Capsicum working group

Robert N.M. Watson
University of Cambridge

FreeBSD Developer Summit
11 May 2012 - Ottawa, Canada

Capsicum

- Lightweight operating system capability and sandbox framework
  - Kernel features appear in FreeBSD 9.0
    - Application compartmentalisation
    - Supplements existing DAC/MAC
- Collaboration between Cambridge, Google
- New work funded (FreeBSD Foundation, Google)

Working group goals November 2011

- Capability ioctl whitelists
- Sandboxed DNS and other services
- Migrating to *privilege upgrade* from *privilege downgrade*
- Updating Chromium patches
- libcapsicum
- 9.1 agenda
- connectat(), bindat()
- pdwait()
- Brainstorm applications to compartmentalise

ioctl white-listing

- Motivation: pjd's work on hastd
- ioctl is message passing
  - Thousands of calls with diverse semantics
  - Can't interpret at system call boundary
- Currently we have CAP_IOCTL
  - Idea: add white list table
  - Only white-listed ioctls available without CAP_IOCTL
- Discussion: masks, groups, …
- Conclusion: keep it simple for now

Missing system calls

- bindat(), connectact()
- [f]chflagsat()
- pdtrace()
  - What about pdwait4()?
  - Can we lose reparenting?

Other system features

- Socket option white-listing
- fcntl white-listing
- cwd as a capability
- close-on-exec for process descriptors
- Kill on ECAPMODE
- Core dumping to descriptor/socket

Services to offer to sandboxes

- Capsicum Services Daemon (CSD) -- launchd?
- Services for (and in) sandboxes
  - sysctl -- think compartmentalised netstat/sockstat/top/...
  - DNS, DNSSEC
  - Sockets, files … ftsat(3)?
  - SSL client/server
  - Authentication services
  - Various nscd services (or not)
- Requires us to clean up Capsicum's fdlist but good starting point

FreeBSD things to sandbox

rpcbind, mountd

SSH client, server

libfetch

libarchive (bsdtar)

zlib (gzip)

ftpd, tftpd

bootpd

inetd

netcat

syslogd

wpa_supplicant

all setuid binaries

snmpd

trsold, routed, …

named

ntpd

lwresd

Kerberos, kadmind

rwhod

file, libmagic

finger, fingerd

openssl

grep, awk, head, tail, …

bc

man

od/hexdump

csup

pkgng

Ports we want sandboxed*

- KDE - IO slaves
- PDF viewer! (xpdf?)
- BIND
- Apache
- Sendmail
- Chromium (update)
- nginx
- Unbound
- Varnish
- Mozilla
- lighttpd
- ImageMagick
- Subversion
- git
- Dovecot
- OpenLDAP
- Cyrus
- Spamassassin
- ClamAV
- mplayer

* notice aggressive use of passive voice

The plan

- 9.0: features present but not in GENERIC
- 9.1: Enable in GENERIC!
  - libcapsicum(3), csd(8)
  - Selected security-critical applications
    - SSH, dhclient, tcpdump, BIND, inetd, …
  - Chromium port uses Capsicum