

Amazon DynamoDB in a Modern Resale Platform on AWS

Executive Summary

Recurate is a Series A startup (\$17.4M total funding amount) with a growing portfolio of clients. The company offers a resale platform for digital commerce and retail businesses. Recurate is growing fastly and with aggressive timelines. Therefore, the executive team has engaged with IO Connect Services to design and implement the next generation of the organization's resale platform offerings.

The Challenge

Recurate's MVP was a self-managed semi-monolithic containerized solution that required constant server rebooting and manual scaling for the proper operation.

The multitenancy approach was rudimentary and delivered content to each client based on a column identifier in the database. This product was a single point of failure, as any application downtime affected all the clients, disrupting the service for everybody. Recurate needed to add more and larger clients to their portfolio. Therefore, a new scalable solution was required to warranty the company's future success.

Why AWS

AWS has been the cloud vendor of choice for Recurate since the beginning. Although the initial architecture leveraged other SaaS to execute some of the use cases, Recurate knew there was more potential in using AWS for these and future use cases. Not only for the infrastructure and hosting point of view but as a cloud platform for application integration.

Why Amazon DynamoDB

Amazon DynamoDB is a fast, flexible NoSQL, and fully managed database service for single-digit millisecond performance. With the support of key-value and document data structures, DynamoDB provides the flexibility to store enriched data models for customization, which is very useful in SaaS solutions. Moreover, DynamoDB offers high durability and availability across multiple data centers natively.

About Costumer

 recurate

Recurate is a tech-enabled resale service that empowers brands and retailers to establish their integrated resale platforms directly on their e-commerce sites. More customers. More sales. More sustainable.

Recurate focuses on designing a superior resale experience for customers that welcomes them into a brand's ecosystem. Their technology integrates with all e-commerce backends and works with each brand to get the look, feel, and functionality just right.

Amazon DynamoDB

“Amazon DynamoDB is a NoSQL database that supports key-value and document data models. Developers can use DynamoDB to build modern, serverless applications that can start small and scale globally to support petabytes of data and tens of millions of read and write requests per second. DynamoDB is designed to run high-performance, internet-scale applications that would overburden traditional relational databases.”

The Solution

At IO Connect Services, we have a cloud-native manifest for startups focusing on business value and product KPIs to build technical solutions. We selected AWS Serverless to shorten development cycles and go to market timelines.

The new platform needed to be a real SaaS B2B (Business to Business) solution. Therefore, we designed the new architecture with the following requirements around the transactional database: 1) Each tenant is isolated from the others. The data and other assets are unique and only accessible to a particular tenant. 2) A new client should be onboarded quickly, regardless of the geographical region, with a target of hours to minutes in future phases. 3) A particular tenant can have custom business logic and composing integrations. 4) Each tenant may enrich the data models with custom data fields and support foreign data structures from the integrating platforms. 5) SLAs and performance metrics must be standard for all tenants; a higher utilization client must not impact other accounts. 6) All transactions must be auditable.

We managed the new solution as a green field scenario to start fresh. The team decided to implement a multi-account mechanism to address the multitenant-related requirements. The solution's architecture was designed with a cloud-native approach using AWS Serverless technologies as much as possible. The objective was to rely on these services' out-of-the-box features and configurations to leverage the AWS investment. The design had an API-led first approach with a Microservices architecture implemented with Amazon API Gateway and AWS Lambda.

We selected DynamoDB as the No-SQL database to keep writing and reading times constant and to keep the schema flexibility of the data stored, as some premium clients could use different data models. Eventual consistency was a crucial factor in designing asynchronous transactions that could help scale the volume substantially.

The partition and sorting keys were selected based on analyzed data access patterns. We designed the tables under the analysis on how users will interact with the data, like retrieving a collection of listings ordered historically by a given seller or the constant read of a given listing identifier. GSIs (Global Secondary Indexes) were also added to fulfill other types of retrievals based on different elements.

We leveraged native DynamoDB Streams to construct an auditable history for all the records inserted in the database. Other applications can obtain historical data without affecting the user transactions' performance. Also, Kinesis Data Streams consume the events from the DynamoDB Tables to share data events with the rest of the business flows.

Amazon DynamoDB features:

- Performance at scale
- Serverless
- Enterprise ready

The solution provides dedicated tables to each tenant to isolate data and to ensure that no other tenant can access others' tables via access control lists at the AWS Account, Amazon DynamoDB Service, and Tables. Also, this approach helped guarantee consistent SLAs for each AWS Account since tenants do not compete for service limit quotas, giving control in throughput and having extra capacity of the service limits for unexpected bursts of read or write transactions with the on-demand capacity feature.

One critical decision in selecting Amazon DynamoDB was that the service is well-supported and available in all AWS Regions. Thus, this global availability allows deploying the solution in one or more regions worldwide. With an IaC (Infrastructure as Code) approach, it can get deployed in minutes.

Results and Benefits

Because all workloads were designed following a serverless approach, the generated code focuses solely on business needs rather than infrastructure or other non-business-related functions. Using serverless shortened development cycles, the project delivery was on time despite the aggressive deadlines and some last-minute changes in requirements.

We ran a load test to provide consistent numbers in throughput, and the results met expectations based on the architectural and design patterns. Each tenant account can handle gracefully more than one thousand transactions per second (1K TPS). Each request besides JSON payload included a couple of images sizing 2 Megabytes per request. These numbers were reached without increasing service limits in any of the services used, and the tables in DynamoDB performed flawlessly. Therefore, the platform can easily handle hyper-scaling volumes.

As the company continues onboarding new clients, the new customizations - especially those from the external systems and applications - have been included in the existing business logic and tables without disruptions. This has been one approach to extending the platform safely and fast.

Recurate is confident to onboard new clients rapidly. The leadership team is optimistic that the new solution is bringing them to the next level as an organization. They are ready to fast-track new funding rounds with better and proven results.

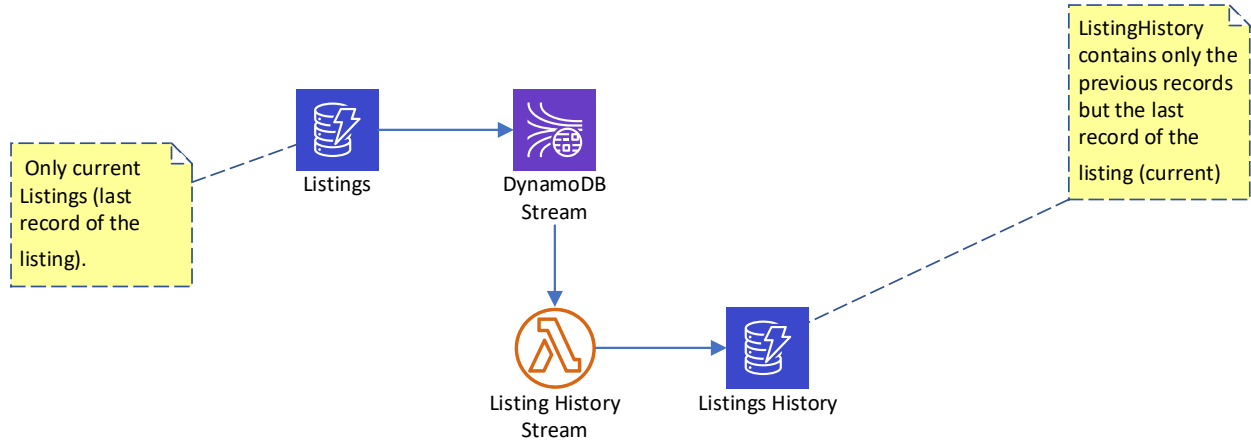


Figure 1. Using DynamoDB Streams to store historical data order by timestamp

Listings Table

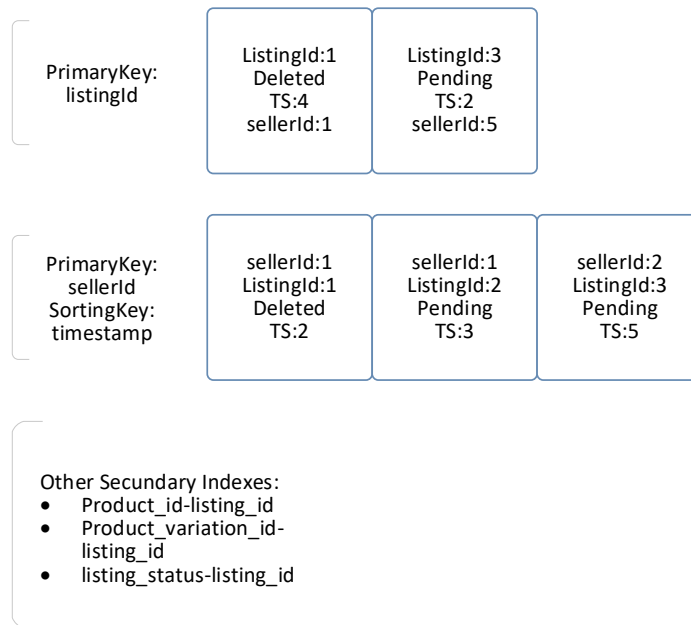


Figure 2. Primary table partition key design and GSIs

Listings History Table




Figure 3. Design of a historical table

Next Steps

There are new requirements to focus on currently and others evolving. The organization has provisions in the roadmap, like providing multi-region active-active functionality for global brands. Amazon DynamoDB features like Global Tables will be crucial to address this and additional future global needs.

Amazon DynamoDB will facilitate passing data records to downstream components for more sophisticated BI (Business Intelligence), data analytics, and potential ML (Machine Learning) use cases.

Superior Performance

DynamoDB scales to more than ten trillion daily requests, with over ten million read and write requests per second and petabytes of data storage.

Low TCO

Save money by replacing physical hardware with expensive license fees with AWS services and only pay for what you use.

Fully Managed

With DynamoDB, there are no servers to provision, patch, or manage, and no software to install, maintain, or operate. DynamoDB automatically scales tables to adjust for capacity and maintains performance with zero administration.

About IO Connect Services

IO Connect Services is a company specializing in Information Technology Consultancy Services. All our team members have one thing in common: our enthusiasm for technology and our passion for customer service excellence. We provide services in all North America, LATAM and Europe. Our headquarters are in NYC metropolitan area, and we also have offices in Guadalajara, Mexico and Madrid, Spain.

