

Application Modernisation The Data Perspective

By Iqbal Goralwalla



Introduction

Without doubt one of the most active areas of IT change is that of modernisation to Cloud. Disruption is deemed positive for organisations as they look ahead to the digitalisation benefits that come with modernisation to Cloud.



If someone asks me what cloud computing is, I try not to get bogged down with definitions. I tell them that, simply put, cloud computing is a better way to run your business.

Marc Benioff
CEO Salesforce.com

It is unsurprising that with quotes like this many companies are looking to gain the suggested benefits of modernisation to Cloud.

Drivers can include the desire to reduce costs; move away from on-premise infrastructure to a more flexible Cloud based service; the ability to 'spin up' new projects/applications quickly; and more.

However, whatever the driver there are also risks and costs to be considered. Most applications today are running 24x7 with large volumes of data underpinning them. Modernisation of an application, whether from/to in-house platforms, to the Cloud, or extracting an application from an existing service to another company' system – needs careful planning. The move of the data is key to the success of the project. How to move large volumes of data efficiently? How is this done in a timely manner without impacting availability or losing transactions? De-risking the move of the data in a modernisation project is key.

There are 4 stages that need to be considered when modernising data:

- 1. Validate** – assess the characteristics of the applications and databases vs the chosen Cloud solution and identify any issues with the move, and actions needed to match the Cloud environment to the application requirements.
- 2. Plan the Migration** – detail the plan of the tasks and timing for a successful implementation, highlighting areas of risk or tasks on the critical path.
- 3. Build, Test and Cutover** – the migration plan should include a robust test plan to ensure the Cloud based solution works and the data migration strategy can meet the Client’s availability requirements. Then, execute the plan to deliver the application live on the Cloud.
- 4. Support** – depending on in-house skills and resources on-going support should be put in place for production monitoring, maintenance and support for the Cloud based system.

Let’s look at each of these in turn

Validation

This stage is understanding the current application and database system requirements to ensure the chosen Cloud solution is configured correctly and the migration plan includes all the elements required to deliver the application into production. You need to validate whether the application and database is suitable for the Cloud and verify whether everything has been considered to make the move successful - hardware and performance requirements, users, licensing, compliance needs and application dependencies.

The following should be analysed and documented:

- **Workload** – Assess the characteristics of the database application workloads – transaction volumes; batch processes; timings; peak volumes; etc. This is critical to understand how the Cloud solution will need to operate. For example, batch processes may be dependent on significant data transfer which will have network capacity and performance implications if the Cloud service is distant.

This exercise will also help determine which applications cannot be transferred to the Cloud. For example, legacy applications which have dependency on several other on-premise applications.

- **Response times and SLA's** – You want to make sure your performance and availability is not compromised with the Cloud move. Applications won't always run faster in a Cloud environment, and can sometimes be significantly slower. Unexpected bottlenecks can occur when changing to a distributed environment especially with the additional communication overhead. Wide area network conditions, such as bandwidth constraints, can significantly slow down an application.
- **Backup/recovery requirements** – The disaster recovery requirements and features in the chosen Cloud service needs to be understood to ensure a successful migration. Your current backup and retention policies such as frequency, backup type, recovery time objectives (RTO) and recovery point objective (RPO), data restoration time, all need to fit with those provided by the Cloud service. Managed Cloud Service Providers may have standard backup policies with some level of customization. It is worth having a look at those to see if they are suitable for your needs before they become a potential roadblock.
- **Security** – Many of your business processes might not apply to Cloud applications. You'll have less control over physical security and other procedures as in many cases you will not own any of the hardware. This means dealing with threats and breaches will be very different than before. To keep your data secured, you'll need to re-evaluate your security procedures and strategies.
- **Application linkages** – Dependencies between applications needs to be addressed. It is important to make sure applications with dependencies can be moved together to the Cloud environment. For example, the application might have an Active Directory dependency or require access to an on-premises database.
- **Data transfer requirements** – The data routes, feeds, and APIs to production and other environments need to be carefully mapped out.

-
- **Infrastructure design** – Cloud server migration involves computing environments that run differently than your current system. Servers, networks and data services all run and interact differently, and you need to make sure your Cloud infrastructure design is ready for this change. Additionally, one has to be wary of hidden costs. While a Cloud service provider can usually scale its offerings to meet even the most demanding workloads, this scalability comes at a price.
 - **ROI** – Assessment of the costs of any recommended changes and the impact this has on the cost case for the move to Cloud.
-

Migration plan

A fully documented migration plan is the most important stage of process of modernisation to Cloud. If done well, the actual migration steps will proceed smoothly.

The plan should include the following considerations:

- **Build of the Cloud database environment** – This will involve building the database environment on the Cloud to mimic the on-premise environment. Pre-requisites such as machine sizes; network capacity; storage size; etc. and how these will be built should be included in the plan.
- **Migrating your data** – This is the heart of the plan and needs careful preparation. Moving a database can be tricky, particularly if the database is large and the business requirement is for little or zero down time during the move. How do you get the data to the Cloud with all current transactions/ updates applied? Use the right tools for the job, and make sure the new structure and limitations of your database will work correctly. Whatever tools your provider has available, make sure you really know how they work and how long the migration will run for. Depending on the size of your databases and applications, you will use different techniques for actually copying everything over. If you don't have too much to migrate, you can just copy the data over your internet connection. This approach isn't ideal for larger databases. You might have very long transfer times or charges from the Cloud provider. To deal with this, you could compress the data before sending it. Alternatively, you could ship your physical drives to the provider to reduce bandwidth costs.

-
- **Applications Move** – Plan to move the applications to the Cloud first that are the easiest and least critical. This way you can get a feel for the process, and deal with unforeseeable problems while the stakes are still low. Choose the data and applications that are most suitable for a Cloud environment, preferably standalone ones, and migrate those first.
 - **Downtime** – What downtime for the application/service is acceptable to your business/users while you migrate? This will inform the overall strategy for the plan, particularly the data migration plan.
 - **Security of the move** – If security is important while on the Cloud, it should be just as important when sending all that information and code to the Cloud. Any temporary storage locations used along the way for sensitive data should be just as secure as your end destination.
 - **Application portability** – You'll want to make sure once you move your application, it doesn't get 'locked in' to one Cloud provider. The main obstacle to this is with metadata. Tools like virtual machines seem like they could move easily between Cloud environments. However, metadata is usually specific to a platform. Understand how you can keep track of this and keep your application portable.
 - **Test** – The plan should indicate how and at what stages you should test that the move is working and has delivered successfully. If needed how to back out the latest step and return to a working system.
 - **Cutover** – Assuming the above plan has been executed successfully, all tests completed – how does the cutover to production take place and the close down of the in-house applications. Output from the stage will be a fully documented plan for the move with tasks and timescales, which is capable of execution. Risks and activities on the critical path will be identified with suggested actions to minimise impact.

Build, Test and Cutover

Time now to execute the plan – building the database environment on the Cloud platform; testing the move at various stages depending on the complexity of the move, and then the final cutover to production.

- **Build** – Things to think about as the build happens:
 - Make sure all OS, software, and storage prerequisites are met.
 - Set up users
 - Install database software
 - Update system and database level configurations and registry parameters to match those on-premise
 - Create the database structure and objects
 - Move the data whilst taking care of security implications whilst transferring data. The most appropriate method should be used. Some examples include:
 - Batch data transfers
 - Load data from a local file stored on premise
 - Load data from a Cloud storage object like SoftLayer swift object store or Amazon S3
 - Use native data movement utilities, remotely
 - Use any Cloud specific data transfer or mass migration service
 - Backup & Restore
 - Set up and configure high availability environments
 - Move the applications starting with ones that are standalone and have no dependencies

- **Test** – The next stage is to test the workload(s) and will involve testing the system in a way that is representative of the final production environment. This will be the proof of concept for validating the performance of the workload and the costs for running on the Cloud. This phase will help get a feel for the resources and capacity that the applications require – such as

storage types and size, the number of VMs needed, and network bandwidth. It will also ensure that you know which network and security controls are needed to replace your legacy firewall systems.

Various aspects need to be looked at in this stage. Did the tests work? Did all of the data make it? Is it all reachable to users? Are all of the internal components communicating properly? Are the timings acceptable to the business? Can all of your admin tools monitor the new cloud app properly? And many more.

- **Cutover** – Migrating data and associated applications during the cutover phase is one of the trickiest parts of a Cloud migration. The location of your data can significantly impact the performance of your application. Moving your data to the Cloud when the applications are still primarily on-premises can significantly impact performance. The same holds true if the data is still on-premise but the application resides in the Cloud. So, when and how do you switch over the production database system from the legacy on-premise solution to the new Cloud version? There are two common approaches:
 - **Big bang.** Move all the data and applications over to the Cloud and test that it works there, and then switch traffic from on-premise database to the Cloud database.
 - **Phased.** Move the data and then move a few applications over first, test that things are still working, and then move a few more applications. Continue this process all applications have been moved to the Cloud.

You need to determine the approach most suited for your environment. Two key factors in that go hand in hand with the above approaches are the amount of downtime and data loss that are acceptable for your applications which are critical to your business.

Options for data migration include:

- Using a bi-directional syncing mechanism between your on-premise and cloud databases. Once all applications of the data have been moved to the Cloud, the on-premise database can be removed.

-
- Use a one-way synchronization between the on-premise database to the Cloud-based database, and allow applications to connect only to the on-premise version. At cutover, disable access to the on-premise version so the Cloud-based version becomes the main database, and enable Cloud-based applications access to the Cloud database.
-

Support

How are the Cloud based applications to be supported? The extent of the on-going support needed will depend on the environment and services chosen from the Cloud provider. In some cases, support requirements will vary little from those currently if the choice is simply provision of a virtual system. Choosing a standard, managed service will reduce the system maintenance, but the application will still need on-going care. Depending on in-house skills and resources on-going support should be put in place for production monitoring, maintenance and support for the Cloud based system.

We have outlined in this article the 4 stages for a successful move to the Cloud from a data perspective. Planning ahead is the key. If you require support with your Application Modernisation project [contact Triton's team of experts](#) or visit our website for details on our [Modernisation service](#). Triton provide consultancy and services to help organisations plan, implement and reduce operational risks associated with moving data to the Cloud.

About Triton

Triton Consulting are experts in Hybrid Data Management and Digital Transformation. The company's team of consultants represent some of the most highly experienced and qualified in the industry, and are able to advise on a range of Data Management solutions including DB2 for z/OZ, DB2 for LUW plus data related infrastructure and transformation services.

As well as expert consultancy in all areas of DB2, Triton Consulting also cover a wider spectrum of high-level consultancy including senior project management, technical planning, technical architecture, performance tuning and systems programming.

Triton Consulting has been providing consultancy services for over 24 years. Triton are internationally recognised for their DB2 expertise with three IBM Gold Consultants and four IBM Champions.

Find out more about Triton Consulting: www.triton.co.uk

About the Author

An IBM Gold Consultant and an IBM Champion for Data Management, Iqbal Goralwalla has over 20 years of experience in all aspects of relational databases on the LUW platform.

Iqbal is a highly experienced Data Management specialist, having been involved in all aspects of data including data architecture, strategy, design, modelling, development and administration. Iqbal has extensive project support and technical liaison experience within DB2 on LUW environments.

As well as leading the DB2 Midrange Consulting team and running the Managed Services and Education team, Iqbal is also an active Principal Consultant and is engaged with customers for short, focused assignments to help them with all DB2 issues.

Contacts

Rob Gould

Business Development Lead
+44 (0) 870 241 1550
+44 (0) 7766 838 904
rob.gould@triton.co.uk

Iqbal Goralwalla

Head of Distributed Data Management
+44 (0)870 241 1550
iqbal.goralwalla@triton.co.uk