



Success Story: Catalina Marketing

Catalina Marketing Stages Data for Warehouse Loads Up to 75% Faster with SyncSort

Organizational Profile

- The leading provider of in-store electronic marketing services
- Collects data from over 10,000 supermarkets

Business Need

- Speed up aggregate processing to improve the performance to load a data warehouse

Environment

- Red Brick Warehouse
- Digital Alpha 8400

Benefits

- Cut nightly load time by almost 75%, from three hours to 50 minutes
- Summed a 3 gigabyte file in just 40 minutes
- Used SyncSort to aggregate data and Red Brick's TMU to load it into the warehouse
- Fully exploited the parallel processing capability on the Alpha box

"I chose SyncSort for our system because of quality and performance..."

Kelly Carrigan
Manager of Information Technology Architecture
Catalina Marketing Corporation

The Challenge:

Catalina Marketing Corporation, a leading provider of in-store electronic marketing services, was in the process of building a terabyte-size data warehouse. After loading a few billion rows into the new warehouse, Kelly Carrigan, manager of Information Technology Architecture, realized that he had to speed up aggregate processing. The size of his daily load was growing, but the size of his batch window wasn't. The data was loaded into a Red Brick Warehouse engine, running on a Digital Alpha 8400.

The Solution:


The company found a solution for the load problem that was caused by the large amount of detailed data that was being posted to aggregate tables. "I chose SyncSort for our system," Carrigan stated, "because of quality and performance, and also because we could leverage the skill set we already had from the mainframe. But I also chose SyncSort because it is fully SMP-enabled, and allowed us to exploit the parallel processing capability of our Alpha box to the fullest."

Red Brick's Table Management Utility (TMU) was a good tool for the load that Carrigan needed to do. "The TMU was one of the best general purpose utilities that I have ever used," says Carrigan enthusiastically. "Unfortunately, it couldn't load and aggregate in parallel. Because we used SyncSort to aggregate, and the TMU to do only the load, we sped up overall processing significantly."

The Benefits:

"Usually we aggregated about 40 million rows a day into about 5 million rows with SyncSort," explained Carrigan, "and then we loaded those 5 million rows into the warehouse with the TMU. It took SyncSort about 40 minutes to sum the 3 gigabyte file, and then we would take the output and post it with the TMU. The alternative was to have the TMU load all 40 million rows. By using SyncSort for the aggregation and then the TMU, we cut our nightly load time by almost 75%, from three hours to about 50 minutes."

Success Story: Catalina Marketing



SyncSort cuts nightly data warehouse load time by almost 75%, from three hours to 50 minutes.

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.