



CASE STUDY
Food Distribution Company
Delivers Faster Reporting
Nicholas and Company

NICHOLAS
AND COMPANY

hvr-software.com

info@hvr-software.com



Background

Nicholas and Company has offered food distribution services since 1939. With operations across middle and western United States, the company distributes food, beverage, and chemical products and equipment to schools, healthcare organizations, and restaurant chains. It also offers services such as employee training and marketing.

Daily Batch Integration

Previously, the company enabled decision support by moving data in a daily batch from several ERP systems to multiple data warehouses and completing extract, translate, load (ETL) operations. But the complexity of the process meant that the data warehouses were straining to produce daily reports in a timely fashion.

The company looked for a real-time data integration solution to dramatically speed up its existing batch business intelligence (BI) reporting. It also created a second solution that would deliver near real-time data access.

Challenge

Nicholas and Company had long been running a business intelligence solution to help it make better business decisions. This system took data from Microsoft Dynamics AX running on SQL Server, a specialized legacy AS/400 ERP running on DB2, a SharePoint solution, as well as vertical industry solutions that included Peoplenet and Roadnet for fleet management and TRUCKBUILDER for cargo load optimization. Data from all these applications was batch loaded into a SQL Server based data warehouse. The company then used a BI solution to create daily reports for vendors, customers, and in-house teams such as sales and procurement.

“We wanted to make more timely business decisions.”

Kanav Puri, Data Engineer on the BI Team at Nicholas and company.

In the Beginning: 16-18 Hour ETL Process

Unfortunately, the complexity of these jobs and the inefficiency of the ETL process meant that they took as long as 16-18 hours to run. Since some reporting jobs relied on the results of previous jobs, the company had to run the jobs sequentially. If an earlier job was not completed quickly enough, reports

CASE STUDY SNAPSHOT

Customer:

Nicholas and Company

Challenge:

- 16-18 hour ETL jobs
- Reduce time of running reports
- Data accuracy

Solution:

HVR speeds data movement and transformation for batch ETL processes

HVR's CDC capability transmits changes from legacy systems to a new data lake in near real-time

Databases:

- Amazon S3 Data Lake
- Amazon Redshift Data Warehouse

Results:

- Reduced time for one ETL job from 10 to 2 hours
- Reduced time for another ETL job from 60 minutes to five minutes
- Data is now centralized in a data lake and no longer needs to be purged
- Reliable and accurate data

Use Case:

- Data Warehousing
- Data Lake

that relied on data from that job might include old data. This meant users did not always trust the data in reports. Nicholas and Company wanted to reduce the time necessary to cleanse and transform data it needed for its BI reports. The company also wanted to publish near real-time metrics.

Need for More Timely Business Decisions

“We wanted to make more timely business decisions,” explains Kanav Puri, Data Engineer on the BI Team at Nicholas and Company. As a result, they decided to undertake a second project to enable near real-time data dashboards. This new near real-time solution would replace an application that previously took data from three different on-premises legacy ERP systems used to operate three separate warehouse facilities: Microsoft Dynamics AX running on SQL Server, a specialized legacy AS/400 ERP running on SQL server, and a legacy system ERP running on DB2. Like the first application, this one also previously used a batch job to trigger an ETL solution that transformed and moved data from the source systems to a data warehouse for reporting.

Optimizing Business Through Fast Data Delivery

The company planned to use the new near real-time data for multiple use cases, including answering questions about sales to determine the most profitable items, which locations or sales representatives are generating the most revenues, as well as trends by location or customer. In the future, Nicholas and Company also wanted to use the system to track food delivery to customers to determine whether the deliveries would be on time.

Solution

One Product. Two Solutions.

In early 2018, the company implemented HVR. For the BI application, it uses HVR to move the data from the data sources to their existing BI solution.

For the near real-time dashboard, Nicholas and Company uses HVR to replicate data from its three ERP solutions to a large-scale cloud-based Amazon S3 data lake. Although the company's data is currently structured, the data lake can also store flat-file, image and unstructured data. By taking advantage of HVR's change data capture (CDC) functionality, Nicholas and Company moves data to the data lake in near real-time. The company then employs a cloud-based Amazon Redshift data warehouse for reporting.

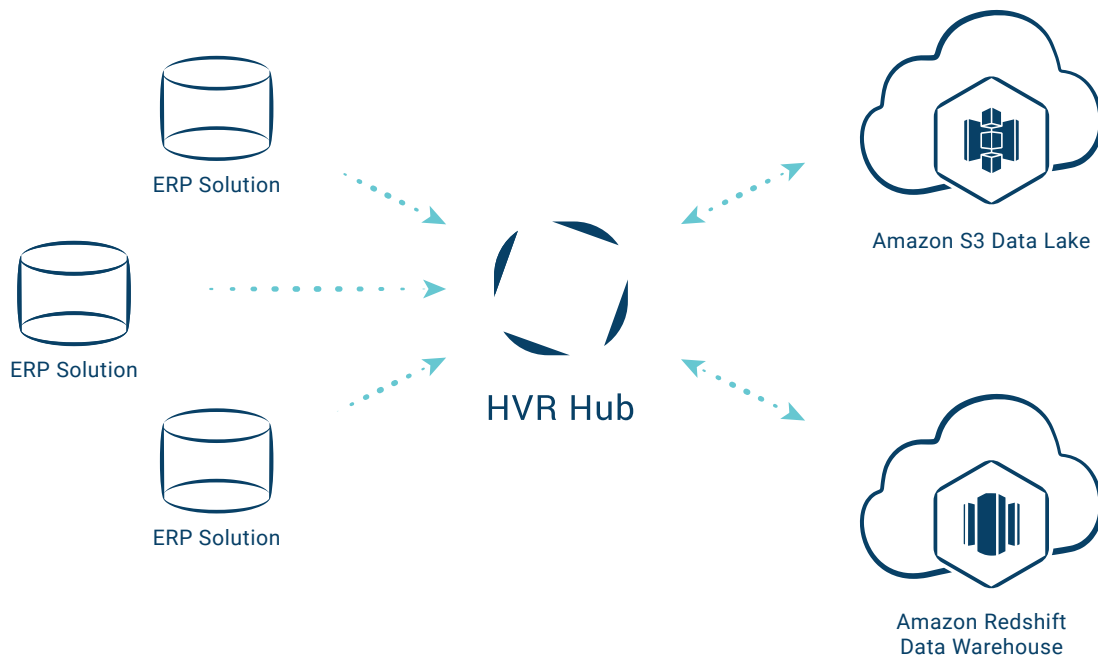
In some cases, HVR transmits data to both the data lake and the data warehouse simultaneously for near real-time reporting. In other cases, an ETL solution transforms the data for use in Redshift. End users take advantage of data visualization tool to analyze the data in the data warehouse as well as drill down for more detailed information.

Why HVR

Before selecting HVR, Nicholas and Company evaluated several solutions. "We selected HVR, says Puri, "because the GUI is user friendly and because of the great service the HVR team provides. HVR engineers know the product inside and out and have been supportive throughout. HVR designed our solution, gave us best practices to follow for our use case, and was very responsive to any issues."

"We selected HVR because the GUI is user friendly and because of the great service the HVR team provides."

Kanav Puri, Data Engineer on the BI Team at Nicholas and company.



Results



10 hrs



2 hrs

| Faster ETL for Batch Reporting

ETL for batch reporting is now much faster. For example, using HVR, Nicholas and Company has reduced the time one of its ETL jobs run from 10 hours a day to 2 hours. Creating a Sales Scorecard report that once took 60 minutes is now completed in five minutes.



| Near Real-Time Analytics

Nicholas and Company also provides users with a near real-time sales dashboard that allows them to ask questions about orders and see the most up-to-date information.

For the new distribution application, says Puri, "We will use sensors on our trucks to track deliveries to different customers. Each truck typically stops at 25-50 different delivery locations along its route. Every time a delivery is made, we will use HVR to capture that data in near real-time and send it to the data warehouse where we will be able to leverage it for reports and metrics."

| A Future-Ready Architecture

The current use cases for the data lake and data warehouse are just the beginning for Nicholas and Company. Says Puri, "Now that our architecture is complete, we have the capabilities that will open the door for us to adopt more advanced capabilities, such as machine learning or using unstructured data from different sources like social media sites."

"Some of our systems are very busy. A table might have 8 million total records and make 1 million updates on top of that."

Kanav Puri, Data Engineer on the BI Team at Nicholas and Company.

| Track Comprehensive Information

HVR also allows Nicholas and Company to keep track of all data changes in the data lake to deliver a comprehensive archive. Explains Puri, "Some of our systems are very busy. A table might have 8 million total records and make 1 million updates on top of that. When our legacy systems were installed in the 1990s, memory was at a premium. The systems could only store 90-days' worth of data. With HVR moving data into the data lake, we can now keep track of everything and we're not forced to delete anything. Our source systems can't keep all these historical changes, but the data lake can."

Comprehensive storage means Nicholas and Company can better track what happens in its business. For example, say a customer returns an order. The order must go through several phases before it can be called a successful return. HVR allows the company to keep track of every change as it happens so they can analyze and improve each step of the process. Says Puri, "We now have data to perform analysis that we couldn't do otherwise."