

Answers.

Version 1	Version 2	Version 3	Version 4
1.D	C	B	A
2.B	A	D	C
3.A	D	C	B
4.A	D	C	B
5.B	A	D	C
6.C	B	A	D
7.C	B	A	D
8.C	B	A	D
9.D	C	B	A
10.D	C	B	A
11.B	A	D	C
12.B	A	D	C
13.C	B	A	D
14.A	D	C	B
15.A	D	C	B
16.B	A	C	B
17.C	B	A	D
18.D	C	B	A
19.D	C	B	A
20.B	A	C	B
21.C	B	A	D
22.C	B	A	D
23.C	B	A	D
24.C	B	A	D
25.B	A	C	B
26.D	D	D	D
27.B	A	D	C
28.A	D	C	B
29.B	A	C	B
30.C	B	A	D

31.

a)

The optimal algorithm requires knowledge of the future. The LRU algorithm approximates this using the recent past.

b)

The reference bit can be used to produce a number representing how recently a page has been accessed. Every so often the reference bits can be checked and pushed into a reference byte (from the most significant end). The greater the value of the referenced byte the more recently (and more often) a page has been accessed. On each check the reference bit is cleared.

32.

a)

So that direct process addresses don't have to be used and maintained in the communication code.

b)

To reduce the need for buffers to hold the messages.

c)

To increase the overall speed of the system. This way writers don't have to wait for their messages to be received before carrying on.

33.

a)

Each memory access that doesn't have the page and frame in the TLB requires 3 memory accesses (the 2 for the levels of the page table and 1 for the data itself).

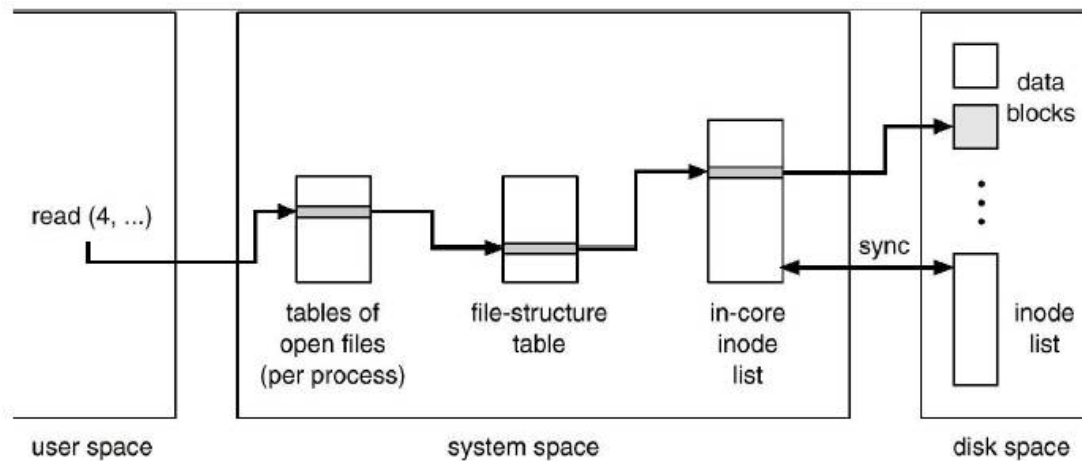
$$\begin{aligned} \text{EAT} &= 0.9 * (5 + 100) + 0.1 * (5 + 3 * 100) \\ &= 94.5 + 30.5 \\ &= 125 \text{ ns} \end{aligned}$$

b)

$$\begin{aligned} \text{EAT} &= 0.8 * 100 + 0.2 * (10,000,000 + 100) \\ &= 80 + 2,000,020 \\ &= 2,000,100 \text{ ns or } 2.0001 \text{ ms also accepted } 2.00008 \text{ ms} \end{aligned}$$

34.

a)



b)

All open devices are shared with the parent process. The child gets a copy of the parent process's open file table.

c)

Maintaining the information for all data is usually regarded as too expensive. Maintaining it for metadata protects the integrity of the file system. The logs are used to return the file system to a consistent state in the event of some accident.

35.

a)

The first line
child
another child
another parent
parent
another child
another parent

b)

All processes get copies of the same output buffers. The child processes therefore display the output generated by their parents.

c)

3

d)

The order of execution of the processes is uncontrolled.

36.

a)

They continue with "init" as their parent process.

b)

To prevent the accumulation of zombie processes.

37.

a)

The usual implementations of these constructs require shared memory, and distributed systems don't usually have shared memory.

b)

A controller process allocates the resource to processes that request it.

A request message is sent to the controller, the requesting process blocks until it receives a reply message from the controller. When it has finished using the resource the process sends a release message to the controller.

It suffers from the controller being a bottle-neck and also problems to do with controllers or currently allocated processes going down.

38.

a)

Things like:

name

access times

owner and privilege information

size

location

b)

Mac HFS and HFS+

NTFS

c)

It helps to structure the information in the file. This makes it easier to work with a particular type of information in the file, without requiring multiple files.

d)

If the file is moved to a file system that doesn't support multiple data streams then the information is lost. Or else it has to be split up into several files.

39.

a)

One of the directories in the path will indicate that its contents is on another machine, via the local mount table. When the remote directory was mounted the client machine received information on the machine and inode associated with the remote directory. This is checked with the Virtual File System layer. The client sends a message across the network to the server for the file. The file is accessed on the server and results are sent back to the client.

b)

(Any two of these.)

Because each client maintains its own information on remote directories, each machine must be updated when directories are moved.

Different clients can mount remote directories at different places, this means the location (pathname) of a file depends on the particular client machine.

There are limits on the number of machines that can appear in export tables.

c)

It uses a location server rather than maintaining the information on each client. This means that changes only need to be made at the location servers.

The standard approach is to mount AFS at /afs on all clients.

All machines can access all files on the network as long as the user is included in the access control lists associated with the directory.

40.

a)

Polling is useful if the device cannot generate interrupts or if the resource becomes available very quickly, saving the overheads of interrupt handling.

Interrupt handling is good when the resources take a while to become available or become available over a wide range of times. This saves the busy waiting involved with polling.

b)

Polling could be used for high speed serial connections.

Interrupt handling could be used for disk devices using DMA.

41.

a)

i. Since D1 owns F2, it could allocate itself the write permission.

ii. It could switch to D3 then to D2 then write F2.

b)

D3 as owner of F1 and D1 via control of D3.

c)

There is a cycle of switch permissions that covers all domains.

d)

Either ACLs or capabilities or key/locks.