


z/TPF File System Updates


Chris Filachek
TPF Development Lab

Database / TPFDF Subcommittee
March 11, 2014




Disclaimer

Any reference to future plans are for planning purposes only. IBM reserves the right to change those plans at its discretion. Any reliance on such a disclosure is solely at your own risk. IBM makes no commitment to provide additional information in the future.



Agenda

- File System Integrity Enhancements
- Performance and Operability
- Logical Record Cache




File System Integrity: Attributes

- Attributes are “name=value” pairs for files and directories
- File attributes are supported by
 - APIs: `tpf_getFileAttribute`, `tpf_setFileAttribute`, etc.
 - ZFILE ATTR command
- System attributes
 - Attributes with specific names recognized by the file system
 - Controls file system behavior with respect to the specific file or directory
 - Some current system attributes
 - shadow collections
 - dir size max, file size max
 - data record id, directory record id, index record id




File System Integrity: Directory Backup

- Backup collection attribute
 - New system attribute named “backup collection”
 - Only recognized on TFS directories
 - `zfile attr -s /etc 'backup collection=1'`
- When backup collection attribute is set
 - File system maintains backup of directory structure
 - Child directories inherit attribute from parent when directories are created
 - Does not backup files in the directory
 - Automatically uses backup directory if the primary directory structure is not accessible
- Only set attribute after all processors in the complex have the APAR applied
- PJ40897 (PUT 10)



File System Integrity: File System Check (FSCK)

- ZFILE FSCK Enhancements
 - ZFILE FSCK scans the TFS more completely
 - “Lost & Found” directory is now included in scanning and recovery
 - Improved error recovery in FSCK
 - PJ41768 (PUT 11)



File System Performance Improvements

- File system performance improvements
 - More efficient locking mechanisms and more intelligent use of locking
 - Improves performance for all file system types
 - TFS, PFS, FFS, MFS
 - PJ40937 (PUT 10)
 - PJ40944 (PUT 10)
- syslog performance improvement
 - Optimize use of file system APIs
 - PJ40578 (PUT 09)



File System Operability: Improved State File Management

- ZFILE State Files
 - Manages file system state for ZFILE command
 - Stored in the .tpfZfileState directory
- Original method
 - New state file created for every LNIATA
 - Most files were identical – same state information
 - Large number of files accumulate over time
- New method
 - Only create file when it differs from a “base” version
 - Reduces number of files and use of file system resources
 - PJ40967 (PUT 10)




Logical Record Cache

- What is a logical record cache?
 - Shared memory structure identified by a unique name
 - MDBF Subsystem shared
 - Processor unique or shared (requires CF)
 - Entries in each named cache
 - Identified by DBI, primary, & secondary keys
 - Referenced through hash table for efficient lookup
- Why use a logical record cache?
 - Reduce I/O by caching frequently referenced data in memory
 - Memory cache for frequently read data
 - Read-only data or high-read, seldom-updated data
 - Data that can be lost across an IPL and rebuilt

Logical Record Cache Application APIs

- Create a new logical record cache
 - newCache()
- Add a new cache entry or updates an existing entry
 - updateCacheEntry()
- Read a cache entry
 - readCacheEntry()
- Delete a cache entry
 - deleteCacheEntry()
- Flush the contents of the cache
 - flushCache()
- Delete a logical record cache from the processor
 - deleteCache()




Logical Record Cache Application APIs (cont.)

- Castout function
 - Application function that is called before an entry is cast out of the cache for any reason
 - Application defines castout function and passes function pointer to function to logical record cache APIs
- APIs supporting castout functions
 - `tpf_newCache_ext()`
 - `tpf_updateCacheEntry_ext()`
 - `tpf_deleteCacheEntry_ext()`



Logical Record Cache Support Enhancement

- Improved logical record cache hashing to reduce hashes to the same location
 - Reduces occurrence of long synonym chains
 - Improves performance of logical record cache when searching for cache entries
- See PJ41767 (PUT 11)



Trademarks

- IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.
- *(Include any special attribution statements as required – see Trademark guidelines on <https://w3-03.ibm.com/chq/legal/lis.nsf/lawdoc/5A84050DEC58FE31852576850074BB32?OpenDocument#Developing%20the%20Special%20Non-IBM%20Tr>)*

Notes

- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
- This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

