



**Dragon|FlyBSD**

# DragonFlyBSD?

---

- The DragonFlyBSD operating system is a fork of the highly respected FreeBSD operating system
- Its goals are to maintain the stability, quality and performance of the FreeBSD 4 branch and further enhance and improve it using new concepts
- Brainchild of Matt Dillon, who started work on DragonFlyBSD in June 2003

# Threading Model

---

- **M:1**                    4BSD (FreeBSD <5)
- **1:1**                     Linux, Solaris
- **M:N**                    KSE (FreeBSD5), SA (NetBSD)
- **M:Ncpus**             DragonFlyBSD

# Architectural advantages

---

- MP support, a departure from the outdated SMP model
- Should scale very well
- Will achieve Single-System Image clustering
- Makes better use of the L1 and L2 cache of all CPUs
- Maintainable code
- Flexible code

# DragonFlyBSD future

---

- Journaling filesystem
- Flexible package management
- SSI clustering
- VFS environments
- Next generation Jails similar to UML (User Mode Linux) in functionality but without the complex drivers needed to access the host hardware
- Ability to choose and switch between userland schedulers and multi-threaded libraries on the fly

# Journaling Filesystem

---

- Will make security auditing easier
- Will make mirroring and local/offsite backups trivial
- Filesystem synchronization in a Clustered environment via cache coherency protocols
- Can revert the FS or part of it to a chosen state
- First stage will be decoupled from the system
- Second stage will provide feedback to journal-aware filesystems to allow them to use the system generic journal instead of rolling their own
- Also intended to provide data transport services in the cache coherency model as a third stage

# Why not Softupdates & background fsck?

---

- Softdeps is very complex code
- Softdeps does not solve the problem of fast reboots after a crash
- Softdeps does not allow replaying changes for off-site synchronization
- Background fsck is still doing massive IO
- Background fsck doesn't scale well with today's increasing storage requirements
- Background fsck can not correct real disk corruption, manual user intervention is required

# Packaging system

---

- Currently using the FreeBSD ports tree with an override mechanism and NetBSD's PKGSRC which supports DragonFly natively
- Will handle most of the base system to reduce bloat
- Will take advantage of varsyms and other DragonFly features to present a source build environment
- Will use a binary-install methodology
- Will allow multiple versions of anything to be installed



# Most advanced networking stack

- Working SACK (RFC 2018)
- NewReno Modification to TCP's Fast Recovery Algorithm (RFC 2582)
- Larger congestion window sizes on connection startup (RFC 2414)
- Eifel detection for spurious retransmits (RFC 3522)
- Enhanced TCP loss recovery (RFC 3042)
- Early retransmit for TCP
- Parallel networking stack
- Cleaned up the BSD routing code. Check /usr/src/sys/net/route.c *rtalloc1()* on *FBSD* and *\_rtlookup()* on *DFlyBSD*

- BSD Installer (<http://www.bsdiinstaller.org>)
- Checkpointing
- Variant Symlinks
- Kernel Trace Routines (KTR)
- Various userland enhancements (fstat etc.)
- Fine-grained nanosleep()

- Abstraction for supporting multiple compilers via CCVER environment variable; ie. `% CCVER=gcc34 cc -v`
- Currently GCC 2.95.4 and GCC 3.4.3 are in the base system
- PF+ALTQ in base
- OpenNTPD in base
- Device driver support is mostly in-sync with other BSDs
- Two additional distributions available, GoBSD (community) and Firefly BSD (commercial)

<http://www.faqs.org/rfcs/rfc2018.html>

<http://www.faqs.org/rfcs/rfc2582.html>

<http://www.faqs.org/rfcs/rfc2414.html>

<http://www.faqs.org/rfcs/rfc3522.html>

<http://www.faqs.org/rfcs/rfc3042.html>

<http://www.icir.org/mallman/papers/draft-allman-tcp-early-rexmt-03.txt>

<http://ftp.cdut.edu.cn/pub/document/book/SunOS50MultithreadArchitecture.pdf>

<http://people.redhat.com/drepper/nptl-design.pdf>

<http://web.mit.edu/nathanw/www/usenix/freenix-sa/freenix-sa.html>

[http://www.asdf.dk/sjov/freebsd\\_kse.pdf](http://www.asdf.dk/sjov/freebsd_kse.pdf)

<http://www.sun.com/software/whitepapers/solaris9/multithread.pdf>

<http://leaf.dragonflybsd.org/mailarchive/kernel/2003-10/msg00204.html>

<http://leaf.dragonflybsd.org/mailarchive/kernel/2003-11/msg00159.html>

<http://docs.freebsd.org/cgi/mid.cgi?20040527120819.B8434>

<http://wiki.dragonflybsd.org>

<http://www.shiningsilence.com/dbsdlog/>

# Thank You

Eduardo Tongson  
[propolice@gmail.com](mailto:propolice@gmail.com)  
Feb 2005