

# Disk and File I/O performance with ETW traces: Miscellany

 [devblogs.microsoft.com/oldnewthing/20201127-00](https://devblogs.microsoft.com/oldnewthing/20201127-00)

November 27, 2020



Raymond Chen

Building upon the [preliminary notes on analyzing Disk and File I/O performance with ETW traces](#), I'll go into some more of the columns in the data.

The *Priority* column represents the I/O priority, where higher numbers are more urgent and lower numbers are less urgent. As a general rule, the system will attempt to satisfy higher priority I/O before low priority I/O, although it will let some lower priority I/O trickle through every so often to avoid starvation.

Normal priority is priority 3, and this is the priority for your typical application-initiated I/O. Applications can also issue I/O at priority zero, which makes it take a back seat to other I/O.<sup>1</sup> Search indexers use priority zero, for example, so that indexing operations do not interfere with normal I/O. Applications do not have access to I/O priority above 3. Those are reserved for system components.

Another interesting column is *Source*. Values for this column are *Original*, *Prefetch*, *VolSnap*, and *Unknown*.

Most I/O is marked as *Original*, which means basically that there's nothing special going on. The I/O occurred because somebody explicitly asked for it.

If the I/O is marked *Prefetch*, then this is I/O was initiated by [ReadyBoot](#), which should not be confused with *ReadyBoost*. (The similarity in the names is quite unfortunate.) *ReadyBoot* (without the s) prefetches data off the disk during boot based on its observations of what data was read from disk during previous boot sessions, hoping to stay one step ahead of the system's actual I/O needs.

Another curiosity of ReadyBoot I/O is that it occurs below the file system level, so the file name in the events will be *Unknown*. You can use the *Source* column to identify that the I/O came from ReadyBoot.

I/O that is marked as *VolSnap* is initiated by volume snapshots. Volume snapshots capture the state of the disk at a moment in time. It would be expensive to make a copy of every byte of the disk when a snapshot is taken, so instead the snapshot is done lazily. The first time a write occurs to a part of the disk that was part of a snapshot, you actually get multiple operations:

- A read is issued to the disk to obtain the data that is about to be overwritten.
- A write is issued to the disk to write the old data to some other part of the disk for safekeeping.
- The original write request is allowed to go through.

The I/O operations resulting from volume snapshot lazy-copies are marked with the *VolSnap* source.

You will see *Unknown* as the source for Flush operations, since they apply to the disk as a whole, not to any individual file.

<sup>1</sup> Notice that disk priority is opposite of bug priority. For bug priority, lower numbers represent greater importance.

Raymond Chen

**Follow**

