Rocket® software

How Much Data Loss Can You Tolerate?

Contents

Introduction

If your system recovery strategy relies on infrastructure-level replication, snapshots, or file sync, this is a must-read for you.

We all want to believe that our SANs, virtualization systems, or OS-level utilities will provide sufficient protection against data loss. The literature suggests it, but if you dive deeper into how these systems work, you'll discover a subtle yet critical weakness.

If you think these solutions fully cover you, you're not alone. Nearly 70% of IT leadership depends on this style of protection, mistakenly believing that it will provide the ability to recover their systems and data it ensures with little or no data loss.

However, the reality is that while these systems do provide a degree of protection, they are fundamentally incapable of meeting the Recovery Point Objective (RPO) standards that a modern computing environment requires.

Most databases overcome this limitation by using a database-specific protection mechanism known as logging. Database logging systems safeguard database tables against data loss and facilitate database-level replication across multiple servers. For example, Microsoft calls their logging system for SQL Server "Transaction Logging", and the use of these transaction logs for replication is known as "Log Shipping."

The Rocket® jBASE® logging system is called "Transaction Journaling." Similar to Microsoft's SQL Server Transaction Logging, it addresses data loss risks that non-database-aware protection mechanisms cannot address. The use of jBASE Transaction Journals for replication is called "jBASE Replication." Implementation of jBASE Transaction Journaling can be easily implemented with no change to your applications.

Back to contents

Data loss risks

Years ago, system backups or “File Saves” were the primary mechanisms to protect against data loss. Due to the slower system and backup device speeds at that time, most organizations performed nightly backups. This meant that any transactions made between the end of one backup and the start of the next would be lost if a system failure occurred. Since backups were typically performed once a day, organizations had an RPO of 24 hours. In simpler terms, if your system crashed 23 hours after the last backup, you'd lose 23 hours worth of data.

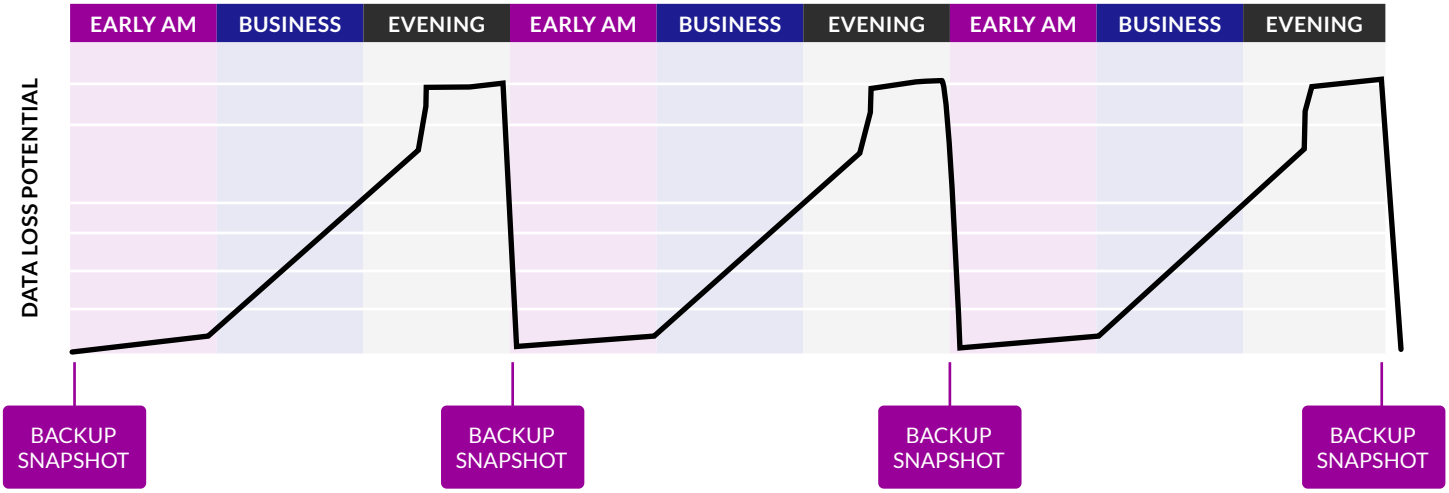

The above table illustrates the risk of data loss over time, assuming the system undergoes a daily backup around midnight. In this example, backups and snapshots are the same. If backups or snapshots occurred more than once per day, the volume of lost data would be reduced proportionally.

Reading the plot from left to right:

1. Right after a backup is completed, there is no potential for data loss.*
2. During the day, transaction volume grows steadily, and the volume of unrecoverable transactions increases. This is shown as an upward curve.
3. In the evening, end-of-day processing might lead to a spike in transaction volume, further increasing the potential for data loss.
4. Just before the completion of the next backup, the risk of data loss is at its highest. If a system failure occurs before this backup completes, all the transactions since the last completed backup will be lost.
5. Once the next backup is completed, the potential for data loss drops, and the cycle repeats.

* This doesn’t consider that if your backup is based on infrastructure level snapshots, your snapshot may contain dirty disk blocks, resulting in “broken” files/tables when recovered. You won’t know this until you attempt a recovery. Furthermore, many organizations will test snapshot recovery by bringing down the server prior to snapshot, thus ensuring a ‘clean’ backup. The real world seldom affords us the opportunity to bring the system down to perform clean snapshots. Therefore, the ‘clean room’ tests often give us a false sense of safety that our recovery methods have been tested, when in truth, they have only been testing under ideal conditions which seldom occur under real operating conditions.

Back to contents

Tolerance for data loss

In the past, losing up to 24 hours of data might have been tolerable. However, today, any amount of loss is intolerable. Now, the goal is to measure acceptable data loss in minutes or moments, certainly not hours or days.

Is more frequent snapshots the solution?

Many organizations attempt to minimize potential data loss by increasing the frequency of backups or snapshots. On the surface, this seems like an effective strategy, but there are hidden costs to this approach.

The most notable issue is that frequent snapshots or backups can significantly impact system performance, which users find intolerable. An effective snapshot should be taken when the system has been quiesced, ensuring there are no dirty disk blocks in the snapshot image. However, achieving this quiet state requires additional time and can further disrupt users, who typically have little patience for such interruptions.

Because of these limitations, snapshots cannot be taken frequently enough to reduce the volume of potential data loss to an acceptable level. Snapshotting can also create a false sense of security; due to the dirty block conundrum, snapshots may or may not be a reliable recovery point.

Introducing Transaction Logging

Rocket jBASE Transaction Journaling (TJ) is a record-level journaling system that provides enterprise recoverability for low RPO (Recovery Point Objective) requirements. It differs from File System, SAN, or VM-based replication strategies because, instead of being based on disk blocks, it's based on entire database records.

Like Microsoft's Transaction Logging for SQL Server, jBASE TJ addresses the shortfalls in disk-level or VM-based snapshot or replication strategies. With jBASE Transaction Journaling, you can reduce your RPO from hours to moments, depending on your transaction volume and system speed. This significantly reduces your exposure to data loss.

jBASE TJ is superior to disk or VM-level replication strategies in several additional areas that are vital for those seeking to substantially reduce the risk of data loss inherent in conventional SAN, File System, or VM-based strategies.

Here's why jBASE TJ stands out:

01 Allows for the recovery of data up to the moment of failure, rather than to the point of the last backup or snapshot.

02 Enables data replication using a technique that is not susceptible to "dirty blocks," which plague SAN or infrastructure snapshot strategies, potentially leading to an unusable replication target.

Back to contents

Minimizing overhead

jBASE Transaction Journaling is built to be efficient and fast. While it does require the system to not only update a database record but to also write that record into the transaction journal. In the past, this double write could have had a perceivable impact on performance. However, with today's high-performance disk technology and proper configuration, you can set up Transaction Journaling with minimal impact on your system.

Straightforward implementation

The implementation of jBASE Transaction Journaling doesn't require changes to application programs or data files. **Here's how to get started:**

1. Install jBASE Transaction Journaling
2. Allocate fast disk space to hold transaction journal files
3. Configure TJ to target specific accounts and files
4. Exclude unnecessary files (work files, transient files, spooler output, etc.)
5. Understand and document data recovery procedures using TJ and train staff
6. Turn it on and monitor
7. Practice a recovery to become familiar

Optionally, you can "play" the transaction logs to another system where they will be consumed, and the transaction updates will be applied to the second system. This is particularly useful in High Availability (HA) and Disaster Recovery (DR) scenarios, where reducing data loss in failover scenarios is very important.

Discover more

Our Rocket jBASE team is here to help you explore how Transaction Journaling can benefit your organization. Please contact your Rocket sales or partner representative and let's create a simple plan to better protect your organization from data loss.

Contact an expert

About Rocket Software

Rocket Software is the global technology leader in modernization and partner of choice that empowers the world's leading businesses on their modernization journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,000 global employees, Rocket Software enables customers to maximize their data, applications, and infrastructure to deliver critical services that power our modern world. Rocket Software is a privately held U.S. corporation headquartered in the Boston area with centers of excellence strategically located around the world. Rocket Software is a portfolio company of Bain Capital Private Equity. Follow Rocket Software on LinkedIn and X (formerly Twitter) or visit RocketSoftware.com to learn more.

Modernization. Without Disruption.™

Visit RocketSoftware.com ›

Rocket software

MAR-10583_RS-MultiValue_HADR_whitepaper_v3