White Paper

Pro2 Replication Suite

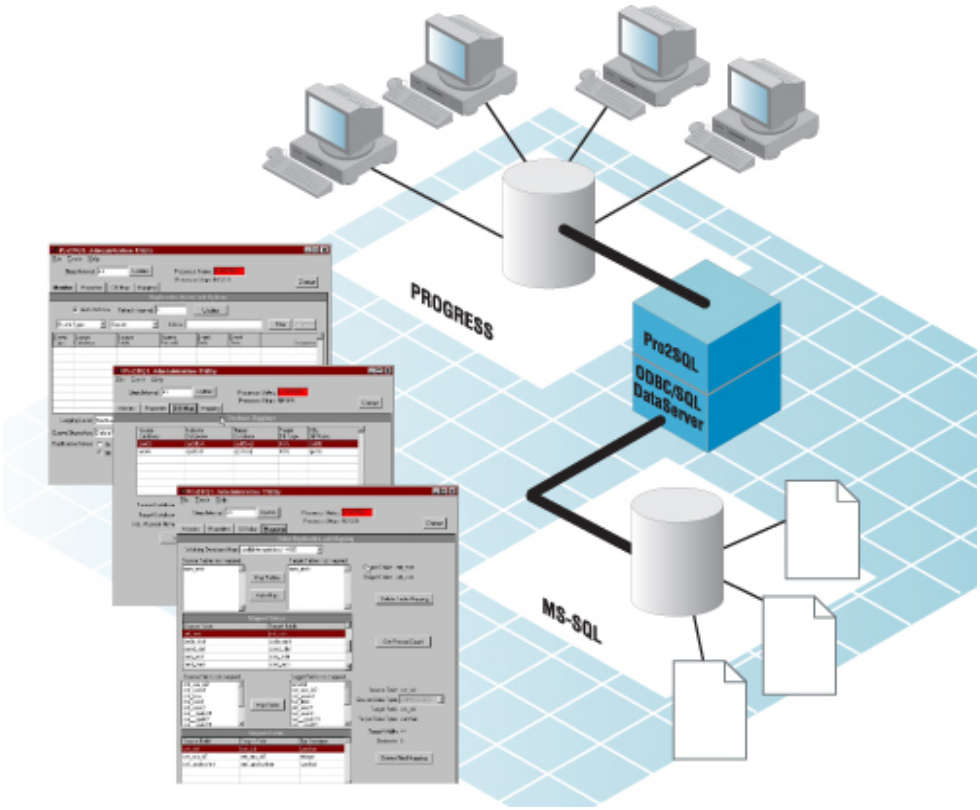

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880
www.BravePoint.com

BRAVEPOINT

Contents

Introduction

The BravePoint Pro2 Replication Suite (Pro2*) is the easiest and most cost effective approach for replicating PROGRESS OpenEdge data to Microsoft SQL Server or to the Oracle RDBMS.

No more slow ODBC pulling from PROGRESS while 4GL users compete for resources or nightly dump and load routines that may not work. Data is in near real-time, instead of 24 hours old. The actual data format and replication time interval is totally configurable by you. Another major advantage of the Pro2 Replication Suite is that all data transformations can be done using the Open Edge 4GL. This provides for maximum flexibility and excellent performance since the solution is totally native to PROGRESS.

Pro2* takes advantage of built in replication support for PROGRESS databases and provides the most reliable platform for consistent performance and data integrity with a minimal footprint.

Pro2* is easy to install and configure and can be up and running within a few days. Additionally, Pro2* is easy to maintain with robust administration tools so that your current support team can easily administer the product. Behind Pro2* is an expert support team that Bravepoint has been recognized for since 1987.

Pro2* can replicate data from PROGRESS v8.x, v9.x and OE 10.x and replicate to Microsoft SQL 2000, 2005 and 2008.

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880 www.BravePoint.com

Executive Summary

There are many reasons why companies want to replicate their PROGRESS OpenEdge data to SQL Server or Oracle; heterogeneous application integration, reporting, data archival and Business Intelligence are just a few. But there are just as many pitfalls and complications that make it a difficult and risky adventure. Pro2* was built with these requirements in mind, while minimizing the risks involved in moving meaningful data from one database platform to another.

The *push vs. pull* question: Many companies have plenty of Microsoft resources; however it is often the situation that the PROGRESS resources are not as available. The common first attempt is to "pull" the data from PROGRESS using Microsoft SQL Server Integration Services (SSIS). This is a perfect solution for a single "one time" data pull to populate a SQL database, but not practical for replication, real-time data or long term needs. Most companies that have tried this approach have decided that this works best for single data pulls for conversions and not repeatable replication tasks.

The next option companies have tried is the "dump and load" method. Microsoft SSIS can cause severe performance degradation on a PROGRESS database because ODBC access to the 4GL/ABL database is somewhat intrusive to online 4GL/ABL users. Companies have overcome this obstacle by only dumping smaller amounts of pertinent PROGRESS data and loading into the target database based on a scheduled batch job. This may work, until users start complaining that the data is not timely enough. Additionally, this option can be administratively expensive when changes to data schema and/or variable data problems require code changes to these customer dump and load routines. Risk factors are more prevalent with this option because you are moving ASCII files around and polling for changes with scheduled tasks with little, if any, error handling if the process should break down.

The Pro2 Replication Suite was created to overcome these challenges. In a perfect world, you would like to replicate data from PROGRESS to the target database in real or near real-time. Utilizing PROGRESS native support for replication, Pro2* can push immediate data changes to SQL with a very small footprint. Furthermore, Pro2* insulates you from the differences between PROGRESS being a row-oriented, variable-length database and SQL being a page-oriented, fixed-length database.

The Pro2* Administrator is a robust interface for managing and maintaining the replication link to SQL Server and Oracle. This tool handles all the hidden issues that present themselves when dealing with PROGRESS to SQL, minimizing risk and possible loss of the replicated data. It is easy to use and your existing staff can use this tool, so when changes need to be made, they can be done internally.

More detailed information and product videos can be found at www.Pro2SQL.com.

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880 www.BravePoint.com

The Pro2* Solution

Pro2* utilizes the strengths of both the PROGRESS OpenEdge database and OpenEdge 4GL/ABL to achieve the most dependable and configurable solution for PROGRESS to SQL replication.

Real-Time Is Better Than Old Data

Because Pro2* utilizes replication triggers, the data is replicated as users make changes to the data by adding, changing or deleting in near real-time. Even if the link between the PROGRESS and SQL Server databases is down, Pro2* queues the activity and will catch up when the link becomes available again. Users will enjoy being able to report on data that is current rather than reporting on yesterday's financial data without negatively impacting your 4GL/ABL application.

Reduced Risk

There are no external applications like SQL Server Integration Services, cron or task schedulers to rely on, no intrusive ODBC connections to PROGRESS, and no dump files getting moved around from machine to machine. When something changes either in the data schema or the required replication objects, there is no need for redevelopment and testing by your operations team. Fewer moving parts mean better reliability, and fewer chances for data integrity issues. Because of the reduced risk, Pro2* will save you money by eliminating lost data, data integrity problems and the need to constantly change programs and batch files as the requirements change. Pro2* insulates your company from the obstacles of replicating PROGRESS data to MS SQL Server or Oracle.

Set It and Forget It

Once Pro2* is configured, there is very little else that needs to be changed. The Pro2 Replication Suite is an enterprise-class application that will allow for flexibility to adapt to changing requirements, and save costs for development efforts and chasing down production issues. Pro2* will work with single or multiple sources and targets regardless of where they reside. Pro2* can also use multiple threads to guarantee the best possible performance when replicating. Pro2* supports all 4GL/ABL data types: logical fields, arrays, dates, etc. Pro2* has a robust administration tool that allows for easy maintenance and enhancements.

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880 www.BravePoint.com

Pro2 Administration Utility:*

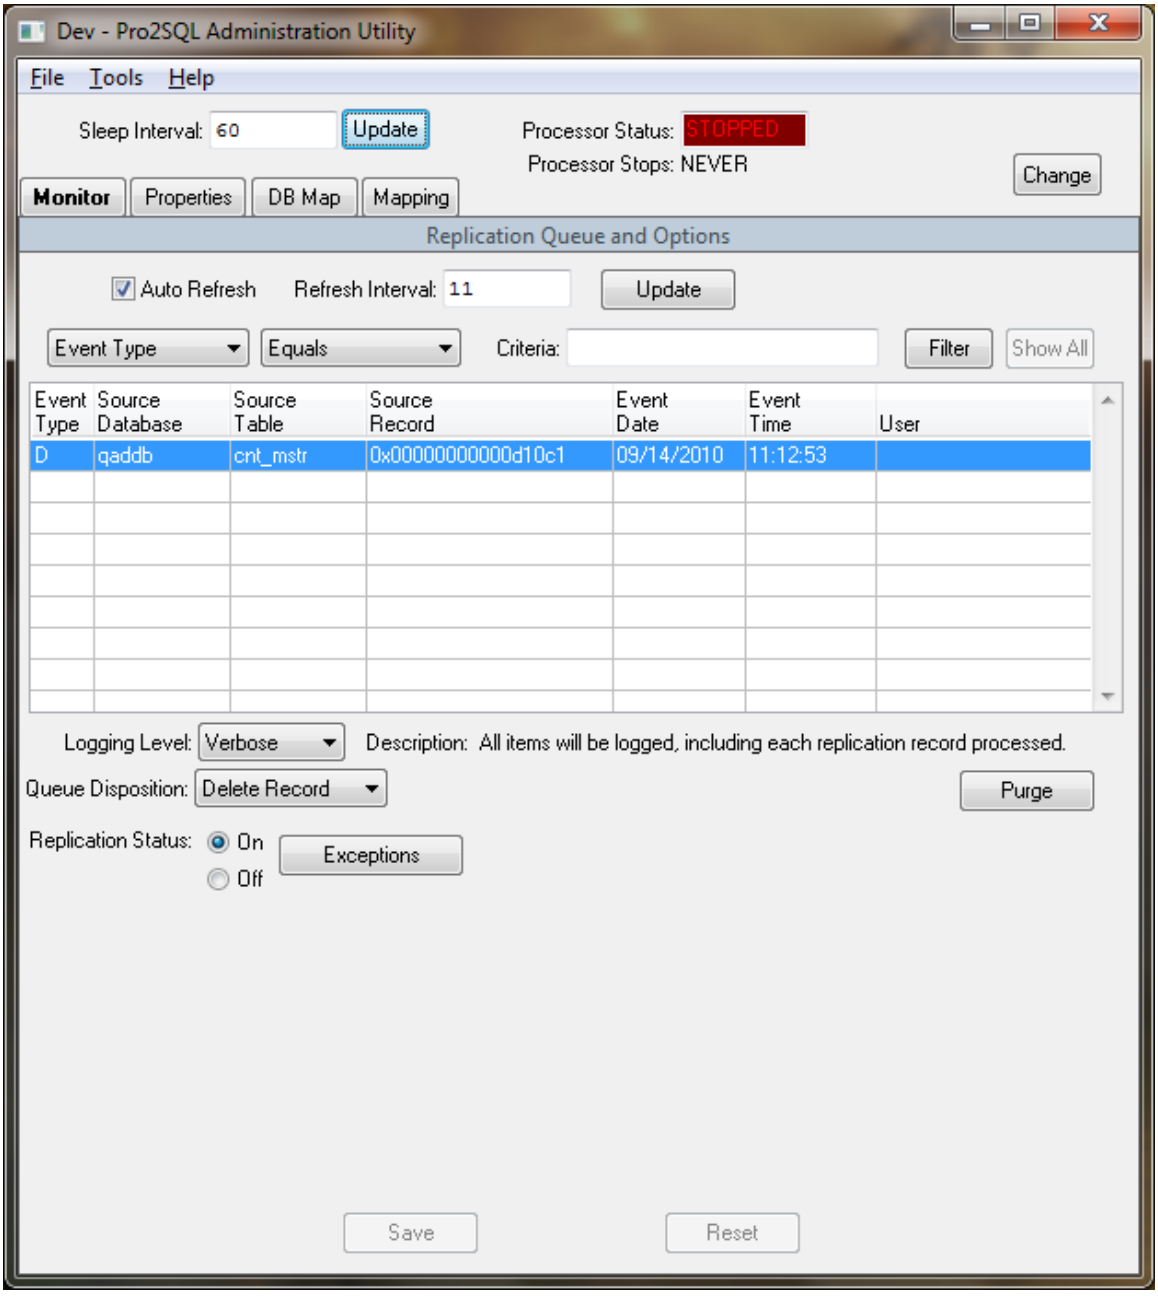

The Pro2* Administration Utility allows you to see what records are in the queue to be replicated.

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880 www.BravePoint.com

The Pro2 Mapping Utility:*

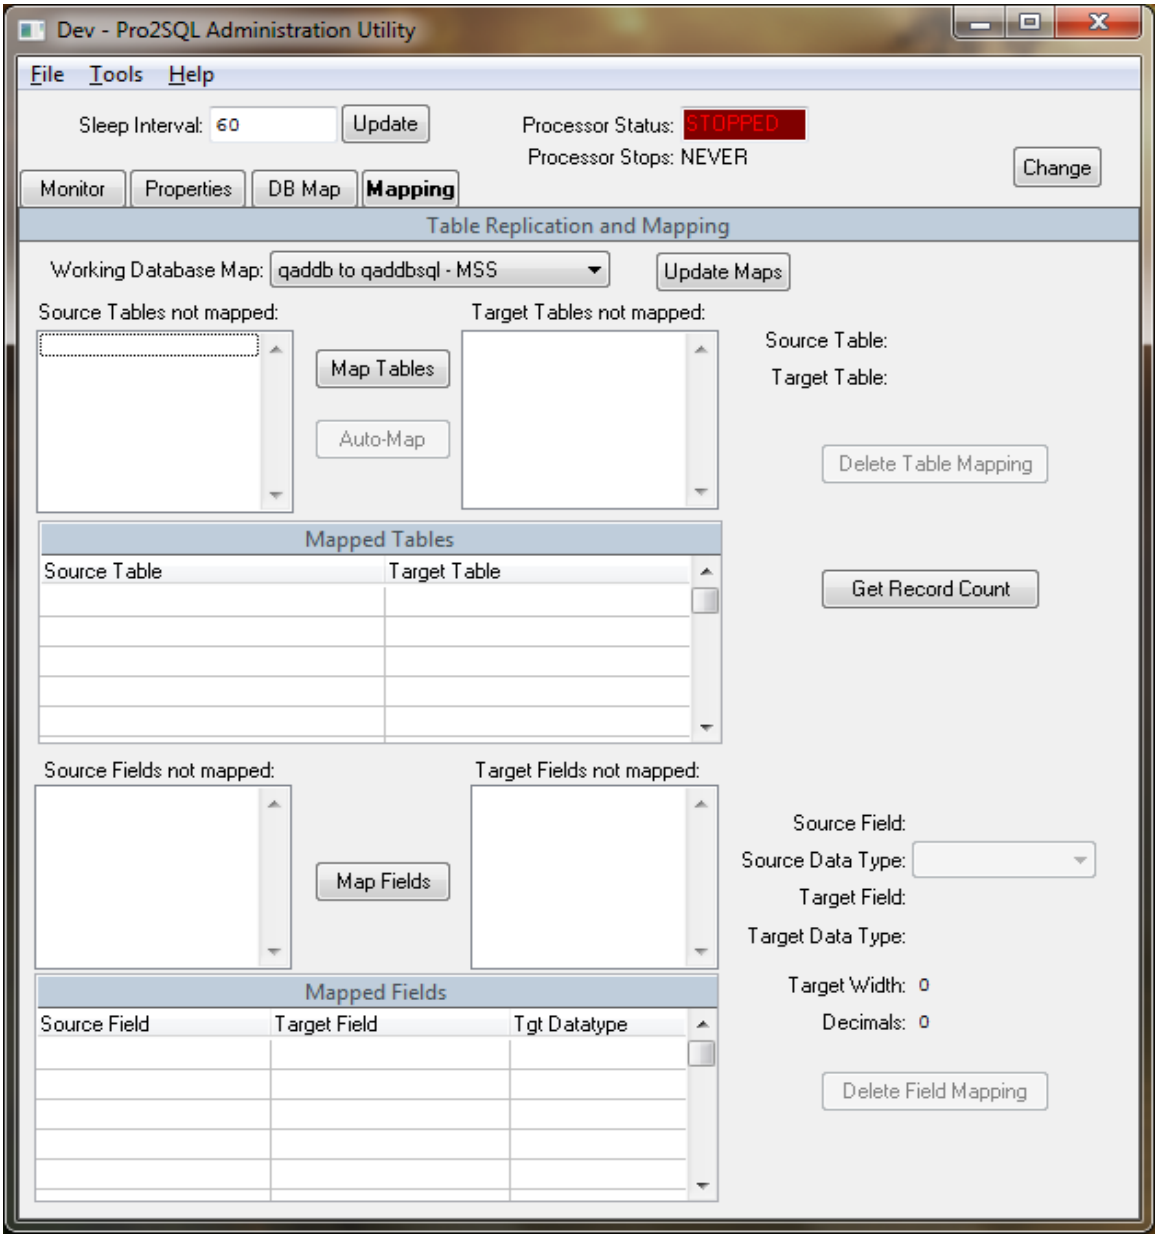

Pro2* has an Auto Map feature that maps the fields from the source to the target automatically.

Mapping features include:

- Map automatically but make changes wherever needed
- Easy interface for adding/changing field maps
- Handles schema differences between SQL and PROGRESS data types

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880 www.BravePoint.com

The Pro2 Database Mapping Tool:*

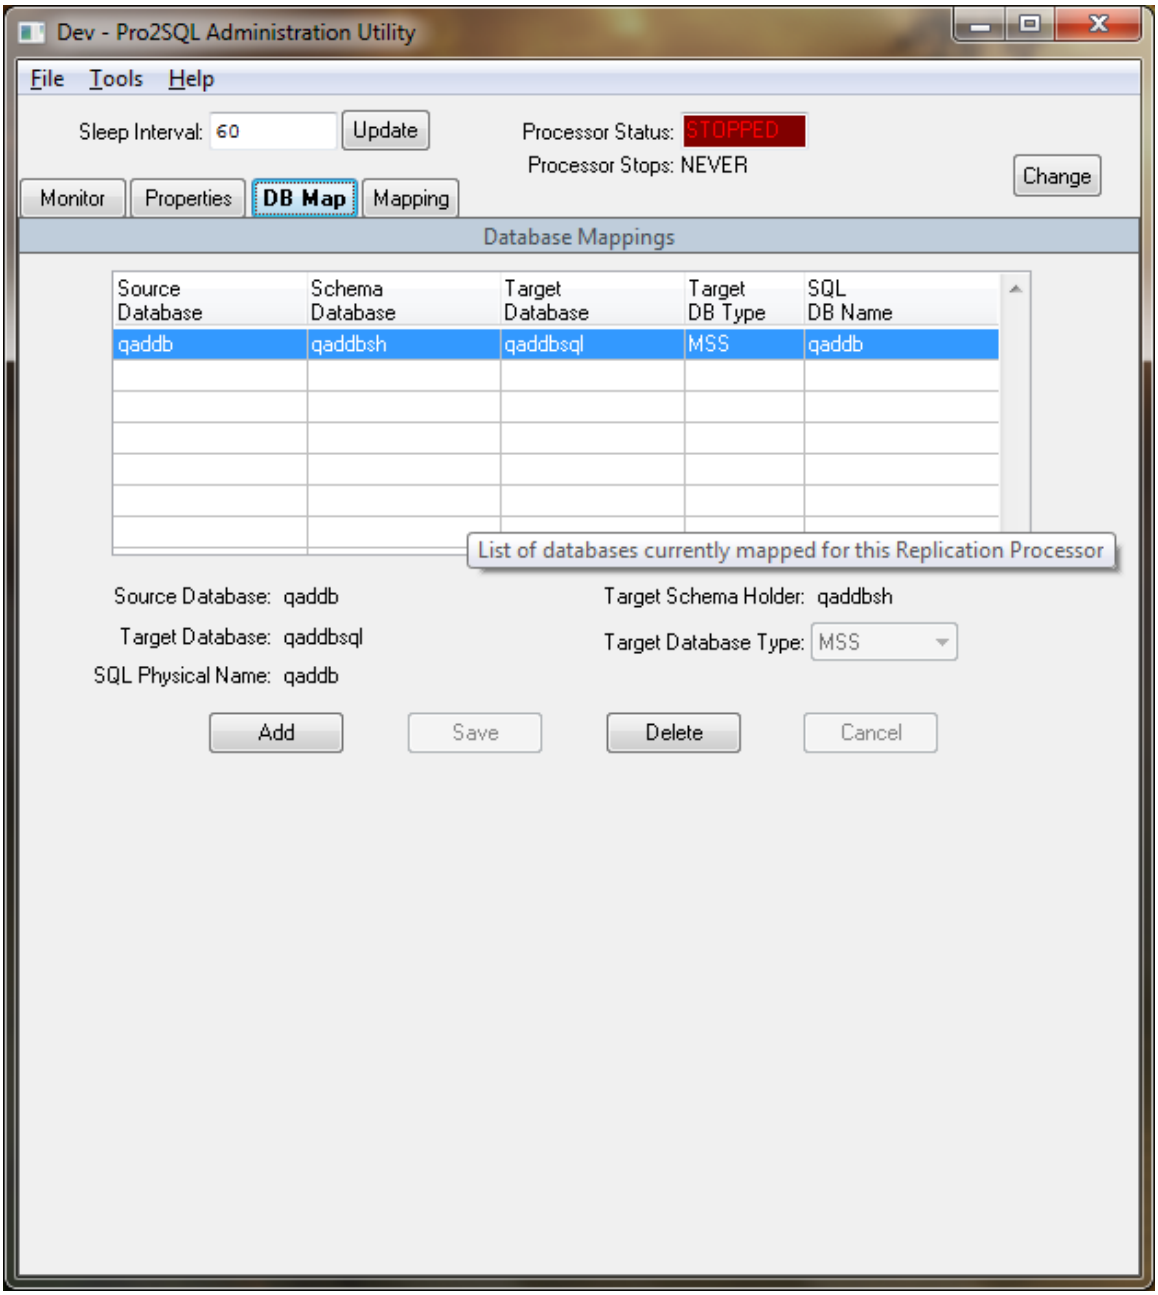

Pro2* can map multiple data sources to a single target, single sources to multiple targets and multiple sources to multiple targets using this utility.

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880
www.BravePoint.com

Implementation

Pro2* has a very small footprint. You will need to add 6 tables to an existing PROGRESS database, or create a small replication database which is used for queue functionality with minimal impact on disk I/O activity.

Implementation also includes time for training and knowledge transfer so your company can maintain Pro2*.

How it works:

- Data changes are captured with replication triggers
- A minimal amount of information is written to the queue to identify the updated record
- The multi-threaded replication process retrieves the updated record
- The data queued in the replication database is moved via the MS or Oracle Data Server to SQL Server to the target database.

Note: Because the replication process is in near real-time, the I/O operation is optimized because the updated record will still reside in cache.

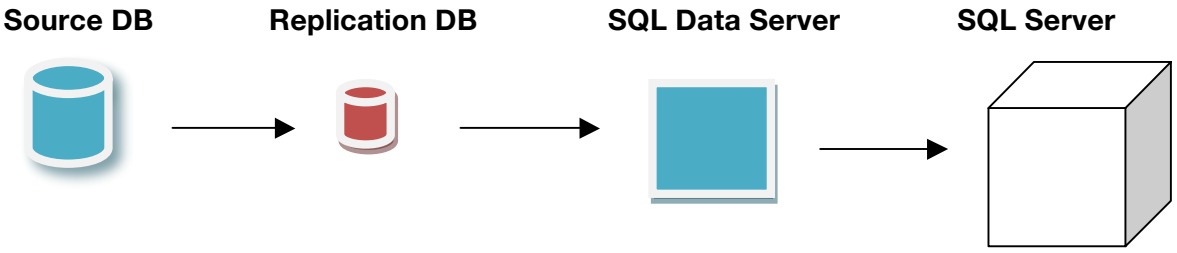

The entire implementation can be completed in as little as 3 days. The actual amount of time depends on three major factors:

1. The number of databases and amount of data that you wish to replicate
2. The number of target databases
3. The complexity of any data transformations

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880
www.BravePoint.com

Summary

The Pro2 Replication Suite makes replicating data from PROGRESS to SQL Server and Oracle easy while minimizing risks and reducing costs.

- Includes GUI administration tools
- Alleviates the obstacles for replication of PROGRESS to SQL / Oracle
- Small footprint and little impact to production applications
- Real-time and near real-time replication
- No need to “re-invent the wheel,” Pro2* is a proven technology being used by over 50 customers
- Easy to implement and maintain
- Adapts to changing environments
- Backed by BravePoint Managed DBA team
- It’s configurable, reliable and fast
- Avoids common issues like ODBC performance, fixed vs. variable-length databases objects and field / row types like logical fields and arrays

5000 Peachtree Industrial Boulevard, Suite 100 – Norcross, GA 30071
888-874-1880
www.BravePoint.com