

Porting ZFS file system to FreeBSD

Paweł Jakub Dawidek
<pjd@FreeBSD.org>



The beginning...

- ZFS released by SUN under CDDL license
- available in Solaris / OpenSolaris only
- ongoing Linux port for FUSE framework (userland); started as SoC project; very, very slowly moving forward



Features...

- ZFS has many interesting features, which make it one of the most wanted file systems
- some of them: dynamic striping, RAID-Z (RAID-5 without “write hole” (more like RAID-3 actually)), RAID-1, copy-on-write model, 128 bits (POSIX limits FS to 64 bits)...



...cheap snapshots, clones
(writable snapshots), snapshots
rollback, build in compression,
checksums, self-healing, storage
pool model (many file systems
out of one pool), soon
per-filesystem encryption and
much more.



Porting...

- very portable code (nice surprise)
- few ugly Solaris-specific details
- few ugly FreeBSD-specific details (VFS, buffer cache)
- ZPL is hell (ZFS POSIX layer);
yes, this is the thing which VFS talks to



Porting order...



Porting order...

- make it compile



Porting order...

- make it compile
- make it load as a kernel module



Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load



Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load
- make it work (properly)



Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load
- make it work (properly)
- make it work fast



Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load
- make it work (properly)
- make it work fast
- make it uncrashable by Kris (possible?)



Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load
- make it work (properly)
- make it work fast
- make it uncrashable by Kris (possible?)

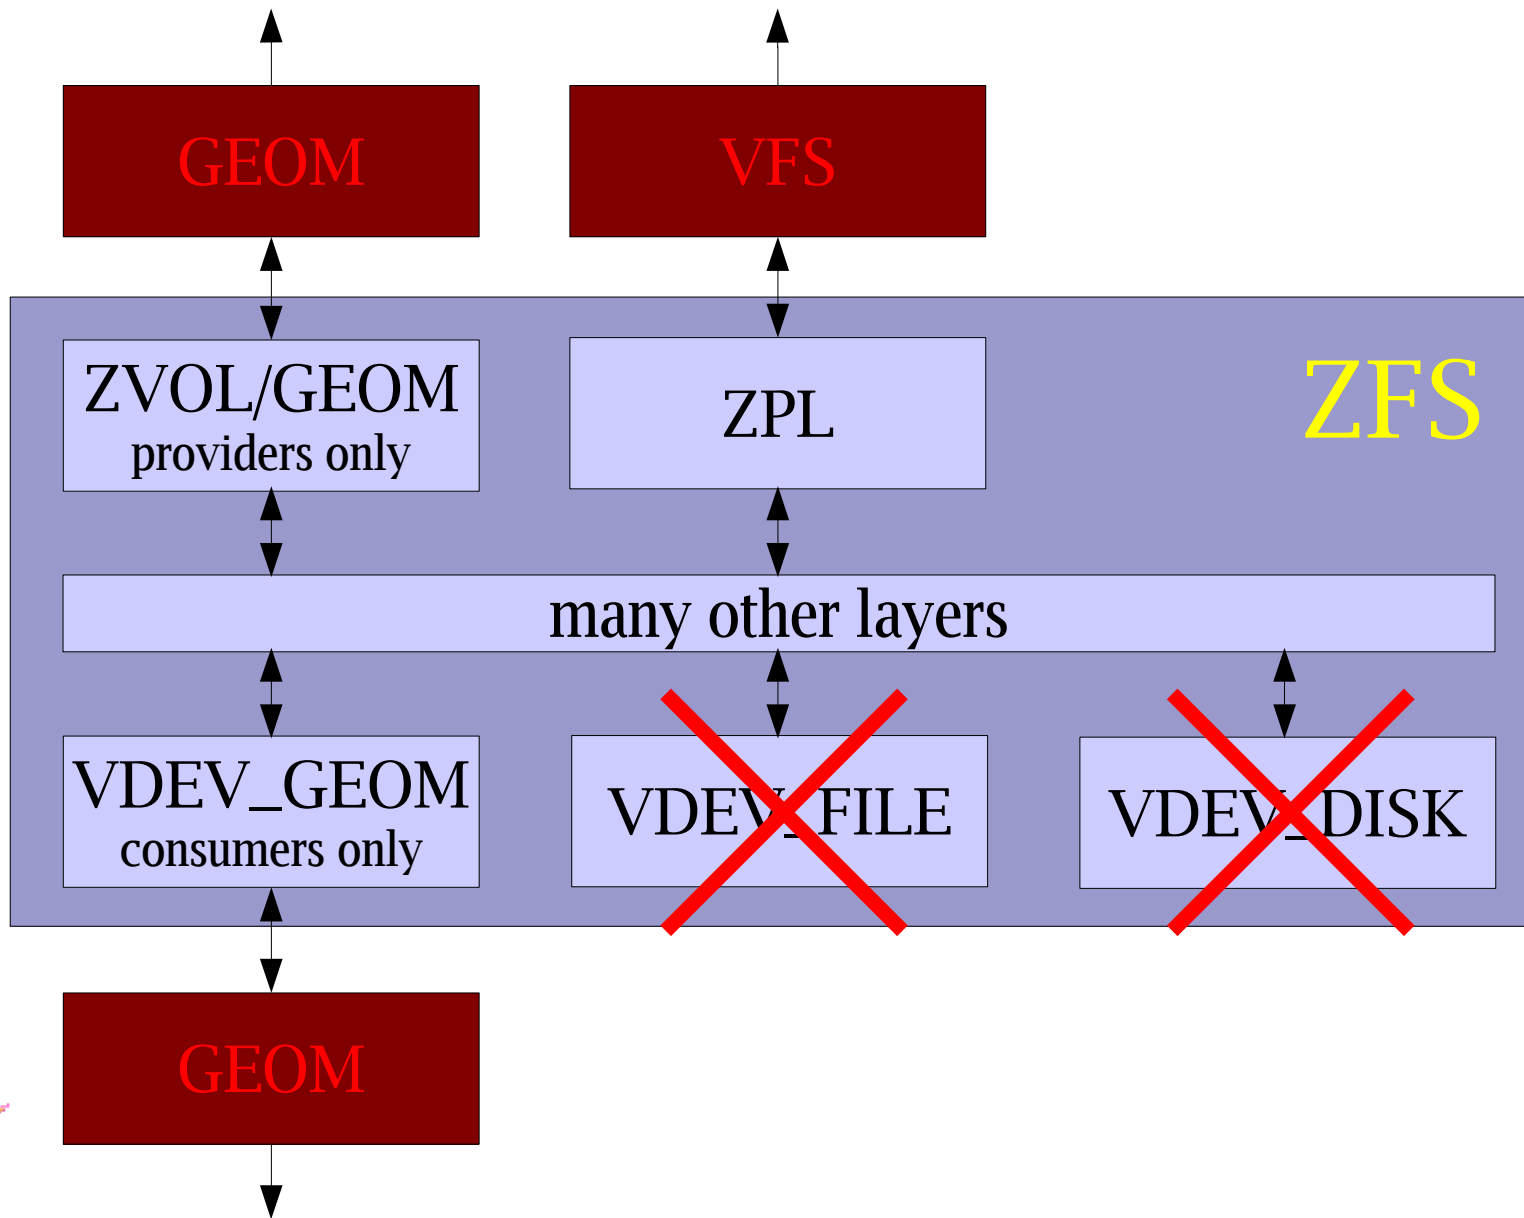


Porting order...

- make it compile
- make it load as a kernel module
- make it not panic on load
- make it work (properly)
- make it work fast
- make it uncrashable by Kris (possible?)



How does it look on FreeBSD...



Proofs...

