NetBSD 2016

BSDCan 2016 quickie

Who am I?

- Masanobu SAITOH
  - msaitoh@netbsd.org
  - (masanobu@iij.ad.jp)
  - msaitoh@jin-magic.com
- What did I? (past)
  - Became a NetBSD developer in 1997
  - NetBSD/sh3
  - Port NetBSD to dreamcast
  - Make NetBSD/arm bi-endian
  - Fix bugs
- What are you doing?
  - Developing NetBSD based routers since 1999.
  - Maintain wm(4) and some Ethernet drivers
  - Some pci(4) common stuff.
  - MP networking
  - Driver for devices that Intel chipset has

GSoC 2016; 7 projects

- NetBSD on MS Azure
- U-Boot improvements
- Improve pkgin
- Ext4 support
- NPF & blacklistd gui
- POSIX test compliance
- pkgsrc split debug symbols

New ports

- Some ARM SoCs
  - Allwinner A31
  - i.MX6 and i.MX7
  - NVIDIA Jetson TK1
- Not yet
  - arm64

New drivers

- NVMe driver
  - Kernel: ported from OpenBSD
  - Userland: ported from FreeBSD
  - It has MSI-X support but uses only one submission queue.
- qat(4): Intel Quick Access Technology driver
  - Written from scratch (Not based on Linux/FreeBSD’s driver)
  - Not merged yet though.

USB 3 improvement

- The code for USB3 is in netbsd-7, but it’s disabled by default
  - Because it’s not stable.
- A lot of bugs were fixed not only xHCI but also USB common part.
  - One of funded project.
- Now xHCI is enabled by default in -current.

PaX

- ASLR: Address Space Layout Randomization
- MPROTECT: Strict w^x: Once a segment has been writable, it cannot be remapped to executable, and vice-versa
- SEGVGUARD: Suspend execution for programs that are DoS'ed into frequent coredumping

PaX controls

- Sysctl
  - Affects individual binaries (that have ELF PaX notes):
    - security.pax.aslr.enabled
    - security.pax.mprotect.enabled
    - security.pax.segvguard.enabled
  - Affects default behavior globally (not overriding ELF PaX notes):
    - security.pax.aslr.global
    - security.pax.mprotect.global
    - security.pax.segvguard.global

PaX controls

- paxctl(8): Allow editing of notes that affect PaX behavior on individual binaries
    - ASLR
        - +a: Disable ASLR (overriding global default)
        - +A: Enable ASLR (overriding global default)
    - MPROTECT
        - +m: Disable MPROTECT (overriding global default)
        - +M: Enable MPROTECT (overriding global default)
    - SEGVGUARD
        - +g: Disable SEGVGUARD (overriding global default)
        - +G: Enable SEGVGUARD (overriding global default)

PaX/ASLR

- Default in current for: i386, amd64, sparc64, evbarm (all PIE binaries)
- Randomizes:

What	BITS32	ALIGN32	BITS64	ALIGN64
TEXT/DATA(PIE)	16	PGSHIFT	32	PGSHIFT
STACK	⅛ of max	varies	⅛ of max	varies
STACKGAP	PGSHIFT	4	PGSHIFT	8
MMAP	16	PGSHIFT	32	PGSHIFT
EXEC_OFFSET	12	PGSHIFT	12	PGSHIFT
RTLD	12	PGSHIFT	12	PGSHIFT

PaX ASLR

- MMAP randomization offset is computed once per binary
- Things that break
  - Emacs because of undumping
- Handled automatically in pkgsrc, no programs affected in base

PaX MPROTECT

- Default in current for: i386, amd64, sparc64, evbarm
- Things that break:
  - JIT (Java, nodejs, bpfjit), gdb
- Handled by base and pkgsrc automatically
- GDB single stepping
  - Needs to modify the program text
  - Enabled via sysctl:
    - security.pax.mprotect.ptrace=1
      - 0: disallow
      - 1: only for programs started ptraced
      - 2: allow ptrace attach to work

PaX SEGVGUARD

- VNODE based, file-system independent - uses fileassoc(9)
- Sysctl: Programs need to crash 5 times in 120 seconds and they get suspended for 600 seconds:
    - security.pax.segvguard.expiry_timeout=120
    - security.pax.segvguard.suspend_timeout=600
    - security.pax.segvguard.max_crashes=5

Infrastructure improvements

- cdn.netbsd.org & nycdn.netbsd.org (hosting by Fastly)
- WWU hosting - untapped dev environments
- “private” hosting after ISC shutdown
- NetBSD now owns (or re-owns): netbsd.org, netbsd.com, netbsd.net, and netbsd.foundation (thanks, gjb@freebsd), pkgsrc.org, and others. -- The fight against domain squatters.
- New core@ team member: martin@
- We have a github

MP Networking

- First goal: Layer 2 forwarding [done]
  - MSI-X, interrupt distribution, hardware multi-queue (Intel 1G NICs), MP-safe bridge
- Second goal: Layer 3 forwarding [ongoing]
  - ARP/NDP cache separation from the routing table
  - Softint-based packet input
  - MP-safe routing table
  - MP-safe other objects: ifnet, if_addr, etc.
- Further tasks
  - MP-safe bpf, gif, vlan, ipsec, opencrypto, etc.
- A lot of L2 and L3 related ATF tests
- Test tool "ipgen" (FreeBSD netmap based)
  - https://github.com/iij/ipgen
  - RFC2544 test
- See also: http://www.netbsd.org/gallery/presentations/

Some others

- blacklistd is making headway
- Dtrace by default

BSDCan 2017

- We will have some NetBSD presentations in the next BSDCan.