


Storage for Startups



Valerie Henson



VAH Consulting

`val@vahconsulting.com`




awesome web 2.0
thing!






storage

Which file system to use?




ZFS!

FAT16!!!



ext3!!



Which file system to use?

Make the file system irrelevant

Every file system goes fast with:

- $O(1000)$ files per directory
- File size a few KB to a few GB
- Read-mostly access
- Infrequent file creation/deletion
- Sequential file read/write patterns
- Shallow directory depth (< 10 levels)
- Total file system size $O(100 \text{ GB})$

Example: Web hosting service

- Bad: One big directory with each domain name as a directory
- Better: One directory per letter of the alphabet with domain names in subdirectory
- Best: Hash domain names into appropriate number of subdirectories
- Awful: Put every file in an expensive database