

Service Name

Amazon DynamoDB

Service Overview

Amazon DynamoDB is a fully managed, serverless NoSQL database service that delivers single-digit millisecond performance at any scale. It supports both key-value and document data models with a flexible schema, allowing applications to store and retrieve large amounts of data without infrastructure management. DynamoDB automatically handles hardware provisioning, setup, configuration, replication, cluster scaling, patching, and monitoring. The service offers built-in security features, continuous backups, multi-region replication, in-memory caching, and seamless integration with other AWS services for building modern, scalable applications.

Key Use Cases

- **Adtech applications:** High-performance ad serving and real-time bidding platforms requiring low latency and horizontal scaling
- **Mobile and web applications:** Session management, user profiles, and gaming leaderboards for applications serving millions of users
- **Gaming applications:** Player data, game state management, and global leaderboards with multi-region support
- **Internet of Things (IoT):** Time-series data collection and device management for sensor networks and IoT applications
- **Real-time analytics:** Event-driven architectures using DynamoDB Streams for data processing and monitoring
- **Financial services:** Transaction processing, fraud detection, and trading platforms requiring ACID transactions
- **E-commerce and retail:** Product catalogs, shopping carts, and inventory management systems
- **Generative AI applications:** Storage for chat history, user preferences, and contextual data for AI-powered chatbots

Features

- **Serverless Architecture:** Zero infrastructure management with automatic scaling and pay-per-use pricing
- **Single-digit Millisecond Performance:** Consistent low-latency performance at any scale with optional microsecond latency via DAX
- **Global Tables:** Multi-active, multi-region replication for globally distributed applications with 99.999% availability

- **ACID Transactions:** Support for complex business logic with atomicity, consistency, isolation, and durability
- **DynamoDB Streams:** Real-time change data capture for event-driven architectures and data processing
- **Point-in-time Recovery:** Continuous backups with recovery to any second within the last 35 days
- **Encryption:** Server-side encryption at rest and in transit with AWS KMS integration
- **Auto Scaling:** Automatic capacity adjustment based on traffic patterns for cost optimization
- **Secondary Indexes:** Global and local secondary indexes for flexible query patterns
- **PartiQL Support:** SQL-compatible query language for familiar data access patterns

Value Proposition

DynamoDB offers a unique combination of serverless convenience, predictable performance, and enterprise-grade reliability that eliminates operational overhead while delivering consistent single-digit millisecond latency. Unlike traditional databases, it provides automatic scaling without capacity planning, built-in security and compliance features, and seamless integration with AWS services:

- The service's global tables enable multi-region active-active replication with up to 99.999% availability, while features like ACID transactions, point-in-time recovery, and DynamoDB Streams support complex business requirements
- Cost optimization through on-demand pricing, multiple table classes, and automatic scaling provides better economics than managing infrastructure, making it ideal for modern applications requiring scale, performance, and reliability

Service Integration

DynamoDB integrates seamlessly with core AWS services to enable comprehensive application architectures. AWS Lambda functions can be triggered by DynamoDB Streams for real-time data processing and event-driven workflows:

- Amazon S3 supports bulk import/export of data for analytics and backup scenarios
- Amazon Kinesis Data Streams provides advanced streaming capabilities for real-time analytics
- AWS CloudWatch offers monitoring, metrics, and alarms for performance optimization
- Amazon OpenSearch Service enables full-text search and analytics through zero-ETL integration
- DynamoDB Accelerator (DAX) provides microsecond latency caching, while AWS Backup manages enterprise backup policies

- Integration with Amazon SageMaker and Amazon Redshift supports machine learning and data warehousing use cases through zero-ETL pipelines

Onboarding and Offboarding

Customers can start using DynamoDB immediately through the AWS Management Console, AWS CLI, NoSQL Workbench for DynamoDB, or AWS SDKs without any upfront setup or infrastructure provisioning. The service is included in the AWS Free Tier with 25 GB of storage, 25 read/write capacity units, and 2.5 million stream reads monthly:

- Getting started involves creating a table with a primary key, choosing between on-demand or provisioned capacity modes, and optionally configuring features like encryption, streams, or global tables
- AWS provides comprehensive tutorials, code samples, and the NoSQL Workbench tool for data modeling and application development
- Migration tools and best practices documentation help customers transition from relational databases or other NoSQL solutions

Getting Started Guide:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GettingStartedDynamoDB.html>

Data Removal

To offboard from DynamoDB, customers can delete tables through the AWS Management Console, CLI, or APIs, which permanently removes all data. Before deletion, customers should create backups using on-demand backup or export data to Amazon S3. Point-in-time recovery backups are automatically deleted 35 days after table deletion. For global tables, customers must remove all regional replicas before deletion. AWS provides data export capabilities and migration tools to facilitate moving to alternative solutions. All associated resources like streams, DAX clusters, and backups should be cleaned up separately to avoid ongoing charges.

Backup/Restore and Disaster Recovery

DynamoDB provides comprehensive backup and disaster recovery capabilities including continuous backups with point-in-time recovery (PITR) for up to 35 days with per-second granularity, and on-demand backups for long-term retention. Global tables offer active-active multi-region replication for disaster recovery with automatic failover capabilities. AWS Backup integration enables centralized backup management with cross-account and cross-region copying. All backups are automatically encrypted and stored redundantly across multiple Availability Zones for durability.

Backup Capabilities

DynamoDB offers native backup capabilities with point-in-time recovery providing continuous backups for 1-35 days, and on-demand backups for manual snapshots. AWS Backup is fully supported, enabling centralized backup policies, scheduling, retention management, and advanced features like cross-region backup copying and lifecycle management:

- Backups are automatically encrypted and can be restored to new tables with the same configuration
- System backups are automatically created when tables are deleted

Disaster Recovery

DynamoDB supports comprehensive disaster recovery through global tables that provide active-active replication across AWS regions with automatic failover and up to 99.999% availability. Point-in-time recovery enables restoration from accidental deletions or corruptions. While AWS Elastic Disaster Recovery is designed for EC2-based workloads, DynamoDB's native global tables and backup features provide superior disaster recovery for database workloads with multi-region active-active architecture and automatic conflict resolution.

Recovery Objectives

DynamoDB helps customers achieve aggressive RTO and RPO objectives through global tables with near-zero RTO for region failures due to active-active architecture, and RPO approaching zero with eventual consistency or zero RPO with strong consistency mode. Point-in-time recovery provides RPO of up to 1 second and RTO measured in minutes for data restoration. On-demand backups offer customizable RPO based on backup frequency, while the multi-AZ architecture within regions provides sub-minute recovery from infrastructure failures.

Service Constraints

Service Quotas:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html>

FAQ: <https://aws.amazon.com/dynamodb/faqs/>

Technical Requirements

Service Documentation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

User Guide: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/>

API Reference:

<https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/Welcome.html>

Pricing Overview

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Public Pricing: <https://aws.amazon.com/dynamodb/pricing/>

Cost Optimization

DynamoDB offers unique cost optimization options including on-demand mode with 50% reduced pricing that scales to zero when unused, eliminating capacity planning overhead. The Standard-IA table class reduces storage costs by up to 60% for infrequently accessed data:

- Auto scaling automatically adjusts provisioned capacity based on traffic patterns, preventing over-provisioning
- Reserved capacity provides up to 76% savings for predictable workloads
- Global tables now offer up to 67% savings with simplified pricing
- Database Savings Plans provide additional discounts with usage commitments
- The service's serverless nature eliminates infrastructure costs, while features like DynamoDB Streams and built-in security reduce the need for additional services and associated costs

Savings Considerations

Major cost savings considerations include choosing the right capacity mode based on traffic patterns, with on-demand being cost-effective for unpredictable or intermittent workloads, while provisioned capacity with auto scaling optimizes costs for steady traffic. Implementing proper data modeling reduces unnecessary read/write operations and storage costs:

- Using appropriate table classes like Standard-IA for infrequently accessed data, removing unused indexes and streams, and leveraging the free tier effectively reduces costs
- Cross-region replication costs should be evaluated against availability requirements, while monitoring tools help identify optimization opportunities
- The elimination of database administration overhead, infrastructure management, and the pay-per-use model provides significant operational cost savings compared to traditional database solutions