

Standard SQL Gap Analysis

Features where PostgreSQL lags behind its competitors

PgCon.org 2018 – @MarkusWinand

MIND THE GAP

Background: Where the data comes from

I run modern-sql.com:



- ▶ Teaching “new” SQL features to developers
- ▶ Showing availability of those features in popular databases

Background: Where the data comes from

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
filter clause	✗	✗	✗	✓	✗	✗
Emulation using case	✓	✓ ⁰	✓	✓	✓	✓

⁰ The `filter_plugin` extension (3rd party) rewrites `filter` to `case` using regular expressions.

The charts are based on test cases.^[0]

The test cases are created while reading ISO/IEC 9075:2016.

The level of detail for different features varies widely at the moment.

^[0] Some “legacy charts” are still based on reading the docs.

One last Word

For brevity, I'm using the word
“wrong”
to mean
“not conforming to the standard”.

This neither implies that it is “bad” nor that it is a bug,
nor that it is worth changing.
I just means that it is not the way I understand the standard.

Less Complete or Conforming Features

EXTRACT

Get a Field from a Date or Time Value

EXTRACT: “Wrong” declared type

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
extract(... from <datetime>)	✓ ⁰	✓ ¹	✓	✓ ²	✗	✗
extract(... from <interval>)	✗	✗	✓	✓ ²	✗	✗
cast(<timestamp> as date)	✓	✓	✗ ³	✓	✓	✗
cast(<timestamp> as time)	✓	✓	✓	✓	✓	✗

⁰ No time zone fields.

¹ No time zone fields. SECOND does not include fractions. Use SECOND_MICROSECOND.

² Returns approximate numeric type.

³ See “Caution: Oracle Database” above.

EXTRACT: “Wrong” declared type

<code>extract(field from timestamp)</code>	double precision
<code>extract(field from interval)</code>	double precision
<code>extract(... from <datetime>)</code>	✓ ₀ ✓ ₁ ✓ ✓ ₂ ✗ ✗
<code>extract(... from <interval>)</code>	✗ ✗ ✓ ✓ ₂ ✗ ✗
<code>cast(<timestamp> as date)</code>	✓ ✓ ✗ ✓ ✓ ✗
<code>cast(<timestamp> as time)</code>	✓ ✓ ✓ ✓ ✓ ✗

⁰ No time zone fields.

¹ No time zone fields. SECOND does not include fractions. Use SECOND_MICROSECOND.

² Returns approximate numeric type.

³ See “Caution: Oracle Database” above.

EXTRACT: “Wrong” declared type

7) If <extract expression> is specified, then

Case:

- a) If <extract field> is a <primary datetime field> that does not specify SECOND or <extract field> is not a <primary datetime field>, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- b) Otherwise, the declared type of the result is an implementation-defined exact numeric type with scale not less than the specified or implied <time fractional seconds precision> or <interval fractional seconds precision>, as appropriate, of the SECOND <primary datetime field> of the <extract source>.

NO time zone fields.

¹ No time zone fields. SECOND does not include fractions. Use SECOND_MICROSECOND.

² Returns approximate numeric type.

³ See “Caution: Oracle Database” above.

<code>extract(field from timestamp)</code>	double precision
<code>extract(field from interval)</code>	double precision

[RESPECT | IGNORE] NULLS

Skip over `null` values in window functions

`lead`, `lag`, `first_value`, `last_value`, `nth_value`

(T616, T618)

Window Functions: null handling, from last

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
LEAD and LAG	✓ ⁰	✓ ¹	✓ ²	✓	✓ ²	✓ ²	✗
FIRST_VALUE, LAST_VALUE	✓ ³	✓ ²	✓ ²	✓	✓ ²	✓ ²	✗
NTH_VALUE	✓	✓ ⁴	✓ ⁴	✓	✓ ⁴	✗	✗
Nested window functions	✗	✗	✗	✗	✗	✗	✗

⁰ No IGNORE NULLS Different syntax: `first_value(<expr>, 'IGNORE NULLS')` (it's a string argument)

¹ No IGNORE NULLS No default possible (3rd argument).

² No IGNORE NULLS

³ No IGNORE NULLS Different syntax: `lead(<expr>, 1, null, 'IGNORE NULLS')` (it's a string argument)

⁴ No IGNORE NULLS. No FROM LAST

Window Functions: null handling, from last

Note

The SQL standard defines a `RESPECT NULLS` or `IGNORE NULLS` option for `lead`, `lag`, `first_value`, `last_value`, and `nth_value`. This is not implemented in PostgreSQL: the behavior is always the same as the standard's default, namely `RESPECT NULLS`. Likewise, the standard's `FROM FIRST` or `FROM LAST` option for `nth_value` is not implemented: only the default `FROM FIRST` behavior is supported. (You can achieve the result of `FROM LAST` by reversing the `ORDER BY` ordering.)

	\downarrow 3	\downarrow 2	\downarrow 2	\downarrow		
<code>NTH_VALUE</code>	✓	✓ 4	✓ 4	✓	✓ 4	✗
Nested window functions	✗	✗	✗	✗	✗	✗

⁰ No `IGNORE NULLS` Different syntax: `first_value(<expr>, 'IGNORE NULLS')` (it's a string argument)

¹ No `IGNORE NULLS` No default possible (3rd argument).

² No `IGNORE NULLS`

³ No `IGNORE NULLS` Different syntax: `lead(<expr>, 1, null, 'IGNORE NULLS')` (it's a string argument)

⁴ No `IGNORE NULLS`. No `FROM LAST`

	Oracle	PostgreSQL	SQL Server	SQLite
✓	✓	✓	✗	
✓ 2	✓ 2	✓ 2	✗	
✓ 2	✓ 2	✓ 2	✗	
✓ 2	✓ 2	✗ 4	✗	✗

COUNT(DISTINCT ...) OVER(...)

Distinct aggregates as window function

(T611)

Window Functions: no distinct aggregates

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Aggregates (count, sum, min, ...)	✓	✓	✓	✓	✓	✓	✗
Distinct Aggregates	✗	✗	✗	✓	✗	✗	✗

FETCH [FIRST|NEXT] ...

The standard's answer to **LIMIT**, but more options

(T866, T867)

FETCH FIRST: no percent, no with ties

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Top-level fetch first	✓ ⁰	✗ ⁰	✓	✓	✗ ¹	✗ ⁰
Subqueries with fetch first	✓	✗ ⁰	✓	✓	✗ ¹	✗ ⁰
Top-level fetch first in views	✗ ²	✗ ⁰	✓	✓	✗ ¹	✗ ⁰
Dynamic quantity	✓	✗	✓	✓ ³	✗	✗
fetch first ... percent	✗	✗	✓	✗	✗ ⁴	✗
fetch first ... with ties	✗	✗	✓	✗	✗ ⁵	✗
SQL State 2201W if quantity < 1	✓ ⁶	✗	✗	✓ ⁶	✗	✗

⁰ Use proprietary limit

¹ Use proprietary top

² Use nested query: CREATE VIEW ... AS SELECT ... FROM (SELECT ... FROM ... FETCH FIRST ...) t

³ Requires parenthesis: (2)

⁴ Use proprietary select top ... percent

⁵ Use proprietary select top ... with ties

⁶ Not for 0 (zero)

FETCH FIRST: no percent, no with ties

Docs: unsupported features:

		F866					
		F867					
	FETCH FIRST clause: PERCENT option						
	FETCH FIRST clause: WITH TIES option						
Top-level fetch first in views		✗ ²	✗ ⁰	✓	✓	✗ ¹	✗ ⁰
Dynamic quantity		✓	✗	✓	✓ ³	✗	✗
fetch first ... percent		✗	✗	✓	✗	✗ ⁴	✗
fetch first ... with ties		✗	✗	✓	✗	✗ ⁵	✗
SQL State 2201W if quantity < 1		✓ ⁶	✗	✗	✓ ⁶	✗	✗

⁰ Use proprietary limit

¹ Use proprietary top

² Use nested query: CREATE VIEW ... AS SELECT ... FROM (SELECT ... FROM ... FETCH FIRST ...) t

³ Requires parenthesis: (?)

⁴ Use proprietary select top ... percent

⁵ Use proprietary select top ... with ties

⁶ Not for 0 (zero)

Functional Dependencies

(T301)

Functional dependencies: only simplest cases

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Base table PRIMARY KEY	✗	✗	✓	✗	✓	✗	✗
Base table UNIQUE	✗	✗	✓	✗	✗	✗	✗
Joined tables	✗	✗	✓	✗	✓ ⁰	✗	✗
WHERE clause	✗	✗	✓	✗	✗	✗	✗
GROUP BY clause	✗	✗	✓	✗	✗	✗	✗

⁰ Not following joins to PRIMARY KEYs or UNIQUE constraints

Docs: unsupported features:

Functional dependencies: only simplest cases

```
SELECT COUNT(*) cnt, t2.b
  FROM t1
INNER JOIN t2 ON (t1.pk = t2.pk)
 GROUP BY t1.pk
```

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Base table PRIMARY KEY	✗	✗	✓	✗	✓	✗	✗
Base table UNIQUE	✗	✗	✓	✗	✗	✗	✗
Joined tables	✗	✗	✓	✗	✗	✗	✗
WHERE clause	✗	✗	✓	✗	✗	✗	✗
GROUP BY clause	✗	✗	✓	✗	✗	✗	✗

⁰ Not following joins to PRIMARY KEYs or UNIQUE constraints

Docs: unsupported features:

Functional dependencies: only simplest cases

4.24	Functional dependencies.....	97
4.24.1	Overview of functional dependency rules and notations.....	97
4.24.2	General rules and definitions.....	98
4.24.3	Known functional dependencies in a base table.....	99
4.24.4	Known functional dependencies in a viewed table.....	99
4.24.5	Known functional dependencies in a transition table.....	100
4.24.6	Known functional dependencies in <table value constructor>.....	100
4.24.7	Known functional dependencies in a <joined table>.....	100
4.24.8	Known functional dependencies in a <table primary>.....	102
4.24.9	Known functional dependencies in a <table factor>.....	103
4.24.10	Known functional dependencies in a <table reference>.....	103
4.24.11	Known functional dependencies in the result of a <from clause>.....	103
4.24.12	Known functional dependencies in the result of a <where clause>.....	104
4.24.13	Known functional dependencies in the result of a <group by clause>.....	104
4.24.14	Known functional dependencies in the result of a <having clause>.....	105
4.24.15	Known functional dependencies in a <query specification>.....	105
4.24.16	Known functional dependencies in a <query expression>.....	106

Functional dependencies: only simplest cases

4.24	Functional dependencies.....	
4.24.1	Overview of functional dependency rules and notation.....	
4.24.2	General rules and definitions.....	
4.24.3	Known functional dependencies in a base table.....	
4.24.4	Known functional dependencies in a viewed table.....	99
4.24.5	Known functional dependencies in a transition table.....	100
4.24.6	Known functional dependencies in <table value constructor>.....	100
4.24.7	Known functional dependencies in a <joined table>.....	100
4.24.8	Known functional dependencies in a <table primary>.....	102
4.24.9	Known functional dependencies in a <table factor>.....	103
4.24.10	Known functional dependencies in a <table reference>.....	103
4.24.11	Known functional dependencies in the result of a <from clause>.....	103
4.24.12	Known functional dependencies in the result of a <where clause>.....	104
4.24.13	Known functional dependencies in the result of a <group by clause>.....	104
4.24.14	Known functional dependencies in the result of a <having clause>.....	105
4.24.15	Known functional dependencies in a <query specification>.....	105
4.24.16	Known functional dependencies in a <query expression>.....	106

Still room for vendor extensions.
e.g. related to **ROW_NUMBER** and
ORDINALITY.

Unsupported features that other DBs have


Row Pattern Recognition

(`match_recognize`)

(R010, R020, R030)

Row Pattern Matching

Since SQL:2016




```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
    FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
  MEASURES
    LAST(ts) - FIRST(ts) AS duration
  ONE ROW PER MATCH
  PATTERN ( any cont* )
  DEFINE cont AS ts < PREV(ts)
                + INTERVAL '30' minute
) t
```

Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Since SQL:2016




```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
    FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
  MEASURES
    LAST(ts) - FIRST(ts) AS duration
  ONE ROW PER MATCH
  PATTERN ( any cont* )
  DEFINE cont AS ts < PREV(ts)
    + INTERVAL '30' minute
```

*define
continued*

Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Since SQL:2016




```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
    FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
  MEASURES
    LAST(ts) - FIRST(ts) AS duration
    ONE ROW PER MATCH
  PATTERN ( any cont* )
  DEFINE cont AS ts < PREV(ts)
                + INTERVAL '30' minute
```

*undefined
pattern variable:
matches any row*

Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Since SQL:2016




```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
  FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
    MEASURES
      LAST(ts) - FIRST(ts) AS duration
      ONE_ROW_PER_MATCH
    PATTERN ( any cont* )
    DEFINE cont AS ts < PREV(ts)
                  + INTERVAL '30' minute
  ) t
```

any number of "cont" rows

Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Since SQL:2016




```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
    FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
  MEASURES
    LAST(ts) - FIRST(ts) AS duration
  ONE ROW PER MATCH
  PATTERN ( any cont* )
  DEFINE cont AS ts < PREV(ts)
                + INTERVAL '30' minute
) t
```

Very much
like GROUP BY

Oracle doesn't support avg on intervals — query doesn't work as shown


Row Pattern Matching

Since SQL:2016



```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
  FROM log
 MATCH_RECOGNIZE(
    ORDER BY ts
    MEASURES
      LAST(ts) - FIRST(ts) AS duration
    ONE ROW PER MATCH
    PATTERN ( any cont* )
    DEFINE cont AS ts < PREV(ts)
           + INTERVAL '30' minute
  ) t
```


*Very much
like SELECT*



Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching


Since SQL:2016



Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Since SQL:2016



```
SELECT COUNT(*) sessions
      , AVG(duration) avg_duration
    FROM log
  MATCH_RECOGNIZE(
    ORDER BY ts
  MEASURES
    LAST(ts) - FIRST(ts) AS duration
  ONE ROW PER MATCH
  PATTERN ( any cont* )
  DEFINE cont AS ts < PREV(ts)
                + INTERVAL '30' minute
) t
```

Oracle doesn't support avg on intervals — query doesn't work as shown

Row Pattern Matching

Endless possibilites

GROUP BY

→ **ONE ROW PER MATCH**

OVER ()

→ **ALL ROWS PER MATCH, FINAL, RUNNING**

HAVING, WHERE

→ **PATTERN** (unmatched, suppressed { - ... - })

Mixing **GROUP BY** and **OVER()**

→ **ALL ROWS PER MATCH** + all-but-one rows suppressed

Data-driven match length

→ **SUM, COUNT, ...** in **DEFINE**

Duplicating rows (to some extend)

→ **ALL ROWS PER MATCH** + **AFTER MATCH SKIP TO ...**


Row pattern matching – match_recognize

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
from clause	✗	✗	✓	✗	✗	✗
window clause	✗	✗	✗	✗	✗	✗
full aggregate support	✗	✗	✗	✗	✗	✗

Free technical report by ISO:
<http://standards.iso.org/ittf/PubliclyAvailableStandards/c065143 ISO IEC TR 19075-5 2016.zip>

Row Pattern Matching

Since SQL:2016



<https://www.slideshare.net/MarkusWinand/row-pattern-matching-in-sql2016>

Stew Ashton has a lot material on this too:

https://stewashton.wordpress.com/category/match_recognize/

Temporal and bi-temporal tables

(T180, T181)

Temporal and bi-temporal tables

First appeared in SQL:2011.

There is an excellent free paper on it:

Temporal features in SQL:2011

<https://sigmodrecord.org/publications/sigmodRecord/1209/pdfs/07.industry.kulkarni.pdf>

If you don't have access to the standard,
this is the next best resource on it.

Temporal and bi-temporal tables

There are two versioning features:

→ System Versioning

Mostly transparent
(done by the system).

Models when changes
happened in the DB.

→ Application Versioning

Managed by the application
(with SQL support).

Can model when changes
happened in the real world.

Both can be applied on per table level as needed.

Temporal and bi-temporal tables

Both require explicit datetime columns and a period:

→ System Versioning

Generated columns

GENERATED ALWAYS

Period name fixed:

SYSTEM_TIME

→ Application Versioning

Arbitrary columns

Arbitrary period names
(but only one per table)

Temporal and bi-temporal tables

System versioning takes care of the DMLs.

→ System Versioning

Datetime columns visible
(not 100% transparent)^[0]

User cannot set them.

Constraints remain
unchanged.

→ Application Versioning

Datetime columns visible

User has to provide values.

Constraints need to
consider periods
(e.g. WITHOUT OVERLAPS).

^[0] Some databases offer invisible or hidden columns for transparency.

Temporal and bi-temporal tables

For queries, they use a different syntax:

→ System Versioning

FROM ...

FOR SYSTEM_TIME
[AS OF | BETWEEN | FROM...TO]

→ Application Versioning

In where clause.

New predicates for periods:
contains, overlaps,
precedes, succeeds,...

Temporal and bi-temporal tables

Recent discussions on -hackers:

→ System Versioning

“AS OF Queries”

Konstantin Knizhnik

Dec 2017 - Jan 2018

→ Application Versioning

“Periods”

Vik Fearing

May 2018

System-versioned tables

(T180)

System-versioned tables

Released
May 25, 2018

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
generate always as row ... ⁰	✓	✓	✗	✗	✗	✓	✗
period for system_time ¹	✓	✓	✗	✗	✗	✓	✗
Add system versioning to table ²	✓	✓	✗	✗	✗	✓	✗
Drop system versioning from table ²	✓	✓	✗	✗	✗	✓	✗
for system_time as of ...	✓	✓	✗	✗	✗	✓	✗
for system_time between ...	✓	✓	✗	✗	✗	✓	✗
for system_time from ...	✓	✓	✗	✗	✗	✓	✗
Immutable transaction time ⁶	✓	✗	✗	✗	✗	✓	✗

Oracle has similar
(yet different) syntax
to access undo data
("flashback").

⁰ Requires row begin instead of row start

¹ Without keyword for (period system_time (...))

² Syntax varies widely

³ Expressions not supported.

⁴ Without between symmetric

⁵ Expressions not supported. Without between symmetric

⁶ Row [start|end] uses statement time, not transaction time.

System-versioned tables

Limitations and gaps in the standard:

- Schema changes are not supported
(Most **ALTER** statements on system-versioned tables fail)
- No functionality for retention
(also: **delete** cannot delete historic rows—GDPR right of erasure ;)
- **FOR SYSTEM_TIME** only works for base tables
(not for views, for example. Also no session setting in the standard).
- Based on “transaction time” (!= commit time)

System-versioned tables

Notes from current implementations:

- **History tables** are most popular
 - Db2 (LUW) and SQL Server use separate tables for old data.
- **Partitions** let the user choose
 - MariaDB 10.3 use a single logical table that can optionally be partitioned so that current and historic data are segregated.
- **Finding history data in UNDO** (data kept for rollback)
 - Oracle uses the UNDO segment to access historic data.
 - Automatic retention, configurable up to 2^{32} seconds (136yrs)^[0].

^[0] Don't know if there is a way to retire selected rows (GDPR)

Application-versioned tables

(T181)

Application-versioned tables – model the real world

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
period for business_time	✓	✗	✗	✓	✗	✗	✗
without overlaps constraint	✓	✗	✗	✗	✗	✗	✗
update ... for portion of	✓	✗	✗	✗	✗	✗	✗
delete ... for portion of	✓	✗	✗	✗	✗	✗	✗

⁰ Use range type.

¹ Use exclusion constraint.

Functionality is available,
only the standard syntax is
missing

“Periods” Patch from
May 26 2018

Period Predicates

(T502)

Period Predicates — like range type operators

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
overlaps ⁰	✓	✗	✗	✗	✗	✗	✗
equals	✗	✗	✗	✗	✗	✗	✗
contains	✗	✗	✗	✗	✗	✗	✗
precedes	✗	✗	✗	✗	✗	✗	✗
succeeds	✗	✗	✗	✗	✗	✗	✗
immediately precedes	✗	✗	✗	✗	✗	✗	✗
immediately succeeds	✗	✗	✗	✗	✗	✗	✗

⁰ Doesn't recognize period names. Use `(start_ts, end_ts)` syntax without keyword period.

¹ Use range type and respective operators.

Period Predicates — like range type operators

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
overlaps	✓ ⁰	✗	✗	✗	✗	✗	✗
equals	✗	✗	✗	✗	✗	✗	✗
contains	✗	✗	✗	✗	✗	✗	✗
precedes	✗	✗	✗	✗	✗	✗	✗
succeeds	✗	✗	✗	✗	✗	✗	✗
immediately precedes	✗	✗	✗	✗	✗	✗	✗
immediately succeeds	✗	✗	✗	✗	✗	✗	✗

Functionality is available,
only the standard syntax is
missing

⁰ Doesn't recognize period names. Use (start_ts, end_ts) syntax without keyword period.

¹ Use range type and respective operators.

Generated Columns

(T175)

Generated Columns

Syntax is shared with system-versioned tables and identity columns.

`generate always as (...)`

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
<code>generate always as (...)</code>	✓	✓ ⁰	✓ ⁰	✓	✗	✓ ¹	✗

⁰ Requires data type declaration.

¹ Requires data type declaration. Without keywords `generated always`.

From standards perspective:

- Generated columns can be used almost like base columns (e.g. in constraint definitions)

Other use cases:

- Function-based indexes (MariaDB/MySQL, SQL Server)

Combined data change and retrieval
Similar to writeable CTEs
(T495)

Combined Data Change and Retrieval

```
INSERT INTO target
SELECT *
  FROM OLD TABLE (DELETE FROM source)
```

```
WITH cte AS (
  DELETE FROM source
  RETURNING *
)
INSERT INTO target
SELECT *
  FROM cte
```

Combined Data Change and Retrieval

```
INSERT INTO demo_t495_c
SELECT *
  FROM OLD TABLE (DELETE FROM demo_t495)
```

Differences to writeable CTEs:

- Three modes: OLD, NEW, FINAL (similar to triggers)
- NEW and FINAL is still before AFTER triggers
- FINAL fails in case the target table is further modified by
 - constraints (cascade)
 - AFTER triggers

Combined Data Change and Retrieval

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
[new final] TABLE (INSERT ...)	✓	✗	✗	✗	✗	✗	✗
[...] TABLE (UPDATE ...)	✓	✗	✗	✗	✗	✗	✗
[old] TABLE (DELETE ...)	✓	✗	✗	✗	✗	✗	✗
[...] TABLE (MERGE ...)	✗	✗	✗	✗	✗	✗	✗
In DML	✗ ₀	✗	✗	✗	✗	✗	✗

⁰ Main statement must be select. Workaround via chained with clause.

Partitioned Join

(not related to partitioned tables)

(F403)


Partitioned Join – Filling gaps in time series

Gap




ts	value
1	...
3	...
4	...
5	...

Partitioned Join – Filling gaps in time series



```
SELECT *
  FROM data
RIGHT JOIN generate_series(..)
    ON ...
```

Partitioned Join – Filling gaps in time series



```
SELECT *
  FROM data
RIGHT JOIN generate_series(..)
    ON ...
```

Partitioned Join – Filling gaps in time series

grp	ts	value
A	1	...
A	3	...
A	4	...
A	5	...
B	2	...
B	4	...
B	5	...

What if you have several time series,
all of them to be padded?

```
SELECT *
  FROM (SELECT DISTINCT grp
        FROM data) dist
  CROSS JOIN LATERAL
  (SELECT *
    FROM data
    RIGHT JOIN generate_series(..)
    ON ...
    AND data.grp = dist.grp
  )
```

Partitioned Join – Filling gaps in time series

grp	ts	value
A	1	...
A	3	...
A	4	...
A	5	...
B	2	...
B	4	...
B	5	...

What if you have several time series,
all of them to be padded?

```
SELECT *
  FROM data PARTITION BY (grp)
  RIGHT JOIN generate_series(...)
```

ON ...

Partitioned Join – Filling gaps in time series

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
LEFT OUTER partitioned join	✗ ⁰	✗ ¹	✗ ¹	✓	✗ ⁰	✗ ⁰	✗ ¹
RIGHT OUTER partitioned join	✗ ⁰	✗ ¹	✗ ¹	✓	✗ ⁰	✗ ⁰	✗
FULL OUTER partitioned join	✗ ¹	✗	✗	✗ ⁰	✗ ⁰	✗ ⁰	✗

⁰ Alternative: Select distinct partition key and join lateral for each partition.

¹ Alternative: join to cross join of distinct partition key and gap-filler.

LISTAGG

Like **STRING_AGG**

(T625)

LISTAGG

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
listagg(... within group (...))	✓	✗	✓	✗	✗	✗
listagg(... on overflow ...)	✗	✗	✓ ⁰	✗	✗	✗
listagg(distinct ...)	✓ ¹	✗	✗	✗	✗	✗
SQLSTATE 22001 on truncation	✗ ²	✗	✗ ³	✗	✗	✗
listagg with grouping sets	✓	✗	✓	✗	✗	✗
listagg... within group... filter...	✗	✗	✗	✗	✗	✗
listagg... within group... over...	✗	✗	✓	✗	✗	✗

⁰ Since 12.2

¹ If ordered by the aggregated values: listagg(distinct x,...) within group (order by x)

² SQLSTATE 54006

³ SQLSTATE 72000

Distinct data types

`CREATE TYPE ... AS <predefined types>`

(S011 - Core SQL)

Distinct Data Types

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
CREATE TYPE..AS <pred. type>	✓	✗	✗	✗	✗	✗	✗

Work in progress

MERGE

(F312, F313, F314)

MERGE — conditional insert/update/delete

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Merge	✓	✗	✗	✓	✗	✓	✗
[NOT] MATCHED AND <condition>	✓	✗	✗	✗ ⁰	✗	✓	✗
Multiple update/insert branches	✓	✗	✗	✗	✗	✗ ₁	✗
WHEN MATCHED DELETE	✓	✗	✗	✗ ₂	✗	✓	✗

⁰ Alternative syntax: WHEN MATCHED (UPDATE|INSERT) ... [WHERE ...]
¹ Two when matched clauses are allowed if one uses update and the other delete.
² Alternative syntax: WHEN MATCHED UPDATE ... [DELETE [WHERE ...]]

MERGE — conditional insert/update/delete

	DB2 LUW	MariaDB	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
Merge	✓	✗	✗	✓	✗	✓	✗
[NOT] MATCHED AND <condition>	✓	✗	✗	✗ ⁰	✗	✓	✗
Multiple update/insert branches	✓	✗	✗	✗	✗	✗ ₁	✗
WHEN MATCHED DELETE	✓	✗	✗	✗ ₂	✗	✓	✗

⁰ Alternative syntax: WHEN MATCHED (UPDATE|INSERT) ... [WHERE ...]
¹ Two when matched clauses are allowed if one uses update and the other delete.
² Alternative syntax: WHEN MATCHED UPDATE ... [DELETE [WHERE ...]]

As of c9c875a
(just before revert)

	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
✓	✓	✓	✓	✗	
✗ ₀	✗	✓	✓	✗	
✗ ₁	✓	✓	✗	✗	
✗ ₂	✗	✗	✓	✓	✗

Multiple update/insert branches	✓	✗	✗	✗	✗
WHEN MATCHED DELETE	✓	✗	✗	✗ ₂	✓

⁰ Alternative syntax: WHEN MATCHED (UPDATE|INSERT) ... [WHERE ...]
¹ Two when matched clauses are allowed if one uses update and the other delete.
² Alternative syntax: WHEN MATCHED UPDATE ... [DELETE [WHERE ...]]

JSON

(T811–T838)

JSON

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

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

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

⁰ No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys].
¹ Defaults to absent on null. No construction by query: json_array(select ...).
² No construction by query: json_array(select ...).
³ Supports comma (,) instead of values or colon (:).
⁴ No colon syntax (T814). No key uniqueness constraint (T830): [with|without] unique [keys]. Sup
⁵ Absent on null is buggy.

JSON

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

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

JSON

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

⁰ Lax mode does not unwrap arrays.
¹ Keyword lax not accepted (only default mode). Lax mode does not unwrap arrays.
² Keyword strict is accepted but not honored.
³ Only in filters. Not supporting size(), datetime(), keyvalue(). type() returns null for arrays.
⁴ Not in json_value. Not in json_query. Not in json_table. Only as last step of expression.
⁵ Not in json_query. Only as last step of expression.

JSON – Preliminary testing of patches

Used 7fe04ce92 as basis, applied those patches on top:

0001-strict-do_to_timestamp-v14.patch

0002-pass-cstring-to-do_to_timestamp-v14.patch

0003-add-to_datetime-v14.patch

0004-jsonpath-v14.patch

0005-jsonpath-gin-v14.patch

0006-jsonpath-json-v14.patch

0007-remove-PG_TRY-in-jsonpath-arithmetics-v14.patch

0010-add-invisible-coercion-form-v13.patch

0011-add-function-formats-v13.patch

0012-sqljson-v13.patch

0013-sqljson-json-v13.patch

0014-json_table-v13.patch

0015-json_table-json-v13.patch } SQL/JSON: JSON_TABLE

SQL/JSON: jsonpath

SQL/JSON: functions

JSON – Preliminary testing of patches

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

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

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

⁰ No colon syntax (T814). No key uniqueness constraint (T830): `[with|without] unique [keys]`.
¹ Defaults to absent on null. No construction by query: `json_array(select ...)`.
² No construction by query: `json_array(select ...)`.
³ Supports comma (,) instead of values or colon (:).
⁴ No colon syntax (T814). No key uniqueness constraint (T830): `[with|without] unique [keys]`. Sup
⁵ Absent on null is buggy.

JSON – Preliminary testing of patches

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

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

JSON – Preliminary testing of patches

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

⁰ Lax mode does not unwrap arrays.
¹ Keyword `lax` not accepted (only default mode). Lax mode does not unwrap arrays.
² Keyword `strict` is accepted but not honored.
³ Only in filters. Not supporting `size()`, `datetime()`, `keyvalue()`, `type()` returns `null` for arrays.
⁴ Not in `json_value`. Not in `json_query`. Not in `json_table`. Only as last step of expression.
⁵ Not in `json_value`. Not in `json_query`.
⁶ Not in `json_query`. Only as last step of expression.
⁷ Not in `json_value`.

Standard SQL Gap Analysis

Incomplete or “wrong”:

- ▶ `extract` (declared type)
- ▶ `ignore nulls`
- ▶ `agg(distinct) over()`
- ▶ `fetch...percent, with ties`
- ▶ Functional dependencies

Missing

- ▶ Row pattern recognition
- ▶ Temporal tables
- ▶ Generated Columns
- ▶ Combined data change and retrieval
- ▶ Partitioned join
- ▶ `listagg`
- ▶ Distinct data types
- ▶ ... (this list is not exhaustive)

Work in progress

- ▶ `merge`
- ▶ `JSON`

How can I help?

- I publish an article on each new version once it is released (pretty late for helpful feedback)
- I start preparing for this article once a public beta is available (but it is often pushed by higher priority tasks -> no guarantee)
- I do not monitor -hackers, but Depesz's “waiting for”
(This is typically the first time I notice a new feature is coming up)
- If you have questions on the standard or would like to get conformance test results at a earlier stage, ping me.

Twitter: @MarkusWinand – markus.winand@winand.at