WRED Weighted RED

Most routers have packet queues, which allow them to hold packets in their buffers during periods of congestion, rather than discarding them.

However, the buffers have limited size and the queue is allowed to fill to its maximum size. If the queue is bigger than the buffer, surely some packets must be discarded; the decision is depended on which algorithm is used. One of the congest management algorithm is Random Early Discard (RED) algorithm whereby random frames are refused admission to the queue once a threshold has been exceeded. Cisco routers do not support RED but it supports the better one: WRED.

The difference between RED and WRED is that WRED can selectively discard lower-priority traffic when the interface begins to get congested. In WRED, a queue may have several different queue thresholds.

By default, WRED uses a different RED profile for each weight. Each queue threshold is associated to a particular IP precedence or DSCP; for example, a queue may have lower thresholds for lower priority packet so that it drops less important packets more aggressively than important packets during periods of congestion.

Global Synchronization:

TCP has automatic recovery from dropped packets (usually when the network is congested). The sender reduces its sending rate for a certain amount of time, and then tries to find out if the network is no longer congested by increasing the rate again. This is known as the slow-start algorithm.

Almost all the senders will use the same time delay before increasing their rates. When these delays expire, at the same time, all the senders will send additional packets, the router queue will again overflow and packets will be dropped, the senders will all back off for a fixed delay… This pattern of each sender decreasing and increasing transmission rates at the same time as other senders is referred to as “global synchronization” or “TCP synchronization” and leads to inefficient use of bandwidth, due to the large numbers of dropped packets, which must be retransmitted.

WRED reduces the chances of tail drop (used by TCP) by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once and minimizes the chances of global synchronization. Thus, WRED allows the transmission line to be used fully at all times -> C is correct.

Note: Tail drop is the simplest technique to limit queue size. When the queue is full, it simply discards any new packets until there is space in the queue again.