

Improving Performance through Increased Memory

A case study of Northern Safety with Incredible Results

D. Scott Pierson

Have you ever been so surprised by an event that it caused you to put pen to paper and tell the world of your experience. Well, we just were and we thought you might find it interesting to hear the story.

Northern Safety recently embarked on a WEB delivery project to expand the availability of their entire safety product line catalog through the Smith Gardner & Associates MACS system. Following an extensive 10 month phase of testing, the Ecometry solution had been running in their environment for the past 8 months, and was now ready for the WEB enabled portion. The WEB access was due to go live in March 2001 but the project was suspended due to system performance issues.

Mrs. Donna Rivett, the HP e3000 system manager, and the WEB enablement team had seen early signs that the current 969/220 with 3.75 GB of memory and 45 GB of mirrored disk across a single Jamaica enclosure would not be adequate. During a 5-user Web test, the system experienced search times of up to two minutes, extended batch-processing times, and occasional slow interactive response times. These were just a few of the signs that the project's success was in question.

Donna had already performed an extensive investigation on her own and utilized the resources of Strategic Memory Solution's Mike Oneffrey. Finally, Lund Performance Solution's Consulting Division (LUND) was enlisted to perform a full in-depth performance analysis and capacity plan to provide an objective view and confirm their suspicions. The goal was to quantify the current environment, extrapolate the potential impact of the WEB access, and then based on the results, put in place a plan to increase the system capacity, thus creating the best potential for the project success.

LUND's capacity plan proposal recommended improvements in every resource including CPU, memory, disk, and data management. First, a processor upgrade from the current 969/220 to a 979/200 was suggested. Although this is only a 15% CPU increase, based on Hewlett Packard's published relative performance guidelines, the PA8xxx chip would allow a higher memory configuration of 8 GB's. Next, LUND would upgrade the operating system to MPE/iX 6.5, power patch 2, with additional performance and memory related patches, required to expand the Strategic Memory Solutions increase to 8 GB's of memory. Finally, an I/O expansion bay and new HVD-10 mirrored 9 GB drives would be installed to double the existing disk capacity and spindle count. Lund Consulting Services would then expand the data sets, reorganize the databases, and evenly distribute the data across all of the drives.

Following the MPE update and memory installation, we were still not ready for the full WEB test. The WEB team was still on track for a May 14-16 test following the disk expansion and database activities. This maintenance would require each of the next three weekends and would complete the final activity required to position the system for the best possible scenario for success.

With excitement, Rivett needed to test some aspect of the partial changes completed thus far including the expansion of the memory to the full 8 GB. Periodically, Rivett is required to load new customer prospect information into the MACS finder file database. It was decided that this would provide an excellent opportunity to test our changes. Robelle's SUPERTOOL product is used to extract records from an upload file. This activity has always been performed on a fully dedicated system with no other active user or batch processes. Past record loads averaged 1000 records every 5 minutes, or 3.3 records per second. A typical 75,000 record import would take a long as 6 ½ hours. A current finder file included 235, 510 records for an estimated 19+ hours.

Anxious to quantify the results of the memory expansion upgrade, we launched the SUPERTOOL extract job Sunday afternoon at 3:00 PM. Surprisingly, 3 minutes and 48 seconds later, the first status message appeared showing an amazing 5%, 10,716 records processed, or an average of 47 records per second. This was an almost unbelievable increase of 1424%. If left to continue, the job would have only required and estimated 1 hour 24 minutes to complete. It was just too good to be true. After 17 minutes, a 20% completion rate, and 45,966 records processed, the import was aborted. We restarted the job once again to confirm our suspicions

and found that the process was now processing the import at 47 records per second. If the data was accurate, it would not be completely illogical to see the increase due to the previous 45, 966 records still loaded in memory. Just over an hour later, when the process reached the 80% completion point, it aborted with a dataset capacity full message. Argh! In our excitement and disbelief of our own success, we neglected to clear the FINDER04 file prior to reinitiating the process a second time.

Finally, after clearing the data from the first two import attempts, the job was launched at 7:00 PM and was allowed to complete an amazing 1 hour and 12 minutes later for a completion rate of 54 records per second. The data was accessible, loaded correctly, and a great test of the efficiency improvements gained from large memory configurations.

The significant majority of the 1424% throughput improvement can be largely attributed to the memory expansion. First, the CPU was only increased by a relative performance factor of 1.3 or 15%. The process ran on a dedicated system with no other processes active, was performed in a single thread mode utilizing only one of the two processors (an MPE imposed design limitation), yet still consumed only 25% to 30% of the processor. For the most part, the new disk drives can also be discounted due to the fact that no information had been populated to the newly added disk. During the test, I/O was observed to only two disk drives of the ten total drives available, one of which was a new HVD-10. This was the location and drive where the new finder information was being placed. Pause for disk averaged between 3% and 10% during the extract process. Next, the update to MPE/iX 6.5 introduced both positive and negative aspects, probably canceling each other out. Many customers have experienced performance degradation when updating to the new OS version; however, the pseudo 64 bit implementation and new transaction manager tuning patch probably reduced the problem to a minimum. As of yet, no database maintenance activity has been performed. Finally, all 3rd party applications were already in a 6.5 compatible state. So, with no changes made to current versions, minimal success can be attributed to whatever 64 bit architecture improvements the 3rd party software may have already implemented in the current Northern Safety installed versions.

With all of the above resources discounted to a great degree, many inferences can be drawn about the efficiencies of large memory arrays and MPE. In this specific circumstance, we can draw attention to the operating system's ability to challenge the CPU with large amounts of data already present in memory. Due to the large amount of memory installed and the lack of other processes occupying and requesting memory resources, the data reflects the systems ability to pre-fetch larger and larger segments of user data in anticipation of the serial read and put into the application. This does not reduce the amount of I/O's required of the disk, but rather improves the efficiency of the operating system to anticipate the request of the process for the data before the data is needed. This is achieved through a feature called "pre-fetch". The more data already present in memory, the less need there is for an I/O, and thus, the more efficient a given execution process will be. With each successful serial read completed from prior pre-fetched data, the operating system is able to infer the success of its work and improve on the next I/O performed with a doubling in size of the next pre-fetched segment of data. Pause for disk is reduced as the system strives for pure uninterrupted CPU.

Obviously, there is great anticipation of the WEB testing scheduled for next month following the completion of the remainder of the system work. Based on the results of this one test, we can only rest a little bit easier as we position the WEB enablement team for the greatest probability for success. Look to next months issue as we follow-up on the conclusion of the project and other statistics that report on the successful partnership of Strategic Memory Solutions, Lund Performance Solutions, and Northern Safety.