NDB - SELECT - Cursor-Oriented

Like the Natural FIND statement, the cursor-oriented SELECT statement is used to select a set of rows (records) from one or more DB2 tables, based on a search criterion. Since a database loop is initiated, the loop must be closed by a LOOP (reporting mode) or END-SELECT statement. With this construction, Natural uses the same loop processing as with the FIND statement.

In addition, no cursor management is required from the application program; it is automatically handled by Natural.

Below is information on:

- OPTIMIZE FOR integer ROWS

- WITH - Isolation Level

- QUERYNO

- FETCH FIRST

- WITH HOLD

- WITH RETURN

- WITH INSENSITIVE/SENSITIVE

OPTIMIZE FOR integer ROWS

```
[ OPTIMIZE FOR integer ROWS ]
```

The OPTIMIZE FOR *integer* ROWS clause is used to inform DB2 in advance of the number (*integer*) of rows to be retrieved from the result table. Without this clause, DB2 assumes that all rows of the result table are to be retrieved and optimizes accordingly.

This optional clause is useful if you know how many rows are likely to be selected, because optimizing for *integer* rows can improve performance if the number of rows actually selected does not exceed the *integer* value (which can be in the range from 0 to 2147483647).

Example:

```
SELECT name INTO
#name FROM table WHERE AGE = 2 OPTIMIZE FOR 100 ROWS
```

WITH - Isolation Level

```
WITH   CS
       RR
       RR KEEP UPDATE LOCK
       RS
       RS KEEP UPDATE LOCKS
       UR
```

This WITH clause allows you to specify an explicit isolation level with which the statement is to be executed. The following options are provided:

Option	Meaning
CS	Cursor stability
RR	Repeatable Read
RS	Read Stability
RS KEEP UPDATE LOCKS	Only valid if a FOR UPDATE OF clause is specified. Read Stability and retaining update locks.
RR KEEP UPDATE LOCKS	Only valid if a FOR UPDATE OF clause is specified. Repeatable Read and retaining update locks.
UR	Uncommitted Read

WITH UR can only be specified within a SELECT statement and when the table is read-only. The default isolation level is determined by the isolation of the package or plan into which the statement is bound. The default isolation level also depends on whether the result table is read-only or not. To find out the default isolation level, refer to the IBM literature.

Note:
This option also works for non-cursor selection.

QUERYNO

```
[ QUERYNO integer ]
```

The QUERNO clause specifies the number to be used for this SQL statement in EXPLAIN output and trace records. The number is used as QUERYNO column in the PLAN_TABLE for the rows that contain information on this statement.

FETCH FIRST

```
┌                                                    ┐
│              ┌  1       ┐ ┌ ROWS ┐                  │
│ FETCH FIRST  │          │ │      │  ONLY            │
│              └ integer  ┘ └ ROW  ┘                  │
│                                                    │
└                                                    ┘
```

The FETCH FIRST clause limits the number of rows to be fetched. It improves the performance of queries with potentially large result sets if only a limited number of rows is needed.

WITH HOLD

```
[ WITH HOLD ]
```

The WITH HOLD clause is used to prevent cursors from being closed by a commit operation within database loops. If WITH HOLD is specified, a commit operation commits all the modifications of the current logical unit of work, but releases only locks that are not required to maintain the cursor. This optional clause is mainly useful in batch mode; it is ignored in CICS pseudo-conversational mode and in IMS message-driven programs.

Example:

```
SELECT name INTO #name FROM table
WHERE AGE = 2 WITH HOLD
```

WITH RETURN

```
[ WITH RETURN ]
```

The WITH RETURN clause is used to create result sets. Therefore, this clause only applies to programs which operate as Natural stored procedure. If the WITH RETURN clause is specified in a SELECT statement, the underlying cursor remains open when the associated processing loop is left, except when the processing loop had read all rows of the result set itself. During first execution of the processing loop, only the cursor is opened. The first row is not yet fetched. This allows the Natural program to return a full result set to the caller of the stored procedure. It is up to the Natural you to decide how many rows are processed by the Natural stored procedure and how many unprocessed rows of the result set are returned to the caller of the stored procedure. If you want to process rows of the select operation in the Natural stored procedure, you must define

```
IF *counter =1 ESCAPE TOP END-IF
```

in order to avoid processing of the first "empty row" in the processing loop. If you decide to terminate the processing of rows, you must define

```
If condition ESCAPE BOTTOM END-IF
```

in the processing loop.

If the program reads all rows of the result set, the cursor is closed and no result set is returned for this SELECT WITH RETURN to the caller of the stored procedure.

The following programs are examples for retrieving full result sets (Example 1) and partial result sets (Example 2).

Example 1:

```
DEFINE DATA LOCAL
. . .
END-DEFINE
*
*  Return all rows of the result set
*
SELECT * INTO VIEW V2
                   FROM SYSIBM-SYSROUTINES
                   WHERE RESULT_SETS > 0
                   WITH RETURN
ESCAPE BOTTOM
END-SELECT
END
```

Example 2:

```
DEFINE DATA LOCAL
. . .
END-DEFINE
*
*  Read the first two rows and return the rest as result set
*
SELECT * INTO VIEW V2
                   FROM SYSIBM-SYSROUTINES
                   WHERE RESULT_SETS > 0
                   WITH   RETURN
WRITE PROCEDURE *COUNTER
IF *COUNTER = 1 ESCAPE TOP END-IF
IF *COUNTER =3 ESCAPE BOTTOM END-IF
END-SELECT
END
```

WITH INSENSITIVE/SENSITIVE

```
┌                                                                           ┐
│ WITH ⎧ ASENSITIVE SCROLL      ⎫ [:] scroll_hv [GIVING [:] sqlcode]        │
│      ⎨ INSENSITIVE SCROLL     ⎬                                           │
│        SENSITIVE STATIC SCROLL                                            │
│        SENSITIVE DYNAMIC SCROLL                                           │
└                                                                           ┘
```

NDB supports DB2 scrollable cursors by using the clauses WITH ASENSITIVE SCROLL, WITH SENSITIVE STATIC SCROLL and SENSITVE DYNAMIC SCROLL. Scrollable cursors allow NDB applications to position randomly any row in a result set. With non-scrollable cursors, the data can only be read sequentially, from top to bottom.

ASENSITIVE scrollable cursors are either INSENSITIVE - if the cursor is READ-ONLY- or SENSITIVE DYNAMIC - if the cursor is not READ-ONLY .

INSENSITIVE and SENSITIVE STATIC scrollable cursors use temporary result tables and require a TEMP database in DB2 (see the relevant DB2 literature by IBM).

INSENSITIVE SCROLL refers to a cursor that cannot be used in Positioned UPDATE or Positioned DELETE operations. In addition, once opened, an INSENSITIVE SCROLL cursor does not reflect UPDATEs, DELETEs or INSERTs against the base table, after the cursor was opened.

SENSITIVE STATIC SCROLL refers to a cursor that can be used for Positioned UPDATEs or Positioned DELETE operations. In addition, a SENSITIVE STATIC SCROLL cursor reflects UPDATEs, DELETEs of base table rows. The cursor does not reflect INSERT operations.

SENSITIVE DYNAMIC scrollable cursors reflect UPDATEs, DELETEs and INSERTs against the base table while the cursor is open.
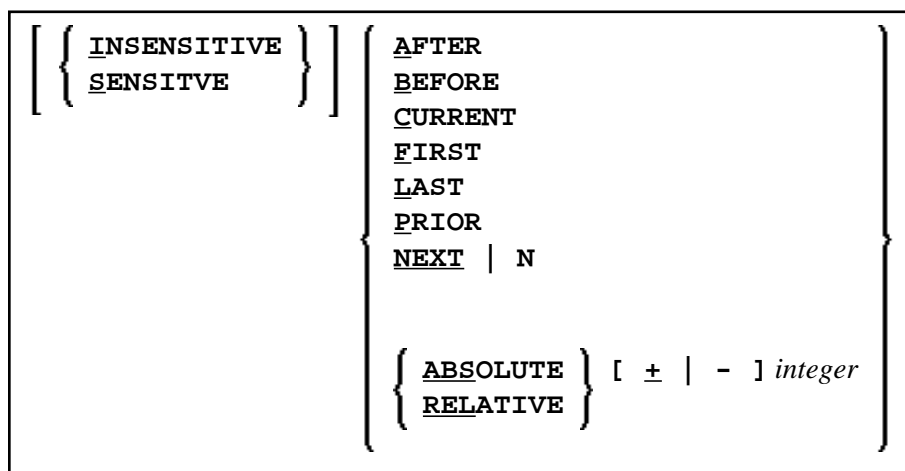
Below is information on:

- scroll_hv

- GIVING [:] sqlcode

scroll_hv

The variable *scroll_hv* must be alphanumeric.

The variable *scroll_hv* specifies which row of the result table will be fetched during one execution of the database processing loop. Additionally, it specifies the sensitivity of UPDATEs or DELETEs against the base table row during a FETCH operation. The contents of *scroll_hv* is evaluated each time the database processing loop cycle is executed.

```
[ { INSENSITIVE } ] { AFTER                              }
  { SENSITVE     }   { BEFORE                             }
                     { CURRENT                            }
                     { FIRST                              }
                     { LAST                               }
                     { PRIOR                              }
                     { NEXT | N                           }
                     {                                    }
                     {                                    }
                     { { ABSOLUTE } [ + | - ] integer     }
                     { { RELATIVE }                       }
```

scroll_hv - Sensitivity Specification

The specification of the sensitivity INSENSITIVE or SENSITIVE is optional.

If it is omitted from a FETCH against an INSENSITIVE SCROLL cursor, the default will be INSENSITIVE.

If it is omitted from a FETCH against a SENSITIVE STATIC/DYNAMIC SCROLL cursor, the default will be SENSITIVE.

The sensitivity specifies whether or not the rows in the base table are checked when performing a FETCH operation for a scrollable cursor.

If the corresponding base table column qualifies for the WHERE clause and has not been deleted, a SENSITIVE FETCH will return the row of the base table.

If the corresponding base table column does not qualify for the WHERE clause or has not been deleted, a SENSITIVE FETCH will return an UPDATE hole or a DELETE hole state (SQLCODE +222).

An INSENSITIVE FETCH will not check the corresponding base table column.

scroll_hv - **Options**

Below is an explanation of the options available to determine the row(s) to fetch, the position from where to start the fetch and/or the direction in which to scroll:

Option	Explanation
AFTER	Positions after the last row. No row is fetched.
BEFORE	Positions before the first. No row is fetched.
CURRENT	Fetches the current row (again).
FIRST	Fetches the first row.
LAST	Fetches the last row.
NEXT	Fetches the row after the current one. This is the default value.
PRIOR	Fetch the row before the current one.
+/- *integer*	Only applies in connection with ABSOLUTE or RELATIVE. Specifies the position of the row to be fetched ABSOLUTE or RELATIVE. Enter a plus (+) or minus (-) sign followed by an integer. The default value is a plus (+).
ABSOLUTE	Only applies in connection with +/- *integer*. Uses *integer* as the absolute position within the result set from where the row is fetched. See the DB2 SQL reference by IBM about further details regarding positive and negative position numbers.
RELATIVE	Only applies in connection with +/- *integer*. Uses *integer* as the relative position to the current position within the result set from where the row is fetched. See the DB2 SQL reference by IBM about further details regarding positive and negative position numbers.

GIVING [:] sqlcode

The specification of GIVING [:] *sqlcode* is optional. If specified, the Natural variable [:] *sqlcode* must be of the Format I4. The values for this variable are returned from the DB2 SQLCODE of the underlying FETCH operation. This allows the application to react to different statuses encountered while the scrollable cursor is open. The most important status codes indicated by SQLCODE are listed in the following table:

SQLCODE	Explanation
0	FETCH operation successful, data returned except for FETCH with option BEFORE or AFTER.
+100	Row not found, cursor still open, no data returned.
+222	UPDATE or DELETE hole, cursor still open, no data returned. The corresponding row of the base table has been updated or deleted, so that the row no longer qualifies for the WHERE clause.
+231	Fetch operation with the option CURRENT, but cursor not positioned on any row, no data returned. This occurs if the previous FETCH returned SQLCODE +100.

If you specify GIVING [:] *sqlcode*, the application must react to the different statuses. If an SQLCODE +100 is entered five times successively and without terminal I/O, the NDB runtime will issue Natural Error NAT3296 in order to avoid application looping. The application can terminate the processing loop by executing an ESCAPE statement.

If you do not specify GIVING [:] *sqlcode*, except for SQLCODE 0 and SQLCODE +100, each SQLCODE will generate Natural Error NAT3700 and the processing loop will be terminated. SQLCODE +100 (row not found) will terminate the processing loop.

See also the example program DEM2SCRL supplied in the Natural system library SYSDB2.