



IBM

Merging QEMU-DM upstream

Taking the '-dm' out of qemu-dm

Anthony Liguori – aliguori@us.ibm.com
IBM Linux Technology Center
XenSummit

April 16th, 2007

IBM



What's qemu-dm?

- HVM “device model”
 - original device model was based on BOCHS
- QEMU provides a lot more than device emulation
 - Integrated VNC server
 - Redirection of serial, parallel, and monitor to a variety of protocols (TCP, domain socket, etc)
- QEMU almost has an integrated mini-Xend



What's different?

- qemu-dm is based on QEMU 0.8.1
 - 73 patches
 - 13 VNC patches
 - 19 bug fix/non-Xen features
 - 41 Xen enablement
- Some added features:
 - map cache
 - asynchronous DMA
 - enhanced ACPI
 - TPM support
 - various VNC enhancements



Step 1

“The first step in recovery is admitting that you have a problem.”



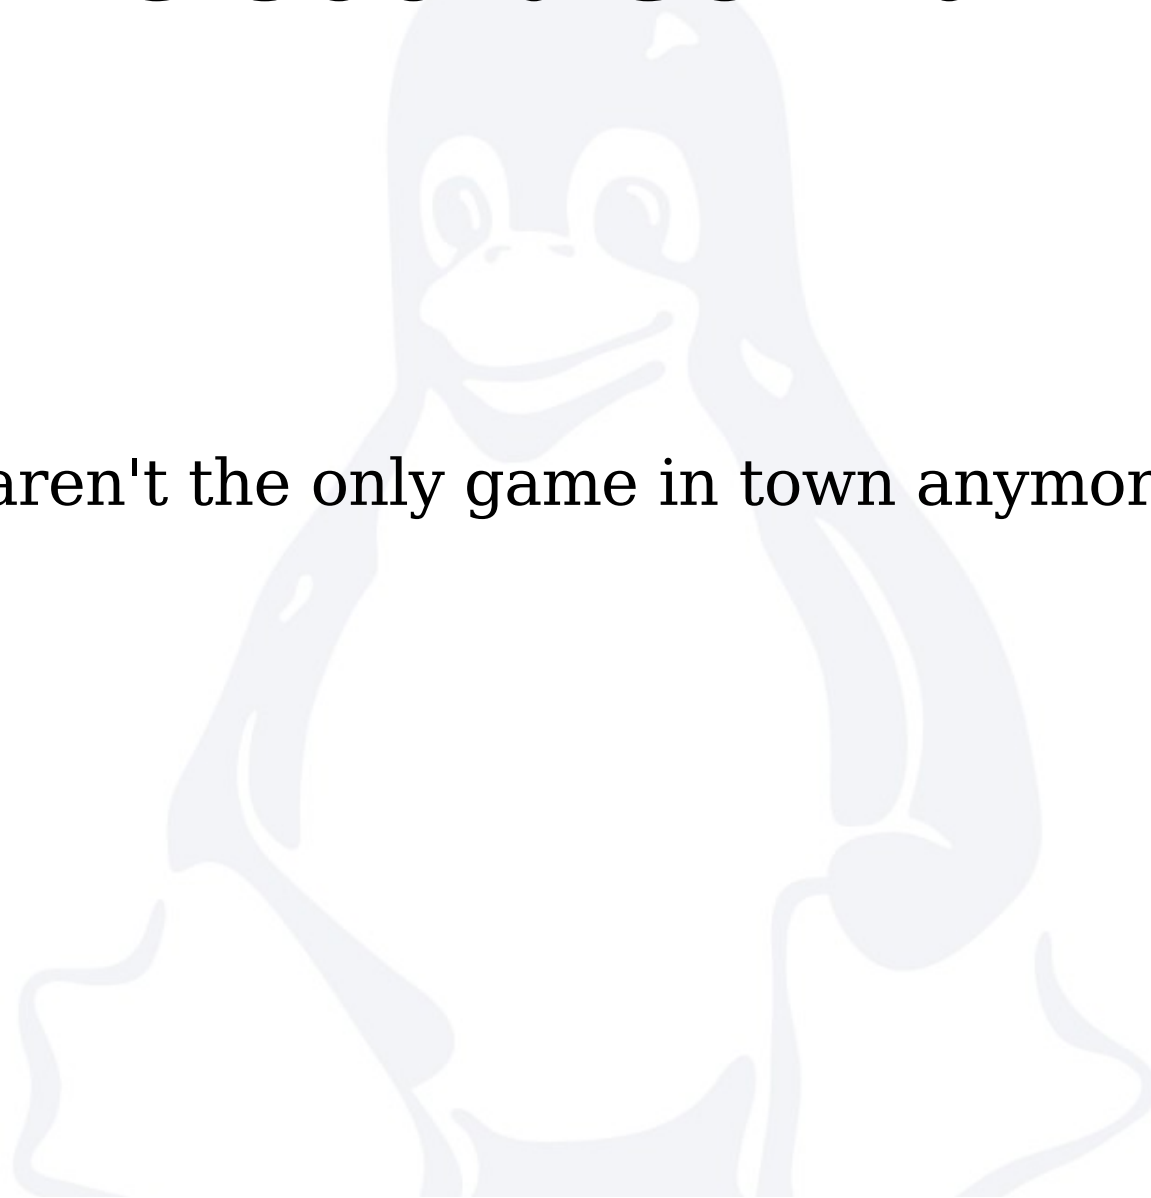
We have a problem

- We are not acting as good members of the community.
- We have been getting better with interacting with the Linux kernel community
 - Jeremy Fitzhardinge and Chris Wright have done a fantastic job here!
- The future of Xen depends largely on how strong its position remains within the Free Software community
- All of the arguments for merging Xen into upstream Linux apply to QEMU



The 800lb Gorilla

We aren't the only game in town anymore.



What can we do?

The least we can do is to stop taking patches for qemu-dm that are not related to core Xen enablement unless they have been submitted to qemu-devel and included in CVS.

A show of good faith would be to submit the patches we currently have that are not Xen related back to qemu-devel.

We should actively work to rearchitect qemu-dm so that it could be included in upstream QEMU.



QEMU community reaction

- The QEMU community is eager to see Xen patches
 - Especially some of the ACPI improvements
- KVM has some similar requirements to Xen
 - Merging Xen should be less painful than we originally thought
- The QEMU community is likely to embrace Xen and help solve some of the problems we face



What will we gain?

- A warm and fuzzy feeling knowing that we've contributed back to a project that played a critical role in Xen's success.
- The distros will not have to ship separate versions of QEMU
- We can make use of all of the new exciting features of QEMU
 - VNC channels
 - VMware VGA emulation
 - vmmouse emulation
 - AIO support
 - Enhanced SCSI support
 - GCC4 support



How it could work...

- Remove the i386-dm target
 - A “target” is absolutely the wrong thing to have here
 - HVM is an i386/x86_64 accelerator
 - At worst, we would need a different machine type
- Remove all of the stuff in qemu-dm that reads configuration from XenStore
 - Configuration should be passed via the command line
- Fold domain creation process into QEMU
 - Users should be able to run QEMU directly in Domain-0 to create an HVM domain
- HVM really doesn't need Xend...



Likely issues

- Some qemu-dm features are fundamentally incompatible with QEMU
 - map cache
- We've broken save/restore/migration compatibility for no obvious gain
 - KVM save/restore images will be compatible with normal QEMU. Migration will also be compatible
 - If we use the QEMU save/restore format, we could migrate from KVM machines to Xen and vice versa.
 - QEMU migration is already richer than Xen migration; migration over SSH, lightweight checkpointing, etc.



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 of 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.



The End.

- Questions?
- Comments?

