



Compute Made For Analytics: The Key to Solving for the Compute Performance Gap

Why Compute Made For Analytics is essential for data analytics processing going forward

AUTHOR

RON WESTFALL
Senior Analyst + Research Director

MAY 2023

IN PARTNERSHIP WITH





Introduction

Getting Valuable Data Insights is Becoming More Costly and Complex

Digital transformation is driving exponential data growth, increasing heterogeneity of data sources, and growing data streaming. To keep pace, many organizations are rapidly expanding server counts, resulting in more costs and management complexity. The data layer is bigger than ever - expected to reach over [180 zettabytes by 2025](#). However, the ability of processing power to keep up with this data growth is increasingly constrained as Moore's Law reaches its outer limits. Data warehouses and data lakes are struggling to perform and meet the demands of today's immense data volumes, with up to two-thirds of data never actually getting analyzed. Practitioners, such as data analysts and data engineers, are having to force-fit their analytics problems into the limitations of current infrastructure instead of advancing design to optimize for their business workload requirements. The impact is that organizations are not realizing the full benefit of the expected efficiencies, cost-savings, or product and service enhancements that analytics were meant to help deliver.

The amount spent on analytics is significant. In 2022 the global data analytics market is projected to be \$41.39 billion, and is expected to exceed around \$346.33 billion by 2030. This gives the market a [projected compound annual growth rate \(CAGR\) of 30.41% during the forecast window 2022 to 2030](#). Database and data management software tools are the second-most procured pieces of software [according to our research](#), as organizations try to keep up with the demands of the growing data layer. As data grows, the cost of running queries grows as well, due to increased latency and more data processed to get the same result. Identifying the obstacles to seeing value from analytics, especially from large datasets, is set to be a top priority in the years to come.

Analytics Infrastructure is Creating Obstacles

Today's analytics infrastructure, that is the physical infrastructure of servers, storage, network, and supporting databases and software, is reaching its limits to efficiently support ever growing data volumes and processing needs required for modern analytics, such as BI, fraud detection, predictive analytics, and more. Challenges include:

- **Spiraling Server-Related Overhead Causes Inconsistent Analytics Adoption.** For most infrastructure teams, it's become common to spin up more servers or new clusters to handle large and growing data analytics volumes. However this comes at a price, as it's not possible for a business to spend an unlimited amount, and this often leads to prioritization of different department's analytics needs based on budgets, SLA, overall power and datacenter capacity. This can lead to shadow IT, that is departments spinning up their own analytics in the cloud. In addition, the investment of time, labor and resources required to stand up new compute and keep it running is considerable, and often comes at the expense of other pressing business priorities. Due to the high cost and long timelines of getting access to large data volumes for analytics, trade-offs are common, with decisionmakers often accepting or assuming that they will not get all of the data they need for their work. This means that data's value is not fully realized.
- **Compute and Storage Dichotomies.** In the cloud, we observe an architectural shift. Through the application of disaggregation methods, compute has become decoupled from storage. This has enabled easy scaling to support big data that is now hitting both budget and efficiency thresholds. Analytics workloads have been unable to keep up with today's increasingly disaggregated cloud architecture. One of the primary factors driving analytical underperformance, especially when workloads run in the cloud, is increased data transfer latency and greater networking cost.
- **Current Hardware and Software-Based Acceleration Approaches Are Showing Their Limits.** Central processing units (CPUs) can only do so much with large volumes of data, and are flattening out in terms of their ability to handle the processing power, memory, and bandwidth required relative to the amount of data and the speed of data coming at them. This makes CPUs inadequate for today's analytics because they are not built with a data-oriented architecture, not purposefully developed with analytics as a top priority, and have insufficient memory resources for optimizing analytics outcomes. Software-based analytics acceleration approaches are being attempted to fill the gap, but are often limited by their siloed, fragmented designs, and we see data warehouses and data lakehouse implementations showing weaknesses related to performance, particularly when it comes to structured data at a very large scale in data warehouses. Some databases vendors are taking time-consuming approaches, such as rewriting their code entirely in C++ to address these challenges.



Compute Made For Analytics: The Comprehensive Solution Key to Accelerating Analytics

Today, most companies are taking a fragmentary approach to their analytics environment (software, compute, storage, network), with little thought to how each element might be optimized for a particular workload.

In contrast, the Compute Made For Analytics approach is holistic, with each software, compute, storage, and network element co-designed to maximize analytical workload throughput throughout the system, from query to data store and back again, and in any environment such as on-prem, hybrid or cloud. In this way, organizations can increase the speed of analytics across the entire infrastructure, instead of in bits and pieces that eventually turn into a logjam where no optimization occurred.

From our perspective, Compute Made For Analytics can provide breakthrough benefits such as:

- **A Unified Way to Remove Data Bottlenecks Across Analytics Infrastructure.** By completely accounting for infrastructure needs, Compute Made For Analytics is key to ensuring that computing, network, and storage requirements don't lead to chokepoints in the design, implementation, or administration of analytics acceleration. This ensures that no data is left behind when completing analytical missions, and that greater insight value is gained with the ability to query over 90% of all data.
- **Consistently Improved Query Performance.** The query data scan rate (QDSR) is an emerging benchmark that measures query processing throughput agnostic of total compute, memory, networking, and storage capacity. This is important, because existing hardware- or software-based analytics acceleration solutions struggle to maintain a consistently measurable query data scan rate over the entire analytics workload since they are limited by the fact that operations take variable time to execute on a CPU. The ideal solution will process data consistently, ignoring query complexity and be limited only to the data throughput. An optimized system could sustain a QDSR that fully utilizes storage IO at full capacity, no matter how complex the query, for as long as there is demand on the system.
- **Lower Costs and Better Economics.** Through Compute Made For Analytics, we expect data platform teams can avoid the need to spin up exorbitant compute resources as data volumes grow – in effect managing cost growth vs. usage growth. Infrastructure teams also gain predictability over costs – such as those for software, compute, storage, and power – and can avoid some of those costs entirely through more efficient computing architectures.





NeuroBlade's Proven Solutions for Ecosystem-Wide Analytics Acceleration

NeuroBlade is the only proven approach on the market today that is optimizing the underlying compute infrastructure that supports analytics workloads. They have developed the SPU, the SQL Processing Unit, a custom processor designed for analytics workloads which delivers orders of magnitude, or 10-100x faster performance improvement over today's CPU's for analytics workloads. The SPU technology is further enhanced by the HEQS, the Hardware Enhanced Query System, which is an acceleration appliance that pre-integrates a number of SPUs, to allow seamless scale and fast deployment for any customers' needs. Finally, NeuroBlade is an easy to integrate solution, providing a no-SQL like data acceleration API, DAXL, which can easily snap into any database engine and immediately begin to add value with no visible change in the UI or datafiles or queries to the end-user.

The immediate benefit is improved query performance – customers can run queries up to 10-100x faster than today's solutions and can co-exist with / accelerate any software based acceleration already existing in a customers environment. NeuroBlade's solutions support a consistent QDSR, or query data scan rate enabling consistently faster query performance than achievable today. It is plug and play, enabling easy integration with analytics such as business intelligence dashboards, running loyalty program offers, and detecting fraud.

Futurum's Take: Compute Made For Analytics Key to Streamlining Analytics Performance

We believe that organizations will benefit substantially from accelerated analytics, which will allow them to fulfill their key objectives of gaining more insights from scaling and administering vaster volumes of data, no matter where the data is located.

Most CIOs and CTOs are now challenged to do more with less, especially during these tough economic times. Reviewing the foundation of their infrastructure for analytics, and looking to take a Compute Made For Analytics approach, is a way that they can simultaneously improve bottom line TCO costs even while handling increased user demand for timely data that drives revenue and competitive value.

Data analytics decision makers should quickly evaluate this emerging Compute Made For Analytics approach to achieve orders of magnitude leaps in their analytics performance. From our research and perspective, only Compute Made For Analytics can deliver the tenfold to hundredfold improvement over existing piecemeal solutions that are critical to streamlining analytics performance and optimizing acceleration layer capabilities.

What would you do when analytics costs are reduced by a factor of 10?

IMPORTANT INFORMATION ABOUT THIS REPORT

CONTRIBUTORS

Ron Westfall
Senior Analyst + Research Director

PUBLISHERS

Daniel Newman
CEO
The Futurum Group

Shelly Kramer
President + Chief Evangelist
The Futurum Group

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

This document, including any supporting materials, is owned by The Futurum Group. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of The Futurum Group.

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT NEUROBLADE

NeuroBlade is reinventing infrastructure for data analytics by setting a new standard for query processing speed and scale. Founded in 2018 and led by veterans of the systems, storage, and data analytics industries, NeuroBlade is rethinking processor design with its proven Compute For Analytics approach, built to maximize throughput at petabyte-scale so that queries can run multiple orders of magnitude faster. NeuroBlade is headquartered in Tel Aviv, Israel, and Palo Alto, California, with operations in Taipei, Taiwan, and New York City. [To learn more, visit www.neuroblade.com.](http://www.neuroblade.com)

ABOUT THE FUTURUM GROUP

The Futurum Group provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing, along with creating, distributing, and amplifying rich media content where all stakeholders read, watch, and listen. Our industry analysts conduct comprehensive qualitative and quantitative assessments of technology solutions, business issues, market drivers, and end-user demand dynamics across industry sectors. We work closely with your Analyst Relations, Product, and Marketing teams to establish and amplify your subject matter expertise and thought leadership while furthering your brand and product recognition by creating, distributing, and amplifying rich content where all stakeholders read, watch, and listen. [Read our disclaimer statement here.](#)

CONTACT INFORMATION

The Futurum Group | futurumgroup.com and futurumresearch.com
720-258-5591 | info@futurumgroup.com
Twitter: [@TheFuturumGroup](https://twitter.com/TheFuturumGroup) and [@FuturumResearch](https://twitter.com/FuturumResearch)

©2023 The Futurum Group and Futurum Research.

Company and product names are used for informational purposes only and may be trademarks of their respective owners.