

HVM IO Performance

Ryan Harper – IBM
Anthony Liguori - IBM






Legal Statement

This work represents the view of the author and does not necessarily represent the view of IBM.

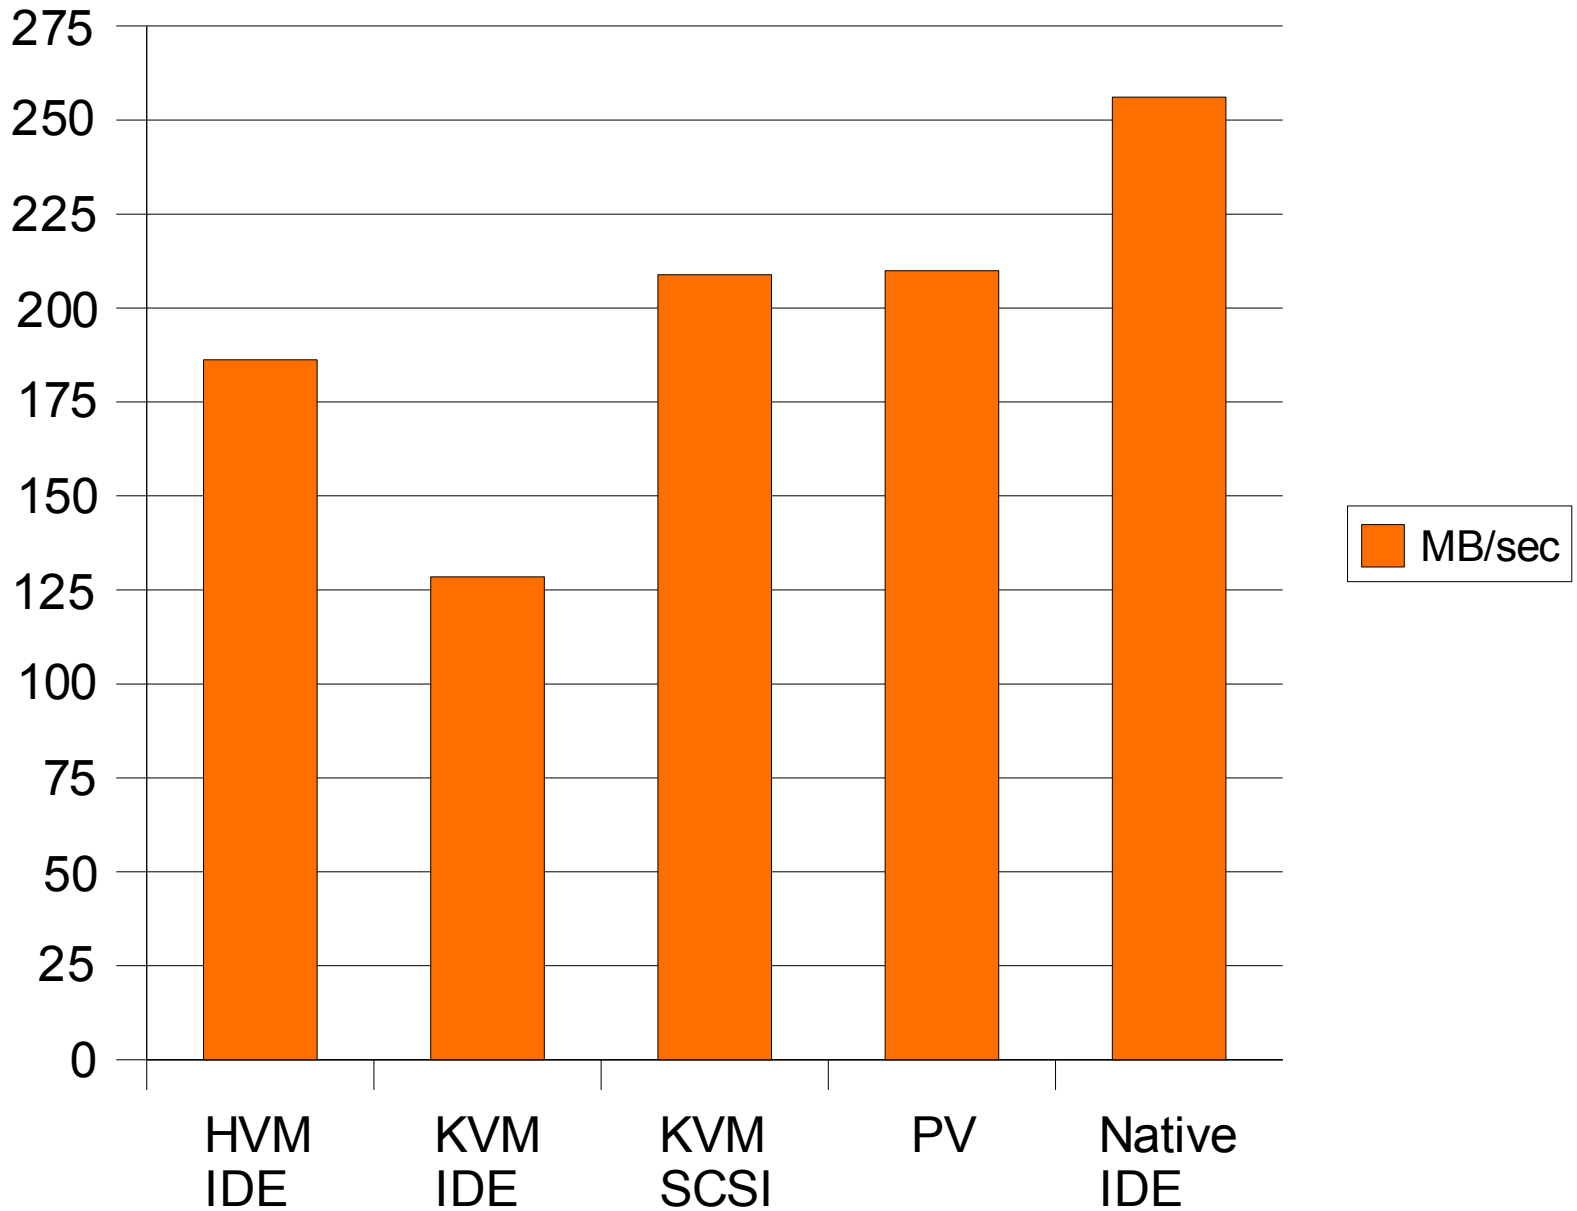
IBM, IBM (logo), e-business (logo), pSeries, e (logo) server, and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds. Other company, product, and service names may be trademarks or servicemarks of others.

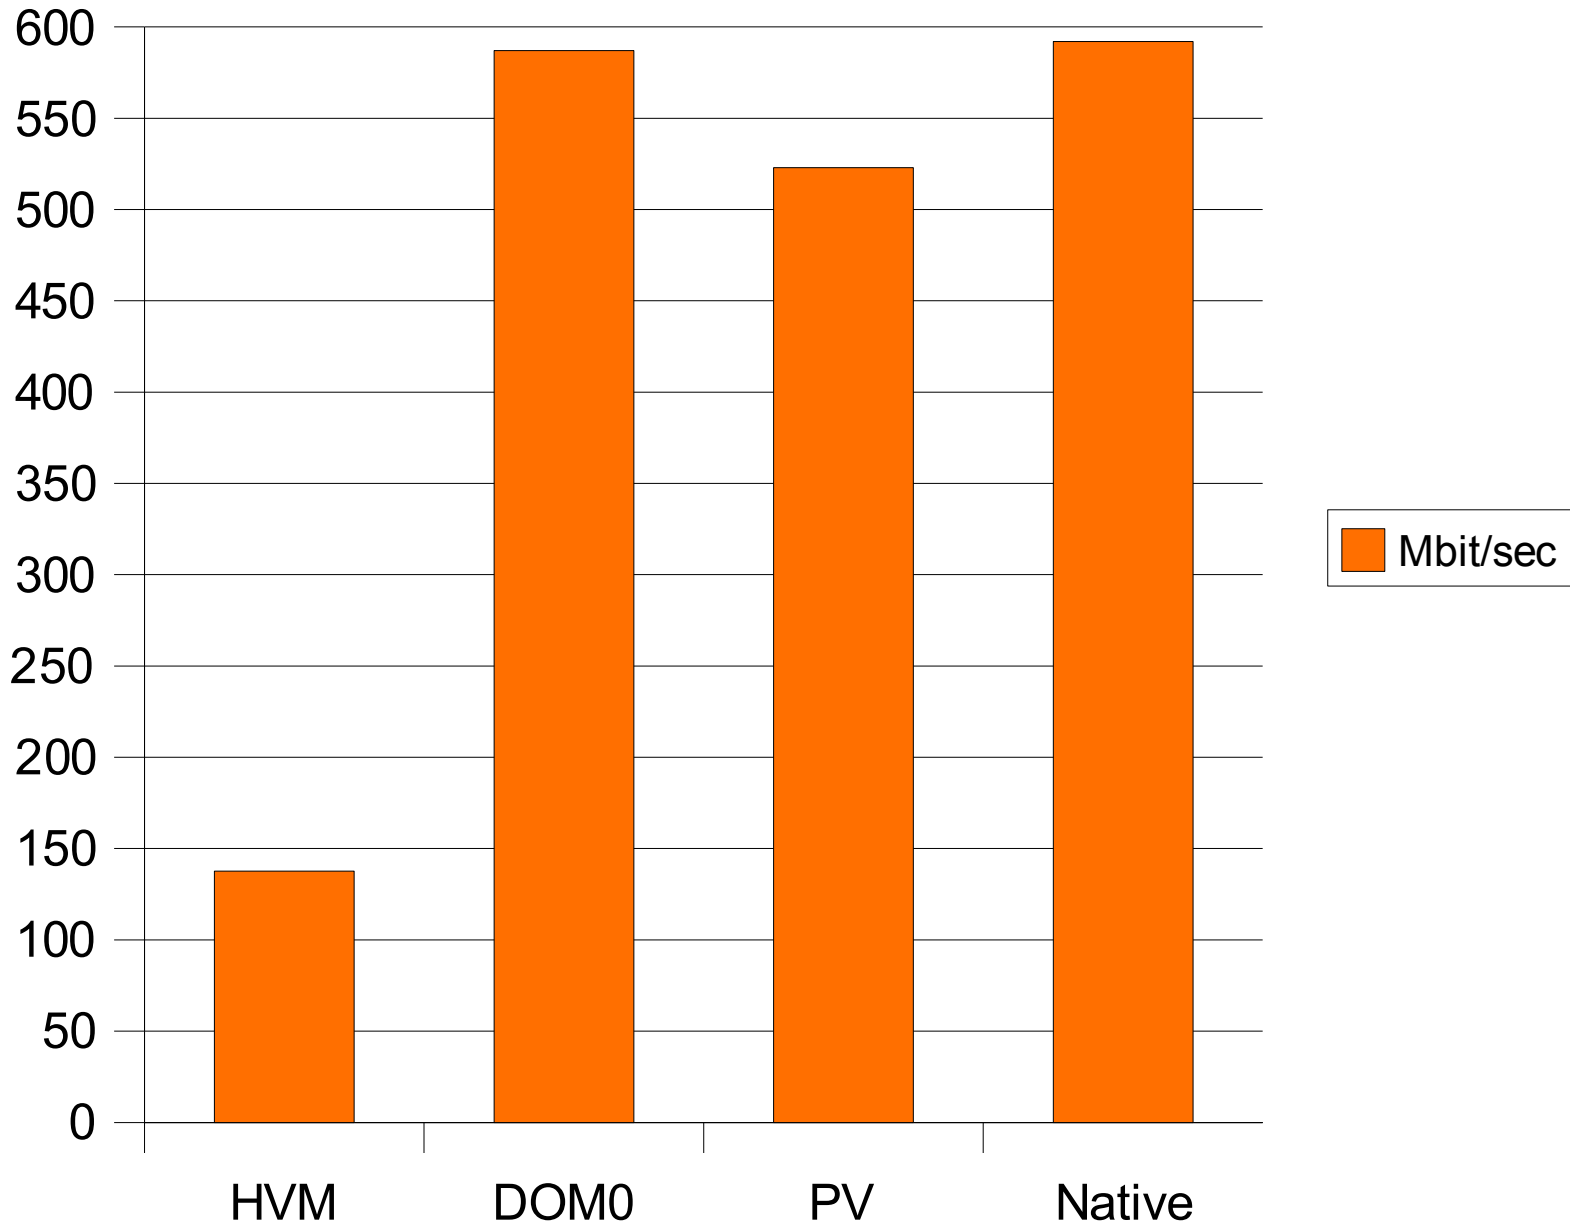





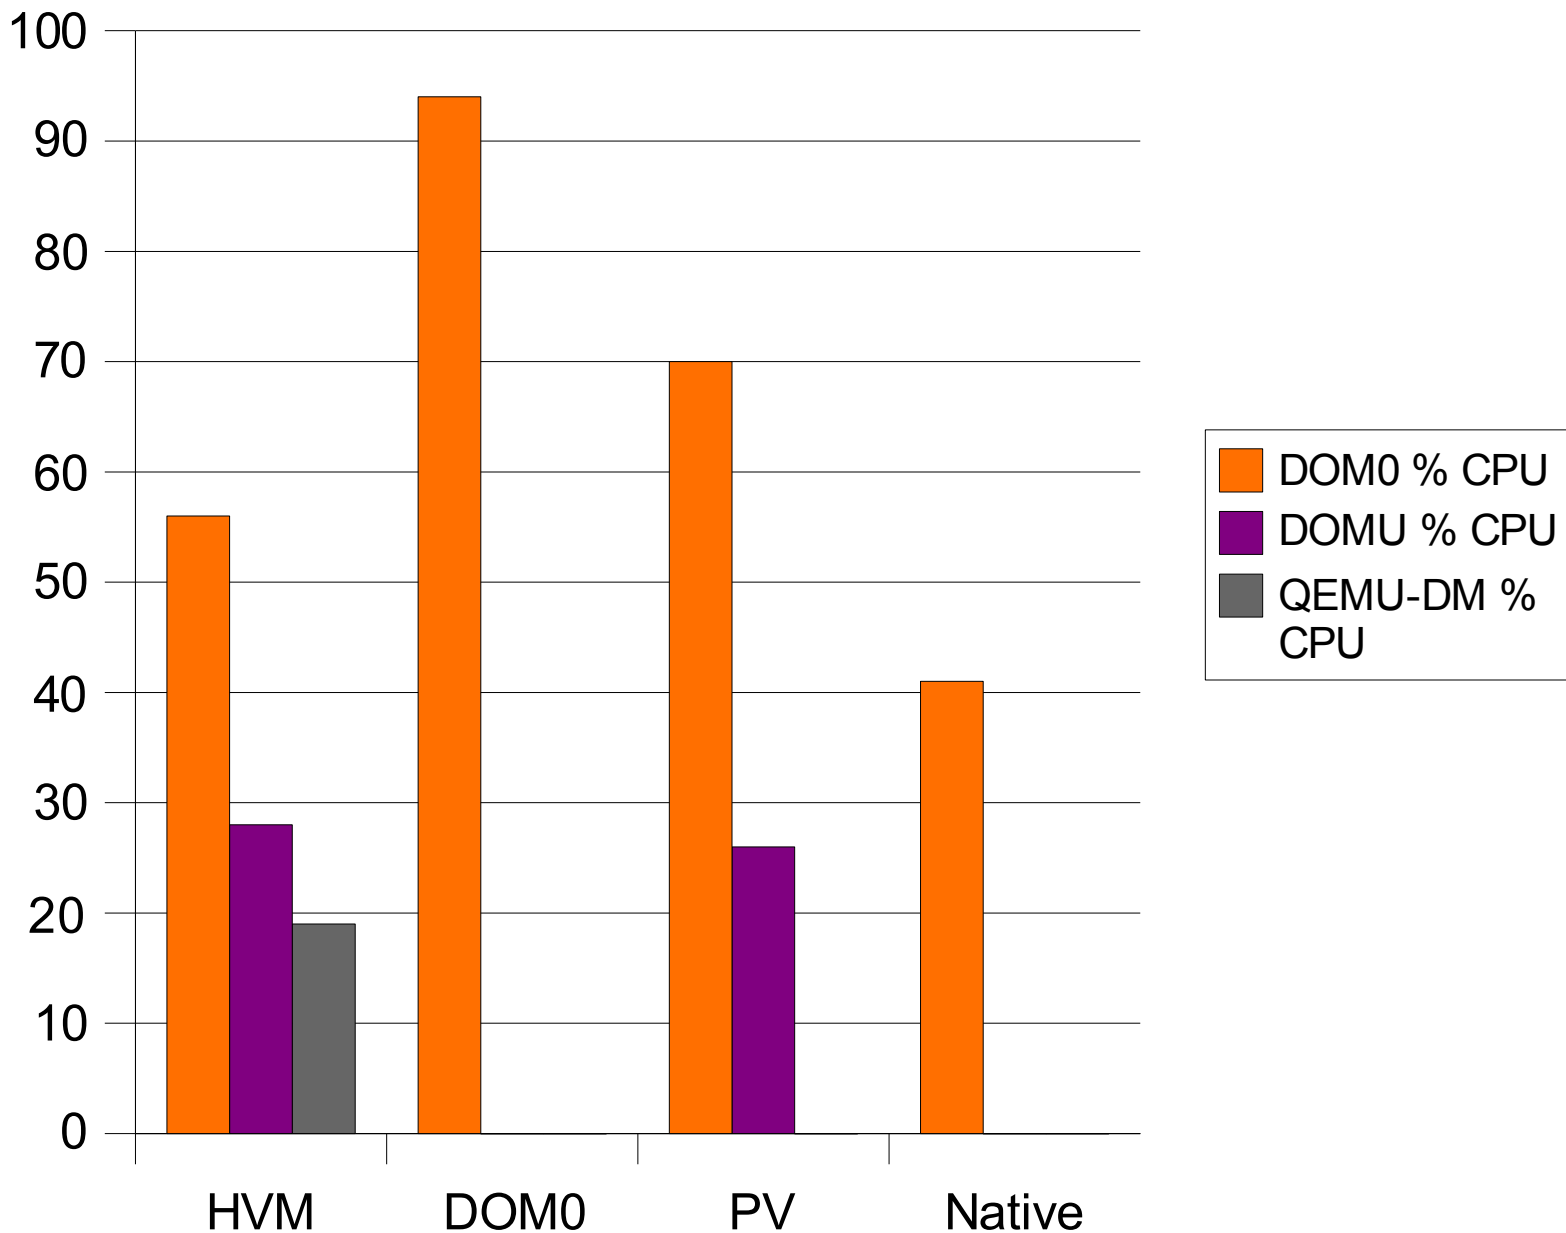
Unmodified Disk IO Performance




Unmodified Network IO Performance




Unmodified Network IO CPU Usage





Dbench3 Scaling 1 to 16 HVM Domains


dbench3 absolute -- hvm-scaling-1-noup






Reaim AIM9 Scaling 1 to 16 HVM Domains


reaim_aim9 absolute -- hvm-scaling-1-noup






SDET Scaling 1 to 16 HVM Domains


sdet2 absolute -- hvm-scaling-1-noup






HVM SCSI Disk

- Add LSI SCSI device to QEMU-DM
- Option ROM for boot support
- DMA batching through LSI “Scripts”



QEMU + Linux AIO + Zero Copy

- QEMU 0.9.0 introduces async IO
 - Linux AIO
 - Important for RAID workloads
- Zero Copy
 - Proper O_DIRECT
 - Avoid page cache
- Sample patch for IDE done for 0.8.2



Questions?

