Automated Deployment

What is automated deployment?

The traditional data warehouse upgrade process involves one or more people declaring a window in which part or parts of the infrastructure will be unavailable, logging in to the command line and making changes until the system can be declared upgraded. Sometimes it goes wrong. If fixes are made as part of the process, sometimes they aren't properly recorded and reflected in the source code used by development, which makes problems more likely the next time.

Deployments can be fully automated. There are techniques which allow a deployment to be done by a single button press, to an environment where processes are running. There are some challenges involved, especially in the context of business intelligence infrastructure such as a data warehouse, but there are proven solutions.

An extension of this is to have deployments done automatically following a pipeline of automated tests. We have used this on a large scale enterprise data warehouse to allow several safe and fully audited deployments to occur each day.

Why should I use it?

The Visible Ops Handbook suggests that without modern processes, you can expect that 80% of outages are caused by a change, and 80% of the MTTR is caused by figuring out what the change was. By using automation, everything has to be done via mechanisms that can be effectively logged, and so you vastly reduce that 80% of MTTR caused by not knowing what the change was.

It then becomes realistic to have mechanisms that can undo changes conveniently and reliably.

Fully automated deployments mean that the same upgrade process can be used in development and test as in production. This saves significant time when there are a lot of environments to set up, and then means that the upgrade process has effectively been done many times by the time it is used on production.

When automated environment creation and configuration is also being done to ensure that these development and test environments match production, then these tests are reliable, and it means upgrades are less likely to fail.

There is trivial cost to an automated deployment, meaning you can do many more of them. This means that production environments can safely be built more frequently, so you can go from idea to production in less time.

The benefits then increase exponentially when you have trivial effort to create and configure more environments - automated configuration and deployment is the key to making efficient use of the vast resources available in the cloud. Thus we were able to add further Redshift clusters to mirror our original one and keep them configured identically.

Can you help us work out how to use it?

We also developed the mechanisms which allow hot deployments to a data warehouse. We have some posts about this in the Cloud BI Limited blog which explain the high level technical concepts. We solved a number of challenges at the detail level which allowed genuine push button deployments to be done at any time, and integrated that with the automated testing mechanisms. We have an integrated BI optimised approach to automated environment creation and configuration, deployment and testing that we can tailor to solve your challenges and maximise the benefits you get from your opportunities.

The Cloud BI Limited Integrated Automation Approach

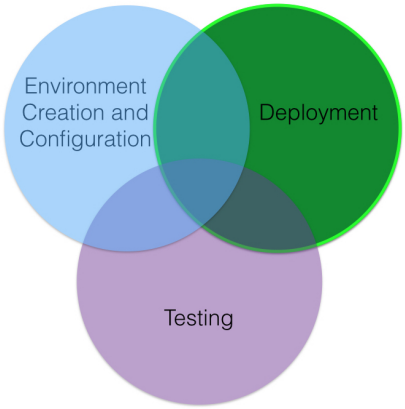

Can you help us implement and operate it?

Yes - we built all these things at a major communications provider for an Amazon Redshift data warehouse project. We can also advise and coach on the people and processes aspects.

What extra value does Cloud BI Limited add?

We have done these things before with great success. We want to share the benefits of a devops influenced approach to data warehousing. We're easy to work with and want to see you succeed.

To explore how we can help you further, please contact info@cloudbilimited.com