

Accelerating Data Analytics for Faster Time-to-Insight and Improved ROI

Article

The world’s data is doubling each year, making the data science and machine learning applications that leverage this data the world’s largest computing segment. A broad swath of enterprise applications are looking to incorporate AI and machine learning capabilities to deliver better performance to their end users. The desired improvements need to come in the form of more accurate and faster analytics and predictions from the machine learning models. Faster and better data analytics performance translates into billions in profits and significantly better return on investment (ROI).

Challenges of end-to-end data processing

The need to develop better data analytics and machine learning models cannot be overstated. Furthermore, this is not an isolated task. A comprehensive view of the end-to-end data processing flow is required to put all the necessary parts in place to address the need for improved models. Such data flow pipelines are made up of multiple tasks and stages, from initial data engineering to acquire, ingest and store ever-increasing amounts of data, to data science for creating and evaluating models, and finally to deployment of best models into production.

Depending on the type and amount of data, the intended application domain and the breadth of available skills, each stage can pose considerable challenges because of the complexity of underlying data engineering and data science workflows.




Figure 1. End-to-end data processing workflow

Lenovo data management solutions

At Lenovo, we are collaborating with industry-leading partners to create data processing and data management solutions suitable for various stages of the data flow. These solutions include:

- Data ingest and storage repositories for structured and unstructured data
- Compute clusters for machine learning and deep learning model training
- Tools for deployment of trained models in inference across data centers and edge IT infrastructures

The [Lenovo Big Data Reference Design for Cloudera Data Platform \(CDP\)](#) on ThinkSystem Servers is one example of a solution building block. Lenovo has been delivering validated designs for big data clusters that combine CDP (and its predecessors) with targeted Lenovo systems to enable a whole host of big data use cases.

Using the NVIDIA platform with GPU-accelerated clusters for Data Analytics and AI model training is another example. Learn about the offerings at the [Lenovo Analytics and AI page](#).

Parallel processing with GPUs

With ever-growing data volume, the need for accelerating data processing has continued to increase. In particular, using parallel computing capabilities of GPUs has been shown to speed up data processing, data management and deep learning model training. While some enterprise applications require deep learning models, a majority of them can be improved using a broad set of traditional and emerging machine learning techniques. This is where attention is required to broaden the applicability of acceleration.

How can teams of data engineers handle large so-called big-data platforms and enable their data scientist colleagues to do faster model development and evaluation?

Fortunately, there is good news on this business challenge front. Lenovo partners NVIDIA and Cloudera have integrated [NVIDIA RAPIDS accelerated data science libraries](#) on Cloudera Data Platform (CDP), enabling GPU-accelerated [Apache Spark 3.0 applications](#). Apache Spark included in CDP has been a workhorse for numerous data analytics tasks such as batch/real-time streaming, data warehouse and machine learning among others. Accelerating Spark with GPU-enabled computation is the next leap forward in helping enterprises achieve the goal of faster and better model development.

Faster Training and ROI with GPU-Accelerated Apache Spark

Lenovo has been collaborating with NVIDIA and Cloudera on validating the performance of GPU-accelerated Apache Spark 3.0 running on CDP on Lenovo systems connected with NVIDIA Networking. The early results are very exciting. Not only are we seeing significant performance improvements, but also price-to-performance benefits.

Ultimately, price-performance is the metric about which customers often care most. Can I get a better return on my investment? That is, can I get more performance if I invest in more capable systems? In fact, the results point to **5X or more reduction in training times** and **3X improvement in ROI**.

All this is surely leading to faster and better insights, improved ROI and ultimately, a much better business outcome for enterprises and their end customers.

Lenovo is proud to enable a leadership hybrid cloud data analytics solution comprised of Lenovo's NVIDIA-certified ThinkSystem servers with NVIDIA A100 and A30 GPUs and Cloudera's CDP Private Cloud Base 7.1.6. These servers have been validated for running accelerated workloads with optimum performance, manageability, scalability, and security.




Figure 2. Lenovo-NVIDIA-Cloudera hybrid cloud data analytics solution

More detailed performance validation work will be published soon. Stay tuned. This is just the start. We are planning to address the aforementioned broad spectrum of enterprise applications that are hungry for data analytics and ML acceleration.

Related product families

Product families related to this document are the following:

- [Artificial Intelligence](#)
- [Cloudera](#)
- [GPU adapters](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1510, was created or updated on August 5, 2021.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1510>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1510>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

Other company, product, or service names may be trademarks or service marks of others.