

Sexy world of Linux kernel pvops project

by Konrad Rzeszutek Wilk

Microsoft Apologizes For Inserting Naughty Phrase Into Linux Kernel

Posted by **timothy** on Thursday July 19, @03:32PM
from the programmers-at-work dept.

Microsoft

[netbuzz](#) writes

"Microsoft has apologized and promised to rectify the fact that one of its developers [slipped a sexist phrase into Linux kernel code](#) supporting Microsoft's HyperV virtualization environment. In that code, the magic constant passed through to the hypervisor reads '0xB16B00B5,' or a slightly camouflaged 'BIG BOOBS.' After Linux developer/blogger Matthew Garrett criticized Microsoft for the stunt, the predictable debate over sexism in the technology world ensued. Microsoft issued a statement to Network World apologizing and added, 'We have submitted a patch to fix this issue and the change will be published in a future release of the kernel.'"

Read the **897** comments



[linux](#) [microsoft](#) [programming](#)

Agenda

- Quick introduction [5 min]
- Past
- Present
- Future
- Surprises
- QA


Paravirtualized Operations (pv-ops)

- Solves the early problem of classic Xen kernels: a distribution needed to ship **two** kernels: one for native, one for Xen.
- Pvops extensions solved the problem - it allows the Linux kernel during runtime to figure out if it is running under any virtualization stacks (Xen, KVM) and swap over to using optimized low-level operations for the specific virtualization stack.

Pvops success


- Its been quite successful - and since its introduction has not had much changes. Lately it has been more of removing some of them as they are not used.
- Majority of Xen-related Linux kernel development is in (next slide)

v3.0-v3.5 kernel Xen development




Development per subsystems

More detailed look



Development by lines of code



What went in v3.X?

<http://wiki.xen.org/wiki/XenParavirtOps#Changelog>

- v3.6: MCE, pcpu, perf fixes, and bug fixes
- v3.5: perf enablement, bug fixes, perf works
- v3.4: ACPI PM work, PV console in HVM, blkback can work in HVM domain
- v3.3: TTM works with Xen, faster IRQ ACK, bug-fixes, netback can work in HVM domain
- v3.2: blkback supporting 'feature-discard', sync wall clock to hypervisor,...

Fast pace of Linux kernel

- Every 2-3 months released!
- Exciting new features
- No stale code!
- Must have code tested two weeks before merge window opens
- Maintainers dream - can concentrate on new features without having to worry about backporting in an distro's kernel.

Surprises! (or not so happy conversations)

- MCE [<https://lkml.org/lkml/2012/4/23/286>]:
"..Because, if you'd hooked into it, just imagine one fine day, when we remove mcelog support, what screaming the xen people will be doing when mcelog doesn't work anymore."
- Followed MCE maintainers desires
- People have misconceptions and sometimes need to explain .. a couple of times.
- We have a responsibility to the general Linux community to help with non-Xen components

Future!

http://wiki.xen.org/wiki/Xen_Development_Projects

- Existing patchsets that apply but are not upstream yet:
 - EFI (only for 3.0)
 - kexec
 - microcode
 - oprofile
 - ACPI S3
 - pv usb
 - pv scsi

Future - new work!

- **blkback faster**
 - persistent grants, larger ring, larger segment
 - DIF/DIX
- **netback faster**
 - persistent grants
 - separate TX/RX rings, large rings
 - separate events, pagepool
 - inline small SKB on ring
- **Hybrid PV aka HVM dom0**
- **perf**
 - sampling the hypervisor
 - sampling the guest

And more future work...

- pv and pvhvm ticketlocks
- other work as it comes from perf analysis.
- .. detailed list on the URL

QA