

# L2 and L5 Commands: Read Physical Sequential

The L2 and L5 commands read records in the sequence in which they are physically located in 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: L2 and L5 Commands
  - ACBX Interface Direct Call: L2 and L5 Commands
  - Buffers
  - Additional Considerations
- 

## Function and Use

The L5 command performs the same function as the L2 command, but places each record read in hold status. If the record to be read and held is currently being held by another user, the user will be placed in wait status until the record becomes available. If the L5 command was issued with the Command Option 1 field set to "R", Adabas returns response code 145 if the record is not available.

The L2 and L5 commands do not read records in any particular logical order unless the records were loaded initially in a particular logical sequence and no subsequent update changed this order.

The L2 and L5 commands may be used to read an entire file at optimum speed since no access is required to the Associator (as with the L3 command), and all physical blocks are read in consecutive sequence.

Specify the file to be read and the fields within each record for which values are to be returned. Specify the fields in the format buffer. Adabas returns the requested field values in the record buffer.

The multifetch/prefetch option allows prior accessing of one or more sequential records, reducing overall operating time and eliminating the time needed for single-record fetches. Multifetching or prefetching can be enabled by specifying "M", "O" (for multifetching), or "P" (for prefetching) in the Command Option 1 field. For more information, read *Using the Multifetch/Prefetch Feature*.

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

## **ACB Interface Direct Call: L2 and L5 Commands**

This section describes ACB interface direct calls for the L2 and L5 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	A
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
	36-44	--	--	--
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 *	F	U

where:

F	Supplied by user before Adabas call
A	Supplied by Adabas
U	Unchanged after Adabas call
*	The ISN buffer and length are required only if the multifetch or prefetch option is specified
**	May contain compress option control characters "C."
--	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)

L2 or L5

### Command ID (ACBCID)

This field must be set to a non-blank, non-zero value. It is used by Adabas to provide the records in the correct physical order and to avoid the repetitive interpretation of the format buffer. The value provided must not be modified during any given sequential pass of a file.

The first byte of this field may not be set to hexadecimal 'FF'.

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 L2 or L5 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 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 indicates that the command was executed successfully. Non-zero 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 indicates an end-of-file (EOF) condition was detected.

### **ISN (ACBISN)**

If this field is set to zero before the initial L2 or L5 call, the sequential pass will begin with the first record contained in the first physical block of the file.

If this field is set to an ISN value before the initial L2 or L5 call, the sequential pass will begin at the first record physically located after the record identified by the ISN specified. The ISN specified must be present in the file.

This field need not be modified by the user after the initial L2 or L5 call. Adabas returns the ISN of the record which has been read in this field.

### **ISN Lower Limit: Multifetch Record Count (ACBISL)**

If either "M" or "O" (multifetching 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.

### **Format Buffer Length (ACBFBL)**

The format buffer length (in bytes). The format buffer area defined in the user program must be as large as (or larger than) the length specified. If either multifetch option "M" or "O" is specified in the Command Option 1 field, this value must be less than 32 kilobytes.

### **Record Buffer Length (ACBRBL)**

The record buffer length (in bytes). The record buffer area defined in the user program must be as large as (or larger than) the length specified. If either multifetch option "M" or "O" is specified in the Command Option 1 field, this value must be less than 32 kilobytes.

### **ISN Buffer Length: Only with Command-Level Multifetch/Prefetch Option (ACBIBL)**

The ISN buffer length (in bytes). The ISN buffer area defined in the user program must be as large as (or larger than) the length specified.

### **Command Option 1 (ACBCOP1)**

To improve performance by eliminating the time needed for single-record fetches, set option M, O, or P to enable multifetch or prefetch processing for the command. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

Option	Description
M	Enable multifetching
O	Use multifetching with the "R" option described below
P	Use prefetching
R (return)	Returns response code 145 if the record to be read and held by an L5 command is not available.

### 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 L2 or L5 command returns a non-zero 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 or Adabas SAF 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.

As long as the high-order (leftmost) bit of the first byte of the Additions 5 field is not set to 1, the value provided in the command ID field will be 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 ID, Format ID, Global Format ID* for more information and examples.

## ACB Examples

### Example 1

File 2 is to be read in physical sequential order. All the values for all the fields within each record are to be returned.

#### Control Block

<b>Command Code</b>	L2	
<b>Command ID</b>	EXL2	nonblank CID required
<b>File Number</b>	2	
<b>ISN</b>	0	all records are to be read
<b>Format Buffer Length</b>	3	or larger
<b>Record Buffer Length</b>	49	or larger
<b>Command Option 1</b>	b	return option not used
<b>Additions 3</b>	password	file 2 is security-protected
<b>Additions 4</b>	bbbbbbbb (blanks)	file is not ciphered

#### Buffer Areas

<b>Format Buffer</b>	RG
----------------------	----

The L2 call is repeated to obtain each successive record. The ISN field need not be modified between calls.

### Example 2

File 2 is to be read in physical sequential order. The values for fields RA, XA, and XB (3 bytes unpacked) are to be returned. Each record read is to be placed in hold status for updating purposes.

#### Control Block

<b>Command Code</b>	L5	
<b>Command ID</b>	EXL5	nonblank CID is required
<b>File Number</b>	2	
<b>ISN</b>	0	all records are to be read
<b>Format Buffer Length</b>	13	or larger
<b>Record Buffer Length</b>	21	or larger
<b>Command Option 1</b>	b	response code 145 option not used
<b>Additions 3</b>	password	file 2 is security-protected
<b>Additions 4</b>	bbbbbbbb (blanks)	file 2 is not ciphered

**Buffer Areas**

<b>Format Buffer</b>	RA,XA,XB,3,U
----------------------	--------------

The L5 call is repeated to obtain each successive record. The ISN field need not be modified between L5 calls.

To complete logical transactions and release held records, ET logic users should issue an ET command. Non-ET users should release held records with an A4, E4, or RI command.

## ACBX Interface Direct Call: L2 and L5 Commands

This section describes ACBX interface direct calls for the L2 and L5 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



Field	Position	Format	Before Adabas Call	After Adabas Call
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	A
	33-36	---	---	---
ISN Lower Limit	37-40	binary	F	U
	41-48	---	---	---
Command Option 1	49	alphanumeric	F	U
	50-64	---	---	---
Additions 2	65-68	binary	---	A
Additions 3	69-76	alphanumeric/ binary	F	A
Additions 4	77-84	alphanumeric	F	A
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 *	F	U

where:

F	Supplied by user before Adabas call
A	Supplied by Adabas
U	Unchanged after Adabas call
*	A multifetch buffer is required only if the multifetch option is specified.
**	May contain compress option control characters "C."
---	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)

L2 or L5

### Response Code (ACBXRSP)

Adabas returns the response code for the command in this field. Response code 0 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 indicates an end-of-file (EOF) condition was detected.

### Command ID (ACBXCID)

This field must be set to a non-blank, non-zero value. It is used by Adabas to provide the records in the correct physical order and to avoid the repetitive interpretation of the format buffer. The value provided must not be modified during any given sequential pass of a file.

The first byte of this field may not be set to hexadecimal 'FF'.

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 L2 or L5 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 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)

If this field is set to zero before the initial L2 or L5 call, the sequential pass will begin with the first record contained in the first physical block of the file.

If this field is set to an ISN value before the initial L2 or L5 call, the sequential pass will begin at the first record physically located after the record identified by the ISN specified. The ISN specified must be present in the file.

This field need not be modified by the user after the initial L2 or L5 call. Adabas returns the ISN of the record which has been read in this field.

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" (multifetching 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 multifetch buffer lengths. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

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

### Command Option 1 (ACBXCOP1)

To improve performance by eliminating the time needed for single-record fetches, set option M, O, or P to enable multifetch or prefetch processing for the command. Refer to the section *Using the Multifetch/Prefetch Feature* for more information.

Option	Description
M	Enable multifetching
O	Enable multifetching with the "R" option, described below
R (return)	Returns response code 145 if the record to be read and held by an L5 command is not available.

### 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 or Adabas SAF 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.

As long as the high-order (leftmost) bit of the first byte of the Additions 5 field is not set to 1, the value provided in the command ID field will be 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 ID, Format ID, Global Format ID* 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 L2 and L5 commands:

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

### Format Buffer

The fields for which values are to be returned are specified 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 indicates that the compressed option is to be in effect. This causes the record to be returned in compressed instead of decompressed format.

The format buffer may not be modified after the initial L2 or L5 call has been issued.

### Record Buffer

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

### ISN Buffer

The ISN buffer is used only if the command-level multifetch option is used and if the command is issued using an ACB interface direct call. For more information, read *Using the Multifetch/Prefetch Feature*. When multifetching is used, the L2 or L5 command returns record 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 used and if the command is issued using an ACBX interface direct call. For more information, read *Using the Multifetch/Prefetch Feature*. When multifetching is used, the L2 or L5 command returns record descriptor elements, each up to 16 bytes long, in the multifetch buffer (see the section *BT/ET Multifetch Processing*).

## Additional Considerations

The following additional considerations are applicable for the L2 and L5 commands:

1. The command ID used with an L2 or L5 command is saved internally and used by Adabas. It is released by Adabas when an end-of-file condition is detected, an RC or CL command is issued, or the Adabas session is terminated. The same command ID may not be used by the user for another read sequential command until it has been released.

2. You are permitted to update or delete records from a file that you are reading with an L2 or L5 command. Adabas maintains information about the last and next record to be provided to you, and is able to provide the correct next record despite any interim record update or deletion you may perform.
3. If another user is updating the file being read with an L2 or L5 command, it is possible that you will not receive one or more records in the file when reading it with the L2 or L5 command. In addition, you may receive the same record more than once during a given sequential pass of the file.