

A dramatic silhouette of an elephant's head and trunk against a vibrant sunset sky. The sun is a bright, glowing orb near the horizon, casting a warm orange and yellow light across the clouds. The elephant's trunk is curved downwards, and its ears are visible on the sides of its head. The overall mood is serene and majestic.

How PostgreSQL's  
SQL dialect stays ahead  
of its competitors

@MarkusWinand



*I'm sorry!*



*I couldn't make it  
to PgConf.EU in 2014!*





*Don't repeat my mistake.*



A nighttime street scene in Lisbon, Portugal. On the left, a yellow and blue tram is visible, with a window showing people inside. The street is narrow, with white buildings on the right and a few people walking. The text "Come to PgConf.EU 2018 in Lisbon! Oct 23-26." is overlaid in white cursive script on the right side of the image.

*Come to  
PgConf.EU 2018  
in Lisbon!  
Oct 23-26.*



A dramatic silhouette of an elephant's head and trunk against a vibrant sunset sky. The sun is a bright, glowing orb near the horizon, casting a warm orange and yellow light across the clouds. The elephant's trunk is curved downwards, and its ears are visible on the sides of its head. The overall mood is serene and majestic.

How PostgreSQL's  
SQL dialect stays ahead  
of its competitors

@MarkusWinand



# PostgreSQL 9.5

2016-01-07

# GROUPING SETS



# GROUPING SETS

Before SQL:1999

---

Only one **GROUP BY** operation at a time:

Monthly revenue

```
SELECT year
      , month
      , sum(revenue)
FROM tbl
GROUP BY year, month
```

Yearly revenue

```
SELECT year
      , sum(revenue)
FROM tbl
GROUP BY year
```



# GROUPING SETS

Before SQL:1999

---

```
SELECT year
      , month
      , sum(revenue)
  FROM tbl
 GROUP BY year, month
UNION ALL
SELECT year
      , null
      , sum(revenue)
  FROM tbl
 GROUP BY year
```



# GROUPING SETS

Since SQL:1999

---

```
SELECT year
      , month
      , sum(revenue)
FROM tbl
GROUP BY year, month
UNION ALL
SELECT year
      , null
      , sum(revenue)
FROM tbl
GROUP BY year
```

```
SELECT year
      , month
      , sum(revenue)
FROM tbl
GROUP BY
GROUPING SETS (
      (year, month)
      , (year)
)
```



# GROUPING SETS

In a Nutshell

---

**GROUPING SETS** are multiple **GROUP BY**s in one go

**()** (empty parenthesis) build a group over all rows

**GROUPING** (function) disambiguates the meaning of **NULL**

(was the grouped data **NULL** or is this column not currently grouped?)

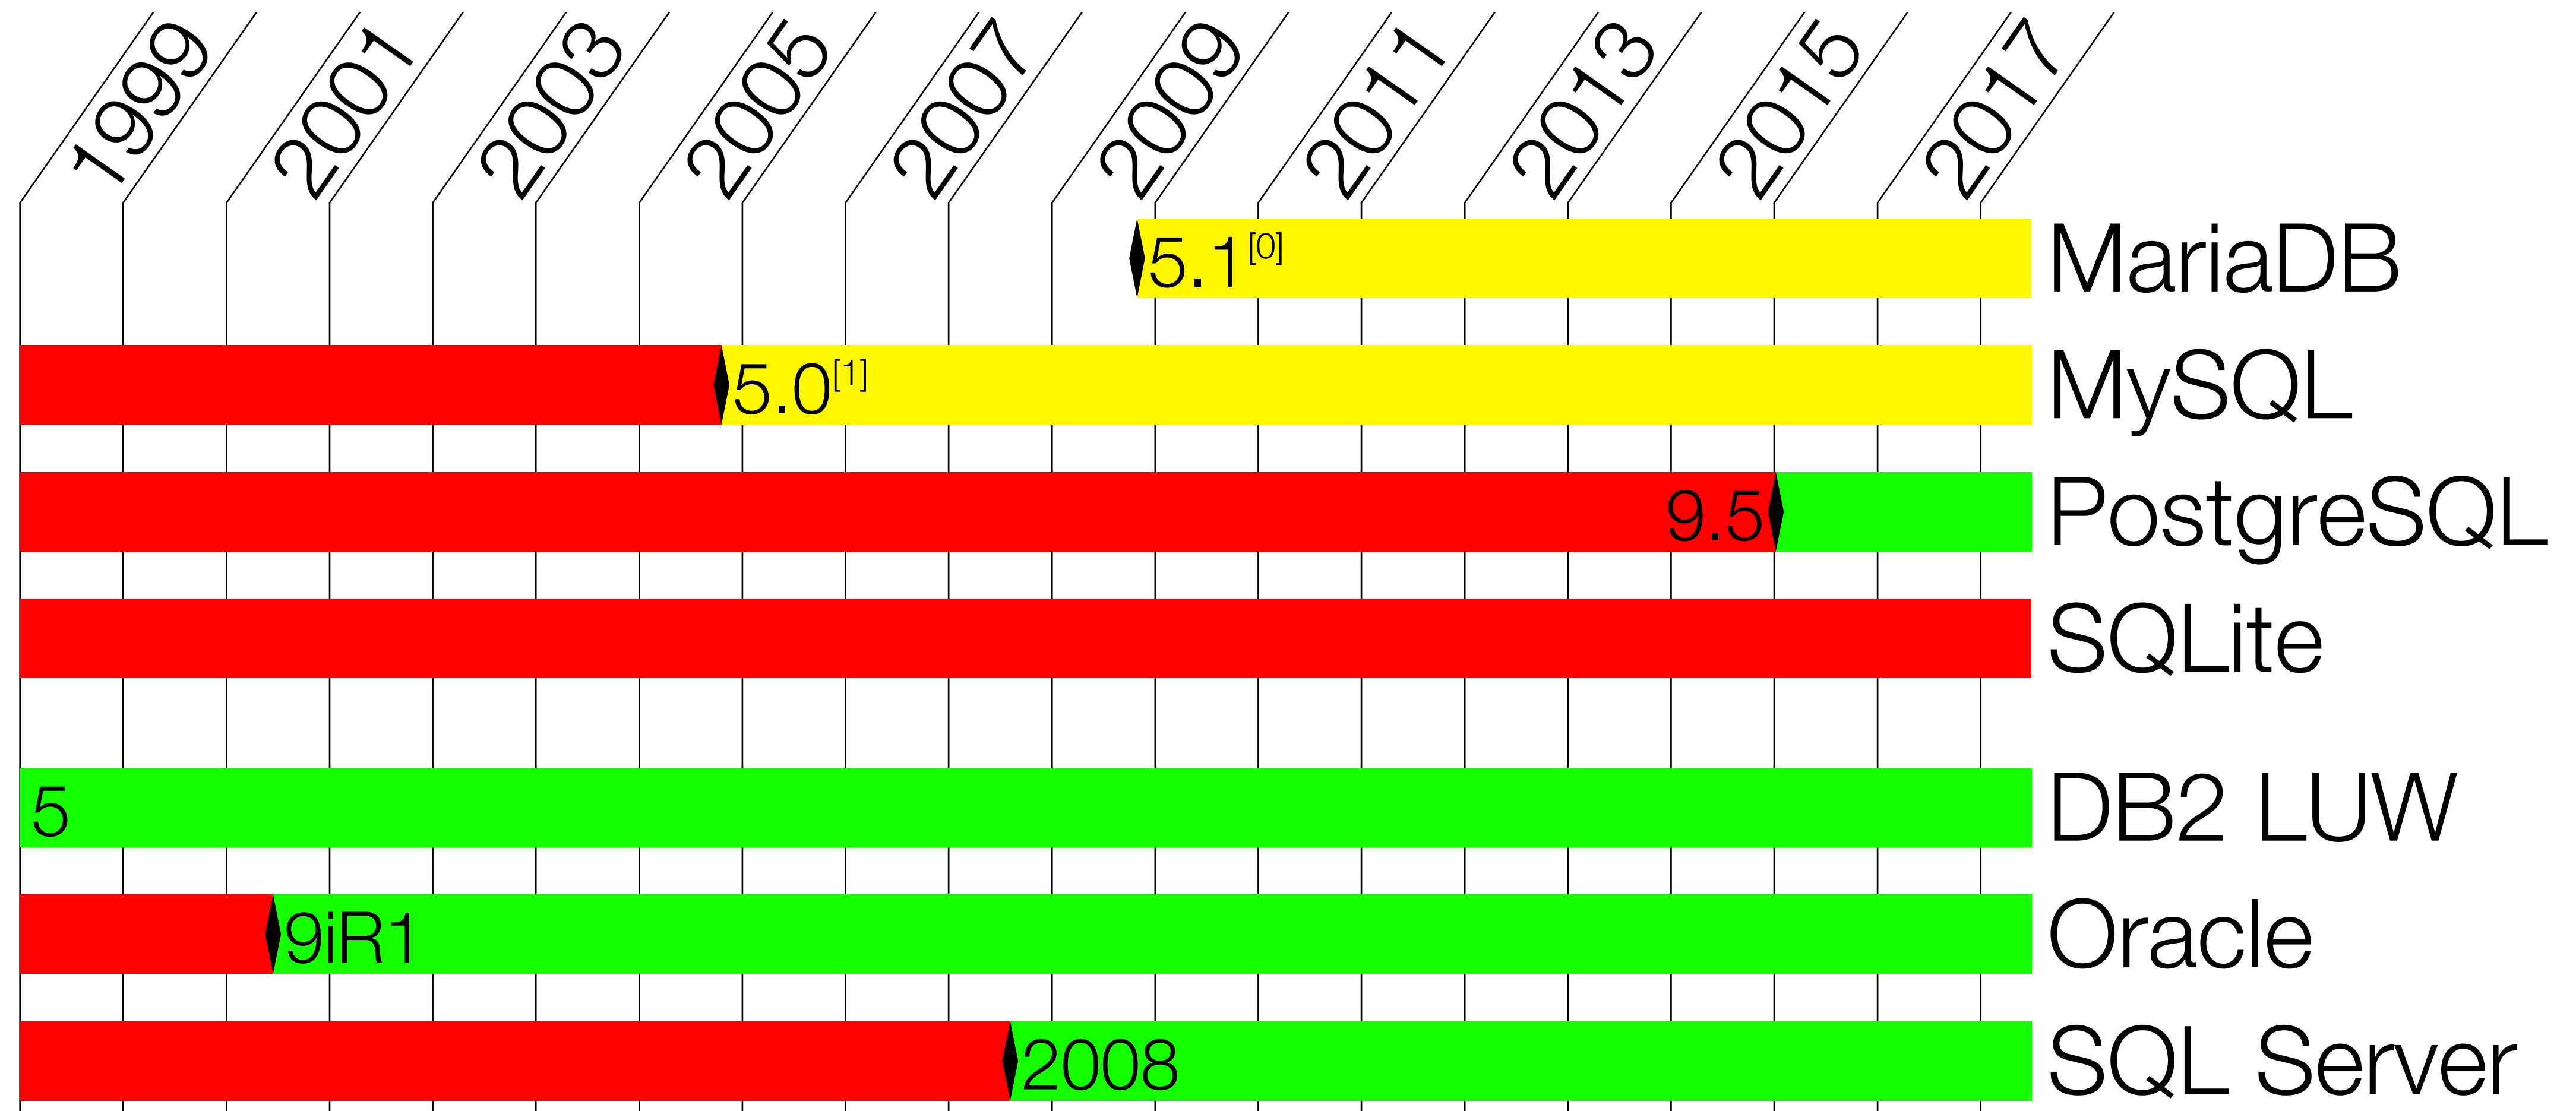
Permutations can be created using **ROLLUP** and **CUBE**

**(ROLLUP(a, b, c) = GROUPING SETS ((a, b, c), (a, b), (a), ()))**



# GROUPING SETS

Availability



[0] Only ROLLUP (proprietary syntax).

[1] Only ROLLUP (proprietary syntax). GROUPING function since MySQL 8.0.



**TABLESAMPLE**



# TABLESAMPLE

---

Since SQL:2003

```
SELECT *  
FROM tbl TABLESAMPLE system(10)
```

*Or:*

*bernoulli(10)  
better sampling,  
way slower.*



# TABLESAMPLE

---

Since SQL:2003

```
SELECT *  
FROM tbl TABLESAMPLE system(10)  
REPEATABLE (0)
```

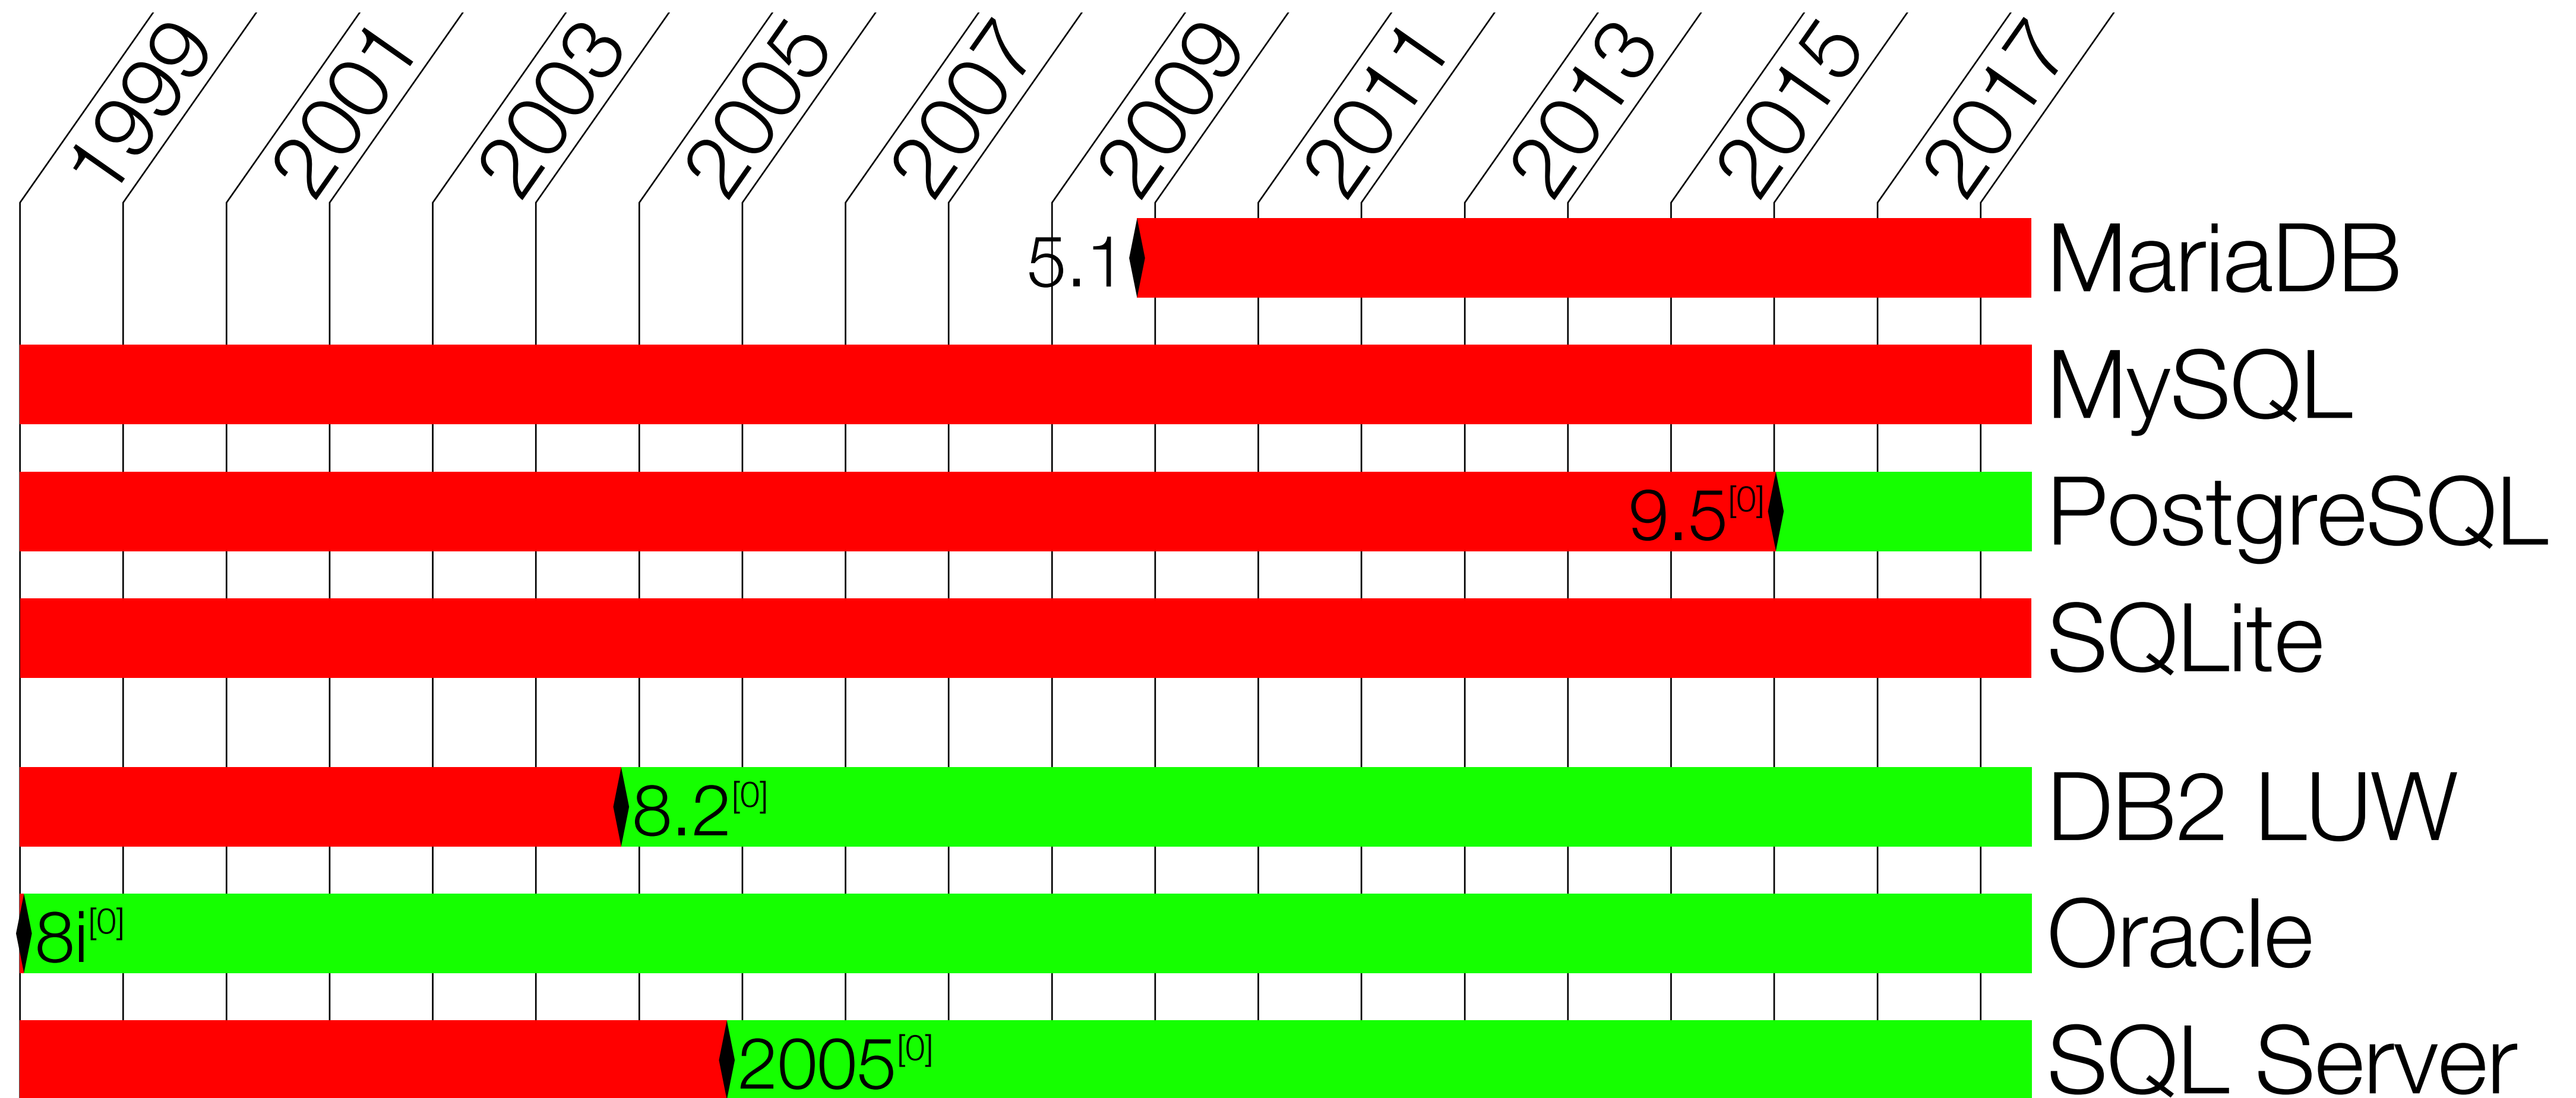
*Or:  
bernoulli(10)  
better sampling,  
way slower.*

*User  
provided seed  
value*



# TABLESAMPLE

Availability



<sup>[0]</sup>Not for derived tables



# PostgreSQL 10

2017-10-05



**XMLTABLE**



# XMLTABLE

Since SQL:2006

```
FROM tbl  
  , XMLTABLE(  
    '/d/e'
```

*XPath\* expression  
to identify rows*

Stored in tbl.x:

```
<d>  
  <e id="42">  
    <c1>...</c1>  
  </e>  
</d>
```

```
) r
```

\*Standard SQL allows XQuery



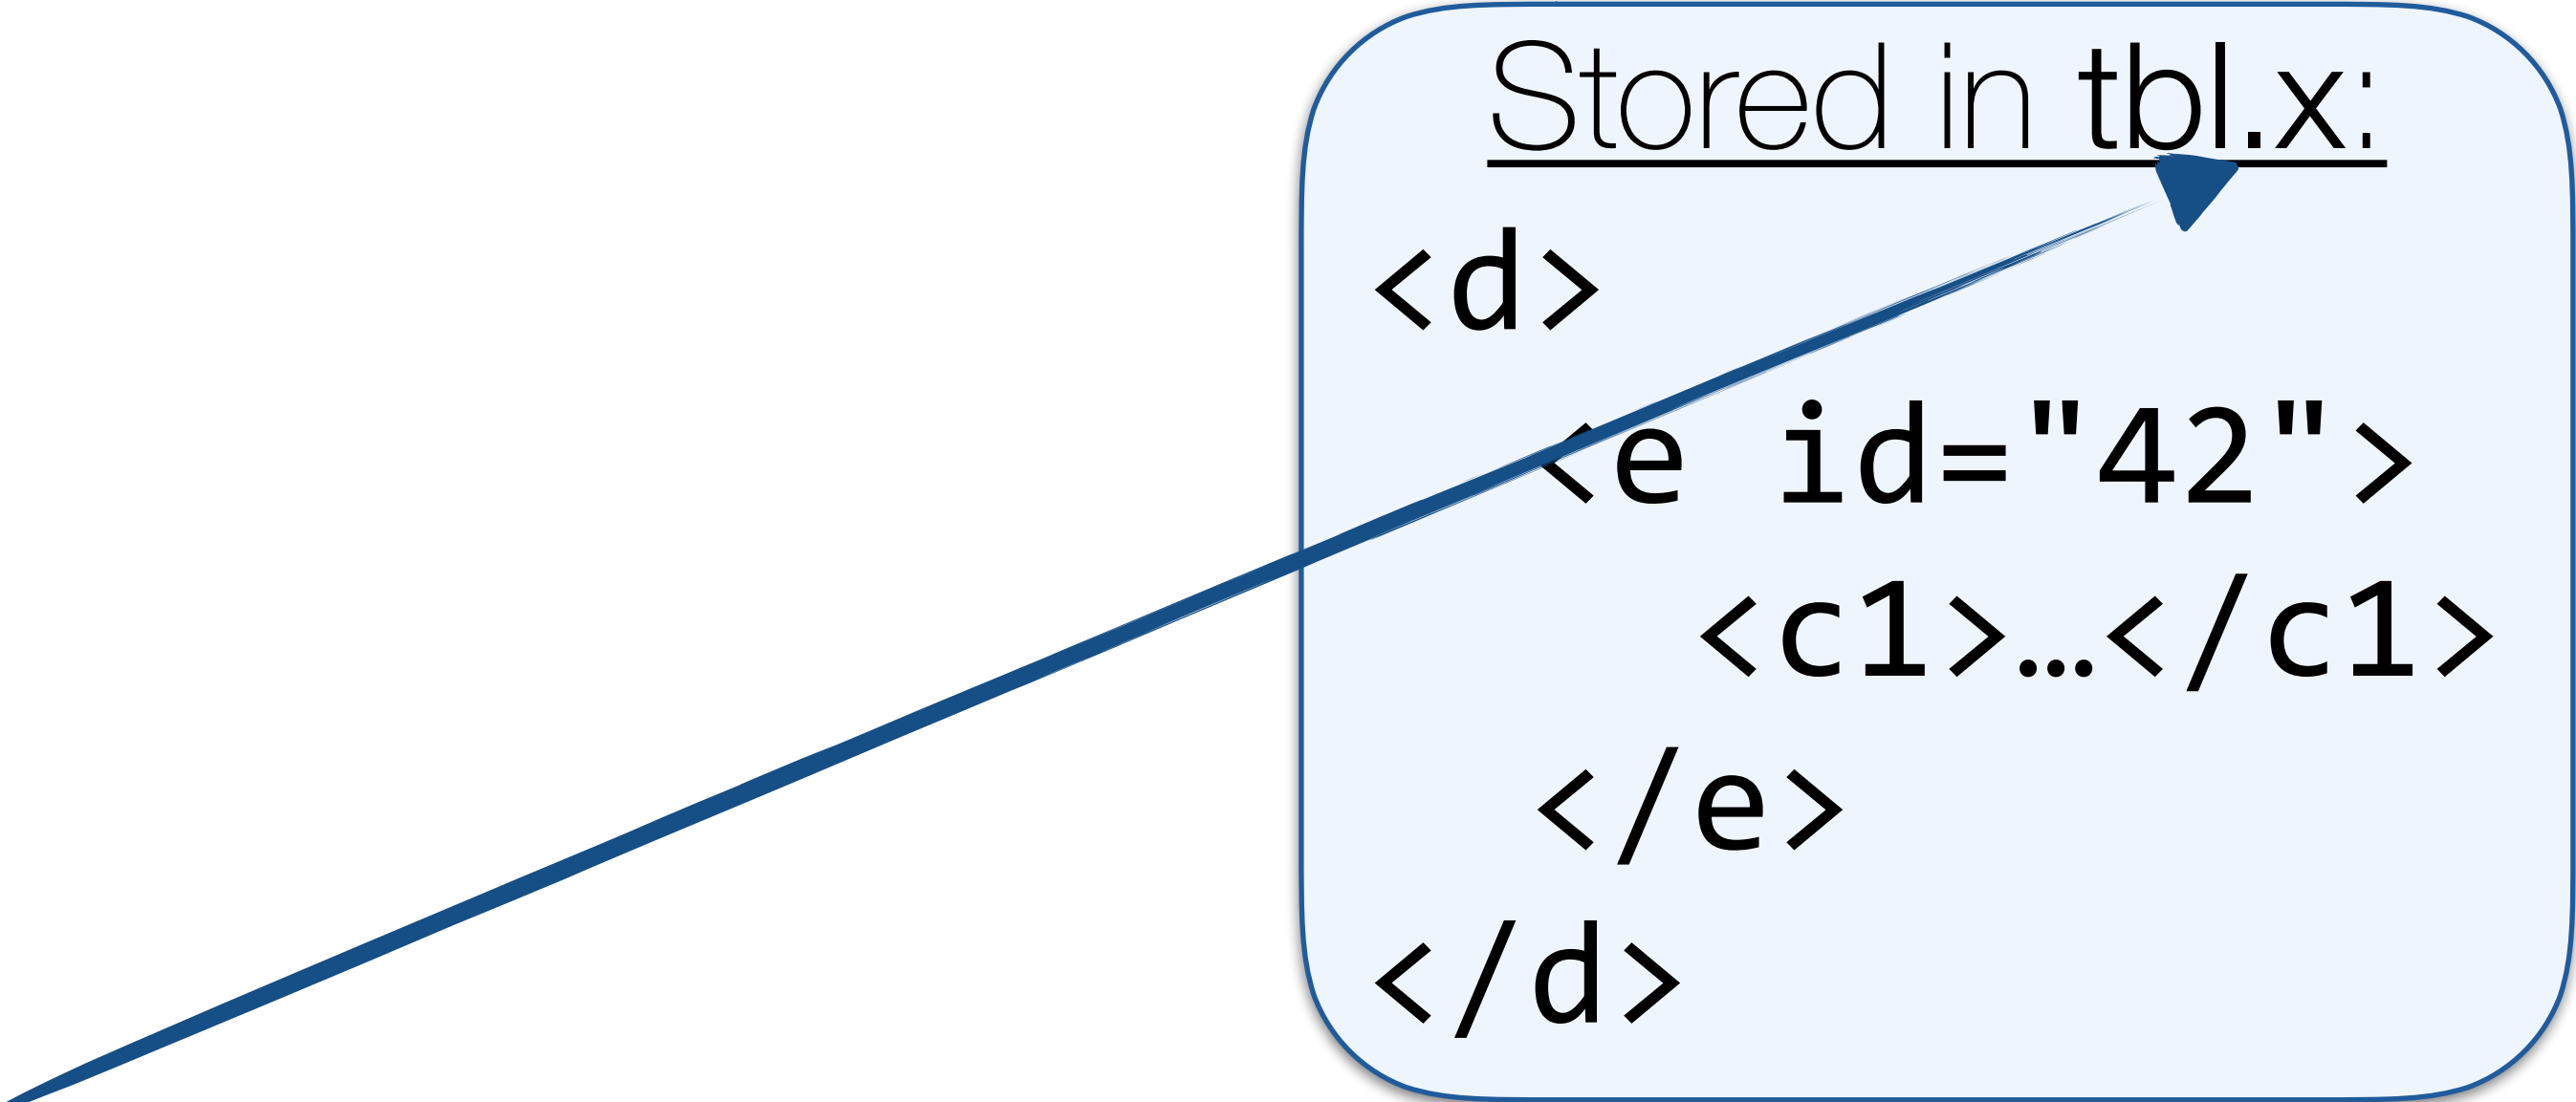
# XMLTABLE

Since SQL:2006

```
FROM tbl  
  , XMLTABLE(  
    '/d/e'  
    PASSING x
```

```
) r
```

Stored in tbl.x:



```
<d>  
  <e id="42">  
    <c1>...</c1>  
  </e>  
</d>
```

\*Standard SQL allows XQuery



# XMLTABLE

Since SQL:2006

```
FROM tbl
  , XMLTABLE(
    '/d/e'
    PASSING x
    COLUMNS id INT PATH '@id'
             , c1 VARCHAR(255) PATH 'c1'

  ) r
```

*XPath\* expressions  
to extract data*

Stored in tbl.x:

```
<d>
  <e id="42">
    <c1>...</c1>
  </e>
</d>
```

\*Standard SQL allows XQuery



# XMLTABLE

Since SQL:2006

```
FROM tbl
  , XMLTABLE(
    '/d/e'
    PASSING x
    COLUMNS id INT PATH '@id'
              , c1 VARCHAR(255) PATH 'c1'
              , n
              FOR ORDINALITY
  ) r
```

Stored in tbl.x:

```
<d>
  <e id="42">
    <c1>...</c1>
  </e>
</d>
```

*Row number  
(like for unnest)*

\*Standard SQL allows XQuery

# XMLTABLE

Since SQL:2006

```
SELECT id
      , c1
      , n
FROM tbl
      , XMLTABLE(
          '/d/e'
        PASSING x
        COLUMNS id INT PATH '@id'
                  , c1 VARCHAR(255) PATH 'c1'
                  , n
                                FOR ORDINALITY
        ) r
```

## Result

id	c1	n
42	...	1

## Stored in tbl.x:

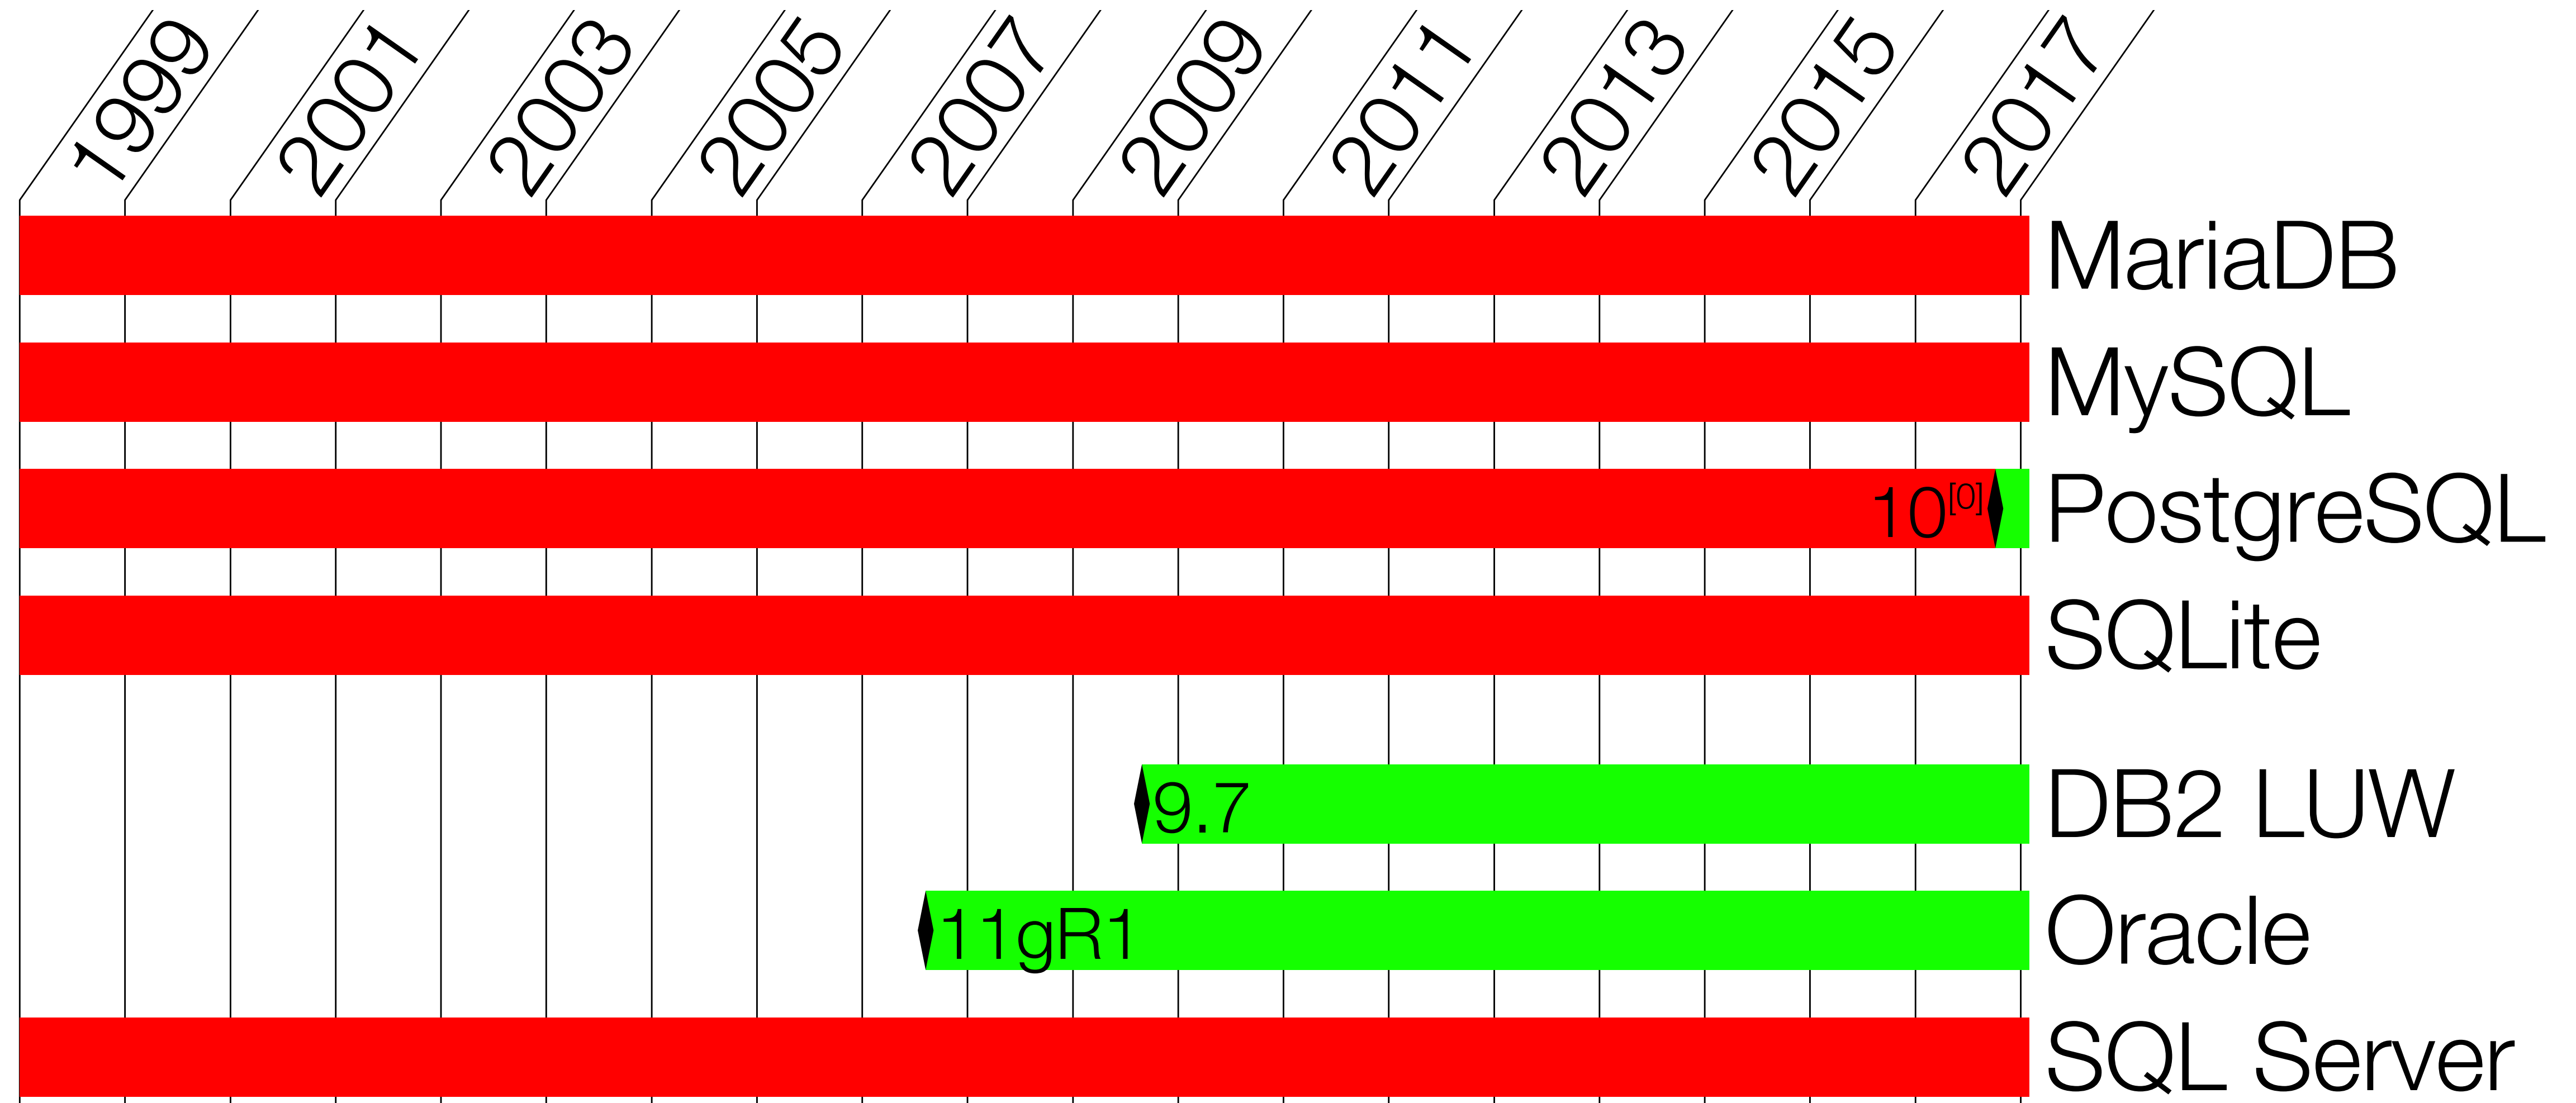
```
<d>
  <e id="42">
    <c1>...</c1>
  </e>
</d>
```

\*Standard SQL allows XQuery



# XMLTABLE

Availability



<sup>[0]</sup>No XQuery (only XPath). No default namespace declaration.

**IDENTITY COLUMNS**



# IDENTITY column

Since SQL:2003

---

Similar to **serial** columns, but standard and more powerful.

```
CREATE TABLE tbl (  
    id INTEGER GENERATED BY DEFAULT AS IDENTITY  
  
    , PRIMARY KEY (id)  
)
```

# IDENTITY column

Since SQL:2003

---

Similar to **serial** columns, but standard and more powerful.

```
CREATE TABLE tbl (  
    id INTEGER GENERATED BY DEFAULT AS IDENTITY  
        (START WITH 10 INCREMENT BY 10)  
  
    , PRIMARY KEY (id)  
)
```

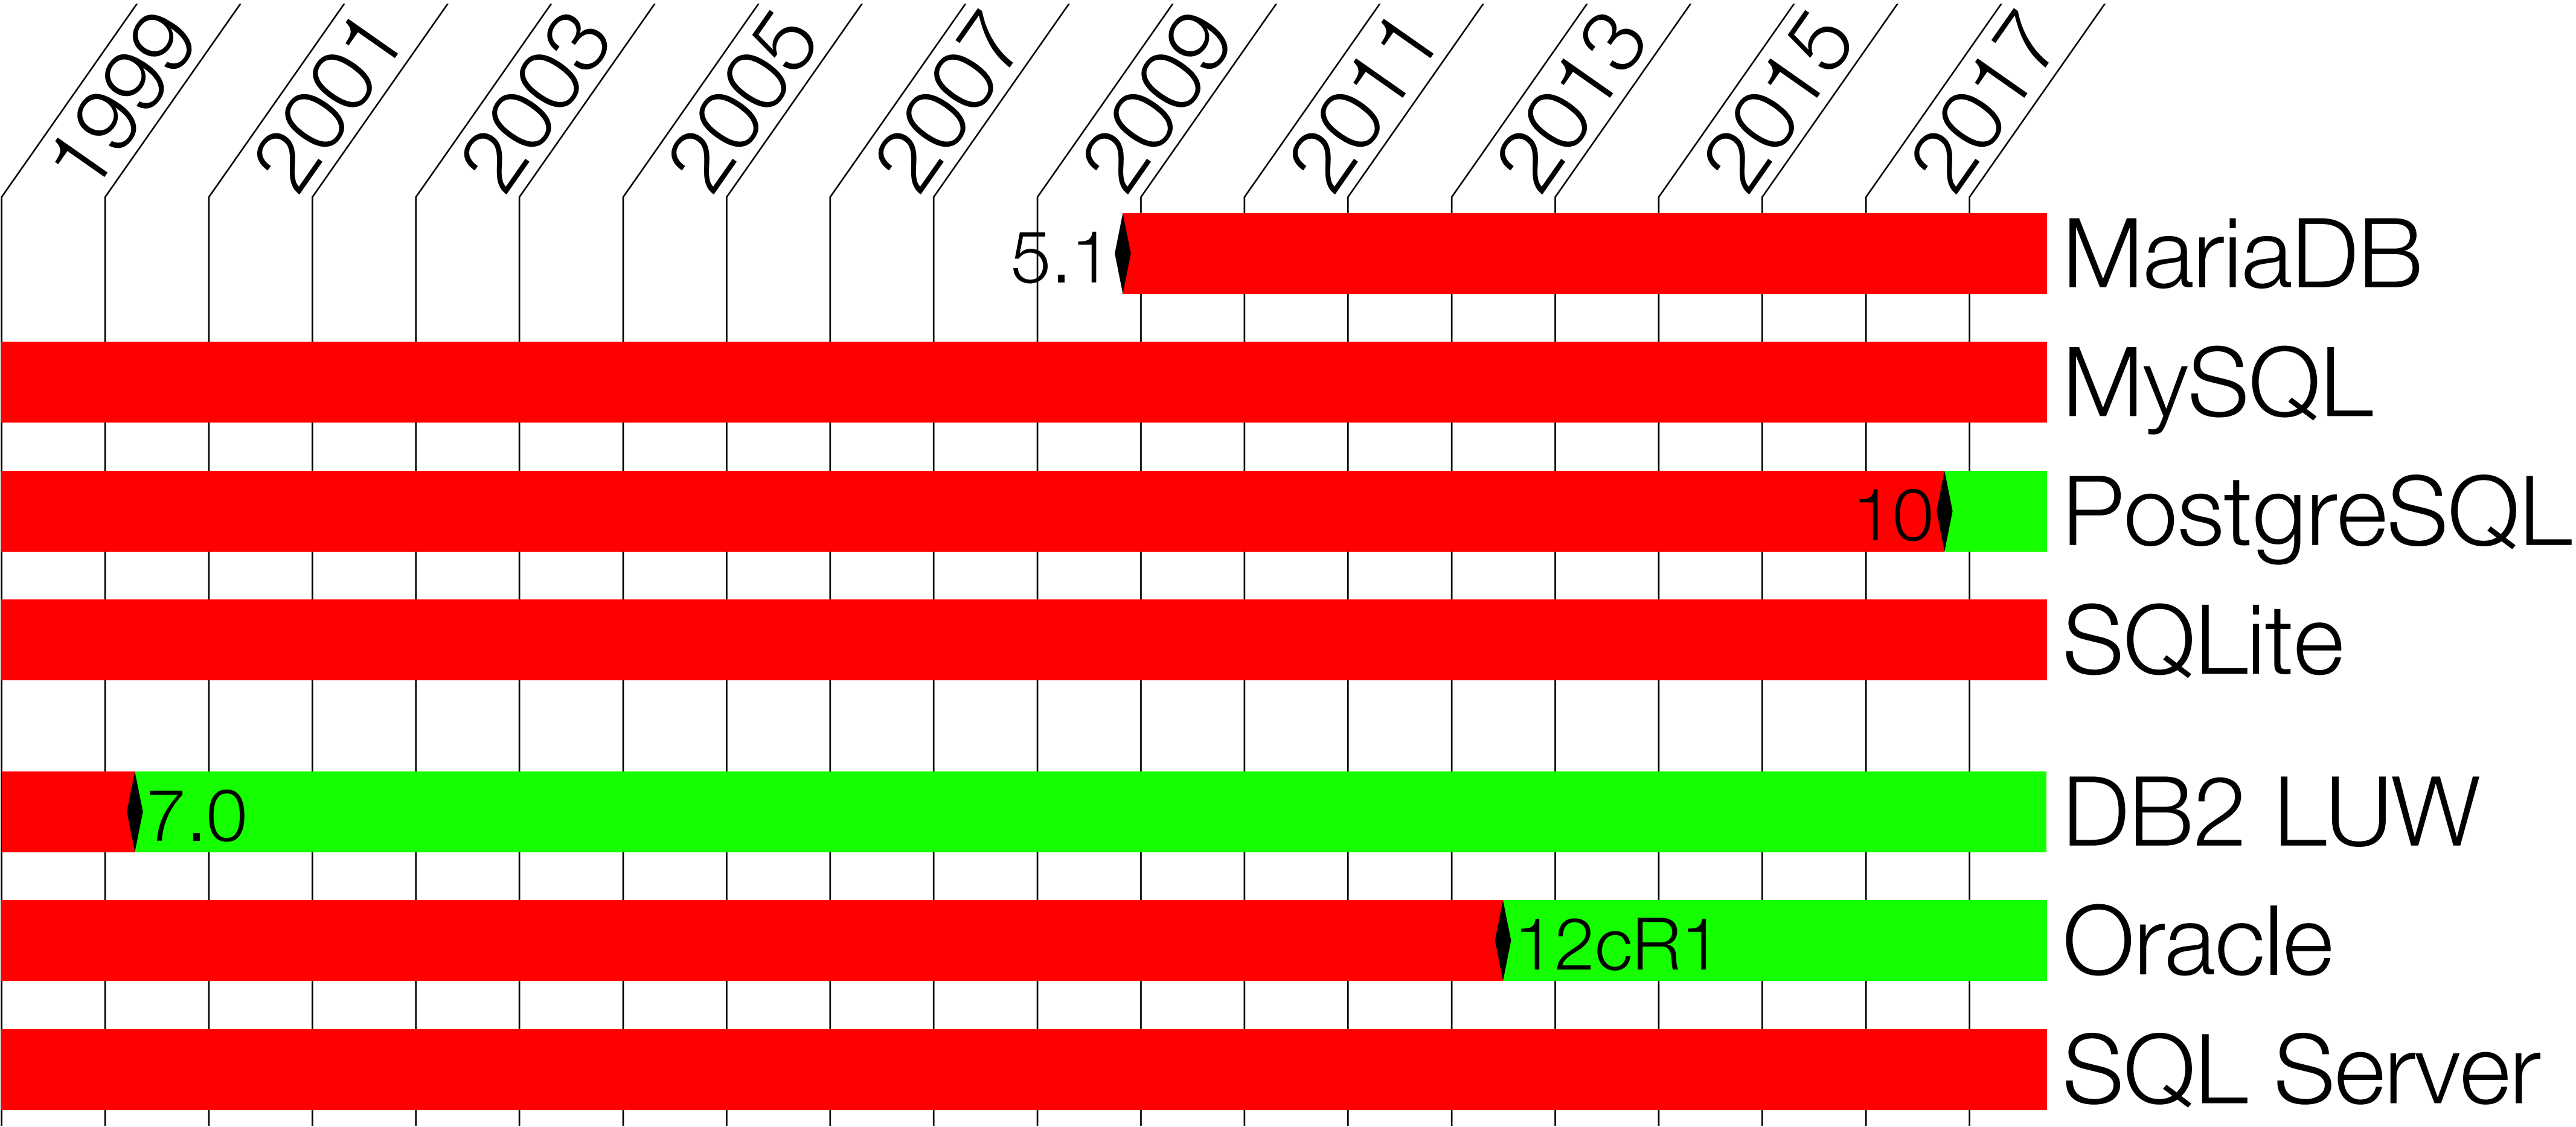
More (esp. vs. serial):

<https://blog.2ndquadrant.com/postgresql-10-identity-columns/>



# IDENTITY column

Availability



# PostgreSQL 11

201?-??-??



OVER

and

PARTITION BY



*Not new, but  
bear with me*

# OVER (PARTITION BY)

## The Problem

---

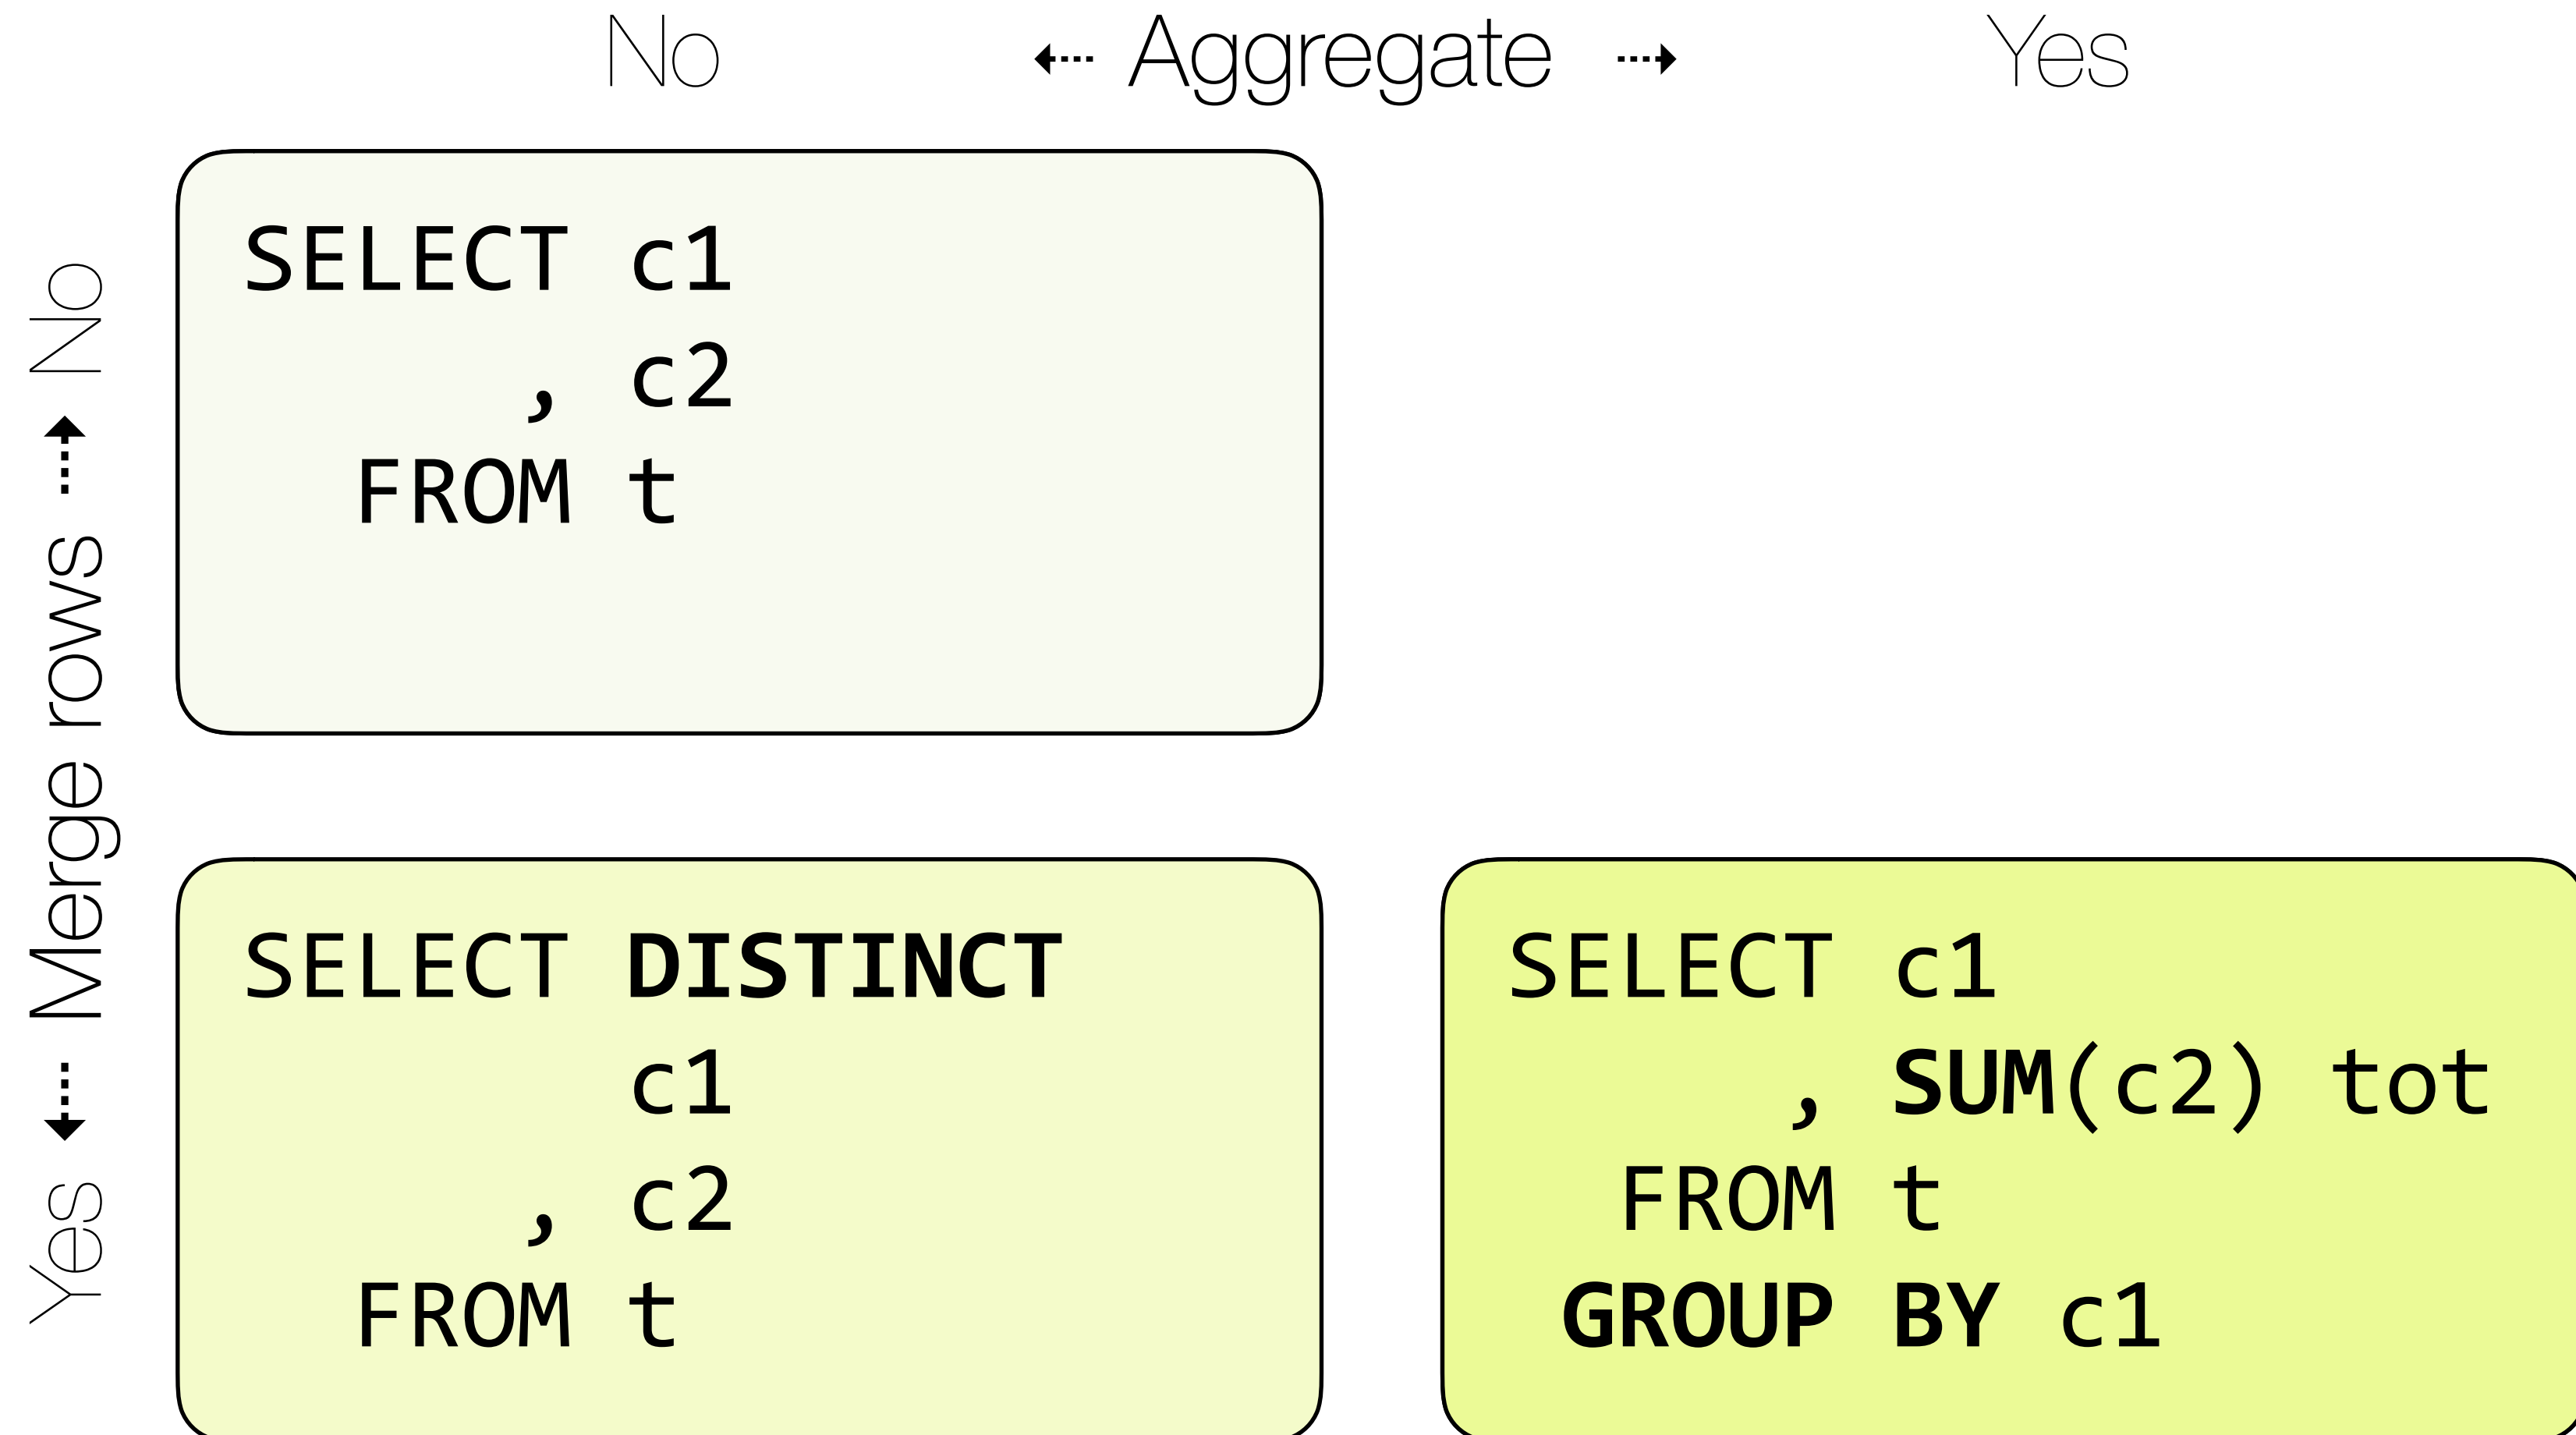
Two distinct concepts could not be used independently:

- ▶ Merge rows with the same key properties
  - ▶ **GROUP BY** to specify key properties
  - ▶ **DISTINCT** to use full row as key
- ▶ Aggregate data from related rows
  - ▶ Requires **GROUP BY** to segregate the rows
  - ▶ **COUNT, SUM, AVG, MIN, MAX** to aggregate grouped rows



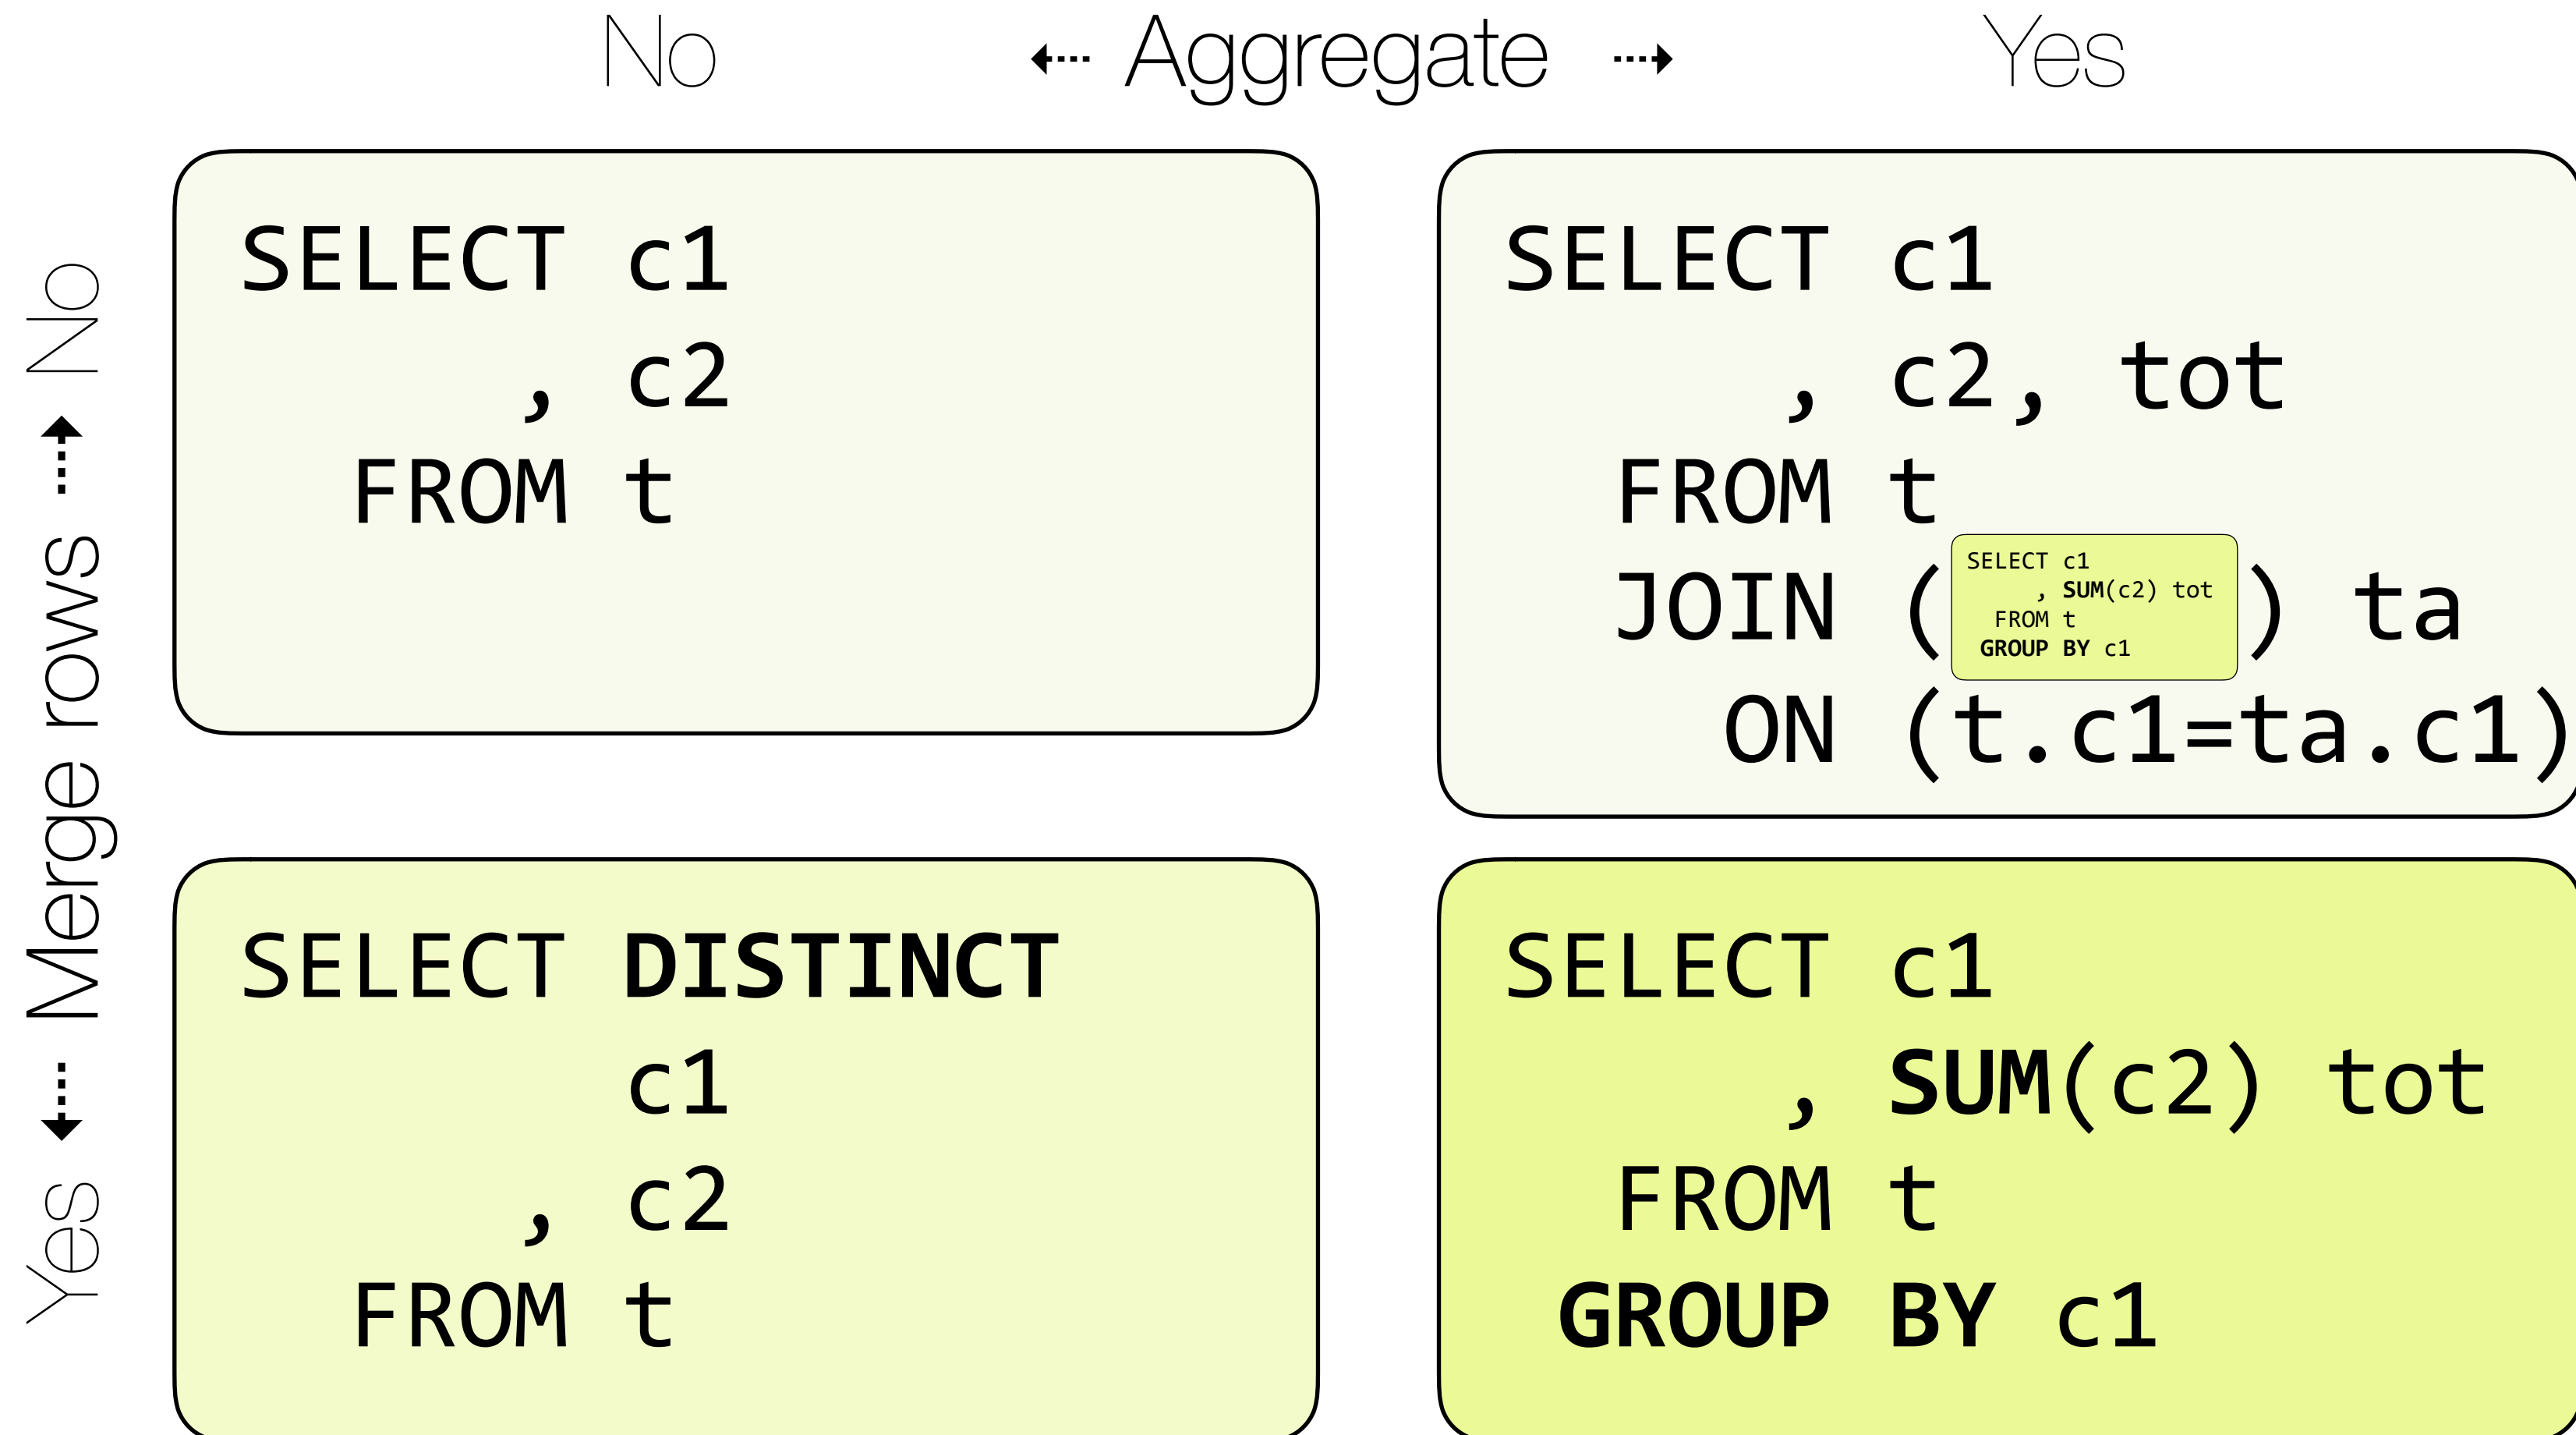
# OVER (PARTITION BY)

## The Problem



# OVER (PARTITION BY)

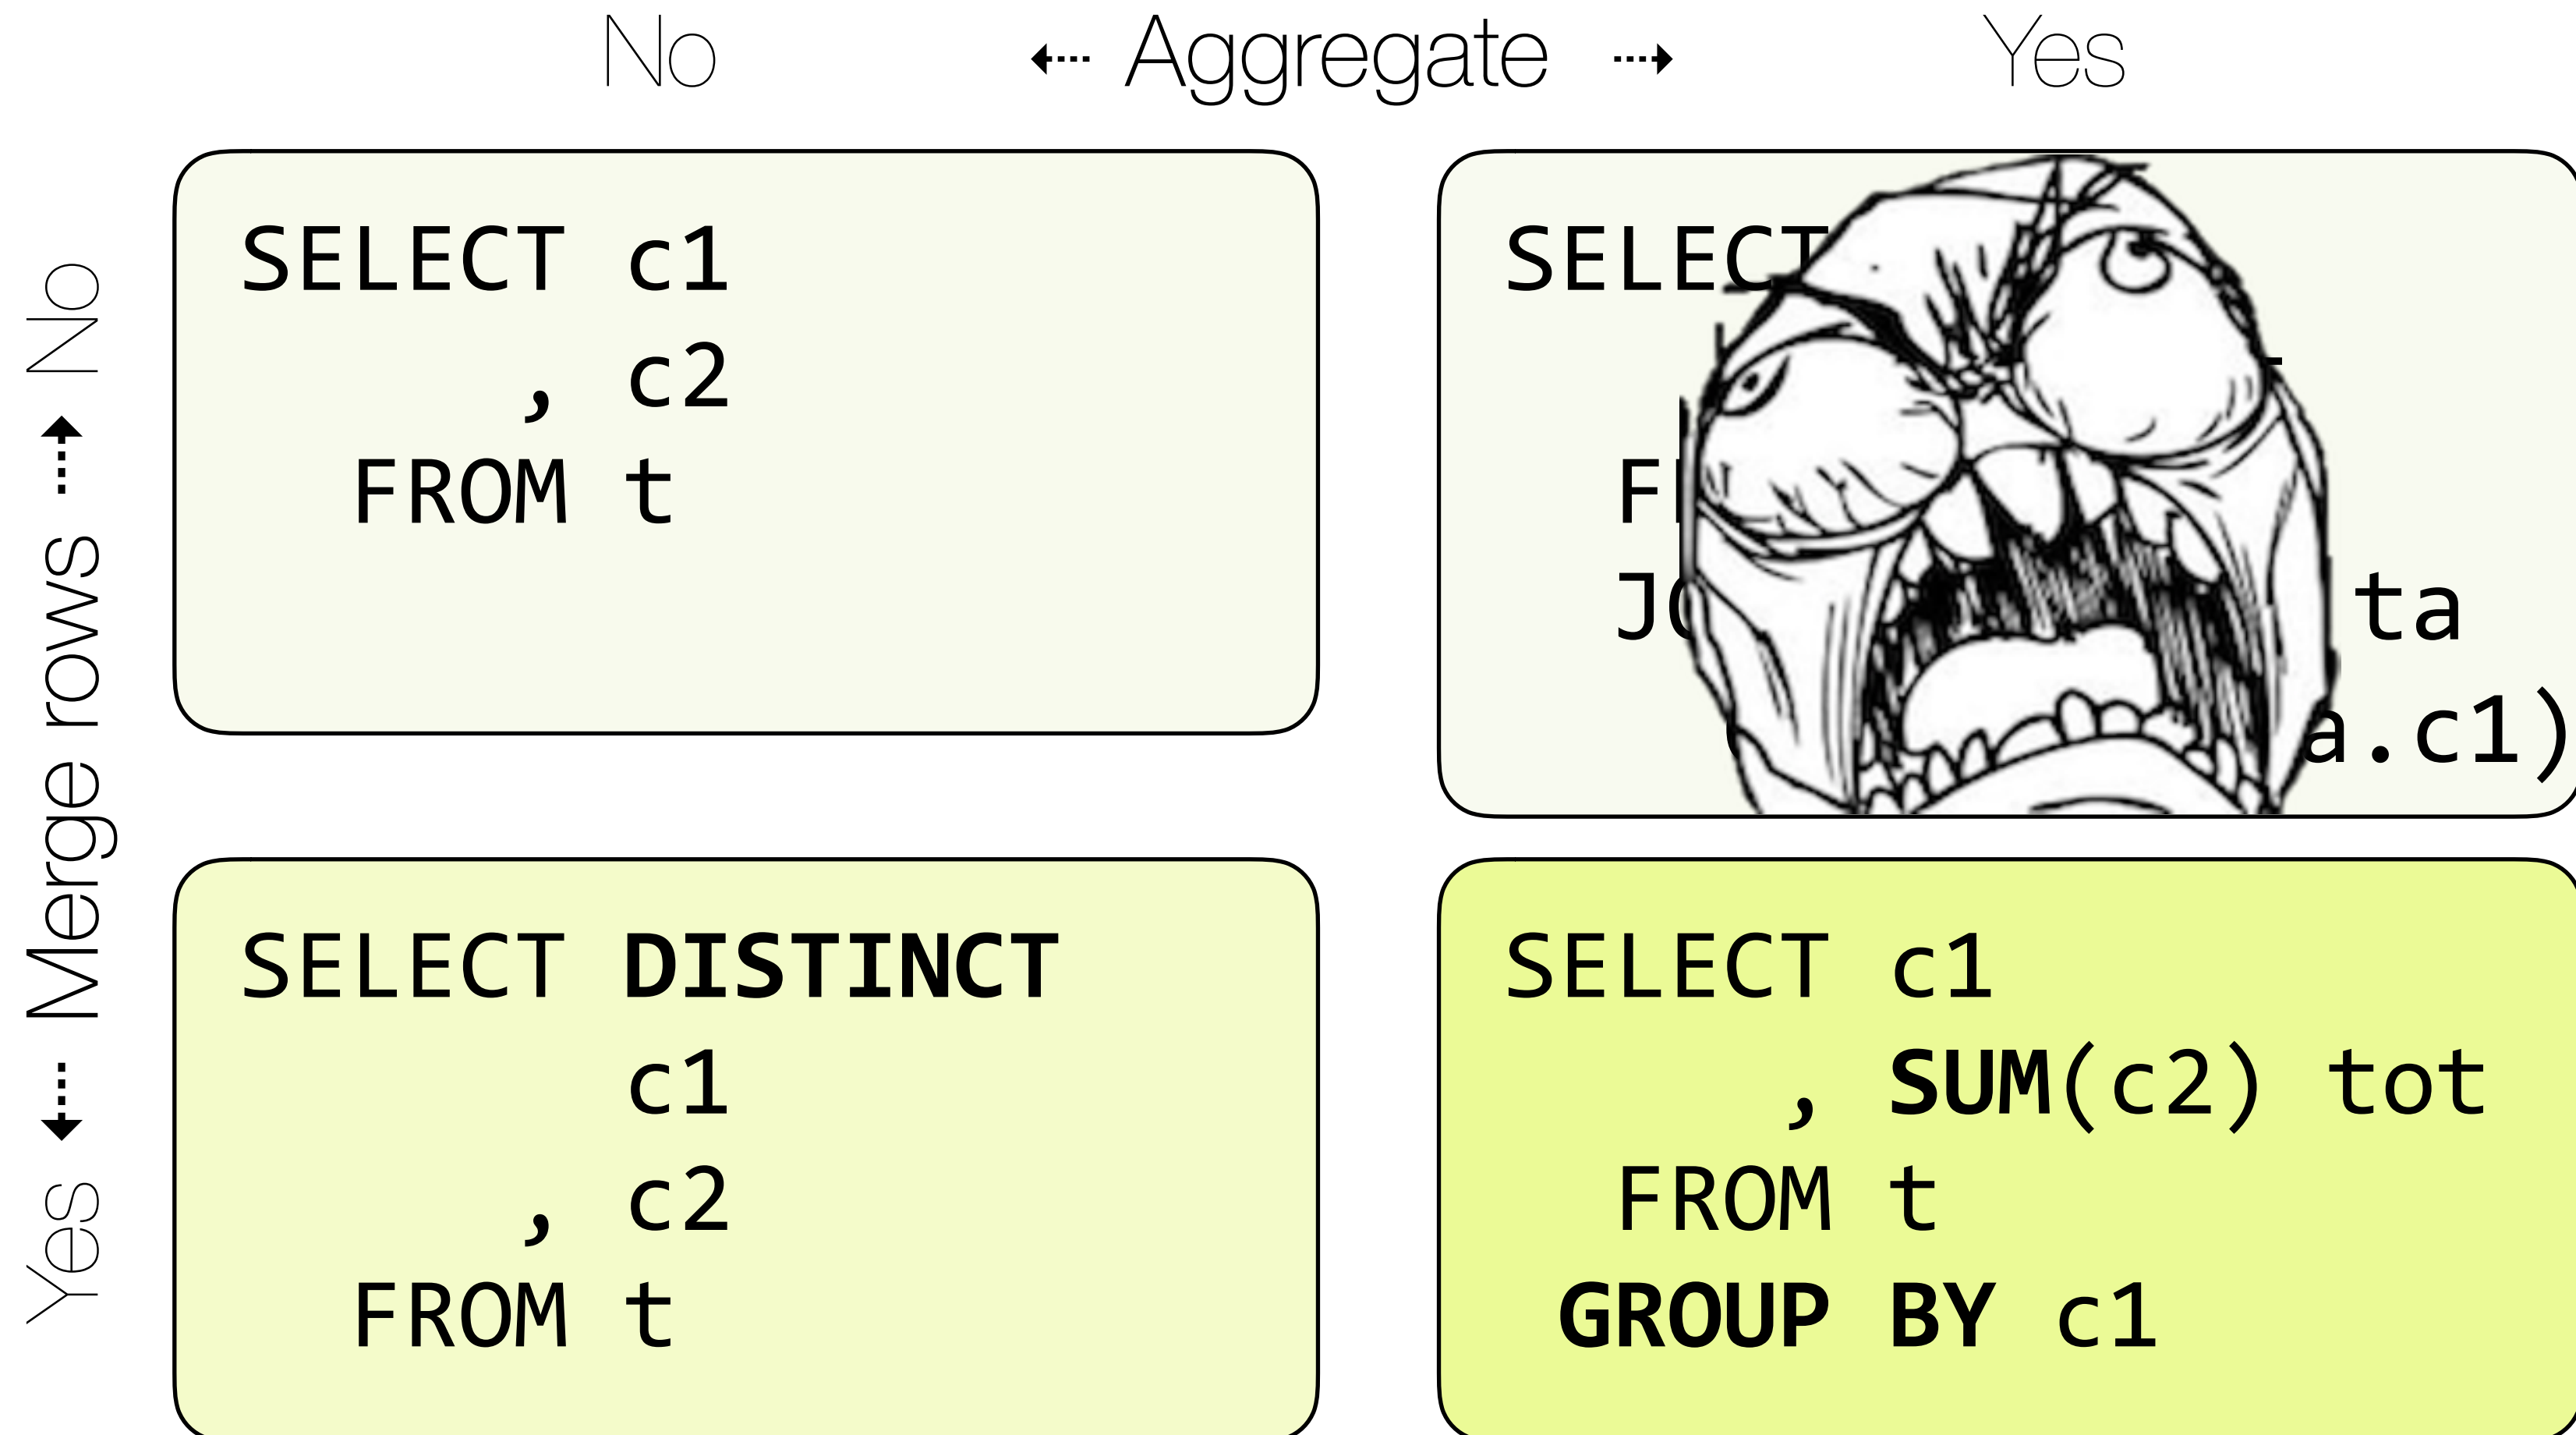
## The Problem





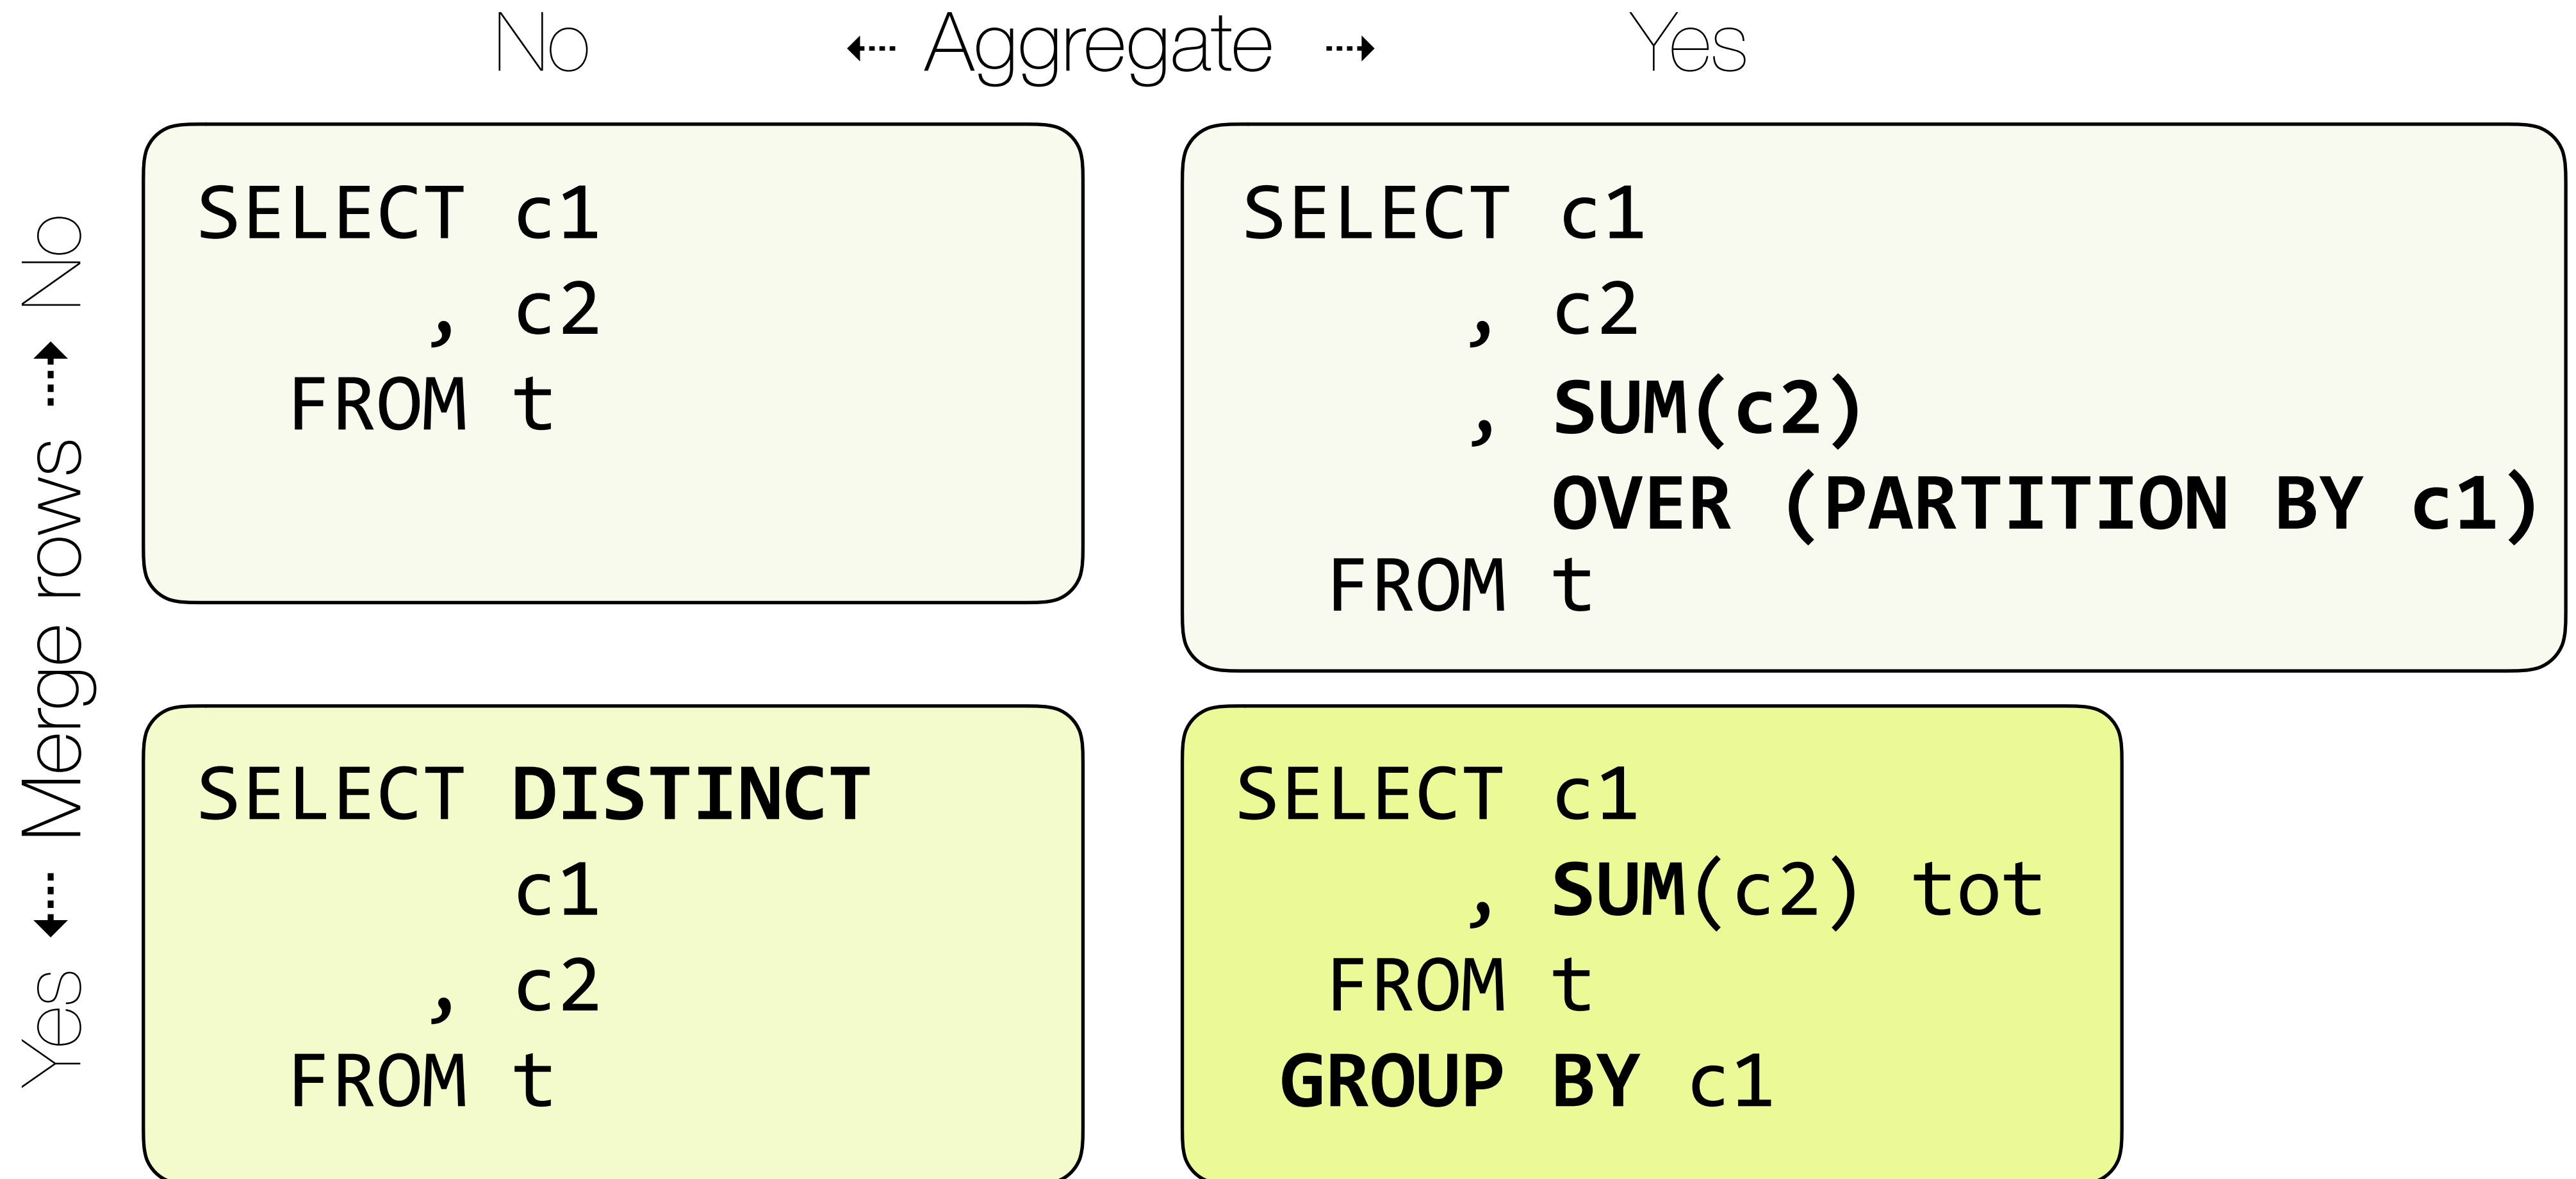
# OVER (PARTITION BY)

## The Problem



# OVER (PARTITION BY)

Since SQL:2003

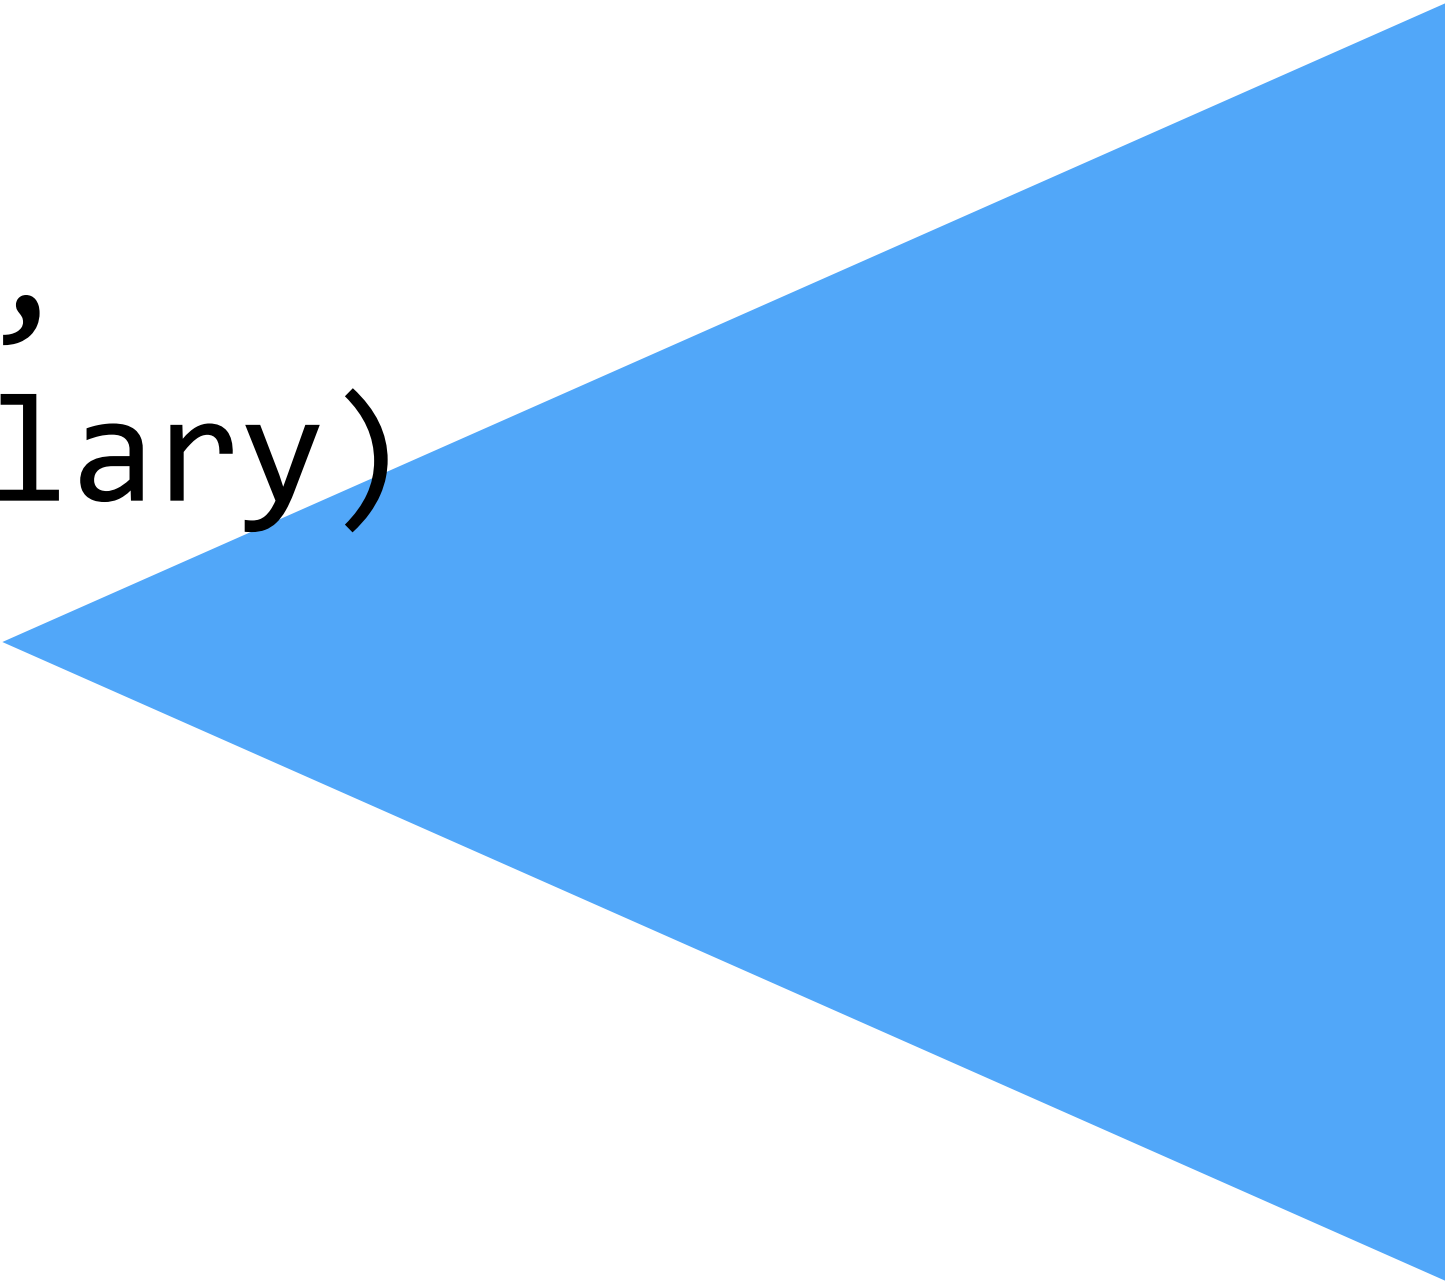




# OVER (PARTITION BY)

How it works

```
SELECT dep,  
       salary,  
       SUM(salary)  
       OVER()  
FROM emp
```

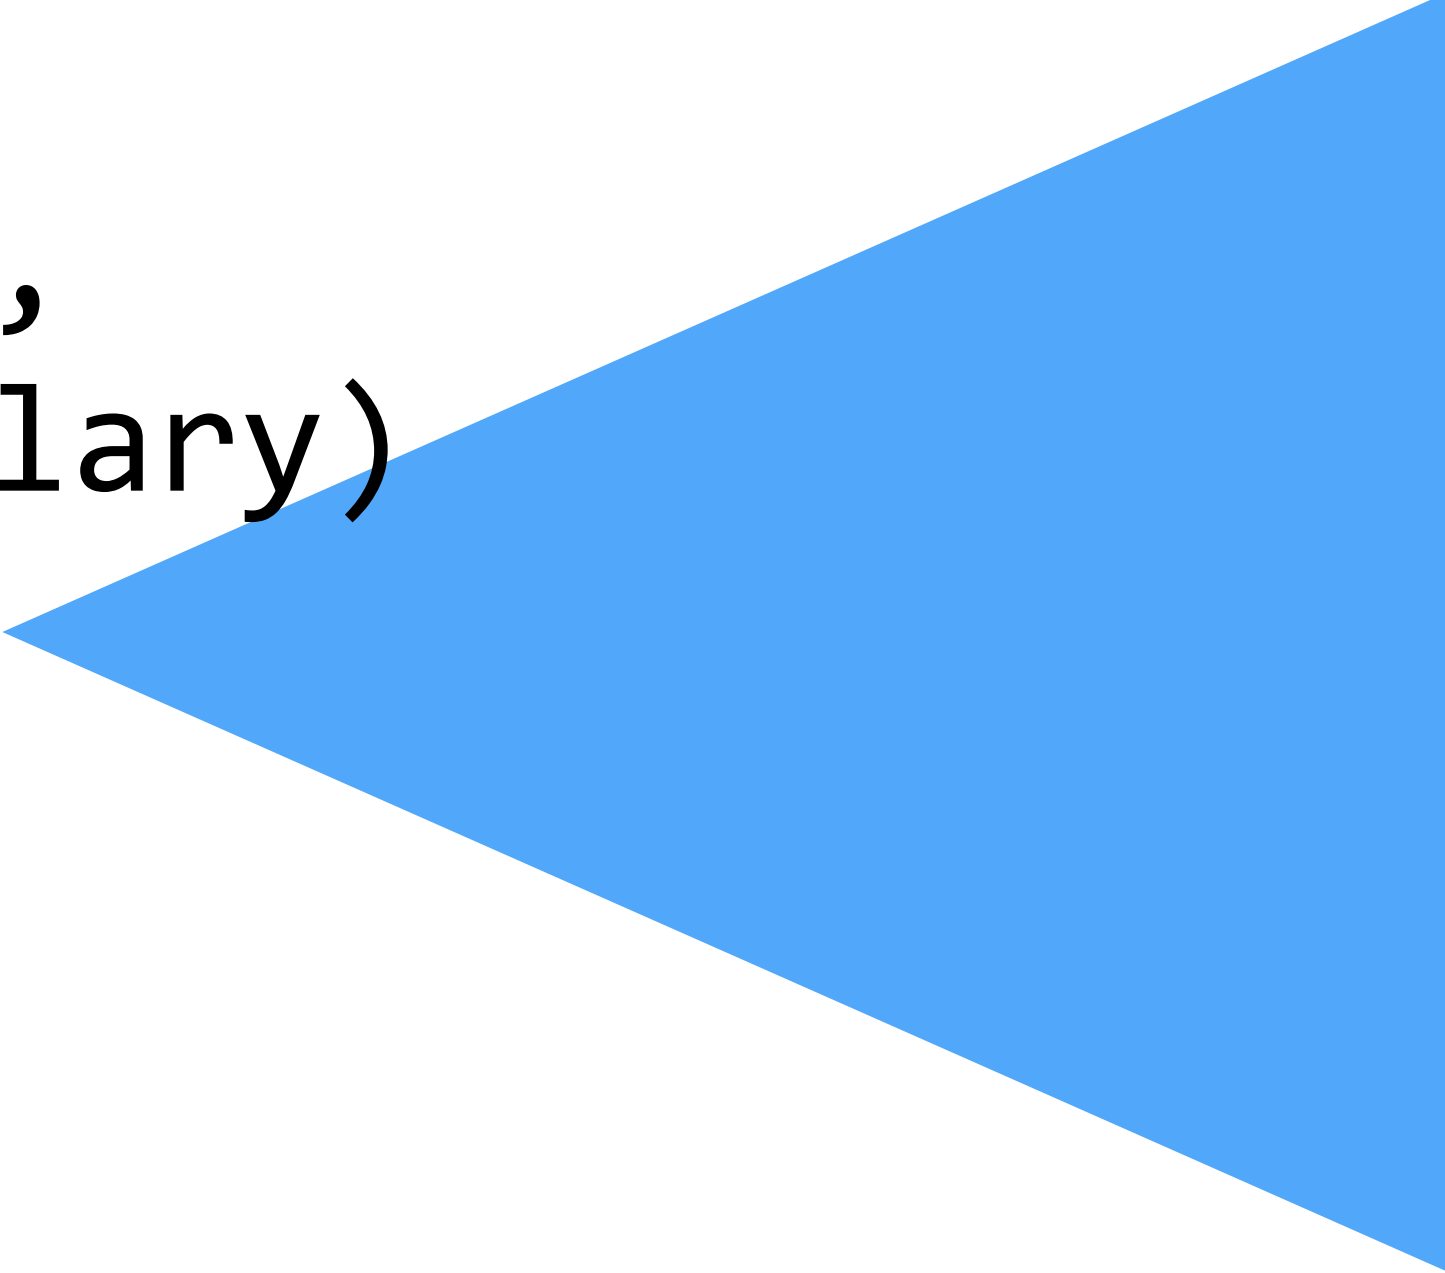


dep	salary
1	1000
22	1000
22	1000
333	1000
333	1000
333	1000

# OVER (PARTITION BY)

How it works

```
SELECT dep,  
       salary,  
       SUM(salary)  
       OVER()  
FROM emp
```



dep	salary	ts
1	1000	6000
22	1000	6000
22	1000	6000
333	1000	6000
333	1000	6000
333	1000	6000



# OVER (PARTITION BY)

How it works

```
SELECT dep,  
       salary,  
       SUM(salary)  
       OVER(PARTITION BY dep)  
FROM emp
```

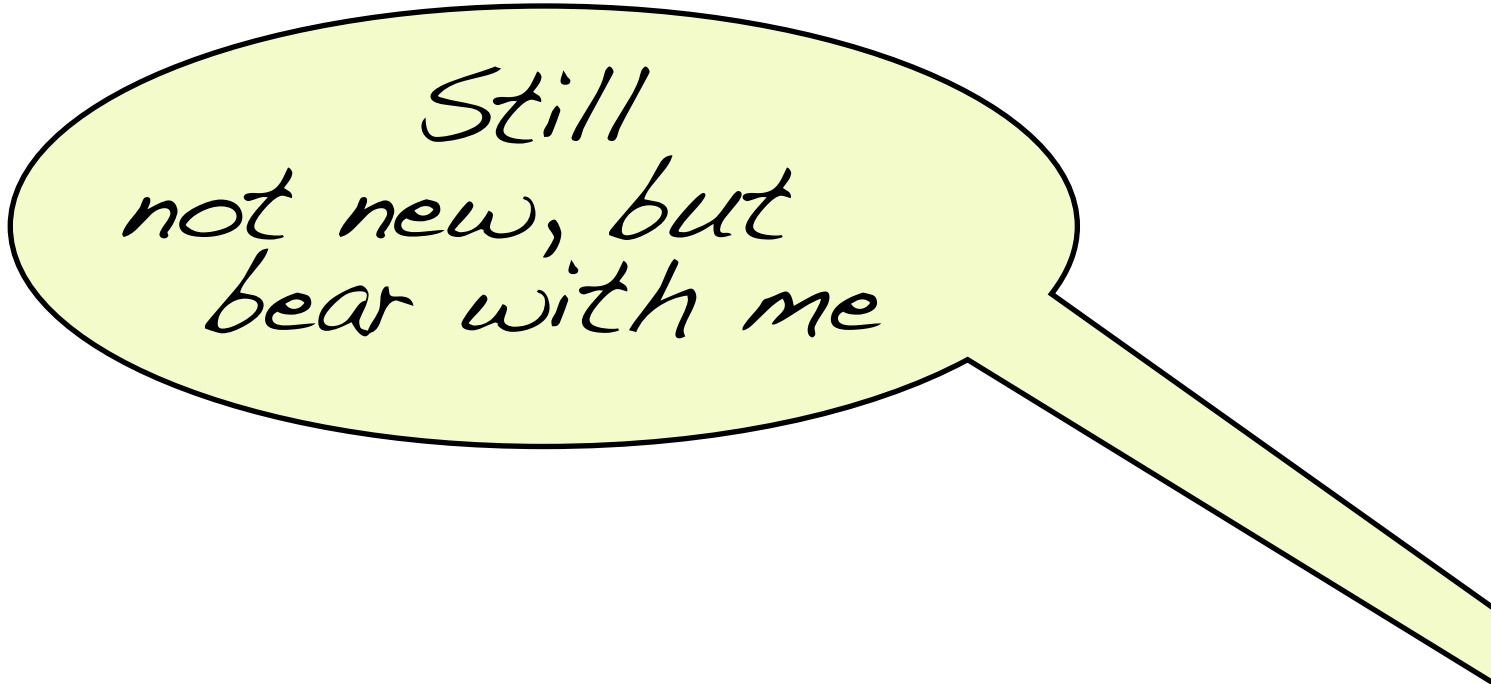
dep	salary	ts
1	1000	1000
22	1000	2000
22	1000	2000
333	1000	3000
333	1000	3000
333	1000	3000

# OVER

and

# ORDER BY

(Framing & Ranking)



*Still  
not new, but  
bear with me*

# OVER (ORDER BY)

## The Problem

```
SELECT id,  
       value,  
       (SELECT SUM(value)  
        FROM transactions t2  
        WHERE t2.id <= t.id)  
FROM transactions t
```

id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20



# OVER (ORDER BY)

## The Problem

```
SELECT id,  
       value,  
       (SELECT SUM(value)  
        FROM transactions t2  
        WHERE t2.id <= t.id)  
FROM transactions t
```


*Range segregation (<=)  
not possible with  
GROUP BY or  
PARTITION BY*

id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
       )  
FROM transactions t
```

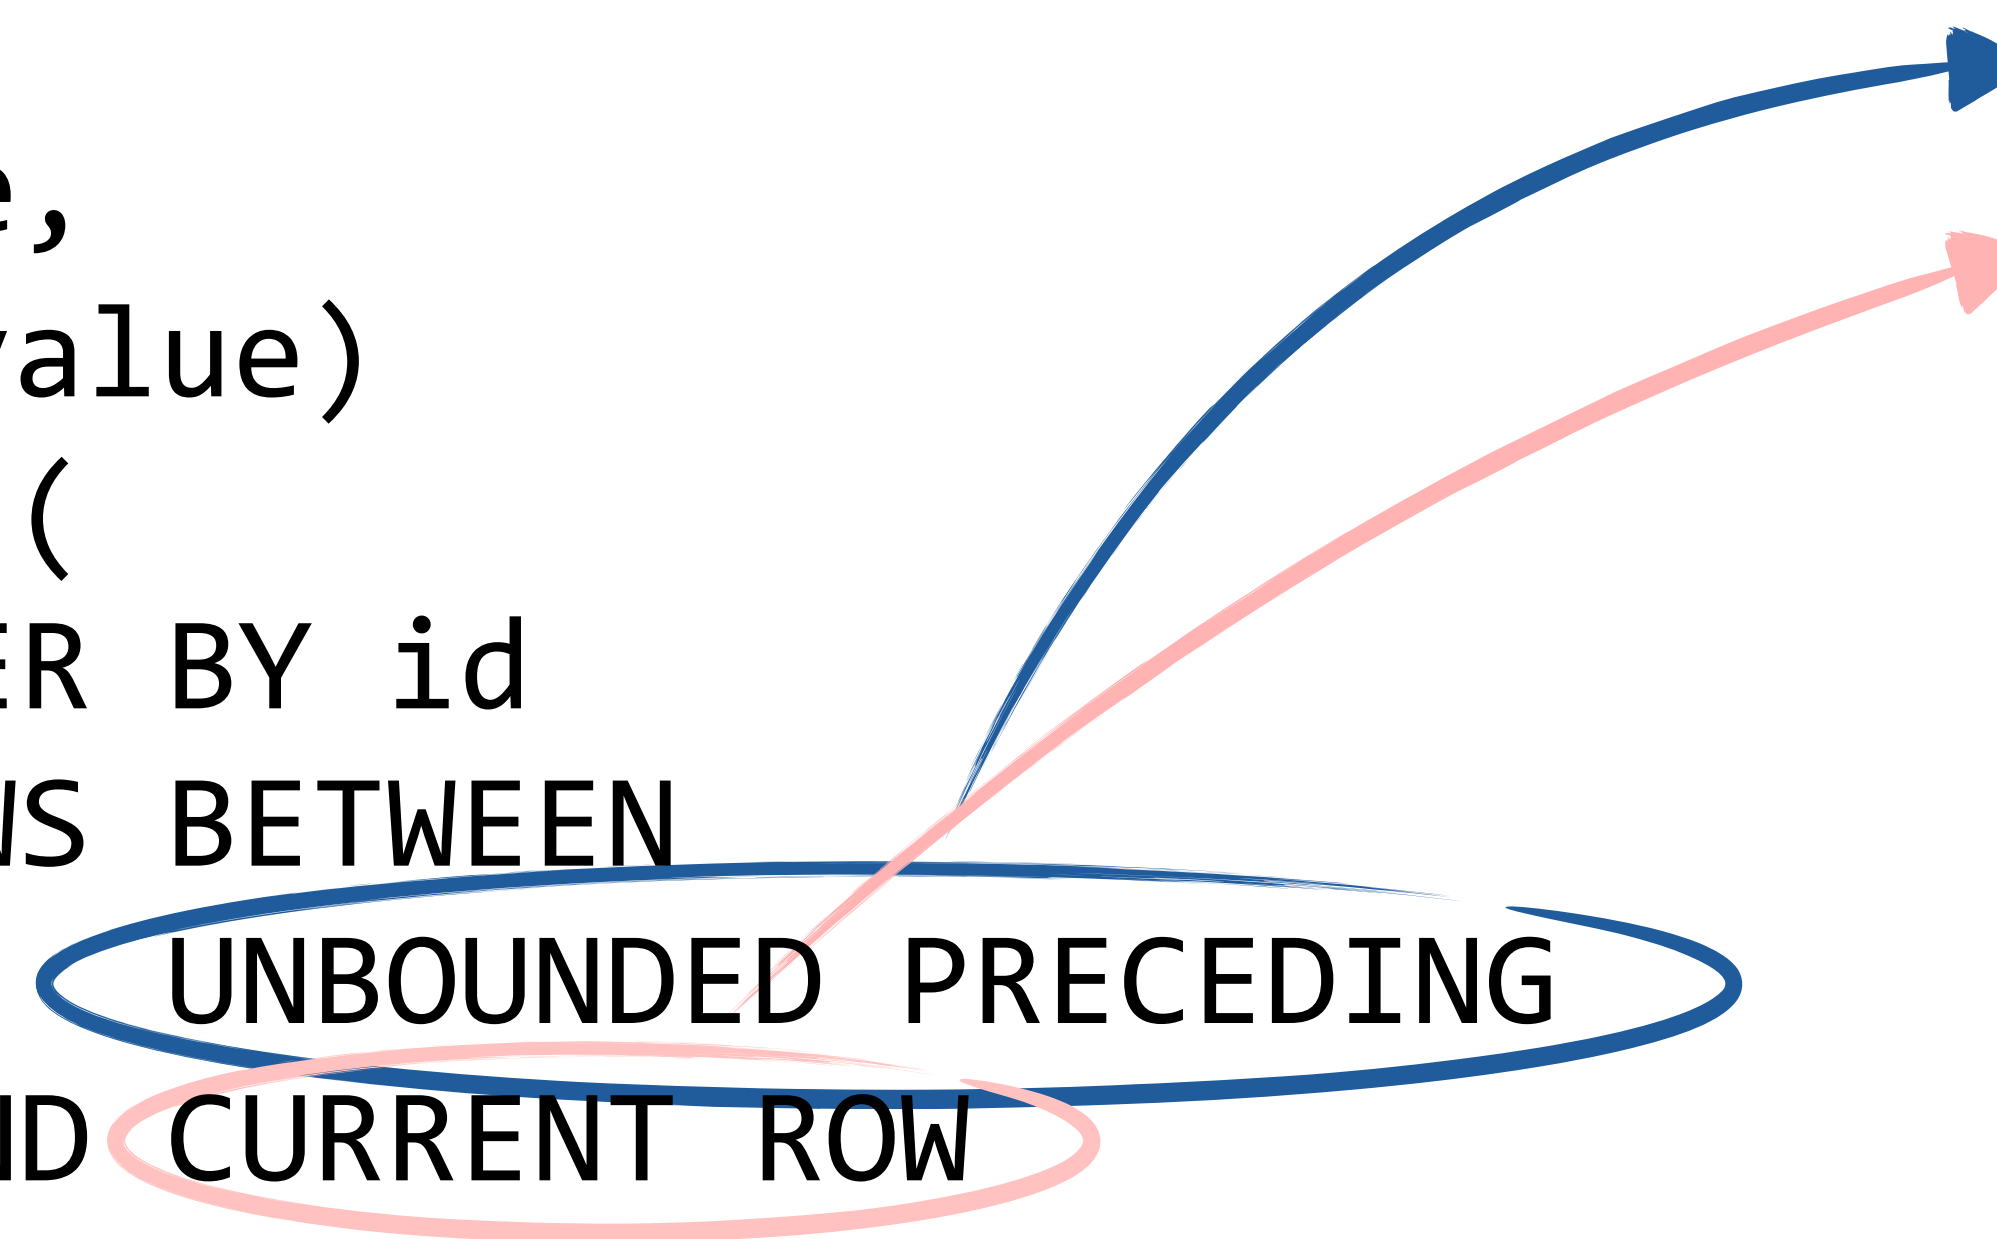


id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```



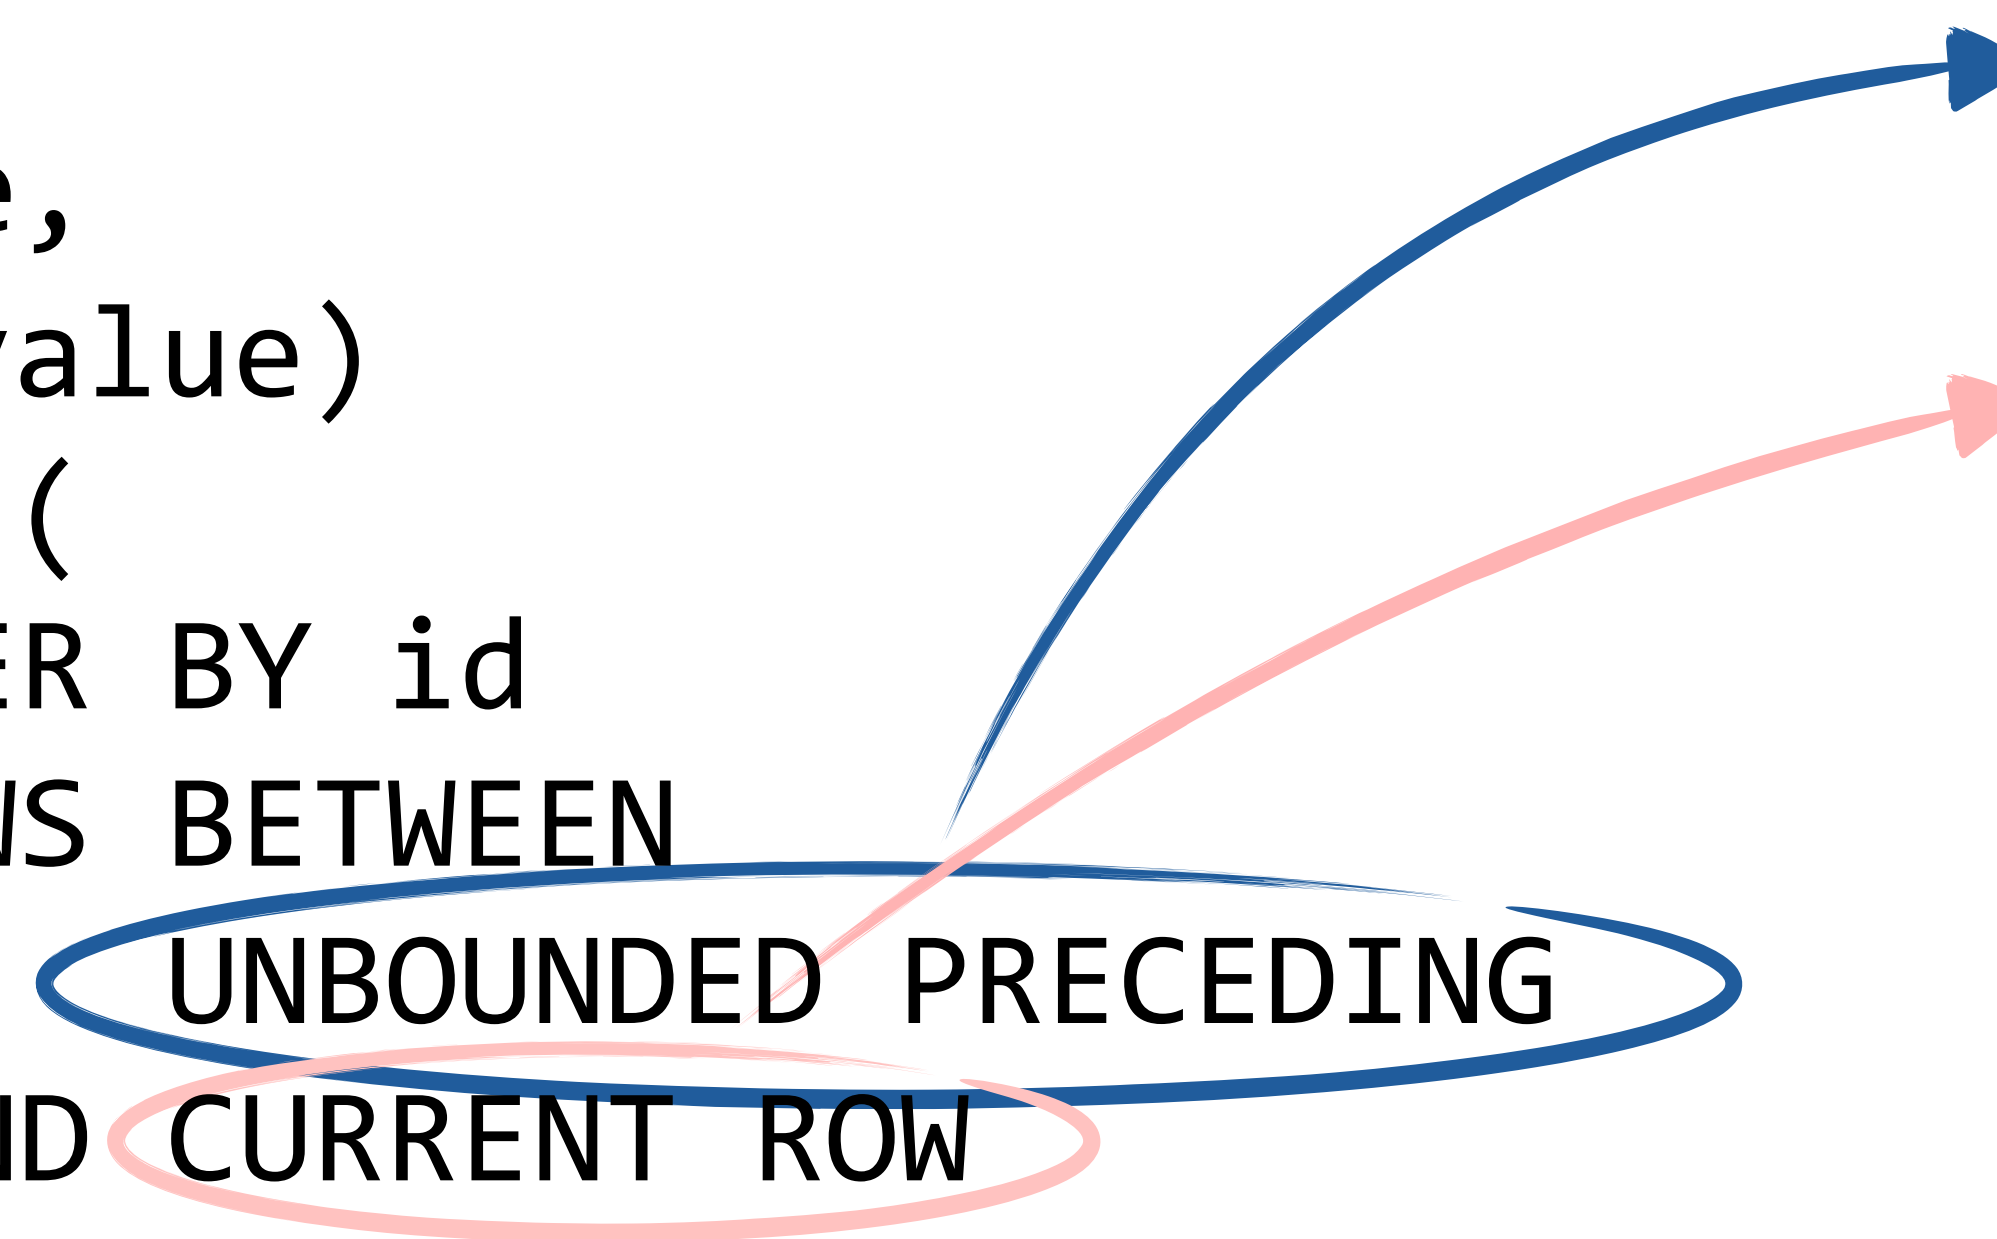
id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20



# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```



id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

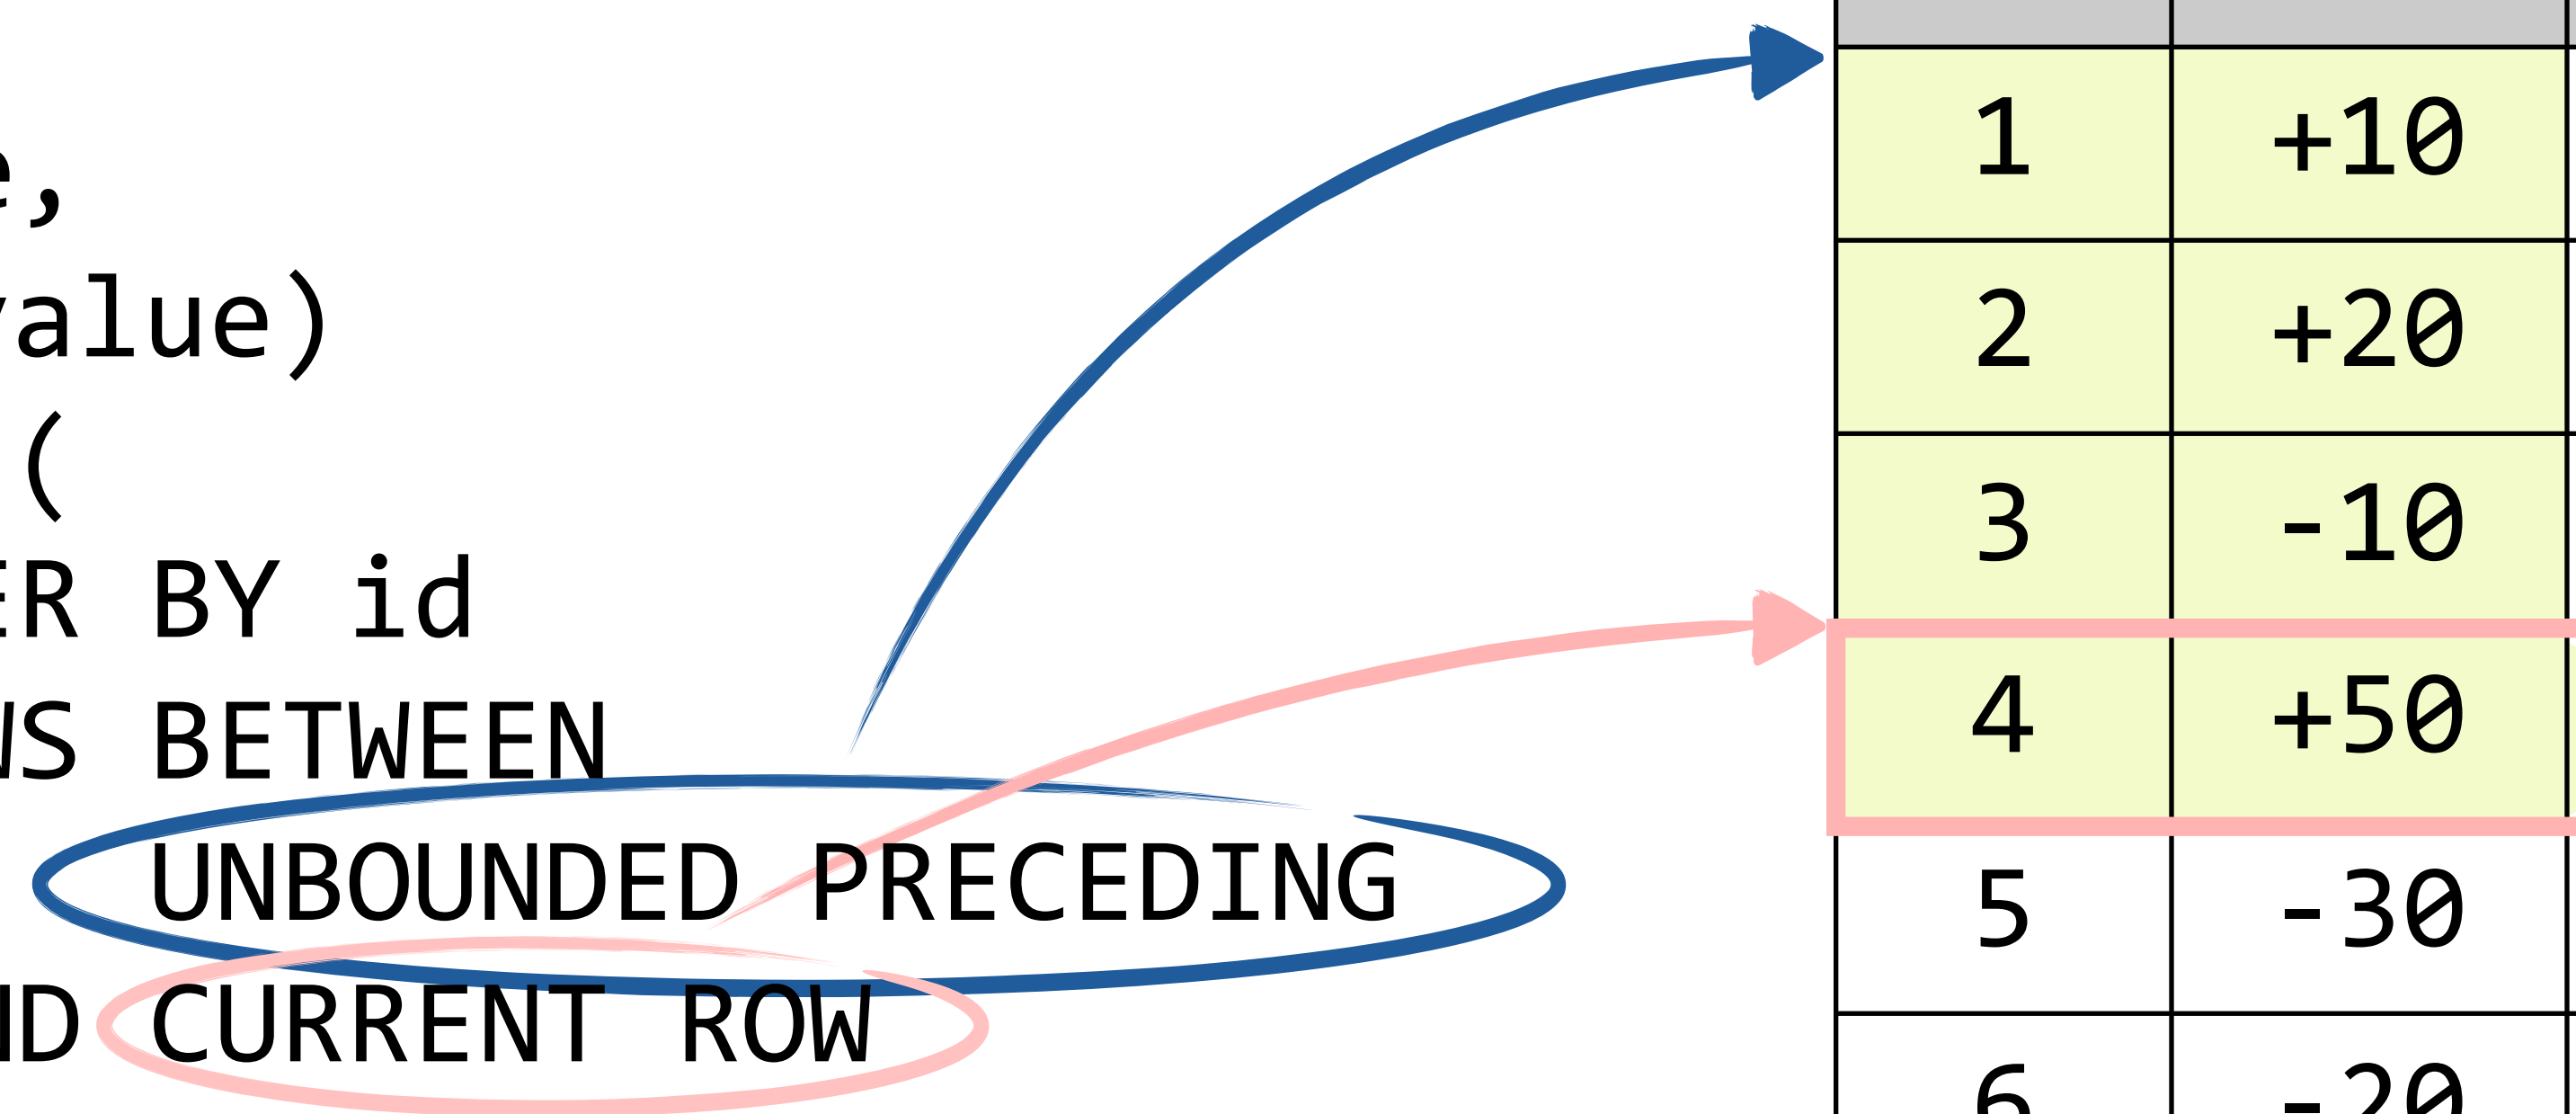
```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```

id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```



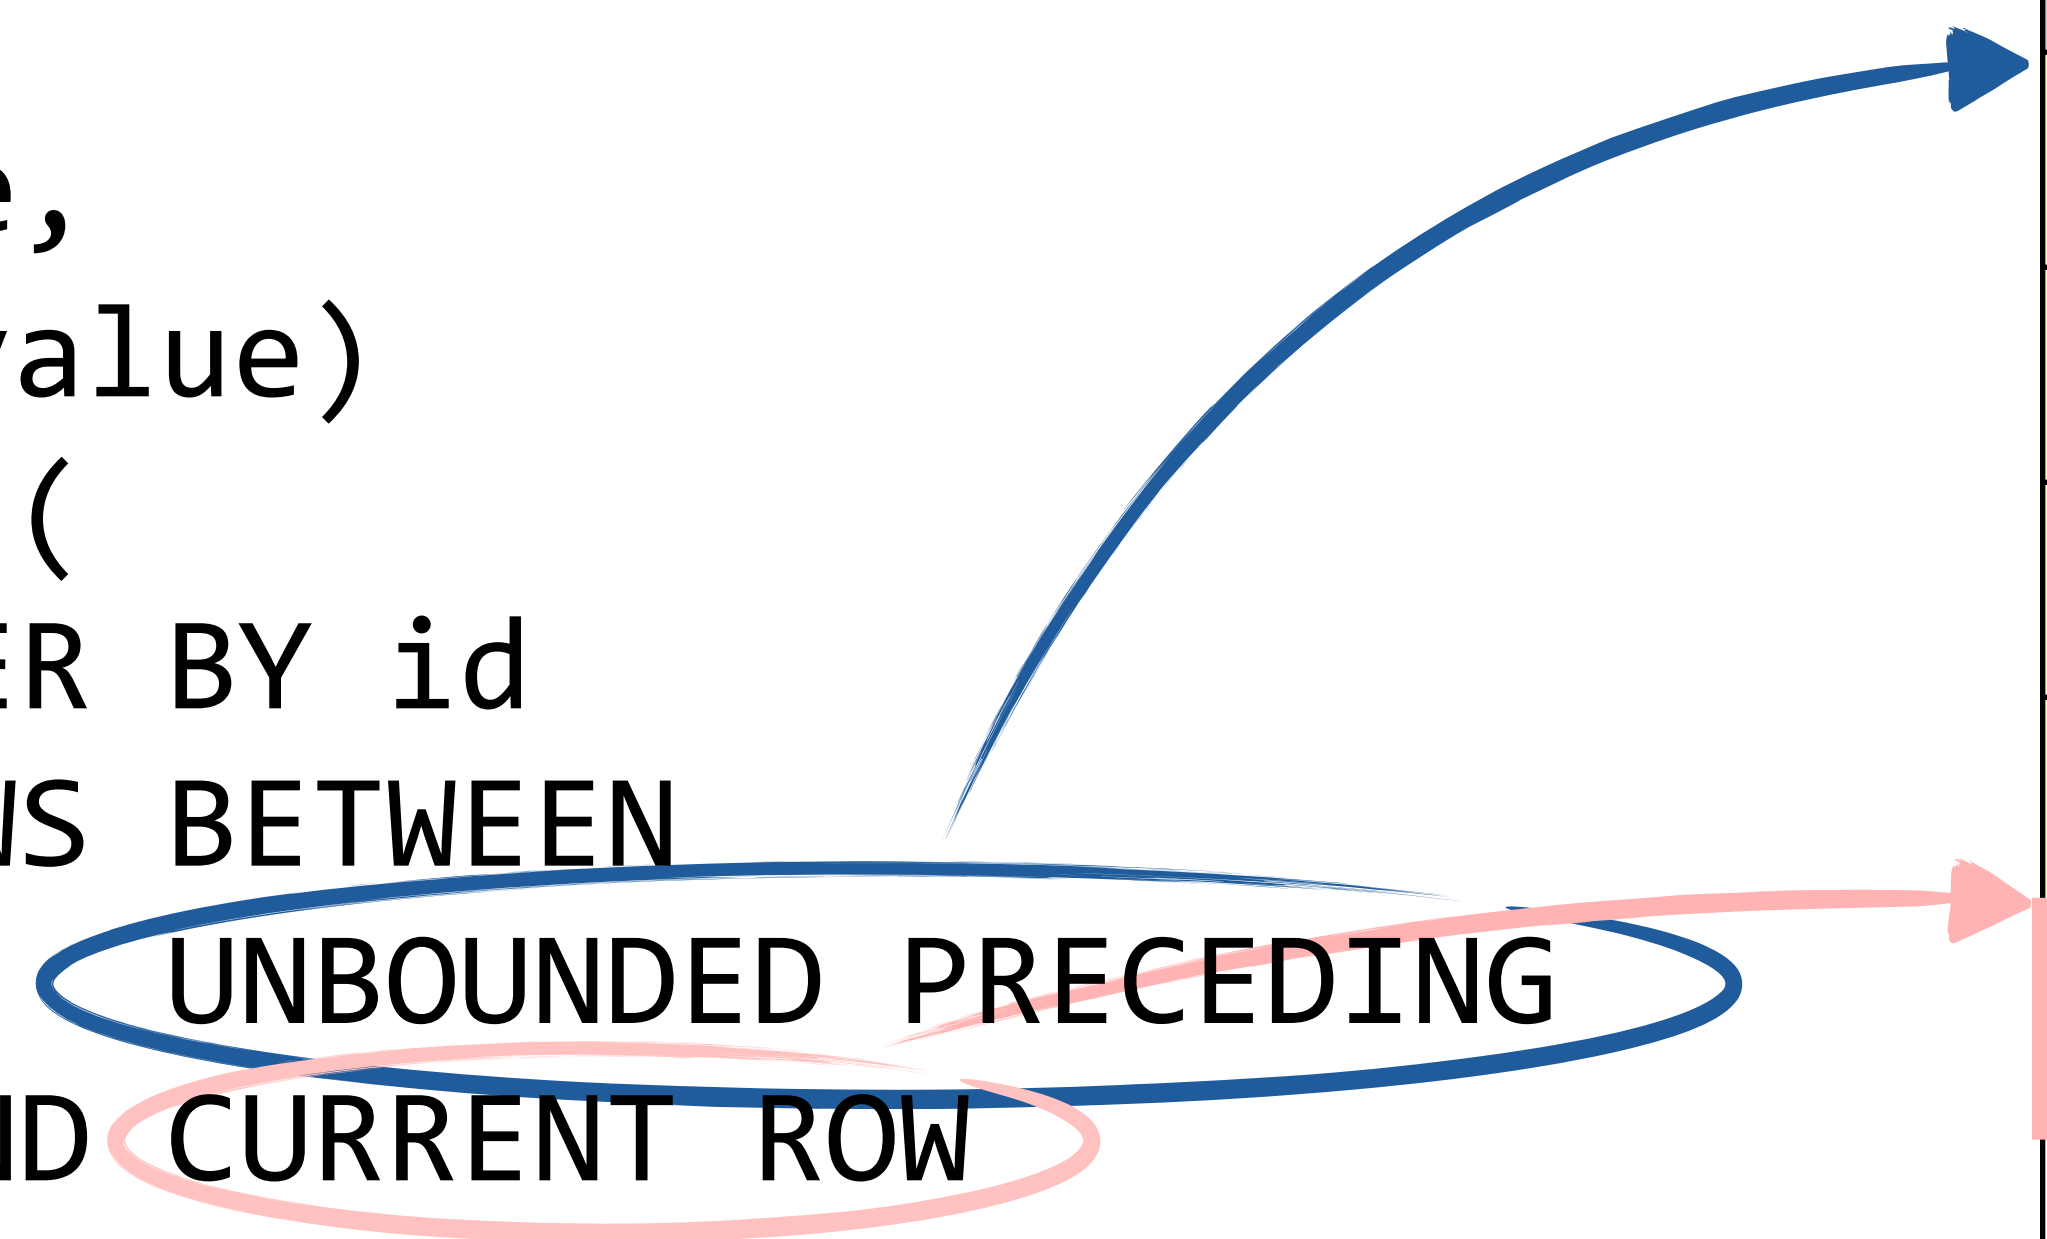
id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20



# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```



id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```

id	value	balance
1	+10	+10
2	+20	+30
3	-10	+20
4	+50	+70
5	-30	+40
6	-20	+20

# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (  
         ORDER BY id  
         ROWS BETWEEN  
           UNBOUNDED PRECEDING  
           AND CURRENT ROW  
       )  
FROM transactions t
```

acct	id	value	balance
1	1	+10	
22	2	+20	
22	3	-10	
333	4	+50	
333	5	-30	
333	6	-20	



# OVER (ORDER BY)

Since SQL:2003

```
SELECT id,  
       value,  
       SUM(value)  
       OVER (PARTITION BY acct  
             ORDER BY id  
             ROWS BETWEEN  
                   UNBOUNDED PRECEDING  
                   AND CURRENT ROW  
             )  
FROM transactions t
```

acct	id	value	balance
1	1	+10	+10
22	2	+20	+20
22	3	-10	+10
333	4	+50	+50
333	5	-30	+20
333	6	-20	0

# OVER (ORDER BY)

Since SQL:2003

With **OVER (ORDER BY n)** a new type of functions make sense:

n	ROW_NUMBER	RANK	DENSE_RANK	PERCENT_RANK	CUME_DIST
1	1	1	1	0	0.25
2	2	2	2	0.33...	0.75
3	3	2	2	0.33...	0.75
4	4	4	3	1	1

# OVER (SQL:2003)

## Use Cases

- ▶ Aggregates without GROUP BY

- ▶ Running totals, moving averages

```
AVG(...) OVER(ORDER BY ...  
              ROWS BETWEEN 3 PRECEDING  
              AND 3 FOLLOWING) moving_avg
```

- ▶ Ranking

- ▶ Top-N per Group

- ▶ Avoiding self-joins

```
SELECT *  
  FROM (SELECT ROW_NUMBER()  
         OVER(PARTITION BY ... ORDER BY ...) rn  
        , t.*  
        FROM t) numbered_t  
 WHERE rn <= 3
```

[... many more ...]



# **OVER** (SQL:2003)

In a Nutshell

---

**OVER** may follow any aggregate function

**OVER** defines which rows are visible at each row

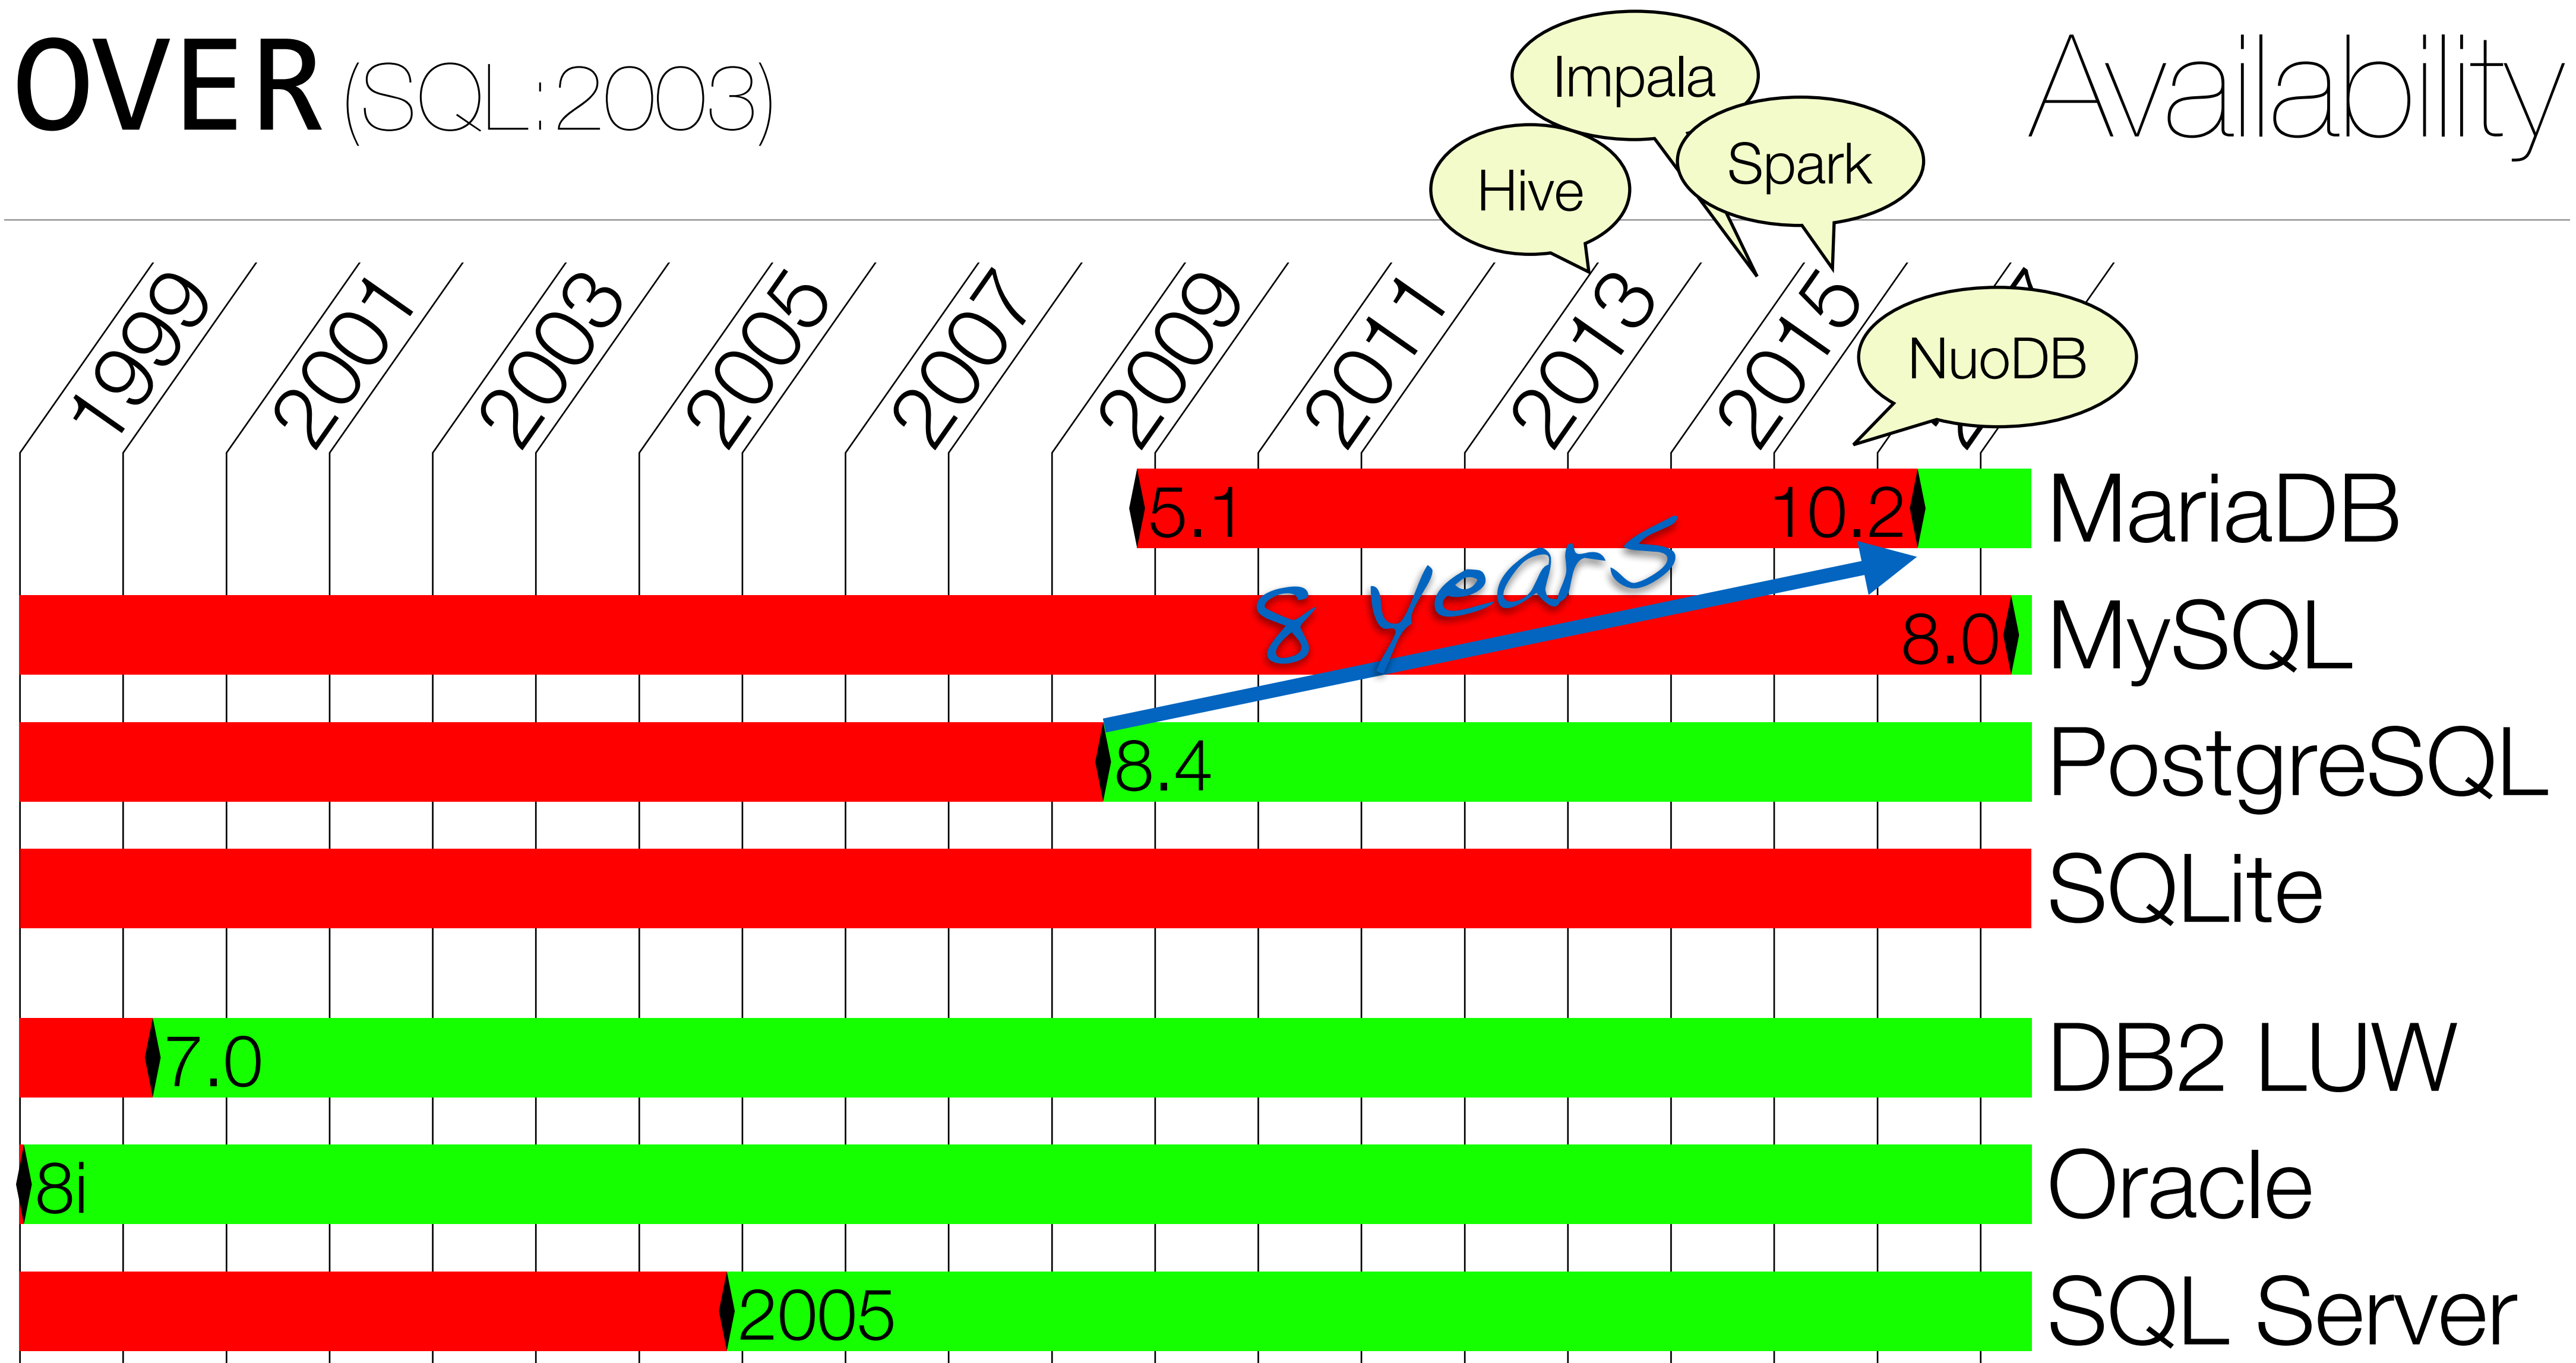
**OVER()** makes all rows visible at every row

**OVER(PARTITION BY ...)** segregates like **GROUP BY**

**OVER(ORDER BY ... BETWEEN)** segregates using **<, >**

# OVER (SQL:2003)

Availability



# OVER

SQL:2011 functions

# OVER (SQL:2011)

## The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
SELECT *  
FROM t
```

balance	...
50	...
90	...
70	...
30	...



# OVER (SQL:2011)

## The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
SELECT *  
      , ROW_NUMBER() OVER(ORDER BY x) rn  
FROM t
```

balance	...	rn
50	...	1
90	...	2
70	...	3
30	...	4

# OVER (SQL:2011)

## The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
WITH numbered_t AS (SELECT *  
                     , ROW_NUMBER() OVER(ORDER BY x) rn  
                     FROM t)
```

balance	...	rn
50	...	1
90	...	2
70	...	3
30	...	4

# OVER (SQL:2011)

## The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
WITH numbered_t AS (SELECT *  
                     , ROW_NUMBER() OVER(ORDER BY x) rn  
                     FROM t)
```

```
SELECT curr.*
```

```
FROM      numbered_t curr
```

curr		
balance	...	rn
50	...	1
90	...	2
70	...	3
30	...	4

# OVER (SQL:2011)

# The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
WITH numbered_t AS (SELECT *  
                    , ROW_NUMBER() OVER(ORDER BY x) rn  
                    FROM t)  
  
SELECT curr.*  
  
FROM      numbered_t curr  
LEFT JOIN numbered_t prev  
ON (
```

curr			prev		
balance	...	rn	balance	...	rn
50	...	1	50	...	1
90	...	2	90	...	2
70	...	3	70	...	3
30	...	4	30	...	4



# OVER (SQL:2011)

## The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
WITH numbered_t AS (SELECT *  
                     , ROW_NUMBER() OVER(ORDER BY x) rn  
                     FROM t)
```

```
SELECT curr.*
```

```
FROM      numbered_t curr  
LEFT JOIN numbered_t prev  
  ON (curr.rn = prev.rn+1)
```

curr			prev		
balance	...	rn	balance	...	rn
50	...	1			
90	...	2	50	...	1
70	...	3	90	...	2
30	...	4	70	...	3

# OVER (SQL:2011)

# The Problem

Direct access of other rows of the same window is not possible.  
(E.g., calculate the difference to the previous rows)

```
WITH numbered_t AS (SELECT *  
                    , ROW_NUMBER() OVER(ORDER BY x) rn  
                    FROM t)  
  
SELECT curr.*  
      , curr.balance  
      - COALESCE(prev.balance,0)  
FROM      numbered_t curr  
LEFT JOIN numbered_t prev  
  ON (curr.rn = prev.rn+1)
```

curr			prev			
balance	...	rn	balance	...	rn	
50	...	1				+50
90	...	2	50	...	1	+40
70	...	3	90	...	2	-20
30	...	4	70	...	3	-40

# OVER Since SQL:2011

---

SQL:2011 can access other rows directly:

```
SELECT *, balance - LAG(balance, 1, 0)
                    OVER(ORDER BY x)
FROM data
```

# OVER Since SQL:2011

---

SQL:2011 can access other rows directly:

```
SELECT *, balance - LAG(balance, 1, 0)
                        OVER(ORDER BY x)
FROM data
```

Available functions:

LEAD / LAG

FIRST\_VALUE / LAST\_VALUE

NTH\_VALUE(col, n) FROM FIRST/LAST  
RESPECT/IGNORE NULLS



# OVER Since SQL:2011

---

SQL:2011 can access other rows directly:

```
SELECT *, balance - LAG(balance, 1, 0)
                OVER(ORDER BY x)
FROM data
```

Available functions:

LEAD / LAG

FIRST\_VALUE / LAST\_VALUE

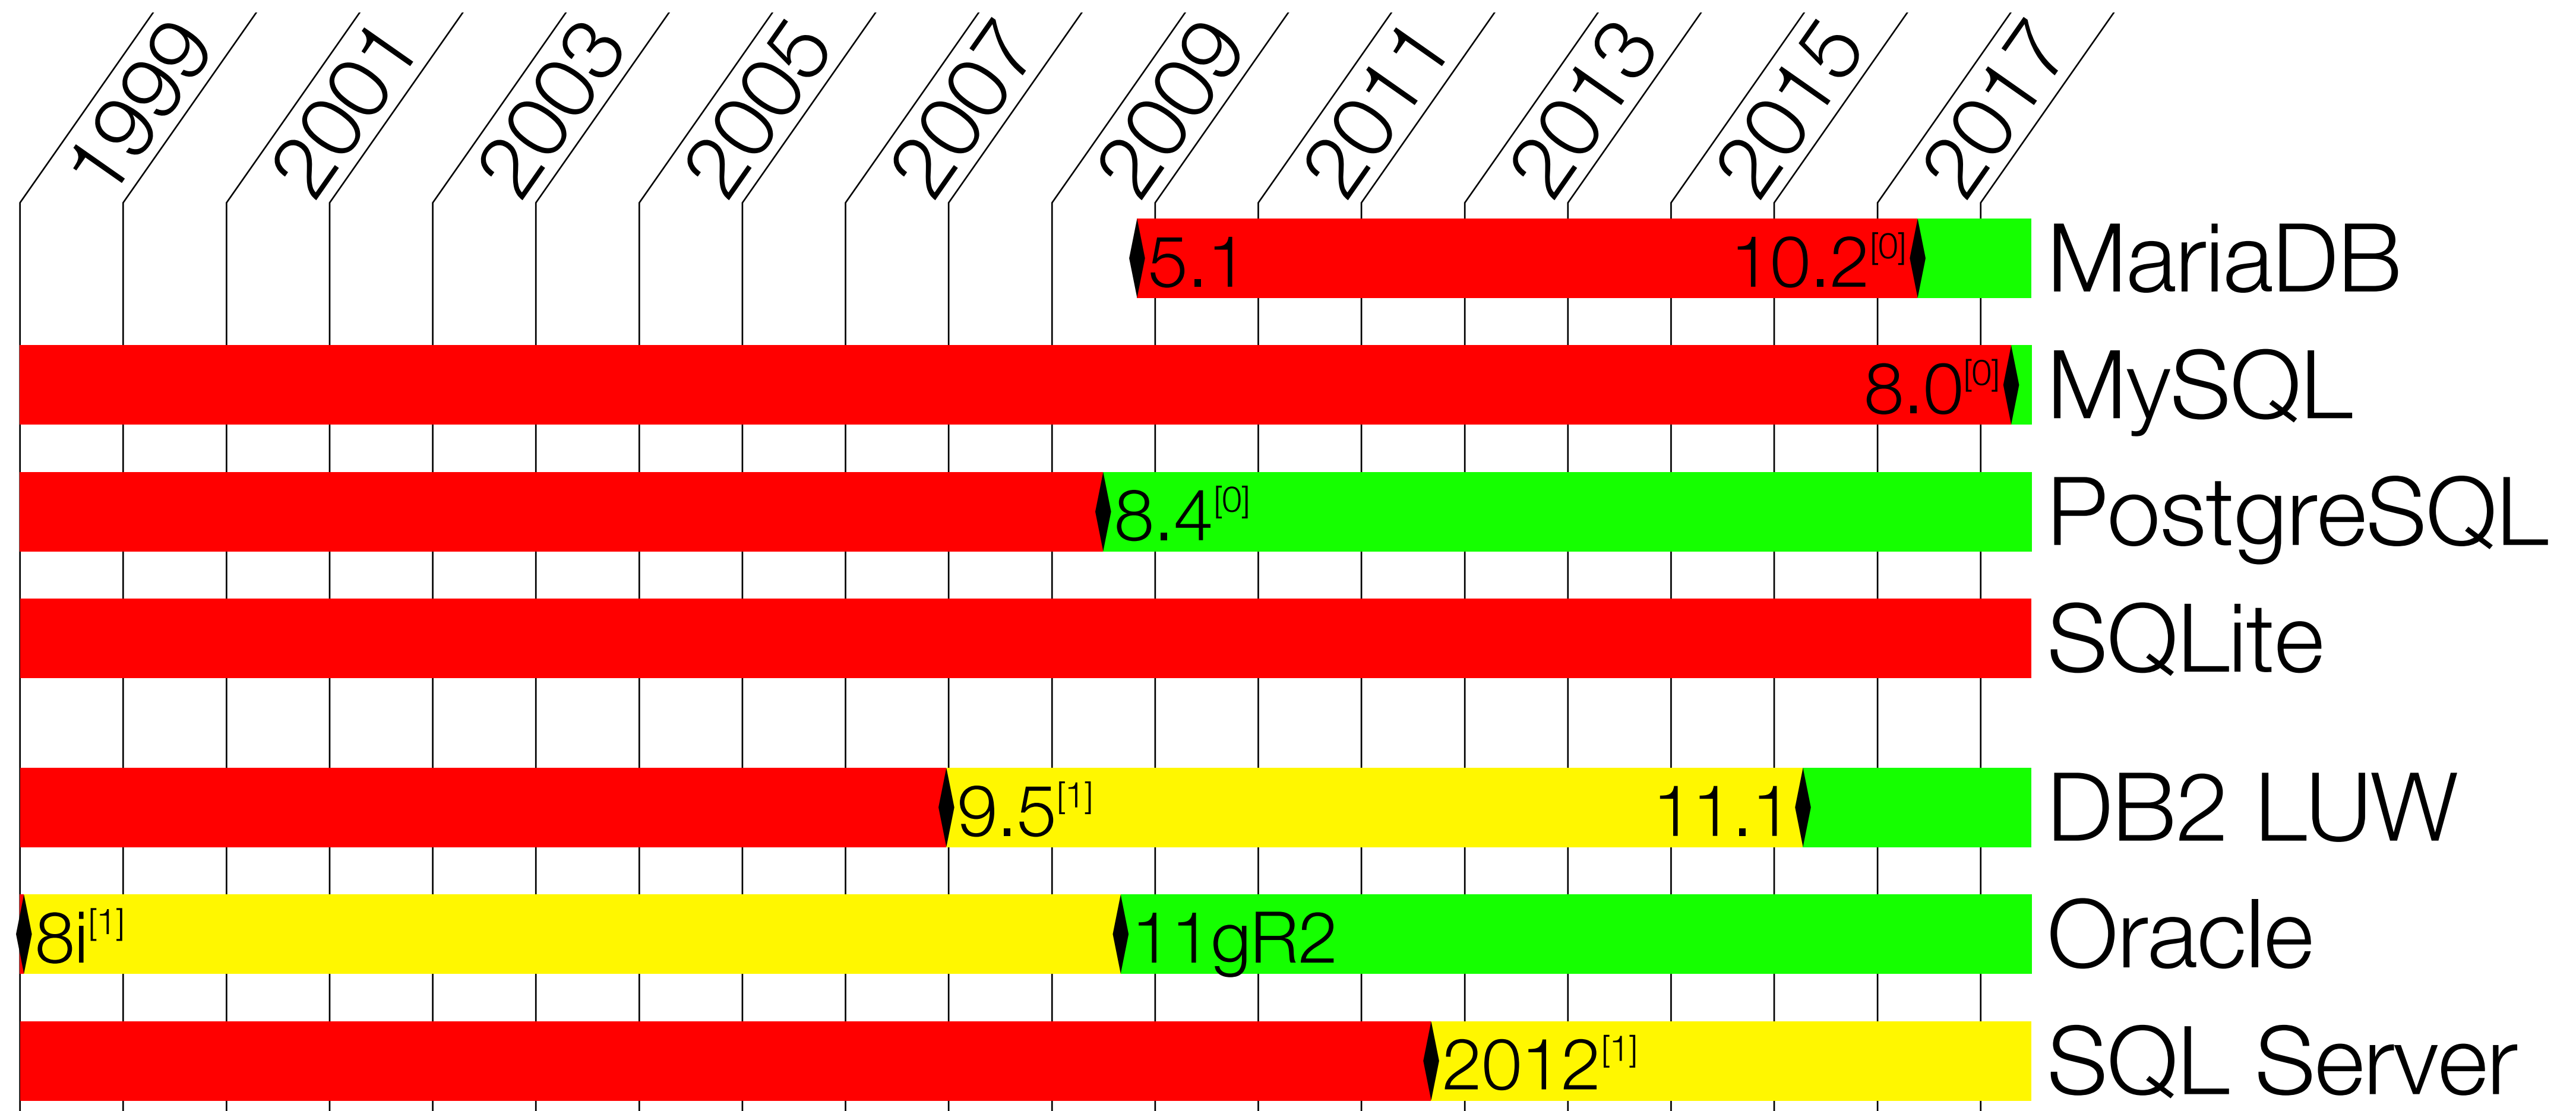
NTH\_VALUE(col, n) FROM FIRST/LAST

RESPECT/IGNORE NULLS

*Not supported  
by PostgreSQL  
(as of 11)*

# OVER (LEAD, LAG, ...)

Since SQL:2011



<sup>[0]</sup>No IGNORE NULLS and FROM LAST

<sup>[1]</sup>No NTH\_VALUE

# OVER

SQL:2011 **groups** option

# OVER (groups option)

Since SQL:2011



rows,  
range

ORDER BY x  
<frame unit> between 1 preceding  
and 1 following

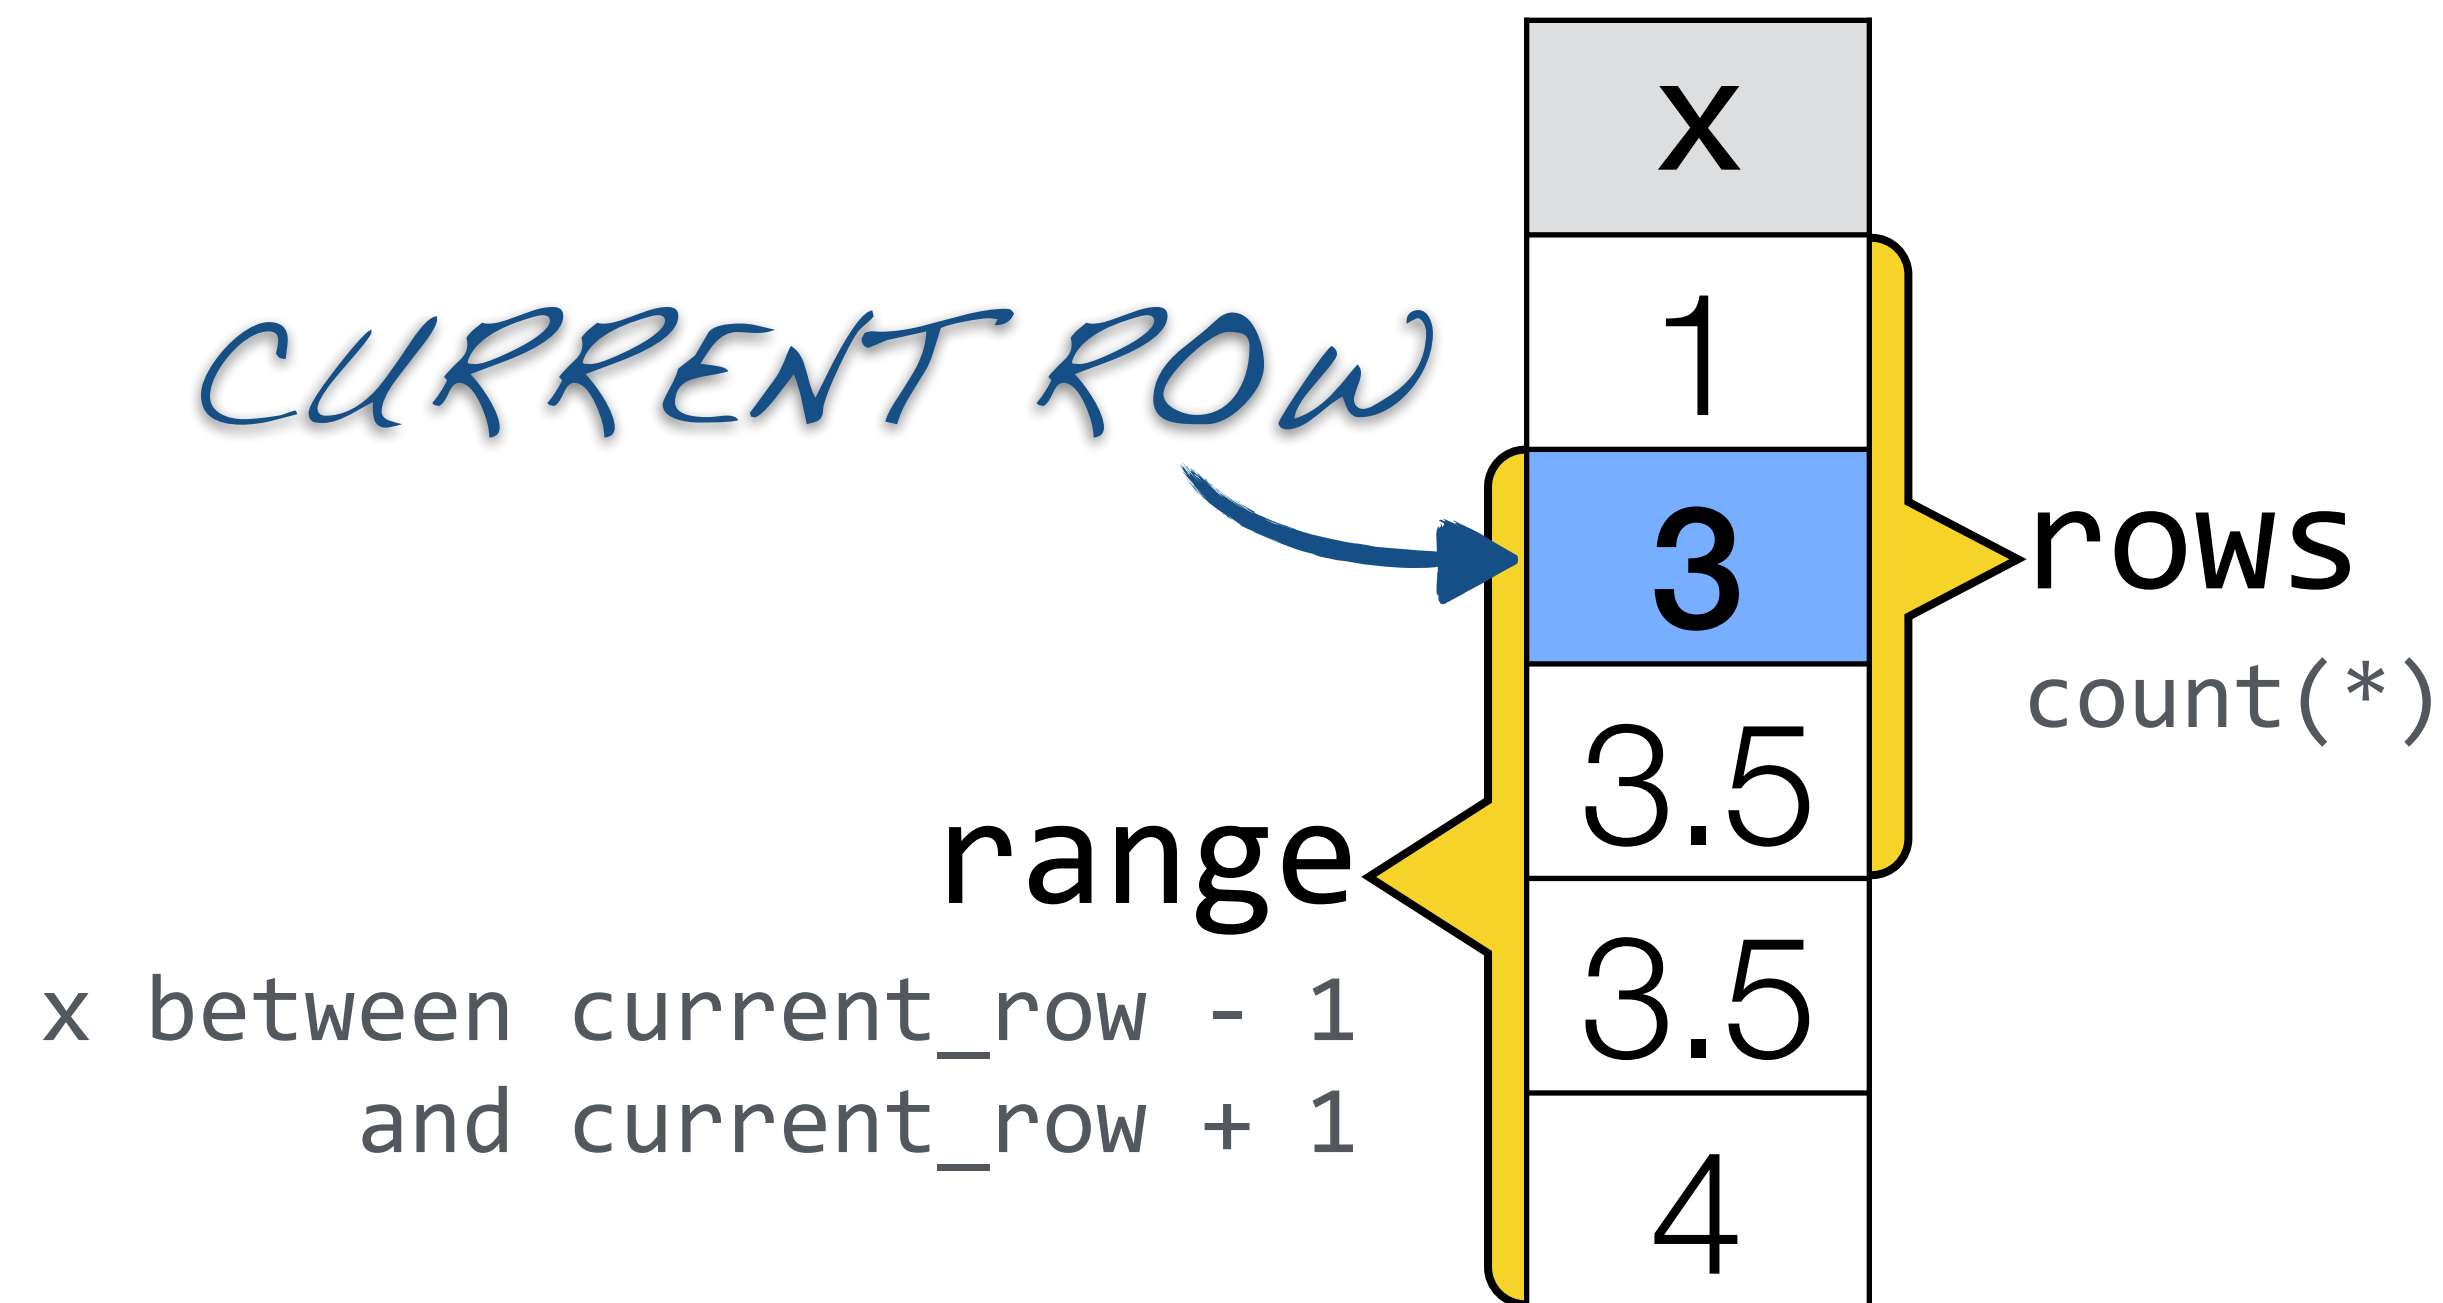


# OVER (groups option)

Since SQL:2011

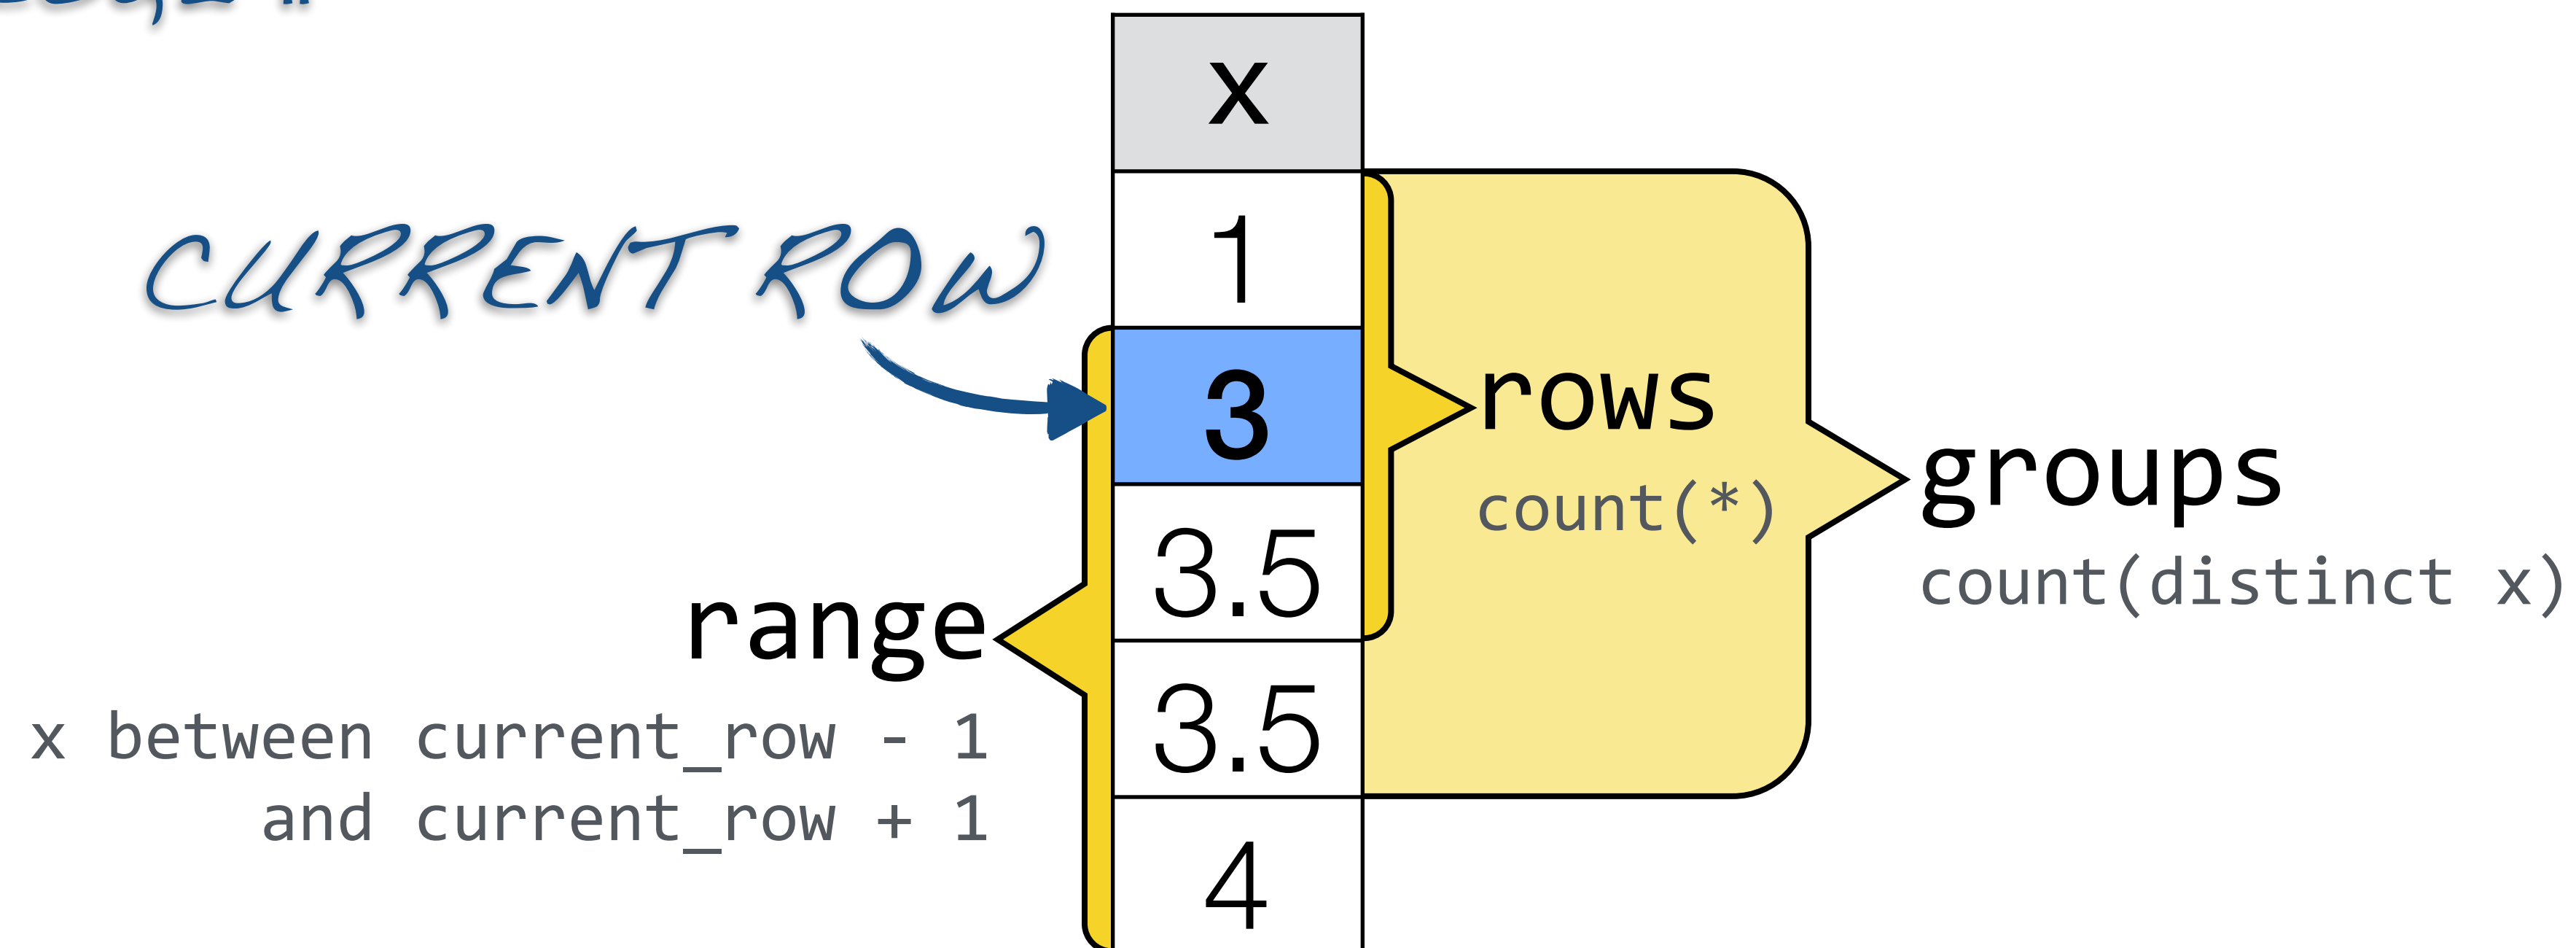
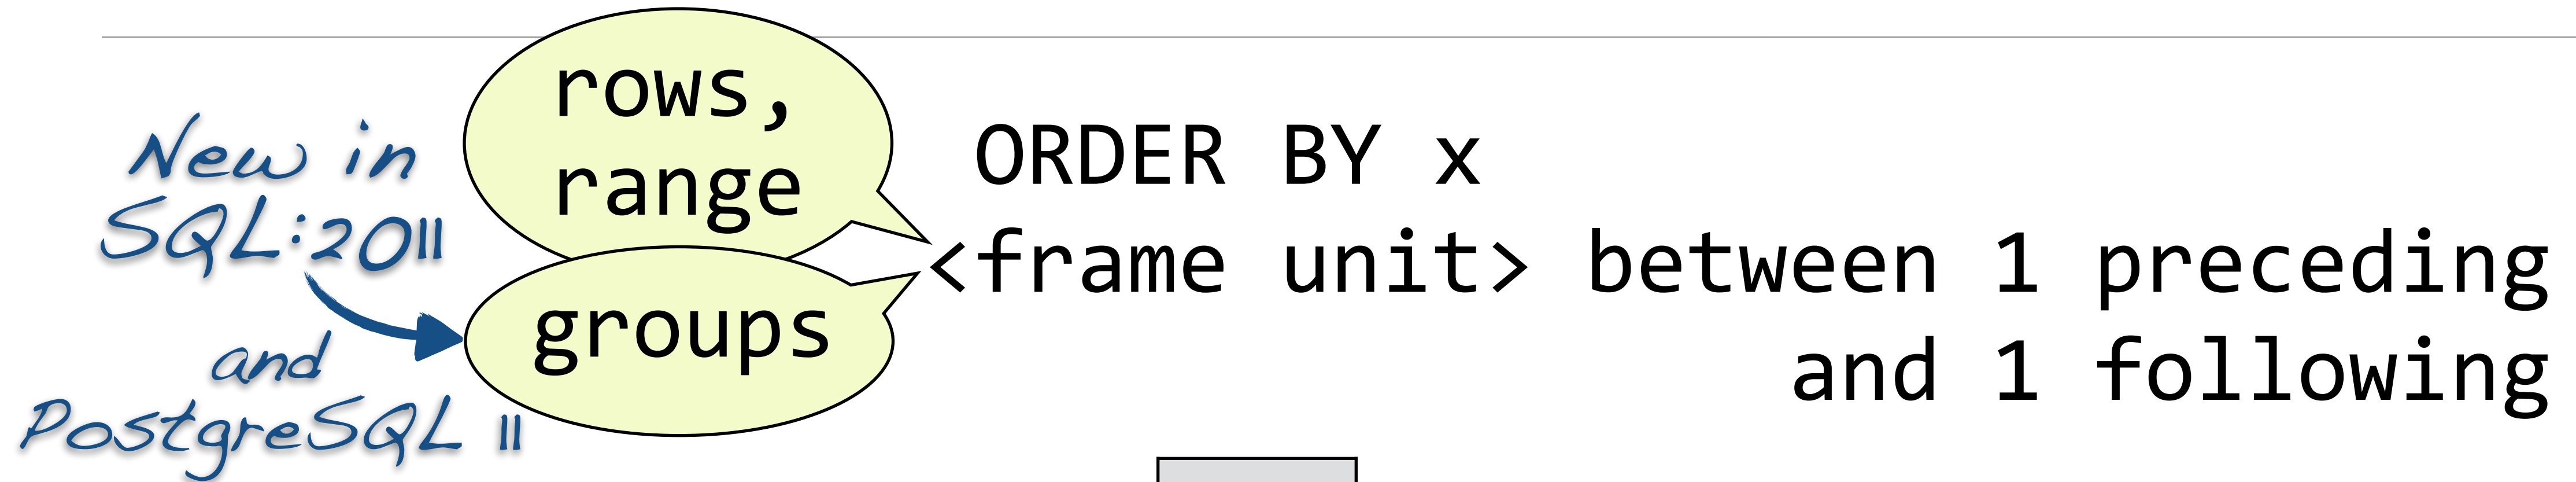
rows,  
range

ORDER BY x  
<frame unit> between 1 preceding  
and 1 following



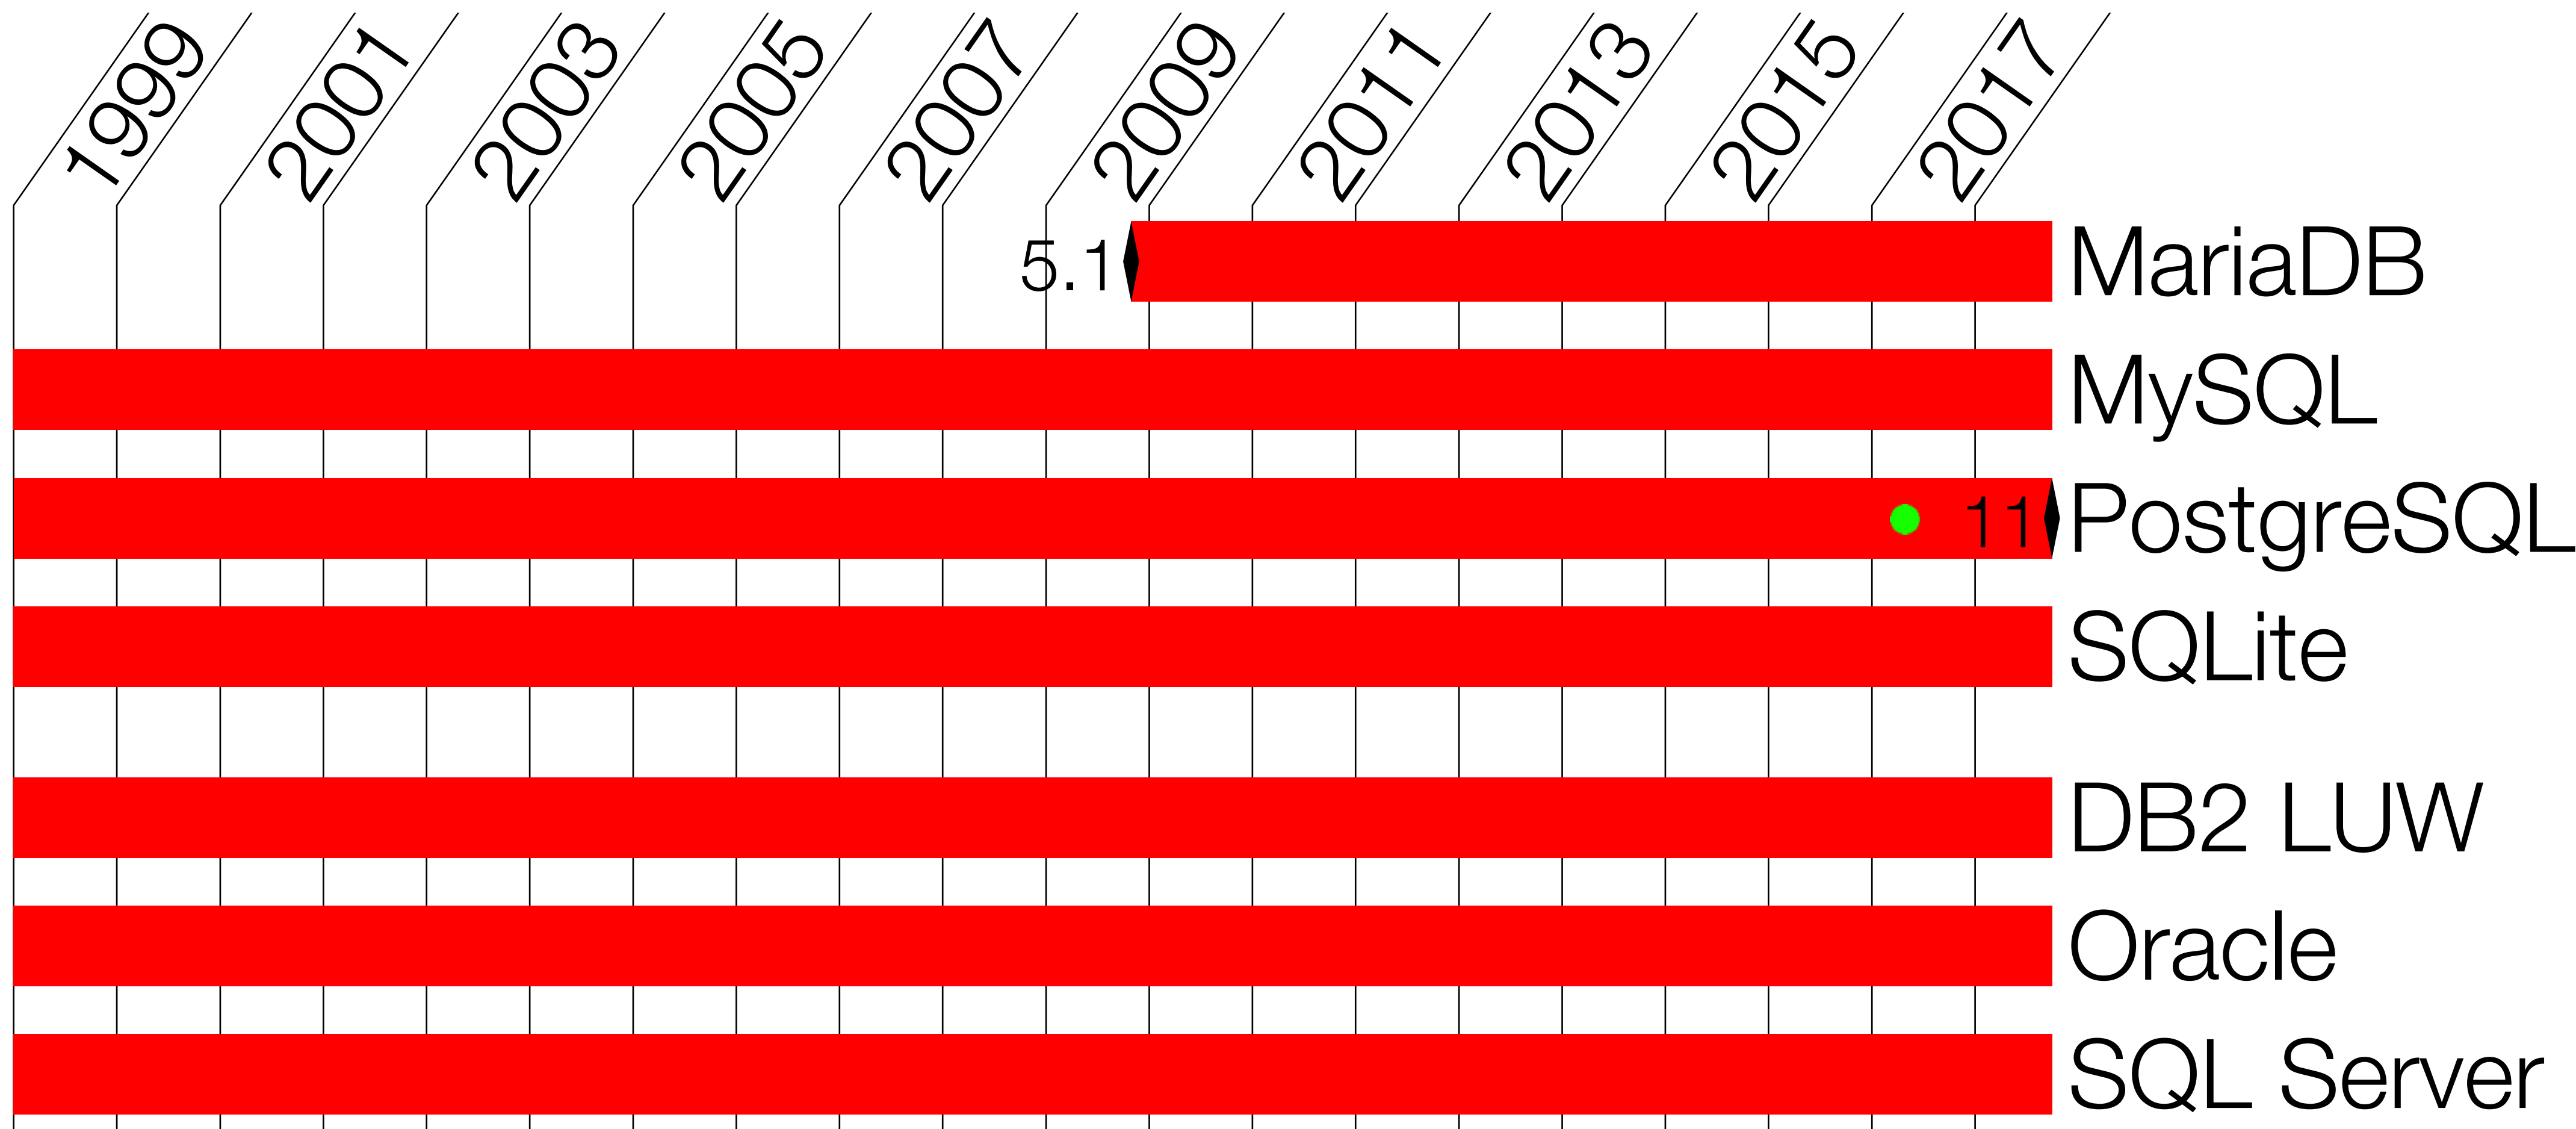
# OVER (groups option)

Since SQL:2011



OVER (groups option)

Since SQL:2011



# OVER

SQL:2003 frame exclusion



# OVER (frame exclusion)

Since SQL:2003

OVER (ORDER BY  
... BETWEEN ...  
exclude [ **no others** | current row  
| group | ties ] )

*default*

*New in  
PostgreSQL 11*

# OVER (frame exclusion)

Since SQL:2003

OVER (ORDER BY  
... BETWEEN ...  
exclude [ no others | **current row**  
| group | ties ] )

*default*

*New in  
PostgreSQL 11*

current row

x
1
2
<b>2</b>
2
3

# OVER (frame exclusion)

Since SQL:2003

OVER (ORDER BY  
... BETWEEN ...  
exclude [ no others | current row  
| **group** | ties ] )

*default*

*New in  
PostgreSQL 11*

group  
x = current\_row

current row

x
1
2
<b>2</b>
2
3

# OVER (frame exclusion)

Since SQL:2003

OVER (ORDER BY  
... BETWEEN ...  
exclude [ no others | current row  
| group | **ties** ] )

*default*

*New in  
PostgreSQL 11*

group  
x = current\_row

current row

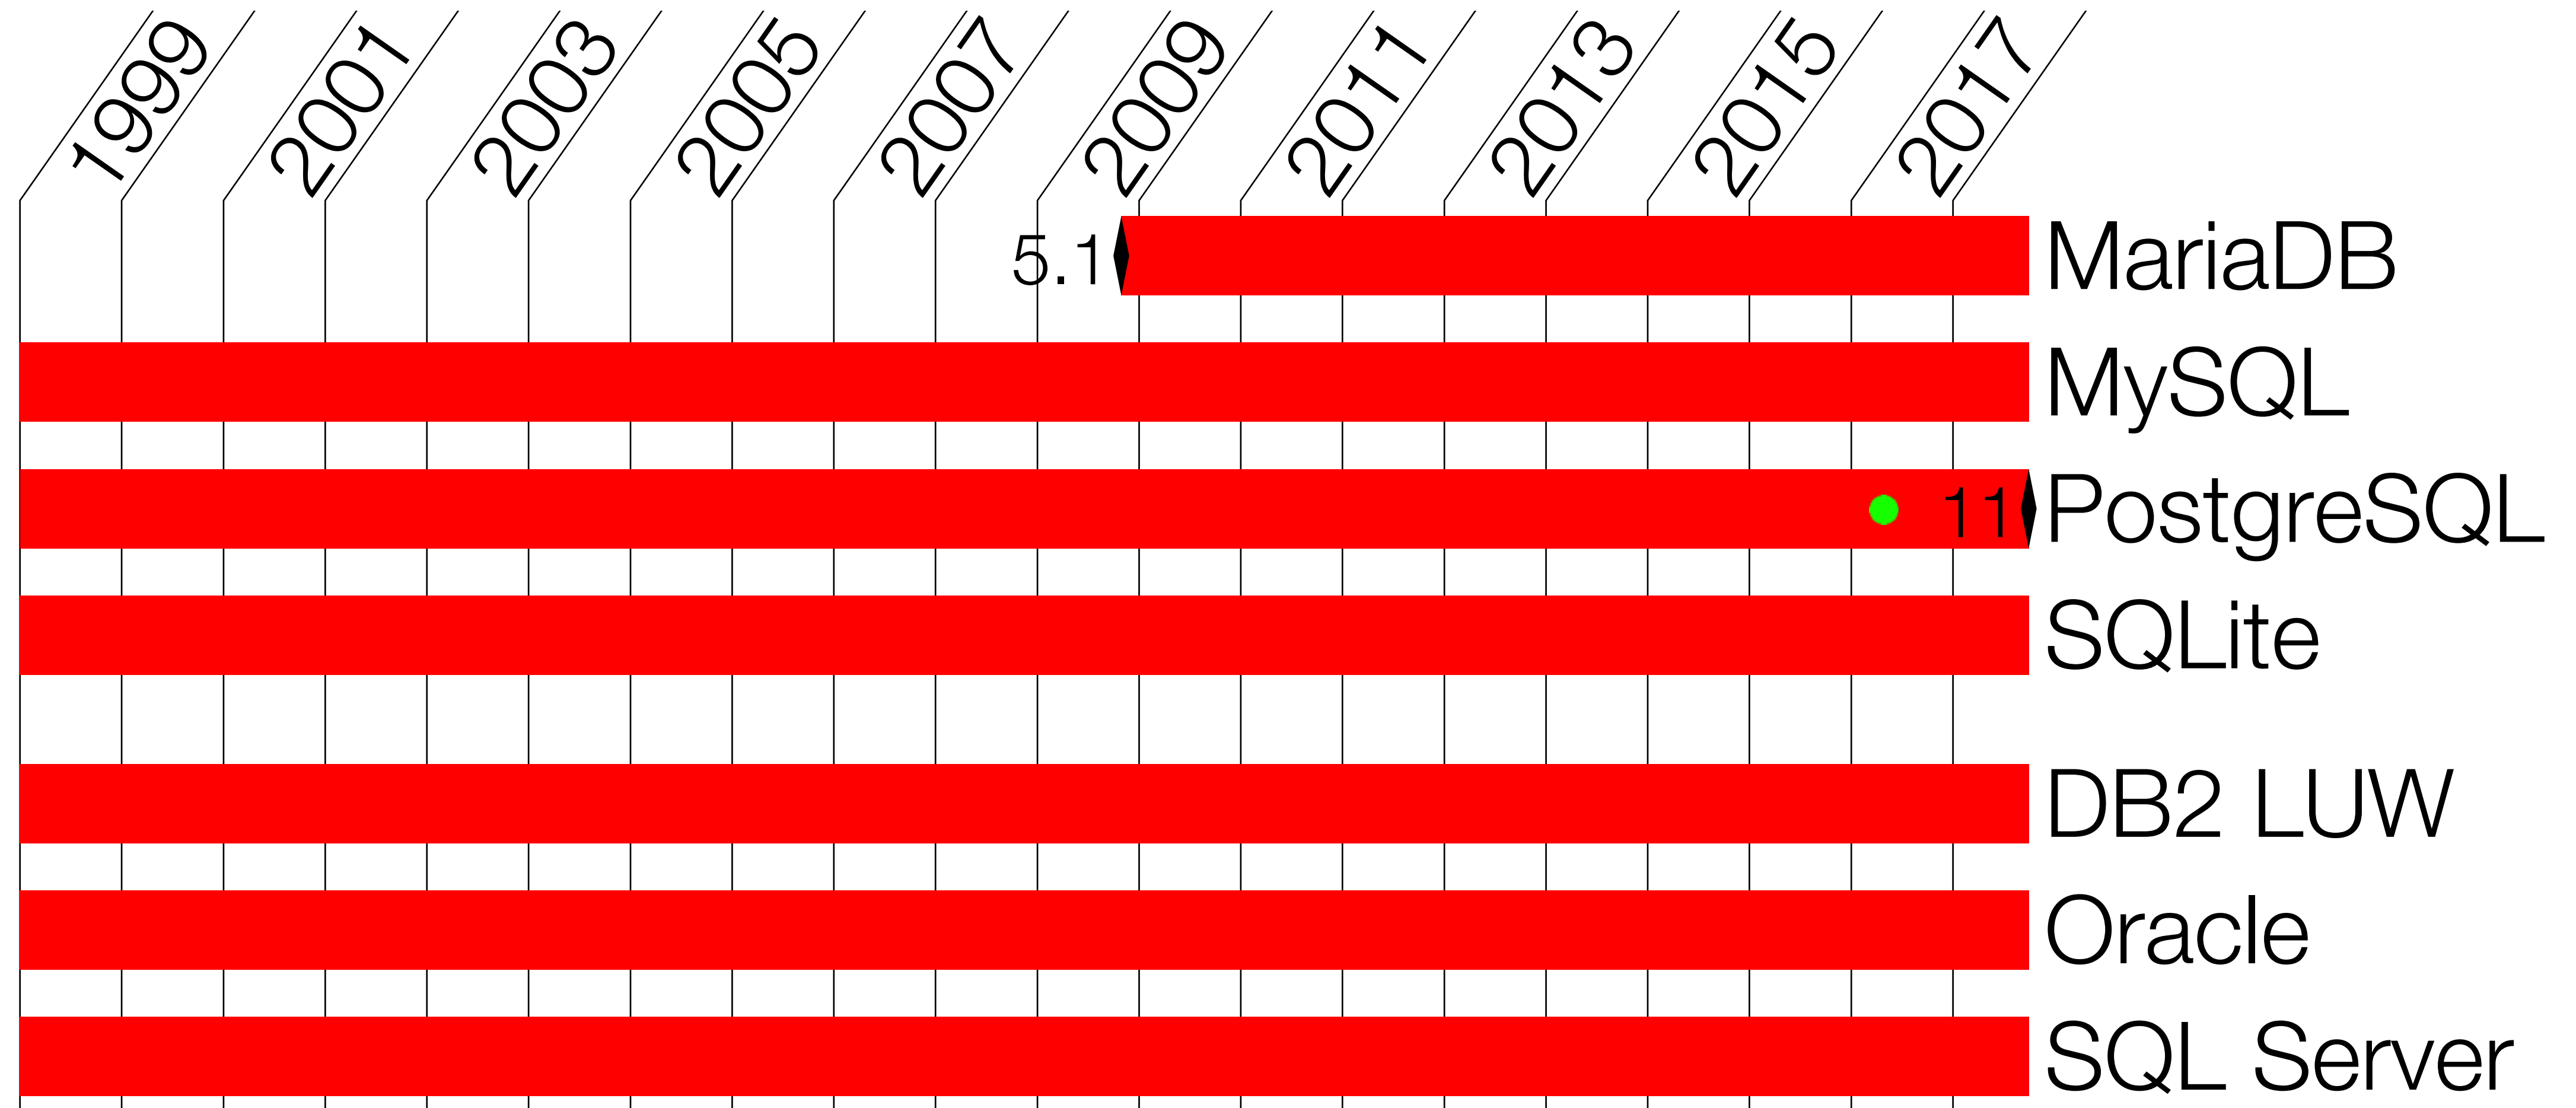
x
1
2
2
3

ties



**OVER** (frame exclusion)

Since SQL:2011



PostgreSQL 12+

20??-??-??

**MERGE**

# MERGE

## The Problem

Copying rows from another table is easy:

```
INSERT INTO <target>
SELECT ...
    FROM <source>
WHERE NOT EXISTS (SELECT *
                   FROM <target>
                   WHERE ...
                   )
```



*Both,  
<target> and <source>  
are in scope here.*



# MERGE

## The Problem

Deleting rows that exist in another table is also possible:

```
DELETE FROM <target>  
WHERE EXISTS (SELECT *  
              FROM <source>  
              WHERE ...  
              )
```



*Both,  
<target> and <source>  
are in scope here.*

# MERGE

## The Problem

Updating rows from another table is awkward:

```
UPDATE <target>  
  SET ...  
 WHERE ...
```

*Bringing another tables  
rows into scope of the  
SET clause is tricky*

*Sometimes, updatable  
views can help.*

*Subqueries are a more  
common choice*

# MERGE

## The Problem

---

Updating rows from another table is awkward:

```
UPDATE <target>  
    SET (col1, col2) = (SELECT col1, col2  
                        FROM <source>  
                        WHERE ...  
                        )  
WHERE ...
```



*Both,  
<target> and <source>  
are in scope here.*

# MERGE

Since SQL:2008

SQL:2008 introduced `merge` to improve this situation two-fold:

- ▶ It has always two tables in scope, the source can be a derived table (subquery).
- ▶ It can do `insert`, `update`, or `delete` in one go.

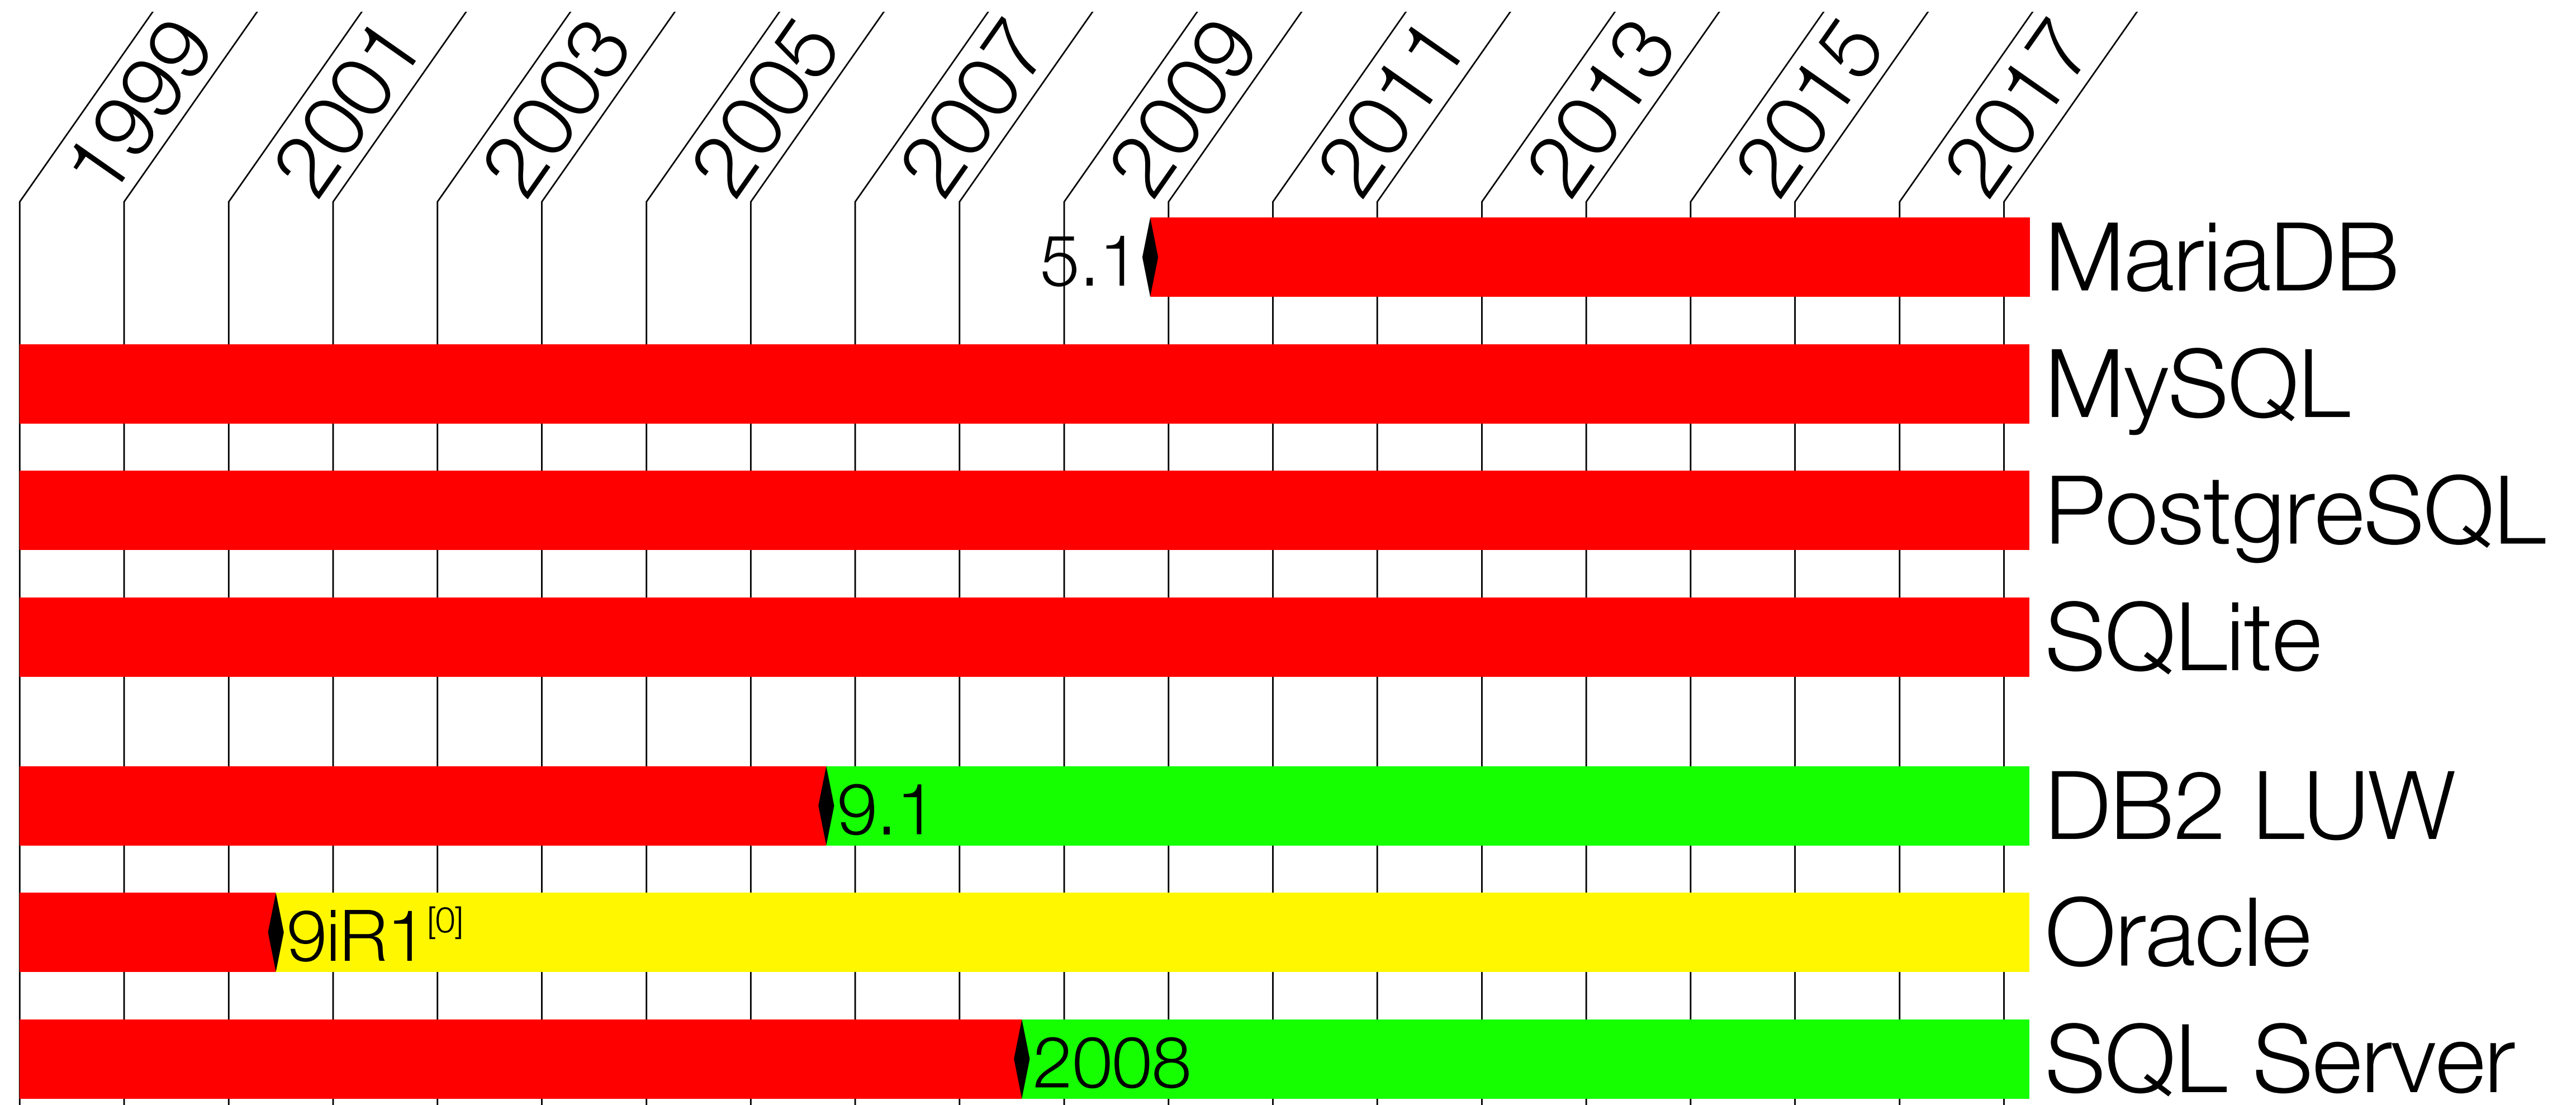
```
MERGE INTO <target>  
USING <source>  
ON <join condition>
```

*WHEN/THEN  
can appear many times*

```
WHEN MATCHED [AND <cond>] THEN [UPDATE...|DELETE...]  
WHEN NOT MATCHED [AND <cond>] THEN INSERT...
```

# MERGE

Since SQL:2008



<sup>[0]</sup>No AND condition.



(standard) **JSON**

# JSON

---

PostgreSQL got the ***JSON data type*** and  
the first JSON related functions  
with 9.2 (2012)

ISO added SQL JSON functions (but ***no type***)  
with SQL:2016

Both approaches follow a different mantra.  
The following slides only reflects the standards perspective.

# JSON

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
is [not] json predicate	✗	✗	✓ <sub>0</sub>	✗	✗	✗
on [ error   empty ] clauses	✗	✗	✓ <sub>1</sub>	✗	✗	✗

<sup>0</sup> No type constraints: is json [ value | array | object | scalar ]. Also unique keys (T822).

<sup>1</sup> No unknown on error. No expressions in default ... on [ error | empty ]

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
json_object	✗	✗	✓ <sub>0</sub>	✗	✗	✗
json_array	✗	✓ <sub>1</sub>	✓ <sub>2</sub>	✗	✗	✓ <sub>1</sub>
json_objectagg	✗	✗ <sub>3</sub>	✓ <sub>4</sub>	✗	✗	✗
json_arrayagg(... order by ...)	✗	✗	✓ <sub>5</sub>	✗	✗	✗

<sup>0</sup> No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys].

<sup>1</sup> Defaults to absent on null. No construction by query: json\_array(select ...).

<sup>2</sup> No construction by query: json\_array(select ...).

<sup>3</sup> Supports comma (,) instead of values or colon (:).

<sup>4</sup> No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys]. Sup

<sup>5</sup> Absent on null is buggy.

# JSON

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
json_exists	✗	✗	✓	✗	✗	✗
json_value	✗	✗	✓ <sub>0</sub>	✗	✓ <sub>1</sub>	✗
json_query	✗	✗	✓ <sub>3</sub>	✗	✓ <sub>2</sub>	✗
json_table	✗	✓ <sub>4</sub>	✓ <sub>4</sub>	✗	✗	✗

<sup>0</sup> Only returning [ varchar2 | number ] — neither is a standard type.

<sup>1</sup> Defaults to error on error.

<sup>2</sup> No quotes behavior: [ ~~keep~~ | ~~omit~~ ] quotes.

<sup>3</sup> with unconditional wrapper seems to be buggy. No quotes behavior: [ ~~keep~~ | ~~omit~~ ] quotes.

<sup>4</sup> Without plan clause.



# JSON

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
JSON path: lax mode (default)	✗	✓ <sub>1</sub>	✓	✗	✓ <sub>0</sub>	✓ <sub>1</sub>
JSON path: strict mode	✗	✗	✗ <sub>2</sub>	✗	✓	✗
JSON path: item method	✗	✗	✓ <sub>3</sub>	✗	✗	✗
JSON path: multiple subscripts	✗	✗	✓	✗	✗	✗
JSON path: .* member accessor	✗	✓	✓	✗	✗	✗
JSON path: filter expressions	✗	✗	✓ <sub>4</sub>	✗	✗	✗
JSON path: starts with	✗	✗	✓ <sub>4</sub>	✗	✗	✗
JSON path: like_regex	✗	✗	✗	✗	✗	✗

<sup>0</sup> Lax mode does not unwrap arrays.

<sup>1</sup> Keyword lax not accepted (only default mode). Lax mode does not unwrap arrays.

<sup>2</sup> Keyword strict is accepted but not honored.

<sup>3</sup> Only in filters. Not supporting size(), datetime(), keyvalue(). type() returns null for arrays.

<sup>4</sup> Not in json\_value. Not in json\_query. Not in json\_table. Only as last step of expression.

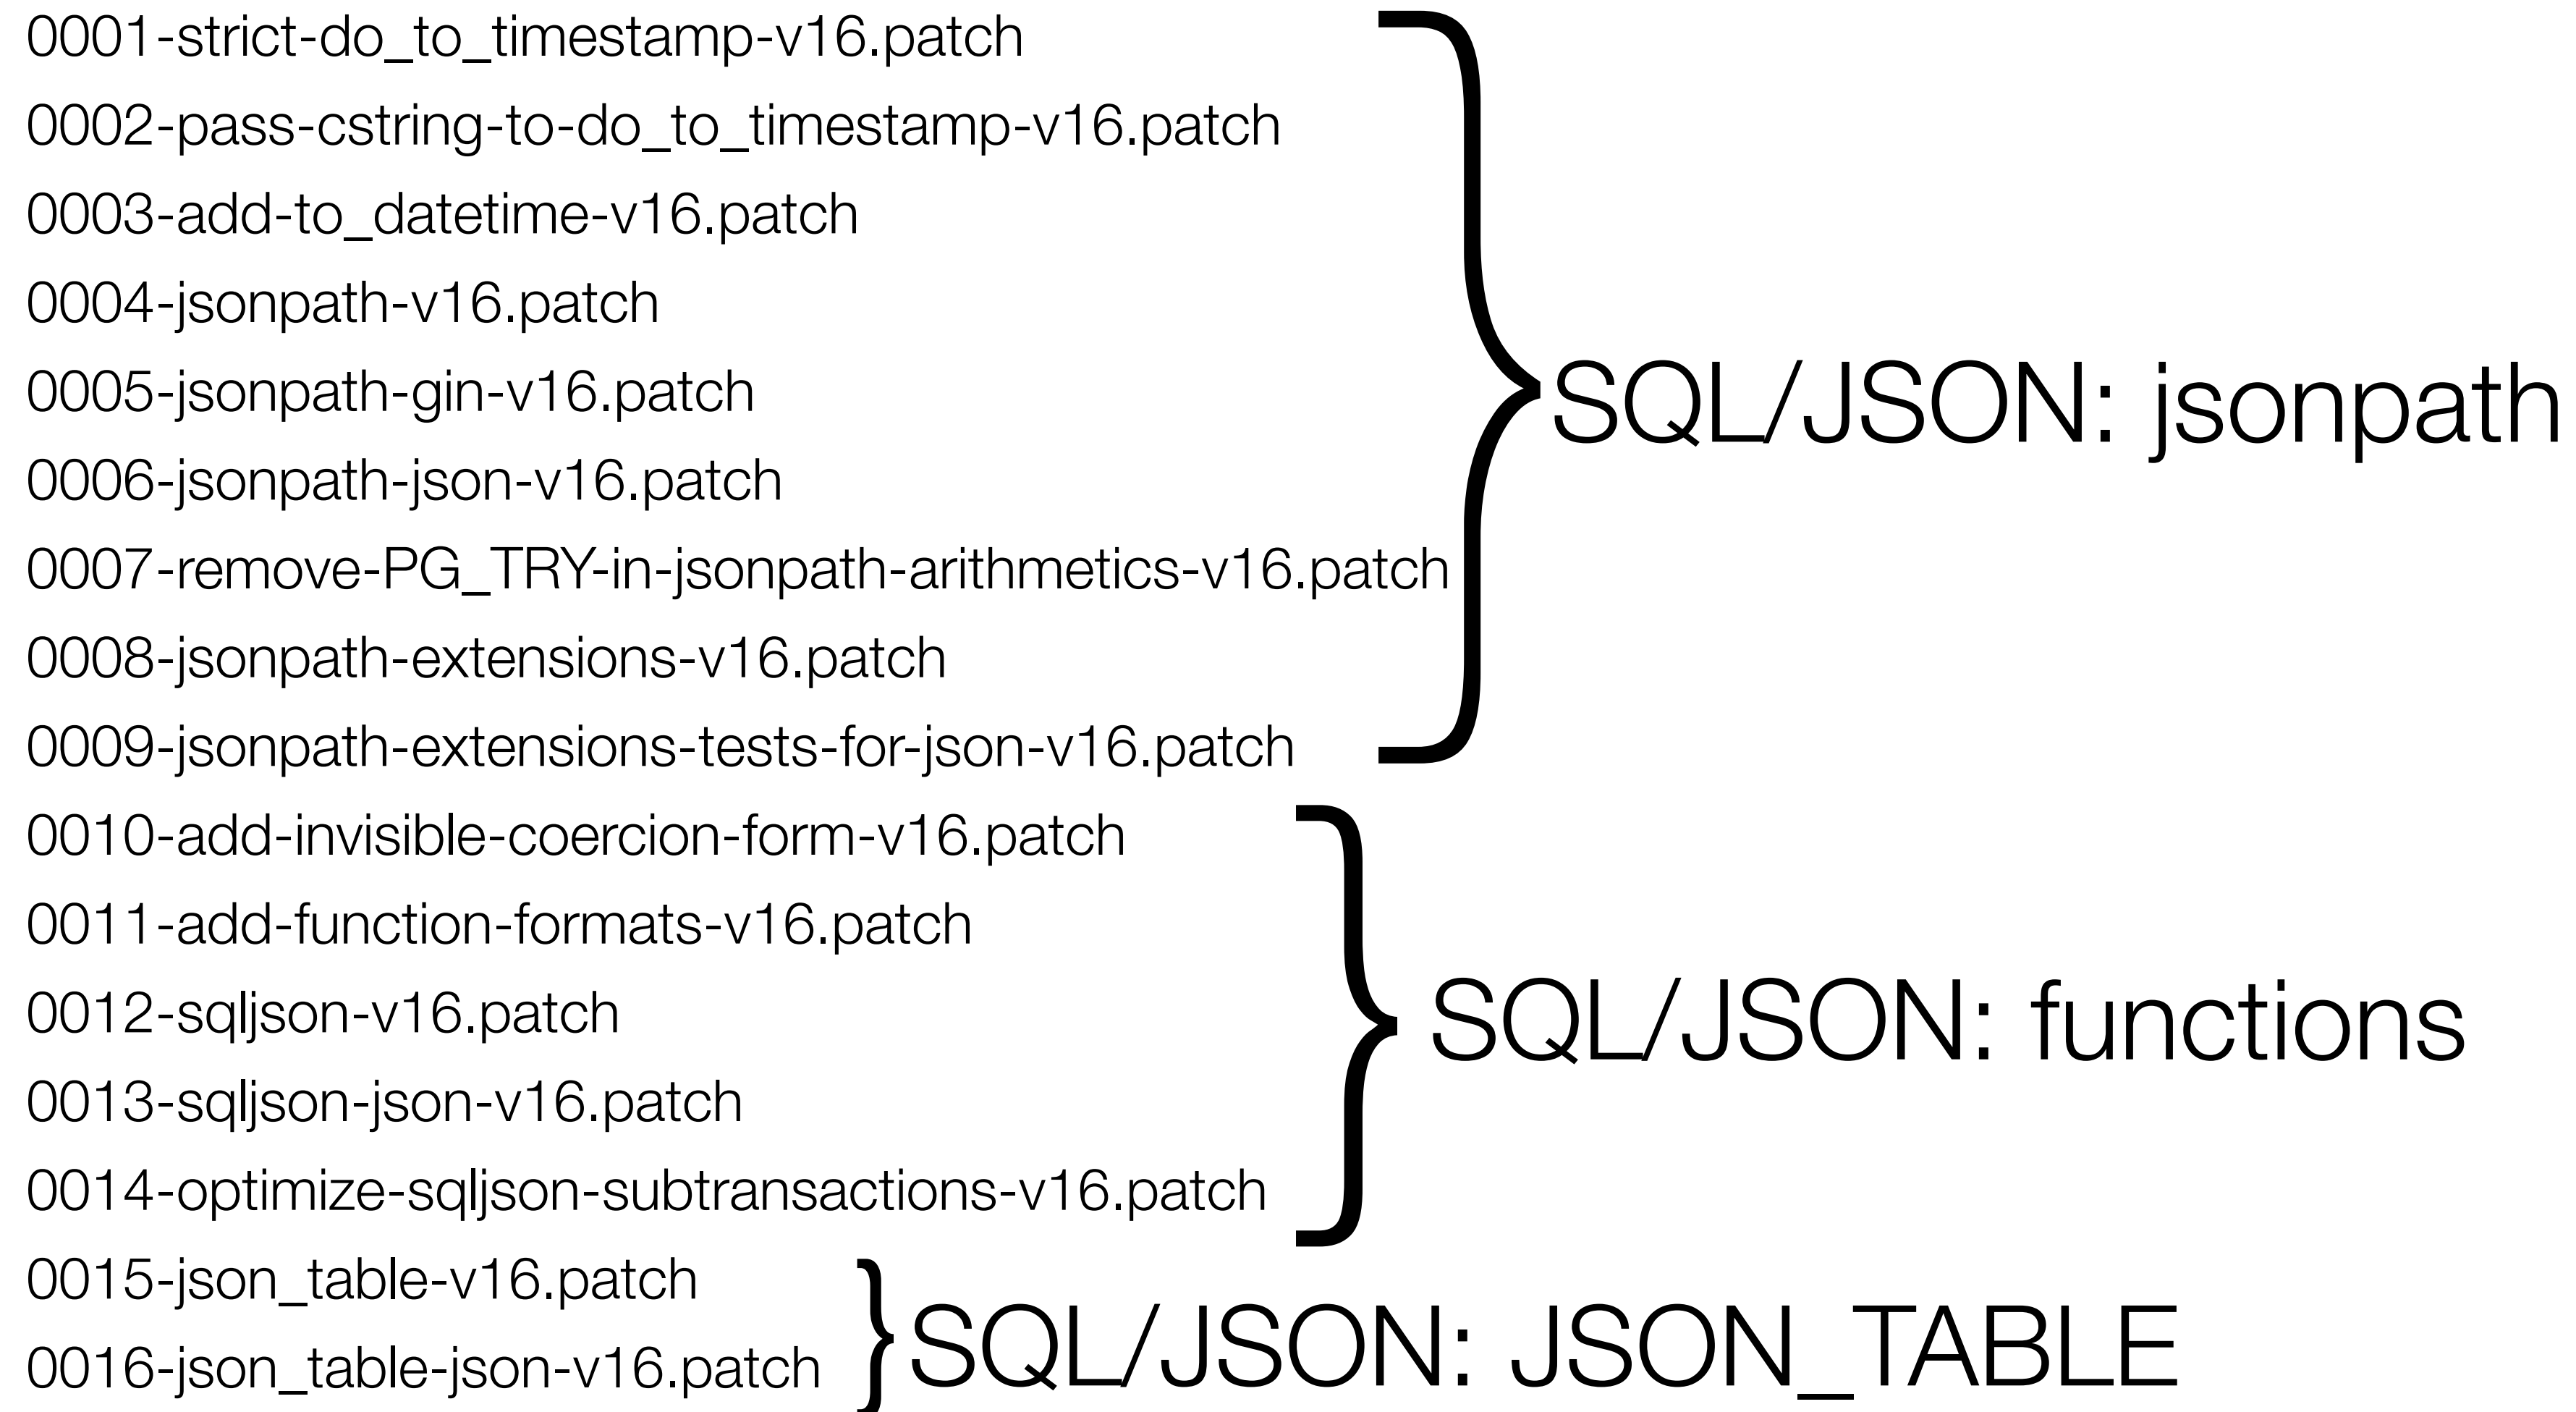
<sup>5</sup> Not in json\_query. Only as last step of expression.



# JSON — Preliminary testing of patches

---

Used a06e56b24 as basis, applied those patches on top:



# JSON — Preliminary testing of patches

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
is [not] json predicate	✗	✗	✓ <sub>0</sub>	✓ <sub>1</sub>	✗	✗
on [ error   empty ] clauses	✗	✗	✓ <sub>2</sub>	✓	✗	✗

<sup>0</sup> No type constraints: is json [ value | array | object | scalar ]. Also unique keys (T822).  
<sup>1</sup> Also unique keys (T822).  
<sup>2</sup> No unknown on error. No expressions in default ... on [ error | empty ]

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
json_object	✗	✗	✓ <sub>0</sub>	✓ <sub>0</sub>	✗	✗
json_array	✗	✓ <sub>1</sub>	✓ <sub>2</sub>	✓ <sub>2</sub>	✗	✓ <sub>1</sub>
json_objectagg	✗	✗ <sub>3</sub>	✓ <sub>4</sub>	✓ <sub>0</sub>	✗	✗
json_arrayagg(... order by ...)	✗	✗	✓	✓	✗	✗

No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys].  
Defaults to absent on null. No construction by query: json\_array(select ...).  
No construction by query: json\_array(select ...).  
Supports comma (,) instead of values or colon (:).  
No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys]. Su

# JSON — Preliminary testing of patches

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
json_exists	✗	✗	✓	✓	✗	✗
json_value	✗	✗	✓ <sub>0</sub>	✓	✓ <sub>1</sub>	✗
json_query	✗	✗	✓ <sub>3</sub>	✓ <sub>2</sub>	✓ <sub>2</sub>	✗
json_table	✗	✓ <sub>4</sub>	✓ <sub>4</sub>	✓ <sub>4</sub>	✗	✗

<sup>0</sup> Only returning [ varchar2 | number ] — neither is a standard type.  
<sup>1</sup> Defaults to error on error.  
<sup>2</sup> No quotes behavior: [ ~~keep~~ | ~~omit~~ ] quotes.  
<sup>3</sup> with unconditional wrapper seems to be buggy. No quotes behavior: [ ~~keep~~ | ~~omit~~ ] quotes.  
<sup>4</sup> Without plan clause.



# JSON — Preliminary testing of patches

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
JSON path: lax mode (default)	✗	✓ <sub>1</sub>	✓	✓	✓ <sub>0</sub>	✓ <sub>1</sub>
JSON path: strict mode	✗	✗	✗ <sub>2</sub>	✗ <sub>2</sub>	✓	✗
JSON path: item method	✗	✗	✓ <sub>3</sub>	✓	✗	✗
JSON path: multiple subscripts	✗	✗	✓	✓	✗	✗
JSON path: .* member accessor	✗	✓	✓	✓	✗	✗
JSON path: filter expressions	✗	✗	✓ <sub>4</sub>	✓ <sub>5</sub>	✗	✗
JSON path: starts with	✗	✗	✓ <sub>4</sub>	✓ <sub>7</sub>	✗	✗
JSON path: like_regex	✗	✗	✗	✓	✗	✗

<sup>0</sup> Lax mode does not unwrap arrays.

<sup>1</sup> Keyword lax not accepted (only default mode). Lax mode does not unwrap arrays.

<sup>2</sup> Keyword strict is accepted but not honored.

<sup>3</sup> Only in filters. Not supporting size(), datetime(), keyvalue(). type() returns null for arrays.

<sup>4</sup> Not in json\_value. Not in json\_query. Not in json\_table. Only as last step of expression.

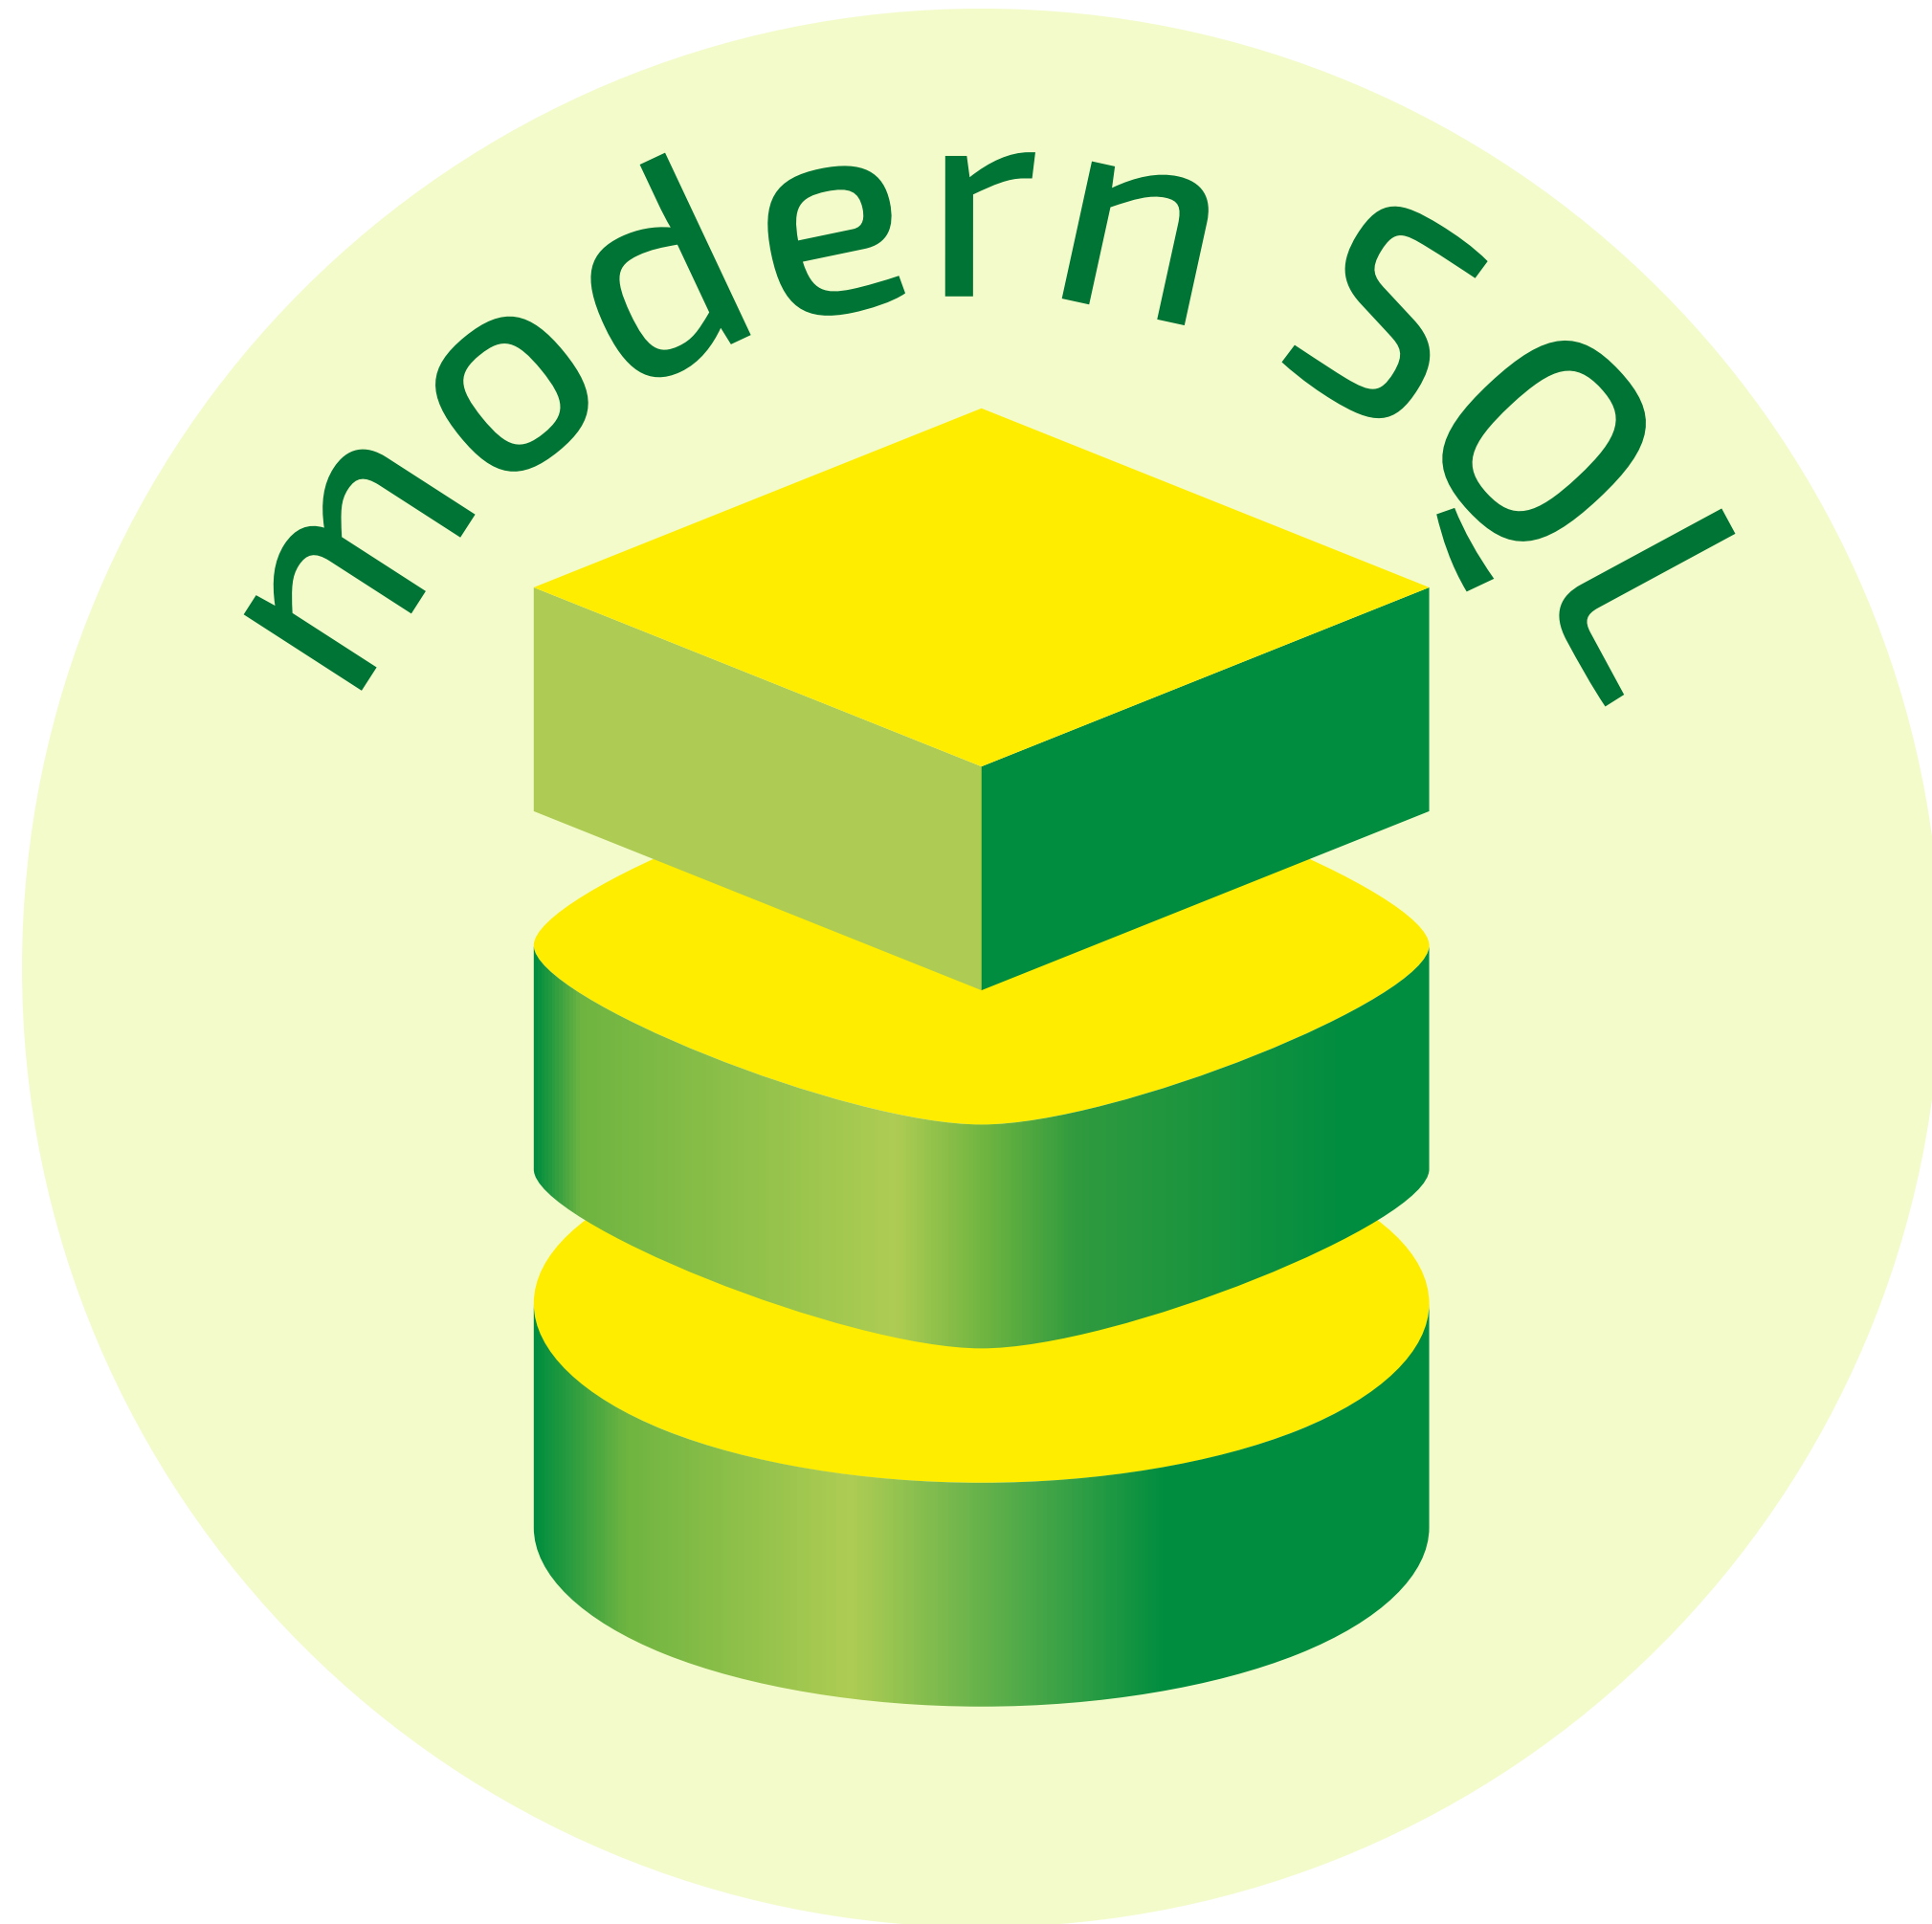
<sup>5</sup> Not in json\_value. Not in json\_query.

<sup>6</sup> Not in json\_query. Only as last step of expression.

<sup>7</sup> Not in json\_value.

@ModernSQL

[modern-sql.com](https://modern-sql.com)



My other website:  
<https://use-the-index-luke.com>



Training & co: <https://winand.at/>



# Visit Vienna!

5-day intensive training — Sept 17-21

---

## SQL Performance Kick-Start

- Indexing, Indexing, Indexing...
- Joining
- Ordering, grouping,
- ...

## Analysis and Aggregation

- OVER, GROUPING SETS
- Avoiding self-joins
- Grouping consecutive events
- ...

## SQL Reloaded

- Type safety, NULL and 3VL
- Modern interpretation of the relational model
- Keeping historic data
- Design guidelines

## Recursive Queries

- WITH
- WITH RECURSIVE
- Use cases and examples

<https://winand.at/sql-training/5-day-training/aug-sept-2018>