

## Understanding the HFS and EXCP Counts in RMF and SDSF

The EXCP count for HFS datasets which are reported in SDSF and RMF is calculated differently than non-HFS EXCP counts. The EXCP count is used for a number of things. From a z/OS perspective, EXCP rates are used by WLM to compute I/O service units. For performance analysts this value is reviewed to understand I/O activity. Some customers often use this value for accounting or charge back purposes.

When IBM introduced the Hierarchical File System (HFS) in z/OS it created a new type of I/O environment. We now have a file system that can be shared by many users, and optimizes performance by caching file activity and often avoiding I/O. By exploiting these capabilities it is possible to allow more than one user to run the same application, reading/writing the same exact file(s), and depending on the activity of other users each user may see completely different amounts of I/O activity. One user may show virtually no I/O while another user may show large amounts of I/O.

From a performance perspective this may be a good thing. However, from the perspective of accounting and WLM management this can become an issue. For example a user who runs the same job two days in a row can be charged two different EXCP counts for each run. Since WLM uses EXCP counts as input to the I/O service units charged to a job, variations in the EXCP rate will be reflected in different service rates being applied to a job. For work in a multiple period service class it would mean the same exact work would transition periods differently based on HFS contention. Work in earlier periods typically have more stringent performance goals, and hence access to more resources. Changes in the amount of time a user is permitted in each period could impact throughput and response time of other workload.

The solution to this issue of variable HFS EXCP counts was to calculate EXCPs for HFS work similar to the way EXCPs are calculated for cached control units. Users are charged for all the I/O the physical file system would do if the data was not cached. Each 4k block which is read/written is charged one EXCP count. Directory lookups on the HFS are charged as well. While this approach produces repeatable results, the HFS does an excellent job of avoiding I/O, therefore the EXCP counts are much higher than the number of actual I/Os required.