

Accelerated BI and Analytics:

**Making data led decisions
in real time with GPU Technology**

A Whitepaper on the future of Business Intelligence

Table of Contents

Intro	3
Real time visualisation with GPUs	4
Ease of use	5
Security	5
Streaming	6
Data loading and Data ingest	7
Serverless architecture	7
Use Cases	8
Energy	
Telecommunications	
Smart Cities	
Conclusion	10

Intro

In 2022, studies have shown that over 40% of IT decision makers fear their infrastructure won't be able to handle future data demands. Enterprises seek to harness customer, company and market data to gain real-time insight that will drive growth. Decision making processes today are experiencing significant lag as enterprises struggle to analyse their evolving datasets at speed with legacy infrastructure. As workloads reach billion-row levels, CPU technology fails to grapple with modern advancements in data.

The immense parallel processing power of GPUs can now extend beyond the use of graphics processing and can accelerate the processing of numeric and general workloads. In the last five years, GPU (Graphics Processor Unit) technology has found its use case for databases.

GPU technology enables data analysts and other BI experts to gather insight from business-critical data in milliseconds.

The Brytlyt platform leverages GPU accelerated technology to supercharge business intelligence. The parallel processing capabilities of GPUs means that they work extremely well with in-depth database processing in milliseconds. With the ability to adapt to new queries and raw data faster than any other system, Brytlyt helps companies proactively perform analytics in real-time.

This paper will explore some of the obstacles companies face in accelerating decision-making processes and how the Brytlyt platform tackles these challenges with GPU accelerated technology.



Real time visualisation with GPUs

Brytlyt enables users to analyse data in motion, as things are happening. The immense parallel processing power of the latest GPU technology combined with Brytlyt's serverless solution, means that analytics can be carried out in real-time, even on billion-row datasets. Brytlyt empowers users with millisecond query response time and powerful real-time visuals that can be manipulated on-the-fly with customisable dashboards. Business experts can gain insight into their critical data instantly, as analysts can avoid waiting around for tedious data processing to complete or for legacy systems to catch up with streaming data.

Third party tools

Brytlyt is built on the open-source database PostgreSQL, a database used by organisations across the globe including Amazon Redshift, CitusDB, IBM, and Greenplum. With this integration, Brytlyt can seamlessly connect to and work alongside any product that has a native PostgreSQL connector.

Users can easily pair it with tools like Tableau and TIBCO Spotfire or integrate with Python and other common programming languages through the Brytlyt platform. This has huge benefits for both performance and usability compared to other GPU accelerated solutions.

PostgreSQL

Brytlyt offers all the great features, functionality and security that PostgreSQL already offers. Having an established and commonly deployed database foundation means essential features are already built in. From security to management processes, users can be up and running in hours without any added hassle. Most importantly, Brytlyt provides GPU-enhanced database functions. This ensures that you can create and run complex workloads and models that can be processed in milliseconds, for powerful new insights.

Ease of use

Brytlyt is easy to use and extremely sustainable. Users familiar with PostgreSQL will be able to use Brytlyt immediately, with no additional training.

AI integration

Many platforms offer an extracted AI or machine learning element to their analytics. Brytlyt takes this to another level with AI functionality fully embedded into the database. This means that, by integrating PyTorch, Brytlyt enables users to run PostgreSQL and AI workloads in situ and simultaneously, on the same data without extract or copy.

Security

PostgreSQL has advanced security options that enable users to ensure their critical data workloads are secured.

Geospatial analytics

Geospatial data is now available 24/7 through an increasing number of intelligent devices that collect location-based data. From site planning and geotargeting to ground and asset movement, to fraud detection and risk assessment, the number of use cases for location-based datasets are wide and varied. The on-demand availability of geospatial data is invaluable to businesses that want to leverage location-based data to optimise decision-making processes.

The traditional way to visualise geospatial data is through mapping with associated charts and graphs.

Until now, analysis tools have been client based, rendering maps by pushing data to the client first. However, this comes with large wait times and the inability to process workloads in real-time due to a lack of processing power on the client side.

The Brytlyt platform uses GPU technology to enable in-depth visual analytics on vast datasets, including geospatial analysis. It is a web-based tool which can be used with Brytlyt's database or any other compatible database. Businesses can exploit their data for increased value, turning raw information into business-ready intelligence. With visualisation graphics also rendered by the latest GPU technology, the dashboard is extremely responsive and intuitive, enabling users to derive insights in seconds. Geospatial mapping further unlocks more meaning and value; with the ability to drilldown into intricate details or discover widespread, complex patterns across different data sources.

By using the Brytlyt platform to run geospatial workloads, teams can:

- Reduce the need for costly site visits
- Ensure more accurate location-based risk assessment
- Increase revenue through targeted sales and optimal store locations
- Optimise logistics and operational procedures

Streaming

Geospatial data is not the only kind of data available 24/7. Today, all manner of data is generated at speed by sources worldwide. Traditional solutions to handling data sources are built with CPU technology to suit the simple pipeline of ingestion and processing data before it can be acted upon. Modern 'streaming data' solutions add further value to the ability to consume, enrich, archive, and analyse the data in motion in real time. If companies want to make the most of these massive datasets and introduce things like deep learning and millisecond response time, CPUs will not cut it.

Brytlyt was built to handle the demanding and rapidly growing data economy of today. At core of the platform sits the world's fastest database. It can consolidate, combine and analyse billion-row datasets in milliseconds to provide critical insights in context, at scale and at speed.

Because the database sits on PostgreSQL, it is highly compatible with commonly used technologies and provides comprehensive database functionality.

Brytlyt is also the only platform to use parallel processing for JOINS. This enables the identification of relationships between tables within datasets, for truly meaningful data analytics.

It can be used on its own, or paired with popular visualisation tools like TIBCO Spotfire, Tableau, MicroStrategy, and Microsoft's PowerBI.

Data loading and ingest

Brytlyt supports a wide range of tools that facilitate loading data to the GPU database. The data can be either ingested directly from the database storage or using external data sources.

The data loading is initiated from the interface of PostgreSQL, so the syntax is already familiar to data engineers. By using PostgreSQL foreign data wrappers (FDWs), Brytlyt also supports loading the data from external SQL data sources like Snowflake, AWS Athena, Google Big Query, Microsoft SQL Server, Redshift, and Databricks. Additionally, any solutions that use the PostgreSQL drivers can use Brytlyt as a data source.

Serverless architecture

Cloud computing provides incredible flexibility to access compute and storage infrastructure. But legacy software running on this infrastructure is not designed to support the on-off profile that a serverless pay-as-you-go needs, so typically requires hardware to be up and running constantly.

Serverless solutions like Brytlyt combine hardware, software, user session state and data using an on-demand 'pay-asyou-go' method. This all happens when the user connects to Brytlyt. When the user is disconnected, the hardware is made inactive by triggering an 'on-off' switch within the software. The framework of serverless computing is cloud-native so can be used to build and run applications without having to manage or maintain hardware, install software or manage security. A serverless solution requires a rewrite of the application software so that it can rapidly switch between an active and inactive state.

Using serverless computing means resources can be allocated and scaled on demand, and only used as needed.

Besides lower cost and an individually economical approach to the physical resources required, other benefits of serverless platforms are:

- Load balancing and capacity management is automated
- Software updates and security patches are immediately available
- Effortless scaling up or down depending on usage

Use Cases

Energy

Oil and gas companies are in the data business as much as they are in the hydrocarbon business. To stay ahead of the competition, they need to derive insights from the massive amounts of public data, private data, sensor, geolocation, drilling and seismic data they generate. Traditional interpretation methods are increasingly challenged by the volume of data, fewer experts in the industry and slow processing power.

As it stands, the industry faces several key problems when it comes to gathering the insight they need from the vast array of data available to them, but the two stand-out issues are:

Speed of analysing large datasets is important for any company leveraging analytics to support multi-million-dollar decisions

Public data solutions are typically disconnected from private data silos. No easy way exists to analyse an entire dataset of all available public data, blended with an operator's proprietary private data, of any type – at scale, with speed of thought analytics.

Brytlyt uses serverless GPU technology to power real-time visualisations and simplify existing tools like PowerBI, Tableau and TIBCO Spotfire. The speed and accuracy of seismic interpretation is critical in the exploration workflow for Oil and Gas companies. MultiGPU and multi-node GPU technology boosts throughput for visualisation and heavy computation, improving calculations of 3D seismic trace attributes and visual analysis of complex regional basins right at the interpreter's desk.

The high compute power, massively paralleled processors and high-speed memory of GPUs allow oil and gas companies to:

- Visualise and analyse petabytes of well location, and associated production data in milliseconds, for example
- Implement advanced algorithms to locate unique insights in seconds
- Use deep learning training on raw seismic trace data to accelerate exploration and discover faults in geology
- The result is dramatically reduced model processing cycle times and sharper images of region-of-interest datasets.

This can lead to more effective lease bidding, higher service revenues and, ultimately greater chances of increasing production whilst reducing environmental impact.

Telecommunications

Many telecom organisations are looking to introduce intelligent benchmarking services using crowd source data. However, this process involves a deep analysis of billion row, complex datasets. With legacy architecture, this is a slow and cumbersome process requiring a lot of manpower. Even with existing business intelligence tools like Tableau, Power BI and TIBCO Spotfire, businesses are struggling to increase their workloads without compromising on speed. Brytlyt is based on PostgreSQL technology, allowing a great deal of flexibility and familiarity for end-users. With its ability to work with existing business intelligence tools, the Brytlyt platform is able to use GPU-accelerated database technology to supercharge the analysis of large crowd source data.

Telecom companies are therefore able to improve their benchmarking services with instant insight and more accurate results displayed in stunning visuals.

Brytlyt also enables:

- Use of scripting and controls for deep insights and visualisations
- The use of pre-configured dashboards based on subject matter expertise from the Brytlyt customer related to telecoms

IoT and Smart Cities

Smart city schemes are an emerging way of monitoring transport and logistics in large built-up areas. However, the schemes typically require a lengthy list of KPI's (Key Performance Indicators) to ensure that any reports generated from location-based data sources are accurate and insightful. For this to be the case, reports have to be calculated in real-time. With legacy-based architecture and more traditional business intelligence models, these schemes can become slow running and ineffective.

Brytlyt provides a platform where organisations facing this problem monitor the same KPIs, but reports are updated in a few seconds, something not possible with a legacy database powered by CPU technology.

Conclusion

Business intelligence workloads are more demanding than ever before because there is more data available. The impact on top and bottom-line companies is bigger than it's ever been, putting immense pressure on business intelligence solutions, both from a visualisation perspective and the backend database functionality that powers these large datasets. Innovative new technologies like GPU databases and next generation analytics tools provides a transformative solution for companies looking to enhance their business intelligence pipelines. GPU technology is driven in part by the continuing growth of massive data workloads, the demand for deep learning, and the expansion into different markets. Larger workloads are pushing memory requirements to their limits, while parallel and deep tech are driving better hardware performance for compute and speed. As connectivity constantly increases, the amount of data produced also inevitably rises. Vast data sets are becoming the norm and companies require more advanced technology to fully exploit this.

Overall, GPU accelerated databases and analytics are quickly moving from a performance preference to a necessity. Their ability to handle data is unmatched and will continue to be enhanced due to massive market and use case demand.

The Brytlyt platform is built to work with the latest GPU technology. It benefits from unrivalled processing power, seamless accessibility, and an incredible user experience.

Its business intelligence capabilities allow for simplified handling of complex queries over billions of rows of data, and is optimised to deliver a responsive and intuitive user experience.

Through detailed and visually impressive graphics, maps and charts, Brytlyt can deliver an incredibly quick, clear visual analysis to better understand, interpret and even predict user and system behaviour in a range of sectors such as customer retention, network monitoring, identifying fraud, and energy – all working towards improving an organisation's decision-making processes. Advanced visualisation tools like Brytlyt are the ideal solution for fast-paced, dynamic markets, where time is of the essence.