OpenEdge® Replication or After Imaging?

That is The Question

Data Management Group

PROGRESS SOFTWARE

OpenEdge® Replication or AI-based Replication? That is the question.

Shakespeare's question "To be, or not to be?" is considerably more famous than ours. However, like Shakespeare's question, ours also poses a stark choice, and the answer can mean the difference between life and death for a company confronted with "the slings and arrows of outrageous fortune".

Data replication has two major attractions. It can distribute copies of information to one or more sites. It can also provide failure recovery to keep data constantly available to users, customers, and partners.

Increasingly, disaster recovery and 24/7 operation is expected and demanded by computer users. As a result, the latter attraction of replication is currently driving replication implementation in systems of all types and sizes.

Here, we outline two approaches to replication which are available to OpenEdge DB administrators. Both are widely used, and each has its own advantages and disadvantages. The intention is to analyze the techniques and technologies, in order to assist you in deciding which is better suited to your particular situation.

This document does not attempt to sell a replication approach. Nor is it biased towards either technology. The intention is to be pragmatic, the better to help you to make the right decision. Your decision should be based on a thorough understanding of what both the business and IT requirements are. We leave the balancing of these considerations entirely in your hands.

Neither is licensing covered in this document. This is obviously a significant factor for some businesses, but such negotiations are best undertaken with your regional sales team

What does your business need? What do YOU need?

In determining the approach to replication, certain expectations have to be established. When Progress Professional Services (PPS) are asked to implement disaster recovery measures for clients, the first thing they do is to ask the business sponsors and the on-site IT staff to quantify a few of their expectations for data recovery.

Being techies ourselves, we know that nothing delights you more than TLAs (three letter acronyms). Therefore, we are providing you with four that PPS endeavor to quantify consider prior to MTR (making the recommendation).

RTO (Recovery Time Objective)

RTO is a measure of the acceptable time it would take to be back up and running in production mode. In this category, you also have to define what "back up and running" means. You must consider the complete application when determining the RTO requirement.

If the business must have the application back up and running within 5 minutes then the RTO is "5 minutes". If the business has the capabilities to run on paper or a long-redundant legacy system for a single day, then your RTO would be "1 day". If a week or more is okay, then you might suggest that your business could save money by scrapping their computer systems altogether!

RPO (Recovery Point Objective)

RPO is a measure of how much data your business can afford to lose. RPO is generally dependant on the type of data on which transactions are being performed. Generally, the business will not be as forgiving for financial transactions and orders taken as they would be for an HR application, for example.

The initial response to this is "none". There was a time that this would be entirely unrealistic, but not so today. If you cannot afford to lose any transactions then your RPO is "zero". If you can lose one day's worth of transactions then your RPO is "1 day".

RMO (Recovery Maintenance Objective)

RMO is a measure made by IT staff. It refers to how much time notionally can be spent maintaining the solution.

Maintaining the solution involves programming, scheduling, managing, and ensuring that the solution is running correctly and effectively. It also involves performing the recovery steps if and when that is needed.

RDO (Recovery Distance Objective)

RDO is the distance your recovery site is from your production site.

If it is within "LAN distance" then there are more alternatives and opportunities. If not, then geography and speed of access come into play. This, naturally, limits the approach.

The Alternatives

Armed with these criteria, your next step is to consider the capabilities of the alternatives to see which more closely fits your requirements.

OpenEdge Replication

OpenEdge Replication provides the following features:

- Standard Progress utility. OpenEdge Replication is developed by Progress and integrates seamlessly into the database broker. Setup and configuration are clearly documented, along with the failover process and the failback process.
- Near real-time standby of the database.
- Zero Latency.
- Automated and managed failover processing.
- Simple setup and management of the replication solution.
- Provides automatic synchronization between databases. OpenEdge Replication performs an automatic synchronization process that insures that both the target and source databases are identical. Synchronization is performed during database and replication startup and during failure recovery processing.
- Provides a "plug and play" technology with no application changes required.
- Allows read-only access to standby . Enhanced Read-only Mode (ERO) allows read-only, multi-user functionality while data replication is being performed. ERO offers full database server capabilities with a buffer pool, shared buffers and read-only private buffers. This allows the standby machine to be used for query, reporting or any other read-only application, thereby also minimizing the load on the production machine.
- Bandwidth requirements known. OpenEdge Replication provides the ability to determine what type of bandwidth the network will need. Additionally, OpenEdge Replication provides a consistent level of network activity, minimizing peaks of data transfer between machines.

OpenEdge Replication poses the following challenges:

- Failing back to the source machine is a manual process that requires significant down time.
- Requires after-image management to be implemented.

OpenEdge Replication Scorecard

RTO	**Minimal** The database must be shut down to transition it to a production database.
RPO	**No data loss** Replication will lose any in memory transactions that have not been sent over the wire to the target database. These can be recovered by using the last AI file to validate that changes have made it across to the target database.
RMO	**Moderate** The replication process itself requires very little maintenance. Maintenance and management of AI files is still required, but the applying of AI files to the target database does not have to be performed.
RDO	**Extensible** Depending on transaction load FR can handle running over a WAN. There are customers that have no problems in this type of configuration. If you are trying to stuff 1 gigabyte worth of information per hour over a 256K line then it will not work for you.

After-Image–based Replication

After-Image Based Replication offers the following features:

- Provides a "plug and play" technology with no application changes required.

After-Image Based Replication poses the following challenges:

- AI-based replication is a complete custom solution. Implementation is on a database by database basis.
- AI-based replication is also known as After Image roll forward or “warm” standby replication. The standby database is not a "hot" standby.
- The database must be down in order to apply After Image files. Therefore, there is a down time involved in preparing the standby database for use. If your standby cannot be reached on the LAN, travel time the standby site might also become a significant factor.
- AI scripts must be manually written and managed.
- Scripts do not provide an automatic failover capability.
- Does not allow read access to the standby database while scripts are being applied – you must bring the database down in order to apply the AI files.
- AI scripts do not provide automatic synchronization between databases.
- Failback processing must also be customized and managed.

AI-based Replication Scorecard

RTO	**Medium** The last After Image file(s) must be rolled forward against the target database before it can be brought up in production.
RPO	**No data loss** After-imaging will lose any in-memory transactions that have not been rolled against the target database. These can be recovered by rolling the last AI file(s) forward against the target database.
RMO	**High** Maintenance and management of AI files is required. This must occur both on the production and on the target database. Transferring of the files between the machines must also be managed, as well as validation of the AI files on the target machine.
RDO	**Extensible** Depending on the RTO, the files could be transferred via FTP or sneaker net. This would make it possible for electronic transfers during off-peak hours. However, if, for example, you are trying to transfer 1 gigabyte of information per hour over a 256K line, then it will not work for you.

Conclusion

As you can see from the scorecards for each option, the main differentiators are the respective recovery time and recovery maintenance objectives. OpenEdge Replication outperforms “manual” AI on both counts. The additional benefit of a read-only DB for reporting, etc has also been a strong argument in favor of OpenEdge Replication in many cases.

This does not, of course, indicate that OpenEdge Replication is the ideal solution – otherwise OpenEdge Replication would be a “no brainer”, and this document would be essentially pointless. Rather, we would encourage you to use the information provided to carefully weigh up the pros and cons of each solution with respect to your particular situation.

Corporate and North American Headquarters
Progress Software Corporation, 14 Oak Park, Bedford, MA 01730 USA Tel: 781 280 4000 Fax: 781 280 4095

Europe/Middle East/Africa Headquarters
Progress Software Europe B.V. Schorpioenstraat 67 3067 GG Rotterdam, The Netherlands Tel: 31 10 286 5700 Fax: 31 10 286 5777

Latin American Headquarters
Progress Software Corporation, 2255 Glades Road, One Boca Place, Suite 300 E, Boca Raton, FL 33431 USA Tel: 561 998 2244 Fax: 561 998 1573

Asia/Pacific Headquarters
Progress Software Pty. Ltd., Level 2, 25 Ryde Road, Pymble, NSW 2073, Australia Tel: 61 2 9496 8439 Fax: 61 2 9498 7498

PROGRESS SOFTWARE

www.progress.com