



ZFS para Administradores de Sistemas



Some ZFS key features

- Self healing: write and read checksum against bit rot.
 - Copy on write transactional.
 - Online scrub (“fsck”) and resilvering (“raid”). Online resizing.
 - Compression: LZJB, gzip1-9, ZLE .
 - Deduplication.
 - Unlimited snapshots and clones.
 - Hybrid storage pool: auto tiered.
 - Variable block size: 512 B to 128 KB, dynamic when compression enabled.
 - Serializes and aggregate writes in transaction groups, TXG.
 - “Unlimited” file size (16 EB).
-



-
- 128 bit scalability, without inode pre-allocation.
 - 2^{48} files per directory.
 - Thin provisioning datasets.
 - File and block protocols: CIFS, NFS, iSCSI. FC using ComStar stmfadm.
 - Includes block device and volume management: RAID-0, RAID-Z, Z2, Z3, mirror, ZVOL, VDEV, pool, hot spares.
 - Remote replication.
-




-
- Multiple filesystems and mountpoints into one pool.
 - Number of file copies beyond RAIDZ or mirror redundancies.
 - User and group quotas dataset filesystems.
 - Per dataset filesystems space reservations in one pool (“thick provisioning”).
 - RAID-Z* does not have RAID write hole.
 - JBOD using SAS/SATA expanders instead of detrimental RAID controllers.
 - ARC with ZIL VDEV reduce fragmentation.
-



-
- General use:
 - Debian GNU/kFreeBSD 7.x
 - FreeBSD 9.x
 - OpenIndiana
 - PCBSD
 - module ZFSonLinux
 - ZFS-fuse
 - Enterprise Data Storage Servers:
 - NexentaStor
 - OpenIndiana + Napp-It
 - Oracle Solaris Data Storage Servers
 - Small Business and Home Storage Servers:
 - TrueNAS
 - FreeNAS
 - Cloud virtualization with ZFS
 - SmartOS
 - FreeBSD BeHyve
 - Base Free kernel
 - Illumos
 - FreeBSD
-

ZFS storage pool

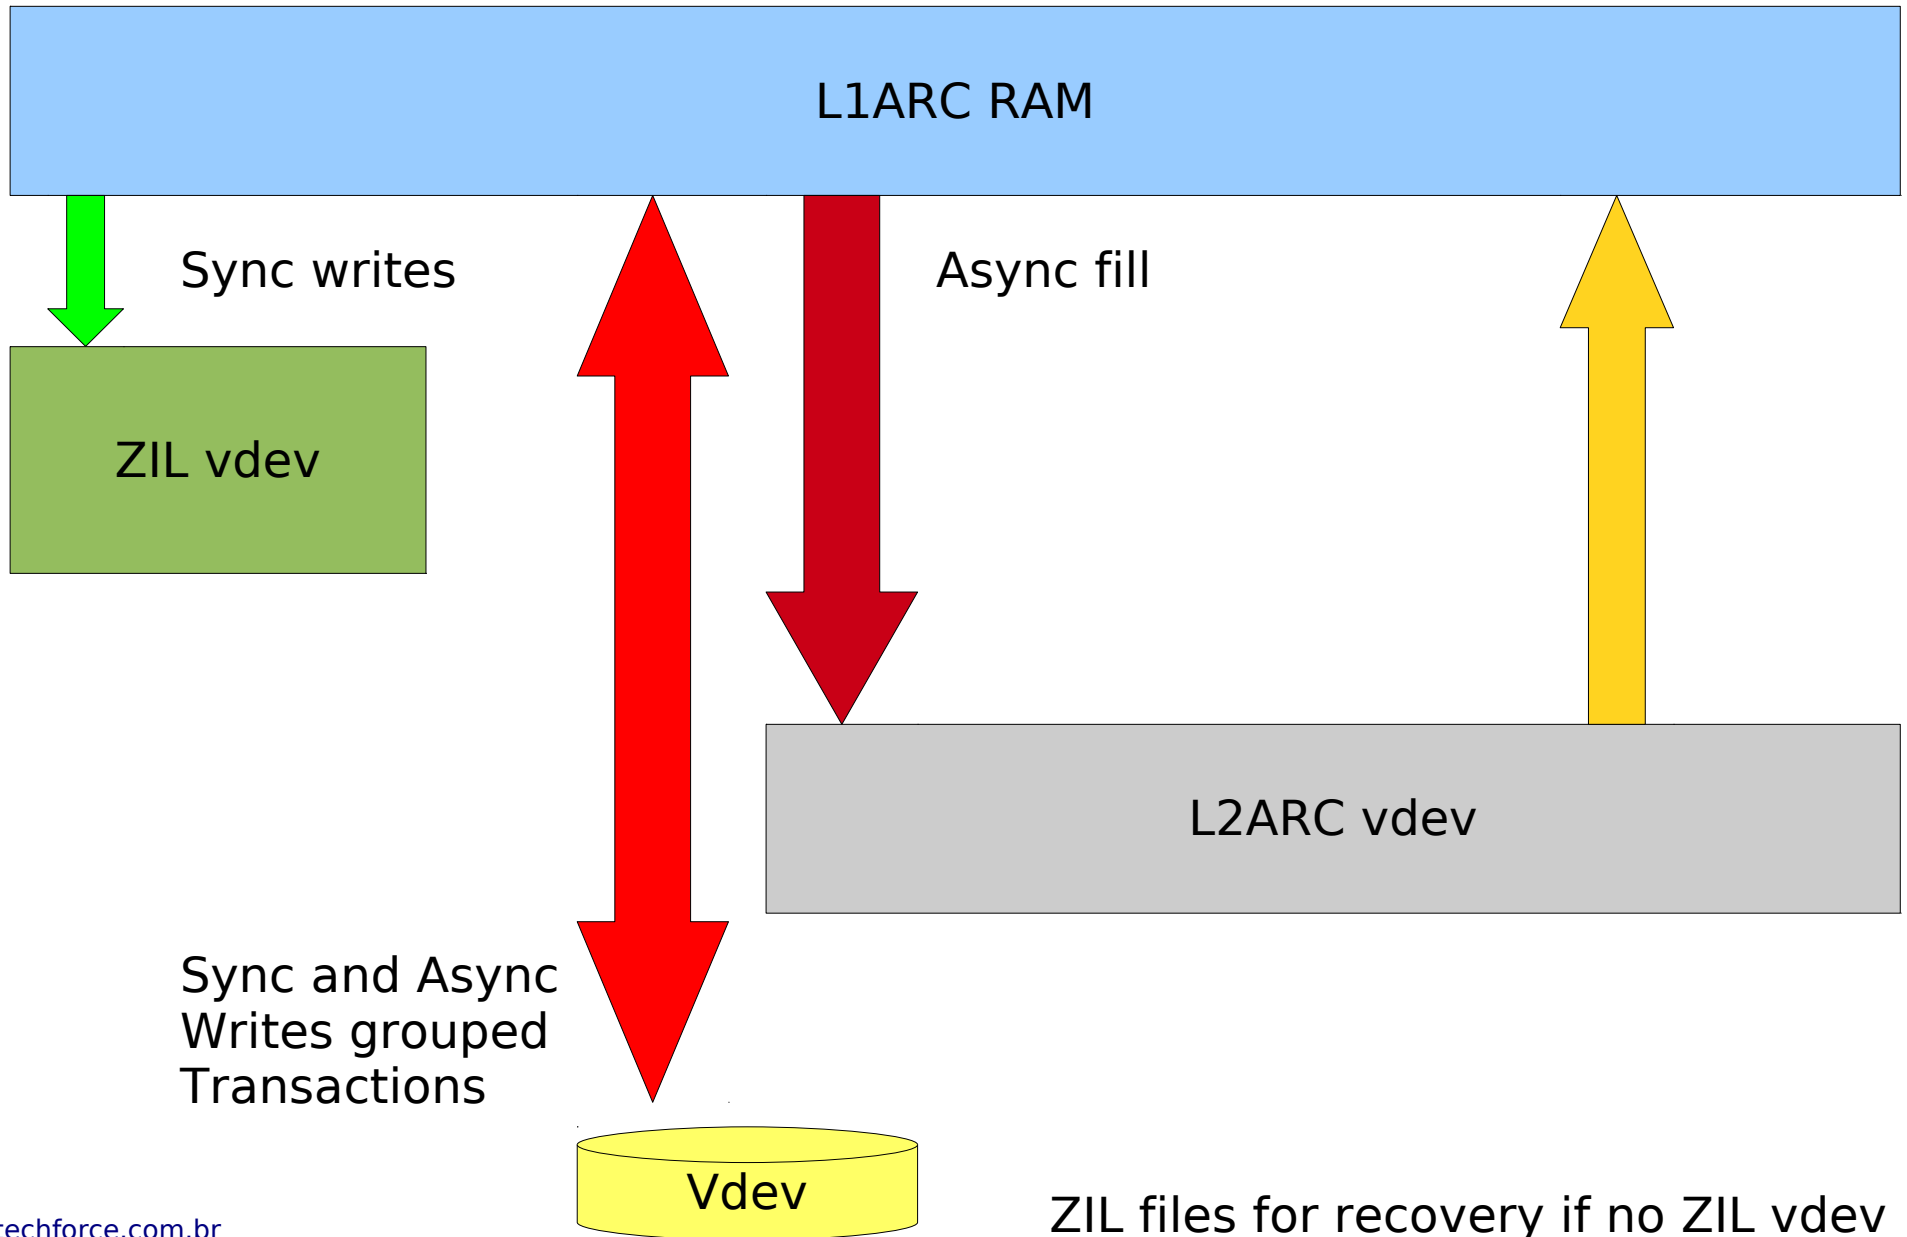


Sync and Async
Writes grouped
Transactions

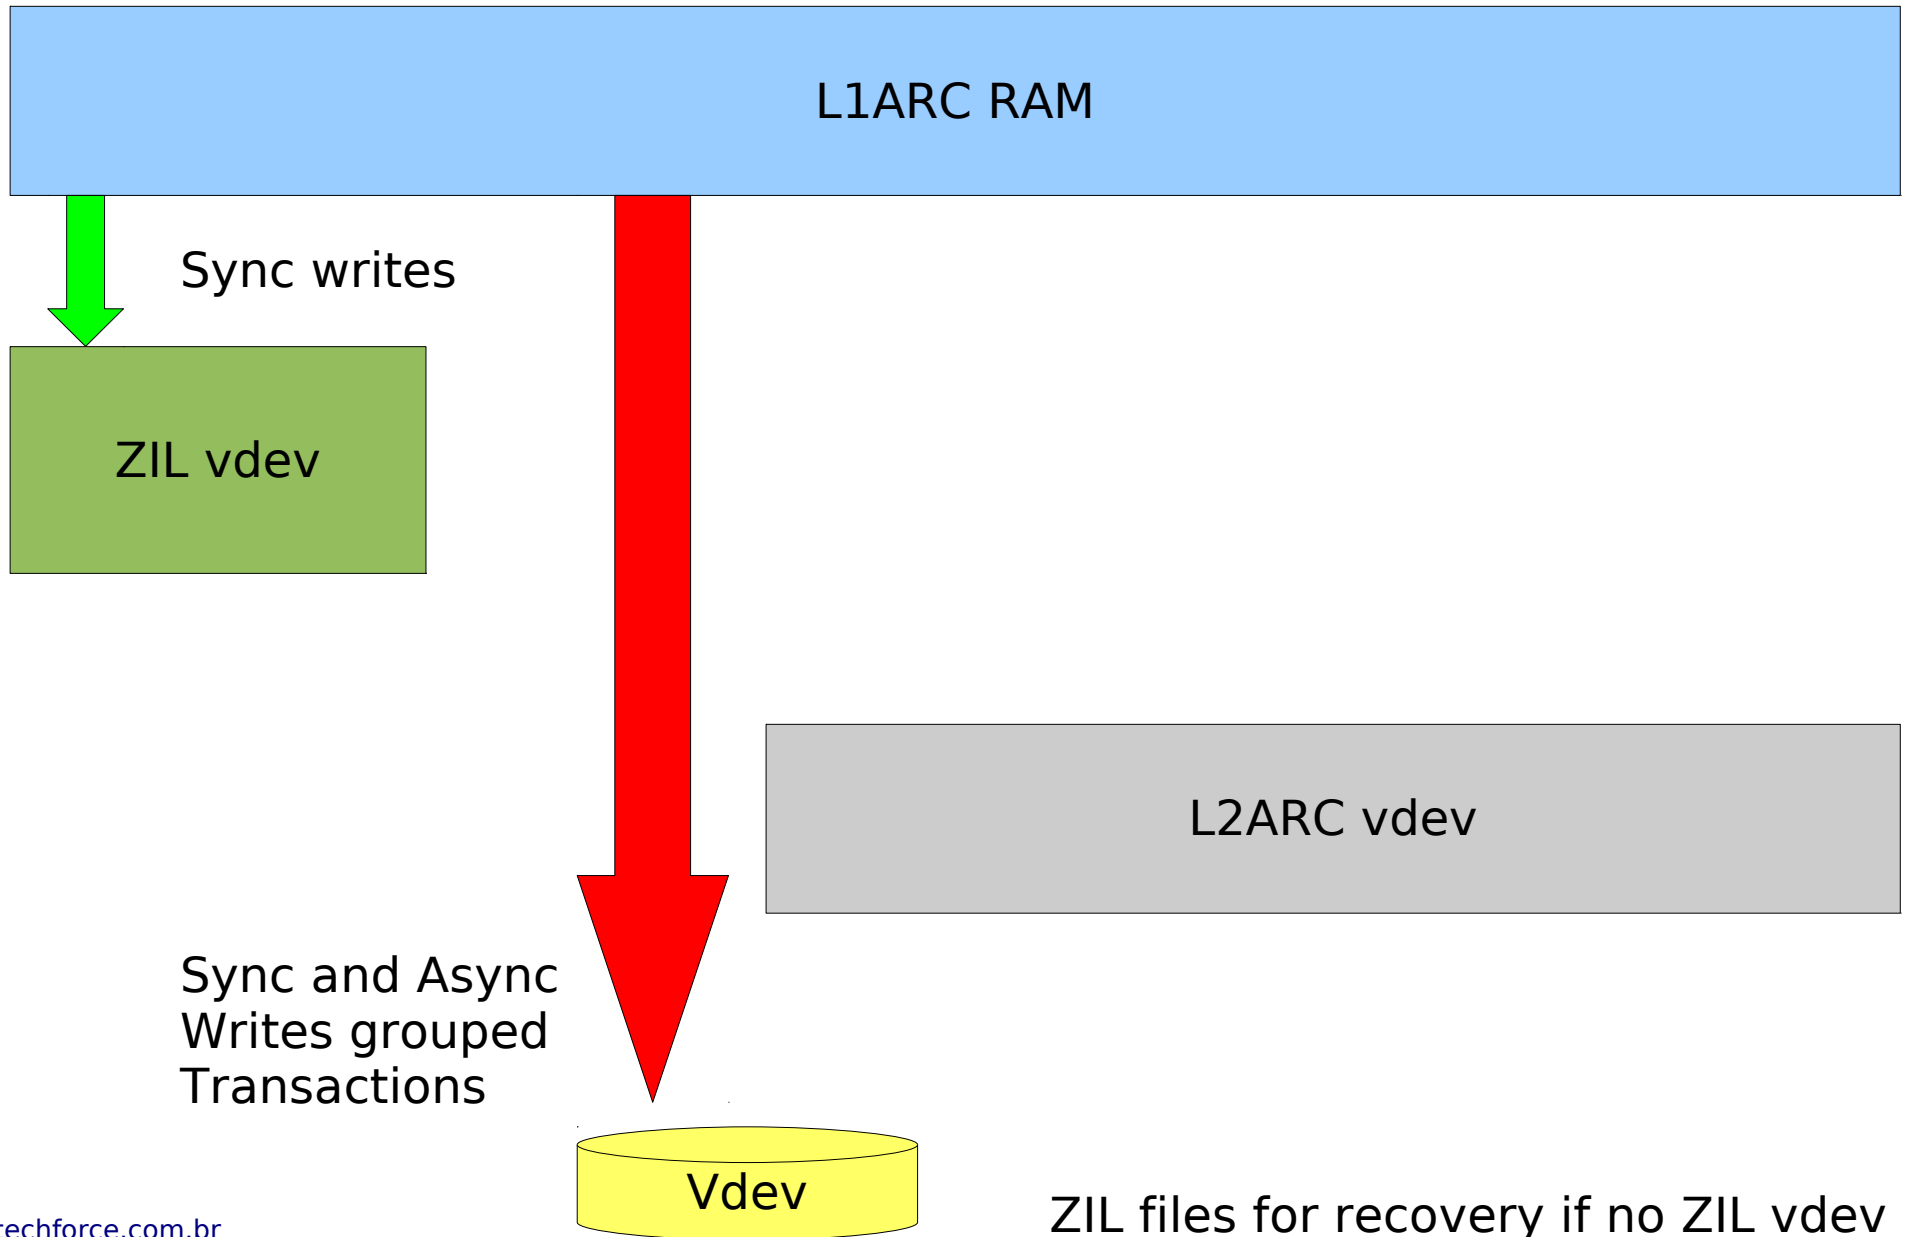


ZIL files for recovery if no ZIL vdev

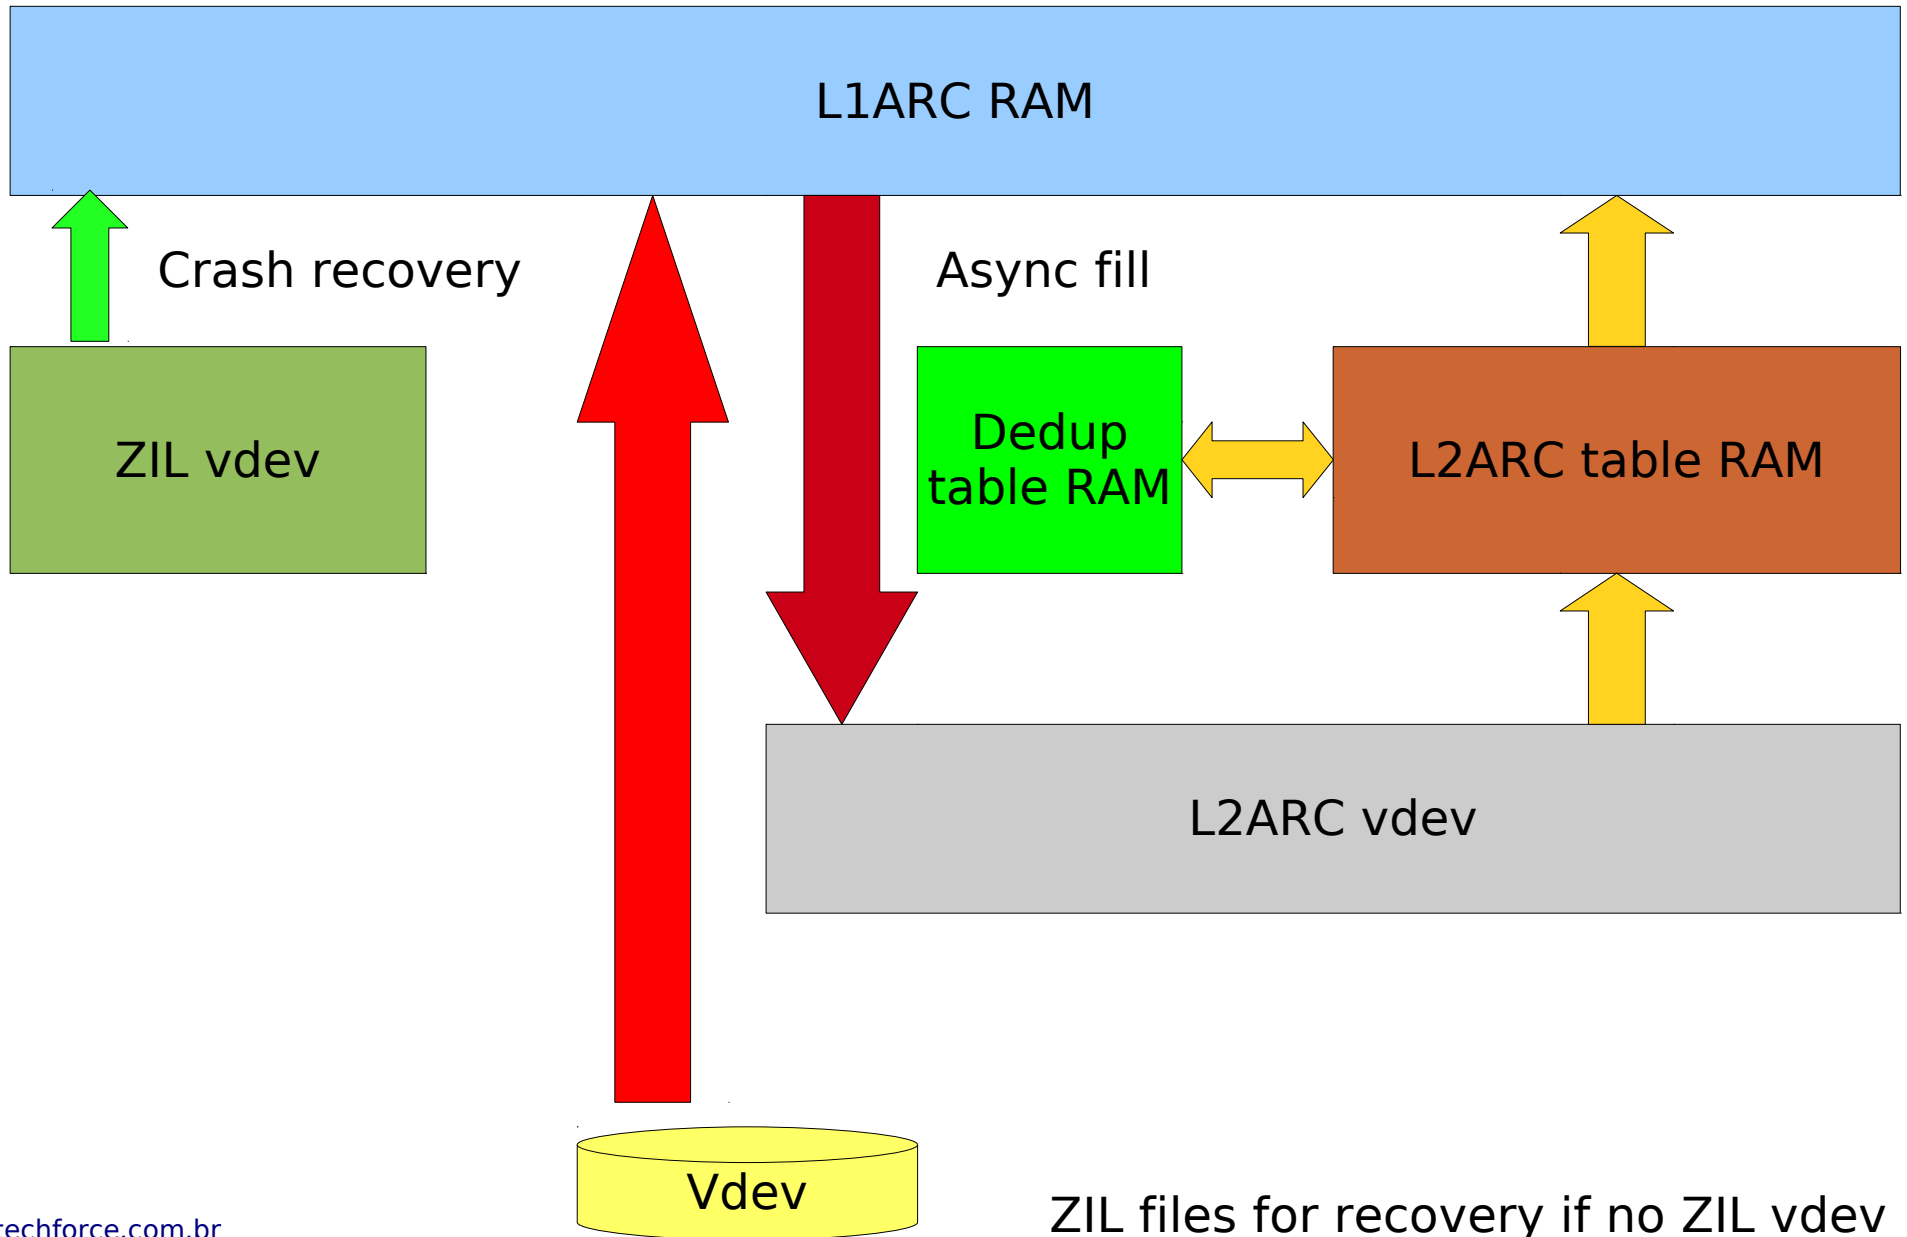
ZFS hybrid storage pool

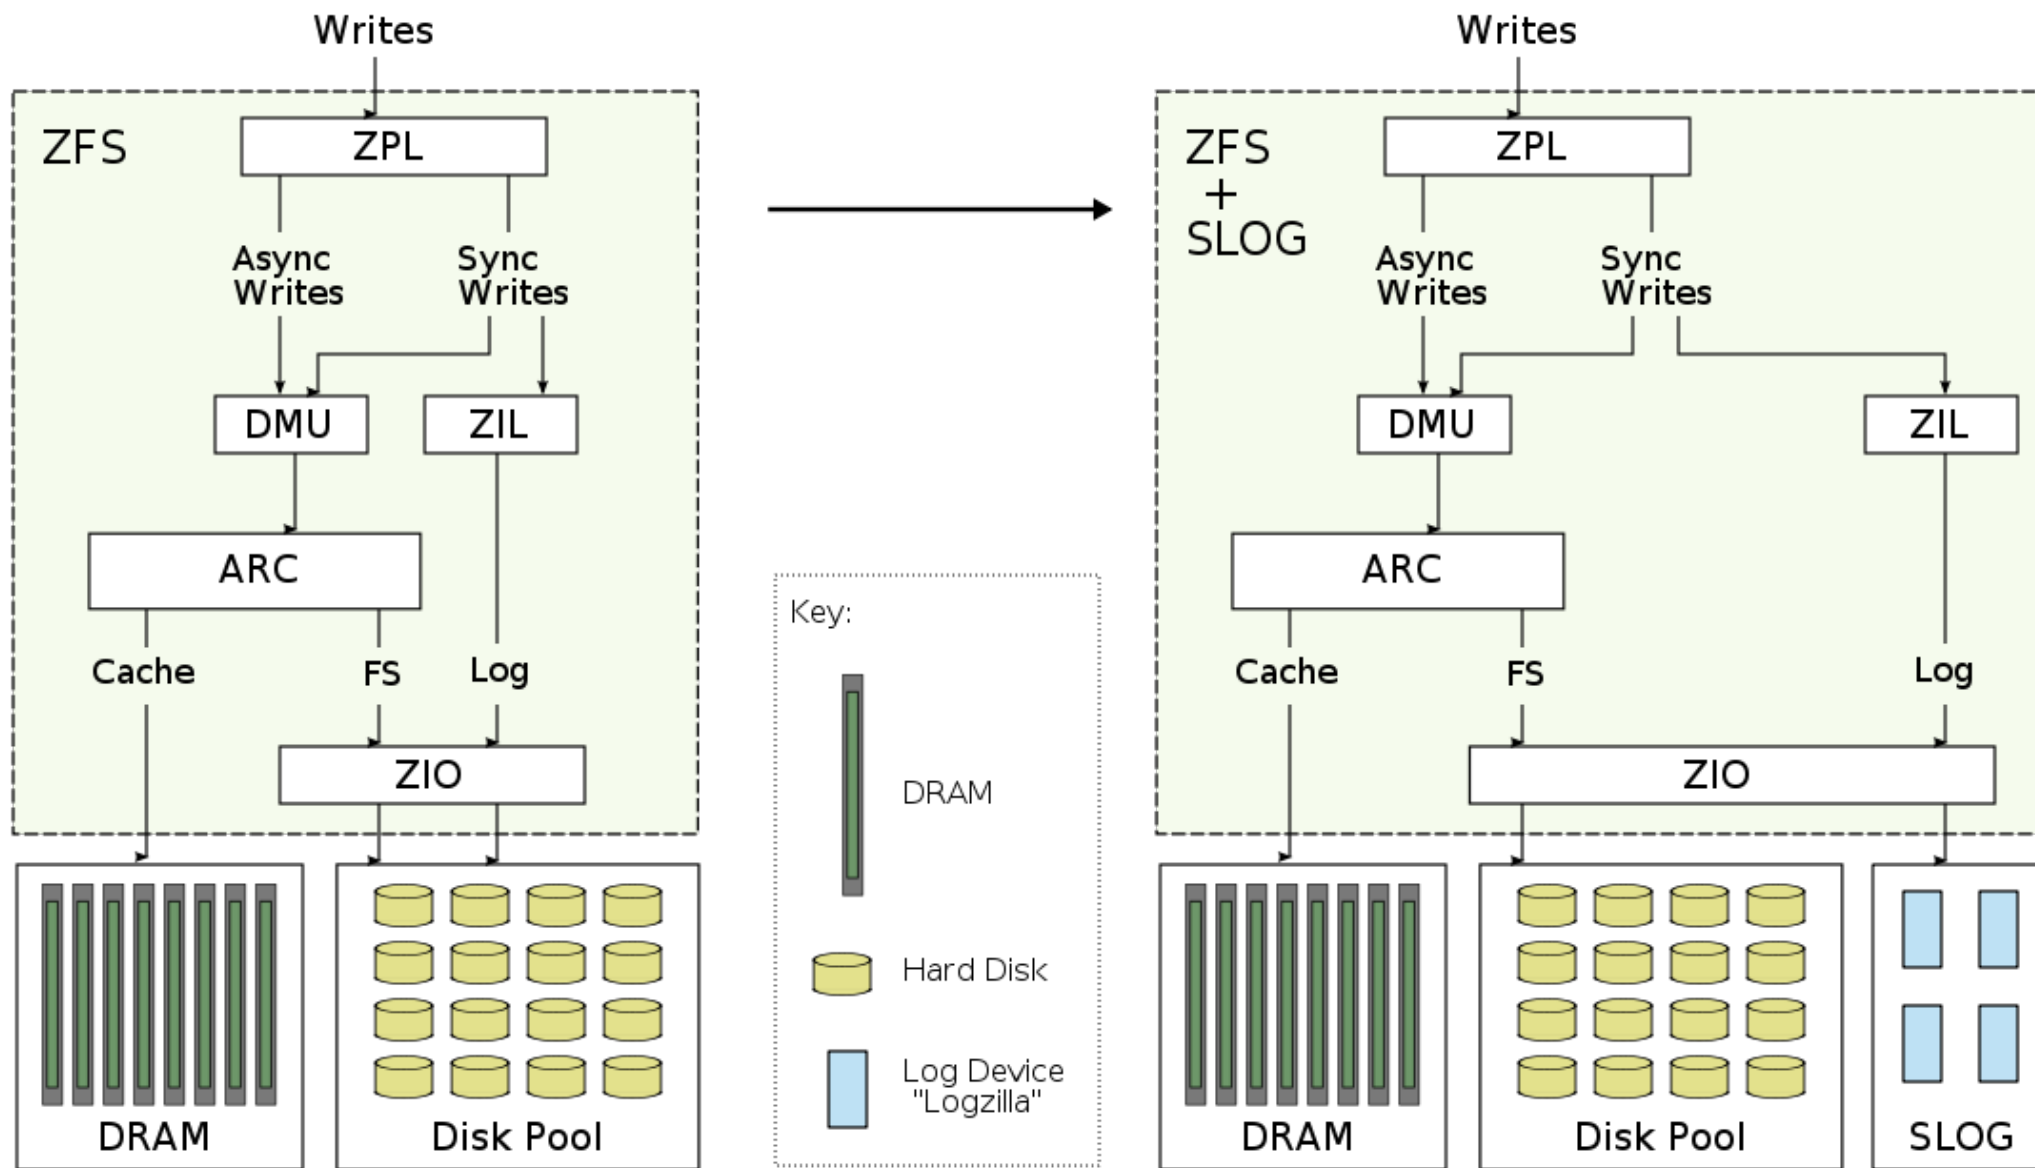


ZFS write data path



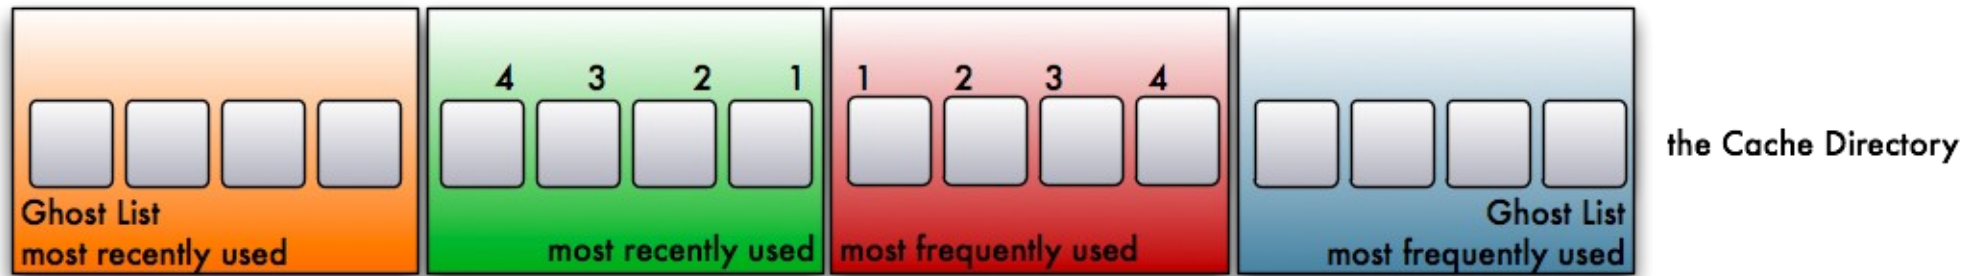
ZFS read data paths





<http://dtrace.org/blogs/brendan/2009/06/26/slog-screenshots>

ZFS L1ARC





<http://wiki.debian.org/WhyDebianForDevelopers>
<http://wiki.debian.org/DebianForNonCoderContributors>
<http://wiki.debian.org/PkgSplit>

André Felipe Machado <andremachado@techforce.com.br>

<http://www.techforce.com.br>

Este texto é licenciado segundo Creative Commons

Atribuição-Uso Não-Comercial-Compartilhamento pela mesma licença 2.5 Brasil

<http://creativecommons.org/licenses/by-nc-sa/2.5/br/>

15set2012



