



Dell EMC PowerEdge Servers for Data Analytics

How to choose the right server to power traditional and emerging data analytics use cases

Select the right server for these use cases:

- Data transactions and reporting
- Advanced analytics
- Distributed data analytics

Make smarter use of data.

Data analytics can help you harness massive amounts of data, learn from it and monetize it. Data-driven use cases can range from keeping processed, structured data in databases to creating data lakes for combining structured and unstructured data in a format that enables more widespread and flexible use of data analytics. Many enterprises also take the next step of extending data analytics capabilities to distributed edge and cloud environments.

No matter where you are with data analytics, Dell Technologies has a complete portfolio of solutions that can help you recognize and take advantage of data to drive value across the enterprise.

Boosting data analytics performance

Several factors will impact the performance of your data analytics server. These include:

- **Memory:** Some databases and data analytics applications carry out computation in-memory, so adding more RAM accelerates performance. Traditional databases will use as much RAM as you can give them. Advanced data analytics workloads benefit from extra RAM, but often benefit more directly from local storage.
- **Storage:** Data lake applications, such as Apache® Spark® and Hadoop®, perform best with large amounts of direct-attached storage.
- **Networking:** Networking can easily become a performance bottleneck for data analytics. Connection speeds of 1Gbps to 10Gbps are suitable for large databases with multiple concurrent users. For data lakes, which are designed for heavy concurrent use, 25Gbps to 100 Gbps is preferred.
- **GPUs:** Most database and analytics applications such as SAP HANA®, Hadoop and Splunk® do not yet leverage GPUs. As your data analytics workloads evolve to include artificial intelligence (AI), deep learning training, GPU-accelerated servers may be more advisable.

Extend data analytics from core to cloud to edge.

The right Dell EMC PowerEdge server for your use case depends on a number of factors, as outlined in the following chart:

Use case	Transactions and reporting	Advanced analytics	Distributed data analytics
Description	Traditional databases store structured data in rows and columns and make it accessible to via queries.	Data lakes hold structured, semi structured and unstructured raw data, which can be combined, queried and processed as needed for advanced analytics.	Containerized microservices enable portability of data and applications across core, cloud and edge, while streaming data analytics enable real-time insights at the point of creation/ consumption.
Applications/ workloads	Popular SQL, NoSQL and in-memory databases, including Microsoft® SQL Server®, Oracle®, Apache® Cassandra®, Microsoft Azure® and SAP HANA.	Data lakes are commonly built on Apache Hadoop® and Apache Spark®.	Kubernetes® and Kubeflow enable containerized microservices, while streaming data analytics can be run on Apache Kafka®, Splunk® and Confluent®.
Key considerations	Traditional database applications are written to perform best on a four-socket server with lots of memory.	Data lake workloads benefit from a two-socket server with more local storage.	Workloads are typically split between a lighter-weight head node and more powerful worker nodes.
Recommended Dell EMC PowerEdge servers	<ul style="list-style-type: none"> PowerEdge R840 and R940 have the most DIMM slots, for the largest memory capacity. PowerEdge R940xa is recommended if a GPU is needed for deep learning. 	PowerEdge R740xd with extra storage is a great fit for data lakes.	<ul style="list-style-type: none"> PowerEdge R440, R650 and XE2420 at the edge, with R750 or R740xd in the data center PowerEdge R750/R740xd for private cloud PowerEdge R750xa or R940xa with GPUs for AI/DL use cases

Learn more

Visit us online at delltechnologies.com/analytcs
delltechnologies.com/servers

Dell Technologies can help you throughout your data analytics journey, enabling you to get started with your first data analytics project and continue to create value with more sophisticated data analytics approaches, including AI.