

# NetBSD Status Report

Ignatios Souvatzis  
University of Bonn, CS Dept., Chair V  
<ignatios@cs.uni-bonn.de>

29th September 2004

## 1 Introduction

NetBSD is an Unix-like open source operating system based on the 4.4BSDlite code and other contributions. One goal is compatibility with the POSIX and SUS specified interface definitions. The complete kernel sources and big parts of libraries and programs are BSD-style licensed, thus easily adoptable for commercial third-party applications.

This paper outlines the major changes in NetBSD-2.0 and beyond.

## 2 Project Structure

When NetBSD was founded in 1993, it was just a few developers cooperating informally.

In the meantime, The NetBSD Foundation was founded to formally represent the project. The NetBSD Foundation is formal owner of (part of the) software, and administers material contributions and the project servers (ftp, cvs, www, ...). Members of the foundation are the project developers, some 300 world-wide, communicating mostly via the Internet.

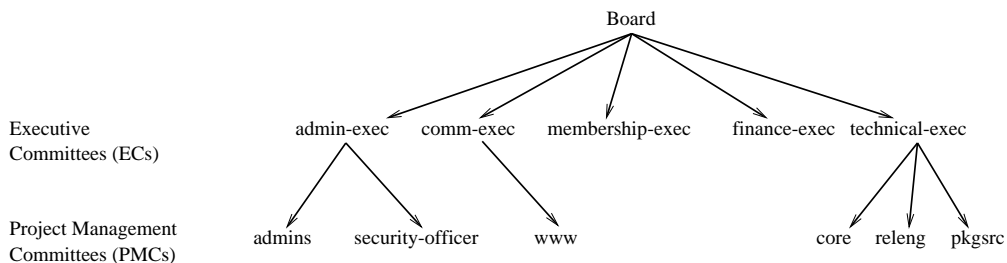


Figure 1: Structure of the NetBSD foundation offices

The NetBSD Foundation is a non-profit organization and has (in the USA) IRC 501(c)(3) status. It has registered the “NetBSD” and the “pkgsrc” trademarks in the USA.

## 3 The Operating System

### 3.1 Releases and Development

The NetBSD release numbering scheme was changed this year.




Figure 2: NetBSD development and release branches.

Formerly, major releases were numbered 1.x, with patch releases 1.x.y containing mostly bug fixes. The latest patch release of the old release branch, NetBSD 1.6.2, was released on March 1, 2004.

The next major release will be called 2.0. Patch releases will be called 2.1, 2.2 ..., and the major release after that will be called 3.0. The 2.0 release branch has reached its first release candidate at the time of writing this paper.

Unfortunately, a pkgsrc-like database for the base system (“syspkgs”) is not yet activated, so tracking binary patches to the base system isn’t yet possible. Thus, security advisories normally contain, per source tree branch affected, a source patch or a date the source repository was fixed, and instructions telling what part of the system to rebuild from sources.

### 3.2 Building The System

Building the system is done using a central script, `build.sh`, that controls building the toolchain — cross-compiler, cross-assembler, linker, auxiliary programs — into a separate space, then building the system into a separate directory structure, than creating the release archives. As of 2.0, this includes building of the X11 window system.

No elevated privileges are needed to do this, which allows for unattended operation. In fact, currently two project machines are bulding both release branches about twice a week (normally, the development branch and the latest releaase branch), for all target architectures, and a WWW page shows the last couple of lines of output for the last build and the last successful build for each architecture. New problems building the system are thus early noticed, even if specific to some architecture with slow (by today’s standards) machines.

In 2.0, most CPU architectures will use the newest set of the GNU toolchain: gcc 3.3.x, gdb 5.3, binutils 2.13.2.

### 3.3 Device Support

Here is a list of newly supported devices:

- The AMD64 CPU. That's the fourth 64bit architecture supported by NetBSD (after Digital Alpha, 64 bit MIPS and 64 bit Sparc).
- IDE drivers were split into the different chipsets.
- Serial ATA
- A generic 802.11 software infrastructure is available, allowing to implement 802.11 access points.
- Wireless LAN cards building on the above.
- some RAID controllers
- Gigabit Ethernet cards
- A TCPA hardware driver is in development, but won't make it into the release 2.0.
- An HPPA (hp700) port is being developed.

### 3.4 Storage

- There is new filesystem code for Apple-UFS (as of MacOS-X) and Kirk McKusicks UFS2. The former is especially interesting for users of PowerPC- Macs who want to use both operating systems. The latter is necessary for people using very large file systems ( $> 1TB$  when using 512 bytes block size).
- There is a SMB file system – it's possible to mount SMB shares.
- The volume manager "VINUM" was integrated.
- There is a new pseudo device "fss(4)" to take snapshots of a file system – useful for taking backups of an active system.

### 3.5 Kernel Internal Changes

- Multi CPU Support

Support for multiple CPU machines was added for AMD64, i386, MacPPC, 32bit and 64bit Sparc. (Multi CPU support was already there for VAX and Alpha in 1.6).

- Multithreading

Up until the 1.6 branch NetBSD didn't have kernel assisted multi threading. After that, an implementation of "Scheduler Activations" was brought into the kernel[15] and an N:M Posix thread library based on this was added.[14]

- Kqueue

Kqueue(2) is a generic interface for the kernel to send notifications to user processes originally written by Jonathan Lemon for FreeBSD [8]. In its NetBSD integration, device and filesystem events are reported. Support for reporting USB and network events is planned.




Figure 3: Scalability of a simple HTTP server (latency in  $\mu\text{s}$  vs. number of simultaneous requests, not counting the `connect(2)` time, using `poll(2)` for NetBSD-1.6.1 and `kqueue(2)` for NetBSD-current. This benchmark was done end of 2003 by Felix v. Leitner[7].

- Log-structured File System(LFS) with Unified Buffer Cache (UBC) support  
LFS uses the UBC[11] now, to allow for more file caching and synchronize `mmap(2)` with `read(2)/write(2)`.

### 3.6 Performance Tuning

These are some of the performance-related NetBSD changes:

- `pmc`: a new API to access performance counters
- Hardware checksumming of some network interface cards is used by their drivers.
- The not so new Virtual Memory Subsystem UVM of NetBSD makes it possible to effectively loan out memory pages to other address spaces[1]. This was used to implement a transmitter side “Zero Copy TCP” (and UDP).

To use this, an application has to `mmap(2)` (or directly create data in a local buffer), then `write(2)`. The buffer may not be changed anymore by the application until the data are sent, else a copy on write fault will be created. [12]

### 3.7 Security

In this section some new features are enumerated, that help the system administrator to build secure systems:




Figure 4: Touch after mmap benchmark, done end of 2003 by Felix v. Leitner[7].

- In the default installation no network services are activated (like on the 1.6 branch); they have to be activated by the system administrator in `/etc/inetd.conf` or `/etc/rc.conf`.
- On some architectures, that provide hardware support for this, stack and heap are not executable per default.
- If so configured, NetBSD only allows to execute programs that match their registered (in the kernel) checksum.
- Systrace[10] allows to control privileges of running programs at the system call + parameter level. E.g., binary programs can be restricted to only write in a specific directory, or not to do network calls, etc.
- There is a new pseudo device driver “cgd(4)” which is a encryption layer for disk partitions. This works with preconfigured keys (for data partitions) and with random keys (to secure swap).

Of course, known security fixes for the integrated OpenSSL, OpenSSH, BIND, Sendmail etc. are incorporated.

### 3.8 Miscellaneous

- Linux binary emulation was enhanced to better support Java and OpenOffice on i386 and PowerPC.
- On PowerPC machines a MacOS-X emulation is available.[4]




Figure 5: The growth of the `pkgsrc` collection between EuroBSDCon 2002 and EuroBSDCon 2004

- The `sysctl` interface is dynamically created: New kernel modules can enhance the `sysctl` tree without requiring a recompilation of `/bin/sysctl`.
- Long host names are supported in `utmpx`, `wtmpx`, `lastlogx`.
- All system binaries (including `/bin` and `/sbin`) are dynamically linked now. For emergency situations, a small `/rescue` directory with a few statically linked binaries is available.
- Lots of externally maintained but integrated software has been upgraded to their latest versions: `pppd`, `tcpdump`, `file`, `named`, `gcc`, `binutils` (`as`, `ld`, etc.), `postfix`, `sendmail`, `cvs`, `routed`, `texinfo`, `diff`, `grep`, `amd`, `openssh`, `openssl`, `less` ...

## 4 `pkgsrc`: The NetBSD 3rd Party Software Installation Infrastructure

Part of NetBSD is “`pkgsrc`”, a system for building and installing third party software on NetBSD (and elsewhere[3]), and for tracking the version of installed packages, the affected files, and dependencies on other packages.

`Pkgsrc` originated with FreeBSD, but was heavily modified for NetBSD and beyond. Currently it contains nearly 5000 files.

Most of `pkgsrc` is not traditionally cross-compilable. However, some effort has been done to do batch builds, and to do cross-compilation by using a mixture of using a cross-toolchain automatically and running parts of the `pkgsrc` tools in an emulator.[13]

Security advisories for `pkgsrc`-installed packages are not issued in a human readable form. Instead a (text form) database is maintained, that connects package version

numbers to security problem types and an URL describing the problem. The (pkgsrc) tool “audit-packages” can be used to automatically check the installed packages against the list, and to decide what packages to upgrade.

Pkgviews([2][6]) — a system that allows to install multiple versions of a package in parallel, thus allowing to upgrade packages and their dependencies on a live system without making it unusable for an extended time — are integrated in the source code, but not yet supported by all packages, and thus not activated.

There has been a pkgsrc development conference in Vienna in the summer of 2004, and the next one is planned for next year.

## 5 (Not only) New Documentation

- NetBSD contains the usual Unix-style Manual Pages. The usual eight chapters known from other Unix-like systems are enhanced with a ninth chapter about kernel functions and interfaces.
- Frederico Lupi’s “NetBSD Guide” is maintained by the NetBSD project now [9].
- A NetBSD Device Driver Guide is being written[5].
- Both documents, and many more design documents and HOWTOs, can be found on the NetBSD WWW servers in the *Documentation* subdirectory.
- The NetBSD WWW pages are available in multiple languages, selectable from the home page.

## Important URLs

[netbsd]	NetBSD – homepage: <a href="http://www.NetBSD.org/">http://www.NetBSD.org/</a>
[changes]	NetBSD – changes: <a href="http://www.NetBSD.org/Changes/">http://www.NetBSD.org/Changes/</a>
[autobuild]	binary snapshots of NetBSD-current and the release branch from the autobuild machines: <a href="http://releng.netbsd.org/ab/">http://releng.netbsd.org/ab/</a>

## References

- [1] Charles D. Cranor und Gurudatta M. Parulkar, *The UVM Virtual Memory System*, in: Proceedings of the USENIX Annual Technical Conference, Monterey, CA, USA 1999, [http://www.usenix.org/events/usenix99/full\\_papers/-cranor/cranor.pdf](http://www.usenix.org/events/usenix99/full_papers/-cranor/cranor.pdf)
- [2] Alistair G. Crooks, *PkgViews* — , 2nd European BSD Conference, Amsterdam 2002, <http://www.NetBSD.org/Documentation/software/pkgviews.pdf>
- [3] Alistair G. Crooks, *A Portable Packaging System*, Proceedings of the European BSD Conference 2004, Karlsruhe, Germany.
- [4] Emmanuel Dreyfus, *MacOS X binary compatibility on NetBSD*, Proceedings of the European BSD Conference 2004, Karlsruhe, Germany.
- [5] Jochen Kunz, *NetBSD Device Driver Writing Guide*, <http://www.unixag-kl.fh-kl.de/~jkunz/NetBSD/>

- [6] Johnny C. Lam, *User's Guide to PkgViews*, <http://mail-index.netbsd.org/tech-pkg/2004/01/06/0004.html>
- [7] Felix v. Leitner, *Unix Scalability Benchmarks*, work in progress, <http://bulk.fefe.de/scalability/>
- [8] Jonathan Lemon, *Kqueue: A generic and scalable event notification facility*, in: Proceedings of the FREENIX track, 2001 USENIX Technical Conference, Boston, MA, USA, [http://www.usenix.org/events/usenix01/freenix01/full\\_papers/lemon/lemon.pdf](http://www.usenix.org/events/usenix01/freenix01/full_papers/lemon/lemon.pdf)
- [9] Federico Lupi et al., *The NetBSD Operating System: A Guide (6 Languages)*, <http://www.netbsd.org/Documentation/#netbsd-guide>
- [10] Niels Provos, *Improving Host Security with System Call Policies*, 12th USENIX Security Symposium, Washington, DC, USA, August 2003.
- [11] Chuck Silvers, *UBC: An Efficient Unified I/O and Memory Caching Subsystem for NetBSD*, in: Proceedings of FREENIX Track: 2000 USENIX Annual Technical Conference, San Diego, CA, USA, [http://www.usenix.org/publications/library/proceedings/usenix2000/freenix/full\\_papers/silvers/silvers.pdf](http://www.usenix.org/publications/library/proceedings/usenix2000/freenix/full_papers/silvers/silvers.pdf)
- [12] Jason Thorpe, *Experimental zero-copy for TCP and UDP transmit-side*, <http://mail-index.netbsd.org/current-users/2002/05/02/0016.html>
- [13] Krister Walfridsson, *Cross-Compiling Packages*, Proceedings of the European BSD Conference 2004, Karlsruhe, Germany.
- [14] Alexandre Wenmmacher, *Symmetrie in NetBSD*, freeX 1/2003–3/2003
- [15] Nathan J. Williams, *An Implementation of Scheduler Activations on the NetBSD Operating System*, in: Proceedings of the FREENIX Track, 2002 Usenix Annual Technical Conference, Monterey, CA, USA, <http://www.usenix.org/events/usenix02/tech/freenix/williams.html>