


Debian GNU/Hurd Status update

Samuel Thibault

2015 August 17th




It's all about freedom #0


“The freedom to run the program, for any purpose”


I.e.:

- Freedom from sysadmin!
 - WTH is fdisk/mke2fs/... hidden in /sbin?
 - I should be able to just work with my disk/network access
- Freedom to innovate
 - Experimental filesystem, personal work-flow, new kind of process combination,...
- Also provide freedom from misbehaving programs and drivers




Micro-kernel layering







Micro-kernel layering

- Server crash? Not a problem
 - “Computer bought the farm” is just an error, not something-of-the-death
- Easier to debug/tune
 - Just run gdb, gprof, ...
- Can dare crazy things
 - The Hurd console has dynamic font support
 - See chinese support in pseudo-graphical mode (actually pure VGA textmode!) of Debian installer.
- Kernel only handles Tasks, memory, IPC



Hurd possibilities





Hurd possibilities

```
€ settrans -c ~/ftp: /hurd/hostmux /hurd/ftpfs /  
(just once for good)
```

```
€ settrans -a ~/mnt /hurd/iso9660fs  
~/ftp://ftp.gnu.org/old-gnu/gnu-f2/hurd-F2-main.iso
```


```
€ ls ~/mnt
```

```
README-or-FAIL
```

```
...
```

- Only downloads what is needed.
- Can be permanently stored in ext2fs

```
€ settrans ~/.signature /hurd/run /usr/games/fortune
```



Example: interpose TCP/IP stack


```
€ settrans -ca $HOME/servers/socket/2  
  /hurd/pfinet -i $HOME/servers/tun0
```

```
€ openvpn ... $HOME/servers/tun0 &
```

```
€ remap /servers/socket/2  
  $HOME/servers/socket/2
```

```
€€€ wget www.gnu.org
```

- My own translator
- Can plug my own VPN software
- Only wget accesses it (well, the shell too :))




But also

```
€ remap /bin/sh $HOME/bin/sh
```


```
€ remap /bin $HOME/unionbin
```


...

- Check out Stow/Nix/Guix!




How does it work?







Rationale

- **Everything** is an (interposable) RPC
- Translators exposed in the FS
- The user gets to decide what/how to interpose
 - Without need for costly ptrace or fragile libc symbols interposition.
 - **Native** fakeroot/chroot
 - Fully virtualized and fine-grained interface
- Just need to use what's provided by the admin, e.g.
 - \$HOME/
 - TCP/IP stackand pile over it




Hurd possibilities (cont'ed)






Hurd possibilities (cont'ed)




i.e. ISO image inside a partitioned disk image
on ftp over a VPN



Porting packages to hurd-i386


- This is essentially a POSIX system
 - So portable programs should just work fine
- Top dumb issues
 - Not linux or BSD? `#include <windows.h>`
 - Have `mach.h`? Must be MacOS
 - `make -j $(grep ... /proc/cpuinfo) → make -j`
 - `#include <linux/limits.h>`
 - hardcoded `errno` values
 - Missing `-lpthread`, `-ldl`, `-IX11`, ...
- See Hurd porter page's developer corner
- And `PATH_MAX` not defined, as allowed by POSIX...



PATH_MAX is evil for you

Fragile semantic in POSIX


- Never meant to mean “reasonable size for a buffer containing a file name”
 - Linux' 4096 : a whole page, a whole TLB entry !
- Paths can actually be longer
 - “hidden”, ”protected” files?!
- Does that include the trailing \0? (not so clear in POSIX...)
 - Does your code actually properly handle that?
- A can of worms that nobody really tests...



Current State

Hardware support


- i686
- start of 64bit support
 - Kernel boots completely, now missing RPC 32/64bit translation
- DDE Linux 2.6.32 drivers layer for network boards
 - In userland netdde translator!
- IDE, Xorg, ...
- AHCI driver for SATA
- Xen PV domU
 - Required GNU Mach changes only
- Preliminary sound support through userland Rump
- No USB yet



Current State

Software support


- Quite stable
 - Have not reinstalled boxes for a decade.
 - Debian buildbots keep building packages, no hang after weeks!
- ~81% of Debian archive builds out of tree
 - XFCE, almost gnome, almost KDE
 - Firefox (aka iceweasel), gnumeric, ...
- Standard *native* Debian Installer



Recent work


Special thanks to Justus Winter!!

- Init system decoupled
 - Allows to use standard Debian sysvinit scripts!
 - Using dmd for Guix & such
- Distributed mtab translator
- Various optimizations
 - Protected payloads
 - Lockless implementations
 - Paging management
 - Message dispatch
- Valgrind start-of-port



Releases

- Nice 0.401 release on April 2011.
- Arch Hurd LiveCD release on August 2011.
- Released Debian-unofficial
 - wheezy/sid snapshot CDs on May 2013 \o/
 - jessie/sid snapshot CDs on May 2015 \o/




Removal from ftp-master

- Due since many years
- Not really useful to mirror all over the world anyway
- But beware of consequences
 - buildd.debian.org able to get fed from debian-ports?
 - Scheduled binNMUs for transitions
 - Losing this would be very tedious for ports maintainers
 - Exposure in e.g. <http://buildd.debian.org/yourpackage>
 - Toolchain version upgrades (gcc, perl)
 - Give \$WHOEVER_DOES_IT accounts just to be able to check those?
- In short, managing to get most support without extra load?


Conversely, those would help ports currently on debian-ports

- Real status for Second-Class Citizens
 - BoF?



Future work

- Xen PVH support, X86_64 support
- Language bindings for translators (ADA?)
- Read-ahead
- {hdd,sound,usb}dde?
- Rump drivers?
- GNU system: Guix/Hurd?
- Startup in scheme?
- Your own pet project?



Thanks!

- <http://hurd.gnu.org/>
- <http://www.debian.org/ports/hurd/>
- <http://people.debian.org/~mbanck/debian-hurd.pdf>
- The increasing irrelevance of IPC performance for microkernel-based Operating Systems

<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.37.9653&rep=rep1&type=pdf>