



**Welcome to the  
Xen Summit**

September 7<sup>th</sup>-8<sup>th</sup>, 2006

# Schedule



Day/Time	Track A	Track B
Thurs 08.30-09.00	Session 1 Intro/plenary (C)	
Thurs 09.00-10.30	Session 2A Xen Core Chair: Keir Fraser	Session 2B Security #1 Chair: Steve Hand
Thurs 11.00-12.30	Session 3A HVM Core Chair: Ian Pratt	Session 3B Security #2 Chair: Reiner Salier
Thurs 13.30-15.00	Session 4A Non x86 Chair: Alex Williamson	Session 4B Performance Chair: Ian Pratt
Thurs 15.30-17.00	Session 5A Testing Chair: James Bulpin	Session 5B HVM Devices Chair: Christian Limpach
Thurs 18.30-+	Dinner and Billiards	

# Schedule



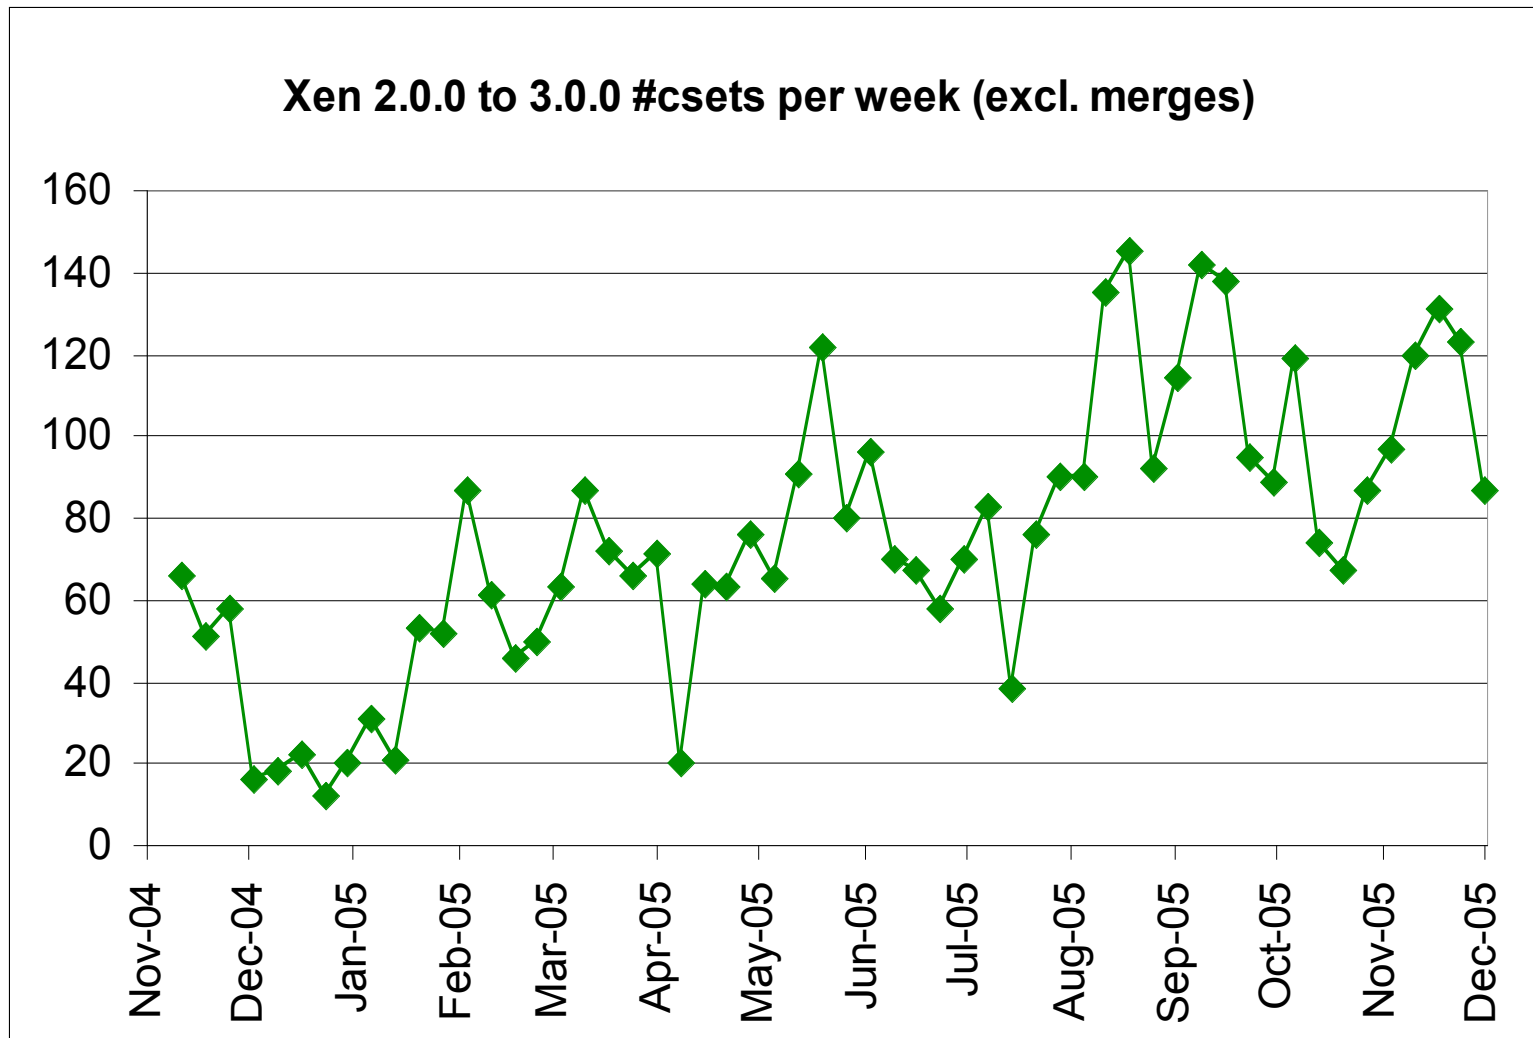
<b>Day/Time</b>	<b>Main Topic</b>	<b>Chair</b>
Fri 09.00-10.30	Session 6A Virtual Devices Chair: Andy Warfield	Session 6B OS Ports Chair: Chris Wright
Fri 11.00-12.30	Session 7A Smart IO H/W #1 Chair: Ian Pratt	Session 7B Control Stack #1 Chair: Ewan Mellor
Fri 13.30-15.00	Session 8A Smart IO H/W #2 Chair: Steve Hand	Session 8B Control Stack #2 Chair: Jim Fehlig
Fri 15.30-16.30	Session 9 Closing/plenary (C)	

# Mini Summit Agenda

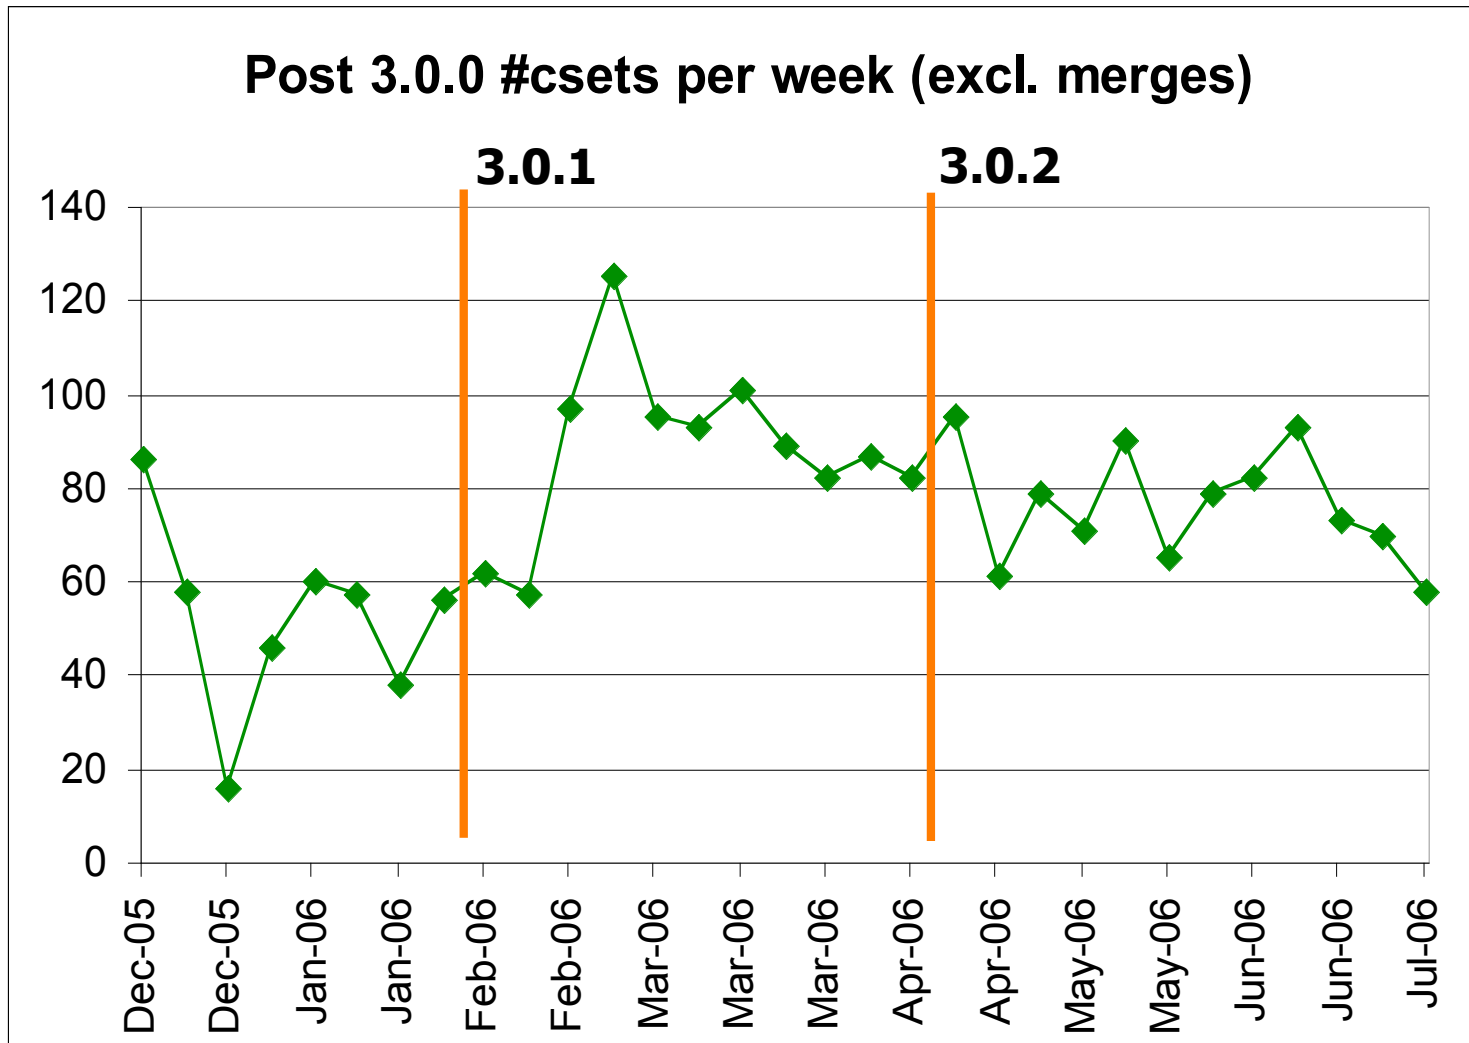


- Facts and figures
- Roadmap presentation/discussion
- Linux upstream merge status (Chris / Keir / Rusty)
- ia64/power status (Alex Williams / Hollis)
- Xen API / xm extensions proposal (Ewan)
- HVM IO development (Jun / Leendert)
- File-based disk images: blktap/qcow (Andrew)
- Misc: Kexec, bootloaders, GSO

# Xen 3-unstable Change Log



# Post-3.0.0 Change Log



# Post-3.0.0 Rough Code Stats

	<b>aliases</b>	<b>checkins</b>	<b>insertions</b>
<b>xensource.com</b>	15	1020	319786
<b>intel.com</b>	24	225	25343
<b>ibm.com</b>	28	193	36331
<b>valinux.co.jp</b>	2	137	10958
<b>bull.net</b>	1	104	6350
<b>hp.com</b>	8	93	18008
<b>fujitsu.com</b>	11	72	2812
<b>amd.com</b>	6	58	2563
<b>novell.com</b>	7	56	2860
<b>virtualiron.com</b>	5	23	1434
<b>apana.org.au</b>	1	21	3080
<b>verge.net.au</b>	1	18	296
<b>ncsc.mil</b>	1	17	5024
<b>redhat.com</b>	4	17	265
<b>cam.ac.uk</b>	1	9	1211
<b>other</b>	24	173	33210

Stats since 3.0.0 Release

# Post-3.0 Development Model



- Stabilize “unstable” tree every 8-12 weeks
  - Sweep unstable into 3.0-testing
  - Release as 3.0.1, 3.0.2 etc
- Bug fixes cherry picked into 3.0-testing until next release
  - After being in unstable for a few days, requests to push into 3.0-testing
  - Release as 3.0.x-y etc.
  - (much like Linux)

# Achievements: 3.0.1/3.0.2



- AMDV support, generic 'HVM' layer
- Linux 2.6.16 upgrade in –unstable
  - subarch of i386/x86\_64
  - Linux tip maintained in linux-2.6.tip-xen.hg
- Progress on linux patch upstreaming
  - Patch queue in linux-2.6-patchqueue.hg
- PCI pass-through is back
  
- Bug fixes, stabilization for SLES10

## 3.0.3 Release



- Planned for before OLS, postponed to synchronize with FC6 freeze and drive more features:
  - New CPU scheduler with CPU migration
  - Blktap/qcow for file-backed virtual disks
  - Upgrade qemu for HVM IO (usb, vnc)
  - New shadow pagetable code
  - PV extensions to HVM guests (net, block IO)
  - *PV framebuffer support*
  - *NUMA memory allocator support*
  - *Dom0 kexec/kdump support*
  - *Xend life-cycle management*
  - *Full Segmentation Offload in netfront/back*
  - *Power architecture merge*

## 3.0.4 Release



- Early Q4 2006 (October)
  - (leftovers from above)
  - Xen control API
  - XML config files, extended 'xm' syntax
  - QEMU 'v2e' integration
  - Performance optimizations
  
- Should we call this release 3.1?

# Performance



- Performance and scalability work
  - Time is right for some close attention
  - 1-4 socket systems the priority
  - Optimizations for bigger systems must not hurt smaller ones (they often help)
    - Onus is on submitter to demonstrate
    - (Patches that clearly hurt larger systems should be rejected too)
- Good performance tools now available
  - s/w perf counters, xen oprofile, tracebuf etc

# API stability



- Guest API stability (hypercall, IO)
  - Backward guaranteed:
    - Old 3.0 guests must run on new xen
  - Need to add forward compatibility as well
    - Important now SLES10 hypervisor in the wild
  - [Linux upstreaming may force API change]
- Privileged domain hypercall API
  - Over time, stabilise dom0 priv domain API
  - Still in flux through 2006
- Xen API control protocol and CLI syntax
  - Finalize wire format, create bindings

# Tools (**xm/xend/CIM**)



- VM life-cycle management
- XML config file format adoption
  - xm syntax changes
- Xen control API
  - C/C++/perl/python bindings
  - Platform for CIM providers
  - https/unixdomain sockets/ssh transport
- Guest coredump support
- Guest bootloader support

# HVM Support



- Further shadow pagetable optimizations
- Allow PV hypercalls from HVM guests
- HVM save/restore support; live relo
- Initial SMP guest support; ACPI support
- Superpage PTE support
- Enhanced device emulation
  - Move emulator to stub domain
  - Change interface to hypervisor ("V2E")
  - Better device models

# Core Xen



- Live relo on PAE/x86\_64; tuning
- Support for 32b PAE PV guests on 64b xen
- NUMA allocation/policy
- Scheduler tuning
- Remove artificial x86\_64 heap sz restriction
- Superpage support
- Scheduler tuning, partic. for SMP guests
- Power management

# Linux



- Extensive benchmarking and tuning
- x86\_64 optimizations
- PV framebuffer (plus X extensions)
- Avoid balloon driver 'crunch' issues
- SMP scalability improvements
- USB virtualization
- OpenGL virtualization (Chromium)

# Storage



- Support for file-based virtual disk storage
  - blktap/ublkback
  - dm-userspace
- Adopt qcow as favoured xen disk format
  - Compatible optimizations vs. v2 format
- Block IO QoS via ionice and/or blkback
- Media change, resize events
- SCSI-level front/back (tape, CD writer)

# Network



- Segmentation Offload support
- Cleanup up checksum offload
- Copy vs. page-flip tuning
- Streamline linux bridge code
- TOE/RDMA support
- Point-to-point inter-VM NIC support

# Misc



- Dom0 kexec/kdump support
- Bootloader support
- xenfs filesystem virtualization