

Case Study:

MANUFACTURING – RECALL SPEED AND COST



How Trace3 helped a global automotive manufacturer reduce recall numbers with Big Data

BENEFITS

- Accelerated time to value for Big Data solution
- Reduced risk through global visibility by VIN
- Faster responses to inquiries about warranty issues and manufacturing defects
- Stronger competitive positioning in an extremely competitive industry
- 50% reduced recall costs, ROI under 1 year

THE CHALLENGE

The global automotive industry is undergoing an upheaval. Lawsuits, recalls, and the entry of Tesla into the marketplace—a startup company with the luxury of documenting every nut and bolt from the ground up—are striking fear into the hearts of the market leaders.

This large automotive manufacturer, a Trace3 client, has recalled hundreds of thousands of vehicles over the past few years because it lacked visibility across its global supply chain, manufacturing plants, dealers, and service network. Processes, parts, and vehicles were documented along the way, but the information existed piecemeal in disparate data warehouses around the world.

Because there was no logical, structured way of bringing information together once a vehicle was out in the world, drawing correlations and determining which vehicles to recall was nearly impossible. Manual discovery was the only option. As a result, large numbers of vehicles had to be recalled, most of which had no defects. This resulted in astronomical costs, loss of customer loyalty, and damage to the brand.

To make matters worse, the company was spending tens of millions of dollars on siloed reporting systems that were not accessible by or integrated with other lines of business.

Without Trace3, this customer would have spun its wheels, so to speak, and spent a lot of time and money on point solutions that would not have given it the desired visibility into global data.

THE SOLUTION

Big Data is a brave new world for most companies, and expertise is scarce. This customer had engaged a large consulting firm to implement a Cassandra-based solution that would only have solved a subset of its problems. However, Trace3 identified opportunities to go farther using an Apache™ Hadoop®-based analytical platform. Trace3 won the services business, and installed and configured the Hadoop cluster on Cisco UCS Servers with internal storage. To automate cluster management, Trace3 provided StackIQ Cluster Manager for Big Data. The deployment was completed on time and on budget.

This project became a showcase for a large global initiative and the solution is now being implemented in every country where the manufacturer has a presence.

By using an enterprise Hadoop distribution to centralize data from thousands of sources, including SAP, IBM mainframes, data warehouses, Microsoft databases, and scanned paper documents, this customer can now identify correlations quickly and much more effectively. Vehicle identification numbers (VINs) are now used as database keys, allowing any authorized user to associate specific vehicles with parts and bills of materials from suppliers, assembly line times and workers, and dealer and service center interactions worldwide.

Instead of recalling an entire model line, this customer can reduce the sample set when looking for variables and narrow recalls

down to specific VINs. It can then schedule spot-checks and visual inspections with those customers, minimizing the impact and publicity around any potential manufacturing defects. Having a centralized repository for all data across the lifecycle of a vehicle gives the customer complete visibility into the reliability of its supply chain, the true quality of its manufacturing, and what happens to every vehicle after it has been sold.

Without Trace3, this customer would have spun its wheels, so to speak, and spent a lot of time and money on point solutions that would not have given it the desired visibility into global data.

FUTURE GOALS

As the usage of this customer's Big Data infrastructure continues to gain traction, Trace3 is suggesting using Hadoop-based in-memory applications to do real-time stream processing and reduce the time it takes for typical Hadoop processes (e.g., ETL functions, join queries, and data counting) from hours to minutes. This project became a showcase for a large global initiative and the solution is now being implemented in every country where the manufacturer has a presence.