


Porting FreeBSD on Xen on ARM

How to support your OS as Xen ARM guest

Julien Grall
`julien.grall@linaro.org`

FOSDEM – February 1, 2014

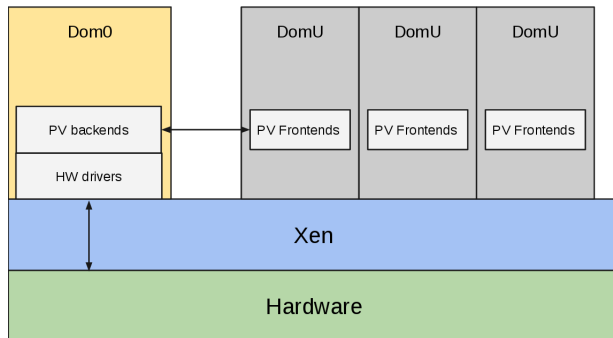


Xen




- ▶ Type-I hypervisor
- ▶ Support for ARM v7 and ARM v8 with virtualization extension
- ▶ Hardware support
 - ▶ Fast Model
 - ▶ Versatile Express Cortex A15
 - ▶ Arndale Board
 - ▶ Allwinner A20/A31 (SunXi)
 - ▶ TI OMAP5
 - ▶ Applied Micro X-Gene
 - ▶ Calxeda "Midway"
 - ▶ ...

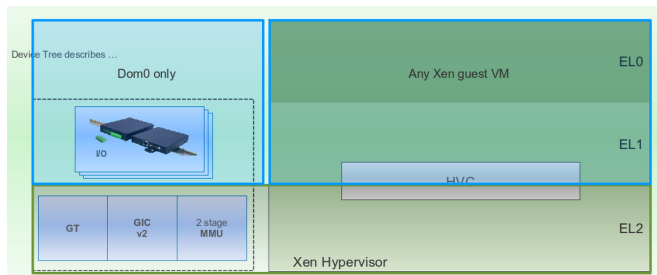
Xen architecture



ARM architecture



Xen on ARM architecture



Requirements



- ▶ Guest boot ABI
- ▶ Device tree support
- ▶ Specific memory attribute
- ▶ Xen PV drivers
- ▶ Copy of *xen/include/public*
 - ▶ *arch-arm.h* provides hypercalls convention

Guest boot ABI



Interface of the virtual machine:

- ▶ Linux zImage
 - ▶ Specific values on some registers
 - ▶ $r0 = 0$
 - ▶ $r1 = 0xffffffff$
 - ▶ $r2 = \textit{Device Tree physical address}$
 - ▶ MMU disabled
 - ▶ Data cache disabled
 - ▶ Instruction cache in an unknown state
- ▶ ELF (in progress)
- ▶ Use of PSCI to bring up secondary CPUs

Device Tree



- ▶ Basic Device Tree generated by the toolstack which contains:
 - ▶ CPUs
 - ▶ Memory
 - ▶ Timer
 - ▶ GIC
 - ▶ Hypervisor
- ▶ The guest should use the values from the Device Tree
- ▶ Working group to decide core bindings

Memory



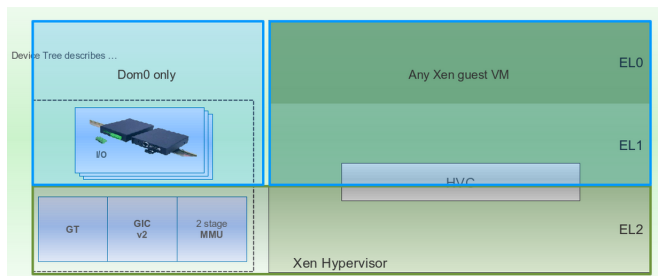
- ▶ Before calling hypercalls the OS must enable:
 - ▶ MMU
 - ▶ Data and Instruction cache
- ▶ RAM attribute should be *Write-Through* or *Write-Back*

Xen PV drivers



- ▶ Xen core architecture
 - ▶ Xenstore
 - ▶ Grant-Table
 - ▶ Event-channel
- ▶ Xen device drivers
 - ▶ Console
 - ▶ Block
 - ▶ Network
 - ▶ Framebuffer (need to recompile QEMU)
- ▶ Drivers already available under BSD license in FreeBSD

DOM0



DOM0 (2)



- ▶ First guest to start
- ▶ Nearly every devices are assigned to DOM0
 - ▶ Serial, IOMMU, Timer and GIC are used by Xen
 - ▶ Some devices can be blacklisted by Xen
- ▶ DOM0 kernel should use the Device Tree to discover the hardware

FreeBSD on Xen



- ▶ Support for x86 PVHVM
- ▶ Experimental support for ARM architecture
- ▶ Patch series to support Xen on ARM guest
 - ▶ New kernel config **XENHVM** created
 - ▶ Non-modular guest configuration (memory, ...)
 - ▶ Only support for guest with 1 VCPU

Device Tree



Device Tree is hardcoded:

- ▶ Missing support to use Device Tree with Linux boot ABI
- ▶ FreeBSD enumerates the devices in DTB order
 - ▶ The interrupt controller should be loaded first
 - ▶ Discussion made on FreeBSD ML ARM
- ▶ Some bindings are different
 - ▶ FreeBSD only supports interrupt with 2 cells
 - ▶ Work in progress to handle 3 cells
 - ▶ On-going work to standardize the bindings

Memory



- ▶ Modify early page table attribute
 - ▶ Use Write-Back instead of Write-Through
 - ▶ On-going patch to resolve the issue in Xen
- ▶ Create a new pmap function to map memory region cacheable
- ▶ FreeBSD requests to be loaded at a specific physical address

Xen PV drivers



- ▶ Update interface headers to Xen 4.4
 - ▶ FreeBSD is based on Xen 4.2 headers
 - ▶ ARM interface was not set in stone
- ▶ Drivers common with x86
 - ▶ Use the right xen type (`xen_pfn_t`, `xen_ulong_t`,...)
 - ▶ Support for HVM in console drivers
- ▶ Rework event channel handling
 - ▶ was x86 specific
 - ▶ still missing features
 - ▶ suspend/resume
 - ▶ pirq

What needs to be done?



FreeBSD port to Xen on ARM has started. Still to come...

- ▶ Add support for Device Tree loading via Linux Boot ABI
- ▶ Uniform Xen drivers across the different architecture
- ▶ Guest SMP support
- ▶ DOM0 support
- ▶ Stability



OS supported by Xen



- ▶ Out-of-box
 - ▶ Linux based distribution
- ▶ Future support
 - ▶ FreeBSD
 - ▶ Erika OS
 - ▶ other ***BSD**

Questions?



- ▶ Xen devel ML: `xen-devel@lists.xenproject.org`
- ▶ Xen user ML: `xen-user@lists.xenproject.org`
- ▶ **#xenarm** on freenode

Fin

