
Fact Sheet

What does VELOCIRAPTOR (VR) do?

VR is an intuitive performance optimization solution designed to dramatically improve application performance by significantly reducing the resources needed for I/O processing.

VR reduces application overhead for both VSAM and non-VSAM files by dynamically implementing buffering strategies that provide the optimal execution environment for the access method.

VSAM data sets

Each time an application opens a data set and requests access to it, VR builds a custom buffering environment (using either NSR (Non-Shared Resources) or LSR (Local Shared Resources)) that provides the VSAM access method with the optimal buffering scheme for that particular access to the data set. The buffering environment provided by VR results in a dramatic reduction in the EXCP activity incurred by VSAM on behalf of the application, along with a decrease in the associated CPU consumption.

VR's solution is unique in that it is based on a real-time analysis of the application and data set each time the data set is accessed by the application, thus allowing VR to dynamically respond to application and data set changes – as they occur.

VR supports all current VSAM macro formats, and it intelligently uses them in conjunction with the selected buffering strategy (NSR/LSR) to accelerate application processing by reducing all unnecessary activity on the part of the access method. When needed, VR can exploit the use of all 256 LSR buffer pools. VR also provides VSCR (Virtual Storage Constraint Relief) by placing VSAM control blocks and buffers above the 16MB line.

Non-VSAM data sets

VR ensures optimal tuning and buffer allocation by selecting a buffer count that will optimize the use of SAM/E chained scheduling for asynchronous sequential access (QSAM).

How does buffer optimization help reduce application and system overhead?

The reduction in overhead is directly proportional to the reduction in the resources (EXCP and CPU) required by VSAM to process the application's request.

An inefficient buffering environment constrains VSAM and causes it to incur EXCP and CPU overhead as it attempts to deal with those constraints. When application files are optimally buffered by VELOCI-Raptor, those constraints are removed and VSAM is allowed to execute as efficiently as possible with its EXCP and CPU requirements being reduced to a minimum.

This reduction in EXCP and CPU consumption can be brought about only by an intelligent optimization of the application's buffering environment.

The resources that are reclaimed by VELOCI-Raptor are then available for other applications. In some cases, a reclamation of these resources has allowed users to defer CPU upgrades, thereby saving both time and money.

What are the primary challenges that VELOCI-RAPTOR addresses?

The primary challenges faced by data centers are usually related to the following:

- The need to reduce batch processing time.
- The need to meet service level agreements.
- The need to maximize system performance without extensive programming changes.
- The need to determine (quantify) system-wide impact of application tuning activities.

VR addresses these challenges by:

- Eliminating CPU time that is wasted during I/O operations.
- Making application run time more consistent, thus ensuring that SLAs are more easily met.
- Automatically reducing system overhead and batch window constraints.
- Eliminating manual tuning efforts.
- Providing real-time monitoring of the tuning efforts.
- Allowing users to defer or eliminate CPU upgrades.

What are the primary benefits of VELOCI-RAPTOR?

- **Reduce batch window processing constraints.** The global reduction in EXCP and CPU resource consumption by VSAM (when VR provides it with the optimal buffering platform) results in a reduction not only in CPU time for the applications but also in elapsed time. The reduction in elapsed time can have a considerable impact on the batch “window.” In some cases, application elapsed time can be reduced by 90%.
- **Eliminate manual tuning efforts.** Technical resources that might otherwise be utilized for manual tuning can be reassigned elsewhere without any loss in performance.
- **Eliminate or defer the need for CPU upgrades.** Fluid application processing reduces CPU requirements while making better use of main memory, thereby extending the life of the CPU and offsetting the cost of expensive upgrades.
- **Meet or exceed SLA expectations.** The result of the reduction in batch overhead positively influences the service provided to users. For example, when batch window constraints are removed, on-line systems come up on time, often with improved application response times due to the reduction in system resources consumed by batch processing.
- **Obtain performance improvements without JCL or application changes.** Manual performance tuning requires significant JCL and/or application code changes, while VR requires neither.
- **Eliminate monitoring of JCL changes.** Once manual changes are made to JCL to implement tuning for a specific data set, those changes must be monitored (as the personality of the data set and access to the data set may change over time). VR’s dynamic processing replaces this monitoring.
- **Make dynamic changes in file processing.** Since VR dynamically adapts to changes in both application behavior and file status, applications can now be modified without having to perform any corresponding update to tuning efforts.
- **Expand virtual storage.** LSR and NSR buffers are built above the 16MB line to alleviate any virtual storage constraints. For LSR, expanded storage using hyperspace buffering is automatically and transparently used when the file size is larger than the LSR pool. This provides up to 2GB of virtual storage for buffers and gives ample room below the line to run larger applications.

How would I use VELOCI-RAPTOR?

VR is simply and transparently installed on any z/OS system. VR’s ISPF interface then allows users to easily define optimization candidates to VR.

Once candidates are defined to VR, it monitors batch requests to open (access) data sets, looking for requests that match the definitions. As they are encountered, they are selected for buffer optimization.

For sites that are looking to migrate from an existing tool to the latest technology offered by VR, the latest release now includes the **Conversion-Assist** feature. Typically, to run two buffering products at the same time, a user must be careful to assign candidates to only one product at a time. In many environments, the **Conversion-Assist** feature can remove that requirement and allow selected products to coexist and run simultaneously without conflict.

For a site that wants to trial VELOCI-Raptor, this means that during the initial validation period and through complete conversion, the other product can remain active and unaltered. VR takes care of ensuring that only one product processes each data set request.

For the final phase of converting entirely to VELOCI-Raptor from another product, the **Conversion-Assist** feature allows a user to perform the conversion as their schedule and plan dictates. Without **Conversion-Assist**, a large part of the conversion effort is centered upon the synchronization of the removal of candidates from one product with the addition of those candidates to the other product. With **Conversion-Assist**, that is not needed. All the user needs to do is identify the candidates to VR - VR will take care of the rest.

Additionally, if needed, upon customer provision of some basic configuration details, **Conversion-Assist** can replicate the instructions of the existing tool for a drop and swap replacement, thereby providing a swift and transparent conversion process.

How does VELOCI-RAPTOR work?

VR utilizes a proprietary data structure to manage the rules and parameters that allow it to define the optimal buffering environment for the access method. VR uses both the application declaration (when the data set is opened) and the current characteristics of the target data set as filters in selecting the applicable optimization rule(s) from that data structure. Those rules are then used to generate the needed buffer environment.

Users can override these rules and definitions by providing their own parameters or by excluding a candidate from VR's optimization process. Most users choose to use the product defaults and achieve outstanding results from VR's I/O optimization.

VR also monitors non-VSAM (QSAM and BSAM) requests to access a data set and, through a similar process, sets the optimal buffer count.

What kind of results should I expect?

VR customers routinely report that they can reduce VSAM batch processing times from 60 percent to 90 percent. Some non-VSAM batch processing times can be reduced by a significant percent. VR can also have an impact on some DB2 Unloads and Reloads.

What operating environment does VELOCI-RAPTOR require?

VR runs on all z/OS platforms and requires 30kb of common storage (CSA) below the 16MB line. The user interface requires ISPF Version 3.3 or above.

What kind of reports can VELOCI-RAPTOR generate?

Detailed statistics can be produced which show the type of VSAM buffering that was applied and the amount of data movement that occurred. Also available is the number and type of program I/O requests for VSAM files. VSAM and non-VSAM optimization particulars may also be recorded in SMF records for later analysis.

