Progress® OpenEdge®

Progress OpenEdge Pro2 Data Replication

WHITEPAPER

Introduction

Progress® OpenEdge® Pro2™ Data Replication (Pro2) is the easiest and most cost-effective approach for replicating Progress OpenEdge data to Microsoft SQL Server or the Oracle RDBMS.

Extracts no longer need to compete for resources with ABL users or nightly dump and load routines that may not work. Data is replicated in near real-time, instead of 24 hours later. The actual data format and replication time interval is totally configurable by the user. Another major advantage of the Pro2 Replication suite is that all data transformations can be completed using the OpenEdge ABL. This provides maximum flexibility and excellent performance since the solution is completely 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.

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 proposition. 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 ABL database is somewhat intrusive to online ABL users. Companies have overcome this obstacle by only dumping smaller amounts of pertinent Progress data and loading them 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.

Progress OpenEdge Pro2 was created to overcome these challenges. In a perfect world, you would like to replicate data from OpenEdge to the target database in real or near real-time. Utilizing OpenEdge 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 OpenEdge being a row-oriented, variable length database and SQL being a page-oriented, fixed length database.

The Pro2 UI Dashboard 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. Your existing staff can easily use this tool, so when changes need to be made, they can be done internally.

More detailed information can be found at progress.com/openedge/features/openedge-pro2

The Pro2 Solution

Pro2 utilizes the strengths of both the Progress OpenEdge database and OpenEdge ABL to achieve the most dependable and configurable solution for Progress to SQL replication.

Real-Time Data is Better Than Old Data

Because Pro2 utilizes replication triggers, or CDC if available, 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 ABL application.

Reduced Risk

There are no external applications like SQL Server Integration Services or task schedulers to rely on, 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 allows for flexibility to adapt to changing requirements and save costs for development efforts and chasing down production issues. Pro2 works with single or multiple sources and targets regardless of where they reside and can also use multiple threads to guarantee the best possible performance when replicating. Pro2 supports all ABL data types: logical fields, arrays, dates, etc. The Pro2 UI Dashboard is a robust tool that allows for easy maintenance and enhancements.

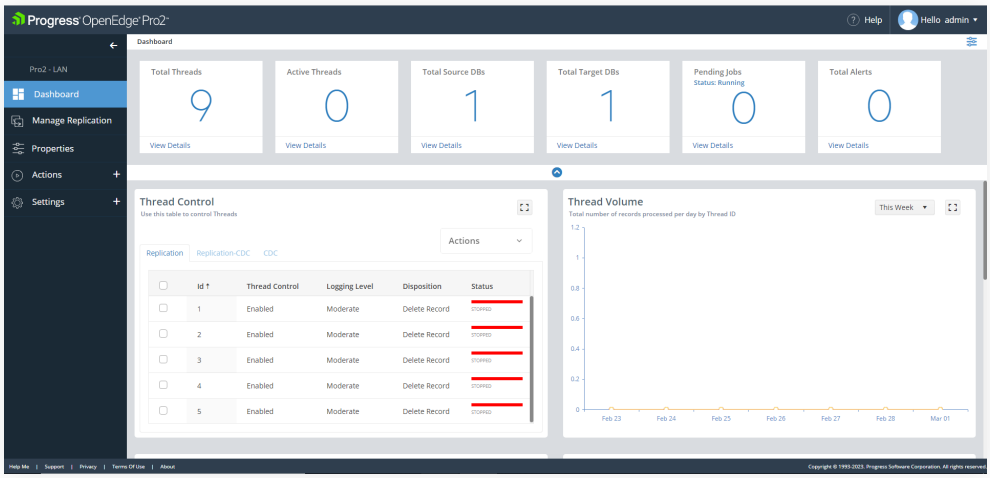

Manage Replication (DB mapping page) is responsible for creating a new replication set up, managing the existing replication setup, and deleting the existing replication setup.

Creating a new replication process consists of the following steps:

- Select Source
- Get Source side schema (WAN configuration only)
- Select Target
- Generate Target Schema
- Mapping
- CDC Mapping
- Advanced Configuration
- Generate Code

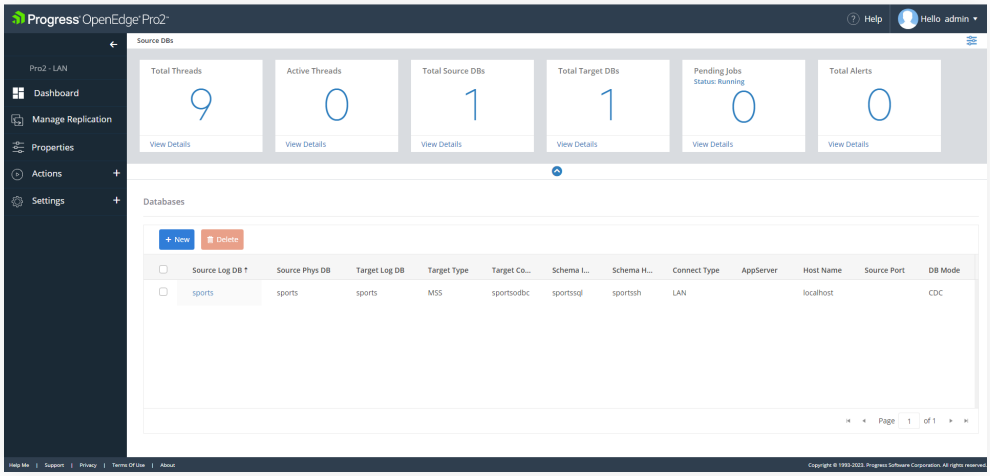

The Bulk Load Running Status pane displays the near real-time status of the ongoing/ completed bulk load. Granular data, such as the count of tables successfully processed, left to process, or failed during the bulk load process, is displayed along with the global status of the bulk load process on the bulk load status cards. The information on these cards is automatically updated every five seconds.

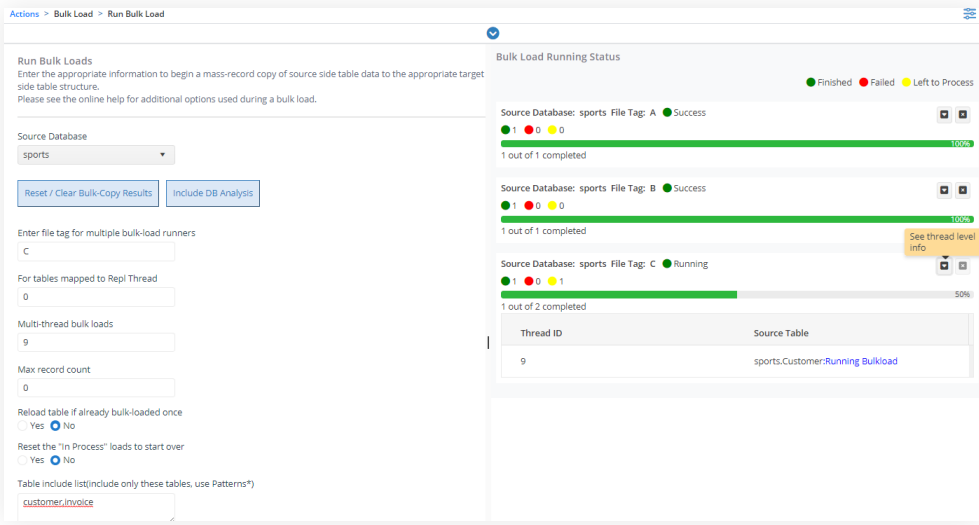

Implementation

Pro2 has a very small footprint. You will need to add ten 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, or via the database change data capture feature, when applicable
- 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 the target database

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

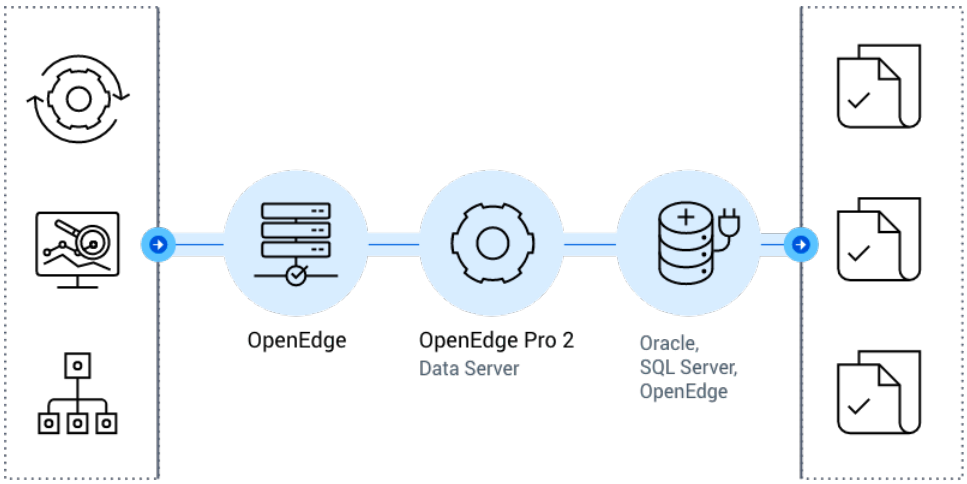

The entire implementation is quick and painless; time is dependent on four major factors:

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

Summary

The Pro2 replication suite makes replicating data from Progress to SQL Server and Oracle easy while minimizing risks and reducing costs.

- Includes WebUI 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 used by over 250 customers with over 600 installations
- Easy to implement and maintain
- Adapts to changing environments
- Backed by a Progress Services professional Pro2 team
- Its 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

More detailed information can be found at https://www.progress.com/openedge/features/openedge-pro2

Learn more

About Progress

Dedicated to propelling business forward in a technology-driven world, Progress (NASDAQ: PRGS) helps businesses drive faster cycles of innovation, fuel momentum and accelerate their path to success. As the trusted provider of the best products to develop, deploy and manage high-impact applications, Progress enables customers to build the applications and experiences they need, deploy where and how they want and manage it all safely and securely. Hundreds of thousands of enterprises, including 1,700 software companies and 3.5 million developers, depend on Progress to achieve their goals—with confidence. Learn more at www.progress.com

Rev 2023/08 RITM

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA01803, USA
Tel: +1-800-477-6473

facebook.com/progresssw
twitter.com/progresssw
youtube.com/progresssw
linkedin.com/company/progress-software
progress_sw_

Progress