

XenSummit Asia

November 2-3, 2011

Seoul, Korea



Xen in Linux (aka PVOPS update)

Ian Campbell

Ian.Campbell@citrix.com

&

Konrad Rzeszutek Wilk

konrad.wilk@oracle.com

Sponsored by:



Agenda

- Current state
- Roadmap
- Distributions
- How you can help
- QA


Current state

Domain 0 Support Added to Linux 3.0

Ongoing work to round out the feature set in Linux
3.1 and onwards

Current state – completed

Feature/Kernel	2.6.37 or earlier	2.6.38	2.6.39	3.0	3.1	future
Xen SWIOTLB	Maintainer happy	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written
ACPI/MSI	Maintainer happy	Needs work	Needs work	Needs work	Needs work	Needs work
Xen PCI	Maintainer happy	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written
KMS(framebuffer),Watchdog	Maintainer happy	Needs work	Needs work	Needs work	Needs work	Needs work
netback	Hasn't been written	Hasn't been written	Maintainer happy	Hasn't been written	Hasn't been written	Hasn't been written
blkback	Needs work	Needs work	Needs work	Maintainer happy	Needs work	Needs work
pciback	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Maintainer happy	Hasn't been written
VGA text	Needs work	Needs work	Needs work	Needs work	Maintainer happy	Needs work
debugfs MMU tracing API	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Maintainer happy	Hasn't been written
Memory hotplug (balloon driver)	Needs work	Needs work	Needs work	Needs work	Maintainer happy	Needs work
Tmem/self-ballooning	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Maintainer happy	Hasn't been written
TTM (3D graphics, or 32bit cards)	Needs work	Needs work	Needs work	Needs work	Needs work	3.3?
Oprofile/ftrace	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Hasn't been written	Needs work
PV ticket spinlocks	Needs work	Needs work	Needs work	Needs work	Needs work	Maintainer happy



Hasn't been written




Needs work



Maintainer happy

Features we would like in future

Kernel/Feature	2.6.37 or earlier	2.6.38	2.6.39	3.0	3.1	future
Netback zero-copy (aka PG_foreign)						3.3?
Netback NAPI						
Netfront SG coalescing						
frontswap						
Blkback/front DIF/DIX (Data Integrity Extensions)						
Blkback/front TRIM/DISCARD						3.2
Blkback/front serial numbers (VPD)						
Hugepages (PV)						
PV Audio (GSOC 2011) – userspace front/black daemons						
PV USB (or use QEMU's USB 3.0 to be written stack)						
ACPI S3 (and hibernate)						
ACPI Cx/Px						
EFI						
Kexec for Dom0 and PvonHVM (GSOC 2011)						



Hasn't been written



Needs work



Maintainer happy

New in 3.1

- Xen-pciback module
 - Last major backend to be included
- Usability improvements
 - e.g. Auto loading of backend modules
 - Helps distros to package / deploy
- Memory Hotplug
- Bug fixes
 - e.g. VGA text console for dom0 fixed

Planned for 3.2

- Hwclock support (date -s)
- Blkback: "feature-barrier" support
- Blkback/front: "feature-discard" support
- PCIback: Support for multi-segmented ("big") boxes
- Kexec/kdump support for PVHVM guests
- Bug fixes

Roadmap 3.3 and onwards

- PV Spinlock support
- ACPI S3
- 3D graphics
- ACPI cpufreq support
- Blkback multiring
- Netback optimisations
- Continue to round out the feature set, usability, rough edges

In other words

- Features to enhance **existing** functionality
- Bug fixes

Pace of development

- Non-merge commits Signed-off-by a Xen MAINTAINER
 - 2.6.39: 145
 - 3.0: 131
 - 3.1: 90
 - 3.2: 122 (5 days into merge window)

Executing roadmap

- No more "dom0" branches in either Jeremy or Konrad's git tree! All development is against Linus's kernel.
- Patches are now against stock kernel.
- Always include the "other" subsystem maintainer:
 - e.g. Block or network subsystem maintainer
 - Use MAINTAINERS file
 - `./scripts/get_maintainer.pl`

OK, so upstream has stuff.

- So I can just install <favorite distro> and use Xen?
 - Yes! (when distributions start shipping 3.0+ kernel)
 - For details visit [XenDom0Kernels](#)
 - Some distros don't enable all backends – please open distro bugs (and let xen-devel know)
- Or you can build a v3.0+ Linux kernel with Xen 4.1.1 on existing distro.
 - Details, explanations, etc: [XenParavirtOps Wiki](#)

How you can help

- Take Linux 3.1 for a spin with Xen 4.1.2...
- ... or even Linux 3.2-rc1 etc if adventurous
- Run it first without Xen to establish a baseline
- Then run it under Xen and see what happens
- Please send e-mail to xen-devel with what works and with what does not.
- As always: XenParavirtOps

Known Bugs in 3.1

- 32-bit graphic cards (ATI ES1000) with Xorg do not work. Has to revert two patches. Ongoing
- 3.1 is still quite new.
- ... Bugs to be documented in: `Linux_31_bugs`
- Please report on xen-devel so we can reproduce, fix and proposed for backport

QA