ITG STATUS REPORT

May 2011

Bottom-line Advantages of IBM InfoSphere Warehouse

Overview

Data warehouse deployments continue to accelerate. Organizations have found that once data warehouses are put in place, end user demand expands rapidly. Businesses that start with small-scale, compact data warehouses may find that, within a few years, they must deal with terabytes of data, increasingly diverse applications and hundreds to thousands of users.

There is a downside. As data warehouses grow in size and sophistication, costs escalate. Projecting outlays over, say, the next three to five years may be a sobering experience. What can be done to reduce, or at least contain growth in expenditures?

Vendor pricing and packaging strategies have a major impact on costs, and these strategies – and their bottom-line implications – can vary widely. This is particularly the case for solutions offered by IBM and Oracle.

Oracle, for example, prices its Oracle Database 11g and data warehouse extensions based on numbers of users or numbers of processor cores in underlying servers. For IBM InfoSphere Warehouse solutions, IBM offers a new pricing option, based on terabytes (TB) of user data.

Per terabyte pricing allows data warehouse costs to be more closely aligned with growth, enables user populations to expand without corresponding increases in license fees, and does not penalize organizations for concurrent or complex query workloads that require significant processing power. Costs may be further reduced if data compression is employed.

IBM also offers a low-cost version, InfoSphere Warehouse Departmental Edition, with a full set of data warehouse capabilities for servers with up to 16 cores and systems with up to 15TB of user data are supported.

In comparison, Oracle's low-end (up to four sockets) Database 11g Standard Edition does not support such functions as partitioning, data mining, online analytical processing (OLAP) and advanced compression. While it may be employed for basic business intelligence applications, Standard Edition is not a serious candidate for data warehousing.

The combined effect is that InfoSphere Warehouse offers significantly lower software costs. For data warehouses with between 1TB and 5TB of user data, for example, initial license and three-year support costs for InfoSphere Warehouse Departmental Edition average 69 percent less than for use of Oracle Database 11g Enterprise Edition. Figure 1 illustrates this picture.

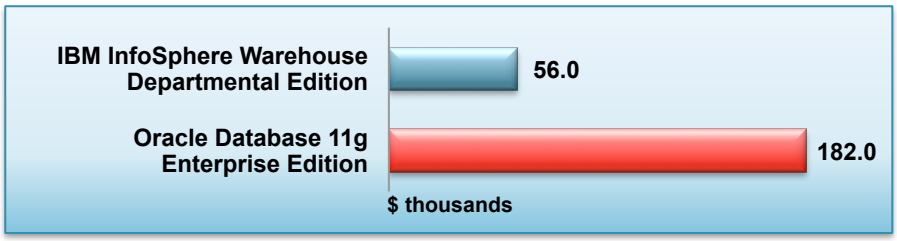

Figure 1: Average Three-year IBM InfoSphere Warehouse and Oracle Database 11g Software Costs (1)

For data warehouses with between 10TB and 50TB of user data, comparable three-year costs for InfoSphere Warehouse Enterprise Edition average 39 percent less than for use of Oracle Database 11g Enterprise Edition. Figure 2 illustrates this picture.

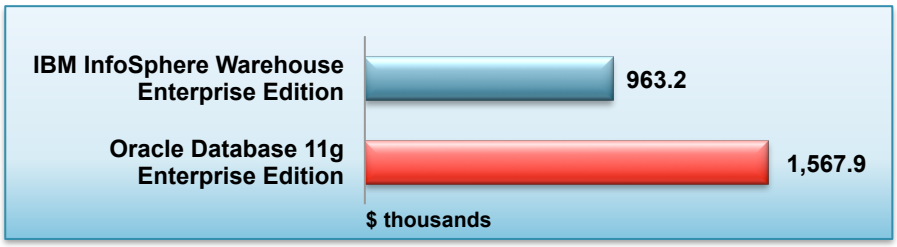

Figure 2: Average Three-year IBM InfoSphere Warehouse and Oracle Database 11g Software Costs (2)

InfoSphere Data Warehouse calculations employ a conservative assumption of 50 percent data compression. Oracle Database 11g also enables data compression, but this does not affect pricing.

Calculations are based on list prices discounted by 80 percent for Oracle Database 11g Enterprise Edition and 40 percent for InfoSphere Warehouse Enterprise and Departmental Editions. Discounts offered by both vendors may vary between customers.

Pricing and packaging issues are discussed below. Bottom-line differences in core technologies, particularly in data compression, that affect comparative costs are addressed later in this report. Details of cost calculations are also presented.

Pricing

Vendors of databases and data warehouse frameworks typically price based on numbers of processor cores in underlying servers, numbers of users, or both. There are a number of potential drawbacks to both approaches:

- ***Core-based pricing***. The trend in microprocessor design is to increase the number of cores per chip. While this improves overall server price/performance, performance per core evolves more slowly. Lower hardware costs do not necessarily translate into lower per core software costs.

  Upgrade increments, moreover, become larger. For an organization replacing, say, a four year-old dual-socket x86 server with latest-generation Intel technology, the number of cores will increase from 4 to 12. If this platform is later upgraded to a 4-socket server, the number may increase from 12 to 32.

  Similarly, adding a latest-generation dual-socket server to a cluster typically means adding 8 (2 x 4) or 12 (2 x 6) cores. Increases in core density can be expected to continue for the foreseeable future.

- ***User-based pricing***. Experience has shown that numbers of data warehouse users increase over time. Even when systems initially support specialist users such as executives and analysts, other groups typically seek access to the same data resources.

An analytical data warehouse designed to support, say, 30 users may expand within a few years to serve hundreds of managers, professionals and front-line workers such as salespeople and customer service representatives. In large organizations, thousands or tens of thousands of users have become common.

Per terabyte pricing offers greater cost-effectiveness. This approach has become the norm among vendors of special-purpose data warehouse "appliances" such as Teradata, IBM (Smart Analytics System and Netezza) and Oracle (Exadata Database Machine). IBM is the first major vendor to apply it to software-only solutions.

IBM per terabyte pricing is, it should be noted, for core user data only. The definition does not include indexes, logs, system temporary spaces and other system data catalog structures, which typically represent 30 to 50 percent of overall warehouse data volumes.

Packaging

In terms of packaging, there are major differences between IBM InfoSphere Warehouse and Oracle offerings that are illustrated in figure 3.

INFOSPHERE WAREHOUSE ENTERPRISE EDITION	ORACLE DATABASE 11g ENTERPRISE EDITION
Included: • Partitioning • Deep Compression • Cubing Services • Intelligent Miner • Management tools List price per TB: $70,000	Separate charge: • Partitioning • Advanced Compression • OLAP • Data Mining • Diagnostics & Tuning Packs List price per user: $2,530 List price per processor: $126,500
INFOSPHERE WAREHOUSE DEPARTMENTAL EDITION	**ORACLE DATABASE 11g STANDARD EDITION**
Included: • Partitioning • Deep Compression • Cubing Services • Intelligent Miner • Management tools List price per TB: $35,000	Not supported: • Partitioning • Advanced Compression • OLAP • Data Mining Separate charge: • Diagnostics & Tuning Packs List price per user: $550 List price per processor: $27,500

Figure 3: IBM InfoSphere Warehouse and Oracle Data Warehouse Packaging

Both the Enterprise and Departmental editions of InfoSphere Warehouse include the DB2 9.7 database as well as partitioning, compression, OLAP, data mining, management and other modules as part of the basic package.

Oracle charges separately for equivalent modules. While Oracle Database 11g Enterprise Edition is list priced at $950 per named user and $47,500 per processor, inclusion of these modules pushes overall list prices for a full-function data warehouse solution to $2,530 and $126,500 respectively.

Technology

Technology differences between IBM DB2 9.7 and Oracle Database 11g affect comparative costs. DB2 9.7 is generally more efficient in its use of system resources.

Differences in data compression capability are particularly significant. Compression allows database size to be significantly reduced, and savings may be realized in processor capacity, main memory, disk and tape storage, and I/O and network bandwidth. Replication, backup and recovery, and other data movement processes may also be accelerated.

DB2 9.7 features one of the industry's most effective across-the-board implementations of compression. It extends to rows as well as indexes, temporary tables, log files, large objects, XML data and other structures. Users have routinely experienced overall compression levels of from 55 percent to more than 85 percent.

Savings in server, storage and network costs in the 20 to 40 percent range have been widely documented among DB2 9.7 users. Some organizations have achieved higher levels.

In comparison, users employing Oracle Database 11g Advanced Compression have found that, while high levels of compression may be realized, unacceptable performance degradation tends to occur at an early stage.

DB2 9.7 compression is synergistic with IBM's new per terabyte pricing. An organization that is able to compress, say, 2TB of user data to 1TB would enjoy significantly lower InfoSphere Warehouse license and support costs. Reductions of two to three times are realistic.

Users have also found that extensive integration and automation result in levels of database administrator (DBA) productivity that are typically higher than for Oracle environments. DBA staffing and personnel costs may thus be lower.

Cost Calculations

The calculations presented in this report are for representative data warehouse installations in large and midsize organizations. Data warehouse sizes are for uncompressed data, while numbers of users are for systems that have been in place for two or more years.

Calculations are based on vendor list prices. Oracle offers two pricing schemes, based on number of named users or number of processors. Costs for both Oracle pricing schemes – named users and numbers of processors – as well as for InfoSphere Warehouse per terabyte pricing are shown in figures 4 and 5.

In the comparative cost calculations summarized in figures 1 and 2, Oracle Database 11g Enterprise Edition costs are based on the lowest-cost option for each installation.

Oracle Database 11g Enterprise Edition costs are adjusted for Oracle core factors – for pricing purposes, the company counts Intel cores as 0.5 processor and IBM POWER7 cores as one processor each. Per user calculations allow for the Oracle pricing minimum of 25 users per core. For the 1TB and 5TB installations, minimum numbers of users are higher than actual levels.

IBM InfoSphere Warehouse pricing is based on the next largest 1TB increment. In figures 4 and 5, license and support costs for 0.5TB, 2.5 TB and 12.5TB of compressed user data are rounded to increments of 1TB, 3TB and 13TB respectively.

DATA WAREHOUSE			
User data	**1TB**	**2TB**	**5TB**
Number of users	60	150	400
Number of processors	2/8 x Intel 5600	2/12 x Intel 5600	2/16 x Intel 7500
ORACLE DATABASE 11g ENTERPRISE EDITION			
Minimum number of users	100	150	500
Per user license cost	$ 253,000	$ 379,500	$ 1,265,000
3-year cost including support	$ 419,980	$ 629,970	$ 2,099,900
Per processor license cost	$ 506,000	$ 759,000	$ 1,012,000
3-year cost including support	$ 839,960	$ 1,259,940	$ 1,679,920
IBM INFOSPHERE WAREHOUSE DEPARTMENTAL EDITION			
UNCOMPRESSED DATA			
Per terabyte license cost	$ 35,000	$ 70,000	$ 175,000
3-year cost including support	$ 56,000	$112,000	$ 280,000
50% DATA COMPRESSION			
Per terabyte license cost	$ 35,000	$ 35,000	$ 105,000
3-year cost including support	$ 56,000	$ 56,000	$ 168,000

Figure 4: Comparative Cost Calculations: 1TB to 5TB Installations

DATA WAREHOUSE			
User data	**10TB**	**25TB**	**50TB**
Number of users	1,000	3,000	5,000
Number of processors	2/16 x 3.2 GHz IBM POWER7	4/32 x 3.86 GHz IBM POWER7	8/64 x 4 GHz IBM POWER7
ORACLE DATABASE 11g ENTERPRISE EDITION			
Minimum number of users	400	850	1,600
Per user license cost	$ 2,530,000	$ 7,590,000	$ 12,650,000
3-year cost including support	$ 4,199,800	$ 12,599,400	$ 20,999,000
Per processor license cost	$ 2,024,000	$ 4,048,000	$ 8,096,000
3-year cost including support	$ 3,359,840	$ 6,719,680	$ 13,439,360
IBM INFOSPHERE WAREHOUSE ENTERPRISE EDITION			
UNCOMPRESSED DATA			
Per terabyte license cost	$ 700,000	$ 1,750,000	$ 3,500,000
3-year cost including support	$ 1,120,000	$ 2,800,000	$ 5,600,000
50% DATA COMPRESSION			
Per terabyte license cost	$ 350,000	$ 910,000	$ 1,750,000
3-year cost including support	$ 560,000	$ 1,456,000	$ 2,800,000

Figure 5: Comparative Cost Calculations: 10TB to 50TB Installations

Additional Information

This ITG Status Report is based upon the preliminary results and methodology for an upcoming Management Brief to be released by the International Technology Group. For copies of this Management Brief, please email requests to Contact@ITGforInfo.com.

International Technology Group
609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: (831) 427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com