

Analyzing Hydra Historical Statistics – Part 1

Fabio Massimo Ottaviani – EPV Technologies

The IBM Virtualization Engine TS7700, code named Hydra, is the new generation of tape virtualization solution for mainframe systems designed to replace the IBM Virtual Tape Server (VTS).

The Hydra architecture has been completely redesigned to make it more modular and scalable and to allow an easy implementation of advanced Disaster Recovery and Business Continuity solutions. The Hydra statistics production system has also been redesigned to provide all the needed metrics..

Two types of statistics are provided:

- Point-in-time statistics are useful to understand what's happening in the TS7700 subsystem at this moment; data provided by this type of statistic is a snapshot of the activity over the last 15-second interval; each new 15-second interval data overlays the previous interval's data;
- Historical statistics are useful to understand how you are using TS7700 resources; the data provided by this type of statistic is captured over a 15-minute interval; each new 15-minute interval data does not overlay the prior interval's data; 90 rolling days of historical statistics are kept in the TS7700 subsystem database.

Historical statistics are designed to support Performance Analysis and Capacity Planning activities so they are essential information to collect, trend and analyze.

In this paper we'll discuss the available historical statistics data types (records) focusing on the most important metrics from the performance analyst perspective.