An ICF Catalog
White Paper
From Mainstar Software Corporation

MAINSTAR

ICF Catalog Caching Performance –

ISC vs. VLF

By Eileen A. McClintock

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

A great deal of performance tuning has to do with eliminating 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.

This White Paper explains the two types of caching and their implementation.

***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.*

In-Storage Catalog Cache

The original – and default – method of caching is called In-Storage Catalog (ISC) Cache. With ISC, each catalog has a fixed amount of storage reserved within CAS that is used to cache individual records as they are retrieved for data set locate operations. Once a data set's catalog records are in ISC cache, future references to the data set can be processed using the information in the cache, thus eliminating I/Os to the BCS and VVDS. As the cached records for a catalog reach the maximum amount of allotted storage, the least-recently referenced record is deleted, allowing for a newly referenced record to be cached.

When you are using ISC, you need a mechanism to guarantee the accuracy of the cached information in a shared catalog environment. Each system has its own CAS, control blocks, buffers and cache for each catalog. VSAM SHAREOPTIONS(3 4) decreases the likelihood of obtaining incorrect information by causing CAS to read the catalog for each request rather than reuse the contents of a buffer to retrieve a different record. Caching involves reusing individual records that were previously read. It is still necessary to determine whether the contents of the cache are accurate.

VVDS-Mode Sharing

A specialized VVR, referred to as the integrity or time-stamp VVR, resides in the VVDS on the volume with the catalog and is used by all systems to communicate changes made to the BCS. At open time on each system, if a catalog is shared (SHR 3,4 and residing on shared DASD), data from the catalog's integrity VVR is copied into control blocks in CAS. When a system changes a catalog record, it notes the change in the VVR and updates a counter and timestamp there. When any system references the catalog, the timestamp and counter in the VVR are compared with the information in the CAS control blocks. If the values agree, there have been no updates since the last time the catalog was referenced and a record found in cache may be used. An I/O to the catalog itself is avoided.

If the timestamps and counters do not match, a change has occurred and some of the contents of cache may no longer be valid. For ISC, the entire cache for the catalog is released, and caching begins anew. An I/O to the catalog must be completed to locate the requested record.

Decreasing the number of I/Os to catalogs obviously improves performance. However, both ISC and the use of the integrity VVR ("VVDS-mode sharing") have their limitations. With ISC, every BCS has the same fixed amount of virtual storage allotted regardless of activity level. Whenever another system has made any change to a shared BCS, the entire cache is released. If a catalog does not utilize ISC well, you must

explicitly turn it off with the command:

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

ISC and VVDS-mode sharing traded one I/O for another – instead of an I/O to the actual catalog, there is now an I/O to the special VVR in the VVDS. In order to ensure the integrity of the information, that operation, whether for read or update, is accompanied by a RESERVE request. GRS can process it as a physical reserve or convert it to a SYSTEMS ENQ but there are obvious performance implications with either technique.

Virtual Lookaside Facility (VLF)

Alternatives to both ISC and VVDS-mode sharing have evolved. Beginning with MVS/ESA, a second type of caching became available to address the limitations of ISC. The Catalog Data Space Cache (CDSC) is managed by the Virtual Lookaside Facility (VLF) and is defined in the COFVLFxx member of SYS1.PARMLIB. Only BCSs named there are eligible for CDSC. To add or remove catalogs from the list, you must stop and restart VLF. CDSC can be dynamically turned on and off for those BCSs already named in COFVLFxx with the command:

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

All catalogs share the data space. When the data space is full, the least-recently referenced record will be deleted to make room for the new record. The default integrity checking prior to accessing a shared catalog utilizes the same integrity VVR checking as described above for ISC. However, for a catalog using CDSC, only individual records that have changed are released from cache. If the reference was for a dataset record that has not been changed, the cached record is used and an I/O to the BCS is avoided. CDSC space is only invalidated if, since the last access, more changes have been made than catalog management can track. (Although there is a limit, I haven't seen that number published.)

Enhanced Catalog Sharing

VVDS-mode sharing evolved into Enhanced Catalog Sharing (ECS), introduced with DFSMS/MVS 1.5. It does not change how the two caching techniques or communication of changes work. The integrity VVR is simply copied into a cache structure defined in the coupling facility when the catalog is first opened. Sharing systems retrieve and update information there rather than in the VVDS so I/Os to the VVDS and the associated reserves on the volume are eliminated.

For more information on ECS, refer to the Mainstar White Paper entitled, “ICF Catalog Shared Access Performance: Experiences with ECS.”

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.

The break-even point (where the overhead of maintaining the cache becomes less than the cost of doing the I/O to the catalog) is a 20% hit rate. If a catalog does not achieve that rate on a given system it can be removed from caching with the command:

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

Statistics are also provided on the number of records that were deleted from cache (VLF) or the number of times cache was purged (mainly ISC, occasionally VLF) because of updates on sharing systems.

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***)**

Conclusion

Monitor performance for each 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 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.
- To eliminate RESERVEs associated with VVDS-mode sharing, consider implementing ECS for catalogs shared only by systems connected to a coupling facility.

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.