

Buffer Pool Tool® vs. BPA4DB2 is there a Comparison?

Buffer Pool Tool

Buffer Pool Tool is the industry standard, and for a good reason. It has proven accuracy, and a client base that is willing to discuss their performance improvements and cost savings. It provides the most comprehensive set of performance information, and **predicts** the most important performance metric - **I/Os per second**.

Hit ratios, miss-ratios, and re-read rates are useless for measuring local buffer pool performance. They cannot be converted into cpu or elapsed time savings. Buffer Pool Tool provides the most comprehensive set of graphic analyses in the industry, and uses a statistical cluster analysis technique to show you which objects belong together based on the industry proven RAMOS and SAMOS methodology.

Buffer Pool Tool also reports upon, and **can predict**, the performance of Group Buffer Pools in Data Sharing environments. There is a vast difference between reporting, and predicting.

BPA4DB2

BPA4DB2 provides a very nice looking user interface, and nice graphics. Unfortunately, the data illustrated on their website does not provide a lot of substance, and some of their highlighted recommendations for threshold settings are blatantly incorrect. (i.e. Set vdwtq = 145 to avoid duplicate dasd writes within one minute). It has absolutely no basis for recommending pool size settings, since it does not predict the performance effect of any changes.

Where does BPA4DB2 get its data? It collects the 198 and 199 records from DB2. The 198 record shows every getpage and release page request for every object. The 199 record provides object summary information, at intervals specified in the DB2 zparm SYSPDPTIME parameter. This is just summary information, and not detailed data, so BPA4DB2 does not know anything whatsoever about the real I/Os taking place on your system. The 199 record shows synchronous and asynchronous average and maximum delay times, and overall I/O counts. It does not provide any information about List Prefetch, or Dynamic Prefetch activity. These usually have a significant impact upon overall pool and object performance. *For analysis purposes, it is not valid to group all types of asynchronous I/O events together.* Additionally, without knowing how many pages of objects have truly been read into a pool *by all types of prefetch activities*, it is impossible to perform any predictions of IO activity from altering pools or moving objects. The DB2 buffer manager has distinct LRU queues for managing random versus sequentially accessed pages. BPA4DB2 does not report anything about sequential pages physically read into a pool.

As another point, the 199 record only provides information for objects with an I/O rate greater than one per second. While this may sound like a good approach initially - only providing data for objects with significant I/O activity - in large systems with many thousands of objects, there may be thousands with very low I/O rates. But they are still being read, generating I/Os, displacing other pages from pools, and - as a sum - have a real impact on both pool and system performance.

Without simulating the interaction of all objects in a pool, at varying pool sizes, there cannot be any basis to recommend pool size settings; or have any idea about the effects of moving an object out of an existing pool, or into a new pool.

Summary

Buffer Pool Tool predicts I/O performance. All other products let you play guessing games with your system, pool, and object performance - can you take a chance?