Performance and Tuning

This chapter covers the following topics:

- ADARUN Parameter Settings
- Allocating Work Data Set Space
- Using Close (CL) Commands
- Deferred Publishing
- Tuning Buffer Flushes
- Optimizing Global Cache and Lock Areas
- Minimizing Communication with Global Areas
- Optimizing Block Sizes

ADARUN Parameter Settings

Software AG recommends that you use your existing Adabas ADARUN parameters (or the default values) for each nucleus in an Adabas cluster, and then tune the values after analyzing the performance of the cluster.

Session statistics can be used to determine the best settings for each parameter. The statistics can be displayed using operator commands during the session; the statistics are also printed automatically at the end of a session or in response to an ADADBS REFRESHSTATS command.

For parameters that allot processing resources to the cluster nuclei (such as NU, NH, LP, etc.), Software AG recommends that you set them large enough that each individual cluster nucleus could handle the entire load on the database if the other nuclei were to terminate abnormally.

Allocating Work Data Set Space

Each Adabas cluster nucleus requires its own Work data set to hold its temporary data.

The individual sizes of the different Work parts (1, 2, and 3) as specified by ADARUN parameters such as LP and LWKP2 can be different among the nuclei; however, the overall size of each Work data set must be the same, because the total Work size is stored in the Adabas general control block (GCB). Software AG recommends that you use the same LP and LWKP2 values on each nucleus active for the same database.

For each nucleus, you need to specify shared access to DD/WORKR1. During an offline or online recovery, a nucleus may access the Work data sets belonging to other nuclei in the cluster.

Using Close (CL) Commands

Users are assigned to a nucleus for their entire sessions and should therefore issue Adabas close (CL) commands as appropriate. The close command ends the user's session, making the user eligible for reassignment to another nucleus when the user again issues an Adabas open (OP) command. This allows Adabas Parallel Services to rebalance the workload over the participating nuclei.

Deferred Publishing

Publication of updated blocks to the global cache area is usually deferred until just before the end of the associated transaction. Multiple updates to a block may produce only a single write of the block to the cache rather than a cache write for each update.

The greater the number of database update in parallel transactions, the greater the expected improvement in performance.

Note:
Deferred publishing creates an asymmetry between users on the update nucleus, who see uncommitted updates (unless they read with hold), and users on other cluster nuclei, who may or may not see uncommitted updates.

- Redo Pool
- ADARUN Parameter LRDP

Redo Pool

Since the write of updated blocks to the cache may fail due to conflicting updates to the same blocks by other nuclei in the cluster, every cluster nucleus must be capable of redoing the updates it has not yet written to the cache. The nucleus maintains information about these updates in the "redo pool".

ADARUN Parameter LRDP

The size of the redo pool is specified by the new ADARUN parameter LRDP. The LRDP parameter is effective only in a cluster nucleus, that is, when a nonzero NUCID is specified.

If LRDP is not specified, the nucleus takes as default the value of the LFIOP parameter. If LRDP is explicitly set to zero, the nucleus writes each update immediately to the cache.

Different nuclei in the same cluster can have different settings of LRDP. It is also possible, although not recommended, to run one nucleus with LRDP=0 and a peer nucleus with LRDP>0.

Note:
If one nucleus runs with LRDP=0 and a peer nucleus runs with LRDP>0 and the different cluster nuclei concurrently update the same Data Storage blocks, incorrect DSST entries may be produced. These are reported by ADADCK. Such errors are harmless and do not affect the results of the application programs.

The nucleus reports on the use (high watermark) of the redo pool in a shutdown statistic and in the response to the DRES command from the operator console or from ADADBS OPERCOM.

Tuning Buffer Flushes

When the update load on the database is so high that the buffer flush becomes the bottleneck, you can improve performance by reducing the duration of buffer flushes.

Instead of starting one I/O per volume, a buffer flush can initially start a predetermined number of I/Os on each volume and then start a new one when another I/O on the same volume finishes. This occurs independently on each volume.

The ADARUN parameters LFIOP and FMXIO (see the *Adabas Operations* documentation for details) can be used to control buffer flushes. The LFIOP parameter enables asynchronous buffer flush operation and sets the I/O pool size. The FMXIO parameter sets the limit on the number of I/O operations that can be started in parallel by LFIOP flush processing.

- Effect of ASYTVS Parameter Setting
- Dynamically Modifying the FMXIO Parameter Setting

Effect of ASYTVS Parameter Setting

The meaning of the FMXIO parameter is affected by the setting of the ASYTVS parameter:

When ASYTVS=YES (buffer flushes occur by volume), FMXIO specifies the number of I/Os to be started in parallel *on each volume*. The minimum and default number is 1; the maximum number is 16. If you specify a number greater than 16, it is reduced to 16 without returning a message.

When ASYTVS=NO (buffer flushes occur in ascending RABN sequence without regard to the distribution of the blocks over volumes), the minimum, default, and maximum values continue to be 1, 60, and 100, respectively.

Dynamically Modifying the FMXIO Parameter Setting

The setting of FMXIO can be modified dynamically using the FMXIO=nn command from the operator console or the Modify Parameter function of Adabas Online System.

Optimizing Global Cache and Lock Areas

As a user, you must allocate and define sizes that are appropriate to your application needs for the global cache and lock areas.

This section provides guidelines for determining optimal sizes for these areas based on current experience.

Note:
There may be sites for which these guidelines are not appropriate.

- Global Cache Area Size
- Global Lock Area Size

Global Cache Area Size

The global cache area must be large enough to retain:

- *directory elements* for all blocks that reside in all the buffer pools; and
- enough data elements to keep the changed blocks between buffer flushes (cast-outs).

Directory elements are used to keep track of the cluster members that have a particular block in their buffer pools so that the block can be invalidated should any member modify it.

If the number of directory elements is insufficient, Adabas Parallel Services reuses existing directory elements and invalidates the blocks associated with those directory elements, because they can no longer be tracked. These blocks must then be reread from the database and registered again the next time they are referenced and validated, even though they did not change.

It is generally better to reassign storage for data elements to keep more ASSO and DATA blocks in the global cache area than to define too many directory elements in the global cache area. More data elements than necessary can be used to keep additional blocks to improve the local buffer efficiency.

The number of directory elements need not be greater than the sum of the sizes of all buffer pools divided by the smallest block size in use for ASSO and DATA.

When connecting to the global cache area during startup, the ADAX57 message reports the number of directory elements and data elements. The ADARUN parameters DIRRATIO and ELEMENTRATIO determine the ratio between the number of directory and data elements.

Global Lock Area Size

All nuclei in a database cluster share the global lock area.

Locks are held for a variety of entities, for example unique descriptor values. These lock types tend to occur with very different frequencies. The amount of lock activity during a session for each lock type is displayed in the shutdown statistics.

It is often the case that ISN locks show the greatest activity. The sum of high-water marks for NH yields an upper limit for the number of ISN locks that were held concurrently during the session.

The global lock manager uses a hash table to allocate and find a specific lock entry.

When the global lock manager receives a lock request (for example, to put an ISN of a file into hold status), it allocates a specific lock entry unless another member of the cluster has already made a conflicting allocation. A conflicting allocation produces lock contention because another member holds the same lock. Depending on its type, the lock request is then rejected or remains pending, waiting for the associated resource to become available.

The minimum lock structure size can be roughly estimated as:

```
(NU*3 + NH + LDEUQP/16 + MAXFILES*4 + 50) * 240 +  500,000 bytes
```

where MAXFILES is the maximum number of files in the database (set in ADADEF or ADAORD) and NU, NH, and LDEUQP are the ADARUN parameters of the cluster nuclei. The formula in parentheses `(NU*3 + NH + LDEUQP/16 + MAXFILES*4 + 50)` is used to calculate the minimum number of

lock records that the cluster nuclei expect to have available.

Minimizing Communication with Global Areas

Most of the additional processing required for Adabas Parallel Services environments compared to a single Adabas nucleus involves communication with the global areas.

For this reason, optimizing the performance of an Adabas Parallel Services environment means minimizing the need for communication with the global areas. It is also important to keep the time required for each communication as short as possible.

- Avoiding the Hold Option
- Reducing Direct Interaction with the Global Cache Area

Avoiding the Hold Option

Lock requests usually depend on application requirements. Under data-sharing, the hold option is more expensive and access with the hold option should be avoided unless records will in fact be updated or must be protected from concurrent updates.

Reducing Direct Interaction with the Global Cache Area

Cache area requests occur when blocks:

- that are referenced do not exist in the local buffer pool;
- exist in the local buffer pool but have become invalid due to concurrent updates by other cluster members or from directory reuse; or
- are updated.

The first and second situation require registering and (re)reading the blocks from the global cache area. This is more expensive than just validating blocks.

The first situation is related to the buffer efficiency in a noncluster environment. In a cluster environment, buffer efficiency represents the combined effect of the local buffer pool and the global cache area. In order to reduce the interaction with the global cache, the local buffer pool (LBP) should not be decreased from what would be used in a noncluster nucleus. A large LBP parameter and the usage of forward index compression are recommended to improve the buffer efficiency in the local buffer pool.

Optimizing Block Sizes

Although earlier versions of Adabas often worked well with large block sizes, the buffer pool manager and forward index compression features introduced with Adabas version 7 make smaller block sizes more attractive, especially in data-sharing mode.

Use the following guidelines when selecting an optimal block size for ASSO and DATA:

Note:
Only general recommendations can be given.

1. Avoid 4-byte RABNs

   If the database is not extremely large, avoid 4-byte RABNs as this increases the number of AC blocks by 33%. When growth considerations are taken into account, this may require larger block sizes or limit reductions in block size. The same holds true for the maximum compressed record length.

2. Use forward index compression

   Forward index compression can significantly reduce the number of index blocks in a database. Apply forward index compression to all frequently accessed files (or to all files, regardless of their frequency of use). Choose the ASSO block size that is as small as possible but large enough to keep the number of index levels down to 3 or 4.

3. Minimize frequently updated descriptors

   When files are updated frequently, the number of blocks that are modified and need to be written to the global cache area often depends on the number of descriptors that have been defined and modified during update processing. Support for additional keys whose descriptor values are subject to frequent modifications becomes even more expensive in a data-sharing environment.