

**Hardware and Software**  
Engineered to Work Together



## Top 10 Data Warehouse Trends for 2013

## What are the most compelling trends in storage and data warehousing that motivate IT leaders to undertake new initiatives? Which ideas, solutions, and technologies are at the top of the CIO agenda?



*This brief highlights the most salient trends in data warehousing and discusses best practices that you can follow to keep your organization on the cutting edge. Whether it's implementing real-time analytics, consolidating information to create more efficient data centers, or utilizing Flash memory to improve performance, forward-thinking companies depend on these technical breakthroughs to take the lead in their respective markets. Many of them are heeding these Top Ten Trends in data warehousing:*

### **1 – Performance Moves to the Forefront for Data Warehouse Initiatives**

Just about every organization is wrestling with a tremendous volume of data. In the context of data warehousing, runaway growth leads to more demanding workloads for reporting, data mining, and statistical analysis activities. Delivering extreme performance for data warehousing workloads enables a wider set of users to conduct richer analytics on bigger data sets, so business leaders can make better decisions. Elavon, the third largest payment processing company in the United States, boosted the performance of routine reporting activities for thousands of internal employees and millions of merchants. Large queries that used to take 45 minutes now run in seconds, and massive transactional summaries that formerly took days to complete now run in a few hours, enabling employees to explore new data sets, create new business opportunities, and focus on more strategic tasks.

### **2 – Real-Time Data Gets Real**

Generating reports with up-to-the-second data was a luxury just two or three years ago. Now many businesses depend on low-latency apps to run their businesses. Industry examples include just-in-time inventory alerts for retailers, fraud detection algorithms for credit card companies, and call-data analyses for telecommunications companies. LinkShare, a leading provider of full-service online marketing solutions, sped up reporting for hundreds of thousands of advertisers and publishers, cutting query response times in half and enabling its clients to instantly analyze campaign performance. The improvement in performance has enabled LinkShare to deliver deeper levels of insight to its clients. For example, advertiser dashboards now feature summaries of the top 10 publishers and their results so they can quickly see how best to



maximize their advertising spend, keeping clients abreast of minute-to-minute changes in their affiliate marketing programs.

### **3 – Engineered Systems Simplify the Data Center**

The next wave of computing infrastructure is more integrated, more flexible and easier to deploy, manage and support. This model can be seen in engineered systems where software and hardware are engineered together to create a cohesive customer experience. Within the data warehouse domain, database machines represent the next wave of computing because they perform better, are fully integrated, and simpler to maintain overall. These purpose-built database machines are optimized at every layer of the technology stack to address performance bottlenecks with technical advancements such as columnar compression, Flash memory, and intelligent caching. RL Polk, a leader in automotive intelligence and marketing solutions, replaced seven floor-to-ceiling equipment racks with a one-rack database machine, simplifying maintenance, conserving energy, and freeing up additional space in its data center.

*Purpose-built database machines are optimized at every layer of the technology stack to address performance bottlenecks with technical advancements such as columnar compression, Flash memory, and intelligent caching.*

### **4 – Consolidating onto Private Clouds Maximizes Flexibility and Reduces Costs**

Forward-looking companies are consolidating standalone servers into private clouds. These shared resource pools simplify management by enforcing centralized administrative procedures. You can expand or contract the capacity of any database to meet the dynamic nature of the workload at hand. CIOs love this self-service model because it gets the IT department out of the day-to-day business of approving project requests, procuring servers, installing software, and configuring unique application environments. Instead, IT creates a virtual pool of resources that the business can share and utilize on demand.

### **5 – Business Analytics Get More Accessible**

Analytic applications provide a significant competitive advantage, especially in fast-moving industries with very large data sets such as telecommunications, finance, and retail. Information is the strategic differentiator, and its value grows in direct proportion to its accessibility. Greater penetration of BI across the business raises the collective intelligence of the workforce and enables more informed decision-making—so every department can operate more efficiently. To support all types of analytic applications, you need an integrated, high performance platform that can easily store, organize, manage, analyze, backup and archive all types of data—and make it available to executives, mid-level managers, business analysts, and operational users within the enterprise, and to customers and partners on the outside.





***Modern data warehouses and big data platforms capture social media, location information, sensor data, email, conversations, and many forms of machine-generated data. When properly analyzed, this information can provide unique insights into market trends, business opportunities, maintenance cycles and many other business issues.***



## **6 – Big Data Infiltrates the Enterprise**

Capturing and analyzing all data, including unstructured, semi-structured and traditional data, allows managers to gain a more thorough understanding of their businesses. Whether its fine-tuning supply chains, monitoring shop floor operations, gauging consumer sentiment, or any number of other large-scale analytic challenges, market leaders mine all of their information for valuable insight. As Philip Russom of The Data Warehouse Institute says, “Data used to be just data. Now there’s big data, real-time data, multi-structured data, analytic data, and machine data, all converging with new paradigms like Hadoop, MapReduce, and personal clouds.” Modern data warehouses and big data platforms capture social media, location information, sensor data, email, audio, images, conversations, and many forms of machine-generated data. When properly analyzed, this information can provide unique insights into market trends, business opportunities, maintenance cycles and many other business issues, enabling organizations to operate more efficiently and competitively.

## **7 – Database Platforms Support Mixed Workloads**

Normally data must be extracted from an OLTP system and loaded into a data warehouse before you can run reports and queries. This process consumes valuable resources and requires the IT staff to maintain two separate database systems. Engineered systems that support mixed workloads bring data warehouses and operational data together to simplify reporting and analysis. What does that mean for customers? Better performance and a smaller data center footprint with lower hardware costs and easier system management. More and more companies are adopting these multi-purpose environments, especially financial services companies, telecommunications companies, online retailers and other organizations that support lots of data, images, and transactions. ValueOptions Inc., the largest independent behavioral healthcare and wellness company in the U.S., implemented a data warehouse and online transaction processing application on a single multi-purpose database machine to provide scalability, reliability, and the redundancy needed to support a growing volume of behavioral health data. The high-performance system enables members to quickly access information about programs and health plans, as well as find doctors, providers, clinics, and services more easily and accurately than ever before, simplifying support processes and providing customized managed-care services to meet customer needs.

## **8 – In-Database Analytics Simplify Management**

As data sets grow larger it becomes harder to move them around for processing. In-database analytics takes the processing to the data, instead of the data to the processor. This is an ideal scenario for high-volume data warehouse applications. It opens up new



***Market-leading companies are enhancing their hardware and software infrastructures to accommodate today's demanding requirements for business intelligence, data warehousing, big data, and OLTP/mixed workloads.***

avenues of inquiry that were formerly the domain of super computers—even as it simplifies administration. It minimizes data center costs as there is no need to implement additional hardware and software to run data mining models that can predict customer churn, analyze social networks, compute the most optimal route between two locations, and many other resource-intensive tasks. SoftBank Mobile, a mobile communications company in Japan, utilized this architecture to analyze its call-detail records more than three times as quickly as before, giving it an edge in Japan's highly competitive mobile communications market.

### **9 – Performance Skyrockets with Flash and DRAM**

Today's modern hardware platforms improve performance by managing data in server memory, reducing latency and eliminating I/O operations that are the bottleneck for large queries. Having a cache of ultra-fast DRAM memory coupled to a Flash infrastructure delivers tremendous performance improvements, far exceeding the mechanical limits of standard disk drive technology. By placing an entire database in memory, you can improve response time for time-sensitive BI applications such as real-time dashboards and event-driven reports. Going forward, IT leaders will focus on achieving the right balance between Flash, DRAM and disk to maximize performance while lowering costs.

### **10 – HA Becomes Essential for Mission Critical Data Warehousing**

Data warehousing workloads are becoming mission critical. Enterprises rely on the information flowing from them to support operational analytics and make business decisions in real time. Many of these organizations are deploying high availability (HA) data warehouse solutions that enable continuous access to information, including during planned and unplanned outages, with zero downtime. To gain the high ground in an extremely competitive market, one leading biotech firm in California deployed a data warehouse with high availability to help its sales and marketing team analyze sales data, review case management activities, confirm customer payments, and spot market trends, enabling round-the-clock support for thousands of customer and partner interactions.

For more information on Oracle Data Warehousing solutions, visit [www.oracle.com/datawarehouse](http://www.oracle.com/datawarehouse).



**Oracle Corporation**  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

[oracle.com](http://oracle.com)



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2012, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

**Hardware and Software, Engineered to Work Together**