An ICF Catalog
White Paper
From Mainstar Software Corporation

MAINSTAR

ICF Catalog Caching –

Monitoring Performance

By Eileen A. McClintock

***Preface**: ICF catalogs must provide fast access to large amounts of data – therefore, high performance is essential!*

A great deal of performance tuning involves the elimination of I/O. To that end, catalog records that are directly retrieved can be cached in either the Catalog Address Space (CAS) or in a VLF managed data space.

The first White Paper in this series, ICF Catalog Caching Performance – ISC vs. VLF, explained the two types of caching and their implementation. This second White Paper continues the discussion of monitoring caching performance and provides recommendations for tuning performance.

***Catalog RecoveryPlus (CR+)**, Mainstar's software product for backup, recovery, diagnosis, and repair of ICF catalogs, provides convenient ISPF interface execution of MODIFY commands to assist in managing your catalog environment. This information can be used to monitor cache performance as well as many other aspects of CAS activity.*

Caching Catalog Records – Review

As individual catalog records are retrieved for data set locate operations, they can be cached either within CAS (In-Storage Catalog, also known as ISC), or in a VLF-managed data space (Catalog Data Space Cache, also known as CDSC). Once a data set's catalog records are in cache, future references to the data set can be processed using the information in the cache, thus eliminating I/Os to the BCS.

Caching involves reusing individual records that were previously read. With either method it is necessary to determine whether the contents of the cache are still accurate in a shared catalog environment.

A specialized VVR, referred to as the integrity or time-stamp VVR, is used by all systems to communicate changes made to the BCS. Enhanced Catalog Sharing (ECS) does not change how the communication of changes works. The integrity VVR is simply copied into a cache structure defined in the coupling facility, where sharing systems retrieve and update information, rather than in the VVDS.

Cache Performance Monitoring

The performance of cache can be monitored using the command:

```
MODIFY CATALOG,REPORT,CACHE
```

The resulting message IEC359I provides information on the number of searches that were satisfied with a cached record. A sample message display is shown below.

```
HIT%- RECORDS- SEARCHES- FOUND- DELETES- SHR UPD- PURGE
CRTSRP.TEST.OA10995                        (ISC)
78%      33        339      266      0       0       0
USERCAT.UCAT1                                (VLF)
98%      66     10,989    10,829   122       0      11
USERCAT.UCAT2                                (VLF)
99%     394     83,689    83,109    66       0       1
USERCAT.VIPLT91                              (ISC)
 6%      46      6,310       415     0       0       0
SYST80.MASTER.CATALOG                        (VLF)
61% 135,380     51,125    31,394    62      12       1
```

Information in the display includes:

- Catalog name and type of caching in effect.
- The percentage of searches that found a record in the cache.
- The number of records currently contained in the cache. If the catalog is using VLF, VLF may have trimmed records and the actual number cached may be less than this.
- The number of searches made against the catalog.
- The number of records found in the caches when searched.
- The number of records currently deleted in the cache.

- The number of updates made to the cache as a result of changes on other systems sharing this catalog.
- The number of times this cache was purged as a result of changes on other systems sharing this catalog.

The PURGE column is of particular interest. All cached records for a catalog are purged from ISC whenever any change has been made by a sharing system. A catalog using ISC with a low hit rate and large number of purges may be a good candidate for CDSC. All of a catalog's records are purged from CDSC when there have been a large number of updates since the system last referenced the catalog. VLF purges the cached records for that catalog rather than handle each change on the record level.

A rule of thumb says that the break-even point (where the overhead of maintaining the cache is less than the cost of doing the I/O to the catalog) is a 20% hit rate. If your system is not constrained in terms of the processor or storage, you may be able to accept a lower rate. If a catalog does not achieve the desired rate on a given system, it can be removed from caching with the command:

MODIFY CATALOG,NOISC(*bcsname***)**
or
MODIFY CATALOG,NOVLF(*bcsname***)**

Relying solely on data in the IEC359I message might be a mistake. Fields (other than PURGE) are reset when a purge occurs and all values are reset if a catalog is closed or in the event of a restart of VLF or CAS or an IPL. The statistics reflect activity since the last reset and may vary widely depending on the workload.

Cumulative and probably more realistic data on CDSC caching performance can be found in SMF Record Type 41 Subtype 3 – VLF Statistics. Records are written for all VLF classes. IGGCAS is the class name for catalog caching. Information in the Type 41 record includes:

- MAXVIRT specified in the COFVLF*xx* PARMLIB member, in 4K blocks.
- Amount of virtual storage currently being used, in 4K blocks.
- Number of times the cache was searched.
- Number of objects found in the cache.
- Number of objects added to the cache.
- Number of objects deleted from the cache.
- Number of objects trimmed from the cache.

VLF trims the least recently referenced records from cache once it has reached 90% of MAXVIRT. Be sure to increase that parameter when you add catalogs to the COFVLFxx member, keeping in mind that a large value may impact paging.

Master Catalogs

When CDSC was new, IBM explicitly said *not* to put the master catalog in the COFVLFxx list (in *MVS/DFP: Managing Catalogs,* SC26-4555). However, the current recommendation is to put it there. Some installations see poorer hit ratios in VLF than ISC, while others work wonderfully, as seen in the sample IEC359I message display shown above. Caching only applies to data set records. The catalog connector and alias tables are still in CAS. Searches in VLF or ISC would be for data sets cataloged in the master, or for data set names that have no alias related to a user catalog.

Tuning Recommendations

- Monitor performance for a catalog on each system. They will not all behave the same.
- Start by putting shared catalogs in the COFVLFxx member of PARMLIB and see how they perform using CDSC.
- If you find an individual catalog does not get a satisfactory hit rate, remove it from the COFVLFxx PARMLIB member on that system. It will revert to using ISC unless you also issue the command:
  MODIFY CATALOG,NOISC(*bcsname***)**
- If a catalog is *never* shared, put it on an unshared device, alter the share options to (3 3), and let it use ISC.
- Catalogs that are not frequently updated may use ISC cache effectively and not require the overhead of a VLF managed data space.
- If all updates to a given catalog are made from one system, it may use ISC cache well, while other systems use CDSC to take advantage of record-level invalidation.
- A catalog that is infrequently accessed by a given system may not be cached there at all. Consider putting the following command in the COMMNDxx member of PARMLIB:
  MODIFY CATALOG,NOISC(*bcsname***)**

Conclusion

Well-planned catalog caching can make a significant contribution to a high-performance catalog environment. Take advantage of it!

For more information, see the IBM publications, *DFSMS: Managing Catalogs* (SC26-7409) and *Enhanced Catalog Sharing and Management* (SG24-5594).

Eileen A. McClintock – Product Support Specialist

Eileen McClintock has a technical background in large-scale OS/390 systems. As a software instructor for 20+ years, she has presented courses on ICF catalogs, VSAM, MVS, and SMP/E. Eileen is a member of the team supporting Catalog RecoveryPlus.