

FUJITSU Enterprise Postgres

Database Indexing Technology – In-Memory Columnar Index



Vertical Clustered Index (VCI) is Fujitsu's proprietary implementation of In-Memory Columnar Index, and can significantly improve aggregation processing – on a single server with 56 CPU cores, **performance is almost 5 times faster!**

What is FUJITSU Enterprise Postgres?

FUJITSU Enterprise Postgres is an exceptionally reliable and robust relational database, created for organisations that require strong query performance and high availability. It is fully compatible with the world-renowned open source database management system—PostgreSQL, with additional enterprise-grade features for enhanced security and better performance.

PostgreSQL index types

PostgreSQL provides several index types: B-tree, Hash, GiST, SP-GiST, GIN and BRIN. Each index type uses a different algorithm that is best suited to different types of queries.

As FUJITSU Enterprise Postgres extends PostgreSQL, it also supports all these index types.

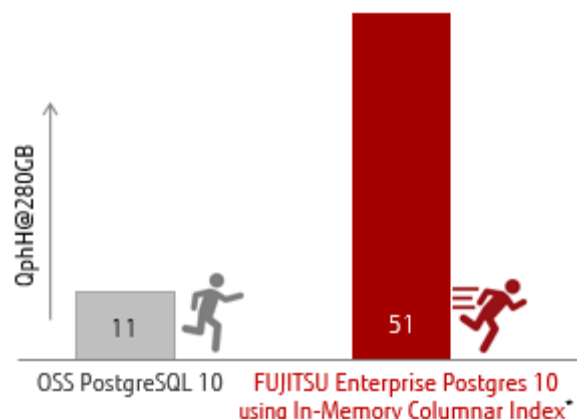
What is Vertical Clustered Index?

Unique to FUJITSU Enterprise Postgres Advanced Edition, Vertical Clustered Index (VCI) is an integrated feature that greatly improves the performance of queries performing aggregate operations such as SUM, AVG, MIN, and MAX on very large tables. VCI stores row-oriented data in columnar format in memory, enabling quick aggregation to be performed against the data.

VCI features

- Data can be aggregated because updated row-oriented data is automatically reflected to its columnar equivalent in real time.
- The impact on other jobs can be avoided by expanding to memory dedicated for columnar data (flushing from the database cache).
- Stable performance can be achieved even immediately after restarting because columnar data is also stored on disk.

Results for a 280 GB dataset
on a 56-core Linux node



* Implemented via Vertical Clustered Index (VCI) by Fujitsu Laboratories Limited

Using VCI

- VCI can be cached into memory dedicated to it by configuring the `reserved_buffer_ratio` parameter introduced in FUJITSU Enterprise Postgres specifically for caching VCI. This means a certain percentage of `shared_buffers` can be dedicated to VCI.
- VCI needs to be pre-warmed into the cache using the FUJITSU Enterprise Postgres-specific pre-warming function called `pgx_prewarm_vci`, which is an enhanced version of the open source feature called `pg_prewarm`. The `pgx_prewarm_vci` function loads the specified VCI data into buffer cache and returns the number of blocks of the loaded VCI data.

Immediately after an instance is started, the aggregation process using VCI may take time, because the VCI data has not been loaded to buffer cache. Therefore, the first aggregation process can be sped up by executing `pgx_prewarm_vci` after an instance is started. The amount of memory required for preloading is the number of blocks returned by `pgx_prewarm_vci` multiplied by the size of one block.

This function can only be executed if the user has reference privilege to the VCI index and execution privilege to the `pg_prewarm` function.

Configuration is as easy as follows.

Install the extensions

```
CREATE EXTENSION vci
CREATE EXTENSION pg_prewarm;
```

Update postgresql.conf

```
shared_preload_libraries='vci, pg_prewarm,...'
reserved_buffer_ratio=50*
```

* Suggested value only

Create VCI index on the relevant columns

```
CREATE INDEX idx_vci_make_model
ON items_sold
USING vci (make, model)
WITH (stable_buffer=true);
```

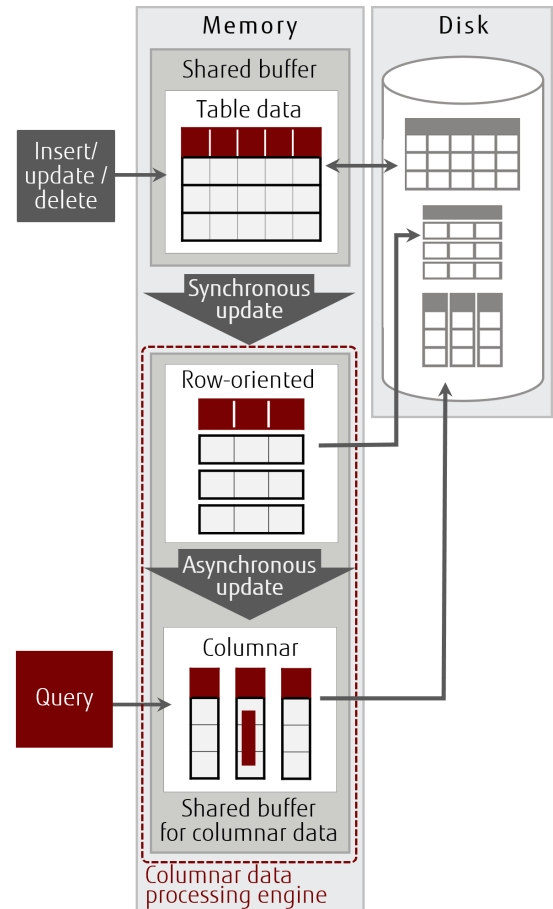
Pre-warm (load) VCI data into the cache

```
SELECT pgx_prewarm_vci('idx_vci_make_model');
```

Run the aggregate operation

```
SELECT make, model, SUM(sales)
FROM items_sold
GROUP BY (make, model);
```

VCI in action



Conclusion

Developed by Fujitsu Laboratories Ltd, Fujitsu's Vertical Clustered Index is a powerful feature that can help you drastically improve aggregation performance with very little configuration effort required. It gives query optimisation the additional option to use columnar data stored in memory instead of row data to suit the query's nature.

Moving Forward with Fujitsu

Now that OSS-based databases are increasingly adopted in leading industries, we believe further development of our relationship with the PostgreSQL community will not only provide our existing customers with the confidence to make a greater investment in open source database technologies, but also encourage other software developers to engage in this process.

About Fujitsu

Fujitsu is the 5th largest IT service provider in the world, offering a full range of technology products, solutions and services. Around 160,000 Fujitsu employees support customers in over 100 countries.

Contact

Fujitsu Australia Software Technology Pty Ltd
Address: 14 Rodborough Rd
Frenchs Forest NSW 2086
Australia
Email: postgresql@fast.au.fujitsu.com
Website: fast.fujitsu.com

Copyright 2020 FUJITSU AUSTRALIA SOFTWARE TECHNOLOGY. Fujitsu, the Fujitsu logo and Fujitsu brand names are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. All rights reserved. No part of this document may be reproduced, stored or transmitted in any form without prior written permission of Fujitsu Australia Software Technology. Fujitsu Australia Software Technology endeavours to ensure the information in this document is correct and fairly stated, but does not accept liability for any errors or omissions.