

L1 / L4 Command: Read / Read and Hold Record

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

This chapter covers the following topics:

- Function and Use
 - Command: L1 / L4
 - Control Block
 - Buffers
 - Examples
-

Function and Use

The user specifies the file number and ISN of the record to be read. The user indicates 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, and, in addition, 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 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* on page 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/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.

Command: L1 / L4

User 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	--
COMMAND TIME	73-76	binary	--	A
USER AREA	77-80	--	--	U

User Buffer Areas

Buffer	Before Adabas Call	After Adabas Call
FORMAT BUFFER	F	U
RECORD BUFFER	--	A
ISN BUFFER **	--	A

where:

- F Filled in by user before Adabas Call
- A Filled in by Adabas
- U Unchanged after Adabas call
- * Except for special options
- ** The ISN buffer and length required only if the multi-/prefetch option is specified
- Not used

Control Block

Command Code

L1/L4

Command ID

If a series of records is to be read with a series of L1/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/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/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/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/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

Specify the binary number of the file to be read in this field. For the 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

Adabas returns the response code for the command in this field. Response code 0 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* documentation.

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

ISN

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/L4 command returns that ISN's data record in the record buffer.

If you specify the "F" option, the L1/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

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.

Format Buffer Length

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

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

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

Option	Description
F	(first unused ISN) returns the next highest unused ISN for the specified file in the ISN field. The "next unused ISN" is determined by referring to the file control block (FCB). Do not use the "F" option when reading Adabas expanded files.
M	(multifetching)
O	(multifetching with the "R" option)
P	(prefetching)
	Specifying one of these options 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 <i>Using the Multifetch/Prefetch Feature</i> for more information.
R	(return) returns response code 145 if the record to be read and held by an L4 command is not available.

Command Option 2

Option	Description
F	(first unused ISN) returns the next higher unused ISN for the specified file in the ISN field. The "next unused ISN" 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 is returned.
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/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 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

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* documentation.

Additions 3: Password

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

Additions 4: Cipher Code

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

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 the leftmost bit is set to 1, The rightmost four (fifth through eighth) bytes of the additions 5 field (additions 5 + 4(4)) will be 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'), the value in all eight bytes of the additions 5 field (additions 5 + 0(8)) is 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.

Buffers

Format Buffer

The user specifies the fields to be read in this buffer. format buffer syntax and examples are provided in the section *Adabas Calling Procedure*.

"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 and/or format in the format buffer.

ISN Buffer

The ISN buffer is used only if the command-level multifetch or prefetch option is used. See the section *Using the Multifetch/Prefetch Feature* for more information. When multifetching is used, the L1/L4 command returns descriptor elements, each up to 16 bytes long, in the ISN buffer (see the section *BT/ET Multifetch Processing*).

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

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	bbbb	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	bbbbbbbb	file not security-protected
Additions 4	bbbbbbbb	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	bbbbbbbb	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	bbbbbbbb	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 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	bbbb	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 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	bbbbbbbb	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	bbbbbbb	file is not security-protected
Additions 4	bbbbbbb	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	bbbb	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
----------------------	---------

2. 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.

1. 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.

2. 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.