



Third Iron Chooses Crunchy Bridge for Analytics Over Amazon Athena

Third Iron is a library technology company, delivering modern digital access solutions that streamline the user experience through AI models based off of large data sets. Third Iron's services are designed to save researchers time, reduce the load on help desks, and reduce unnecessary ILL requests. Their products are used by researchers from over 1,600 libraries in more than 35 countries, including corporations, hospitals, universities, and governments.

Third Iron recently migrated their Postgres workloads to Crunchy Bridge. Third Iron was looking for a fully managed host with a better price to performance ratio than their previous host, and they needed a no-downtime migration. Crunchy Bridge was able to help them with a cutover, replicating data ahead of time, and working together 1-on-1 during the final migration. Since migrating they have enjoyed working with the excellent support team, both during their trial and since their migration.

“We were looking for a better Postgres experience. The most important drivers for us were better scalability, improved functionality including Logical Replication, and better performance. Crunchy Bridge was consistently 2X to 4X faster than our previous Postgres hosting. Crunchy was a smashing success in every way. We are paying less money but getting more performance and more features.”

Karl Becker, Third Iron

Third Iron's data strategy heavily relies on PostgreSQL. In addition to their relational database application needs, they were searching for an analytics solution for a separate initiative involving receiving large amounts of data.

Looking for Fast, Easy Analytics on JSON

Third Iron wanted to do some investigation and data analysis on one of their large data sets. They knew they'd be bouncing back and forth between the developers writing SQL and the teams asking questions and performing further analysis. They also knew that they might not need their analytics server all the time; there might be days or weeks in between use.

Third Iron

“Before Crunchy Bridge for Analytics, every time we asked the dataset a question, it would take hours to get an answer. After we imported our data, the same query took minutes.”

Karl Becker

Co-Founder
Third Iron



crunchydata

crunchydata.com



Third Iron Chooses Crunchy Bridge for Analytics Over Amazon Athena



Third Iron set out to find an analytics solution that met these basic needs:

- Easily store JSON more than 1TB
- Store files somewhere cost effective, like S3, rather than the production database
- Query data directly with straightforward SQL
- Scale analytics system down when not in use

Crunchy Bridge for Analytics vs Amazon Athena

Crunchy Bridge for Analytics connects to your S3 data in a similar way as Amazon Athena, but Third Iron found that Athena had a steeper learning curve given that the query language is slightly different. Third Iron's data team isn't huge, and they wanted to get up and running pretty quickly. Since their data analytics stack is more for answering business questions, they didn't want to invest a large amount of engineering time learning a new tool. "It seemed like it would take a fair bit of time for our developers to learn and get comfortable with Athena, plus an ongoing cognitive load of keeping familiar with that syntax," said Karl.

"I started trying out Amazon Athena but it was a pain. The querying was just different enough from what we're used to to be annoying. Whereas I was able to get things working using Crunchy Bridge for Analytics unbelievably quickly. We have a lot of in-house experience and knowledge with Postgres, so the fact that it's still Postgres is extremely appealing."

John Muth, Third Iron

The team's familiarity with Postgres tools made Crunchy Bridge for Analytics a much easier solution. The data could still live in the remote object store, but the SQL and tools to work together as a team on iterating on the data analytics, could stay in their existing workflow, based in Postgres.

A Better Data Analytics Experience

After doing more research and choosing Crunchy Bridge for Analytics for their analytics project, Third Iron is really happy with their move. They've been able to take advantage of saved queries to share and iterate on SQL.

We looked into various data pipelines and had a plan in place, but when Crunchy Bridge for Analytics was released, it solved our problem without all the complexity we were expecting.

Karl Becker

Co-Founder
Third Iron



crunchydata
crunchydata.com



Third Iron Chooses Crunchy Bridge for Analytics Over Amazon Athena

They're using the saved queries features to organize queries into folders and are impressed with the UI over the other Postgres providers. Third Iron is also using the suspend and resume instances feature to manage costs and suspend their cluster when it's not in use.

"Today we're able to easily analyze 1000s of raw files without any ETL complexity. We get all the benefits we're used to with Postgres on Crunchy Bridge, but without writing any custom JSON parsing code or specifying a schema for the big chunks of JSON. Instead of causing Third Iron to create engineering solutions to new data analysis problems, Crunchy Data solutions solved that complexity, and now we can focus on solving problems for our customers."

Karl Becker, Third Iron

Third Iron is also really impressed with the performance out of the box. When compared with tools they created in house, performance for their analytics queries is measured in minutes rather than hours.

We had previously tried to write our own parser to analyze our data. We were parsing big XML files which are larger and more clumsy to work with than JSON. Before Crunchy Bridge for Analytics, every time we asked the dataset a question, it would take hours to get an answer. After we imported our data, the same query took minutes.

Karl Becker, Third Iron

Overall Third Iron has found Crunchy Bridge for Analytics to be:

- Easy to implement, especially when compared to similar tools like Amazon Athena
- Integrated with their Postgres database tools
- Cost effective and easy to control costs
- Performant, running queries in minutes vs hours
- A simple solution to data analytics backed by Postgres

"We looked into various data pipelines and had a plan in place, but when Crunchy Bridge for Analytics was released, it solved our problem without all the complexity we were expecting."

Karl Becker, Third Iron

Third Iron

"Today we're able to easily analyze 1000s of raw files without any ETL complexity"

Karl Becker

Co-Founder
Third Iron



crunchydata
crunchydata.com