



Range Types: Your Life Will Never Be The Same

Jonathan S. Katz
CTO, VenueBook
October 24, 2012

What's in a Range?

- Conference schedule
- Pick a number from 1-10
 - Integer or real?
- Budget for buying a new laptop

Ranges are Everywhere...

- Scheduling
- Probability
- Measurements
- Financial applications
- Clinical trial data
- Intersections of ordered data

How Do We Deal With Ranges?

```
CREATE TABLE employee_schedule (  
    id serial,  
    employee_id integer REFERENCES  
    employees(id),  
    start_time timestamptz,  
    end_time timestamptz  
);
```

Who is on duty at...

```
SELECT *
FROM employee_schedule
WHERE
    employee_id = 24 AND
    CURRENT_TIMESTAMP BETWEEN start_time AND end_time;

-- start_time <= CURRENT_TIMESTAMP <= end_time
```

Can I schedule an employee shift?

- Easy!

```
SELECT EXISTS(id)
FROM employee_schedule
WHERE
    employee_id = 24 AND
    (
        '2012-09-18 10:00' <= start_time AND
        '2012-09-18 11:00' >= start_time AND
        '2012-09-18 10:00' <= end_time AND
        '2012-09-18 11:00' <= end_time
    ) OR ( --...wait this is really hard
```

Why Overlaps Are Difficult



In PostgreSQL 9.1, Can I...

- Use a built-in function to determine if my ranges overlap?
- Easily create a composite type and add logic to recognize the ranges?
- Change to a different database software that makes the problem easier?

...can someone smarter than me
make my life easier?


...Yes!!!

[projects](#) / [postgresql.git](#) / [commit](#)

[summary](#) | [shortlog](#) | [log](#) | [commit](#) | [commitdiff](#) | [tree](#)
(parent: [4334289](#)) | [patch](#)

Support range data types.

```
author      Heikki Linnakangas <heikki.linnakangas@iki.fi>  
            Thu, 3 Nov 2011 11:16:28 +0000 (13:16 +0200)  
committer   Heikki Linnakangas <heikki.linnakangas@iki.fi>  
            Thu, 3 Nov 2011 11:42:15 +0000 (13:42 +0200)  
commit      4429f6a9e3e12bb4af6e3677fbc78cd80f160252  
tree        a2e272129e5515f7ef2f4e09989bddf0fd8158ea      tree | snapshot  
parent      43342891861cc2d08dea2b1c8b190e15e5a36551      commit | diff
```

Support range data types. 

Selectivity estimation functions are missing for some range type operators, which is a TODO.

Jeff Davis

Built-In Ranges

- `INT4RANGE` (integer)
- `INT8RANGE` (bigint)
- `NUMRANGE` (numeric)
- `TSRANGE` (timestamp without time zone)
- `TSTZRANGE` (timestamp with time zone)
- `DATERANGE` (date)

Range Bounds

- Ranges can be inclusive, exclusive or both
- Math review:
 - $[2, 4] \Rightarrow 2 \leq x \leq 4$
 - $[2, 4) \Rightarrow 2 \leq x < 4$
 - $(2, 4] \Rightarrow 2 < x \leq 4$
 - $(2, 4) \Rightarrow 2 < x < 4$
- Can also be empty

Ranges...Unbound

- Ranges can be infinite
 - $[2,) \Rightarrow 2 \leq x < \infty$
 - $(, 2] \Rightarrow -\infty < x \leq 2$
- **CAVEAT EMPTOR**
 - “infinity” has special meaning with timestamp ranges
 - $[CURRENT_TIMESTAMP,) = [CURRENT_TIMESTAMP,]$
 - $[CURRENT_TIMESTAMP, 'infinity') \langle \rangle [CURRENT_TIMEAMP, 'infinity']$

Constructing Ranges

- Simple!

```
test=# SELECT ' [1,10] ' :: int4range;
```

```
int4range
```

```
-----
```

```
[1,11)
```

```
(1 row)
```

Constructing Ranges

```
test=# SELECT '[2012-03-28, 2012-04-02]'::daterange;
```

```
      daterange
```

```
-----
```

```
[2012-03-28,2012-04-03)
```

```
(1 row)
```

Constructing Ranges

- Constructor functions too
 - Defaults to '['

```
test=# SELECT numrange (9.0, 9.5);
```

```
numrange
```

```
-----
```

```
[9.0, 9.5)
```

```
(1 row)
```


Constructing Ranges

```
test=# SELECT tsrange('2012-04-01 00:00:00', '2012-04-01
12:00:00', '[]');
```

tsrange

```
-----
["2012-04-01 00:00:00", "2012-04-01 12:00:00"]
```

```
(1 row)
```

Just For Oleg...

- Can have arrays of ranges

```
test=# SELECT ARRAY[int4range(1,3),  
    int4range(2,4), int4range(3,8)];  
          array
```

```
-----  
{ "[1, 3)", "[2, 4)", "[3, 8)" }
```

Using Ranges

- Normal comparison operations

```
SELECT int4range(100,200) = int4range(100,200);
```

```
-- true
```

```
SELECT int4range(100,200) <> int4range(200,300);
```

```
-- true
```

```
SELECT int4range(100,200) < int4range(200,300);
```

```
-- true
```

```
SELECT int4range(100,200) <= int4range(200,300);
```

```
-- true
```

```
SELECT int4range(100,200) >= int4range(200,300);
```

```
-- false
```

```
SELECT int4range(100,200) > int4range(200,300);
```

```
-- false
```

Why Your Life Will Change

- Let's see the magic with an example
- Shopping for a used car
 - Cars listed with a price range
 - Have a min/max budget

Inspect Our Data

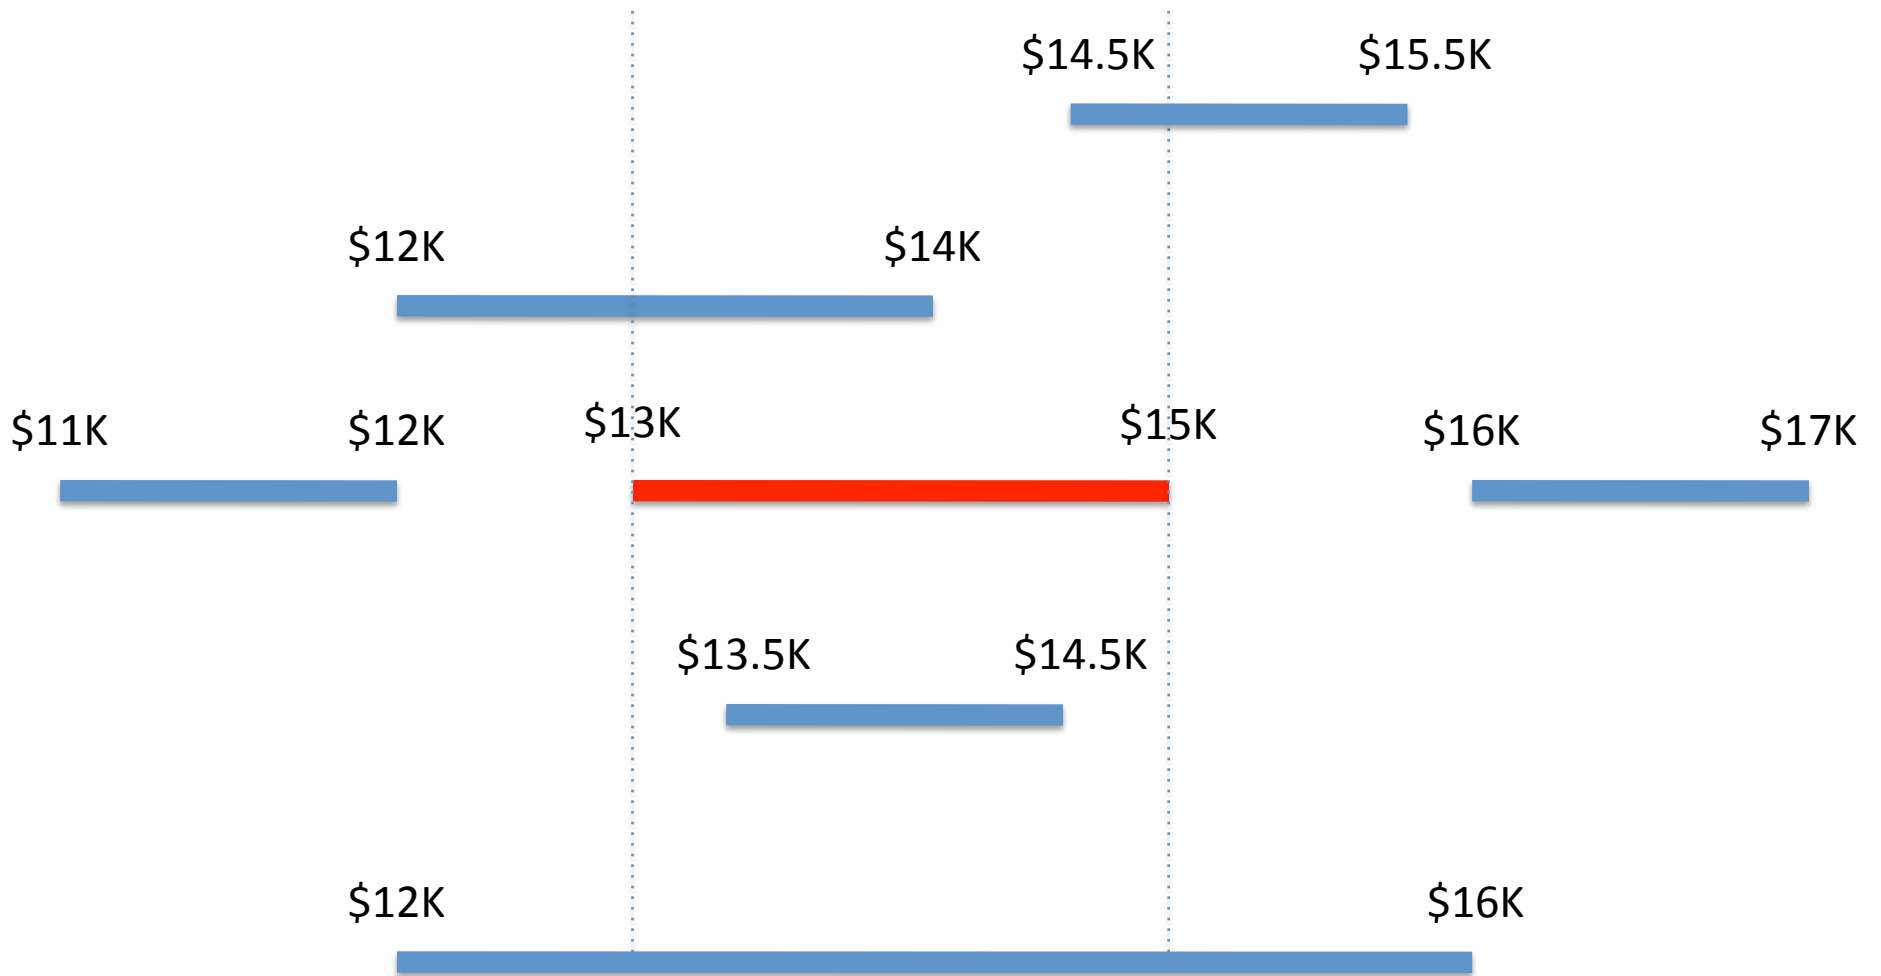
- Sort by range lower bound

```
test=# SELECT * FROM cars ORDER BY lower(cars.price_range) ;
```

id	name	price_range
2	Buick Skylark	[2000,4001)
3	Pontiac GTO	[5000,7501)
4	Chevrolet Camaro	[10000,12001)
5	Ford Mustang	[11000,15001)
6	Lincoln Continental	[12000,14001)
7	BMW M3	[35000,42001)
8	Audi RS4	[41000,45001)
9	Porsche 911	[47000,58001)
10	Lamborghini LP700	[385000,400001)

(9 rows)

Car Shopping: Conceptually Simple



Car Shopping: Nightmarishly Complicated

- Budget of \$13,000 - \$15,000, find cars price in that range

```
SELECT *
FROM cars
WHERE
  (
    cars.min_price ≤ 13000 AND
    cars.min_price ≤ 15000 AND
    cars.max_price ≥ 13000 AND
    cars.max_price ≤ 15000
  ) OR
  (
    cars.min_price ≤ 13000 AND
    cars.min_price ≤ 15000 AND
    cars.max_price ≥ 13000 AND
    cars.max_price ≥ 15000
  ) OR
  (
    cars.min_price ≥ 13000 AND
    cars.min_price ≤ 15000 AND
    cars.max_price ≥ 13000 AND
    cars.max_price ≤ 15000
  ) OR
  (
    cars.min_price ≥ 13000 AND
    cars.min_price ≤ 15000 AND
    cars.max_price ≥ 13000 AND
    cars.max_price ≥ 15000
  )
ORDER BY cars.min_price;
```

Car Shopping: Magically Painless

- Budget of \$13,000 - \$15,000, find cars price in that range

```
SELECT *
FROM cars
WHERE cars.price_range && int4range(13000, 15000, '[')
ORDER BY lower(cars.price_range);
```

id	name	price_range
5	Ford Mustang	[11000,15001)
6	Lincoln Continental	[12000,14001)

(2 rows)

In more details

- **&&**

- the “overlap” operator
- take two ranges: $[x,y]$ and $[a,b]$

$(a \leq x \text{ AND } a \leq y \text{ AND } b \geq x \text{ AND } b \leq y) \text{ OR}$

$(a \leq x \text{ AND } a \leq y \text{ AND } b \geq x \text{ AND } b \geq y) \text{ OR}$

$(a \geq x \text{ AND } a \leq y \text{ AND } b \geq x \text{ AND } b \leq y) \text{ OR}$

$(a \geq x \text{ AND } a \leq y \text{ AND } b \geq x \text{ AND } b \geq y)$

(Math for the win: inverse only two lines)

The Saver

- Find cars whose price does not exceed \$13,000

```
SELECT *
FROM cars
WHERE cars.price_range << int4range(13000, 15000)
ORDER BY lower(cars.price_range);
```

id	name	price_range
2	Buick Skylark	[2000,4001)
3	Pontiac GTO	[5000,7501)
4	Chevrolet Camero	[10000,12001)

The Cautious

- Budget of \$13,000 - \$15,000, but want to see cheaper options

```
SELECT *  
FROM cars  
WHERE cars.price_range <& int4range(13000, 15000)  
ORDER BY lower(cars.price_range);
```

id	name	price_range
2	Buick Skylark	[2000,4001)
3	Pontiac GTO	[5000,7501)
4	Chevrolet Camaro	[10000,12001)
5	Ford Mustang	[11000,15001)
6	Lincoln Continental	[12000,14001)

(5 rows)

The Dreamer

- Budget of \$13,000 - \$15,000, but want to see what lies beyond...

```
SELECT *  
FROM cars  
WHERE cars.price_range >> int4range(13000, 15000)  
ORDER BY lower(cars.price_range);
```

id	name	price_range
7	BMW M3	[35000,42001)
8	Audi RS4	[41000,45001)
9	Porsche 911	[47000,58001)
10	Lamborghini LP700	[385000,400001)

(4 rows)

Determine Negotiating Window

- For cars in my budget, what prices am I looking at?

```
SELECT *,
       cars.price_range * int4range(13000, 15000) AS price_window
FROM cars
WHERE cars.price_range && int4range(13000, 15000)
ORDER BY lower(cars.price_range);
```

id	name	price_range	price_window
5	Ford Mustang	[11000,15001)	[13000,15000)
6	Lincoln Continental	[12000,14001)	[13000,14001)

(2 rows)

Are Range Queries Fast?

- QUERY PLAN
-

```
Sort (cost=11.76..11.77 rows=1 width=552)
```

```
Sort Key: (lower(price_range))
```

```
-> Seq Scan on cars (cost=0.00..11.75 rows=1 width=552)
```

```
Filter: (price_range && '[13000,15001)')::int4range)
```

Range Indexes

- Creating a GiST index on ranges speeds up queries with these operators:

=

&&

<@

@>

<<

>>

-|-

&<

&>

Range Indexes

```
CREATE INDEX cars_price_range_idx ON cars USING gist (price_range);
```

```
-- EXPLAIN $PREVIOUS_QUERY
```

```
QUERY PLAN
```

```
-----  
-----  
Sort  (cost=129.66..129.87 rows=84 width=49)  
  Sort Key: (lower(price_range))  
    -> Bitmap Heap Scan on cars2 (cost=4.95..126.97 rows=84 width=49)  
      Recheck Cond: (price_range && '[13000,15000)')::int4range)  
        -> Bitmap Index Scan on cars2_price_range_idx  
          (cost=0.00..4.93 rows=84 width=0)  
            Index Cond: (price_range && '[13000,15000)')::int4range)  
(6 rows)
```

- Note: I used a more populous table to make the index scan to occur

But How Fast Is It?

```
test=# SELECT count(*) FROM ranges;
```

```
count
```

```
-----
```

```
2000000
```

```
(1 row)
```

Well?

```
test=# EXPLAIN ANALYZE SELECT * FROM ranges WHERE  
int4range(500,1000) && bounds;
```

```
QUERY PLAN
```

```
-----
```

```
Bitmap Heap Scan on ranges
```

```
(actual time=0.283..0.370 rows=653 loops=1)
```

```
  Recheck Cond: ('[500,1000)')::int4range && bounds)
```

```
    -> Bitmap Index Scan on ranges_bounds_gist_idx
```

```
        (actual time=0.275..0.275 rows=653 loops=1)
```

```
          Index Cond: ('[500,1000)')::int4range &&
```

```
bounds)
```

```
Total runtime: 0.435 ms
```

What If the Range is Larger?

```
test=# EXPLAIN ANALYZE SELECT * FROM ranges WHERE
int4range(20000,100000) && bounds;
```

```
QUERY PLAN
```

```
-----
Bitmap Heap Scan on ranges
  (actual time=19.285..26.914 rows=86096 loops=1)
    Recheck Cond: ('[20000,100000)')::int4range &&
bounds)
  -> Bitmap Index Scan on ranges_bounds_gist_idx
      (actual time=19.212..19.212 rows=86096 loops=1)
        Index Cond: ('[20000,100000)')::int4range &&
bounds)
Total runtime: 30.749 ms
```

What If the Range is Much Larger?

```
test=# EXPLAIN ANALYZE SELECT * FROM ranges WHERE
int4range(10000,1000000) && bounds;
```

```
QUERY PLAN
```

```
-----
```

```
Bitmap Heap Scan on ranges
```

```
(actual time=184.028..270.323 rows=993068 loops=1)
```

```
Recheck Cond: ('[10000,1000000)')::int4range &&
bounds)
```

```
-> Bitmap Index Scan on ranges_bounds_gist_idx
```

```
(actual time=183.060..183.060 rows=993068 loops=1)
```

```
Index Cond: ('[10000,1000000)')::int4range &&
bounds)
```

```
Total runtime: 313.743 ms
```

Scheduling

- ...now is super easy*
- Unique constraints to save the day!

Scheduling

```
CREATE TABLE travel_log (  
    id serial PRIMARY KEY,  
    name varchar(255),  
    travel_range daterange,  
    EXCLUDE USING gist (travel_range WITH &&)  
);  
  
INSERT INTO travel_log (name, trip_range) VALUES  
    ('Boston', daterange('2012-03-07', '2012-03-09'));  
INSERT INTO travel_log (name, trip_range) VALUES  
    ('Chicago', daterange('2012-03-12', '2012-03-17'));
```

Scheduling

```
test=# INSERT INTO travel_log (name, trip_range)
VALUES ('Austin', daterange('2012-03-16',
'2012-03-18'));
```

```
ERROR:  conflicting key value violates exclusion
constraint "travel_log_trip_range_excl"
```

```
DETAIL:  Key (trip_range)=([2012-03-16,2012-03-18])
conflicts with existing key
(trip_range)=([2012-03-12,2012-03-17]).
```

- Easy.

And That's Not All!

- Ranges can be extended – I kid you not

```
CREATE TYPE inetrange AS RANGE (  
    SUBTYPE = inet  
);
```


interange

```
SELECT '192.168.1.8'::inet <@  
       inetrange('192.168.1.1', '192.168.1.10');
```

?column?

t

```
SELECT '192.168.1.20'::inet <@  
       inetrange('192.168.1.1', '192.168.1.10');
```

?column?

f

In the Wild?

Home > Venue Search

Type of Event
Choose Event Type

Date
9 18 2012

Guests Budget
40 2000

Filter Search Results


Advanced Search

- Neighborhoods ✓
- Type of Venue ✓
- Venue Features ✓

Search Results

Found 214 Venues


1 2 3 4 5 6 7 8 9 10 11 12 13 14 »



Village Pourhouse - Downtown
\$0 – \$3,840

Private Rooms: Yes
Total Capacity: 200
Location: SoHo, TriBeCa
Type: Bar, Restaurant

[View Venue](#) [Save Venue!](#)



Ben and Jack's Steakhouse-44th
\$1,600 – \$5,760

Private Rooms: Yes
Total Capacity: 250
Location: Midtown East
Type: Restaurant

[View Venue](#) [Save Venue!](#)

For More Information

- <http://www.postgresql.org/docs/9.2/static/rangetypes.html>
- <http://www.postgresql.org/docs/9.2/static/functions-range.html>
- <http://www.postgresql.org/docs/9.2/static/sql-createtype.html>
- <http://wiki.postgresql.org/wiki/RangeTypes>

Conclusion

- If you are not completely smitten by range types, then I have failed at explaining them
- Upgrade to PostgreSQL 9.2.x – now.
 - (currently 9.2.1)

Thanks To...

- Jeff Davis for implementing range types
- Alexander Korotkov for GiST improvements for handling range type data

Contact

- Jonathan S. Katz
- jonathan@venuebook.com
- @jkatz05