Whitpaper

Using IBM with OmniSci GPU-accelerated Analytics

Benchmark Testing for IBM Power System S822LC for High Performance Computing with NVLink

Table of Contents

What would be possible if you could explore your data 100x faster?

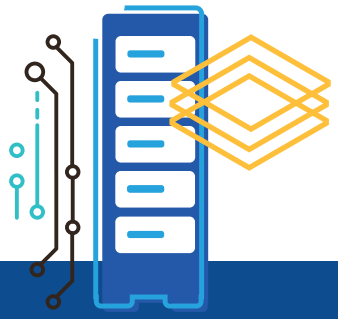

Executive Summary

Graphics processing units (GPUs) can make Big Data analytics 100 times faster than systems based entirely on Central Processing Units (CPUs). While the latest CPU can have up to 20 cores, today's GPU has more than 3,500. OmniSci, Inc. harnesses the parallel processing power of GPUs to deliver the world's fastest data queries and visualization. However, most analytic workflows still move data between CPUs and GPUs, and so much of the speedup provided by a GPU can be lost during transfers to a CPU. NVIDIA created NVIDIA® NVLink™ to reduce that bottleneck, and it is only available for CPU-to-GPU transfers on IBM POWER systems. This white paper provides OmniSci performance benchmark results on IBM Power System S822LC for High Performance Computing and shows how NVLink speeds data science workloads shared across NVIDIA® Tesla® P100 GPUs and IBM POWER8 CPUs. This paper is ideal for those looking to add general-purpose computing with GPUs to their existing analytic architectures that include CPUs.

IBM OpenPOWER LC Servers: Designed for Open Source Cognitive Computing

IBM designed OpenPOWER LC servers to power cognitive computing analytics that redefine how businesses operate. Mobile, cloud and Big Data give data scientists a flow of new opportunities for insight from social media, mobile gaming platforms, industrial IoT streams, virtual reality and artificial intelligence.

OpenPOWER LC servers were designed to capture that business value and smooth the IT challenges of integrating their data center and cloud resources. From the beginning, IBM architects designed OpenPOWER servers for data scientists and business analysts trying to keep up with that accelerating influx of data. Those end users needed a much faster feedback loop between their queries and the data they were analyzing. OpenPOWER servers shorten that feedback loop, allowing data scientists to make much faster queries, speeding both ad hoc discovery and the training of AI and machine learning models.

Much of that analytic acceleration comes from a new chip: the POWER8® processor with NVIDIA® NVLink™. The IBM POWER8 is a state-of-the-art CPU in its own right, but because it comes with NVLink, it works together seamlessly with the NVIDIA® Tesla® P100 GPU. NVLink accelerates CPU-to-GPU analytic pipelines up to 2.5 times compared to systems without NVLink, and that gives businesses far faster insight.

IBM OpenPOWER
LC Servers

IBM Power System S822LC for High Performance Computing: The Fastest GPUs and CPUs with the NVLink Interconnect

Of the four OpenPOWER LC servers on the market, this paper will focus on IBM Power System S822LC for High Performance Computing.

With the rise of GPU-powered analytics, data architects sought ways to connect GPUs to each other in 4-GPU systems to solve bigger problems faster. They also began connecting their new GPU compute resources to their existing CPU systems. As the speed and size analytics increased, the load exposed a clear latency choke point: bandwidth for CPU-to-GPU data transfers.

NVIDIA's NVLink breaks that bottleneck. It is a high-speed interconnect for communication between POWER8 CPUs and NVIDIA® Tesla® P100 GPU accelerators. Tesla P100 GPUs are comprised of 15.3 billion transistors and 3,584 CUDA cores. IBM Power System S822LC for High Performance Computing combines the parallel processing power of four Tesla P100s with two Power8 CPUs. S822LC servers are the first to provide the NVLink interconnect between CPU and GPU, and they deliver up to 2.5x the data transfer bandwidth of PCI-E x16 3.0. IBM designed IBM Power System S822LC for High Performance Computing with NVLINK to provide the highest computational throughput for analytic and machine learning workloads, ensuring that GPUs can be fed data as quickly as possible.

In fact, this is the only configuration on the market that combines NVLink, Power8 processors and Tesla GPU processors. The following image shows the power of that relationship. NVIDIA's NVLINK technology transfers data between CPUs and GPUs at 80 gigabytes per second, which almost matches the 115 gigabyte transfer rate between each CPU and system memory.

IBM Power System S822LC for High Performance Computing is the first to provide the NVLink interconnect between CPU and GPU, and it delivers up to 2.5X the bandwidth of PCI-E x 16 3.0.

Detailed Diagram of IBM Power System S822LC for HPC

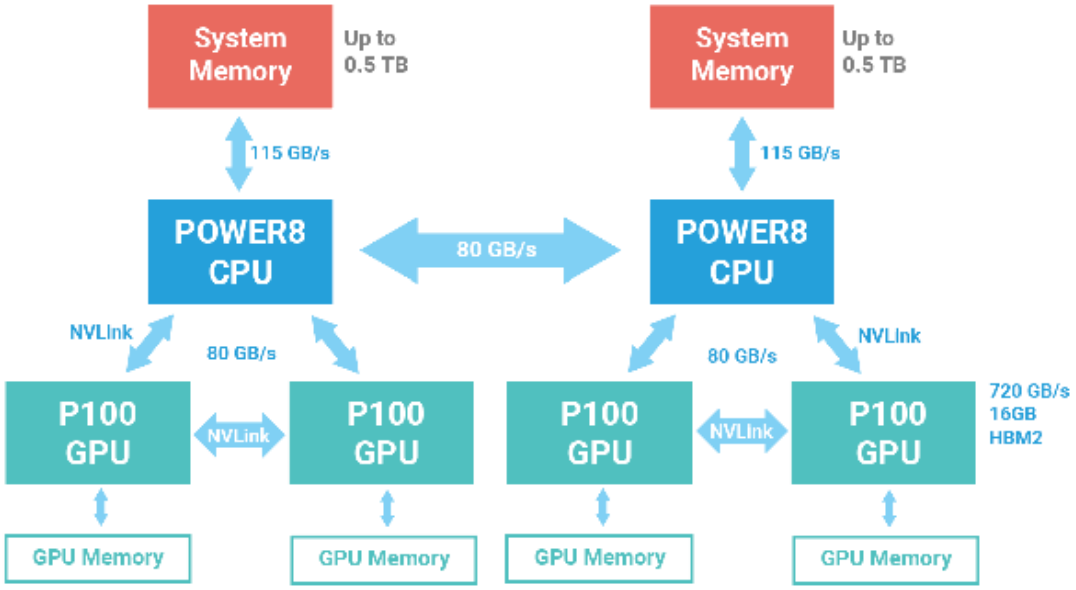

Source: https://www.ibm.com/blogs/systems/ibm-nvidia-present-nvlink-server-youve-waiting

NVLink between CPUs and GPUs enables fast memory access to large datasets in system memory

IBM Power System S822LC for High Performance Computing servers give enterprises the type of speed at scale that data scientists and business leaders need to stay ahead of new business challenges. These analytics require both the scale of big data and sub-second query speed, for use cases like:

- Telecommunications Network Reliability Analysis
- Vehicle Telematics Analysis
- Investment Banking Alternative Data Insights
- Utility Smart Meter Analysis
- Oil & Gas Well Log Analysis
- Pharmaceutical Clinical Trial Analysis
- Cyber Incident Investigation
- Defense and Intelligence GEOINT

OmniSci with IBM Power System S822LC for High Performance Computing: A Powerful Coupling for Interactive Analytics

IBM Power System S822LC for High Performance Computing delivers big data processing power at truly interactive speeds, over a hybrid GPU-CPU system. Because OmniSci makes Big Data analytics far faster on both GPUs and CPUs, OmniSci running on those systems delivers the fastest analytic capabilities even as it abstracts the underlying hardware that makes that possible. As far as the analyst or data scientist is concerned, it's just very fast and easy to use.

Todd Mostak, OmniSci founder and CEO, first understood the need for general-purpose computing with GPUs when he was a graduate student at Harvard. Todd wanted to analyze hundreds of millions of tweets for his research

on the Arab Spring, but his queries took hours to run. Driven by his end-user demand for fast and interactive data discovery, Todd then joined the MIT Computer Science and Artificial Intelligence Lab (CSAIL) as a research fellow in the database group. That research led Todd to create an analytic database (OmniSci Core) with an integrated visualization client (OmniSci Immerse). From the beginning, Todd architected the two platform components to seamlessly work together to make optimal use of a GPU's unique speed and visualization advantages. That work became the foundation of OmnicSci, Inc.

OmniSci Core is open source, designed from the ground up to leverage the massive parallelism of GPUs. By efficiently combining the parallel power of multiple GPUs per server, OmniSci Core can execute queries over billions of records in milliseconds. That power and speed can scale with distributed computing across multiple highly-available (HA) servers. Since OmniSci Core was open sourced in May 2017, partners like IBM and NVIDIA have rallied behind OmniSci Core as a standard for GPU analytics.

OmniSci Core is an open source platform designed from the ground up to leverage the massive parallelism of GPUs.

OmniSci Immerse is the GPU-powered visualization client, designed from the beginning to take advantage of the lightening-fast SQL and in situ rendering capabilities of OmniSci Core. Every time a data scientist interacts with the OmniSci Immerse dashboard, it generates a set of SQL queries to OmniSci Core and the dashboard refreshes in milliseconds—even with tables containing billions of rows. Its speed comes from the fact that a GPU node contains more than a thousand times more processing cores than a CPU node, and those cores can process data in parallel.

In addition, OmniSci Immerse can leverage the native rendering engine built into OmniSci Core. For example, instead of having to send a billion points from server-to-client to render a pointmap, it can request a rendered PNG from the server. Unlike typical business intelligence systems, OmniSci Immerse lets the analyst instantly generalize massive datasets at whichever level of detail is useful, without downsampling or pre-aggregating any individual datapoint.

That GPU horsepower makes OmniSci Immerse a far superior window through which data scientists or business analysts can work with their data. They can explore data in its raw form. Any time an analyst changes one data parameter, all other charts, graphs and maps cross-filter to redraw the entire visualization dashboard in that new context. As the analyst becomes more familiar with the data, she can easily update her dashboard and add new visual objects. Compare that speed of iteration to the status quo on a CPU system. It might take minutes or hours to script and run each query on a billion-row dataset and then more time to visualize and interpret the results.

OmniSci Benchmark Testing of IBM Power System S822LC for High Performance Computing Demonstrates the Power of NVLink

The OmniSci collaboration around IBM Power System S822LC for High Performance Computing demonstrates the superior performance and real-world impact of coupling two of the highest-performing processors: the Tesla P100 and POWER8.

In June 2017, OmniSci ran preliminary benchmark testing to quantify the acceleration of SQL queries on Power processors, versus X86 CPUs that do not include NVLink.

The results showed an average 65% acceleration over the group of 18 benchmark SQL queries, with an acceleration range from 25-126% across the entire group.

Benchmark System and Query Specifications

The following table shows the characteristics of the two systems tested.

Test system and configurations with NVLink ("S822LC"):

- IBM Power System S822LC for High Performance Computing
- GPUs: 4 x NVIDIA Tesla P100
- CPUs: IBM Power8, 2 x 10 cores
  Memory: 512GB

Test system and configurations without NVLink ("X86"):

- Intel® Xeon® CPU E5-1650 v3 @ 3.50GHz
- GPUs: 4 x NVIDIA Quadro P6000
- CPUs: Intel X86, 6 cores
- Memory: 256GB

Query	Query Text
Q1	select count(*) from ##TAB##
Q2	select carrier_name, count(*) from ##TAB## group by carrier_name
Q3	select carrier_name, avg(arrdelay) from ##TAB## group by carrier_name
Q4	select origin_name, dest_name, avg(arrdelay) from ##TAB## group by origin_name, dest_name
Q5	select date_trunc(month,dep_timestamp) as ym, avg(arrdelay) as del from ##TAB## group by ym
Q6	select dest_name, extract(month from dep_timestamp) as m, extract(year from dep_timestamp) as y, avg(arrdelay) as del from ##TAB## group by dest_name,y,m
Q7	select count(*) from ##TAB## where origin_name = 'Lambert-St Louis International'
Q8	select count(*) from ##TAB## where origin_name='Lambert-St Louis International' and dest_name = 'Lincoln Municipal'
Q9	select uniquecarrier,flightnum,dep_timestamp,dest_lat from ##TAB## where origin_name = 'Lambert-St Louis International' and flightnum=586
Q10	select origin_name, dest_name, avg(arrdelay), avg(depdelay), avg(arrdelay * depdelay) from ##TAB## group by origin_name, dest_name
Q11	select uniquecarrier,flightnum,dep_timestamp,dest_lat from ##TAB## where origin_name = 'Lambert-St Louis International' and flightnum=586 limit 5000
Q12	SELECT ##TAB##.carrier_name as key0,AVG(##TAB##.depdelay) AS x,AVG(##TAB##.arrdelay) AS y,COUNT(*) AS size FROM ##TAB## WHERE ((##TAB##.dep_timestamp >= TIMESTAMP(0) '1996-07-26 16:30:06' AND ##TAB##.dep_timestamp < TIMESTAMP(0) '1997-05-16 16:30:06')) GROUP BY key0 ORDER BY size DESC LIMIT 50
Q13	SELECT COUNT(*) as val FROM ##TAB## WHERE ((##TAB##.dep_timestamp >= TIMESTAMP(0) '1996-07-28 00:00:00' AND ##TAB##.dep_timestamp < TIMESTAMP(0) '1997-05-18 00:00:00'))

Benchmark Results

The following table shows the benchmark results for "first-run" queries on IBM Power System S822LC for High Performance Computing, versus an x86 system. We measured an average speedup of 65% across the 13 queries, with a range of speedups between 25% for Q1 and 126% for Q12. Much of the speed is due to the use of NVIDIA NVLink on IBM Power System S822LC for High Performance Computing, versus the slower GPU-to-CPU data transfers on the traditional PCIe bus used on x86 system.

Query #	Minsky First Run (ms)	X86 First Run (ms)	Speedup
Q1	312	390	25%
Q2	282	542	92%
Q3	692	962	39%
Q4	846	1121	33%
Q5	1110	1600	44%
Q6	202	336	66%
Q7	481	644	34%
Q8	258	339	31%
Q9	81	156	93%
Q10	200	359	80%
Q11	170	336	98%
Q12	786	1775	126%
Q13	108	197	82%
AVERAGE	425	674	65%

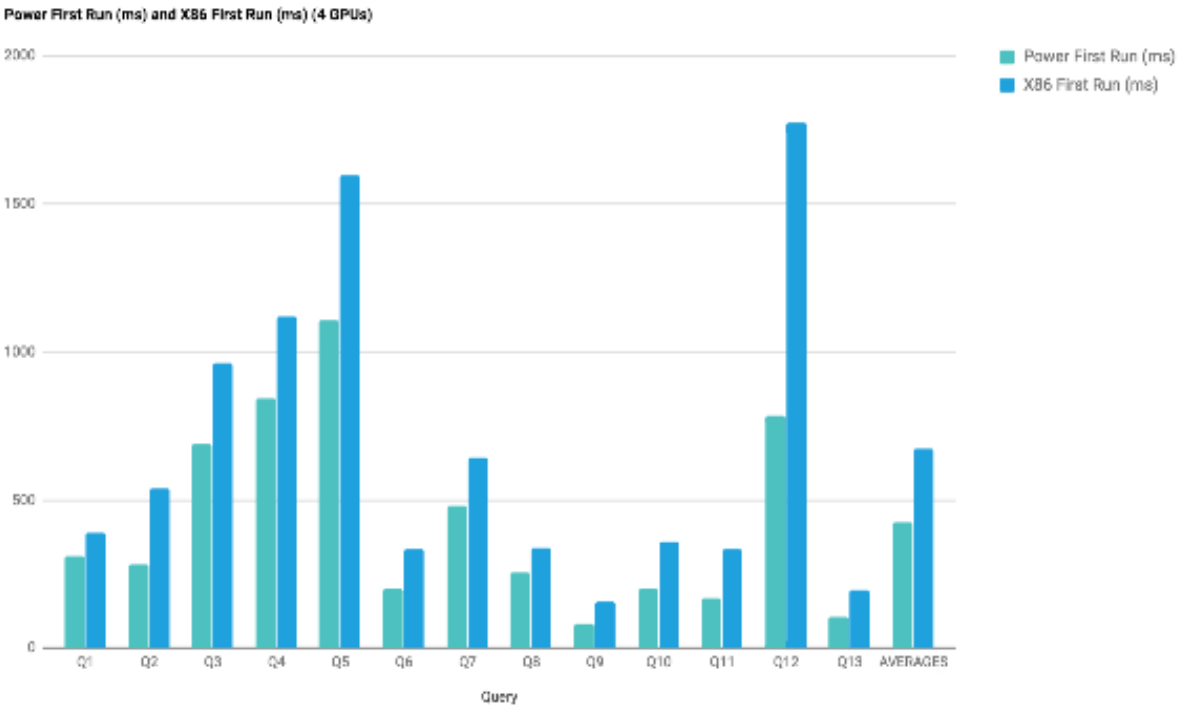

The GOAI Initiative: A Coalition for GPU Analytics, Standardizing on OmniSci Core

Only rarely do advances in hardware and software come about in well-timed, well-coordinated ways. IBM created OpenPOWER LC servers in anticipation of future progress from the open source coalition striving to make analytics easier on GPUs. NVIDIA anticipated that momentum and developed a new breed of GPU that could seamlessly work within CPU platforms, and NVIDIA created NVLink to remove the GPU-to-CPU communication bottleneck that was slowing that progress.

OmniSci, Inc. was founded to provide the software for those extreme, interactive analytics on the GPU. This paper focuses primarily on the interaction between the hardware components. There is a need for integrated cooperation between the various software solutions needed for GPU analytics. For this, OmniSci partnered with H2O.ai and Anaconda (formerly Continuum Analytics) to found the GPU Open Analytics Initiative (GOAI). The purpose of GOAI is to create common data frameworks enabling developers and statistical researchers to accelerate data science on GPUs. Along with founding members, others have joined the GOAI, including: BlazingDB, Graphistry, and Gunrock.

Along with cofounding GOAI, OmniSci open sourced OmniSci Core in May 2017. As the only open source analytic database architected on GPUs, OmniSci Core provides a consistent, open standard for new software development to take advantage of GPU-accelerated analytics. The parallel processing power of GPU hardware enables OmniSci Core to query billions of rows in milliseconds using familiar SQL queries.

OmniSci Immerse is the world's fastest visual analytics client. Natively integrated with the record-breaking query performance and rendering capabilities of OmniSci Core, OmniSci Immerse allows analysts to visually interact with their biggest SQL datasets to discover correlations, trends and anomalies. With immediate visual feedback and cross-filtering across maps, charts and graphs, data scientists and business analysts experience time-to-insight they've never had before. Those speeds aren't possible with a similarly sized CPU system. With faster time-to-insight comes the ability to answer far more questions in a much shorter period of time.

Another paradigm shift comes from a dramatic reduction in the time required to prepare machine learning and artificial intelligence models. GPUs are much faster than CPUs, but most data scientists still work using CPU systems. One of the most tedious and time-consuming parts of building a machine learning or AI model is feature engineering—"the process of using domain knowledge of the data to create features that make machine learning algorithms work."

The first GOAI project is the GPU Data Frame (GDF). The GDF is a common API that enables efficient interchange of data between processes running on the GPU. Users of the OmniSci Core database can output the results of a SQL query into the GDF, which then can be manipulated by the Anaconda's NumPy-like Python API or used as input into the H2O.ai suite of machine learning algorithms—all without additional data manipulation. Finally, the results of the best ML model can be visualized again in OmniSci Immerse, serving as an ongoing dashboard to visually compare predicted outcomes to actual results.

Researchers use OmniSci Immerse to explore data on OmniSci Core. Once they find a promising list of features, they move that information easily into H2O which generates hundreds of machine-learning models and chooses the best one.

As the GOAI open source collaboration grows in the coming years, much of that innovation will come back to data scientists analyzing data with OmniSci Core and OmniSci Immerse, running on the IBM Power System S822LC.

Bring IBM Power System S822LC for High Performance Computing, NVIDIA NVLink and OmniSci to Your Team

Learn how your team can take advantage of GPU-accelerated analytics to answer questions that have been out of reach before the rise of GPUs. Contact IBM and OmniSci today to learn more about the unprecedented power of OmniSci on IBM Power System S822LC for High Performance Computing.