Adabas Vista

Parameters

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.

Table of Contents

1 Parameters

This document describes the parameters which are used to control Adabas Vista processing. These parameters are defined and maintained using Adabas Vista Online Services.

Job Parameters	General Parameters	File Partitioning Parameters	File Translation Parameters
Clustered Application Service Name	Command Limit	Access	Target Database Number
Distributed Lock Mode	Enable API Override	Adabas TOPISN	Target File Number
Enable Multiple Database Updates	Source Database Number	Critical	
Enable Profile Overrides	Source File Number	Enable ISN Positioning	
Error Reporting	Source Name	Maximum Number Partitions	
Error Response Code	Source Profile ID	Partition Database Number	
Estimated Client Sessions	Target Category	Partition File Number	
ET Data Database Number		Partition ID	
Fixed Memory Pool Size		Partition ID Assignment	
Global Format IDs		Partitioning Field	
Manage Sessions		Partitioning Field High Value	
Profile ID Validation		Shared Partition	
Profile Overrides Selection		Store Control Option	
Read-Ahead Pool Size		User Partition Concurrency	
Read-Ahead Process			

Job Parameters	General Parameters	File Partitioning Parameters	File Translation Parameters
System Coordinator Group Name			
Trace			
Vista ON/OFF for Job			

2 Job Parameters

This section describes the Adabas Vista job parameters.

Clustered Application Service Name

Description	Possible Values	Default
The service name used to link together the instances of the clustered application system. For clustered applications such as job types CICS Cluster, IMS, or UTM, the service name is required to link all the instances of the same application system. If the job is running in a single operating system image (for example, UTM), the service name is still required. Refer to the *Adabas System Coordinator* documentation for more information.	see text	none

Distributed Lock Mode

Description	Possible Values	Default
Controls the type of record hold processing to be used in a partitioned environment. This parameter is applied when a distributed `READ(L6)` or `FIND(S4)` command is issued against a partitioned file. Possible values are: ■ 0: Normal processing. ■ 1: The hold option is removed from the command and a `L3` or `S1` command is issued instead. If a record is to be modified, the corresponding ISN for the record is placed in hold status before it is modified. ■ 2: The hold option is removed from the command and a `L3` or `S1` is issued instead. Record collating-sequence processing occurs immediately before returning a record to the user. After this processing, the record is placed in hold status. ■ 3: The same as value '2' with an added integrity check during collating-sequence processing. ■ 4: `READ (L6)` and `FIND (S4)` commands result in response code 249.	0 | 1 | 2 | 3 | 4	0

See also section Distributed Lock Mode.

Enable Multiple Database Updates

Description	Possible Values	Default
Controls whether or not a user can update more than one database in a single transaction. Using Adabas Vista for file partitioning and translating increases the probability of a single transaction spanning multiple databases. In such circumstances, the Adabas Transaction Manager (ATM) should be used to ensure the integrity of the transaction. If this parameter is set to YES, Adabas Vista processes ET commands serially. Any failure during this serial process jeopardizes the integrity of the transaction. It is therefore recommended that this parameter be set to YES *only* when the Adabas Transaction Manager is installed. If Adabas Transaction Manager is not used, this parameter can be used to allow or disallow multiple database updates within the same transaction: ■ If allowed, ET (or BT) commands are issued serially to the relevant databases. Under such conditions, the programmer should not assume any particular sequence of ETs. Furthermore, any transaction data specified with the ET may appear anywhere in this sequence. ■ If disallowed, an Adabas Vista response code is generated whenever an attempt is made to modify a second database within one transaction.	YES | NO	YES

Enable Profile Overrides

Description	Possible Values	Default
Controls whether or not profile overrides are to be allowed for the job. Each job defined in the Adabas Vista configuration file may have job profile overrides specified. These override parameters influence the target category for a given source profile, and thus the actual database and file number for the appropriate commands issued for the duration of this job. Any job profile overrides specified for the job will be ignored if this parameter is set to NO. Note that only published overrides will be processed.	YES | NO	NO

Error Reporting

Description	Option	Possible Values	Default
Controls whether or not the WTO option is to be used during Adabas Vista error message processing. The WTO option identifies the subcode of an Adabas Vista error for those programs that do not have ON ERROR processing to display the subcode. The subcode is needed to uniquely identify the reason for the error. See also section Error Handling. If YES is specified for the WTO option, an error message written to the job log identifies the subcode associated with each Adabas Vista error. This error message has the prefix AVI-0018-10.	WTO	YES | NO	YES

Error Response Code

Description	Possible Values	Default
The Adabas response code to be used for identifying Adabas Vista processing errors. A subcode in the Additions 2 field of the Adabas Control Block uniquely identifies the actual error. Possible values for this subcode are detailed in the section Messages. See also section Error Handling. If you need to use a response code other than the default 249, refer to the *Adabas Messages and Codes* documentation for unused response codes.	see text	249

Estimated Client Sessions

Description	Possible Values	Default
Determines the approximate size of the Adabas System Coordinator user pool where ■ type "a" represents batch, TSO, CMS, and TIAM jobs ■ type "b" represents Com-plete, CICS, CICSplex, IMS, and UTM jobs Refer to the *Adabas System Coordinator* documentation for more information.	number	type a: 2 type b: 1000

ET Data Database Number

Description	Minimum	Maximum	Default
Indicates the database number to be used for ET data. ET data requests are processed according to the value of this parameter: ■ value 0: the request is issued with no modification by Adabas Vista. It is the user's responsibility to ensure that the database number specified for the request is valid. ■ value not 0: Adabas Vista modifies the request by issuing it to the database number specified by this parameter. **Note:** In a Natural environment, the Natural parameter ETDB specifies the database where transaction data is to be stored (refer to the section Profile Parameters in the *Natural Installation and Operations* documentation). If this parameter is used dynamically, it is recommended that you set the Adabas Vista job parameter to zero to maintain the dynamic capability in an Adabas Vista environment.	0	65535	0

Fixed Memory Pool Size

Description	Possible Values	Default
Determines the initial size of all fixed pools managed by the Adabas System Coordinator. Refer to the *Adabas System Coordinator* documentation for more information.	size	256

Global Format IDs

Description	Language	Possible Values	Default
Indicates whether or not global format IDs are to be used. This parameter enables a user-supplied global format ID to be propagated appropriately when issued against a partitioned file. **Note:** Adabas Vista supports Natural global format IDs.	Natural	YES | NO	Natural: YES 3GL: NO

Description	Language	Possible Values	Default
Note: If this parameter is set to YES for a 3GL language, the programmer must ensure that the uniqueness of the supplied global format ID is maintained in the last 5 bytes of the 8-byte Additions 5 field of the Adabas Control Block. This allows Adabas Vista to use the first 3 bytes internally.			

Manage Sessions

Description	Possible Values	Default
Determines type of session management to be used for clustered application sessions For clustered applications (job types CICS Cluster, IMS, UTM) where dynamic user movement is possible, you can choose to manage only terminal tasks. This is more efficient and is possible when non-terminal tasks do not move between regions in a clustered application. Refer to the *Adabas System Coordinator* documentation for more information.	see text	Manage Terminal Sessions only

Profile ID Validation

Description	Possible Values	Default
Indicates whether or not a Profile ID is to be validated against other translation rules and partition definitions that have been defined in the Adabas Vista configuration file. Adabas Vista checks each command to determine if it conforms to a published translation rule or partition definition. This check consists of Profile ID, the source database number, and the source file number. If no matching translation rule or partition definition is found, this parameter is then evaluated according to the parameter value specified: ■ NO: the command is allowed to continue without any change. ■ YES: a further check against the configuration file is performed to see if the Profile ID has been defined to any published translation rule or partition definition. If so, the command is allowed to continue; if not, an error code is returned.	YES | NO	NO

Profile Overrides Selection by Profile ID

Description	Possible Values	Default
Indicates whether or not profile overrides for this job are to be selected based on the Profile ID. Separate overrides may be specified for the same source database and file number, but with different Profile IDs. This parameter enables selection of the appropriate overrides based on the current Profile ID. Possible values are: ■ NO: only overrides with a blank Profile ID are considered for processing. ■ YES: only overrides with a matching Profile ID will be processed. Other overrides, including those with a blank Profile ID (unless the current Profile ID is blank) are ignored. **Note:** This parameter is ignored unless the parameter `Enable Profile Overrides` is set to YES.	YES | NO	NO

Read-Ahead Pool Size (Batch Only)

Description	Minimum	Maximum	Default
Indicates the size (in kilobytes) of memory to be allocated to the read-ahead pool. The Adabas Vista read-ahead feature provides the benefits of multifetch processing for batch jobs in an Adabas Vista environment, thus reducing the amount of interprocess communication and command selection within the database. The read-ahead pool is used only during read-ahead processing. This parameter defines the maximum amount of storage to be used. This pool is not extended. Because all commands are dynamically analyzed to determine their suitability for read-ahead processing, it is not necessary to specify files or commands to be excluded from the process. For each suitable command, the read-ahead factor is dynamically modified for best performance. Adabas Vista read-ahead is controlled at the job level by specifying the Adabas parameter `ADARUN PREFETCH`. No other Adabas prefetch parameter is necessary.	32	99999	128

Read-Ahead Process (Batch only)

Description	Possible Values	Default
Indicates whether read-ahead processing is to be handled by Adabas Vista or by Adabas. The read-ahead process itself is controlled by specifying the Adabas parameter `ADARUN PREFETCH`. If Adabas Vista read-ahead is to be used, no other Adabas prefetch parameter is required. Possible values are: ■ V: Adabas prefetch is disabled and all calls are satisfied using Adabas Vista read-ahead processing. ■ A: Adabas prefetch is to be used.	V | A	V

System Coordinator Group Name

Description	Possible Values	Default
Identifies the Adabas System Coordinator group that is to manage the job. For clustered applications (job types CICS Cluster, IMS, UTM), this name is required to link all the instances of the job in the cluster. If the job is running in a single image (for example, UTM), the name is still required. Refer to the *Adabas System Coordinator* documentation for more information.	see text	none

Trace

Description	Option	Possible Values	Default
Indicates whether or not user command tracing is to be used.	Active	YES | NO	NO
The number of commands for which trace entries are to be used. **Note:** Information from these trace entries can be displayed by using the CLOG function of the Adabas Vista API. A sample use of this API can be found in the program AVICLOG in the INPL dataset. **Caution:** Do not modify this program as it may be required for support purposes.	Commands	0 - 1000	128

Vista ON/OFF for Job

Description	Possible Values	Default
Indicates whether or not Adabas Vista is to be enabled for a job. Adabas Vista checks each command which may include checking against the configuration file to determine if the command should be translated, or whether it accesses a partitioned file. If it is known that a particular job makes no access to files that need to be under Adabas Vista control, this parameter can be used to disable Adabas Vista processing in order to enhance throughput. **Caution:** If Adabas Vista processing is disabled, there will be no recognition of commands that need to be targeted elsewhere. Therefore, care should be taken when using this parameter.	YES | NO	YES

3 File Parameters

This section describes the Adabas Vista file parameters.

General

Command Limit

Description	Possible Values	Default
The maximun number of commands permitted for a source profile. Any command that exceeds this number will result in an error code. The default value 0 indicates no limit.		0

Enable API Override

Description	Possible Values	Default
Indicates whether or not the target(s) for a translation rule or partition definition may be modified by the use of an Adabas Vista API function.	Y | N	N

Source Database Number

Description	Minimum	Maximum	Default
The number of a database which identifies an Adabas Vista file. This number is used in conjunction with a source file number. **Note:** In Natural systems, this is the database number defined in a DDM.	1	65535	0

Source File Number

Description	Minimum	Maximum	Default
The file number for an Adabas Vista file. This number is used in conjunction with a source database number. **Note:** In Natural systems, this is the file number defined in a DDM.	1	65535	0

Source Name

Description	Possible Values	Default
A unique name identifying a particular Adabas Vista file without referring to the source database and file number. The name must be unique across all published partitioned file definitions. The name must also be unique across all published translation rules. (optional)		none

Source Profile ID

Description	Possible Values	Default
An ID which can be used to combine translation rules and partition definitions that are active for a given user or group of users. In conjunction with the source database number and source file number, this is known as the source profile. (optional)		blank

Target Category

Description	Possible Values	Default
A mandatory name that is used to differentiate between destinations (target database and file numbers for commands) for user groups. A default target category must be identified. The Adabas Vista 7.4.2 automatic conversion process defines a default target category *DEFAULT in the absence of other instructions. The default value for a given configuration file is indicated on the Maintain Target Categories screen. See Adabas Vista Online Services, Maintain Target Categories, for more information.	The specified name must be defined in the configuration file before it can be used against a translation rule or partition definition.	See text

Partitioning Parameters

Access

Description	Possible Values	Default
Controls the type of access for each partition of a partitioned file. The following values are permitted: ■ FULL: read/write access is permitted ■ READ: read-only access is permitted ■ ONLY: single partition focus is enabled for the partition **Note:** single partition focus may not be used in conjunction with a shared partition. ■ NONE: no access is permitted See section Partition Restriction for more information.	FULL | READ | ONLY | NONE	FULL

Adabas TOPISN

Description	Minimum	Maximum	Default
The maximum Adabas ISN permitted for a partition. Adabas Vista enforces the specified maximum during read and store processing. The initial maximum is determined by the Maximum Number of Partitions / Default Partition TOPISN, established when the partitioned file is first defined. This parameter may be used to provide a lower value in order to maintain specific partition sizes that may have been derived from recovery and maintenance considerations. **Note:** The Adabas TOPISN value is a limit on the maximum ISN for a partition and not the maximum number of records.	1	2,147,483,647	16,777,215

Critical

Description	Possible Values	Default
Indicates the action to be taken whenever a partition becomes unavailable. Possible values are: ■ YES: the user cannot tolerate the partition's unavailability. Normal user operation is interrupted when access is attempted (with the corresponding Adabas response code). ■ NO: the user can tolerate the partition's unavailability. Data in that partition is ignored until the partition becomes available again. The partitions that return one of the partition unavailable response codes (17, 48, 148) when accessed are logged and can subsequently be identified using the CRITREP function of the Adabas Vista API. See section Partition Outage for more information.	YES | NO	YES

Enable ISN Positioning

Description	Possible Values	Default
Indicates whether or not ISN positioning is to be used when an Adabas Vista ISN is provided as an optional start ISN for `L2/5` and `L3/6` commands or as a minimum ISN value for Sx commands. If an `L1` command with Command Option 2 set to I is issued with a starting Adabas Vista ISN of zero and this parameter is set to YES, reading begins from the first partition.	YES | NO	YES

Maximum Number of Partitions (Default Partition TOPISN)

Description	Minimum	Maximum	Default
The maximum number of partitions available for use by a partitioned file. Only those partitions relevant to current processing requirements need to be defined. However, the future requirements of the partitioned file must also be considered when determining the value for this parameter. Because of the structure of the Adabas Vista ISN, this parameter directly affects ■ the amount of space reserved in the ISN field for the Partition ID; and ■ the default Adabas TOPISN for each partition. The default Adabas TOPISN imposed by this parameter is calculated and displayed in order to help determine the correct balance between the maximum number of partitions and the resulting Adabas TOPISN partition limit.	1	65535	255

Partitioning Field

Description	Possible Values	Default
The Adabas name, length, and format of the field used to distribute the data into separate partitions. The partitioning field of a partitioned file can be a standard Adabas field, a descriptor, a superdescriptor, a subdescriptor, or a dummy field. It may *not* be ■ a multiple value field; ■ an item in a periodic group; ■ a field with format F, G, or W; ■ a variable length field (length of zero in the FDT); or ■ a field with the long alphanumeric (LA) attribute. The Adabas UQ attribute is supported only for fields defined as the partitioning field, and only if the appropriate Adabas field name is defined with the UQ option in the Adabas FDT for each partition. The specified Adabas short name, length, and format must be identical to its field definition in the Adabas FDT with one exception: if the partitioning field is a superdescriptor with format A (that is, one of its parent fields is defined with format A), then a format of B must be specified to enable correct specification of the Partitioning Field High Value. The field formats may be obtained from either the full or demo version of Adabas Online System, or by running the Adabas ADAREP utility. **Note:** An optional name for the partitioning field may be provided in the "display as" field to make the partitioning field easier to identify.	see text	none

Partitioning Field High Value

Description	Possible Values	Default
This parameter is mandatory for each partition. It is used to specify the highest value of the partitioning field that can exist in the partition. The value must be specified in accordance with the defined partitioning field's format and length. If the partitioning field's format has been defined as ■ alphanumeric (A), normal alphanumeric values may be specified. ■ packed decimal (P) or unpacked decimal (U), the decimal value must be preceded, if applicable, by a '-' or optionally a '+' sign.	see text	none

Description	Possible Values	Default
■ binary (B), the value must be specified in hexadecimal format (that is, two digits for each byte). In the case of a superdescriptor with a U or P format parent, the sign F or D must be used to indicate a positive or negative value, respectively.		

Some examples:

Defined Partitioning Field	Adabas Field Type	Specified High Value
AA,2,A	standard field	ZZ
BB,2,B	standard field	FFFF
PP,2,P	standard field	999
UU,2,U	standard field	99
S1,4,B	SUPDE=UU(1-2),AA(1-2)	F9F9E9E9
S2,4,B	SUPDE=BB(1-2),PP(1-2)	FFFF999F

Note: The physical data in each Adabas file must be consistent with the implied range specified by the Partitioning Field High Value defined for the partition. This can be achieved by using:

- external sort; or
- the Adabas ADAULD utility and the `SELCRIT/SELVAL` selection criteria parameters.Refer to the *Adabas Utilities* documentation for information.

Partition Database Number

Description	Minimum	Maximum	Default
A database number to which Adabas commands issued against an Adabas Vista partitioned file may be directed.	1	65535	0

Partition File Number

Description	Minimum	Maximum	Default
A file number to which Adabas commands issued against an Adabas Vista partitioned file may be directed.	1	65535	0

Partition ID

Description	Minimum	Maximum	Default
A number that uniquely identifies a partition within a partitioned file. The Partition ID is used together with the Adabas ISN to form an Adabas Vista ISN. This Adabas Vista ISN is returned to the application in place of the Adabas ISN. The allocation of a Partition ID depends on the Partition ID Assignment of the partitioned file. The actual maximum value of the Partition ID is determined by the Maximum Number of Partitions established when the partitioned file is first defined. When such an Adabas Vista ISN is received from the application (for example, in the case of an update), Adabas Vista interprets the ISN and is able to redirect the update to the correct partition. Refer to the section Adabas Vista ISN for more information.	1	65535	none

Partition ID Assignment

Description	Possible Values	Default
Indicates whether Partition ID assignment is to be performed by Adabas Vista automatically or by the user manually. User assignment, which requires that the user specify a Partition ID for each partition, may be useful when an application stores the Adabas Vista ISN as data and the likelihood exists that the structure of the partitioned file may change; for example, inserting new partitions or splitting current partitions. Refer to the section Adabas Vista ISN for more information.	VISTA | USER	VISTA

Shared Partition

Description	Possible Values	Default
This parameter can be used to enable the Adabas Vista shared partition feature (also referred to as the multipart feature). See also the section Partition Sharing.	YES | NO	NO

Example:

Adabas Vista file partitioning normally maps each partition to a unique Adabas file:

Partition 1:	`DBID=1,FNR=10,Partitioning Field High Value=A`
Partition 2:	`DBID=1,FNR=11,Partitioning Field High Value=B`
Partition 3:	`DBID=1,FNR=12,Partitioning Field High Value=C`

The shared partition feature can be used to share an Adabas file between partitions:

Partition 1:	`DBID=1,FNR=10,Partitioning Field High Value=A,Shared Partition=YES`
Partition 2:	`DBID=1,FNR=11,Partitioning Field High Value=B`
Partition 3:	`DBID=1,FNR=10,Partitioning Field High Value=C,Shared Partition=YES`

The above example shows the partition definitions necessary to split all records with a partitioning field value of 'B' from the main file (database 1, file 10) onto a new file (database 1, file 11).

Notes:

1. The partition definitions must still reflect collating sequence.
2. The single partition focus feature may not be used in conjunction with the shared partition feature.

Store Control Option

Description	Possible Values	Default
Controls the placement of new records into a partitioned file. When storing a record to a partitioned file, 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 provide a null value for it, `Store Control Option` may be used to direct the record to a partition.	1 | 2 | F | L	1

The value provided with this option determines the placement of new records according to the presence, absence, or value of the partitioning field within the Adabas format/record buffer. The following table indicates the actions performed for each possible value:

Note: If your requirements are not provided for in the table, contact Software AG support for further assistance.

Value	Partitioning Field with non-null value	Partitioning Field with null value	No Partitioning Field
1	normal placement	normal placement	reject
2	normal placement	reject	reject
F	normal placement	directed to first partition	directed to first partition
L	normal placement	directed to last partition	directed to last partition

Note: In cases where a null value is provided for the partitioning field and the record is subsequently stored, retrieval of the record using the partitioning field depends on the null value suppression (NU) option of the field as defined in the Adabas FDT.

Note: It is recommended that only distributed access (that is, access not based on the partitioning field) be performed on those partitioned files defined with options F or L. Otherwise, records may be retrieved out of sequence.

User Partition Concurrency

Description	Minimum	Maximum	Default
The number of concurrent Adabas Command ID sequences that a user may have outstanding for each partition.	4	255	8

Translation Parameters

Target Database Number

Description	Minimum	Maximum	Default
The number of the database to which an Adabas command that is subject to an Adabas Vista translation rule is to be re-directed.	1	65535	0

Target File Number

Description	Minimum	Maximum	Default
The number of the file to which an Adabas command that is subject to an Adabas Vista translation rule is to be re-directed.	1	65535	0

Index