

---

---

**Information technology database  
languages — Guidance for the use of  
database language SQL —**

Part 9:  
**Online analytic processing (OLAP)  
capabilities**

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020



IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020



**COPYRIGHT PROTECTED DOCUMENT**

© ISO/IEC 2020

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office  
CP 401 • Ch. de Blandonnet 8  
CH-1214 Vernier; Geneva  
Phone: +41 22 749 01 11  
Email: [copyright@iso.org](mailto:copyright@iso.org)  
Website: [www.iso.org](http://www.iso.org)

Published in Switzerland

<b>Contents</b>	<b>Page</b>
Foreword.....	vi
Introduction.....	viii
<b>1 Scope.....</b>	<b>1</b>
<b>2 Normative references.....</b>	<b>2</b>
<b>3 Terms and definitions.....</b>	<b>3</b>
<b>4 Example data.....</b>	<b>4</b>
4.1 Introduction to example data.....	4
4.2 Table sales history.....	4
4.3 Table stock1.....	5
4.4 Table stocks.....	6
4.5 Table homes.....	6
<b>5 Windows.....</b>	<b>8</b>
5.1 Introduction to windows.....	8
5.2 Window definitions.....	9
5.2.1 Introduction to Window definitions.....	9
5.2.2 Window partitioning.....	9
5.2.3 Window ordering.....	10
5.2.3.1 Introduction to window ordering.....	10
5.2.3.2 Null ordering and treatment.....	12
5.2.4 Window frames.....	12
5.2.4.1 Introduction to window frames.....	12
5.2.4.2 Physical window frames.....	13
5.2.4.3 Logical window frames.....	15
5.2.4.3.1 Introduction to logical window frames.....	15
5.2.4.3.2 RANGE window frames.....	15
5.2.4.3.3 GROUPS window frames.....	16
5.2.4.4 Window frame exclusions.....	17
5.3 Explicit vs implicit window definitions.....	19
5.4 Multiple window definitions.....	19
<b>6 Window functions.....</b>	<b>21</b>
6.1 Introduction to window functions.....	21
6.2 Rank functions.....	21
6.3 Distribution functions.....	22
6.4 Row number function.....	23
6.5 Window aggregate functions.....	24
6.6 Ntile function.....	27
6.7 LEAD and LAG functions.....	28
6.8 FIRST_VALUE and LAST_VALUE functions.....	30

6.9	NTH_VALUE function.....	31
6.9.1	Null treatment.....	33
<b>7</b>	<b>Nested window functions.....</b>	<b>34</b>
7.1	Introduction to nested window functions.....	34
7.2	Row markers.....	35
7.3	Offsets.....	36
7.4	FRAME_ROW.....	37
7.5	Nested ROW_NUMBER function.....	38
7.6	Effects of EXCLUDE.....	40
<b>8</b>	<b>Enhanced aggregate functions.....</b>	<b>41</b>
8.1	Introduction to enhanced aggregate functions.....	41
8.2	Unary statistical aggregate functions.....	41
8.3	Binary statistical aggregate functions.....	41
8.4	Hypothetical rank and distribution aggregate functions.....	43
8.5	Inverse distribution functions.....	43
	<b>Bibliography.....</b>	<b>46</b>
	<b>Index.....</b>	<b>47</b>

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020

## Tables

Table		Page
1	Table sales_history .....	4
2	Table stock1 .....	5
3	Table stocks .....	6
4	Table homes .....	7
5	Result of window clause .....	10
6	Result of window clause ordering .....	11
7	Result of physical window frame .....	13
8	Result of physical window frame .....	14
9	Result of logical window clause .....	15
10	Result of window frame .....	16
11	Result of window frame exclusion 1 .....	17
12	Result of window frame exclusion 2 .....	18
13	Result of RANK and DENSE_RANK function .....	22
14	Result of PERCENT_RANK and CUME_DIST functions .....	23
15	Result of ROW_NUMBER function .....	24
16	Result of aggregate function (SUM) ordered .....	25
17	Result of aggregate function (SUM) unordered .....	25
18	Result of aggregate function (AVG) .....	26
19	Result of aggregate function (NTILE) .....	27
20	Result of aggregate function (NTILE) .....	28
21	Result of aggregate function (LEAD) .....	29
22	Result of aggregate function (LAG) .....	29
23	Result of aggregate function (FIRST_VALUE) .....	30
24	Result of aggregate function (LAST_VALUE) .....	31
25	Result of aggregate function (NTH_VALUE) .....	32
26	Result of row markers .....	35
27	Result of row markers (offsets) .....	36
28	Result of window frames with row markers .....	37
29	Result of EXCLUDE .....	40
30	Result of hypothetical aggregate functions .....	43
31	Result of inverse distribution functions .....	44
32	Result of inverse distribution functions with ordering .....	44

## Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see [www.iso.org/directives](http://www.iso.org/directives)).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see [www.iso.org/patents](http://www.iso.org/patents)), or the IEC list of patent declarations received (see <http://patents.iec.ch>).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see [www.iso.org/iso/foreword.html](http://www.iso.org/iso/foreword.html).

This document was prepared by Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 32, *Data management and interchange*.

This document is intended to be used in conjunction with the following editions of the parts of ISO/IEC 9075:

- ISO/IEC 9075-1, sixth edition or later
- ISO/IEC 9075-2, sixth edition or later
- ISO/IEC 9075-3, sixth edition or later
- ISO/IEC 9075-4, seventh edition or later
- ISO/IEC 9075-9, fifth edition or later
- ISO/IEC 9075-10, fifth edition or later
- ISO/IEC 9075-11, fifth edition or later
- ISO/IEC 9075-13, fifth edition or later
- ISO/IEC 9075-14, sixth edition or later
- ISO/IEC 9075-15, second edition or later
- ISO/IEC 9075-16, first edition or later

A list of all parts in the ISO/IEC 19075 series can be found on the ISO website.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at [www.iso.org/members.html](http://www.iso.org/members.html).

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020

## Introduction

This document discusses the syntax and semantics for including online analytic processing (OLAP) capabilities in SQL, as defined in ISO/IEC 9075-2.

The organization of this document is as follows:

- 1) **Clause 1, “Scope”**, specifies the scope of this document.
- 2) **Clause 2, “Normative references”**, identifies standards that are referenced as part of requirements by this document.
- 3) **Clause 3, “Terms and definitions”**, defines the terms and definitions used in this document.
- 4) **Clause 5, “Windows”**, discusses Feature T611, “Elementary OLAP operations” and Feature T612, “Advanced OLAP operations”, introducing the concept of a window in an SQL query.
- 5) **Clause 6, “Window functions”**, further discusses Feature T611, “Elementary OLAP operations” and Feature T612, “Advanced OLAP operations”, as well as Feature T614, “NTILE function”, Feature T615, “LEAD and LAG functions”, Feature T616, “Null treatment option for LEAD and LAG functions”, Feature T617, “FIRST\_VALUE and LAST\_VALUE functions”, and Feature T618, “NTH\_VALUE function”.
- 6) **Clause 7, “Nested window functions”**, discusses the additional window functionality in Feature T619, “Nested window functions”.
- 7) **Clause 8, “Enhanced aggregate functions”**, discusses Feature T621, “Enhanced numeric functions” and its introduction of enhanced aggregate functions in SQL.

# Information technology — Guidance for the use of database language SQL —

Part 9:

## Online analytic processing (OLAP) capabilities

### 1 Scope

This document discusses the syntax and semantics for including online analytic processing (OLAP) capabilities in SQL, as defined in [ISO/IEC 9075-2](#).

It discusses the following features regarding OLAP capabilities of the SQL language:

- Feature T611, “Elementary OLAP operations”,
- Feature T612, “Advanced OLAP operations”,
- Feature T614, “NTILE function”,
- Feature T615, “LEAD and LAG functions”,
- Feature T616, “Null treatment option for LEAD and LAG functions”,
- Feature T617, “FIRST\_VALUE and LAST\_VALUE functions”,
- Feature T618, “NTH\_VALUE function”,
- Feature T619, “Nested window functions”,
- Feature T620, “WINDOW clause: GROUPS option”,
- Feature T621, “Enhanced numeric functions”.

## 2 Normative references

There are no normative references in this document.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020

### 3 Terms and definitions

No terms and definitions are listed in this document.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <http://www.iso.org/obp>
- IEC Electropedia: available at <http://www.electropedia.org/>

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020

## 4 Example data

### 4.1 Introduction to example data

The examples in this document are based on several tables.

The order in which the rows of all sample tables are displayed is immaterial.

### 4.2 Table sales history

Subclause 4.2, “Table sales history”, contains information on a business spread over several territories with total sales accumulated monthly in each territory. Table 1, “Table sales\_history”, shows sample data for Subclause 4.2, “Table sales history”:

Table 1 — Table sales\_history

Territory	Month	Sales
East	199812	11
West	199811	12
West	199901	11
East	199811	4
East	199810	10
West	199810	8
East	199902	10
East	199901	7
West	199812	7
West	199902	6

SQL to create and populate Subclause 4.2, “Table sales history”.

```
CREATE TABLE Sales_History
(Territory CHARACTER (10),
 Month INTEGER,
 Sales INTEGER)

INSERT INTO Sales_History VALUES ('East', 199812, 11)
INSERT INTO Sales_History VALUES ('West', 199811, 12)
INSERT INTO Sales_History VALUES ('West', 199901, 11)
INSERT INTO Sales_History VALUES ('East', 199811, 4)
INSERT INTO Sales_History VALUES ('East', 199810, 10)
INSERT INTO Sales_History VALUES ('West', 199810, 8)
INSERT INTO Sales_History VALUES ('East', 199902, 10)
INSERT INTO Sales_History VALUES ('East', 199901, 7)
INSERT INTO Sales_History VALUES ('West', 199812, 7)
INSERT INTO Sales_History VALUES ('West', 199902, 6)
```

```
INSERT INTO Sales_History VALUES ('East', 199901, 7)
INSERT INTO Sales_History VALUES ('West', 199812, 7)
INSERT INTO Sales_History VALUES ('West', 199902, 6)
```

### 4.3 Table stock1

The next examples are two variants of a stock table containing information on stock transactions for a particular account. Columns in Table 2, “Table stock1”, include transaction ID, trade day, and type, as well as the share amount and ticker symbol. Subclause 4.4, “Table stocks”, covers the columns ticker, tradeday, and price.

Table 2 — Table stock1

Acno	Tid	Tradeday	TType	Amount	Ticker
123	1	1	buy	1000	csc0
123	2	1	buy	400	inpr
123	3	2	buy	2000	symc
123	4	2	buy	1200	csc0
123	5	2	buy	500	inpr
123	6	4	buy	200	csc0
123	7	4	buy	100	csc0
123	9	5	buy	400	inpr
123	10	5	buy	200	goog
123	11	5	buy	1000	inpr
123	12	5	buy	4000	inpr
123	13	8	buy	2000	hpq

SQL to create and populate Table 2, “Table stock1”.

```
CREATE TABLE Stock1
(Acno INTEGER,
 Tid INTEGER,
 Tradeday INTEGER,
 TType CHARACTER (10),
 Amount INTEGER,
 Ticker CHARACTER (10))

INSERT INTO Stock1 VALUES (123, 1, 1, 'buy', 1000, 'csc0')
INSERT INTO Stock1 VALUES (123, 2, 1, 'buy', 400, 'inpr')
INSERT INTO Stock1 VALUES (123, 3, 2, 'buy', 2000, 'symc')
INSERT INTO Stock1 VALUES (123, 4, 2, 'buy', 1200, 'csc0')
INSERT INTO Stock1 VALUES (123, 5, 2, 'buy', 500, 'inpr')
INSERT INTO Stock1 VALUES (123, 6, 4, 'buy', 200, 'csc0')
```

4.3 Table stock1

```
INSERT INTO Stock1 VALUES (123, 7, 4, 'buy', 100, 'csc0')
INSERT INTO Stock1 VALUES (123, 9, 5, 'buy', 400, 'inpr')
INSERT INTO Stock1 VALUES (123, 10, 5, 'buy', 200, 'goog')
INSERT INTO Stock1 VALUES (123, 11, 5, 'buy', 1000, 'inpr')
INSERT INTO Stock1 VALUES (123, 12, 5, 'buy', 4000, 'inpr')
INSERT INTO Stock1 VALUES (123, 13, 8, 'buy', 2000, 'hpq')
```

4.4 Table stocks

Table 3 — Table stocks

Ticker	Tradeday	Price
ZYX	1	10
ZYX	2	11
ZYX	3	12
ZYX	4	12
ZYX	5	12
ZYX	6	11
ZYX	7	12
ZYX	8	12

SQL to create and populate Subclause 4.4, “Table stocks”.

```
CREATE TABLE Stocks
(Ticker CHARACTER (10),
Tradeday INTEGER,
Price INTEGER)

INSERT INTO Stocks VALUES ('ZYX', 1, 10)
INSERT INTO Stocks VALUES ('ZYX', 2, 11)
INSERT INTO Stocks VALUES ('ZYX', 3, 12)
INSERT INTO Stocks VALUES ('ZYX', 4, 12)
INSERT INTO Stocks VALUES ('ZYX', 5, 12)
INSERT INTO Stocks VALUES ('ZYX', 6, 11)
INSERT INTO Stocks VALUES ('ZYX', 7, 12)
INSERT INTO Stocks VALUES ('ZYX', 8, 12)
```

4.5 Table homes

The final example is Table 4, “Table homes”, containing data concerning house prices and locations.

**Table 4 — Table homes**

Area	Address	Price
Uptown	15 Peekaboo St.	456,000
Uptown	27 Primrose Path	341,000
Uptown	44 Shady Lane	341,000
Uptown	23301 Highway 61	244,000
Uptown	34 Desolation Rd.	244,000
Uptown	77 Sunset Strip	102,000
Downtown	72 Easy St.	509,000
Downtown	29 Wong Way	201,000
Downtown	45 Diamond Lane	201,000
Downtown	76 Blind Alley	201,000
Downtown	15 Tern Pike	199,000
Downtown	444 Kanga Rua	102,000

SQL to create and populate Table 4, "Table homes".

```
CREATE TABLE Homes
(Area CHARACTER (10),
Address CHARACTER (20),
Price INTEGER)

INSERT INTO Homes VALUES ('Uptown', '15 Peekaboo St.', 456000)
INSERT INTO Homes VALUES ('Uptown', '27 Primrose Path', 341000)
INSERT INTO Homes VALUES ('Uptown', '44 Shady Lane', 341000)
INSERT INTO Homes VALUES ('Uptown', '23301 Highway 61', 244000)
INSERT INTO Homes VALUES ('Uptown', '34 Desolation Rd.', 244000)
INSERT INTO Homes VALUES ('Uptown', '77 Sunset Strip', 102000)
INSERT INTO Homes VALUES ('Downtown', '72 Easy St.', 509000)
INSERT INTO Homes VALUES ('Downtown', '29 Wong Way', 201000)
INSERT INTO Homes VALUES ('Downtown', '45 Diamond Lane', 201000)
INSERT INTO Homes VALUES ('Downtown', '76 Blind Alley', 201000)
INSERT INTO Homes VALUES ('Downtown', '15 Tern Pike', 199000)
INSERT INTO Homes VALUES ('Downtown', '444 Kanga Rua', 102000)
```

## 5 Windows

### 5.1 Introduction to windows

SQL in Feature T611, “Elementary OLAP operations” and Feature T612, “Advanced OLAP operations” of ISO/IEC 9075-2 adds support for online analytical processing (OLAP). The extensions are parts of the SELECT command.

OLAP is concerned with data aggregation across grouping criteria to generate values such as subtotals and totals on multiple levels. Grouping criteria are often called dimensions. OLAP is based on the concept of multiple dimensions and navigation across the aggregation levels as well as the accumulated data. OLAP is used in applications such as analytics and reporting. It can be used in iterative fashions trying out different grouping criteria and different subsets of the data to be analyzed.

To deal with these requirements, the features mentioned above introduce a “WINDOW” facility that can present aggregated content as a rolling window, ordered and grouped by the specified criteria. See the example in Subclause 5.2.2, “Window partitioning”, Table 5, “Result of window clause”.

The data can be treated like other result sets, and/or can be used for further processing.

A query contains a select list and a table expression. The table expression produces a result set; call it *RT*. The select list is evaluated by applying its expressions to each row in *RT*. Without the OLAP features, the select list expressions can only “see” that one current row, making it impossible to compute values that rely on values from other rows in *RT*. Such computations can be simulated only by arranging for the necessary values to be included as additional columns in *RT*, which may be inconvenient or impossible.

The concept of the windowed table alleviates this limitation by adding windows to the result of the table expression. One way to think of a window might be to imagine that it is a transient copy of *RT*, including an indication of the current row. (This is just a conceptual device; an implementation need not actually copy *RT*.) This transient copy can be logically re-arranged according to a sort ordering, a partitioning, or both; it can also be limited to a subset of rows, via window framing. These rearrangements of the window's rows take place without affecting *RT* itself. Multiple windows may exist, each with its own independently applied ordering, partitioning, and/or framing specifications. Windows may be defined either in the new window clause, or in-line in individual window function specifications in the select list.

Partitioning and/or ordering may be used to compute such results as ranking, Ntiles, and other analytic functions. The frame specifies which rows of a partition, relative to the current row, should participate in the calculation of an OLAP function. Through frames, windows support such important OLAP capabilities as cumulative sums and moving averages.

Ordering in windows is specified with the same sort specification list used by cursors and elsewhere in the SELECT statement, and with the same semantics. Ordering as enhanced for the OLAP capabilities of SQL also includes user specified control of the ordering of nulls. Although ordering of rows may be non-deterministic within a window, the same nondeterministic ordering is used in windows that have equivalent partitioning and ordering clauses specified.

## 5.2 Window definitions

### 5.2.1 Introduction to Window definitions

The <window clause> is an additional syntax element of the query expression and, if specified, follows the <from clause>, <where clause>, <group by clause>, and <having clause>. As with the other clauses of the query expression, the window clause applies to the result of the preceding clauses. The window clause consists of a comma separated list of window definitions. Each of these has a name, for reference by OLAP functions in the select list, and a window specification. For convenience, a window definition may also be coded inline in the OLAP function specification. Clause 6, “Window functions”, describes the various OLAP functions available for use, and which combinations of window specification detail clauses are permitted for each.

The window specification may contain any or all of a partitioning, ordering, and frame definition. For example:

```
WINDOW tms AS
(PARTITION BY territory -- window partitioning
ORDER BY month, sales) -- window ordering
```

A window function can refer to the defined window by name, for instance:

```
SELECT RANK() OVER tms AS the_rank, ...
FROM ...
WINDOW tms AS ... as above ...
```

Or, implicitly by specifying the window definition directly in-line:

```
SELECT RANK() OVER (PARTITION BY territory
ORDER BY month, sales) AS the_rank, ...
FROM ...
```

which has the identical result.

### 5.2.2 Window partitioning

The optional “window partition clause” specifies a partitioning of the result set generated by the preceding from, where, group by and having clauses. Like the group by clause, the window partition clause is a comma-separated list of column references used to group rows for subsequent processing. However, unlike the group by clause, each input row to a window partitioning is retained in the result set. This permits the introduction of analytical functions that operate on the individual rows of a partition. The “collate clause” option allows character columns to be partitioned based on a named collation. If there is no window partition clause, then the entire result set of the containing query constitutes a single partition.

Notice that although the window partition clause is similar to the group by clause, it is not the same thing. The difference is that the grouping specified by a group by clause collapses each group to a single row in its result set. The partitioning specified by a window partition clause does not collapse the partitions to a single row. Rather, the window partition of a row *R* is the collection of rows that are not distinct from *R*, for all columns enumerated in the window partitioning clause.

**Example 1, “window clause”**, shows the effect of partitioning Sales\_history using the Territory column in the definition.

**Example 1 — window clause**

```
SELECT Territory, Month, Sales FROM Sales_history
WINDOW w_eg1 AS (PARTITION BY Territory)
```

- WINDOW `w_eg1` identifies the window definition so that it may be referenced in one or more OLAP functions elsewhere in the query ( `w_eg1` is the window's name).
- PARTITION BY introduces the partitioning. A partitioning is simply a list of one or more columns on which the data is partitioned.

**Table 5 — Result of window clause**

Territory	Month	Sales	
East	199901	7	⌋
East	199811	4	
East	199810	10	⌋ "East" partition
East	199812	11	
East	199902	10	⌋
West	199811	12	⌋
West	199810	8	
West	199902	6	⌋ "West" partition
West	199901	11	
West	199812	7	⌋

The rows of Table 5, “Result of window clause”, are clustered to show the effect of the partitioning, but in reality, the result rows may still appear in any sequence. If a result set is desired to be ordered on one or more columns, an order by clause must be used in the containing query expression.

**5.2.3 Window ordering**

**5.2.3.1 Introduction to window ordering**

The next optional element of a window definition is the “window order clause”. It consists of a “sort specification list” that is syntactically the same as the one found in an “order by” clause of the query expression.

Whether in a window or a query expression, a sort specification list specifies an ordering of rows. The difference is that, in a query expression, the ordering determines the sequence of rows in the result set of the query expression. In a window, the ordering helps to determine the value of order-dependent OLAP functions such as ranking. Another difference is that the query expression ordering applies to all result rows, whereas a window ordering is applied to each partition separately.

As the rows of each partition are ordered, it may turn out that multiple rows are peers; that is, they have the same values in each of the elements of the sort specification list according to ordering semantics. A window ordering group is a maximal set of rows (perhaps just one) in a partition that are peers according to the window ordering. Some analytic functions (RANK, for instance) operate on all rows in an ordering group or set of ordering groups, whereas others (such as NTH\_VALUE) may operate on individual rows regardless of their participation in an ordering group. The ordering of rows within a single window ordering group is implementation-dependent, hence introducing the potential for non-determinism in some OLAP functions. Distinct ordering groups within a partition are of course ordered according to the sort specification list of the window order clause.

If there is no window order clause, then each partition contains a single window ordering group consisting of all the rows in the partition and all rows are peers.

Example 2, “window clause ordering”, shows the effect of ordering Sales\_history using the Sales column in combination with the partitioning of Example 1, “window clause”.

**Example 2 — window clause ordering**

```
SELECT Territory, Month, Sales FROM Sales_history
WINDOW w_eg2 AS (PARTITION BY Territory
ORDER BY Month ASC)
```

- w\_eg2 is the name of the window defined by this example.
- In addition to the partitioning, ORDER BY introduces the ordering of rows within each partition. The syntax is identical to the order by clause of a query expression, including the optional null ordering specification.

**Table 6 — Result of window clause ordering**

Territory	Month	Sales	
East	199810	10	⌋
East	199811	4	
East	199812	11	⌋ "East" partition
East	199901	7	
East	199902	10	⌋
West	199810	8	⌋
West	199811	12	
West	199812	7	⌋ "West" partition
West	199901	11	
West	199902	6	⌋

Once again, the rows of Table 6, “Result of window clause ordering”, are organized to show the effect of the window specification and how they are operated on by the various OLAP functions. However, the result of the application of the window specification and subsequent OLAP functions doesn't assure any particular result row sequence. If a result set is desired to be ordered on one or more columns, an order by clause must be used in the containing query expression.

### 5.2.3.2 Null ordering and treatment

Feature T611, “Elementary OLAP operations” also introduced the “null ordering” option to the sort specification list of both the window order clause and the existing order by clause. Without it, null values may be sorted ahead of all valued instances of a column, or after all valued instances. Which option is chosen by default is implementation-defined. With the null ordering option, the user may explicitly request using NULLS FIRST or NULLS LAST that null values of an ordering column be sorted ahead or following the valued instances. In some cases the ability to explicitly define the null ordering is necessary for producing the correct result of an analytical function.

Unlike ordinary aggregate functions, window functions do not necessarily ignore NULL values. The ROW\_NUMBER function, for instance, is fundamentally a row counting function over a window and as such does not care about NULL values. Some window functions, such as LAG, LEAD, FIRST\_VALUE, LAST\_VALUE, and NTH\_VALUE, permit explicit choice of null treatment: one can ignore nulls or respect nulls, as required by the desired outcome.

## 5.2.4 Window frames

### 5.2.4.1 Introduction to window frames

The window frame of a row  $R$  in a window partition is a multiset of rows, defined relative to  $R$  in the ordering of the rows of  $R$ 's partition. Window frames are used to specify multisets of rows on which to perform window or aggregate functions, such as SUM or AVG or one of several other OLAP-specific functions.

There are three ways to specify the scope of a window frame, given a current row  $R$ : row counting, value offset, or ordering group counting.

- 1) The keyword ROWS says to include the specified number of rows before or after  $R$ . This defines a physical window frame.
- 2) The keyword RANGE requires that the partition be ordered on a single ordering key, of a type that can have a value added or subtracted (thus, numeric, datetime, or interval); one specifies a sort key value offset, and the window frame extends to any ordering groups containing values in the range between that offset and the sort key value in row  $R$ . This defines a logical window frame.
- 3) The keyword GROUPS requires that the partition be ordered (possibly on multiple sort keys), and says to include the specified number of ordering groups before or after the ordering group containing  $R$ . This also defines a logical window frame.

Thus there are two components in the specification of a frame, both mandatory:

- 1) The choice of the keyword ROWS, RANGE, or GROUPS, to indicate how the extent of the window frame is to be determined.
- 2) An indication of the starting row and ending row of the frame. Note that the ending row of the frame may follow the row currently being processed. This enables functions to be computed over “look ahead” data, a very valuable property for OLAP processing. The window frame extent syntax may consist of any of the following elements:
  - UNBOUNDED PRECEDING to request all rows, ranges or groups preceding the current row.
  - $n$  PRECEDING to request the  $n$  rows, ranges or groups preceding the current row.
  - $n$  FOLLOWING to request the  $n$  rows, ranges or groups following the current row.

- UNBOUNDED FOLLOWING to request all rows, ranges or groups following the current row.
- BETWEEN start AND end to request all rows, ranges or groups between the start and end designations.
- start or end may also be CURRENT ROW to start or end the frame with the current row.

If the frame specification would extend the frame past the beginning or end of a partition, the frame will be truncated to the partition boundary.

If there is neither a window frame clause nor a window ordering clause, then there is an implicit window frame consisting of all rows in the partition. If there is no window frame clause, but a window ordering clause is present, then there is an implicit window frame which contains all the rows of the window partition of *R* that precede *R* or are peers of *R* in the window ordering of the window partition defined by the window ordering clause.

### 5.2.4.2 Physical window frames

A physical window frame is defined by counting some number of rows from the current row. Physical window frames are defined by the keyword ROWS. For physical framing to produce a sensible result from many window functions, there has to be a known and predictable number of input rows for each set of values of interest. Gaps or unexpected multiple rows would throw off the result or make it non-deterministic.

For instance, suppose one wanted a moving average of 3 calendar months from the sales\_history example table. One could define a frame as "ROWS BETWEEN 2 PRECEDING AND CURRENT ROW", but that will only produce a meaningful running average if there is exactly one sales\_history row for each calendar month. If there are missing months, the calendar month moving average will be incorrect.

Similarly, physical framing applied to data with multiple rows with the same ordering value may produce non-deterministic results. Rows within an ordering group are ordered in an implementation-dependent way, so a physical frame that does not cover the entire ordering group is not guaranteed to return deterministic results.

On the other hand, physical framing does not place any restrictions on window ordering (it does not even require window ordering), and may be the simplest way to frame data that has a well defined progression of values.

Examples of physical window frames follow.

#### Example 3 — physical window frame 1

```
SELECT Territory, Month, Sales FROM Sales_history
WINDOW w_eg3 AS (PARTITION BY Territory
ORDER BY Month ASC
ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW)
```

- w\_eg3 is the window name of this example.
- For each row of each partition, the window frame syntax uses the keyword ROWS to designate the frames as physical. Each frame consists of the current row and all preceding rows in the partition.

Table 7 — Result of physical window frame

Territory	Month	Sales	
East	199810	10	}

ISO/IEC TR 19075-9:2020(E)  
5.2 Window definitions

Territory	Month	Sales	
East	199811	4	} frame for "East", 199901
East	199812	11	
East	199901	7	J
East	199902	10	
West	199810	8	}
West	199811	12	J frame for "West", 199811
West	199812	7	
West	199901	11	
West	199902	6	

**Example 4 — physical window frame 2**

```
SELECT Territory, Month, Sales FROM Sales_history
WINDOW w_eg4 AS (PARTITION BY Territory
ORDER BY Month ASC
ROWS BETWEEN 2 PRECEDING AND 1 FOLLOWING)
```

- w\_eg4 is the window name of this example.
- For each row of each partition, the window frame syntax uses the keyword `ROWS` to designate the frames as physical. Each frame consists of the current row, the 2 rows that precede it and the row that follows it.

**Table 8 — Result of physical window frame**

Territory	Month	Sales	
East	199810	10	}
East	199811	4	} frame for "East", 199812
East	199812	11	
East	199901	7	J
East	199902	10	
West	199810	8	}
West	199811	12	} frame for "West", 199811
West	199812	7	J
West	199901	11	
West	199902	6	

Note that the frame for "West", 199811 only contains 3 rows since there is only 1 preceding row in the partition.

### 5.2.4.3 Logical window frames

#### 5.2.4.3.1 Introduction to logical window frames

Logical window frames do not count rows, rather, the frame is defined based on values in the data. A logical framing is either looking for a value range, or counting ordering groups. As such, logical window frames require a window ordering. Framing by value range uses the keyword RANGE, while framing by ordering groups uses the keyword GROUPS.

#### 5.2.4.3.2 RANGE window frames

A logical window frame defined by the RANGE keyword may have only a single sort key in its ordering clause. There must be a quantity which can be added to and subtracted from the sort key to define the starting or ending point of a frame. A frame is then defined by value offsets preceding and following the current row. This means that the data type must be numeric, datetime or interval, and that the offset must be a numeric literal or a compatible interval literal.

Consider a similar Sales\_history table, with some months missing from the data, and the Month column stored as a date value containing the first day of each month. The following example demonstrates the benefits of RANGE window frames:

#### Example 5 — logical range window frame

```
SELECT Territory, Month, Sales FROM Similar_sales_history
   WINDOW w_eg5 AS (PARTITION BY Territory
                    ORDER BY Month ASC
                    RANGE INTERVAL '02' MONTH PRECEDING)
```

- w\_eg5 is the window name in this example.
- For each row of each partition, the window frame syntax uses the keyword RANGE to designate the frame as logical defined by a range of values in the ordering column (there may only be one ordering column). INTERVAL '02' MONTH is the literal that defines the range.
- The absence of a following clause means that each frame will include the CURRENT ROW.

Table 9 — Result of logical window clause

Territory	Month	Sales	
East	1998-10-01	10	} frame for "East", 1998-12-01
East	1998-12-01	11	
East	1999-01-01	7	
East	1999-02-01	10	

5.2 Window definitions

Territory	Month	Sales	
West	1998-10-01	8	
West	1998-11-01	12	frame for "West", 1998-12-01
West	1998-12-01	7	
West	1999-02-01	6	

Note that the frame for "East", 1998-12-01 only contains 2 rows since there is no row for "East", 1998-11-01.

5.2.4.3.3 GROUPS window frames

A logical window frame defined by the GROUPS keyword may have any number of sort keys in its ordering clause. Likewise, there is no limitation on the data types of the sort keys. A frame is defined by counting some number of ordering groups from the ordering group that contains the current row. The use of GROUPS syntax overcomes the weaknesses of physical window frames in that the ordering data does not have to be dense and the results are deterministic. Rows within an ordering group may have a non-deterministic order, but the groups as a whole are deterministically ordered. Results of any functions executed on grouped windows are then also deterministic. Likewise, they overcome the weaknesses of RANGE defined logical windows by permitting multiple sort keys with no type restrictions.

The following example demonstrates the use of GROUPS window frames:

**Example 6 — logical window frame with GROUPS**

```
SELECT Acno, Tid, Tradeday, TType, Amount, Ticker FROM Stock1
WINDOW w_eg6 AS (PARTITION BY Acno
ORDER BY Tradeday
GROUPS BETWEEN 2 PRECEDING AND CURRENT ROW)
```

- w\_eg6 is the window name in this example.
- For each row of each partition, the window frame syntax uses the keyword GROUPS to designate the frame as logical defined by 3 ordering groups — the one containing the current row and the 2 preceding ordering groups.
- Note that, despite the CURRENT ROW designation, rows that are in the same ordering group as the current row but have yet to be processed, are considered to be in the frame.

**Table 10 — Result of window frame**

Acno	Tid	Tradeday	TType	Amount	Ticker	
123	1	1	buy	1000	csc0	
123	2	1	buy	400	inpr	
123	3	2	buy	2000	symc	frame for Tday 2 (same frame for Tid's 3, 4 and 5)
123	4	2	buy	1200	csc0	

Acno	Tid	Tradeday	TType	Amount	Ticker	
123	5	2	buy	500	inpr	J
123	6	4	buy	200	csc0	] frame for Tday 8
123	7	4	buy	100	csc0	
123	9	5	buy	400	inpr	
123	10	5	buy	200	goog	] frame for Tday 8
123	11	5	buy	1000	inpr	
123	12	5	buy	4000	inpr	
123	13	8	buy	2000	hpq	J

#### 5.2.4.4 Window frame exclusions

Syntax also allows the identification of rows to be specifically excluded from window frames in which they would otherwise participate. Options may be used to exclude the current row, the current ordering group or all rows that are ties of the current row. This adds a degree of fine tuning to the specification of window frames.

The following examples show the effect of window frame exclusion syntax:

##### Example 7 — window frame exclusion 1

```
SELECT Territory, Month, Sales FROM Sales_history
WINDOW w_eg7 AS (PARTITION BY Territory
ORDER BY Month ASC
RANGE INTERVAL '02' MONTH PRECEDING
EXCLUDE CURRENT ROW)
```

- w\_eg7 is the window name in this example,
- This is the same as the earlier example defining w\_eg5, with the exception that the current row is excluded from each frame.

Table 11 — Result of window frame exclusion 1

Territory	Month	Sales	
East	1998-10-01	10	-- frame for "East", 199812
East	1998-12-01	11	] frame for "East", 199902
East	1999-01-01	7	J
East	1999-02-01	10	

ISO/IEC TR 19075-9:2020(E)  
5.2 Window definitions

Territory	Month	Sales	
West	1998-10-01	8	} frame for "West", 199812
West	1998-11-01	12	J
West	1998-12-01	7	
West	1999-02-01	6	

**Example 8 — window frame exclusion 2**

```
SELECT Acno, Tid, Tradeday, TType, Amount, Ticker FROM Stock1
WINDOW w_eg8 AS (PARTITION BY acno
ORDER BY Tradeday
GROUPS BETWEEN 2 PRECEDING AND CURRENT ROW
EXCLUDE TIES)
```

- w\_eg8 is the window name in this example.
- This example is the same as earlier example defining w\_eg6 with the exception that only the current row is included in the frame from the ordering group that contains it (Tid's 4 and 5 are omitted from the frame for Tid 3 and Tid's 10, 11, and 12 are omitted from the frame for Tid 9).

**Table 12 — Result of window frame exclusion 2**

Acno	Tid	Tradeday	TType	Amount	Ticker	
123	1	1	buy	1000	csc0	}
123	2	1	buy	400	inpr	} frame for Tid 3
123	3	2	buy	2000	symc	} J
123	4	2	buy	1200	csc0	
123	5	2	buy	500	inpr	} frame for Tid 9
123	6	4	buy	200	csc0	
123	7	4	buy	100	csc0	
123	9	5	buy	400	inpr	J
123	10	5	buy	200	goog	
123	11	5	buy	1000	inpr	
123	12	5	buy	4000	inpr	
123	13	8	buy	2000	hpq	

The EXCLUDE GROUP option is used to exclude the entire ordering group of the current row from a frame, and applies to any type of window frame (ROWS, RANGE, or GROUPS).

### 5.3 Explicit vs implicit window definitions

The window definitions discussed thus far have been explicitly declared using the window definition clause in the table expression. Analytic functions may reference a defined window explicitly by name or may include the elements of the window definition syntax in the function specification. Such windows are implicitly defined and the only difference in their specification is the absence of a window name. Either form of window definition may be used, although explicit window definitions are a useful feature when more than one analytic function is to be executed on the same window. In either case, the semantics of function execution are the same.

Example for an explicit window definition:

#### Example 9 — explicit window definition

```
SELECT Territory, Month, Sales,
       SUM (Sales) OVER Win1 AS Sums
FROM Sales_history
     WINDOW Win1 AS (PARTITION BY Territory ORDER BY Month)
```

Example for the same windows definition as an implicit window definition:

#### Example 10 — implicit window definition

```
SELECT Territory, Month, Sales,
       SUM (Sales) OVER (PARTITION BY Territory ORDER BY Month) AS Sums
FROM Sales_history
```

### 5.4 Multiple window definitions

A given query may have any number of window definitions, explicit and/or implicit. Moreover, one window definition may reference another explicit window definition, with the resulting window containing a combination of the specifications of each. This is true of purely explicit window definitions or of mixtures of explicit and implicit window definitions. So, a window function may reference an explicit window definition, then fill in some of the window properties as they apply to the particular function. More specifically, if a window definition x references an existing window definition y:

- x may not include a window partition specification.
- If y includes a window ordering specification, x may not include a window ordering specification.
- y may not include a window frame specification.

This capability permits the easier definition of multiple windows that have some common elements. For example:

#### Example 11 — multiple window definitions

```
SELECT Territory, Month, Sales FROM Sales_history
     WINDOW w_eg9a AS (PARTITION BY Territory),
     w_eg9b AS (w_eg9a ORDER BY Month),
     w_eg9c AS (w_eg9b ROWS 2 PRECEDING),
     w_eg9d AS (w_eg9b GROUPS 1 PRECEDING EXCLUDE TIES)
```

- Any of windows w\_eg9a , w\_eg9b , w\_eg9c , and/or w\_eg9d may be used in different functions in the containing query.
- w\_egc and w\_egd use the same partitioning and ordering, but define different window frames.

5.4 Multiple window definitions

- Another window might also be defined from `w_eg9a` with a different ordering and subsequently, different frame specifications, as well.

IECNORM.COM : Click to view the full PDF of ISO/IEC TR 19075-9:2020

## 6 Window functions

### 6.1 Introduction to window functions

In addition to the window definition, Feature T611, “Elementary OLAP operations” and Feature T612, “Advanced OLAP operations” of ISO/IEC 9075-2 introduce new functions that exploit the window definition to produce results that are commonly used in OLAP applications. This clause discusses these functions, in addition to those introduced by Feature T614, “NTILE function”, Feature T615, “LEAD and LAG functions”, Feature T616, “Null treatment option for LEAD and LAG functions”, Feature T617, “FIRST\_VALUE and LAST\_VALUE functions”. and Feature T618, “NTH\_VALUE function”. Syntax and examples are used to explain the functions, as well as the features of window definitions that are relevant to their execution.

A window function is one of the following:

- rank function.
- distribution function.
- row number function.
- window aggregate function.
- Ntile function.
- lead or lag function.
- first, last, or Nth value function.

### 6.2 Rank functions

The rank functions compute the ordinal rank of each row *R* within a partition as defined by a window. The window definition must include a window order clause but no window frame clause. There are 2 rank functions included in OLAP support:

- RANK() assigns an ordinal rank to each row in a partition according to its position in the ordering of the partition. If a row *R* in a partition has peers according to the ordering, the peers will be assigned the same rank as *R* and there will be a gap in the assignment of rank values to the rows of the partition,
- DENSE\_RANK() also assigns an ordinal rank to each row in a partition according to its position in the ordering of the partition. If a row *R* has peers according to the ordering, the peers will be assigned the same rank as *R*. However, for DENSE\_RANK(), there are no gaps in the assignment of rank values.

The following sample shows rows of sales\_history with RANK and DENSE\_RANK computed for sales within each territory:

#### Example 12 — rank functions with explicit window

```
SELECT territory, month, sales,
       RANK() OVER w_eg13 AS rank,
       DENSE_RANK() OVER w_eg13 AS dense_rank
FROM sales_history WINDOW w_eg13 AS (PARTITION BY territory
                                     ORDER BY sales DESC)
```

The computed values can be seen below. Both RANK() and DENSE\_RANK() show ties, but RANK() also shows gaps when ties are present, unlike DENSE\_RANK() .

Example 13, “rank functions with implicit window”, shows the use of a single explicitly defined window `w_eg13`. The functions could also be specified as

**Example 13 — rank functions with implicit window**

```
SELECT territory, month, sales,
       RANK() OVER (PARTITION BY Territory
                   ORDER BY Sales DESC) AS rank,
       DENSE_RANK() OVER (PARTITION BY Territory
                          ORDER BY Sales DESC) AS dense_rank
FROM Sales_history
```

Clearly the abbreviated form is easier both to write and to understand.

As with numerous examples, the results of this query are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

**Table 13 — Result of RANK and DENSE\_RANK function**

Territory	Month	Sales	Rank	Dense_Rank
East	199812	11	1	1
East	199810	10	2	2
East	199902	10	2	2
East	199901	7	4	3
East	199811	4	5	4
West	199811	12	1	1
West	199901	11	2	2
West	199810	8	3	3
West	199812	7	4	4
West	199902	6	5	5

**6.3 Distribution functions**

The distribution functions compute a relative rank of each row within a partition as defined by a window. The window definition must include a window order clause but no window frame clause. The rank computed by a distribution function is an approximate numeric ratio between 0.0 and 1.0, inclusive. There are two distribution functions included in OLAP support:

- PERCENT\_RANK() is defined the same as the function of the same name in popular spreadsheet applications. Its value for each row in a partition is defined as the rank of the row minus 1 divided by the number of rows in the partition minus 1. This results in values ranging from 0.0 to 1.0 in each partition.

- CUME\_DIST() is defined as the statistical cumulative distribution function. It is computed as the number of rows prior to or peer with the current row divided by the total number of rows in the partition.

An example of the use of RANK, PERCENT\_RANK, and CUME\_DIST is as follows:

**Example 14 — distribution functions**

```
SELECT territory, month, sales,
       RANK() OVER w_eg14 AS rank,
       PERCENT_RANK() OVER w_eg14 AS percent_rank,
       CUME_DIST() OVER w_eg14 AS cume_dist
FROM sales_history
     WINDOW w_eg14 AS (PARTITION BY territory
                       ORDER BY sales DESC)
```

**Table 14 — Result of PERCENT\_RANK and CUME\_DIST functions**

Territory	Month	Sales	Rank	Percent_Rank	Cume_Dist
East	199812	11	1	0.0	0.2
East	199810	10	2	0.25	0.4
East	199902	10	2	0.25	0.4
East	199901	7	4	0.75	0.8
East	199811	4	5	1.00	1.0
West	199811	12	1	0.00	0.2
West	199901	11	2	0.25	0.4
West	199810	8	3	0.50	0.6
West	199812	7	4	0.75	0.8
West	199902	6	5	1.00	1.0

As with numerous examples, the results of this query are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

**6.4 Row number function**

The row number function computes a sequential number (starting with 1 for the first row) for each row *R* within a partition as defined by a window according to the window ordering. A window order clause is not required for the row number function, but no window frame clause is permitted.

An example of the use of ROW\_NUMBER() is as follows:

**Example 15 — row number function**

```
SELECT Territory, Month, Sales,
       ROW_NUMBER() OVER (ORDER BY Month, Territory) AS row_number
FROM sales_history
```

**Table 15 — Result of ROW\_NUMBER function**

Territory	Month	Sales	Row_Number
East	199810	10	1
West	199810	8	2
East	199811	4	3
West	199811	12	4
East	199812	11	5
West	199812	7	6
East	199901	7	7
West	199901	11	8
East	199902	10	9
West	199902	6	10

As with numerous examples, the results of this query are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

**6.5 Window aggregate functions**

Window aggregate functions perform the standard aggregations on frames of rows within a partition of a window. Such “moving” aggregations are commonplace in OLAP applications. These functions are the major beneficiaries of window frames since the aggregations are performed on the rows within a defined frame. Aggregate functions may also be executed on windows without an implicit or explicit frame specification. (Recall that a window ordering without explicit frame specification implies a frame from the start of the partition through the current row and all its ordering peers.) In this case, the aggregate is computed over the entire partition and the same value is the result for each row in the partition.

Examples of window aggregate functions follow:

**Example 16 — window aggregate function (SUM) ordered**

```
SELECT Territory, Month, Sales,
       SUM (Sales) OVER (PARTITION BY Territory) AS sums
FROM Sales_history
ORDER BY Territory, Sales DESC, Month
```

This is an example of an aggregate without a window frame. Each row in a partition computes the same sum over all the rows in the partition.

**Table 16 — Result of aggregate function (SUM) ordered**

Territory	Month	Sales	Sums
East	199812	11	42
East	199810	10	42
East	199902	10	42
East	199901	7	42
East	199811	4	42
West	199811	12	44
West	199901	11	44
West	199810	8	44
West	199812	7	44
West	199902	6	44

In this case, the ordering of the results is guaranteed by the query ORDER BY clause.

The code in Example 17, “window aggregate function (SUM) unordered”, has been altered with the addition of a window ordering clause, that implies a window frame, producing cumulative results over each partition.

**Example 17 — window aggregate function (SUM) unordered**

```
SELECT Territory, Month, Sales,
       SUM (Sales) OVER (PARTITION BY Territory
                        ORDER BY Month) AS sums
FROM Sales_history
```

The window ordering clause defines an implicit window frame that for some row *R*, consists of all rows preceding *R* in the partition, plus all rows in *R*'s ordering group.

**Table 17 — Result of aggregate function (SUM) unordered**

Territory	Month	Sales	Sums
East	199812	11	25
East	199810	10	10
East	199902	10	42
East	199901	7	32
East	199811	4	14
West	199811	12	20

ISO/IEC TR 19075-9:2020(E)  
 6.5 Window aggregate functions

Territory	Month	Sales	Sums
West	199901	11	38
West	199810	8	8
West	199812	7	27
West	199902	6	44

This example has no query ORDER BY clause, so the ordering of the output rows is not guaranteed.

An explicit window framing might be used to calculate a moving average:

**Example 18 — window aggregate moving average**

```
SELECT Territory, Month, Sales,
       AVG (Sales)
         OVER (PARTITION BY Territory
              ORDER BY Month
              ROWS BETWEEN 1 PRECEDING AND 1 FOLLOWING) AS avgs
FROM Sales_history
```

- This example computes the moving average of the sales column taking into account the current row, preceding row, and following row based on the specified ordering. This is very common and useful in analytical applications.
- As with numerous examples, the results of this query are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

**Table 18 — Result of aggregate function (AVG)**

Territory	Month	Sales	Avg
East	199810	10	7.0
East	199811	4	8.3
East	199812	11	7.3
East	199901	7	9.3
East	199902	10	8.5
West	199810	8	10.0
West	199811	12	9.0
West	199812	7	10.0
West	199901	11	8.0
West	199902	6	8.5

## 6.6 Ntile function

The NTILE function assigns a quantile ranking to the rows of a partition (or all rows of a query, in the absence of a partitioning), according to its ordering. The sole parameter of the function is an exact numeric literal or dynamic parameter that produces a positive value indicating the number of quantiles into which the ordered rows are to be distributed. If a partition contains fewer rows than the requested number of quantiles, the function simply returns values 1, 2, ... The window definition for a NTILE function must include a window ordering, but must not include a window frame.

Examples of NTILE follow:

### Example 19 — ntile function 1

```
SELECT Territory, Month, Sales,
       NTILE(3)
       OVER (PARTITION BY Territory
            ORDER BY Sales DESC) AS quantile
FROM Sales_history
```

- This NTILE function requests that the rows be divided into 3 quantiles according to their Sales rank.

**Table 19 — Result of aggregate function (NTILE)**

Territory	Month	Sales	Quantile
East	199812	11	1
East	199810	10	1
East	199902	10	2
East	199901	7	2
East	199811	4	3
West	199811	12	1
West	199901	11	1
West	199810	8	2
West	199812	7	2
West	199902	6	3

### Example 20 — ntile function 2

```
SELECT Territory, Month, Sales,
       NTILE(3)
       OVER (ORDER BY Sales DESC) AS quantile
FROM Sales_history
```

- This NTILE function omits the window partition specification so that the quantile ranking is done over the entire table.

## 6.6 Ntile function

- In this query (as in the previous), some rows are placed in different quantiles, in spite of having the same Sales values (see 'East', 199810 and 'East', 199902). Such rows can legitimately be assigned to either quantile. Multiple rows which contain the same values in the ordering columns can result in non-deterministic results from the `NTILE` function.

Table 20 — Result of aggregate function (NTILE)

Territory	Month	Sales	Quantile
West	199811	12	1
East	199812	11	1
West	199901	11	1
East	199810	10	1
East	199902	10	2
West	199810	8	2
East	199901	7	2
West	199812	7	3
West	199902	6	3
East	199811	4	3

As with numerous examples, the results of these queries are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

## 6.7 LEAD and LAG functions

The LEAD and LAG functions provide the capability to reference rows preceding or following the current row being processed. Both functions take 3 parameters, the first is a value expression, the second is an exact numeric literal with a positive value, and the third is another value expression whose result data type is the same as the first. The result of LEAD is the value of the expression from the first argument evaluated on the row that follows the current row by as many rows as specified by the second argument. If there is no row that corresponds to the specified offset, the result is the value of the third argument. The result of LAG is the same, except that the row on which the first argument is evaluated is the row that precedes the current row by as many rows as specified by the second argument. As with LEAD, if there is no row that corresponds to the specified offset, the result is the value of the third argument. The window definition for LEAD or LAG functions must include a window ordering, but must not include a window frame.

Examples of LEAD and LAG follow:

## Example 21 — lead function

```
SELECT Territory, Month, Sales,
       LEAD (Sales, 1, CAST (NULL AS INTEGER))
       OVER (PARTITION BY Territory
```

```
ORDER BY Month) AS next_sales
FROM Sales_history
```

- This query displays Sales for the current and next month.
- The last row of each partition returns the NULL value for the LEAD() function. Note the NULL value is cast to the same type as Sales.

**Table 21 — Result of aggregate function (LEAD)**

Territory	Month	Sales	Next_Sales
East	199810	10	4
East	199811	4	11
East	199812	11	7
East	199901	7	10
East	199902	10	
West	199810	8	12
West	199811	12	7
West	199817	7	11
West	199901	11	6
West	199902	6	

**Example 22 — lag function**

```
SELECT Territory, Month, Sales,
LAG (Sales, 1, CAST (NULL AS INTEGER))
OVER (PARTITION BY Territory
ORDER BY Month) AS prev_sales
FROM Sales_history
```

- This query displays Sales for the current and previous month.
- The first row of each partition returns the NULL value for the LAG() function. Note the NULL value is cast to the same type as Sales.

**Table 22 — Result of aggregate function (LAG)**

Territory	Month	Sales	Prev_Sales
East	199810	10	
East	199811	4	10
East	199812	11	4
East	199901	7	11

Territory	Month	Sales	Prev_Sales
East	199902	10	7
West	199810	8	
West	199811	12	8
West	199817	7	12
West	199901	11	7
West	199902	6	11

As with numerous examples, the results of these queries are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

## 6.8 FIRST\_VALUE and LAST\_VALUE functions

FIRST\_VALUE and LAST\_VALUE are used to compute an expression using the first or last row of a window frame in a given partition. If there is no row in the window frame, then the result of the function is NULL. If the window on which the FIRST\_VALUE or LAST\_VALUE function is being executed has no window frame specification, the function uses the first or last row of the entire partition.

Examples of FIRST\_VALUE and LAST\_VALUE follow:

### Example 23 — first\_value function

```
SELECT Territory, Month, Sales,
       FIRST_VALUE (Sales)
       OVER (PARTITION BY Territory
            ORDER BY Month) AS first_sales
FROM Sales_history
```

This query displays Sales beside the first Sales value of each partition.

**Table 23 — Result of aggregate function (FIRST\_VALUE)**

Territory	Month	Sales	First_Sales
East	199810	10	10
East	199811	4	10
East	199812	11	10
East	199901	7	10
East	199902	10	10
West	199810	8	8
West	199811	12	8

Territory	Month	Sales	First_Sales
West	199817	7	8
West	199901	11	8
West	199902	6	8

**Example 24 — last\_value function**

```
SELECT Territory, Month, Sales,
       LAST_VALUE (Sales)
         OVER (PARTITION BY Territory
              ORDER BY Month
              ROWS BETWEEN 1 PRECEDING AND 1 FOLLOWING) AS last_sales
FROM Sales_history
```

This query displays Sales beside the last Sales value taken from the frame consisting of the preceding, current, and following row being processed.

**Table 24 — Result of aggregate function (LAST\_VALUE)**

Territory	Month	Sales	Last_Sales
East	199810	10	4
East	199811	4	11
East	199812	11	7
East	199901	7	10
East	199902	10	10
West	199810	8	12
West	199811	12	7
West	199817	7	11
West	199901	11	6
West	199902	6	6

As with numerous examples, the results of these queries are ordered for the reader's convenience. An order by clause for the entire query is required to guarantee the result set ordering.

## 6.9 NTH\_VALUE function

The NTH\_VALUE function is a generalization of the FIRST\_VALUE and LAST\_VALUE functions. The first parameter of NTH\_VALUE is the expression to be evaluated using the specified row. The second parameter is an exact numeric literal whose value is used as an offset from the first or last row of the frame (or

6.9 NTH\_VALUE function

the entire partition if there is no window frame specification) to identify the row to be used in evaluating the first argument. The value of the second parameter must be positive.

NTH\_VALUE also includes an optional direction specification to indicate whether the offset is to be taken from the first row of the frame or the last row of the frame. The direction specification immediately follows the NTH\_VALUE function. The option values are FROM FIRST and FROM LAST, with FROM FIRST being the default.

As is the case with the FIRST\_VALUE and LAST\_VALUE functions, if there is no row in the window frame, then the result of the function is NULL.

Examples of NTH\_VALUE follow:

**Example 25 — nth\_value function 1**

```
SELECT Territory, Month, Sales,
       NTH_VALUE (Sales, 2)
       OVER (PARTITION BY Territory
            ORDER BY Month) AS third_sales
FROM Sales_history
```

This query displays Sales beside the third Sales value in each partition.

Because the partitions each have 5 rows, the same result can be achieved by:

**Example 26 — nth\_value function 2**

```
SELECT Territory, Month, Sales,
       NTH_VALUE (Sales, 4) FROM LAST
       OVER (PARTITION BY Territory
            ORDER BY Month) AS third_sales
FROM Sales_history
```

**Table 25 — Result of aggregate function (NTH\_VALUE)**

Territory	Month	Sales	Third_Sales
East	199810	10	4
East	199811	4	4
East	199812	11	4
East	199901	7	4
East	199902	10	4
West	199810	8	12
West	199811	12	12
West	199812	7	12
West	199901	11	12
West	199902	6	12

### 6.9.1 Null treatment

LEAD and LAG (and the FIRST\_VALUE, LAST\_VALUE, and NTH\_VALUE functions) also include an option to describe how to handle NULL values in the selected row. The null treatment options are RESPECT NULLS and IGNORE NULLS and follow the function specification. If neither RESPECT NULLS nor IGNORE NULLS is specified, RESPECT NULLS is implicit.

When RESPECT NULLS is in effect, the entire window frame of the current row is used to evaluate the function. If IGNORE NULLS is in effect, only those rows from the window frame of the current row for which the FIRST\_VALUE, LAST\_VALUE, or NTH\_VALUE expression evaluates to a non-null value are considered in determining the result.

The following are syntax examples of null treatment specification:

#### Example 27 — null treatment with lead function

```
SELECT Territory, Month, Sales,  
       LEAD (Sales, 1, CAST (NULL AS INTEGER))  
       IGNORE NULLS OVER (PARTITION BY Territory  
                          ORDER BY Month)  
FROM Sales_history
```

This function will only display non-null Sales values. However, if all Sales values in a given partition are null, the third argument of the function will still return a null value.

#### Example 28 — null treatment with first\_value function

```
SELECT Territory, Month, Sales,  
       FIRST_VALUE (Sales)  
       IGNORE NULLS OVER (PARTITION BY Territory  
                          ORDER BY Month)  
FROM Sales_history
```

This function will display the first non-null value of Sales in each partition.

#### Example 29 — null treatment with last\_value function

```
SELECT Territory, Month, Sales,  
       LAST_VALUE (Sales)  
       RESPECT NULLS OVER (PARTITION BY Territory  
                           ORDER BY Month  
                           ROWS BETWEEN 1 PRECEDING AND 1 FOLLOWING)  
FROM Sales_history
```

This function will display the last value of Sales in each partition, regardless of whether it is null. RESPECT NULLS is the default specification and may be omitted in this example.

## 7 Nested window functions

### 7.1 Introduction to nested window functions

Feature T619, “Nested window functions” introduces additional functionality to the windowed OLAP functions described in [Clause 6, “Window functions”](#), above.

Window definitions and functions as described in [Clause 5, “Windows”](#) and [Clause 6, “Window functions”](#), above provide a powerful mechanism whereby ranking, cumulative, moving, and reporting aggregate calculations can be specified succinctly in SQL and evaluated efficiently by the RDBMS. Using this specification, the user or an analytic application can define a window frame (fixed or variable in size) for each row and have an aggregate applied on all rows in the window frame.

For example, assuming the table Stocks (Ticker, TradeDay, Price), a query to find, for each stock ticker and each trade day, the number of times a stock price has exceeded \$30 in the past 30 days of trading can be written like this:

#### Example 30 — Q1: case expression in a window query

```
SELECT Ticker, TradeDay, Price,
       SUM(CASE WHEN Price > 30 THEN 1 ELSE 0 END)
       OVER (PARTITION BY Ticker
            ORDER BY TradeDay
            ROWS BETWEEN 30 PRECEDING AND 1 PRECEDING) AS Freq
FROM Stocks
```

Though the window specification is powerful and allows efficient evaluation of ranking, moving, cumulative, and reporting aggregates by the RDBMS, it lacks the ability to solve what could be called “comparative window function” calculations. These are calculations wherein the window function value for the current row is the result of a comparative analysis involving one or more rows in the window frame.

For example, window functions as described in [Clause 5, “Windows”](#) and [Clause 6, “Window functions”](#) are not useful in answering queries like “find the number of times the stock price has exceeded a day's price in the preceding 30 days”, instead of a fixed \$30 value as in the above example. SQL to answer this query will have to resort to self-joins and becomes difficult to write, understand, and optimize:

#### Example 31 — Q2: complex join

```
SELECT s1.Ticker, s1.TradeDay, s1.Price,
       SUM(CASE WHEN s2.Price > s1.Price THEN 1 ELSE 0 END) AS Freq
FROM Stocks s1, Stocks s2
WHERE s1.Ticker = s2.Ticker
AND s2.TradeDay <= s1.TradeDay-1
AND s2.TradeDay >= s1.TradeDay-30
GROUP BY s1.Ticker, s1.TradeDay, s1.Price
```

To address queries like Q2, Feature T619, “Nested window functions” extends window aggregates, making them more powerful and suitable for “comparative analysis”. The extensions include:

- Row markers that give access to certain interesting rows.
- Offsets that provide ability to access a row at any offset from the marker rows.
- A row marker that moves across the rows in the window frame during calculation of a window aggregate.

## 7.2 Row markers

Row markers give access to interesting rows within a partition and frame for comparative analysis. The row markers that are defined in Feature T619, “Nested window functions” are:

- BEGIN\_PARTITION — the first row of the window partition
- BEGIN\_FRAME — first row of the window frame
- CURRENT\_ROW — the row for which a window frame is created
- END\_FRAME — the last row of the window frame
- END\_PARTITION — the last row of the window partition
- FRAME\_ROW — the variable row that ranges over all rows of a window frame (between BEGIN\_FRAME and END\_FRAME) when computing a window aggregate.

For example, consider the following data in the Stocks table, partitioned by Ticker, ordered by TradeDay, and with the window framing clause ROWS BETWEEN 4 PRECEDING AND 1 PRECEDING. Assume that “(ZYX, 6, 11)” is the row for which we are currently computing the window function. The following illustrates where the row markers other than FRAME\_ROW point:

**Table 26 — Result of row markers**

Ticker	Tradeday	Price	
ZYX	1	10	-- BEGIN_PARTITION
ZYX	2	11	-- BEGIN_FRAME
ZYX	3	12	
ZYX	4	12	
ZYX	5	12	-- END_FRAME
ZYX	6	11	-- CURRENT_ROW
ZYX	7	12	
ZYX	8	12	-- END_PARTITION

To access individual columns from the marker rows, the function VALUE\_OF has been introduced. This function is only permitted nested in an aggregated argument of a window aggregate. For example “VALUE\_OF (Price AT CURRENT\_ROW)” gives the value of column “Price” from the row indicated by the row marker CURRENT\_ROW. Using the CURRENT\_ROW marker, the self-join query Q2 can be rewritten elegantly as:

### Example 32 — Q3: value\_of function

```
SELECT Ticker, TradeDay, Price,
       SUM (CASE WHEN Price > VALUE_OF (Price AT CURRENT_ROW)
              THEN 1
              ELSE 0 END)
       OVER (PARTITION BY Ticker
            ORDER BY TradeDay
```

```
ROWS BETWEEN 30 PRECEDING AND 1 PRECEDING) AS Freq
FROM Stocks;
```

Note that Q3 uses ROWS BETWEEN 30 PRECEDING AND 1 PRECEDING, whereas the preceding sample data showed ROWS BETWEEN 4 PRECEDING AND 1 PRECEDING for brevity.

The CASE expression in Q3 illustrates the use of VALUE\_OF, nested in the window aggregate SUM. The SUM is adding a collection of 1's and 0's; i.e., counting the number of times the WHEN condition in the CASE expression is true. The WHEN condition involves a comparison of the Price column in two rows. The first Price is found in the row that varies throughout the window frame in the course of evaluating the aggregate. The second Price, in VALUE\_OF (Price AT CURRENT\_ROW), is the value of Price in a fixed row, the row that “owns” the window frame and serves as the basis for the window frame's formation via BETWEEN 30 PRECEDING AND 1 PRECEDING. Thus, Q3 computes the number of times in the preceding 30 days that the value of Price has exceeded the current row's Price. By eliminating the self-join, the query Q3 is easier to write and understand. Moreover, the query should also be easier to optimize by the RDBMS.

For convenience, expressions may be the first argument of VALUE\_OF function. For example, instead of having to write multiple functions as in:

```
VALUE_OF(sal AT CURRENT_ROW) +
VALUE_OF(commission AT CURRENT_ROW) -
VALUE_OF(revenue AT CURRENT_ROW),
```

one can code:

```
VALUE_OF(sal + commission - revenue AT CURRENT_ROW)
```

### 7.3 Offsets

Positive and negative integer offsets are also supported to access rows at a specified offset from the marker rows. For example one can say “VALUE\_OF (Price AT BEGIN\_FRAME + 1)” to get the value of column “Price” from the second row in the window frame. Similarly, “VALUE\_OF(Price AT END\_FRAME - 1)” gives access to Price from the next to last row in the window frame.

When an offset specifies a row outside of the window partition (i.e., before BEGIN\_PARTITION or after END\_PARTITION), then VALUE\_OF will return NULL. To illustrate the semantics, consider a window frame that is “BETWEEN 5 PRECEDING AND 2 PRECEDING” operating on the following data (Ticker, TradeDay, Price). Assume that the CURRENT\_ROW is at (ZYX, 7, 12) and the Price column is being accessed using the VALUE\_OF function.

Table 27 — Result of row markers (offsets)

Ticker	Tradeday	Price	
ZYX	1	10	-- BEGIN_PARTITION
ZYX	2	11	-- BEGIN_FRAME
ZYX	3	12	
ZYX	4	12	
ZYX	5	12	-- END_FRAME

Ticker	Tradeday	Price	
ZYX	6	11	
ZYX	7	12	-- CURRENT_ROW
ZYX	8	12	-- END_PARTITION

Then:

VALUE\_OF (Price AT BEGIN\_FRAME+1) = 12  
 VALUE\_OF (Price AT BEGIN\_FRAME-1) = 10  
 VALUE\_OF (Price AT BEGIN\_FRAME-2) is NULL  
 VALUE\_OF (Price AT CURRENT\_ROW) = 12

Of course, the user may wish a different default than NULL when a row outside the window partition is referenced; for that an optional third parameter with the same data type as the first is supported (the second parameter is the row marker defined location of the desired row):

VALUE\_OF (Price AT BEGIN\_FRAME-2, 5) = 5  
 VALUE\_OF (Price AT CURRENT\_ROW+2, 6) = 6

The VALUE\_OF function can be nested within itself, which means that the default for an outer VALUE\_OF can be computed by an inner VALUE\_OF. For example:

VALUE\_OF (Price AT CURRENT\_ROW - 5, VALUE\_OF (Price AT BEGIN\_PARTITION))

In this example, the user wants the Price in the row 5 rows prior to the current row. This row might fall outside the partition, in which case the default applies. The default is another VALUE\_OF expression. In this example, the default is the Price in the row at the beginning of the partition.

## 7.4 FRAME\_ROW

Given a particular current row, most of the row markers have fixed positions within the window partition. FRAME\_ROW, on the other hand, is a variable row that moves from BEGIN\_FRAME to END\_FRAME in the window frame as the window aggregate is computed. In Table 28, "Result of window frames with row markers", using a window frame BETWEEN 4 PRECEDING AND 1 PRECEDING, FRAME\_ROW moves from row (ZYX, 2, 11) to (ZYX, 5, 12), when the CURRENT\_ROW is at (ZYX, 6, 11).

Table 28 — Result of window frames with row markers

Ticker	Tradeday	Price	
ZYX	1	10	
ZYX	2	11	-- BEGIN_FRAME -- FRAME_ROW
ZYX	3	12	-- FRAME_ROW
ZYX	4	12	-- FRAME_ROW
ZYX	5	12	-- END_FRAME -- FRAME_ROW
ZYX	6	11	-- CURRENT_ROW
ZYX	7	12	

Ticker	Tradeday	Price	
ZYX	8	12	-- END_PARTITION

As with other markers, VALUE\_OF notation can be used to get column values or more general complex expressions. Note that VALUE\_OF (value expression AT FRAME\_ROW+1) gives access to the row after FRAME\_ROW and VALUE\_OF (value expression AT FRAME\_ROW-1) gives the row previous to FRAME\_ROW. It may be helpful to note that VALUE\_OF (value expression AT FRAME\_ROW) is just the same as value expression. For example, if W is a window name, then AVG (VALUE\_OF (Price AT FRAME\_ROW)) OVER W is just the same as AVG (Price) OVER W.

This is because AVG (Price) is computed as the average of Price as the moving row varies from BEGIN\_FRAME through END\_FRAME. Thus, the benefit of FRAME\_ROW is the ability to access rows that are offset from the FRAME\_ROW. However, for consistency with other row markers, VALUE\_OF is also supported at FRAME\_ROW with no offset. Using FRAME\_ROW, queries like “for each day D, how many times in the preceding 30-day period, a stock price has exceeded D’s value and maintained (either stayed the same or has increased) on the next day” can be answered efficiently using window function. The SQL would be:

#### Example 33 — Q4: frame\_row and current\_row in value\_of function

```
SELECT Ticker, TradeDay, Price,
       SUM(CASE WHEN VALUE_OF(Price AT FRAME_ROW) >
                 VALUE_OF(Price AT CURRENT_ROW)
                 AND VALUE_OF(Price AT FRAME_ROW+1) >=
                 VALUE_OF(Price AT FRAME_ROW)
                 THEN 1 ELSE 0 END)
       OVER (PARTITION BY Ticker
            ORDER BY TradeDay
            ROWS BETWEEN 30 PRECEDING AND 1 PRECEDING) AS Freq
FROM Stocks;
```

In Q4, the first comparison in the CASE expression looks for a FRAME\_ROW whose Price is greater than the Price in CURRENT\_ROW. And the second comparison looks for a FRAME\_ROW whose Price is equaled or exceeded in the next row. Both conditions being satisfied indicates a “hit” and adds 1 to the sum. The final result is the number of pairs (R1, R2) in which the first row R1 beats the current row's Price, and the second row R2 has a Price at least as high as R1.

## 7.5 Nested ROW\_NUMBER function

Subclause 7.3, “Offsets”, describes an optional third parameter to VALUE\_OF to override the default null value when a row marker plus or minus offset that points to a non-existent row (outside the window partition) is used in VALUE\_OF. This will handle most requirements to deal with boundary situations. However, the user may wish more sophisticated control, for example, to specify different handling at the beginning of the window partition as opposed to the end of the window partition. To cover such cases, the ROW\_NUMBER function may be nested in the window aggregate function with a row marker argument. The syntax is:

```
ROW_NUMBER (row marker)
```

Unlike VALUE\_OF, offsets are not permitted for nested ROW\_NUMBER, since the ROW\_NUMBER function commutes with addition and subtraction, e.g., ROW\_NUMBER (CURRENT\_ROW + 1) = ROW\_NUMBER (CURRENT\_ROW) + 1. Nested ROW\_NUMBER provides complete information about where the row markers are within the window partition. Specifically: