

Adabas Vista

Glossary

Version 7.4.2

September 2009

This document applies to Adabas Vista Version 7.4.2 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © Software AG 2009. All rights reserved.

The name Software AG, webMethods and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA, Inc. Other company and product names mentioned herein may be trademarks of their respective owners.

Table of Contents

Glossary	1
A	1
C	1
D	1
F	2
M	2
N	2
P	3
R	4
S	4
T	4
U	5

Glossary

Glossary entries exist beginning with the following letters:

A

Adabas Vista ISN For an Adabas Vista partitioned file, all Adabas ISNs are modified with a partition identifier before returning the ISN to the application. Such a modified ISN is referred to as an Adabas Vista ISN. Correspondingly, any ISN received from the application program for such a file is interpreted as an Adabas Vista ISN.

C

collating sequence Adabas Vista processes the sequence of data in a file once it is partitioned to ensure that the user always receives the data in the correct sequence across partitions.

context During application logon, the job parameters `Profile ID` and `Target Category` define the client's current context. These parameters may be modified dynamically by use of APIs. The current context may also be temporarily suspended, to enable processing in a different context, and then resumed as required by the application.

D

distributed access Adabas Vista provides access using data that is not related to a partitioning field. Access by field value cannot be focused onto one particular partition. Adabas Vista automatically distributes the access to all partitions before collating the results in the desired and expected order.

distributed lock mode The distributed lock mode facility provides greater control of record hold logic in a partitioned environment. It applies when a hold-based `READ (L6)` or `FIND (S4)` is issued against a partitioned file.

F

file partition A partitioned file comprises one or more physical partitions. Each partition is an Adabas file in its own right identified by a database and file number that tell Adabas Vista where to direct a command in order to satisfy a request for data. A partition is uniquely identified by a Partition ID.

Even while using the single partitioned file image, you may choose to have certain programs access a partition directly by its real file number. See [Mixed-mode Access](#).

file translation File numbers in Adabas applications can be changed dynamically at runtime based on translation information supplied to Adabas Vista.

focused access Applications generally access the file using search data that is based, in whole or in part, on the primary key field. If such a key field has been used to partition data, Adabas Vista is able to focus on the specific partition(s) needed to satisfy the access.

M

mixed mode access You may choose the applications for which Adabas Vista provides a single file image for all the partitions. Additionally, you may choose to have certain programs access a partition directly by its real file number, even while using the single file image. This is referred to as mixed mode access.

multipart (partition sharing) The partition sharing feature (also referred to as multipart) allows multiple partitions of the same Adabas Vista file to share the same Adabas file while preserving the collating sequence of the single file image. In this way, specific portions of a file can be split from the main file with minimal data manipulation.

N

navigation Generally, the main overhead for partitioning is monitoring database requests to determine whether the target is a partitioned file and if so, which partition(s) needs to be used. Overall, this is referred to as navigation.

normal placement Normal placement occurs when the value for the partitioning field is extracted from the Adabas record buffer and used to direct a new record to the correct partition for storage.

P

Partition ID	The Partition ID (PartID) uniquely identifies a partition and, together with the Adabas ISN, forms the Adabas Vista ISN.
partition outage	<p>Partition outage makes it possible to take a partition offline. Processing continues as required for the other partitions, but data in the offline partition is unavailable until returned online. This allows you to manage and mostly eliminate the various planned outages your operation requires.</p> <p>Tolerance levels for partition availability are provided. Application access to the partitioned file can be maintained transparently even though all the partitions of the file may not be available.</p>
partition restriction	Partition restriction provides access control for each partition of a partitioned file. It allows you to 'hide' partitions even though the data is available. You can implement this feature unilaterally or you can allow applications to override it for each user at runtime.
partition sharing (multipart)	The partition sharing feature (also referred to as multipart) allows multiple partitions of the same Adabas Vista file to share the same Adabas file while preserving the collating sequence of the single file image. In this way, specific portions of a file can be split from the main file with minimal data manipulation.
partitioned file	<p>The data in a single file image is partitioned into one or more physical files based on a partitioning field. A partitioned file is uniquely identified by a source database and file number. In Natural systems, this is the database and file number defined to the DDM.</p> <p>Each partitioned file definition may a published version or a draft version. Only the published version is executed; only the draft version is maintained.</p>
partitioning field	The Adabas field used to partition the data.
partitioning field high value	The highest value of the partitioning field that can exist in the appropriate partition. The value must be specified in accordance with the defined partitioning field's format and length.
profile ID	A specific client's Profile ID supplied to Adabas Vista, either through the job parameter or via API calls, combined with the Profile ID defined to a translation rule or a partitioned file definition, enables the calling application to influence file translation dynamically.

For example, a Profile ID for a translation rule may be for accounts or manufacturing. Multiple clients from within the same application can access either the accounts or the manufacturing data based on matching the client Profile ID and the translation rule Profile ID. See also [target category](#).

R

read-ahead feature The Adabas Vista read-ahead feature offers the benefits of multifetch for batch jobs in an Adabas Vista environment, thus reducing the amount of interprocess communication and command selection within the database.

Values for the job parameters `Read-ahead Pool Size` and `Read-ahead Process` define the maximum amount of memory and the type of read-ahead processing to be used, respectively.

S

store control When storing a record, the value for the partitioning field is extracted from the Adabas record buffer and used to direct the new record to the correct partition: this is termed normal placement.

For store operations that do not specify a partitioning field, or do and have null value for it, the `Store Control Option` may be used to direct the record to a relevant partition. It defines the placement of new records according to the presence, absence, or value of the partitioning field within the Adabas format/record buffer.

T

target category Each client has a target category supplied to Adabas Vista, either through the job parameter or via API calls. This gives the client a subset of translation rules and partitioned file definitions that influence file access, which may be changed dynamically.

For example, a translation rule may exist for target categories `test` or `production`. Multiple clients from within the same application can access either the test or production data based on matching the client target category and the translation rule target category. See also [Profile ID](#).

translation rule A translation rule is used to re-direct a command from a source environment to a target environment. It has two versions: a published

version and a draft version. Only the published version is executed; only the draft version is maintained.

U

user partition concurrency The number of concurrent Adabas Command ID sequences that a user may have outstanding on each partition.
