

VSAM Data Sets and ICF Catalogs with IMBED and REPLICATE Attributes

By Ron Ferguson

Recently, we became aware of a discussion in the z/OS user community around IMBED and REPLICATE attributes. As an IBM Business Partner since 1996, we would like to clarify the situation and update you on our support within Catalog RecoveryPlus (CR+) for this requirement.

Historically, IMBED and REPLICATE were valuable performance tuning parameters for VSAM data sets and ICF catalogs. However, as a result of advances in cache memory in today's disk storage devices, IMBED and REPLICATE have become performance problems.

For several years, IBM has recommended the removal of these attributes and Mainstar supports this recommendation. At DFSMS 1.5 release level, IBM modified Access Method Services (IDCAMS) code so that you could no longer define new VSAM data sets or catalogs with IMBED or REPLICATE. However, any existing data sets or catalogs with either attribute continue with the attribute until the next reorg – these are not 'ALTER-able' attributes, so a physical DELETE/DEFINE of the data set or catalog is necessary to remove them. For application data sets, this typically wouldn't be a problem, as they tend to be reorged quite frequently. For static data sets or archival copies of data sets that still exist, this could be a problem at some point in the future. ICF catalogs fall into this latter category, as many installations haven't reorged some or all of their catalogs in years (dating back to the early 1990s in many situations).

On August 10, 2004, IBM released a z/OS and z/OS.e Statement of Direction that IMBED and REPLICATE will not be supported at some point in the future (http://www-1.ibm.com/servers/eserver/zseries/zos/zos_sods.html). After describing actions that will be effective in 2004 and 2005, IBM explains their plans for future releases.

IBM plans to take the following actions in a future release:

- Support for the VSAM IMBED, REPLICATE, and KEYRANGE attributes will be withdrawn. No supported release of z/OS or OS/390 allows you to define new VSAM data sets with these attributes. Using them for existing data sets can waste DASD space and can often degrade performance. When this support is withdrawn, you will not be able to process data sets with these attributes.

While removing the IMBED and REPLICATE attributes is a vital task, customers can complete this task within a planned timeframe that is not determined by the release of z/OS 1.7. Mainstar will strive to stay up to date with changes IBM implements and consequently have allowed for it in our CR+ development plan.

- For ICF catalogs, you can remove IMBED and REPLICATE *automatically*. **CR+** REORG While Open removes IMBED and REPLICATE on any catalog it encounters. This capability is available via PTF right now and will be in the upcoming release of **CR+**.
- For VSAM data sets, **CR+** EXPLORE and REPORT quickly identify and report on VSAM data

VSAM Data Sets and ICF Catalogs with IMBED and REPLICATE Attributes

sets with IMBED or REPLICATE. Using this report, it is then quite easy to remove the attributes.

If you have any questions regarding this topic, please attend our webcast at 11:00 AM Pacific Time on Thursday, February 17th. We will also offer it on at 11:00 AM Pacific Time on Thursday, March 10th. Register online at: <https://www.mainstar.com/events/onlinepres/webcastreg/index.asp>.

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 and Catalog RecoveryPlus is a trademark of Mainstar Software Corporation.

IBM is a registered trademark of the International Business Machines Corporation.