DELL Technologies

Solution Overview

# Maximize the value of your data estate with Dell unstructured storage and Microsoft SQL 2022 S3 integration

As the data landscape continues to grow exponentially, it drives the question: how can data stewards efficiently store and present data to end users to help them make critical business decisions? Real-time and historical data can provide valuable insights, but due to the overwhelming size, format, and complexity of the data set, most data estates cannot effectively support the business and deliver the insights needed.

## Benefits of SQL Server coupled with Dell unstructured data storage

- **Internet-invented storage:** use purpose-built S3 storage to move massive volumes of data
- **Geo-distribution:** span multiple locations with strong consistency
- **Scale:** gain unlimited storage within a single namespace
- **Cloud compatibility:** embrace a wide range of S3 compatible applications; move data with ease within a multi-cloud solution architecture
- **Data virtualization and tiering:** support cold and hot data, increasingly important for larger structured and unstructured data systems

Object-based storage answers those demands with a flexible architecture that manages data as object within a storage pool, rather than in a file hierarchy or block storage architecture. Object storage provides immense scale while helping to alleviate many public cloud limitations. The object data needs to be fluid, possess geo-redundancy attributes, reside close to the application and be protected.

## Dell and Microsoft: Better together for unstructured data

Dell ECS has long been a proven leader and trusted name for unstructured data storage, with rich S3 (Simple Storage Services) compatibility. Dell PowerScale also delivers multi-protocol support, including S3 object. Both solutions offer on-premises and multi-cloud architectures, supporting the evolving object solution attributes required by application architects and data architects – including the ability for on-premises sources of external data stores to be consumed by any relational or non-relational database.

The world's most widely used database, Microsoft SQL Server, is now embracing S3 object storage connectivity with the next-gen release of SQL Server 2022. This Data Virtualization allows for a logical join query, authored with T-SQL, to be executed against relational and external data (both structured and unstructured) residing in another location. Integration is via a new S3 connector for the SQL engine uses the S3 REST API to connect to S3-compatible storage.

Dell Technologies teams have validated both ECS and PowerScale (S3 object protocol) as valid choices for SQL 2022 external data stores. This validation includes CSV and Parquet files, extending into the billons of rows. Query response times are more than adequate for these large data sets, showing that this combined topology is a true data hub.

## Backup and Restore scenario, benefits and requirements

Backup and restore operations for SQL Server are often mission critical. Accessing Dell Technologies S3 connected object storage uses the same native T-SQL backup and restore scripting for traditional backup and restore scenarios. The only difference when reading from or writing to object storage with S3 is a TO URL (Backup) and FROM URL (Restore) command.

Freeing up in-memory data page space for SQL Server also simplifies the process of moving colder, read-only data, to S3. The SQL engine is designed to do work and computations. Freeing up SQL Server to perform these operations can have a very beneficial result on your SQL Server data estate.

These are the minimum technical requirements for SQL backup and restore, along with data visualization, to an S3-compatible destination:

- SQL Server scoped credential for backup and restore
- SQL Server database scoped credential for data visualization from S3 Object
- TLS configured, HTTPS URL and the Dell S3 Object endpoint validated by a certificate installed on the SQL Server OS host or container
- At least one bucket (object container) on Dell S3 Object storage that has been configured from the administration console
- S3 permissions of ListBucket and WriteOnly for the SQL user or SQL Server service
- The URL prefixed with s3:// to denote the S3 connector is being used
- Support for filegroup backups and piecemeal restores

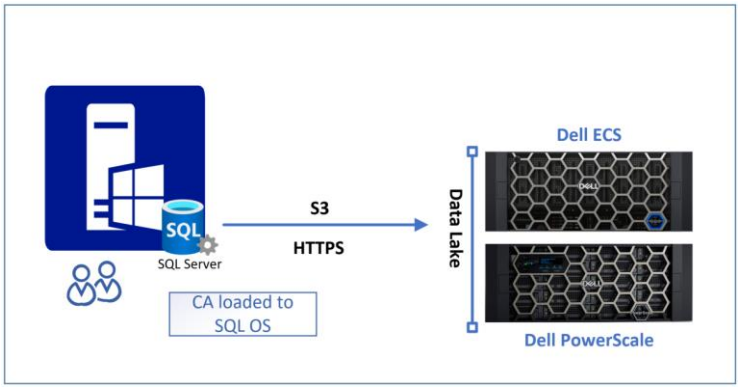


***Solution Architecture: Dell Unstructured Storage with SQL Server 2022 S3 Integration***

## Conclusion

As organizations large and small seek to gain an edge with their intelligent data estate, they also need access to all types of data sets. SQL Server 2022 and Dell unstructured data storage (including ECS and PowerScale) is the preferred technology for querying the data Lake via a T-SQL surface area. This combination of products and tools yields modern opportunities to store and manage different types of data on-premises and at public cloud scale.

Embracing both Dell unstructured data storage and Microsoft SQL Server 2022 can provide limitless scale to your data estate with a true data hub architecture. With both products in your solution architecture, your data stewards will do even more with your data, provide deeper insights and foster next-gen solutions.

Learn more about Dell Unstructured Data Solutions

Contact a Dell Technologies Expert

View a demo

Join the conversation with #DellStorage