software AG

Adabas Parallel Services

Glossary

Version 7.5.1

September 2009

Adabas Parallel Services

Table of Contents

Glossary

Glossary entries exist beginning with the following letters:

D

database, single physical	A single physical database is one set of Associator and Data Storage data sets identified by a single database ID number (DBID).
	Using Adabas Parallel Services, a single physical database is accessed simultaneously by up to 31 Adabas nuclei running under a single operating system image.

E

ECSA	Extended Common Service Area (ECSA) is a virtual storage area above the 16MB line that can be accessed from all address spaces in the operating system image.

G

global areas	Adabas Parallel Services uses global areas in shared data spaces to ensure that current data is available to all nuclei.
	The global areas (data spaces) are accessible by all nuclei in a Parallel Services cluster and are used to synchronize their processing and to share data.
global cache area	Each Adabas Parallel Services cluster uses a global cache area to hold ASSO/DATA blocks that have been updated during the session. The cache area provides the cluster nuclei with up-to-date block images to ensure data integrity.
	The size of the global cache area is determined by the ADARUN `CLUCACHESIZE` parameter setting of the first cluster nucleus to start.

global lock area	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 size of the global lock area is determined by the ADARUN `CLULOCKSIZE` parameter setting of the first cluster nucleus to start.

I

initialization module (ADACOM)	The ADACOM task is the first component of an Adabas Parallel Services cluster to be started and the last component to be stopped. ADACOM sets up the Adabas Parallel Services environment on an operating system image. Parameters of the ADACOM control program define the characteristics of any number of Adabas Parallel Services clusters.
	ADACOM must be active on the image before any Adabas Parallel Services nucleus can be started.
	ADACOM creates and owns the global areas. After initialization, it is responsible for holding the global areas and performing operator commands.
local buffer pool	Each Adabas Parallel Services nucleus has its own local buffer pool, the size of which is determined by the ADARUN parameter `LBP`.

N

NUCID	An ADARUN parameter that identifies a particular Adabas nucleus as a cluster nucleus.
nucleus cluster	Adabas Parallel Services allows a maximum of 31 nuclei in one cluster. All nuclei in the cluster work on the same database. There is no predefined limit on the number of Adabas Parallel Services clusters that can be used.

P

PPT	The parallel participant table (PPT) is an area of 32 blocks in the database's Associator. It is used to track the Work, PLOG, and CLOG data sets/files used by the nuclei running in a cluster (as well as by a noncluster nucleus).

Each cluster nucleus reserves one PPT block for itself and records in this block the names of the Work, PLOG, and CLOG data sets/files that it uses, together with pertinent status indicators.

Individual cluster nuclei use the PPT information to access the Work data sets/files of the peer nuclei when performing recovery after a nucleus failure.

The ADARES utility uses the PPT information to access the PLOG or CLOG data sets/files of the cluster nuclei when performing the automated PLOG or CLOG merge, respectively.

Furthermore, cluster and noncluster nuclei use the PPT information to warn about changes in the Work or PLOG JCL, if the previously used Work or PLOG data sets/files still contain data.

S

SVCCLU

As a component of the Adabas SVC/Router, SVCCLU routes commands to the cluster nuclei. It uses the nucleus and user tables in common storage, which are updated based on nucleus or ADACOM information, to decide to which nucleus to assign a new user and route the user's commands.

There is no predefined limit on the number of Adabas clusters, each servicing a separate database, that can run under a single SVC/Router.