BUSINESS WHITE PAPER / MANUFACTURING

A CAR COMPANY POWERED BY DATA

Volvo Car Corporation drives product design, quality, cost reduction, and customer satisfaction through data-driven decision-making.

TERADATA®

TABLE OF CONTENTS

Executive Summary

Today's car manufacturers face a hyper-competitive global market fueled by stratospheric development costs; rapidly changing customer preferences; and escalating regulation of safety, environmental and fuel efficiency performance. As a result, the automotive lifecycle is increasingly becoming a global feedback loop in which data produced by vehicles on the road directly shapes the design and manufacture of those still on CAD screens and assembly lines.

Manufacturers worldwide are trying to leverage the data produced by vehicles in the field, but finding actionable insights in that flood tide requires an evolved combination of sophisticated IT infrastructure and agile, data-driven business processes. An organization must be able to aggregate, integrate, and analyze very large data volumes from many disparate sources, then act nimbly on the results. It's a formidable challenge, and no company is further up the learning curve than Volvo Car Corporation.

At Volvo, a Teradata® system integrates product configuration, warranty, and vehicle diagnostic data to support technical and business analyses throughout the product lifecycle. The ready availability of current, quality-controlled data is transforming decision-making processes throughout the organization, with beneficial impacts on quality, warranty costs, customer satisfaction, and bottom line profitability.

TERADATA®

Volvo and the Value-driven Vehicle

Volvo Car Corporation (VCC) is the well-known auto manufacturer founded in 1927 in Gothenburg, Sweden. In 2010 the company sold 373,525 cars and currently employs about 20,000 people. The company was acquired by Zhejiang Geely Auto Group from Ford Motor Company in 2010.

One of Volvo's distinguishing characteristics has always been the way its core values have visibly shaped the form and function of its automobiles. Founders Assar Gabrielsson and Gustaf Larson set that expectation in 1927. "Cars are driven by people," they said. "Therefore, the guiding principle behind everything we make at Volvo is – and must remain – safety."

The company has kept its founders' promise, leading the automotive industry in safety engineering for nearly eighty uninterrupted years. Anti-lock braking systems, collapsible steering columns, front and rear crumple zones, energy-absorbing bumpers, side-impact protection systems, three-point seatbelts – all appeared first in Volvo cars. The company's visionary long-term safety objective – to design cars that won't crash – reflected in the shorter-term goal that by 2020, no one will be killed or injured in a Volvo car.

Over time, other values have been added to the Volvo brand identity – quality, innovative design, environmental sustainability – and each has made its mark on the company and its products.

A second distinguishing Volvo characteristic, less obvious from the outside but equally differentiating, is its systematic use of operating data from vehicles in the field to improve the quality and performance of those in production and design. The company has long compiled product configuration and warranty information as an engineering resource, and in 1999 it began collecting diagnostic read-out (DRO) data as a window into performance and mechanical failure under actual field conditions.

> ***"Cars are driven by people, therefore, the guiding principle behind everything we make at Volvo is – and must remain – safety."***
>
> Volvo founders Assar Gabrielsson and Gustaf Larson

Today's Family Wagon Is Wired for Data

Today's automobiles are replete with electronics: sensors, controllers and up to 60 microprocessors per vehicle. Onboard computers perform a wide range of control, monitoring and diagnostic functions in the engine, transmission, braking and traction control systems, cruise control, climate control, passenger compartment, and instrument panel. Each of these systems generates a diagnostic trouble code (DTC) when it detects some sort of fault – sometimes correctly, sometimes not. Faults that require immediate attention trigger the infamous Service Engine Soon light.

Trouble codes are typically stored in the engine control units (ECU) until the vehicle requires scheduled maintenance or repair. At the dealership, a service technician connects an analyzer to the vehicle and reads out the stored codes from the ECU's to guide his troubleshooting and repair efforts. At a Volvo dealership, the codes are then uploaded to a central database at Volvo headquarters, where they form a global reference on all mechanical and electrical failures occurring in all Volvo models over time.

And DTCs aren't the only information Volvo collects. Because trouble codes are simple binary signals, they provide no context for the fault conditions they identify. But the sensors that measure for fault thresholds are perfectly capable of generating incremental measurements, so Volvo equips its vehicles with a variety of data loggers to record a wide range of variable measurements. These include wear factors such as catalyst deterioration, and operating parameters such as engine speed and load.

TERADATA

Much of this data is available in the standard outputs of outsource components that are used throughout the auto industry, yet most of it goes unused. Industry regulators require only four or five of these measurements to be collected and reported. Volvo, however, harvests this data aggressively as a resource for improving product design and manufacturing quality. Uniquely, the company has extended its systems to collect nearly 400 discrete measurements.

The aggregate data volume is large and growing rapidly. Each vehicle in the field generates 100-150 kB of interpreted data per year. In keeping with then-prevailing standards for IT architecture, the company originally began collecting this data in a dedicated data mart. But analysts soon realized they couldn't begin to realize its potential value in isolation.

"We talked to Ford, of course, and to many other references. Everyone told us that the Teradata performance was outstanding. And Teradata was the only vendor we evaluated that would promise us satisfactory performance without a long list of qualifiers."

Bertil Angtorp, senior business analyst at Volvo

Capacity Crunch at the DRO Data Mart

"We had a warranty data warehouse on an Oracle database," recalls Bertil Angtorp, senior business analyst at Volvo, "and we were starting to build a new, separate warehouse for the diagnostic readout information. Then we realized that what we really needed was to integrate those data sets. We knew there was business value to be gained if we could easily match warranty claims against diagnostic data from actual service records, especially if we could do it without manual integration. We tried to bring the diagnostic data onto the existing warranty platform, but the performance was terrible."

"The analyses we do are very CPU-demanding," Angtorp says, "and our old system could only handle about 20 jobs at a time. We'd have SQL jobs queued up and you'd have no idea when you'd get CPU time to process your query. It was a standard joke that when you finally got a response to your database inquiry you'd have forgotten what your question was. We really wanted to bring in additional data from the hardware and software specification archives, but the performance issues made it impossible."

In late 2003, Volvo analysts began building a business case for an integrated data warehouse, and researching possible technical solutions. Performance was their primary search criteria, and they found it at then parent company Ford, who's Inventory Monitoring Alert System (IMAS) is built on a Teradata system.

"We talked to Ford, of course, and to many other references," Angtorp recalls. "Everyone told us that the Teradata performance was outstanding. And Teradata was the only vendor we evaluated that would promise us satisfactory performance without a long list of qualifiers." In September 2006, Volvo began migrating its data to a new 2-node Teradata 5450 Warehouse that went live in July 2007.

TERADATA

A Knowledge Engine Powered by Teradata

The Volvo Cars Data Warehouse brings together data from four primary sources: a system for managing vehicle and hardware specifications, one for managing on-board software specifications, the system that collects vehicle diagnostic data from service centers worldwide, and the warranty claims system. Data access and analysis are enabled through a variety of standard reports and ad hoc analytics, implemented using BI tools and in-house developed applications.

The new warehouse immediately increased the raw data available to Volvo analysts from 364 gigabytes to 1.7 terabytes, and dramatically improved query response times.

A daily fleet mileage calculation that had taken two hours in the previous environment now ran in five minutes, and a comprehensive report of diagnostic failure codes by model /year was reduced from two weeks to 15 minutes. Where performance constraints had restricted access to a handful of users, the new Teradata system extended access to more than 300 users in product design, manufacturing, quality assurance and warranty administration departments. Where the previous data mart had struggled to process a single query and hour, the new Teradata platform completed one a minute.

The reaction of users is perfectly captured in the voice of one astonished engine software analyst who asked, “How can you sleep at night when you know that all the data is there, and you know what you can do with it?”

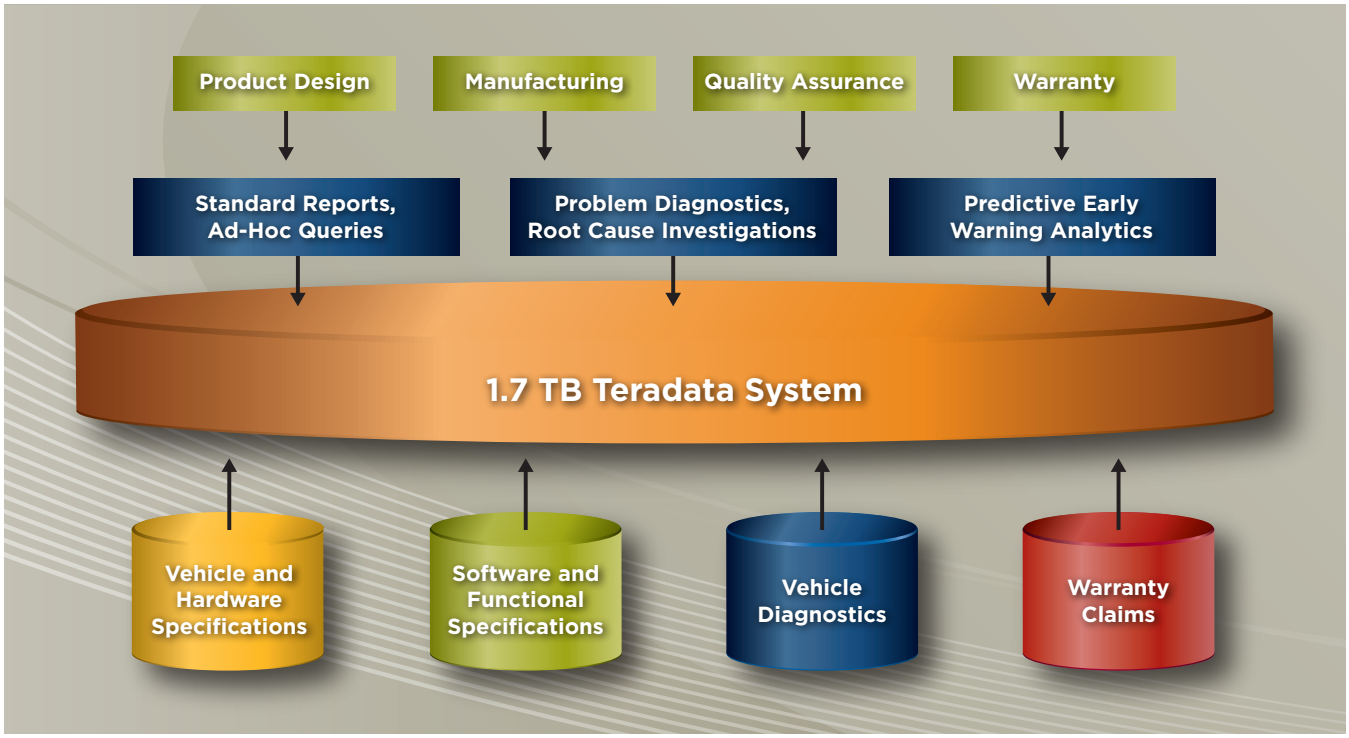

Figure 1.

TERADATA.

Immediate Cost Reduction Impact

Beyond the improvements in query performance and user access, the new data warehouse immediately reduced operating costs by eliminating three single-purpose data marts. "When you consider operation, maintenance, licenses and other resource costs, we've reduced infrastructure expense for data management by approximately two thirds," says Angtorp. An impact analysis performed after the warehouse went live revealed a time-adjusted return on initial project costs of more than 135 percent, significantly better than the 115 percent projected in the original business case.

The warehouse also began generating significant new business value through a growing variety of process improvements and accelerations, in operational areas that range from product design and engineering to manufacturing, quality assurance, warranty administration, and regulatory compliance.

Improving Warranty Reimbursement Accuracy

Volvo analysts had long suspected quality issues in some of the warranty claim data originating in its dealer organization. Small-scale sample comparisons of diagnostic readout data against warranty claim information for the same vehicles did suggest patterns of disparity between actual and reported data, particularly in the mileage reported in association with warranty claims. It seemed certain there must be an opportunity to significantly improve reimbursement accuracy if the data quality issues could be resolved.

With all diagnostic readout data – including mileage data read directly from the vehicle ECU – now in the same warehouse with warranty claim data input by dealers, it was a simple matter to compare values and identify sources of high error frequency. The analysis that followed was once called a "million dollar query". Today, however, Volvo managers carefully describe it as a successful exercise in data quality improvement.

"We did have differences in the quality of data that was delivered by dealers," admits Barbro Forsgren, director of warranty. "And we did discuss our findings with dealers on a case-by-case basis. Today we have very high quality in the data that dealers input. It's always possible to have an isolated error, of course; but there are no patterns of poor data quality, at least not with regard to warranty mileage reporting."

When asked about warranty cost reductions that may have resulted, Forsgren is pleasantly but firmly discrete. "How can I put a cost on a problem we don't have?" she asks with a smile.

Improving Quality and Functionality throughout the Product Lifecycle

Quality has always been a quantitative discipline at Volvo. It's no coincidence that more than 400 members of the design, engineering, manufacturing, and quality organizations are Six Sigma certified. One of the most important impacts of the Volvo Data Warehouse has been vastly improved analytical support for the cross-functional projects these investigators conduct to understand product defects and trace them to their origins in the production process.

"The Volvo Data Warehouse helps us find the cherries in the data," explains Malte Isaksson, Six Sigma black belt and head of Volvo's Six Sigma organization. "We're overwhelmed with data today. We've been using computer aided design and manufacturing since the 1970s, and collecting DRO data for more than a decade. Extracting business value

"Extracting business value from that data depends on our ability to transform it into information, to connect chains of events, to see and understand complex relationships that are often hidden by the sheer volume of data coming at us."

Malte Isaksson, Six Sigma black belt and head of Volvo's Six Sigma organization

TERADATA®

from that data depends on our ability to transform it into information, to connect chains of events, to see and understand complex relationships that are often hidden by the sheer volume of data coming at us. I think the ability to be able to extract business important information from internal and external data sources will be key to understanding customer behavior, needs and wants, hence posing a competitive advantage in the ability of doing so. More and more this will become an important aid in decision making."

Volvo analysts use the Teradata system in a range of sequential investigations aimed at improving quality and customer satisfaction across a product lifecycle that often approaches 20 years. Among the most important focal points are:

> **Prioritizing, targeting, and expediting problem response efforts** – The first order of business in analyzing large volumes of DRO data is deciding which problems are significant and require high priority response. "Diagnostic data include many different alarms created by embedded systems within our vehicles," explains Mikael Krizmanic, senior engine diagnostic engineer. "There's a hierarchical structure within that alarm data, but it usually isn't evident. Sometimes the alarms are related, sometimes not. So our first step is correlation mapping to understand the relationships. We do a type of discriminant analysis looking first at mileage and time to identify patterns of occurrence. We put those in dendrograms to identify correlation strength, then group the related alarms. Comparing those groupings gives us our initial prioritization."

"If we can find and fix a fault in week 26 of a 62-week production run, that's 36 weeks of production that we don't have to address in the field," Krizmanic observes. "It's a problem that none of those customers ever experience. It's a huge advantage both in cost reduction and improved ownership experience."

Mikael Krizmanic, senior engine diagnostic engineer

"Then we calculate failure rates both as a percentage of production and as an absolute number. If you have a fault that occurs at a high frequency but only in a low-production model, that may be much less urgent than one where the failure rate is low but the number of affected vehicles is high. When we're prioritizing issues for remediation we look at the actual number of occurrences."

> **Tracing mechanical faults to their root causes** – "Correlating and grouping failure alarms also gets you part way to root cause analysis," Krizmanic continues. "It's rare that you can get all the way to a root cause from diagnostic data, but you can rule things out and narrow down the possibilities quickly. Once you've eliminated the non-contributors you can usually find it pretty quickly, and we usually have between 50 and 100 parameters to work with."

> **Modeling failure rates over time** – Volvo analysts also calculate predictive failure rates over time using a type of hazard rate analysis based on aging. Each month they look at the number of vehicles that that have reached a certain service age, and the number of those that have experienced a particular failure. They do the same calculations the next month, and accumulate the results.

"This gives us a cumulative hazard function that tells us how many cars in a given population have experienced a particular failure, and how many are at risk," Krizmanic explains. "It's a sort of density function that describes the failure rate over time, and that's what we use for predictive modeling. It's a way of looking at how a failure rate is developing over time, and what will happen if we don't intervene. It gives us a basis for comparing the costs over time with and with correction. Most importantly, it helps us understand which faults will produce large warranty impacts if not addressed systematically."

"These analyses are done for every component of every DTC in the car, and we can view them in many different ways – by model, year, and system. It lets every engineer keep track of the things that are important to him."

TERADATA

> **Correlating mechanical failures with location-specific conditions** – The vulnerability of a vehicle to a specific fault can vary considerably depending on environmental factors and patterns of use, both of which correlate strongly with geography. A car sold in urban China will probably experience different driving conditions and behaviors than one sold in rural Germany – differences in average vehicle speed, engine load, operating temperature, time at idle.

"Because we have both DRO error codes and operational log data in the warehouse, we can understand the relationships between geography, patterns of use and mechanical failure," Krizmanic points out. "A problem may be a high priority for remediation in one geography but not in another."

> **Resolving quality issues within the current production run** – The faster you can find and fix a significant design or manufacturing fault the better – for your costs, profitability, and customer relationships. The holy grail of remedial quality control is to identify a significant issue early and quickly enough to correct it in the eliminate it completely from much of the production run.

"If we can find and fix a fault in week 26 of a 62-week production run, that's 36 weeks of production that we don't have to address in the field," Krizmanic observes. "It's a problem that none of those customers ever experience. It's a huge advantage both in cost reduction and improved ownership experience."

Data-driven Decision-making

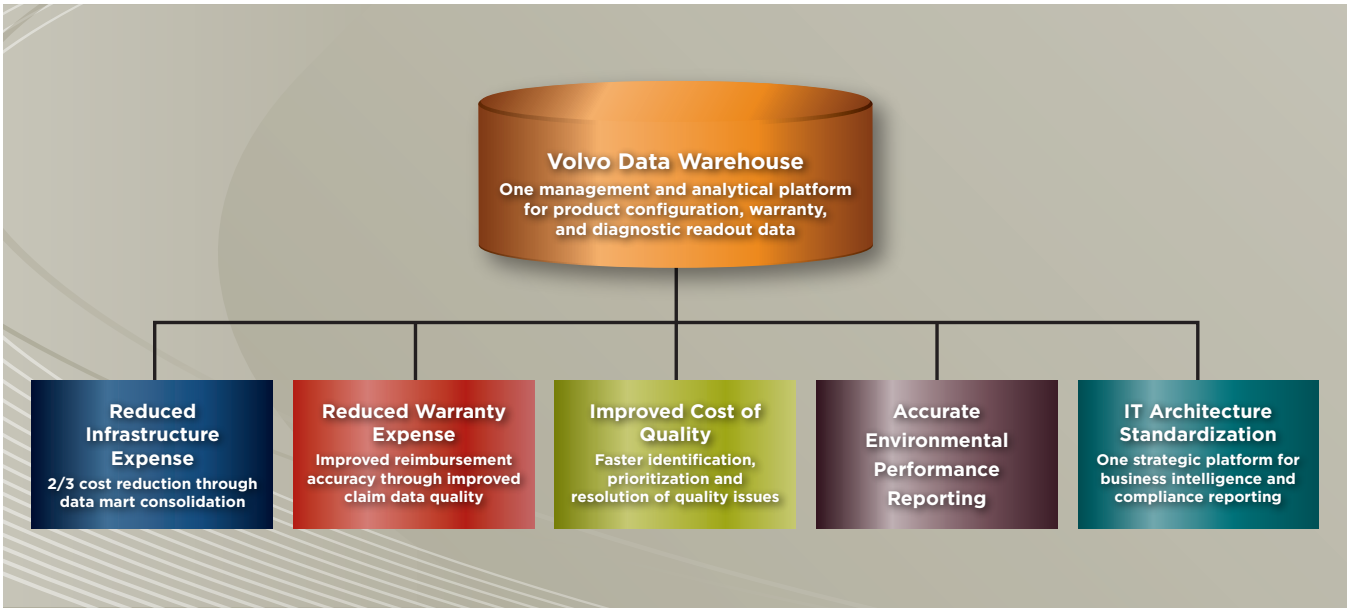

Figure 2. Business benefits of the Volvo data warehouse.

TERADATA

One high-level benefit of the warehouse has been to connect the various groups who contribute to quality issue resolution. "The warehouse has really been something of a silo breaker," Angtorp reflects. "Product Design has become more involved in solving engineering problems. They're now responsible for design throughout the product lifecycle. Our new product development process now includes a formal evaluation of past problems, and an analysis of opportunities to resolve or avoid those issues in the next design cycle. That wasn't the case a few years ago, and it's better reporting that has made it possible."

Another benefit has been a significant reduction in time-to-resolution. "Because Volvo now has detailed data on all its cars, they can scope problems more accurately," explains Teradata industry consultant Torbjorn Rosenquist. "Because they don't need to integrate that data manually they can act on it faster. And because all the functional teams within Volvo are working with the same data, they can act as one."

The importance of those capabilities to vehicle quality and safety is self-evident. But they are also proving essential as Volvo pursues a core value with global relevance – environmental sustainability.

Documenting Environmental Innovation

Beginning in late 2008, Volvo has been rolling out environmentally-optimized configurations of its products. With features that include a lower chassis, aerodynamic underbody panels, an engine start-stop system to minimize idling, higher gear ratios, and specially selected wheels and tires the DRIVe line is designed to maximize fuel efficiency and minimize emissions. All models emit less than the 120g/km limit of carbon dioxide that is the Eurozone threshold for special tax treatment, and the C30 has tested as low as 99g/km. Volvo designers are using the Teradata system to collect and analyze diagnostic information from DRIVe cars in the field to track actual performance against design objectives.

"Because Volvo now has detailed data on all its cars, they can scope problems more accurately, because they don't need to integrate that data manually they can act on it faster. And because all the functional teams within Volvo are working with the same data, they can act as one."

Torbjorn Rosenquist, Teradata industry consultant

"We want to make sure that our customers actually receive the fuel efficiency performance that we've certified,"Mikael explains. "For instance, we know that the start-stop system should engage and stop the engine in 95 percent of all vehicle stops. By collecting and analyzing the diagnostic data from DRIVe vehicles, we can verify that we're achieving this performance. And in any cases where we don't reach that target, we can look at the variable data to understand why. We also have instrumentation on board to calculate average fuel consumption, so we collect that data at each service interval. And we can see the average fuel efficiency performance across the model year fleet."

The warehouse is also providing critical design support for future environmental optimizations. To cite just one example, charging data from the electrical system is being studied to develop an algorithm for balanced use of engine braking to recharge the battery without overcharging.

Enabling Regulatory Compliance

The VDW team recently released a major compliance reporting application onto the Teradata system. The Transportation Recall Enhancement, Accountability and Documentation (TREAD) Act is a U.S. federal law that requires automotive manufacturers to report certain data related to safety recalls, defects, injuries, and deaths related to their products. Volvo previously used a reporting solution

TERADATA.

developed and maintained by Ford, but has been preparing to migrate to a platform of its own following its sale to Khejiang Geely Holding Group.

Because much of the required data was already loading into the VDW, moving the remainder and rebuilding the reporting application to run against the Teradata system was an obvious solution that was quickly endorsed by the Volvo Cars IT organization. Implementation, however, has had to wait for the recently completed upgrade that has more than doubled processing capacity and lifted a temporary freeze on new users and applications.

TREAD Act reporting on the VDW is currently in quality assurance, with the first production report delivered in February 2011.

The Case of the Phantom Vapor Leak

California is one of the world's premier automotive markets. It's also the world's most demanding regulatory environment, especially for environmental regulations. When Volvo engineers began designing the emission monitoring and control systems required under California's On Board Diagnostics (OBD) II rules, it added a fuel system pressure sensor intended to detect leaks as small as 0.5mm.

In the field however, the new evaporative leak detection system proved alarmingly trigger-happy. Out of every 16 leaks detected only 1 proved real. Remediation efforts focused first on the pressure sensor itself, but an analysis of log data captured in the VDW showed the sensor outputs to be accurate and consistent.

Instead of mistakenly replacing perfectly good sensor components, Volvo resolved the issue with a software modification that reduced false alarms by 95 percent and cut the cost of repair by half.

A Designated Platform for Business Intelligence Development

The impact of Volvo's new data warehouse in reducing costs, improving analytical performance, expanding user access and accelerating problem response has impressed management on the business side organization. The VDW steering committee recently implemented a policy that all future business intelligence initiatives should be based on the Teradata platform.

"It is certainly our ambition that when it comes to handling large amounts of data, we use the data warehouse," says Åke Bengtsson, VDW steering committee co-chairman and Vice President of Quality and Customer Satisfaction.
"I would say that today we have only scratched the surface; I don't think we understand yet, from a business point of view, this tool's true potential. And that's one of the main purposes of the steering committee – to understand and communicate the possibilities. I believe that we can better use the data to provide early indications, so we can act more quickly and precisely. That's the purpose really. In today's competitive environment we must be able to act quickly, to reduce the number of steps to an accurate, proactive response."

"Because every car we produce with a fault costs the company money. And every minute, hour, and day by which we can expedite a solution saves money for the company. The earlier we can resolve an issue the better it is for the customer and the company. So I think our direction is clear. We've recently implemented a hardware and software upgrade that should take us several years ahead, with the performance and capacity we need to really utilize the data we have, and to continue developing new solutions and business opportunities. So far our focus with VDW has been on quality but there are also opportunities in sales and service for business enhancement using VDW."

TERADATA

The VDW as Strategic IT Platform

The Teradata system has also earned enthusiastic support from the IT side of the Volvo Cars organization. Jonas Rönnkvist, head of enterprise architecture, recently helped orchestrate its approval as a strategic IT platform within the corporate environment.

"For us, the purpose of a strategic platform is to drive standardization, commonality, and simplification within the IT landscape," Rönnkvist explains, "to standardize and streamline the way we build and deliver solutions to the business. In particular, we look for things that will scale to meet new demands, which hasn't always been the case in the analytical space at Volvo Cars."

"The idea is to build and host an extensible set of services or solutions on a core set of reusable platforms, so that we stop building separate solutions for every kind of need. The Teradata architecture was something we adopted as a way to decrease the number of independent data mart solutions. To standardize the way we build and upgrade analytical needs by using a single data warehouse in a much more scalable fashion than we have done in the past. It's very clear that Teradata is now a key component of the overall enterprise information architecture that my team is driving."

Lessons from the Volvo Data Warehouse

It's all about business value

- Win and keep management support with a strong business case
- Business value is always the highest priority
- IT cost savings are a bonus

It's all about people

- Find the people with strong statistical and mathematical skills (6-Sigma)
- Insight into numbers leads to improvements
- Involve the business at the pilot stage to create ownership

It's all about data

- An enterprise data model based on detailed data saves time and supports the EDW
- Save all your data – new uses will arise
- Plan for capacity, demand WILL grow

"Every new development project at Volvo Cars now follows a standard governance process, including reviews by my team. If the requirements include data consolidation and integration and the design doesn't leverage the VDW platform, the project will be redmarked – it will not be able to proceed without CIO approval."

"The benefit for the business is that we now all believe and agree that the data we're using is accurate, reliable, and true. That is not the case when you have a siloed data mart architecture, because there are always different dialects of the same information. The Teradata system gives us a single view of the business that every business unit trusts and accepts. That's the biggest impact."

"The benefit for the business is that we now all believe and agree that the data we're using is accurate, reliable, and true. That is not the case when you have a siloed data mart architecture, because there are always different dialects of the same information. The Teradata system gives us a single view of the business that every business unit trusts and accepts. That's the biggest impact."

TERADATA.

"Looking at what we've accomplished so far, it was relatively easy to decide on the technical platform. It was much harder to get the process and the support around it up and running. That's why the journey is long. Something like this is always a long-term endeavor, not a one-time project."

Creating a Culture of Data-Driven Design

For all its impact on vehicle design and quality, Volvo's data warehouse may have had its greatest impact on the company's decision-making processes. "Our decision-making has become more fact-based," says Bertil. "Now, whenever a question arises, people invariably ask 'what is the data telling us?' We test our assumptions against the data before we act. Once we've verified the existence of a problem we use the data to determine the scope, to prioritize and scale our response. It helps us make sure that we're focusing on the things that are most likely to affect the customer experience."

The VDW at a Glance

Phase One (2007)

- Two node Teradata 5450 system with about 2.9 TB user data
- Teradata Database V2R6 for MP-RAS (old NCR Unix)
- Backup and restore solution
- FastLoad and MultiLoad for loading data

Phase Two (2010)

- One node Teradata 5555H system, 7.3 TB user data
- Teradata Database version 12, base edition for LINUX
- Backup and restore solution
- FastLoad and MultiLoad for loading data
- Teradata Dynamic Workload Manager
- Teradata Manufacturing Logical Data Model
- Teradata Data Warehouse Miner
- Teradata Viewpoint