

Re: Difference between ServerAliveInterval and TCPKeepAlive?

Source: <http://www.derkeiler.com/Newsgroups/comp.security.ssh/2005-01/0252.html>

From: Richard E. Silverman (*res_at_qoxp.net*)

Date: 01/23/05

Date: 23 Jan 2005 01:47:10 -0500

>>>> *"David" == David Deutsch <sptest@71broadway.info> writes:*

David> Could someone give me a quick explanation of the difference
David> between these two (I have already read the man page)? I'm
David> trying to a) keep my NAT from timing out my connection, and B)
David> terminate the process asap if the NAT **does** time it out (so
David> that I can detect it and restart the process).

TCPKeepAlive is not SSH-specific; it simply sets the "keep-alive" option on the underlying TCP connection. TCP keepalive is not intended to keep anything alive, as in avoiding NAT timeouts. Rather, it detects and reaps half-dead TCP connections to prevent them from building up over time. It typically operates over too long a time scale to be useful in defeating NAT timeout, and that scale is usually only tunable for the whole TCP stack, not per-connection.

ServerAliveInterval controls keepalive messages which are sent within the SSH protocol; since they are application-level and hence tunable per-connection, this is what you want to use.

--

Richard Silverman
res@qoxp.net