

Platform Support in Xen 3.0



Keir Fraser

XenSource Inc.

Broader support for x86 platforms



- Xen 2.0 was a big step forward
- Support for full range of Linux device drivers in domain0
- Successfully field-tested on a wide range of workstation and server hardware
- **But** there are still limitations...
 - ACPI BIOS
 - Large memory systems (PAE)
 - X86/64

Supporting ACPI



- ACPI BIOSes are increasingly prevalent
- Require an AML (ACPI Machine Language) interpreter for certain key operations during platform bringup
 - IRQ routing
- We don't want to include full ACPI support in Xen itself
 - The interpreter in Linux is 70k lines of code
 - Extensive interface to user space which we want to keep
- Solution: platform bringup (IRQ, PCI, ACPI) moved to domain0
- Xen brings up SMP and local APICs. Leaves the rest to dom0.
- Preliminary patch provided by Intel
 - Imminent check-in to the 3.0 tree

Large memory systems



- Xen 2.0 is limited to addressing 4GB physical RAM
- Critical limitation in the data centre
 - Installed base of 32-bit systems with 8-16GB RAM
- Increasing demand for Xen to support **PAE**
 - Physical Address Extension
 - Extends 32-bit x86 to allow big-memory configurations
- Under active investigation by Gerd Knorr @ SuSE Labs
 - Receiving incremental patches
- Other groups are also interested in participating

X86/64



- PAE is a hack: the long-term solution is x86/64
- Xen support has been in the 3.0 tree since January
- Preliminary Linux support has just been checked in
 - Thanks to Intel
- Undergoing testing and stabilisation
- Requires further work to boot secondary VMs
- Also: shadow page tables, save/restore, live relocation, full virtualisation