The Industry speaks
about
Buffer Pool Tool®

Chuck Michalski, VP, Northern Trust Company

Prior to using *Buffer Pool Tool* at Northern Trust, we had already implemented the standard tuning recommendations we'd picked up from IDUG, i.e. isolate the catalog, separate pool for sortworks, separate pool for small heavily accessed objects... We were also making use of some hiperpool buffers, but probably not very effectively. The Buffer Pool Tool enabled us to take full advantage of hiperpools. The Collection Reports and Simulations greatly facilitated our analyses for the assignment of more than 1200 active objects in our Bufferpools.
Over the past year, while Northern's DB2 workload has increased by about 30%, thanks in large part to Buffer Pool Tool, we have actually reduced our RIO rate. At TNT, we estimate an average saving of 100 I/Os per second has been realized during prime time hours. For peak periods, the saving is even higher! The *Buffer Pool Tool* was an essential ingredient in our tuning effort to maximize throughput on our system.

Richard Gaunt & Patrick Morrison, INTRIA/CIBC(Canadian Imperial Bank)

Projected a cost saving from I/O elimination greater than $1,300,000 over three years by using *Buffer Pool Tool*. We eliminated 2,000,000 I/Os per day on one system, and 750,000 I/Os per day on a second system. We were also able to identify and correct application performance problems that provided substantial CPU cost savings. As presented at IDUG 1997 (Chicago, Barcelona, Sydney).

Marla Synder, Missouri Gas

Before tuning with *Buffer Pool Tool* we could not get our online system up until almost 10AM. After tuning, we are able to have the online system up by 4AM.

Chas Strong, Yellow Services

We found a major application problem during the *Buffer Pool Tool* trial, that was not shown as a problem from any other reporting or monitoring system. Fixing this problem recovered more than 15% of our processors CPU, and allowed us to cancel our processor upgrade.

Scott McLeod, Diversified Pharmaceuticals

We always suspected we had more space allocated to bufferpools than we needed, and the paging indicated we had exceeded the amount of real storage on the machine. Using *Buffer Pool Tool* simulations we were able to determine how much we could reduce the pools without impacting performance. We were able to save 350 Megabytes! The elimination of paging saved us 10 MIPs from our daily transaction load. Since we are outsourced, and our FM vendor charges us for everything we do and use, the savings of memory and MIPs are actual *bottom line dollar savings* well into the six figure range.

The Huntington National Bank, Database Administration Staff

For Data Sharing capacity planning, we had to be as accurate as possible without a production implementation. Originally we sized the coupling facility using our best overall estimate of table usage to be used in a Data Sharing environment. When we ran actual numbers from our non-Data Sharing environment through the *Buffer Pool Tool* CF Sizing component, we were able to start the sizes of the CF policies at a smaller storage usage, instead of making them larger and having to tune them down once Data Sharing was implemented.

Richard Yevich

Buffer pool performance is not good at most installations even when they think it is, and just throwing memory at it does not optimize performance - in fact it may hurt performance. The only current way to optimize pool performance and reduce application I/O wait time is with the *Buffer Pool Tool*. For installations moving to a data sharing environment, the tuning of the virtual pools and proper grouping of objects into multiple pools is a *primary requisite for*

success.