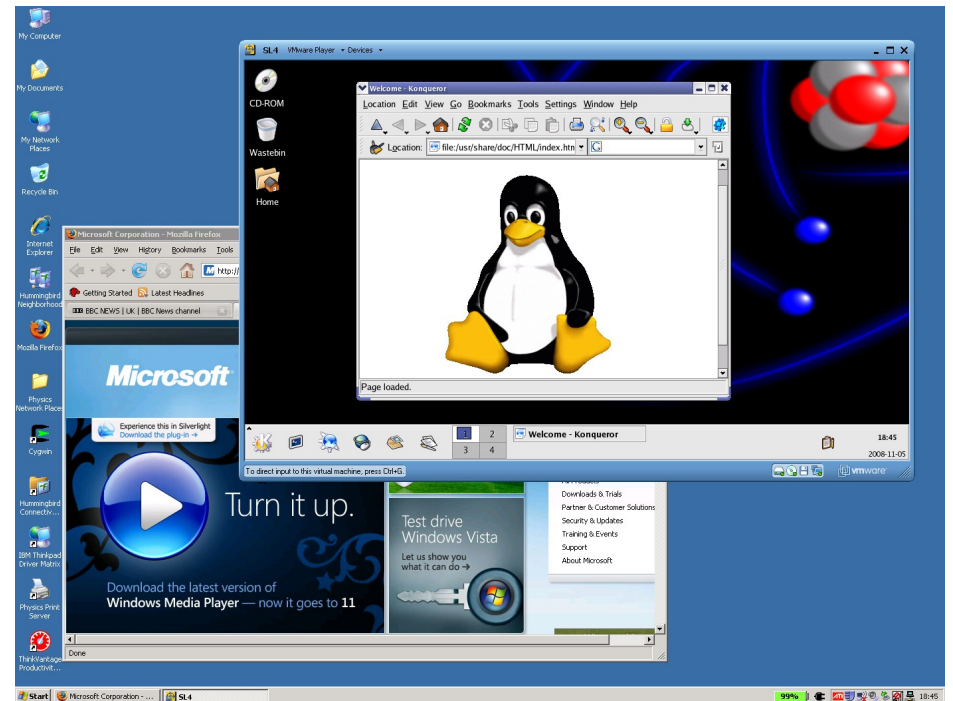
- Virtualisation of service nodes definitely a good thing.
- VMware server a good way to get started.
- Newly available alternatives may be better now, or in the future. Notably:
 - VMware ESX
 - KVM/Libvirt/OVirt

VMware player on Laptops

- Why?
 - Historically we'd used WinXP / SL4 dual boot systems to mostly good effect.
 - But increasingly newer hardware wasn't supported by SL4

VMware player on Laptops

- What's Player?
 - Free of charge download
 - Very limited interface; can only 'play' existing VMs, can't create, can't reconfigure. But that's an interface thing, not a security thing.
- Nifty bits:
 - Associates itself with .vmx files so they can be run directly. We add ours to the Start Menu.
 - NATted networking means that the VM network works transparently as the host moves between wired and wireless.
 - The VM can be suspended when it's not in use; the whole thing can just be used like an application.



VMware player on Laptops

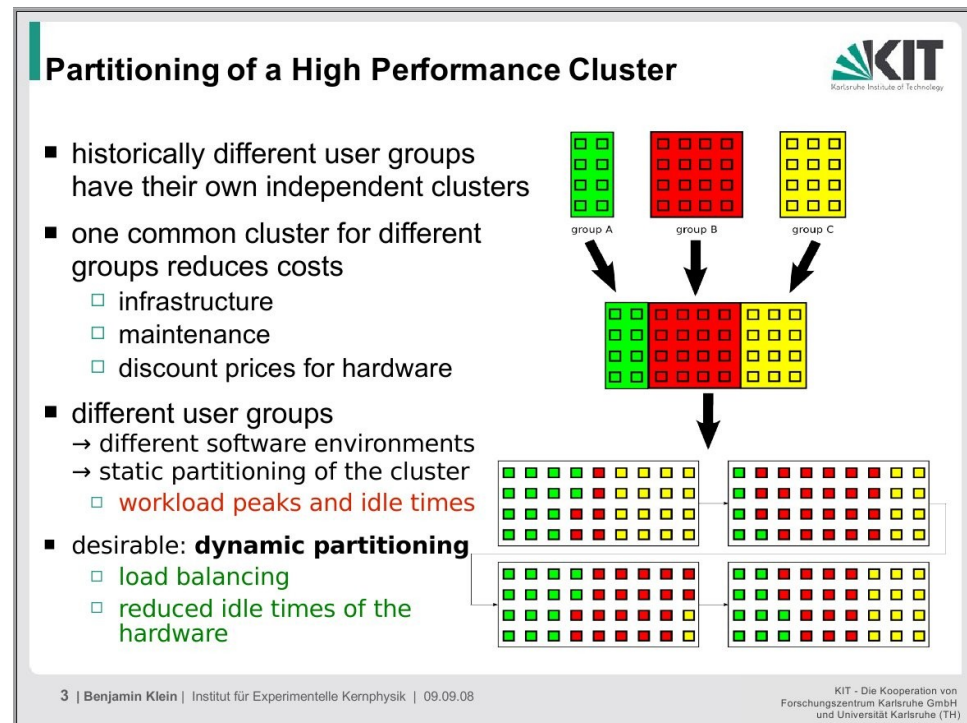
- Nifty bits:
 - The host PM still works, with full hibernate and suspend support, including with the VM running.
 - Display resizing.
 - HGFS (Host-Guest FS) – allows the guest to see the host's filesystem.
- Catches:
 - Player only supports single CPU, 32bit systems.
 - We've had some reports of slipping time sync.
 - HGFS is only one way; we could make the guest visible from the host over the 'network', but we haven't.

Conclusions

- Virtual Linux on laptops useful in the right context(s)
 - Need for old distribution
 - RAL experience with PointSec.
- Arguably complicates things. Arguably doesn't.
- VMware Player a very neat solution. Possible alternatives:
 - VirtualBox

Virtual Worker nodes

- We're not doing this yet, but some people are.
 - http://gks08.fzk.de/Talks/2008_09_09_Tuesday/Virtualization_workshop_batch_system_virtualization.pdf
- Reasons why we'd want to:
 - We've got two clusters, the local one and the grid one.
 - With 8core WNs there's a lot of scope for jobs to interfere with each other.
 - To run on clusters with the 'wrong' OS.



Virtual Worker nodes - approaches

- There are two basic approaches:
 - One batch system, one queue. Each job that's submitted is actually a wrapper that starts a VM, runs the job in it, and shuts it down.
 - Allows full fairsharing etc. between users
 - Requires a single username space
 - Multiple virtual clusters. Have several schedulers (e.g. Our local and grid clusters) and start different VM WNs to take jobs from each. Still requires some sort of metascheduler.
 - Has better isolation
 - Harder to schedule VMs since the metascheduler doesn't get information about the individual jobs.

Virtual WNs – possible technology

- VMware. Server or ESX.
 - Would likely require one of the pay-for versions
- Xen
 - Yuck
 - On the way out of Fedora/RH/SL
 - Actually rather capable, now.
- KVM
 - Clearly the RH/SL future
 - (Mostly) good interface via RH libvirt
- RH ET's OVirt
 - KVM/Libvirt based tool for running VMs over multiple generic hosts.

Conclusions

- Virtual WNs have some significant potential upsides:
 - Good job isolation
 - Improved use of hardware
 - Ability to offer different environments (e.g. 32Bit vs 64bit, local vs grid)
- Has likely downsides too:
 - Some (probably trivial) overhead
 - Could get *really* complicated.