

The background features a dark blue gradient with abstract geometric shapes. On the left, a large triangle is formed by a vertical orange line and a diagonal orange line. On the right, a large curved shape in shades of blue and orange sweeps across the frame. The text is positioned in the upper right area.

AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

DAT210

What's new with Amazon DynamoDB

Chad Tindel

Principal NoSQL Specialist SA
Amazon Web Services



Agenda


Do more with Amazon DynamoDB using new features

Customer examples

Recap of DynamoDB releases in 2021


Key benefits of Amazon DynamoDB

FAST AND FLEXIBLE NOSQL DATABASE AT ANY SCALE



Performance at scale

- Consistent, single-digit millisecond read and write performance
- Nearly unlimited throughput and storage




Enterprise ready

- Data encryption at rest
- Global replication
- Up to 99.999% availability SLA



No servers to manage

- Fully managed serverless database
- Massive scalability




Built-in integration with other AWS services


- Logging, monitoring, and analytics
- Applications that span multiple AWS services

Optimize DynamoDB costs for infrequently accessed data

Data is growing fast



Growing
exponentially




Increasingly
diverse



Source of new
revenue

Data lifecycle




- Data volume is growing fast
- Data relevance decreases over time
- Older data gets less frequently accessed
- Storing data can get more expensive at scale

Common use cases for infrequently accessed data




Social media

Active users expect older posts to be available whenever they want, immediately



Data analytics

Businesses need to capture and refine billions of data points to deliver the most accurate and actionable data analytics



Retail

Online shoppers sometimes want to look up their past orders, re-order the same item, or get product information anytime

Table costs over time

- Data storage grows, but older data is accessed less frequently
- Storage cost grows to exceed throughput (reads and writes) cost






Table costs over time

- Data storage grows, but older data is accessed less frequently
- Storage cost grows to exceed throughput (reads and writes) cost




Option 1: Deleting old data



- Limits the amount of data retained
- Impacts customers' experience
- Miss out on valuable insights

Option 2: Splitting the data



- Tradeoffs: latency and performance
- Build and manage custom solutions to move data
- Access data in two service points using different APIs

DynamoDB Standard-Infrequent Access

NEW

NEW DYNAMODB TABLE CLASS




Announcing the new Amazon DynamoDB
Standard-Infrequent Access
(DynamoDB Standard-IA) table class

Reduce your total costs by up to 60%
while retaining all your data in DynamoDB

Key benefits of DynamoDB Standard-IA

NEW



Lower storage costs

The Standard-IA table class offers 60% lower storage costs than existing DynamoDB Standard tables



No performance tradeoffs


Standard-IA tables offer the same performance, durability, data availability, and massive scalability as existing DynamoDB Standard tables



No developer overhead


Switch between table classes with a single click in the DynamoDB console or using AWS CLI or AWS SDK

More flexibility using table classes




Optimizing costs with table classes

Before optimization




After optimization




Optimizing costs with table classes

Before optimization



After optimization



When to use DynamoDB Standard-IA?



CHECK THAT YOUR USE CASE CAN BENEFIT FROM COST SAVINGS



When storage exceeds 50% of your throughput (reads and writes) cost, DynamoDB Standard-IA can help you reduce your table's cost

Analyze costs Log in to AWS Management Console; use AWS Cost and Usage Reports and AWS Cost Explorer to analyze your table's cost structure

Look at AWS Cost Explorer





“ Amazon DynamoDB Standard-IA will provide us with the ability to store our users’ infrequently accessed data at a significant cost savings while continuing to deliver for our users by maintaining the same high performance, accessibility, and reliability we’ve come to expect from Amazon DynamoDB. ”

Oscar Mullin

Director of IT – Core Services SRE & DBA Head, Mercado Libre



“ The ability to simplify the management and access to our long-term data storage while still benefiting from DynamoDB performance, durability, and data availability with the DynamoDB Standard-IA table class could help us further optimize costs and provide an even better user experience to our customers. ”

Ofir Nir

Head of Data Infrastructure, Singular

JungleScout

“ The ability to switch between DynamoDB table classes without any code changes will allow us to easily optimize our costs as we scale and focus our engineering efforts on the features our customers require as they grow their business. ”

Regan Wolfrom

DevSecOps and Builder Tools Manager, Jungle Scout

Customer examples

Sample customer scenario

Challenge: A dataset that is large at initial ingestion or has grown steadily over time

Solution: The DynamoDB Console now offers the option to create tables with the Standard-IA table class, but this is not the default and almost all use cases will be best starting out with the Standard table class

For this scenario, we're creating a new table with DynamoDB Standard table class

Create table: DynamoDB Standard table class

Create table

Table details [Info](#)

DynamoDB is a schemaless database that requires only a table name and a primary key when you create the table.

Table name

This will be used to identify your table.

Between 3 and 255 characters, containing only letters, numbers, underscores (`_`), hyphens (`-`), and periods (`.`).

Partition key

The partition key is part of the table's primary key. It is a hash value that is used to retrieve items from your table and allocate data across hosts for scalability and availability.

String ▼

1 to 255 characters and case sensitive.

Sort key - optional

You can use a sort key as the second part of a table's primary key. The sort key allows you to sort or search among all items sharing the same partition key.

String ▼

1 to 255 characters and case sensitive.

Create table: DynamoDB Standard table class

Settings

☐ **Default settings**
The fastest way to create your table. You can modify these settings now or after your table has been created.

☒ **Customize settings**
Use these advanced features to make DynamoDB work better for your needs.

Table class

Select table class to optimize your table's cost based on your workload requirements and data access patterns.

Choose table class

☒ **DynamoDB Standard**
The default general-purpose table class. Recommended for the vast majority of tables that store frequently accessed data, with throughput (reads and writes) as the dominant table cost.

☐ **DynamoDB Standard-IA**
Recommended for tables that store data that is infrequently accessed, with storage as the dominant table cost.

Data volume grows . . .

The data volume might grow quickly (such as an initial ingestion)
or can grow steadily over an extended period

We now have ~40M items and ~40 GB of data

Items summary DynamoDB updates the following information approximately every six hours.		Get live item count	View items
Item count 40,000,000	Table size 39.4 gigabytes	Average item size 984.88 bytes	

But throughput is stable . . .

Read and write throughput consumption has stabilized into a relatively predictable pattern and we expect this to continue

Read/write capacity

The read/write capacity mode controls how you are charged for read and write throughput and how you manage capacity.

Edit

Capacity mode	
Provisioned	
Table capacity	
Read capacity auto scaling	Write capacity auto scaling
On	On
Provisioned read capacity units	Provisioned write capacity units
35	5
Provisioned range for reads	Provisioned range for writes
30 - 50	3 - 10
Target read capacity utilization	Target write capacity utilization
70%	70%

Cost analysis (us-west-2)

DynamoDB Standard-IA offers 60% lower storage costs than DynamoDB Standard where throughput is the dominant table cost

DynamoDB Standard offers 20% lower throughput costs than Standard-IA, for tables where throughput is the dominant table cost

	Units	Standard Rate	Standard-IA Rate	Standard	Standard-IA
Storage (GB)	40	\$ 0.25	\$ 0.10	\$ 10.00	\$ 4.00
Read Throughput	35 units*720hrs	\$ 0.00013	\$ 0.00016	\$ 3.28	\$ 4.03
Write Throughput	5 units*720hrs	\$ 0.00065	\$ 0.00081	\$ 2.34	\$ 2.92
Total				<u>\$ 15.62</u>	<u>\$ 10.95</u>

Cost analysis (us-west-2)

DynamoDB Standard-IA offers 60% lower storage costs than DynamoDB Standard where throughput is the dominant table cost

DynamoDB Standard offers 20% lower throughput costs than Standard-IA, for tables where throughput is the dominant table cost

	Units	Standard Rate	Standard-IA Rate	Standard	Standard-IA
Storage (GB)	10	\$ 0.25	\$ 0.10	\$ 2,500	\$ 1,000
Read Throughput	3,000 units*720hrs	\$ 0.00013	\$ 0.00016	\$ 280	\$ 345
Write Throughput	1,000 units*720hrs	\$ 0.00065	\$ 0.00081	\$ 468	\$ 583
Total				<u>\$ 3,248</u>	<u>\$ 1,928</u>


Cost analysis (us-west-2)

DynamoDB Standard-IA offers 60% lower storage costs than DynamoDB Standard where throughput is the dominant table cost


DynamoDB Standard offers 20% lower throughput costs than Standard-IA, for tables where throughput is the dominant table cost

	Units	Standard Rate	Standard-IA Rate	Standard	Standard-IA
Storage (GB)	100	\$ 0.25	\$ 0.10	\$ 25,000	\$ 10,000
Read Throughput	3,000 units*720hrs	\$ 0.00013	\$ 0.00016	\$ 280	\$ 345
Write Throughput	1,000 units*720hrs	\$ 0.00065	\$ 0.00081	\$ 468	\$ 583
Total				<u>\$ 25,748</u>	<u>\$ 10,928</u>

Will table class update disrupt operations?



No, we already switched table classes . . .



Meet your business continuity and regulatory compliance goals

Backup options on DynamoDB

Use on-demand backups for compliance and long-term data archival and PITR (continuous backups) to protect against accidental deletions and updates

- **Easy to use:** fully automated, can be enabled with a single click or an API call
- **No performance impact:** does not consume table capacity and has zero impact on the performance or availability of production applications
- **Scalable:** back up tables of any size


Common use cases

Disaster recovery



Protect your critical data by copying backups across AWS Regions and accounts

Compliance




Prove compliance of organizational data protection practices to auditors

Cost optimization



Easier cost allocation and lifecycle backups to lower cost tier

Current customers' challenges



Complexity

- Custom scripts to automate and copy table backup across Regions and accounts
- Ongoing development and maintenance to lifecycle backups to optimize cost



Management overhead

- Proving regulatory compliance to auditors is not standardized across AWS services
- Manual policy enforcement for data protection



Cost

- Time and resources allocated to building, maintaining, and supporting data protection tools
- Requires continuous oversight and enhancement


Enhanced integration with AWS Backup

NEW

AWS Backup: a fully managed, policy-based backup service that makes it easy to centrally manage and automate the backup of data across AWS services

Enhanced features with AWS Backup

NEW



Disaster recovery

- Continuous backups for point-in-time recovery (PITR)
- Protect critical data by copying backups across AWS Regions and accounts (NEW)
- WORM model to safeguard backups from inadvertent or malicious actions (NEW)



Compliance and security

- On-demand backups for long-term data archiving and compliance
- Centrally enforce backup policies and audit backup activity across AWS services (NEW)
- Secure backups with separate backup encryption (NEW)



Simplified management

- Automate backup scheduling and retention management with backup plans
- Simplify backup cost allocation with tags (NEW)
- Lifecycle backups to cold storage to reduce costs (NEW)

* To use these enhanced backup features, you need to opt in to have AWS Backup service manage your DynamoDB on-demand backups

Example

Create new on-demand backup

Create on-demand backup

Create a one-time snapshot backup of your table. Schedule automatic backups of your table in [AWS Backup](#) 

Source table [Info](#)

Source table

 reinvent_demo



Backup settings [Info](#)

A backup name will be created automatically.

☐ Default settings

Create a backup that stays in warm storage.

☒ Customize settings

Create a backup that can transition to cold storage and be deleted as it ages.

Backup management [Info](#)

☒ Backup with AWS Backup

Creates a backup with AWS Backup encryption and ARN. Includes options for cross-Region and cross-account copy, tags and cold storage.

☐ Backup with DynamoDB

Creates a backup with DynamoDB encryption and ARN. Additional features not supported.

Backup window

Specify when the backup begins.

☒ Create backup now

Starts within 1 hour

☐ Customize backup window



Some custom settings options

Transition to cold storage

Specify how long the backup is kept in warm storage before moving to cold storage.

Days ▼

7 ▲▼

Retention period

Specify how long to store your backups before automatically deleting them.

Years ▼

1

Backup vault

Specify the backup vault in AWS Backup where the backup is stored.

Default ▼

↺

Create backup vault in AWS Backup ↗

IAM role

Specify the IAM role that AWS Backup will assume when creating and managing backups on your behalf.

- ☒ **Default role**
If the AWS Backup default role is not present, it will be created for you.
- ☐ Choose an IAM role

Add tags for cost allocation

Tags - *optional*

AWS Backup copies tags from the DynamoDB table to the recovery point upon creation. You can specify additional tags to add to the recovery point.

Key

🔍 Project



Value - *optional*

🔍 ProjectFoo



Remove

Add new tag

You can add 49 more tags.

Copy backup to another Region or account

Copy 96ffec40-0d0f-4496-ae4c-1f8e33455a57 -
DynamoDB [Info](#)

Copy configuration

Copy to destination [Info](#)

US East (N. Virginia) ▼

☒ Copy to another account's vault

Destination Backup vault

The vault to which your backup copy will be made.

Default ▼

Create new Backup vault

Retention period [Info](#)

Always ▼

IAM role [Info](#)

Specify the IAM role that AWS Backup will assume when creating and managing backups on your behalf.

☒ Default role

If the AWS Backup default role is not present, one will be created for you with the correct permissions.

☐ Choose an IAM role

Use built-in reports

Create report plan [Info](#)

Audit and gain insight into your backup activity. Reports refresh every 24 hours.

Report templates [Info](#)

Set up automatic daily reports on your backup activity.

Choose a report template

Choose a template ▲

Q |

Backup activity reports

Backup jobs report

Backup jobs details for the past 24 hours

Restore jobs report

Restore jobs details for the past 24 hours

Copy jobs report

Copy jobs details for the past 24 hours

Compliance reports

Control compliance report

Control compliance status and details

Resource compliance report

Resource compliance status and details

Cancel

Create report plan

Key takeaways


Key takeaways

Use new DynamoDB Standard-IA table class to optimize your storage costs, with no impact on performance and no management overhead

Use the feature capabilities of AWS Backup to improve your DynamoDB business continuity plans and meet your regulatory compliance goals, with few clicks


Recap of DynamoDB 2021 releases

Recap of DynamoDB 2021 Releases




Security enhancements with detailed logging

Recap of DynamoDB 2021 Releases



Improved developer experience

Recap of DynamoDB 2021 Releases



Integration with CloudFormation for ease of use

Get started today

Learn more about these new features on our website or in the DynamoDB Developer Guide

aws.amazon.com/dynamodb/features

Use AWS Cost and Usage Reports and AWS Cost Explorer to analyze your table's cost structure and determine the best table class for your use case

Get started using Standard-IA table class and the backup capabilities in the AWS Management Console, AWS CLI, or AWS SDK



Thank you!

Chad Tindel

Twitter: @ctindel

