

L1 and L4 Commands: Read / Read and Hold Record

The L1 and L4 commands read a single record from Data Storage.

We recommend that you set unused ACB and ACBX fields to binary zeros before the direct call is initiated.

This chapter covers the following topics:

- Function and Use
 - ACB Interface Direct Call: L1 and L4 Commands
 - ACBX Interface Direct Call: L1 and L4 Commands
 - Buffers
-

Function and Use

Specify the file number and ISN of the record to be read. In addition, indicate in the format buffer which fields are to be read. Adabas returns the requested field values in the record buffer.

The L4 command performs the same function as the L1 command, but places the record in hold status. If the record to be held is currently being held by another user, the command is placed in wait status until either the record becomes available or the transaction times out. If the L4 command is issued with the Command Option 1 field set to "R", Adabas returns response code 145 (ADARSP145) if the record to be read and held is unavailable.

The first-unused-ISN (F) option specified in the Command Option 2 field returns the next highest unused ISN for the specified file in the ISN field. The F option cannot be used for expanded files.

The multifetch/prefetch option can improve performance by eliminating the time needed for single-record fetches. Multifetching/prefetching can be enabled by specifying "M", "O" (for multifetching) or "P" (for prefetching) in the Command Option 1 field. Refer to the section *Using the Multifetch/Prefetch Feature*, for more information.

The GET NEXT (N) option specified in the Command Option 2 field reads the records identified by the ISNs contained in an ISN list (which was created previously by an Sx command) without the user having to provide the ISN of the record to be read with each L1 or L4 call. Adabas selects the ISN from the list and reads the record identified by that ISN.

The read by ISN sequence (I) option specified in the Command Option 2 field reads a record identified by the ISN you specify in the ISN field. If the ISN specified is not present in the file, the record of the next higher ISN is read, and that record's ISN is returned in the ISN field.

The compressed option (set by specifying "C." in the format buffer) can be used to request that the record read is to be returned in compressed format as it is stored internally by Adabas.

ACB Interface Direct Call: L1 and L4 Commands

This section describes ACB interface direct calls for the L1 and L4 commands. It covers the following topics:

- Control Block and Buffer Overview
- Control Block Field Descriptions
- ACB Examples

Control Block and Buffer Overview

Control Block

Field	Position	Format	Before Adabas Call	After Adabas Call
	1-2	--	--	--
Command Code	3-4	alphanumeric	F	U
Command ID	5-8	alphanumeric	F	U
File Number	9-10	binary	F	U
Response Code	11-12	binary	--	A
ISN	13-16	binary	F	U *
ISN Lower Limit	17-20	binary	F	U
	21-24	--	--	--
Format Buffer Length	25-26	binary	F	U
Record Buffer Length	27-28	binary	F	U
	29-32	--	--	--
ISN Buffer Length **	33-34	binary	F	U
Command Option 1	35	alphanumeric	F	U
Command Option 2	36	alphanumeric	F	U
	37-44	--	--	--
Additions 2	45-48	binary / binary	--	A
Additions 3	49-56	alphanumeric	F	A
Additions 4	57-64	alphanumeric	F	A
Additions 5	65-72	alphanumeric	F	U
Command Time	73-76	binary	--	A
User Area	77-80	--	--	U

Buffer Areas

Buffer	Before Adabas Call	After Adabas Call
Format	F	U
Record	--	A
ISN **	--	A

where:

F	Supplied by user before Adabas call
A	Supplied by Adabas
U	Unchanged after Adabas call
*	Except for special options
**	The ISN buffer and length are required only if the multifetch or prefetch option is specified
--	Not used

Control Block Field Descriptions

We recommend that you set unused ACB fields to binary zeros before the direct call is initiated.

Command Code (ACBCMD)

L1 or L4

Command ID (ACBCID)

If a series of records is to be read with a series of L1 or L4 calls, and the same fields are to be specified in the format buffer for each call, this field should be set to a non-blank, non-zero value. This results in a reduction of the time required to process each L1 or L4 call.

If the GET NEXT option is to be used, the command ID of the ISN list to be used must be specified in this field. The format buffer may not be changed between successive L1 or L4 calls when the GET NEXT option is used.

If only a single record is to be read, or if the format buffer is to be modified between L1 or L4 calls, this field should be set to blanks.

If the command ID value is X'FFFFFFFF', automatic command ID generation will be in effect. In this case, the Adabas nucleus will generate values for command ID beginning with X'00000001', and will increment the value by 1 for each L1 or L4 call. When specifying user-defined command IDs, the user must ensure that each command ID is unique.

See also the Additions 5 field for separate format ID and/or global format ID usage.

File Number (ACBFNR)

Specify the binary number of the file to be read in this field. For physical direct calls, specify the file number as follows:

- For a one-byte file number, enter the file number in the rightmost byte (10); the leftmost byte (9), should be set to binary zero (B'0000 0000').
- For a two-byte file number, use both bytes (9 and 10) of the field.

Note:

When using two-byte file numbers and database IDs, a X'30' must be coded in the first byte of the control block.

Response Code (ACBRSP)

Adabas returns the response code for the command in this field. Response code 0 (ADARSP000) indicates that the command was executed successfully. Nonzero response codes, which can also have accompanying subcodes returned in the rightmost half of the Additions 2 field, are described in the *Adabas Messages and Codes Manual* documentation.

Response code 3 (ADARSP003) indicates an end-of-list condition (applicable only if the GET NEXT option is used).

ISN (ACBISN)

Specify the ISN of the record to be read.

If you specify the GET NEXT (N) option, Adabas selects ISNs from the ISN list identified by the command ID, reads the record of that ISN, and returns the ISN for the next record in this field; any specified ISN value is ignored.

If you specify the ISN sequence (I) option, you must also specify an ISN value in this field. The L1 or L4 command returns that ISN's data record in the record buffer.

If you specify the "F" option, the L1 or L4 command returns the next higher available ISN recorded in the file control block (FCB) in this field.

When a record is read, Adabas returns that record's ISN in this field, regardless of the option selected. With the option "F", this field returns the next higher unused ISN.

ISN Lower Limit: Multifetch Record Count (ACBISL)

If either "M" or "O" (multifetch option) is specified in the Command Option 1 field, a non-zero value in this field determines the maximum number of records to be multifetched. If this value is zero, the number of records to be multifetched is limited by the record and ISN buffer lengths. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

If a partial LOB field is requested by this command and the Command Option 2 field is set to L, the ACBISL field is used to keep track of the current position in the LOB value when multiple L1 or L4 calls are made for the field.

Format Buffer Length (ACBFBL)

The format buffer length (in bytes). The actual format buffer area defined in the user program must be at least as large as the length specified.

Record Buffer Length (ACBRBL)

The record buffer length (in bytes). The actual record buffer area defined in the user program must be at least as large as the length specified.

ISN Buffer Length: With Command-Level Multifetch/Prefetch Option Only (ACBIBL)

The ISN buffer length (in bytes). The ISN buffer defined in the user program must be at least as large as the length specified.

Command Option 1 (ACBCOP1)

Option	Description
F (first unused)	Returns the next highest unused ISN for the specified file in the ISN field. The <i>next unused ISN</i> is determined by referring to the file control block (FCB). Do not use the F option when reading Adabas expanded files.
M (multifetching)	Multifetch processing is performed for this command.
O (multifetching with the R option)	Multifetch processing and R option processing (see below) is performed for this command.
P (prefetching)	Prefetch processing is performed for this command.
R (return)	Returns response code 145 (ADARSP145) if the record to be read and held by an L4 command is not available.

Specifying the M, O, or P option indicates that the prefetch or multifetch option is to be used for the command. The selected option is active if either the ISN sequence (I) or GET NEXT (N) option is specified in the Command Option 2 field. The multifetch/prefetch option can improve performance by eliminating the time needed for single-record fetches. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

Command Option 2 (ACBCOP2)

Option	Description
F (first unused ISN)	Returns the next higher, unused, ISN for the specified file in the ISN field. The <i>next unused ISN</i> is determined by referring to the file control block (FCB). Do not use the F option when reading Adabas expanded files.
I (read by ISN sequence)	Reads the record identified by the ISN specified in the ISN field if the ISN is present in the file. If the ISN is not present in the file, the record with the next higher ISN is read and that record's ISN is returned in the ISN field. If the ISN is not present and no higher ISN is present in the file, no record is read and response code 3 (ADARSP003) is returned.
L (track LOB field position)	An L in the Command Option 2 field indicates that the position within a large object (LOB) field be tracked in the ACBISL (ISN lower limit) field. This option is useful when multiple L1/L4 commands are being issued in a series for a LOB field. This option can only be used if a LOB field is specified in the format buffer using LOB segment notation with asterisk notation in the <i>bytenum</i> specification.
N (GET NEXT)	Reads the records identified by the ISNs in an ISN list without the user having to provide the ISN of the record to be read with each L1 or L4 call. Adabas selects the ISN from the list and reads the record identified by that ISN. The ISN list to be used must be identified by the command ID field and must have been created previously by an Sx command. Response code 3 (ADARSP003) is returned when all the ISNs in the list have been selected. Refer to the section <i>ISN List Processing</i> for more information.

Additions 2: Length of Compressed and Decompressed Record (ACBADD2)

If the command is processed successfully, the following information is returned in this field:

- If the record buffer contains at least one valid field value, the leftmost two bytes contain the length (in binary form) of the compressed record accessed;
- The rightmost two bytes contain the length (in binary form) of the decompressed fields selected by the format buffer and accessed.

Note:

This length information is not returned when the prefetch feature is being used.

If the L1 or L4 command returns a nonzero response code, the rightmost two bytes may contain a subcode defining the exact response code meaning. Response codes and their subcodes are defined in the *Adabas Messages and Codes Manual* documentation.

Additions 3: Password (ACBADD3)

This field is used to provide an Adabas security password. If the database, file, or fields are security-protected, the user must provide a valid security password. Adabas sets the Additions 3 field to blanks during command processing to enhance password integrity.

Additions 4: Cipher Code (ACBADD4)

This field is used to provide a cipher code. If the file is ciphered, the user must provide a valid cipher code. If the file is not ciphered, this field should be set to blanks.

Adabas sets any cipher code to blanks during command processing, and returns a version code and database ID in the rightmost (low-order) three bytes of this field. For more information, see the section *Control Block Fields*.

Additions 5: Format ID, Global Format ID (ACBADD5)

This field may be used to provide a separate format ID which is to be used to identify the internal format buffer used for this command, or to provide a global format ID.

If the high-order (leftmost) bit of the Additions 5 field is not set to 1, the value provided in the command ID field is used as the format ID as well.

If, however, this bit is set to 1, the fifth through eighth bytes of the Additions 5 field are used as the format ID.

If the two high-order (leftmost) bits of the first byte of Additions 5 field are set to one (B'11'), all eight bytes of the Additions 5 field are used as a global format ID (that is, the format ID can be used by several users at the same time).

See the section *Command, Format, and Global Format IDs* for more information and examples.

ACB Examples

For the Adabas file definitions used in all the examples in this section, see *File Definitions Used in Examples*.

- Example 1: Reading a Single Record
- Example 2: Reading a Set of Records
- Example 3: Reading a Set of Records Using the GET NEXT Option
- Example 4: Read with Hold
- Example 5: Read Using the Read ISN Sequence Option
- Example 6: Reading Multiple-Value Fields and Periodic Groups

Example 1: Reading a Single Record

ISN 4 in file 1 is to be read. The values for fields AA and AB are to be returned.

Control Block

Command Code	L1	
Command ID	<i>bbbb</i> (blanks)	only 1 record is to be read
File Number	1	
ISN	4	
Format Buffer Length	6	or larger
Record Buffer Length	10	or larger
Command Option 2	b	GET NEXT or read ISN sequence options not used
Additions 3	<i>bbbbbbbb</i> (blanks)	file not security-protected
Additions 4	<i>bbbbbbbb</i> (blanks)	file is not ciphered

Buffer Areas

Format Buffer	AA,AB
----------------------	-------

Example 2: Reading a Set of Records

A set of records for which the ISNs have been obtained previously by a find command are to be read from file 2. The values for fields RA and XB are to be returned with the value for field XB to be returned with length 3 and format U.

Control Block

Command Code	L1	
Command ID	ABCD	a nonblank CID is recommended for a series of read operations in which the same fields are being read in each record
File Number	2	
ISN	n	ISNs are taken from the ISN list created by the find command*
Format Buffer Length	10	or larger
Record Buffer Length	11	or larger
Command Option 2	b	GET NEXT or read ISN sequence options not used
Additions 3	password	file is security-protected
Additions 4	<i>bbbbbbb</i> (blanks)	file is not ciphered

Buffer Areas

Format Buffer	RA,XB,3,U
----------------------	-----------

* n indicates an ISN from the ISN list which resulted from the find command. The L1 call is repeated for each ISN in the ISN list.

Example 3: Reading a Set of Records Using the GET NEXT Option

The requirement as stated for example 2 may also be satisfied by using the GET NEXT option.

Control Block

Command Code	L1	
Command ID	ABCD	CID of ISN list to be used
File Number	2	
ISN	0	the entire ISN list is to be selected starting with the first ISN in the list
Format Buffer Length	10	or larger
Record Buffer Length	11	or larger
Command Option 2	N	GET NEXT option to be used
Additions 3	password	file is security-protected
Additions 4	<i>bbbbbbb</i> (blanks)	file is not ciphered

Buffer Areas

Format Buffer	RA,XB,3,U
----------------------	-----------

The L1 call is repeated for each ISN in the ISN list. No changes to the control block are required between L1 calls. Response code 3 (ADARSP003) will be returned when all the ISNs in the list have been selected.

Example 4: Read with Hold

ISN 5 in file 2 is to be read and held for updating. The values for fields XC and XD are to be returned.

Control Block

Command Code	L4	read with hold
Command ID	<i>bbbb</i> (blanks)	only 1 record to be read
File Number	2	
ISN	5	
Format Buffer Length	6	or larger
Record Buffer Length	14	or larger
Command Option 1	b	response code 145 (ADARSP145) option not used
Command Option 2	b	GET NEXT or read ISN sequence options not used
Additions 3	password	file is security-protected
Additions 4	<i>bbbbbbb</i> (blanks)	file is not ciphered

Buffer Areas

Format Buffer	XC,XD
----------------------	-------

Example 5: Read Using the Read ISN Sequence Option

File 1 is to be read using the read ISN sequence option. The values for fields AA, AB, and AC are to be returned.

Control Block

Command Code	L1	
Command ID	BCDE	nonblank CID is recommended when a series of records for which the same fields are to be returned is to be read
File Number	1	
ISN	1	if ISN 1 is not present, the record of the next higher ISN in the file is read, and the ISN is returned in this field
Format Buffer Length	6	or larger
Record Buffer Length	30	or larger
Command Option 2	I	read ISN sequence option invoked
Additions 3	<i>bbbbbbb</i> (blanks)	file is not security-protected
Additions 4	<i>bbbbbbb</i> (blanks)	file is not ciphered

Buffer Areas

Format Buffer	GA,AC.
----------------------	--------

Adabas returns the ISN of the record which has been read in the ISN field of the control block. The record with the next higher ISN may be read by adding 1 to the ISN field and repeating the L1 command.

Example 6: Reading Multiple-Value Fields and Periodic Groups

The record identified by ISN 2 in file 1 is to be read. The value for field AA, all values for the multiple-value field MF, and all occurrences of the periodic group GB are to be returned.

Alternative 1: For Six or More Occurrences

1. Issue a L1 call to obtain the count of the number of values which exist for MF, and the highest occurrence count for GB.

Control Block Field

Command Code	L1	
Command ID	<i>bbbb</i> (blanks)	only 1 record is to be read
File Number	1	
ISN	2	
Format Buffer Length	8	or larger
Record Buffer Length	2	or larger
Command Option 2	b	GET NEXT or read ISN sequence option not used

Buffer Areas

Format Buffer	MFC,GBC
----------------------	---------

- Assuming that the result of the above L1 call was the record containing four values for MF and six occurrences for GB, repeat the L1 call using the following format buffer:

AA, MF01-04, GB01-06

For each call, the length of the format and record buffers must be large enough to hold all entries and values.

When using this procedure to read a series of records, a valid command ID should be used in step 1, and no command ID (blanks) should be used for step 2 since the content of the format buffer may vary with each step 2 call.

Alternative 2: For Fewer Than Six Occurrences

An alternative solution for example 5 usually provides better performance if the number of values/occurrences is small (less than six) in a large percentage of the records.

- Issue a L1 call in which the counts for MF and GB are requested, plus the expected number of values and occurrences of MF and GB, plus field AA.

Assuming that the expected number of values for MF is 2, and the expected number of occurrences of GB is 3, the format buffer for step 1 would be

AA, MFC, GBC, MF01-02, GB01-03.*

* Maximum performance is normally achieved if a number which will retrieve all of the values/occurrences in 90 per cent of the records is specified.

- If either the count received for MF exceeds 2 or the count received for GB exceeds 3, a format buffer similar to that in step 2 of the previous example could be used to obtain the additional values and/or occurrences. Otherwise, no additional call is required.

ACBX Interface Direct Call: L1 and L4 Commands

This section describes ACBX interface direct calls for the L1 and L4 commands. It covers the following topics:

- Control Block and Buffer Overview
- Control Block Field Descriptions

Control Block and Buffer Overview

Control Block

Field	Position	Format	Before Adabas Call	After Adabas Call
	1-2	---	---	---
Version Indicator	3-4	binary	F	U
	5-6	---	---	---
Command Code	7-8	alphanumeric	F	U
	9-10	---	---	---
Response Code	11-12	binary	---	A
Command ID	13-16	alphanumeric/ binary	F	U
Database ID	17-20	numeric	F	U
File Number	21-24	numeric	F	U
	25-28	---	---	---
ISN	29-32	binary	F	U*
	33-36	---	---	---
ISN Lower Limit	37-40	binary	F	U
	41-48	---	---	---
Command Option 1	49	alphanumeric	F	U
Command Option 2	50	alphanumeric	F	U
Command Option 3 (L4 only)	51	alphanumeric	F	U
	52-64	---	---	---
Additions 2	65-68	binary	---	A
Additions 3	69-76	alphanumeric/ binary	F	A
Additions 4	77-84	alphanumeric	F	A

Field	Position	Format	Before Adabas Call	After Adabas Call
Additions 5	85-92	alphanumeric/ binary	F	U
	93-114	---	---	---
Error Subcode	115-116	binary	---	A
	117-144	---	---	---
Command Time	145-152	binary	---	A
User Area	153-168	not applicable	---	U
---	169-193	do not touch	---	---

ABDs and Buffers

ABD and Buffer	Before Adabas Call	After Adabas Call
Format	F	U
Record	--	A
Multifetch **	--	A

where:

F Supplied by user before Adabas call

A Supplied by Adabas

U Unchanged after Adabas call

* Except for special options

** The multifetch buffer is required only if the multifetch option is specified.

--- Not used

Control Block Field Descriptions

We recommend that you set unused ACBX fields to binary zeros before the direct call is initiated.

Version Indicator (ACBXVER)

F2

Command Code (ACBXCMD)

L1 or L4

Response Code (ACBXRSP)

Adabas returns the response code for the command in this field. Response code 0 (ADARSP000) indicates that the command was executed successfully. Non-zero response codes, which can also have accompanying subcodes returned in the Error Subcode (ACBXERRC) field, are described in the *Adabas Messages and Codes Manual* documentation.

Response code 3 (ADARSP003) indicates an end-of-list condition (applicable only if the GET NEXT option is used).

Command ID (ACBXCID)

If a series of records is to be read with a series of L1 or L4 calls, and the same fields are to be specified in the format buffer for each call, this field should be set to a non-blank, non-zero value. This results in a reduction of the time required to process each L1 or L4 call.

If the GET NEXT option is to be used, the command ID of the ISN list to be used must be specified in this field. The format buffer may not be changed between successive L1 or L4 calls when the GET NEXT option is used.

If only a single record is to be read, or if the format buffer is to be modified between L1 or L4 calls, this field should be set to blanks.

If the command ID value is X'FFFFFFFF', automatic command ID generation will be in effect. In this case, the Adabas nucleus will generate values for command ID beginning with X'00000001', and will increment the value by 1 for each L1 or L4 call. When specifying user-defined command IDs, the user must ensure that each command ID is unique.

See also the Additions 5 field for separate format ID and/or global format ID usage.

Database ID (ACBXDBID)

Use this field to specify the database ID. The Adabas call will be directed to this database.

This field is a four-byte binary field, but at this time only two-byte database IDs are supported. Therefore, the database ID should be specified in the low-order part (rightmost bytes) of the field, with leading binary zeros.

If this field is set to binary zeros, the Adabas API uses either the database ID from the ADARUN cards provided in DDCARD input data or the default database ID value provided in the LNKGBLS module linked with or loaded by the link routine.

File Number (ACBXFNR)

Use this field to specify the number of the file to which the Adabas call should be directed.

This field is a four-byte binary field, but the file number should be specified in the low-order part (rightmost bytes) of the field, with leading binary zeros.

ISN (ACBXISN)

Use this field to specify the ISN of the record to be read.

If you specify the GET NEXT (N) option, Adabas selects ISNs from the ISN list identified by the command ID, reads the record of that ISN, and returns the ISN for the next record in this field; any specified ISN value is ignored.

If you specify the ISN sequence (I) option, you must also specify an ISN value in this field. The L1 or L4 command returns that ISN's data record in the record buffer.

If you specify the F option, the L1 or L4 command returns the next higher available ISN recorded in the file control block (FCB) in this field.

When a record is read, Adabas returns that record's ISN in this field, regardless of the option selected. With the option "F", this field returns the next higher unused ISN.

The ACBXISN field is a four-byte binary field embedded in the eight-byte ACBXISNG field, which is not yet used. Set the high-order part of the ACBXISNG field to binary zeros.

ISN Lower Limit: Multifetch Record Count (ACBXISL)

If either "M" or "O" (multifetch option) is specified in the Command Option 1 field, a non-zero value in this field determines the maximum number of records to be multifetched. If this value is zero, the number of records to be multifetched is limited by the record and ISN buffer lengths. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

If a partial LOB field is requested by this command and the Command Option 2 field is set to L, the ACBXISL field is used to keep track of the current position in the LOB value when multiple L1 or L4 calls are made for the field.

The ACBXISL field is a four-byte binary field embedded in the eight-byte ACBXISLG field, which is not yet fully used. Set the high-order part of the ACBXISLG field to binary zeros.

Command Option 1 (ACBXCOP1)

Option	Description
F (first unused)	Returns the next highest unused ISN for the specified file in the ISN field. The <i>next unused ISN</i> is determined by referring to the file control block (FCB). Do not use the F option when reading Adabas expanded files.
M (multifetching)	Multifetch processing is performed for this command.
O (multifetching with the R option)	Multifetch processing and R option processing (see below) is performed for this command.
R (return)	Returns response code 145 (ADARSP145) if the record to be read and held by an L4 command is not available.

Specifying the "M" or "O" option indicates that the multifetch option is to be used for the command. The selected option is active if either the ISN sequence (I) or GET NEXT (N) option is specified in the Command Option 2 field. The multifetch option can improve performance by eliminating the time needed for single-record fetches. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

Command Option 2 (ACBXCOP2)

Option	Description
F (first unused ISN)	Returns the next higher, unused, ISN for the specified file in the ISN field. The <i>next unused ISN</i> is determined by referring to the file control block (FCB). Do not use the F option when reading Adabas expanded files.
I (read by ISN sequence)	Reads the record identified by the ISN specified in the ISN field if the ISN is present in the file. If the ISN is not present in the file, the record with the next higher ISN is read and that record's ISN is returned in the ISN field. If the ISN is not present and no higher ISN is present in the file, no record is read and response code 3 (ADARSP003) is returned.
L (track LOB position)	An L in the Command Option 2 field indicates that the position within a large object (LOB) field be tracked in the ACBXISL (ISN lower limit) field. This option is useful when multiple L1/L4 commands are being issued in a series for a LOB field. This option can only be used if a LOB field is specified in the format buffer using LOB segment notation with asterisk notation in the <i>bytenum</i> specification.
N (GET NEXT)	Reads the records identified by the ISNs in an ISN list without the user having to provide the ISN of the record to be read with each L1 or L4 call. Adabas selects the ISN from the list and reads the record identified by that ISN. The ISN list to be used must be identified by the command ID field and must have been created previously by an <i>Sx</i> command. Response code 3 (ADARSP003) is returned when all the ISNs in the list have been selected. Refer to the section <i>ISN List Processing</i> for more information.

Command Option 3: Shared Hold Status (ACBXCOP3)

The following command options are available for L4 commands only.

Option	Description
C	Puts the record in shared hold status for the duration of the read operation.
Q	Puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated, whichever happens first. This option is available for the L4 command when command option 2 is set to "N".
S	Puts the record in shared hold status until the end of the transaction.

If the same record is placed in shared hold status more than once (using the C or S options or the Q option for different read sequences), it stays in shared hold status until all of the specified hold lifetimes have expired.

For complete information about shared hold updating, read *Shared Hold Status*.

Additions 2: Length of Compressed and Decompressed Record (ACBXADD2)

If the command is processed successfully, the following information is returned in this field:

- If the record buffer contains at least one valid field value, the leftmost two bytes contain the length (in binary form) of the compressed record accessed;
- The rightmost two bytes contain the length (in binary form) of the decompressed fields selected by the format buffer and accessed.

Note:

This length information is not returned when the prefetch feature is being used.

Additions 3: Password (ACBXADD3)

This field is used to provide an Adabas security password. If the database, file, or fields are security-protected, the user must provide a valid security password. Adabas sets the Additions 3 field to blanks during command processing to enhance password integrity.

Additions 4: Cipher Code (ACBXADD4)

This field is used to provide a cipher code. If the file is ciphered, the user must provide a valid cipher code. If the file is not ciphered, this field should be set to blanks.

Adabas sets any cipher code to blanks during command processing, and returns a version code and database ID in the rightmost (low-order) three bytes of this field. For more information, see the section *Control Block Fields*.

Additions 5: Format ID, Global Format ID (ACBXADD5)

This field may be used to provide a separate format ID which is to be used to identify the internal format buffer used for this command, or to provide a global format ID.

If the high-order (leftmost) bit of the Additions 5 field is not set to 1, the value provided in the command ID field is used as the format ID as well.

If, however, this bit is set to 1, the fifth through eighth bytes of the Additions 5 field are used as the format ID.

If the two high-order (leftmost) bits of the first byte of Additions 5 field are set to one (B'11'), all eight bytes of the Additions 5 field are used as a global format ID (that is, the format ID can be used by several users at the same time).

See the section *Command, Format, and Global Format IDs* for more information and examples.

Error Subcode (ACBXERRC)

If the command returns a nonzero response code, this field contains a subcode defining the exact response code meaning. Response codes and their subcodes are defined in the *Adabas Messages and Codes Manual* documentation.

Buffers

The following buffers apply to L1 and L4 commands:

- Format Buffer
- Record Buffer
- ISN Buffer
- Multifetch Buffer

Format Buffer

Qualify the fields to be read in this buffer. Descriptions of the syntax and examples of format buffer construction are provided in *Defining Buffers*.

A "C." in the first two positions of the format buffer causes the record to be returned in compressed instead of decompressed format.

Record Buffer

Adabas returns the requested field values in this buffer. All values are returned according to the standard format and length of the field unless the user specifies a different length or format in the format buffer.

ISN Buffer

The ISN buffer is used only if the command-level multifetch option is requested and only if the direct call is made using the ACB interface. For more information, read *Using the Multifetch/Prefetch Feature*. When multifetching is used, the L1 and L4 commands return descriptor elements, each up to 16 bytes long, in the ISN buffer (see the section *BT/ET Multifetch Processing*).

Multifetch Buffer

The multifetch buffer is used only if the command-level multifetch option is requested and only if the direct call is made using the ACBX interface. For more information, read *Using the Multifetch/Prefetch Feature*. When multifetching is used, the L1 and L4 commands return descriptor elements, each up to 16 bytes long, in the multifetch buffer (see the section *BT/ET Multifetch Processing*).