Rows and Range, Preceding and Following

SQL Server 2012 adds many new features to Transact SQL (T-SQL). One of my favorites is the Rows/Range enhancements to the over clause. These enhancements are often times referred to as the windowing functions.

Overview:
ROWS PRECEDING, FOLLOWING, UNBOUNDED, refers to the current row and those before or after based on preceding or following.
RANGE PRECEDING, FOLLOWING, UNBOUNDED, means all values in the current range and those before or after.

An example of a need for rows preceding would be to calculate a 3 year trailing average, which needs to look at the current year and the three previous years in the calculation.

Terminology:
ROWS or RANGE- specifying rows or range.
PRECEDING – get rows before the current one.
FOLLOWING – get rows after the current one.
UNBOUNDED – when used with PRECEDING or FOLLOWING, it returns all before or after.
CURRENT ROW

To start out we need a database to work on, so we will create a tsql2012 database. you can use your own database if you wish.

Next we create a table to use for the over clause enhancments of rows and range preceding and

following.

Take a look at the revenue table to see whats there before we start the demo.

Editor | Results | Messages

	DepartmentID	Revenue	Year
1	1	10030	1998
2	2	20000	1998
3	3	40000	1998
4	1	20000	1999
5	2	60000	1999
6	3	50000	1999
7	1	40000	2000
8	2	40000	2000
9	3	60000	2000
10	1	30000	2001
11	2	30000	2001
12	3	70000	2001
13	1	90000	2002
14	2	20000	2002
15	3	80000	2002
16	1	10300	2003
17	2	1000	2003
18	3	90000	2003
19	1	10000	2004
20	2	10000	2004
21	3	10000	2004
22	1	20000	2005

Contents of the REVENUE table.

Then a quick review of sum and avg.

Editor | Results | Messages

	TotalRevenue	AverageRevenue	NumRows
1	1991330	44251	45

OVER Clause Pre SQL Server 2012

The OVER clause before SQL Server 2012 is really handy to perform aggregates over a different range than your standard grouping.

Editor | Results | Messages

	DepartmentID	Reven...	Year	AverageDeptRevenue	TotalDeptReven...
1	1	10030	1998	35355	530330
2	1	20000	1999	35355	530330
3	1	40000	2000	35355	530330
4	1	30000	2001	35355	530330
5	1	90000	2002	35355	530330
6	1	10300	2003	35355	530330
7	1	10000	2004	35355	530330
8	1	20000	2005	35355	530330
9	1	40000	2006	35355	530330
10	1	70000	2007	35355	530330
11	1	50000	2008	35355	530330
12	1	20000	2009	35355	530330
13	1	30000	2010	35355	530330
14	1	80000	2011	35355	530330
15	1	10000	2012	35355	530330
16	2	20000	1998	41400	621000
17	2	60000	1999	41400	621000
18	2	40000	2000	41400	621000
19	2	30000	2001	41400	621000
20	2	20000	2002	41400	621000
21	2	1000	2003	41400	621000
22	2	10000	2004	41400	621000
23	2	20000	2005	41400	621000
24	2	30000	2006	41400	621000
25	2	40000	2007	41400	621000
26	2	50000	2008	41400	621000
27	2	60000	2009	41400	621000
28	2	70000	2010	41400	621000
29	2	80000	2011	41400	621000
30	2	90000	2012	41400	621000
31	3	40000	1998	56000	840000

ROWS PRECEDING

Now the new features, ROWS PRECEDING specifies the the aggregate functions in the current partition in the OVER clause will consider the current row, and a specific number of rows before the current row.

Editor | Results | Messages

	Year	DepartmentID	Revenue	CurrentAndPrev3
1	1998	1	10030	10030
2	1999	1	20000	30030
3	2000	1	40000	70030
4	2001	1	30000	100030
5	2002	1	90000	180000
6	2003	1	10300	170300
7	2004	1	10000	140300
8	2005	1	20000	130300
9	2006	1	40000	80300
10	2007	1	70000	140000
11	2008	1	50000	180000
12	2009	1	20000	180000
13	2010	1	30000	170000
14	2011	1	80000	180000

ROWS FOLLOWING

Now rows following... The ROWS FOLLOWING option specifies a specific number of rows in the current partition to use after the current row.

Editor | Results | Messages

	Year	DepartmentID	Revenue	CurrentAndNext3
1	1998	1	10030	100030
2	1999	1	20000	180000
3	2000	1	40000	170300
4	2001	1	30000	140300
5	2002	1	90000	130300
6	2003	1	10300	80300
7	2004	1	10000	140000
8	2005	1	20000	180000
9	2006	1	40000	180000
10	2007	1	70000	170000
11	2008	1	50000	180000
12	2009	1	20000	140000

Both PRECEDING and FOLLOWING

Using both ROWS PRECEDING and ROWS FOLLOWING allows you to do things like calculate an average including the current year and years both before and after the current year.

Editor | Results | Messages

	Year	DepartmentID	Revenue	BeforeAndAfter
1	1998	1	10030	30030
2	1999	1	20000	70030
3	2000	1	40000	90000
4	2001	1	30000	160000
5	2002	1	90000	130300
6	2003	1	10300	110300
7	2004	1	10000	40300
8	2005	1	20000	70000
9	2006	1	40000	130000
10	2007	1	70000	160000
11	2008	1	50000	140000
12	2009	1	20000	100000
13	2010	1	30000	130000
14	2011	1	80000	120000

UNBOUNDED

UNBOUNDED PRECEDING tells the windowing function and aggregrates to use the current value, and all values in the partition before the current value.

Editor | Results | Messages

	Year	DepartmentID	Revenue	MinRevenueToDate
1	1998	1	10030	10030
2	1999	1	20000	10030
3	2000	1	40000	10030
4	2001	1	30000	10030
5	2002	1	90000	10030
6	2003	1	10300	10030
7	2004	1	10000	10000
8	2005	1	20000	10000
9	2006	1	40000	10000
10	2007	1	70000	10000
11	2008	1	50000	10000
12	2009	1	20000	10000
13	2010	1	30000	10000
14	2011	1	80000	10000

When using UNBOUNDED FOLLOWING as follows, this means the use the current row, and consider all rows after the current row in the current partition.

Editor | Results | Messages

	Year	DepartmentID	Revenue	MinRevenueBeyond
1	1998	1	10030	10000
2	1999	1	20000	10000
3	2000	1	40000	10000
4	2001	1	30000	10000
5	2002	1	90000	10000
6	2003	1	10300	10000
7	2004	1	10000	10000
8	2005	1	20000	10000
9	2006	1	40000	10000
10	2007	1	70000	10000
11	2008	1	50000	10000
12	2009	1	20000	10000
13	2010	1	30000	10000
14	2011	1	80000	10000

ROWS vs RANGE

ROWS vs RANGE appears to cause much confusion.

ROWS means the specific row or rows specified, and RANGE refers to those same rows plus any others that have the same matching values.

Here is an example of ROWS and RANGE producing the exact same results because the years and department IDs don't have any duplicates. In this case the ROWS and RANGE are identical.

Editor | Results | Messages

	Year	DepartmentID	Revenue	RowsCumulative	RangeCumulative
1	2003	1	10300	10300	10300
2	2004	1	10000	20300	20300
3	2005	1	20000	40300	40300
4	2006	1	40000	80300	80300
5	2007	1	70000	150300	150300
6	2008	1	50000	200300	200300
7	2003	2	1000	1000	1000
8	2004	2	10000	11000	11000
9	2005	2	20000	31000	31000
10	2006	2	30000	61000	61000
11	2007	2	40000	101000	101000
12	2008	2	50000	151000	151000
13	2003	3	90000	90000	90000
14	2004	3	10000	100000	100000

Now if we consider the same query of ROWS vs RANGE with duplicate values. Here we insert a duplicate set of values for 2005, and the results below show that for the year 2005 there are two rows in the result set, and the RowsCumulative column is different form the RangeCumulative. The rows counts the specific rows it the aggregate calculation, the range counts the current row, plus the other 2005 rows for that department id.

Editor | Results | Messages

	Year	DepartmentID	Revenue	RowsCumulative	RangeCumulative
1	2003	1	10300	10300	10300
2	2004	1	10000	20300	20300
3	2005	1	20000	40300	50300
4	2005	1	10000	50300	50300
5	2006	1	40000	90300	90300
6	2007	1	70000	160300	160300
7	2008	1	50000	210300	210300
8	2003	2	1000	1000	1000
9	2004	2	10000	11000	11000
10	2005	2	20000	31000	51000
11	2005	2	20000	51000	51000
12	2006	2	30000	81000	81000
13	2007	2	40000	121000	121000
14	2008	2	50000	171000	171000

Summary

With all the great new features in SQL Server 2012, the windowing functions, ROWS and RANGE, PRECEDING and FOLLOWING updates to the OVER clause are a great addition.