



# ACCELERATE YOUR DATA-DRIVEN OPERATIONS WITH EASE

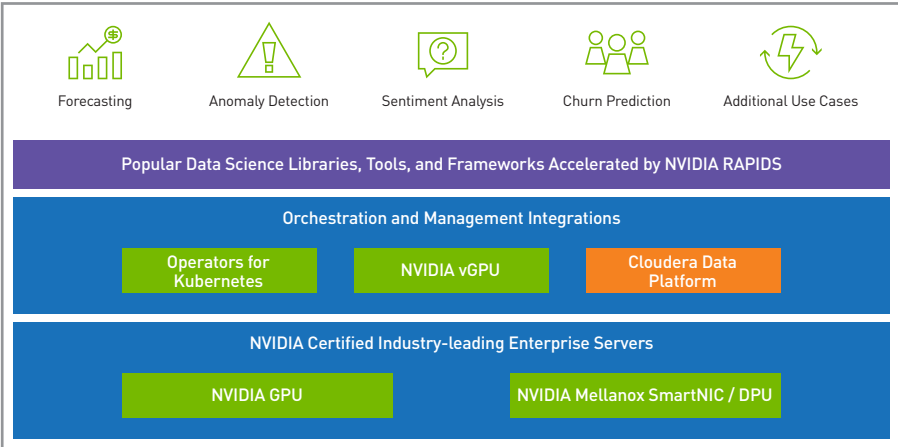
## Addressing the Challenges of Modern Data-Driven Enterprises

Today, data science has become the world's largest compute segment. Modest improvements in the accuracy of analytics models translate into billions for the bottom line. But while it's transformative, data science is also complex and time consuming. From extract, transform, and load (ETL) operations to inferencing, data science pipelines rely on large-scale infrastructure to power critical business operations.

These operations are bottlenecked by the serial nature of CPU-only computing, which is compounded when scaling out for large processes. With the NVIDIA EGX platform, enterprises can easily leverage parallel GPU computing to remove bottlenecks and quickly improve performance. This increased performance significantly improves time to insight and the return on investment for data-driven enterprises.

## NVIDIA EGX Platform for Accelerated Computing

The NVIDIA EGX™ platform enables enterprise IT to deliver diverse applications on high-performance and cost-effective infrastructure. The platform is a combination of high-performance GPU computing and highspeed, secure NVIDIA® Mellanox® networking in NVIDIA-Certified data center servers, built and sold by our partners. The NVIDIA EGX platform allows customers to prepare for the future while driving down costs by standardizing on a single unified architecture for easy management, deployment, operation, and monitoring.



Enterprise Support

- NVIDIA
- Hardware, Management, and ISV Partners
- Cloudera
- RAPIDS Accelerated Software

The EGX platform supports a vast suite of accelerated data science applications that deliver immediate benefits to users.

Easily deploy end-to-end data science pipelines on NVIDIA accelerated infrastructure to improve your data-driven operations.

### KEY APPLICATIONS / PLATFORMS

- > Anaconda Enterprise
- > Cloudera Data Platform
- > H2O.ai
- > Informatica
- > Jupyter Notebooks
- > Plotly
- > Oracle Machine Learning

### PROOF POINTS

- > Running an industry-standard Apache Spark workload, NVIDIA EGX demonstrated a 5X speedup—with some queries showing up to a 10X speedup relative to the same servers without GPUs.

### NVIDIA-CERTIFIED SYSTEMS

- > These systems enable enterprises to confidently deploy scalable hardware and software solutions that securely and optimally run accelerated workloads.
- > Learn more about accelerated servers at [nvidia.com/certified-systems](https://nvidia.com/certified-systems)



# Deploy High-Performance Data Science Solutions in Your Data Center Today

Historically, analytics workflows have been slow, cumbersome, and inefficient, often relying on CPU-only servers for data preparation, training, and deployment. The NVIDIA EGX platform dramatically boosts the performance of your end-to-end analytics workflows, speeding up the time to value while reducing overall cost.



## Reduced Costs

Maximize your budget with NVIDIA EGX rather than adding costs buying, deploying, and managing more inefficient CPU-only servers.



## Faster Insight

Leverage all of your data to make more informed business decisions, improve organizational performance, and better meet customer needs.




## Seamless Scaling

Effortlessly scale from a single server to multi-node, multi-GPU clusters with a consistent, intuitive architecture.

## Lightning-Fast Performance on Big Data

The NVIDIA EGX platform accelerates the end-to-end data science life cycle, whether your organization needs to reduce processing time of ETL pipelines or accelerate terabyte-scale machine learning workflows. With RAPIDS, data practitioners can accelerate pipelines on NVIDIA GPUs, reducing data operations like data loading, processing, and training from days to minutes. RAPIDS abstracts the complexities of accelerated data science by building upon popular Python and Java libraries, enabling users to see benefits immediately.



### Data Prep + ETL

Churn through ETL pipelines on NVIDIA GPUs using RAPIDS + Apache Spark 3.0 or Dask to arm your practitioners with high-quality datasets.

### Training

Develop, iterate, and refine business-enabling models to support your operations with RAPIDS and Dask.

### Visualize

Gain a deeper understanding of your data through massive-scale visualizations with RAPIDS + Plotly.

### Inference

Generate business insights fast to bolster operations and decision making with RAPIDS.