

**Contact:     Stephan Seyfert**   Decision Support  
Email:       seyfert2@uillinois.edu   2001 S. First St., MC 657  
Phone:       217.265.6532            Champaign, IL 61820  
Fax           217.265.6501                   www.ds.uillinois.edu



# Press Release

## **Decision Support beats the clock**

Changes to Data Warehouse processing significantly improve efficiency, availability, and timeliness of data

**Champaign, IL – April 17, 2008:** Kim Nystrom noticed the difference when she updated her local databases first thing in the morning. Nystrom, Database Administrator for the College of Education at the University of Illinois Urbana-Champaign, previously had to wait until mid- to late-afternoon – after the University's Enterprise Data Warehouse had finished loading new data – to update the local databases that enable her team to run current reports. That changed a month ago, and she is much happier.

"The early completion times of the EDW refreshes rock," she exclaimed. "It's truly made my data one day old instead of what ended up being two when refreshes couldn't start until 3 or 4 in the afternoon. I appreciate Decision Support investing time, energy, and money into this."

Decision Support, the business intelligence competency center for the University, manages the daily data loads for the EDW. Since its inception, DS has targeted an 8:00 a.m. load completion time to make the full data warehouse available to all users throughout their normal workday. For numerous reasons, that goal seemed elusive until a change in the processing was implemented.

"We keep the data current by processing and loading new or changed data from the source applications on a nightly basis," explained Mark Cumbow, Assistant Director of Data Architecture. "Previously, all the data went through a three-stage process of first extracting the data from the source tables, transforming it from the source tables format to the target tables format, and finally loading the formatted data to where the users can access it. In most cases, we eliminated the third step and are instead loading the data directly to the final destination as part of the second step."

The payoff was immediately noticeable. The processing changes significantly reduced the overall data processing time and freed computer resources to perform other work as well, such as allowing data processing jobs to run simultaneously. As a result, the Data Warehouse processing has

**For Immediate Release**

consistently completed before 8:00 a.m. on nearly every day since mid-March; an improvement of more than 7 hours for some tables.

The project's success resulted largely from having the right team assembled and dedicating their resources to the initiative. DS staff had worked to address the situation for years, but progress was limited by having more demands than resources, competing project priorities, and staff turnover. To make the desired achievement reachable, DS approached the situation as a specific, independent project and brought in the needed additional resources. ETL Developers from both DS and Ciber, Inc., joined forces with Database Administrators from the University's Administrative Information Technology Services (AITS) department and Richard Buck, a project management consultant from Ciber, Inc., was hired to manage the project.

"This has been one of the best consulting engagements that I have experienced," Cumbow said.

The hybrid team focused solely on improving the data load completion times and applied the processing changes in batches on nearly a weekly basis. The impact was noticed immediately, even before the 8:00 a.m. goal was reached.

"I noticed some significant improvements since the middle of January in the finance ETL times," said John Tolar, Assistant Vice President for Administrative Services in the University Office of Business and Financial Services. "I'm sure this is a great help for the finance users. It certainly is to me!"

The impact on the Data Warehouse users was what Andrea Ballinger, Assistant Vice President for Decision Support, found most rewarding.

"I was very well aware of the frustration this issue had been for our end users and for us," she said. "So it was with great satisfaction and pride that I thanked the members of the project team for a job well done; in fact, a terrific job."

Of course, with the success, greater expectations follow.

"[The team] raised the expectation for giving our end users what they need, when they need it," Ballinger said. "That is an opportunity we gladly embrace; it's a good problem to have!"

Indeed, though the project is completed and the hybrid team disassembled, continued work remains for the DS staff. With additional data constantly added to the EDW and more demand from an increasing number of users, continuing to improve the processing efficiency will be essential to maintaining the desired load completion time.

"We have a committed focus in place," Cumbow said. "We're continuing to eliminate the processing step and optimize the scheduling, and we are also working on tuning long-running

programs. Unfortunately, it's not a 'one size fits all' scenario when it comes to applying technology changes.”

In fact, DS has applied the processing change to only 509 of 700 table load processes. Additionally, the job scheduling realignment has only been applied to 77% of the identified opportunities for simultaneous activity. Finally, DS has identified the 20 longest-running programs and is analyzing each for areas to improve efficiency. To date, changes have been implemented to 8 of those 20 programs.

With room for continued improvement, the Decision Support team expects to maintain their performance objective and keep the University's data available and timely to the thousands of people on each campus who depend on it daily. That allows users like Kim Nystrom, John Tolar, and their teams to stay on top of the information and guide the University's future in a more effective and timely manner, first thing in the morning.

###