



Microsoft Analytics Platform System Delivers Best TCO-to-Performance

Published by: Value Prism Consulting

Sponsored by: Microsoft Corporation

Publish date: May 2016

Abstract: Data Warehouse with large data workloads specific appliances may be difficult to compare, given that each of the competing solutions come preconfigured with their proprietary compute and storage configurations, traditional databases and open source file systems, and other varying specifications. Value Prism Consulting, a management consulting firm, was engaged by Microsoft® Corporation to review and contrast data warehouse offerings from five leading vendors. In this updated whitepaper aimed at IT decision makers, the firm compared each vendor's appliances based on publicly-available costs and specification data. On a TCO-to-performance scale, Microsoft Analytics Platform System was seen as the most cost-effective appliance providing high performance and great value.

Disclaimer

Every organization has unique considerations for economic analysis, and significant business investments should undergo a rigorous economic justification to comprehensively identify the full business impact of those investments. This analysis report is for informational purposes only. VALUE PRISM CONSULTING MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

©2016 Value Prism Consulting, LLC. All rights reserved.

Product names, logos, brands, and other trademarks featured or referred to within this report are the property of their respective trademark holders in the United States and/or other countries.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this report may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this analysis. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this analysis does not give you any license to these patents, trademarks, copyrights, or other intellectual property of Microsoft.

CONTENTS

Executive Summary 1

Introduction..... 3

Total Cost of Ownership 4

 Appliance Costs 5

 Installation Costs 6

 Maintenance and Support Costs 7

 Labor Costs 7

 Facilities Costs 8

TCO-to-Performance Comparisons..... 8

 TCO per Terabyte (TB) of Compressed User-Available Storage 10

 TCO per Terabyte (TB) of Uncompressed User-Available Storage 11

 TCO per Database Core 11

 TCO per Gigabyte (GB) of Memory 12

 TCO per I/O Bandwidth 12

Microsoft Benefits 13

Conclusion 14

Appendix..... 15

EXECUTIVE SUMMARY

The top four Enterprise Data Warehouse (EDW) management software competitors in terms of market share – Oracle, IBM, Microsoft, and Teradata – represented 85.3% of the overall data warehouse management software revenues, with Microsoft showing the largest growth amongst these at 9.7% (2013-2014).ⁱ By way of this intense competition and with vendors offering preconfigured hardware plus software appliance solutions, acquiring a data warehouse platform has become less expensive and easier to install and manage.

This whitepaper is an update to Value Prism Consulting's 2014 data warehouse price-to-performance whitepaper,ⁱⁱ and is meant to be an aid to organizations' IT decision makers looking to compare and contrast similar appliances from these leading vendors. As data warehouse market continues to evolve, this update focuses solely on price-to-performance comparisons of the five year Total Cost of Ownership (TCO) of EDW appliances. To provide fairground comparison non-relational data solutions based on Hadoop are no longer discussed in this update. Other notable changes include Pivotal (EMC joint venture) open sourcing Greenplum in 2015 followed by Dell announcement to acquire EMC.

Appliance offerings from Pivotal, IBM, Microsoft, Teradata, and Oracle were reviewed and compared. TCO-to-performance¹ comparisons have been collected and summarized across each vendor – two based on storage (compressed and uncompressed user-available storage, as a factor of total costs) and two based on hardware compute power (number of cores and amount of standard memory in GB, again as a factor of total costs). In Figure 1 results closer to the center show lower cost per specification. From Figure 1, in all four cases Microsoft has the lowest ratio, showing they are a high-performing and economic data warehousing appliance. Oracle was, in most cases, not just last, but much more expensive based on both total cost and performance ratios.

Care should always be taken in assessing the best solution for your situation. This comparison is based on publicly available costs and specification metrics. Individual vendors offer different discounts and volume price breaks, so results may be different than the ones listed here.

¹ Refers to the five-year net present value (NPV) of the overall costs of ownership. This includes one-time appliance hardware costs, software license purchases, and installation costs; and annual maintenance and support agreements, and management labor costs.

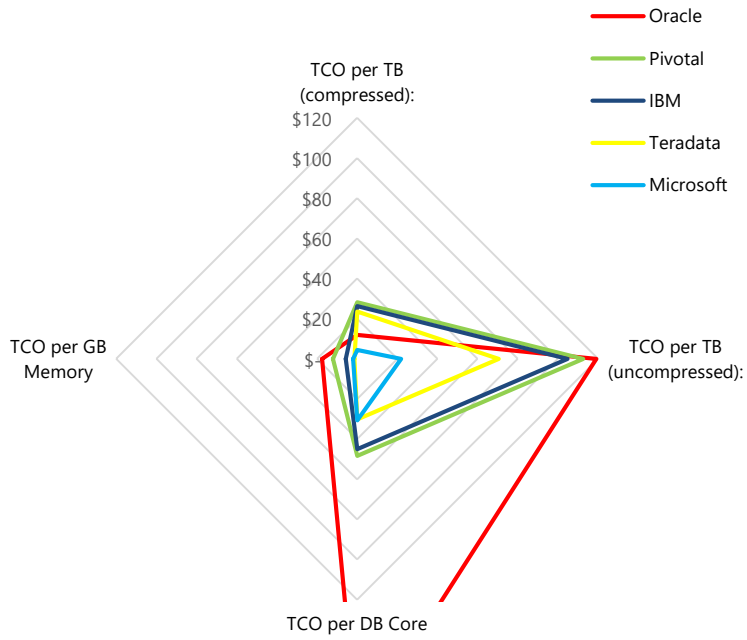


Figure 1: TCO-to-Performance Ratios across Multiple EDW Platforms (Costs in U.S. dollars, in Thousands)

INTRODUCTION

Enterprise Data Warehouse (EDW) solutions currently represent a multi-billion dollar market with the business analytics software market expected to grow at an 8.0% CAGR over the next four years.ⁱⁱⁱ Data warehouse solutions focus on storage and analysis of large structured, time-variant, and non-volatile operational data. The choice of optimal platform will be influenced by an organization's business and technological needs. Most vendors now offer EDW solutions as a pre-configured, pre-optimized, and often single-priced appliances, which include:

- Hardware components required to run the appliance, including the box, disk drives, memory, network connectivity, and processors.
- Software required to run the appliance, including server operating system, database software, and data management tools, and
- Installation services for system and software.

By-and-large, purchase of the appliance is a simple process as a fixed set of software and hardware – if not a single SKU, at least a short list of scalable hardware and software modules. Customers can often pick an appliance and expect it will be nearly ready to plug and play with much less setup and configuration than a custom or build-it-yourself solution, which could take many months.

However with the appliance model, it has also become harder to compare and contrast similar solutions, especially when vendors make various claims in each of their public datasheets. IBM says it is *“a low cost option”* with *“low total cost of ownership.”* Pivotal claims their appliance has an *“industry leading TCO.”* Teradata says they offer *“the best price for performance platform in the marketplace.”* And, Oracle claims that their appliance has *“a low overall total cost of ownership.”*

In this study commissioned by Microsoft®, several primary EDW appliances, as listed in the sidebar, have been reviewed, summarized, and compared. Each vendor provides via its Website an appliance datasheet that has been used as the primary source for specification data (such as storage, cores, etc.). List pricing and other annual costs details are cited specifically, and are also taken from public sources. For each vendor, one leading appliance (if they offer more than one) was selected for comparison. Full rack pricing and specifications were used to ensure consistent comparison.

Data Warehouse Appliances included:

- Microsoft Analytics Platform System (hardware based on HPE CS300)^{iv}
- Oracle Exadata Database Machine X4-2^v
- Pivotal (formerly EMC) DCA v2 with Greenplum Database (GPDB) Standard Modules^{vi}
- IBM PureData System for Analytics N2001-10^{vii}
- Teradata Data Warehouse Appliance 2750^{viii}

To make a more accurate comparison, these solutions were also measured against several TCO-to-performance ratios:

- TCO per terabyte of compressed and uncompressed user space is used as a storage value approximation that can help provide more comparison details when the amount of user space and compression ratios are not the same across all appliances.
- TCO per database core and gigabyte of memory were also used, as more compute performance-related metrics.

TOTAL COST OF OWNERSHIP

The total cost of ownership (TCO) over a five-year period was calculated to include both initial investments in hardware, software licenses, and installation; and annual investments in appliance maintenance and support, facilities – power, cooling, and space, and labor costs associated with management of the databases. The TCO summary is presented as a five year NPV of costs, which assumes a discount rate of 10%.

NOTE: This comparison is based on the list price and publicly available cost and specification metrics. Each unique customer discount situation will be different, and more information (and customization) is likely available from each vendor so the comparison may be different than the one using only list prices and specifications. All prices are listed in U.S. dollars and rounded to the nearest hundreds for sub-categories and nearest thousands for totals.

Costs	Microsoft	Oracle	Pivotal	IBM	Teradata
Appliance	\$1,599,300	\$18,810,000	\$6,385,300	\$2,235,000	\$2,798,200
Installation	\$12,200	\$10,500	\$11,900	(Included)	\$8,000
Total one-time costs	\$1,612,000	\$18,487,000	\$6,397,000	\$2,235,000	\$2,806,000
Maintenance + Support	\$471,600	\$4,116,600	\$1,370,400	\$345,000	\$502,900
Labor	\$127,400	\$319,800	\$137,800	\$378,300	\$444,600
Facilities	\$13,800	\$22,100	\$17,600	\$14,200	\$17,300
Total annual costs	\$613,000	\$4,459,000	\$1,526,000	\$738,000	\$965,000
5 Year NPV of Costs	\$3,936,000	\$35,724,000	\$12,394,000	\$5,033,000	\$5,806,000

Table 1: TCO summary for EDW Appliances (Costs in U.S. dollars)

Appliance Costs

The total price of each full rack appliance includes cost of hardware and software and is based on publicly-available information directly from the vendor, from a reseller that has listed appliance pricing, or if necessary from news or blog articles that have published price estimates. Total retail price for each appliance, along with the pricing source, is listed below.

Vendor	EDW	Price Source and Assumptions
Microsoft	\$1,599,300	Vendor furnished pricing data for the whitepaper.
Oracle	\$18,810,000	Oracle's hardware ^{ix} and software ^x price lists.
Pivotal	\$6,597,300	NASPO ValuePoint EMC price list ^{xi} for hardware and vendor 2012 price list for software.
IBM	\$2,235,000	Gemini licensing website, an IBM partner (includes install and 12 months support). ^{xii}
Teradata	\$2,139,800	Teradata's Pricing Brochure (unspecified support time plus installation are included). ^{xiii}

Data Warehouse Appliance Costs

The price estimates for EDW appliances are derived based on the following sources and assumptions:

- Microsoft's pricing is based on retail software pricing for Microsoft SQL Server PDW, Microsoft Windows Server 2012 Standard, and System Center 2012 Standard licenses totaling to U.S. \$897,300. The hardware appliance cost for a HPE CS300 full rack appliance with 3TB hard drives is approximately U.S. \$702,000. The pricing was furnished by the vendor for the purpose of this whitepaper. Customers also have an option of purchasing APS hardware from Dell and Quanta Cloud Technology at different prices.
- Oracle provides an Engineered Systems price list^{ix} covering the appliance hardware, which is priced at U.S. \$1,098,000 for the full rack. The Exadata X4-2 Datasheet^v identifies a number of software licenses required to run the appliance. This list was used along with the software price list^x to total up software license pricing, as show in the sidebar table.
 - Oracle lists software license prices per core, and then references the *Oracle Processor Core Factor Table*^{xiv} to adjust the number of cores based on the specific processor type/family to be considered for software license purposes. The core factor ranges from 0.25 to 1.0; and is 0.5 for Xeon processors that is used in the Exadata machine.
 - The Oracle Exadata X4-2 appliance includes 192 database cores, so the total software license cost is calculated as: U.S. \$163,500 x 192 x 0.5 = U.S.

Table 3: Appliance Costs, includes Hardware and Software (Costs in U.S. dollars)

Oracle Database Software and Add-ons Pricing:

Oracle Database EE	\$47,500
Real Application Clusters	\$23,000
Active Data Guard	\$11,500
Advanced Compression	\$11,500
Advanced Security	\$15,000
Real Application Testing	\$11,500
Diagnostic Pack	\$7,500
Tuning Pack	\$5,000
Partitioning	\$11,500
Data Integrator EE	\$23,000
Total Software Cost per Core	\$167,000

Oracle pricing, as listed in its Oracle Technology Global Price List, provides details on each software application and add on. The column for Processor License is listed above (though actually the pricing is per-core). The above per-core price is multiplied by the number of cores, which itself is adjusted based on the Oracle Processor Core Factor Table.

\$15,696,000. Additionally, customers need to purchase Exadata Server Software, which is priced per disk for the 168 disk drives in the X4-2 appliance, totaling: U.S. \$10,000 x 168 = U.S. \$1,680,000.

- Pivotal does not publish pricing information; however, a publicly available NASPO ValuePoint price list for EMC Corporation^{xi} lists the pricing for Pivotal DCA v2 standard module at U.S. \$408,100, which gives us a total hardware price of 4 modules per full rack x U.S. \$408,100 = U.S. \$1,632,400. The Greenplum Database (GPDB) software is priced per ingested TB, or as defined by Pivotal, “the amount of TB that would be filled by uncompressing all the data contained in all tables in a GPDB system.”^{xv} A 2012 price list indicates that the GPDB perpetual license is tiered per ingested TB range, starting at U.S. \$30,000 per TB for the first 14TB and lowers down to U.S. \$4,500 per TB beyond 1,000TB+. Using this table the effective software price was calculated at U.S. \$4,964,900 for the 440TB of compressed available user data.
- IBM also does not publish pricing information; however, Gemini, an IBM partner, sells the IBM PureData System for Analytics N2001-10 Appliance and lists IBM List pricing as shown above.^{xii} While Gemini’s site lists slightly discounted pricing, the IBM List retail pricing has been used since all other vendor appliances are listed with retail pricing. Gemini’s pricing does state that 12 months of support plus installation services are included; it is assumed this is also included in IBM’s list price.
- Teradata does provide pricing, but only very generally. For the Teradata Data Warehouse Appliance 2750 (the 2000 Series), Teradata lists a price of U.S. \$26,000 per TB of uncompressed user available storage,^{xiii} which gives us U.S. \$26,000 x 82.3TB = U.S. \$2,139,800.

Installation Costs

Appliance installation are fixed costs of system installation and software configuration provided as a SKU by the vendor. Some vendors, e.g. IBM and Teradata, have the installation costs combined as part of the appliance costs.

Vendor	EDW	Price Source and Assumptions
Microsoft	\$12,200	Vendor furnished pricing data for the whitepaper.
Oracle	\$10,500	Fact Point Group whitepaper ^{xvi}
Pivotal	\$11,900	GP DCA Installation from NASPO EMC price list. ^{xi}
IBM	\$0	Included in appliance costs.
Teradata	\$8,000	Implementation Services and System Installation at \$8,000 per rack. (Source: Teradata SME)

Table 4: Appliance Installation Costs
(Costs in U.S. dollars)

NOTE: Installation is only part of the overall deployment cost considerations. Customers should carefully evaluate overall project management, migration, and testing costs, which are highly variable and may not be similar across platforms. For example, Microsoft APS customers can expect their appliances to be ready to run in 20 to 40 days (based on SMEs and Microsoft partner feedback), compared to a typical deployment period of 4-7 months for other vendors, such as Oracle (5 months for half rack Exadata + Exalogic appliance^{xvii}) and Teradata (196 days average for data warehouse appliances^{xviii}). Although not calculated for the purpose of this whitepaper, these delays in deployment also translate to significant lost opportunity costs where the appliances could have been put to use in many revenue generating BI activities.

Maintenance and Support Costs

Most vendors require customers to pay an annual hardware and software support fee as a percentage of the initial acquisition costs. The table below summarizes these costs along with the source and assumptions:

Vendor	EDW	Price Source and Assumptions
Microsoft	\$471,600	Typical Software Assurance fee at 25% of total license costs + \$40,700 and \$8,200 Microsoft Premier Support at the most for full rack APS appliance. Numbers are average considering 5% YoY increment. (Source: Vendor)
Oracle	\$4,116,600	Software Update and License Support fee at 22% of total license costs ^x + Premier Support for Systems at 12% of hardware costs ^{ix} + Premier Support for OS at 8% of hardware costs. ^{ix}
Pivotal	\$1,370,400	Annual maintenance fee at 23% of perpetual license costs (Source: 2016 Price List) + Annual support fee at 14% of hardware cost (Source: Partner interview)
IBM	\$345,000	Annual Appliance Maintenance + Subscription and Support Renewal at 15% of appliance costs. ^{xii} Error! Bookmark not defined.
Teradata	\$502,900	Annual Software Subscription at 6% and Premier Appliance Support at 14-21% of appliance net selling price. (Source: Teradata SME)

Labor Costs

Appliances offer a simplified architecture requiring little maintenance; still a significant amount of time is invested in building data models, managing and tuning databases, and other related database administration tasks. Full time staffing for database administration is significantly lower in case of Microsoft Analytics Platform System, since highly skilled SQL Server resources are readily available at wages that are typically 10%-15% lower compared to DBAs for other database technologies. For IT departments, being able to use existing SQL Server skills to implement and manage a solution on the APS is a big advantage.

“With [Microsoft] Analytics Platform System you rely on commodity hardware and proven SQL Server database technology. The advantages are two fold – one you lower the risk of getting skilled people, two you lower the costs of getting people. Moreover you have the ability to do things yourself rather than rely on more expensive professional services or consultants.” – Managing Director for data services partner in Asia-Pacific region.

Table 5: Appliance Maintenance and Support Costs (Costs in U.S. dollars)

The table below summarizes the annual DBA labor costs. A 30% gross up for taxes and benefits was added to all annual wages to estimate a fully burdened labor cost to the organization.

Vendor	EDW	Price Source and Assumptions
Microsoft	\$127,400	1 FTE per rack at \$98k ^{xix} annual salary for DBAs with SQL Server skills and \$97k ^{xx} annual salary for DBAs with T-SQL skills.
Oracle	\$319,800	2 FTEs per rack ^{xi} at \$123k ^{xxi} annual salary for DBAs with Oracle Exadata skills.
Pivotal	\$137,800	1 FTE per rack at \$106k ^{xxii} annual salary for DBAs with PostgreSQL skills.
IBM	\$378,300	3 FTEs per rack at \$97k ^{xxiii} for DBAs with PureData / Netezza skills.
Teradata	\$444,600	3 FTEs per rack at \$114k ^{xxiv} for DBAs with Teradata skills.

Table 6: Administration Labor Costs
(Costs in U.S. dollars)

Facilities Costs

Facilities include power, cooling, and floor space requirements for the appliance. All vendors list out the typical power consumption either in KW, W, or VA; cooling requirements in BTU/hour; and floor space in inches. When power was expressed in VA, it was converted to W by multiplying with a unity (= 1) power factor. Further the cooling requirements of BTU/hour was converted to W by multiplying with a factor of 0.293. The effective energy costs were calculated using the U.S. average retail price of 10.2 cents per KWH (kilowatt-hour).^{xxiii} For all floor space calculations, an additional 3 feet clearance was assumed for front and rear panel access. A real-estate average rate of U.S. \$23.23/sq.ft./year was used for calculation of floor space costs.^{xxiv} The total facilities costs are summarized below:

Vendor	EDW	Price Source and Assumptions
Microsoft	\$19,200	10,508W; 35,832BTU/hr; 24" x 42" (Source: Vendor)
Oracle	\$22,100	12,100W; 41,300BTU/hr; 23.62" x 47.24" ^v
Pivotal	\$17,600	9,600W; 32,750BTU/hr; 24" x 42" ^{vi}
IBM	\$14,200	7,500W; 27,000BTU/hr; 25.51" x 40" ^{vii}
Teradata	\$17,300	10,300W; 29,100BTU/hr; 24" x 47" ^{viii xxv}

Table 7: Appliance Facilities Costs,
includes power, cooling, and floor
space (Costs in U.S. dollars)

TCO-TO-PERFORMANCE COMPARISONS

The tables in the following pages summarize key specifications of the Enterprise Data Warehouse appliances as obtained from their data sheets, along with the five-year Total Cost of Ownership (TCO).

As seen in the tables, the appliances have varying numbers of drives, cores, and user storage, so simply looking at total costs is not a reasonable comparison. To provide additional comparison perspectives, scaled metrics have been calculated and presented. The first two are related to total appliance size, the other two are related to appliance power and performance. The first (TCO per terabyte of compressed, user-available storage) is the standard data warehouse comparison metric.

Specifications	Microsoft	Oracle	Pivotal	IBM	Teradata
Total Cores	128	192	256	112	192
Memory (GB)	2,048	2,048	1,024	896	4,096
Compression	5 to 1	10 to 1	4 to 1	4 to 1	3 to 1
Usable Storage Uncompressed (TB)	203	300	110	48	82
Usable Storage Compressed (TB)	900	3,000	440	192	247
I/O Bandwidth (Scan Rate) (GB/s)	36	100	40	31	Unknown
5 Year NPV of Costs	\$3,936,000	\$34,909,000	\$11,966,000	\$5,442,000	\$7,579,000

Table 8: Data Warehouse Appliance Specifications and 5 Year NPV of Costs (Costs in U.S. dollars)

Performance Metrics:

- **TCO per terabyte of compressed user-available storage**, as a measure of the value of useable storage space. This includes any compression factors that each vendor provides – compression shrinks the amount of space data takes, allowing organizations to make better use of existing hardware so that more data can be stored on each drive and I/O transactions are smaller.
- **TCO per terabyte of uncompressed user-available storage**, as an additional measure of the value of useable storage. Since each organization provides different compression technologies and recommendations (and each organization decides on their own compression factor), this allows a comparison regardless of compression.
- **TCO per database core**, as one measure of performance. As more cores provide capacity for more concurrent database and compute tasks.
- **TCO per gigabyte of memory**, as a second measure of performance. Since memory is able to store information in a quickly-accessible place, more memory means more information can be stored without having to write to a disk, which means transactions can be processed more quickly.

Microsoft APS provides the best TCO-to-performance ratio across all above metrics. Oracle’s appliance was seen as the worst performing across all except the TCO per

terabyte compressed metrics mainly because of the high appliance software and annual software subscription costs. The TCO per terabyte compressed was offset by the 10x compression ratio for Oracle data warehouse appliance compared to 3x – 5x for the remaining appliances. A detailed breakdown of these metrics is discussed below.

TCO per Terabyte (TB) of Compressed User-Available Storage

The estimate of reasonable compression within the data warehouse appliance varies across vendors. Microsoft appliance by HPE based on the 1TB drives claims a 5:1 compression ratio.^{iv} Similarly, Pivotal^{vi} and IBM^{vii} data sheets suggest a compression ratio of 4:1 for their appliances. Oracle lists a 5x – 20x, with 10x typical compression based on proprietary compression processes.^v An efficient compression (but not detrimental to performance) of 10:1 for Oracle is used in this whitepaper.

Teradata does not list a specific compression metric in their datasheet specs and marketing materials, but in its pricing datasheet it lists a compression range of “50% to 80%,” or in ratio form, from 2:1 to 5:1.^{xiii} Various other sources including Teradata forum,^{xxvi} technical whitepapers,^{xxviii} and news articles^{xxvii} suggest a range of 1.7:1 to 5.7:1. As an average, an efficient (i.e., that won’t likely overly impact performance) compression of 3:1 is used in this whitepaper. The compression estimates for each vendor are listed in the table to the right.

For TCO-to-compressed storage for full rack appliances, Microsoft APS has the lowest ratio at a little less than U.S. \$4,373 per TB. That is nearly three times better than the next option and Oracle, and five times better than the average. Oracle’s high software price is greatly offset by its 10:1 compression ratio (at least twice as high as any other appliances), but is still a more expensive option on a TCO per user-available storage (compressed) comparison to Microsoft.

Data Warehouse Appliance Compression Ratios:

Vendor	Ratio
Microsoft	5:1
Oracle	10:1
Pivotal	4:1
IBM	4:1
Teradata	3:1




Exhibit 1: Summary of TCO per TB of Compressed Usable Storage (Costs in U.S. dollars)

TCO per Terabyte (TB) of Uncompressed User-Available Storage

Uncompressed user-available storage for a full rack appliance is included to provide perspective of total useful storage provided outside of compression. For TCO-to-compressed storage, Microsoft data warehouse appliance again outdistanced the competition with a much better TCO-per-terabyte ratio, about three to six times better than the competing appliance configurations.




Exhibit 2: Summary of TCO per TB of Uncompressed Usable Storage (Costs in U.S. dollars)

TCO per Database Core

Database processors and cores provide the engines for running data management and analysis tasks. With more cores, database applications and virtual machines run more efficiently, meaning that databases can handle a large user load and manage new data collection, while still providing high performance results. So more is better, but with appliances that can scale by adding more modules or racks, even more important is the TCO per core. Note that the number of cores is not exactly correlated with higher price – some licenses are not required for every core, but TCO-per-database core provides a view into the performance of each appliance as a factor of total costs. When viewed in relation to TCO for full rack appliances, Microsoft and Teradata data warehouse appliances lead the pack, nearly six times better than Oracle.




Exhibit 3: Summary of TCO per Server Core (Costs in U.S. dollars)

TCO per Gigabyte (GB) of Memory

Like cores, the amount of memory included in a full rack appliance is a significant indicator of potential performance. Large amounts of memory allow for faster processing of data by keeping more information in memory to be processed, instead of having to make read/write calls to the hard drive. Memory is much more expensive than disk drives (for equal units of storage), so in addition to the storage ratios above, it is important to include TCO-per-memory ratio comparison as well. Teradata and Microsoft data warehouse appliances provide better value over others, at six times better than average, and nine times better than Oracle, based on complete TCO.




Exhibit 4: Summary of TCO per GB Memory (Costs in U.S. dollars)

TCO per I/O Bandwidth

Many of the data warehouse vendors now offer Solid State Drives (SSDs) and other higher performance storage caches with their appliances, and tout their scan rates as a measure of high storage performance. While I/O Bandwidth data was only available for 4 of 5 vendors, and thus not considered in the data warehouse TCO-to-performance overview above, it is an important enough performance metric to include and discuss, as an indicator of throughput performance capabilities. Added to the point that Microsoft Analytics Platform System is rarely I/O bound and that the SQL Server Buffer Manager is very efficient at keeping hot data in RAM, the data warehouse platform also showed two times better TCO-to-performance ratio compared to the average for all vendors.




Exhibit 5: Summary of TCO per I/O Bandwidth in GB/s (Costs in U.S. dollars)

MICROSOFT BENEFITS

Based on this comparison, Microsoft's Analytics Platform System (APS) was seen to provide the most cost-effective solution in terms of total cost, as well as being the top performer in the TCO-to-performance metrics comparison. In addition to being cost-effective, APS offers several other advantages:

- High performance, true scale-out, MPP data warehouse architecture with SQL Server PDW.
- Reduced storage footprint and costs with Windows Server Storage Spaces and SQL Server In-Memory Columnstore.
- Lowest number of database vulnerabilities and greater uptime.
- Industry standard hardware choice options with software defined storage from multiple vendors: Quanta, HPE, and Dell.

When preparing a business case or determining the Return on Investment (ROI) for data warehouse investments, it is important not only to look at the TCO and direct cost savings, but also any indirect benefits that the appliance provides. Some of the indirect benefits achieved from Microsoft Analytics Platform System are discussed below.

Query Performance

Improvements in query performance directly translates to greater employee productivity and time spent on revenue generating activities. The combination of MPP technology and In-Memory Columnstore in Microsoft Analytics Platform System allows employees to achieve up to 100x faster query performance compared to traditional SMP data warehouse solutions.

Security Improvements

A recent report by Ponemon Institute and Symantec indicates that the organizational cost of data breaches averaged U.S. \$4 million per year.^{xxviii} The losses may include costs related to legal and penal fee, revenue loss due to lost customers, and loss of competitive advantage. The high cost of breach thus puts a greater emphasis on having a secure data appliance that protects the business' data against common threats to IT systems.

Featuring SQL Server, Microsoft APS customers are more protected against database vulnerabilities. When it comes to high severity vulnerabilities due to software flaws, the National Vulnerability Database reports that Oracle database averages 20 vulnerabilities per year compared to SQL Server's less than 1 per year.^{xxix}

CONCLUSION

Overall, Microsoft not only provides competitively priced data warehouse compared to other vendor options, but also the best TCO-to-performance value.

Query Performance Customer

Evidence:

- Reduce complex query execution time from 2 days to 7 minutes. Improve employee productivity by 40%. – *Royal Bank of Scotland*
- Realize 10x performance improvement with SQL Server. – *SBI Liquidity Market*
- Create reports 50% to many times faster when compared to previous system. – *Crossmark*

APPENDIX

Methodology

In a study commissioned by Microsoft, similar Enterprise Data Warehouse (EDW) appliances from five leading vendors have been reviewed, summarized, and compared, based on publicly-available costs and specification information. The five vendors compared in this study were:

- Microsoft
- Oracle
- Pivotal (formerly EMC)
- IBM
- Teradata

These solutions were chosen as some of the top vendors and appliances discussed in the December, 2015 Forrester Wave report on enterprise data warehousing platforms.^{xxx}

High Capacity appliance options were selected for consistent comparison. Full racks were compared, to ensure standard comparisons across large-capacity appliances.


Summary metrics, total costs, total user storage space (compressed and uncompressed), and performance were compared. TCO per terabyte of user space is used as a price/value approximation that can help provide more comparison details when the amount of user space is not the same across all appliances.

Comparisons are based on the list price and publicly available cost and specification metrics. In addition to the appliance price, which is only part of the cost considerations, the overall TCO analysis included facility, installation, maintenance, and support costs.

Each discussion with vendors and their customer discount situation will be different, and more information (and customization) is likely available from the vendors. Individual results may thus be different than the ones listed here. Customer should carefully evaluate installation, migration, and on-going management costs specific to their solution requirements before making a decision.

About Value Prism Consulting

Value Prism Consulting is a valuation and management consulting firm that provides services to a broad range of clients worldwide. Our Management Consulting Practice assists clients with business case development and decision support analysis. Our solutions measure the results of various options related to business process improvement, capital and major budget-spending decisions, including infrastructure upgrades, and the value of new product introductions. Visit www.valueprism.com for more information.



Sources

- ⁱ IDC. (2015, August). Worldwide Business Analytics Software Market Shares, 2014 <http://bit.ly/25uypc1>
- ⁱⁱ Value Prism Consulting. (2014, September). *Microsoft Analytics Platform System Delivers Best TCO-to-Performance*. <http://bit.ly/1UT2zSI>
- ⁱⁱⁱ IDC. (July 2015). *Worldwide Business Analytics Software Forecast, 2015–2019* <http://bit.ly/1MQAUtv>
- ^{iv} HPE. (2015, June 1). *Quick Specs – HPE ConvergedSystem 300 for Microsoft Analytics Platform System*. <http://bit.ly/1TSrogB>
- ^v Oracle. (2013). *Datasheet – Oracle Exadata Database Machine X4-2*. <http://bit.ly/19Jf87k>
- ^{vi} Pivotal. (2015). *Datasheet – Pivotal Data Computing Appliance*. <http://bit.ly/1qOCN63>
- ^{vii} IBM. (2013, January). *Datasheet – IBM PureData System for Analytics N2001*. <http://bit.ly/25Yewds>
- ^{viii} Teradata. (2013). *Datasheet – Teradata Data Warehouse Appliance 2750*. <http://bit.ly/1qtBSQz>
- ^{ix} Oracle. (2016, April 5). *Oracle Engineered Systems Price List*. <http://bit.ly/1v2nk1a>
- ^x Oracle. (2016, March 22). *Oracle Technology Global Price List*. <http://bit.ly/1dSoiWw>
- ^{xi} Naspo.org. (2015, March). *NASPO Value Point – EMC Corporation*. <http://bit.ly/1RSwhFa>
- ^{xii} Gemini eStore. (2016, April 12). *PureData System Anal N2001-010*. <http://bit.ly/1ABW2Pp>
- ^{xiii} Teradata. (2016). *Teradata Workload-Specific Platform Pricing*. <http://bit.ly/1SeaBp0>
- ^{xiv} Oracle. (2015, October 27). *Oracle Processor Core Factor Table*. <http://bit.ly/1jfqnxl>
- ^{xv} Pivotal. (2016, March 22). *Pivotal Product Guide*. <http://bit.ly/1tvGvMU>
- ^{xvi} The Fact Point Group. (2012, October). *Cost Comparison for Business Decision Makers: Oracle Exadata Database Machine vs. IBM Power Systems*. <http://bit.ly/1whjJxN>
- ^{xvii} Forrester. (2013, September). *The Total Economic Impact of Oracle Exadata and Oracle Exalogic*. <http://bit.ly/1tsy85D>
- ^{xviii} International Technology Group. (2014, May) *Cost/Benefit Case for IBM PureData System for Analytics: Comparing Costs and Time to Value with Teradata Data Warehouse Appliance*. <http://bit.ly/20BJWSY>
- ^{xix} Payscale.com. (Accessed 2016, April 12). *Average Salary for Skill: Microsoft SQL Server*. <http://bit.ly/1BL6wO6>
- ^{xx} Payscale.com. (Accessed 2016, April 12). *Average Salary for Skill: Transact-SQL*. <http://bit.ly/YJMeFo>
- ^{xxi} Payscale.com. (Accessed 2016, April 12). *Average Salary for Skill: Oracle Exadata*. <http://bit.ly/1pRlw7X>
- ^{xxii} Payscale.com. (Accessed 2016, April 12). *Average Salary for Skill: PostgreSQL*. <http://bit.ly/1sbDrVg>
- ^{xxiii} Energy Information Administration's (EIA). (2016, January). *Table 5.6.B. Average Retail Price of Electricity to Ultimate Customers by End-Use Sector*. <http://1.usa.gov/1pmRSYc>
- ^{xxiv} Entrepreneur.com. (2013, July 8). *The Best and Worst U.S. Cities for Renting Office Space*. <http://bit.ly/1ttdane>

- ^{xxv} Teradata. (2014, June). *2750 Platform: Product and Site Preparation Guide*. <http://bit.ly/1sccG2V>
- ^{xxvi} The Teradata Forum. (2016, January 1). *Teradata Compression (V2R4) – Mark Morris*. <http://bit.ly/1t0NiBi>
- ^{xxvii} ZDNet. (2013, October 21). *Teradata announces cloud platform, competitive appliance economics*. <http://zd.net/Vliqao>
- ^{xxviii} Ponemon Institute, LLC. (2011, May). *2010 Annual Study: Global Cost of a Data Breach*. <http://bit.ly/1tBiowu>
- ^{xxix} National Institute of Standards and Technology. (2016). *National Vulnerability Database – CVE and CCE Statistics Query Page*. <http://1.usa.gov/VLL1vj>
- ^{xxx} Forrester. (2015, December 7). *The Forrester Wave™: Enterprise Data Warehouse, Q4 2015*. <http://bit.ly/1Sackp1>