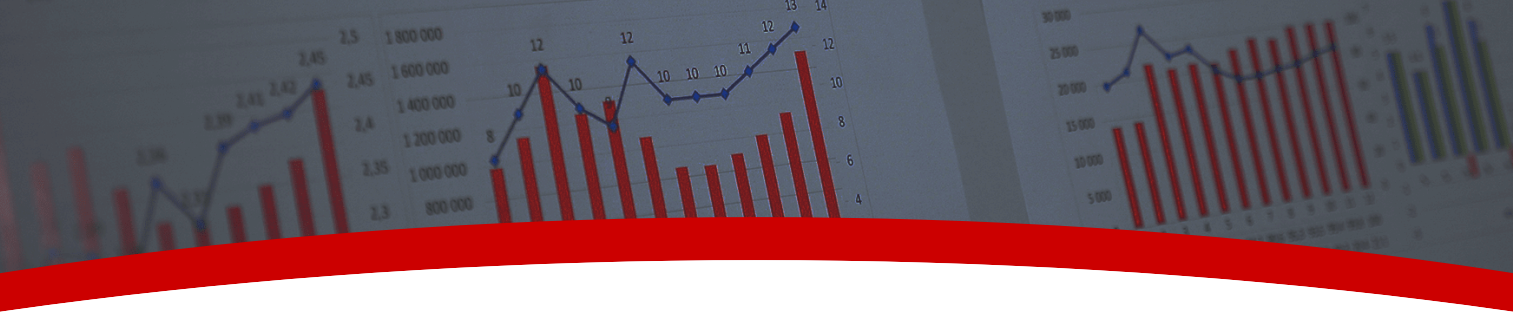

CASE STUDY

How Redgate Flyway Enterprise can automate database deployments with Oracle

redgate

Contents

The Customer

One of the largest, most stable and progressive credit unions in the USA. This member-owned, not-for-profit, cooperative financial institution has 120+ locations across five states and over 1.2 million members.

The Challenge

Creating database migration scripts to release changes was a problem and, without a proper continuous delivery pipeline in place, releases were taking a week, often longer, hindering the speed at which the team worked.

The Solution

Redgate Flyway Enterprise was introduced as it could plug into existing application development pipelines, deploy database changes using a migrations or state approach, and automate CI workflows.

The Results

The Data Science Team now release database changes multiple times a day, depending on cadence and what people are working on. The only limit is how many developers they have available, so it's basically on demand.

"We were tired of waiting for database deployments, that was the main driver. We'd have to schedule out a week and it was just terrible."

The Customer

One of the largest, most stable and most progressive credit unions in the US, they remain a member-owned, not-for-profit, cooperative financial institution. With over 120 locations across five states, assets in excess of $16 billion, and more than 1.2 million members.

As well as growing in size, it has also grown in its use of technology to deliver a range of advanced financial services to its members in the easiest, smartest and most secure ways possible. Close to a hundred developers, DBAs, Analysts, Machine Learning Engineers and Data Scientists work behind the scenes to develop and maintain the many applications and databases in use.

The Data Science Team is part of the wider analytics division, and is run by the Data Science Manager. He outlines that over the last few years, they have been expanding the use and value of business intelligence and data science across the business. His team of full stack Data Scientists and Machine Learning Engineers work on everything from lending automation to check clearing, and are responsible for both the application layer, AI/ML models, and database development.

1,500

EMPLOYEES

20

DATA ENGINEERS

1.2 MILLION

MEMBERS

“We saw a massive amount of efficiency, and we were deploying faster, more often, with way less issues. It's pretty much your archetypal cool Agile success story.”

The Challenge

When the Data Science Manager first joined the credit union, they were using the waterfall model for software and database development, breaking down projects into linear, sequential steps, with each step dependent on the previous one being completed before it could go ahead.

The growing need to deliver features and insights faster prompted the team to move to a DevOps approach. This involves adopting Agile development practices and working in short, two week sprints to release smaller features and updates more often. It also means automating as many laborious and repeatable parts of the development process as possible, using third party tools like Jira, Bitbucket and Bamboo.

While this enabled the team to release AI/ML model changes in a quicker and easier way, database changes were a different story. The team use Oracle as their main database system and all of the custom SQL code is on Oracle Exadata. There are many databases and instances in use, with development, testing, UAT and production environments required for each.

Creating the database migration scripts to release changes was a problem because of the way Oracle handles permissions. Without a proper continuous delivery pipeline in place, there were also differences between the databases and development environments which caused further issues in the release process. Each of those releases took a week, often longer, so database deployments were only made around once a month.

This was hindering the speed at which the Data Science Team could work, so they looked for a solution that would allow them to include the database in their DevOps and Agile practices, and ease the process with automation. As the Data Science Manager outlines, "*That's how my team works: we find something that we really hate and figure out how to automate it and make it better.*"

“A nice byproduct of this is that it’s so auditable. It’s really nice to say, here's the model we deployed in production, and here are all the downstream changes.”

The Solution

There were three major requirements on the checklist for the Data Science Manager and his team that any solution would need to satisfy.

Firstly, they wanted it to integrate with the new tech stack that was working so well with application development. They didn’t want to compromise the advantages they had already gained.

Secondly, it had to be capable of deploying database changes in both a migrations-based and a state-based way, running a series of migration scripts for smaller changes, or auto-generating a single script for major changes.

And thirdly, it had to be relatively easy to implement and use for both the Machine Learning Engineers and the Business Intelligence Analysts on the team.

They looked at many solutions available and quickly narrowed down to Redgate Flyway Enterprise which could plug into the existing pipelines, deploy changes using a migrations or state approach, and simplify the complexity of continuous integration workflows by introducing automation.

The strategy of the Data Science Manager with any new tool is to implement it with the least bells and whistles and then build on top of it. So while it took four months to fully embed Redgate Flyway Enterprise into the infrastructure, team members were deploying database changes within two weeks. He outlined that getting everybody on board was straightforward: *“Because Flyway hides the complexity away, I could say just create your stuff and do your code review and we’ll handle the rest.”*

As a result, the nightmare of database deployments has now disappeared, and the team like to say: *“Our shortest window is four hours, but we can get a hotfix deployed in 15 to 30 minutes.”*

The Results

The Data Science Team now release database changes multiple times a day, depending on cadence and what people are working on. The only limit in terms of how many deployments they can do is how many developers they have available.

In the future Redgate Flyway Enterprise has the potential to be utilized with other teams easily. The only tool they need to learn to include the database in continuous integration and continuous delivery is version control. Once database changes are in version control, the tools and workflow in place take care of testing them and moving them through the pipeline automatically.

This has brought a big advantage to the team. As the Data Science Manager outlines *“One of the key things is developer confidence and enabling them to take risks because the pipeline is going catch it. With all of the tests in place, they can make big changes and they don’t have to worry about remembering all of the database migrations. It just all flows through, so developers deliver faster and they’re more confident.”*

Finally, and importantly for this US based financial institution, Redgate Flyway Enterprise makes compliance with auditing and regulatory requirements far easier. The IT team can demonstrate to auditors and other stakeholders what changes were deployed to the production databases, when they were made, and who made the change for what reason. The Data Science Manager summarizes this, saying *“I hate this term, because nothing’s ever self-documenting, but Flyway really is the best form of self-documenting system you can get for that type of stuff.”*

The Data Science Team are now continuing their quest to identify problems in the software delivery process and fix them. They’ve found that if they aim to automate just one small task a week instead of trying to automate everything, they get a lot more done.