

Operational Impact of Virtualization

Alejandro Iribarren
CERN IT/FIO
HEPiX Fall 2007
7 November 2007

The Dark Side of Virtualization

- You've heard all the hype before
- Types of problems:
 - Scalability
 - Invalidated assumptions
 - New features, new problems
- Disclaimer:
 - Many use-cases, not all problems apply
 - Most problems have a solution
 - Not trying to discourage anyone

Fabric
Infrastructure
and
Operations

Many more nodes

- Standard scalability issues
 - Can we manage them all?
 - Is our automation good enough?
- Infrastructure planning
 - Power and cooling, but especially network
 - Does our planning depend on load estimates?
- Nodes need IP addresses
 - Do we have enough public IPs?
 - Can our applications use private IPs?
 - Can we solve this with IPv6?

Can't kick a node anymore!

- A node is no longer a physical machine
 - but we've always assumed it was!
- Hardware configuration
 - VMs don't have one
 - ... or at least not a complete one
- VMs don't have MAC addresses
 - ... but the networking infrastructure needs them

Can't kick a node anymore!

- Asset tracking
 - Where is my node? Do I even care?
 - No serial numbers
 - No vendor, warranty information, etc.
- Accounting information
 - How much of the machine is this VM using?
- Scheduling interventions
 - ... suddenly became way more complicated

Can't kick a node anymore!

- Tools and workflows have to change
 - How do I install this thing?
 - Where's the console?
 - Some workflow operations no longer make sense
 - Others have changed
- Sysadmin/operator procedures
 - Node is down, can't find power button
 - What do you mean “move this node”?

Fabric
Infrastructure
and
Operations

New features, new problems

- New configurations
 - New host OS, who supports it?
 - Debugging is more difficult
 - Special “virtualized” configurations
 - New monitoring requirements
- VM Migration
 - Do we have to change our network infrastructure?
 - We need shared storage
 - Do we have the necessary infrastructure?

Fabric
Infrastructure
and
Operations

New features, new problems

- Dynamic provisioning
 - Careful: big paradigm shift!
 - Do we want to execute anyone's VMs?
 - Or do we just provide several choices?
 - What about the security risks? Who do we trust?
(No, VMs are not more secure. Really.)
 - We need new tools and interfaces
 - Who creates and supports them?
 - How do you allocate these machines?

New features, new problems

- Image management
 - Who creates, updates, certifies and supports them?
 - What about security updates?
 - Are we just moving the certification problem elsewhere?
 - Where are images stored?
 - How are they transmitted?
 - What does that do to our infrastructure?

Summary

- Virtualization *does* have its benefits
- ... but it also has problems.

Stop, plan carefully and think it through.

Fabric
Infrastructure
and
Operations

Summary

- Virtualization *does* have its benefits
- ... but it also has problems.

Stop, plan carefully and think it through.

When all you have is a hammer....
... think about getting a screw driver.