

# Dissimilarity search: implementing in-memory vector search algorithms for PostgreSQL

**Jonathan Katz**

(he/him)

Principal Product Manager – Technical  
AWS








$[0.5, 0.5]$

Magnitude



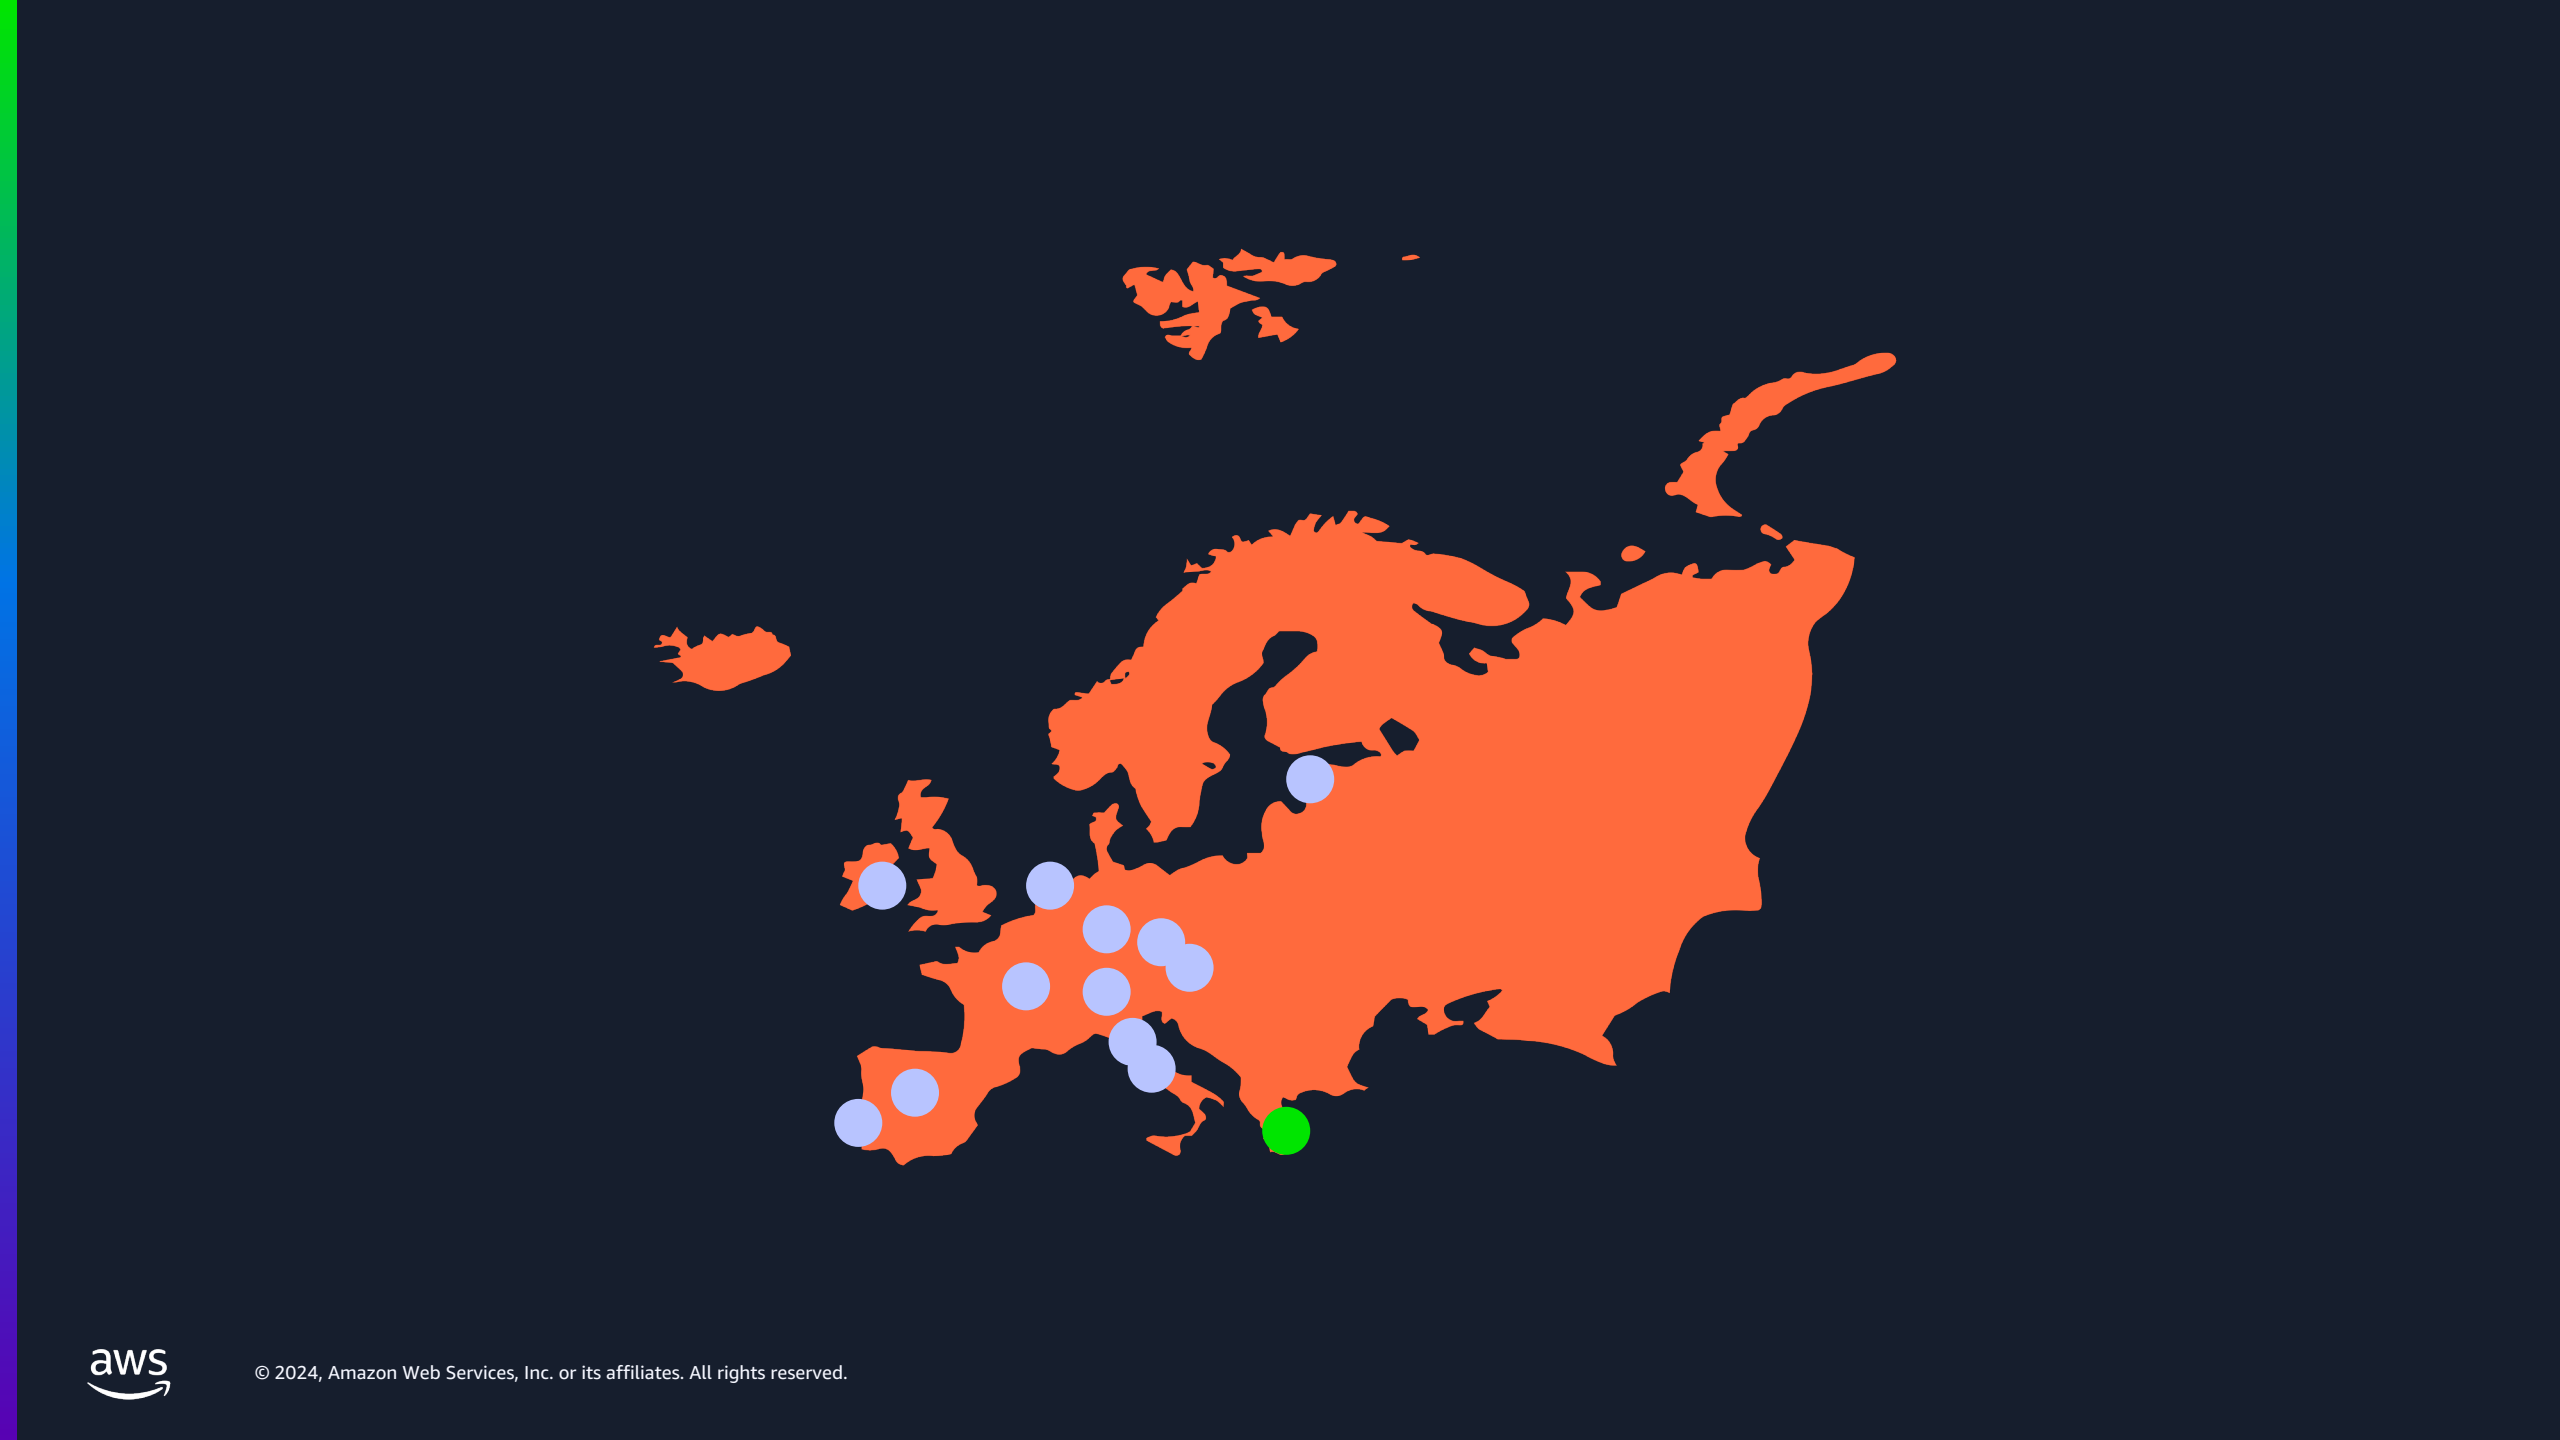
$$\| [0.5, 0.5] \| = \sqrt{0.5^2 + 0.5^2} = \mathbf{0.70710}$$

Magnitude

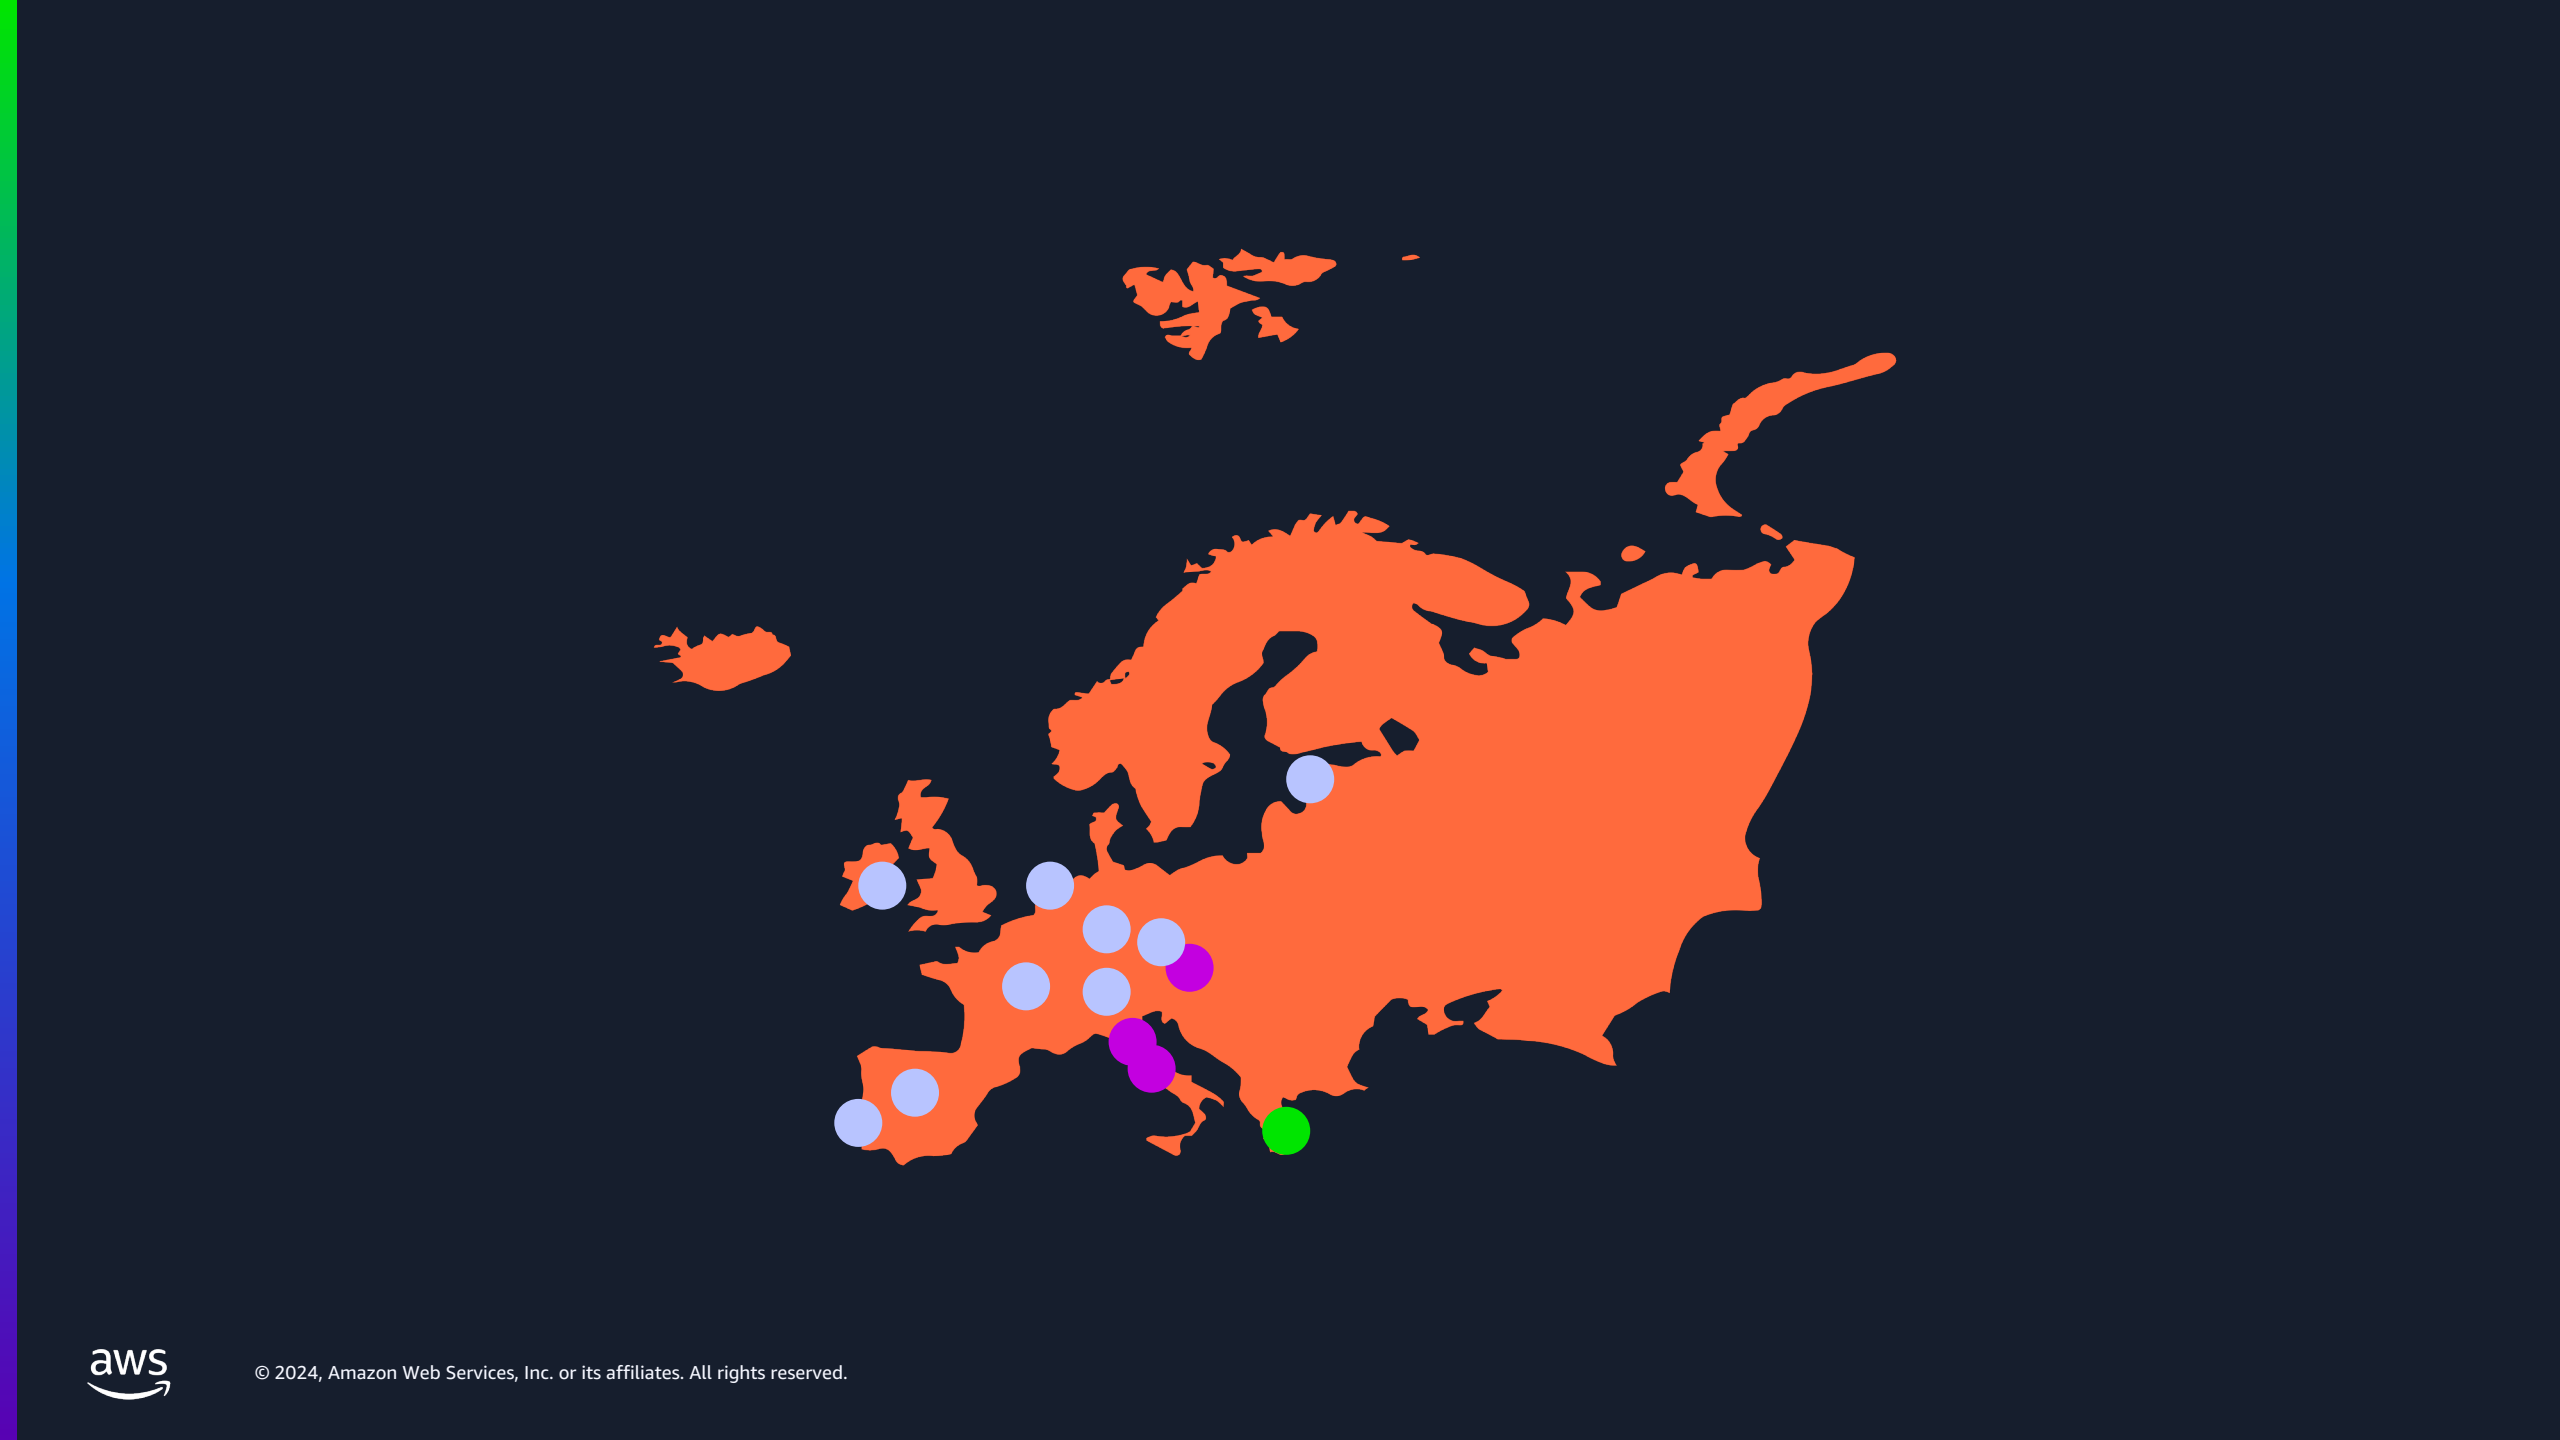


Direction

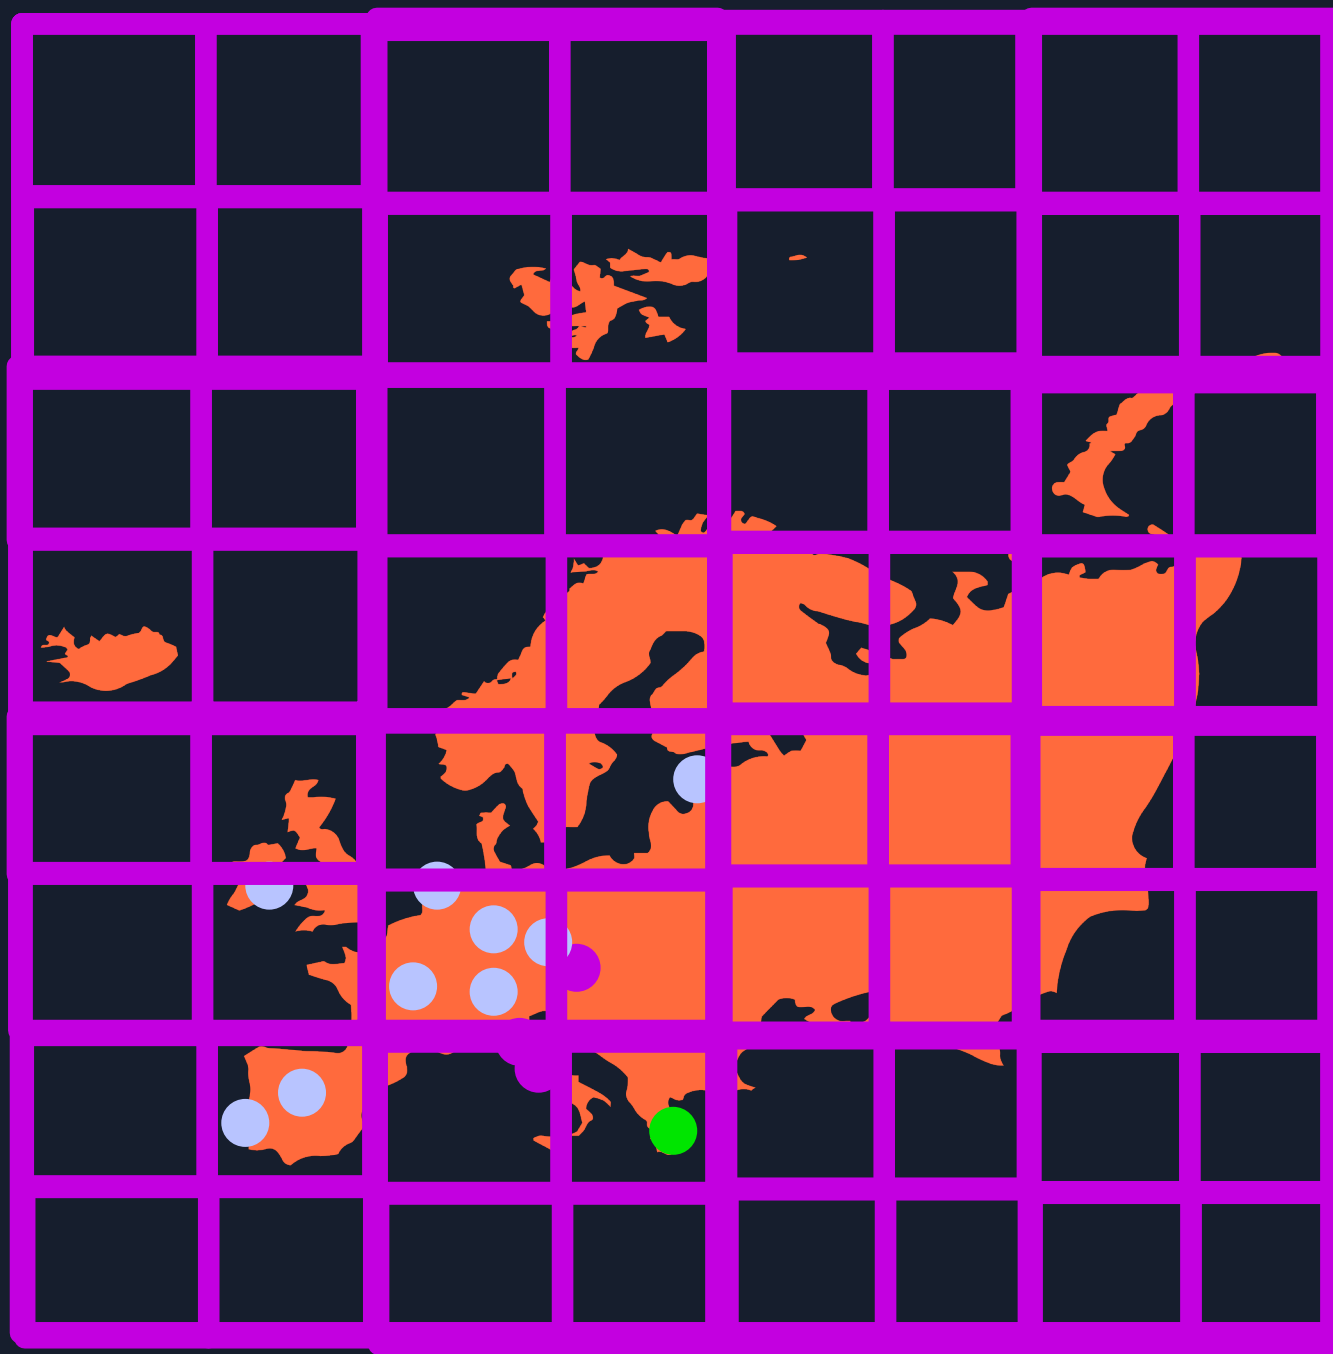
$[0.5, 0.5]$



```
SELECT city_name
FROM conferences
WHERE conference_name LIKE 'PGConf EU%'
ORDER BY
    conference.geocode <-> '(38.0004,23.7195)::point
LIMIT 3;
```








# Foundations of vector search

- Vectors in "vector space" (search area) must all have the same number of dimensions
  - Each dimension should be comparable to each other
- Distance function defines proximity
  - Distance is always  $\geq 0$
  - Distance from a vector to itself is 0



### CUSTOMER

Can you find 5 distributors of green olives near PGConf.EU 2024?

### DEVELOPER CREATED AGENT

Yes, here are a list of distributors based on proximity...



Tags

**Human:** You are an agent who manages orders and returns orders by executing the set of APIs in order to fulfill user input.

Emphasis (capitalized)

Valid "api" values are GetOrderHistory::GetProductCatalogue, GetOrderHistory::GetProductCatalogue, GetOrderHistory::GetProductCatalogue  
- **DO NOT** return an api if all required parameter values are not provided  
- **DO NOT** replace the placeholders in the api\_name with api\_inputs  
- Return available parameters in api\_inputs ONLY.

Valid "verb" is HTTP verb used in "APIs" e.g. GET, PUT etc.

Valid "api\_input" as json from "User Input", "Observation" or "Output"  
- **NEVER** assume value for any parameter, mark the value as "null"

Convergence criteria

**DO NOT** go into a loop and return exact same apis with same parameters

Format (JSON)

Provide only ONE action per \$JSON\_BLOB, as shown:

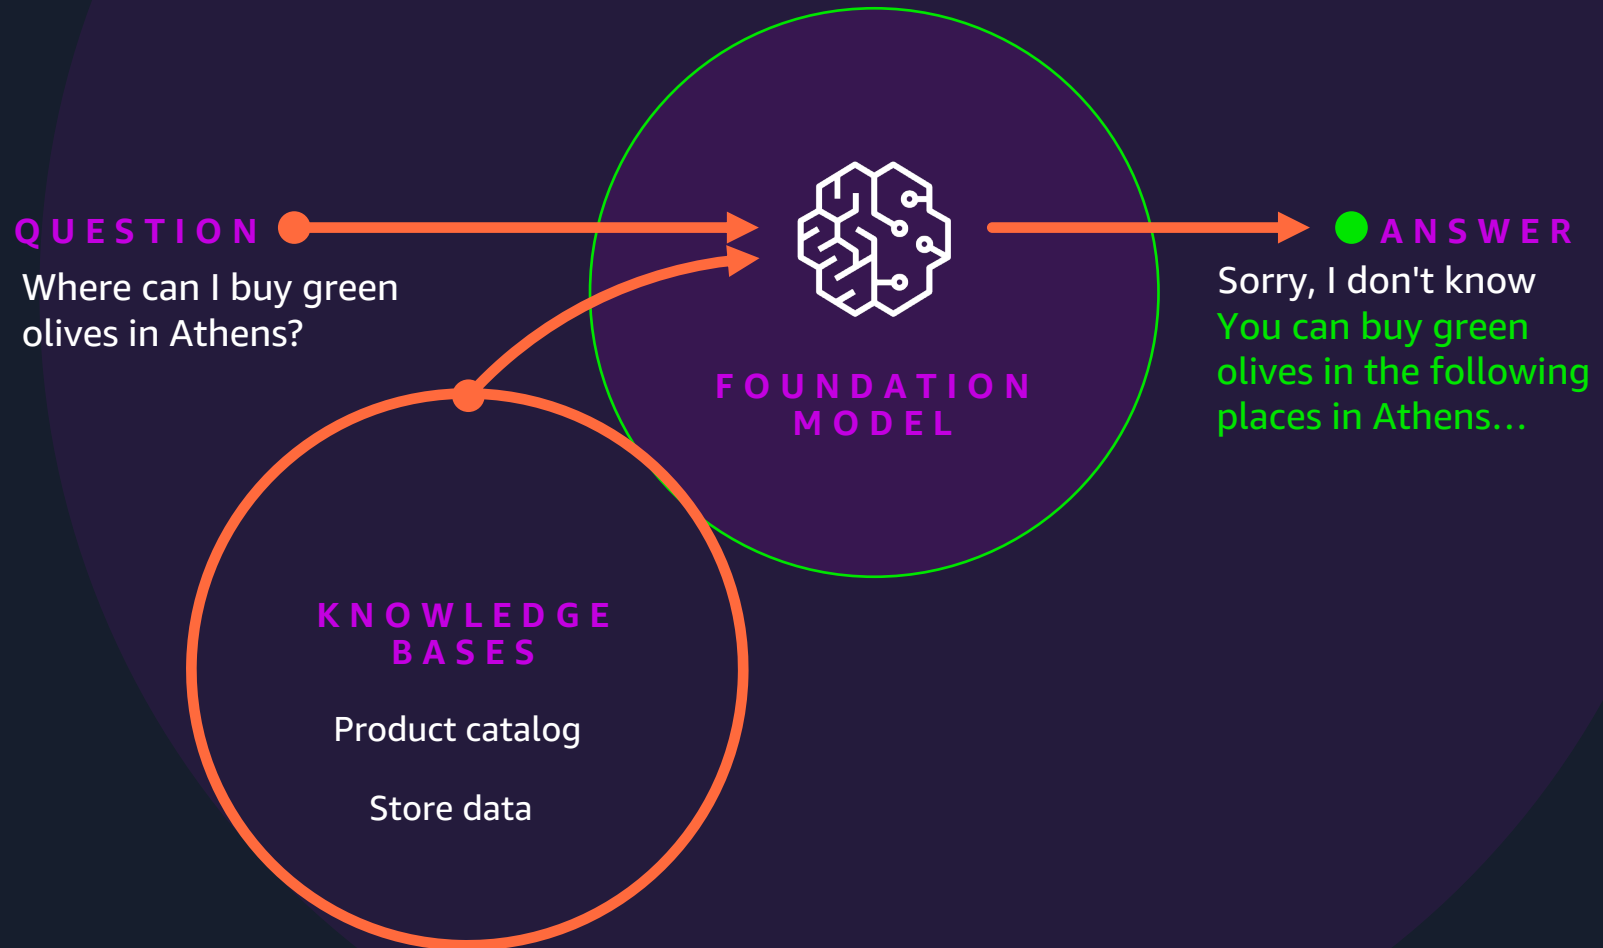
```
{ "api": "$API_NAME", "verb": "$HTTP_VERB", "api_input": { "$PARAM": "$PARAM_VALUE" } }
```

History format

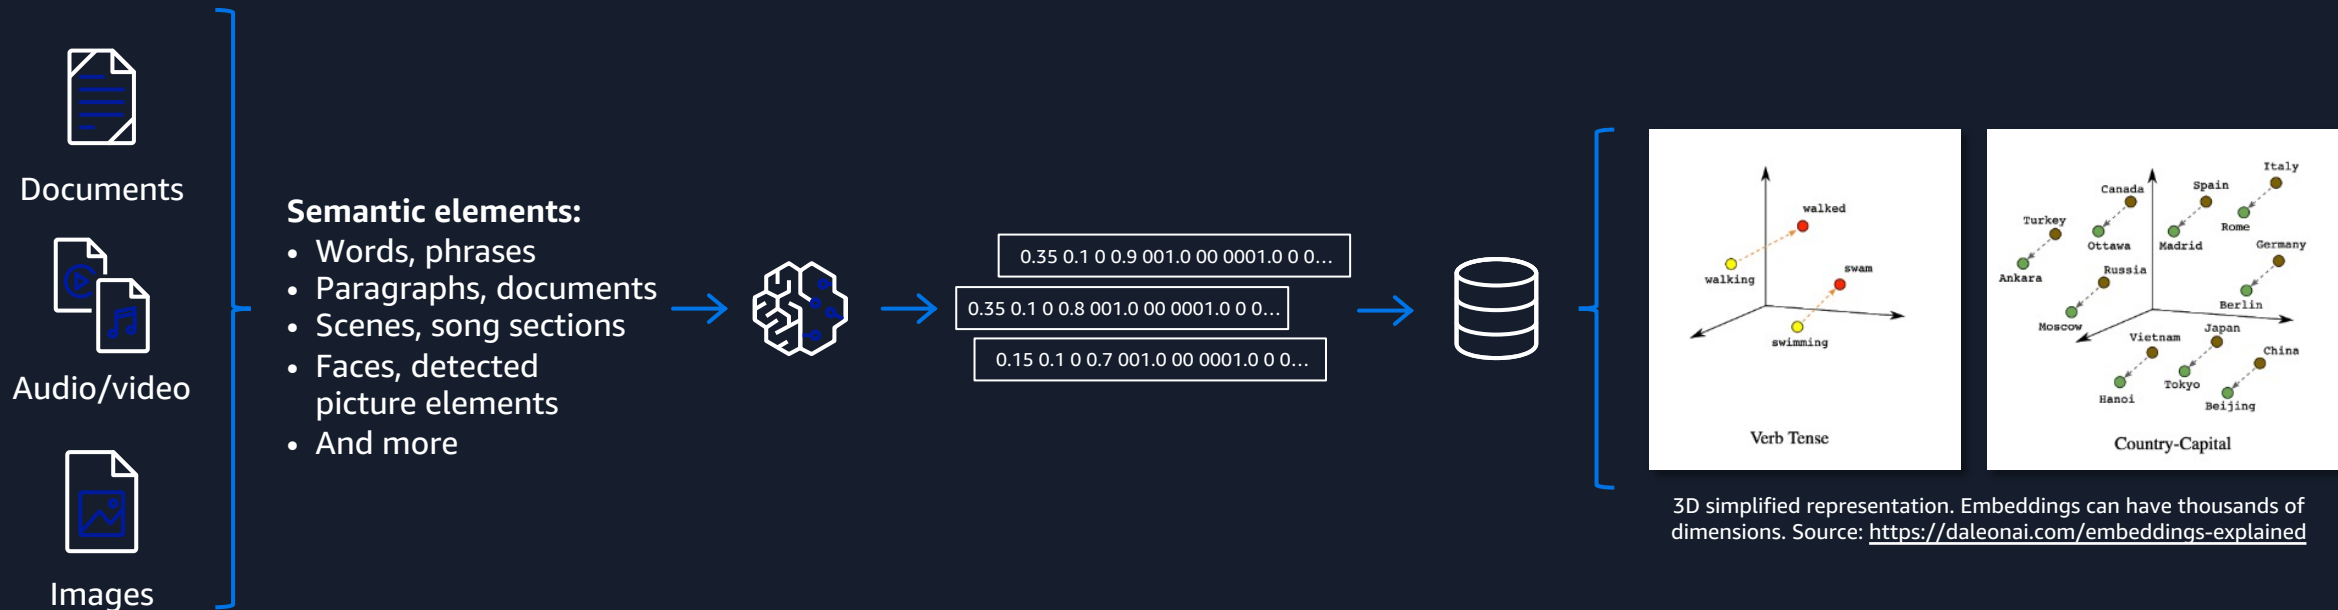
**Conversation History:** Below is the history of the conversation:

# Retrieval-augmented generation (RAG)

Configure foundation model to interact with your data



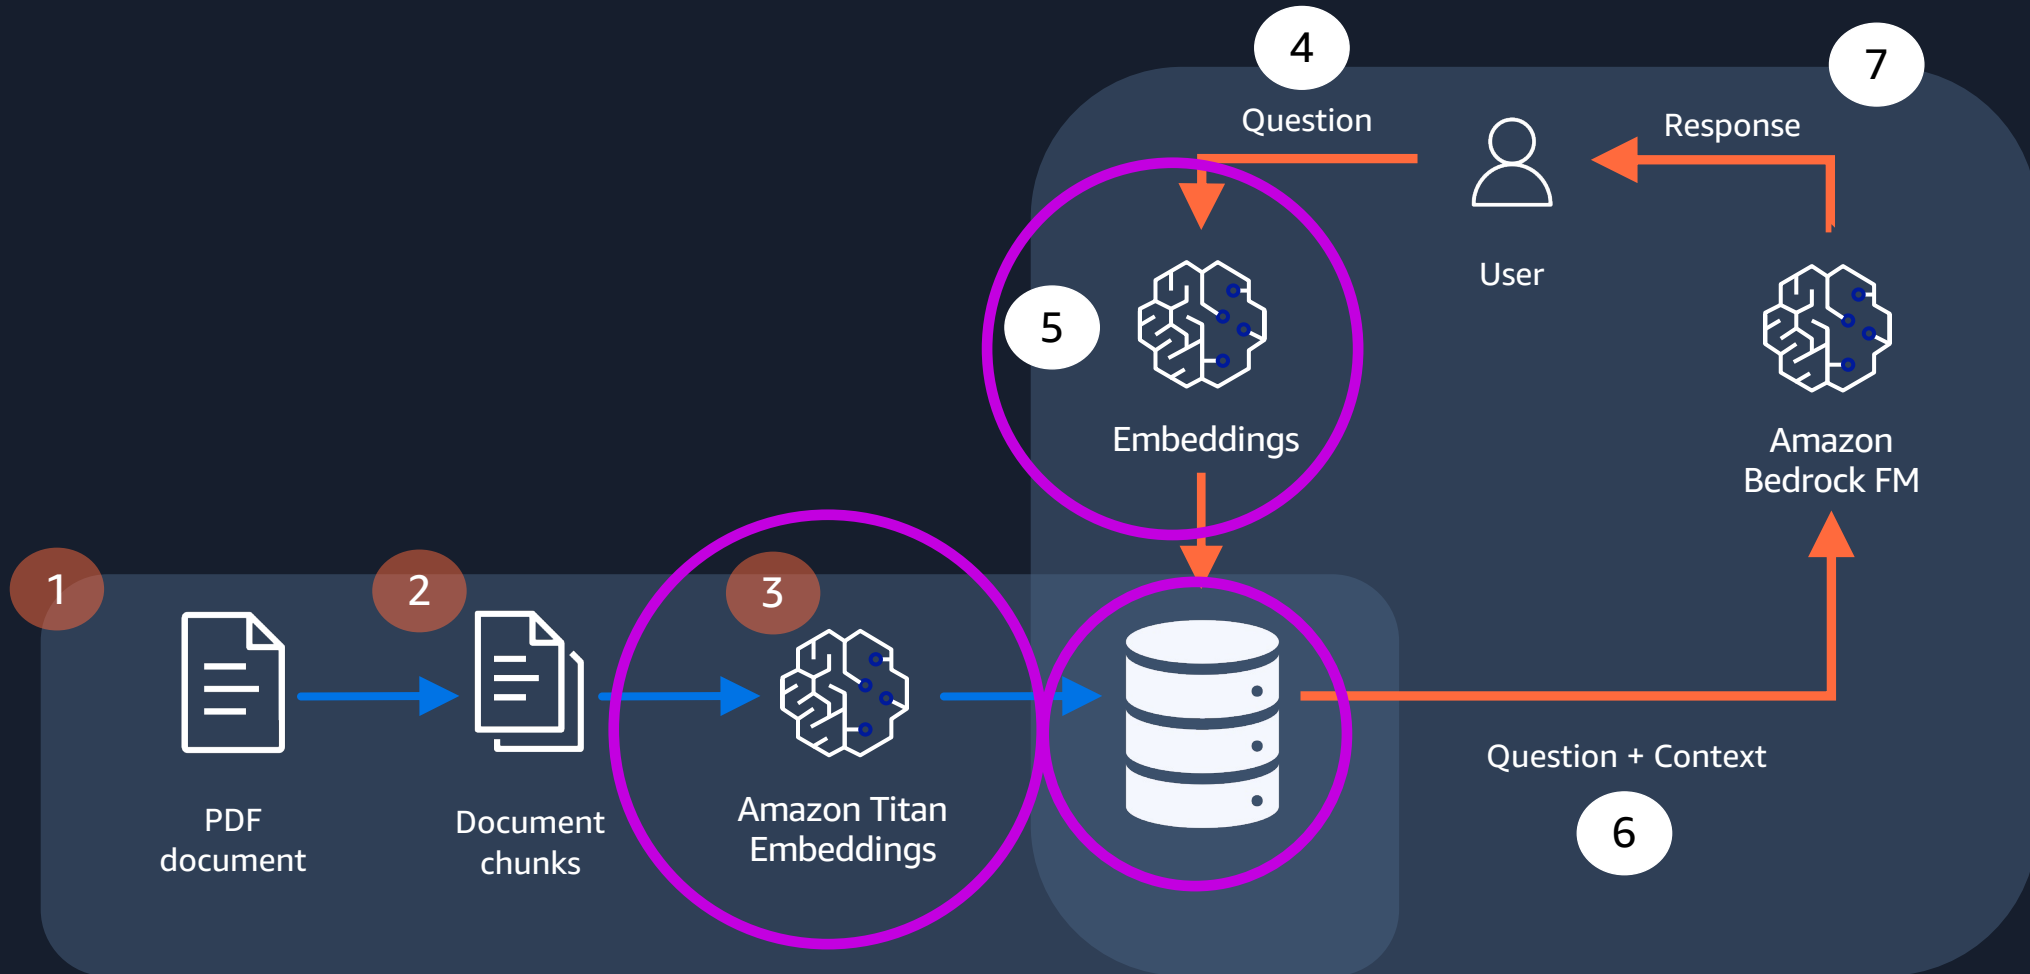
# What are embeddings?



3D simplified representation. Embeddings can have thousands of dimensions. Source: <https://daleonai.com/embeddings-explained>

**Embeddings:** When vector elements are semantic, used in generative AI

# How embeddings are used



# Challenges with larger vectors

- Generation time
- Size
- ~~Compression~~
- Query time



# Approximate nearest neighbor (ANN)

- Find similar vectors without searching all of them
- Faster than exact nearest neighbor
- “Recall” – % of expected results



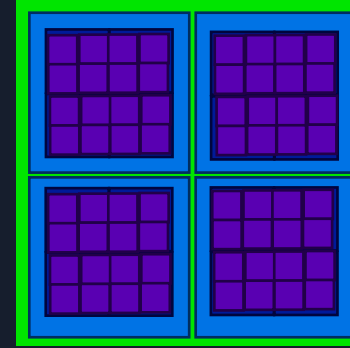
Recall: 80%




# ANN indexing algorithm types




Hash



Tree



Cluster




Graph


# Index layout in memory




# Index layout in a database



# Index layout in a database



# Index size exceeds available memory



# Index size exceeds available memory



# Index size exceeds available memory




# Where memory and storage diverge

- Continuous allocations vs. pages
- Data layout on disk
- Percentage of index in memory
- Hardware acceleration strategies (CPU vs. GPU)



# Vector search design principles for PostgreSQL

 Take shortcuts, where applicable

 Design for 8KiB blocks (page size)

 Leverage PostgreSQL infrastructure

 Understand your tradeoffs

# What is pgvector?

Adds support for **storage, indexing, searching, metadata** with choice of **distance**

**vector** data type

Co-locate with embeddings

Exact nearest neighbor (K-NN)  
Approximate nearest neighbor (ANN)

Supports **HNSW & IVFFlat** indexing, with options for **scalar and binary quantization**

Distance operations include **Cosine, Euclidean/L2, Manhattan/L1, Dot product, Hamming, Jaccard**

[github.com/pgvector/pgvector](https://github.com/pgvector/pgvector)



# Example pgvector query

```
SET hnsw.ef_search TO 60;
```

```
SELECT id, text_chunk
```

```
FROM documents
```

```
ORDER BY
```

```
embedding <=> '[0.003421, -0.23053, 0.402153, ...]'::vector
```

```
LIMIT 10
```

# What do we need to define?

1. Data type
2. Distance functions and operators
3. Indexing strategy

# What do we need to define?

1. Data type
2. Distance functions and operators
3. Indexing strategy

# Data type

```
typedef struct Vector
{
    int32      vl_len_; /* varlena header (do not touch directly) */
    int16      dim;     /* number of dimensions */
    int16      unused;  /* reserved for future use, always zero */
    float      x[FLEXIBLE_ARRAY_MEMBER];
} Vector;
```

# PostgreSQL Infrastructure: TOAST

- TOAST (The Oversized-Atttribute Storage Technique) is a mechanism for storing data larger than 8KB
  - By default, PostgreSQL “TOASTs” values over 2KB (510d 4-byte float)
- Storage types:
  - PLAIN: Data stored inline with table
  - EXTENDED: Data stored/compressed in TOAST table when threshold exceeded
    - pgvector default before 0.6.0
  - EXTERNAL: Data stored in TOAST table when threshold exceeded
    - pgvector default 0.6.0+
  - MAIN: Data stored compressed inline with table





# ↔ Tradeoffs: Impact of TOAST on vector data

- Traditionally, TOAST data is not on the "hot path"
  - Impacts query plan and maintenance operations
- Compression is ineffective
- Unable to use for index pages



# Space utilization on a page

Dimensions	Vectors / Page	Wasted Space (B)
128	15	308
256	7	916
384	5	428
512	3	1,988
768	2	2,000
1,024	1	4,060
1,536	1	2,012
2,000	1	156

$$\text{PAGE\_SIZE} - \text{PAGE\_HEADER} - (\text{VECTORS} * 4) - \text{VECTORS} * (4 * \text{DIMS} + 8)$$

# What do we need to define?

1. Data type
2. Distance functions and operators
3. Indexing strategy

# Distance functions / operators

Euclidean / L2

<->

```
CREATE FUNCTION  
  l2_distance(vector, vector)  
RETURNS float8
```

Cosine

<=>

```
CREATE FUNCTION  
  cosine_distance(vector,  
vector) RETURNS float8  
IMMUTABLE STRICT  
PARALLEL SAFE,  
LANGUAGE C
```

Inner Product

<#>

```
AS 'MODULE_PATHNAME'  
LANGUAGE C  
IMMUTABLE STRICT  
PARALLEL SAFE;
```

Manhattan / Taxicab / L1

<+>

```
IMMUTABLE STRICT  
PARALLEL SAFE;
```

Hamming

<~>

Jaccard

<%>

# PostgreSQL Infrastructure: Function definitions

```
FUNCTION_PREFIX PG_FUNCTION_INFO_V1(l2_distance);
```

```
Datum
```

```
l2_distance(PG_FUNCTION_ARGS)
```

```
{
```

```
    Vector    *a = PG_GETARG_VECTOR_P(0);
```

```
    Vector    *b = PG_GETARG_VECTOR_P(1);
```

```
    checkDims(a, b);
```

```
    PG_RETURN_FLOAT8(sqrt((double)
```

```
        VectorL2SquaredDistance(a->dim, a->x, b->x)));
```

```
}
```

# Shortcut: SIMD using compiler autovectorization

```
VectorCosineSimilarity(int dim, float *ax, float *bx)
{
    /* ... */
    /* Auto-vectorized */
    for (int i = 0; i < dim; i++)
    {
        similarity += ax[i] * bx[i];
        norma += ax[i] * ax[i];
        normb += bx[i] * bx[i];
    }
    /* Use sqrt(a * b) over sqrt(a) * sqrt(b) */
    return (double) similarity / sqrt((double) norma * (double) normb);
}
```

# Shortcut: CPU dispatching (AVX-512)

```
TARGET_AVX512_POPCOUNT static uint64
```

```
BitHammingDistanceAvx512Popcount(uint32 bytes, unsigned char *ax, unsigned char *bx, uint64 distance)
```

```
{
```

```
    __m512i    dist = _mm512_setzero_si512();
```

```
    for (; bytes >= sizeof(__m512i); bytes -= sizeof(__m512i))
```

```
    {
```

```
        __m512i    axs = _mm512_loadu_si512((const __m512i *) ax);
```

```
        __m512i    bxs = _mm512_loadu_si512((const __m512i *) bx);
```

```
        dist = _mm512_add_epi64(dist, _mm512_popcnt_epi64(_mm512_xor_si512(axs, bxs)));
```

```
        ax += sizeof(__m512i);
```

```
        bx += sizeof(__m512i);
```

```
    }
```

```
    distance += _mm512_reduce_add_epi64(dist);
```

```
    return BitHammingDistanceDefault(bytes, ax, bx, distance);
```

# What do we need to define?

1. Data type
2. Distance functions and operators
3. Indexing strategy



# PostgreSQL index interfaces

- GiST (Generalized Search Tree)
  - Supports K-NN queries
- SP-GiST (Space-partitioned Generalized Search Tree)
  - Supports K-NN queries
- GIN (Generalized Inverted Index)
- BRIN (Block Range Index)
- B-tree
- Hash

# Example: Interfacing with GiST

- consistent
- union
- penalty
- picksplit
- same
- compress
- decompress
- distance
- fetch

# ↔ Index access methods ("custom indexes")

- Let you define indexes that don't fit existing interfaces
  - Properties
  - Methods
- "More work"
  - Responsible for vacuum, WAL, locking, planning, et al.
  - (More) responsible for impact due to upstream changes

# Key index access method properties for pgvector

- `amcanorder` => `false`
- `amcanorderbyop` => `true`
- `amcanbuildparallel` => `true`

Reference: <https://www.postgresql.org/docs/current/index-api.html>

# Key index access method functions for pgvector

- `ambuild`
- `aminsert`
- `ambulkdelete`
- `amcostestimate`
- `ambeginscan`
- `amrescan`
- `amgettuple`

Reference: <https://www.postgresql.org/docs/current/index-functions.html>

# Key index access method functions for pgvector

- `ambuild`
- `aminsert`

Reference: <https://www.postgresql.org/docs/current/index-functions.html>

# Key index access method functions for pgvector

- `ambulkdelete`

Reference: <https://www.postgresql.org/docs/current/index-functions.html>

# Key index access method functions for pgvector

- `amcostestimate`

Reference: <https://www.postgresql.org/docs/current/index-functions.html>




# Key index access method functions for pgvector

- `ambeginscan`
- `amrescan`
- `amgettup1e`


Reference: <https://www.postgresql.org/docs/current/index-functions.html>

# pgvector index methods: IVFFlat and HNSW

- IVFFlat
  - K-means based
  - Organize vectors into lists
  - Requires prepopulated data
  - Insert time bounded by # lists



- HNSW
  - Graph based
  - Organize vectors into "neighborhoods"
  - Iterative insertions
  - Insertion time increases as data in graph increases





# Shortcut: Store normalized vectors in index

$$\text{L2 normalization} = \mathbf{v} / \|\mathbf{v}\|$$

0.0234  
0.093  
-0.9123  
0.1055



- ✓ Same dimensions?
- ✓ Magnitude > 0?



🔧 If not, normalize

0.0253  
0.1007  
-0.9880  
0.1142



# Shortcut: skip operations with normalization

```
VectorCosineSimilarity(int dim, float *ax, float *bx)
```

```
{
```

```
    /* ... */
```

```
    /* Auto-vectorized */
```

```
    for (int i = 0; i < dim; i++)
```

```
    {
```

```
        similarity += ax[i] * bx[i];
```

```
        norma += ax[i] * ax[i];
```

```
        normb += bx[i] * bx[i];
```

```
    }
```

```
    /* Use sqrt(a * b) over sqrt(a) * sqrt(b) */
```

```
    return (double) similarity / sqrt((double) norma * (double) normb);
```



# Shortcut: skip operations with normalization

```
VECTOR_TARGET_CLONES static float
VectorInnerProduct(int dim, float *ax, float *bx)
{
    float        distance = 0.0;

    /* Auto-vectorized */
    for (int i = 0; i < dim; i++)
        distance += ax[i] * bx[i];

    return distance;
}
```



# Shortcut: skip operations with normalization

```
FUNCTION_PREFIX PG_FUNCTION_INFO_V1(vector_negative_inner_product);
```

```
Datum
```

```
vector_negative_inner_product(PG_FUNCTION_ARGS)
```

```
{
```

```
    Vector    *a = PG_GETARG_VECTOR_P(0);
```

```
    Vector    *b = PG_GETARG_VECTOR_P(1);
```

```
    checkDims(a, b);
```




```
    PG_RETURN_FLOAT8((double)-VectorInnerProduct(a->dim, a->x, b->x));
```

```
}
```

# What do we need to define?


1. Data type
2. Distance functions and operators
3. Indexing strategy
  1. IVFFlat
  2. HNSW

# Building an IVFFlat index


1. Sample the overall vectors in the table (`MAX(50*lists), 10,000`)
  - Uses BlockSampler (ANALYZE method) 
2. Calculate K-means (`ivfflat.lists`)
3. Assign vectors to lists in memory
4. Sort vectors in lists 
5. Save index to storage 



# IVFFlat: sampling




# IVFFlat: sampling



# IVFFlat: K-means


pgvector Elkan's K-means algorithms



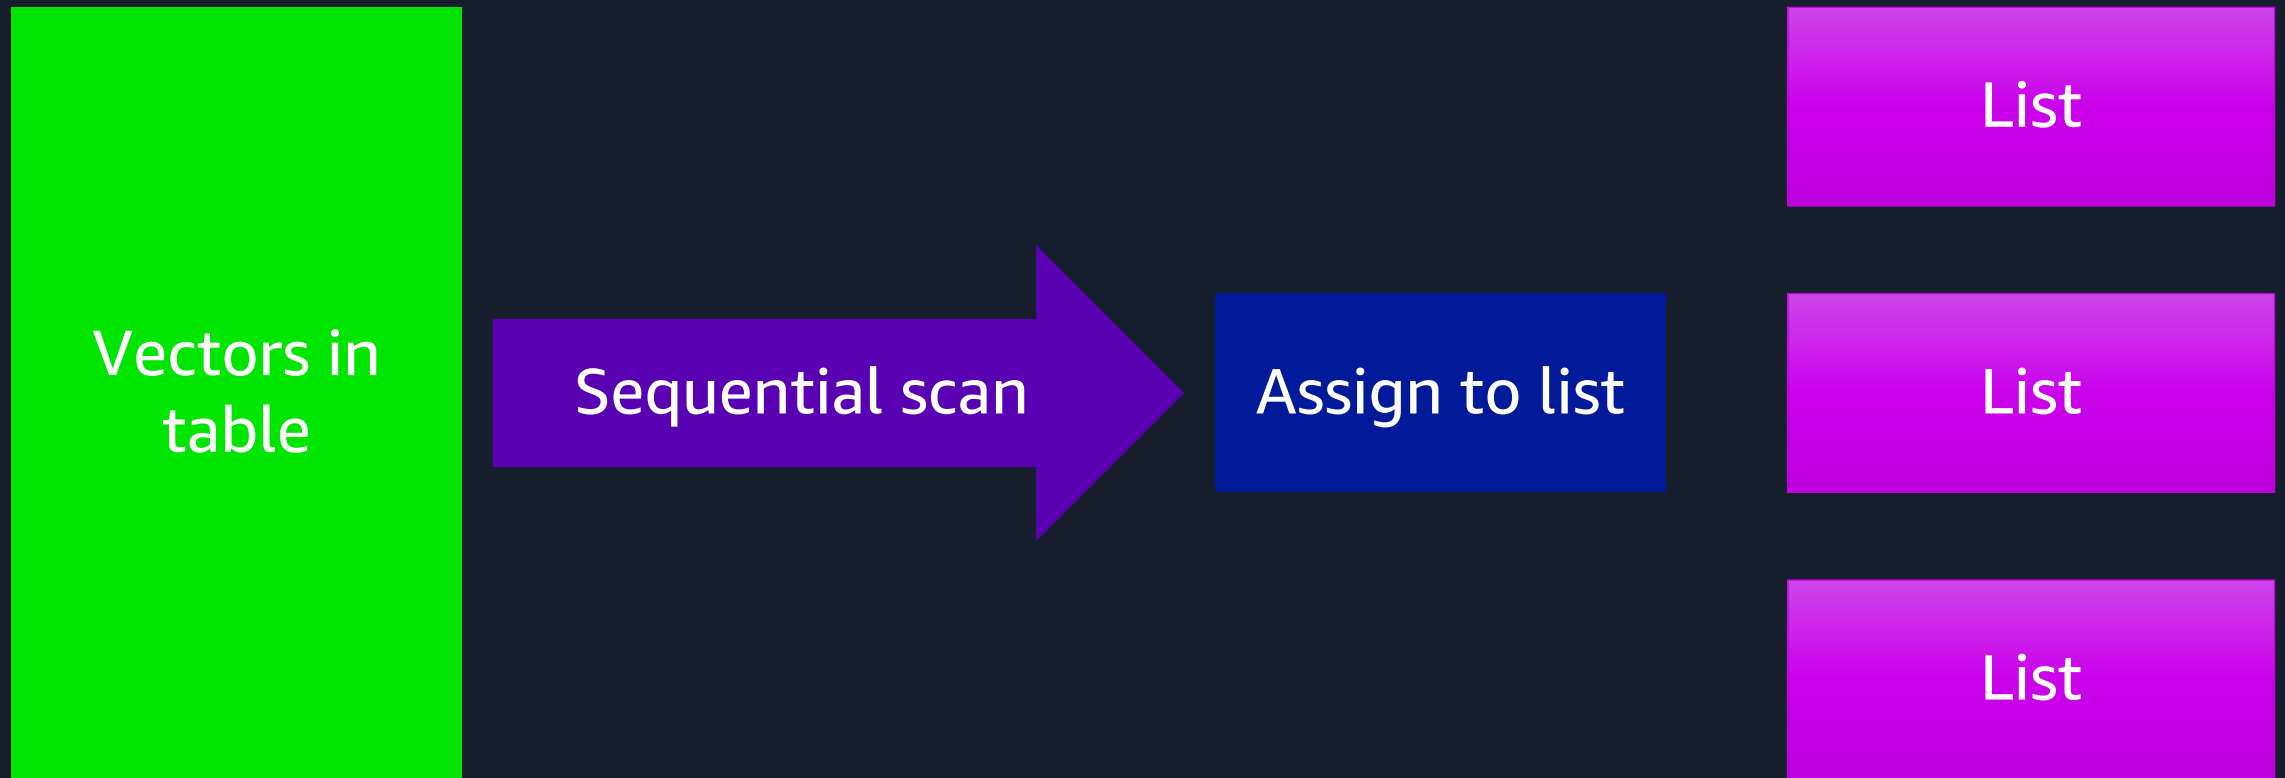
`ivfflat.lists = 3`



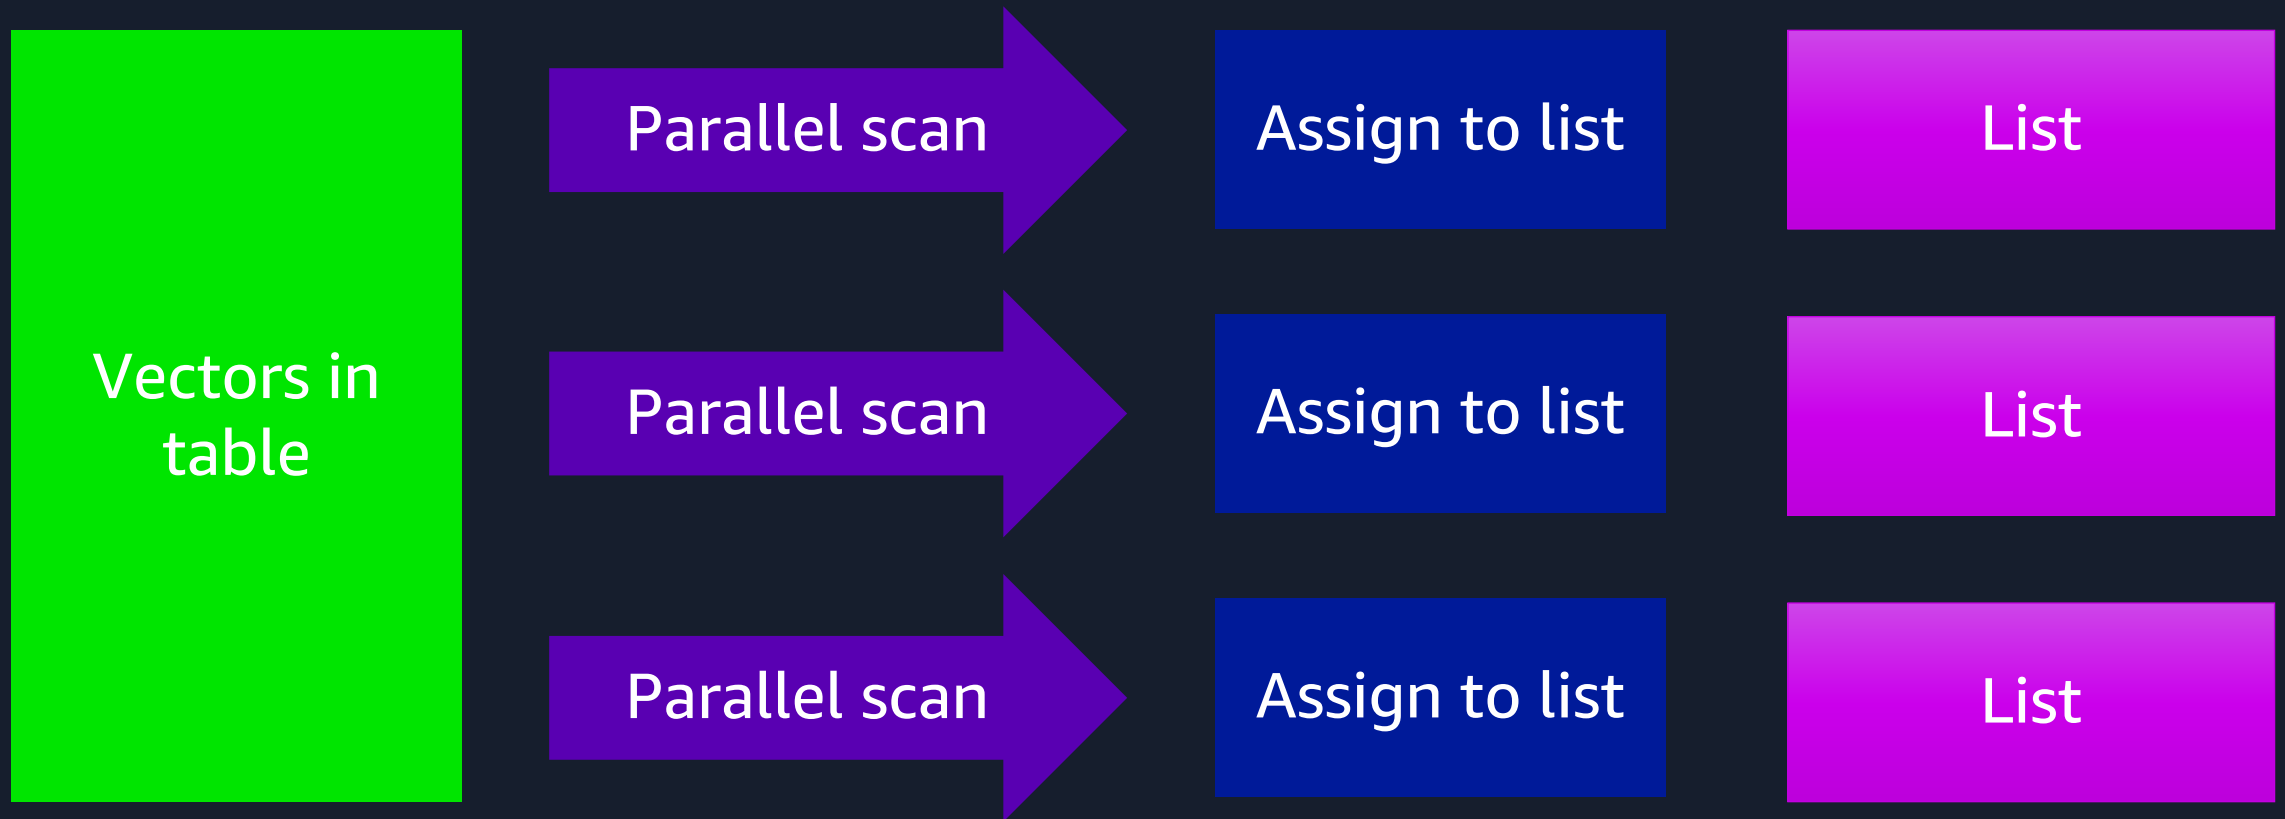
# IVFFlat: list assignment



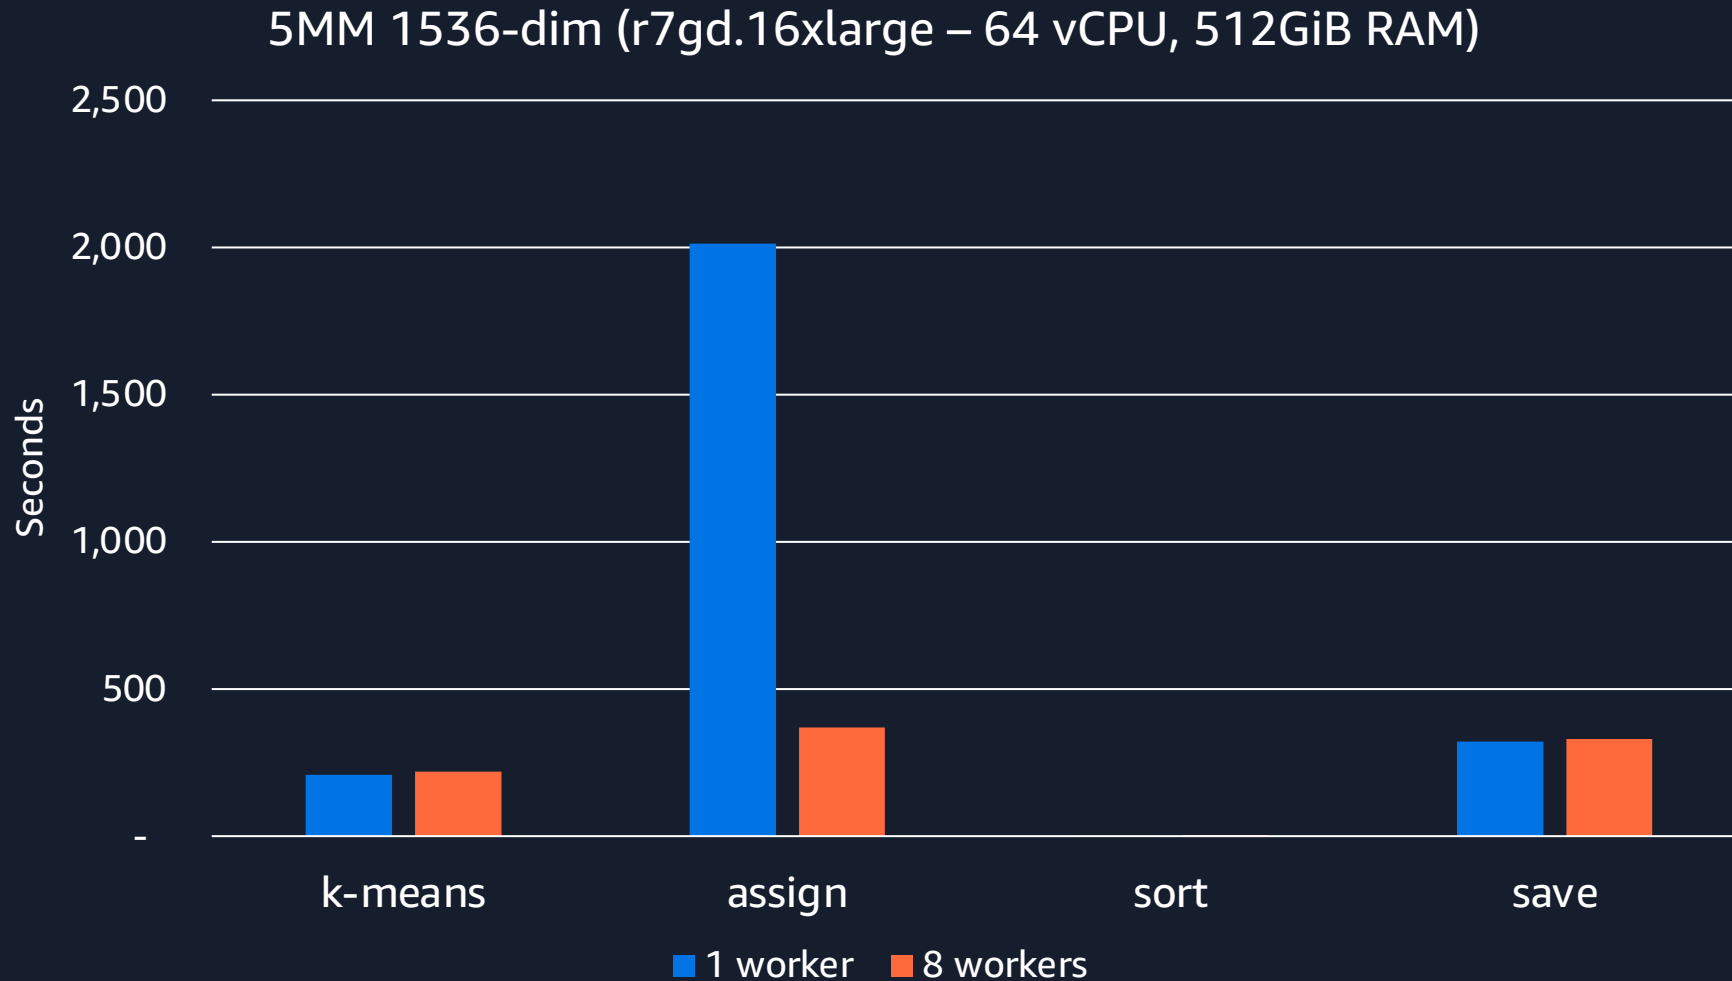
# Parallelism and list assignment



# Parallelism and list assignment




# Parallelism and list assignment




# Parallelism and list assignment

1 worker



■ k-means ■ assign ■ sort ■ save

8 workers




■ k-means ■ assign ■ sort ■ save



# IVFFlat: Save index to storage



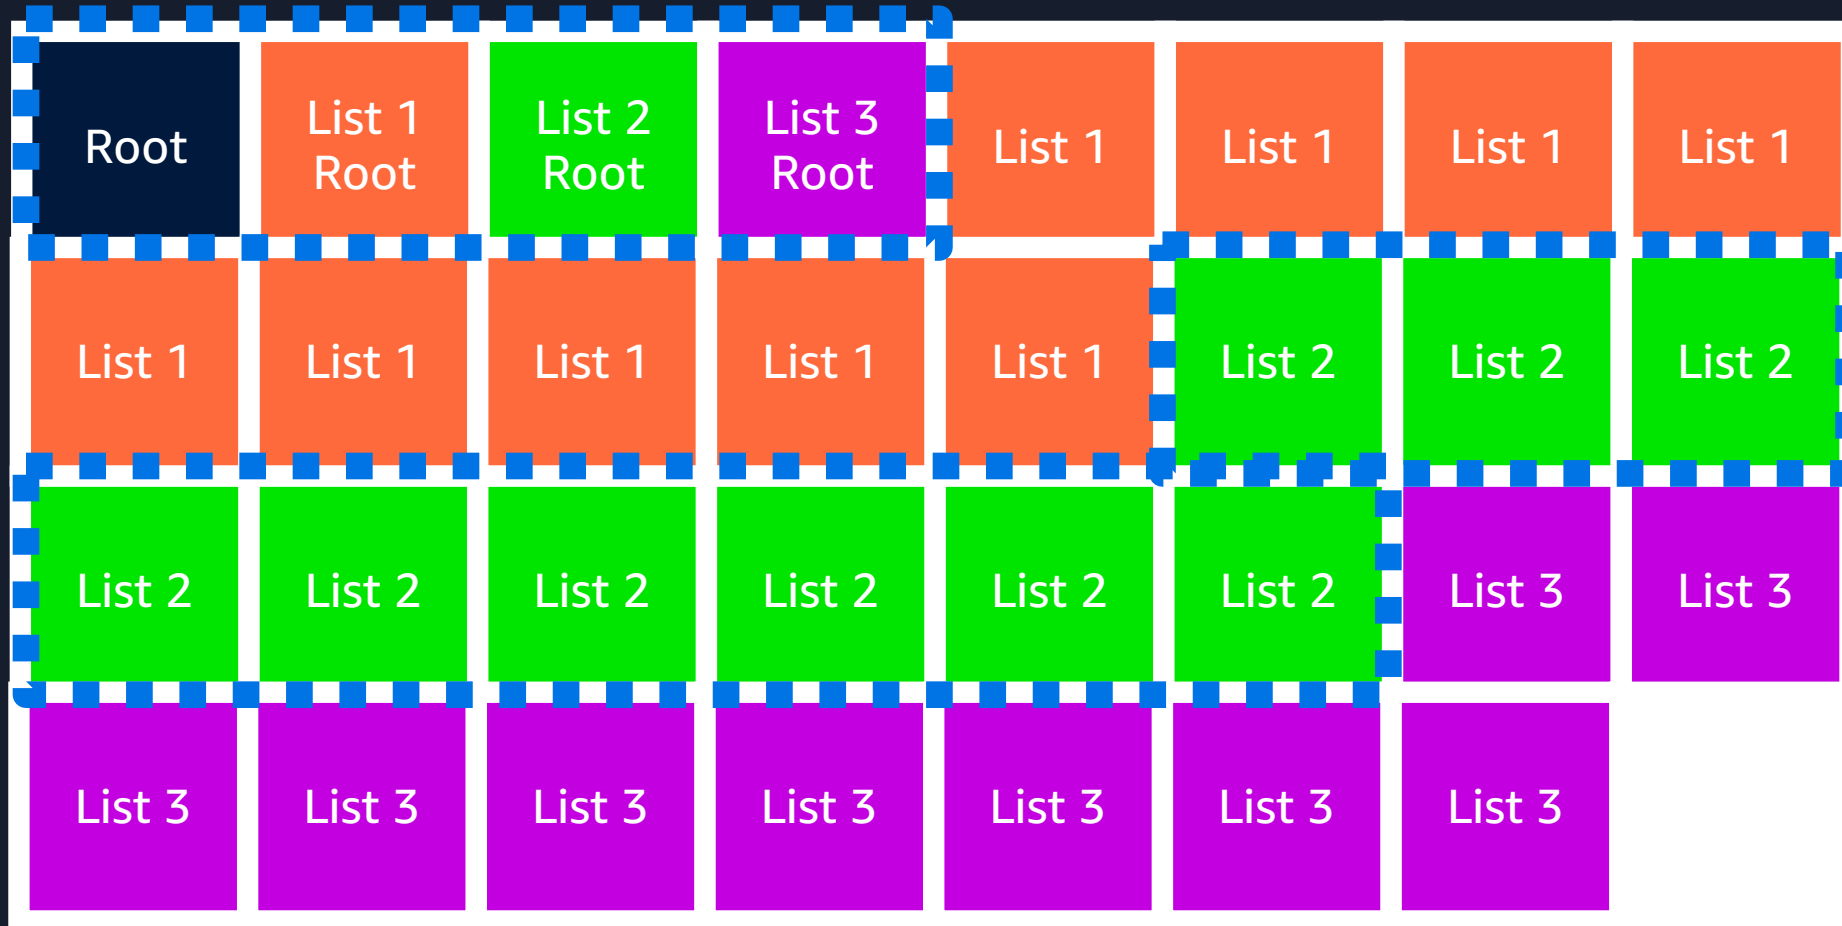
# Querying an IVFFlat index




SET ivfflat.probes TO 1

```
SELECT id FROM products ORDER BY $1 <-> embedding LIMIT 3
```

# Querying an index an IVFFlat index (1 probe)



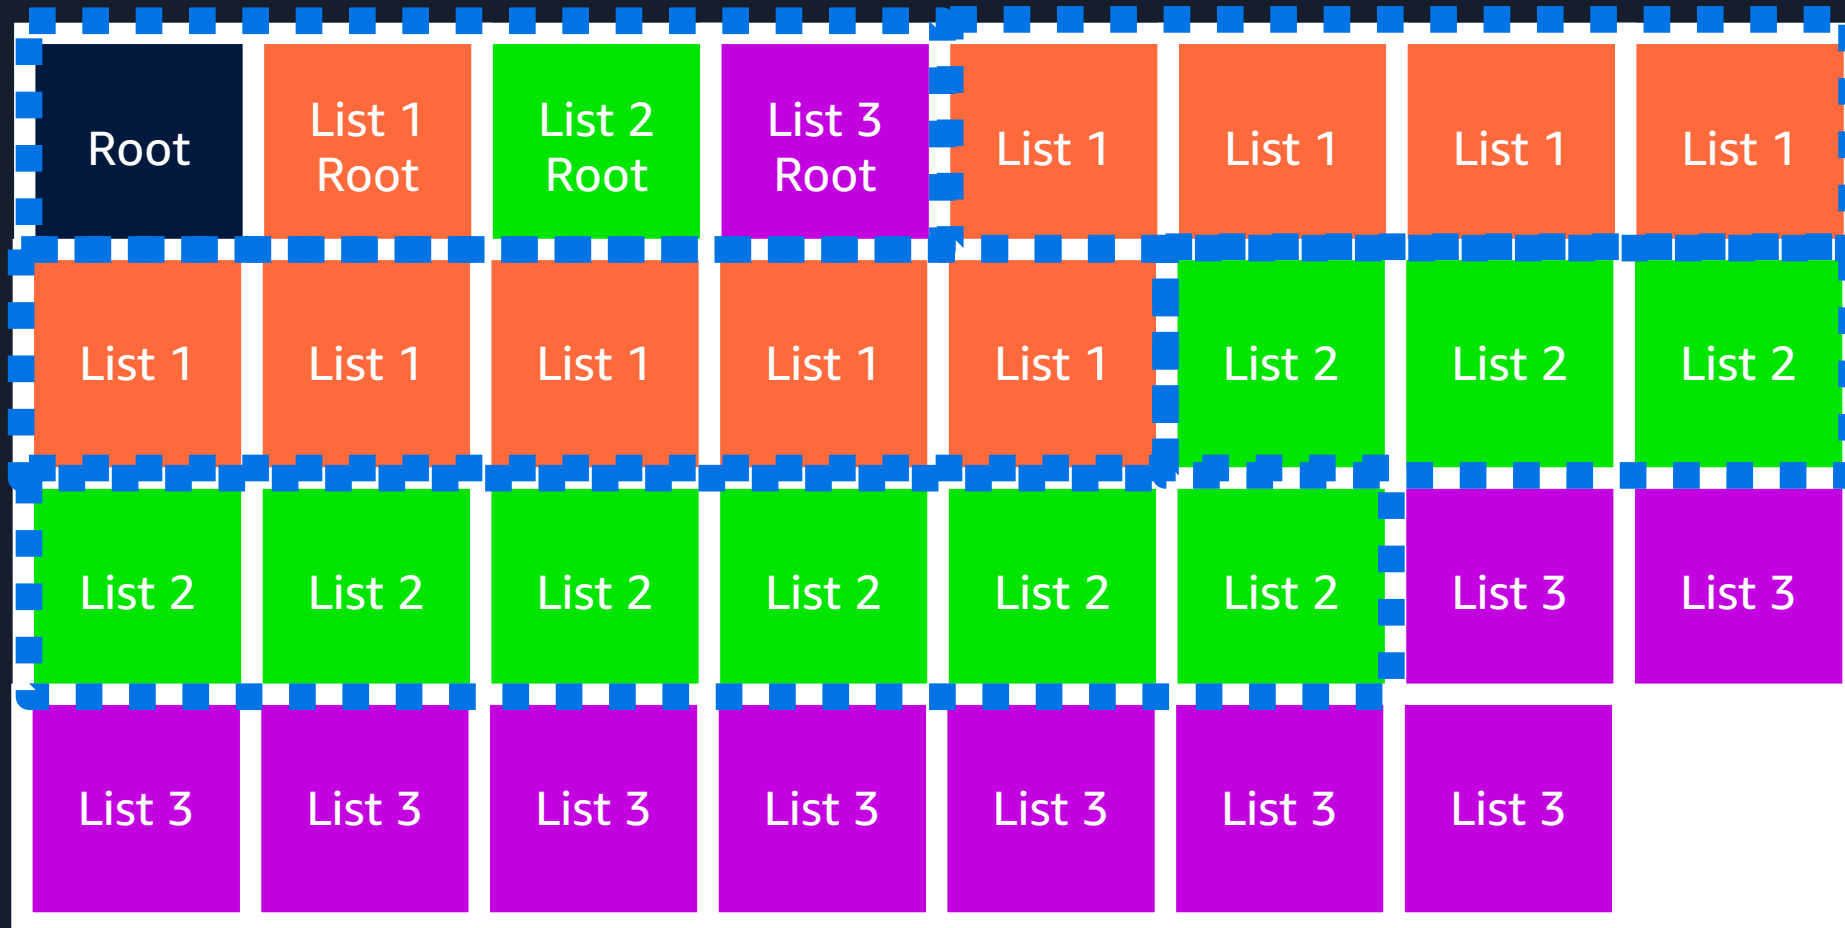
# Querying an IVFFlat index



SET ivfflat.probes TO 2

```
SELECT id FROM products ORDER BY $1 <-> embedding LIMIT 3
```

# Querying an IVFFlat index (2 probes)



# IVFFlat considerations

- Temporal locality is directly impacted by both cluster quality and query patterns
- Latency grows linearly with probes
- Lookups outside of memory can be very expensive
- Insertions / updates can skew lookups and query quality
- Opportunities
  - Streaming I/O
  - Quantization (available, requires more evaluation)
  - Additional algorithmic improvements (e.g. SPANN)

# What do we need to define?

1. Data type
2. Distance functions and operators
3. Indexing strategy
  1. IVFFlat
  2. HNSW

# Hierarchical navigable small worlds (HNSW)

- Each vector organized into "microclusters" ("neighborhoods")
- Spend minimal time in "upper layers" – most search in bottom layer ("Layer 0")



# HNSW index building parameters

m

Maximum number of bidirectional links between indexed vectors

Default: 16

ef\_construction

Number of vectors to maintain in “nearest neighbor” list


Default: 64

# HNSW query parameters


## `hnsw.ef_search`

- Number of vectors to maintain in “nearest neighbor” list
- Before v0.8, must be greater than or equal to `LIMIT`
- v0.8+, can use `hnsw.iterative_search` to satisfy unmet `LIMIT`


# Querying an HNSW index




# Querying an HNSW index




# Querying an HNSW index




# Querying an HNSW index



# Querying an HNSW index



# Querying an HNSW index






# 🚀 What happens when searching a HNSW index?


- Maintain a list of visited vectors (tuple IDs / TIDs)
- Maintain an ordered list of candidates with distances
- `ef_search` is 1 at Layer 1+
- `ef_search` is `ef_search` (default 40) at Layer 0

Visited
0x0102030405060708
0x0102030405060709
0x0102030405060710


Candidates	
0x0102030405060708	0.0123
0x0102030405060709	0.0434
0x0102030405060710	0.0845




# Building an HNSW index




# Building an HNSW index




# Building an HNSW index



# Building an HNSW index



# Building an HNSW index



Layer 0

# Building an HNSW index

1. Determine entry level
2. Determine insertion method ("in memory" or "on-disk")
3. Find neighbors (similar to querying)
  - Layer 0:  $m * 2$
  - Otherwise:  $m$
4. Add vector to graph
5. Update neighbors' bidirectional links

```
entryLevel = (int) (-log(RandomDouble()) * m1);
```

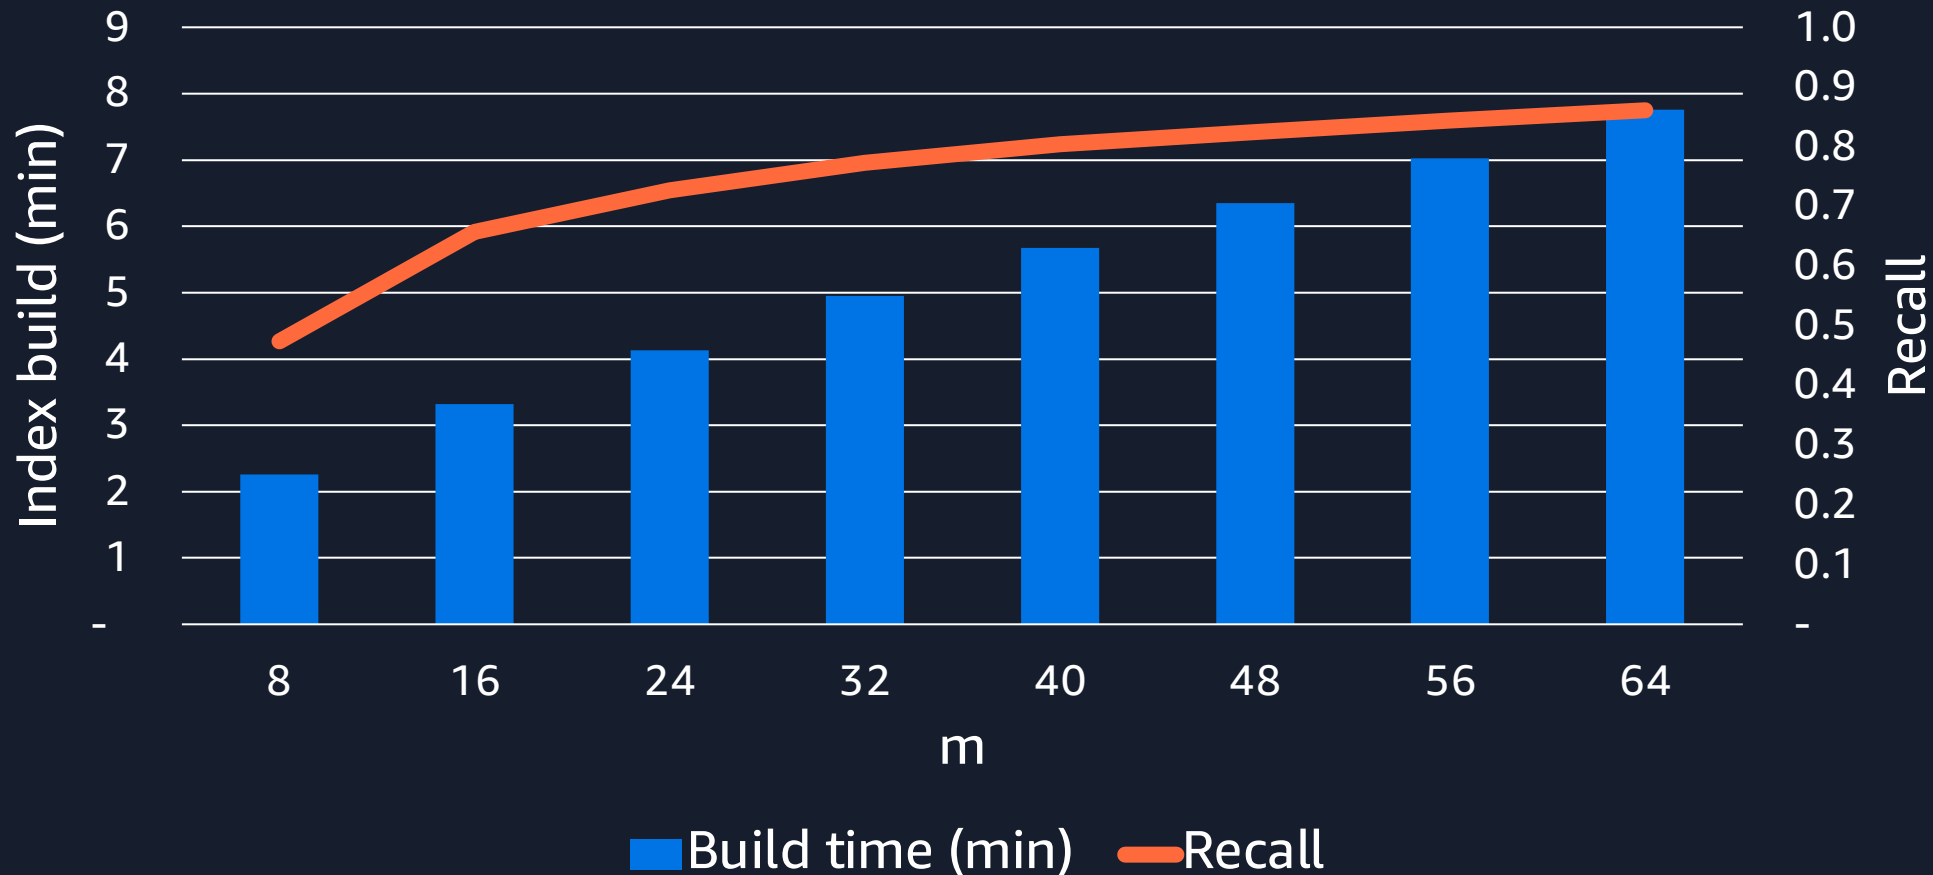
# HNSW entry level distribution

m	Layer 0 Entry Level	Layer 1 Entry Level
2	50%	25%
4	75%	19%
8	87%	11%
12	92%	8%
16	94%	6%
20	95%	5%
24	96%	4%
32	97%	3%
36	97%	3%
48	98%	2%
64	98%	2%



# How “m” impacts index build time & search quality

GIST960 1M 960-dim vectors, ef\_construction=256,  
hnswef\_search=20, max\_parallel\_maintenance\_workers=63



# How "m" impacts query time via tuples scanned

m=16, ef\_construction=64

# tuples scanned

ef	SIFT (N=1M)	GIST (N=1M)	GLoVE25 (N=1.1M)	1536d (N=5M)	768d (N=10M)
10	427	512	438	456	498
20	643	779	652	650	695
40	1044	1272	1049	1005	1050
80	1774	2212	1761	1629	1762
120	2438	3099	2420	2214	2449
200	3638	4755	3629	3328	3833
400	6247	8402	6303	5836	7190
800	10619	14706	10938	10563	13258

# How "m" impacts query time via tuples scanned

ef	1536d (N=5M,m=16)	1536d (N=5M,m=64)	768d (N=10M,m=16)	768d (N=10M,m=64)
10	456	605	498	1425
20	650	1257	695	2038
40	1005	2292	1050	3246
80	1629	4049	1762	5691
120	2214	5728	2449	8046
200	3328	8601	3833	12664
400	5836	15158	7190	23284
800	10563	27249	13258	42200


# HNSW scan cost estimation

```
/*
 * HNSW cost estimation follows a formula that accounts for the total
 * number of tuples indexed combined with the parameters that most
 * influence the duration of the index scan, namely: m - the number of
 * tuples that are scanned in each step of the HNSW graph traversal
 * ef_search - which influences the total number of steps taken at layer 0
 *
 * The source of the vector data can impact how many steps it takes to
 * converge on the set of vectors to return to the executor. Currently, we
 * use a hardcoded scaling factor (HNSWScanScalingFactor) to help
 * influence that, but this could later become a configurable parameter
 * based on the cost estimations.
 *
 * The tuple estimator formula is below:
 *
 * numIndexTuples = entryLevel * m + layer0TuplesMax * layer0Selectivity
 */
```

# Why is cost estimation important?


- Guides PostgreSQL query planner to select "best path"
- Filtering (WHERE clause)
  - A different index (B-tree) or a sequential scan may be a better choice based on selectivity

# Why index build speed matters (serial build)




# Why index build speed matters (parallel build)

990K 1536-dim vectors, m=16, hnsw.ef\_search=20,  
max\_maintenance\_workers=63



# pgvector and HNSW index maintenance

- Innovation: pgvector HNSW implementation supports updates and deletes



Phase 2: Repair



# HNSW considerations

- Embedding model impacts overall query time
- Filtering
  - Iterative scans vs. using other search methods
  - Bitmap scans(?)
- Opportunities to accelerate time spent in Layer 0
- Opportunities
  - Streaming I/O
  - Parallel vacuum
  - "Smart" graph repair to improve clustering

# What is quantization?

## Flat

[0.0435122, -0.2304432, -0.4521324,  
0.98652234, -0.1123234, 0.75401234]



## Scalar quantization (2-byte float)

[0.0432, -0.234, -0.452, 0.986,  
-0.112, 0.751]



## Scalar quantization (1-byte uint)

[129, 99, 67, 244, 126, 230]



## Binary quantization

[1, 0, 0, 1, 0, 1]



# pgvector and scalar quantization (2 byte)

```
CREATE INDEX ON documents USING  
    hnsw((embedding::halfvec(3072)) halfvec_cosine_ops);
```

```
SELECT id  
FROM documents  
ORDER BY embedding::halfvec(3072) <=> $1::halfvec(3072)  
LIMIT 10;
```

# pgvector and binary quantization

```
CREATE INDEX ON documents USING  
    hnsw ((binary_quantize(embedding)::bit(3072)) bit_hamming_ops);
```

```
SELECT i.id FROM (  
    SELECT id, embedding <=> $1 AS distance  
    FROM items  
    ORDER BY  
        binary_quantize(embedding)::bit(3072) <~> binary_quantize($1)  
    LIMIT 40 -- set to hnsw.ef_search  
) i  
ORDER BY i.distance  
LIMIT 10;
```



1536d 5MM (r7i.16xlarge, m=16, ef_construction=256)			
	Flat	2-byte float	Binary (rerank)
Index Size (GB)	38.15	19.07	2.34
Index build time (min)	21	13	4
Recall @ ef_search = 40	0.931	0.929	0.811
QPS @ ef_search = 40	24,216	27,084	33,984
Recall @ ef_search = 80	0.965	0.961	0.900
QPS @ ef_search = 80	11,057	12,759	20,410
Recall @ ef_search = 220	0.989	0.983	0.963
QPS @ ef_search = 220	5,242	5,983	7,856



1536d 5MM (r7i.16xlarge, m=16, ef_construction=256)			
	Flat	2-byte float	Binary (rerank)
Index Size (GB)	38.15	19.07	2.34
Index build time (min)	21	13	4
Recall @ ef_search = 40	0.931	0.929	0.811
QPS @ ef_search = 40	24,216	27,084	33,984
Recall @ ef_search = 80	0.965	0.961	0.900
QPS @ ef_search = 80	11,057	12,759	20,410
Recall @ ef_search = 220	0.989	0.983	0.963
QPS @ ef_search = 220	5,242	5,983	7,856



1536d 5MM (r7i.16xlarge, m=16, ef_construction=256)			
	Flat	2-byte float	Binary (rerank)
Index Size (GB)	38.15	19.07	2.34
Index build time (min)	21	13	4
Recall @ ef_search = 40	0.931	0.929	0.811
QPS @ ef_search = 40	24,216	27,084	33,984
Recall @ ef_search = 80	0.965	0.961	0.900
QPS @ ef_search = 80	11,057	12,759	20,410
Recall @ ef_search = 220	0.989	0.983	0.963
QPS @ ef_search = 220	5,242	5,983	7,856

1536d 5MM (r7i.16xlarge, m=16, ef_construction=256)			
	Flat	2-byte float	Binary (rerank)
Index Size (GB)	38.15	19.07	2.34
Index build time (min)	21	13	4
Recall @ ef_search = 40	0.931	0.929	0.811
QPS @ ef_search = 40	24,216	27,084	33,984
Recall @ ef_search = 80	0.965	0.961	0.900
QPS @ ef_search = 80	11,057	12,759	20,410
Recall @ ef_search = 220	0.989	0.983	0.963
QPS @ ef_search = 220	5,242	5,983	7,856







# Ongoing work

# Areas to further explore

- "Multi-column" vector indexes
- Efficient batch queries
- Recall boosting techniques (statistical binary quantization, hybrid search)
- Demonstrably improved algorithms
- Upstream PostgreSQL changes that help vector search patterns

# Conclusion

- What works in memory may or may not work with storage-based systems 
- Extensible framework of PostgreSQL simplifies adding new search systems 
  - "You have vector search...and every other PostgreSQL feature"
- Rapidly evolving space, including open areas of research (e.g., filtering)

# Thank you!

**Jonathan Katz**

[jkatz@amazon.com](mailto:jkatz@amazon.com)

[@jkatz05](https://twitter.com/jkatz05)

