WHITE PAPER

WC

www.wintercorp.com

SPONSORED RESEARCH PROGRAM

WINTER CORPORATION

FIELD EXPERIENCE WITH LARGE SCALE DATAWAREHOUSING ON ORACLE

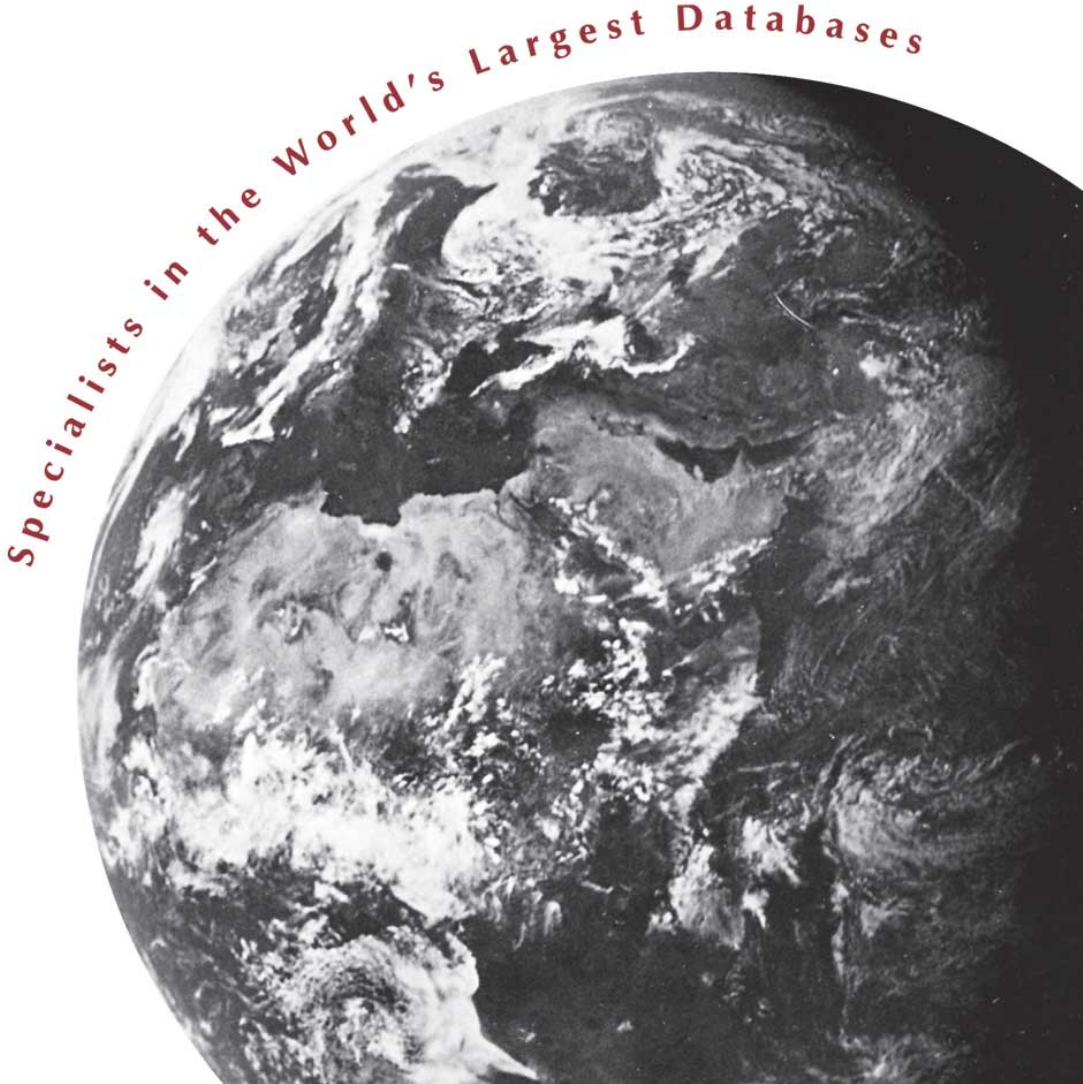

Field Experience with Large Scale Datawarehousing on Oracle

Richard Winter

Winter Corporation
411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

EXECUTIVE SUMMARY

Data warehouses have not only grown to be very large in many enterprises: their growth is accelerating. A data warehouse occupying a terabyte of disk – a rarity a few years ago – is now quite common. And many enterprises are now planning or implementing data warehouses occupying ten terabytes of disk or more.

But the rapid, accelerating growth of a data warehouse can easily become too much of a good thing. If your architecture cannot support the growth, the data warehouse can become unresponsive, unstable and fundamentally unusable. To user executives, then, it is terribly important to learn much headroom they have. The strongest proof is what other users have accomplished.

Oracle, as a widely used, general purpose database management product, has not always had full market recognition for its advances in scalability. Otherwise well informed professionals have come to believe that Oracle is not suited to data warehouse requirements involving more than a terabyte of disk. In fact, Oracle is in successful operation with rapidly growing data warehouses much larger than a terabyte.

As part of an investigation of this phenomemon, Richard Winter, Founder and President of Winter Corporation, personally interviewed the users described here and wrote this paper. While Oracle sponsored this process, it had no active participation in the interviewing process. Winter Corporation retained full editorial control over the content and presentation of this paper and is responsible for the facts, findings and conclusions presented here.

Winter Corporation has become convinced that Oracle is capable of managing data warehouses containing multiple terabytes of *data* and handling thousands of users[1]. This paper describes four such cases: BestBuy (4 TB of data); Telecom Italia Mobile (6 TB of data); France Telecom (15 TB of data); and, Amazon.com (4 TB of data). The largest of these Oracle instances manages a 32 TB database that is expected to double in size in the next year. These systems are all in production, under heavy use, supporting multiple applications and feeding multiple downstream systems. In addition, three of them have online user populations in the hundreds to thousands of users. The evidence presented here, therefore, should lay to rest any lingering questions about Oracle's ability to handle a data warehouse involving more than a terabyte of data. In fact, the evidence shows that Oracle has a significant presence in the management of very large data warehouses and that it has customers in that arena who are confidently moving forward to use Oracle on a yet larger scale.

[1] The volume of disk under management by the database system – which is more frequently quoted in the industry – is usually several times larger than the volume of data. Thus, the systems mentioned here range in size up to 32 TB of disk.

TABLE OF CONTENTS

FIELD EXPERIENCE WITH LARGE SCALE DATA WAREHOUSING ON ORACLE

FIELD EXPERIENCE WITH LARGE SCALE DATA WAREHOUSING ON ORACLE

A WINTER CORPORATION WHITE PAPER

1. The Need for Proven Large Scale Architectures

Users building a very large database – one that will manage a terabyte of data or more – almost always want to do so on a proven architecture. That is, they generally want to know that the particular combination of products they plan to employ has been successfully used in production on a scale that equals or exceeds their own requirements. Having such proof means that someone has gone where they are trying to go, roughly speaking. Such proof indicates that their own project is substantially less risky than if they were going to be the first.

Why does this matter? Users want some indication that the product architecture does actually function correctly, with stability and with satisfactory performance, on the scale they plan to use it. This is a rational concern because there are no guarantees of scalability beyond proven levels. All vendors seek to engineer their products so that they will scale as consistently as possible. All vendors hope that their products will work just as well on 50 terabytes of data as they do on 50 gigabytes of data. And, there is no question that the leading vendors work diligently toward that ideal.

But,to paraphrase Yogi Berra, informed users realize that "it ain't been done until its been done".

Of course, the size of the database is not the only important factor in such assessments. To gain any comfort from what has been done before, the prospective implementor must also consider the nature of the use, the application, the workload, the performance requirements and other factors.

Oracle is the most widely used database product in the world today. Yet, it is not well known among practitioners whether Oracle is in successful use on data warehouses containing a terabyte or more of data. For this reason, Winter Corporation has herein compiled four case studies of Oracle users running substantial, production data warehouses, each of which holds three terabytes or more of data.

Readers of this paper are advised to take note of the difference between the volume of disk storage and the volume of data. Both figures are cited in this paper. In data warehouses, disk storage is usually several times as large as the data itself. Many people refer to "terabyte" data warehouses that employ a terabyte of disk but store only 100 gigabytes (GB) of data.

The data warehouses in this paper all contain more than three terabytes of *data* and occupy at least five terabytes of disk.

2. Process

Winter Corporation asked Oracle Corporation to provide names and contact information for several customers who are in production with an Oracle database containing more than three terabyte of data. Richard Winter, Founder and President of the company, personally interviewed

each of the four users by telephone. Winter also followed up via email and telephone, where he deemed necessary, to obtain additional details and/or validation of what he had been told. In general, the interviews were conducted completely independently of Oracle. In one call, an Oracle representative was present in the interview to facilitate communication with an overseas customer but did not participate in the substance of the conversation.

Oracle sponsored the creation of this White Paper and was provided an opportunity to comment on it prior to publication. By written agreement, however, Winter Corporation retained complete and final control of content and presentation. All facts, findings and conclusions in the paper are the responsibility of Winter Corporation.

3. Some Key Requirements for the Large Scale Data Warehouse

Before reviewing the individual cases, it is worth asking what users should look for when examining an installation for indications that the products being used are suitable to support a large scale data warehouse.

3.1. Data Volume

Data size matters. Algorithms that work efficiently on 50 GB data warehouses don't necessarily succeed on 5 TB data warehouses, which are 100 times as large. Even though all vendors engineer for scalability, some aspects of system behavior can be overlooked and not actually surface until the products are used in production conditions. So, if you plan to implement a data warehouse with 5 TB of data, you want to use a product – if at all possible – that has proven it can operate successfully on 5 TB of data or more. For this purpose, production use is stronger proof than benchmark results because most real applications are so much more complex and demanding than standard benchmarks.

3.2. Stability Under Heavy Load

Its one thing to have a data warehouse that contains terabytes of data and supports a few batch reports or a few online queries – especially if there are weeks or months between updates. Its quite another to have a large data warehouse that supports a demanding, dynamic workload consisting of many reports, many queries, many analyses and many updates – all going on at the same time. Whether such a system can exhibit *stability* under heavy load is a crucial issue in large scale data warehousing. What is stability? It is consistent availability and consistently good performance in the presence of continuing, demanding, diverse and changing use.

3.3. Large User Population

Much larger user populations have become a fact of life in data warehousing in the last few years. Now nearly everyone in the enterprise has access to the intranet. In large organizations this means tens of thousands to hundreds of thousands of potential users. Further, data warehouses are increasingly available to business partners on extranets and to customers via the internet. Open internet access in consumer businesses can mean user

populations in the millions. Many data warehouse products were initially created to support either single users or just a few users. Expanding them to service hundreds, thousands or even millions of users requires an architecture appropriate to the task. Supporting a large user population in combination with a large data volume is a challenge all its own. Users need to quantify their requirements as best they can and look for proof that the products they are considering have been proven with respect to the scale of user population they expect to service.

3.4 Multiple Applications

Where as data marts are often single purpose, data *warehouses* are usually expected to support a diverse and growing set of applications. It may be a noteworthy feat to support ten terabytes of data organized into a handful of tables that are only queried in one characteristic fashion. But this is not the requirement of most data warehouses. Most data warehouses must support multiple applications sharing an enterprise copy of the data. Delivering performance to multiple applications, each of which has its own queries and access patterns, is a far more difficult proposition, especially when data volumes are very large.

3.5 Adaptability

Enterprises build data warehouses in part to deal with rapid change in their business environments, strategies and activities. That means the data warehouse itself must often change rapidly as well. New applications are implemented frequently, changing workloads and usage patterns. The business reacts to market changes and suddenly the queries are focused on a different subject. Mergers and acquisitions occur, sometimes multiplying volumes and workloads – sometimes changing everything. As a result, the data warehouse must be more adaptable than any other business system. The data warehouse is there to help *manage* the business and hence must deal with the sum total of all other changes in the environment.

3.6. Scalability

There seems to be only one constant in data warehousing today: volumes will grow. If the warehouse is successful, then next month there will be more data. There will be more tables, more rows, more columns. There will be more users. The users will generate new and more demanding work based on going more deeply into the meaning of the data and what it says about how the enterprise can perform better.

So the data warehouse must be able to scale: in data volume; in user population; in workload; and, in other dimensions. And the more successful the data warehouse is, the greater will be its demands for scalability.

3.7. Manageability

As data warehouses become larger and more complex, the difficulty of managing them also rises. Data warehouses must deal with large volumes of data; large numbers of tables; huge populations of users; thousands of disk drives; and an ever increasing population of applications and tools to service. The multi-terabyte data warehouse is far too large, complex

and dynamic an entity to be managed by unaided human beings. Instead, the requirement today is for increasingly functional, comprehensive and integrated capabilities for management. And these requirements can be addressed only partly by the creation of external tools and management facilities. To a substantial degree, they must be addressed by building manageability into the database engine – and by making facilities, to the extent practical, self-managing. An important indicator of manageability is whether a large scale data warehouse requires a disproportionately large staff of database administrators (DBAs) to maintain it.

4. Case Studies

4.1. Best Buy

Best Buy Co, Inc. (NYSE: BBY), with headquarters in Eden Praire, Minnesota, is the largest volume specialty retailer of consumer electronics, personal computers, entertainment software and appliances. Best Buy operates over 1900 stores in the US, Canada, Puerto Rico and the US Virgin Islands.

In 1996, Best Buy decided to implement a data warehouse to enhance its merchandising and supply chain management. IT management explicitly sought an open system architecture, anticipating that this would provide it with long term advantages in cost, flexibility, technology and hardware vendor independence. An evaluation anticipating a large scale implementation, heavy use and rapid growth resulted in the selection of the Oracle database, a Sun Microsystems E10000 server and EMC Symmetrix data storage to comprise its strategic architecture. Today the data warehouse includes full detail of every point of sale transaction in every store for the last two years, as well as extensive detail on suppliers and products.

According to Randy Mattran, leader of the data warehouse group at BestBuy, the principal data warehouse contains 4.2 TB of data, stored on a total of 6.4 TB of disk. This is under the control of a single Oracle instance that services approximately 3000 users. Among the major applications are: Business Performance Measurement (BPM), Vendor Performance Management (VPM) and Retail Business Performance Management (R-BPM). R-BPM, the largest of the applications, accounts for 100,000 queries and reports per morning and supports 2500 users in the stores.

Detailed data on every transaction in every store is loaded nightly into the Oracle database. As a result, retail managers are able to respond immediately to changes in customer behavior with respect to product purchase patterns, model preferences, returns, requests for service and warranty claims. About 70% of the use of the data warehouse on Monday mornings is for ad hoc query and reporting. Many of these requests involve scans of large tables. Best Buy reports that Oracle is able to manage a mix of large and small queries, while providing appropriate levels of service to each.

Via the VPM application, 350 vendors are able to monitor the timeliness and accuracy of their deliveries to all Best Buy stores. This application recovered the cost of its implementation within weeks of its initial launch; more importantly, it served to increase the timeliness of deliveries by suppliers and therefore the satisfaction of customers.

Best Buy employs 4 database administrators (DBAs) to support and maintain its principal data warehouse. The database design has been successful in supporting multiple applications; in supporting the addition of new applications and subject areas over time; and, in accommodating changes to the schema.

Much of the distinctive business value of Best Buy's data warehouse results from its ability to keep comprehensive, timely, online records of the full detail of every retail transaction. This need for comprehensive, timely detail meant that Best Buy wanted a data warehouse

company that was committed to continuing investment in scalability and new technology. Best Buy management appreciated that its data warehouse requirements were large scale to begin with; would grow; and, would change. So, Best Buy wanted a vendor that would be enhancing its product in anticipation of customer needs for large scale performance. This has proved to be a critical factor in the last five years. During that period, Oracle was enhanced with *materialized views*, a feature that accelerates query performance by a large factor in many applications, often eliminating or reducing the need for joins, sums, averages and similar calculations at query processing time.

At a recent MicroStrategy Conference, Best Buy received an award for operating the world's most scalable BI application – based on the extraordinary volume of reporting activity supported in its R-BPM system. Data for the 100,000 daily reports, of course, is delivered by Oracle8*i*.

Over the past five years, Best Buy has achieved average annual growth in net income of over 52%, according to the online edition of the Wall Street Journal. Best Buy stock has outperformed the Dow Jones specialty retailers as a group by a factor of more than 20 during that same period. Best Buy's award winning data warehouse program has clearly been a significant factor in its remarkable business performance, contributing to the company's ability to manage its supply chain, manage its stores and merchandise its products.

4.2. France Telecom

France Telecom is one of the world's leading telecommunications carriers, with over 91 million customers on the five continents in 220 countries and territories.

In 2000, France Telecom decided to expand its data warehouse program, creating a new system named Symphonie. Within Symphonie, the company planned to create an online repository for all telecomm traffic, to support fraud detection, customer service, network traffic analysis for both the fixed and wireless network, operations and marketing. A key issue was the volume of data involved: 500 million call detail records (CDRs) per day! At the time, France Telecom had a data warehouse implemented on Teradata but wanted more flexibility in configuration and hardware options.

After an evaluation in which DB2 and Oracle were considered, France Telecom decided to implement the new data warehouse on Oracle, on an HP V2500, with the help of Cap Gemini, a system integrator.

As implemented in production today, Symphonie includes a 32 TB Oracle database that contains approximately 180 *billion* CDRs! In both database size and records, this is larger than the largest single database in the world in the most recent Winter Corporation Database Scalability Survey, completed approximately 12 months earlier. This Oracle data warehouse services 8000 end users, of whom an estimated 600 are concurrently online at peak times. Remarkably, new CDRs are added to the data warehouse hourly – as many as 65 million in a peak hour. Each day, Symphonie acquires over 100GB of new data and distributes tens of TB of data to users and downstream systems. This extraordinary database is thus one of the largest in the world; services thousands of end users; and, is updated with new data hourly.

According to John-Luc Cochennec, Architect of Symphonie at France Telecom, the company has found this enormous Oracle-based system to be very stable in production use. It has also met performance objectives. France Telecom requires that all standard online CDR queries

be completed in 4 seconds and this criteria has been met in production. These standard queries are required to specify a date and either an originating or a receiving phone number. It is noteworthy that the system can process many concurrent queries of this type against a database of over 180 billion CDRs, with response time of under 4 seconds. In addition to standard queries, ad hoc queries for analysts and power users are supported.

On the basis of its experience to date, France Telecom has decided to expand the system further in 2002. Plans for this year include:

- migration from Oracle8.1.7 to Oracle9*i*R2, the latest release of Oracle;
- migration from the HP V2500 to an HP Superdome;
- expansion of storage capacity from 40 TB to 80 TB;
- migration from hourly batch update to continuous update; and,
- at least doubling the volume of online data.

The volume of data to be added this year will depend in part on experiments under way in data compression. With the expansion of the system this year, France Telecom will have consolidated 50 databases into Symphonie, resulting in a large, long term cost savings.

In addition, Symphonie has been used successfully to prevent fraud, by providing company representatives with timely, detailed information about customer usage of services. Thus, if a the activity in a customer account increases suddenly, the company can examine the usage to see whether it is characteristic of fraudulent or legitimate use. It can then make an informed contact with the customer, if appropriate, to investigate further.

Symphonie is now the *only* data warehouse at France Telecom for users who deal with telecommunications traffic. It is thus used to feed data marts for marketing, for network needs, finance and other uses.

According to Christian Becht of Cap Gemini, two features of Oracle that were particularly valuable at France Telecom are composite partitioning and transportable table spaces. Time based partitioning is used at the top level for manageability and query efficiency. The CDRs are partitioned by day, yielding 365 top level partitions for the data that is between one day and one year old. Data for the present day is partitioned by hour. The most recent hour of data is loaded into a separate tablespace, which can be on a separate Oracle instance on a separate machine. Once loaded, the new partition can be switched into the online system and instantly become an integral part of the online data. In this way, loading of new data does not interfere with online query or compete with it for resources.

As the data warehouse program has progressed, France Telecom has added new data and applications to the data warehouse, demonstrating that the schema is in fact application neutral and capable of accommodating growth and change in its usage.

4.3. Telecom Italia Mobile

Telecom Italia Mobile (TIM) was founded in 1995 to create a division of Telecom Italia that would focus exclusively on the rapidly growing wireless market. The company went from zero to one million customers in its first year. It now has more than 50 million lines in Europe and Latin America and is still growing rapidly.

In 1996, TIM decided to create a data warehouse in which it could keep comprehensive data on its network traffic. TIM knew that businesss was growing rapidly but could not predict the

extent of the growth that would occur. Therefore, it needed a scalable data warehouse platform with the potential to manage growth in data far beyond any known requirements.

TIM selected Oracle database software and a Compaq Alphaserver to comprise the key elements of its data warehouse architecture. For the last two years, TIM has used EMC data storage to behind the Oracle/Alphaserver system. TIM executives felt that Oracle, with its financial strength and its commitment to large scale data warehousing, would deliver the capabilities that would be needed as the company and the data warehouse grew.

According to Claudia Scaroni, responsible for the traffic data warehouse at TIM, the Oracle database occupies about 10 TB of EMC storage. WinterCorp estimates this contains about 6 TB of user data. This includes call detail records (CDRs), 150-200 million of which are added to the data warehouse in a typical day. On peak days, as many as 250 million CDRs may be loaded. TIM has the largest community of users on a single telecomm network in the world.

After the traffic data warehouse was in production, TIM expanded its program to also maintain information regarding its customers: demographics, contract information and usage history.

The traffic data warehouse feeds data marts on which several applications are run in such areas as marketing, finance, billing and network operations. Wireless telecom is a very fast moving business, especially in Europe, where new technology and services are introduced at a rapid pace. The competitive picture changes rapidly and TIM users rely on the data warehouse for a clear picture of how customer behavior is changing in response to product introductions, promotions and other events. In addition to supporting applications on data marts, the traffic data warehouse also supports some user inquiry. For example, marketing analysts have many questions that cannot be anticipated in the development of the application data marts. Such questions may be submitted directly to the data warehouse.

This large and rapidly changing data warehouse is supported by approximately 3 full time equivalents of database administrator effort.

TIM had a large increase in data volume following the introduction of SMX (cell phone based messaging), and expects similar large increases as other new value added services are offered in the next few years. TIM is successfully using many of the advanced business intelligence features of Oracle on the traffic data warehouse. These include materialized views with aggregates, bitmap indices and composite partitioning.

4.4. Amazon.com

Amazon.com, the leading online shopping site, is one of the largest and fastest growing businesses on the world wide web. In 2001, its sales exceeded $3 billion. Its online inventory of books and other products includes many millions of items. These are purchased by millions of customers in over 200 countries around the world.

In 1999, Amazon.com determined that it was outgrowing its existing system, implemented with Red Brick Warehouse (now a product of IBM). Amazon's complex database schema was proving unwieldy in Red Brick and the product was not well suited to the volume and frequency of update required by Amazon. Further, Amazon's data and usage requirements were expanding rapidly. In 1999, Amazon already had about one TB of data and its data volume was doubling annually. Amazon needed a more flexible, scalable and robust data

warehouse platform. According to Mark Dunlap, who manages the data warehouse program, the company also wanted an open strategy – one that would leave them a wide range of hardware options. Amazon.com therefore evaluated Oracle and Informix.

Following an evaluation and benchmark -- in which both vendors had to load Amazon data into a warehouse; run Amazon queries; and, unload the data – Amazon chose Oracle. Oracle demonstrated superior query performance and comparable load performance. Both products were tested on a large Sun E10000 server.

Today, the Amazon data warehouse contains 5 TB of data and continues to grow at a 100% annual rate. It contains data on customers, inventory, orders, products, supply chain, web site activity (clickstream), pricing, financial activity and other subjects. The database is updated 5 times a day. The data warehouse is used by about 500 users covering every department within the Amazon organization. On the average, users submit about 2000 queries a day. Growth rates are formidable: while data volume is doubling annually, the number of queries is growing at 2.5 times per year. Even with Amazon's rapid business growth, this indicates that usage is growing faster than business volume – an indication of the extent to which the organization is investing in usage, presumably because business value is being developed. Mr. Dunlap has made several efforts to quantify the business value that is being realized. While the nature of these is confidential, he stated that in one subject area alone, he was able to establish that the business value realized in one year was more than 10 times the cost of the data warehouse.

Recently, Amazon moved the data warehouse to a different data center and a new hardware platform. This entailed moving the database across the country to an instance of Oracle running on an HP Superdome. The migration occurred with essentially no disruption to user activities. The data was actually exported from the existing Oracle system; transferred via high speed network to the new data center; and, imported into the new Oracle system *without ever writing it into a flat file.* This was possible because of partitioning features of Oracle, in which partitions of the database can be independently moved from one system to another. About 5 partitions were moved in each export operation. Over a period of about two weeks, the database was created incrementally on the new system. At the end of the transfer process, two weeks of updates were applied and it was arranged that new signons would connect to the new location. The process was essentially transparent to end users.

Mr. Dunlap also credits Oracle composite partitioning with contributing in a major way to performance. First, the many subpartitions enable highly parallel query. Second, the time based major partitions create a "free" index. The Oracle optimizer uses the information in the partition key to avoid accessing partitions that cannot contain data relevant to a given query. Since many queries are about the most recent month, week or day of data the partition elimination is a powerful optimization.

5. Findings and Conclusions

These five data warehouses between them illustrate some salient points about Oracle:

- *Oracle has a proven capability to support large multi-terabyte databases*: the data warehouses reviewed in this report range in volume from 4.2 TB of *data* at BestBuy to approximately 15 TB of *data* at France Telecom. Measured by total database size, these systems range up to 32 TB.
- All of these very large data warehouses are growing rapidly and all of their owners expect to continue using Oracle as their strategic data warehouse platform. This tells us that *these Oracle customers have confidence in the continued scalability of Oracle – in the near term, well up into the tens of terabytes.* Amazon.com's data volume is doubling in size annually. France Telecom's data volume will double in size this year, presumably resulting in a total database size of 64 TB prior to data compression.
- Most of these data warehouses already *service large user populations*. For example, France Telecom cites 8,000 users; and, BestBuy has 3,000 users.
- Most of these data warehouses *handle substantial workloads*. The BestBuy data warehouse processes 100,000 reports and queries every morning for one application.
- These data warehouses are *updated at least daily, in high volume*. Amazon updates 5 times per day. France Telecom updates hourly and is heading toward continuous update. Thus, these are not old style, static data warehouses, but dynamic entities. The telecomm users in this group each add hundreds of millions of CDRs to their data warehouse daily.
- These data warehouses *support multiple applications* and are therefore leveraged across substantial areas of, if not all across, the owning organization. Thus, the Amazon warehouse is used in every department; and, the France Telecom warehouse is used by everyone in the organization who needs access to traffic data for decision support.
- These data warehouses features complex schemas -- generally too complex for the star schema paradigm. These schemas have been able to evolve and grow as new applications are added. Amazon deals successfully with slowly changing dimensions.

Thus, these case studies lay to rest some myths about Oracle: Oracle is in successful use on a large scale on production data warehouses with over 3 TB of data – in one notable case, well over 10 TB of data. In fact, these data warehouses are growing so fast that, by the time this paper is read, they will be considerably larger than is indicated here. These data warehouses are used by substantial populations of users, ranging up into the thousands. They handle substantial workloads, in one case over 100,000 reports and queries a day. They are updated daily or more than daily. They support multiple applications and complex schema. So, Oracle has won a place in the arena of large scale data warehousing and the users interviewed for this report – and others with similar agendas -- are confidently moving forward and rapidly advancing the frontier of scalability for Oracle.

A leading center of expertise in very large databases,
Winter Corporation provides services in
research, consulting and education.

We help users and vendors understand their opportunities;
select their database and data warehouse platforms;
define and measure the value of their strategies, architectures and products;
plan, architect and design their implementations;
and manage their scalability, performance and availability issues.

Our focus is databases near, at and beyond the VLDB frontier.

WINTER CORPORATION
411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

visit us at www.wintercorp.com