

Starburst technology, facilitating the discovery, querying, and processing of enterprise-wide data assets regardless of their physical locations. By significantly reducing data movement requirements and enhancing query efficiency, the Dell Data Lakehouse sets a new benchmark in data platform optimization and performance.

## Dell Data Lakehouse

A Modern Data Platform with an integrated data lakehouse built on Dell Hardware and a full-service software suite.

May 2025

[H19955.3]

Tech Book

## Copyright

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.

Copyright © 2025 Dell Inc. or its subsidiaries. Published in the USA July 2025 [\[H19955.3\]](#).

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

# Contents

Executive Summary .....4

Introduction .....6

Business challenges.....7

Data preparation approaches .....11

Dell Data Lakehouse .....12

Partner technology overview.....18

Dell Data Lakehouse architecture .....19

Dell Data Lakehouse Performance .....31

Dell Data Lakehouse sizing.....36

Conclusion.....40

References .....42

# Executive Summary

## Overview

The Dell Data Lakehouse is a fully integrated modern data platform built on powerful Dell hardware and advanced software, enabling organizations to manage, analyze, and act on data efficiently. Combining distributed query processing and exceptional SQL performance with new features like a lightning-fast SQL Query Engine, a Distributed Processing Engine for large-scale analytics, a Unified Metadata and Governance Layer for streamlined compliance, Scalable S3 Storage for effortless data growth, and an AI Data Preparation Layer for optimized insights, it offers robust data solutions. Whether centralizing or federating data, the platform minimizes movement, enhances performance, and empowers businesses to unlock the full potential of their data for smarter, faster decision-making.

The Dell Data Lakehouse includes:

- Lakehouse Compute cluster comprising Compute hardware, Dell Data Analytics Engine Software and Dell Data Lakehouse System Software
- Lakehouse Storage cluster

## Audience

This document is intended for enterprises with data lakes, or a data lake strategy interested in empowering their organizations to act more quickly, effectively, and efficiently on their data, as well as modernize to a data lakehouse. Audience roles include:

- Data and application administrators
- Data engineers
- Data scientists
- Hadoop administrators
- IT decision-makers

A data lakehouse can not only assist traditional analytics customers looking to modernize their data collection but also help analytics systems to get more value from their data or standardize their data for modern analytics workloads.

## Revisions

| Date         | Part number/<br>revision | Description                  |
|--------------|--------------------------|------------------------------|
| March 2024   | H19955                   | Initial release              |
| July 2024    | H19955.1                 | Updated version 1.1 features |
| October 2024 | H19955.2                 | Updated version 1.2 features |
| May 2025     | H19955.3                 | Updated version 1.4 features |

**Note:** This document may contain language from third-party content that is not under Dell Technologies' control and is not consistent with current guidelines for Dell Technologies' own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

---

**Note:** This document may contain language that is not consistent with Dell Technologies' current guidelines. Dell Technologies plans to update the document over subsequent future releases to revise the language accordingly.

---

**We value your  
feedback.**

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Author:** Kirankumar Bhusanurmah, DA / AI Specialist | TME, Dell Technologies

---

**Note:** For links to other documentation for this topic, see [dell.com/datamanagement](https://dell.com/datamanagement).


---

## Introduction

### Dell Data Lakehouse

The Dell Data Lakehouse is a state-of-the-art, turnkey solution designed to revolutionize how organizations manage, analyze, and act on their data. It combines Dell's trusted, high-performance hardware with an advanced software ecosystem to create a truly integrated data platform. At its heart is the Dell Data Analytics Engine, powered by Starburst, which now includes a lightning-fast SQL Query Engine for seamless, efficient querying, delivering insights with speed and precision across distributed and centralized datasets. A Distributed Processing Engine further enhances the platform's capabilities, enabling it to handle large-scale analytic workloads with ease, ensuring performance and scalability as data demands grow.

To support consistent and reliable data governance, the Unified Metadata and Governance Layer provides a structured approach to metadata management and compliance, fostering trust and ensuring operational alignment across the enterprise. On the storage front, Dell's leading Scalable S3 Storage solutions, including ECS, ObjectScale, and PowerScale, provide unmatched scalability, durability, and security, offering optimal performance even as datasets expand. Complementing these features, the new AI Data Preparation Layer leverages advanced automation to simplify data cleansing and structuring, accelerating the preparation pipeline for analytics and machine learning use cases.



Whether organizations need to federate data for analytics with minimal movement or centralize it for optimized estate management, the Dell Data Lakehouse delivers tailored solutions. By gathering data securely, delivering faster insights, and enabling smarter decision-making, it empowers businesses to unlock the immense value within their data and drive their next wave of innovation.

**Figure 1. Dell Data Lakehouse Diagram**

## Business challenges

### Market environment

Rapid advancements in artificial intelligence (AI) and machine learning (ML) have transformed the business landscape. Given the growing popularity of hybrid cloud computing, businesses face the challenges of establishing an enterprise data platform that integrates with both on-premises and public cloud providers. They also seek self-service analytics and data mesh principles in a cost-effective way. While these technologies offer exciting opportunities for innovation and data-driven decision-making, they also introduce complexities in storage, integration, and processing.

The types of data generated and collected by businesses have expanded. Organizations have historically dealt with structured data from traditional databases. Today, multinational organizations now deal with semi structured and unstructured data, such as text, logs, images, audio, video, social media content, and sensor data that spans across borders. These types of data cannot be placed in the same repository as structured data or in the same vicinity due to sovereignty regulations. Managing and analyzing this diverse data requires a modern data stack. The stack consists of a data lakehouse storage and lakehouse compute engine called Dell Data Analytics Engine powered by Starburst. It federates queries across multiple data sources while maintaining performance.


As data volumes and processing demands increase, traditional data management systems may struggle to scale efficiently to meet the growing needs of businesses. It is evident that the need for an enterprise-wide data management solution is a must have. However, integrating these disparate approaches can be complex and impeded by siloed data environments. These challenges require adopting new scalable and high-performance data management solutions that complement the traditional data stack. Technologies like data mesh are new data architectures that treat data as a product and each domain or business unit becomes responsible for managing and owning its data.

By implementing a data mesh architecture, organizations can associate distinct data architecture designs across various Hyperscalers or on-premises. Doing so enables each team to leverage the best technology for their domain while ensuring compatibility and cohesiveness at the company level. This modern data stack addresses the issues of data silos, streamlines the data sharing process, and enables the organization to fully use its data assets. It encourages cooperation, innovation within diverse functional areas, and minimize costs.

Addressing these challenges requires adopting modern data management strategies, embracing innovative data technologies, and empowering organizations with the necessary tools and skills for effective data utilization. Businesses that can navigate these complexities and demands successfully are better positioned to leverage data-driven insights and gain a competitive advantage in their respective markets.

### Data analytics processing platforms

Broad data architectures have evolved over time to address large-scale data collection and analysis as diverse types are being generated. Their uses have grown exponentially and diversified. See Data storage methodologies for analytics.



**Figure 2. Data storage methodologies for analytics**

NOTE: The graphic above was adapted from "[Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics](#)," 11th Annual Conference on Innovative Data Systems Research, January 11–15, 2021,p. 2.

## Data warehouse

The data warehouse was designed to handle massive amounts of information. Data warehouses were optimized for business intelligence and decision support, dealing primarily with well-defined and structured data.

Using the foundational extract-transform-load (ETL) process, data could be presented to applications in the consistent manner that they require. ETL does have some degree of overhead and lacks a certain amount of flexibility. Businesses adapted to data warehouse requirements because of the valuable insights that a unified collection of many structured datasets could deliver.

Large organizations continue to use data warehouse methodology. It provides a consolidated and structured repository for diverse data sources and analysis of data. It offers optimized query performance for complex analytics to help deliver rapid insights for business stakeholders.

## Data lake

The data lake technology came into existence driven by the evolving data landscape and the need for more flexible and scalable data storage and processing solutions. It has become popular for several reasons:

- Data lake is designed to store data from multiple sources and multiple data structures in the same repository. They include structured, semi structured, and unstructured data. Data lake also eliminates the need for data modeling at the time of ingestion. It provides a unified platform capable of storing and managing diverse datatypes without the need for predefined schemas.
- Data lake can use both cloud-based and on-premises storage options. It can scale horizontally to accommodate massive amounts of data, making them into cost-effective solutions for storing and processing vast datasets. This scalability is important in the era of big data, where organizations need to handle data growth efficiently.
- Data lake integrates well with various big data processing frameworks and tools, such as Hadoop, Spark, Trino, and more. This capability enables complex data processing and analytics, making it attractive to organizations looking to leverage their existing infrastructure.

## Data lake and data warehouse

To continue using data warehouse technology while maintaining the flexibility and the scalability of data storage, organizations turned to both data lake and data warehouse. These two architectures serve different needs and use cases. This two-tier approach enables organizations to have Atomicity, Consistency, Isolation, and Durability (ACID), dynamic, and flexible operations on increasing quantities and varieties of data.

Data lakes excel at handling diverse and raw datatypes including structured, semi structured, and unstructured data. This capability enables them to store large amounts of data without the requirement of strict schemas. Data warehouses, however, are optimized for querying and analyzing structured data while maintaining their strong data governance and compliance features.

This strategy, however, results in less reliability and the potential for stale data. Users who work in analytics now must deal with two sources, which increases the level of complexity.

## Data lakehouse

A data lakehouse is a modern data architecture that combines the structure and performance of a data warehouse with the flexibility of a data lake. It is designed to provide a unified and scalable platform for storing, processing, and analyzing large volumes of structured, semi structured, and unstructured data.

Traditionally, data lakes and data warehouses have been separate concepts with distinct purposes. A data lake is a centralized repository that stores raw and unprocessed data from various sources in its original format. It allows for the storage of diverse datatypes and provides flexibility for data exploration and analysis. However, a data warehouse is a structured and organized repository that stores processed and transformed data optimized for querying and analytics.

The data lakehouse concept emerged as a response to the limitations and challenges of these traditional architectures. It aims to combine the strengths of data lakes and data

warehouses, bridging the gap between raw data storage and structured analytics. In a data lakehouse, data is stored in its raw form, but it is also curated and organized to enable efficient analytics.

Advancements in certain technologies make a data lakehouse possible. These technologies include:

- Metadata layers for data lakes
- New query engine designs providing high-performance SQL execution on data lake storage
- Advanced analytics and machine learning tools

The primary benefits for adopting a data lakehouse can be summarized as:

- **Flexible and simpler architecture**—A data lakehouse provides flexibility in terms of storage solutions like object storage, S3, or HDFS. It provides fast queries without the need to copy or move data to meet BI requirements. This architecture minimizes data extracts and reduces the challenge of managing multiple copies of data and streamlines cost while improving agility to support changing business requirements.
- **Workloads consolidation**—A data lakehouse supports both BI and data science enabling organizations to consolidate workloads. This consolidation eliminates the need of having two separate platforms, and still maintains an open architecture to interoperate with other tools.

Data lakehouse core architecture layers illustrates the core architectural layers of a data lakehouse.




Figure 3. Data lakehouse core architecture layers

## Data preparation approaches

### Data Pipelines


Data pipelines with Extract, Transform, Load (ETL) and Extract, Load, Transform (ELT) are two common techniques used in data analytics. They prepare and move data from source systems to a target data warehouse or data storage for analysis.

The ETL process is typically used in a data warehouse architecture. The data is first extracted from source systems and then goes through the transformation process to cleanse the data to prepare for analysis. Once the data has been transformed, it is loaded into a target data warehouse or a data mart. This load process typically maps the transformed data into tables and structures in the target systems.

The ELT process, typically used in a modern data lakehouse, starts with data extraction from source systems. This technique works well in a data lakehouse architecture. Data is extracted and loaded into the target storage as-is, with minimal transformation. The transformation step is performed directly on the target systems using integrated processing capabilities without cleansing the data first. This approach takes advantage of compute power of a data lakehouse to transform the data as needed during the analysis phase.

The transformed data is stored in the open table formats, it could be Apache iceberg, Delta or Apache Hudi. The Dell Data Lakehouse is validated to support iceberg and Delta Lake Table formats. Any of the data modeling techniques can be used such as Medallion Architecture or One Big Table (OBT).

[ELT compared to ETL](#) shows the differences between the ETL and ELT processes.



**Figure 4. ELT compared to ETL**


Both ETL and ELT methods can be used simultaneously, and each has its own benefits. The choice between them depends on factors such as data volume, data transformation complexities, and capabilities of the data warehouse or storage system being used for analysis.

## Dell Data Lakehouse

### Overview

In today's data-driven world, organizations face the challenge of efficiently managing and delivering value from vast and diverse datasets. To address this challenge, a comprehensive and robust data management solution is essential.

The Dell Data Analytics Engine is the core engine of the Dell Data Lakehouse. Powered by Starburst, it is a distributed query engine that can serve as both a single point of access to all data across the enterprise, as well as the data lakehouse query engine. The Dell Data Lakehouse leverages Dell's industry proven PowerEdge R660 Servers tailor-made for Lakehouse workloads. They are referred to as the Dell Data Analytics Engine 660 (DDAE660) nodes. The entire stack is managed and orchestrated by a new Dell Data Lakehouse System Software based on Kubernetes. This software ensures maximum scalability, control, and simplified lifecycle management.



**Figure 5. Dell Data Lakehouse Compute Nodes and Lakehouse Software Components**

The Dell Data Lakehouse delivers key benefits for the entire organization.

- Separate compute from storage, easily scale with more DDAE compute nodes

The cluster can be scaled easily from a 2 worker node minimum to up to 20 nodes in the first release. Storage can be scaled independently based on the workload requirements.

- Out-of-the-box automation and orchestration for day to day IT Ops

Dell Data Lakehouse comes with automation out of the box that helps automate the installation and the ongoing lifecycle management. The turnkey stack, when deployed by Dell experts, includes all software components, the management

control plane, user management and associated databases including the Hive Metastore.

Updates, security patches, and feature additions for any component within the Dell Data Lakehouse stack will be delivered as a consolidated payload. This approach eliminates the complexity arising from sourcing drivers, firmware, software, patches, updates, etc.

- Reduced effort and cost of managing middleware such as operating systems, virtualization or container runtimes, orchestration tools, etc.

The turnkey solution delivers data lakehouse in a box with everything built-in so IT teams don't have to incur costs of managing multiple layers of the stack, while maintaining flexibility and openness in data formats.

- Best-in-class integration with Dell storage (ECS, ObjectScale, PowerScale)

Based on thorough validation and integration from Dell engineering teams, organizations can receive an enhanced support experience when using Dell Data Lakehouse with Dell storage.

- Single vendor for E2E transaction

Procurement teams can take advantage of Dell's broad portfolio and procure the complete solution including all hardware (including server, storage, networking), software, deploy and support services as well as curated professional services.

- Single vendor for E2E guidance across the lifecycle

Dell Services provide deep expertise throughout the lifecycle. From aligning a winning strategy, validating data, quickly implementing a data platform or ensuring secure, optimized operations, trusted experts help accelerate time to value in effectively leveraging enterprise data to power AI projects.

Strategize—establish a winning strategy to drive secure business outcomes

- **ProConsult Advisory for GenAI** build a strategy and roadmap, gain consensus on high-priority use cases and align a plan to achieve them.
- Leverage **Accelerator Workshop for GenAI** as a great first step
- And with **Advisory Services for GenAI Data Security** we can help mitigate data security risks throughout the lifecycle

IMPLEMENT: Prepare data for GenAI and implement the platform

- **ProDeploy for Infrastructure**—deploy faster with less effort and more control
- **Accelerator Services for Dell Data Lakehouse**—implement a fully operational Data Lakehouse platform to accelerate AI and data analytics
- **Implementation Services for Data Preparation** validate data sets and align data for use by AI

ADOPT & Scale—improve operations and extend capabilities

- **Advisory Services Subscription for Data Analytics Engine** provides a dedicated expert to maximize value from the data analytics engine within the lakehouse (powered by Starburst), delivered over a pre-defined set of time
- **ProSupport Infrastructure Suite**—24/7 support from trained experts to maximize value from investments

Dell Data Lakehouse Components

The Dell Data Lakehouse solution is made up of four key components.

1. Dell Data Lakehouse Compute Nodes (control plane, coordinator, worker)
2. Dell Data Analytics Engine and Processing Engine powered by Starburst
3. Dell Data Lakehouse System Software
4. Dell Lakehouse storage cluster (fully validated with Dell ECS, compatible with other S3 compliant storage like Dell ObjectScale and Dell PowerScale)

Dell Data Lakehouse Compute Nodes

Built on Dell PowerEdge R660 server (1U), the Dell Data Lakehouse compute cluster is tailormade for data lakehouse workloads. The cluster of nodes is connected to data lake storage as well as the rest of the external environment via customer-provided network equipment. The compute nodes are of two types:

- i. Control Plane Nodes: these are fixed 3 nodes that run the essential management software and services such as a Hive Metastore.
- ii. Coordinator and Worker Nodes: based on a coordinator-worker node architecture, these nodes run the Dell Data Analytics Engine and Processing Engine, and scale with additional worker nodes to handle complex or high volume queries.



Based on architecture for today and tomorrow




Figure 6. Dell Data Lakehouse Compute Nodes

## Dell Data Analytics Engine powered by Starburst

The Dell Data Analytics Engine, powered by Starburst and enhanced with Apache Spark as the Distributed Processing Engine, is an enterprise-grade platform designed to deliver high-performance analytics across diverse and complex data landscapes. This advanced engine now supports a myriad of applications, extending its capabilities to include Streaming, Batch processing, AI/ML workloads, Graph analytics, and Geospatial analytics. These additional functionalities amplify the existing distributed SQL query engine's ability to create a robust, unified platform for comprehensive data analysis.

With the Dell Data Analytics Engine, users can seamlessly query large datasets across relational databases, NoSQL systems, object storage solutions, and more, all with standard SQL syntax. The integration of Apache Spark allows organizations to process real-time streaming data, train machine learning models, analyze graph-based data structures, and extract insights from geospatial datasets. This versatility enables businesses to tackle previously complex analytics tasks with precision and speed. Additionally, users can perform federated queries, accessing and analyzing data spread across on-premises systems, cloud platforms, and edge environments in a single query, reducing data movement and increasing efficiency.

Core functionalities of the Dell Data Analytics Engine include its ability to query data directly from open storage systems like Dell ECS, ObjectScale, and PowerScale, accommodating open formats like Parquet, AVRO, and ORC, as well as metadata in table formats such as Delta Lake and Iceberg. By leveraging its built-in abstraction layer, the platform simplifies connectivity and data integration, eliminating the need to centralize data, which can be time-intensive and resource-heavy. This layer ensures users can work with diverse data sources while administrators maintain uniform access control policies, including row and column-level security and data masking.

Key advantages include:

- **Data Analysts, Data Scientists, and Engineers** can explore and analyze data without moving it, rely on fault-tolerant execution for guaranteed workloads, and access a streamlined query interface that reduces dependency on IT teams.
- **Data Engineers** can build reusable data products and experiment with data models without setting up cumbersome pipelines.
- **Governance and Administrative Teams** benefit from consistent metadata and governance standards, which enable secure and compliant data management across the enterprise.
- **IT and Infrastructure Teams** can quickly grant data access, optimize high-value workloads, and scale operations flexibly to accommodate diverse workloads ranging from AI/ML to operational queries.

Unique additional features, such as table redirection, materialized views, and smart cache, further enhance efficiency for both interactive and ETL workloads. The integration of Spark expands workload isolation and scaling capabilities, ensuring high-performance analytics tailored to the needs of modern businesses. The Dell Data Analytics Engine empowers organizations to unlock deeper insights faster, enabling innovation and operational excellence in a competitive data-driven world.

## Dell Data Processing Engine powered by Apache Spark

The Dell Data Processing Engine, powered by Apache Spark, is a secure, enterprise-grade data processing platform engineered for the demands of modern analytics, AI, and real-time operations. Building on the powerful foundations of Spark, the Dell Data Processing Engine adds integrated security, governance, and resource management features to deliver a production-ready Spark environment that is easy to scale, easy to govern, and optimized for hybrid data architectures.

Designed to support the full spectrum of data-driven workloads — from high-speed ingestion and real-time streaming to advanced machine learning and large-scale batch processing — the Dell Data Processing Engine helps organizations unlock faster insights and smarter operations across their data landscape.

With its advanced integration capabilities, users can ingest, transform, and enrich data from diverse sources including databases, file systems, and object storage like Dell ECS, ObjectScale, and PowerScale. Native support for open table formats such as Delta Lake and Iceberg ensures that enterprises can maintain flexibility and interoperability across their data ecosystem.

Built-in streaming capabilities enable real-time data processing and event-driven architectures, while the platform's MLlib and external ML framework support allows users to build, train, and deploy machine learning models at scale — directly on top of enterprise data. Advanced graph processing, geospatial analytics, and large-scale ETL pipelines are all natively supported, empowering teams to handle complex analytics workloads with ease.

Key enterprise enhancements include:

- **Secure Access and Governance:** Native integration with SSO, LDAP/AD, and built-in role-based access control (RBAC) ensures secure user authentication, fine-grained data access policies, and strong auditability out-of-the-box.
- **Resource Pool Management:** Fine-tuned workload isolation and dynamic resource scheduling ensure critical applications — including ingestion pipelines, streaming jobs, and machine learning workflows — get the right resources at the right time without contention.
- **Enterprise Resilience and Fault Tolerance:** Built-in high availability, fault-tolerant execution, and monitoring integrations enable organizations to operate Spark-based pipelines with enterprise-grade reliability.
- **Unified Metadata and Storage Access:** Consistent metadata management and seamless access to open storage formats eliminate data silos and simplify data engineering across hybrid environments.

Key benefits:

- **Data Engineers** can rapidly ingest and transform massive data volumes, build reusable data products, and automate streaming pipelines with minimal operational overhead.
- **Data Scientists** can train ML models using distributed compute power while accessing fresh, real-time data without data movement or duplication.
- **Business Analysts and Developers** can access prepared datasets and enriched features quickly, enabling faster analytics, smarter applications, and improved decision-making.
- **Governance and Security Teams** benefit from unified control over data access, user rights, and data compliance across the entire enterprise data estate.

- **IT and Infrastructure Teams** can scale compute and storage independently, optimize resource usage, and maintain operational resilience for diverse analytics and AI workloads.

By combining the unmatched versatility of Apache Spark with Dell's enterprise-ready innovations, the Dell Data Processing Engine transforms raw, fast-moving, and complex data into real-time intelligence — securely, reliably, and at massive scale.

## Dell Data Lakehouse System Software

This new system software is the central nervous system of the Dell Data Lakehouse. It simplifies lifecycle management of the entire stack, drives down IT OpEx with pre-built automation and integrated user management, provides visibility into the cluster health and ensures high availability, enables easy upgrades and patches and lets admins control all aspects of the cluster from one convenient control center. Based on Kubernetes, it's what converts an otherwise DIY data lakehouse into an easy button for enterprises of all sizes.

The Dell Data Lakehouse offers three user interfaces tailored to different users' needs.

### Dell Data Lakehouse System Software user interface

The Dell Data Analytics Engine offers a web-based administrator user interface known as the Dell Data Lakehouse System Software UI. The Lakehouse System Software UI allows users to manage and monitor system features and settings remotely from any location over a network.

### Dell Data Analytics Engine user interface

Dell Data Analytics Engine user interface is powered by Starburst Enterprise Platform (SEP) which is a commercial distribution of Trino.

### Dell Data Processing Engine user interface

Dell Data Processing Engine user interface is powered by Apache Spark, and is embedded into the Dell Data Analytics Engine user interface

### Dell Data Lakehouse System Software - User Management user interface

Administrators use the Dell Data Lakehouse System Software's User Management interface to manage local users and LDAP-configured users. These users gain access to both the Dell Data Lakehouse System Software and Dell Data Analytics Engine interfaces through this management system.

## Dell Lakehouse storage cluster

An object storage system is connected to function as data lake storage to store large amounts of data using an array of industry standard formats. Dell ECS has been fully validated with Dell Data Lakehouse. Other S3-compliant storage like Dell ObjectScale and Dell PowerScale are also compatible with the Dell Data Lakehouse.

### Dell ECS

Dell ECS, the world's most cybersecure object storage<sup>1</sup>, delivers unmatched scalability, performance, resilience, and economics. ECS delivers rich S3-compatibility on a globally distributed architecture, empowering organizations to support enterprise workloads such as AI, analytics, and archiving at scale. ECS customers are reducing TCO up to 76% over public cloud.

### [Dell ObjectScale](#)

ObjectScale is high-performance containerized object storage built for the toughest applications and workloads—Generative AI, analytics, and more. Innovate faster, at any scale, with a global namespace, strong S3 compatibility, and enterprise-class security that is ready on day one. Expanding its software-defined options, ObjectScale is now also available as the world's most powerful object storage appliance purpose-built for Kubernetes.

### [Dell PowerScale](#)

The world's most flexible, secure, and efficient scale-out file storage.

## Partner technology overview

### Delta Lake from the Linux Foundation

A Linux Foundation project since 2019, Delta Lake by the Linux Foundation is an independent project controlled by a development community rather than any single technology vendor. The Dell Data Lakehouse enables reliable deployment and operation of Delta Lake in the solution. More than 150 developers from over 50 different organizations associated with Delta Lake working on multiple storage repositories are all engaged daily to help push the program goals forward. Key Delta Lake capabilities include:

**ACID transactions**—Ensure data consistency and reliability and isolating it at the strongest isolation level, the serializable level.

**Time travel and data versioning**—Each data write to a Delta table creates a version number. This feature enables users to query a Delta Lake table as of a specific time. Users can view and revert to previous versions of the data using a timestamp or a version number.

**Scalable metadata**—Leverages Spark's distributed processing power to easily handle all the metadata for petabyte-scale tables with billions of files.

**Schema evolution and enforcement**—Perform automatic schema validation by checking against a set of rules to determine the compatibility of a write from a DataFrame to a table.

**DML operations**—Support Data Manipulation Language (DML) operations like updates, deletes, and merges by using transaction logs. They enable easy handling of complex use cases like change-data-capture, slowly changing dimension (SCD) operations and streaming upserts.

### Apache Iceberg

Apache Iceberg is a new open-source, high-performance data table format designed for large-scale data platforms. Its primary goal is to bring the reliability and simplicity of SQL tables to big data while providing a scalable and efficient way to store and manage data. This performance is especially important in the context of big data workloads. Iceberg makes it possible for engines such as Spark, Trino, Flink, Presto, Hive, and Impala to safely work with the same tables, simultaneously. It addresses some of the limitations of traditional data storage formats such as Apache Parquet and ORC.

Some key Apache Iceberg features and concepts include:

**ACID transactions**—Iceberg tables guarantee atomic, consistent, isolated, and durable (ACID) transactions for write operations, ensuring data integrity.

**Catalog**—A mechanism that houses metadata pointers for Iceberg tables.

**Table metadata**—Maintains comprehensive metadata information about tables, including schema details, partitioning information, and datafile locations. The metadata layer consists of three categories:

- Metadata files that define the table
- Manifest lists that define a snapshot of the table
- Manifests that define groups of datafiles that may be part of one or more snapshots

**Time travel**—Supports the ability to query data at different points in time, enabling historical analysis and data rollback to specific versions.

**Schema evolution**—Enables schema evolution by adding, renaming, or deleting columns without requiring expensive and time-consuming data migrations.

**Partitioning**—Enables data to be partitioned into logical segments based on specific columns, such as date or region. Partitioning helps improve query performance by enabling efficient data pruning and reducing the amount of data scanned during queries.

## Dell Data Lakehouse architecture

### Dell Data Lakehouse overall framework

Designing for today and tomorrow, the Dell Data Lakehouse architecture consists of five layers:

**Storage layer** – This layer is meant to keep all types of data in high performance, scalable and secure object storage. By integrating other layers with this data lake storage as shown in the below figure, client tools can query (via a query engine) the objects directly residing on object storage and eliminating ETL.

**Compute layer** – This layer provides scale-out compute horsepower for running the distributed query and data processing engine as well as the associated essential services such as a metadata service (Hive Metastore).

**Lakehouse System Software layer** – While a data lakehouse stack can be assembled with individual pieces, it is hard to deploy, maintain and secure. The Lakehouse System Software delivers a managed compute experience with automated deployment and simplified lifecycle management of the entire lakehouse stack.

**Distributed Query and Processing Engine layer** – This layer pulls data from various sources, including databases, IoT devices, applications, logs, and many more, and deliver

it to the end user as a query outcome, or write the data into the storage layer in an open format where it can be further used to transform data (ELT).

**Metadata layer** – The Distributed Query Engine layer relies on the metadata layer to understand the schema and layout of the data stored in the data lakehouse to query and update efficiently. Commonly referred to as a metastore or catalog, this layer provides information on directory structure, file format, and metadata about the stored data and enables ACID compliance, data versioning, caching, indexing, and more. Increasingly popular open source projects such as Iceberg and Delta Lake are now the de-facto sources of metadata that defines the structure of the data stored in open file formats.

**Consumption layer** – This layer hosts different tools and applications such as BI tools (e.g., Tableau, Power BI), or AI tools (Python, Spark, ML tools) or other data ecosystem tools like data catalog, data security, access control, etc. This layer is where the users perform analytics activities such as running SQL queries, BI, data prep for AI modeling or govern data access and rely on data stored in the data lakehouse.

Dell Data Lakehouse stack diagram illustrates the essential layers of a data lakehouse.





Figure 7. Dell Data Lakehouse stack and marketecture diagram

## Infrastructure Overview

Dell provides the necessary infrastructure (compute, memory, storage, and network resources) for the platform. [Dell Data Lakehouse infrastructure](#) illustrates the required infrastructure components and their roles in the Lakehouse cluster.



**Figure 8. Dell Data Lakehouse infrastructure**

The Dell Data Lakehouse includes tailor-made nodes that are configured specifically to handle the data lakehouse workloads. They come in two different configurations to support the control plane and coordinator/worker node architecture. Individual node configurations include node-specific guidance as their role. More guidance for cluster level sizing and scaling is provided in [Sizing the Dell Data Lakehouse solution](#).

**Server Infrastructure**

The server infrastructure provides compute, memory, and some of the storage resources that are required to run customer workloads. Dell Data Lakehouse comes with two types of nodes, control plane nodes and analytics engine nodes (coordinator and worker).

Details of the node configurations can be found in Table 1. Node Configurations for Dell Lakehouse.

**Table 1. Node Configurations for Dell Lakehouse**

| Node Type                 | Control Plane x 3                                                                                | Worker / Coordinator x n                                                                                                       |
|---------------------------|--------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Server Model              | Dell Data Analytics Engine Compute 660 based on PowerEdge Rack Server                            | Dell Data Analytics Engine Compute 660 based on PowerEdge Rack Server                                                          |
| CPU (dual-socket)         | Intel Xeon Gold 5416S 2G, 16C/32T, 16GT/s, 30M Cache, Turbo, HT (150W) DDR5-4400                 | Intel Xeon Gold 5416S 2G, 16C/32T, 16GT/s, 30M Cache, Turbo, HT (150W) DDR5-4400                                               |
|                           |                                                                                                  | Intel Xeon Platinum 8568Y+ 2.3G, 48C/96T, 20GT/s, 300M Cache, Turbo, HT (350W) DDR5-5600 QTY2                                  |
| Total CPU Cores (Threads) | 32 (64 Hyper threads)                                                                            | 32 (64 Hyper threads) OR 96 (192 Hyper threads)                                                                                |
| Memory DIMM               | 128 GB at 5,600 MT/s                                                                             | 256 GB at 5,600 MT/s<br>OR<br>512 GB at 5,600 MT/s                                                                             |
| Hard Drives               | 1.6TB SSD SAS Mixed Use up to 24Gbps SED FIPS-140 512e 2.5in Hot-Plug 3DWPD AG Drive (x4) RAID 5 | 480GB SSD SATA (RAID 1)<br><br>2 X 3.2TB SSD SAS Read Intensive up to 24Gbps SED FIPS-140 512e 2.5in Hot-Plug (For Warp Speed) |
| Network OCP (SFP)         | Intel E810-XXV Dual Port 10/25GbE SFP28, OCP NIC 3.0                                             | Intel E810-XXV Dual Port 10/25GbE SFP28, OCP NIC 3.0                                                                           |
| Network PCIe (SFP)        | Intel E810-XXV Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile                                | Intel E810-XXV Dual Port 10/25GbE SFP28 Adapter, PCIe Low Profile                                                              |

## Storage infrastructure

Dell Data Lakehouse is fully validated with Dell ECS and is compatible with S3-compliant storage like Dell ObjectScale and Dell PowerScale.

Dell ObjectScale, ECS, and PowerScale are deployed as cluster-level systems. The recommendations here can be used as guidance for new clusters, verification of compatibility with existing clusters, or expansion of existing clusters.

### ECS Node

Dell Technologies recommends using the ECS EX500 node configuration or ECS EXF900 node configuration for their primary S3 Object lakehouse.

**Table 2. ECS EX500 and EXF900 node configuration**

| Model                                | ECS EX500                               | ECS EXF900                               |
|--------------------------------------|-----------------------------------------|------------------------------------------|
| Model ECS                            | EX500                                   | EXF900                                   |
| Chassis                              | 2U node                                 | 2U node                                  |
| Nodes per rack                       | 16                                      | 16                                       |
| Node storage                         | 384 TB (twenty-four 16 TB NLSAS drives) | 184 TB (twenty-four 7.68 TB NVMe drives) |
| Node cache                           | 960 GB SSD                              | N/A                                      |
| Usable capacity per chassis          | Slightly less than 384 TB               | Slightly less than 184 TB                |
| Front-end networking                 | Two 25 GbE (SFP28)                      | Two 25 GbE (SFP28)                       |
| Infrastructure (back-end) networking | Two 25 GbE (SFP28)                      | Two 25 GbE (SFP28)                       |

The ECS EX500 configuration provides a good balance of storage density and performance for lakehouse usage. And the ECS EXF900 configuration is an all-flash configuration and provides the highest performance for lakehouse usage.

Two Ethernet network ports per node included for connection to the Cluster data network or an ECS storage network. Two additional network ports included for connection to the ECS back-end network.

### ObjectScale infrastructure

Dell Technologies recommends the configuration in ObjectScale all flash configuration for storage clusters that use ObjectScale for primary lakehouse storage using the s3a:// protocol.

Customers can use either of the following options to deploy ObjectScale:

- Software-defined storage to deploy ObjectScale software on Red Hat OpenShift
- An XF960 appliance that is based on the latest generation of Dell PowerEdge servers

**Table 3. ObjectScale all flash configuration**

| Machine Function                  | Component                                                                                                                                                        |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Platform                          | PowerEdge R760 server                                                                                                                                            |
| Nodes per rack                    | 16                                                                                                                                                               |
| Chassis                           | 2.5" chassis with up to 24 NVMe Direct Drives, two CPUs                                                                                                          |
| Chassis configuration             | Riser configuration 3, half-length, two 2-channel full-height slots (Gen4), two 16-channel full-height slots (Gen5), and two 16-channel low-profile slots (Gen4) |
| Power supply                      | Dual, hot-plug, fully redundant (1+1) 1100 W power supplies                                                                                                      |
| Processor                         | Intel Xeon Gold 6426Y 2.5 G, 16 C/32 T, 16 GT/s, 38 M                                                                                                            |
| Memory capacity                   | 512 GB (sixteen 32 GB RDIMM, 4800 MT/s, dual rank)                                                                                                               |
| Internal RAID storage controllers | C30, no RAID for NVMe chassis                                                                                                                                    |
| Disk – NVMe                       | 24 6.4 TB Enterprise NVMe, mixed-use agnostic drive, U.2                                                                                                         |
| Boot-optimized storage cards      | BOSS-N1 controller card + with two M.2 960 GB SSDs (RAID 1)                                                                                                      |
| Network interface controllers     | NVIDIA ConnectX-6 Lx dual port 10/25 GbE SFP28 adapter, PCIe low profile                                                                                         |
| Node storage                      | 153.6 TB (24 6.4 TB NVMe drives)                                                                                                                                 |
| Front-end networking              | Two 25 GbE (SFP28)                                                                                                                                               |

The ObjectScale configuration is an all-flash configuration and provides the highest performance for lakehouse usage.

Two Ethernet network ports per node included for connection to the Cluster data network or an ECS storage network.

### PowerScale infrastructure

Dell Technologies recommends the configuration in PowerScale configuration for storage in clusters using PowerScale for their primary lakehouse storage using both S3 and HDFS protocols.

**Table 4. PowerScale configuration**

| Machine Function                     | Component                                                       |
|--------------------------------------|-----------------------------------------------------------------|
| Model                                | PowerScale H7000 (hybrid)                                       |
| Chassis                              | 4U node                                                         |
| Nodes per chassis                    | Four                                                            |
| Node storage                         | Twenty 12 TB 3.5-inch, four native sector size SATA hard drives |
| Node cache                           | Two 3.2 TB SSDs                                                 |
| Usable capacity per chassis          | 600 TB                                                          |
| Front-end networking                 | Two 25 GbE (SFP28)                                              |
| Infrastructure (back-end) networking | Two InfiniBand QDR or two 40 GbE (QSFP+)                        |
| Operating system                     | OneFS 9.5.0.2                                                   |

The recommended configuration is sized for typical usage as lakehouse storage.

Two Ethernet network ports per node included for connection to the Cluster data network or a PowerScale storage network. Two additional network ports are included for connection to the PowerScale back-end network. These additional ports can be either InfiniBand QDR or 40 GbE, depending upon on-site preferences.

One PowerScale H7000 chassis supports four PowerScale H7000 nodes. This configuration provides approximately 720 TB of usable storage. At 85% utilization, 600 TB of storage is a good guideline for available storage per chassis.

This configuration assumes that the PowerScale nodes are primarily used for lakehouse storage. If the PowerScale nodes are used for other storage applications or clusters, those applications must be taken into account in the overall cluster sizing. Other PowerScale H7000 drive configurations may also be used.

## Network infrastructure

The network is designed to meet the needs of a high performance and scalable cluster, while providing redundancy and access to management capabilities. The architecture is a leaf and spine model that is based on Ethernet networking technologies. It uses PowerSwitch S5248F-ON switches for the leaves and PowerSwitch Z9432F-ON switches for the spine.

Physical networking in this architecture is straightforward since most of the advanced capabilities of the system are implemented using software-defined networking. The logical network is described in the [Container platform implementation](#). This architecture has three physical networks, as shown in [Physical network infrastructure](#).

### iDRAC (or BMC) network

The iDRAC or BMC network is isolated. For utmost security, there is no connectivity to iDRAC via iDRAC interfaces. All BMC activity is secured within the Dell Data Lakehouse software stack and happens securely via the Dell Data Lakehouse OS directly accessing the iDRAC via Dell Data Lakehouse OS passthrough mechanism.

### Cluster data network

The Cluster data network is the primary network for internode communication between all server and storage nodes. Each server node on this network is assigned a single IP address on this network.

### Core data center network

The Core data center network is the existing enterprise network. The Cluster data network is interfaced with this network through switching and routing allowing cluster services to be exposed to system users.




Figure 9. Physical network infrastructure

### Management network fabric

The management traffic and the data traffic share a unified network involving two 25Gbps per node configured into a link aggregated group.

### Cluster network fabric

The Cluster network uses a scalable, resilient, nonblocking fabric with a leaf-spine design as shown in [Cluster data network connections](#). Each node on this network is connected to two S5248F-ON leaf switches with 25 GbE network interfaces.

The switches run Dell SmartFabric OS10. SmartFabric OS10 enables multilayered disaggregation of network functions that are layered on an open-source Linux-based operating system.


On the server side, the two network connections are bonded and have a single IP address assigned.

On the switch side, the network design employs a Virtual Link Trunking (VLT) connection between the two leaf switches.

VLT technology enables a server to uplink multiple physical trunks into more than one S5248F-ON switch by treating the uplinks as one logical trunk. In a VLT environment, a connected pair of switches acts as a single switch to a connecting server while all paths are active. It is possible to achieve high throughput while still providing resiliency against hardware failures. VLT replaces Spanning Tree Protocol (STP)-based networks, providing both redundancy and full bandwidth utilization using multiple active paths.

The VLT configuration in this design uses four 100 GbE ports between each Top of Rack (ToR) switch. The remaining 100 GbE ports can be used for high-speed connectivity to spine switches, or directly to the data center core network infrastructure.

The above illustration is Dell specific, Dell Data Lakehouse solution also allows Bring Your Own network, in that case the network architecture must be comparable to the Dell network architecture illustrated below.



**Figure 10. Cluster data network connections**

## Backup and Restore

### Overview

The backup and restore process is essential for preserving data integrity during upgrades, ensuring business continuity, and mitigating risks associated with data loss or corruption. The Dell Data Lakehouse offers robust backup capabilities through Dell ProSupport, enabling customers to protect their valuable data assets effectively and maintain operational resilience.

### Backup

During the initial Dell Data Lakehouse installation, with the help of Dell's ProDeploy personnel, customers can trigger a backup instantly, or customize the backup frequency to align with their organizational policies, ensuring regular backups for mitigating the risk of data loss. The backups are currently supported on NFS export and CIFS share locations. Customers need to provide a reliable backup location during the installation process. Such backups can optionally be encrypted as well for even greater security.

The backup functionality comprehensively captures critical data elements, including database schemas, content repositories, configuration settings, and user profiles. Since this operation can run on any node in the lakehouse cluster, it is necessary that the IP addresses of all nodes are allowed access to the NFS export or CIFS share, and no firewall be enabled on the NFS server for accessing the backups. And finally, these

backups are timestamped. Therefore, sufficient capacity must be available in the NFS or CIFS location enabling more flexible data retention.

## Restoration


Restoring the Dell Data Lakehouse is streamlined and efficient, allowing customers to select the desired backup point and initiate the process seamlessly via the help of ProSupport. However, a successful restore depends on the availability of the latest backup files, system compatibility, and data integrity. By adhering to these dependencies and Dell ProSupport service, customers can efficiently recover their data and maintain continuity in the operations.

## Active Passive Disaster Recovery

This feature intends to place infrastructure into the DDLH system that enables customers to maintain business continuity between two different clusters, possibly maintained at two separate sites.

Active-passive architecture, or a standby or failover system, involves a primary active system and a secondary passive system that remains inactive until the primary system fails. The primary system handles all incoming requests and operations while the passive system remains on standby, ready to take over if the primary system encounters a failure or becomes unavailable.

- In the event of a failure, the passive system is activated to assume the workload and responsibilities of the primary system, ensuring continuity of service and minimizing downtime.
- This architecture prioritizes redundancy and reliability, making it suitable for applications where uninterrupted service is critical, such as disaster recovery systems or mission-critical infrastructure.



# Dell Data Lakehouse Performance

The Dell Data Lakehouse performance is benchmarked with a cluster configuration described below. The specification of the nodes is described in the node configuration details table.

## Compute cluster

1. 3 Control Plane Nodes
2. 1 Coordinator Node
3. 6 Worker Nodes, or 11 Worker Nodes, or 16 Worker Nodes

## Storage cluster

4. ECS EX500, eight nodes

## Test Setup

In this section, we describe the test setup, including the queries we have used, how we generated data and a summary of the methodology we followed.

### Test Data and Queries

Performance validation is done from the queries derived from the TPC-DS benchmark. The TPC-DS benchmark data is modeled on the decision support functions of a retail product supplier and consists of seven (7) fact tables and seventeen (17) dimensions. We use the Scale Factor of 1000 which corresponds to 1TB of data. The 99 queries are divided into four broad classes:

- Reporting queries
- Ad-hoc queries
- Iterative OLAP queries
- Data Mining queries

### Test Data Generation

TPC-DS dataset is generated using TPC-DS dsdgen utility in CSV format. TPC-DS Dataset size is 1TB, that is Scale Factor 1K. The dataset is stored in an ECS bucket, the bucket is mounted as an NFS mount to the host where the dsdgen utility is run during the creation of the dataset. The SQL Create Table as Select (CTAS) statement is used to convert the generated CSV data and write it to ECS in Iceberg table format and Parquet file format using the Iceberg connector. The data is generated in both partitioned and nonpartitioned layout. Partitioned by date on all the seven (7) fact tables. Further partitioned data are compacted for higher read performance.

### Test Methodology

The Apache JMeter test suite is used to run benchmarking tests. Starburst Iceberg connector is used to connect to TPC-DS data in stored in Dell ECS.

### Tests Executed

Execute 3 iterations of 99 queries for each of the below mentioned test conditions.

### Metrics Collected

**Total execution time** = total time taken to execute 99 TPS-DS queries from the 3 runs

**Per query execution time** = time taken to execute each of the 99 TPC-DS queries.

**No. of queries in 10mins** = Number of Queries executed in 10mins

### Metrics Computed

**Average execution time for 99 queries** = average of execution time to run all 99 queries from the 3 runs

**Average execution time** = average execution time for each query from the 3 runs

### Observations Made

Variation in the number of queries executed in 10minutes for Partitioned and non-partitioned dataset.

Variation in Average execution time for 99 queries and Average execution time per query with increasing number of concurrent users 1, 5, 10. Each concurrent thread corresponds to a User and will execute the 99 TPC-DS queries in its entirety and independently of other test threads. For example, 10 concurrent users will run 10 parallel sessions of 99 queries each. Multiple threads are used to simulate concurrent users to Dell Data Analytics Engine. Apache JMeter Thread groups are used to achieve these concurrent users.

Variation in the execution time with the increase in the number of worker nodes, each iterations number of worker nodes is the DDAE cluster are scaled from default 6 to 11 to 15 worker nodes.

---

Note: Following testing is based on TPC-DS SF1K dataset


---

## DDAE Performance testing on Partitioned and Non-Partitioned Dataset

### Below Bar Chart illustrates:

Concurrency in the Data Analytics Engine didn't return any failures and didn't require rewrites to any SQL query indicating strong compliance with industry standard ANSI-SQL.

Partitioned tables are more performant than non-partitioned tables, it is recommended to partition the tables on Dell Data Lakehouse to get best performance.



## Concurrent users' impact on the Query execution


**Number of USERS Scaling test, below bar chart results illustrate:**

The concurrency is inversely proportional to the execution time.

Increase in the number of users proportionately decreases the execution of the query due to the resource sharing.

On the same cluster size, the number of queries executed for 10 minutes reduces with increases in number of users.

Recommended to take concurrency in the cluster sizing to adjust the number of nodes for number of users to meet the SLA.



## DDAE Nodes Scaling Impact on Execution Time


**Number of DDAE Worker Nodes Scaling test, below results illustrate:**

Compute Resources or No of DDAE Worker Nodes (CPU and RAM) directly proportional to the execution time.

Increase in the number of DDAE workers proportionately increases the execution of the query due to increase in the compute resource.

On the same cluster size, the number of queries executed for 10minutes increases with increases in number of DDAE Worker Nodes (increase in CPU and RAM).  
Recommended to take No of Worker Nodes in the sizing for number of users to meet the SLA.

DDAE Warp Speed



Dell Data Analytics Engine’s Warp Speed transparently adds an indexing and caching layer to enable higher performance. You can take advantage of the performance improvements by enabling Warp Speed. Your DDAE cluster nodes are provisioned with suitable hardware and configurations to setup Warp Speed utility connector for any catalog accessing object storage with the Hive, Iceberg, or Delta Lake connector.

Default acceleration

When a query accesses a column that is not accelerated, the system performs data and index materialization on the cluster to accelerate future access to the data in the column. This process of creating the indexes and caches is also called warmup. Warmup is performed individually by each worker based on the processed splits and uses the local high-performance storage of the worker. Typically, these are SSD drives.

When new data is added to a table or the index and cache creation are in progress, the new portions of the table that are not accelerated are served from the object storage. After the asynchronous indexing and caching is complete, query processing accessing that data is accelerated, because the data is available directly in the cluster from the indexes and caches, and no longer must be retrieved from the remote object storage.

This results in immediately improved performance for recently used datasets. In addition to the automatic default acceleration, advanced users can create specific warmup rules. The default acceleration has a lower priority than a user-created warmup rule.

Default acceleration is performed for SELECT \* FROM <table\_name> queries that are commonly used to explore a table rather than to retrieve specific data. The maximum number of accelerated columns in SELECT \* clauses is dynamic and depends on the column type.

Refer Starburst Warp Speed documentation for detailed information about the Warp Speed from [Starburst Warp Speed documentation](#).

## Warp Speed Performance Testing

As discussed above, the same TPC-DS benchmarking tool will be used to evaluate the performance of the Warp Speed-enabled DDAE cluster. In this testing TPC-DS dataset size is increased to 10TB that is TPC-DS Scale Factor 10,000 or 10K. The test setup and benchmarking process will remain consistent to ensure comparability. Observations will be made using the same metrics and criteria as before and will be compared against the performance of the **DDAE** cluster **without Warp Speed**. This approach will provide a clear and objective assessment of the improvements brought by the Warp Speed feature.

---


Note: Following testing is performed on TPC-DS SF 10K

---

### Warp Speed impact on the execution time


**Number of DDAE Worker Nodes Scaling test with and without Warp Speed, below bar chart results illustrate:**

- Partitioning tables impacts the execution time and number of workers required to execute on the dataset. Our test show 6 DDAE worker nodes are sufficient to execute all 99 queries of the TPC-DS for Scale Factor 10K that is 10TB dataset size.
- For the non-partitioned dataset, Warp Speed helps in the boosting the query execution significantly by indexing and caching the data locally.
- Compute Resources or No of DDAE Worker Nodes (CPU and RAM) directly proportional to the execution time.
- Warp Speed enabled cluster with same number of worker nodes shows significant performance increase compared to the cluster without Warp Speed.
- The 6 Worker Nodes execution time with Warp Speed is almost equal to 11 worker node DDAE cluster without warp speed.
- The Warp Speed enabled cluster performance boost impacts the cluster sizing, now cluster sizing of the worker nodes can be reduced by ~40% for the same number of workloads.
- Increase in the number of DDAE workers proportionately increases the execution of the query due to increase in the compute resource irrespective of Warp Speed feature.
- On the same cluster size, the number of queries executed for 10minutes increases with increases in number of DDAE Worker Nodes (increase in CPU and RAM), and further increases with Warp Speed by approximately 70% to 100%. The top 3 accelerated queries run 4x-5x faster.



**Warp Speed performance consistency across all 99 TPC-DS queries.**

Warp Speed shows consistent performance improvement across all 99 queries of TPC-DS. The bar chart below compares the performance of a 6-worker-node DDAE cluster with and without Warp Speed enabled. Majority of queries demonstrate enhanced performance in the Warp Speed-enabled cluster, highlighting the significant benefits of this feature.



## Dell Data Lakehouse sizing

Cluster sizing and scaling are important to understand when architecting the Dell Data Lakehouse solution. Sizing is concerned with ensuring the cluster meets the workload requirements for storage and processing throughput. Scaling is concerned with cluster growth over time as capacity needs increase.

### Sizing and scaling overview

The Dell Data Lakehouse architecture is a parallel scale-out system with decoupled compute and storage. Some sizing requirements can be addressed through scaling while others must be addressed through node level sizing.

Sizing and scaling of a cluster are complex topics that require knowledge of the workloads. This section emphasizes the main considerations that are involved but does not provide detailed recommendations for workload sizing. Design guides for specific workloads running on the platform include workload specific sizing guidance. A Dell Technologies or authorized partner sales representative can help with detailed sizing calculations.

There are many parameters that are involved in cluster sizing. The primary parameters are:

## Data volumes and growth rates

Data volume and growth rate of data volume has a significant impact on cluster sizing. The cluster must be sized to have sufficient memory to process any given query at a given user concurrency. This memory requirement may grow as queries become larger, or underlying datasets grow (e.g., additional months/years of data gets added). Data ingestion also impacts network utilization. Since the lakehouse storage is external to the cluster nodes, network bandwidth is required to access it. The processing throughput requirements must be considered as well as the data size.

## Memory and processor capacity

Memory and processor requirements for jobs running on the cluster must be considered when sizing. Memory and processor capacity increases as nodes are added to the cluster.

## Service-level agreements

Production cluster sizing must meet any performance requirements that SLAs specify. Critical path jobs that must meet a specific execution time or throughput may require adjusting the cluster sizing and balance between compute and storage accordingly. Overall cluster throughput is as important as storage capacity, and often influences the number of nodes independent of the required storage capacity.

## Storage capacity

While the Dell Data Lakehouse compute cluster does not include object storage, it relies on data lakehouse storage externally that also must be sized in sync with the compute cluster. Sizing the storage capacity relies on understanding how much raw or fresh data is expected to be ingested into the data lakehouse storage, as well as how much data will be materialized from other data sources. In addition to raw capacity, it is also important to understand performance of the storage needed to meet SLAs. The available network bandwidth between the compute and storage clusters must also be considered.

Bandwidth on the storage and compute clusters scales in direct proportion to the number of nodes. However, dense storage capacity is possible with ECS, ObjectScale, and PowerScale. Such density can result in a large storage capacity without enough bandwidth to support the lakehouse data transfer requirements. An analysis of workload data transfer requirements is necessary to correctly size the storage for both capacity and bandwidth. The network architecture allows both compute and storage clusters to use the same fabric. This configuration enables the network bandwidth to scale as either storage or worker nodes are added. The bandwidth available to the external storage systems should also be considered when referring to external storage that is not connected to the core cluster data network.

The good news is that the processing capacity of the lakehouse is decoupled from storage and can be adjusted at any time during the life of the cluster.

## Sizing guidelines Cluster sizing

Below table lists some cluster level starting points for possible deployments.

**Table 5. Example cluster configurations**

| Configuration | Size         |         |         |            |                        |
|---------------|--------------|---------|---------|------------|------------------------|
|               | PoC/<br>Demo | Entry   | Growth  | Enterprise | Advanced<br>Enterprise |
| Control Plane | 3 Nodes      | 3 Nodes | 3 Nodes | 3 Nodes    | 3 Nodes                |
| Coordinator   | 1 Node       | 1 Node  | 1 Node  | 1 Node     | 1 Node                 |
| Worker        | 2 Nodes      | 3 Nodes | 6 Node  | 11 Node    | 16 Node                |

We assume the average TPC-DS query fetches ~20% of the dataset. The below table for SF10K, which is equivalent to 10TB total dataset, highlights estimated query count and dataset size that a given size of the cluster could be expected to handle.

Table 6 highlights the results for the example configurations based on the performance tests highlighted in the previous section.

It is evident that WarpSpeed enabled clusters provide a higher level of performance both in terms of more queries executed, as well as a larger data set processed within the same resource context.

It is important to note that these results are highly dependent on network latency, throughput of the underlying S3 storage, and conflicts with any other workload being run on the underlying storage. Such results may also be different when the data source is federated (e.g., a relational database, a data warehouse, etc.) instead of an S3 storage.

**Table 6. Performance results of example configurations**

| Configuration         |                      | Size          |         |                |             |                        |
|-----------------------|----------------------|---------------|---------|----------------|-------------|------------------------|
|                       |                      | PoC /<br>Demo | Entry   | Growth         | Enterprise  | Advanced<br>Enterprise |
| Control Plane         |                      | 3 Nodes       | 3 Nodes | 3 Nodes        | 3 Nodes     | 3 Nodes                |
| Coordinator           |                      | 1 Node        | 1 Node  | 1 Node         | 1 Node      | 1 Node                 |
| Worker                |                      | 2 Nodes       | 3 Nodes | 6 Nodes        | 11 Nodes    | 16 Nodes               |
| Without Warp<br>Speed | Avg. # of<br>queries | N/A           | N/A     | 1.7<br>queries | 2.7 queries | 3.6 queries            |

|                                   |                              |     |     |             |             |             |
|-----------------------------------|------------------------------|-----|-----|-------------|-------------|-------------|
| (concurrency = 1)                 | per minute                   |     |     |             |             |             |
|                                   | Data processed per min       | N/A | N/A | 3.4 TB      | 5.4 TB      | 7.2 TB      |
| With Warp Speed (concurrency = 1) | Avg. # of queries per minute | N/A | N/A | 2.6 queries | 3.6 queries | 4.8 queries |
|                                   | Data processed per min       | N/A | N/A | 5.2 TB      | 7.2 TB      | 9.6 TB      |

### Cluster Size PoC/Demo for Proof of concept

The proof of concept (POC) configuration is a minimal configuration for basic evaluation. In this scenario, three control plane nodes are for Dell Data Lakehouse system software and 2 worker nodes are for the analytics engine. This configuration provides limited resources for workloads but is adequate for basic functionality evaluation.

The Dell Technologies Customer Solution Centers can be easily engaged for a demo of the solution, or work with the Dell Account Representative for a POC.

### Cluster Size Entry or Growth

The entry configuration is the smallest production grade configuration. In the Entry cluster, three dedicated control plane nodes host the Dell Data Lakehouse System Software, and 3 to 6 worker nodes (+ 1 coordinator node) run the Data Analytics Engine.

Dell Technologies recommends it for preproduction or development and test usage. This configuration provides enough resources to support one or two teams running analytics workloads.

A PoC/Demo cluster can be scaled to an Entry or Growth cluster simply by adding additional worker nodes.

### Cluster Size Enterprise and Advanced Enterprise

The Enterprise and Advanced Enterprise configuration is a medium to large production grade configuration. Just like Entry / Growth, three control plane nodes host the Dell Data Lakehouse System Software, and 11 to 16 worker nodes (+1 coordinator node) run the Data Analytics Engine.

This configuration provides substantial resources for running analytics workloads supporting multiple teams. This cluster can scale even further with more worker nodes as needed to support petabyte-scale workloads.

## Scaling guidelines

### Scaling overview

During the lifetime of the system, it is typical to scale the platform to support larger workloads or increase compute and storage capacity. The Dell Data Lakehouse architecture is designed to scale at the compute, storage, and workload levels. The

design incorporates network scaling as part of the infrastructure scaling. Compute and storage can be scaled independently.

### Compute scaling

Compute scaling is accomplished by adding additional DDAE nodes to the cluster. The DDAE node installation and provisioning the nodes become parts of the cluster.

### Storage scaling

Storage scaling is accomplished by adding or upgrading nodes in the ECS, ObjectScale, or PowerScale storage cluster, using the storage cluster management tools.

### Network scaling

The Dell Data Lakehouse architecture scales network bandwidth as compute or storage nodes, such as ObjectScale nodes, are added. When scaling either compute or storage, the impact to the network performance must be considered to maintain consistent SLAs and avoid any bottlenecks.

Substantial changes in expected data transfer volumes should be considered to ensure that the available bandwidth on the compute and storage clusters is aligned.

## Conclusion

The Dell Data Lakehouse is a state-of-the-art platform tailored to meet the rigorous needs of advanced analytics and AI/ML workloads. By combining its established strengths with newly added capabilities, it offers an industry-leading solution for modern organizations. Its key features include:

- **Dell Data Analytics Engine powered by Starburst:** A high-performance distributed query engine that allows seamless querying across diverse data sources using standard SQL, eliminating the need for extensive data movement.

- **Dell Data Processing Engine powered by Apache Spark:** Apache Spark-based distributed processing capabilities that support Streaming, Batch, AI/ML, Graph, and Geospatial analytics, enabling advanced and real-time data operations.
- **Warp Speed Performance:** Accelerates query execution to handle complex and high-volume workloads with exceptional speed and efficiency.
- **Unified Metadata and Governance Layer:** Offers robust, enterprise-level metadata management and governance, supporting row and column-level access control and data masking for secure operations.
- **AI Data Preparation Layer:** Simplifies workflows by preparing data for AI and ML models, allowing data scientists to unlock actionable insights quickly and effortlessly.
- **High-Performance Dell Hardware Infrastructure:** Combines PowerEdge R660 servers, PowerScale, ECS, and ObjectScale storage solutions with optimized networking using PowerSwitch for robust reliability and scalability.
- **Scalable and Validated Cluster Architecture:** Includes predesigned and validated guidance for the cluster server configuration, physical and logical networks, and scaling practices to support evolving workloads.
- **Smart Cache and Materialized Views:** Streamlines data access and enhances efficiency with built-in caching mechanisms and materialized view functionalities for faster data insights.

With the Dell Data Lakehouse's blend of intelligent design, cutting-edge technology, and unmatched scalability, organizations are empowered to streamline operations, drive innovation, and make data-driven decisions with confidence. This platform sets a new standard for unlocking value from data and shaping the future of analytics and AI-driven transformation.

# References

## Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

Additional information can be obtained at the [Dell Technologies Info Hub for Data Analytics](#). If you need additional services or implementation help, contact your Dell Technologies sales representative.

| Document type                | Location                                                           |
|------------------------------|--------------------------------------------------------------------|
| Dell Data Lakehouse          | <a href="#">Dell Data Lakehouse Sizing and Configuration Guide</a> |
|                              | <a href="#">Dell Data Lakehouse Validation Guide</a>               |
|                              | <a href="#">Dell Data Lakehouse Spec Sheet</a>                     |
| Server specification sheets  | <a href="#">PowerEdge R660 Spec Sheet</a>                          |
| Storage Specification Sheets | <a href="#">ECS EX500 Spec Sheet</a>                               |
|                              | <a href="#">ECS EXF900 Spec Sheet</a>                              |
|                              | <a href="#">ObjectScale Solution Overview</a>                      |
|                              | <a href="#">PowerScale H7000 Spec Sheet</a>                        |
| Switch Specification Sheets  | <a href="#">PowerSwitch S3100 Series Spec Sheet</a>                |
|                              | <a href="#">PowerSwitch S5200-ON Series Spec Sheet</a>             |
|                              | <a href="#">PowerSwitch Z9264F-ON Spec Sheet</a>                   |
| Server manuals               | <a href="#">PowerEdge R660 Manuals and Documents</a>               |
| Storage manuals              | <a href="#">ECS EX500 Manuals and Documents</a>                    |
|                              | <a href="#">ECS EXF900 Manuals and Documents</a>                   |
|                              | <a href="#">ObjectScale Overview and Architecture</a>              |
|                              | <a href="#">PowerScale H7000 Manuals and Documents</a>             |
| Switch manuals               | <a href="#">PowerSwitch S3100 Manuals and Documents</a>            |
|                              | <a href="#">PowerSwitch S5200-ON Series Manuals and Documents</a>  |
|                              | <a href="#">PowerSwitch Z9264F-ON Manuals and Documents</a>        |

## Delta Lake documentation

The following documentation on the [Delta Lake documentation website](#) provides additional and relevant information.

**Table 7. Delta Lake documentation**

| Document type                             | Location                                                                                                         |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| Lakehouse architecture introductory paper | <a href="#">Lakehouse: A New Generation of Open Platforms that Unify Data Warehousing and Advanced Analytics</a> |
| Delta Lake project                        | <a href="#">Delta Lake Project website</a>                                                                       |
| Delta Lake documentation                  | <a href="#">Delta Lake documentation website</a>                                                                 |

## Apache Iceberg documentation

The following documentation on the Apache Iceberg documentation website provides additional and relevant information.

**Table 8. Apache Iceberg documentation**

| Document type        | Location                       |
|----------------------|--------------------------------|
| Iceberg table format | <a href="#">Apache Iceberg</a> |

## Prior Reference Architecture for Starburst with Dell Infrastructure

The following reference architecture was posted in 2023 highlighting the integration between Starburst and Dell infrastructure. This document also includes performance validation results.

**Table 9. Prior Reference Architecture for Starburst with Dell Infrastructure**

| Document type          | Location                                                                                                     |
|------------------------|--------------------------------------------------------------------------------------------------------------|
| Reference Architecture | <a href="#">Reference Architecture—Multicloud Data Analytics with Dell Technologies Powered by Starburst</a> |

## Dell Technologies Info Hub

The [Dell Technologies Info Hub](#) is your one-stop destination for the latest information about Dell Solutions products. New material is frequently added, so browse often to keep up to date on the expanding Dell portfolio of cutting-edge products and solutions.

## More information

For more information, including sizing guidance, technical questions, or sales assistance, email [Analytics.Assist@dell.com](mailto:Analytics.Assist@dell.com), or contact your Dell Technologies or authorized partner sales representative.