
Handling Database Migrations with Embarcadero Solutions

Embarcadero Technologies
September 2004

Corporate Headquarters
Embarcadero
Technologies
100 California Street,
12th Floor
San Francisco, CA 94105
www.embarcadero.com

EMEA Headquarters
Embarcadero Technologies,
Ltd.
Thames House
17 Marlow Road
Maidenhead
Berkshire
SL6 7AA
United Kingdom
www.embarcadero.co.uk

Table of Contents

Introduction	3
Database Migration Considerations	3
Embarcadero Data Lifecycle Management Solutions	4
Migrating Database Schemas with Models.....	4
Moving Data with Models	5
Ensuring the Success of a Migration Effort.....	7
Planning for High Performance.....	7
Guaranteeing Performance Success.....	8
Reducing the Database Management Learning Curve.....	8
Easing Business Logic Migration.....	9
Simplifying Cross-Platform Administration	9
Effectively Managing Migration Changes	9
Automation Made Easy.....	10
Conclusion.....	10

Introduction

Database professionals can count on being involved in a database migration at least once – and probably more than once – in their careers. Unless the migration involves a small, trivial system, such operations are not insignificant to say the least. Why do migrations occur? While there can be many reasons for performing a database migration, the motivation behind such a thing usually boils down to one of three issues:

- **Scale** — the current database configuration cannot support the current or upcoming workload and so a new database platform must be acquired.
- **Cost** — the current system has become cost prohibitive given its contribution to the overall business, so a more cost effective database platform must be found.
- **Integration** — for business reasons such as a corporate acquisition, the current database needs to be consolidated and integrated into another system that is not the same (but can be the same) database platform.

As has already been stated, the movement to and from even similar platforms can be challenging, but when a dissimilar database platform move is attempted (like Oracle to Microsoft SQL Server) or a new operating system is utilized, then things can quickly become complex. This paper discusses the issues surrounding database migrations and focuses on how data lifecycle solutions from Embarcadero Technologies can take the bite out of such endeavors.

Database Migration Considerations

The most common database migrations involve migrating a database to a new database engine type (like Sybase to DB2), migrating a database to a new operating system platform (such as from HP-UX to Red Hat Linux), or consolidating one database into another database that may or may not be of the same engine type. When addressing a migration effort, the team charged with tackling the migration project generally has three core mandates:

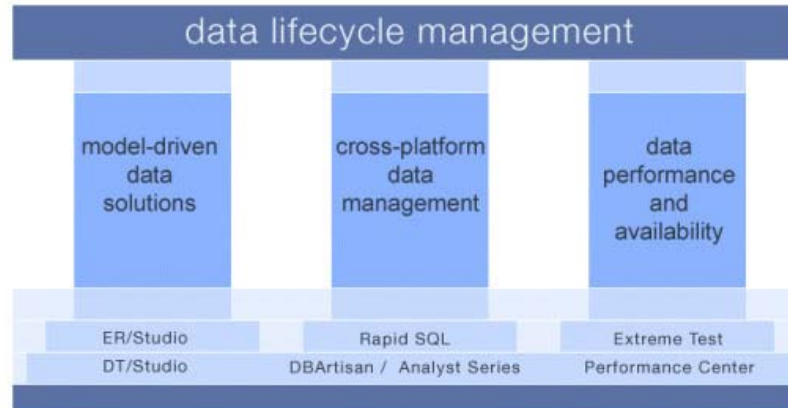
1. **Thoroughly understand the source and target systems** — the team has to become very intimate with the design of the source and target database as well as the underlying operating system if different O/S platforms come into the mix.
2. **Define the physical migration scope** — are just database structures being moved, or will data be included? What about the business logic embedded into the source database? These questions and more will have to be addressed and understood, and a plan to accomplish the actual physical migration must be crafted.
3. **Overcome the personnel learning curve** — unless experts are brought in who understand the internals of a new database or operating system platform, the existing staff must determine how they will become productive on the new configuration in the shortest time possible.

Without question, these three tasks can involve a lot of complexity. Fortunately, tools from Embarcadero Technologies can make short work of any migration effort, no matter the size of the system being converted.

Embarcadero Data Lifecycle Management Solutions

Embarcadero Technologies offers solutions for every area of the data lifecycle. All Embarcadero lifecycle management products offer near-instant ROI and dramatic cost savings because of their ease-of-use and powerful, timesaving productivity aids.

Embarcadero's integrated product set consists of six core products that handle the bulk of data lifecycle management tasks, along with a number of supporting products that round out the complete data lifecycle suite. Embarcadero offers complete cross-platform coverage, so it does not matter if Oracle, Microsoft SQL Server,



DB2 UDB, DB2 z/OS, Sybase, or MySQL is being used; everything looks and feels the same. Plus, all database platforms can be seamlessly managed from within the same core consoles, so a data professional does not have to jump in and out of separate products to manage their heterogeneous environment.

The first steps of a migration project involve understanding the source and target platforms and then migrating the existing structures and data to the new platform. Let's look at how each step can easily be accomplished using tools from Embarcadero Technologies.

Migrating Database Schemas with Models

It's obvious that database professionals need to practice formal logical and physical data modeling, however such activities take expertise, time, and attention. Modeling tools can not only assist with the creation of a new database design, they can also play a key role in migrating existing database designs to new systems.

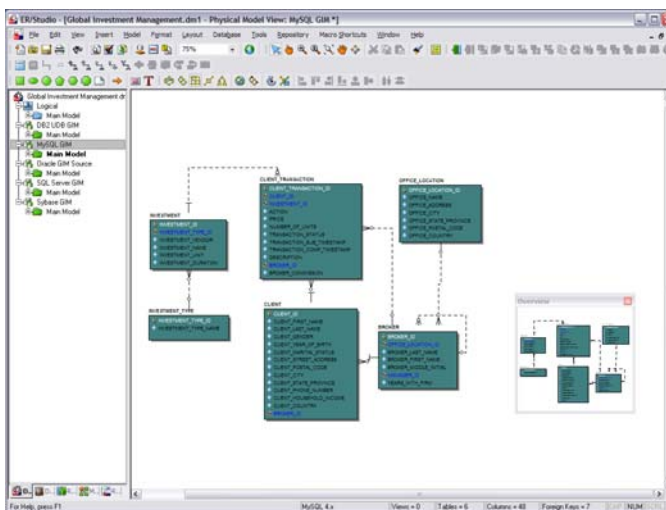
Migrating models and physical designs are not functions to be carried out with pen and paper; instead, they should be facilitated through the use of strong modeling tools that assist in the proper creation of logical models as well as in the building of physical designs that target actual database engines. There is a vast array of features that data architects, modelers, and DBAs should consider when choosing a data modeling tool and such a complete list is beyond the scope of this document. However, a core set of must have features include:

- **Multi-level design environment** — the modeling environment should support both logical and physical modeling with intelligent mechanisms to sync either way.
- **Data dictionary support** — to support data standardization and data element reuse, data dictionary support should be present.

- **Reverse/Forward engineering** — the tool should be able to perform changes against an existing database through reverse engineering an existing design, as well as be able to build a completely new database through DDL generated by a physical model.
- **Documentation strengths** — to promote understanding and communication of logical and physical models, the modeling environment should be able to document designs in a variety of ways (HTML, Microsoft Word, etc.)
- **Team based collaborative abilities** — large models and designs are normally constructed through a team of architects or modelers, so the modeling environment should support the ability for team members to work on parts of a global model and then merge them back into a single, consolidated model.

Embarcadero's ER/Studio meets all the requirements listed above for a robust modeling solution plus offers much more. Its powerful combination of logical modeling and physical design capabilities make it the ultimate solution for serious data architects, modelers and database administrators. Data architects and modelers will appreciate the complete support for graphical logical design and data dictionary support in ER/Studio, and DBAs will enjoy the near instantaneous transformation of a logical model to a physical design that the tool can produce.

ER/Studio makes short work of understanding a source or target database. Its powerful reverse-engineering capabilities can visually model even the most complex database systems, and with one mouse click, a database can be completely documented in an MS Word document or an HTML-based intranet site. Converting databases to new database engine types is a snap as well. A data architect can reverse engineer, for example, an Oracle database and then instantly convert it to a physical SQL Server database with all datatype transformations being invisibly handled.



Once a DBA or data architect has put the finishing touches on their physical design in ER/Studio, they are only a few mouse clicks away from generating all the DDL necessary to create a physical database along with all accompanying objects (tables, indexes, procedures, triggers, etc.) And should the design effort require a group of modelers, ER/Studio's optional repository provides the mechanisms necessary for a team environment to operate with complete check in/check out and merge capabilities being present.

Once an architect or modeler builds a target physical database through ER/Studio, the next step is to migrate data from an existing system to the newly created database. To assist with this process, models once again come to the rescue.

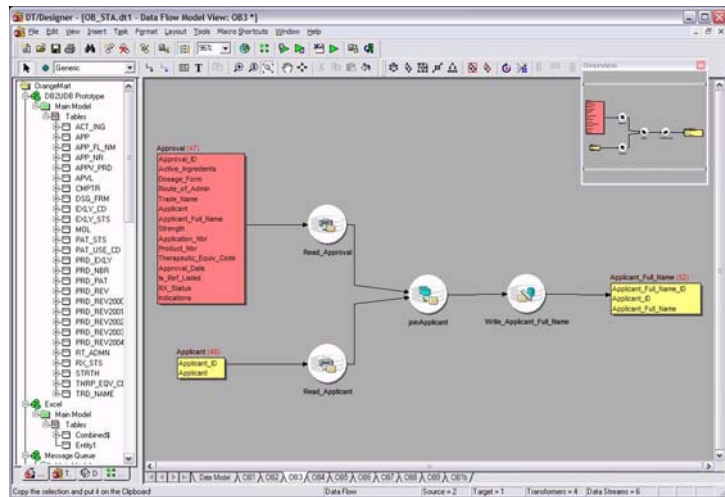
Moving Data with Models

Data movement and migration projects are rarely simple assignments. Developers, DBAs, and Extract, Transform, and Load (ETL) specialists oftentimes struggle to wrap their arms around

complex migration efforts. Some use custom scripts coupled with bundled database utilities, but when it comes to serious and ongoing data movement challenges, the need for more capabilities is often expressed. Large migration projects are especially difficult for the following reasons:

- **Data source and target understanding** — this fact has already been covered, but bears repeating. The need to deeply understand the source and target systems is oftentimes complicated by the sheer number of source or target destination objects (or both) and all their inter-relationships.
- **Large data volumes** — moving a few hundred megabytes is one thing, but moving half a terabyte is many times a different matter.
- **Complex data transformations** — rare is it that data is moved without modification from a source to a target database. More often than not, the data is sliced and diced in many different ways so it can become usable on the new target system.

To cope with complex and large data migration projects, DBAs, developers, and ETL specialists can turn to Embarcadero's DT/Studio. DT/Studio is a model-driven ETL solution that uses models to facilitate the understanding of source and target systems. The visual modeling interface of DT/Studio allows those involved in ETL efforts to represent their source and target systems in models found in data design tools like ER/Studio, which takes all the guesswork out of knowing how data is to be moved. Users of DT/Studio can visually design their dataflows between systems and not have to worry about the underlying complexities of how data is read or written, although they can control every aspect of such operations if they wish. DT/Studio supports all the major database engines as well as being able to deal with data that is housed in flat files, XML files, and more, so if another platform is involved in a data migration effort, DT/Studio will have it covered.



Database professionals needing to move large amounts of data can do so confidently with DT/Studio because the solution is able to scale to meet the challenges of very large database installations. DT/Studio's data movement engine is written to run on many different operating system platforms (MS Windows, Linux, UNIX, etc.) and take advantage of multi-CPU machines so even large data movement tasks are completed in the time window desired.

When it comes to transforming incoming data, it's hard to beat DT/Studio's extensible transformation framework. Common transformation tasks (horizontal, vertical partitioning, etc.) can be designed through visual wizards. Detailed transformations can be carried out through the use of DT/Studio's over 1,000 built-in data transformers or through customized data alteration tasks that can be written by ETL specialists and saved for later use. Built-in scheduling abilities let ETL professionals create data migration jobs for recurring execution.

Migration tasks that are built in DT/Studio can be versioned, which allows rollbacks to previous migration tasks to be performed as well as comparisons between versions. Task reporting is also available, which makes the tracking of ETL operations from a business perspective clear and easy. These features and more leave no room for doubt that DT/Studio has what it takes for serious ETL projects to be successful.

Ensuring the Success of a Migration Effort

A successful migration to a new platform means more than just converting or consolidating a database schema and data over to a new server. For a migration project to truly succeed, the IT team performing the migration has to do more than just hope that a new application will be able to hold up under a serious production load or handle ‘N’ number of transactions per second. Instead, they need to properly stress test and monitor the new database-driven system to ensure that the application will perform as expected on Day One of production.

Planning for High Performance

Proper performance testing catches the showstopper and performance vampire problems that inadequate user and quality assurance testing miss. In a nutshell, performance testing simulates what is expected from real world use. It stresses a database server and application in ways that could otherwise only be accomplished by opening the floodgates of the production user community.

Smart performance testing uses the following elements to pull off a realistic simulation of what a database will experience during expected production usage:

- **Anticipated user presence** — it is critical that the test simulate the number of user connections that are expected during peak times and normal working hours. This is the major area where manual methods that pick a subset of users to test a database and application fail. The database server may run just fine with 15 or so user connections, but may fall over when 300 connect to the system.
- **Repetitive user activity** — once the anticipated user sessions have connected to the database, they obviously have to “do something” for the system to be stressed. And they cannot just “do something” once. Either all or a portion of the connected sessions need to repetitively perform tasks as they would occur during a normal workday. For an OLTP system, this may mean entering multiple orders. For a data warehouse, this may mean issuing long running analytical queries. The key is that the work is repetitive so repeated blows are dealt against the database.
- **Extended duration** — once a DBA sets a number of sessions performing repetitive work, they next need to ensure that the work continues for a period of time that makes the test meaningful. What a DBA is looking for is to unearth problems that take time to develop. For example, a database table may not become fragmented after 30 minutes of OLTP work, but may surprisingly fragment in a dramatic fashion after 2 or more hours of repeated action.

So how is this practically accomplished? How can performance testing address both the hit-the-wall problems along with the silent killers of database performance? How should a DBA arrange tests to flush out all the major headaches that threaten a database server? Even database professionals who understand the need for real world testing of their database and know what

areas need to be stressed are oftentimes at a loss for practically creating the scenarios they need for pulling off a true performance test. This is not surprising when one considers all the manual work that has to go into such an effort for it to yield the answers that will give the Go/No go signal to a database project team. It's oftentimes much easier to perform a small test with a selected group of willing users and then cross your fingers on the first day of production or when a massive upgrade begins.

This does not have to be the case. Embarcadero's performance testing tool, Extreme Test, is designed to bring peace of mind to a database staff by allowing them to know what areas of a database server and accompanying application/Web system will break prior to it going into production. By simulating production or accelerated user and data volume, Extreme Test can make a database server unbreakable by letting a staff understand the weaknesses of a database server before the end users find them, with the end result being ironclad availability and performance for production systems.

Extreme Test allows a DBA to capture and simulate the activity of hundreds or even thousands of users, and then apply the resulting workload in varying amounts against a target database. While generating database load, Extreme Test monitors the underlying database components to identify the ones that are struggling to handle the stress. It captures and stores test results in a central repository that a DBA can query for further analysis and reporting. Extreme Test makes it exceptionally easy for an administrator and QA team to implement repeatable testing, to isolate bottlenecks, and to measure the impact of corrective actions on overall database performance. And for those wishing to test system components other than a database (application server, Web server, etc.), Extreme Test can build a single performance test against all defined components to produce a real world picture for an entire system.

So to ensure a DBA does not have an unwelcome surprise waiting for them in a database server, they need to apply proper performance testing to a system to uncover all problem situations before they impact a business. The best way to do this is to use intelligent and powerful testing solutions like Extreme Test that can provide the answers needed before a database server sees the first day of production or before a large upgrade occurs. Nothing serves as a better confidence builder than knowing what to expect.

Guaranteeing Performance Success

Over time, a newly migrated system can break down and quickly begin to show signs of reduced performance. Therefore, a new system should be constantly monitored so that a database staff is instantly notified if performance falls below certain thresholds or if availability is threatened. Embarcadero's Performance Center is the perfect weapon to use for such critical systems. Performance Center monitors cross-platform database applications around the clock and proactively notifies personnel of any performance abnormality. With Performance Center on guard, the continued success of any new migration effort is ensured.

Reducing the Database Management Learning Curve

IT organizations are oftentimes hesitant to switch to another database or operating platform for fear that their personnel who are experienced on an existing RDBMS will not be productive on a new platform. The learning curve, they fear, may be too steep. However, database management software from Embarcadero Technologies can greatly ease the process of managing one or many unfamiliar systems. There are three core sets of software that database developers and administrators will find valuable.

Easing Business Logic Migration

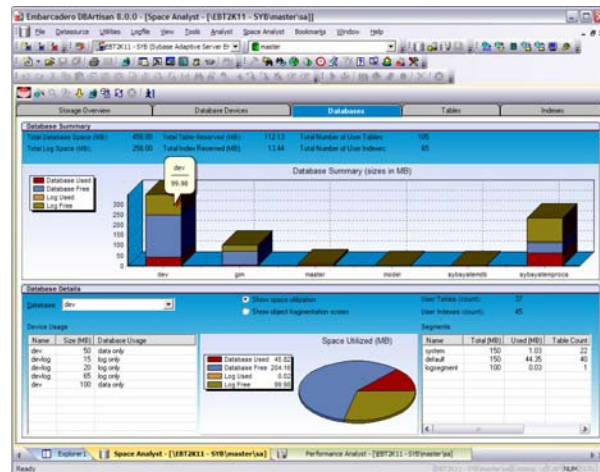
The movement of business logic to dissimilar platforms can be tough. Oracle PL/SQL is much different than Microsoft Transact-SQL, for example. If business logic migration is involved in a migration effort, then there are two plans of attack that are offered through Embarcadero.

First, Embarcadero's partner Adventnet offers a number of database logic migration tools that can transform database syntax from one database engine into that of another. A developer can input a SQL Server stored procedure into Adventnet's tools and have it produce an Oracle (or other platform) complement. Adventnet supports a wide range of platform-to-platform possibilities, so database professionals will likely find their situation is covered.

For custom business logic development, developers will use Embarcadero Rapid SQL, which is a cross-platform development IDE that contains a powerful SQL programming environment and supersedes anything found in the industry. Developers will enjoy powerful code generation abilities, many coding aids (cross-database search utility, visual object creation wizards, data imports from Excel spreadsheets, etc.), along with built-in version/source code control integration with all the popular source code software vendors. Rapid SQL is just what the doctor ordered when it comes to creating or modifying database-driven business logic.

Simplifying Cross-Platform Administration

Administrators will find DBArtisan the answer to their need for a full, cross-platform database administration solution that supports all the major database platforms. DBArtisan greatly reduces the learning curve of switching from one platform to another because everything looks and feels the same across all platforms. Everything from storage and security to object management is handled via visual wizards and editors that guide those unfamiliar with a new database engine's syntax through the process of creating and managing databases and objects. Powerful utilities make short work of extracting object DDL, running object checks, documenting databases, and monitoring performance. With all this and more, DBArtisan can make any DBA productive on a new platform in no time.



Effectively Managing Migration Changes

During migration efforts, the target system oftentimes needs to be tweaked to make things just right. This can involve things like testing new indexing schemes and handling programming changes to business logic that needs to be moved from a test system to a production server. Change control is a difficult task to accomplish manually, but Embarcadero Change Manager simplifies the change control process, by visually automating every change control task and taking away the burden that comes from migrating to unfamiliar platforms. DBAs and developers can easily capture and track changes to the most complex database schemas. Migrating schema changes between platforms is easy too. And finally, powerful compare and synchronization

functions provide a safety net that instantly pinpoints and rolls back unwanted changes. There's no doubt that Change Manager can play a big part in ensuring the success of any database migration project.

Automation Made Easy

Finally, developers and administrators alike often struggle with the process of scheduling and monitoring database and system jobs. Creating and managing jobs across multiple servers can be a real challenge, but not with Embarcadero Job Scheduler. Job Scheduler is an enterprise-wide job scheduling tool that dramatically reduces the time it takes to create, run, manage, and monitor cross platform database jobs. A DBA or developer can easily create SQL or operating jobs from their Windows desktop and then immediately deploy them to any Windows, Linux, or UNIX machine running in their network. A central desktop console allows everything to be controlled and managed from one place. Job Scheduler also offers tight integration with DBArtisan and Rapid SQL so DBAs and developers can create and schedule any SQL job without leaving their favorite development or administration tool. And for those moving to new operating system platforms (like from UNIX to Linux), Job Scheduler offers CRON import abilities that make moving from UNIX to Linux a breeze.

Conclusion

Whether database migrations are small and straightforward or large and complex, Embarcadero Technologies offers the shortest path to a successful migration effort. With its model-driven data solutions, Embarcadero can quickly and accurately migrate anyone to a new (or consolidated) database platform. Embarcadero can also ensure that the new system will hold up under expected production loads with its performance testing and monitoring solutions. Finally, Embarcadero dramatically reduces the new platform learning curve with its visual database administration and development tools.

The end results of utilizing tools from Embarcadero Technologies in a migration project are dramatic time savings, reduced costs, high performing information systems, and a well-equipped and productive database staff.

Embarcadero, the Embarcadero Technologies logos and all other Embarcadero Technologies product or service names are trademarks of Embarcadero Technologies, Inc. All other trademarks are property of their respective owners.