

PlasmaENGINE® on PowerEdge R940xa

Abstract

This whitepaper evaluates the performance and efficiency of running PlasmaENGINE®, a GPU-based real-time stream processing software, on Dell EMC PowerEdge R940xa server. The objective of this whitepaper is to look at how the PowerEdge R940xa with its unique architecture is suitable for real-time processing and how customers can implement this solution for real-time streaming.

April 2019

Revisions

Date	Description
April 19	Initial release

Acknowledgements

This paper was produced by the following members of the Dell EMC SIS team and FASTDATA.io:

Author: Bhavesh Patel [Dell EMC], Tim Livingston, Amir Szekley [FASTDATA.io]

Support: Ritika Saini, April Berman, Ravo, Giacomo [Dell EMC]

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© April 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
1 Overview.....	5
2 Dell EMC PowerEdge R940xa server	6
3 Performance Benchmark.....	7
4 Conclusion.....	10
A Related resources	11

Executive summary

In this age of exponential data growth, organizations that can transform data into real-time insights will have the competitive edge in the market. Dell EMC's PowerEdge R940xa is tailor-made to accelerate data analytics by deploying GPUs for real-time stream processing for massive data sets. The PowerEdge R940xa's powerful quad socket design in tandem with associated GPUs is optimized to avoid PCIe bottlenecks to unleash uncompromised performance.

The R940xa offers up to 112 processing cores and up to 6TB of memory for consistently fast response times and achieves extreme acceleration by combining four CPUs with four GPUs in a 1:1 ratio. (Four doublewide GPUs or eight singlewide GPUs or FPGAs can be accommodated).

PlasmaENGINE® by FASTDATA.io is an incredibly innovative GPU-native real-time stream processing software. Utilizing the thousands of cores on a GPU, PlasmaENGINE® processes structured or unstructured streaming data up to three orders of magnitude faster than CPU-based alternatives like Apache Spark. Built to provide CTOs, product managers, developers, data engineers, scientists and architects with highly efficient real-time data analytics for business intelligence, PlasmaENGINE® aims to give companies the ability to accelerate operations and grow their business through unprecedented data processing speed and efficiency.

1 Overview

In order for every industry to take advantage of GPU platforms, an efficient software has to be written for it. FASTDATA.io developed exactly that: an equally performant and efficient software technology to transform data processing from collect > store > process to collect > process > store.

GPUs are exceptionally efficient in processing vast amounts of data because of the thousands of additional cores they have compared to CPUs (for comparison, NVIDIA V100 GPUs have 5,120 cores, compared to a high-end Intel(R) Xeon(R) Platinum CPUs, which has 28 cores per socket). GPUs provide greater throughput for operations that need to be performed on a tremendous amount of data concurrently.

PlasmaENGINE harnesses the power of the GPU and its many cores to process data streams in real-time at scale.

“Think of the GPU as a coin press machine, which can punch out 100 coins with a single operation every four seconds, whereas a CPU is a coin press which can punch out 1 coin per operation every one second. While the CPU has a faster “punch time”, the GPU can punch more coins per minute. This is the key difference between the GPU and CPU. The GPU is throughput oriented, while the CPU is latency oriented.”

Often in the data processing pipeline, bottlenecks arise when moving data. When using GPUs to process large amounts of data, it's especially important to be able to move data quickly from the CPU to the GPU. Dell EMC's PowerEdge R940xa and its groundbreaking 1:1 CPU-to-GPU architecture moves data with speed and efficiency, which PlasmaENGINE® brings to life in the benchmark below.

2 Dell EMC PowerEdge R940xa server



Figure 1: Front and rear views of the PowerEdge R940xa

The rapid increase in machine learning and artificial intelligence applications is changing everything about the way enterprise does business. With a powerful 4-socket and highly scalable 4U design, the Dell EMC R940xa Server is a great solution to power real-time GPU-powered stream processing for massive data sets.

The R940xa offers up to 112 processing cores and up to 6TB of memory for consistently fast response times. Add up to 12 NVDIMMs of memory or up to 4 direct-attached NVMe drives to maximize performance and minimize latency. In the R940xa each CPU has direct PCIe connectivity to a single GPU which results in extremely fast data processing in all 4 CPU-GPU subsystems and minimizes latencies to a level that was impossible to achieve earlier. Four doublewide GPUs or six singlewide GPUs or FPGAs can be accommodated.

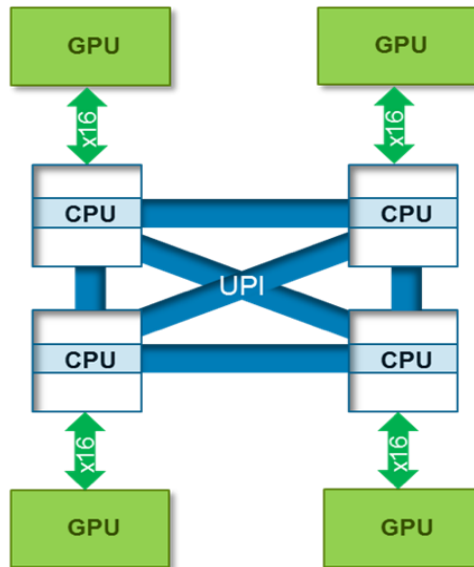


Figure 2: CPU: GPU High Level block diagram

3 Performance Benchmark

PlasmaENGINE® was up and running in minutes on the R940xa thanks to an easy to use Docker image.

From there, FASTDATA.io used the Haversine Benchmark -- a SparkSQL query used to calculate and compare distance between two GPS points on Earth -- to test the data processing capabilities of the R940xa.

The Haversine Benchmark simulates pipelines common in large Telecom companies, which takes gigabytes of CDR rows, each containing the location of both ends of a cell phone call, and then filters those calls based on distance as calculated by the haversine function.

The test was run with the following specs:

Machine: Dell R940xa

GPU: 4x V100 with 16GB each

CPUs: 4x Intel(R) Xeon(R) Platinum 8180M CPU @ 2.50GHz

Memory: 3TB

The benchmark measures how many records can be processed per second. We compared Spark running on the CPU with PlasmaENGINE® running on GPUs, and as you can see, PlasmaENGINE® processed a whopping 2.128 billion rows per second on the R940xa, compared to Apache Spark’s 10.49 million rows.

PlasmaENGINE® Rows Per Second (r/s)

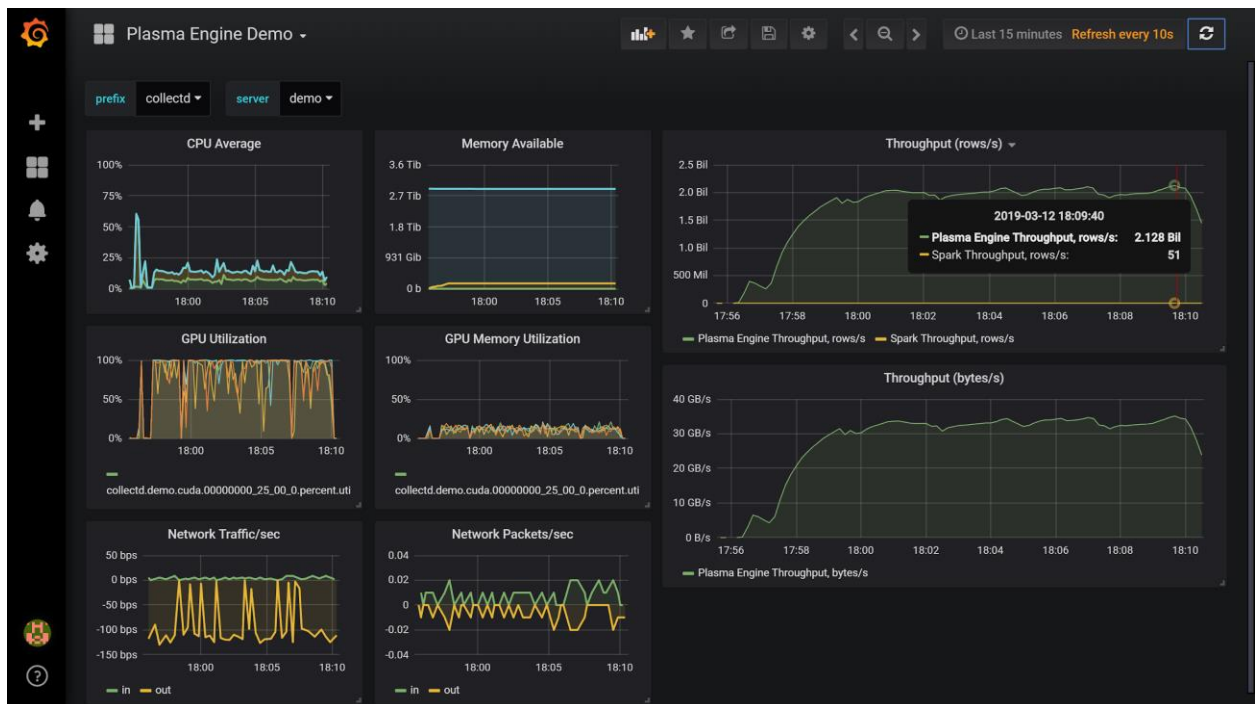


Figure 3: Rows per second

Graph Plot Points:

- **Top Left (CPU Average):** Average utilization of CPUs during benchmark test. (Note that CPU utilization is much higher on Apache Spark graph)
- **Middle Left (GPU Utilization):** Average utilization of GPU, one line per GPU. (Note that on PlasmaENGINE® it's heavily used, whereas on Apache Spark it's not used at all)
- **Bottom Left (Network Traffic/sec):** Network traffic (ignore because the test is local)
- **Top Middle (Memory Available):** System memory, not GPU memory.
- **Middle (GPU Memory Utilization):** Note that it's used when PlasmaENGINE® is running, but not when Spark is running.
- **Bottom Middle (Network Packets/sec):** Network traffic (ignore because the test is local)
- **Top Right (Plasma Engine and Apache Spark Throughput rows/sec):** Yellow line is Spark, Green line is PlasmaENGINE®. Throughput is the rows/sec processed by PlasmaENGINE® on the R940xa.
- **Bottom Right (Plasma Engine Throughput bytes/sec):** Throughput of bytes/sec processed on PlasmaENGINE® on the R940xa.

PlasmaENGINE® Bytes Per Second

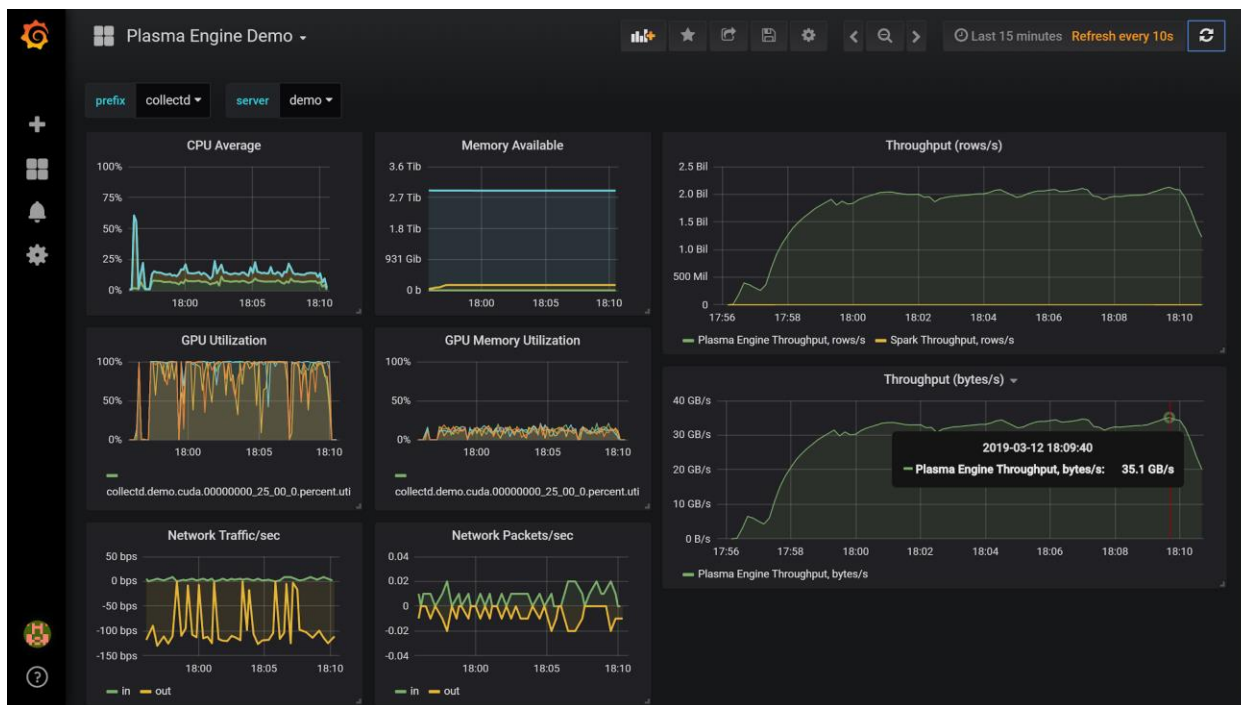


Figure 4: Bytes per second

Graph Plot Points:

- **Top Left (CPU Average):** Average utilization of CPUs during benchmark test.
- **Middle Left (GPU Utilization):** Average utilization of GPU, one line per GPU.
- **Bottom Left (Network Traffic/sec):** Network traffic (ignore because the test is local)
- **Top Middle (Memory Available):** System memory, not GPU memory.
- **Middle (GPU Memory Utilization):** Note that it's used when PlasmaENGINE® is running, but not when Spark is running.

- **Bottom Middle (Network Packets/sec):** Network traffic (ignore because the test is local)
- **Top Right (Plasma Engine and Apache Spark Throughput rows/sec):** Green line is PlasmaENGINE®. Throughput is the rows/sec processed by PlasmaENGINE® on the R940xa.
- **Bottom Right (Plasma Engine Throughput bytes/sec):** Throughput of bytes/sec processed on PlasmaENGINE® on the R940xa.

Apache Spark Rows Per Second

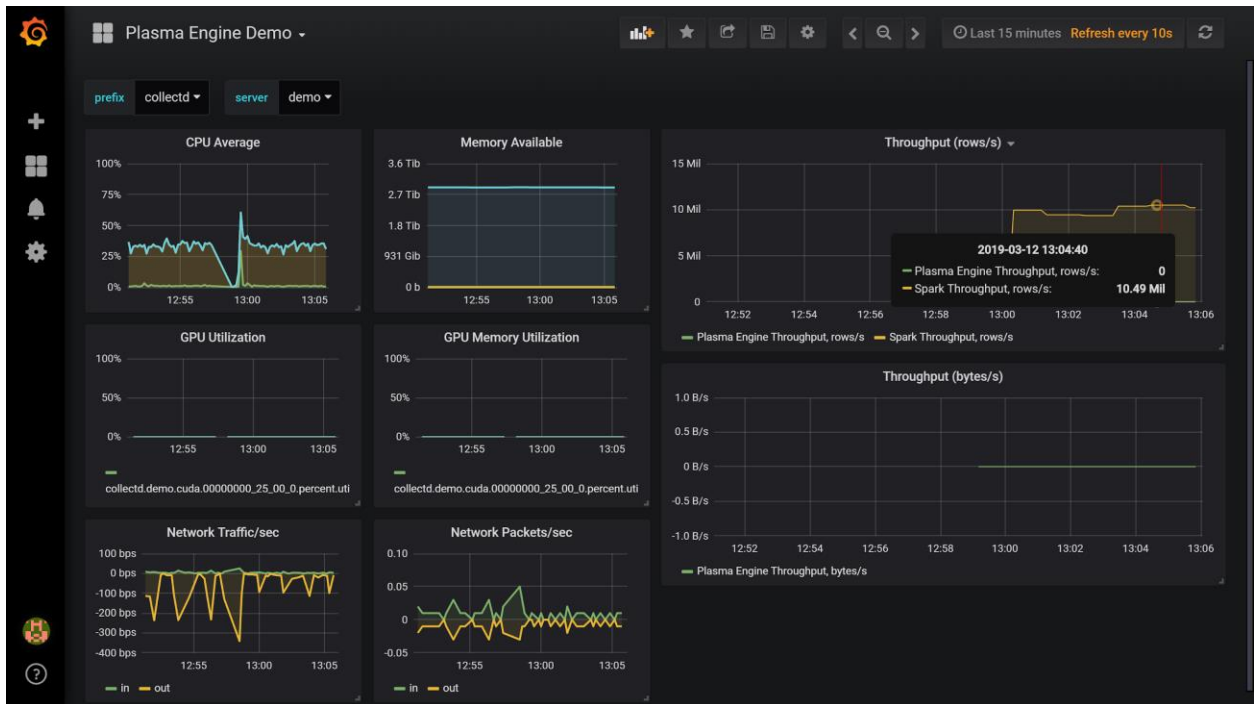


Figure 5: Rows per second

Graph Plot Points:

- **Top Left (CPU Average):** Average utilization of CPUs by Apache Spark during benchmark test.
- **Middle Left (GPU Utilization):** Average utilization of GPU. Note Spark doesn't use the GPU.
- **Bottom Left (Network Traffic/sec):** Network traffic (ignore because the test is local)
- **Top Middle (Memory Available):** System memory, not GPU memory.
- **Middle (GPU Memory Utilization):** Note that it's not used when Spark is running.
- **Bottom Middle (Network Packets/sec):** Network traffic (ignore because the test is local)
- **Top Right (Plasma Engine and Apache Spark Throughput rows/sec):** Yellow line is the throughput (rows/sec) processed by Apache Spark on the R940xa.
- **Bottom Right (Plasma Engine Throughput bytes/sec):** N/A

4 Conclusion

By leveraging the powerful capabilities of GPUs, the Dell EMC PowerEdge R940xa allows PlasmaENGINE® to process data with unparalleled performance. The 1:1 CPU to GPU ratio of the PowerEdge R940xa enables Dell EMC powerful server to maximize the performance between CPU and GPU and avoid PCIe bottleneck, which in turn allows PlasmaENGINE® to process a new-record 2.128 billion rows/second.

PlasmaENGINE® also processed over 35GB/s of data on the PowerEdge R940xa, or 8.75GB/s per card, which is very close to the 10GB/s limit of PCIe. These results are directly attributed to the industry-leading unique CPU and GPU implementation in the PowerEdge R940xa coupled with PlasmaENGINE's unparalleled data processing capabilities. It's important to note that PlasmaENGINE® achieved these results on a single server -- imagine how much additional data could be streamed with multiple Dell EMC PowerEdge R940xa servers.

A Related resources

- a) To see a PlasmaENGINE® demo, email sales@fastdata.io or visit <https://fastdata.io/>.
- b) To run a Proof of Concept(POC) on PowerEdge R940xa please contact the Customer Solution Centers: [internal](#) | [external](#) and you can schedule it [[HERE](#)].