User-Mode Linux

An Introduction to UML

William Stearns
ISTS/Dartmouth College

Introduction to UML

- Virtual Machine
- Port of Linux kernel to system call architecture
- System call proxy

Functional block diagram

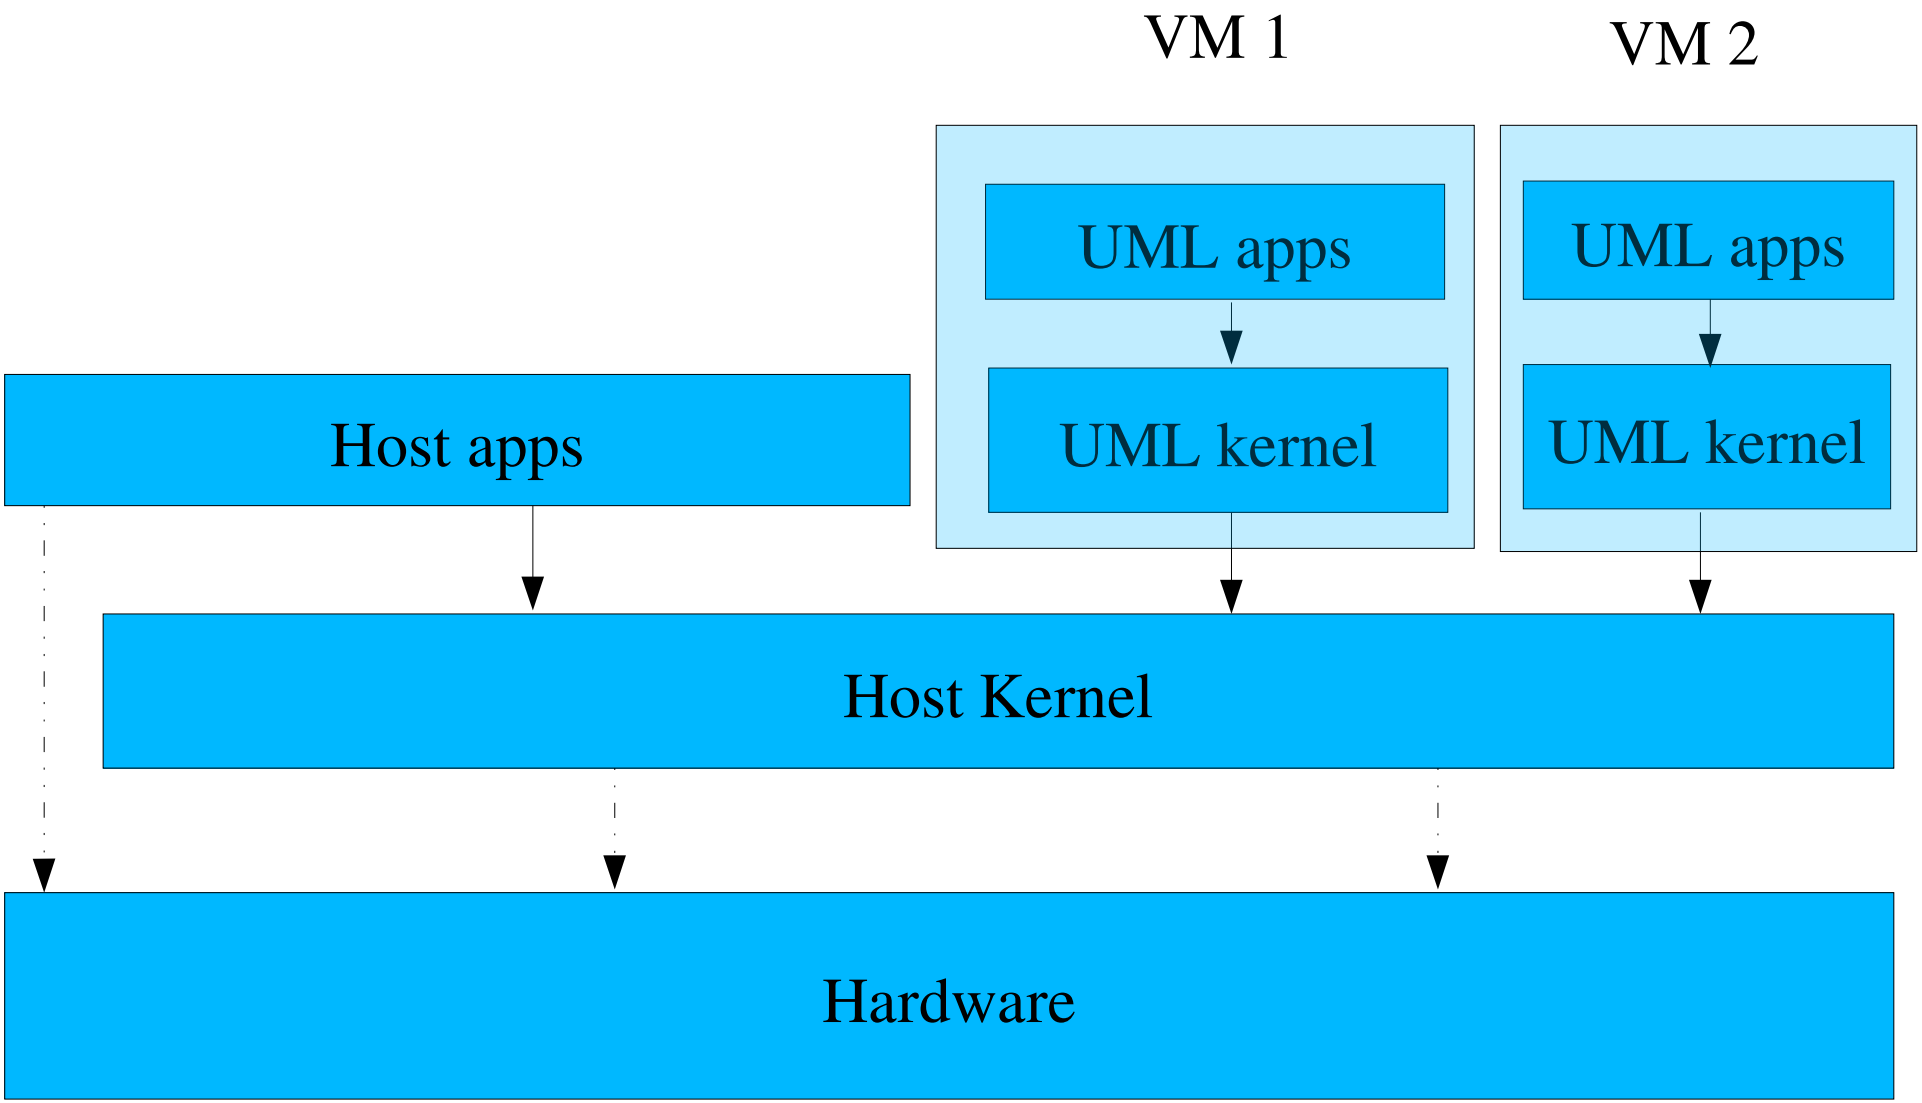

Why?

- Compartmentalization
  - By user/customer/team
  - By task
  - ISP Virtual hosting services
- Safe root access
  - Classes
  - Ever wanted to see how long a system lasts with “rm -rf /”?

Why?

- Controlled malware analysis
- Distribution of complete work environment
- Running different/newer/older Linux distributions
- Kernel development and debugging
  - Except hardware device drivers

Comparison to...

- VMWare
- Bochs/Qemu

Operation

- ELF binary, commonly run under screen
- Runs as non-root user
  - One-time root setup of networking helpful, can work around with slirp
- Need compiled UML kernel and root filesystem

Memory

- Provided as a block at startup
- Not allocated until needed
- Double caching/buffering inefficient
- May be a limit ~600M
  - Host can still cache VM swap
  - Probably gone in x86_64
- linux ... mem=384M

Virtual block device

- File, partition, disk, ramdisk on host
  - Floppy RAID array, anyone? :-)
- Presented as a block device to VM
- Swap space is just a swap-formatted file on host
- Files can be sparse
- linux ... ubdN=/path/to/host_file

Hostfs

- Access to host files
- Synchronization not perfect
- Can even be used for /
  - Privacy loss
- UML kernel needs to support it
- VM /etc/fstab
  - none /pub hostfs defaults,/home/pub 0 0

Copy On Write

- One read-only (immutable!) underlying file
- One Copy-On-Write layer for each VM
- Save COW file for system snapshot
- linux ... ubdN=/path/to/cow,/path/to/pristine

Networking

- UML hooks into host kernel
  - tun or tap device
  - uml_switch acts as a router/switch
  - pcap interface
- All upper level stuff just works
- linux ... eth0=daemon,C0:FF:EE:C0:FF:EE

Audio proxy

- OK, this is getting insane. :-)
- Audio plays out to host sound card

Performance

- Lose 3%-15% CPU
  - Faster CPU :-)
- No penalty except when performing system calls
- Slight increased latency

Hardware requirements

- Almost nothing, all the way up to insane... :-)
- Works fine on old, slow hardware
- Enough disk, ram for host + all VMs
- Add more RAM!

Controlling VM from host

- uml_mconsole
  - pause/start/halt/reboot/c-a-d
  - Sysrq-N
  - Add/Remove virtual device
  - Log to VM kernel log
  - Show /proc/N
- nice

Host considerations

- tmpfs on /tmp
- Clean shutdown
  - ssh preferred
  - uml_mconsole
  - kill -9 :-)
- Thundering herd at boot, 4AM

Application Considerations

- Almost none!
  - Disable /lib/tls/
- Essentially no direct hardware access
  - hwclock
  - Usual X servers

Ways to run applications

- Console
  - Screen, host *TY, port
- SSH
- X apps over SSH
- X Desktop tricky but possible
  - Xnest

UML operating modes

- TT: Tracing thread
- TT/Jail
- SKAS
  - Requires host patch
  - Patch requirement may be going away

State of UML kernel

- Generally stable
  - > month uptimes common
- Patches available for 2.4 kernel
- In 2.6 kernel
  - Add-on patches needed at the moment

UML Architectures

- Stable on i386
- Functions on x86_64
- PowerPC port just finished
- Two cygwin ports partially finished
  - One works up until a fork needed. :-)

UML distributions

- Redhat, Fedora, Slackware, Debian, Gentoo, Toms Root/Boot, Mandrake roots available
- Included in SuSE/Novell, Debian

Slartibartfast/Zaphod

- 26 virtual machines
- Dual PIII, 4G, 360G
- zaphod.stearns.org:1500
- SMTP, HTTP, DNS, SQL, SSH, Honeypots, RTCW/Team Fortress, PHP, SA-Blacklist build...
- URLs at the end

Ford

- Dual Opteron, 6G, 500G
- Testbed for x86_64
- www.stearns.org/ford/ford-project.html

More information

- user-mode-linux.sf.net
  - Lots of documentation
  - Mailing lists, IRC
  - Patches, kernels, distributions
- www.stearns.org
  - Patches, kernels, distributions
  - This paper in /doc/

Articles

- www.stearns.org/slartibartfast/
- /uml-coop.current.html
- www.stearns.org/slartibartfast/
- /zaphod-users-guide.current.html
- www.whoopis.com/howtos/
- /uml-admin-howto.html

Preview of coming attractions

- Honeypots and security work
  - George Bakos, ISTS
  - 2/3/2005

Credits

- Jeff Dike, UML author
- UML contributors
- ISTS Honeypot team
- Zaphod participants
- Bill Stearns
  - wstearns@pobox.com
  - www.stearns.org

Questions?