

White paper


FUJITSU Enterprise Postgres 13 on IBM LinuxONE™ – Improving your organization's data performance

Organizations are finding it increasingly challenging to be responsive to business requirements with the growth in volume and variety of data they need to capture every year.

While database management systems have provided a number of solutions to address these challenges, they fall short in their ability to deliver satisfactory performance of both business transactions and analytical queries against typical data architecture.

FUJITSU Enterprise Postgres 13 on IBM LinuxONE™ delivers an exceptionally reliable and robust relational database for organizations that require strong query performance and high availability.

This white paper quantifies how FUJITSU Enterprise Postgres on IBM LinuxONE™ can deliver significant performance gains by providing the means for fast aggregation of large data sets and improved high-speed data loads.

A composite image featuring a female athlete in mid-stride, wearing a black sports bra and shorts, and purple running shoes. She is positioned in front of a dark background where binary code (0s and 1s) is visualized as glowing blue lines and numbers, forming a circular track around her. The overall theme is the speed and efficiency of data processing.

| Content | |
|--|---|
| Executive summary | 2 |
| FUJITSU Enterprise Postgres on IBM LinuxONE..... | 2 |
| In-Memory Columnar Index for HTAP workloads | 2 |
| Benefits..... | 2 |
| Benchmarking results..... | 2 |
| High-speed Data Load | 3 |
| Benefits..... | 3 |
| Benchmarking results..... | 3 |
| Conclusion..... | 3 |

Executive summary

One of the driving forces behind business innovation is utilization of information to gain insight on your business and respond accordingly. This means that IT systems will need to adapt, not only to keep pace with exponential data generation, but also to be able to quickly store and analyses larger and larger amounts of data.

The pressure is on for organizations to be able to quickly transfer and aggregate data for analysis, to make faster business decisions and quickly identify business opportunities in order to remain competitive.

FUJITSU Enterprise Postgres on IBM LinuxONE

In-Memory Columnar Index for HTAP workloads

The demand to quickly perform data analysis has been a crucial element for information utilization in recent years, and as result IT systems are under pressure to provide high-volume data manipulation. It is no longer acceptable to have to wait for hours while data aggregation/manipulation processes run before the result can be used for business intelligence.

Being able to obtain real-time information quickly and regularly, and swiftly reflect that to the business often impacts the ability to quickly perform transactions, but as we will see in this paper, this need not be the case.

Our solution is to have a single database system efficiently handle both large amounts of data transactions and data analysis of large amounts of data, bridging the gap between the solutions tailored for either operation.

Fujitsu's Vertical Clustered Index (VCI) engine is a performance feature of FUJITSU Enterprise Postgres. It provides significantly faster analytical query processing by storing a columnar representation of row-oriented data in memory, since this orientation is best suited for data analysis.

The engine reflects updated row-oriented data to its columnar equivalent, stores it without being dependent on memory capacity, and quickly conducts analysis of that data. Massive volumes of column-oriented data can be stored by taking advantage of this technique for managing data.

Benefits

The Vertical Clustered Index provides performance benefits for a variety of use cases. It allows OLAP workloads to run side by side with OLTP workloads in a single system, when normally each type would require a separate database instance.

This means that the data does not have to be transferred from an operational database to a data warehouse for large analytical SQL queries, thereby eliminating the need of an ETL data pipeline connecting OLTP instances with OLAP instances.

Hence, customers can bypass ETL processes, which can save several hours of data transfer data to data warehouses for critical time-consuming queries.

Eventually this feature allows reduction in operational costs by consolidating database servers and eliminating the need of multiple copies of the same data, while further reducing the storage, compute, network and maintenance costs.


Additional benefits include:

- Reduction in system complexity, as transactional and analytical tables reside in one location
- Savings in hardware, since there is no need for a separate reporting database
- Savings in operation costs, as ETL processes do not need to be maintained or revised
- Clearer/cleaner data lineage to analytical tables (less error-prone)
- Analytical results are more up to date, as there is no need to move data to a separate reporting environment, allowing business decisions to be made sooner


Benchmarking results

Our recent HTAP benchmark tests on IBM LinuxONE III show that for a 120 GB dataset on a 32 Integrated Facility for Linux (IFL) Linux node, this results in almost **4 times the throughput of analytical queries** while maintaining equivalent transaction volumes when compared to open source PostgreSQL running on an equivalent x86 32-core Linux node.

Vertical Clustered Index HTAP processing:
FUJITSU Enterprise Postgres on IBM LinuxONE III
vs Open source PostgreSQL on x86



TPC-C throughput:
FUJITSU Enterprise Postgres on IBM LinuxONE III
vs Open source PostgreSQL on x86



High-speed Data Load

FUJITSU Enterprise Postgres offers the ability to load large volumes of data using as many parallel processes as possible, based on the number of cores and CPU availability.

Data load features of other database systems can only use a fixed number of parallel process, regardless of CPU availability. In an era of mission-critical systems which are able to take advantage of multi-core technology, this is not the optimal solution. FUJITSU Enterprise Postgres High-speed Data Load sends data from the input file to several parallel workers, each of which will simultaneously perform data conversion, table creation, and index creation.

Benefits


Benefits include:

- Less outage time in upgrades and other activities dependent on data loading
- New information can be loaded faster, making it available for analysis sooner
- Ability to configure degree of parallelism, allowing the impact on other systems to be managed, and thereby making scheduling easier
- Improved batch processing helps maximize available computing power and resources

Benchmarking results

Our recent Data Load performance benchmark on IBM LinuxONE III with an 8 Integrated Facility for Linux (IFL) Linux node using FUJITSU High-speed Data Load has shown **8.4 times better performance** for a similar workload tested on an equivalent 8 core x86 Linux node using the open source PostgreSQL COPY utility for data load.

Fujitsu High-speed Data Load on IBM LinuxONE III
vs Open source PostgreSQL COPY utility on x86



Conclusion

FUJITSU Enterprise Postgres on IBM Linux ONE embraces open source with improved security, performance, scalability and resiliency. Fujitsu's Vertical Clustered Index (VCI) engine and high-speed data load capabilities offer a world-class solution for data analysis of large data sets to support business intelligence while helping reduce operational costs.

Read more

For more information on FUJITSU Enterprise Postgres on IBM LinuxONE™ capabilities, we recommend our following resources:

- [FUJITSU Enterprise Postgres on IBM LinuxONE™ webpage](#)
- [FUJITSU Enterprise Postgres on IBM LinuxONE™ resources](#)

Contact us

If you have any questions about FUJITSU Enterprise Postgres on IBM LinuxONE, please feel free to contact us at enterprisepostgresforibm@fujitsu.com.

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
Email: postgresql@fast.au.fujitsu.com
Website: fast.fujitsu.com
2021-04-01 WW EN

Copyright 2021 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 endeavors to ensure the information in this document is correct and fairly stated, but does not accept liability for any errors or omissions.