Where is this CRC that is allegedly invalid on my hard drive?

devblogs.microsoft.com/oldnewthing/20130711-00

July 11, 2013

Raymond Chen

If you're unlucky, your I/O operation will fail with `ERROR_CRC`, whose description is "Data error (cyclic redundancy check)." Where does NTFS keep this CRC, what is it checking, and how can you access the value to try to repair the data? Actually, NTFS does none of that stuff. The CRC error you're getting is coming from the hard drive itself. Hard drives nowadays are pretty complicated beasts. They don't just plop data down and suck it back. They have error-checking codes, silent block remapping, on-board caching, sector size virtualization, all sorts of craziness. What's actually happening is that the file system asks the hard drive to read some data, and instead of handing data back, the hard drive reports, "Sorry, I couldn't read it back because of a CRC error." NTFS itself doesn't do any CRC checking. "Well, that's awfully misleading. If NTFS is reporting a CRC error, then that makes the user think that NTFS is maintaining CRCs. Shouldn't it just report 'general I/O error' instead of a more specific error?" NTFS is just bubbling upthe error message from the hard drive. This dates back to the old MS-DOS days, where the BIOS reported hard drive error codes, and those error codes were returned all the way back to the application. Who knows, maybe the end-user knows enough about drive technology that they can tell the difference between a CRC error and a seek error. (For example, a seek error may be fixed by removing the floppy disk and reinserting it, or by recalibrating.) What about the converse? If an I/O operation completes successfully, does that provide metaphysical certitude that the data read back exactly matches the data that was originally written? No. It only provides metaphysical certitude that *the hard drive reported that* the data read back exactly matches the data that was originally written, *as far as it could tell*.

Generally speaking, upper layers of a system trust that a lower layer is functioning properly (and often they have no way of detecting a malfunction in the lower layer, anyway). If the hard drive says that it read the data successfully, well, the hard drive is the expert at this sort of thing, so who are we to say, "Nuh uh, I think you're wrong"?

Raymond Chen

Follow