A comparison between xen and kvm

Andrea Chierici
Riccardo Veraldi
INFN-CNAF

Outline

- Introduction
- xen
- kvm
- Test description
- Benchmarks
- Conclusions

Introducing virtualization

- Several benefits delivered
  - Reduced server number
  - Power saving
  - Maximization of hardware resources
  - Server isolation
  - Flexibility (migration, load-balancing, disaster recovery)

Introducing virtualization

- Virtualization widely adopted already in HEP community
- Xen proved to be a reliable tool
- KVM is an emerging technology worth to be investigated
  - We will show benchmark results on machines running O.S. compatible with EGEE grid middleware
    - This causes some limitations but gives the idea of what can be **today** the best choice for our farm

Virtualization approaches

- Full virtualization
  - Unmodified O.S., soft migration approach, slow
- Para virtualization
  - Requires modified O.S., fast
- Hardware virtualization support (hvm)
  - Requires modern CPUs, trade-off between previous approaches
    - Intel VT and AMD-V technologies

What we have now: Xen

- Open source “industry standard” for virtualization
- Included in all recent linux distributions
- Supports para and hvm approaches
- Widely used in HEP community
- Won’t go into details during presentation

What could be the future: KVM

- Kernel-based Virtual Machine
- Open source
- included in latest linux kernels: implemented as a module
  - A user space program uses /dev/kvm interface to set up VMs (qemu-kvm)
- Supports hvm approach
- Rather new to HEP community
- Qumranet now owned by Red Hat
  - Foreseen boost in development

KVM seen by system admin

- rpms: basically the kernel module and a modified version of qemu (*qemu-kvm*)
- Network configuration to be done by hand in order to get public IP
  - Used bridge-utils and tunctl rpms to set-up tap interface
  - A modified init.d script to configure software bridge has been developed at CNAF
- No VM configuration file direct support
  - VMs are launched via a standard UNIX command + command line options

KVM seen by system admin

- No direct interface to manage VMs by command line (e.g. “xm”)
  - Libvirt support
- Virt-manager working
- Monitor interface
  - Very powerful, lots of options available
- qemu already well documented, support available on-line

KVM introduced at CNAF

- We use **quattor** to configure and install grid nodes and virtual machines
  - Profiles describe machine configuration
    - See my poster today!
  - Network boot for installation working
    - Keeps booting from network, need to restart with *"-boot c"*
  - No need for any modification, just like xen-hvm

Qualitative test

- CNAF is running LHCb tier2 site entirely on xen VMs (2 CEs, 1 SE)
- Changed one Computing Element with a KVM machine
  - Has been working for more that 3 weeks flawlessly
  - Quattor machine profile unmodified, no effort for sysadmin
- CMS secondary squid server installed on the same host with same result (2 weeks ago)
- KVM executed (and live migrated!) a win7 VM
- Hardware used: 1 node, dual E5420, 16GB ram, sata disks via Areca controller

Quantitative test: description (1)

- Need some measures to understand what is the best solution
- Tested 3 classic parameters
  - CPU → hep-spec06 (v1.1)
  - Network → iperf (v2.0.4)
  - Disk access → bonnie++ (v1.94)
- Compared Xen (para-virtualized and hvm) with KVM, using non virtualized machine as a baseline

Quantitative test: description (2)

- Hardware used: 1 blade, dual E5420, 16GB ram, 10k sas disk via LSI logic raid controller (raid0)
- Xen-para VM specs: 1 vcpu, 2 GB ram, disk on a file
- Xen-hvm VM specs: 1 vcpu, 2GB ram, disk on a file, “netfront” network driver
- KVM VM specs: 1 vcpu, 2GB ram, disk on a file, e1000 network driver emulation
- Host OS: SL 5.2 x86_64, kernel 2.6.18-92.1.22.el5
- VM OS: SLC 4.5 i386, kernel 2.6.9-67.0.15.EL.cern
- KVM version: 83
- Xen version: 3.2.1

Benchmarks: HEP-Spec06

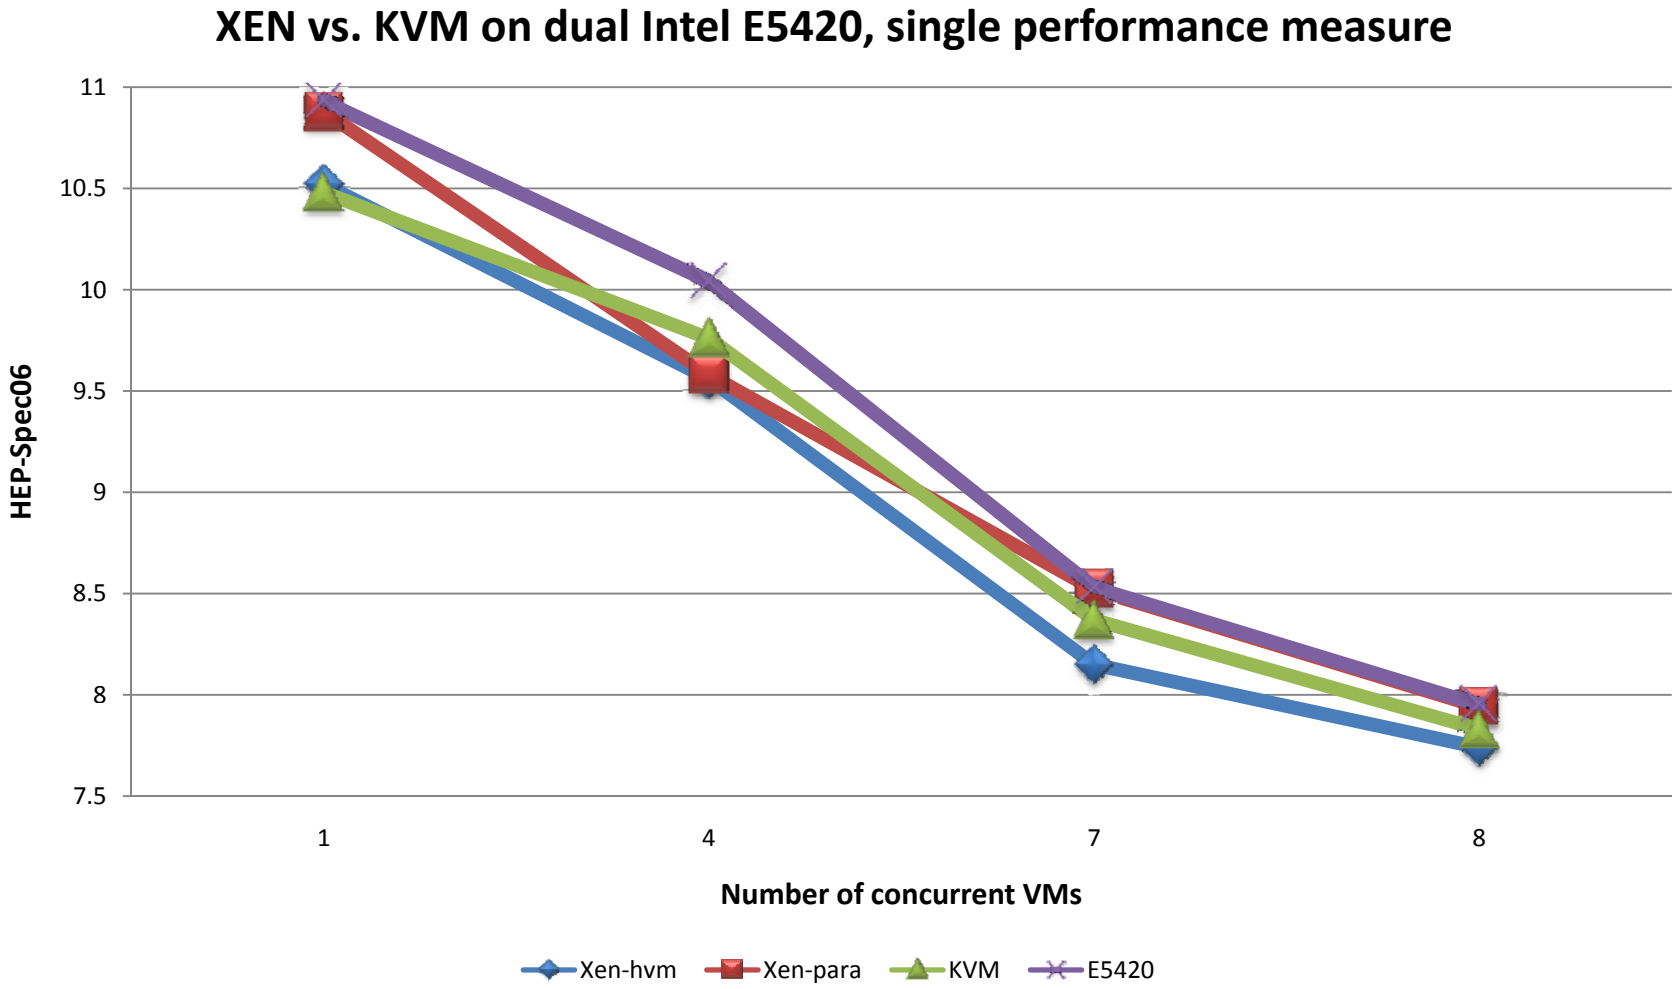

Andrea Chierici, CHEP 2009

Benchmarks: HEP-Spec06

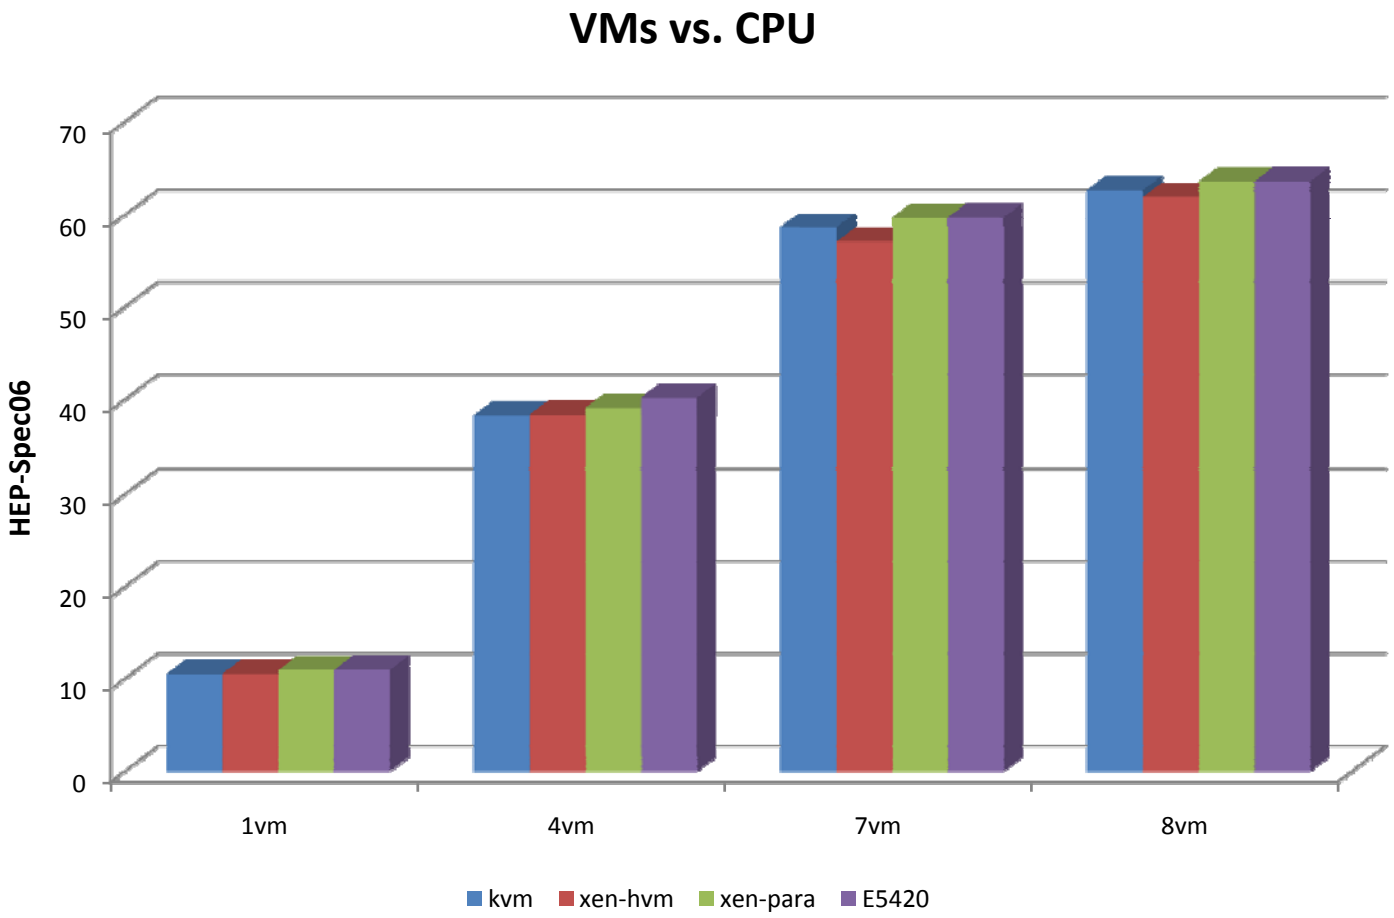

Andrea Chierici, CHEP 2009

Benchmarks: HEP-Spec06

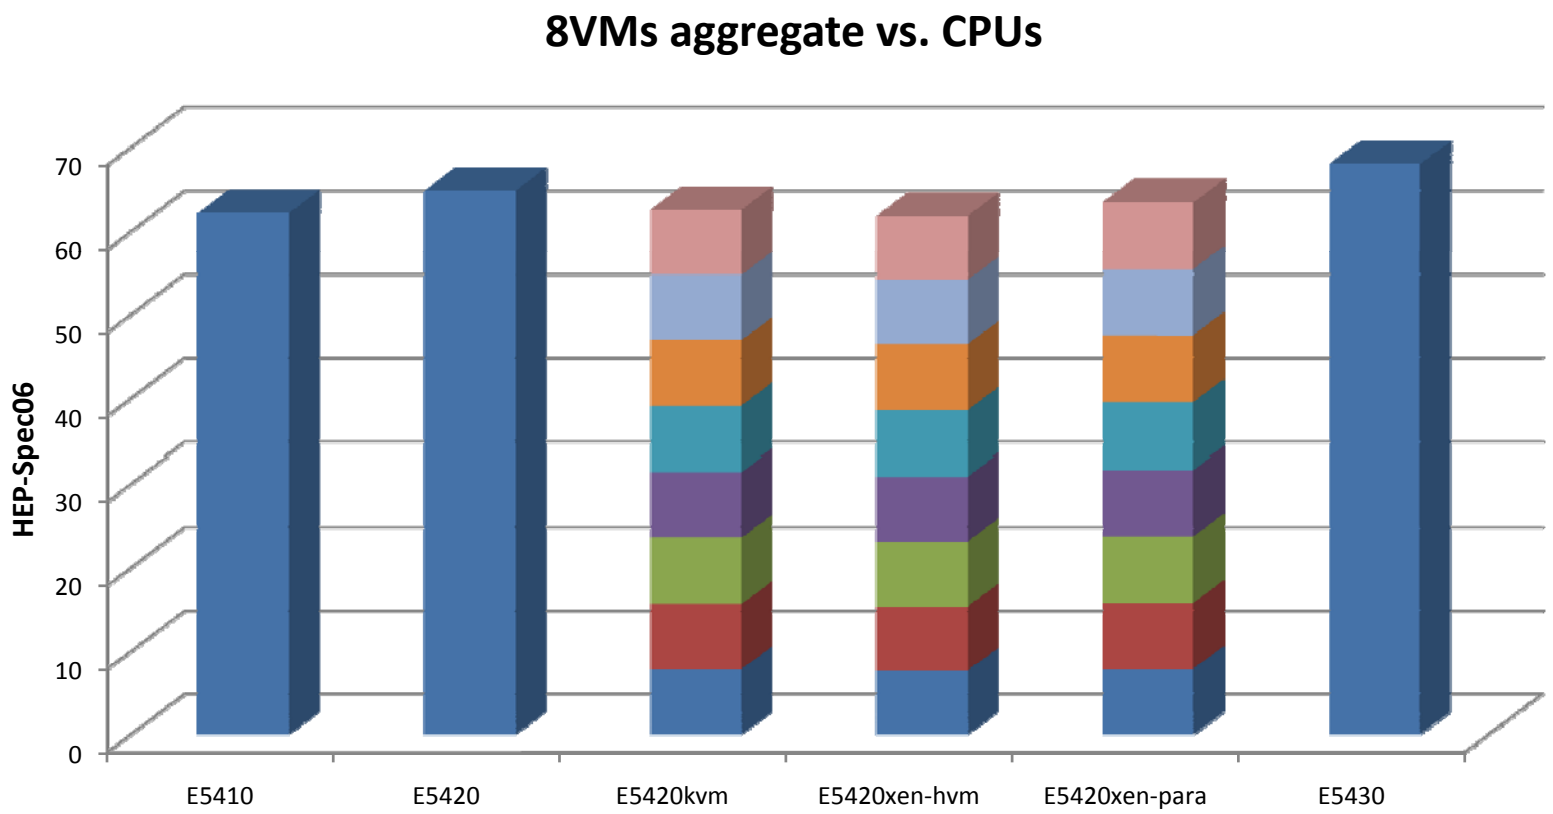

Andrea Chierici, CHEP 2009

Benchmarks: HEP-Spec06

Virtualization Technology	% loss from non emulated CPU (E5420, 8vm)
E5420kvm	3,42
E5420xen-hvm	4,55
E5420xen-para	2,02
E5410 vs. E5420	4,07

Benchmarks: Iperf

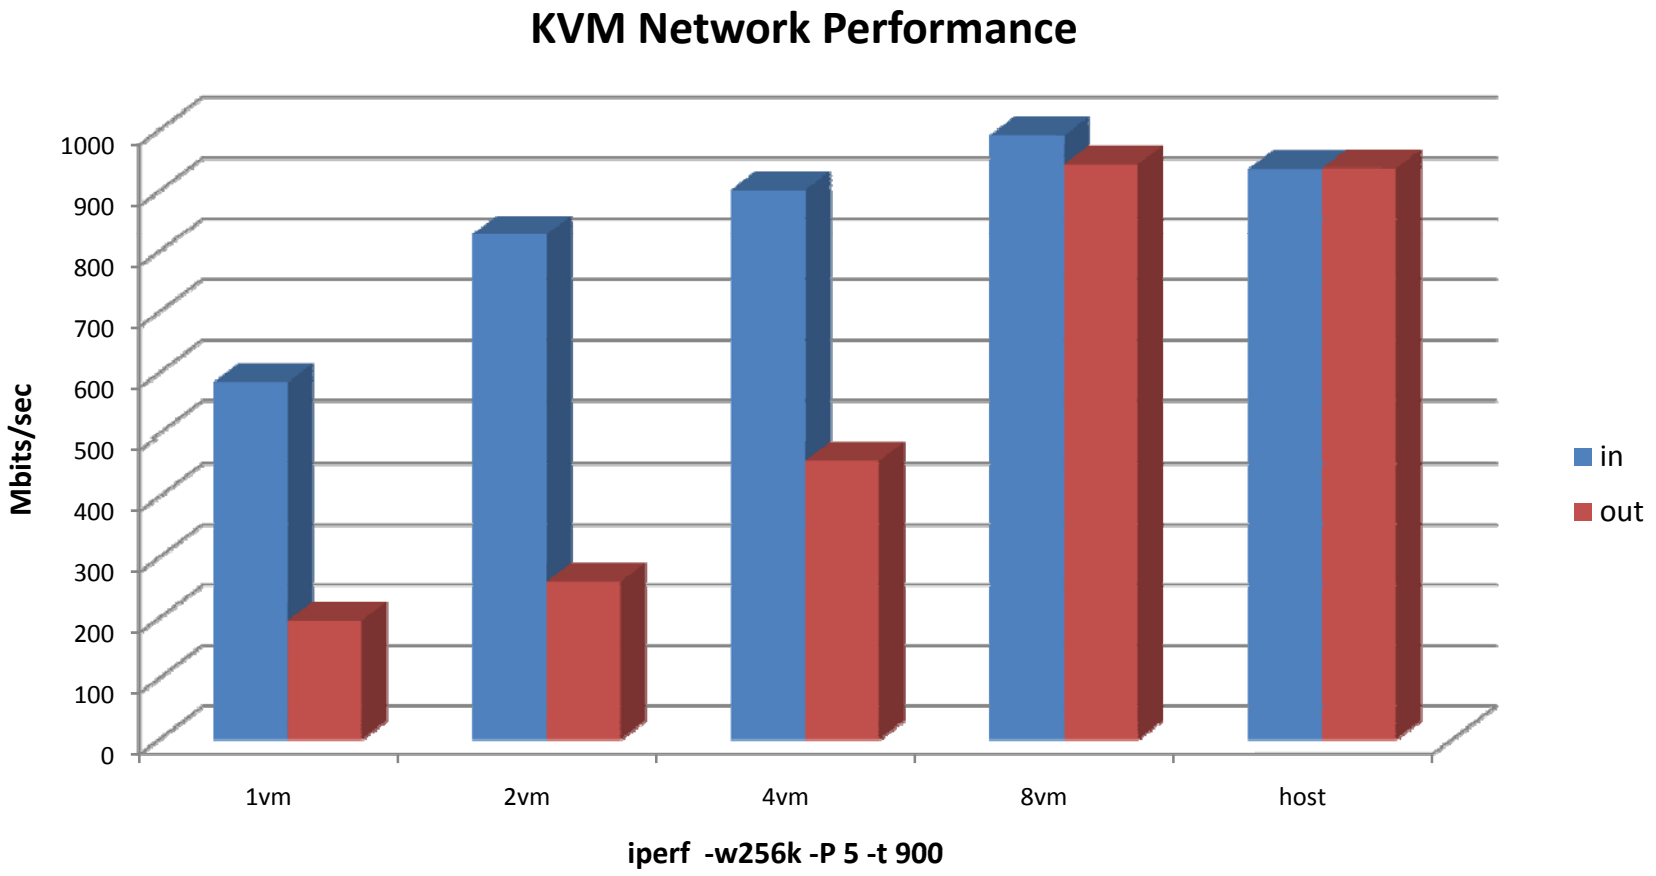

Benchmarks: Iperf

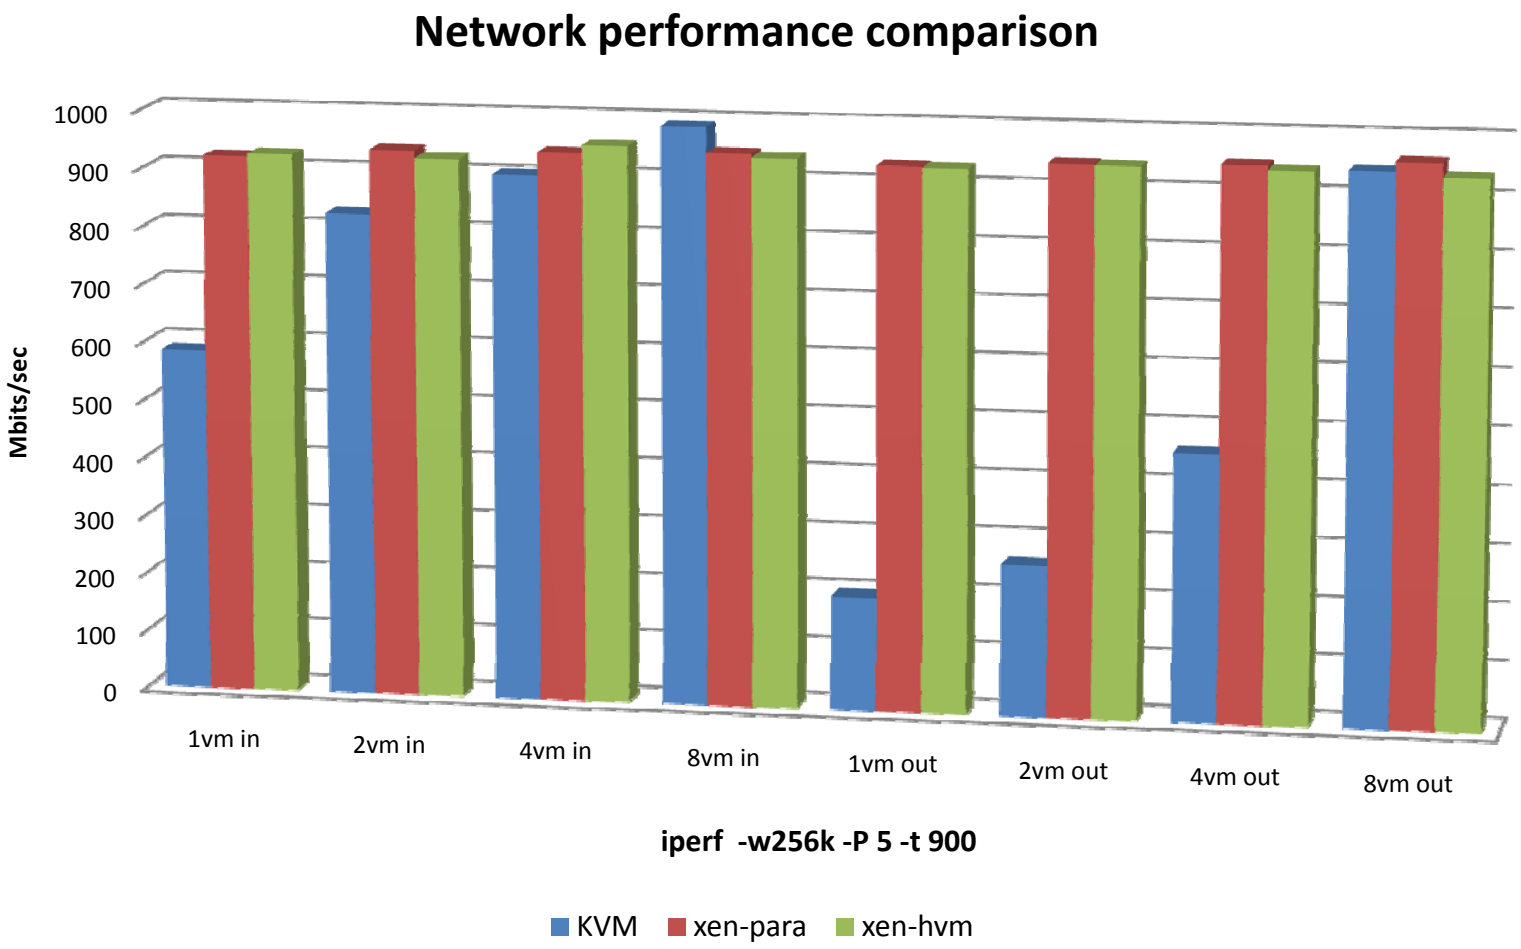

Andrea Chierici, CHEP 2009

Benchmarks: bonnie++

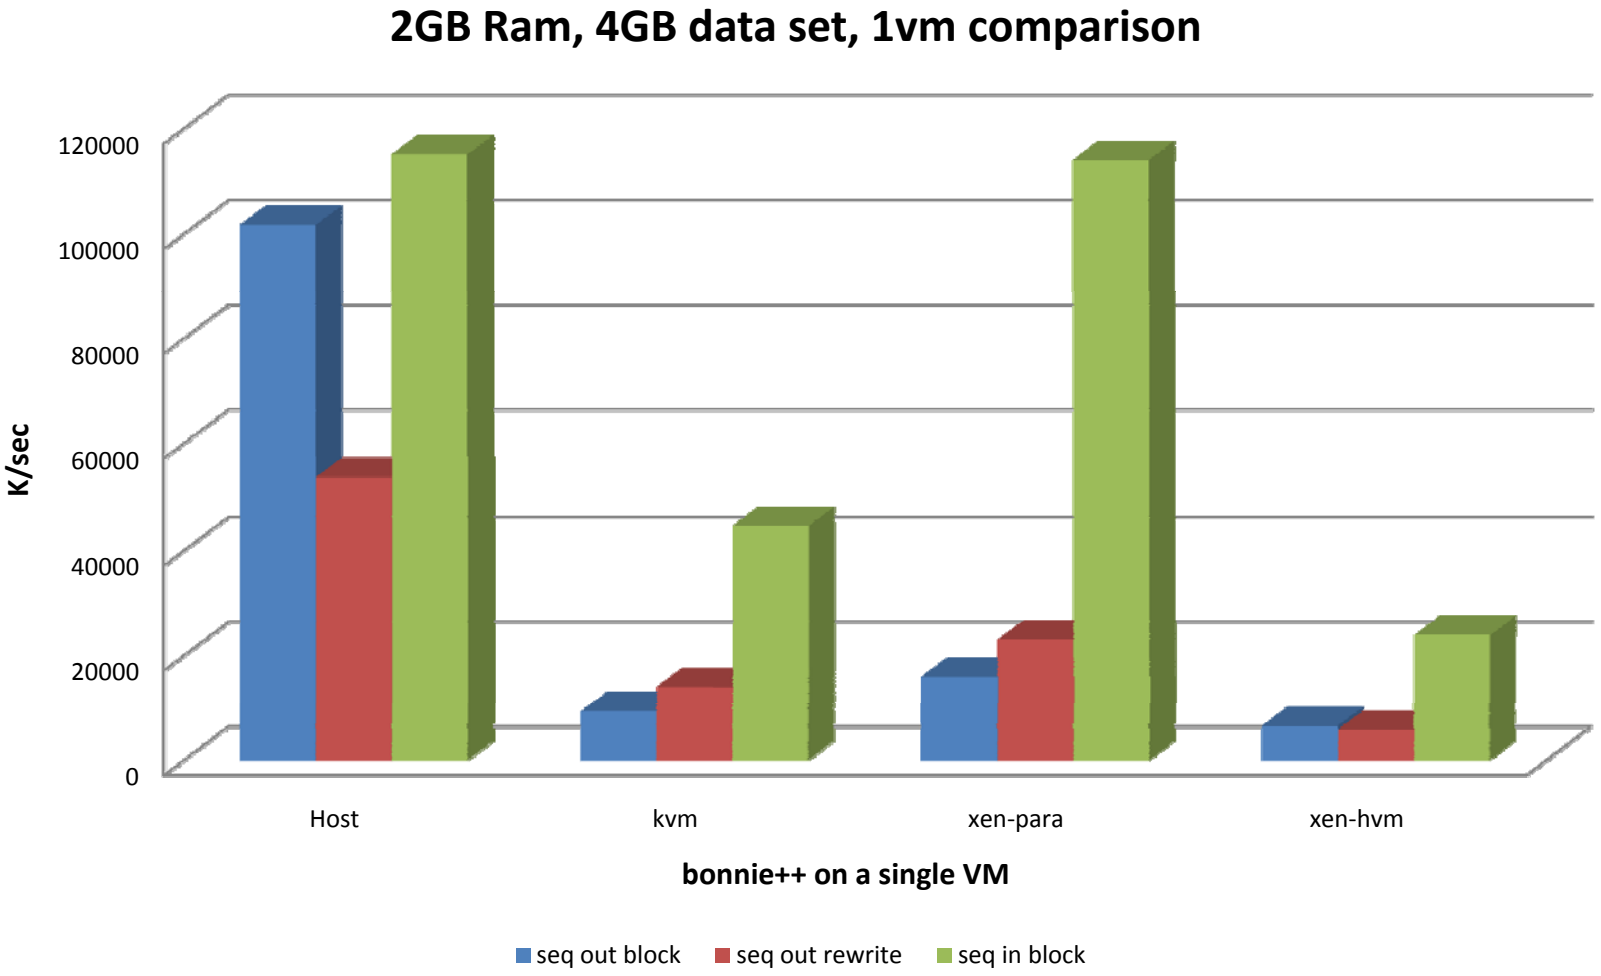

Benchmarks: bonnie++

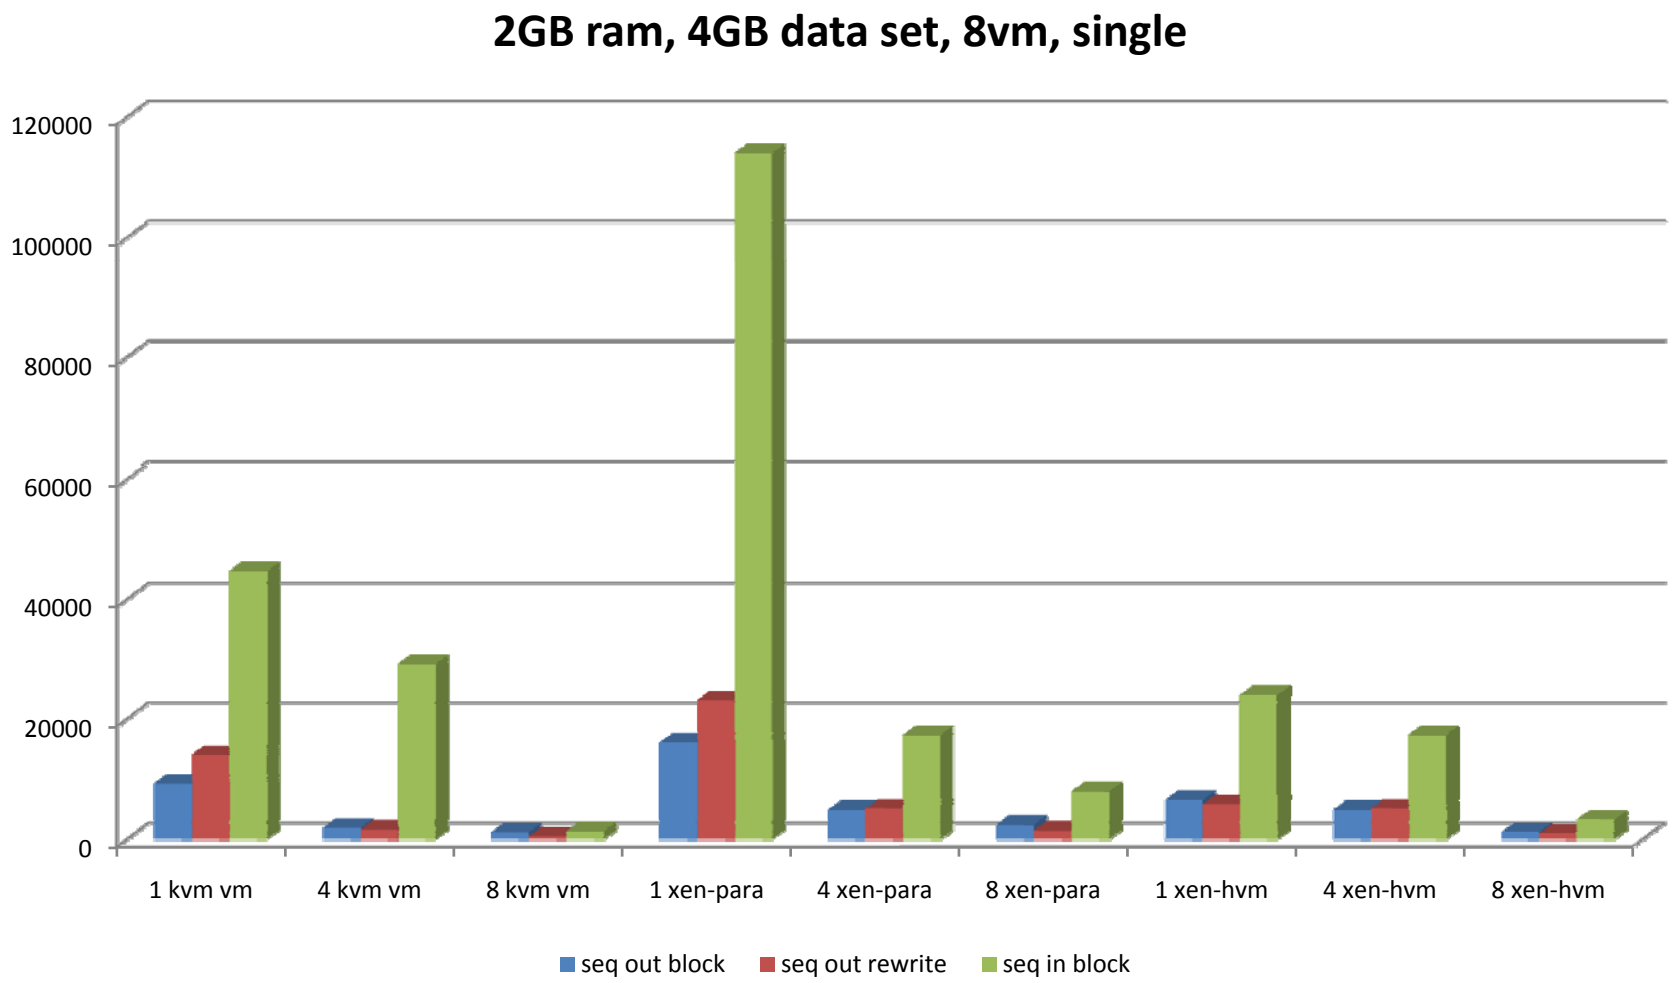

Benchmarks: bonnie++

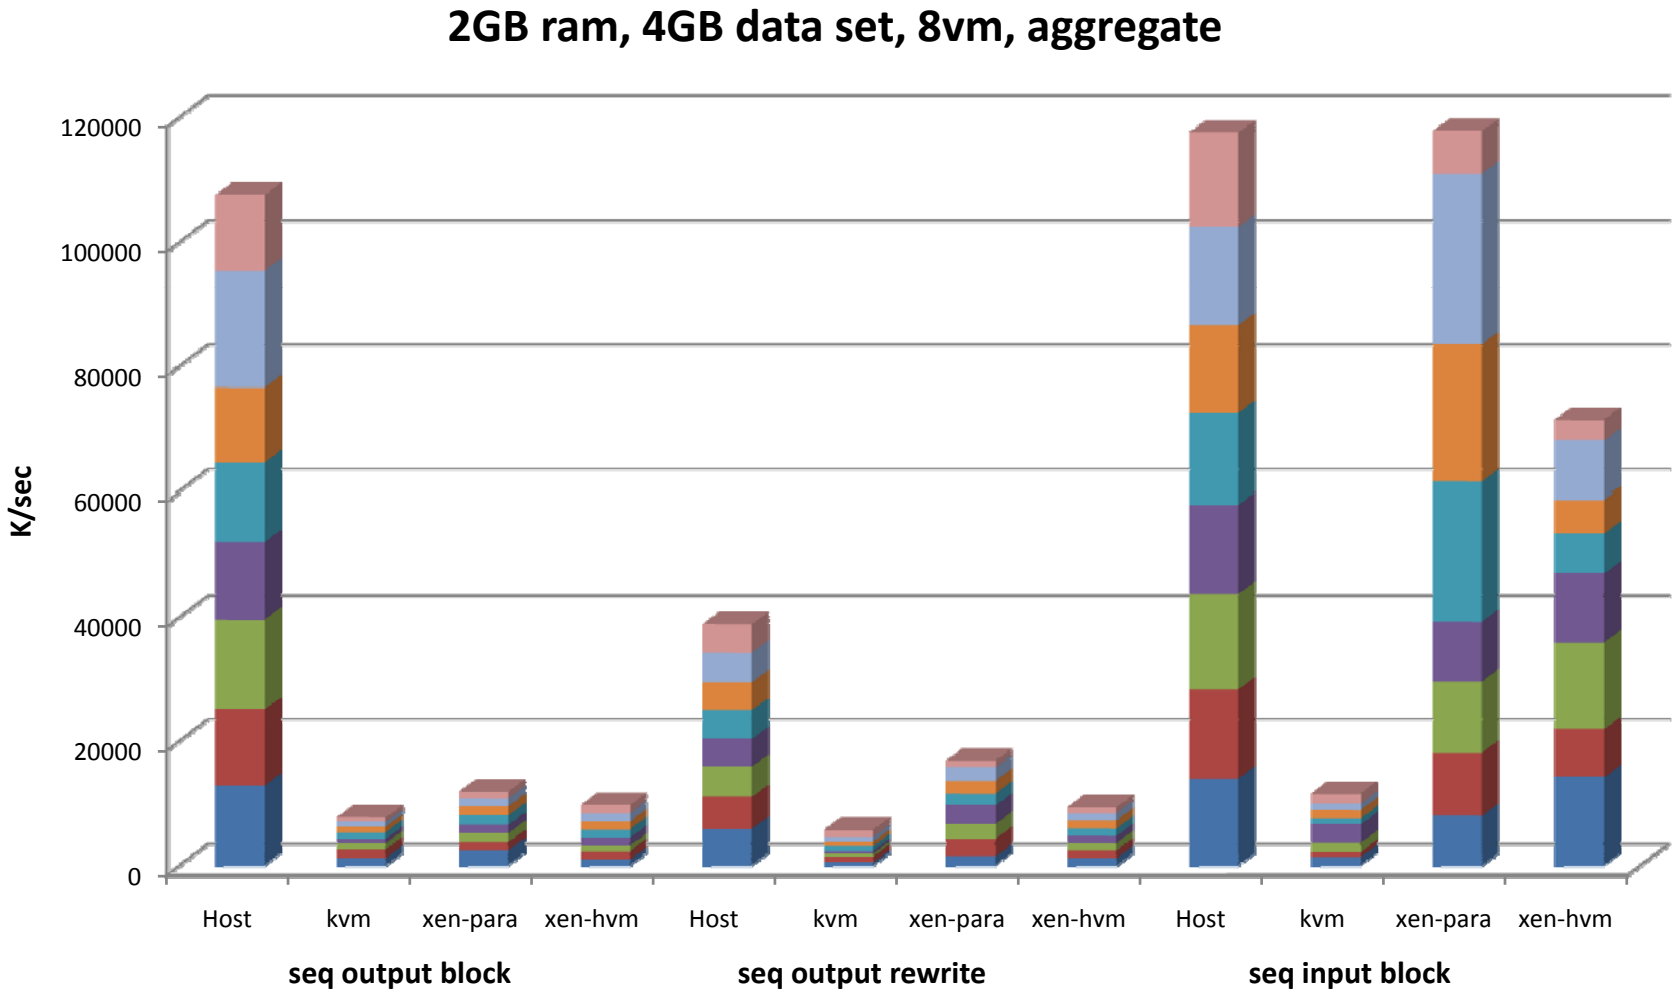

Andrea Chierici, CHEP 2009

Conclusions

- KVM proved good stability and reliability
  - No problems on running production machines for more than 3 weeks
  - CPU performances are extremely good
    - Modern CPU virtualization technologies help!
  - Network performances are fair
    - Could not test virtio drivers because sl4 kernel is not supported
  - Disk I/O seems the most problematic aspect
    - Other solutions have problems too
  - Requires small effort from sys admins
- Even if looking promising, right now xen is the most performing solution

Future work

- I/O performance with disk partition vs. file
- KVM virtio drivers (kernel v2.6.25)
  - Currently not suitable for grid nodes
    - SL5 worker nodes, back port drivers?
- qemu snapshot features
- high-level VM managers
  - Ovirt, enomalism, ganeti

Bibliography and links

- Xen
  - Xen repository
- KVM
  - Kvm repository
- Quattor
- HEP-SPEC