

A One-Time Solution for the Shrinking Mainframe Batch Window

By Allan Zander

Mainframe Batch Window Solutions | www.dkl.com/fst

Industry Issue: Working within the Batch Window

The Importance of Batch Processing

Time is of the essence for mainframe operators in the banking, credit, and data processing industries. Much of their processing is time sensitive and a big portion of it is done in the off-hours, either at night or on weekends. In the mainframe world, this period of time is known as “the batch window.” Because online systems are frequently dependent on the completion of the batch window before they can be ready for customer transactions, it is vital to keep batch processing times within this window. Thus, the impact on an application’s batch processing time is a major consideration when adding new functionality, introducing new products, updating customer or product accounts, creating new reports, and meeting new financial requirements.



The Shrinking Batch Window

Many IT organizations have problems getting all of their batch processing done during their off-hours period or during the weekend. The scenario goes like this: During the day, client transactions are processed with results being saved on disk for final processing at the end of business.

This scenario worked very well for many companies for years. However, with increased demand put on mainframe systems from the explosive growth of online and web-based transactions, and 24x7 operations, IT organizations have had difficulty finishing their batch processing before the start of business the next day—their batch windows are filling up.

Worse, most financial institutions operate under SLAs (Service Level Agreements) that stipulate that batch processing be completed in a given time frame—with financial penalties for SLA non-compliance.

What can organizations do about this? An entire industry has grown around batch monitoring and tuning—these solutions can help, but only to a point—eventually, the batch window fills up no matter how much tuning you do. Also, they often rely heavily on technical personnel and/or consultants, and can be exceedingly costly in the long run.

What else? You can radically modify your applications, or go for a costly hardware upgrade, but these are extremely expensive solutions, and don't necessarily address the issue. You can switch over completely to a distributed system solution, but this requires tremendous architectural and management changes. In addition, large IT organizations in the financial sector often want to retain their mainframe infrastructure for its inherently superior throughput performance, reliability, scalability, and especially security.

There is another option that is available now, and does not bring with it the drawbacks that come with these other solutions: In-memory optimization.

In-Memory Optimization of the Batch Window

Large, successful organizations employ in-memory performance improvement solutions because they enable significantly improved batch processing performance, without the need for hardware upgrades, and without the need for extra monitoring and tuning. It is a one-time fix that does not require increased ongoing personnel expense. It also allows as much as an order of magnitude improvement in batch processing time.

This makes it easier for financial institutions to address market issues, industry changes, and reporting requirements. They no longer have to be overly concerned with new workloads adversely affecting their batch window, because their applications are running at optimum performance. Exactly how can in-memory optimization deliver these benefits for mainframe operators? By replacing repetitive VSAM calls and repetitive SQL calls to DB/2 with more efficient calls to data stored in-memory.

Here is an example of how an in-memory solution was implemented by a DataKinetics customer.

Reducing Disk Access Calls

In-memory optimization turned out to be the perfect solution at the perfect time for the IT organization of a large bank on the west coast of the US. Using their current applications and the

same mainframe hardware as before, they were able to dramatically improve processing times within their batch windows.

The bank was finding that its increased customer activity was putting a heavy burden on its data processing systems. The bank's evening batch processing took almost 7½ hours to complete. The available capacity remaining in their batch window was measured in minutes, so they had to limit the time that they could accept files from their other locations.

This cost both the bank and their customers a day's interest if the files didn't make the cutoff time and had to wait an additional day for posting. The reason for the lengthy processing time was the method used for accessing files during processing. Transactions were pulled from disk files and summarized. The summarized transaction files were then compared with the master files one account at a time, resulting in millions of I/Os a day to complete the transaction posting. See Figure 1.

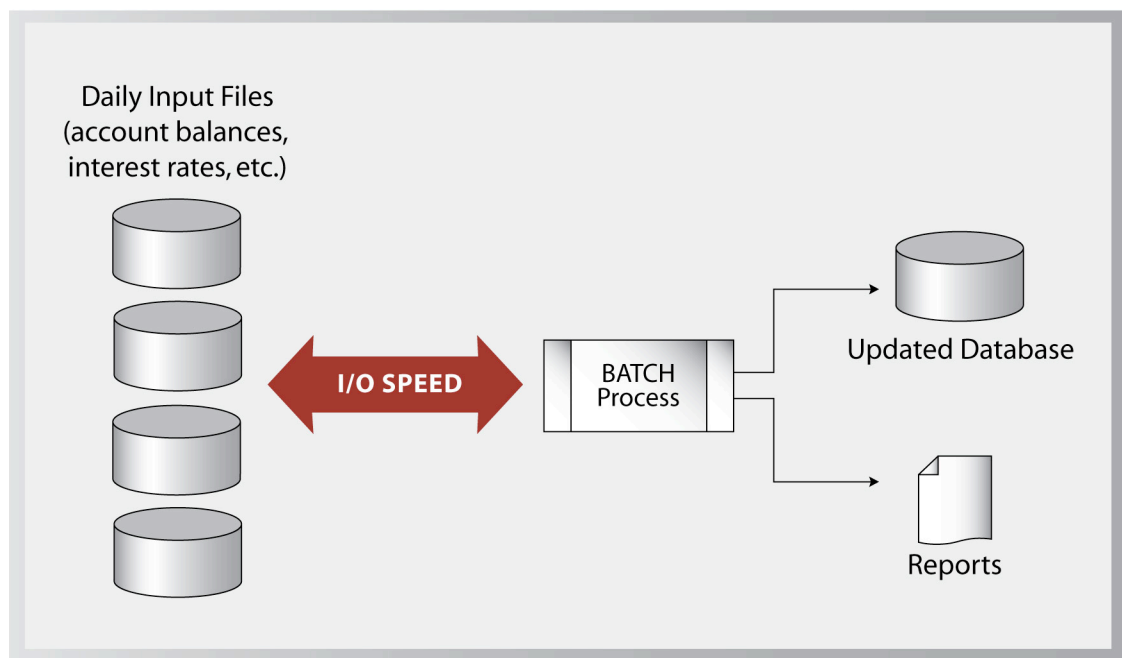


Figure 1: Data stored in disk files, accessed at I/O speed

The first solution they tried was hardware-related; they used a dedicated disk pack on the most powerful mainframe they had—the result was unimpressive. The possibility of redesigning the software was not an option, because it was too risky, expensive, and time consuming.

Their conclusion then was to use DataKinetics tableBASE® in conjunction with their current software package (see Figure 2). tableBASE exploits the mainframe's architecture by storing and manipulating data in resident memory, thus reducing the I/O calls

and processing time. VSAM does not provide such a feature, and while DB/2 can be cached, it is not optimized for fast access.

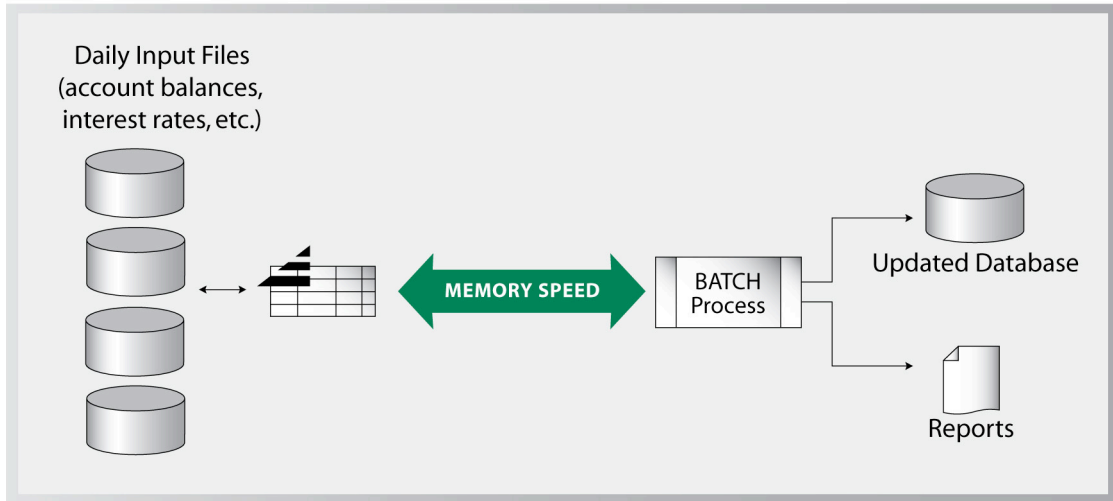


Figure 2: Data stored in-memory, accessed at memory speed

The result: Elapsed processing time was cut dramatically. The nightly process now took 15 minutes, down from 7½ hours and saved a tremendous amount of money and improved customer satisfaction. Equally important was that the I/O-to-memory conversion was made without changing any application logic.

A Well-Established Standard

For years, DataKinetics has been providing mainframe batch window performance improvement solutions to the largest banks and financial services companies running mission-critical applications on mainframe systems. A proven performance optimization technology, DataKinetics tableBASE provides rapid return on investment, coupled with long-term cost savings.

More information is available at www.dkl.com/fst.