

Checking Database Size and Related Issues

As your Hydra database grows over time, it's important to monitor its size and identify what parts of the database consume the most space. Unmanaged growth can lead to performance issues, disk exhaustion, and backup delays. On Elestio, where Hydra is hosted in a managed environment, you can use SQL and command-line tools to measure database usage, analyze large objects, and troubleshoot storage problems. This guide explains how to check database size, detect bloated tables and indexes, and optimize storage usage efficiently.

Checking Database and Table Sizes

Hydra provides built-in functions to report the size of the current database, its individual schemas, tables, and indexes. These functions are useful for understanding where most of your storage is being used and planning cleanup or archiving strategies.

To check the total size of the active database:

```
SELECT pg_size_pretty(pg_database_size(current_database()));
```

This returns a human-readable value like "2 GB", indicating how much space the entire database consumes on disk.

To list the largest tables in your schema:

```
SELECT relname AS table, pg_size_pretty(pg_total_relation_size(relid)) AS total_size
FROM pg_catalog.pg_statio_user_tables
ORDER BY pg_total_relation_size(relid) DESC
LIMIT 10;
```

This helps you identify which tables take up the most space, including indexes and TOAST (large field) data.

To break down table vs index size separately:

```
SELECT relname AS object,  
       pg_size_pretty(pg_relation_size(relid)) AS table_size,  
       pg_size_pretty(pg_indexes_size(relid)) AS index_size  
FROM pg_catalog.pg_statio_user_tables  
ORDER BY pg_relation_size(relid) DESC  
LIMIT 10;
```

This distinction allows you to assess whether most space is used by raw table data or indexes, which can inform optimization decisions.

Identifying Bloat and Inefficiencies

Database bloat occurs when Hydra retains outdated or deleted rows due to its MVCC model. This is common in high-write tables and can lead to wasted space and degraded performance. Bloated indexes and tables are often invisible unless explicitly checked. To estimate bloat at a table level, you can use a community query like this:

```
SELECT schemaname, relname, round(100 * (pg_total_relation_size(relid) -  
pg_relation_size(relid)) / pg_total_relation_size(relid), 2) AS bloat_pct  
FROM pg_catalog.pg_statio_user_tables  
ORDER BY bloat_pct DESC  
LIMIT 10;
```

This query calculates how much of a table's total size is not accounted for by its base data—higher percentages suggest unused or dead space. You can also check dead tuples directly:

```
SELECT relname, n_dead_tup  
FROM pg_stat_user_tables  
ORDER BY n_dead_tup DESC  
LIMIT 10;
```

A high count of dead tuples suggests that autovacuum might not be keeping up and that a manual VACUUM could help.

Optimizing and Reducing Database Size

Once you've identified large or bloated objects, the next step is to optimize them. Hydra offers tools like VACUUM, REINDEX, and CLUSTER to reclaim space and improve storage efficiency. These commands must be run with care to avoid locking critical tables during active hours. To reclaim dead tuples and update statistics:

```
VACUUM ANALYZE;
```

This command removes dead rows and refreshes query planning statistics, which helps performance and frees up storage. To shrink large indexes that aren't cleaned automatically, use:

```
REINDEX TABLE <table_name>;
```

This recreates the table's indexes from scratch and can free up disk space if indexes are fragmented or bloated. If a table is heavily bloated and full table rewrites are acceptable during maintenance, use:

```
CLUSTER <table_name>;
```

This rewrites the entire table based on an index order and reclaims space similar to VACUUM FULL, but with more control.

Additionally, removing or archiving old data from large time-based tables can reduce total size. Consider partitioning large tables to manage this process more efficiently.

Revision #1

Created 2025-05-08 11:16:06 UTC

Updated 2025-05-08 11:17:27 UTC