Maximize IT Infrastructure — not budgets — with proven Micron® DDR4 Server DRAM

Maximize IT Infrastructure — not budgets — with proven Micron® DDR4 Server DRAM

Even as companies brace for economic headwinds, experts forecast that the amount of data created globally will more than double by 2025 — up to 180 zettabytes[1]. This means that more performance will be demanded from a smaller budget. The answer to the challenge is Micron® DDR4 Server DRAM. Data analytics workloads experience performance gains of up to 60% by simply doubling DDR4 Server DRAM capacity[2].

Adding memory capacity accelerates database analytics

Adding more capacity (while keeping the memory channels the same) or adding more capacity with more channels results in increased workload performance for data analytics applications such as Microsoft® SQL/TPC-H, Apache Spark™ TPC-DS and Spark Scalable Vector Machines (SVM).

MS SQL TPC-H capacity expansion

The popular TPC-H data analytics workload, using the MS SQL database, witnessed a 2.2x performance increase for a 3,000-scale factor problem size, when the available memory was doubled from 512GB to 1TB (spread over two sockets) on a 3rd generation Intel® Xeon® 32-core dual-socket system. The test setup compared 16x DDR4 32GB DIMMs for a total of 512GB against 16x DDR4 64GB DIMMs for a total of 1TB.*

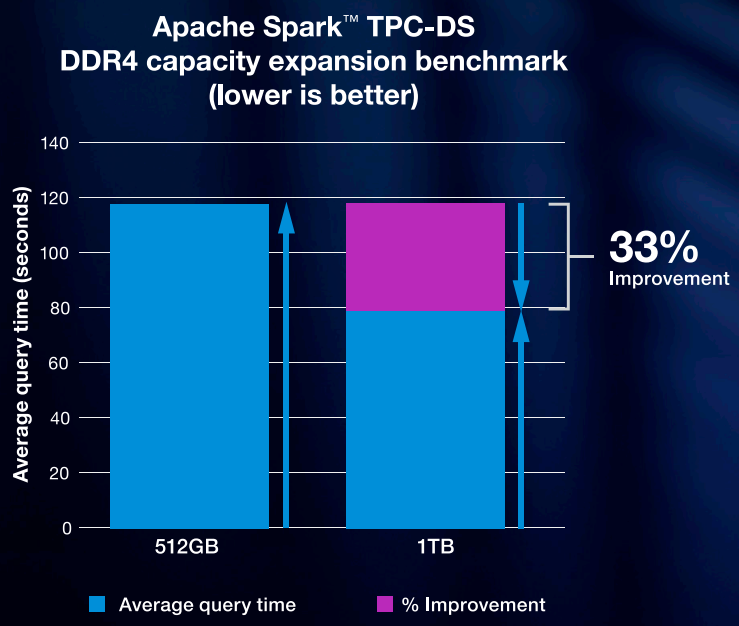

Results:

2.2x

Performance increase

when server's memory capacity was increased from 512GB to 1TB*

Benchmark notes:

TPC-DS is a popular benchmark for data warehouse analytics environments. TPC-DS with Apache Spark™ SQL achieved a 33% performance improvement when memory capacity was increased from 512GB to 1TB. The test was performed on AMD EPYC™ 7713 64 core dual-socket system for a TPC-DS problem size of 3,000 scale factor.

Apache Spark™ SVM – capacity expansion

Support Vector Machine (SVM) is a widely used supervised algorithm that learns to separate the input data into different classes by finding the hyperplane that maximizes the margin between the classes. SVM can be used in classification, data pre-processing, and anomaly detection pipelines. Benchmarks have shown that by increasing the server's memory capacity from 512GB (8x64GB DIMMs) to 1TB (16x64GB DIMMs) with Micron's DDR4, Spark's SVM execution witnessed a significant performance increase of around 3.2x for a 360GB input.**

Interestingly, this benchmark also revealed that Spark's SVM did not benefit from memory interleaving, suggesting that the performance improvement of SVM may not be solely due to the extra bandwidth available from the second socket's additional memory channels.

Apache Spark™ SVM (360GB input)
DDR4 capacity expansion benchmark
(lower is better)

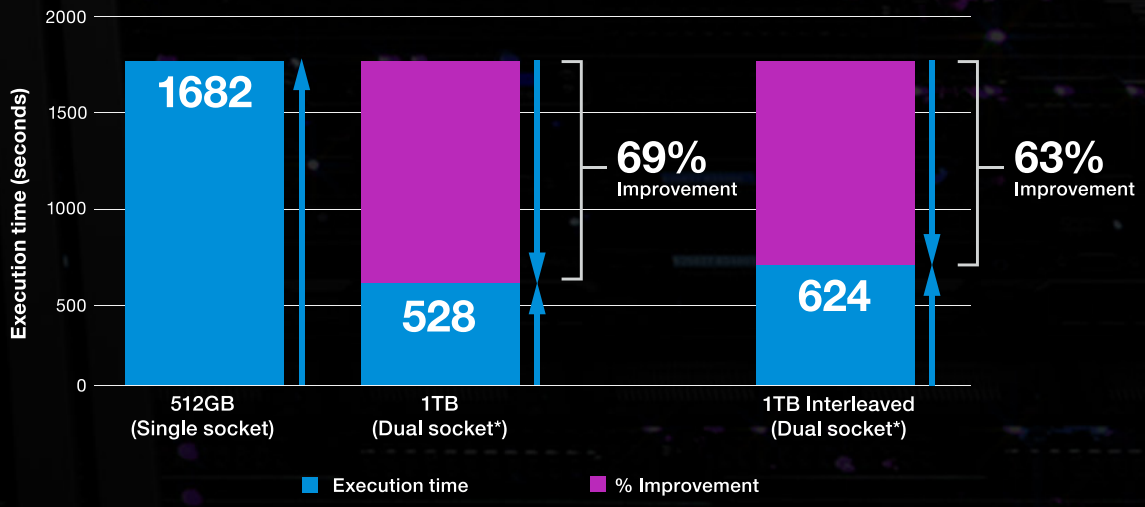

*Second CPU socket is disabled

Results:

3.2x

Performance increase

when server's memory capacity was increased from 512GB to 1TB**

Benchmark notes:

Support Vector Machine (SVM) is a common technique of supervised machine learning used in various classification, data pre-processing, and data anomaly detection pipelines.

Why it's useful to you

Given the experiment results, adding DDR4 memory capacity enables SVM to process larger inputs. SVM is suitable for dealing with high-dimensional data, which can be typically associated with large data inputs (100GB+). Also, the space complexity of SVM is okay (O(n) or lower (grows linearly with the input); which means it can handle large inputs. Therefore, it can be used in finance to process large historical data of stock prices; in healthcare for disease diagnosis based on body sensor data; in retail for customer segmentation and product recommendation; in marketing for customer churn prediction; in transportation for traffic flow prediction and vehicle detection; and energy for demand forecasting and load prediction.

Many of these applications can also be solved by state of the art (SOTA) deep neural networks (DNN). However, SVM can still be used for preprocessing the data before using a SOTA DNN solution, such as outlier detection, data cleaning, and dimensionality reduction.

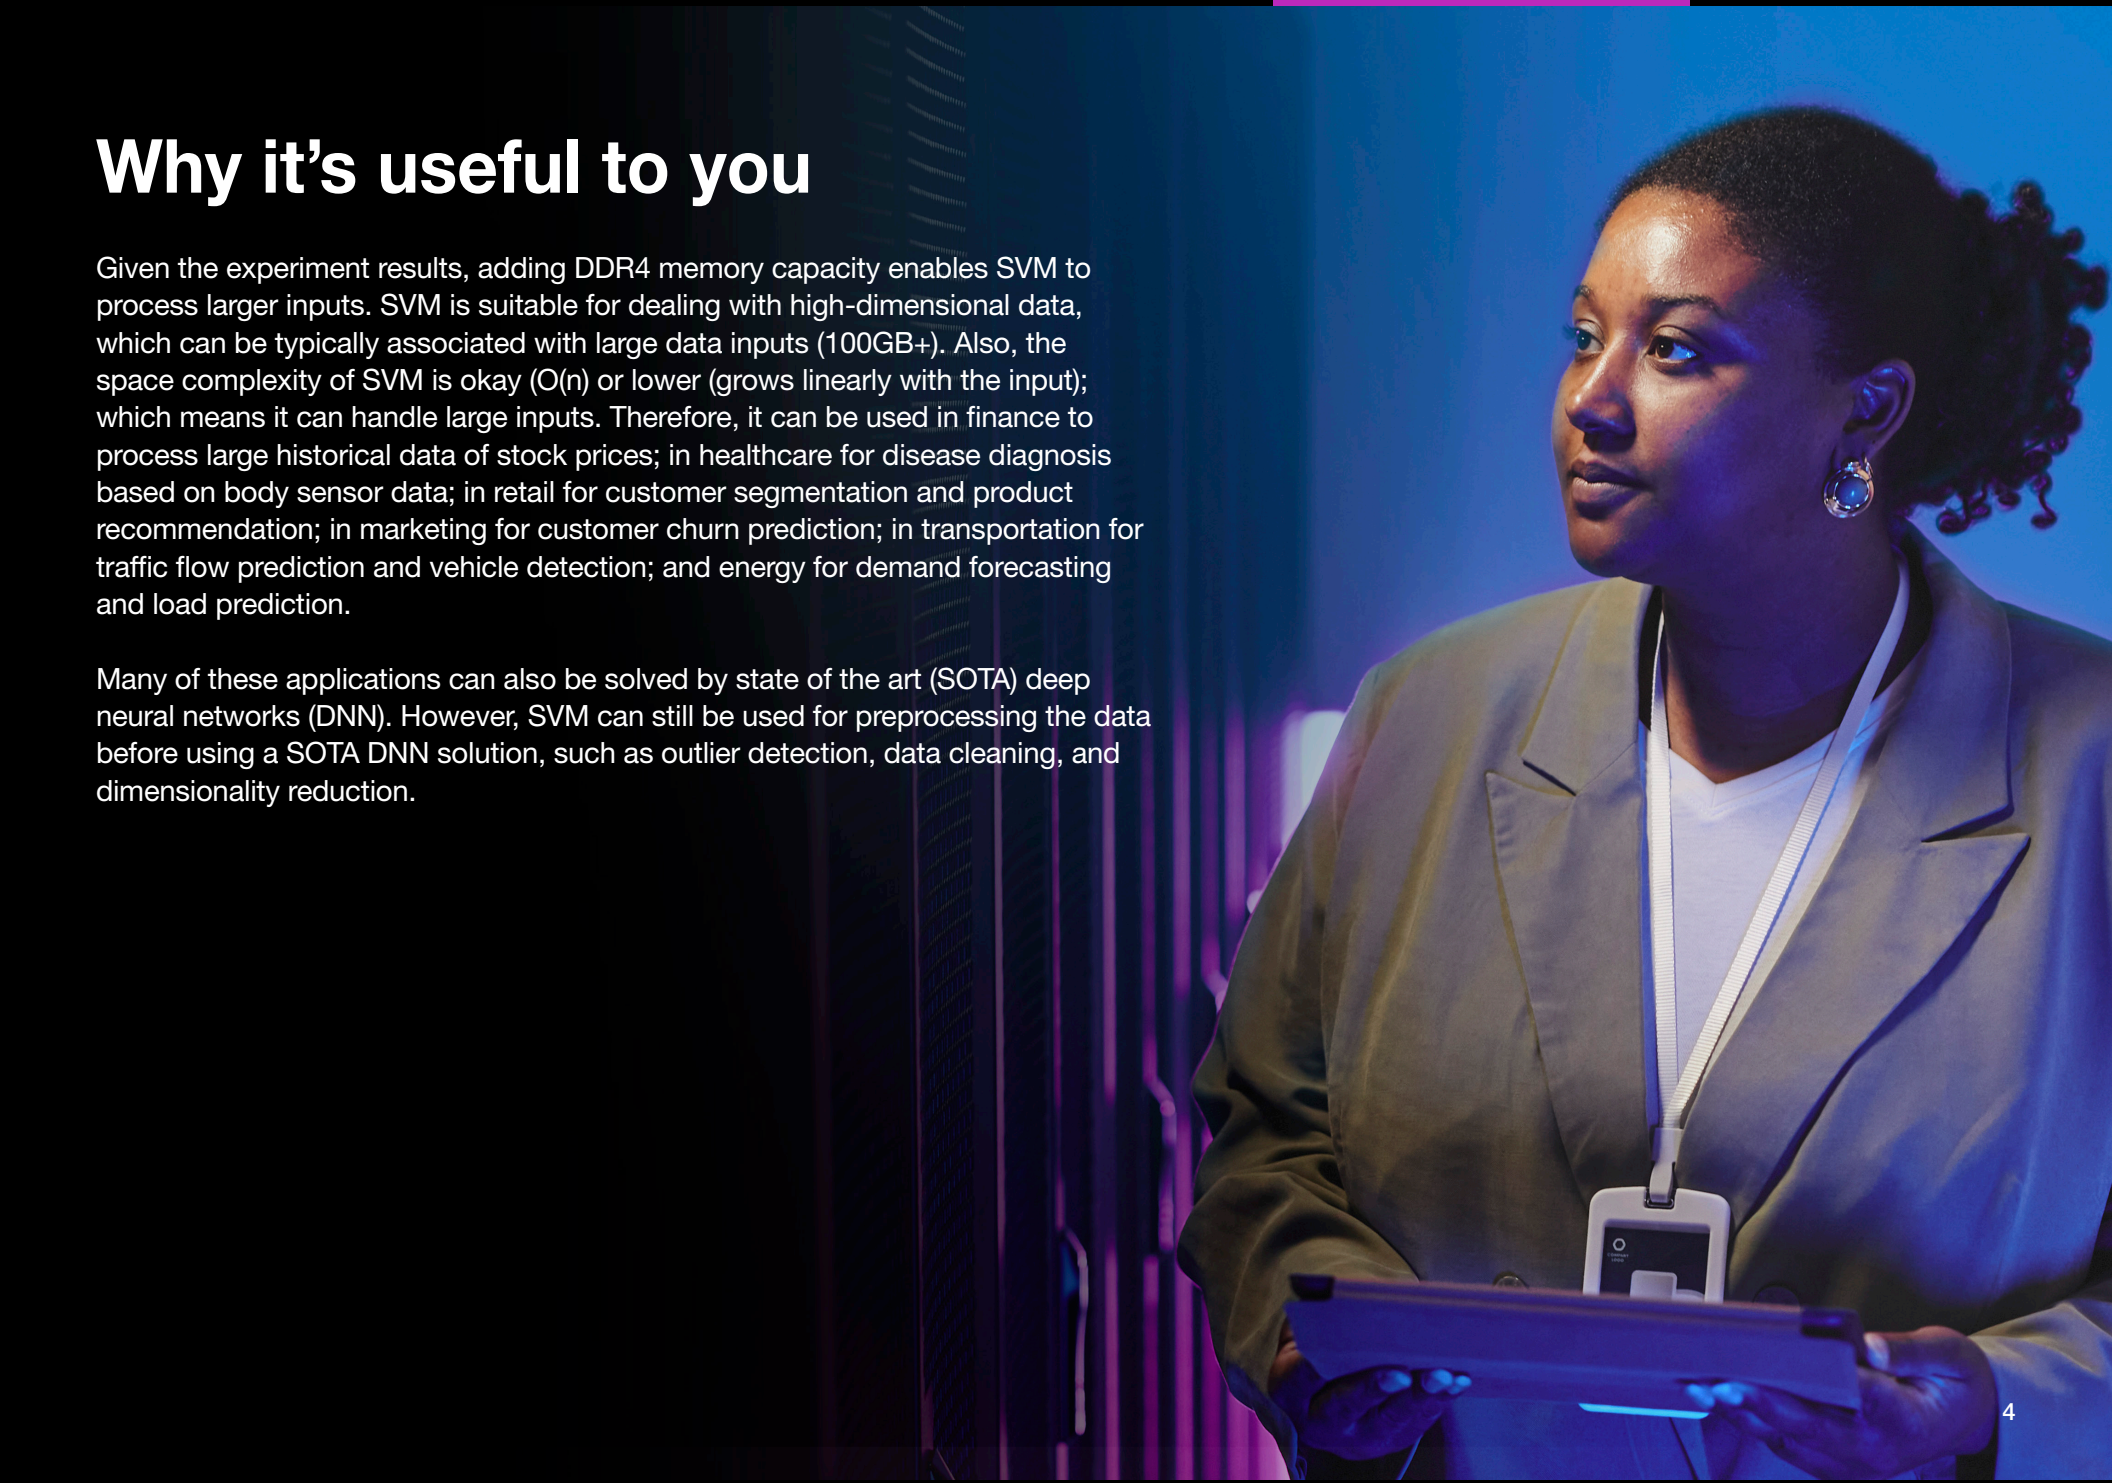

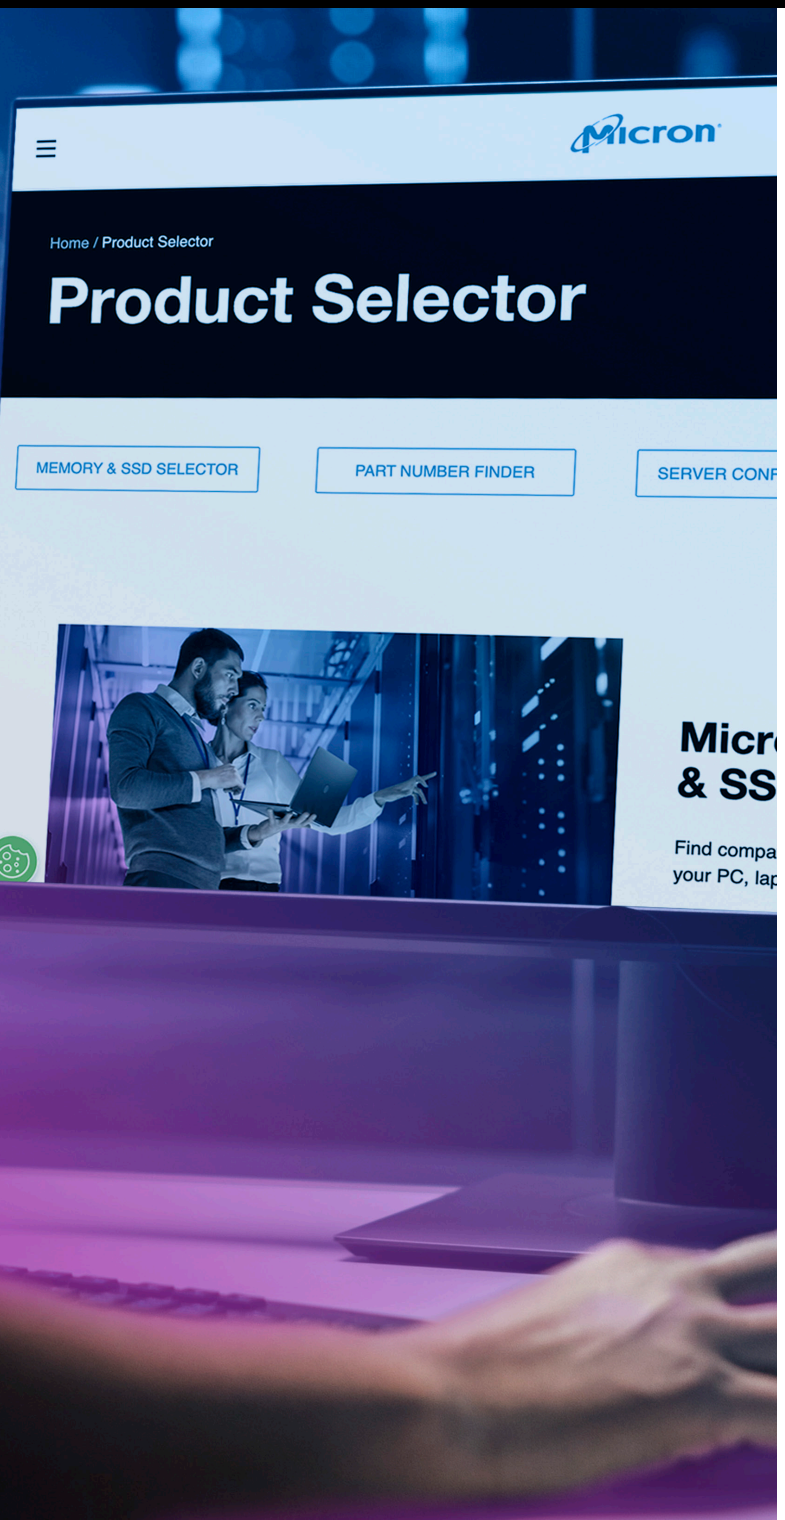

Micron DDR4 Server DRAM also extends value when purchasing new legacy platforms

The value of existing infrastructure can be extended with proven Micron DDR4 Server DRAM performance. In fact, data analytics workloads experience performance gains of up to 60% by simply doubling DDR4 Server DRAM capacity[2].

And Micron's improved DRAM and SSD Selector tool takes the guesswork out of memory and SSD upgrades, making it easier and quicker to find compatible server upgrades.

As economic headwinds blow, you can't afford to ignore the value of data, yet you can't strain your budget, either. Improving on last-gen tech with Micron DDR4 Server DRAM will get you through the storm not only intact but improved. Combine that with Micron's DRAM and SSD Selector, and you'll find the solution surprisingly easy.

Learn more at micronpg.com/serverDDR4

Maximize your IT infrastructure

Sources

[1] According to analyst firm IDC's Worldwide Global DataSphere Forecast, 2021–2025, business and consumer data has been amassing at a compound annual growth rate (CAGR) of about 23% since last year, with a 28% CAGR attributed to enterprises, and is expected to reach 180 zettabytes by 2025.

[2] Based on Micron internal benchmarks using TPC-H on a dual-socket, 32-core server equipped with a 3rd generation Intel® Xeon® Scalable Processor (SP) running MS SQL Server. DDR4 Server DRAM capacity was doubled for a 3,000-scale factor problem size. Actual results may vary.

Rev. 04/2023