

# VSAM/ICF Catalog

## White Paper

From Mainstar Software Corporation



## ICF Catalog Shared Access Performance: *Experiences with ECS*

By Ronald K. Ferguson

**Preface:** IBM introduced Enhanced Catalog Sharing (ECS) for ICF catalogs several years ago for DFSMS 1.5. As ICF catalog sharing became more prevalent in a sysplex environment, ECS was designed to improve performance for the integrity checking aspects of catalog management. For quite a number of system programmers who support this area of the operating system, this valuable feature has become a lightning rod, with enough FUD (fear, uncertainty, and doubt) surrounding it that many are now afraid to use it. Those who are successfully using it are quietly in the minority.

This technical paper is the result of two unofficial surveys that explored the use of ECS, and asked for anecdotal experiences – good or bad – about problems or performance benefits. The survey responses came from Mainstar's users of **Catalog RecoveryPlus (CR+)**, as well as volunteers from the IBM Main Listserv. Since it is a small subset of the OS/390 and z/OS user base, it may not be statistically accurate; nevertheless, the intent is to present a cross section of users and non-users of ECS, in the hope that this feature might become more widely used.

**CR+**, Mainstar's software product for backup, recovery, diagnostic, and repair of ICF catalogs, provides ISPF interface execution of MODIFY commands to help users identify Catalog Address Space (CAS) information about ICF catalog buffering and caching. This information can be used to determine performance benchmarks of catalog management.

### Catalog Management Improvements

The number of systems sharing catalogs has grown dramatically in recent years – leading to a corresponding growth in access

to data sets through those shared catalogs. The result is increased demand for better catalog performance, and better and more reliable catalog integrity.

In recent years, IBM has introduced several new features to improve ICF catalog performance and integrity. These include a common CAS to support all open catalogs within a system. (Prior to the introduction of CAS, each user's catalogs opened within their own private area.)

Two types of cache buffering were introduced: In-storage Catalog Cache (ISC) and Catalog Data Space Cache (CDSC). ISC is the default, and sets up its cache buffers within CAS itself. CDSC caches within a data space associated with CAS, and uses the Virtual Lookaside Facility (and therefore, CDSC is most commonly referred to by the acronym VLF). One of the most recent improvements is ECS.

### Background on ECS

ICF catalogs, since their introduction in 1983, have always supported shared system access. A catalog is eligible for shared system access when it is defined with SHR(3 4), resides on a shared DASD volume, and is "connected" to the master catalog on two or more MVS systems.

With shared catalogs, each CAS has its own cache, which creates a potential for integrity exposures. Each of the sharing CASs also has its own copy of the catalog's control blocks defining it. On each access to the catalog, CAS must ensure that its copy of the catalog's records in its buffers are current before it can attempt a lookaside, and it must know that its control blocks correctly reflect the physical catalog at that moment.

## **VVDS-Mode Sharing**

Since the original days of ICF catalogs, there's been an integrity mechanism – it never had a name before, but it's now called VVDS-mode sharing. Here's how it works.

When one of the sharing MVS systems makes an update to the catalog, the system timestamp value is stored in a "catalog integrity record" in the VVDS on the volume that the catalog resides on. That same timestamp value is kept in one of the catalog's control blocks in CAS on the system making the update. Before any sharing system attempts to use its cache buffers or control blocks for that catalog, catalog management on that system will read the "catalog integrity record" in the VVDS, comparing the last update timestamp in the record to the timestamp in its control blocks. If it's equal, it knows it was the last system to update the catalog, and therefore its buffers and control blocks are current. If it's unequal, it knows some other system was the last to update the catalogs, and therefore, it must flush its cache and update its catalog control blocks before accessing the catalog. To protect the catalog integrity record from being updated while the catalog's records are being processed, a RESERVE is issued against the volume containing the catalog.

The cost of this integrity management, while necessary, is quite expensive (in terms of performance) for any high-activity catalog.

## **ECS-Mode Sharing**

In 1999 IBM introduced ECS support in OS/390 V2R7. ECS doesn't change any of the basic ways that each CAS determines if its cache and catalog control blocks are current. Essentially, ECS caches the catalog integrity record in a Coupling Facility structure – which means that Parallel Sysplex is required for *all* systems sharing access to the catalog.

When an ECS-defined catalog is first opened by CAS, the catalog integrity record is read into the Coupling Facility (CF) structure from the VVDS. Subsequent systems that open the catalog "connect up" and utilize the copy of the catalog integrity record within the CF. Systems sharing the catalog do not have to read the catalog integrity record from the VVDS to detect changes, but rather, use

cross-memory services to compare their control block timestamp to the value in the CF. The RESERVE to the volume containing the catalog is also removed, because the CF architecture incorporates the necessary locking.

---

## **ECS in the Real World**

A search of the IBM Main Listserv on ECS produced 208 postings. When you peruse the posts, time and again you see a query, asking if anyone has experience with ECS. Typically, the post is then followed by someone who says they tried it, had some type of problem shortly thereafter, and immediately pulled their catalogs out of ECS. While others claim to have had absolutely no problems with ECS, it seems to be one of those system features that's hard to "see" what's going on, and even harder to know if it's working right and saving you time and money. For many, there is just an unnamed fear that it's somehow going to hurt more than help, and so they shy from ECS.

In Mainstar's ICF Catalog Seminars, I often ask the attendees, "Do you use ECS?" The responses are hard to quantify, but out of a typical group of 10 to 20 people, sometimes I hear 1 to 2 positive responses, and sometimes almost all of the attendees respond positively. It's very difficult to gauge the level of usage, and for the ones who do take advantage of ECS, it's difficult to know if it was deemed useful.

Over a two year period, we queried the ICF catalog development group at IBM Storage Systems Group (SSG) in San Jose to determine if they were seeing problems in the real world with ECS. In response to our first query in March 2003, we received the following response:

*We currently have one PTF (actually 3, one for each supported release) to correct a problem in ECS. We have no other problems open that can be attributed to ECS. We recommend customers implement ECS in stages - that is, they can implement it for a catalog at a time - it's not an "all or nothing" approach. And of course, since catalogs can be shared without ECS, it's strictly a customer choice*

*whether or not they want to use it (and how extensive the use should be). They'll have to make the decision based on their needs.*

In response to our second query, in September 2004, we received the response:

*We currently have no open problems for ECS. We have one customer who is having several rather strange but probably related problems with it but we have as yet been unable to diagnose the cause.*

## The Survey

To help us determine the extent that ECS is being used, and hopefully to identify problems that people have experienced, we sent a survey to **CR+** customers and posted it on the IBM Main Listserv.

Here is a summary of the survey:

1. Do you currently use ECS (yes or no)?
2. If yes,
  - How many of your user catalogs are under ECS?
  - How many of your master catalogs are under ECS?
  - What release of OS/390 or z/OS are you on?
3. If no,
  - Are you in a parallel sysplex environment?
  - Are you afraid of it?
  - Is it in the plan for future?
4. Have you used ECS in the past, but turned it off due to real or perceived problems?
5. Please provide a brief description of any problems you've encountered with ECS while it was turned on.

## Survey Results

| <b>Do Not Use ECS</b>                                    | <b>77%</b> |
|--|------------|
| No reason given  | 40%        |
| No plans to use ECS                                      | 12%        |
| Not at system level to use ECS                           | 2%         |
| Planning to implement ECS in future                      | 11%        |
| Tried ECS in the past, had problems, and discontinued it | 3%         |
| Have never tried ECS, and don't trust it                 | 8%         |
| Catalogs shared across sysplexes, so can't use ECS       | 6%         |
| Do not have a parallel sysplex environment               | 18%        |

| <b>Use ECS</b>  | <b>23%</b> |
|---|------------|
| Use ECS extensively, no problems to report  | 53%        |
| Use ECS extensively, only problem encountered was with TDMF during volume migrations                                  | 16%        |
| Use ECS extensively, only problems being with day-to-day management – re-orging and deleting catalogs while under ECS | 5%         |
| Use ECS on a limited basis, with catalogs that have known performance problems  | 5%         |
| Use ECS only on the master catalog  | 5%         |
| Use ECS only in test, and have plans to roll it into production   | 11%        |
| Tested ECS, found a problem that was possibly related to ECS, turned it off, and haven't tried it since               | 5%         |

## Conclusion

This was a limited survey; not large in numbers, given the total number of MVS sites around the world.

As one person mentioned in his response,

“ECS is still bleeding edge and I have discouraged our Storage Admin team from using it. Maybe your survey will convince me otherwise”.

Since ECS was defined over 5 years ago, it's interesting that anyone would still call it “bleeding edge”, but this may be an indicative opinion that is widely shared. Nevertheless, I don't know if the information in this White Paper will change those opinions, but I do hope it provides you with another data point that is more reliable than once-in-a-while anecdotes.

Excellent information about ECS can be found in the IBM Redbook, *Enhanced Catalog Sharing and Management*, SG24-5594-00. This Redbook was written in 1999, soon after ECS first became available, and hasn't been updated since. The information about how ECS works and how to implement it is still valid and very useful, and I would highly recommend this Redbook.

If you wish to query the “real world” yourself about their latest experiences, I would also recommend joining the IBM Main Listserv. You can subscribe by sending an e-mail to: [LISTSERV@BAMA.UA.EDU](mailto:LISTSERV@BAMA.UA.EDU). In the subject line, specify: SUBSCRIBE. In the text body, specify: SUBSCRIBE IBM-MAIN. The response e-mail will contain all necessary instructions to help you complete the subscription. Since the topic of ECS has been raised many times in the past, it would be wise to peruse the listserv archives back to 1999 before you submit your query.

**Ronald K. Ferguson – Founder, President, & CEO of Mainstar Software Corporation.**

*Ron Ferguson has a technical background in large-scale OS/390 systems. As a software instructor for 20+ years, he has presented over 600 courses on VSAM and ICF catalogs, and is recognized worldwide as an expert in these areas. Ron travels widely, meeting with customers and presenting at national and international conferences.*

Mainstar is a registered trademark of Mainstar Software Corporation. Catalog RecoveryPlus is a trademark of Mainstar Software Corporation. OS/390 is a registered trademark and Redbook and MVS are trademarks of IBM.