

bhyve topics

grehan@freebsd.org

BSDCan 2013

AMD SVM support

- Demo'd at the March BAFUG
- `/projects/bhyve_svm`
- “last” patch submitted by Anish Gupta
 - should be submitted shortly
- Developed on Phenom II
 - would like coverage on other models

Modularization

- Move to loadable modules
 - BHYVE_SYSINIT
- Allows external modules
 - device models
 - e.g. non base-system code (graphics)
 - non BSD-licensed code
- But, requires stable APIs

Configuration

- Command line unable to express deep hierarchy
- Thinking of using YAML config file
 - human readable
 - better for libvirt integration?
- Provide script front-ends for ease-of-use
e.g. `bhyve.small`

Memory overcommit

- Guest memory wired at VM create
- Should be done via VM object
 - Accounting, swap, superpages etc
 - path to live migration
- May require VM system hooks
 - e.g. Linux mm-notifiers

Networking

- have per-dev thread for tx; shared kqueue thread for rx
- 1 Gbps throughput, high host CPU %
- virtio: TSO, multiqueue
- backend: netmap, open-vswitch, better tap
- alt dev models e.g. e1000

Storage

- currently virtio-block, sync backend
- have async patch
- AIO if available ?
- want help with VM image formats - VMDK etc, BSD-licensed

PCI pass-thru

- Currently hard-coded to Tylersburg IOH
- Could extract VT-d registers from ACPI DMAR table
- or, hard-code other IOH ID's/offsets
- AMD work in progress
- FreeBSD IOMMU framework ?

Logging

- Difficult to diagnose issues
- Requires high-speed logging
 - kernel: KTR (style) ?
 - user-space: varnish-like ?

Security

- Requires root privilege: shouldn't
- network: "user" ifnet ? Slirp-style NAT ?
- requires memory overcommit code
 - accountable/swappable guest mem

Non FreeBSD guests

- UEFI (OVMF) + CentOS 6.2
 - contributed by Tycho Nightingale @ Pluribus Networks
 - UP-only
- Partial OpenBSD
 - loaded with hacked grub 1.98

Open Floor

- What do people want to see ?