WHITEPAPER / JULY 14, 2017 / QL218356

Cost/Benefit Case for IBM Db2 11.1 Compared to Oracle Database 12*c* for OLAP Deployments

Sponsored by: IBM

Market Situation

Today, organizations looking to grow their businesses have the opportunity to adopt innovative analytics applications to facilitate the process. New generation applications, coupled with modern database technologies and cutting-edge hardware, will enable organizations to rapidly generate actionable business insights to increase competitive advantage. At the same time, organizations must not only evaluate the performance of these business intelligence (BI) and analytics applications and the systems required to run them, but also their cost-effectiveness, compatibility with existing IT assets, capacity to meet growing organizational requirements, and ability to ensure business continuity.

An increasing number of businesses are investing in cost-effective emerging technologies as a way to boost competitive advantage. However, despite the spotlight placed on BI and analytics offerings, the database foundation remains the crucial engine driving businesses processes and critical decision making for most organizations. A database platform that efficiently manages data and allows users to extract insights with ease will improve an organization's efficiently and agility in confronting the challenges faced by modern businesses.

IBM Db2, Oracle Database, MySQL, and Microsoft SQL Server continue to be dominant enterprise-ready databases, each supported by various hardware from established vendors. This paper presents a cost/benefit case for two industry-leading database platforms for analytics workloads—IBM Power Systems with AIX running IBM Db2 11.1 with BLU Acceleration, and Oracle Exadata Database Machine configured with Oracle Linux and Oracle Database 12*c* Release 2.

Comparisons are of database installations in the manufacturing, retail, and financial service industries. For comparable deployments, costs of ownership for Db2 11.1 with BLU Acceleration are—by wide margins—lower than those for Oracle Database 12*c* with Oracle Database In-Memory. Three-year costs average 43 percent less for use of Db2 11.1 on Power Systems compared to Oracle Database12*c* R2 on Exadata (Figure 1).

Cost disparities between the use of Oracle Database on Oracle Exadata Database Machine and Db2 on Power Systems are largely driven by higher enterprise and system software costs as a result of Exadata's less granular licensing options and higher core densities. In addition, features offered in Db2 must be purchased as extra cost options for Oracle Database. Lastly, Power Systems with PowerVM offers flexibility for organizations to consolidate workloads whereas Oracle Exadata primarily functions as a database appliance, requiring businesses to dedicate additional hardware for enterprise applications.

Assumptions employed in constructing composite company profiles were based on information supplied by organizations worldwide that employ either Db2 and/or Oracle databases, or have migrated from one to the other. Detailed functional comparisons of IBM Db2 11.1 and Oracle Database 12*c* that highlight relationships between technologies and costs are also discussed.

OLAP Technology Differentiators

The development of in-memory technologies has provided new foundations upon which organizations may build critical BI applications. As more organizations make plans to integrate high performance in-memory databases to support online analytical processing (OLAP) workloads for accelerated and agile decision making, it remains prudent to consider not only the performance of a database solution and the hardware that supports it, but also the reliability and ability of these to integrate with existing infrastructures.

IBM's database offering, Db2 11.1 with BLU Acceleration, enables industry-leading OLAP performance through technologies such as in-memory columnar processing, data skipping, actionable compression, and Single Instruction Multiple Data (SIMD) processing. Db2 11.1 may be deployed on Linux or Windows on x86 servers, and Linux or AIX on Power Systems servers. Due to IBM's tightly integrated database, operating system, and hardware, Db2 on Power Systems with AIX is able to achieve high levels of performance without relying on the engineered hardware design found in appliances such as the Oracle Exadata Database Machine. Table 1 provides an overview of notable Db2 features that optimize performance for analytical workloads.

FIGURE 1: Average Three-year IT Costs of Ownership for IBM Db2 11.1 with BLU Acceleration on Power Systems with AIX Compared to Oracle Database 12*c* on Oracle Exadata Database Machine for Analytics Processing

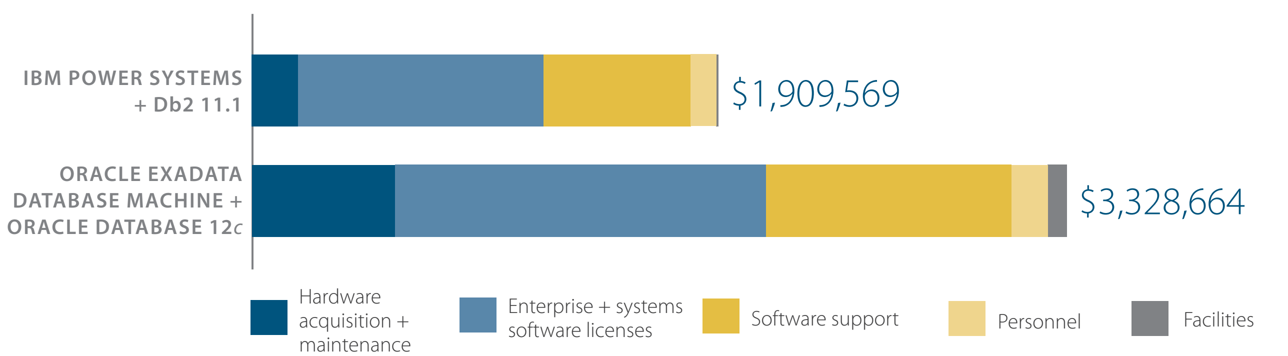

Released in March 2017, Oracle Database 12*c* Release 2 (12.2) for on-premises deployment provided added functionality designed to address the escalating data workload requirements facing today's enterprises. As with previous updates, Oracle enhanced or added features such as higher workload densities, increased SQL performance, and security improvements to supplement legacy core functionality.

Oracle Database runs on Linux, Windows, and several UNIX-based operating systems. Various hardware architectures, such as those based on Intel's x86 and Itanium and IBM POWER processors, from multiple vendors are also supported. The Oracle Exadata Database Machine, running Oracle Database, uses Intel x86 processor and hardware technology specifically modified to address high-volume, scan-intensive, OLAP operations.

TABLE 1: Performance Optimization Features of Db2 with BLU Acceleration

PERFORMANCE OPTIMIZATION FEATURE	DESCRIPTION
In-Memory Columnar Technology	Column-organized tables can be created using BLU Acceleration. Db2 BLU is in-memory optimized, & is not limited by RAM limitations such as those imposed by other in-memory technologies, including Oracle's. BLU dynamically loads data into memory & moves colder data to storage, optimizing use of memory resources based on query workload. Memory is automatically tuned by the Db2 Self Tuning Memory Manager (STMM).
Actionable Compression	Enables automatic storage saving without the need for DBA intervention. Minimize performance impacts through use a growing library of sort, join, & other operations that can be applied to encoded columns without needing to decompress the data.
Data Skipping	Contributes to faster, actionable compression analytics by detecting which encoded data are not required for processing a query, & ignoring those pages. Results in lower I/O & CPU requirements.
Massively Parallel Processing (MPP)	Compatibility with Database Partition Feature (DPF) environments enables Db2 BLU to scale near-linearly with hardware. Parallel processing across partitions improves OLAP performance by leveraging horizontal scaling of resources.
Single Instruction Multiple Data (SIMD)	Hardware instructions are leveraged by BLU to apply a single instruction (e.g., predicate evaluation, join, grouping, arithmetic, R script) simultaneously to multiple data elements.
Multidimensional Clustering (MDC)	Reduces data maintenance & improves performance in large database environments. MDC clusters data & improves query performance through more efficient prefetching & queueing of data.
Continuous Data Ingest (CDI)	Allows users to move large amounts of data into tables without requiring a repeated unlocking & locking of the tables during the process. The ingest utility can continuously process data output from extract, transform, load (ETL) tools to populate large databases stored in partitioned environments. It uses row-level locking, populating tables without affecting other user activities conducted on the same tables.

Although crucial for business growth, the importance of OLAP workloads can overshadow online transaction processing (OLTP) workloads that typically constitute the bulk of daily operations for most organizations. The characteristics of transactional workloads are significantly different from those of analytics. The former involve frequent reads and writes to a few rows of data at a time, rather than the sequential scans of large volumes of data utilized by OLAP applications. Due to these differences, systems that can only be optimized for one type of workload may, in the long term, contribute to inflexibility and increase costs of ownership.

PERFORMANCE OPTIMIZATION

Key advances in the latest Db2 release, version 11.1, build upon high-performance features introduced in earlier versions, and established Db2 strengths such as simplification and automation, workload management, clustering, and SQL compatibility with Oracle Database. Db2 on Power Systems also leads in such areas as multi-core parallelism (thread-to-core ratios on symmetric multiprocessing are adjusted automatically to optimize performance), intra-partition parallelism, and enhanced algorithms for a wide range of processes. Due to these features, users that have migrated to Db2 have reported an average of 33 times faster performance for report generation.

IBM BLU Acceleration is a fully integrated feature set of Db2 that was introduced in Db2 10.5 as an advanced bundling of SIMD CPU-acceleration technology with in-memory columnar data processing, actionable compression, and data skipping. BLU Acceleration has been reported to boost analytic workload processing, while significantly improving data compression and eliminating conventional database management complexity. Advanced compression techniques and data skipping reduce memory and storage requirements to ensure optimal use of resources.

Massively Parallel Processing

Db2 11.1's BLU Acceleration enhancements are tightly integrated and enable high performance analytics processing across massively parallel processing (MPP) network clustering implementations without the need for indexes, aggregates, or tuning. In Db2 11.1, for example, support for MPP greatly increases the scalability of workloads harnessing the capabilities of BLU. BLU MPP support of both scale-up and scale-out strategies opens a field of new possibilities.

The introduction of BLU Acceleration to environments using the Database Partitioning Feature (DPF) enables horizontal scaling over a multitude of partitions and machines to deliver either more hardware to increase the numbers of queries served per hour or to enable petabyte-scale OLAP workloads, without sacrificing performance (Figure 2). In addition to allowing organizations to scale BLU performance beyond previous limitations, users can now more easily integrate BLU benefits to existing DPF data warehouses.

Although it is possible to use BLU on a single machine with hundreds of CPU cores and terabytes of RAM, this becomes increasingly cost prohibitive as systems are scaled up. Leveraging the resources provided by multiple servers for a multi-partition DPF environment, such as increased total buffer pool space, improves the efficiency of BLU. Inter-partition and intra-partition parallel processing of queries enables near-linear scaling of BLU capabilities with additional partitions and hardware.

BLU Acceleration in Db2 11.1 also introduces enhancements such as native column sort and OLAP capabilities using the columnar run-time engine, enabling these operations to be executed without the overhead of converting to row format. These Db2 11.1 capabilities improves OLAP query processing times by a factor of 4.3 compared to Db2 10.5.

Oracle Parallel Execution

Although Oracle Database's parallel execution feature set has been positioned as a MPP technology, it is more limited than BLU MPP in practice. Oracle Database has a shared-everything architecture, in which all nodes have access to all data. Parallel execution for Oracle Database requires the initialization of a pool of parallel execution servers that can be dispatched by the parallel execution (PX) coordinator to execute parallel processes (Figure 2). However, the PX coordinator and servers can only service one request at one time. If all parallel execution servers are occupied and the number of parallel execution servers is at the maximum, Oracle Database reverts to serial processing for incoming tasks. Using parallel execution for operations that do not benefit from parallel processing may negatively affect overall performance.

FIGURE 2: Oracle Parallel Execution and IBM Db2 BLU Massively Parallel Processing (MPP) Technologies for OLAP Workloads

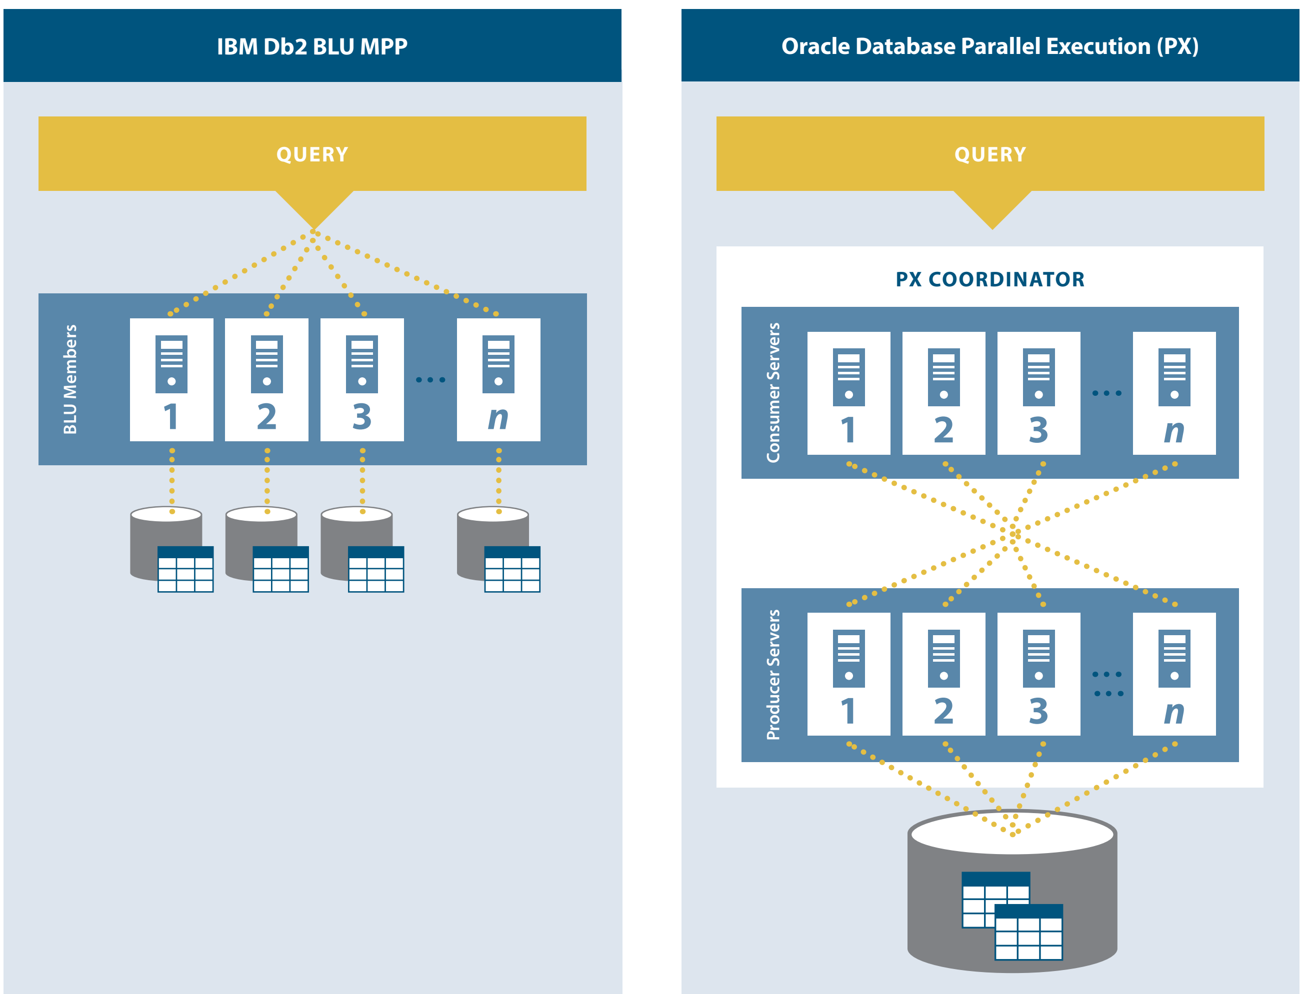

Parallel execution uses a producer/consumer server model, in which producer servers are used to generate data required and consumer servers use the generated data for subsequent operations. In practice, this essentially doubles the number of servers used for a certain degree of parallelism, compared to shared-nothing environments where both data retrieval and OLAP tasks can be completed by each node or member independently, reducing the amount of internodal communication required. This type of split processing between consumer and producer servers also requires optimizing the method of data distribution between server sets for workload balance and efficiency.

Oracle Database partitioning techniques help streamline parallel execution processes, but does not alleviate all performance inefficiencies. For OLAP workloads involving retrieval of large data sets for analysis, shared-everything systems can experience issues such as inefficient caching. Due to each node having access to all of the data, cache misses are more likely, contrasted to shared-nothing systems where each node only needs to cache the subset of information it is responsible for.

In-Memory Technology

Oracle's Database In-Memory option addresses the need to perform transaction and analytical processing within the same database instance, but this solution comes with additional financial and performance overhead costs. To deliver combined in-memory access, an Oracle DBA must select which tables to store as columnar tables (Figure 3). The row form of the selected tables are stored in persistent storage whereas the column-organized tables are in-memory only. Oracle Database maintains transactional consistency between row-based tables and their in-memory columnar counterparts.

FIGURE 3: In-Memory Technologies in IBM Db2 11.1 and Oracle 12*c*

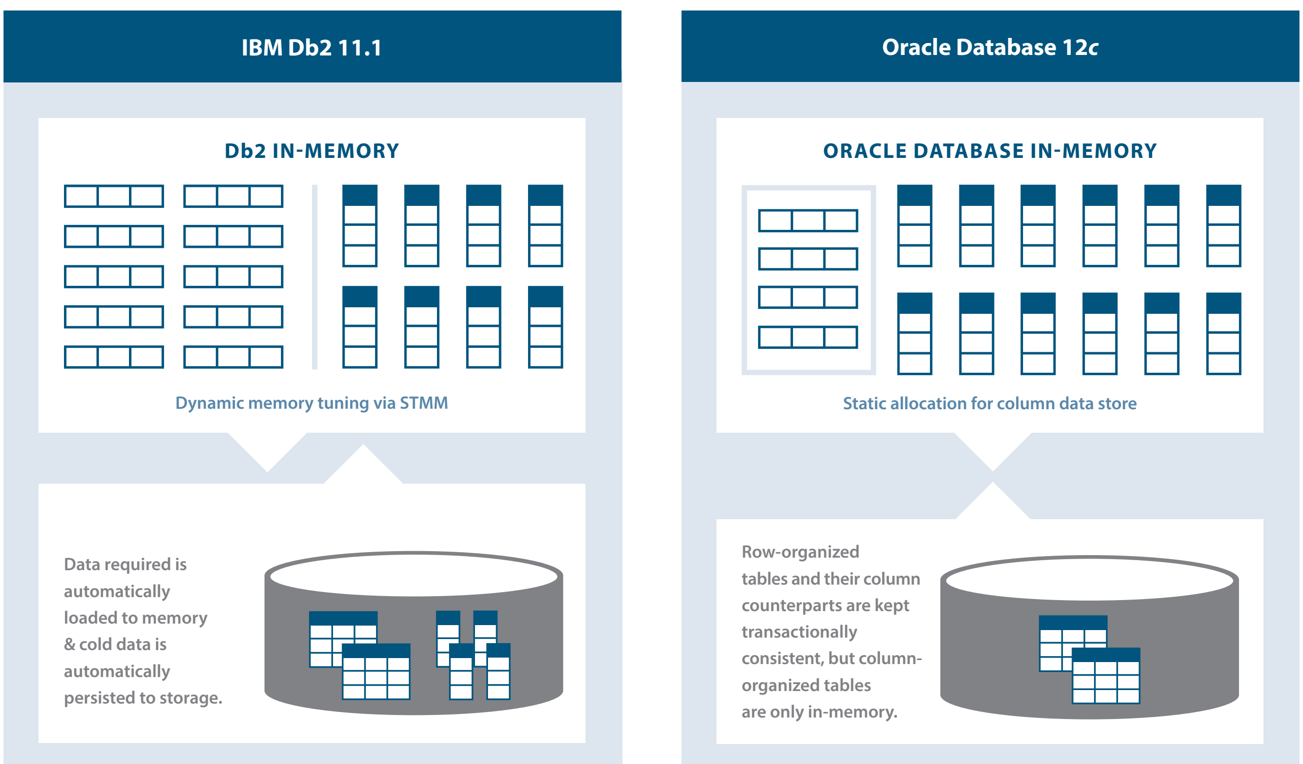

Oracle Database's duplication of row- and column-based data in memory not only increases overall sizing requirements, but also overhead by requiring frequent compression and decompression. Although Oracle claims that this dual-format architecture imposes less than 20 percent additional memory overhead, the company still strongly recommends larger tables be partitioned and use in-memory compression to maximize capacity. Oracle also recommends sizing in-memory columns to accommodate the other objects that must be stored in memory. However, the optimal compression level for columns or partitions must be manually determined by DBAs in order to balance between space-spacing and query performance. According to Oracle's own documentation, high levels of compression optimized for space saving may negatively impact query performance, due to use of "heavier-weight compression algorithm with a larger penalty on decompression."

Oracle DBAs also must decide which type of data (hot or cold) is to be put in-memory, and how much and which form of memory is available for which use. For example, because this in-memory columnar data is not made persistent, In-Memory Column Store (IMCS) mirroring or similar strategy must also be setup and managed to allow for recovery of the data should a failure occur. As part of this contingency planning, the DBA must also set priorities for columnar objects to manage the database repopulation process.

Oracle Database's multi-step tuning process typically requires DBA assistance, in contrast to the various automated features used by Db2. With no need for database tuning, indices, or aggregates, Db2 DBAs can just load the data and go. Users have estimated a 10 to 15 times reduction in database size with Db2 BLU Acceleration, since the need to create aggregates or indices is eliminated.

Db2 BLU's columnar technology automatically optimizes the use of memory resources. Column-organized tables can scale beyond the limitations of RAM because BLU Acceleration dynamically loads data required by queries into memory and moves colder data to persistent storage (Figure 3). Data skipping in BLU improves query performance, and reduces memory requirements, by automatically detecting and skipping sections of data that is not relevant for queries. Oracle Partition Pruning offers similar capabilities, but only allows for less granular skipping of Oracle Database partitions. Actionable compression in BLU, along with the lack of indexes and aggregates required, further maximizes the capacity of memory and storage resources. In Db2, system memory is automatically tuned via the Self Tuning Memory Manager (STMM), eliminating the need for manual tuning.

Compression

IBM Db2 11.1 implements adaptive compression for data stored in rows. This capability, introduced in Db2 10.1, incorporates a set of compression algorithms that supplement table and page-level compression techniques.

Compression has also been optimized for column-organized tables in Db2 11.1. Although different columns are uniquely encoded, encoding for a column across all partitions remain the same. BLU exploits this common compression encoding across data slices to allow for transfer of compressed data to other members without the need to decompress/recompress, reducing bandwidth and CPU requirements.

The Oracle equivalent, Advanced Compression, remains essentially as implemented in Oracle Database 11g R1 in 2007. It employs block-level techniques that tend to be most effective in compressing indexes. Oracle Advanced Compression will compress full pattern matches for column patterns, whereas Db2 Adaptive Compression can compress both full and partial pattern matches.

A further differentiator should be noted. Oracle Advanced Compression tends to be most effective when databases are highly structured, and undergo few changes over time. In more dynamic environments, compression effects tend to be less efficient unless frequent database reorganizations are conducted—which may not always be feasible.

Users who have migrated from Oracle Database to Db2 have typically reported additional compression of up to 50 percent. In practice, compression levels may vary according to data and workload characteristics.

SYSTEMS CONFIGURATION AND MANAGEMENT

For most of Db2's history, simplification and automation have been major IBM design emphases. Enhancements in automation and system administration and management serve to eliminate overhead and reduce costs. Higher levels of DBA productivity than for Oracle have been the norm for more than a decade. Oracle databases typically require more manual DBA intervention than Db2 equivalents. The widening technology gap between IBM Db2 and Oracle Database is not only due to differences in database designs, but also in the different enhancement strategies adopted by the two companies.

Core Db2 features allow DBAs to perform tasks with fewer, simpler actions, in less time than their Oracle counterparts. Monitoring and management tools are more centralized to allow for more efficient database administration. High levels of automation reinforce these characteristics. Autonomic (i.e., artificial intelligence) technologies are employed in numerous features, such as the wide variety of tools used for database tuning and performance optimization.

The Db2 self-tuning memory manager (STMM), in particular, is one of the industry's most advanced self-tuning technologies. Db2 automated storage management, database maintenance, installation and other processes are more extensively automated than Oracle equivalents. Db2 automation contributes not only to DBA efficiency, but also to performance (system parameters may be adjusted more rapidly and efficiently than with manual techniques) and availability (risks of performance bottlenecks and human error are reduced).

IBM Data Server Manager (DSM) is an integrated tool that consolidates monitoring, tuning, configuration, and administration of the Db2 database. DSM allows administrators to manage one or multiple Db2 instances with one centralized tool, using a browser-based, graphical user interface (GUI), thus increasing efficiency for novice to expert administrators. New to Db2 11.1 is the replacement of an earlier text-based tool, db2top, with a DSM text-based utility, *dsmtop*. Both the GUI DSM tool and the new dsmtop can monitor and modify on-premises, hybrid, or cloud database instances. DSM currently supports Db2 and IBM BigInsights. DSM support for more solutions can be expected.

Additional tools help increase operational efficiency. For example, *Configuration Advisor*, which automatically sets database configuration and management parameters to optimize performance at database creation, reduces the amount of database tuning that is typically required later. The *health monitor* tool provides database health information through snapshots and does not degrade performance. *Utility throttling* further regulates utilities and tools to ensure performance is not negatively impacted.

Oracle enhancements tend to focus on adding functionality through add-ons and overlays. This approach has increased both the complexity and inefficiency of overall code structures, resulting in higher processor overhead and lower DBA productivity. In contrast, simplification and automation of database administration tasks continue to be high priorities for IBM.

For example, SQL compatibility is a functionality in which IBM has invested heavily over the course of the last several Db2 releases to reduce personnel overhead. Native support for Oracle Procedural Language/Structured Query Language (PL/SQL) and open source PostgreSQL, along with a wide range of code, tools, and functions commonly employed by developers, is built into the core database engine rather than implemented as a software overlay. Thus, organizations that migrate to Db2 do so with ease, and typically require few changes to applications and personnel skills.

In most cases, Db2 is most advantageous in dynamic environments where integrated, centralized tools and autonomic capabilities allow for more efficient management and administration, and reduces the overhead that can be required with use of other databases. Using Oracle Database in these environments generally requires manual solutions and additional extra cost features, configured using elaborate administrative processes, at considerable additional expense.

Although Oracle offers similar autonomic configuration and tuning features such as automatic memory management, automatic workload management, and automatic storage management, these tools tend to be less integrated and have more rigid operating parameters. Db2 on Power Systems with AIX and PowerVM is an integrated platform where software and hardware features were designed to operate with synergy. In contrast, Oracle Database on x86 with Linux or UNIX operating systems tend to involve multiple vendors, resulting in lower levels of integration and higher complexity when seeking support for any issues that may arise.

SECURITY ADVANCEMENTS

Vulnerabilities and exposure to risk have increased as the availability and variety of enterprise applications continue to expand across IT landscapes. There is no question about how critical security is when choosing the best database platform. Both Oracle and IBM Db2 support similarly impressive lists of encryption and authentication compliance standards; however, between January 2016 and June 2017, Oracle Database had 13 vulnerability alerts whereas IBM Db2 had only one, according to the National Institute of Standards and Technology (Figure 4).

Db2 11.1 improvements in security strategies are responsible for these kinds of results. Db2 native encryption provides secure encryption and key management, which is transparent to applications and schemas, and protects both the database and any backup images created. Db2's symmetric encryption scheme also exploits hardware acceleration for cryptographic operations using server and storage system resources.

Db2 native encryption, included with all 11.1 editions, leverages a two-tier model referred to as envelope encryption, and combines local and master key (MK) mechanisms to ensure that sensitive data is securely encrypted within internal and external communications and repositories at all times. Since the 10.5 release, Db2 utilizes FIPS (Federal Information Processing Standard) 140-2 certified cryptographic libraries and cryptographic algorithms that meet the requirements of NIST (National Institute of Standards and Technology) SP 800-131a.

The expansion of native encryption in Db2 11.1 includes the centralizing of master encryption key management. As a result, Db2 can now leverage the same centralized encryption keys through which other enterprise systems are being managed via the Key Management Interoperability Protocol (KMIP) 1.1 industry standard, such as the IBM Security Key Lifecycle Manager (ISKLM).

User identity authentication within Db2 can be managed via the operating system, through a Lightweight Directory Access Protocol (LDAP) server, or using the Kerberos protocol.

Db2 also manages which operations a user is authorized to conduct on specific data or resources, either directly through a user ID or through the creation and use of database roles, which can be used to combine several database-level authorities and object privileges together.

Db2 also provides a transparent means of limiting the rows or columns that a user can see or access—known as Row and Column Access Control (RCAC)—so the existence of unauthorized row or column data is not detectable. Label Based Access Control (LBAC) provides additional multilevel security for managing access to classified data.

HARDWARE PLATFORMS

Oracle's Exadata Database Machine, heavily marketed for OLAP workloads, involves a significant design modification of the conventional server hardware architecture to optimize it for analytical workloads.

FIGURE 4: Comparative Database Vulnerability Data—January 2016 through June 2017

DATA SOURCE: NIST Computer Security Division, National Vulnerability Database, CVSS Metrics Version 3

The Exadata Database Machine hardware design splits and distributes the Oracle Database between two separate subsystems, each with its own processors and supporting memory infrastructure. The Storage Server nodes are responsible for Oracle's proprietary I/O processing software while Oracle Database resides in the Database Server nodes.

In April 2016, the latest generation of x86-based Exadata machines became available, offering the two-socket X6-2 model with Intel Xeon E5 v4 processors and the eight-socket X6-8 model featuring Intel Xeon E7 v3 processors. A SPARC-based Exadata system, SL6, has been available since February 2017. The two-socket SL6 model contains SPARC M7-based database servers and Intel Xeon E5 v4 storage servers and is offered at a comparable price point to the x86-based models.

The Exadata storage servers include a collection of specialized technologies created over the years to assist with I/O processing communication challenges. Specifically, Hybrid Columnar Compression (HCC), Smart Scan, Storage Indexes, Smart Flash Cache, and Write Back Flash Cache are storage server-based features that together play a critical role in working to minimize traffic over the InfiniBand fabric.

The Exadata Database Machine runs Oracle Linux, a Linux distribution based on Red Hat Enterprise Linux that has been modified to accommodate specific Oracle software and hardware requirements. Oracle Linux is available free and open source; however, a support contract with Oracle is required for related implementation, performance, and operating issues. Third-party support services for Oracle Linux are also available.

Exadata has maintained its original bifurcated database software/storage software design from the X2 through the current X6 generation. The hybridization of data storage and processing in storage nodes using Smart Scan while other operations execute in compute nodes contributes to improving OLAP performance (Figure 5). This hybrid architecture may optimize analytical workloads, but the overall result may be limiting for organizations with diverse workloads.

This specialized hybridization also undermines the ability to host database-dependent enterprise applications on the Exadata platform directly. Oracle recommends that applications accessing the Oracle Database on Exadata utilize additional application servers and/or appliances, incurring substantial costs.

In contrast, IBM's Power Systems design originated as a general-purpose UNIX server, and has been refined and enhanced since its introduction. The company's expertise in processor manufacturing, supercomputer and mainframe designs, and operating system and compiler optimization has been focused within each new generation of Power Systems on providing broad support for mixed workloads with industry-leading, high-performance computing speeds and throughput.

The POWER8 processor has been developed with flexibility and scalability in mind, designed to handle a variety of high-performance computing workloads such as artificial intelligence (AI), in-memory OLTP, and analytics. POWER8 processors, for example, support larger numbers of threads per core (by a factor of four), greater memory bandwidth, a higher clock speed, and more cache (compared to the Intel E7-8895 v3 processor and the Intel E5-2699 v4 processor used in the X6-8 and X6-2 models, respectively). In the POWER8 design, L1 through L3 cache are on-chip, while up to 128 megabytes (MB) of L4 cache is added off-chip. Intel offers no equivalent to POWER8's L4 cache.

A focus on improving Power Systems workload management and performance has resulted in the close integration of PowerVM with the AIX operating system. In a mixed workload environment, the amount of work a server can perform over time depends not only on processor, memory, and I/O throughput, but also on the mechanisms that allocate and reallocate system resources as demands change. Power Systems are equipped to host development, test, and production requirements, efficiently, in robust and protected hardware environments as the need arises.

FIGURE 5: Oracle Database OLAP Workloads on IBM Power Systems and Oracle Exadata Database Machine

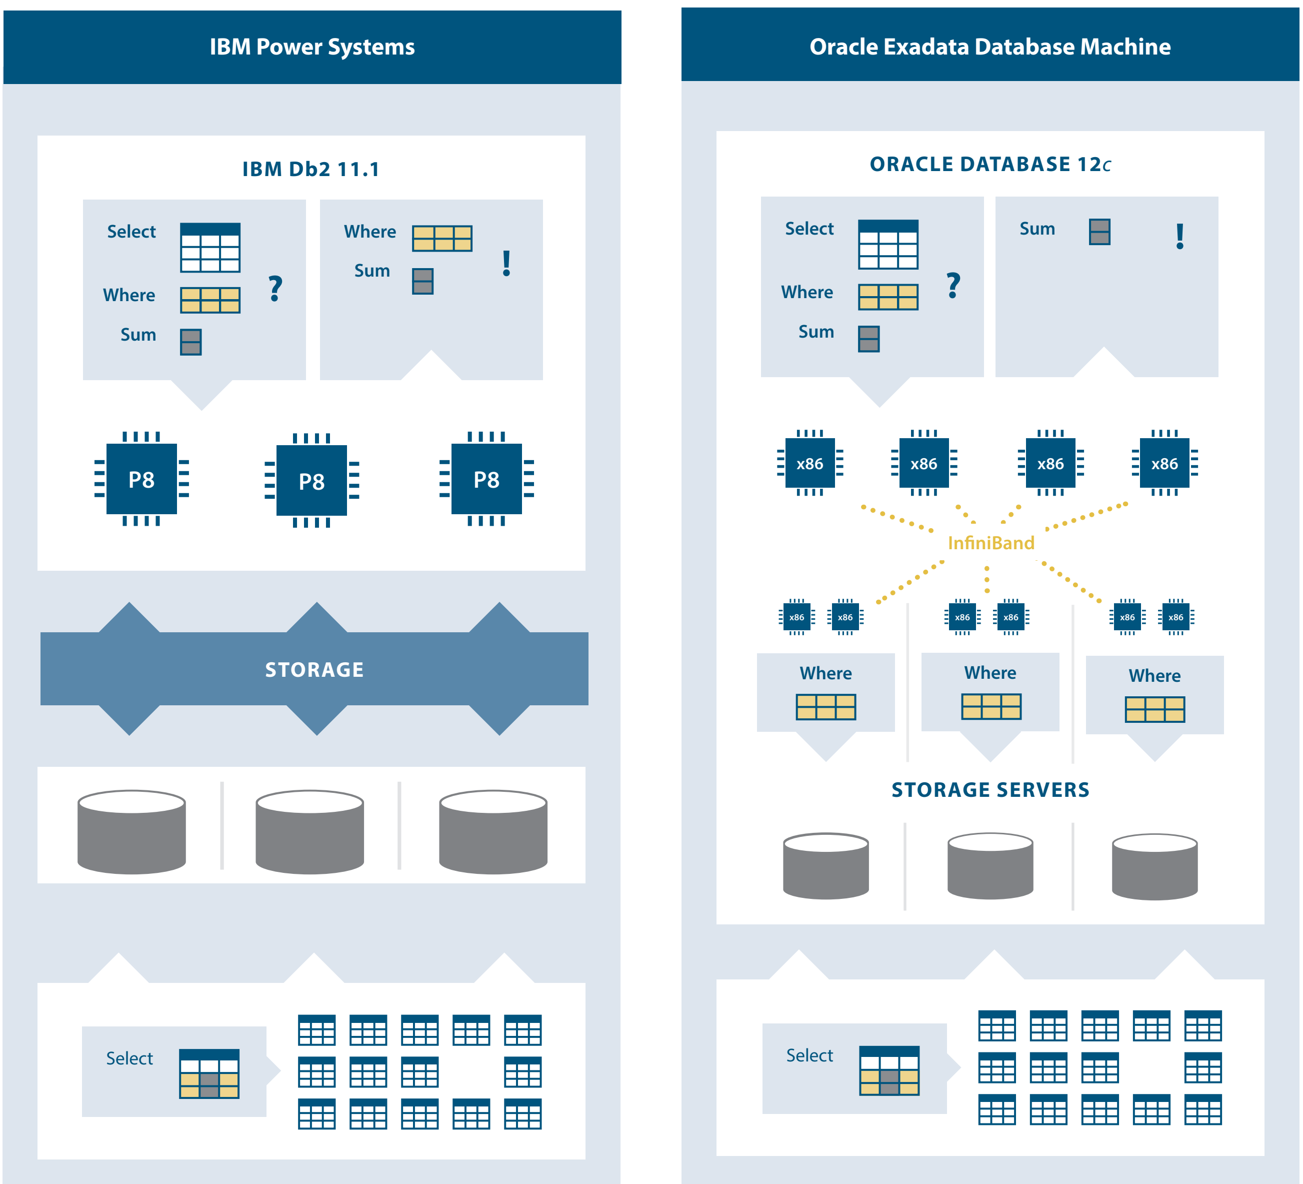

Cost Details

IBM Db2 can be deployed on premise, in the cloud, or in a hybrid environment. Db2 Release 11.1 editions, each with a variety of features, range from the no-charge Db2 Express-C recommended for small businesses and academia to the feature-rich Db2 Advanced Enterprise Server Edition (AESE). Pricing models include Processor Value Unit, per Authorized User Install, and per Terabyte.

For on-premises deployments, Db2 offers exceptional scalability and flexibility. BLU Acceleration works with available capacity to automate in-memory data allocation, allowing Db2 to expand beyond the limits imposed by memory. Db2 Enterprise Server Edition (ESE) and AESE have no limits on processors, memory, or database size.

Cost estimates are based on use of Db2 on Power Systems with the AIX operating system compared to Oracle Database on the Exadata X6 Database Machine with Oracle Linux. Use of Db2 on Power Systems averaged 43 percent less than use of Oracle Exadata for the profile installations used in this paper.

Cost details for OLAP profile installations are presented in Tables 2 and 3.

PACKAGING AND PRICING

Cost differences are also reflected in the packaging and deployment options for Oracle Database on Oracle Exadata and Db2 on Power Systems.

Oracle Exadata X6-2 eighth-, quarter-, half-, and full-rack models are configured with 44, 88, 176, 352 Intel database server cores, respectively. Exadata X6-8 half-and full-rack models contain 288 and 576 database cores, respectively. Capacity on Demand (CoD) can be used to deactivate a fraction of Exadata database cores as a means of reducing total cost of ownership (TCO). However, CoD has limitations with virtualization.

The lack of CoD for Exadata storage servers, coupled with Oracle's mirroring recommendation for high availability, contribute to increased licensing costs. Oracle recommends the triple-mirroring high redundancy setting of Oracle Automatic Storage Management (ASM) to ensure optimal availability. If Data Guard is enabled, ASM normal redundancy (double mirroring) can be used. When using either double or triple mirroring, organizations are required to license the full physical capacity of their Exadata storage. This requirement significantly drives up costs based on the usable mirrored capacity, even if the used storage capacity is effectively a fraction of the total storage capacity.

Storage software licensing prices for Exadata storage servers are listed as $10,000 per disk and $20,000 per flash drive. A full X6 rack can contain up to 14 storage servers with either 8 NVMe flash drives (EF) or 12 disks and 4 flash drives. Storage licensing costs, before discounts, for a full rack can reach from $2.2 to $2.8 million for a single Exadata Database Machine.

In contrast, Power Systems servers offer more granular configurations of 4-, 6-, 8-, 12-, 16-, 24-, 32-, 48-, and 192-cores. Besides database workloads, Power Systems servers can be application-serving as well. Power Systems compatibility with a wide range of storage solutions also affects pricing. Storage

TABLE 2: Profile Installations for Total Cost of Ownership Comparisons

MANUFACTURING COMPANY	RETAIL COMPANY	FINANCIAL SERVICES COMPANY
BUSINESS PROFILE		
Employees: 8,500 Revenue: $4+ billion Locations: 5 production centers	Employees: 25,000 Revenue: $3+ billion Locations: 2,500+ stores	Employees: 10,000 Revenue: $3+ billion Assets: $300+ billion
APPLICATIONS		
Sales & promotions analysis Customer segmentation Production & inventory forecasting Logistics planning Management reporting	Marketing & sales analysis Promotions & pricing analysis Loyalty program management Customer profiling Logistics analysis Multichannel data consolidation	Customer contribution analysis Customer behavior modeling & scoring Operational CRM Risk assessment & reporting Asset & liability management Cost optimization
ORACLE DATABASE 12*c* RELEASE 2 ON ORACLE EXADATA DATABASE MACHINE		
24 cores enabled 1 GB RAM *Oracle Linux, Oracle VM* *0.15 FTE sysadmin*	48 cores enabled 3 TB RAM *Oracle Linux, Oracle VM* *0.25 FTE sysadmin*	96 cores enabled 6 TB RAM *Oracle Linux, Oracle VM* *0.5 FTE sysadmin*
IBM Db2 WITH BLU ACCELERATION ON POWER SYSTEMS		
2 x Power System S824 4.08 GHz *AIX 7.2, PowerVM* *0.15 FTE sysadmin*	1 x Power System E850 3.72 GHz *AIX 7.2, PowerVM* *0.25 FTE sysadmin*	2 x Power System E850 3.72 GHz *AIX 7.2, PowerVM* *0.25 FTE sysadmin*

TABLE 3: Three-Year Cost Breakdown—Use of IBM Db2 11.2 Advanced Enterprise Server Edition (AESE) for Analytics Processing Compared to Oracle Exadata Database Machine

	IBM Db2 ON POWER SYSTEMS			ORACLE EXADATA DATABASE MACHINE		
	MANUFACTURING COMPANY	RETAIL COMPANY	FINANCIAL SERVICES COMPANY	MANUFACTURING COMPANY	RETAIL COMPANY	FINANCIAL SERVICES COMPANY
Hardware acquisition + maintenance	85,163	170,326	340,653	941,600	293,040	529,760
Enterprise + system software licensing	428,419	856,838	1,713,677	490,499	1,362,899	2,692,499
Software support	257,052	514,103	1,028,206	323,729	899,513	1,777,049
Personnel	70,578	117,629	117,629	74,303	123,839	247,678
Facilities	4,062	8,124	16,248	32,797	65,595	131,190
TOTAL ($)	845,274	1,667,021	3,216,413	1,862,929	2,744,886	5,378,176

hardware and software can be supplied by IBM or other vendors and tailored to customers' needs. An organization's existing storage solutions may also be used.

Higher Exadata X6 core counts, compared to Power Systems, as well as previous generation Exadata systems, results in higher system and enterprise software licensing costs. Power Systems can also support additional operating systems, such as enterprise Linux distributions from Red Hat, SUSE, and Canonical as well as system management and tuning tools from IBM.

Database licensing costs also exacerbate the higher cost associated with the use of Oracle Database on Exadata Database Machine compared to Db2 on Power Systems. Many high value features are included in Db2, such as BLU Acceleration and HADR. In contrast, similar features are often extra-cost options to be added on to Oracle Database.

Installations

Cost comparisons for use of Oracle Database 12*c* on Exadata Database Machine and IBM Db2 11.1 with BLU Acceleration on Power Systems for high-performance analytics applications were based on the installations summarized in Table 2.

Costs were calculated as follows:

- **Oracle Database 12*c* Release 2** costs for were calculated based on per core licensing for Oracle Enterprise Edition, Oracle RAC, and other features; along with Exadata hardware acquisition and maintenance; and three-year support costs for Oracle Linux and Oracle VM.
- **Db2 11.1 with BLU Acceleration** costs were calculated based on IBM Db2 11.1 AESE licenses plus three years of support; along with Power Servers hardware acquisition and maintenance; and licenses for AIX and PowerVM.
- **Personnel costs** were calculated based on estimated annual salaries of $112,097 for Oracle Database DBAs, and $106,476 for Db2 DBAs. These estimates were based on industry standard salaries for each database that were available online. Salaries were increased by 43.7 percent to allow for bonuses, benefits and other per capita costs, and multiplied for three years. DBA employment across organizations is affected by variations in applications supported, types of administrative tasks performed, and job description differences. For this paper, FTEs required for each company profile are based on user surveys. Organizations that have conducted migrations from Oracle to Db2 report two to three times reductions in FTE staffing levels.
- **Facilities costs** are for energy consumption and data center occupancy. Energy calculations assume near-24/365 operations over a three-year period. Power usage costs were calculated based on national averages per kilowatt-hour (kWh). Occupancy costs were calculated using a conservative assumption for annual average cost per square foot for existing facilities, but do not include costs for acquisition of facilities.

Conclusions

As the analytics marketplace continues to diversify with new offerings, it becomes more difficult for organizations to select a solution that provides industry-leading performance for OLAP workloads, integrates into existing infrastructure with ease, and provides flexibility to meet growing requirements.

IBM Db2 11.1 with BLU Acceleration on Power Systems servers delivers high speed business insights, while enabling operational efficiency through autonomous tools and providing robust reliability, availability, and serviceability (RAS) features through the hardware and operating system.

Three year costs for IBM Db2 11.1 deployment in an analytics environment average 43 percent less than Oracle Database 12*c* Release 2 on Exadata. Cost differences can be attributed to lower acquisition and support costs for Db2 deployments. Enterprise versions of IBM Db2 come prepackaged with advanced capabilities, such as BLU Acceleration. In contrast, advanced Oracle Database capabilities, such as RAC, are priced as an extra cost option. Higher core density for Exadata Database Machine also drive up software licensing and support costs.

Lower costs are not the only attribute that distinguish Db2 11.1 from Oracle Database 12*c*. There are also major differences in architecture and technology. Db2 is better geared to handle the way in which the database world is evolving regarding hardware, software, and infrastructure communication strategies. Db2 is more effective in handling complex mixed workloads supporting large, diverse user populations. Its design optimizes high performance analytical workloads, and its structure more easily allows introduction of new technologies.

IBM Db2 is, moreover, better designed to deliver higher levels of efficiency and administrator productivity across the entire infrastructure of storage systems, servers, and software that support databases. Equally, if not more important, gains may also be realized in availability, response time, and critical future upgrading of service-level parameters.

Such capabilities will play an increasingly important role in determining how well organizations use information to achieve a business advantage. The choice between Db2 and Oracle, from a business as well as a technology perspective, is between leadership and legacy commitment.

Index

LIST OF FIGURES

LIST OF TABLES

LIST OF REFERENCES

Oracle Database In-Memory with Oracle Database 12*c* Release 2 Technical Overview found at www.oracle.com/technetwork/database/in-memory/overview/twp-oracle-database-in-memory-2245633.pdf

Next-generation Db2 release highlights BLU Acceleration found at www.ibmbigdatahub.com/blog/next-generation-db2-release-highlights-blu-acceleration

Db2 V11.1 Technology Overview found at www.nedb2ug.com/sites/default/files/DB2_V11_Overview_NEDB2UG_2016-05-15.pdf

National Institute of Standards and Technology found at nvd.nist.gov/

Quark + Lepton™

CORPORATE OFFICE
3550 Frontier Avenue, Unit A2
Boulder, Colorado 80301 USA

www.quarkandlepton.com
info@quarkandlepton.com

© 2017 Quark + Lepton LLC. All rights reserved.

Quark + Lepton and the Quark + Lepton logo are trademarks or registered trademarks of Quark + Lepton LLC. This publication may not be reproduced or distributed in any form without Quark + Lepton's prior written permission. The information contained in this publication has been obtained from sources believed to be reliable. Quark and Lepton disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This publication consists of the opinions of Quark + Lepton's research organization and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice. Although Quark + Lepton research may include a discussion of related legal issues, Quark + Lepton does not provide legal advice or services and its research should not be construed or used as such.

IBM sponsored this publication, however, the information and conclusions contained in this publication do not necessarily represent the positions of IBM or other referenced sources.