Virtualization: *Beyond the Physical Machine Limits*

Summer Student Sessions

18 August 2010

Tomas Tauber IT-DB-IMS

Virtualization

What is it?

"Machine abstraction by separating an instantiation of an operating system from physical hardware..."

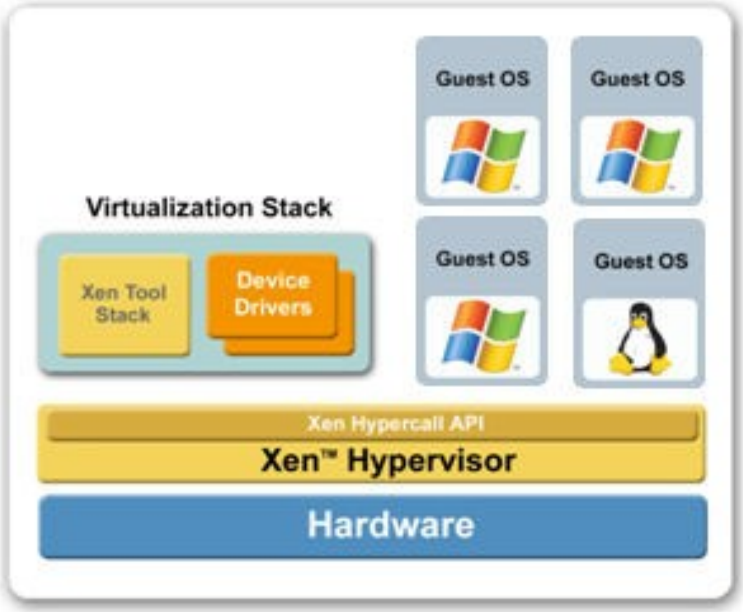

Why is it useful?

- Workload consolidation
- Application compatibility
- Higher availability

...

My tasks

- Automation
  - Integration of OracleVM (Xen based) into the current IT-DB infrastructure
  - Virtual machine management tools: physical host search, recovery, creation…
- Optimization
  - Memory overcommitment – ballooning
  - Stress tests

Motivation

"Memory is increasingly becoming a bottleneck in virtualized system" *

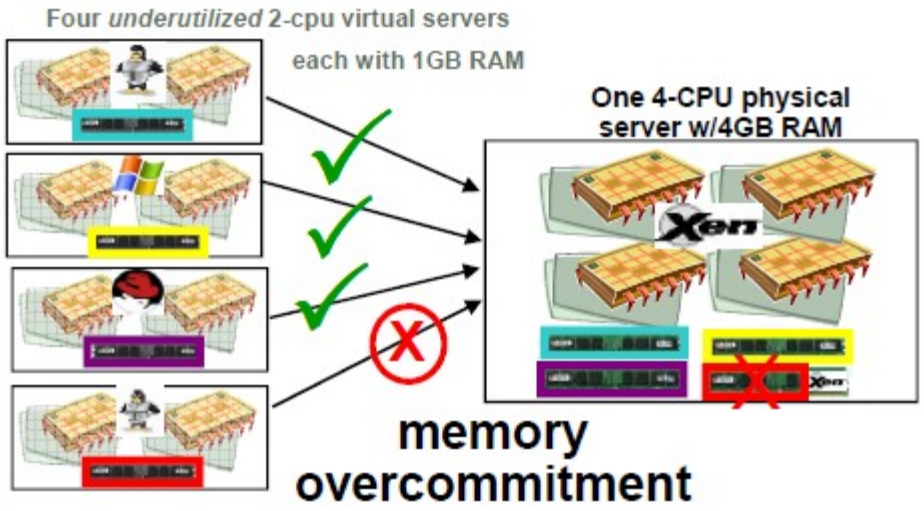

* Following slides are borrowed from *"Update on Transcendent Memory on Xen"* (Xen Summit 2010) by Dan Magenheimer (Oracle Corp.)

Physical Memory Management

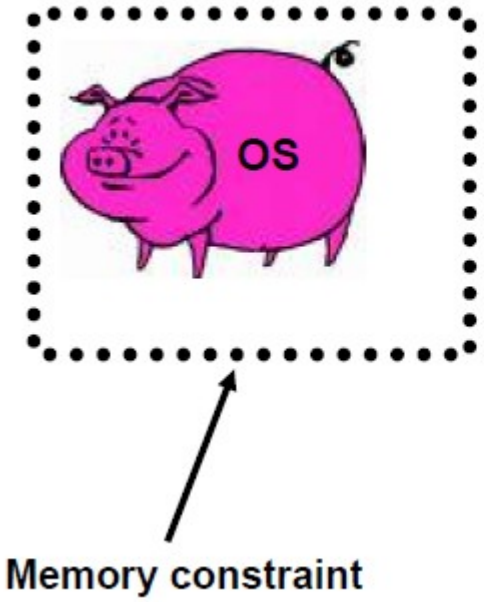

- Operating systems are memory hogs!

Physical Memory Management

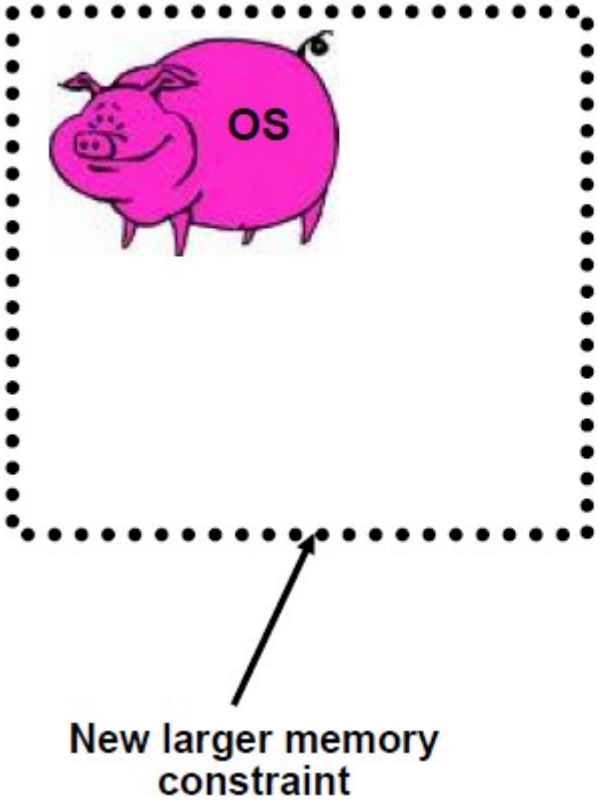

- Operating systems are memory hogs!

If you give an operating system more memory…..

Physical Memory Management

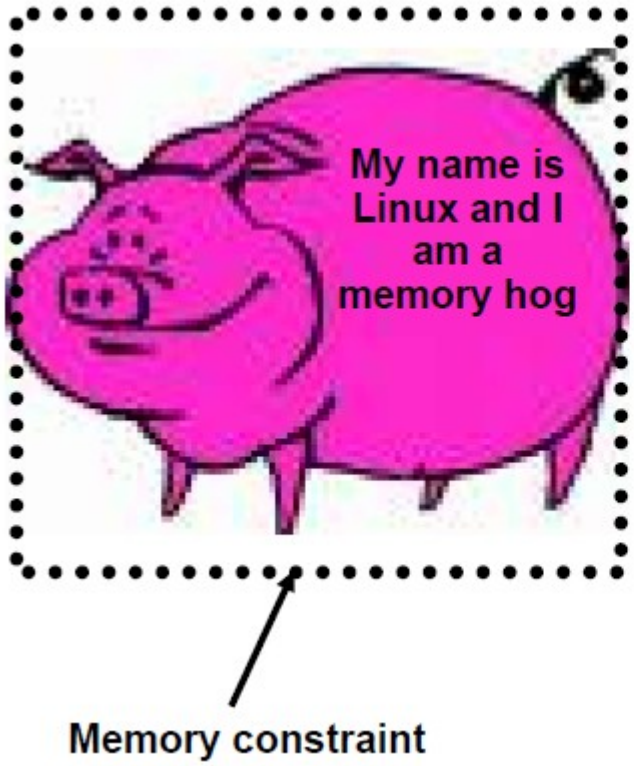

- Operating systems are memory hogs!

If you give an OS more memory

…it uses up any memory you give it!

Physical Memory Management

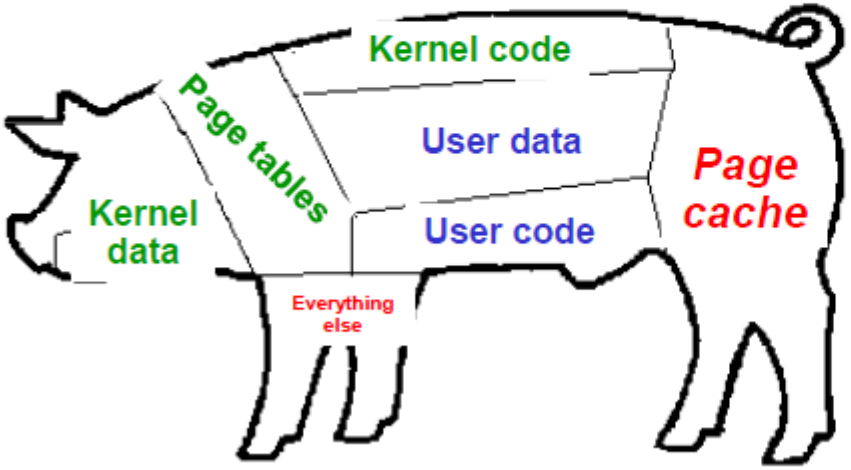

- What does an OS do with all that memory?
  - Kernel code and data
  - User code and data
  - *Page cache!*

Physical Memory Management

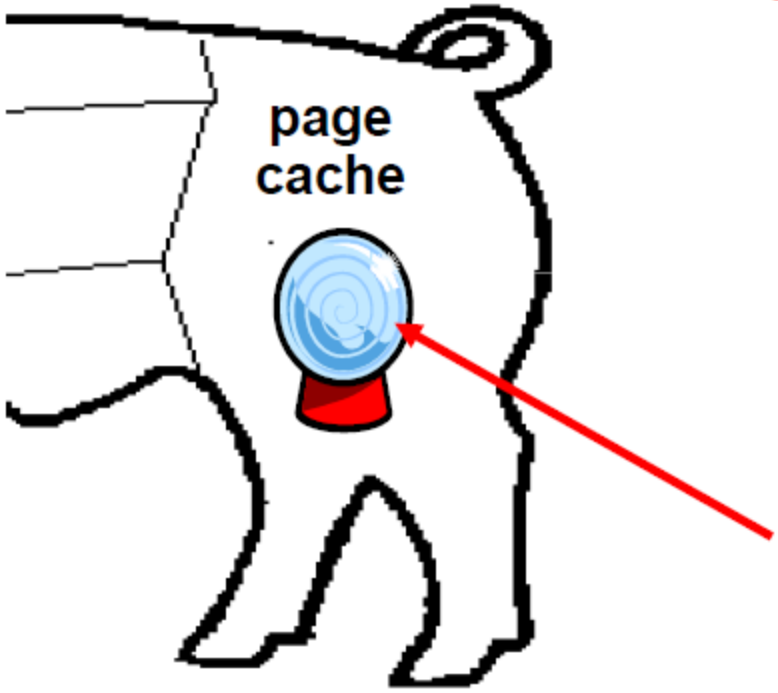

- What does an OS do with all that memory?

Page cache attempts to predict future needs of pages from the disk…

sometimes it gets it right
→ "good" pages

Physical Memory Management

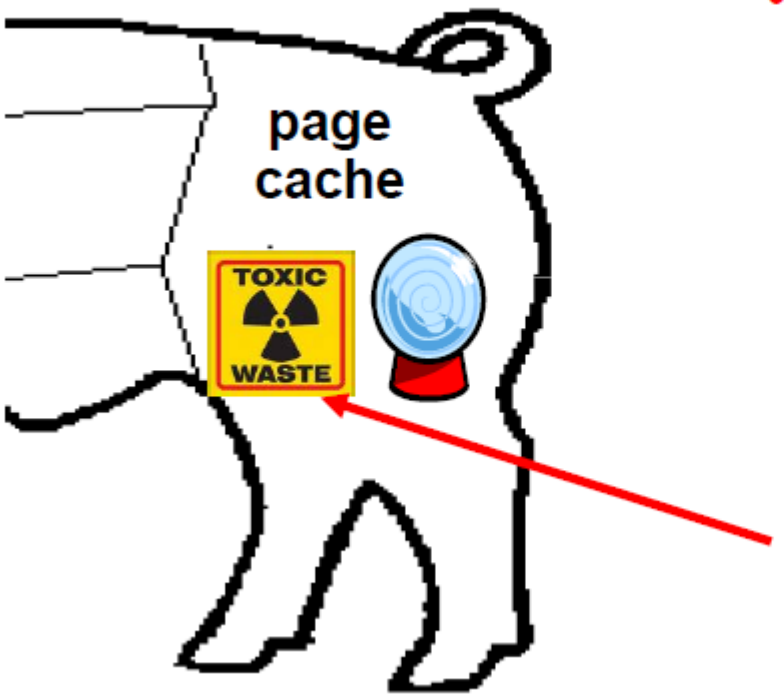

- What does an OS do with all that memory?

Page cache attempts to predict future needs of pages from the disk…

sometimes it gets it wrong
→ "wasted" pages

Memory Ballooning

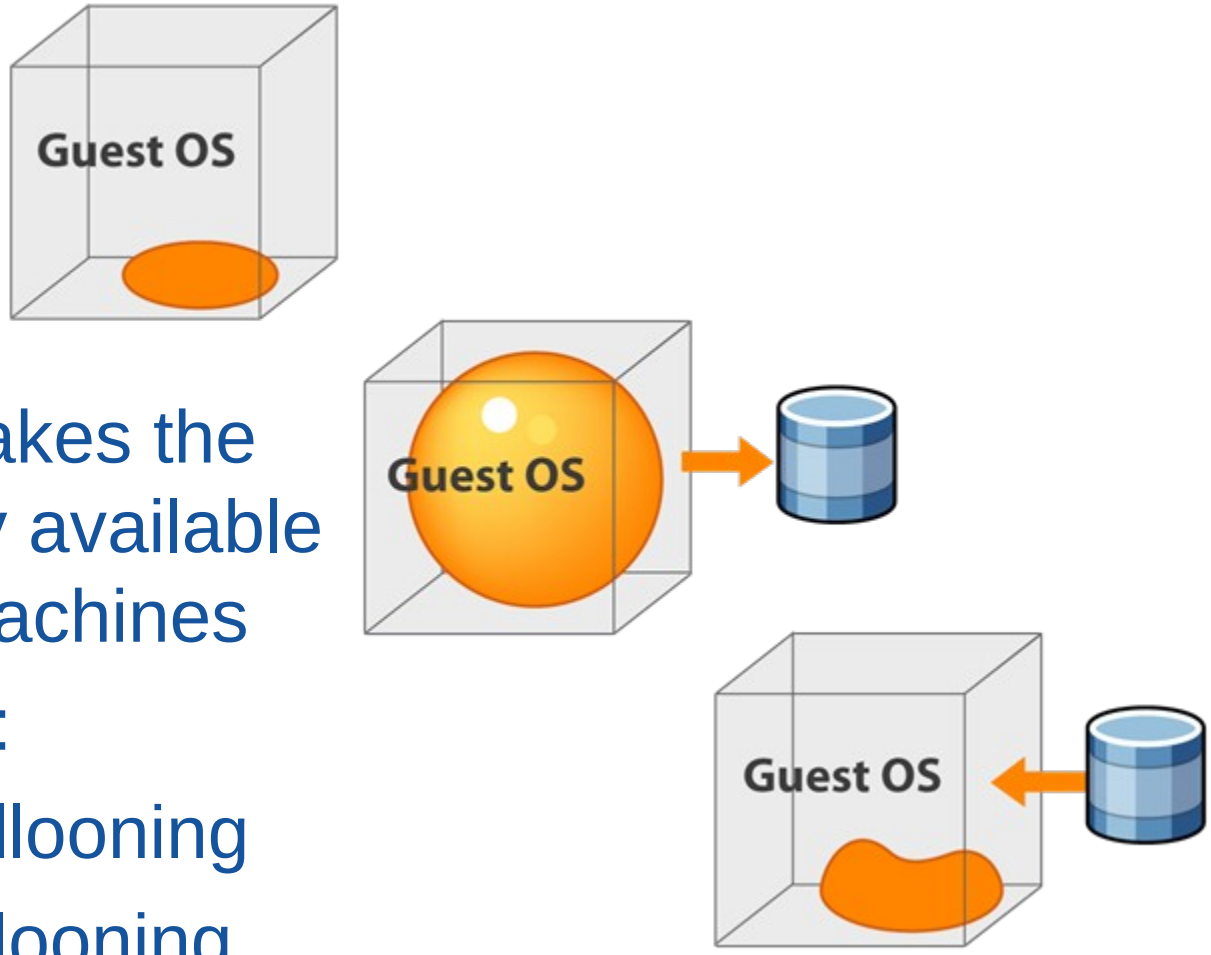

- Balloon driver makes the potential memory available to other virtual machines
- Used techniques:
  - *Dom0* auto-ballooning
  - *DomU* self-ballooning

Memory stress tests

- Physical machine: Intel Xeon CPU L5520 @ 2.27GHz (1/8), **48 GB RAM**
- Virtual machines: 2 VCPU, **4 GB RAM**; *"clean"* and *"stress"* images
- Free physical memory measured

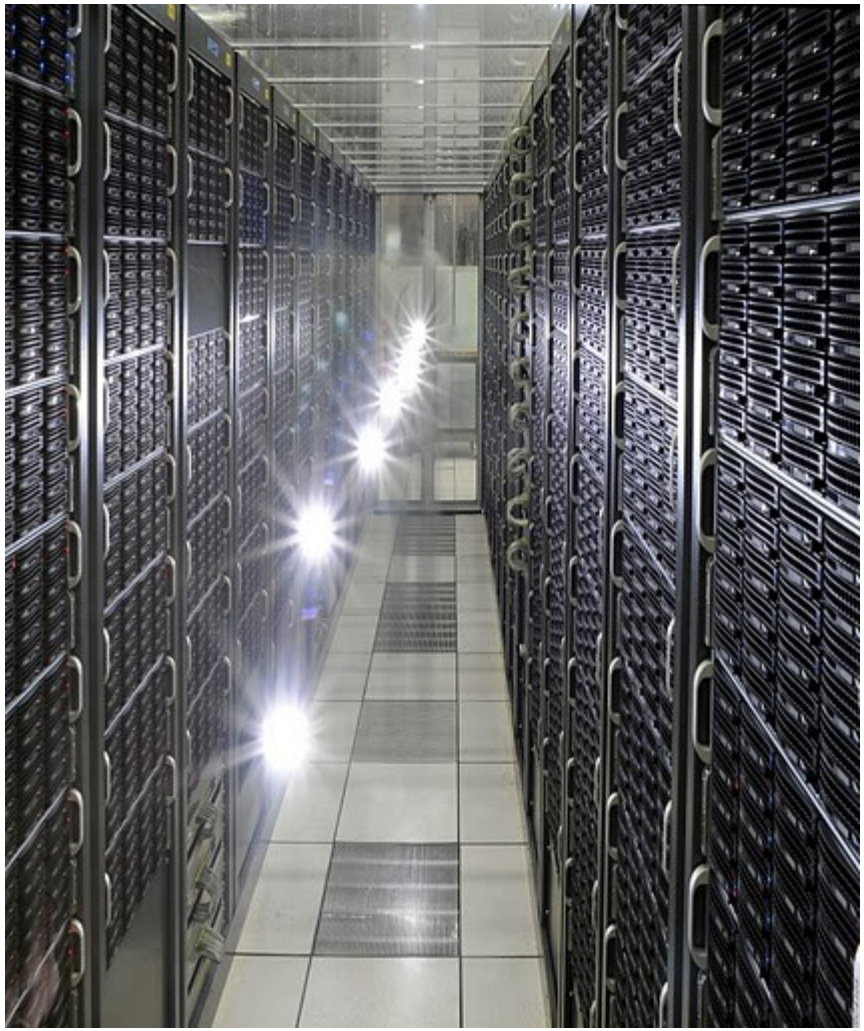

Results: rough behaviour

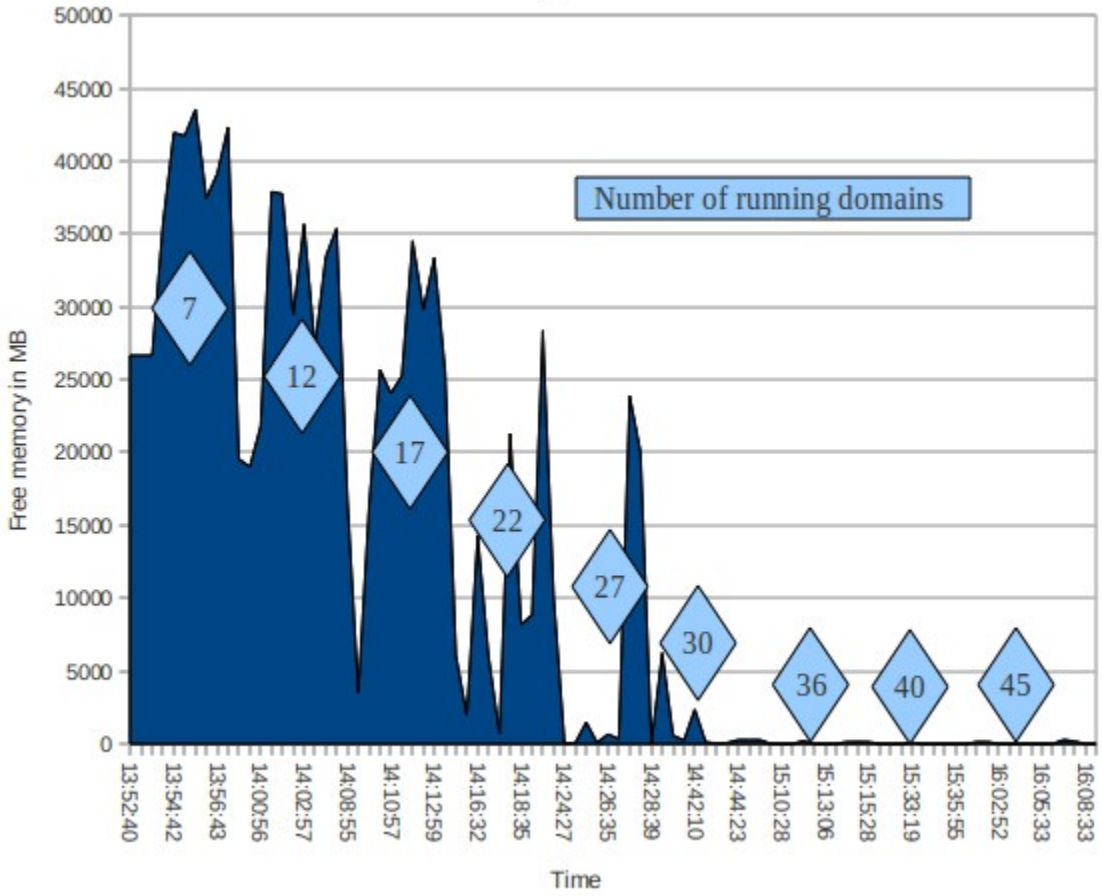

Results: memory comparison

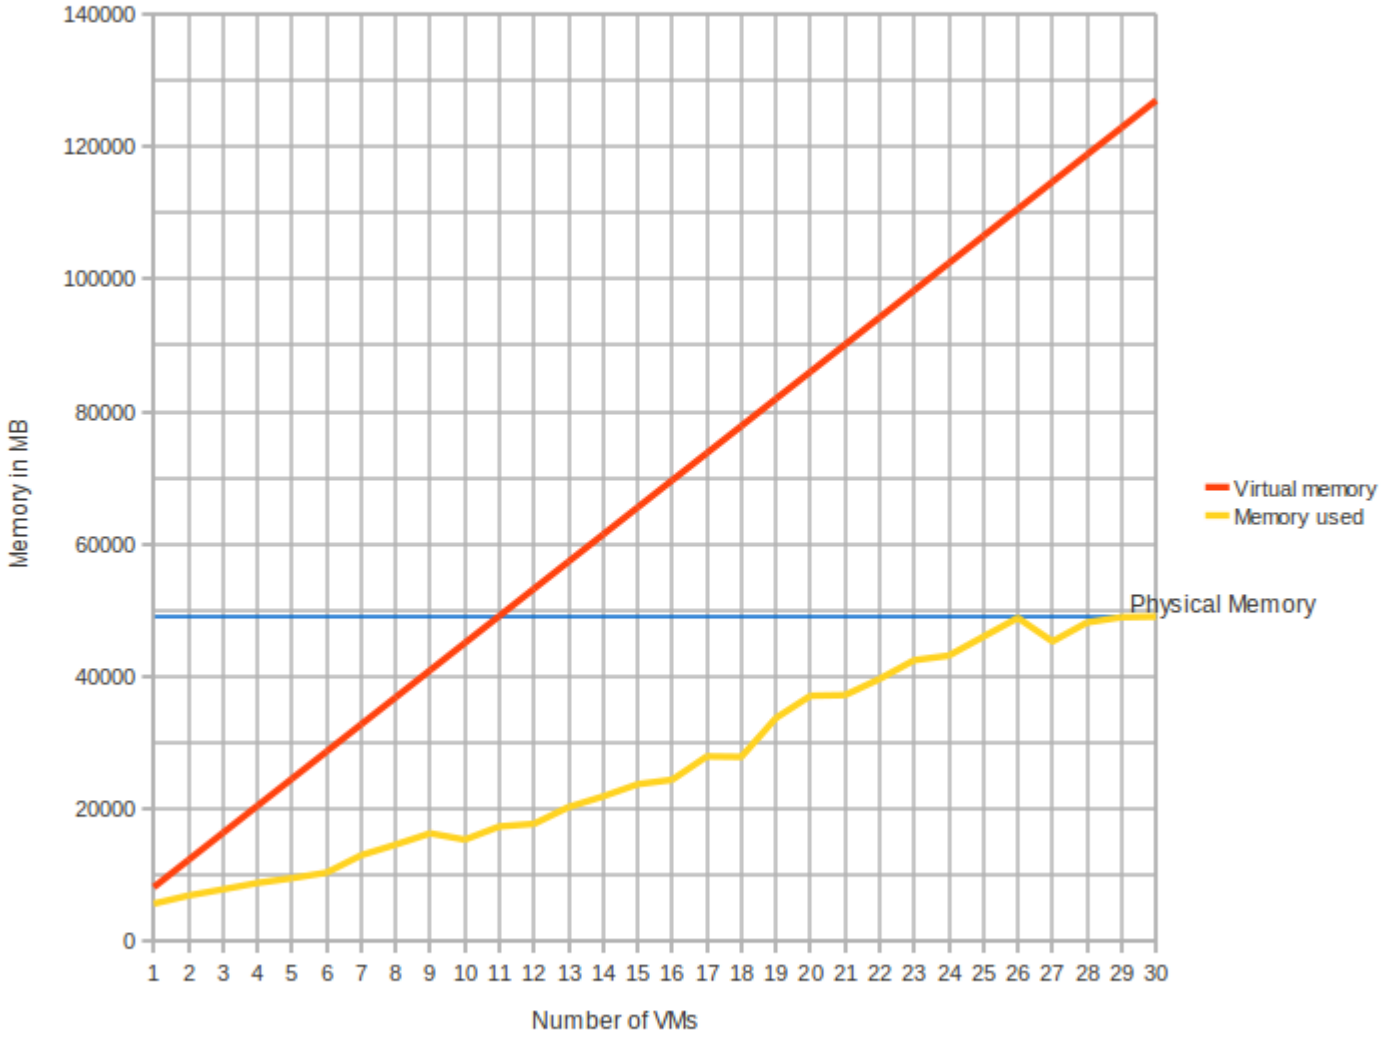

Conclusion

- In our tests, we can run:
  - ~30 “stress” virtual machines → ~2.7 more
  - ~120 “clean” virtual machines (the minimal used memory ~400 MB) → ~10.9 more
- With more compact configuration (< 400 MB needed), 200+ VMs per one physical host
- CPU was underutilized during the memory stress tests (< 5 %), only when hot-swapping, its usage jumped up to ~40 %

Questions

- More information:
  - www.cern.ch/openlab
  - “Memory Overcommit… without the commitment” by Dan Magenheimer (Oracle Corp.)
  - wiki.xensource.com
- My e-mail: tomas.tauber@gmail.com

ANY QUESTIONS?