What is Adabas Parallel Services?

Adabas Parallel Services implements multinucleus, multithread parallel processing and optimizes Adabas in a multiple-engine processor environment on a single operating system image.

Up to 31 Adabas nuclei in an Adabas Parallel Services cluster are distributed over the multiple engines provided by the system.

All nuclei in the cluster access a single physical database simultaneously. A "single physical database" is one set of Associator and Data Storage data sets identified by a single database ID number (DBID). Each nucleus writes its protection data to its own private work data set and protection logs (PLOGs). Each nucleus also records the commands it executes in its own command logs (CLOGs).

The nuclei communicate and cooperate with each other to process the users’ work. Compression, decompression, format buffer translation, sorting, retrieving, searching, and updating operations can all occur in parallel.

In addition to the increased throughput that results from parallel processing, Adabas Parallel Services increases database availability during planned or unplanned outages: the database can remain available when a particular cluster nucleus requires maintenance or goes down unexpectedly.

A practically unlimited number of Adabas Parallel Services clusters can operate in the same operating system image under the same or different SVCs or IDTNAMEs; that is, a practically unlimited number of separate databases can be processed, each with its own Adabas Parallel Services cluster of up to 31 nuclei.

Applications see only one database target; no interface changes are required. Applications still communicate with their intended databases and communicate with an Adabas Parallel Services cluster of nuclei without modification.