

Adabas Vista

Adabas Vista Introduction

Version 8.2.2

November 2016

This document applies to Adabas Vista Version 8.2.2.

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

Copyright © 2016 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

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

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Use, reproduction, transfer, publication or disclosure is prohibited except as specifically provided for in your License Agreement with Software AG.

Document ID: AVI-INTRO-822-20161128

Table of Contents

Preface	v
1 File Partitioning	1
When is Partitioning Required?	2
Implementing Partitioning within Application Systems	2
Implementing Partitioning using Adabas Vista	3
Typical Partitioning Deployment Scenarios	5
Accessing Partitioned Files	7
Additional Partitioning Features	9
2 File Translation	15
Standard File Translation	16
Target Categories	16
Pages	17
3 Adabas Vista Tools	19
File Analysis (AVILOOK)	20
Online Services (SYSAVI)	21
4 Cluster Support	23

Preface

This document describes the benefits, concepts and features of Adabas Vista.

Adabas Vista:

- partitions data without increasing the complexity of client applications;
- translates access to target files without the need to change application programs;
- handles large volumes of data with increased flexibility and availability;
- provides for rapid growth by permitting the addition of new partitions or splitting partitions that have grown too large.

The following topics are provided:

File Partitioning

File Translation

Adabas Vista Tools

Cluster Support

1 File Partitioning

■ When is Partitioning Required?	2
■ Implementing Partitioning within Application Systems	2
■ Implementing Partitioning using Adabas Vista	3
■ Typical Partitioning Deployment Scenarios	5
■ Accessing Partitioned Files	7
■ Additional Partitioning Features	9

This section provides an overview of the benefits and features provided when using file partitioning as provided by Adabas Vista.

When is Partitioning Required?

File partitioning may be useful for any of the following reasons:

- data volume is too large to be managed within a single file;
- data partitions are needed for separate groups of users who use the same application but usually access different data;
- separate files need a consolidated view as well as an individual view;
- data archiving is needed where recent data is moved to partitions and Data Archiving for Adabas is used to archive out the historical data when it is no longer needed (so often). Data Archiving for Adabas can also be used to recall data to be made available again as new partitions too.

Implementing Partitioning within Application Systems

A partitioned file never provides the same performance levels as a standard file, whether the data management software is Adabas or some other system. The reason is that some form of navigation must occur to determine the location and relative content of partitions.

A number of performance options are available for computer systems. Some require application changes on a massive scale. The cost of these changes is not normally measured in the day-to-day operation of the computer, but they do exist and they are significant.

The burden of partitioning has fallen traditionally on application systems, making them extremely complex. For example, the application has often been required to decide the partition targeted by an access or update. The bigger problem is, however, that the usual level of application independence from the physical data model is lost. The cost of development and ongoing maintenance in these situations can be enormous.

Implementing Partitioning using Adabas Vista

- [Benefits](#)
- [Data Partitioning Example](#)

Benefits

Using Adabas Vista, a partitioned file can be spread across databases and computers, providing excellent scope for load balancing based on application requirements. In addition, navigation processing is performed in the client process, which increases the opportunity for parallel processing.

Adabas Vista can be used to partition data into separate Adabas files without having to re-construct the client applications which use these files. The applications continue to refer to one (simple) Adabas file entity that is accessed and updated as a whole. The physical data model is partitioned and can be distributed across a wide-ranging computer complex.

Adabas Vista addresses the problem of managing and productively using the massive physical capacity of an Adabas file, given commercial constraints on the amount of data that can be put in a single file. For example, a commercial requirement to limit outage to 15 minutes focuses attention on restore times and thus on the amount of data maintained in each file.

Because Adabas Vista processing is performed predominantly in the client process, any overheads associated with partitioning have minimal impact on the database service. Adabas is usually the busiest process by far within the computer. Normally, many thousands of clients concurrently use the same Adabas service. Adabas has proven its worth in these situations, providing excellent response times. The limiting factor for Adabas is often the processing capacity of the CPU. New ways are continually being sought to minimize database CPU consumption and to spread processing over as many CPU engines or even computers as possible.

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 .

The cost of navigation

- is the same for 1 or 100 partitioned files because database requests must be monitored in any case; and
- does not increase with the addition of new partitioned files and the processing does not affect the capacity of the database.

In summary, using Adabas Vista to partition data

- ensures the independence of programs and data; and

- saves the time and money required for program changes.

Data Partitioning Example

The following example shows a typical use of data partitioning.

A company which has offices in various countries maintains a single file which contains information about all employees regardless of location. Each country is represented in the file by a 3-character code (D: Germany, UK: United Kingdom, USA: United States).

Country	Employee No.	Name
USA	100	Smith
D	200	Kenji
USA	300	Barker
UK	125	Smith
D	350	Matsui
USA	150	Wills
D	175	Smith
UK	325	Adams
UK	425	Fern

Using Adabas Vista, the same data may be partitioned into three separate files, organized by a partitioning field, in this case Country:

Country	Employee No.	Name
UK	125	Smith
UK	325	Adams
UK	425	Fern

Country	Employee No.	Name
USA	100	Smith
USA	300	Barker
USA	150	Wills

Country	Employee No.	Name
D	200	Kenji
D	350	Matsui
D	175	Smith

The organization may also have details of all the vehicles it has produced. Access to this data is usually for the current year, with occasional access to previous years. New records are only added for the current year and updates to the data are only permitted for the current year.

If this Vehicles file is partitioned on the field Production Year, updates (new records and updates to existing records) will only take place against a single partition. Regular backup and recovery procedures can be targeted to this single file, with an occasional backup for the entire data.

Typical Partitioning Deployment Scenarios

This section shows typical deployments of Adabas Vista:

Very Large Files

Very large files present an obvious use for partitioning. Although Adabas can hold a massive number of records within a file, the difficulties of handling volume remain. The key to managing mass data is the ability to divide the data into more easily managed units.

By using Adabas Vista to partition a large file across multiple databases, the processing load can be spread across the computer service. If the computer has more than one CPU engine, advantages are gained for all users by making greater use of the parallel availability of the CPU engines.

More and more sites are clustering computers to operate as one large service. Partitions can be placed on specific computers within the overall computer complex, thereby localizing the majority of data usage while maintaining the overall large file availability for management information.

The Adabas Vista **partition outage** feature can also be used to greatly increase the overall availability of data. This feature provides tolerance levels for partition availability. Application access to the partitioned file can be maintained transparently even though all the partitions of the file may not be available.

Archiving

Many organizations have a requirement to keep data for a specified period of time; often, many years. The sheer volume of data forces many organizations to archive, just to manage the data. In either case, an archiving and a recall process must be established and maintained.

Archiving is often a case of managing date-related data. However, types of data that are not related to date may also require archive processing.

A date-ordered file eventually contains old data that may not be used by most or all users. The old data can have an adverse effect on outage time because of additional data volumes, or perhaps because index sizes are too large making search times longer. Archiving is needed.

Most archiving systems access old data, copy it to archive, and delete it from the current file. This can take a considerable amount of time.

Adabas Vista offers a more flexible approach. Partitioning by date means that new data can be directed towards new partitions. Partition criteria (and sizing) might, for example, be based on handling a complete year, or month. Using the Adabas Vista **partition restriction** feature, old partitions can be left intact but made unavailable to all but a few users who require archived data.

The need for an archiving operation that interferes with normal operation as it searches through masses of data is thereby greatly reduced.

Should an emergency arise where older data is required urgently, the old database or file can simply be restored and the partition definition altered to make it available to the application again.

Using this approach, it is much easier to keep more old data available for longer, perhaps on remote (connected) computers.

Merger

Some organizations merge their operations for efficiency as well as growth. The merged units are often in the same business and therefore have similar systems and data, although perhaps not identical.

Using the Adabas Vista **consolidation** feature, a single file image may be imposed upon multiple, previously unrelated files. Although the files may be different, they support the same consolidated view.

For example, two organizations merge and each has its own accounting applications and files. Both support a common view. The new management requires information that is consolidated for the overall operation. However, the individual operations still need their old systems to operate as they did prior to the merge.

Using the Adabas Vista feature **mixed mode access**, both the individual and the consolidated views can be used by all applications. A new management information system can therefore be developed using the consolidated view without losing the operability of the previous systems.

Service Bureau

Smaller organizations often use a central computing service. Sometimes a common application such as payroll or accounts is used, but the independence of the data is mandatory. More often, a complete, separate copy of the application is made. This means, of course, that maintenance has to be rolled out to multiple services or repeated.

Adabas Vista makes it possible for multiple, separate parties to simultaneously share an application with separate data files managed as individual partitions of a simple, single file image. This can be accomplished using the previously mentioned features **partition restriction** and partition outage.

The cost saving alone is enough to justify using a single application for multiple, separate user groups.

The same effect can be achieved by using Adabas Vista definitions to partition a single, massive file without physically splitting the data. One physical file is used transparently as multiple, discreet entities.

Accessing Partitioned Files

This section describes the access modes which can be used to access partitioned files.

Access Modes

Adabas Vista supports the following modes of access to a partitioned Adabas file:

- *focused mode* is based on the partitioning field and allows direct access to a particular partition.
- *distributed mode* is not based on the partitioning field and requires access to all partitions of a file.
- *mixed mode* allows a partition to be addressed directly by its real file number, even while using the single file image.

- **Focused Mode**
- **Distributed Mode**

■ Mixed Mode

Focused Mode

A file often has a field which is used primarily for access to the file. In fact, many sites regularly re-order files according to their dominant key field. When using Adabas Vista, this key field is most commonly used as the *partitioning field*.

Applications generally access a file using search data that is based, in whole or in part, on the primary key field. In these cases, Adabas Vista is able to detect the specific partition(s) needed to satisfy the access. This is referred to as *focused mode access*.

For example, a Natural application contains the following:

```
0250  FIND EMPLOYEES WITH COUNTRY = 'USA'
0260      DISPLAY NAME EMP-NUMBER COUNTRY
0270  END-FIND
```

Adabas Vista

- detects access explicitly or implicitly based upon the partitioning field Country;
- directs the access to the corresponding partition (USA) based on the search argument.

Using this access mode, Adabas Vista can focus access to only one partition, which results in a low processing overhead..

Performance Notes:

Selecting the correct partitioning field is very important when first partitioning a file. It is often the case that files are regularly re-ordered according to a key in order to achieve greater I/O performance. Such a key is generally suitable as the Adabas Vista partitioning field.

If certain groups of users focus on specific portions of a file, the **partition restriction** feature may be beneficial, especially if access can be completely restricted to a partition. Restricting access increases the amount of focused access achieved.

Distributed Mode

Adabas Vista also provides *distributed mode access* using data that is not related to the partitioning field. For example:

```
0250    FIND EMPLOYEES WITH NAME = 'SMITH'  
0260        DISPLAY NAME EMP-NUMBER COUNTRY  
0270    END-FIND
```

Access by Name cannot be focused onto one particular partition. Adabas Vista distributes the access to all partitions before collating the results in the expected order.

Mixed Mode

Adabas Vista also provides mixed mode, in which applications can not only access a single file image for all partitions, but also can access a partition directly using its actual Adabas file number. This is referred to as *mixed mode access*.

Additional Partitioning Features

- [Partition Outage](#)
- [Independent Partition Maintenance](#)
- [Partition Restriction](#)
- [Consolidation](#)
- [Partition Sharing \(MultiPart\)](#)
- [Distributed Lock Mode](#)
- [Extreme files](#)

Partition Outage

This section describes the Adabas Vista partition outage feature.



Note: Earlier versions of Adabas Vista referred to the partition outage feature as partial data operation (PDO). This method of specifying this feature is no longer supported.

- [Description and Use of Partition Outage](#)
- [Defining Outage Tolerance using the Critical Parameter](#)

- **Dynamic Modification of the Critical Parameter**

Description and Use of Partition Outage

The partition outage feature can be used to manage planned outages for a partitioned file.

Using this feature, tolerance levels can be defined which control the action to be taken when a partition becomes unavailable. Application access to the partitioned file can be maintained transparently even though all the partitions of the file may not be available. Data is returned from the available partitions.

Sensitivity to partition outage can be set unilaterally or overridden on a user basis by an application. This provides flexible control over what happens when a partition becomes unavailable.

The suitability of this feature should be assessed for each partitioned file. For suitable files, using this feature:

- maximizes Adabas Vista file availability (24 * 7);
- maintains partitions transparently; and
- recovers partitions transparently.

Partition unavailability is determined by the following Adabas response codes:

Response Code	Description
17	partition locked
48	partition locked by an Adabas utility
148	database not active

Defining Outage Tolerance using the Critical Parameter

The partition parameter `Critical` is used to specify the outage tolerance level for a partition.

As an example, assume that a file is partitioned by country.

Using partition outage, it will be possible to take only the UK partition offline. Access may continue as required for the other countries, but UK data is unavailable until returned online.

In that users in the UK work predominantly on UK data, partition outage can be set so that only UK data is critical to UK users. UK users tend not to use or need data from the other countries.

Consequently, UK users are unaffected by USA outages. However, UK outages do interrupt UK users because their critical data is affected. Similar arrangements can be made for the users in other countries.

Dynamic Modification of the Critical Parameter

Adabas Vista API function PARTOPTS can also be used to to update or extract a user's current Critical parameter setting for a specific partition. The API function CRITREP can be used to list those partitions that are not currently available.

Independent Partition Maintenance

Since Adabas Vista partitions are actually independent Adabas files, it is possible to maintain them individually.

Independent partition maintenance can be used to size, order, and restore according to the needs of the individual partition. It is not necessary that all partitions adhere to the same physical constraints. The Adabas Associator space required by the partition can be individually fine tuned as well.

Separating partitions into independent Adabas files also allows you to perform parallel maintenance.

Partition Restriction



Note: Earlier versions of Adabas Vista referred to this feature as restricted segment option (RSO).

The partition restriction feature provides access control for each partition of a partitioned file. Partition restriction can either be implemented for all of the partition users or it can be defined so that certain applications are permitted to override the restriction at runtime.

Partition restriction:

- dynamically restricts user views of partitioned files,
- provides access based on user role controlled by the application,
- maximizes focused access.

This feature therefore provides the following benefits:

- Security: the availability of data can be restricted to authorized users
- Performance: the amount of focused access can be increased by limiting the number of partitions available to each user based upon role, location, or other criteria.

The type of access to a partition (none, read-only, read/write, or only) is specified using the parameter `Access`.

Consolidation

Adabas Vista partitions do not have to be identical. Provided all the partitions support all of the views to be used by Adabas Vista, the files can operate with different physical layouts (FDTs). Of course, the Adabas source fields that are common to all partitions must be defined identically in each FDT.

For example, if an organization combines with a similar trading company, it is likely that the two companies hold similar data. In this case, a new file for another country could be added to the partitioned definition.

Partition Sharing (MultiPart)

The partition sharing (multipart) feature can be used to define multiple partitions of the same partitioned file in order 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 away from the main file with minimal data manipulation.

This feature is activated using the parameter `Shared Partition`.

Distributed Lock Mode

The distributed lock mode feature provides greater control of record hold logic in a partitioned environment.

When multiple partitions are targeted for distributed access and hold processing is involved, it is possible for successive accesses within the same process to be satisfied from the same partition. However, the other partitions still have another record within the same process on hold because that record is taking part in collating-sequence checking to ensure that the user receives the data in the correct sequence across partitions. If as a result, records are held for longer periods than expected, the frequency of the following response codes may increase:

- response code 145 because of record contention; and
- response code 9, subcode 2 because of Adabas transaction timeout (TT) expiration.

The `GET` statement should be used to hold records immediately before modification. This minimizes hold queue sizes and contention. This is even more relevant when using partitioned files.

If this recommendation cannot be followed, it may be useful to adjust the runtime control `Distributed Lock Mode`.

Extreme files

The Extreme partitioned file feature enables applications to use files that are extremely large in size (up to 255 times 4,294,967,294 records). The feature is designed for 100% application transparency, however there may be circumstances where minimal changes to the application may be necessary in order to adopt the feature.

For more information, refer to *Using Extreme partitioned files* in the *Programming Guidelines*.

2 File Translation

■ Standard File Translation	16
■ Target Categories	16
■ Pages	17

This section provides an overview of the benefits and features provided when using file translation as provided by Adabas Vista.

Standard File Translation

Many organizations have a requirement to have an application use either test or production data without requiring that the application itself be modified. The ability to do so reduces significantly the risk involved in modifying applications immediately prior to their being placed into production mode.

Adabas Vista's file translation feature can be used to change file numbers dynamically within Adabas applications at runtime. The file numbers are translated based on the translation information supplied to Adabas Vista.

For example, a Natural program is compiled in a test environment using a test DDM to access file 25 in database 10. After testing has been completed, the program is to be placed into production using file 67 in database 41. By providing this information to Adabas Vista using translation rules, the switch to production mode can be made without any further application modification

Target Categories

An application file generally has many different versions of its data. These different versions normally relate to the promotion stages that an application follows as it progresses into production. In order to avoid application changes, as each promotion stage is completed, the Vista translation rules for that file would need to be updated to reflect the new data to be used at that particular promotion level. With standard file translation, updating the translation rule will cause everybody using the application to use that particular version of data which is clearly not ideal if there are a number of application changes each of which may be at a different promotion stage.

However, by defining each of your promotion stages as a Target Category, Adabas Vista will allow the specification of translation rules for each promotion stage. Because the Target Category is runtime configurable, you can then have the same application being run from any number of jobs, each being at a different promotion stage, yet each using data relevant for that promotion stage.

Pages

When different business groups share the same applications, each business group's data requirements have to be met. For example, groups may share the same infrastructure files (for example, FNAT) but each group may also have their own independent files (for example, Accounts).

File translation pages are a means of providing this flexibility. Each page can contain one or more translation rules. Any number of pages can be defined and any job which is configured for Adabas Vista can specify up to 8 of these pages. For example, one page could define all the infrastructure translation rules which are shared across the applications, another page could be those translation rules pertinent to one business group, another page those pertinent to another business group, etc. Jobs run by any particular business group need only ensure the appropriate pages are defined to it.

At runtime, Vista will merge the specified pages into a single set of rules. Rule attributes can be established to determine rule selection in the event that an application file is defined to more than one page.

3

Adabas Vista Tools

■ File Analysis (AVILOOK)	20
■ Online Services (SYSAVI)	21

This section describes the tools provided with Adabas Vista.

File Analysis (AVILOOK)

The Adabas Vista file analysis tool AVILOOK provides statistics on the type and quantity of key-based accesses against an Adabas file. These statistics can be used to identify the Adabas files that may benefit from the Adabas Vista partitioning option.

The following screen is an example of the output provided by AVILOOK.

```
16:56:07      ***** A D A B A S   VISTA SERVICES 8.1.1 *****      2006-05-29
              -   Avilook File Statistics   -                          V14120M1

DBID: 231      (TEST-V7-DB)
File: 2                               Started: 2006-05-29 09:09:20
                                      Paused:

      CC      Desc      Command Count      CC      Desc      Command Count
      L3      AA              2836          L3      AA              2836
      S1      AB              1324          S1      AB              1324
      L3      BC               24          L3      BC               24
      L9      S1               26          L9      S1               26

Other Commands not listed above: 1558

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10---PF11---PF12---
      Help      Exit  Refr                                Menu
```

This screen shows statistics on the command type and quantity of key-based accesses against the file (for example, S1, L3, and L9 commands). The statistics are displayed in descending order by command code (CC) and Adabas two-character field name (Desc).

In the example, there are 2836 accesses to file 2 using an L3 command with the Adabas field AA as the primary sequence field. Such a file, where the predominant access is by a single key, may benefit from being partitioned using the Adabas field AA as the Adabas Vista partitioning field.

For more information, see section Using AVILOOK.

Online Services (SYSAVI)

Adabas Vista Online Services (SYSAVI) provides a wide variety of online operational and administrative services.

The following services are provided:

Option	Description
System Settings	Maintain Adabas Vista configuration file setting.
Maintenance	Maintain Adabas Vista definitions and parameters.
Current Activity Displays	Display the activity of the current session.
Special Services	Provide installation and applied zap information for the Adabas Vista components.
AVILOOK	File analysis tool.

For more information, see section Adabas Vista Online Services.

4 Cluster Support

A clustered operating system is a collection of independent operating system images working together as one. The objective is to spread work around the cluster to achieve better load balancing, throughput, and availability.

Clustered applications appeared before their operating system counterpart. Application clusters operated over multiple processes within a single operating system image. Now clustered applications are enhanced to operate throughout an operating system cluster.

Adabas Cluster Services allows multiple instances of the Adabas server to operate against the same database across multiple operating system images, which enhances the general clustering approach for high-end systems.

For more information, see the section [Using Cluster Services](#).

