WHITEPAPER / APRIL 14, 2020 / QL065299

Cost/Benefit Case for IBM Db2 11.5 Compared to Oracle Database 19c for OLTP Deployments

Sponsored by IBM

Market Situation

Digital transformation and infrastructure modernization are driving organizations to explore a wide range of new technologies. On-premises, cloud, containers, microservices and hybrid alternatives are enabling organizations to implement a modern architecture that allows for customization to specific needs and existing assets. The corporate database platform, a primary driver of enterprise-wide planning, operations, and data governance, remains a central business decision that will affect organization-wide business processes. Understanding the impact of the database choice while selecting a secure, reliable platform that is readily compatible with new innovations will allow you to minimize cost of ownership as you enable modernization.

Online transaction processing, or OLTP, can be characterized as workloads that facilitate and manage transaction-oriented applications. Typically involving inserting, updating and deleting small amounts of data in a database, OLTP deals with large numbers of transactions from a large number of users. Large commercial transaction processing applications can be found across all industries—from automated teller machine (ATM) for consumer banking, retail point-of-sale (POS) processing, to reservation systems for hospitality and electronic health records for healthcare—along with a range of cross-industry business process applications such as enterprise resource planning (ERP) and e-commerce.

The underlying goals of such systems are typically defined in terms of availability, speed, concurrency, and recoverability. Users expect immediate response and organizations require accurate and timely execution of business processes. OLTP solutions continue to evolve to meet the growing needs of enterprises, while a plethora of new applications are being deployed to enable web and mobile transactions. Companies aiming to improve the efficiency and performance of essential OLTP systems should adopt a database platform that is cost effective, flexible, and scalable to more effectively meet the increasingly diverse and demanding business-process requirements that will affect digital transformation.

This paper presents a cost/benefit case for two leading enterprise database solutions—IBM Db2 11.5 and Oracle Database 19c—with regard to delivering effective security capabilities, high-performance OLTP capacity and throughput, and efficient systems configuration and management automation. Both databases are running on IBM Power Systems E980 servers in this analysis.

Comparisons are of database installations in the telecommunications, healthcare, and consumer banking industries. For OLTP workloads in these environments, three-year costs average 30 percent less for use of Db2 11.5 compared to Oracle 19c. (Figure 1).

Assumptions employed in constructing composite company profiles were based on information supplied by organizations worldwide that employ either Db2 or Oracle Database. Detailed functional comparisons of Db2 11.5 and Oracle Database 19c that highlight relationships between technologies and costs are also discussed.

There are compelling reasons to pay close attention to database investments. Databases determine, in no small measure, how effectively an organization can leverage information for business advantage—whether to mine and nurture new and/or existing customer sales opportunities, verify customer satisfaction, or to discover operational inefficiencies. To a much greater extent than is generally realized, the database platform influences an organization's year-over-year IT cost trajectories. A highly available database remains as important as ever in ensuring business continuity, and the ubiquity of databases makes their effective configuration and efficient administration crucial to an organization's productivity.

FIGURE 1: Average Three-year Costs for IBM Db2 11.5 Advanced Edition and Oracle Database 19*c* Enterprise Edition for Transaction Processing Installations

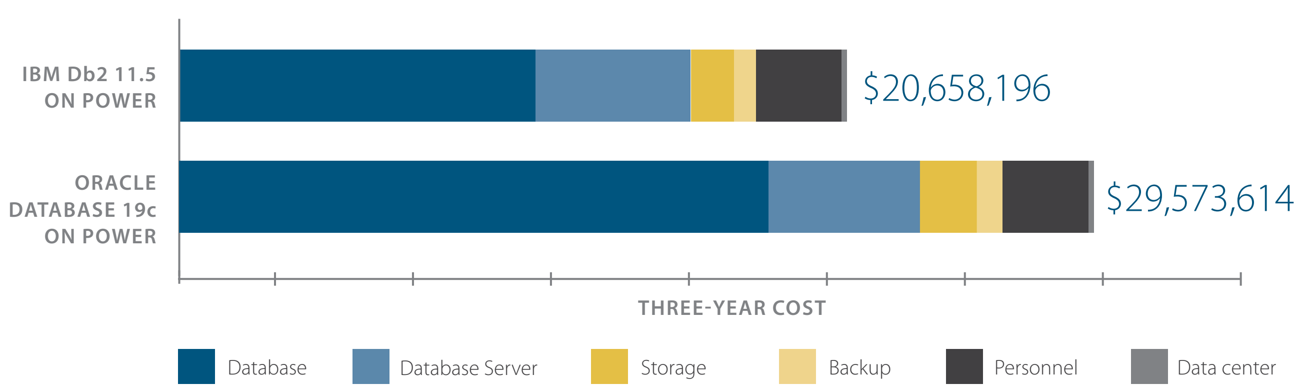

SOURCE: Quark + Lepton (March 2020)

OLTP Workload Characteristics

Commercial transaction systems are typically designed to meet a certain service level. The expected service level will vary by type of transaction but is most often defined in terms of response time and system availability. For example, 20 milliseconds (ms) to find available flights from Denver to New York, and 40 ms to purchase a ticket. Or perhaps 15 ms to check your flight status before heading to the airport and 25 ms to check into your flight. Depending on the particular system, there could be hundreds or possibly thousands of concurrent transactions active at any time.

From a system perspective, these workloads initiate queries and movement of small amounts of data. Read statements such as SELECT SQL and write statements such as INSERT, UPDATE and DELETE SQL are executed by multiple active applications. OLTP workloads typically have the following characteristics in common:

- Relatively simple queries that do not include complex joins
- Concurrent applications that affect disjointed datasets, resulting in high levels of random I/O that may stress the I/O subsystem
- Only a minimal amount of serial I/O operations
- Disjointed datasets and minimal reuse of data often leading to large working dataset size, high buffer pool hit ratio (BPH) and frequent page cleaning
- High demand on transaction logging devices can cause a bottleneck
- Strict application of isolation levels can place high demands on locking infrastructure
- Stringent availability/uptime requirements often approaching 24/7

The characteristics of transactional workloads are significantly different from those of analytics. The former involves frequent reads and writes to a few rows of data at a time, rather than sequential scans of large volumes. Such databases are also typically smaller than those for analytics systems, and I/Os involves smaller blocks of data.

OLTP Technology Differentiators

IBM Db2 and Oracle Database compete as two of the most established enterprise database platforms. In comparing Db2 11.5 and Oracle Database 19c, it is important to distinguish between what a particular feature has been designed to accomplish, and how efficiently it functions in practice.

Business-critical applications require high levels of availability and disaster recovery capability. Performance is largely determined by the efficiency of core data structures. Inefficiencies of these core data structures may increase processor, memory, and disk overhead; reduce throughput; increase risks of bottlenecks and outages; and produce other undesired effects.

In all of these areas, there are significant differences between IBM Db2 11.5 and Oracle Database 19c. Db2's automatic database tuning and management features are beneficial when databases undergo frequent changes in size, schemas, underlying data structures, and workloads.

IBM has invested significantly in SQL compatibility to simplify the adoption of and transition to Db2. Organizations that have migrated Oracle Database applications to Db2 have typically found that 98 to 99 percent of code remains unchanged, and few changes to development skills are required. 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 administrators and developers, is built into the core database engine rather than implemented as a software overlay. Thus, organizations will experience the same levels of performance as native Db2 users.

Table 1 contains a summary of some of the advanced technology features found in the latest version of Db2, Db2 11.5. The importance of these will be discussed in greater detail in the following sections.

DEPLOYMENT FOUNDATIONS

Shared data scale-out clusters are typically implemented as a means of delivering both continuous availability and scalable transaction processing support required for enterprise database systems. This approach allows the distribution of transactions across multiple servers for applications serving large numbers of concurrent users.

In a cluster approach, a database is deployed across multiple servers, or nodes, that effectively function as a single database server. Additional processing power is achieved by adding more nodes as needed, with all nodes cooperating as a single system. The database server may support a single database or multiple databases.

The shared data deployment model supports database instances running on multiple, networked compute nodes, each with equal access to all data storage in the cluster. The compute workload is distributed across nodes with a load-balancing scheme. This shared data approach is considered better for transactional database workloads, where update activity is essentially continuous, and requests are relatively small.

A shared data scale-out architecture thus provides the foundation for the deployment of modern enterprise transactional systems. Applications view the database as a single system that scales as needed.

IBM Db2, when configured with pureScale, functions as a shared data cluster. This allows applications to interact with the cluster as a single database system, without any customization or special configuration. Db2 pureScale is designed to provide extreme scalability and continuous availability of transaction processing across enterprise environments, enabling IT managers to affordably meet the requirements of growing transaction volumes.

TABLE 1: Summary of Db2 Advanced Edition 11.5 Capabilities and Technologies

FEATURE	DESCRIPTION
Common SQL Engine	Allows businesses to write SQL queries once & deploy them anywhere against any data form factor.
Machine Learning Query Optimizer (ML Optimizer)	Improves the performance & efficiency of queries using ML algorithms for significantly faster query performance.
Data Federation	Eliminates data movement & provides a single view of all data, making queries across multiple data sources fast & easy.
IBM Db2 Augmented Data Explorer	Makes it easy for business users to instantly pull key insights through automatically generated visualizations & natural language summaries & act upon them.
BLU Acceleration	Delivers breakthrough performance of in-memory columnar processing without the cost or limitations of in-memory-only systems, dramatically simplifying & speeding the delivery of business insights.
Compression	Helps reduce storage needs & increase performance using multiple techniques including table & index compression with page-level compression as well as Db2 BLU Acceleration with advanced encoding to maximize compression of columnar tables.
Continuous data ingest	Loads data continuously from multiple sources throughout the organization to support faster decision-making.
IBM Database Partitioning Feature (DPF)	Enables massively parallel processing by transparently splitting the database across multiple partitions & using the power of multiple servers to satisfy requests for large amounts of information.
Db2 pureScale	Delivers high availability & exceptional scalability to applications transparently, utilizing shared-disk, cluster technology to free them from the complexities of the underlying database architecture.
Db2 Workload Management	Enables fine-grained resource allocation, monitoring & management of workloads based on service classes, workload characteristics, elapsed time, time of day & more.
Materialized query tables (MQTs)	Improves the performance of complex queries with precomputed results of the whole query or parts of queries.
Multi-temperature data management	Helps maximize performance & reduce overall media costs with storage tiering & the ability to transfer data in real time between different types of storage media.
Column store	Improves performance & reduces consumption of processor, memory & I/O resources for analytics workloads by directing scans to values in a specific column or columns, avoiding the need to process all data in a table.
Data skipping	Reduces processor, memory, & I/O resource consumption by automatically avoiding the processing of data that is not needed for a query.
Shadow tables	Provides the performance benefits of BLU Acceleration to analytic queries that must be executed in OLTP environments with column-organized copies of row-organized tables implemented as MQTs that are maintained by replication.
Data virtualization	Designed to enable users to easily search across diverse data sources.
Hybrid Data Management (HDM)	Provides a platform to manage all data types across all sources & destinations.

SOURCE: Quark + Lepton (March 2020)

Db2 pureScale utilizes architectural concepts and technology from the mainframe version of Db2 along with technologies from Tivoli, providing a cluster environment that is easy to administer and largely autonomic in operation.

All nodes in a pureScale cluster operate cooperatively—database buffers and locks are synchronized by the cluster caching facility or CF. This centralized mechanism allows workloads to be balanced across nodes and optimizes database performance, while ensuring consistent database images on all nodes. In the event of a node failure, a pureScale cluster provides transparent failover of the affected workload to another node. In addition, a standby caching facility is available should CF failover be needed. This results in a resilient database server environment that provides continuous service for business-critical applications.

The Db2 pureScale design enhancements included in Db2 11.1 and 11.5 achieve substantial performance and scaling improvements, while simplifying installation and deployment.

With Db2 11.5, a pureScale cluster can include up to 128 nodes. Administrators can provision the resources needed for current workload and performance requirements. Additional nodes can be added to the cluster as workload demands increase. This is accomplished without interruption to database operations, and Db2 pureScale automatically balances workloads across the additional resources.

Cluster data communications with Db2 11.5 can be accomplished over 10 Gigabit Ethernet (10 GbE), 40 Gb RDMA (remote direct memory access) over Converged Ethernet (RoCE), or InfiniBand networks. IBM Spectrum Scale software-defined storage offers integration with Db2 pureScale, allowing customization and optimization of the storage infrastructure.

Oracle Real Application Clusters (RAC) with Oracle Clusterware functions as a shared data cluster as well, but with a key architectural difference. Rather than utilizing a CF as Db2 does, Oracle RAC uses a peer-to-peer model for coordinating between clusters nodes. With all nodes communicating with each other, and no central arbiter, the number of coordination operations rise sharply as nodes are added to the cluster. This effect is further compounded when the database cluster is servicing transactional applications, where there is a high volume of not only read operations, but also write operations.

The overhead of this coordination will slow the performance of the Oracle cluster and will inhibit efficient scalability beyond a few nodes. While these effects are dependent on the characteristics of specific workloads, IT managers should expect to see a degradation of cluster performance and scalability for large enterprise transactional applications.

The *Db2 Configuration Advisor* automatically sets database configuration parameters, buffer pool size, and database manager configuration parameters to optimize performance at database creation, thus reducing the amount of database tuning that is typically required later. Utility throttling further regulates utilities and tools to ensure performance is not negatively impacted.

Db2 Data Management Console (DMC) is a new integrated database management tools platform that consolidates administration and monitoring, along with performance management and optimization into a modern browser-based console. It replaces the previous generation Data Server Manager (DSM) and is packaged with all Db2 editions at no additional cost. DMC provides real-time and historic monitoring, alerts and notifications, database object exploration, an integrated SQL editor, and support for RESTful service APIs. Use of the DMC tool can simplify data administration and management while increasing operational efficiency.

AI AND DATA SCIENCE

With Db2 11.5, IBM is taking steps to introduce artificial intelligence (AI) and data science capabilities into Db2. This will enable users to more easily build machine learning (ML) into applications, hence allowing greater insights from data. It also offers the potential to optimize complex queries and transactional processes.

Data science has gained prominence in recent years. A multidisciplinary field that uses scientific methods, processes, algorithms, and systems to extract knowledge and insights from both structured and unstructured data, data science leverages powerful systems and programming tools.

A new series of drivers is available for the open source programming languages and frameworks that are commonly used for AI development. Supported languages include Python, GoLang, PHP, Ruby, Java, Node.js, and sequelize. There is also support for popular frameworks such as Jupyter notebooks and Visual Studio. This will allow developers to more easily build applications using Db2 and provides a more durable development and support environment while encouraging innovation. The resulting applications should require less management and provide increased resilience.

Db2 *Augmented Data Explorer* is a new natural language tool that provides developers with a traditional search-engine like user experience. Users can ask questions of their data and receive an answer in both natural language summaries and data visualizations. The tool is designed to be an aid in building database applications, as it enables developers to easily explore datasets during the development process. Data scientists and engineers will also find it useful as a means of easily exploring data, since it can quickly lead to a better understanding of the dataset.

Db2 *Machine Learning (ML) Optimizer* is included in Db2 11.5 as a technology preview of a machine learning-based cost optimizer. ML Optimizer is intended to complement the current Db2 Optimizer with further opportunities to improve query performance. Initial focus is on improving cardinality cost estimation and the breadth of application should be expected to broaden overtime.

IBM Cloud Pak for Data is a integrated data and AI platform designed to infuse AI capabilities throughout an organization. Built on Red Hat OpenShift Container Platform, Cloud Pak for Data supports all leading commercial cloud environments along with private clouds. Integrating IBM Watson AI technology with hybrid data management, DataOps, and governance and business analytics, organizations can quickly modernize the collection, organization, and analysis of data.

OPTIMIZATION AND AUTOMATION

Databases are growing dramatically and can often be measured in the tens of terabytes. At the same time, systems and the database functionality have grown in complexity, driving growth in the number of staff required for design, administration, and analyst activities.

Database automation addresses administrative tasks for the database administrator (DBA). The facilitation of unattended processes and self-updating procedures will lead to more successful deployments with fewer errors and more reliable change management activities. Automation can often reduce the amount of staff resources utilized for tasks such as applying code updates or patches, provisioning of system resources, and backup/recovery processes.

Database optimization generally refers to strategies for reducing database response time. Maximizing the speed and efficiency with which data is retrieved can result from good database design, a strong understanding of the target dataset, and well-designed queries or applications.

Simplification and automation have long been a core component of Db2. There are a variety of tools that allow Db2 DBAs to perform tasks with fewer, simpler actions, and in less time than their Oracle counterparts. High levels of automation reinforce these capabilities.

Autonomic technologies are employed with numerous features, such as the wide variety of tools used for database tuning and performance optimization. Queries against a database in a Db2 pureScale environment are automatically offloaded to members (database servers) with high resource availability. Scaling out with Db2 pureScale is simple. Members can be added without requiring repartitioning, application changes, or performance tuning, because the database topology is transparent to applications. In contrast, Oracle RAC's shared data architecture is not transparent to applications and requires DBAs to extensively test applications for scalability.

Routine file and storage management within a Db2 pureScale instance is handled by Spectrum Scale, the high-performance distributed file system installed with Db2 pureScale. Oracle Automated Storage Management (ASM) and Oracle ASM Cluster File System (ACFS) provide comparable functions for Oracle Database 19c. However, Oracle ACFS is less integrated than Spectrum Scale, only supporting database files in the most recent releases under certain conditions. Oracle's dynamic rebalancing of a disk group simply involves moving data between disks to ensure that every file is evenly spread across all of the disks in the group, which does not help optimize workload processing performance. Neither workload characteristics nor I/O statistics are considered.

Db2's self-tuning memory, in particular, is an industry leading autonomic technology. Available for each member in a Db2 pureScale environment, it automatically and iteratively adjusts configuration parameters to allocate resources optimally for memory performance. Db2 automated storage management, database maintenance, installation, and other processes are more extensively automated than Oracle equivalents. This 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).

Compared to Oracle RAC, Db2 pureScale generates lower levels of cluster overhead. Requirements for DBA intervention are also lower for such tasks as initial configuration, expansion, addition of new applications, workload balancing, performance optimization, and testing.

Other examples of Db2 self-managing technologies that assist in managing a database system once initial parameters are set up include:

- *Automatic storage* is used by default by databases and simplifies storage management for table spaces.
- *Data compression* reduces the size of tables, indexes, and backup images, thus reducing storage requirements.
- *Automatic database backup* ensures that a recent, full backup of a database is performed as needed and can be an integral part of a disaster recovery strategy.
- *Automatic reorganization* periodically evaluates tables and indexes that have had their statistics updated to see if reorganization is required, and schedules such operations as necessary.
- *Automatic statistics* collection helps improve database performance by ensuring that up-to-date table statistics are available. Which statistics are required by the workload and which statistics must be updated is determined by the database manager.
- *Automatic maintenance* capabilities for performing database backups, keeping statistics current, and reorganizing tables and indexes as necessary are provided by the database manager.
- *Health monitor* proactively monitors changes in the database environment that could result in a performance degradation or a potential outage. Should a health risk be encountered, the database manager informs and advises on how to proceed

COMPRESSION

Data compression relies on various techniques to reorganize database content in order to reduce physical storage space, improve performance, and reduce costs. Compression methods typically vary depending on the nature of the data and patterns of redundancy.

Db2 11.5 Adaptive Compression offers the most dramatic storage savings from the various data compression techniques available for Db2. This is accomplished by incorporating classic row compression along with a page-by-page approach to achieve further data compression.

Adaptive Compression utilizes a table-level compression dictionary used in classic row compression to affect data based on redundancy within a sampling of data from the table. A page-level dictionary-based compression algorithm is then utilized to compress data based on data redundancy within each page of data.

Oracle Database Advanced Compression uses Heat Map to determine the activity level of data blocks and compresses the tiered data accordingly. Advanced Compression tends to be most effective when databases are highly structured and undergo few changes over time. Advanced Compression may not be as effective for highly transactional workloads involving frequent changes to data.

In addition, Advanced Compression tends to be recommended for systems with high CPU availability, due to its potential to generate high levels of overhead.

Storage compression technology in Db2 is generally more effective than with Oracle Database. This results in greater savings in storage hardware and software tools, in backup systems and media, and in data center energy and occupancy costs.

Db2 data compression also gets a boost with the on-chip Nest Accelerator NX842 hardware first available with POWER7+ servers. Compression of Db2 backup and archive logs are offloaded to the accelerator, thus achieving greater compression ratios than possible with software compression.

Due to these advanced compression and autonomic features, users who migrated to Db2 have reported an average of 43 percent reduction in backup time. Users who have migrated to Db2 from Oracle reported an average of 47 percent reduction in storage requirements. On average, Db2 users have also experienced an additional 39 percent increase in compression rate.

ENCRYPTION AND SECURITY

Db2 encryption capabilities are integrated into core processing and apply to data-at-rest in storage as well as data-in-transit. For data-in-transit between client and database server, Secure Sockets Layer (SSL)—sometimes referred to as Transport Layer Security (TLS)—technology is supported.

For data at rest, Db2 *native encryption* provides secure encryption and key management, which are transparent to applications and schemas, and protects both the database and/or any backup images created. This symmetric encryption scheme also automatically detects and exploits hardware acceleration for cryptographic operations, such as the Intel AES-NI. Db2 on AIX can further leverage the AIX encrypted file system (EFS) to protect all files in the system.

Db2 10.5 first provided support for native encryption by using a per-instance, local file keystore mechanism. Since that release, Db2 has continued to utilize FIPS (Federal Information Processing Standard) 140-2 certified cryptographic libraries and meet the requirements of NIST (National Institute of Standards and Technology) SP 800-131a compliant cryptographic algorithms.

Support for a centralized keystore in Db2 11.5 allows for use of remote, centralized management of all encryption keys. Db2 11.5 ability to use the IBM Security Key Lifecycle Manager (SKLM) keystore, for example, ensures master key management industry standard compliance with the Key Management Interoperability Protocol (KMIP) 1.1. In addition, Db2 11.5 offers direct support for hardware security modules (HSMs) such as the Gemalto SafeNet HSM (formerly Luna) version 6.1 and higher, and nCipher's nShield HSM, security world software version 11.50 and higher.

An additional, fine-grained, access control (FGAC), row and column access control (RCAC), can be used to manage user access to tables at the row- or column-level, thus restricting a user's ability to only see the subset of data required for them to perform their specific job. Another security measure, label-based access control (LBAC), allows a security administrator to use an organization's security policy to define criteria that will determine exactly who has write access and who has read access to individual rows and individual columns.

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. This is controlled either directly through database-level authorities and privileges that are associated with a specific 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.

Oracle Advanced Security offers similar capabilities, but only as an extra-cost option for Oracle Database Enterprise Edition. Advanced Security is comprised of two features: Transparent Data Encryption (TDE) enables data to be encrypted so only the authorized user can read it; and Oracle Data Redaction, which enables column data to be redacted (masked) in compliance with industry regulations. TDE encrypts sensitive data stored in data files and stores the encryption keys in a security module external to the database (a keystore). Oracle Key Vault allows for central management of TDE keystores, but requires a separate license for each server installation. SSL is also supported by Oracle Database.

Vulnerabilities and potential exposure to risk have increased as the universe of hybrid and managed cloud services continues to expand. There is no question about how critical security is when choosing the best database platform. Both Oracle Database and IBM Db2 support similarly impressive lists of encryption and authentication compliance standards.

CONTINUOUS AVAILABILITY

IBM's Db2 11.5 pureScale is designed to provide continuous operations for business-critical transactional applications. In the case of either planned system maintenance or unexpected component failure, database processing continues without interruption. Workload is automatically balanced across all active nodes with no impact to applications.

In the case of an unplanned failure, heartbeat detection quickly identifies and isolates any failed system component. Should a member failure occur during database processing, it is immediately isolated from active system resources while Db2 pureScale initiates and completes an automated crash recovery process. Any data that was in-flight on the affected member is temporarily locked while awaiting member recovery. The in-flight data is then released from lock and processing continues normally.

For a software failure, the member is re-started on the original host where recovery is performed. However, with a hardware failure, the member is re-started on a different host allowing for the data to be recovered. When the original host is available, the member moves back to that host and processing continues. During these failure/recovery processes, only the in-flight data is locked. The database instance remains available and accessible to all other members of the cluster.

There are important differences to these processes in an Oracle Database environment. If an Oracle RAC node should fail during database processing, all data (rather than in-flight only) is locked while the recovery process is completed. The database must be re-instantiated in passive standby node(s) during this process.

Planned systems maintenance is accomplished with only minimal disruption to the Db2 pureScale cluster. Each member is simply temporarily taken offline for any required maintenance and then restarted when ready. During this time, any processing workload is handled by the other active members of the cluster and there is no effect on database availability. This minimizes planned downtime for mission critical workloads by enabling online rolling maintenance, such as patches, hardware and firmware updates, and operating system fixes.

CFs also contribute to rapid failover and recovery by providing a common group buffer pool (GBP) that coordinates caching across cluster members. CFs contain data from members' local buffer pools and can act as a shared communication area, which members may use as cluster-wide shared memory. In addition to facilitating buffer pool management, CFs provide global locking capabilities to ensure data concurrency and to prevent data corruption. These centralized attributes result in faster failover and recovery, because navigating through local caches and locking mechanisms is not required. (Figure 2)

FIGURE 2: IBM Db2 pureScale and Oracle Real Application Clusters Architecture

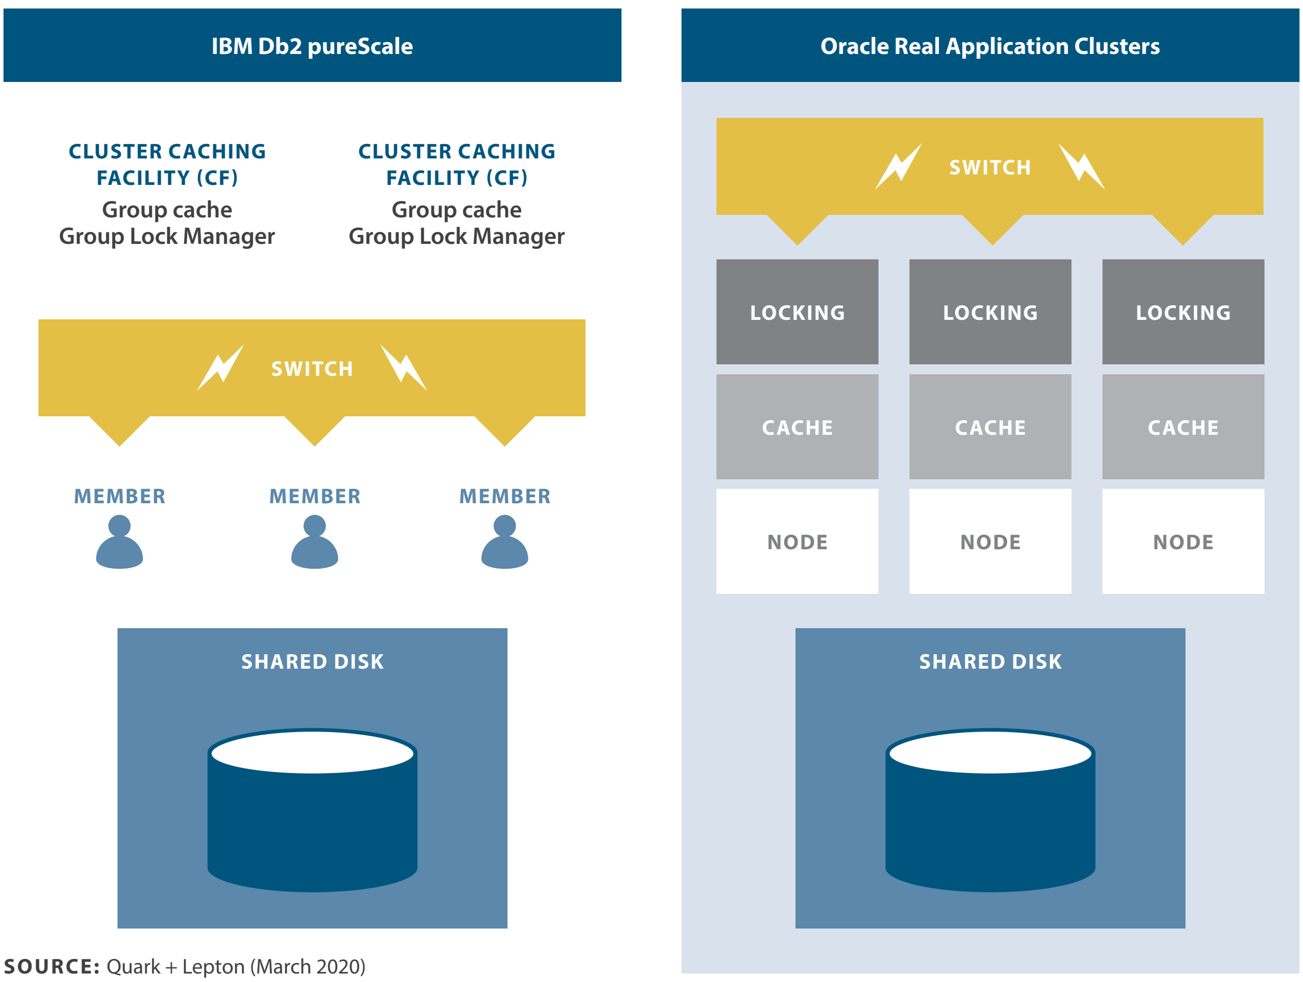

SOURCE: Quark + Lepton (March 2020)

In addition, Tivoli System Automation for Multiplatform (SA MP) cluster managing software provides self-healing processes through the detection of improper operations of systems, transactions, and processes, and initiates corrective actions without disrupting application users. SA MP automatically and transparently moves applications within the cluster to facilitate continuous availability, eliminating the need for manual monitoring.

In comparison, Oracle RAC clusters with failed nodes must go through a lock redistribution process that may take minutes to complete. Database processes are frozen until the system determines what needs to be recovered. For large-scale, business-critical systems, a large number of transactions and records may be impacted. Db2 pureScale mitigates this complex, protracted process, because its centralized locking and cache management capability is not affected by individual member failures. CFs are also redundant to protect against CF failure.

DISASTER RECOVERY

Db2 high availability disaster recovery (HADR) provides a disaster recovery (DR) solution that protects against data loss in the case of a partial or complete site failure. HADR replicates the primary database to a standby database and supports up to three remote standby servers.

The standby database to be hosted at a different physical location. This approach is preferred for critical transactional systems to safeguard against catastrophic system failures. Should applications lose access to the primary database, a standby database can take over rapidly. Transactions are re-routed automatically until the system failure is resolved and the primary database is back online.

A combination of HADR in a Db2 pureScale environment offers not only a disaster recovery solution, but also the high and continuous availability, manageability, and scalability of Db2 pureScale. HADR benefits from the use of a geographically dispersed Db2 pureScale cluster (GDPC) configuration that allows a Db2 pureScale cluster to be distributed, with members of the cluster located at different sites.

HADR is a widely used disaster recovery solution and is comparable to Oracle's Data Guard and Active Data Guard. Both vendors' solutions allow users to combine central HA clusters with remote site failover and recovery.

PACKAGING AND PRICING

With the introduction of Db2 11.5, IBM has greatly simplified the portfolio. Db2 11.5 Community, Standard, and Advanced editions share a common code base and users can easily upgrade as needed.

Db2 11.5 Community Edition is a no-cost download intended for single-developer use for design, test, and prototype applications. Db2 11.5 Standard Edition is intended for deployment in medium-size organizations and departmental usage within large organizations and capped at 16 cores and 128 GB RAM. Db2 11.5 Advanced Edition is designed for large enterprises, and can be deployed in physical and virtual environments with no server or storage limitations.

Db2 pureScale and other high-value features, such as encryption, are included in Db2 at no additional cost. This contrasts with Oracle Database, where comparable features must be purchased as add-on products. A summary of these differences is shown in Table 2.

Customers who have migrated to Db2 typically report up to 40 percent reductions in software and middleware licensing costs. For large deployments, savings can be significant, and free budget resources for other productive uses.

With Oracle Database 19c, Standard Edition usage is limited by CPU sockets and cannot be used with Oracle RAC. This effectively requires Oracle Database 19c Enterprise Edition deployment for all database cluster implementations, including development, testing, and prototyping applications.

Cost Details

The cost comparisons presented in this paper are for three installations of Db2 11.5 Advanced Edition and equivalent Oracle Database 19c Enterprise Edition software stacks in telecommunications, healthcare, and consumer banking companies. These installations are summarized in Table 3.

Configurations reflect input from user surveys and others with similar business profiles, application portfolios, and database characteristics. Performance, compression, FTE staffing and other variables employed are based on information from users that had deployed Db2 and/or Oracle databases for similar applications.

TABLE 2: Db2 11.5 Advanced Edition Features and Oracle Database 19c Enterprise Edition Equivalents

	Db2 11.5 Advanced Edition	Oracle Database 19c Enterprise Edition	
FUNCTION	*No-cost Features*	*Extra Cost Options*	*Processor Core License + Support Cost ($)*
Disaster recovery	HADR	Data Guard	Included*
Data compression	Adaptive Compression	Advanced Compression	14,030
Advanced security	Label-based Access Control (LBAC)	Label Security	14,030
Machine learning	Machine Learning Query Optimizer & Augmented Data Explorer (ADE)	Advanced Analytics	28,060
Spatial data & analysis	Spatial Extender	Spatial & Graph	21,350
Data partitioning	Table Partitioning	Partitioning	14,030
Replication	Q Replication & SQL Replication	Active Data Guard	14,030
Multitenant	Architectural design	Multitenant	21,350
Column organized	BLU Acceleration	In-Memory	28,060
Cluster	pureScale	Oracle RAC	28,060
TOTAL COST			$183,000

SOURCE: Quark + Lepton (March 2020), Oracle Technology Price List (March 2020)

Resulting configurations and FTE staffing levels are as shown in Tables 4 and 5.

Cost details for installations are presented in Tables 6 and 7.

Costs were calculated as follows:

Database software costs are based on use of Db2 11.5 Advanced Edition and Oracle Database 19c Enterprise Edition. To provide equivalent functionality to Db2, applicable extra cost options for Oracle Database shown in Table 2 are included.

Db2 pureScale and Oracle RAC are employed as clustering solutions. Calculations include three-year software licensing and support fees. Db2 licenses are based on IBM's virtual processor core (VPC) values and Oracle Database licenses are based on Oracle's per processor metric as defined in its Processor Core Factor Table.

Storage and backup costs were based on use of FlashSystem 900 and Storwize V7000, respectively. FlashSystem 900 arrays were deployed with 18 TB IBM MicroLatency Modules in RAID 5 configurations. For backup, Storwize V7000 enclosures with were configured with either 19.2 TB 2.5-inch NVMe FlashCore Modules or 7.78 TB 2.5-inch NVMe Flash Drives. Redundant storage and backup systems were sized to estimated installation requirements.

Database server calculations were based on use of IBM Power Systems E980 servers. The AIX 7.3 operating system and PowerVM virtualization were employed for all installations. Calculations for the telecommunications, healthcare, and consumer banking organizations include IBM PowerHA SystemMirror failover clustering software.

TABLE 3: Transaction Processing Installations Summary

TELECOMMUNICATIONS COMPANY	HEALTHCARE ORGANIZATION	CONSUMER BANKING
Business Profile		
Diversified communications services provider, including wireless & Internet $8.5+ billion sales 8+ million subscribers 20,000 employees	Healthcare network with 7 hospitals, and 177 clinics, doctors' offices & other specialty services sites $2.5 billion revenue 1,500 licensed beds 1,100 physicians 2.1 million patient encounters per year 14,500 staff	Diversified retail bank $34+ billion assets $1.5+ billion operating revenue 144 branches 800,000+ customers 4,300+ employees
Applications		
Billing, CRM, operational systems, data store	EHR / EMR, HIS, PACS, business & departmental	Core banking systems

SOURCE: Quark + Lepton (March 2020)

TABLE 4: Configurations and FTE Staffing Summary for Db2 11.5 Advanced Edition on IBM Power E980 Systems

TELECOMMUNICATIONS COMPANY	HEALTHCARE ORGANIZATION	CONSUMER BANKING
Databases + Tools		
Db2 11.5 Advanced Edition including pureScale	Db2 11.5 Advanced Edition including pureScale	Db2 11.5 Advanced Edition including pureScale
Database Servers		
2 x E980 132 x 3.9 GHz *AIX, PowerHA SystemMirror*	6 x E980 432 x 3.9 GHz *AIX, PowerHA SystemMirror*	2 x E980 120 x 3.9 GHz *AIX, PowerHA SystemMirror*
Storage		
FlashSystem 900 x 360 TB *Storage Insights*	FlashSystem 900 x 720 TB *Storage Insights*	FlashSystem 900 x 360 TB *Storage Insights*
Backup		
Storwize V7000 *Spectrum Virtualize, Storage Insights*	Storwize V7000 *Spectrum Virtualize, Storage Insights*	Storwize V7000 *Spectrum Virtualize, Storage Insights*
Personnel		
3.5 FTEs	4.5 FTEs	3.25 FTEs

SOURCE: Quark + Lepton (March 2020)

TABLE 5: Configurations and FTE Staffing Summary for Oracle Database 19c Enterprise Edition on IBM Power E980 Systems

TELECOMMUNICATIONS COMPANY	HEALTHCARE ORGANIZATION	CONSUMER BANKING
Databases + Tools		
Oracle Database 19c Enterprise Edition including RAC	Oracle Database 19c Enterprise Edition including RAC	Oracle Database 19c Enterprise Edition,including RAC
Database Servers		
2 x E980 144 x 3.9 GHz *AIX, PowerHA SystemMirror*	4 x E980 456 x 3.9 GHz *AIX, PowerHA SystemMirror*	2 x E980 130 x 3.9 GHz *AIX, PowerHA SystemMirror*
Storage		
FlashSystem 900 x 360 TB *Storage Insights*	FlashSystem 900 x 720 TB *Storage Insights*	FlashSystem 900 x 360 TB *Storage Insights*
Backup		
Storwize V7000 *Spectrum Virtualize, Storage Insights*	Storwize V7000 *Spectrum Virtualize, Storage Insights*	Storwize V7000 *Spectrum Virtualize, Storage Insights*
Personnel		
4.75 FTEs	6.5 FTEs	4.5 FTEs

SOURCE: Quark + Lepton (March 2020)

TABLE 6: Three-year Cost Breakdown—Use of IBM Db2 11.5 Advanced Edition

	TELECOMMUNICATIONS COMPANY	HEALTHCARE ORGANIZATION	CONSUMER BANKING
IBM Db2 11.5 with pureScale on IBM Power Systems			
DB2 11.5	4,752,000	15,552,000	4,320,000
Support	1,900,800	6,220,800	1,728,000
Database Total ($)	6,652,800	21,772,800	6,048,000
Server	1,548,456	4,564,680	1,548,456
Licensing + Support	1,371,314	4,261,668	1,279,814
Server Total ($)	2,919,770	8,826,348	2,828,270
FlashSystem 900	1,384,400	2,768,800	1,384,400
Licenses + Support	Included	Included	Included
Storage Total ($)	1,384,400	2,768,800	553,760
Storwize V7000	141,280	141,280	141,280
Media	372,960	803,520	266,400
Backup Total ($)	514,240	944,800	407,680
THREE-YEAR SYSTEM TOTAL ($)	11,471,210	34,312,748	9,837,710
Personnel cost	1,868,476	2,402,326	1,735,013
Facilities	75,528	196,046	75,528
THREE-YEAR TOTAL COST ($)	13,415,214	36,911,121	11,648,252

TABLE 7: Three-year Cost Breakdown—Use of Oracle Database 19*c* Enterprise Edition

	TELECOMMUNICATIONS COMPANY	HEALTHCARE ORGANIZATION	CONSUMER BANKING
Oracle Database 19c with Oracle Real Application Clusters on IBM Power Systems			
Oracle Database 19*c*	6,768,000	21,432,000	6,110,000
Support	4,466,880	14,145,120	4,032,600
Database Total ($)	11,234,880	35,577,120	10,142,600
Server	1,549,284	4,564,680	1,549,284
Software + Support	1,431,646	4,412,312	1,361,646
Server Total ($)	2,980,930	11,976,198	2,901,992
FlashSystem 900	1,384,400	2,768,800	1,384,400
Software + Support	Included	Included	Included
Storage Total ($)	1,384,400	2,768,800	1,384,400
Storwize V7000	141,280	141,280	141,280
Media	372,960	892,800	319,680
Backup Total ($)	514,240	1,034,080	460,960
THREE-YEAR SYSTEM TOTAL ($)	16,114,450	48,356,992	14,898,890
Personnel cost	2,709,671	3,707,970	2,567,056
Facilities	80,206	210,079	75,528
THREE-YEAR TOTAL COST ($)	18,904,326	52,275,041	17,541,474

SOURCE: Quark + Lepton (March 2020)

Personnel costs were calculated based on estimated annual salaries of $121,348 for Oracle Database with RAC DBAs, and $113,561 for Db2 DBAs with pureScale certification. These estimates were based on industry standard salaries for each database that were available online. Salaries were increased by 56.7 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.

Data center cost calculations include data center occupancy and energy consumption, as well as allowance for acquisition, maintenance, and operational costs for data center infrastructure equipment, such as uninterruptible power supplies (UPS), power distribution systems (PDS), and cooling systems. Costs for power usage were calculated based on national averages per kilowatt-hour (kWh). Costs were calculated using a conservative assumption for annual average cost per square foot for existing facilities (i.e., costs do not include new facilities construction).

All costs are in shown in U.S. dollars.

Conclusions

Large commercial transaction processing applications are growing in complexity and data volumes while meeting demanding availability and security requirements. A shared data scale-out cluster architecture provides a compelling database environment allowing processing capacity to grow as needed while meeting the needs of the applications.

IBM Db2 11.5 provides an enterprise-ready and cloud-enabled OLTP database platform that continues to meet such business requirements. Db2 11.5 pureScale clusters deliver near linear performance scalability, transparent access to applications, and continuous availability with unfaltering disaster recovery capabilities.

Three-year costs of transaction processing environments in telecommunications, healthcare, and consumer banking companies, average 30 percent less for Db2 11.5 compared to Oracle 19c. Much of the cost variance can be attributed to resource requirements driven by differing data compression techniques, workload management mechanisms, differing degrees of administrative complexity, and vendor database packaging and pricing.

Db2 11.5 compares favorably to Oracle Database 19c when implementing an enterprise OLTP shared data clustered database solution.

Index

LIST OF FIGURES

LIST OF TABLES

Quark + Lepton®

CORPORATE OFFICE
Boulder, Colorado USA

www.quarkandlepton.com
info@quarkandlepton.com

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