



# Xen Hypervisor Update

Keir Fraser, Citrix Systems Inc.

- Current stable release: 3.4.2
  - Released just last week
- Next releases: 3.4.3 and 4.0.0
  - Both anticipated early in the New Year
- Ongoing strategy:
  - Maintain two stable branches until the later one has matured enough for switchover
  - Quarterly releases from stable branches
  - Six to nine months between major releases

- Efficient support has been upstream for a couple of years now
  - Running well as 'domU' VM
- Jeremy Fitzhardinge manages the tree for further development
  - Path for further upstreaming to kernel.org
  - Default kernel for Xen 4.0
- Still to do:
  - Port your favourite backend driver
  - Test!

- Support VHD storage format
- Snapshot virtual disks
  - Checkpointing, gold images, ...
- Simplified kernel support
  - Leverage blkback
- Simpler invocation model
  - Command-line driven rather than xenstore
  - More generic, easier test and debug

- Maintain an up-to-date replica of a VM on a backup server
  - Regular (50ms) checkpoints
  - Buffer I/O over this interval (IMQ, blktap2)
- Fail over to backup if primary fails
  - Simple heartbeat protocol
  - But scope for using more featureful external monitors such as Linux-HA
- Simple and transparent

- Network virtualisation is particularly hard
  - High packet rates; latency sensitive
- Old netfront/back drivers have limitations
  - High cost for packet receive
  - Not designed for next-generation NICs
- Ongoing work on netchannel2 to address this
- Lazy copy in the guest (reduces dom0 load)
  - Provide guest a copy-only, sub-page, revocable grant
- Support multi-queue NICs
  - DMA directly to guest buffers
- Smart NICs, SR-IOV
- Reusable extensible ring architecture

- Potential for reducing memory pressure by sharing identical pages across VMs
  - Significant savings in ‘ideal’ cases
  - Rather smaller gains in typical heterogeneous scenarios (10-20%)
- How to find identical pages?
  - Identical disc blocks
- Allows memory overcommit
  - Hence requires demand paging
  - Or don't give spare pages directly back to guests

- A range of toolstacks now in use
  - Xend, xapi, xenvm
- Share common code at the lowest level
  - Libxenctrl, libxenguest
- But still a good deal of shareable infrastructure is implemented separately in each toolstack
- Libxenlight: higher-level operations which can be shared among all toolstacks
  - Implemented in C, licensed under LGPL
  - Domain creation, destruction, migration, PCI passthru, ...
- Plans to move all toolstacks over to this library (and delete equivalent but non-shared code)



- Use the developer mailing lists
  - [xen-devel@lists.xensource.com](mailto:xen-devel@lists.xensource.com)
  - {xci-devel,xen-api}@lists.xensource.com
  - Requests for comments
  - Announcements of dev plans
  - Patches posted for comment, review and checkin
- Use the wiki
  - <http://wiki.xensource.com/xenwiki/XenRoadMap>