AVOIDING THE BIG-BANG:

A GUIDE TO ITERATIVE ATLASSIAN MIGRATION.

Clearvision

CONTENTS

ABSTRACT

Leading Atlassian applications such as Jira and Confluence have evolved over the years, but the method of moving data from one application instance to another has been limited by the full-backup and restore mechanisms within these products. This has led to high-risk Big-Bang data migrations for example Server to Cloud and vice versa, becoming the norm. The good news is, there are lower-risk and more agile alternatives available that can deliver greater business value and benefits.

DO WE NEED TO MOVE ALL OF OUR DATA?

According to Carnegie Mellon University's Software Engineering Institute, there are some points to consider when moving data:

o Does it make sense to migrate the legacy data?
o What legacy system data can be reused?
o What changes to data are required to accomplish the migration?

o What migration strategies are most appropriate?
o What are the preliminary estimates of cost and risk?
o What is an ideal pilot project that can help address certain risks?

It's important to consider these points to ensure Atlassian data migrations result in the best business value and at the lowest possible cost and risk.

DOES IT MAKE SENSE TO MIGRATE ALL OF THE LEGACY DATA? WHAT WILL WE REUSE? WHAT CHANGES TO THE DATA ARE NEEDED?

The quality of Atlassian application data, significantly influences the risk and effort associated with migrating to another platform. Essentially, the further we go back in history, the bigger mess we discover. This is compounded by the complexities of various Atlassian Marketplace Apps' data or bespoke hidden scripting. To mitigate this risk, it's vital to decide on how much history to transfer into the new system.

Without a thorough understanding of the source and target system, transferring data into a newer platform nearly always amplifies the negative impact of inaccurate, inconsistent and irrelevant data, i.e. unused workflows, custom fields in Jira, unread/unwanted pages in Confluence etc.

It's important not to perpetuate these hidden legacy problems, and to ensure that the data

that populates the new system is fit for purpose, and delivers quantifiable improvement on what we had before.

Starting afresh with the minimum viable data affords the opportunity to focus on the data and workflows that offer the most business value, and that deliver efficiency and productivity improvements.

WHAT MIGRATION STRATEGIES ARE MOST APPROPRIATE? WHAT ARE THE ESTIMATES OF COST AND RISK?

Whether you're moving essential Atlassian application data to a new platform such as Atlassian Cloud, or there's been a recent acquisition resulting in the need to consolidate Atlassian application data from several different legacy instances, the headache is, how do I ensure access to the data whenever and wherever I need it?

A Big-Bang data migration execution is the process of moving an entire dataset from the legacy to the target system in one operation. This is typically carried out over a weekend or planned downtime period. On a Friday the legacy system might be in use but come Monday morning the business in question could be switched to the target system.

In order to mitigate as many change risks as possible, there are often several test migrations conducted. This often results in an increased cost and effort to complete days and weeks worth of remedial action and retesting. A typical Big-Bang data migration and consolidation of a Jira instance may take up to 30 days worth of additional effort.

The risk of Big-Bang migrations is accentuated by Atlassian application data volumes which have increased dramatically over the last 5 years. It is becoming more difficult to migrate volumes of data to platforms such as Atlassian Cloud; even with 48 hours of downtime to migrate data, a finite amount of time to perform testing with a rollback strategy will still be required.

Finally, while the traditional Big-Bang method of Atlassian application data migration or consolidation may be considered in the first instance, it results in moving everything including the problematic legacy data (identified previously), which won't necessarily deliver business value.

The business case for moving legacy data in a Big-Bang migration is often weak considering the risk, cost, and effort that's associated. An alternative method to significantly reduce risk and cost is the elimination of a Big-Bang strategy in favour of a more agile, iterative, and pragmatic approach.

THE ITERATIVE MIGRATION STRATEGY

If a Big-Bang is too risky/costly, what are the other options?

Iterative data migration or phased data migration, trickle-feed data migration, synchronised data migration, etc., all mean the same thing — moving only the valuable data in managed smaller increments until there's nothing left.

Historically, Atlassian focused on a Big-Bang mechanism using the application full-export and import. More recently, they've started developing Atlassian Marketplace Apps to allow iterative data migrations of Jira projects and Confluence spaces to the Atlassian Cloud platform.

Confluence Cloud Migration Assistant is currently available, while the Jira Cloud Migration Assistant is in early access beta.

There are two main challenges with an iterative Atlassian data migration strategy:

How can we keep our target and source systems data operable, until the migration is complete? How can we coordinate the migration of distinct elements of the business users and functionality, without breaking overall business continuity?

In order for an iterative data migration to succeed, we need to effectively plan the running of two systems in tandem for a transitional period, without impacting either.

Take a pragmatic approach: Move teams or business units one by one, starting with new teams and projects on the new system, and decommission old data on the legacy system.

For example:
o The IT Helpdesk could be updated so that if a customer needs a particular service, they're rerouted to a new Jira Service Management on the target system.

o When migrating project data from regional teams to a new central Jira instance, migrating essential data one team at a time may be the preferred option. When all teams have eventually migrated over to the central instance, they can all be shut down even if they're on the target system.

The iterative strategy, and migrating just the inflight or minimum data, allows effort that would otherwise be spent mitigating a Big-Bang approach, to be put toward delivering tangible business value instead.

For example:
o Improving workflows and business processes.
o Streamlining data and improving status reporting.
o Training and coaching on new ways of working.

HOW TO MANAGE DATA ARCHIVAL AND DECOMMISSIONING

Using the iterative data migration strategy, commonly results in 70-80% of data not needing to be migrated or consolidated to deliver business value, which may still be required for governance or record management. This remainder is the quantity of data that needs to be archived, decommissioned, and retrieved.

Defining required data seems obvious, but can lead to an entire analysis on its own. Is all the data I want archiving important? Is there enough detail to be valuable?

A Jira database containing information on 20,000 issues may seem important due to the volume, but what if the only valuable records are the few hundred identified and resolved defects?

It is not prudent financially or operationally to carry on a legacy of incomplete or irrelevant data just because it's there, otherwise, it's a waste of time, resources, and storage capacity.

One model, is to plan a managed iterative data migration, and decommissioning of the legacy system over a 12 month period:

- In the first months, select initial pilot teams and early adopters to migrate minimal essential data, and move to the target system first to shake-out any issues or challenges.
- As teams move to the target system, archive and decommission the old valuable data to back it up or use a raw data export such as CSV.
- Often, the vast majority of teams will have moved or planned to have moved to the target system within the first 3-6 months.
- At 9 months, plan how to manage the move of the final teams and provide additional support or remedial action for them.
- After 12 months, identify the remaining legacy record data that has objective value, and export and store it in a form that could be imported if and when required, i.e. a CSV data export from Jira, or a PDF Space export from Confluence.

AVOIDING THE BIG-BANG USING AN ITERATIVE DATA MIGRATION STRATEGY

Clearvision has many years of experience in migrating Atlassian application data between platforms such as Server to Atlassian Cloud and so on. We have successfully managed many Big-Bang data migrations, and know first-hand that there can be extended downtime and significant effort associated with the large volumes, complexity, and clean-up of unwanted legacy data.

The complexity and sheer scale of modern Atlassian application data migrations may mean having to think outside the box for a data migration project.

At Clearvision, we know there are agile alternatives offering iterative data migration strategies that can deliver not only greater benefits at a lower cost, but at a significantly lower impact and risk to the business too. ◻

If you're thinking of moving to the Cloud, but are concerned about the risk involved, learn about our Atlassian Cloud Enterprise Migration for organisations with multiple tools and instances.

For more information, please visit www.clearvision-cm.com.

08