

Case study: Database integration by Hokuriku Coca-Cola using a database appliance

Hokuriku Coca-Cola Bottling (Hokuriku Coca-Cola) integrated data used in 3 types of databases into Oracle Exadata and thus achieved low total cost of ownership (TCO) and high-response environment.

Gartner

© 2010 Gartner, Inc. and/or its Affiliates. All Rights Reserved. Reproduction of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

Summary

This research note is an analysis of the database integration at Hokuriku Coca-Cola. New technologies are introduced on the market and there are a lot of companies that are waiting for these technologies to mature and the success stories to accumulate. However, it should be reaffirmed that only new technologies may be the key to cost reduction and strategic initiatives.

Major observations

- When it comes to database, performance and operation burden can be significantly improved by adopting Appliance.
- Appliance is not a magic box. Performance and cost can get worse than those in other environments depending on the treatment, that is why it is necessary to establish an usage policy once the Proof of Concept (PoC) is completed and the nature of the product is understood.

Recommendations

- Perform a PoC with data and SQL used in the actual operating environment during the introduction of Data Warehouse Appliance aimed at resolving the performance issues of the current system.
- When considering the adoption of HP Oracle Database Machine (Oracle Exadata), (see note 1) please bear in mind to add indexes and to perform constant performance identical with the regular Oracle Database.

Overview

Hokuriku Coca-Cola used an enterprise system aimed at understanding sales trends and improving customer service, but there were still issues like poor performance and high operation burden. HISCOM is a 100% subsidiary company of Hokuriku Coca-Cola, which provides IT system services, such as planning, design, structure and operation. When Oracle Exadata was announced in 2008, September HISCOM considered that Oracle Exadata could solve the above-mentioned problems. As a result of the PoC it was confirmed that performance can be significantly improved and Oracle Exadata was adopted. Furthermore, low TCO, timely data provision, as well as improved site productivity were achieved.

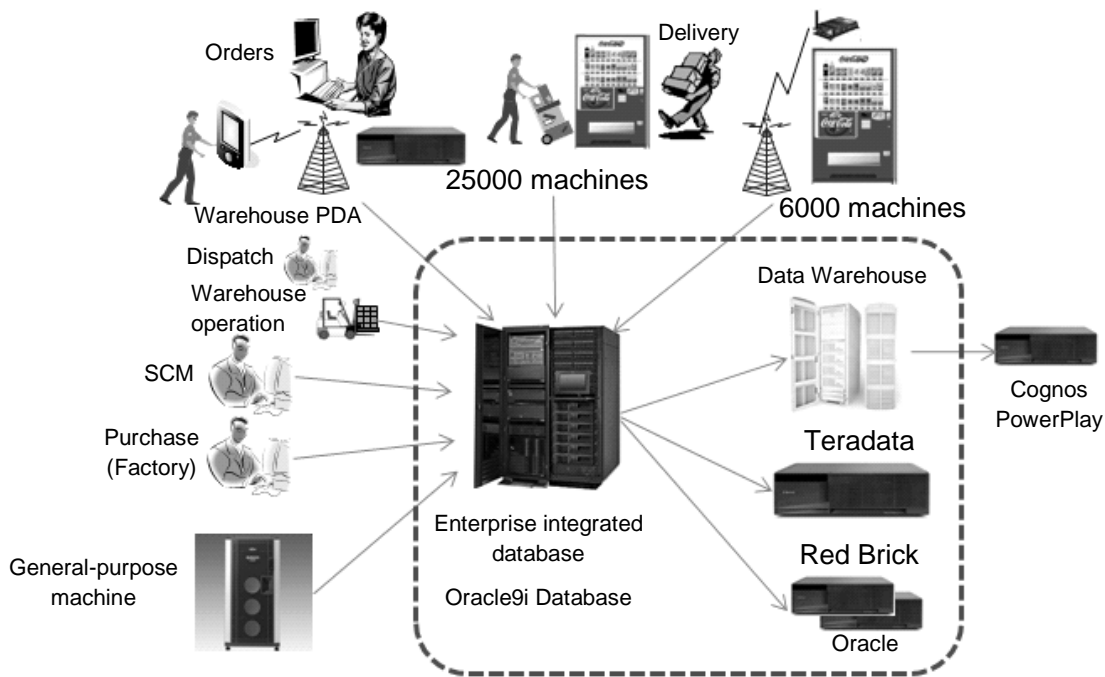
Case study

Introduction

Hokuriku Coca-Cola is a company that produces and sells soft drinks as Coca-Cola, Georgia and Aquarius in the prefectures Toyama, Ishikawa, Fukui and Nagano. As of April 2010 there are 711 employees. This is the number of) the full-time employees. The proceeds of sales for 2009 were JPY 52,3 billion 36 million.

Hokuriku Coca-Cola used an enterprise system aimed at understanding sales trends that change in accordance with the changes in economic environment, weather, diversity of beverage products and improving customer service. More than 5 million reports containing sales details data from around 6000 vending machines are sent wirelessly every month, as well as order data from the stores and order data input by the operators. The enterprise system concentrated the enormous amount of data into an enterprise integrated database. Furthermore, in order to understand and analyze sales trends, Data Warehouse was prepared along with the enterprise integrated database, the necessary data was extracted from the enterprise integrated database and was loaded after processing. Oracle 9i Database was adopted into the enterprise integrated database, Teradata was adopted into Data Warehouse, Red Brick was used for a part of the analysis (see Figure 1).

Figure 1 Summary of the former enterprise system of Hokuriku Coca-Cola



Source: partially amended at Gartner on the basis of HISCOM materials

Issues

A series of systems has been used since 2003 and the analyzed data has been increasing year by year along with the progress of the system use. As a result, issues like poor performance and increased system operational load became serious and the demands for real-time data were growing. The issues are arranged below.

Poor performance: The burden for data extraction, processing, transfer, load processing increased, batch processing time became insufficient and problems such as poor performance of the enterprise integrated database appeared.

Increased system operation burden: Each system operated maximally because of the heterogeneous environment, respective technical experts had to be provided and as a result the operation burden increased.

Increased demand for real time: Experienced employees were about to retire, tasks performed by part-time employees were increasing and although the decision support in the IT field was adequate, the simultaneous searching in the enterprise integrated database and data warehouse performed everyday by a large number of users was not expected. Data about budget and performance management, profit management, analysis of vending machines sales was taken and inserted in Excel and Business Intelligence (BI) tools.

Approach

HISCOM is a 100% subsidiary company of Hokuriku Coca-Cola, which provides IT system services for Hokuriku Coca-Cola such as planning, design, structure and operation. In 2007 HISCOM started collecting information and considering specific measures for improving the performance and reducing the operation burden. The adoption of Oracle Database into the enterprise database was almost decided and Data Warehouse investigation began. HISCOM announced the following 4 conditions for the selection of the new Data Warehouse.

- Possibility to apply an Entity Relationship design (ER) identical with the one of the enterprise system (possibility to normalize the structure of the star schema adopted in Red Brick).
- Possibility to reduce the data loading time.
- Possibility for an adequate search speed realized at a lower TCO.
- Remarkable decrease in response does not occur when accessed by many people.

The first things to be examined when investigating the new Data Warehouse were Oracle, which was almost adopted as an enterprise database and the adoption of General-Purpose Relational Database (RDB). However, it was quickly removed from the list, because although improvements in function and server performance were achieved, it was still necessary to invest significantly in a hardware and to perform a heuristic tuning in order to achieve an adequate search speed.

The next things to be analyzed were Teradata, which was adopted as a Data Warehouse and the so-called Data Warehouse specialized products. It was already proven at Hokuriku Coca Cola that searching and batch update at a specific terminal can be performed at high speed and it was possible to estimate the extent to which the server should be introduced in order to achieve the required performance when it comes to Data Warehouse specialized products. However, there were concerns that the configuration would be very expensive if expansion of users and real-time data browsing were considered as a use case.

In the process of gathering such information, Oracle Exadata was introduced by Oracle in September 2008. The competitive Data Warehouse Appliances of Oracle Exadata were Teradata and Netezza, but Oracle pointed out that specific high-speed mechanisms that improved significantly the search performance were added to Oracle Database. That is why, Mr. Watanabe, the responsible official of HISCOM decided that the Enterprise Integrated Database and Data Warehouse could be integrated into Oracle Exadata, which provided not only a high-speed search, but also a high credibility of the Online Transaction Processing (OLTP).

However, Oracle Exadata was a recently presented product and was not as matured as the regularly used Data Warehouse Appliances. In the Gartner Research Note APP-09-109 2009 November "The Hype Cycle⁵ of used information in Japan: 2009" Data Warehouse Appliance is a technology positioned just before the information campaign and 2 ~ 5 years are necessary to adopt the mainstream. That is why HISCOM participated in Oracle Exadata PoC program and performed a verification. At the PoC were prepared 9 types of SQL for searching data and were compared to non-normalized data model (more than 70 million cases) and normalized data model (more than 20 million cases). This data and SQL were actually used in Hokuriku Coca-Cola and the performance in the environment at that time was poor. PoC results were good, the search results that could not be obtained in Red Brick were obtained in several seconds and there were sections that even exceeded the original expectations. What is more, it was logically possible to obtain similar performance with Teradata as well, but the necessary investments for that were relatively large. The technical experts of Oracle promised a relatively easy maintenance and that is why Oracle Exadata was adopted as Data Warehouse and the traditionally divided database was unified and integrated.

Results

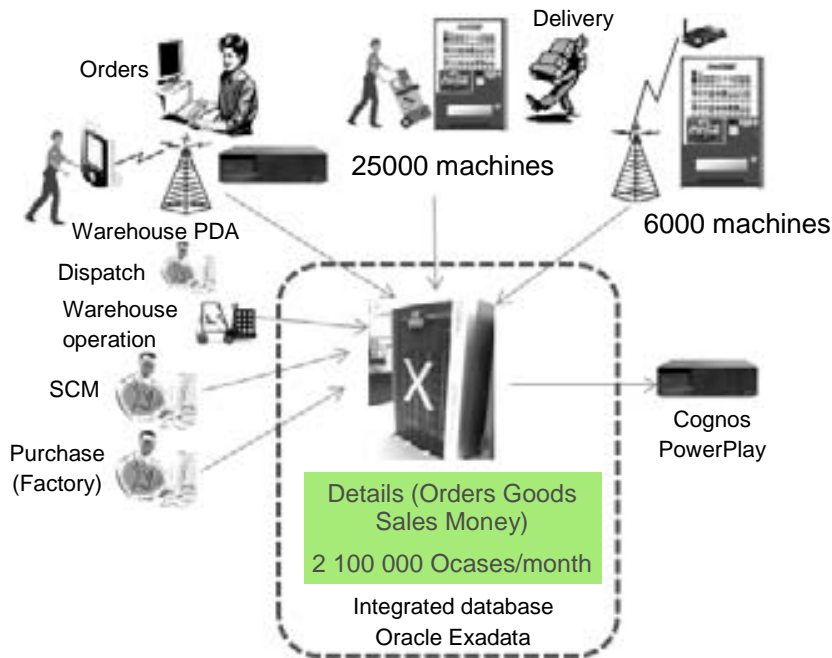
The system was launched in 2009, October after 9 months of surveys and design/development tests of the enterprise system and after the verification of the Data Warehouse that was performed at the same time. Since May 2010 the following results have been accomplished without occurrence of major problems low TCO, timely data provision, improved site productivity.

Acquisition of flexible augmentability: This environment offers a much lower price than the estimated one for a separate upgrade of the enterprise integrated database and Data Warehouse. In addition, if the performance is insufficient in the future, it will be possible to solve the problem by minimal additions of the necessary servers and storage and by purchasing a software license in accordance with the structure of the added hardware, instead of re-purchasing the whole system as in the former configuration.

Timely Data Provision: Up to now it was necessary to perform operations such as data extraction, transfer and load from the enterprise integrated database to the Data Warehouse, but since the consolidation of the database, the processes of transfer and load now can be substituted by standard functions of the database, such as Insert and Update. That is why the reduction the development costs and the early release of the data analysis are possible. Furthermore, the nightly batch process which used to take several hours was speeded up to several minutes at most and the re-run of batch in case of an error can now be completed within the time limit.

Improved site productivity: Contribution to data extraction and analysis process PoC results-the search process that used to take 3 minutes in the old system system now takes around 5 seconds, high response performance is realized, the productivity of the sales representative has increased. The analysis and operations have become much more efficient and as a result the labor unions are also grateful because the employees do not have to work overtime in order to finish the processing.

Figure 2 Summary of the enterprise system of Hokuriku Coca-Cola



Notes: The accounting data indicates the following: orders-received orders, goods-products distribution, sales-number of sales, money-proceeds.

Source: partially amended at Gartner on the basis of HISCOM materials

The speeding-up of the Data Warehouse and the batch processing is starting to bring benefits in terms of business as well. For example, in the previous environment it was necessary to prepare the analyzed data in a nightly batch and that is why the analyzed data was data from more than 1 day ago. However, now it became possible to perform analysis of a data closer to the present situation in less time than before (fresh data) and inventory level could be reduced from 5 days to 3 days. Furthermore, by obtaining such database, very important results in terms of business are achieved- more sophisticated demand forecasting, area marketing combined with map data, readiness to hit the next move which was almost given up before.

Major Success Factors

PoC performance: In this case there were problems with the data base performance and the solution of these problems was the most clear point. May be the effects in the environment of our company could have been predicted from the white papers of the product or from the occurring cases. However, through the verification performed with data used in an actual operation environment and SQL, it was found that the necessary performance could be obtained and the minimum configuration could be determined.

Understanding why it is not a magic box: HISCOM understood that Oracle Exadata eliminates the common bottleneck in Oracle Database due to a large amount of data transfer from the storage to the database server by combining a new software and hardware and without any special tuning. When it comes to Database Appliances, the non-required tuning is a sales point, but in Oracle Exadata it is possible to perform a performance tuning as in a common Oracle Database by adding an index. However, HISCOM evaluated that when it comes to tuning operations, it is much easier to use the Tuning Advisor than Oracle 9i Database, but when it comes to an atypical search, it is executed in different environments without being an object of integration in terms of cost effectiveness.

Decision and scheme: In Oracle Exadata the price of the whole system has been reduced by adopting a commodity hardware. There were many concerns regarding the adoption of this hardware as an enterprise system, such as: will the necessary support be provided if any problem occurs, won't there be frequent break downs, etc., but having in mind our experience, we decided to adopt the system at certain risk. However, in order to reduce the risk of troubles, JOB control, system operation/ monitoring, application functions, such as Shell were installed into an environment, isolated from Oracle Exadata.

Instructions

In Japan, when it comes to enterprise systems, Unix is adopted in most of the cases. On the other hand, all over the world there are more and more cases of obtained high performance at an affordable price through adoption of the Linux-based database and commodity hardware. There were many concerns outside the company regarding this case and the adoption of a new product based on commodity hardware.

If you have to choose between the commonly-used technology that has already proved its efficiency and a new technology that offers only risk without any proven achievements, it is understandable to choose the commonly-used technology. However, we should not forget that by accepting the risks and giving a chance to the new technology, we can achieve results that are much higher than expected.

Related research

- “Release of the first hardware product of Oracle, Japan” (INF-08-71, 2008, October20th);
- “Presentation of Oracle RAC into the major adoption phase” (APP-09-84, 2009, August31st);
- “Hype Cycle of used information in Japan: 2009” (APP-09-109, 2009, October5th);

Remarks

HP Oracle Database Machine is a database appliance that consolidates InfiniBand network and Exadata Storage Server with a grid configuration into an HP server/grid that operates with Oracle Database 11g Real Application Clusters. This product is no longer available for sale because on September 15, 2009 Oracle and Sun Microsystems presented their next version-Sun Oracle Database Machine (Oracle Exadata Version 2).

Abbreviations

PoC	Proof of Concept	Proof of Concept
-----	------------------	------------------