

I/O Models

1. In the I/O subsystem, explain the two stages by which data is delivered from an I/O device to a user-level process.
2. What are the five I/O models? How do they handle the two stages of data reception?
3. Why is process blocking generally not a concern for *write* I/O operation?
4. When can a process block on a write I/O operation?
5. Compare the blocking I/O model with the I/O Multiplexing model? How are they similar and how are they different?
6. Which I/O model is used for event-driven programming? Explain why.
7. What is asynchronous I/O? In practice, why is it hard for any OS to support true asynchronous I/O?
8. Broadly we have learnt about two I/O classification systems. One is a device-level I/O classification, namely programmed I/O, interrupt-driven I/O, and DMA-based I/O. The other classification is for application-level "I/O Models", namely, blocking, non-blocking, signal-driven, and asynchronous I/O models. Explain the analogies and differences between these two I/O classification systems.
9. Explain the two key steps in delivering data in packets from a network card to a user-level process reading from a network socket?
10. Explain the difference between (a) signal-driven I/O and asynchronous I/O, (b) blocking I/O and I/O multiplexing.