Looking at the world through kernel-colored glasses

devblogs.microsoft.com/oldnewthing/20110512-00

May 12, 2011

Raymond Chen

During a discussion of the proper way of cancelling I/O, the question was raised as to whether it was safe to free the I/O buffer, close the event handle, and free the `OVERLAPPED` structure immediately after the call to `CancelIo`. The response from the kernel developer was telling.

> That's fine. We write back to the buffer under a try/except, so if the memory is freed, we'll just ignore it. And we take a reference to the handle, so closing it does no harm.

These may be the right answers from a kernel-mode point of view (where the focus is on ensuring that consistency in kernel mode is not compromised), but they are horrible answers from an application point of view: Kernel mode will write back to the buffer and the `OVERLAPPED` when the I/O completes, thereby corrupting user-mode memory if user-mode had re-used the memory for some other purpose. And if the handle in the `OVERLAPPED` structure is closed, then user mode has lost its only way of determining when it's safe to continue! You had to look beyond the literal answer to see what the consequences were for application correctness. (You can also spot the kernel-mode point of view in the clause "if the memory is freed." The developer is talking about freed from kernel mode's point of view, meaning that it has been freed back to the operating system and is no longer committed in the process address space. But memory that is logically freed from the application's point of view may not be freed back to the kernel. It's usually just freed back into the heap's free pool.) The correct answer is that you have to wait for the I/O to complete before you free the buffer, close the event handle, or free the `OVERLAPPED` structure. Don't fall into this trap. The kernel developer was looking at the world through kernel-colored glasses. But you need to look at the situation from the perspective of your customers. When the kernel developer wrote "That's fine", he meant "That's fine *for me*." Sucks to be you, though. It's like programming an autopilot to land an airplane, but sending it through aerobatics that kill all the passengers. If you ask the autopilot team, they would say that they accomplished their mission: Technically, the autopilot did land the airplane. Here's another example of kernel-colored glasses. And another.

Epilogue: To be fair, after I pointed out the kernel-mode bias in the response, the kernel developer admitted, “You’re right, sorry. I was too focused on the kernel-mode perspective and wasn’t looking at the bigger picture.”

Raymond Chen

Follow