LTM: Dueling Timeouts

Deb Allen, 2008-15-04

LTM has a number of timeouts that can be set to promote active connection management. This article provides a high level explanation of LTM timeout options and a few guidelines on configuring them appropriately.

LTM Connection Management Overview

LTM manages each load balanced connection explicitly by keeping track of it in the connection table while it is active. The connection table contains state information about clientside flows and serverside flows, the relationships between them, and the last time traffic was seen on each.

Each connection in the connection table consumes system resources maintaining the table entry and monitoring the connection's status. LTM (like any other IP system) must determine when a connection is no longer active and retire it in order to avoid exhaustion of critical system resources which are at risk if the connection table grows unchecked. For instance, excessive memory and processor cycles may be consumed managing the table itself, and ephemeral ports required for the LTM end of serverside flows may be exhausted if they are not recycled.

LTM Connection Reaping

Connections that close normally (4-way close) or are reset by either side are retired from the connection table automatically. A significant number of connections often go idle without closing normally for any number of reasons. These connections must be "reaped" by the system once they have been determined to be inactive.

In order to promote proactive connection retirement or recycling (also known as "reaping"), several different timeouts may be configured in LTM to tear down connections that have seen no active traffic after a specified period of time. Most of these timeouts are configurable to meet the needs of any application. Some are not.

Since you can set timeouts in multiple places, it's important to understand that sometimes more than one timeout will affect the same connection. The optimal timeout configuration is one that retains idle connections for an appropriate amount of time (which of course will vary by application) before deciding they are inactive and retiring them to conserve system resources.

LTM Timeout Options

LTM connections may be timed out by protocol profiles or SNATs associated with the virtual server handling the connection. Connections not handled by a virtual server may be timed out based on SNAT automap or VLAN group settings.

Here is a list of the possible LTM connection timeouts, their default values, and whether that value is configurable:

Option	Default in seconds	Configurable?
Protocol profiles (tcp, fastHTTP, or fastL4)	300	Y
Protocol profile (udp)	60	Y
SNAT / SNATpool	300	Y
SNAT automap	300	N
VLAN group	300	N

The shortest timeout that applies to a connection will always take effect. In some cases, that's not desirable. For example, when configuring a forwarding virtual server that's intended to carry long-standing connections that may go idle for long periods of time (such as SSH sessions), you can configure a long idle timeout on the related protocol profile (tcp in this case), but the 300 second static timeout will still take effect if SNAT automap is also enabled. You may be able to avoid the non-configurable global SNAT automap timeout by creating and applying a similar but more specific object with a configurable timeout (SNAT or SNATpool) as detailed in SOL6017 (see below for link.)

Other LTM Timeouts

There are a couple of other timeouts that can be configured in LTM that do NOT cause connections to be reaped: OneConnect and persistence.

The OneConnect timeout controls only how long an idle serverside flow will be available for re-use, and may cause a serverside connection to be closed after it goes idle for a time. Since that connection will never have been actively in use, no active clientside connections will be affected, and a new serverside flow will transparently be selected or established for new connections. OneConnect timeout settings need not be coordinated with other idle timeouts.

Persistence timeouts are actually idle timeouts for a session, rather than a single connection. With that in mind, persistence timeouts should typically be set to a value slightly larger than the applicable connection idle timeouts to allow sessions to continue even if a connection within it is timed out.

More Information About LTM Timeouts

SOL7606: Overview of BIG-IP LTM idle session timeouts

SOL1638: Reasons why all SNATs do not use the global TCP idle timeout for SNATs

SOL5401: Idle connections may be allowed to exist after the idle timeout expires

SOL6017: BIG-IP LTM SNAT automap has a static timeout value of 300 seconds

(Please disregard the note in this solution indicating that SNAT automap timeout does not affect VS-handled connections. The information in SOL7606 above is more current and accurate. A correction has been requested.)

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

F5 Networks, Inc.	F5 Networks	F5 Networks Ltd.	F5 Networks
Corporate Headquarters	Asia-Pacific	Europe/Middle-East/Africa	Japan K.K.
info@f5.com	apacinfo@f5.com	emeainfo@f5.com	f5j-info@f5.com