BSD Operating Systems

Fred Morcos

FHLUG – April 2014

Contents

Introduction

- Berkeley Software Distribution is a UNIX operating system
  - Computer Systems Research Group (CSRG)
  - University of California, Berkeley
- Direct UNIX derivative
  - As opposed to a clone like GNU/Linux
  - From AT&T UNIX, distributed with source code
- Originally was software distributed for AT&T UNIX
- Later incorporated the kernel and base system utilities
- Some proprietary files had to be rewritten for a full distribution
  - AT&T Licensing

History

- 1977, Bill Joy assembles 1BSD
- 1979, 2BSD with vi and csh
- End of 1979, 3BSD with virtual memory
- 1980, 4BSD with job control in csh, delivermail (sendmail precedent), signals, curses
- 1981, 4.1BSD with performance tuning
  - 4.1a TCP/IP
  - 4.1b Berkeley Fast File System (FFS) by Marshall Kirk McKusick
  - 4.1c
- 1983, 4.2BSD
  - Bill Joy leaves to co-found Sun Microsystems
  - Mascot

History (2)

- 1986, 4.3BSD, tuning of many external contributions
- 1989, Net/1, networking stack under BSD License without AT&T code
- 1990, 4.3BSD-Reno, interim release between 4.3 and 4.4BSD
  - POSIX compliance
  - NFS
  - Support for HP9000 machines, from the HPBSD port
- 1991, Net/2, more and more under the BSD License
  - Rewritten UNIX utilities, nvi (new vi) because vi was based on UNIX ed
  - Only few files remained in the kernel that were AT&T licensed
  - Basis for ports to Intel 80386
    - Free 386BSD ⟶ basis for FreeBSD and NetBSD
    - Proprietary BSD/386 (BSD/OS) by BSDi

History (3)

- 1992, AT&T's Unix System Labs (USL) filed a lawsuit against BSDi
  - Stop distribution of Net/2 until copyright claims on source code could be determined
  - Slowed development of free descendants of BSD (FreeBSD, NetBSD) for 2 years
  - Systems based on the Linux kernel gained better support due to no legal ambiguity
    - Linus Torvalds wouldn't have created Linux if 386BSD had been available
  - Lawsuit was settled in 1994 in Berkeley's favor
- 1994, 4.4BSD
  - 4.4BSD-Lite without any AT&T source code
  - 4.4BSD-Encumbered for AT&T licensees
- 1995, 4.4BSD-Lite Release 2
  - BSD development at Berkeley ceased

License

- Very permissive
- Take the source, do whatever you want with it
- Give credit to the original author
- Proprietary/corporate-friendly
- Examples:
  - Starting from Microsoft Windows 2000
    - TCP/IP stack and command-line networking tools are taken from FreeBSD
    - Hotmail used to run on FreeBSD
  - Apple's Darwin (which Mac OS X is based on)
    - Derived from 4.4BSD-Lite2 and FreeBSD
  - Solaris also contains BSD code
  - Juniper Networks' JunOS is a FreeBSD redistribution
  - QNX uses the NetBSD network stack

Current BSD Flavors

- NetBSD
  - BSD Net/2
  - 386BSD
  - 4.4BSD and 4.4BSD-Lite2
- OpenBSD
  - NetBSD
  - 4.4BSD and 4.4BSD-Lite2
- FreeBSD
  - 386BSD
  - 4.4BSD and 4.4BSD-Lite2
- DragonFlyBSD
  - FreeBSD
  - Independent Fork
- PC-BSD
  - FreeBSD
  - Redistribution
- GhostBSD
  - FreeBSD
  - Redistribution
- PfSense
  - FreeBSD
  - Firewall

Base System vs 3rd Party Software

- BSDs are full distributions of a computer operating system
- Unlike the Linux kernel, GNU userland and distros
- The entire system is developed, tested, documented & released as one
- Very well integrated systems, very cohesive and monotonic
- Highly stable due to conservative and well thought development processes

- Ports/package systems offer tools to audit, build & install 3rd party SW
- Ports are instructions for downloading, patching, auditing, building, installing and uninstalling 3rd party software and their dependencies
- Packages are pre-built and ready for install

Documentation

- Documentation is very important in the BSD community
- Essential part of development
  - Lack of documentation is sometimes a release blocker
- Handbooks for installation and setting up systems
- Man pages are very extensive and detailed
  - Also provide many useful examples
- Always available offline
  - Not much requirement for online search, forums or wikis
- Documentation is also a very easy way to contribute :-)

Development Model

- FreeBSD has a specialized “release engineering” team
- Follows a “development”, “stable” and “release” flow
- Other BSDs are more or less similar
- -CURRENT
  - SVN “head”, unstable development branch
  - Sometimes may not even build
- -STABLE
  - Untested, but builds and passes some automated tests
  - Stable/9, stable/10, ...
- -RELEASE
  - Well tested, after many alphas and betas
  - Release/9.0, release/9.1, release/10.0, ...

NetBSD

- Initial release in 1993
- Focus on clean design, is very portable and easily cross compiled ⟶ build.sh
- Pkgsrc ports system
- Android userland is from NetBSD
- Extensive use of HW abstraction layers ⟶ drivers are split into machine-independent and machine-dependent counterparts
- Porting drivers across architectures requires very low effort
- Proof: NetBSD **does** run on a kitchen toaster!
- “Of course it runs NetBSD!”

OpenBSD

- Initial release in 1996, forked from NetBSD
- Theo de Raadt
- Focused on security
  - Simplicity, correctness, performance
  - Auditability (free and open source software)
  - NDAs are not acceptable (as opposed to FreeBSD)
- Portability and standardization
- Home of OpenSSH, OpenSMTPD, OpenNTPD, pf, cwm, tmux, spamd, libressl, ...
- X and cwm are developed and distributed as part of the base system ⟶ X does not run as root anymore!
- Not much in SMP and virt ⟶ security issues
- “Only two remote holes in the default install, in a heck of a long time!”

FreeBSD

- Initial release in 1993
- Large application support pool, fancy hardware
- UFS with soft-updates, journaling, snapshots
- ZFS, the self healing filesystem → killer feature!
- GEOM – pluggable, stackable geometry → RAID, caching, journaling, encryption (gbde, geli), networked storage, decompression, compression (?)
- Ports system, pkg system, poudriere
- Bhyve hypervisor, jails, Linux binaries emulation layer
- Clang, Dtrace, USB stack, Capsicum, CARP, HAST
- Pf, ipfw, IPFilter → firewalls that don't suck
- Handbooks for users, developers, documentation writers
- “man tuning”, “man firewall”, ...
- “The Power to Serve”

FreeBSD

New in FreeBSD 10

- Clang → No more GPL GCC in base system!
- 25-30% speed improvement in IO workloads
- Raspberry PI support
- Bhyve, Hyper-V, VirtIO, Xen
- ARM Superpages
- AES-NI support, RDRAND
- AMD KMS
- SMP pf
- SMP wifi
- ZFS improvements → TRIM, LZ4
- FUSE in base system

FreeBSD

Coming in FreeBSD 11 (So far)

- C blocks from Apple
- UDP-Lite in IPv4 and IPv6
- New NetMap
- ARM, PowerPC improvements
- Virtualization improvements
- Apple Macbook trackpad support
- NFS performance improvements
- Various upgrades ⟶ Clang/LLVM 3.4, groff, bmake, OpenSSH, OpenSSL, sendmail, libc++, ...
- UEFI boot
- newcons

FreeBSD

DragonFlyBSD

- Initial release in 2004, forked from FreeBSD
- Threading and SMP in FreeBSD wasn't good enough
- Both projects work together closely
- DPorts system was recently made to be highly parallel
- Work on high-performance computing, goal is a truly distributed computing system
  - Highly-parallel kernel with tokens for synchronization
- New lightweight kernel threads and message passing
- HAMMER high-performance FS with online history, instant crash recovery, FS undo cmd, stream FS to remote machines
- Virtual kernel ⟶ emulated kernel in isolation
- Resident applications ⟶ snapshot a virtual memory space of a process for later use (possibly on a different system)
- Many security features from OpenBSD

PC-BSD

- Initial release in 2006, FreeBSD redistribution
- Pre-configured and pre-installed applications
- Graphical user interface & graphical installer
- Additional PBI software package system
  - Mac-like single bundle software packages
- AppCafe ⟶ GUI for search and installation of PBIs
- Life Preserver ⟶ GUI for local and remote backups and restoration of files (rsync and ZFS frontend)
- Warden ⟶ GUI for FreeBSD jails for easy sandboxing of software
- Ports Jail ⟶ System to use Warden for the FreeBSD ports tree by isolation package builds from the rest of the system

FreeNAS

- Initial release in 2005
- NAS system based on FreeBSD and ZFS
- Web interface for configuration and management
- Supports Windows, OS X and Unix clients
- Supports virt with Xen and Vmware
- Full disk encryption
- Plugin architecture
- Variety of protocols
- Local and remote replication
- Windows Backup, Time Machine (Mac), rsync, Life Preserver (PC-BSD)
- Plugins for DLNA, OwnCloud, Plex, BTSync, CouchPotato, ...

ZFS

- Storage pooling, CoW, snapshots
- Data integrity protection
  - Checksums
  - Automatic “healing” - Mirroring/RAID/Copies
- No fsck ⟶ Online scrub
  - Data **and** meta-data
- Quotas/reservations
- Adaptive Replacement Cache
  - No Least Recently Used
  - Frequently Used + Recently Used
- ZFS Intent Log ⟶ Journal writes
- Dynamic Striping across vdevs
- Live compression
- Deduplication
- Encryption (Oracle Solaris)
- OpenZFS Initiative

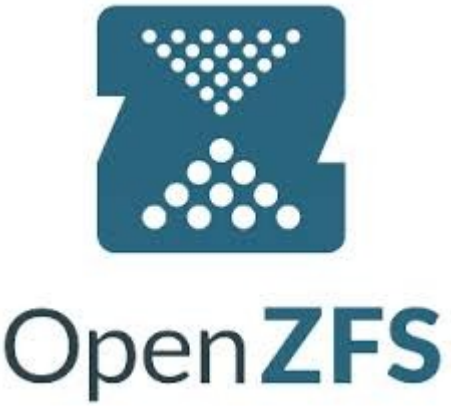

Jails

- Lightweight kernel-level virtualization ⟶ based on chroot
- Compartmentalization ⟶ No extra running kernels
- Full system
  - Users, groups, processes, files, ...
  - Networking
  - Ports/packages
- Limitations/Security
  - Cannot modify host kernel (modules, sysctls, ...)
  - No interaction with processes from other jails or host system
  - No mounting/unmounting or creating device nodes
- Cross-version jails are possible
- Linux jails are possible
- X11 jails are somewhat possible

Pros and Cons

- Stable development model and releases
  - Minimal breakage between releases (including drivers), if at all
- 20-30% increase in RAID0 + UFS read/write performance
- Documentation, monotonic and consistent system
- Leading technologies in storage (FreeBSD) and security (OpenBSD)
- Coherent networking tools, powerful package management
- **Very good** hardware support for what is supported

- Low interactivity, but scalable high throughput system
  - System is still responsive under high load, though

- Power management
- Kernel mostly following, especially in graphics
- Skype, Steam, ...

Links

freebsd.org

openbsd.org

netbsd.org

pcbsd.org

freenas.org

dragonflybsd.org

pfsense.org

freshports.org

bsdnow.tv