

BSD BoF

Asia BSDCon 2004

Robert Watson
FreeBSD Project
rwatson@FreeBSD.org



Introduction

BSD BoF

- BoF
 - Birds of a Feather
- Nominally a BSD BoF
 - But I'll talk mostly about FreeBSD
 - A brief talk on FreeBSD status
 - Q&A, problems, etc



Obligatory About FreeBSD Slide

- FreeBSD is a production-quality operating system derived from BSD UNIX
 - Widely deployed as a component, and in its own right
 - Workstation, server, and high-end embedded markets
 - Berkeley permits broad commercial re-use in open and closed source products
 - i386, ia64, amd64, sparc64, alpha



Releases over the Past 13 Months

- FreeBSD 4.x-STABLE continues incremental feature, performance, stability development
 - FreeBSD 4.8 (April, 2003)
 - FreeBSD 4.9 (October, 2003)
- FreeBSD 5.x-CURRENT continues higher risk development, approaches -STABLE
 - FreeBSD 5.0 (January, 2003)
 - FreeBSD 5.1 (June, 2003)
 - FreeBSD 5.2 (January, 2004)
 - FreeBSD 5.2.1 (February, 2004)



FreeBSD 4.8

- Port of OpenBSD crypto framework, drivers
- FAST_IPSEC; IP fragment DoS resistance
- Firewire, more USB serial drivers; more device polling support for ethernet drivers
- Hyper-threading hits 4.x
- Various security fixes
- Third party software upgrades, etc



FreeBSD 4.9

- Incremental development continues
 - PAE – Intel's >4GB physical memory extensions
 - Hardware accelerated crypto enhancements, USB ethernet device drivers, etc.
 - O_DIRECTIO
 - Various contributed software updates (OpenSSL, OpenSSL, sendmail, etc)
 - Various security fixes



FreeBSD 5.0

- First cut off of the 5.x Technology Branch
 - SMPng, sparc64, ia64, UFS extended attributes, ACLs, Mandatory Access Control, GEOM, GBDE, OpenPAM, ACPI, TIRPC, IPFW2, Firewire, UFS2, bgfsck, ufs snapshots, bluetooth, atapicam, devd, FAST_IPSEC, gcc 3.2, CardBus, devfs, ...
- Motivate stability, performance improvements
- Provide early access to consumers, product developers, etc.



FreeBSD 5.1

- Features
 - PAE, NSS, libkse, libthr, HTT, amd64, NFS locking, jail management, lazy context switch, SCHED_ULE, no major numbers, EHCI, various USB network drivers, reduced tcp state on connection close, volume labels on UFS and UFS2, / on vinum, DIRECTIO, dump snapshots, thread-safe rtld, ACPI updates, ...
- More stability, performance
 - Lock pushdowns
 - Much broader exposure, testing




FreeBSD 5.2

- Features
 - Full Tier-1 support for UP/SMP AMD64 systems
 - Dynamically linked root partition
 - IDE SATA, 802.11a/b/g devices, ACPI upgrades
 - NFSv4 client
 - IP forwarding plane now runs without Giant lock
 - Gnome 2.4, KDE 3.1
- And more stability, performance



FreeBSD Development Tree



FreeBSD 4.10 in 2004Q2

- FreeBSD 4.10 incremental release
 - Stability, performance enhancements
 - Bug fixes
 - Additional driver updates
 - Contributed software updates
 - ACPI power management
 - KDE 3.2



FreeBSD 5.3 in 2004Q3

- FreeBSD 5.3-RELEASE: Maturity!
 - KSE “M:N” threading the default
 - SMP enhancements
 - SCHED_ULE scheduler the default
 - Network stack running Giant-free
 - VM running entirely Giant-free
 - NDISulator
 - Use Windows network device drivers
 - OpenBSD's 'pf' packet filter integrated
 - Experimental support for Vinum in GEOM



KSE Status

- KSE: M:N threading implementation similar to Scheduler Activations
- libthr: 1:1 threading using KSE kernel primitives (not Linux semantics)
- Default in -CURRENT (shipped in 5.3)
- Runs KDE, Java, OpenOffice, ...
- SCHED_ULE offers thread affinity bits, especially in HTT context, per-cpu queues



SMPng Status

- Process locking complete
- File descriptor locking about done
- Pipe locking except VM optimizations
- VM is about 2/3
- GEOM storage
- IP stack in progress
- Kernel memory allocation
- Support for HTT
- SMP scheduler (ULE)
- MAC Framework
- Buffer cache $\frac{3}{4}$
- VFS cleanups
- Lazy TLB

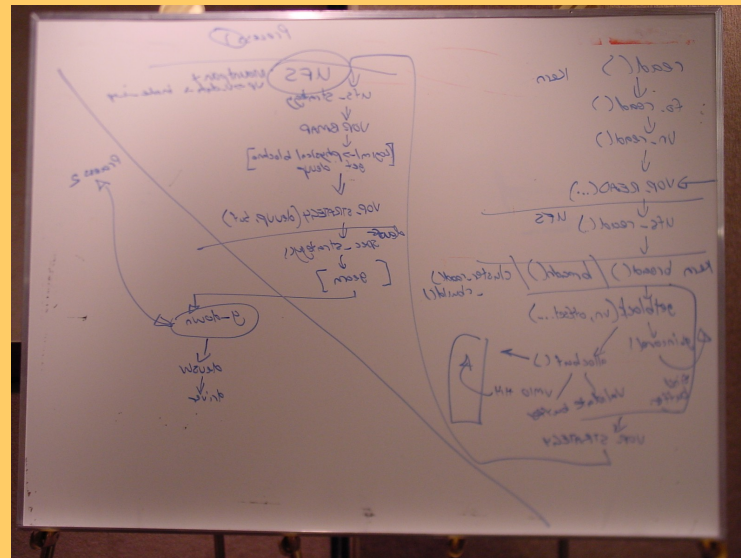


Developer Summit IV

- Two-day developer summit
 - June 10-11 2003 in San Antonio
 - June 10: unstructured working groups, hacking...
 - June 11: structured presentations, moderated discussion
- Very productive
 - Roadmap towards 5.2, 6.x clarified; 4.x life cycle
 - Removal of buffer cache, network stack/storage lockdown, prebinding, bridge burning, HTTP scheduler optimization, lazy context switches, ...



Dev Summit IV



Conclusion

- Things are very exciting
 - 5.x branch is maturing well
 - Aggressive feature set
 - SMPng work progressing; optimization, stability
 - M:N/1:1 threading support increasingly productionable
 - 4.x branch continues to offer high performance and stability
 - Will continue to see incremental improvements, optimizations; will remain most widely deployed



Q&A

- Any questions

