

Parameter Types

Adabas Fastpath optimization and operation is controlled by the following types of parameters:

- *Buffer parameters* control the operation of the Adabas Fastpath buffer.
- *File parameters* control the optimization of specific files. Some file parameters can override buffer parameter settings when applied to a specific file.
- *Job parameters* define which jobs participate in Adabas Fastpath optimization. Some job parameters can override buffer parameter settings when applied to a specific job.

Adabas Fastpath parameters can be maintained using Adabas Fastpath Online Services, function Parameter Maintenance.

Parameter Descriptions

This section provides a description of each Adabas Fastpath parameter:

Buffer Parameters	File Parameters	Job Parameters
Average Item Size	Direct Access Optimization	Clustered Application Service Name
Fixed Memory Pool Size	End Time	Direct Access Optimization Control
Freespace Index	Expanded File	End Time
Global Operational Control	Initial Status	Estimated Client Sessions
Keep	L1/L2/L3 Read; L9 Histogram	Job End Statistics
Log	RB Length Limit	Read-Ahead Memory Limit
Manage Sessions	Read-Ahead Optimization	Read-Ahead Memory Sizes
Maximum Jobs	Set Concurrency	Read-Ahead Optimization Control
Minimum Buffer Size	Set Expansion	Start Time
RB Length Limit	Set ID Length Limit	System Coordinator Group Name
Restart	Set Limit	
Set Concurrency	Start Time	
Set Expansion	S1/S2 FIND; S8/S9 Sort ISN List	
Set ID Length Limit	Update Sensitivity	
Size		
Synchronous Remote Updates		

Average Item Size

Parameter Type	Use	Minimum	Maximum	Default
Buffer	All memory allocations within the Adabas Fastpath buffer are recorded in an index. This parameter is used as a divisor into the buffer size to decide how many items will be expected within the first index block.	64	32768	1024

Clustered Application Service Name

Parameter Type	Use
Job	<p>The service name to be used to link together the instances of the application system.</p> <p>For clustered applications (job types CICS Cluster, IMS, UTM), the service name is required to link all the instances of the same application system. If the job is running in a single image (for example, UTM), this name is still required.</p> <p>Refer to the <i>Adabas System Coordinator</i> documentation for more information.</p>

Direct Access Optimization

Parameter Type	Use	Possible Values	Default
File	<p>Defines the direct access optimization settings at the file level for Adabas commands:</p> <ul style="list-style-type: none"> ● L1: Get ISN ● L3: Read Logical ● L9: Histogram ● S1: Find ● S2: Find Sorted <p>Each command type may have optimization set ON or OFF. In addition, this setting may have up to seven field level overrides which can be achieved by inserting the Adabas two-character field name(s) on the appropriate command line.</p> <p>For L3 , L9 , S1 , S2 commands, when optimization is set</p> <ul style="list-style-type: none"> ● ON, field level overrides indicate that commands using that field <i>are not to be optimized</i>. ● OFF, field level overrides indicate that commands using that field <i>are to be optimized</i>. 	ON OFF	OFF

Direct Access Optimization Control at Job Level

Parameter Type	Use	Possible Values	Default
Job	<p>Controls direct access optimization at the job level.</p> <p>If this job parameter is set to OFF, direct access optimization is not performed, regardless of the settings in the buffer and file parameters.</p>	ON OFF	ON

End Time

Parameter Type	Use	Minimum	Maximum	Default
File or Job	<p>Indicates the time that optimization is to be ended.</p> <p><i>File</i> start and end times define the period for which the file parameters are to be used for optimization. The start and end times are continuously monitored and the file parameters are switched on/off dynamically. It is also possible to define various file parameters each with a different start and end time period.</p> <p><i>Job</i> start and end times are only checked when a job starts. If the current time is between the limits, the job's optimization parameters will be used. No additional checks are made.</p> <p>Note: The values 00:00 and 24:00 have the same meaning; that is, midnight.</p>	00:00	24:00	none

Estimated Client Sessions

Parameter Type	Use	Default
Job	<p>This parameter is used to determine the approximate size of the Adabas System Coordinator user pool where</p> <ul style="list-style-type: none"> ● type "a" represents batch, TSO, CMS, and TIAM jobs ● type "b" represents Com-plete, CICS, CICSplex, IMS, and UTM jobs <p>Refer to the <i>Adabas System Coordinator</i> documentation for more information.</p>	type a: 2 type b: 1000

Expanded File

Parameter Type	Use	Possible Values	Default
File	<p>This parameter must be defined (no default) and tells Adabas Fastpath whether or not the Adabas file is defined as 'expanded'. If this parameter is set to Y, direct access optimization for L3 or L9 commands for the file are not permitted.</p> <p>For more information on expanded files, refer to the Adabas documentation.</p>	Y N	none

Fixed Memory Pool Size

Parameter Type	Use	Default
Buffer	<p>Determines the initial size of all fixed pools managed by the Adabas System Coordinator.</p> <p>Refer to the <i>Adabas System Coordinator</i> documentation for more information.</p>	256

Freespace Index

Parameter Type	Use	Minimum	Maximum	Default
Buffer	All free memory in the Adabas Fastpath buffer is classified into a fixed-size freespace index. This parameter sets the number of different freespace categories (entries) in the index.	16	32768	512

Global Operational Control

Parameter Type	Use	Possible Values	Default
Buffer	<p>Defines the optimization settings at the buffer level. The following optimization settings can be defined:</p> <ul style="list-style-type: none"> ● direct access ● sequence processing ● FIND (Sx/L1) ● READ PHYSICAL (L2) ● READ LOGICAL (L3) ● HISTOGRAM (L9) <p>If OFF is specified, no optimization will be in effect regardless of any file or job settings.</p>	ON OFF	ON

Initial Status

Parameter Type	Use	Possible Values	Default
File	<p>Controls the activation of file parameters.</p> <p>If the setting is OFF, the file parameters are inactive. The setting must be changed to ON and SYSAFP must be used to stop/start the file.</p>	ON OFF	ON

Job End Statistics

Parameter Type	Use	Possible Values	Default
Job	<p>Controls the printing of optimization statistics for a job.</p> <p>If this parameter is set to YES, at the end of a job, optimization statistics are output to the operator console by the Adabas Fastpath asynchronous buffer manager using the AFP-0040 to AFP-0042 messages for the named job(s).</p>	YES NO	NO

Keep

Parameter Type	Use	Minimum	Maximum	Default
Buffer	<p>Determines the number of days that the statistics log is to be retained.</p> <p>When automatic statistic logging is enabled (see the parameter <code>LOG</code>), an old log that exceeds the n days is deleted. This limits the number of log records kept (per buffer) in the Adabas Fastpath configuration file.</p>	1	32768	30

Log

Parameter Type	Use	Minimum	Maximum	Default
Buffer	<p>Controls how often the statistics log is to be written (specified in minutes).</p> <p>A statistics log record is written to the Adabas Fastpath configuration file by the Adabas Fastpath asynchronous buffer manager every n minutes. A value of 0 implies no automatic logging but does not prevent logs being requested by using the Adabas Online Special Services function. See also the parameter <code>Keep</code>.</p>	0	32768	60

L1/L2/L3 Read; L9 Histogram

The use of these parameters is described in the descriptions of the following parameters:

- Direct Access Optimization
- Read-Ahead Optimization
- Global Operational Control

Manage Sessions

Parameter Type	Use
Buffer	<p>Determines type of session management to be used for clustered application sessions.</p> <p>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.</p> <p>Refer to the <i>Adabas System Coordinator</i> documentation for more information.</p>

Maximum Jobs

Parameter Type	Use	Minimum	Maximum	Default
Buffer	<p>This parameter indicates the maximum number of optimized jobs which are to be run concurrently. The corresponding number of job areas are then reserved. If an area is not available, additional (concurrent) jobs are not optimized.</p> <p>TP Monitors such as CICS, Com-plete, and Shadow require only a single job area. UTM and IMS/DC require a job area for each task being used (plus 1). TSO, TIAM, and CMS require a job area for each user. Batch jobs use one area each.</p>	1	9999	12

Minimum Buffer Size

Parameter Type	Use	Minimum	Maximum	Default
Buffer	<p>Defines the minimum buffer memory size (in k) to be requested from the operating system for the Adabas Fastpath buffer.</p> <p>If the value is 0, the buffer parameter <code>Size</code> is the only value used.</p>	0	value	0

RB Length Limit

Parameter Type	Use	Minimum	Maximum	Default
Buffer or File	Restricts the size of buffer used for evaluating data for direct access optimization. A file parameter value can be used to restrict the buffer parameter setting. A file parameter with a value of zero is ignored.	16	32768	1024

Read-Ahead Memory Limit

Parameter Type	Use	Minimum	Maximum	Default
job	Limits the amount of memory that can be used for read-ahead optimization for a job.	0		0

Read-Ahead Memory Sizes

Parameter Type	Use	Minimum	Maximum	Default
Job	This parameter defines the sizes of areas (in k) to be allocated by individual jobs to perform read-ahead optimization using the Adabas multifetch function. The sizes (4 categories) must be defined in ascending order.	1	32	4, 8, 16, 32

Read-Ahead Optimization

Parameter Type	Use	Possible Values	Default
File	<p>Defines the read-ahead optimization settings at the file level for Adabas commands:</p> <ul style="list-style-type: none"> ● L1: Get ISN ● L2: Read Physical ● L3: Read Logical ● L9: Histogram ● S1: Find ● S2: Find Sorted ● S8: Process ISN List ● S9: Sort ISN List <p>Each command type may have optimization set ON or OFF. In addition, this setting may have up to seven field level overrides which can be achieved by inserting the Adabas two-character field name(s) on the appropriate command line.</p> <p>When optimization is set</p> <ul style="list-style-type: none"> ● ON, field level overrides indicate that commands using that field <i>are not to be optimized</i>. ● OFF, field level overrides indicate that commands using that field <i>are to be optimized</i>. 	ON OFF	OFF

Read-Ahead Optimization Control at Job Level

Parameter Type	Use	Possible Values	Default
Job	<p>Controls read-ahead optimization at the job level.</p> <p>If this job parameter is set to OFF, read-ahead optimization is not performed, regardless of the settings in the buffer and file parameters.</p> <p>In addition, you can choose between the batch (BAT) or TP monitor (TP) algorithm for setting the incremental multifetch factor. The batch algorithm accelerates the read-ahead rate more than the TP algorithm.</p>	OFF BAT TP	<p>ON</p> <p>Note: Default settings are Batch=BAT and TP Monitor=TP.</p>

Restart

Parameter Type	Use	Minimum	Maximum	Default
Buffer	<p>Indicates after how many hours a buffer is to restarted automatically after the last buffer start or restart.</p> <p>If the value is not zero, the buffer is restarted <i>n</i> hours after the last buffer start or restart time.</p>	0	999	0

Set Concurrency

Parameter Type	Use	Minimum	Maximum	Default
Buffer or File	<p>This parameter indicates the number of concurrent retrievals which can occur from each set of direct access data created by Adabas Fastpath.</p> <p>A file parameter value can be used to restrict the buffer parameter setting. A file parameter with a value of zero is ignored.</p>	1	16	2

Set Expansion

Parameter Type	Use	Minimum	Maximum	Default
Buffer or File	<p>This parameter defines the number of free areas to be allocated by the Adabas Fastpath asynchronous buffer manager when a direct access set is expanded. This is also referred to as the growth rate for direct access sets.</p> <p>A file parameter value can be used to restrict the buffer parameter setting. A file parameter with a value of zero is ignored.</p> <p>Note: An efficiency rating for growth is calculated by Adabas Fastpath. Expansion is accelerated for sets which have above average growth rates, and expansion is decelerated for sets which have below average growth rates.</p>	1	256	8

Set ID Length Limit

Parameter Type	Use	Minimum	Maximum	Default
Buffer or File	<p>Adabas Fastpath uses direct access set identifiers which contain search and format data for direct access command models. This parameter limits the size of data items within each set.</p> <p>A file parameter value can be used to restrict the buffer parameter setting. A file parameter with a value of zero is ignored.</p>	256	16384	1024

Set Limit

Parameter Type	Use	Minimum	Maximum	Default
File	Adabas Fastpath creates sets of direct access command models. This parameter limits the size (in KB) of data items within each set.	0		0


Size

Parameter Type	Use	Minimum	Maximum	Default
Buffer	This parameter defines the size (in k) of the memory to be allocated to the Adabas Fastpath buffer. When requesting the memory from the operating system, the value actually obtained depends on the setting of the parameter <code>Minimum Buffer Size</code> .	1MB		4MB

Start Time

Parameter Type	Use	Minimum	Maximum	Default
File or Job	<p>Indicates the time that optimization is to be started.</p> <p><i>File</i> start and end times define the period for which the file parameters are to be used for optimization. The start and end times are continuously monitored and the file parameters are switched on/off dynamically. It is also possible to define various file parameters each with a different start and end time period.</p> <p><i>Job</i> start and end times are only checked when a job starts. If the current time is between the limits, the job's optimization parameters will be used. No additional checks are made.</p> <p>Note: The values 00:00 and 24:00 have the same meaning; that is, midnight.</p>	00:00	24:00	none

Synchronous Remote Updates

Parameter Type	Use	Possible Values	Default
Buffer	<p>This parameter indicates whether synchronous or asynchronous update communication is to be used for remote databases and for remote cluster services instances.</p> <p>When a database is remotely located from an Adabas Fastpath asynchronous buffer manager, it communicates update information by waiting in the nucleus thread for the reply before allowing the update transaction to complete (this mode is called synchronous mode; which is the default).</p> <p>If this parameter is set to N, update communication does not wait for the reply before allowing the update transaction to complete (asynchronous mode).</p> <p> Warning: Although the asynchronous option may improve performance in certain situations, it may also compromise data integrity in the Adabas Fastpath buffer and must therefore be used with caution.</p> <p>If this parameter is set to Y for any file in a database, the database defaults to synchronous mode updates. In order to change to asynchronous mode, all files for the database must be set to N.</p>	Y N	Y

System Coordinator Group Name

Parameter Type	Use
Job	<p>Identifies the Adabas System Coordinator group that is to manage the job.</p> <p>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.</p> <p>Refer to the <i>Adabas System Coordinator</i> documentation for more information.</p>

S1/S2 FIND; S8/S9 Sort ISN List

The use of these parameters is described in the descriptions of the following parameters:

- Direct Access Optimization
- Read-Ahead Optimization
- Global Operational Control

Update Sensitivity

Parameter Type	Use	Possible Values	Default
File	<p>This parameter is used to control how update commands are to be processed for a file.</p> <p>The possible settings are:</p> <p>N (none): Ignore update type commands for data held in the Adabas Fastpath buffer R (record level): Remove data held in the Adabas Fastpath buffer when update type commands are found F (file): Remove data for the file held in the Adabas Fastpath buffer when update type commands are found D (distributed record level): Use only on advice from Software AG.</p>	N R F D	R