




Migrating from Legacy to Modern Data Tools:

The 5 Things You **Must Know** Before Getting Started

—
eBook





The Arctic tern is a remarkable species of bird, known for its long-distance migrations. It travels from Greenland to Antarctica and back every year—some 44,000 miles in total annually. It makes your garden-variety data tool migration look like a walk in the park, right? Nah. All the Arctic tern has to do is fly from one of Earth's poles to the other—difficult, for sure, but not complicated.

But migrating from a legacy data tool to a modern one, such as the transition from SQL Server Reporting Service (SSRS) to Microsoft Power BI, or from SSIS to Azure Data Factory? Now that's both difficult and complicated, and it usually involves an extraordinary number of tedious, manual tasks just to lay the groundwork for what you hope is a smooth transition.

What is “Migration” Anyway?

In IT, the term “migration” refers to any transition from one system, platform, standard, or physical location to another. Server hardware can migrate from one data center to another, or the application hosted on a physical server can migrate to a virtual server or to the cloud. Program code can migrate from one version of .NET Framework to a more recent version. A company can migrate from one ERP system to another.

Because migrations are almost always painful, complex, time-consuming, high-risk projects, no one ever migrates anything just for fun. Usually, there is a circumstance that forces a migration (as when migrating an application from a deprecated operating system to a currently supported one) or a compelling business need.

In this eBook, “migration” refers to moving from an older data tool to one that is more modern, capable, responsive, and easy to use. Many enterprises are embarking on this journey because even though it's a huge project, the expected return on investment is too good to pass up.

Data migrations are no less complicated than any other kind, but there are actions you can take to make the transition relatively painless.

Here are 5 crucial things you can (must!) do to make your migration smoother:

1. Discard

Have you ever completed a physical move of your home or office and wondered, “I haven’t even looked at half of this stuff in years. Why did I pay good money to move a bunch of useless junk?”

It’s much the same with a data tool migration: Don’t migrate anything you don’t need. One of the first steps in a data tool migration is to take stock of your existing ETL processes/report portfolio. Chances are good that some—perhaps many—of your processes or reports have one or more of the following characteristics:

- It was needed for a one-time or short-term purpose and is no longer relevant
- The data underlying the process or report has changed in some way, and they can no longer be relied upon
- The process or report includes business processes or products that are no longer in use
- The process or report definition is narrowly tailored—perhaps for one consumer who is no longer with the company—and could provide misleading information to anyone else
- The process or report is based on legacy data that is no longer being updated, so they might not be useful anymore
- The process or report mostly or completely duplicates one or more other processes or reports in the portfolio



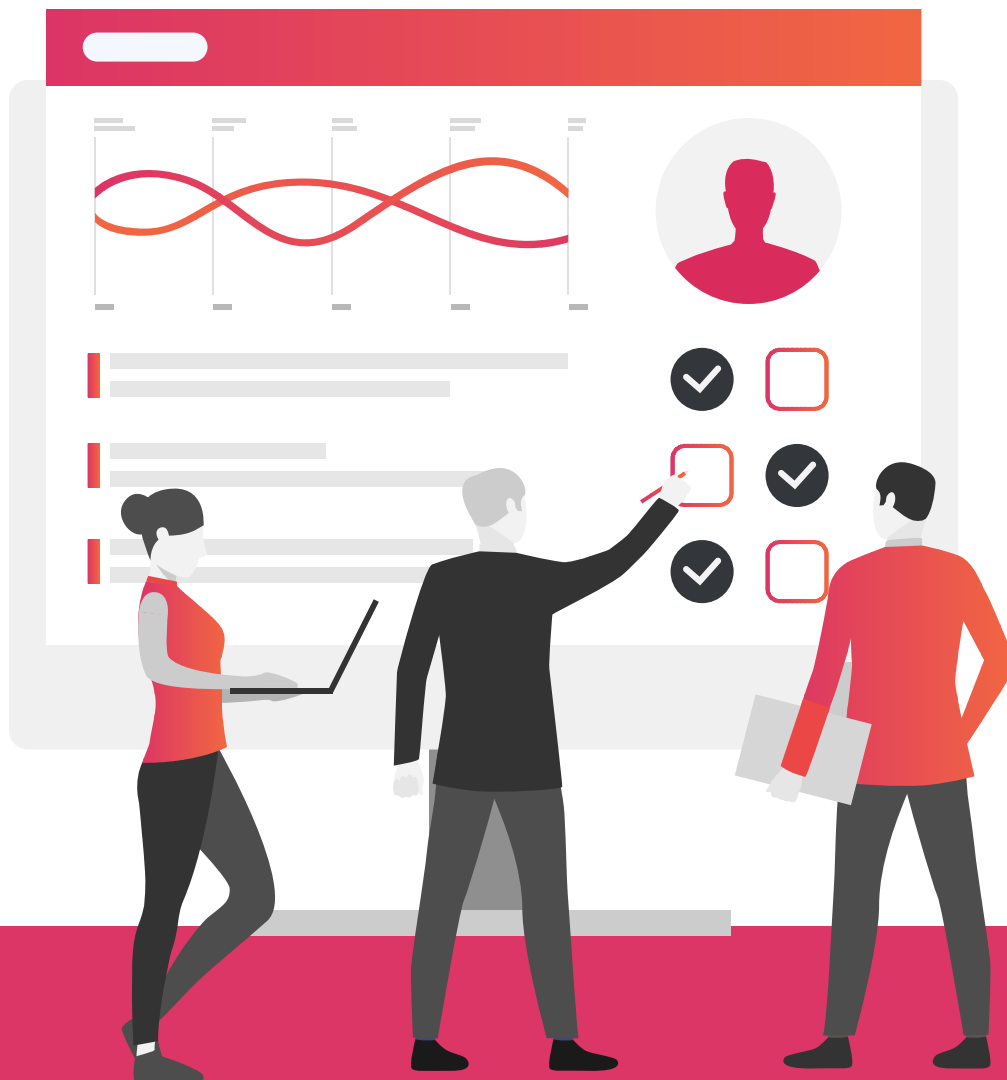
You get the idea. It’s time to pick out the processes or reports that are still useful and discard those that are unnecessary.

Depending on the size and age of the portfolio, this exercise could take a considerable amount of time. In a large organization, it’s not always easy to track down who is using each report or what the active processes are. Some reports may be consulted once or twice a year and some processes may run just once a year, but they are still useful and relevant. It may require some detective work, but it will be worth the effort because reproducing irrelevant processes or reports in the new tool is a total waste of time.

2. Compare

When migrating a data environment, there are many moving parts—databases, tables, views, ETL processes, reports and more—and it's easy to leave a component out of the migration or to include something you didn't intend to.

How can you verify that everything you intended to migrate—and only what you intended to migrate—actually made it across? You need a comparison tool—one that represents the two environments graphically and not a spreadsheet file that requires you to eyeball object names row by row. It doesn't take long before all those object names start to look alike.



3. Simplify

Nineteenth-century author Henry David Thoreau of Walden fame had it right: “Simplify, simplify.” He didn’t have to deal with a data tool migration—lucky guy—but his words are still sage advice for a data migration project. Many data environments are littered with complicated components: ETL processes that are based on other ETL processes, which are fed by still more ETL processes, and on and on as in a decision tree, or you might encounter reports that are using a very long and complicated native SQL query. The result is a fragile environment in which a failure of any one of those processes (or leaves of the decision tree) could blow up the whole system—or worse, go undetected for a long time while your reports are providing bad or incomplete information.

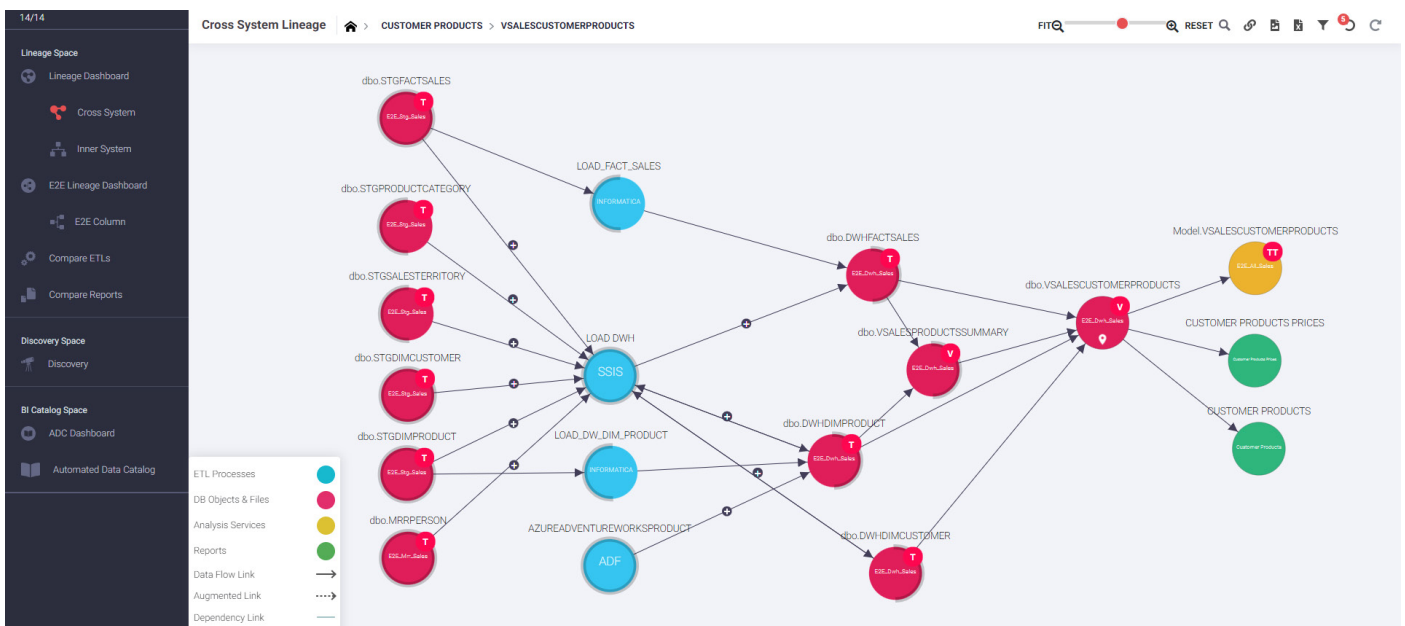



Figure 1: Cross-system lineage for the view named VSALESCUSTOMERPRODUCTS. This lineage screen shows that the view is being used by the Tabular table (yellow), as well as two reports (green). It sources data from various database objects (pink) including a view and different tables which are being populated through various ETLs (blue) including Informatica, SSIS and ADF.

To the extent possible, you should simplify these processes or reports and reduce the interdependencies among various types of data (finance, marketing, sales, and so on). This is extra important if the migration task is divided up this way, with different people handling each piece.

However, that’s difficult to do without a good way to visualize the environment. Complicated data environments are usually difficult to visualize—the “big picture” soon resembles a bowl of spaghetti.



4. Improve

As arduous and time-consuming as a data tool migration project can be, it affords a rare opportunity to improve not only the tools but the underlying data environment.

Most data environments that modern businesses want to migrate were built many years ago and tinkered with ever since. Lots of fingerprints have been left on them by people who didn't always know what they were doing. They may have known what they wanted to accomplish but not the "right" way to do it.

Tools have evolved over the years as well. Components that were written as complex SQL code and long-stored procedures are still around, though, because no one wants to touch them for fear of breaking something. Now is the chance to upgrade those components using modern ETL tools or new capabilities in reporting tools.

In short, an important part of any data tool migration project is to use the opportunity to clean up the mess, get rid of the poorly written (and poorly documented) code and kludgy workarounds, adopt modern best practices, enforce some consistency, and leave an easy, logical, working environment for the next crew that comes along.



5. Protect Your Data Critical Path

In project management, a “critical path” is the series of project tasks, each of which must be completed on time, or the entire project schedule will slip. Critical path data, if not managed properly, can not only lead to poor/misled decision making by management, but also to potentially detrimental mistakes with direct impact on customers such as money wasted on sending technical teams to the wrong location.

Data environments have a similar “critical path” concept. It’s the series of components (tables, ETL processes, reports, and other artifacts) that must be migrated correctly; failure of any one of these components can put the entire migration project at risk. Your most important reports—the ones that upper management depends on, and the ones that end up in regulatory filings—are naturally based on data in the critical path, so getting the critical-path pieces right is an absolute must.

Good data migration tools can help you understand the sources and processes associated with the critical path, and they can show you the impact of a change to any one of these components. The two related techniques that good tools enable are root cause analysis and impact analysis.

- **Root cause analysis:** A good data migration tool enables you to determine which ETL processes are loading a particular table or report; you need to be extra careful when migrating these processes to ensure you haven’t missed even a single lookup, SQL override, or SQL command. You can use root cause analysis to find all of the reports that rely on this critical path and double check them before you approve the migration process. If you’re migrating one or more ETL processes or from one reporting tool to another, this step is mandatory - unless you like doing things over and over again.
- **Impact analysis:** This tool shows you the paths from database tables, through any ETL processes, to all the reports that depend on the data, and it enables you to double-check all the reports to make sure they have everything they need.



The Key to a Successful Migration: Metadata Management Automation

You might be thinking to yourself at this point, “Great tips, but how the heck am I going to be able to do all that? I can’t even find the source of a reporting error without spending hours digging around my systems.”

Well, metadata management automation is key to be able to do all five of these things properly and effectively ahead of migrating systems. At the end of the day, everything in data depends on metadata—discovering it, cataloging it, tracking it, understanding it, and ultimately using it to create the reports and dashboards that businesses need to be successful and data-driven. Migration is no different.

Need to verify that everything you wanted to migrate made it across? Octopai provides views of both the old and new environments in one screen so you can do an apples-to-apples comparison and learn whether the environment you wanted at the far end of the migration is the one you got.

Want to simplify prior to migrating? Octopai’s visualization tools can help

Those lucky old Arctic terns just need their instincts to successfully migrate. You need intelligence. Data intelligence. With Octopai’s Data Intelligence Platform, you’ll be able to migrate your data tools wisely, effectively, and in a manner that will deliver cost-effective data-driven insights for years to come.

here too, clarifying the dependencies and enabling team members to work together so that everything that needs to migrate is included and nothing is overlooked.

Octopai’s metadata management automation can reduce or eliminate many of the tedious, time-consuming, and error-prone activities that have been performed manually. The visualization, root cause analysis, and impact analysis tools already mentioned can shorten the preparation phase of any data tool migration project in a significant way.

Wish to facilitate your users’ adjustment to the new data systems? Octopai’s automated data catalog, which can be automatically built and populated as part of the migration, can help users locate the data assets in their new locations, communicate with data owners, stewards and other users about the assets and improve the usability of and trust in the data.

