



**White Paper**  
**CSI's V-Share - Overcoming the VSAM Bottleneck**

By: John Mycroft  
Office of the CTO  
CSI International

For years, VSAM has been the backbone of many application systems. Properly implemented, it is efficient, flexible and easy to use. It can share files between applications, with certain limitations, and can let your files grow with no effort on your part by inserting records and allocating extra extents. Did you notice the “with certain limitations” in that last sentence? Those limitations are VSAM’s Achilles’ heel, and have probably caused more database conversions and platform changes than were really necessary.

## Sharing datasets with VSAM

One of VSAM’s strengths is its ability to share datasets between jobs – “with certain limitations”. If you’ve ever searched for an example of “there’s no such thing as a free lunch,” VSAM’s share options provide it. VSAM offers four ways to share data via the SHAREOPTIONS parameter on the dataset’s definition.

**Option 1** allows multiple simultaneous readers or one writer for a dataset. Integrity is guaranteed, the dataset is sharable and performance is superb. Unless you are the one guy who wants to write to the dataset and someone else has it open for reading.

**Option 2** opens the door a bit further by allowing multiple readers and one writer. Again, integrity and performance and altogether lovely. Unless you are the second guy who wants to write to the dataset. So it is tempting to go for ...

**Option 3** which allows for any number of readers and writers at the same time. This option sometimes gets chosen while ignoring what the IBM VSAM manual says “VSAM does nothing to ensure either the integrity of information written in the file or the integrity of the data retrieved from the file.” Meaning it’s your gun and it’s your foot.

**Option 4** allows any number of jobs in the same system to have full read and write access. For which there is a hefty performance price to be paid.

## The problem grows

Long ago, it was feasible to wait to update a dataset until CICS was down. But as business needs dictated longer network hours, the batch window got shorter and desperate measures were called for. In many cases, VSAM was reluctantly abandoned in favor of a database system or even the nightmare of conversion to a different platform to provide access to critical data at any time.

## Stop-gap measures to avoid the problem

Many stop-gap measures were taken to satisfy the need to access data both online and in batch. Online systems were shut down regularly to allow batch updates to run. CICS update transactions were disabled for a while or VSAM files were closed while batch work ran. Disruption to the business was, naturally, unacceptable and a number of these “solutions” created problems of their own, such as the CICS transaction that successfully updated one dataset before failing because another dataset was closed, leaving half the transaction done. Use of “shadow” files was tried but failed because of difficulties in keeping the files truly in step and because of the heavy I/O load and hence sluggish performance they produced.

## **A successful solution**

What was really needed was a “traffic cop” – a file server to handle all of the I/O to the datasets needed both in CICS and batch at the same time, without impacting the availability or performance of either. It needed to benefit from and build on an organization’s significant investment in VSAM rather than call for a costly move to a different platform for no other reason than to cram more hours into the day. CSI’s V-Share satisfies these needs by acting as a database manager providing full read/write integrity by routing all file I/O, both for batch and for online, through a central point within CICS.

This brings other immediate and lasting benefits to the organization. Firstly, no program changes are required – there is nothing for programmers to do to make use of V-Share’s file sharing. It also brings a tiered recovery solution that journals all updates and automatically backs out any updates made by an abending task at the field level. So a program failure causes no data loss to either batch or online work. Naturally, operations support people have full control over which jobs can use V-Share and which datasets it is to manage. V-Share interacts with security software to make sure that there are no back doors into the system and performance can be controlled dynamically so that, for example, a heavy batch job can be throttled back or suspended to make sure that the online system gets top priority.

## **Full high-performance file sharing**

As mentioned earlier, VSAM’s Share Option 3 is an accident looking for somewhere to happen (and I have seen that accident happen spectacularly at least twice, and they were just the ones that got noticed). V-Share provides exactly the file sharing mechanism that is needed with Share Option 3 (as well as working with the other share options, of course). So it becomes possible to share datasets with full integrity between jobs. Even better, because each dataset has a single set of data and index buffers under CICS’s control, all the performance gains associated with large BUFND and BUFNI allocations are immediately available to all the jobs sharing that dataset. This contrasts sharply with the constricting performance limitation placed on a dataset by using Share Option 4.

## **Cross-memory and cross-system support**

Besides the usual performance-draining effects of sharing a dataset between jobs, VSAM has always been a bottleneck in companies that run multiple processors, as the need to share data between processors has meant that VSAM constantly had to reserve data paths and refresh buffers to make sure it was dealing with current records.

V-Share provides a single point of control for each dataset it manages. Records are sent between CICS and batch jobs in the same processor via XPCC and between remote processors via LU 6.2, removing the need for mirroring data across data centers or across domains in the same data center.

## **Conclusion**

CSI’s V-Share solves a long-standing problem for VSE and z/OS data centers with a high-performance and secure tool that requires virtually no work to install. It can save the expense of a costly database or platform conversion to solve a business need for higher availability and more responsive batch systems.