

Investigating the LTM TCP Profile: Slow Start



Jason Rahm, 2008-09-12

Introduction

The LTM TCP profile has over thirty settings that can be manipulated to enhance the experience between client and server. Because the TCP profile is applied to the virtual server, the flexibility exists to customize the stack (in both client & server directions) for every application delivered by the LTM. In this series, we will dive into several of the configurable options and discuss the pros and cons of their inclusion in delivering applications.

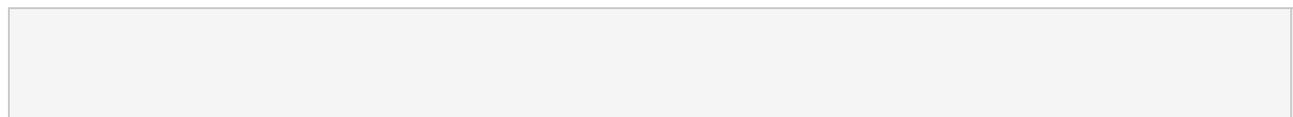
1. [Nagle's Algorithm](#)
2. [Max Syn Retransmissions & Idle Timeout](#)
3. [Windows & Buffers](#)
4. [Timers](#)
5. [QoS](#)
6. [Slow Start](#)
7. [Congestion Control Algorithms](#)
8. [Acknowledgements](#)
9. [Extended Congestion Notification & Limited Transmit Recovery](#)
10. [The Finish Line](#)

Quick aside for those unfamiliar with TCP: the [transmission control protocol](#)(layer 4) rides on top of the [internet protocol](#)(layer 3) and is responsible for establishing connections between clients and servers so data can be exchanged reliably between them.

Normal TCP communication consists of a client and a server, a 3-way handshake, reliable data exchange, and a four-way close. With the LTM as an intermediary in the client/server architecture, the session setup/teardown is duplicated, with the LTM playing the role of server to the client and client to the server. These sessions are completely independent, even though the LTM can duplicate the tcp source port over to the server side connection in most cases, and depending on your underlying network architecture, can also duplicate the source IP.

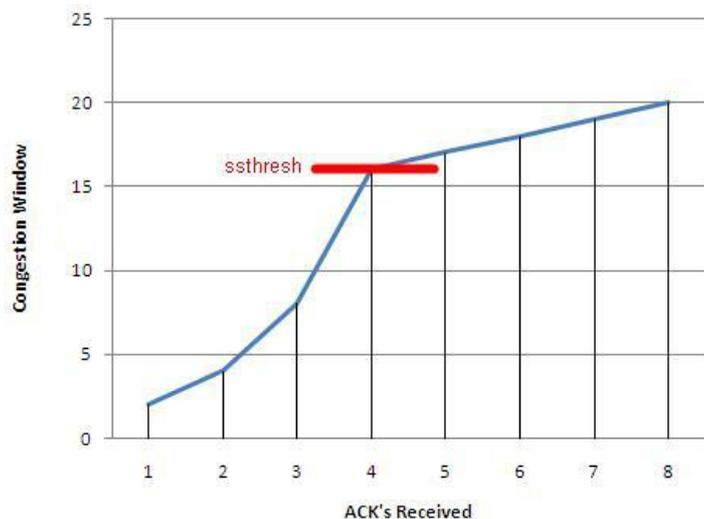
TCP Slow Start

Refined in [RFC 3390](#), slow start is an optional setting that allows for the initial congestion window (cwnd) to be increased from one or two segments to between two and four segments. This refinement results in a larger upper bound for the initial window:



```
If (MSS <= 1095 bytes)
then win <= 4 * MSS;
If (1095 bytes < MSS < 2190 bytes)
then win <= 4380;
If (2190 bytes <= MSS)
then win <= 2 * MSS;
```

The congestion window (cwnd) grows exponentially under slow start. After the handshake is completed and the connection has been established, the congestion window is doubled after each ACK received. Once the congestion window surpasses the slow start threshold (ssthresh, set by the LTM and dependent on factors like the selected congestion algorithm), the tcp connection is converted to congestion avoidance mode and the congestion window grows linearly. This relationship is represented in the following graph.



Slow Start is triggered at the beginning of a connection (initial window), after an idle period in the connection (restart window), or after a retransmit timeout (loss window). Note that this setting in the profile only applies to the initial window. Some advantages of increasing the initial congestion window are eliminating the wait on timeout (up to 200ms) for receivers utilizing delayed acknowledgements and eliminating application turns for very short lived connections (such as short email messages, small web requests, etc). There are a few disadvantages as well, including higher retransmit rates in lossy networks. We'll dig a little deeper into slow start when we cover the congestion control algorithms.

An excellent look at slow start in action can be found [here](#).

[Get the Flash Player](#) to see this player.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com