Improving the Performance and Durability of SD and eMMC Components with ESS

Abstract

Most embedded devices have simple Flash Translation Layers (FTLs) in order to keep their manufacturing costs low. Such devices experience extremely high write amplification and slow random write performance. Enterprise Storage Stack (ESS) is system level FTL software which fixes these performance and durability problems, letting OEMs build better products at lower cost.

How Embedded Device FTLs Commonly Work

While there are exceptions, the FTLs of Legacy embedded devices typically perform one update per program/erase (PE) cycle. This update can range in size from 512 bytes to the full length of the PE block, but what ever the size of the write, the process was more or less the same:

1. Erase a new PE block,
2. Write the data of the existing PE block up to the point where new logical content begins.
3. Write the new logical content, whether it is one or many blocks long.
4. Write the remaining content in the old PE block.

Early Flash modules were built with 32 kilobyte PE blocks. Accordingly, write amplification of 4KB blocks was 8:1. Practical wear was often only 2:1 at a time when Flash commonly had 100,000 PE cycles.

But as Flash has grown from megabyte modules to gigabyte densities, and fallen from SLC to TLC durability, the effective erase block size has also grown dramatically. Testing of SD cards reveals that current generation effective erase blocks can be as large as 64 megabytes.

There are many reasons for larger PE blocks. Among these are enhanced linear speed and reduced manufacturing costs. But the problem remains that increasing PE block size inherently increases the wear amplification of simple FTLs. If someone has an embedded device with a simple FTL controller and a 64 megabyte erase block and wishes to update a 4KB block, then the write amplification for that single write may be as high as 16,000:1.

Not all large scale embedded devices have such high wear amplification. Testing of SD cards indicates that some – particularly brand-name devices – will operate at random write rates ranging from 25 4KB IOPS to several hundred IOPS. This indicates that such devices have gone beyond the simple FTL described earlier. Similarly, there are a few SD and eMMC modules which do much better, particularly over some restricted test ranges, which will execute thousands of random writes a second at performance levels closer to those of SSDs. Obviously, such devices must have enhanced FTLs which also reduce wear while increasing speed. Even so, these enhanced devices still do not reach the wear and performance levels of ESS.

How ESS Works and Performs

Patented ESS is software which installs in your system and filters the block IO to your Flash storage device. ESS typically resides below volume manager or a file system. ESS executes on the host CPU, usually an ARM or x86, although other CPU architectures are possible.

Instead of random writing, the ESS FTL converts clusters of random writes into metadata bounded, FIFO, atomic, linear writes on PE block boundaries. This linear writing method produces extremely high "random write" performance with extremely low wear amplification.

ESS "random writes" at the linear write speed of your Flash device. For instance, ESS will write to a Class 10 SD card at 10MB/second, or 2,500 4KB random writes a second. By comparison, Class 10 SD cards commonly random write at rates ranging from 1 to 200 4KB IOPS.

Wear amplification with ESS is minimal. When tested using the JEDEC 219A client specification and trace, ESS with 10% free space generates wear amplification of only 1.2:1. By comparison, high performance consumer SSDs with 10% free space typically generate wear amplification of 10:1 in this test. The write amplification of embedded devices tends to be significantly higher than that of consumer SSDs.

Common Occurrences

In the following sections we look at common product development problems experienced by OEMs and examine how ESS can fix these problems.

Phantom phone Calls

Many cell phone users experience missed calls. Even though the phone is functioning properly, and is "on network", the phone never rings. Often this is because the phone storage is busy writing blocks from other applications. A good example of this occurs during a software update, which can often require tens of megabytes of writes.

During these writes, the embedded Flash storage starts stacking up IO latencies, delaying reads. In this case, the reads are the phones "dialer" applications and its support files like the sound file for the ringtone. Often, thousands of IO operations must complete before the dialer can display the first sign of a call, or emit the first ring. If the Flash device is busy with long latency writes, then the phone call can be gone before the dialer ever wakes up.

ESS dramatically lowers write latency, so that reads can make it through. Even on 10 MB/sec media, average 4KB write latency for ESS is under 0.5 milliseconds. By comparison, a device with only 10 write IOPS will have a 100 millisecond latency. So if an application requires 200 random writes, the Flash is dominated by the writes for nearly 20 seconds.

WildFire has measured IO "wait" times longer than 10 seconds on some Android phones, all caused by a single update.

Surveillance Cameras

Surveillance cameras typically collect continuous video streams either in traditional or high definition formats. Video capture commonly encounters two problems. The first is that the storage module cannot keep up with the feed rate of the source, producing lost content. The second is unexpected high levels of write amplification, causing rapid wear out of the Flash component.

One would normally presume that journaled file systems cause little data activity because the metadata and journal data consume just a few bytes relative to the massive amount of video data. But each of these tiny writes can fire off a PE cycle even though the data written is only a few dozen kilobytes in size. Accordingly, developers may find that their design has 6:1 write amplification, or even more. Similarly, they may find that data transfers stutter as devices manage heavy internal demands on the Flash. The bigger erase blocks get, the greater this problem becomes.

ESS fixes these problems by writing all data of the linear stream in the order received, with cross-index to a logical reference memory table. Writing all content linearly gets rid of write amplification and chatter. ESS typically delivers write amplification very near 1:1 for video recording applications with stock file systems in place.

Secondary Storage Problems

Many phones and tablets are sold with support for user-provided secondary storage. ESS provides optimal storage responsiveness on these devices, without requiring specialized controllers. Even better, the device can typically achieve it's "speed rating" in real-world tests, rather than just referring to a never achievable theoretical value.

Decades of Durability

The ability to generate all of your random write activity with write amplification near 1:1 means that finished solutions will likely have dramatically more useful life than would otherwise be the case. Components which might before have lasted months will now last years; and those with years of life expectancy may last decades.

Retrofitting Under-Engineered Products

If the storage component of your product does not meet specifications, ESS can solve your problems. ESS can be compiled into your operating system in minutes. And it is a low cost fix to a broad range of durability and performance problems.

A Cost-reducing Solution to Problems

ESS software is extremely affordable in its own right. But its ability to radically improve the write performance and durability of any embedded Flash component gives you the capacity to either increase quality or to decrease cost while maintaining quality standards. ESS is so fast, and Flash wear is so well managed, that Flash becomes practical for DRAM swap, a nearly unthinkable option with stock Flash performance and wear.

Sam Anderson
EasyCo LLC
3/9/16

EasyCo LLC
DBA WildFire Storage
220 Stanford Drive
Wallingford PA 19086 USA
Tel: (+1) 610-237-2000
Free: 888-473-7866
Email: sales@EasyCo.com
Web: http://EasyCo.com