

NAME

xen - Xen Hypervisor Guest (DomU) Support

SYNOPSIS

To compile hardware-assisted virtualization (HVM) Xen guest support with para-virtualized drivers into an amd64 or i386 kernel, place the following lines in your kernel configuration file:

```
options XENHVM  
device xenpci
```

DESCRIPTION

The Xen Hypervisor allows multiple virtual machines to be run on a single computer system. When first released, Xen required that i386 kernels be compiled "para-virtualized" as the x86 instruction set was not fully virtualizable. Primarily, para-virtualization modifies the virtual memory system to use hypervisor calls (hypercalls) rather than direct hardware instructions to modify the TLB, although para-virtualized device drivers were also required to access resources such as virtual network interfaces and disk devices.

With later instruction set extensions from AMD and Intel to support fully virtualizable instructions, unmodified virtual memory systems can also be supported; this is referred to as hardware-assisted virtualization (HVM). HVM configurations may either rely on transparently emulated hardware peripherals, or para-virtualized drivers, which are aware of virtualization, and hence able to optimize certain behaviors to improve performance or semantics.

FreeBSD supports hardware-assisted virtualization (HVM) on both i386 and amd64 kernels.

Para-virtualized device drivers are required in order to support certain functionality, such as processing management requests, returning idle physical memory pages to the hypervisor, etc.

Xen DomU device drivers

These para-virtualized drivers are supported:

balloon Allow physical memory pages to be returned to the hypervisor as a result of manual tuning or automatic policy.

blkback Exports local block devices or files to other Xen domains where they can then be imported via **blkfront**.

blkfront Import block devices from other Xen domains as local block devices, to be used for file systems, swap, etc.

- console** Export the low-level system console via the Xen console service.
- control** Process management operations from Domain 0, including power off, reboot, suspend, crash, and halt requests.
- evtchn** Expose Xen events via the `/dev/xen/evtchn` special device.
- netback** Export local network interfaces to other Xen domains where they can be imported via **netfront**.
- netfront** Import network interfaces from other Xen domains as local network interfaces, which may be used for IPv4, IPv6, etc.
- pcifront** Allow physical PCI devices to be passed through into a PV domain.
- xenpci** Represents the Xen PCI device, an emulated PCI device that is exposed to HVM domains. This device allows detection of the Xen hypervisor, and provides interrupt and shared memory services required to interact with the hypervisor.

Performance considerations

In general, PV drivers will perform better than emulated hardware, and are the recommended configuration for HVM installations.

Using a hypervisor introduces a second layer of scheduling that may limit the effectiveness of certain FreeBSD scheduling optimisations. Among these is adaptive locking, which is no longer able to determine whether a thread holding a lock is in execution. It is recommended that adaptive locking be disabled when using Xen:

```
options NO_ADAPTIVE_MUTEXES
options NO_ADAPTIVE_RWLOCKS
options NO_ADAPTIVE_SX
```

HISTORY

Support for **xen** first appeared in FreeBSD 8.1.

AUTHORS

FreeBSD support for Xen was first added by Kip Macy <kmacy@FreeBSD.org> and Doug Rabson <dfr@FreeBSD.org>. Further refinements were made by Justin Gibbs <gibbs@FreeBSD.org>, Adrian Chadd <adrian@FreeBSD.org>, and Colin Percival <cperciva@FreeBSD.org>. This manual page was written by Robert Watson <rwatson@FreeBSD.org>.

BUGS

FreeBSD is only able to run as a Xen guest (DomU) and not as a Xen host (Dom0).

As of this release, Xen PV DomU support is not heavily tested; instability has been reported during VM migration of PV kernels.

Certain PV driver features, such as the balloon driver, are under-exercised.