Database DevOps Best Practices: Build Once, Deploy Many

Implementing Safe Database Release Automation

DBmaestro
DevOps for Database

Build Once, Deploy Many: Repeatability

Modern DevOps processes and best practices suggest that you follow a very basic rule:

Build your code once, and deploy it many times to any relevant environment.

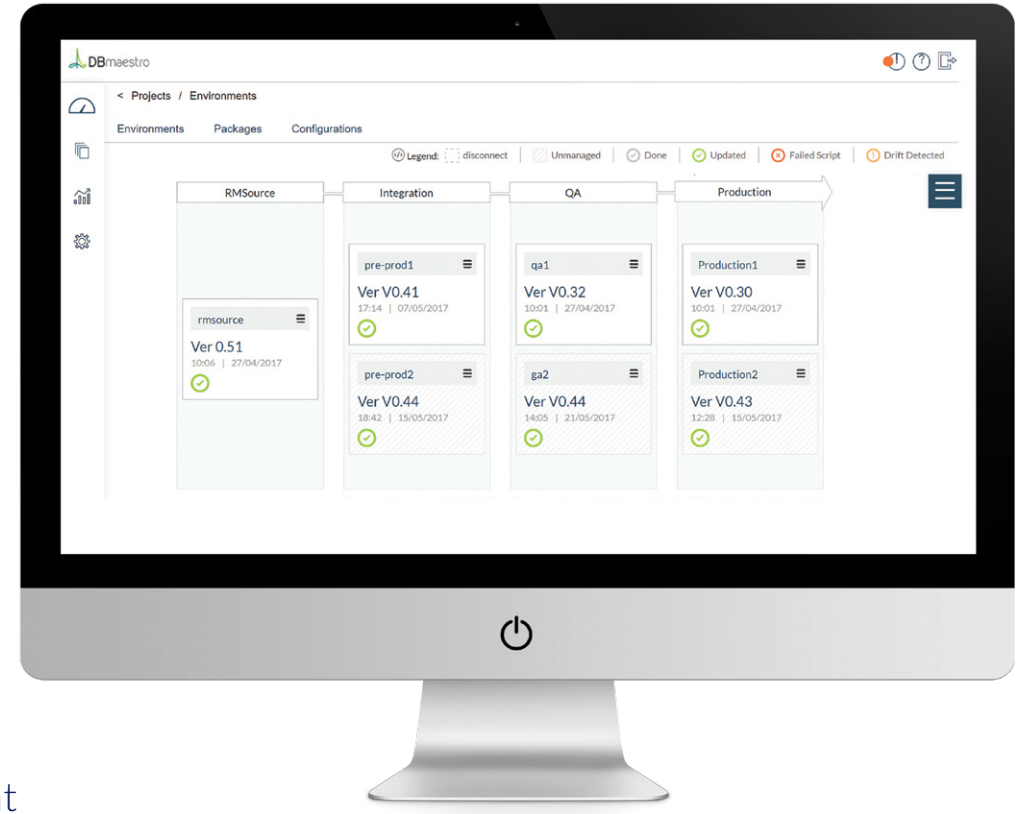

The logic is very clear:

- **Build once:** Compiling code should only happen once. This eliminates the risk of introducing differences due to varying environments, third party libraries, different compilation contexts, and configuration differences.
- **Deploy Many:** Once the binaries are created, they should be applicable to any environment. The same automated release mechanism should be deployed for each environment, making sure the deployment process itself is not a source of potential issues. Deployment is frequent to lower environments (Integration, QA, etc.), less frequent to higher environments and—ideally—once to Production. There is considerable risk when the PROD deployment is different than the other environments. You can't afford the only untested deployment to be to PROD.

The Database Challenge

The principles above are tried and true. The challenge now is to fit database release and deployment into that model. Releasing application code is far simpler than releasing database changes; you can override your changes with newer or older releases, you can override configuration gaps, and you can reinstall your application from scratch.

Databases, however, are different, they have both configuration and data; the data is persistent and accumulating. In most cases, you can't override QA with DEV data, and you shouldn't override PROD data with anything else. As a result, the way to introduce changes is to alter the state of the database from its older structure and code, to its desired new structure—in sync with your application code.

This requires a transition code that alters the database's structure—the SQL script. Again, the app code is just the app code. The database code is the result of running the delta scripts.

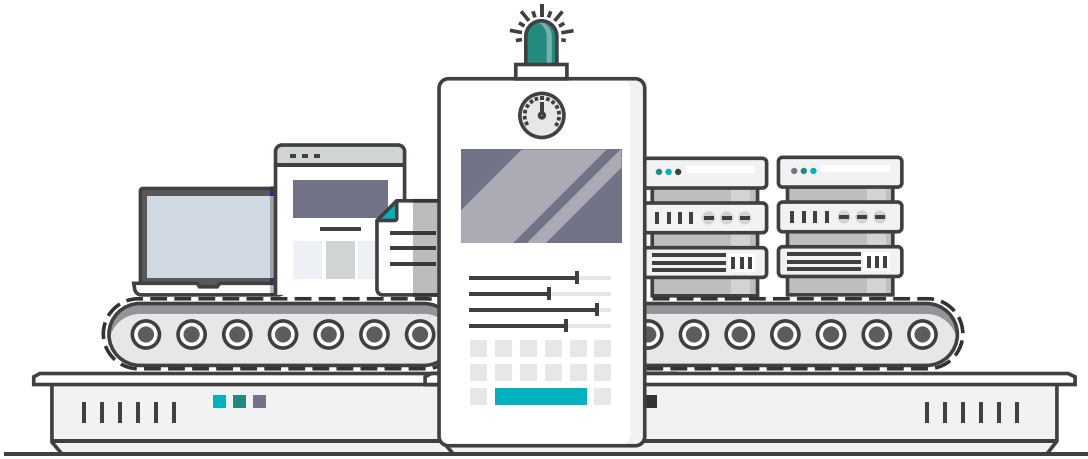

In theory, upgrading your database with SQL scripts just works; it's true, the sum of every change to the database represented as a script is the final configuration. In practice, however, there's usually something missing: a database configuration tends to drift. Someone performs a maintenance task, a performance optimization is implemented, or a different team's work overlaps with your own. As a result, that script that worked in one environment might not work in the next, or worse—it ends up creating damage and downtime. Configuration drift can present itself as different schema configuration, different code, good code that was introduced by other teams, or plain production hotfixes, that might be blown away by the SQL script you are introducing. Therefore, working with just the change scripts is effectively a faith-based methodology. All the stars must be aligned 100% of the time for it to work.

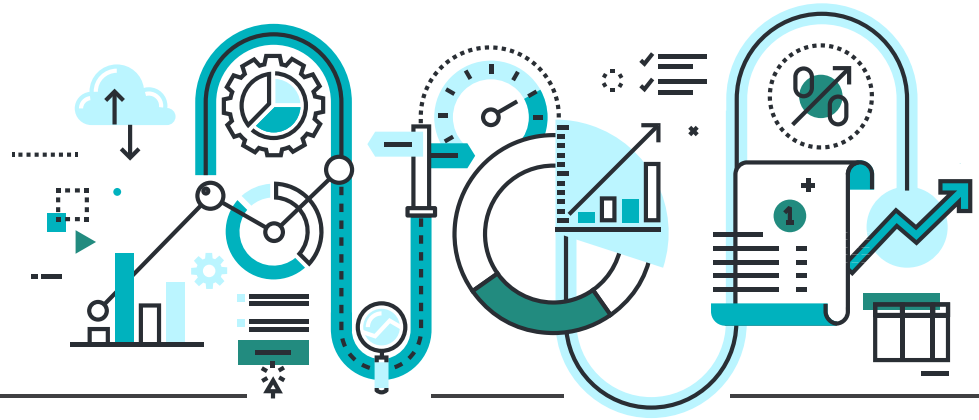

Clearing the Confusion: State-Driven and Migration-Driven Deployments

To create the transition or upgrade code, we can employ the concepts of state-driven or migration-driven deployments:

State-driven database delivery (generally referred to as a compare and sync tool). The idea is simple and very appealing: all we need to do is to use a compare tool to auto-generate the scripts required to upgrade any existing database to the next environment. That tool will always push changes from lower environments upward, such as from DEV, to QA, to PRE-PROD and finally PROD.

A model-based delivery method is another variant for state-driven delivery. Defining a model—with either UI diagrams (a designer) or XML representation (translator)—enables you to define the desired database structure, and then compare and sync to the target environment.

Pros:

The compare or model tools do the heavy lifting: offering a script to make the target database look like your source database or your model.

Cons:

The change process is not repeatable: the script generated to one environment might be different from the one generated for the next environment, as the environments might be different (drifted). This means that the process is not repeatable and does not follow a 'build once' model.

The tool will push changes forward (from source to target) unless manually reviewed and explicitly directed to do otherwise. This can present high risk in an automated process. In this scenario, changes originating from the target environment will be overridden with older code or conflicting code from the source (e.g. dropping an index that was added to PROD to deal with performance). This is a great example of a mistake that would be very easy to make and very costly to fix. The script is generic and designed for a broad audience. You must review it, adjust it, and change or fine-tune it to fit your specific case with every iteration and environment (useless you fall under the generic case).

Bottom line:
This approach creates a situation where you 'build many' for each environment and eliminate repeatability.

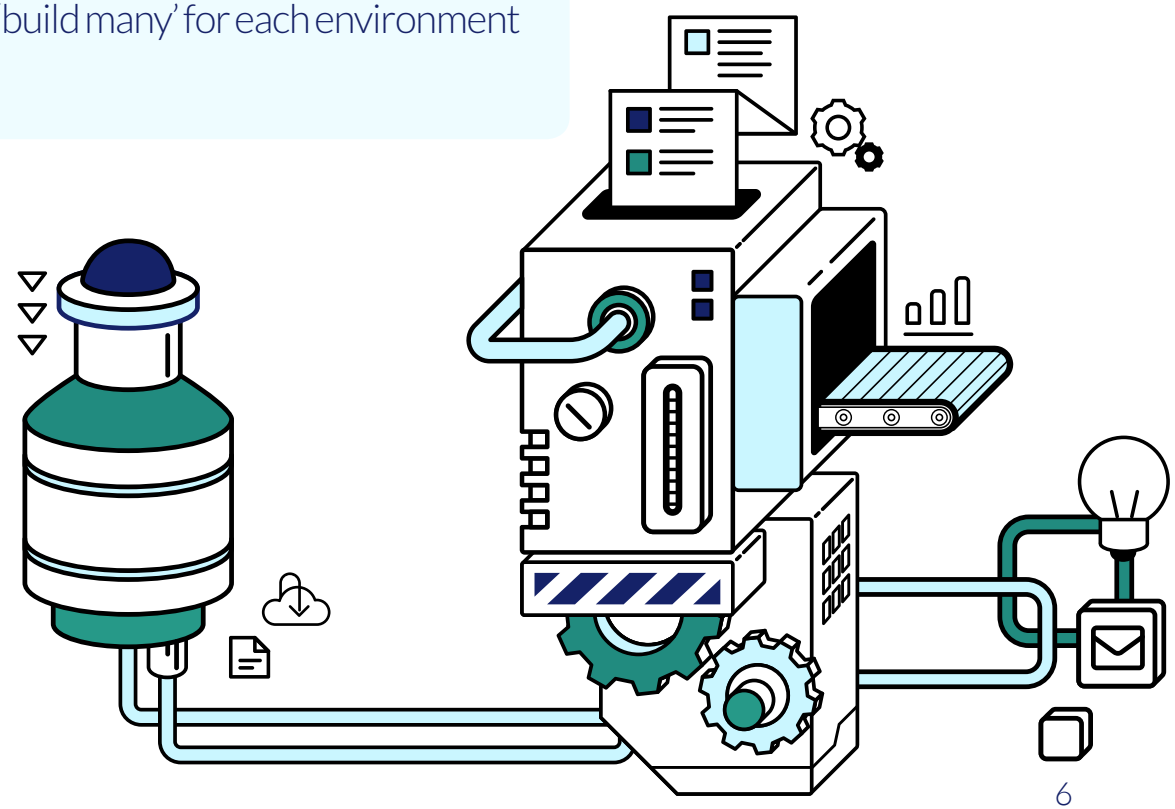

Migration-driven database delivery – migration steps are created to transition a database from one version to the next, mostly implemented as plain SQL scripts. Upgrade scripts can be executed on the database.

Pros:

Perfect repeatability. The same code is executed in all environments. A classic 'build once' approach.

Detailed and fine-tuned to your needs. The script is honed to fit your coding style, specific needs, performance variations and your application requirements.

Cons:

1. In most cases, you'll have to build the migration code manually.
2. In the face of drifted database environments, the SQL script might produce results that vary from undesirable to catastrophic—these can include anything from overriding someone else's changes (if the script is not properly synced with other development efforts), to overriding hotfixes targeted for production.

Bottom line:

Build once, deploy many - but with some risk at the deployment phase.

DBmaestro – Database Release Automation

DBmaestro Release Automation helps you automate and get full visibility and insights into your database release pipelines—at enterprise scale. It allows you to easily define and run continuous delivery pipelines with high security and compliance with organizational policies. The product supports seamless integration with all sources of database changes while predicting the success of database deployments and alerting for configuration drift or non-policy actions. DBmaestro Release Automation enables you to publish releases quickly, prevent accidental overrides of changes, and reduce application downtime caused by database-related errors, all while maintaining stability, scalability, and security.

Fast, safe, repeatable and scalable.

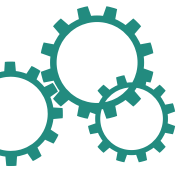

About DBmaestro

DBmaestro is a leading DevOps for database solution provider. Our flagship product, DBmaestro DevOps Suite, introduces DevOps and automation best practices to databases for the enterprise, dramatically simplifying, accelerating, and improving release processes, while modernizing database development via release pipelines, long enjoyed elsewhere in the industry.

We provide both database source control and database release automation capabilities across the board for developers, DBAs, security, and operations in multi-database enterprise environments.

For more information please visit us:

DBmaestro
DevOps for Database