Subject: Re: Socket: how to check wether new data arrived? timeout = NULL?
Posted by rylek on Fri, 15 Dec 2006 20:45:01 GMT
View Forum Message <> Reply to Message

Hello!

You are right, current implementation doesn't support reporting the reason for termination of the last Socket::Wait call. The motivation here was that, in practice, select can return due to several reasons - data available on read sockets, ready to send on output sockets, broken connections, EOFs and other network-related situations (OOB data, for instance). I was afraid that, in order to be able to report all combinations of these events, a very complex result structure would be necessary, which seemed to me an overshot. Honestly it still does, although it is true that repeated calls to select (called internally from the PeekXXX method group) might slow down the connection somewhat. My personal reasoning behind letting it be (at least for the time being) is that under normal circumstances the physical network connection is typically orders of magnitude slower than the processing unit of the computer so that a few nanoseconds spent within the call to the select function can hardly substantially hamper performance.

However, the current interface enables (and it is quite possible we'll support this in the future) the introduction of internal status flags which would allow for optimization of the PeekXXX calls so that, after a call to Socket::Wait, the following Peek wouldn't have to call the select function again but just report the socket status as detected by the previous call to Socket::Wait.

It is also true that you cannot Wait infinitely long (although some other methods of the Socket interface, like PeekCount or ReadCount, support Null for their timeout argument and interpret it as the request to wait for an unlimited time until the requested condition is met); however, I personally see not much difference between waiting forever and waiting e.g. for a day (86400000 milliseconds) at a time; honestly, all servers I have written so far usually Wait for just a single second; it seems to me that being able to cycle through the main server loop every few seconds or so it a good idea in almost any case, as it lets you do hosts of other housekeeping actions (like cleaning up idle caches or connections, checking for externally triggered server shutdown or displaying a watchdog message indicating the server is alive).

Regards

Tomas