

INFINITE LOOP / THE APPLE ECOSYSTEM

A ZFS developer's analysis of the good and bad in Apple's new APFS file system

Encryption options are great, but Apple's attitude on checksums is still funky.

by Adam H. Leventhal - Jun 26, 2016 1:00pm UTC

Share Tweet Email 51



Two hours or so of WWDC keynoting and Tim Cook didn't mention a new file system once?

Andrew Cunningham

This article was originally published on Adam Leventhal's blog in multiple parts.

Apple announced a new file system that will make its way into all of its OS variants (macOS, tvOS, iOS, watchOS) in the coming years. Media coverage to this point has been mostly breathless elongations of Apple's developer documentation.

The overview is divided into several sections. I'd encourage you to jump around to topics of interest or skip right to the conclusion (or to the tweet summary). Highest praise goes to encryption; ire to data integrity.

Table of Contents

- Basics
Paying down debt
Encryption
Snapshots and backup
Management
Space efficiency
Performance
Data Integrity
Redundancy
Crash Consistency
Checksums
Scrub
Summing Up

Basics

APFS, the Apple File System, was itself started in 2014 with Giampaolo as its lead engineer. It's a stand-alone, from-scratch implementation (an earlier version of this post noted a dependency on Core Storage, but Giampaolo set me straight in this comment).

Giampaolo praised the APFS testing team as being exemplary. This is absolutely critical. A common adage is that it takes a decade to mature a file system, and my experience with ZFS more or less confirms this.

Paying down debt

HFS was introduced in 1985 when the Mac 512K (of memory! Holy smokes!) was Apple's flagship. HFS+, a significant iteration, shipped in 1998 on the G3 PowerMacs with 4GB hard drives.

APFS first and foremost pays down the unsustainable technical debt that Apple has been carrying in HFS+. (In 2001 ZFS grew from a similar need where UFS had been evolved since 1977.)

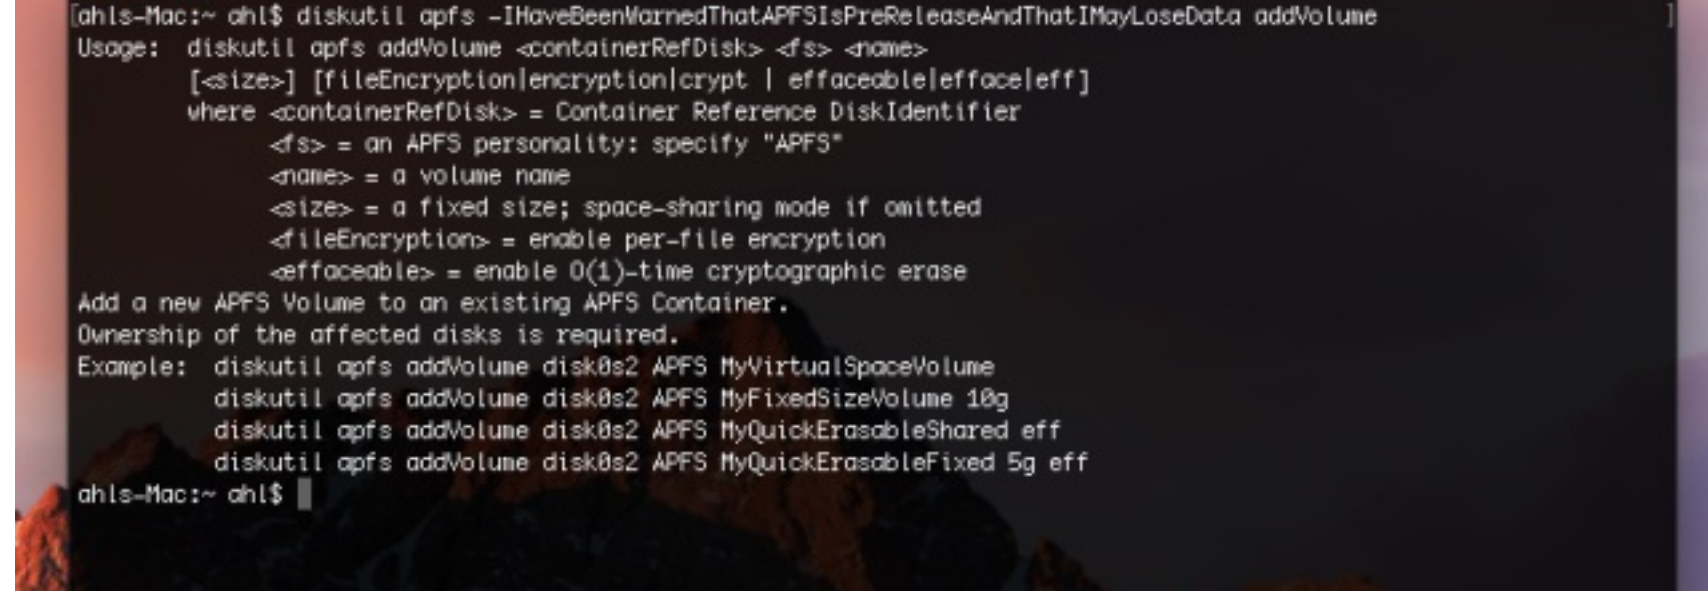
Compression is an obvious common feature that's missing in the APFS feature list. It's conceptually quite easy, I told the development team (we had it in ZFS from the outset), so why not include it?

Encryption

Encryption is clearly a core feature of APFS. This comes from diverse requirements from the various devices; for example, multiple keys within file systems on the iPhone, or per-user keys on laptops.

- Unencrypted
Single-key for metadata and user data
Multi-key with different choices for metadata, files, and even sections of a file ("extents")

Multi-key encryption is particularly relevant for portables where all data might be encrypted, but unlocking your phone provides access to an additional key and therefore additional data.



Enlarge / Can't even make this up!

Related to encryption, I noticed an undocumented feature while playing around with diskutil (which prompts you for interactive confirmation of the destructive power of APFS unless this is added to the command-line: -IHaveBeenWarnedThatAPFSIsPreReleaseAndThatIMayLoseData).

Snapshots and backup

APFS brings a much-desired file system feature: snapshots. A snapshot lets you freeze the state of a file system at a particular moment and continue to use and modify that file system while preserving the old data.

ZFS includes snapshots and serialization mechanisms that make it efficient to back up file systems or transfer file systems to a remote location. Will APFS work like that? Probably not, answered Giampaolo.

While APFS dev manager Eric Tamura demonstrated snapshots at WWDC, the required utilities aren't included in the macOS Sierra beta.

FURTHER READING
BITROT AND ATOMIC COWS: INSIDE "NEXT-GEN" FILESYSTEMS
We look at the amazing features in ZFS and btrfs—and why you need them.

LATEST FROM ARS TECHNICA/UK

- Brexit in Brussels: Juncker's mic drop and political brexploitation
UK votes to leave the European Union [Updated]

KEEPING ON TRUCKING
Sweden trials electrified highway for trucks

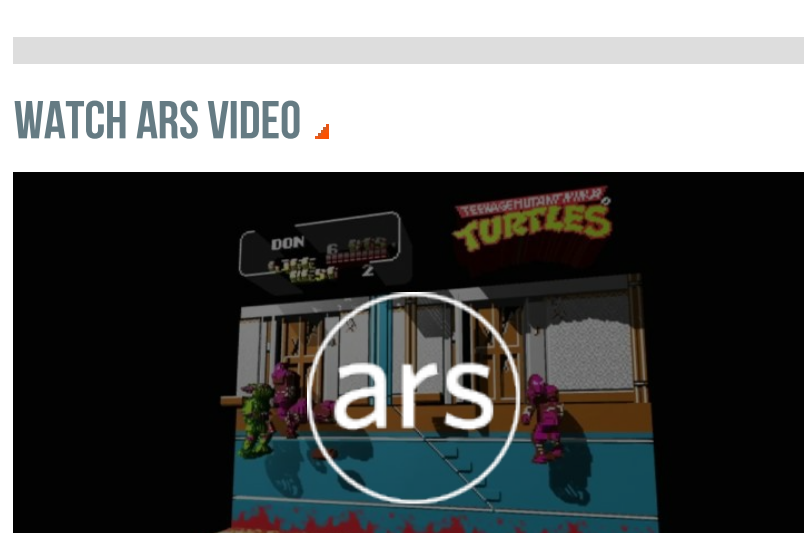
INDEPENDENCE. WHAT INDEPENDENCE?
Brexit: Trade deals will force UK to follow EU data privacy laws anyway

UK voters don't understand Brexit, Google searches suggest

LATEST FEATURE STORY


FROM FILE-SHARING TO PRISON: A Megaupload programmer tells his story
Programmer Andrew Nömm: "I had to be made an example of as a warning to all IT people."

WATCH ARS VIDEO



Hands-on with the emulator that adds depth to old 2D NES games
Standalone 3DNES emulator is a promising starting effort with some rough edges.

STAY IN THE KNOW WITH



LATEST NEWS

- Tokyo Mirage Sessions: #FE: Auto-tuned expression
Y Combinator's Xerox Alto: restoring the legendary 1970s GUI computer
World going to hell? Here are the Solar System's five most livable places
The auto engineering marvels at the Lane Motor Museum are ridiculously cool
To mitigate poverty, Y Combinator set to launch minimum income plan

THE LONG MARCH
China's powerful new rocket makes a successful debut launch

READER COMMENTS 51

Share Tweet Email Google+ Reddit

OLDER STORY NEWER STORY

YOU MAY ALSO LIKE

SITE LINKS

- About Us
Advertise with us
Contact Us
Reprints

SUBSCRIPTIONS

Subscribe to Ars

MORE READING

- RSS Feeds
Newsletters
Visit Ars Technica UK

CONDÉ NAST SITES

- Reddit
Wired
Vanity Fair
Style
Details

Visit our sister sites
Subscribe to a magazine



VIEW MOBILE SITE