

Global Areas

Adabas Parallel Services uses global areas in dataspace for cluster caching and locking functions. Under z/OS version 1 release 5 or later releases, cache can also be maintained in shared 64-bit addressable virtual storage.

On BS2000/OSD systems, the global areas are also in dataspace and will restrict this operation to ESA machines for the time being.

The global cache area ensures that current data is available to all nuclei in a cluster. It helps keep Associator and Data Storage blocks in the local buffer pools up to date.

The global lock area is used to protect the resources needed during command execution against conflicting use by multiple cluster nuclei.

Global cache and lock areas are sized in each nucleus using the ADARUN parameters CLUCACHESIZE and CLULOCKSIZE, respectively.

The first cluster nucleus that starts provides its CLUCACHE_{xxxx} and CLULOCK_{xxxx} parameters to ADACOM, which uses a subtask to allocate the global areas for the cluster in which the nucleus participates. In addition, ADACOM dynamically allocates a data set to each cluster for cluster-related message output if the file is not defined in the ADACOM startup JCL.

Note:

Read *Performance and Tuning* for a formula for estimating cache and lock size.

ADACOM maintains the global areas. It refuses to terminate normally as long as it owns any global areas, because doing so would cause all active nuclei in all clusters using these global areas to fail.

For each Parallel Services cluster, ADACOM prints dataspace-related messages to an output data set/file with the DD name/link name *Dssddddd*, where *ss* is the last two digits of the SVC number and *dddd* is the DBID. (On BS2000, the *ss* number is derived for each separate IDTNAME as the first available, counting backward from 255.) On z/OS systems, ADACOM automatically allocates this data set in the spool with SYSOUT=X, if it is not explicitly specified.

The global areas are not persistent; that is, they disappear when the last active nucleus terminates, whether normally or abnormally. No manual cleanup is required.

This chapter covers the following topics:

- Global Cache Area and Manager
- Global Lock Area and Manager

Global Cache Area and Manager

A global cache area is allocated to each Adabas Parallel Services cluster for ASSO and DATA blocks that have been updated during the session.

This section covers the following topics:

- Local Buffer Pool and Manager
- Global Cache and Manager
- Buffer Flush
- Global Cache Storage Options

Local Buffer Pool and Manager

Every nucleus in an Adabas cluster has a local buffer pool and manager.

The buffer pool manager oversees all nucleus requests for reading and writing Associator and Data Storage blocks. For each block in its local buffer pool, the buffer pool manager:

- registers its interest in the block with the global cache manager;
- checks the global cache manager for the status of the registered block to ensure that its nucleus always has the most current copy of the block; and
- writes changed blocks to the global cache area.

The size of the local buffer pool of each nucleus is determined by the ADARUN parameter LBP. It must be large enough to hold the active working set of database blocks being used by the nucleus at any one time.

Global Cache and Manager

The global cache area must be large enough to hold the active working set of blocks from the local buffer pool of each nucleus in the cluster. The global cache manager oversees all requests for Associator and Data Storage blocks and copies changed blocks between the local buffer pools and the global cache area to maintain data integrity.

When a nucleus requests a block, it checks the global cache area as well as its local buffer pool to locate the block. If an up-to-date copy of the block is already in the local buffer pool, it is used straight away. Otherwise, if the block is in the global cache area, it is copied to the local buffer pool. If neither is the case, the block is read in from the database. In any case, the global cache manager keeps track of the existence of the block in the local buffer pool and invalidates the local copy if another nucleus updates the same block.

The global cache manager also handles the deletion of blocks in the global cache area when it becomes necessary to reclaim the space they occupy.

Buffer Flush

Any active nucleus may perform buffer flushes.

The buffer flush accommodates the fact that all updated blocks are located in the global cache area. From the global cache area, modified blocks are "cast out" to the flush I/O pool (FIOP) buffer before they are written to disk. The FIOP buffer is sized using the ADARUN LFIOP parameter, and the frequency of buffer flushing depends on the limit set. Until such blocks are written to the database, the global cache

area holds more current information than the database.

Global Cache Storage Options

Global cache data can be maintained in a dataspace. On z/OS systems running version 1 release 5 or later, global cache data can also be maintained in shared 64-bit addressable virtual storage. On z/OS systems running version 1 release 9 or later, global cache data can also be maintained in shared 64-bit virtual storage that is backed by large pages (provided the large page pool has been configured to a sufficient size and is available in the system).

Note:

Virtual 64-bit storage backed by large pages can only be used on IBM z10 mainframes running z/OS Version 1.9 or above and for which IBM large page support has been enabled via IBM APAR OA20902. You must also allocate the size of the large page pool (use the LFAREA parameter in the IEASYSxx member of SYS1.PARMLIB).

When a dataspace is used for global cache, the maximum size is 2 GB, a limit imposed by the operating system. Using shared 64-bit addressable virtual storage removes this virtual storage constraint and extends the maximum size of the cache from 2 GB to tens of GB.

Use the CLUCACHETYPE parameter of ADARUN to specify the virtual storage type for the global cache. The default is "DSP", indicating a dataspace of the size specified by the CLUCACHESIZE parameter of ADARUN will be used for both control structures and cached data.

The other two values (valid only on z/OS systems) for the CLUCACHETYPE parameter are:

- "V64" (on z/OS v1.5 or later), which indicates that the CLUCACHESIZE parameter should specify the amount of shared 64-bit virtual storage that will be used for both control structures and cached data.
- "L64" (on z/OS v1.9 or later), which indicates that the CLUCACHESIZE parameter should specify the amount of shared 64-bit virtual storage, backed by large pages, that will be used for both control structures and cached data.

To use the 64-bit global cache, your systems programmer must enable shared 64-bit virtual storage in SYS1.PARMLIB. To use large pages, your systems programmer must enable large pages in SYS1.PARMLIB.

Global Lock Area and Manager

Each Adabas Parallel Services cluster uses a global lock area to manage the setting, status, and release of various locks imposed during multiple update nucleus processing. The global lock manager synchronizes the nuclei, users, and transaction processing to ensure data integrity.

The global lock manager is used to:

- connect to and disconnect from the global lock area;
- obtain (conditional or unconditional), release, and alter the ownership (shared or exclusive) of resource locks; and

- read recovery information about a failed peer nucleus.

Lock manager calls may be asynchronous, meaning that the nucleus may continue processing in other threads before a call has completed.

If a lock request is conditional, it is rejected if the lock is not free; if the lock request is unconditional, the nucleus thread waits until the requested lock is free before continuing.

In general, locks are used to prevent two cluster nuclei from using the same resource at the same time. Such resources include:

- data records, which are protected by hold queue element (HQE) locks;
- unique descriptor values, which are protected by unique descriptor element (UQDE) locks;
- end-transaction IDs, which are protected by the ETID lock; and
- various other single-instance resources.