


The Future of File Systems

Bill Bolosky

October 11, 2009

The “Absurd” 10x (=4 year) Disk

- 2.5 hr scan time
(poor sequential access)
- 1 aps / 5 GB
(VERY cold data)
- It’s a tape!



Slide from Jim Gray’s “Rules of Thumb in Data Engineering” talk, 7 May 2002.

Was Jim Wrong?

- No: The technology changed in much the way that he expected.
- No: Those changes mean that you can't use a 2TB disk today the same way that you used a 100GB disk in 2002, just like he predicted.
- Yes: We found applications for the disks that weren't the same as tape applications back then. The objects just got bigger, so you get Maps instead of Kaps.

So What Happened?

- The same thing that happens with technology pretty much all the time:
 - Underlying technologies change based on advances in science and engineering
 - Someone figures out how to use the new technologies to make something useful
- Technological change drives features, not the other way around (usually)
 - Invention is the Mother of Necessity, at least for consumer products.

So, What's the New Tech?

- More of the same in disks: They'll come preloaded since disk size/home net BW (10TB/3 Mb/s \approx 1 year).
- Flash in more devices, still has a lifetime problem for enterprises. Basically low power, high cost disk with great read IOP/s.
- PCM is a game changer: not a block device means file system API is wrong.