NetBSD 2019

AsiaBSDCon 2019 WIP session

Rev. 1.0

Sad news

- A NetBSD developer Eric Schnoebelen passed away a few days ago
    - schnoebe@n.o

About me

- maya@NetBSD.org
- Bugs, pkgsrc, drivers (graphics, wireless, ...)

NetBSD 9 and 8.1

- We are going to make netbsd-9 branch in the next month.
- Also 8.1 will be released.

Security

- Kernel ASLR: randomize the kernel image and the kernel memory areas
  - GENERIC: randomizes by default **all** the dynamic memory areas (direct map, PTE space, etc)
  - GENERIC_KASLR: adds randomization on the kernel image
  - Most advanced KASLR implementation to date
- Audited Network Stack
  - More security and more robustness in the network components, cleaner code, safety measures
- Kernel Heap Hardening
  - More difficult to exploit use-after-frees and double-frees on kernel pools
  - More difficult to exploit buffer overflows on kernel pools
  - WIP
- Sanitizer/Instrumentation Support (next slide): quality assurance, detected dozens of bugs and weaknesses, thanks to advanced kernel support

Sanitizer/Instrumentation Support

- Allow to detect several kinds of bugs: buffer overflow, undefined behavior, etc
- Kernel: NetBSD is one of the few OSes to have extensive kernel support
  - KASAN: detect kernel memory corruptions
  - KUBSAN: detect kernel undefined behavior
  - KCOV: ease fuzzing
  - KLEAK: detect kernel memory disclosures (developed in and for NetBSD)
- Userland:
  - ASAN
  - UBSAN + micro-UBSAN (in libc)
  - TSAN
  - MSAN
- SyzBot fuzzing: 24h/24 fuzzing of the NetBSD kernel

Syzkaller fuzzing the NetBSD kernel in Google Cloud Engine (GCE)

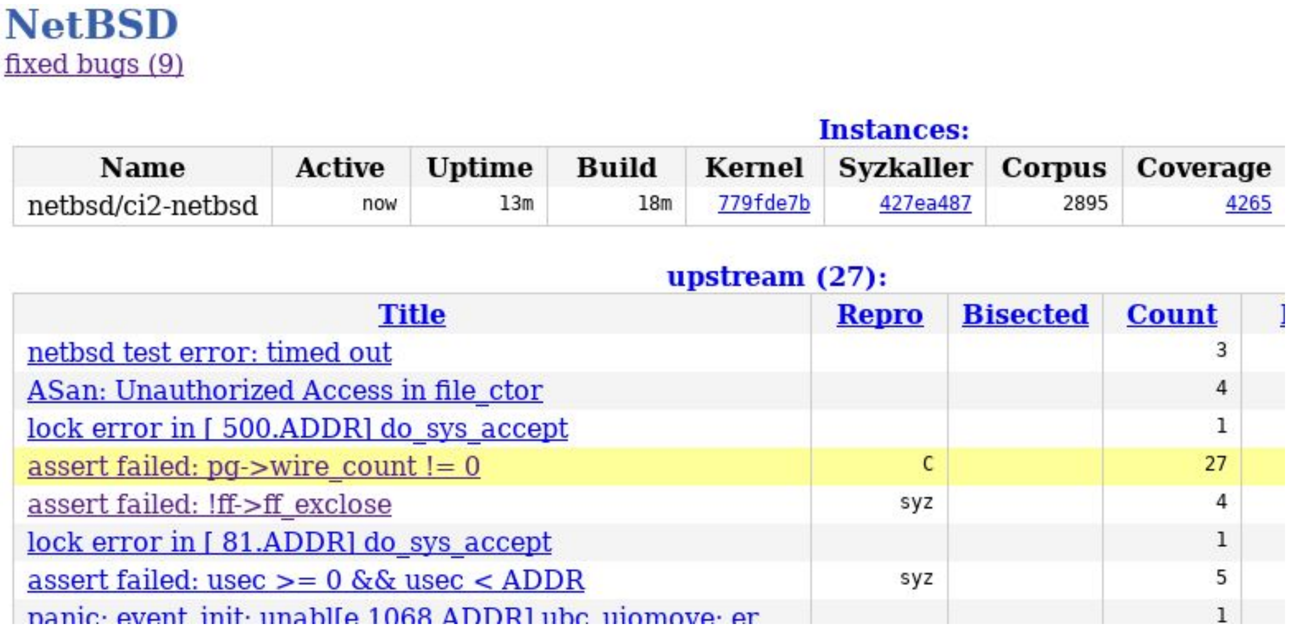

NetBSD

fixed bugs (9)

Instances:

Name	Active	Uptime	Build	Kernel	Syzkaller	Corpus	Coverage
netbsd/ci2-netbsd	now	13m	18m	779fde7b	427ea487	2895	4265

upstream (27):

Title	Repro	Bisected	Count	I
netbsd test error: timed out			3	
ASan: Unauthorized Access in file_ctor			4	
lock error in [500.ADDR] do_sys_accept			1	
assert failed: pg->wire_count != 0	C		27	
assert failed: !ff->ff_exclose	syz		4	
lock error in [81.ADDR] do_sys_accept			1	
assert failed: usec >= 0 && usec < ADDR	syz		5	
panic: event_init: unabl[e 1068.ADDR] ubc_uiomove: er			1	

NVMM

NetBSD x86_64

- Comprehensive virtualization solution for NetBSD
- CPUs supported: x86 AMD, x86 Intel, more can be added
- Fully MP-safe, fine-grained locking, up to 128 VMs with 256 VCPUs
- NVMM **does not implement an emulator**, but provides ***libnvmm***, a virtualization API that allows to add NVMM support in **already-existing emulators** (like Qemu, VirtualBox, etc)
- Advanced features: ***libnvmm*** can emulate certain guest operations on behalf of the emulator
- Security: the complex emulation machinery **runs in userland**, and not in the kernel! This means limited attack surface for the host

NVMM

NetBSD x86_64

- NVMM support in QEMU, fully functional

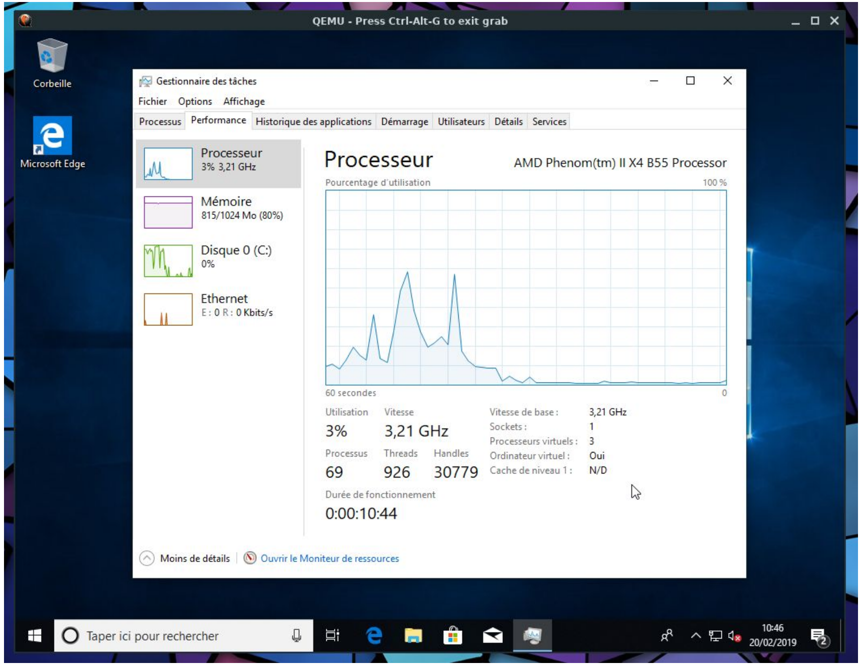

Processeur AMD Phenom(tm) II X4 B55 Processor

Utilisation 3%
Vitesse 3,21 GHz
Processus 69
Threads 926
Handles 30779
Durée de fonctionnement 0:00:10:44

Vitesse de base : 3,21 GHz
Sockets : 1
Processeurs virtuels : 3
Ordinateur virtuel : Oui
Cache de niveau 1 : N/D

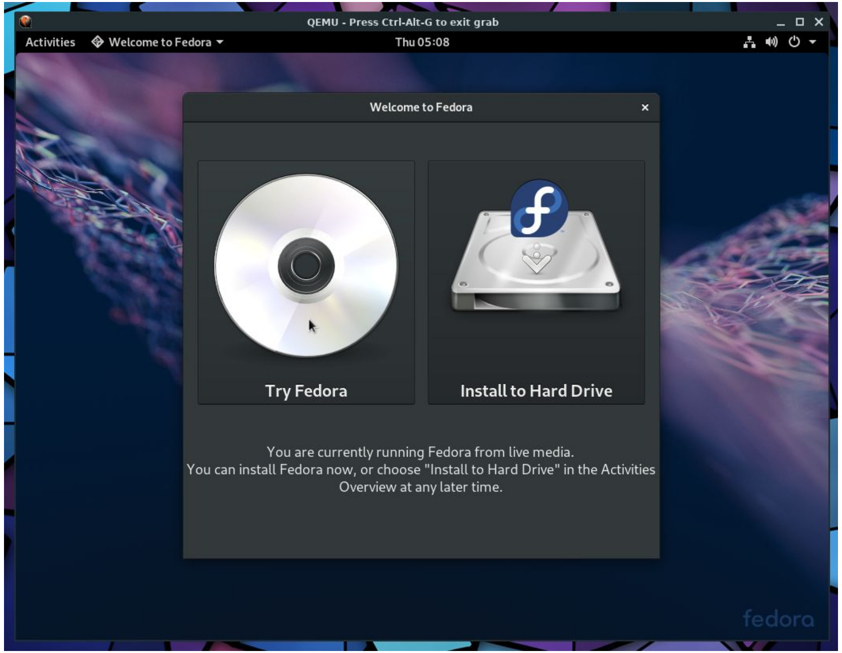

You are currently running Fedora from live media.
You can install Fedora now, or choose "Install to Hard Drive" in the Activities Overview at any later time.

ARM improvements - NetBSD -current

- AArch64 support
- Single kernel for multiple SoC families - GENERIC (armv7), GENERIC64 (aarch64)
- ACPI support - ARM Server Base System Architecture (SBSA), Server Base Boot Requirement (SBBR)
- UEFI bootloader
- Automatic root device detection on live images (armv7.img, arm64.img)
- Performance event monitoring support - tprof(4)
- Multiple cluster / big.LITTLE support
- GICv3 interrupt controller driver
- MSI and MSI-X support added to GIC (v2m) and GICv3 (ITS) drivers
- QEMU ARM Virtual Machine (“virt”) and virtio-mmio support
- Loadable kernel module support
- COMPAT_NETBSD32 support
- kernel address sanitizer (kASan)
- New SoCs
    - Allwinner A10, A13, A64, A83T, GR8, H5, H6, R8
    - Amlogic S805 (converted to FDT), S905
    - NVIDIA Tegra X1
    - Rockchip RK3328, RK3399
    - Samsung Exynos 5422
- Cloud providers
    - Amazon AWS EC2 a1
    - Scaleway

aarch64 TODO

- COMPAT_LINUX
- kernel preemption
- interrupt affinity (intrctl)
- Thumb mode support for COMPAT_NETBSD32
- TLB ASID randomization
- DTrace

The Attic Museum

- Removal of old unmaintained/buggy components:
  - vm86, ipkdb, NDIS, NATM, ISDN, compat_svr4, compat_ibcs2, and the list goes on
  - Reduces the maintenance burden, simplifies the kernel code, sometimes also reduces the attack surface (security)
  - Cleanup still ongoing...
- See the full list
  - https://wiki.netbsd.org/attic_museum/

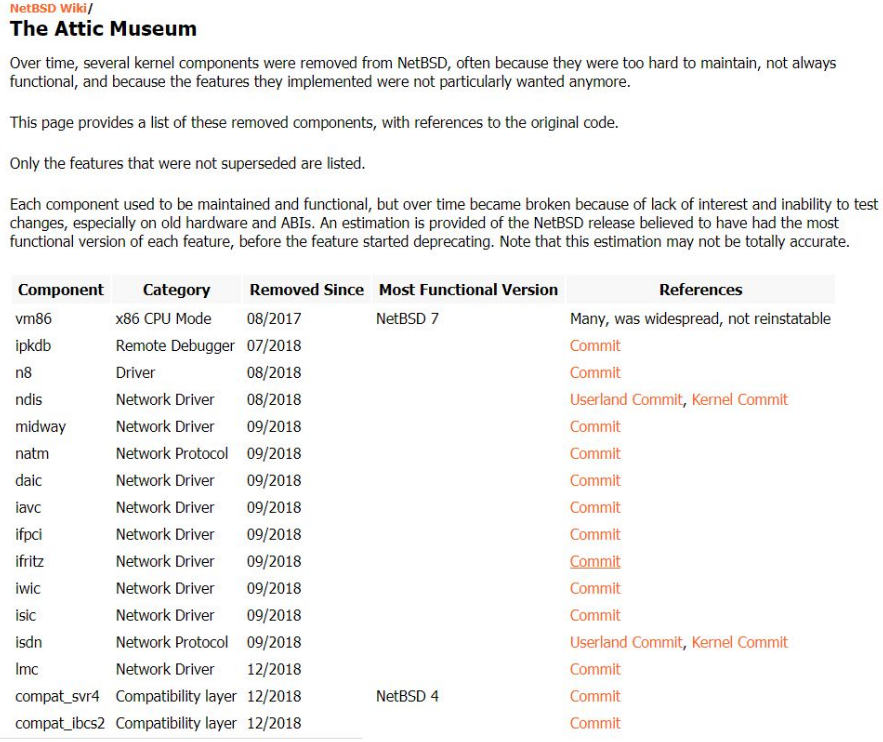

NetBSD Wiki/

The Attic Museum

Over time, several kernel components were removed from NetBSD, often because they were too hard to maintain, not always functional, and because the features they implemented were not particularly wanted anymore.

This page provides a list of these removed components, with references to the original code.

Only the features that were not superseded are listed.

Each component used to be maintained and functional, but over time became broken because of lack of interest and inability to test changes, especially on old hardware and ABIs. An estimation is provided of the NetBSD release believed to have had the most functional version of each feature, before the feature started deprecating. Note that this estimation may not be totally accurate.

Component	Category	Removed Since	Most Functional Version	References
vm86	x86 CPU Mode	08/2017	NetBSD 7	Many, was widespread, not reinstatable
ipkdb	Remote Debugger	07/2018		Commit
n8	Driver	08/2018		Commit
ndis	Network Driver	08/2018		Userland Commit, Kernel Commit
midway	Network Driver	09/2018		Commit
natm	Network Protocol	09/2018		Commit
daic	Network Driver	09/2018		Commit
iavc	Network Driver	09/2018		Commit
ifpci	Network Driver	09/2018		Commit
ifritz	Network Driver	09/2018		Commit
iwic	Network Driver	09/2018		Commit
isic	Network Driver	09/2018		Commit
isdn	Network Protocol	09/2018		Userland Commit, Kernel Commit
lmc	Network Driver	12/2018		Commit
compat_svr4	Compatibility layer	12/2018	NetBSD 4	Commit
compat_ibcs2	Compatibility layer	12/2018		Commit

Miscellaneous

- DRM/KMS update to Linux 4.4 (Intel up to Kaby Lake)
- x86: Kernel support for 16TB of physical memory, 32TB of virtual memory
- ZFS update
- DTrace update

GSoC

- NetBSD will participate Google Summer of Code in this year, too. Yay!
- Submit your application!

Resources

Papers and slides include this talk will be available in the following page

http://www.netbsd.org/gallery/presentations/#2019