

UML: User-Mode Linux

PLUG presentation
by
Frank Sorenson

June 9, 2004



User-Mode Linux

- Introduction and Architecture
- Booting and Running a UML
- UML in-depth
- Development
- Resources



UML: Introduction

- UML is a port of Linux to Linux
 - OS system call interface, not hardware interface
 - Full-blown Linux kernel with scheduler, VM, networking, fs, etc.
 - Binaries execute natively on host CPU
- 'official architecture'
 - Developed separately
 - Requires patches to run correctly

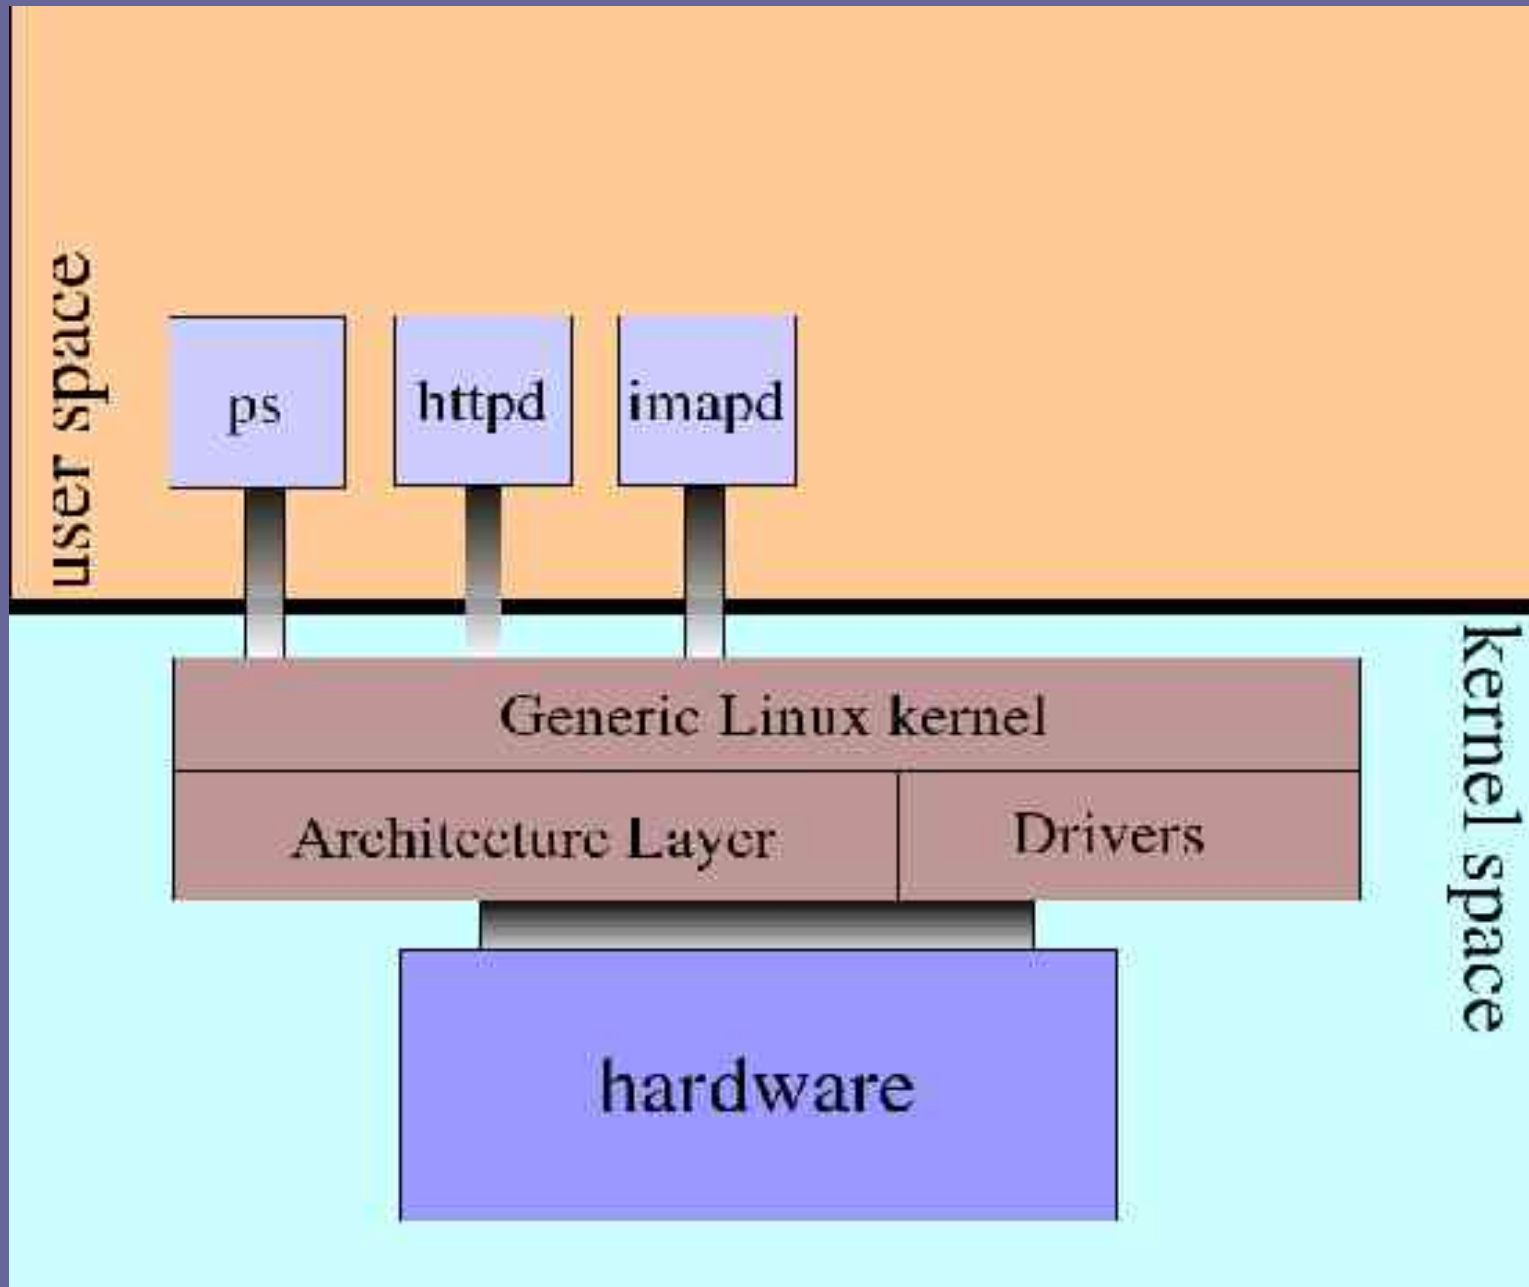


History

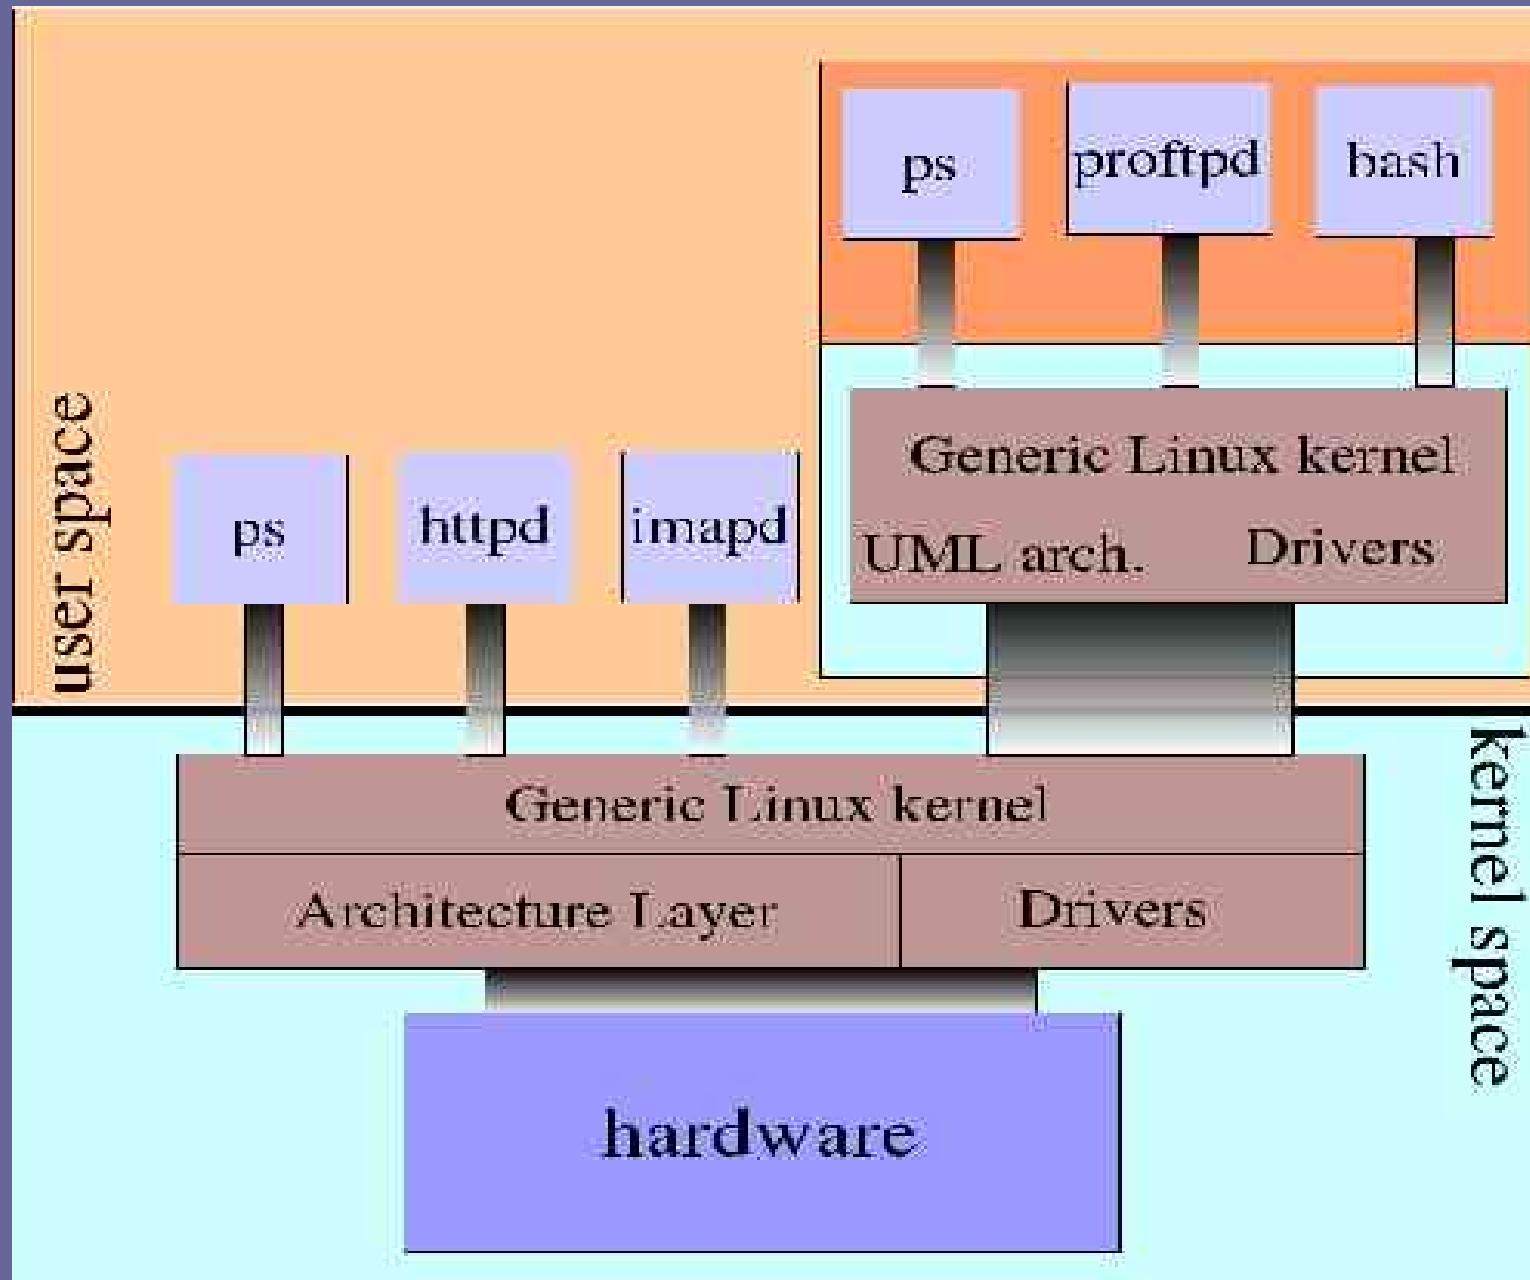
- UML development began in February 1999 by Jeff Dike
- Linux architecture/target
- Under active development
- Win32 and Sparc platforms started (status unknown)



Kernel architecture



UML Kernel Architecture



UML Modes: Tracing Thread

- Tracing Thread (TT)
 - Each UML process has corresponding host process
 - Host kernel switches context to TT, TT switches context to process, then back again
 - Four context switches, plus signal delivery and return



UML Modes: SKAS mode

- Separate Kernel Address Space (SKAS) mode
 - UML kernel has separate process and address space (processes can't address kernel memory)
 - Two context switches (much faster than TT mode)
 - Requires modification to ptrace on host kernel



ptrace

- Process debugger
 - Kernel capability allowing one process to intercept signals and syscalls of another process
 - used by gdb, strace, etc.
- TT mode
 - TT calls `PTRACE_SYSCALL` on new UML process and waits
 - Each process will switch to kernel mode for host calls, then back to userspace for tracing



- SKAS mode
 - Kernel process handles tracing and switches the single userspace process between UML processes
 - Creates empty address space using /proc/mm, not using new host process
 - Fewer processes appear on host



Running/Booting UML

- Host kernel requires patches for SKAS mode and other useful fixes
- UML kernel binary
- Filesystems (root, swap, others)
- Networking
- Build and execute UML startup command line



Compiling UML kernel binary

- Official kernel.org sources
 - UML patches (2.4.24-1 most stable)
 - other patches
- Configure and build
 - make xconfig ARCH=um
 - make dep ARCH=um
 - make linux ARCH=um
- Produces 'linux', a binary to run on your Linux machine



Root Filesystems

- Download existing filesystems
- Use root filesystem utility
 - rootstrap, UML Builder, etc.
- Create your own filesystem
 - dd can create empty file
 - mkfs on the empty file
 - loopback mount filesystem file
 - copy in files
- Single-service root filesystem
 - bind chroot



Booting a UML

- Create command line
 - executable: `./linux`
 - allocate memory: `mem=64M`
 - misc options: `devfs=mount fakehd
umid=myuml initrd=<initrd image>`
 - add disks: `ubd0=root_fs ubd1=swap_fs
ubd2=/dev/cdrom ubd3=/dev/fd0`
 - set up console: `con0=fd0,fd1
con=xterm`
 - networking:
`eth0=tuntap,tapdev,MAC,IP`
 - runlevel: `single`



- Filesystems
 - Mount something from the host with `hostfs/humfs`
 - COW filesystems
 - Using sparse files
- Security
 - Run UML as a user
 - `chroot (jail_uml)`
 - Single-service UMLs
 - No modules
 - Read-only filesystem (`'chmod -w'`)



UML in-depth

- Set 'nice' value when starting
- cpucap
- I/O token buckets
- Startup UML in background using screen or conserver
- `uml-utilities` (`uml-mconsole`, etc.)



UML console

```
md: ... autorun DONE.
VF3: Mounted root (ext2 filesystem) readonly.
Mounted devfs on /dev
INIT: version 2.84 booting
Activating swap.
Adding 262136k swap on /dev/udf/1. Priority:-1 extents:1
System time was Mon Jun 7 19:30:05 UTC 2004.
Setting the System Clock using the Hardware Clock as reference...
hwclock is unable to get I/O port access: the ioctl(3) call failed.
System Clock set. System local time is now Mon Jun 7 19:30:05 UTC 2004.
Checking all file systems...
fsck 1.27 (8-Mar-2002)
Setting kernel variables.
Mounting local filesystems...
nothing was mounted
Cleaning: /etc/network/ifstate.
Setting up IP spoofing protection: rp_filter.
Configuring network interfaces: Unrecognized kernel version
done.

Setting the System Clock using the Hardware Clock as reference...
hwclock is unable to get I/O port access: the ioctl(3) call failed.
System Clock set. Local time: Mon Jun 7 15:30:07 EDT 2004

Cleaning: /tmp /var/lock /var/run.
Initializing random number generator... done.
Starting OpenBSD Secure Shell server: sshd.
Recovering nvi editor sessions... done.
INIT: Entering runlevel: 2
Starting system log daemon: syslogd.
Starting kernel log daemon: klogd.
Starting internet superserver: inetd.
Starting OpenBSD Secure Shell server: sshd.
Starting periodic command scheduler: crond.

Debian GNU/Linux 3.0 (none) vs/0
(none) login: █
```



UML modes

TT mode

```
1970 ?      S      0:00 SCREEN -S um11 -t um11 console -h 1000 -T vt100 -d -m /uml
1971 pts/1  R      0:00 \_ /uml/test/linux-2.6.6-um1 (um11) [(tracing thread)]
1975 pts/1  T      0:00 | \_ \_ /uml/test/linux-2.6.6-um1 (um11) [(kernel thread)]
2061 ?      S      0:00 \_ /usr/lib/uml/port-helper -uml-socket /tmp/xterm-pipeBv
1981 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1983 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1985 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1986 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1988 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1990 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1992 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1994 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1996 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
1998 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
2000 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
2002 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
2004 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [(kernel thread)]
2006 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/sbin/init]
2062 pts/2  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [hwclock]
2354 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/usr/sbin/sshd]
2394 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/sbin/syslogd]
2406 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/sbin/clogd]
2434 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/usr/sbin/inetd]
2472 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/usr/sbin/cron]
2484 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/sbin/getty]
2488 pts/1  S      0:00 /uml/test/linux-2.6.6-um1 (um_1) [/sbin/getty]
```

SKAS mode

```
27E4 ?      S      0:00 SCREEN -S um_1 -t um11 console -h 1000 -T vt100 -d -m /uml/
27E5 pts/1  S      0:02 \_ /uml/test/linux-2.6.6-um1 (um11) [/sbin/getty]
27E7 pts/1  T      0:00 | \_ \_ [linux-2.6.6-um1]
2778 pts/1  Z      0:00 | \_ \_ [uml:-startvt <defunct>]
27E1 pts/2  S      0:00 | \_ \_ /uml/test/linux-2.6.6-um1 (um_1) [hwclock]
27E0 ?      S      0:00 \_ /usr/lib/uml/port-helper -uml-socket /tmp/xterm-pipeZCr
```



Top on the host, and the UML console

```
13:57:08 up 5:52, 3 users, load average: 0.13, 0.06, 0.01
56 processes: 51 sleeping, 3 running, 1 zombie, 1 stopped
CPU states: 31.6% user 22.0% system 0.0% nice 0.0% iowait 46.4% idle
Mem: 515644k av, 493708k used, 21936k free, 0k shrd, 126812k buff
      162616k active, 122676k inactive
Swap: 1052216k av, 1420k used, 1050796k free 101500k cached
```

PID	USER	PRI	NI	SIZE	RSS	SHARE	STAT	%CPU	%MEM	TIME	CPU	COMMAND
2930	sorenson	15	0	11536	11M	11356	R	43.1	2.2	0:02	0	linux-2.6.6-uml
2932	sorenson	11	0	472	472	472	T	9.7	0.0	0:00	0	linux-2.6.6-uml
7	root	10	0	0	0	0	SJ	0.1	0.0	0:00	0	kupdated
2943	sorenson	3	0	292	292	296	S	0.1	0.0	0:00	0	port-helper
1	root	3	0	476	476	428	S	0.0	0.0	0:04	0	init
2	root	3	0	0	0	0	SJ	0.0	0.0	0:00	0	keventd
3	root	3	0	0	0	0	SJ	0.0	0.0	0:00	0	kapmd
4	root	13	-9	0	0	0	SJN	0.0	0.0	0:00	0	ksoftirqd_CPU0
5	root	3	0	0	0	0	SJ	0.0	0.0	0:00	0	ksuend

```
sorenson@macaque:/uml/test
```

```
hwclock is unable to get I/O port access: the iopl(3) call failed.
System Clock set, _local time: Mon Jun 7 15:57:05 EDT 2004

Cleaning: /tmp /var/lock /var/run.
Initializing random number generator... done.
Starting OpenBSD Secure Shell server: sshd.
Recovering rvi editor sessions... done.
INI: Entering runlevel: 2
Starting system log daemon: syslogd/etc/rc2.d/S10syslogd: line 91: 122 Termin
eted start-stop-daemon --start --quiet --exec $binpath -- $$SLOGD

Starting kernel log daemon: klogd.
Starting internet superserver: inetd.
Starting OpenBSD Secure Shell server: sshd.
Starting periodic command scheduler: cron.

Debian GNU/Linux 3.0 (none) vc/0
(none) login: █
```



TuxRocks.com (my UML server)

```
13:59:19 up 20 days, 21:53, 1 user, load average: 0.01, 0.03, 0.00
40 processes: 39 sleeping, 1 running, 0 zombie, 0 stopped
CPU states:  0.0% user  0.5% system  0.0% nice  0.0% iowait 99.4% idle
Mem:   91348k av,  84536k used,   7412k free,    0k shrd,   6636k buff
       24680k active,    44788k inactive
Swap: 263160k av,   5796k used, 257364k free      22328k cached
```

PID	USER	PRI	NI	SIZE	RSS	SHARE	STAT	%CPU	%MEM	TIME	CPJ	COMMAND
1	root	8	0	108	80	80	S	0.0	0.0	0:00	0	init
2	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	keventd
3	root	19	19	0	0	0	SWN	0.0	0.0	0:00	0	ksuflirq_CPU0
4	root	9	0	0	0	0	SW	0.0	0.0	0:29	0	kswapd
5	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	bdflush
6	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	kupdated
7	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	jfsIO
8	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	jfsCommit
9	root	9	0	0	0	0	SW	0.0	0.0	0:00	0	jfsSync
10	root	-1	-20	0	0	0	SWK	0.0	0.0	0:00	0	mdrecoveryd
11	root	9	0	0	0	0	SW	0.0	0.0	0:28	0	kjournalc
556	root	9	0	0	0	0	SW	0.0	0.0	0:02	0	kjournalc
827	root	8	0	648	464	456	S	0.0	0.5	0:00	0	dhclient
877	root	9	0	328	284	284	S	0.0	0.3	0:29	0	syslogd
881	root	9	0	164	116	116	S	0.0	0.1	0:00	0	klogd
926	root	8	0	456	320	280	S	0.0	0.3	0:13	0	sshd
936	root	9	0	168	28	28	S	0.0	0.0	0:00	0	xinetd
930	root	9	0	1056	576	484	S	0.0	0.6	0:01	0	sendmail
938	smmsp	9	0	828	356	356	S	0.0	0.4	0:00	0	sendmail
1018	root	8	0	152	124	104	S	0.0	0.1	0:07	0	crond
1036	daemon	9	0	192	172	172	S	0.0	0.1	0:00	0	atd
1047	mnt	9	0	72	58	58	S	0.0	0.0	0:00	0	mingetty
30822	root	9	0	4528	2140	2112	S	0.0	2.3	1:18	0	httpd
16644	root	9	0	160	8	8	S	0.0	0.0	0:00	0	safe_mysqld
16669	mysql	9	0	1100	792	792	S	0.0	0.8	0:00	0	mysqld
16735	root	9	0	1448	1448	1092	S	0.0	1.5	0:00	0	bash
16771	root	14	0	1116	1116	816	R	0.0	1.2	0:01	0	top



UML Development

- Numerous contributors to UML
- `uml-devel` and `uml-users` mailing lists
- Improvements
 - `hostfs/humfs`
 - AIO (Asynchronous I/O)
 - I/O token buckets
 - `sysemu` context switch reduction
 - kernel debugging
 - **PERFORMANCE**



My projects

- UML performance
- yucon
 - Web-based management for both user and admin
 - Easy console access and UML control through ssh
- UML debugging (kernel panic on devel versions of UML)
- UML add-ons (uml status, exec, etc.)
- Networking improvements



[logout](#) | [passwords](#) | [start refresh \(10 sec.\)](#) | [refresh now](#)

User: sorenson

UML	UML Status	Info
uml1 Frank	Running (3) <small>(as of 2004-05-28 03:47:44)</small> Reboot Shutdown Connect to console: <code>ssh -l uml1 uml-host1</code>	Host: uml-host1 123MB (512MB available) 2.4.24-2uml up 3 days, 04:17:19.19 load: 0.10 0.04 0.01
		Host: uml-host1 512MB available

```
User: root (3) Host: uml-host2 _load: 0.09 0.04 0.01 2004 06 07 13:50:26
  UML_      UML Status      UML Load
-->uml5    Running (3)      0.07 0.04 0.00
  uml6      Off
  uml7      Off
  uml8      Off

Options: b - Boot UML, r - Reboot UML, s - Shutdown UML, p - ps,
         t - top, <enter> - Console (C-a, C-d to return), q - Quit
```



Resources

- Web pages
 - `user-mode-linux.sf.net`
 - `usermodelinux.org`
 - `http://www.user-mode-linux.org/~blaisorblade/`
 - `http://perso.wanadoo.fr/laurent.vivier/UML/`
 - `http://www.theshore.net/~caker/uml/tips_and_notes.html`
- `irc: irc.oftc.net #uml`
- mailing lists
 - `user-mode-linux-devel`
 - `user-mode-linux-user`



- Filesystems
 - Debian, Gentoo, RH9, FC1
- UML Kernels
 - 2.4.26, 2.6.6
- Presentation
- Source for compiling
- UML startup scripts
 - run 'install' from CD to start
 - edit 'config', and run 'setup_uml'
- Live - Run a UML right from CD



UML: User-Mode Linux

PLUG presentation
by
Frank Sorenson

June 9, 2004

