

FreeBSD on Xen

Doug Rabson

What is this Xen thing

- Virtualization platform
- Supports i386, amd64 and others
- 64 bit Xen supports both 32 bit and 64 bit guests
- Full hardware virtualization to ease porting
- Paravirtualization for extra performance

How does it work?

Hypervisor

- A cut down operating system that boots and runs on the real hardware
- Manages guest virtual machines and their memory
- Supports shared-memory and event channels between guests
- Starts the special Dom0 guest

How does it work?

Dom0

- Privileged guest virtual machine
- Has direct access to physical hardware devices (interrupts become Xen events)
- Responsible for starting further guests and allocating resources for them
- Handles the 'back end' of guest virtual devices

How does it work?

DomU

- Each regular guest is a DomU virtual machine
- Virtual devices use hypervisor shared memory and events to talk to back end drivers running in Dom0

How does it work?

Paravirtualization

- Modified operating system uses hypervisor services to help manage virtual address spaces and communicate with other guests.
- Most communication is with Dom0 for device access
- Dom0 itself must be paravirtualized
- FreeBSD supports i386 only

How does it work?

Hardware virtualization

- HVM guest looks like real hardware to the guest kernel
- Emulated BIOS and devices for a simple hardware platform
- FreeBSD supports i386 and amd64
- Uses Intel and AMD extensions to help

How does it work?

Hardware virtualization

- Modified guest kernel can detect HVM
- Can use paravirtualized device drivers to improve performance
- Can use hypervisor to support suspend, resume and live migration

FreeBSD on Xen Paravirtualized

- FreeBSD/i386 supports paravirtualized mode
- Driver support for virtual block and network interfaces, suspend, resume, migrate
- All the code is in FreeBSD-current
- Use the provided XEN kernel config

FreeBSD on Xen

Hardware virtualization

- FreeBSD/i386 and FreeBSD/amd64 work fine
- FreeBSD/amd64 has support for paravirtualized drivers, suspend, resume, migrate
- All the code is in FreeBSD-current
- Use the provided XENHVM kernel config

Limitations

Paravirtualized

- Paravirtualized kernel only supported on 32bit FreeBSD
- Only a few drivers currently ported (console, virtual disk, virtual network interface)
- Dom0 not supported

Limitations

Hardware virtualization

- Hardware support required (SVM or VMX)
- Recent version of Xen required (3.3?)
- VM emulation is relatively expensive (shadow page tables)
- Adaptive mutexes don't work well - disabled in sample kernel config

Further work

- Paravirtualized drivers for FreeBSD/i386 on HVM (easy)
- Full paravirtualization for FreeBSD/amd64 (harder)
- Add Dom0 support to FreeBSD (hard)

Thanks to:

- Kip Macy for the original Xen PV support
- Citrix for funding the HVM work

Questions?