

TCP Timers

Overview

The TCP keepalive timeout defines the time interval for the TCP connection to check whether the FCIP connection is working correctly. This ensures that FCIP connection failures are quickly detected, even if the connection is inactive. When the TCP connection is idle for the required time interval, a TCP keepalive packet is sent to verify if the link is active. The MDS switch uses the "TCP keepalive timeout" command to specify the number of keepalive timeout seconds (the range is between 1 and 7200 seconds, and the default value is 60). During the time interval when the connection is inactive within the configured time interval, eight keepalive probes are sent at an interval of 1 second. If no response to these eight probes is received and the connection remains idle, the FCIP connection will be automatically closed. TCP uses multiple timers to ensure that there is no excessive delay in communication. Some of these timers are elegant and can handle problems that were not obvious during the initial analysis. Every timer used by TCP helps to ensure that data is sent correctly from one connection to another.

Timer Types

The four central timers used by TCP implementation are

1. Timeout Timer
2. Time Waiting Timer
3. Keep Alive Timer
4. Persistent Timer

Time Out Timer

TCP uses an outgoing timer to retransmit lost segments.

- The sender starts the out-of-office timer after transmitting the TCP segment to the receiver.
- If the sender receives an acknowledgment before the timer appears, they will stop the timer.

- If the sender does not receive a declaration, the timer bursts, a TCP retransmission occurs.
- The sender resends the equivalent segment and resets the timer.
- The value of the timeout timer is dynamic and will change with the amount of traffic on the network. The timeout timer is also called the retransmission timer.

Time Wait Timer

TCP uses a time wait timer during connection termination.

- The sender starts the time waiting for the timer after sending the ACK of the second FIN segment.
- If the final confirmation is lost, retransmission is allowed.
- It can prevent the port that has just been closed from being opened again quickly by another application.
- It ensures that any segment that flows to the port that was just closed is discarded.
- The value of your time waiting for the timer is usually set to twice the lifetime of the TCP segment.

Keep-Alive-Timer

TCP uses Keep-Alive-Timer to stop TCP connections that have not been used for a long time.

- Each time the server receives a message from the client, it resets the keepalive timer to 2 hours.
- If the server does not receive any message from the client for two hours, it will send ten detection segments to the client.
- These probe segments are sent within 75 seconds.
- If the server does not receive a response after sending 10 test segments, the client has failed.
- The server then automatically terminates the connection.

Persistent Timer

- TCP uses a persistent timer to affect a deadlock of Zerowindow size.
- Even if the peer closes its receiver window, it still keeps the flow of window size information.