

# AWS DATABASE MIGRATION SERVICE IS NOW SERVERLESS

## Serverless DMS brings a host of benefits

### AT A GLANCE

AWS Database Migration Service (DMS) was first announced in the fall of 2015 and launched in 2016 as a tool that facilitated relational data store replication or migration to the AWS cloud with minimal or no downtime. At the time, support was limited to just a handful of databases.

Fast forward to 2023, AWS DMS now supports an expanded list of data sources. AWS DMS is great for homogeneous data migrations (where the data source and target source engines are the same). It has enhanced support for heterogeneous data migrations using DMS Schema Conversion or AWS Schema Conversion Tool (AWS SCT) that can run against your target data source objects, application code, and stored procedures. There is also a serverless offering of the product that automatically provisions and scales capacity based on your transactional migration and data replication needs.

### SETTING THE STAGE

According to recent studies, the amount of data generated worldwide has been growing exponentially. Businesses depend on these valuable data to drive value and business decisions.

Several data trends have surfaced, including:

- **Cloud migration:** Organizations are increasingly moving their data to the cloud for improved scalability, agility, and cost efficiency. Cloud providers offer robust migration services that simplify the process and ensure data security.
- **Data lake migration:** With the rise of big data and analytics, organizations are consolidating their data into data lakes for centralized storage and analysis. Data migration to data lakes involves moving data from various sources, transforming it, and integrating it into a unified schema.
- **Real-time data replication:** Organizations require real-time access to their data for analytics, decision-making, and business intelligence. Real-time data replication allows continuous syncing of data from multiple sources to a target system, ensuring up-to-date and accurate information. (Continued on next page.)

### WHAT THIS MEANS FOR YOU

One of the key benefits of AWS DMS is how simple it is to use, allowing users to set up and use AWS DMS with just a few steps and clicks within the console. Proper research and planning should be considered before taking these steps. It is recommended to research and plan appropriately to determine the required capacity before migrating, while realizing that it can be a challenge to plan ahead of time against changing variables, especially when migrating multiple workloads in addition to continuously replicating data. Once AWS DMS has been set up and is running, monitor usage and scale capacity to ensure optimal performance against cost.

Now with the release of AWS DMS Serverless, the initial research and planning have been made easier. Manual monitoring and adjusting or scaling running workloads is minimized, giving users the freedom to focus on other differentiating tasks. AWS DMS Serverless automatically provisions, monitors, and scales migration resources to the optimal capacity needed to meet demand. (Continued on next page.)

## SETTING THE STAGE (CONT'D.)

- Database modernization: Many organizations are modernizing their legacy databases by migrating them to more advanced and scalable database technologies, such as NoSQL or cloud-native databases. This trend enables better performance, scalability, and flexibility.
- Data privacy and compliance: Especially in highly regulated industries, an organization's data migration processes must adhere to strict data privacy regulations, such as General Data Protection Regulation (GDPR), Payment Card Information (PCI), or Health Insurance Portability and Accountability (HIPAA).

At the heart of these trends, two shifts have emerged to the top of consideration that involve data residency, databases, and data store consideration. The first is moving and embracing open-source solutions to reduce vendor dependency and/or lock-ins as well as costly licensing. This allows users to develop and leverage community development within their specific timelines of priority. The second is the shift to purpose-built databases or data stores. Relational databases are no longer viewed as the only solution or answer to all data needs.

## AWS DATA MIGRATION SERVICES (AWS DMS)

Data migration initiatives that result from these drivers can be challenging. To address the challenge and to facilitate the vast amount of data being moved and created, and to address their reliability, security, availability, and performance requirements, AWS offers AWS Database Migration Service (AWS DMS).

AWS DMS is a cloud service that allows users to migrate relational databases, data warehouses, NoSQL databases, popular open-source databases, and other data stores. It enables data migration into the AWS cloud or between combinations of cloud and on-premises setups by defining the source and target endpoints. The only requirement for the combination potential is that an AWS service is defined within the combination as a source or target endpoint.

AWS DMS is a logical replication server, allowing users the freedom to identify specific ranges of data to entire schemas and databases to migrate, which is identified in a task that can consist of three major phases:

- Migration of existing data (Full load)
- The application of cached changes
- Ongoing replication (Change Data Capture)

AWS DMS integrates with other tools such as DMS Fleet Advisor and AWS SCT. DMS Fleet Advisor is used to discover the source data infrastructure by collecting data from on-premise databases and analytic servers. DMS Schema Conversion can be used to convert source schemas to a new target engine, or users can use the AWS Schema Conversion Tool (AWS SCT) to perform schema conversions locally. Once the source schemas have been converted, AWS DMS can be used to migrate the data, either as a one-time migration or ongoing replication to keep the sources and targets in sync.

If migrating to a new target database, continuous replication exists until application switchover is achieved. If the use case requires ongoing data replication, AWS DMS is used to keep a source and target in sync indefinitely.

## WHAT THIS MEANS FOR YOU (CONT'D.)

It starts by scanning the data source to determine the optimal capacity initially to set up and changes this based on the volume of data transactions over time. DMS Serverless supports popular DMS use cases, including continuous data replication, database consolidation, and migrations, even if the source and target database engines differ. It aids even the most complex migrations, including migrating tens or even hundreds of workloads simultaneously or completing ongoing data replications on AWS.


AWS DMS Serverless is cost-effective, as one no longer needs to overprovision resources to peak demand. The user only pays for the data migration capacity used, with the user specifying the minimum. Review the AWS DMS Serverless documentation to learn more.

For homogeneous migrations between like-to-like or compatible engines, AWS DMS also offers the option to use built-in native tooling with automatic scaling for seamless database migration. Only the hours used are billed.

## NORDIC IN ACTION

Switching to AWS DMS serverless is done within the AWS DMS console. There is a new option called “Serverless replications” under the Migrate data menu.

**FIG. 1** AWS DMS serverless replication process

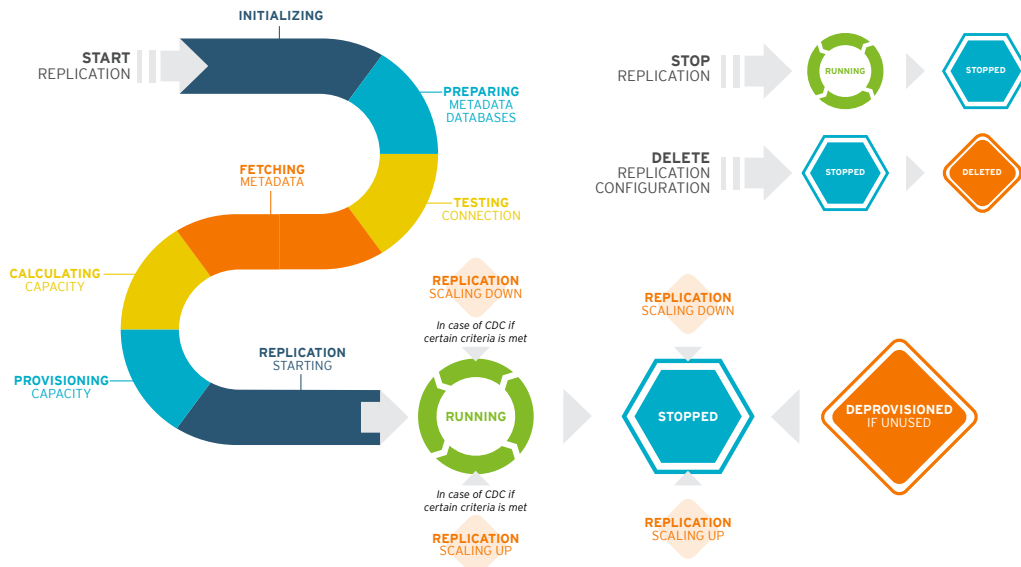


There are many similarities between standard and serverless, as they target the same goal and use cases. It is more important to understand the differences. For example, while one would need to first create and define the endpoints with both solutions, the current release of AWS DMS Serverless does not support all source and target endpoint types that the standard version supports. It is recommended to consult the AWS DMS Serverless components resource for updates on endpoint types currently supported. If existing endpoints are already defined by supported endpoint type, their use can be continued to define the serverless replications.

For Serverless replications, replications are specified instead of replication instance types or tasks to define a replication as well as the minimum and maximum DMS capacity units (DCUs). A DCU is 2GB of RAM and is a unit that is billed. DMS Serverless will scale between, but not exceed, these table mappings based on predicted size of the workload. It is recommended that to identify the highest DCUs that would ever be approved and needed and set this from the beginning to allow DMS serverless the maximum flexibility to ensure optimal performance as DMS will not exceed the defined thresholds on either end.

Autoscaling can result in as much as a 5-minute delay or pause in the replication as the scaling is happening. This should be a set expectation to the autoscaling toward any concerns if this pause is detected. Monitoring can be done directly in the DMS console. Figure 2 is a representation of the replication state transitions that can be reflected in the console while monitoring the status. (Continued on next page.)

**FIG. 2** Replication state transitions



**LIEN HUYNH**  
Solution Architect, AWS



**SARAVANAN SONAISAMY**  
Cloud Solutions Lead, Digital Health



**KEVIN ERDAL**  
Practice Leader, Digital Health

CloudPress@NordicGlobal.com

## NORDIC IN ACTION (CONT'D.)

It is a best practice to automate whenever and wherever possible, including when leveraging serverless options. Switching to AWS DMS Serverless is no exception, as it automates the scanning and scaling of capacity based on an organization's transactional needs. However, limitations to this serverless release do exist and must be considered when determining if this is the appropriate time to start or transition to AWS DMS serverless. Over time it is expected that the limitations or gaps between the standard and serverless offering will shrink. Therefore, it is recommended to review the current version's limitations often and at each consideration.

## REFERENCES

1. Amazon Web Services. (2023). *AWS Database Migration Service Features*. <https://aws.amazon.com/dms/features/>
2. Amazon Web Services. (2023). *AWS Database Migration Service Serverless*. <https://aws.amazon.com/dms/features/#Serverless>
3. Amazon Web Services. (2023). *Homogeneous data migrations in AWS Database Migration Service (AWS DMS)*. <https://docs.aws.amazon.com/dms/latest/userguide/data-migrations.html>
4. Amazon Web Services. (2023). *New - AWS DMS Serverless: Automatically Provisions and Scales Capacity for Migration and Data Replication*. <https://aws.amazon.com/blogs/aws/new-aws-dms-serverless-automatically-provisions-and-scales-capacity-for-migration-and-data-replication/>
5. Amazon Web Services. (2023). *Modifying AWS DMS serverless replications*. [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Serverless.Components.html#CHAP\\_Serverless.modify](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Serverless.Components.html#CHAP_Serverless.modify)
6. Amazon Web Services. (2023). *New - AWS DMS Serverless: Automatically Provisions and Scales Capacity for Migration and Data Replication*. <https://aws.amazon.com/blogs/aws/new-aws-dms-serverless-automatically-provisions-and-scales-capacity-for-migration-and-data-replication/>
7. Amazon Web Services. (2023). *AWS DMS Serverless components*. [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Serverless.Components.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Serverless.Components.html)
8. Amazon Web Services. (2023). *Serverless limitations*. [https://docs.aws.amazon.com/dms/latest/userguide/CHAP\\_Serverless.Limitations.html](https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Serverless.Limitations.html)
9. IDC, & Statista. (2021, June 7). *Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025 (in zettabytes) [Graph]*. Statista. <https://www.statista.com/statistics/871513/worldwide-data-created/>