Discardability in drivers has nothing to do with discardability in user-mode (which has nothing to do with discardability, really)

devblogs.microsoft.com/oldnewthing/20160901-00

September 1, 2016

Raymond Chen

Some time ago we discussed what the `DISCARDABLE` keyword means. In summary: It has no effect in user mode, and in kernel mode, it means that the memory should be thrown away after initialization is complete.

The two uses for the `DISCARDABLE` flag aren't really all that related to each other. They are just two components (the 16-bit Windows memory manager and the 32-bit Windows device driver loader) that evolved independently, both of which saw a bit and said, "Hey, I can use this for something!" but they each used it for something different.

Once upon a time, back in the days when two megabytes was a decent amount of memory, someone came up with an optimization: Since driver initialization code and driver initialization data are both thrown away at the end of initialization, we can merge them into the same page and save 4KB of memory. Multiply this by the number of drivers in the system, and that's a lot of memory being saved at system boot, which in turn means that we can boot in less memory. This is important when you are trying to minimize your system requirements.

The 32-bit device driver folks needed a bit in the segment attributes to say "This memory should be thrown away once initialization is complete." They saw a bit lying around with the sticker `DISCARDABLE` written on it and said, "Discardable, yeah, that perfectly describes what we want. Thanks for reserving that bit for us!"

Okay, so that explains the `DISCARDABLE` bit, and the important-at-the-time 4KB memory savings explains why driver initialization code and data are merged into a single page. But this results in a dreaded W|X page, which negates any benefit of DEP! Why are drivers still using this optimization that isn't that useful any more?

Inertia, probably.

You had a driver originally written in the days when 4KB was a lot of memory, so it used this one weird trick to save 4KB of memory. The driver then evolves over time, but the merging of driver initialization code and data hangs around because things stay the same until something makes them change.

Who knows, maybe there will someday be evolutionary pressure to get all the old drivers to change their section attributes. (I suspect pressure is low because the W|X page is not in memory for very long, so the attack window is hard to hit reliably.)

Raymond Chen

Follow