

Success Story: Allied Insurance

Allied Insurance Speeds Data Warehouse Loads and Cuts Data Mart Processing Time by up to 75% Using SyncSort

Organizational Profile

- The largest provider of insurance to small and medium sized farms in California

Business Need

- Improve overall database performance and deliver timely data access for data mart users

Environment

- DEC Alpha 4100
- Windows NT

Benefits

- Reduced processing time windows by up to 75%
- Minimized elapsed time for data warehouse tasks, including loads and data mart reporting
- Saved programming, testing and debugging time, and reduced resource usage
- Used SyncSort's parallel processing capabilities to run multiple processes at the same time

"Because of its performance and its efficient use of resources, we can run half a dozen processes at the same time and still maintain performance for each one"

Fred Treffinger
Director of Data Warehousing
Allied Insurance

The Challenge:

Allied Insurance, formerly CalFarm Insurance, the largest provider of insurance to small and medium sized farms in the state of California, decided to improve the information available within the organization, according to Fred Treffinger, who joined as director of data warehousing in 1996. Allied wanted to better understand the underlying economics of its business: "In the insurance business, you don't know how much your product costs when you sell it," says Treffinger. "You only find out several years later when loss information becomes available. Actuaries look at claims for prior years, so they can predict what future losses will be. That's how they determine reserves and appropriate pricing for current policies."

To take care of its information needs, Allied built a data warehouse that includes every premium, billing, and loss transaction in Allied's property and casualty lines. That's about 40-50 gigabytes of raw data, which turns into over 100 gigabytes when indices, summaries and downstream data marts are included. From the outset, Treffinger's idea was to give users point-and-click access to the data. That meant moving to a LAN environment.

The Solution:

After evaluating many solutions, Treffinger chose SyncSort to improve overall database performance, BrioQuery (Brio Technology, Palo Alto, CA) as the query tool and Oracle 8 as the database engine. Windows NT was chosen as the database platform because of the need for dedicated processing capability, decreased DASD and administration costs, and the targeted size for the data warehouse. A DEC Alpha 4100 was chosen as the initial hardware platform.

The Benefits:

SyncSort now helps with a wide variety of data warehouse tasks, such as loading, according to Treffinger. Allied stages data in DB2 on the mainframe with COBOL, SAS, and SyncSort. Then the data is *ftp*-ed to an NT server, where it is pre-processed with SyncSort for loading into Oracle databases. COBOL copybooks are also downloaded to NT, and SyncSort uses them to automate the process of defining data for its input files. "We knew we would get better performance on our database loads and on our query access if we sorted the data before loading it into the Oracle warehouse," said Treffinger. "We also felt it would be more efficient to process the data on the LAN servers instead of within DB2 on the mainframe. SyncSort lets us process the extracted data on an NT server, freeing up the mainframe and allowing us to get excellent processing performance at the same time."


Allied also uses SyncSort for specific data processing tasks that allow the warehouse team to work with greater speed and flexibility. For example, Allied

Success Story: Allied Insurance



wanted to make sure that its employees could not view other employee owned policies, but at the same time maintain the view for the appropriate agents. Treffinger's team added a field to each warehouse record to implement this security feature with SyncSort, avoiding time consuming C programming. SyncSort also proved to be the perfect tool for tasks like making the form codes consistent. "SyncSort can do a lot of these quick fixes or changes," adds Treffinger, "and that saves us a lot of programming, debugging, and testing time."

Treffinger's team needs to process a large volume of data, so performance is important. Treffinger exploits parallel processing, noting that SyncSort can run many processes simultaneously. "Because of its performance and its efficient use of resources, we can run half a dozen processes at the same time and still maintain performance for each one," Treffinger says, adding, "We don't see degradation of the individual processes occurring." Treffinger also uses SyncSort's copy and merge features to reduce resource usage. "We've used the copy capability to process data without sorting and we've used merge so we don't have to re-sort."



SyncSort speeds data warehouse loads and cuts data mart processing time by up to 75%

About SyncSort

SyncSort is a high-performance application accelerator that improves the performance of multiple applications and reduces elapsed time for a broad range of applications. It speeds ETL transactions by up to 90% and facilitates data mining and click-stream processing. SyncSort merges, aggregates, cleanses, and converts data. Other features include filtering, pattern matching, and partitioning. SyncSort will save you time in operations like data warehousing, data mining, data marts, CRM, ERP, DSS, BI, Oracle Financials, and legacy migration. Using Visual SyncSort, SyncSort applications can be created through a Windows-like GUI. For more information or to arrange for a free trial, call Syncsort at (201)930-8200 or visit www.syncsort.com.