

1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28

Methods for handling Deadlocks

Generally speaking, we can deal with the deadlock problem in one of three ways:-

- We can use a protocol to prevent or avoid deadlocks, ensuring that the system will never enter a deadlocked state.
- We can allow the system to enter a deadlocked state, detect it, and recover.
- We can ignore the problem altogether and pretend that deadlocks never occur in the system.

→ The third solution is the one used by most operating systems, including Linux and Windows.

To ensure that deadlocks never occur, the system can use either a ~~dead~~ deadlock prevention or a deadlock-avoidance scheme.

S	M	T	W	T	F	S
30	31					1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29

→ Deadlock prevention provides a set of methods to ensure that at least one of the necessary conditions cannot hold.

→ These methods prevent deadlock by constraining how requests for resources can be made.

→ Deadlock avoidance requires that the operating system be given additional information in advance concerning which resources a process will request and use during its lifetime.

→ With this additional knowledge, the operating system can decide for each request whether or not the process should wait.

→ To decide whether the current request can be satisfied or must be delayed, the system must consider the resources currently available, the resources currently allocated to each process, and the future requests and releases of each process.