quark lepton

**WHITEPAPER** / MARCH 13, 2018 / QL914158

# Comparing IBM Db2 11.1 and Microsoft SQL Server 2017 for Analytics Processing

**Sponsored by: IBM**

## Market Situation

How to plan for and achieve agile decision making and operational resilience are two of the most pressing concerns faced by organizations of all sizes. The IT environments of modern businesses are characterized by increasing complexity and size, and extracting business insights from ever-expanding datasets requires more than just raw computing power.

Legacy technologies cannot sufficiently support the next-generation data mining and artificial intelligence (AI) technologies that drive business intelligence (BI) discovery. To meet these increasing requirements, a plethora of high-performance features have been added to popular databases to boost online analytical processing (OLAP) capability. In particular, in-memory tables and columnar processing techniques have been introduced to reduce latencies, eliminate I/O bottlenecks, and lower storage footprints associated with the large volumes of data used for analytics workloads. IBM Db2, Microsoft SQL Server, Oracle Database, and PostgreSQL continue to be dominant, enterprise-ready databases capable of analytics processing.

When introducing new technologies, decision makers must consider how these will integrate into existing IT assets. It is important to consider both flexibility and scalability to ensure growing workloads are supported. In terms of operational efficiency, the trend toward built-in tools, modern interfaces, and intuitive dashboards improves user productivity and drives opportunities for new insights.

Many analytics environments start off with an amalgamation of datasets collected by various online transaction processing (OLTP) systems using extract, transform, and load (ETL) tools. As these data warehouses increase in size, complexity, and diversity, the data must still be accessible across vendor-disparate platforms and environments. IBM Db2's leading data virtualization capabilities enable joining of Db2 tables with other heterogeneous data sources, such as Oracle Database, SQL Server, Hadoop, and others, to maximize analytics workload potential. Thus, organizations deploying high performance

databases for data warehousing and analytics to support agile decision making must consider all aspects of a solution, including integration and scalability as well as reliability, functionality, and cost.

This paper presents a cost/benefit case for use of two database platforms for analytics workloads—IBM Db2 11.1 and Microsoft SQL Server 2017 on x86 servers with Linux. Comparisons are of database installations in the financial services, retail, and discrete manufacturing industries. For comparable deployments, three-year costs of ownership for SQL Server 2017 are higher than Db2 11.1 with BLU Acceleration by 13 percent (Figure 1). Additional analysis for seven-year costs of ownership indicate SQL Server is 16 percent more costly compared to Db2. Higher cost disparities over time are largely driven by the increased ongoing costs such as software support and personnel requirements associated with SQL Server. Cost breakdowns for individual instillations may be found in the Cost Details section of this document.

Cost estimates are for use of x86-hardware for all installations. Both Db2 and SQL Server installations include use of Red Hat Enterprise Linux 7.2 (RHEL), VMware vSphere 6.5, and VMware vCenter Server 6.5. Enterprise software licenses and support costs were based on database costs whereas system software and support costs include operating system and virtualization.

Assumptions employed in constructing composite company profiles were based on information supplied by organizations worldwide that employ either Db2 and/or SQL Server databases, or have migrated from one to the other.

## OLAP Technology Differentiators

Perhaps the most integral part of any organization's IT infrastructure is the database platform, yet many businesses continue to rely on antiquated systems incapable of high performance workloads. Today's in-memory, columnar databases, optimized for analytics, can deliver BI faster than ever, allowing users to transform insights into competitive advantages. Organizations that build their business-critical applications on databases with these features are more prepared for and responsive to industry and market changes as a result.

**FIGURE 1:** Average Three-year IT Costs for Use of IBM Db2 11.1 and Microsoft SQL Server 2017 on Intel x86-based Systems

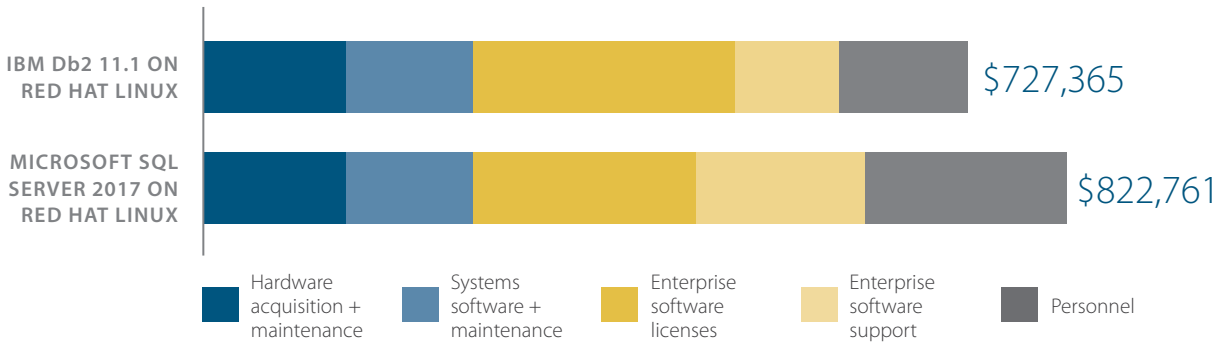


**SOURCE:** Quark + Lepton (March 2018)

Db2 11.1 with BLU Acceleration, the latest version of IBM's database offering, facilitates industry-leading OLAP performance through features such as in-memory columnar processing, data skipping, massively parallel processing (MPP), actionable compression, Single Instruction Multiple Data (SIMD) processing, and continuous data ingest (CDI). Table 1 provides an overview of these features.

IBM Db2 11.1—formerly Db2 for Linux, Unix, and Windows (LUW)—may be deployed on Linux and Windows on x86 servers; Linux and AIX on IBM Power Systems servers; Linux on IBM Z; or in Docker containers. No matter the hardware platform, organizations deploying Db2 benefit concurrently from high-performance as well as the advantages provided by an open source ecosystem. Db2 supports development using programming languages including Python, Java, Ruby, and PHP, and integrates with the IBM Cloud platform or the open source OpenStack platform.

Released in October 2017, the latest version of SQL Server brought the database platform, previously only available for Windows Server, to Linux, thus reaching a wider user base and delivering choice to

**TABLE 1:** Performance Optimization Features of IBM Db2 with BLU Acceleration

| PERFORMANCE OPTIMIZATION FEATURE | DESCRIPTION |
|---|---|
| **In-Memory Technology** | **Column-organized tables** can be created using BLU Acceleration. Db2 BLU is in-memory optimized, & is not limited by RAM limitations imposed by other in-memory technologies, including Oracle's. BLU dynamically loads data into memory, optimizing use of memory resources based on query workload. Memory is automatically tuned by the Db2 Self Tuning Memory Manager (STMM).<br><br>**Actionable Compression** minimizes performance impacts through use of a growing library of sort, join, & other operations that can be applied to encoded columns without needing to decompress the data.<br><br>**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.<br><br>**Single instruction, multiple data (SIMD)** allows BLU to apply a single instruction (e.g., predicate evaluation, join, grouping, arithmetic, R script) simultaneously to multiple data elements in parallel. |
| **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. |
| **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. 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. |

**SOURCE:** Quark + Lepton (March 2018)

customers. SQL Server 2017 may be deployed on-premises on Windows Server; Linux distributions such as RHEL, SUSE Linux Enterprise Server (SUSE), and Ubuntu; as well as in Docker containers.

Notable changes introduced in SQL Server 2017 include graph data processing, Adaptive Query Processing, automatic database tuning, and machine learning capabilities, summarized in Table 2.

## PERFORMANCE OPTIMIZATION

Although analytics workloads are drivers of business growth, transactional workloads typically constitute the bulk of daily operations for most organizations. Data access patterns for transactional workloads are distinct from those of analytics, thus organizations are likely to have systems that are only optimized for one type of workload. OLTP workloads 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. When choosing a database solution for analytics workloads, it is important to select one that can be easily integrated into existing infrastructure and will meet organizational needs.

Key advances in Db2 11.1 build upon established Db2 strengths such as automation and workload management. BLU Acceleration columnar tables coexist with traditional row-based data and improve OLAP performance, without the need for database administrators (DBAs) to create additional indexes, partitioned tables, or aggregates.

**TABLE 2:** Notable Features of Microsoft SQL Server 2017

| FEATURE | DESCRIPTION |
|---|---|
| **Adaptive Query** | Improves performance in SQL Server through three features:<br><br>Batch mode memory grant feedback—calculates & collects memory requirements for certain queries to optimize memory allocation for future queries.<br><br>Batch mode adaptive joins—enables dynamic switching of hash or nested loop joins after scanning of initial input to optimize performance for workloads that oscillate between small & large joins.<br><br>Interleaved execution—mitigates inefficient optimization caused by fixed cardinality estimates by enabling dynamic optimization using updated, more accurate cardinality estimates for multi-statement table-valued functions. |
| **Automatic Tuning** | Continuously monitors & analyzes queries to capture workload characteristics. Provides users with recommendations to alleviate query performance issues &/or enables automatic correction of these problems. |
| **Graph Database** | Enables creation of graph databases from complex interconnected data to model many-to-many relationships. Allows users to conveniently extract meaningful insights. |
| **Machine Learning** | Users can run R or Python through Machine Learning Services & access additional open source libraries. |

**SOURCE:** Quark + Lepton (March 2018)

For organizations using SQL Server for data warehousing, the SQL Server Integration Services (SSIS) tool can be used to migrate and integrate data. Analytical workloads on SQL Server are typically handled by SQL Server Analytics Services (SSAS), a data modeling engine that supports a variety of data visualization tools and client applications including Power BI and Excel. When installed on-premises, SSAS instances can run in one of three modes: tabular; multidimensional and data mining; or Power Pivot mode. For BI reporting, SQL Server Reporting Services (SSRS) can be implemented, offering mobile and web-based report generation and viewing.

## MASSIVELY PARALLEL PROCESSING

A key BLU Acceleration feature includes support for high performance analytics processing across MPP scale-out implementations. MPP greatly increases the scalability of workloads harnessing the capabilities of BLU, and support of both scale-up and scale-out strategies opens a field of new possibilities for expanding organizations.

MPP environments have evolved from symmetric multiprocessing (SMP) environments, in which growing requirements are met by scaling up systems. SMP systems contain multiple processors that share computing resources, whereas MPP environments typically contain multiple interconnected systems, or processing nodes, with independent processors and hardware resources (Figure 2). Data warehouse, or OLAP, workloads in MPP environments can concurrently leverage the processing power of multiple systems, with incremental performance scaling through addition or removal of nodes.

**FIGURE 2:** IBM Db2 11.1 BLU MPP Architecture vs Microsoft SQL Server 2017 SMP Architecture in On-Premises Deployment

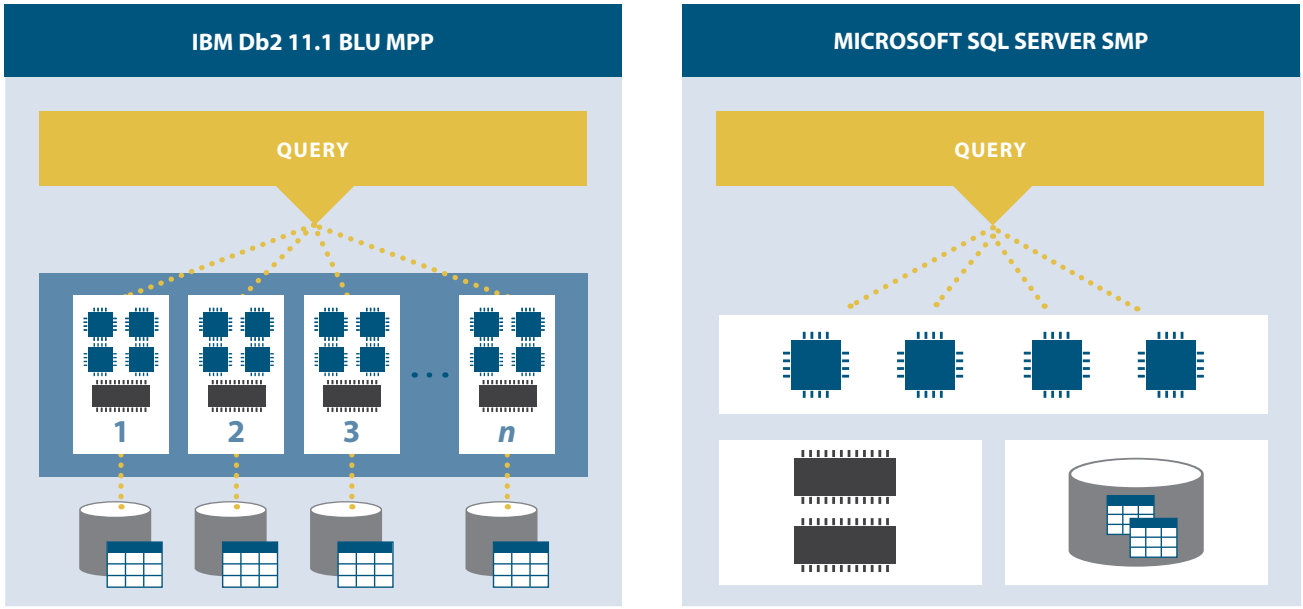


**Massively Parallel Processing (MPP) benefits from multi-node parallelism as well as processor parallelism**

**Symmetric Multi-Processing (SMP) only leverages processor parallelism**

**SOURCE:** Quark + Lepton (March 2018)

Db2 11.1's support of BLU Acceleration in environments using the Database Partitioning Feature (DPF) enables horizontal scaling of up to 1000 virtual or physical database servers. DPF scaling delivers more hardware to increase the numbers of queries served per hour and enables petabyte-scale OLAP workloads, without sacrificing performance. In addition to allowing organizations to scale BLU performance beyond previous limitations, users may now more easily leverage 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 and difficult to manage 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- and intra-partition parallel processing of queries enables near-linear MPP scale-out of BLU capabilities with additional partitions and hardware. Organizations with growing data assets may incrementally scale the hardware supporting BLU MPP. This approach tends to be more cost efficient than implementing SMP solutions.

BLU Acceleration in Db2 11.1 also introduced 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.

In contrast, SQL Server 2017 does not have comparable MPP capabilities. Although Microsoft offers MPP data warehousing through its Azure SQL Data Warehouse cloud service and the Microsoft Analytics Platform System (APS) appliances, SQL Server deployed on-premises is a SMP solution. APS appliances are offered by several hardware vendors—Dell, HPE, and Quanta Cloud Technology—with Microsoft acting as a single point of contact for support. Customers who wish to upgrade can choose to scale up their existing SQL Server system or migrate to an appliance or the cloud.

## IN-MEMORY COLUMNAR PROCESSING

In-memory, column-based databases have become a standard solution for analytical workloads for enterprises, drastically improving query processing performance compared to traditional row-based databases. Db2 BLU's columnar technology dynamically optimizes performance, minimizing latency associated with storage access by intelligently and automatically streamlining memory usage through data skipping and actionable compression techniques.

Db2's column-organized tables can scale beyond the limitations of RAM because BLU Acceleration dynamically loads data required by queries into memory. 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. Actionable compression in BLU further maximizes the capacity of memory and storage resources. Figure 3 depicts how various Db2 data reduction capabilities synergize to improve query performance and reduce memory requirements.

For analytics, SQL Server relies on in-memory columnstore indexes to improve performance and efficiency for storing, retrieving, and managing data. A deltastore accumulates small loads and inserts until the threshold for compressing into the columnstore has been reached. Through use of updateable

nonclustered columnstore indexes, SQL Server supports real-time operational analytics in hybrid workload environments. OLTP workloads utilize traditional row-based tables whereas OLAP workloads run against an automatically updated selective columnstore index of the same data.

For data warehousing, row-based tables can be converted to column-based tables through clustered columnstore indexes. For datasets with low cardinality, this approach improves disk usage as well as analytical query performance. Clustered columnstore indexes are not typically used with memory-optimized tables.

**FIGURE 3:** Data Reduction Capabilities Featured in Db2 BLU Acceleration for a Sample Query

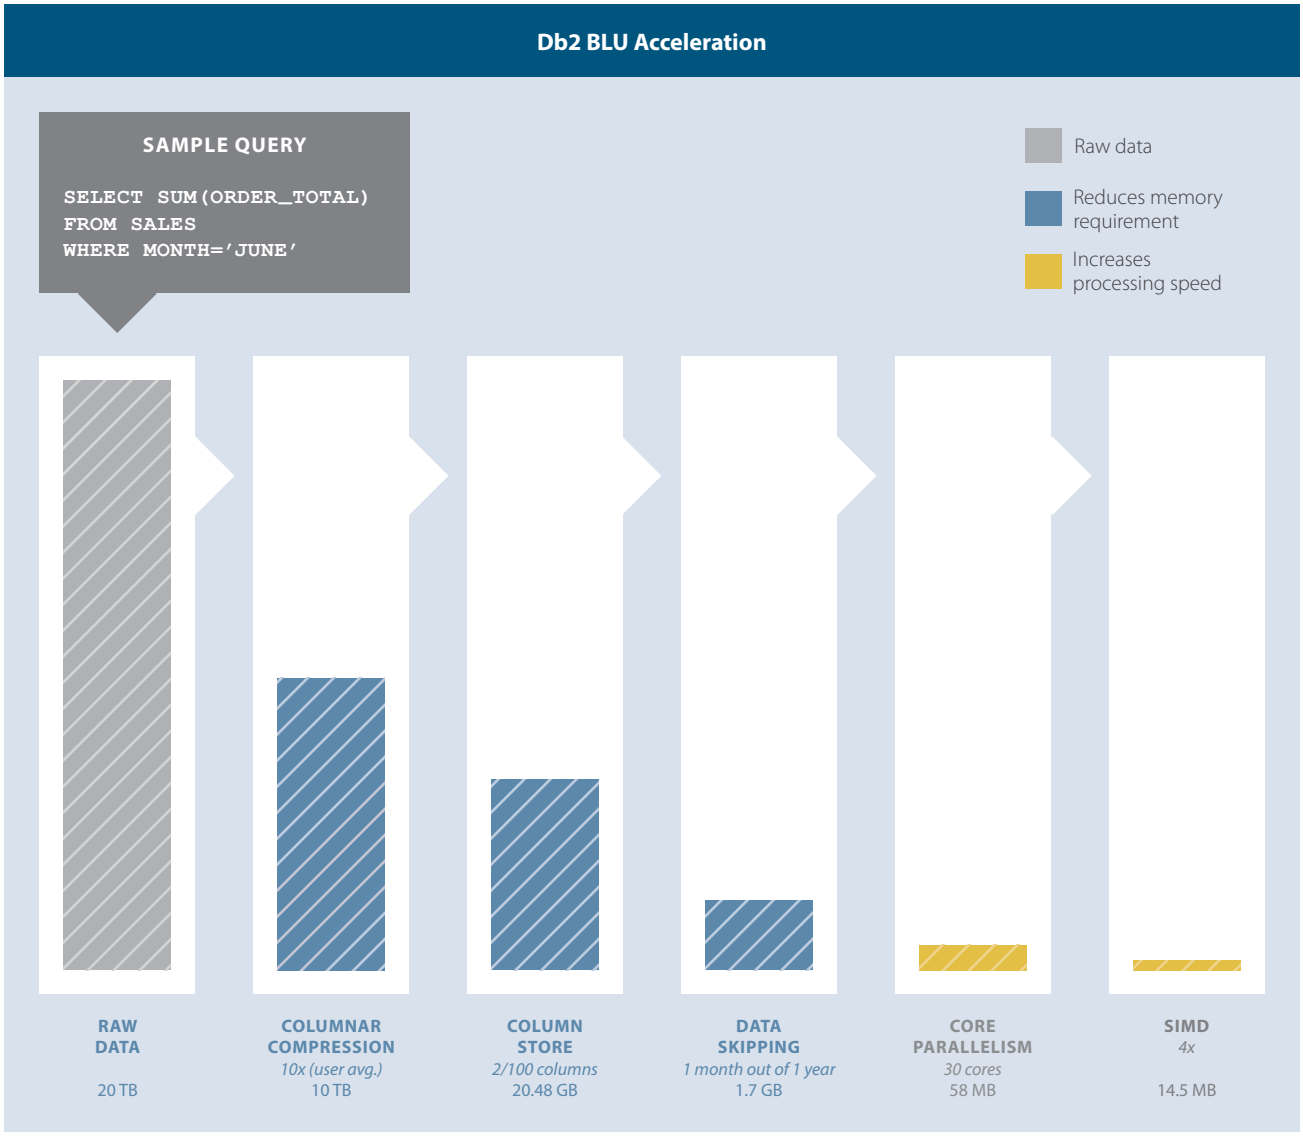


**SOURCE:** Quark + Lepton (March 2018)

## COMPRESSION

Introduced in Db2 10.1, adaptive compression implements a dual approach to data compression, utilizing both table- and page-level techniques. The table-level dictionary is used for classic row compression and a separate page-level dictionary is used to account for data repetition within pages of data. Adaptive compression is the default compression technique for row-organized tables in Db2, but administrators may specify which tables to exclude from page-level compression. SQL Server 2016 and 2017 offer row- or page-level compression for row-based tables. However, unlike Db2, SQL Server does not benefit from the ability to dynamically use both row- and page-level compression at the same time.

For column-organized tables, actionable compression is fully automatic and enables certain operations, such as data comparison, without incurring decompression overhead. Although different columns are uniquely encoded, encoding for a column across all database partitions remains 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, thus improving processing efficiency.

## SYSTEMS CONFIGURATION AND MANAGEMENT

With each iteration of Db2, simplification of DBA tasks has remained a central focus. Enhancements in automating system administration and management serve to reduce personnel overhead and costs.

Core database features allow Db2 DBAs to perform tasks with fewer, simpler actions, and in less time compared to SQL Server DBAs. 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 tools used for database tuning and performance optimization.

Db2 self-tuning memory manager (STMM) is one of the industry's most advanced self-tuning technologies. STMM is a memory-tuning capability that dynamically allocates memory resources among a database's memory consumers. STMM continuously determines database memory requirements, and reallocates resources to and from database shared memory when possible.

Db2 also uses automatic storage by default, simplifying DBA workloads by eliminating the need to manage explicit containers at the tablespace level. Databases created with automatic storage are associated with storage paths, enabling the database manager to automatically create, extend, and add storage containers in the storage path to meet the database's changing resource needs.

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.

Additional DSM capabilities, such as automatically setting database configuration and management parameters to optimize performance at database creation, reduce 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. Db2 can also collect database performance statistics in real-time, and optimize workloads using these statistics.

In addition to simplifying tasks for Db2 administrators, IBM has invested heavily in enabling organizations to migrate from other databases to Db2 rapidly, with ease. Enhancements in SQL compatibility, such as native support for Oracle's Procedural Language/Structured Query Language (PL/SQL) and open source PostgreSQL, along with the provision of a wide range of code, tools, and functions commonly employed by developers, has helped reduce personnel overhead.

In most cases, Db2 is advantageous in dynamic environments where integrated, centralized tools and autonomic capabilities allow for more efficient management and administration, and reduce the manual intervention required by other databases.

As SQL Server does not offer the breadth and sophistication of autonomic tools found in Db2, deploying SQL Server in dynamic environment tends to generate additional overhead because of necessary manual configuration and management requirements.

Automatic tuning is the primary feature used to boost query performance in SQL Server. The SQL Server database engine monitors query performance and attempts to correct any performance issues associated with SQL plan executions. It is a continuous process that collects performance statistics and applies tuning actions that improves performance for known workloads, reverting any actions that have not resulted in increased performance.

Although DBAs can mitigate some manual configuration and tuning tasks via automatic tuning, many capabilities found in Db2, such as those featured in STMM, do not have equivalents in SQL Server.

In addition, Db2's industry-leading integrated data virtualization and data federation features allow users to query from disparate data sources, such as Oracle Database, SQL Server, and others, along with Db2 data for more effective BI generation. In contrast, SQL Server's data virtualization capabilities are much more limited, and often require use of additional cloud-based tools such as Azure HDInsight.

## RELIABILITY, AVAILABILITY, AND SERVICEABILITY (RAS)

Businesses rely heavily on OLAP workloads to support decision making, and any interruptions to systems that support analytics applications may cripple an organization's resilience. As an integral component of IT infrastructure, a reliable database helps ensure business continuity and minimize downtime costs.

Numerous built-in Db2 high availability (HA) features incorporate autonomic technology to continuously protect against database downtime, even without DBA intervention. Automatic Client Reroute (ACR) redirects client applications to alternate servers in case of server failure. Db2 fault monitor facilities automatically monitor and restart database instances that have exited prematurely.

SQL Server offers comparable HA features through Always On availability groups. Availability groups are failover environments that consist of a set of primary databases, which failover together, and up to eight sets of active standby replicas.

## HARDWARE PLATFORMS AND OPERATING SYSTEMS

Vulnerabilities and exposure to risk have increased as organizations continue to add new applications across IT landscapes. More than ever, security is a critical concern when choosing a database platform, especially for workloads involving sensitive data. Although SQL Server and IBM Db2 both include impressive security features, the choice of the underlying operating system for any database may become a driver of significant security differences.

Db2 has been and continues to be available for highly-secure operating systems such as AIX on Power Systems or Linux on IBM Z. A comparison of statistics from the National Vulnerability Database for selected operating systems is detailed in Table 3. Until SQL Server 2017 on Linux was released, SQL Server was only available for Windows, which has received significantly more vulnerability advisories compared to other operating systems.

When expanding IT infrastructure, investing in integrated systems with long term support from a trusted vendor is critical in minimizing operational downtime as well as the costs associated with downtime. As the diversity of the IT environment increases, so does the difficulty in balancing interoperability with adopting innovative technologies.

Db2 provides optimal scalability and reliability to organizations of any size without the need to sacrifice flexibility. IBM Db2 can be deployed on-premises, in the cloud, or in a hybrid environment. Db2 11.1 is available in several editions, each with a variety of features, ranging from the Workgroup Server—recommended for small businesses—to the feature-rich Db2 Advanced Enterprise Server Edition (AESE).

**TABLE 3:** Comparative Operating System Vulnerability Data—January 2010 through September 2017

| CVSS SEVERITY LEVEL | Microsoft Windows Server | | Red Hat Enterprise Linux Server (RHEL) | | SUSE Linux Enterprise Server (SLES) | | IBM AIX | |
|---|---|---|---|---|---|---|---|---|
| *Version*<br>*Release date* | *2012*<br>*Oct '12* | *2016*<br>*Oct '16* | *6*<br>*Nov '10* | *7*<br>*Apr '14* | *11*<br>*Mar '09* | *12*<br>*Oct '14* | *7.1*<br>*Sept '10* | *7.2*<br>*Dec '15* |
| Critical | 7 | 4 | 12 | 11 | 1 | 10 | 0 | 0 |
| High | 404 | 196 | 70 | 39 | 10 | 13 | 3 | 0 |
| Medium | 221 | 152 | 21 | 18 | 24 | 12 | 2 | 0 |
| Low | 41 | 6 | 3 | 3 | 2 | 1 | 3 | 0 |
| TOTAL VULNERABILITIES | 673 | 358 | 106 | 71 | 14 | 36 | 8 | 0 |

**SOURCE:** Quark + Lepton (March 2018), NIST Computer Security Division, National Vulnerability Database, CVSS Metrics Versions 2 & 3

Licensing models include Processor Value Unit (PVU) or per Terabyte licensing for data warehousing workloads.

Along with being more operating system agnostic than SQL Server, Db2 is supported on multiple hardware platforms. Operating systems supporting Db2 11.1 include AIX, Microsoft Windows Server, as well as a variety of enterprise Linux distributions. In addition to x86-based hardware, Db2 can run on high performance hardware such as IBM Z and Power Systems.

Users of a synergistic stack of IBM solutions can benefit from industry-leading reliability and expert support from IBM. For example, Db2 is tightly integrated with AIX or Linux running on Power Systems, and can leverage PowerVM's robust virtualization technology. PowerVM is extremely scalable, offering highly granular workload management capabilities with minimal overhead. PowerVM enables users to run multiple operating systems in different partitions on the same server.

On the other side of the spectrum, businesses with smaller workloads can deploy Db2 on commodity x86-based hardware with open source software to benefit from the various tools and applications offered by the open source ecosystem.

In contrast, SQL Server has historically been restricted to Windows Server until the most recent release, which brought SQL Server 2017 to Linux. As with most first-generation technologies, SQL Server on Linux is less feature-rich than its Windows counterpart, lacking features such as SQL Server Analysis Services, Reporting Services, and mirroring for HA. Table 4 lists some of the features not available in the General Availability (GA) release of SQL Server 2017 on Linux.

**TABLE 4:** Standard Microsoft SQL Server Features Currently Unavailable in SQL Server 2017 for Linux

| FUNCTIONAL AREA | UNSUPPORTED FEATURE & SERVICES FOR SQL SERVER 2017 FOR LINUX |
|---|---|
| **Database Engine** | Transactional replication, merge replication, Stretch Database (DB), Polybase, distributed query with 3rd-party connections, system extended stored procedures, Filetable, FILESTREAM, buffer pool extension |
| **High Availability** | Database mirroring |
| **Security** | Extensible Key Management, Active Directory (AD) Authentication for Linked Servers, AD Authentication for Availability Groups, 3rd Party AD tools |
| **Services** | SQL Server Browser, SQL Server R services, StreamInsight, Analysis Services, Reporting Services, Data Quality Services, Master Data Services |
| **SQL Server Agent** | Subsystems such as CmdExec, PowerShell, Queue Reader, SSIS, SSAS, & SSRS; alerts, Log Reader Agent, Change Data Capture, Managed Backup |

**SOURCE:** Quark + Lepton (March 2018)

## Cost Details

Installation and cost details across three industries, financial services, retail, and discrete manufacturing, are presented in Tables 5 and 6 respectively.

**TABLE 5:** Profile Installations for Total Cost of Ownership Comparisons

| FINANCIAL SERVICES COMPANY | RETAIL COMPANY | DISCRETE MANUFACTURING COMPANY |
|---|---|---|
| **BUSINESS PROFILE** | | |
| Employees: 400+<br>Assets: $5 billion<br>Locations: 10 branches | Employees: 15,000<br>Revenue: $2 billion<br>Locations: 500+ stores | Employees: 10,000<br>Revenue: $3+ billion<br>Locations: 5 manufacturing facilities |
| **APPLICATIONS** | | |
| Customer contribution analysis<br>Customer behavior modeling & scoring<br>Operational CRM<br>Risk assessment & reporting<br>Asset & liability management<br>Cost optimization | Marketing & sales analysis<br>Promotions & pricing analysis<br>Loyalty program management<br>Customer profiling<br>Logistics analysis<br>Multichannel data consolidation | Sales & promotions analysis<br>Customer segmentation<br>Production & inventory forecasting<br>Logistics planning<br>Management reporting |
| **IBM Db2 11.1 WITH BLU ACCELERATION ON x86** | | |
| 16 cores enabled<br>Db2 11.1 AESE<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.15 FTE DBA* | 64 cores enabled<br>Db2 11.1 AESE<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.5 FTE DBA* | 32 cores enabled<br>Db2 11.1 AESE<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.3 FTE DBA* |
| **MICROSOFT SQL SERVER 2017 ON x86** | | |
| 16 cores enabled<br>SQL Server 2017 Enterprise Edition<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.3 FTE DBA* | 64 cores enabled<br>SQL Server 2017 Enterprise Edition<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.75 FTE DBA* | 32 cores enabled<br>SQL Server 2017 Enterprise Edition<br>*RHEL, VMware ESXi 6.5, VMware vCenter Server 6.5*<br>*0.5 FTE DBA* |

**SOURCE:** Quark + Lepton (March 2018)

Db2 11.1 with BLU Acceleration costs were calculated using Terabyte licensing for AESE, along with two years of software support in addition to the included 12 months of support. System software and support costs account for use of Red Hat Enterprise Linux 7.2 and VMware ESXi hypervisor.

SQL Server 2017 costs were based on per core licensing for Enterprise Edition with three-year Microsoft Software Assurance. System software and support costs include use of Red Hat Enterprise Linux 7.2 and VMware ESXi hypervisor.

Hardware costs for all installations include latest generation x86-based servers and three-year maintenance costs.

Personnel costs were calculated based on estimated annual salaries of $85,991 for SQL Server DBAs, and $88,966 for Db2 DBAs. These estimates were based on industry standard DBA 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.

Facilities costs are not included in this paper. Facilities cost differences were insignificant between Db2 and SQL Server on x86 deployments of similar sizes.

**TABLE 6:** Three-Year Cost Breakdown—Use of IBM Db2 11.1 Compared to Microsoft SQL Server 2017 for Analytics Processing

| | IBM Db2 11.1 ON x86 SYSTEMS WITH RHEL | | | MICROSOFT SQL SERVER 2017 ON x86 SYSTEMS WITH RHEL | | |
|---|---|---|---|---|---|---|
| COST CATEGORIES | FINANCIAL SERVICES COMPANY | RETAIL COMPANY | DISCRETE MANUFACTURING COMPANY | FINANCIAL SERVICES COMPANY | RETAIL COMPANY | DISCRETE MANUFACTURING COMPANY |
| Hardware acquisition + maintenance | 59,643 | 234,410 | 119,286 | 59,643 | 234,410 | 119,286 |
| System software licenses + support | 51,811 | 207,245 | 103,622 | 51,811 | 207,245 | 103,622 |
| Enterprise software licenses + support | 173,600 | 520,800 | 347,200 | 159,667 | 638,669 | 319,334 |
| Personnel | 57,549 | 191,830 | 115,098 | 111,212 | 278,030 | 185,354 |
| TOTAL ($) | 342,603 | 1,154,285 | 685,206 | 382,333 | 1,358,354 | 727,596 |

**SOURCE:** Quark + Lepton (March 2018)

## Conclusions

The introduction of SQL Server 2017 on Linux signifies Microsoft's commitment to bringing the database solution to a new user base as well as opening the platform to next-generation applications and the open source ecosystem.

Organizations with complex, expanding data assets need a database that reliably works with different types of data, supports a variety of applications on different types of devices, interoperates with existing IT infrastructure, and scales with ease to accommodate future workloads growth.

IBM Db2 11.1 with BLU Acceleration supports high-performance analytics workloads to rapidly deliver business intelligence. Db2's industry-leading workload management boosts operational efficiency while maintaining industry-leading reliability and availability through built-in RAS features. BLU MPP allows for scale-out of Db2 environments to seamlessly support growing business requirements. Autonomic capabilities built into Db2 simplify DBA tasks, reducing personnel requirements and costs. Three-year costs for use of SQL Server 2017 for OLAP workloads averaged 13 percent more than use of Db2 11.1 for the profile installations used in this paper.

In contrast, SQL Server 2017 requires higher levels of DBA intervention for manual configuration and tuning. SQL Server is a SMP solution without true scale-out capability and works best for smaller OLAP workloads.

Although the latest version of SQL Server has been released for Linux and Docker container implementations, the technology is far less mature than that of Db2. The current version of SQL Server on Linux does not include some key features available for the Windows Server version.

Microsoft's commitment to user flexibility should be lauded, but organizations considering a database solution to support next-generation analytics should look elsewhere. Db2 is not only an industry-leading database supporting analytics and mixed workloads, but it also enjoys a history of user satisfaction as well as consistent support and investment by IBM.

# Index



## LIST OF TABLES



## LIST OF FIGURES



## LIST OF REFERENCES


Release Notes for SQL Server 2017 on Linux found at docs.microsoft.com/en-us/sql/linux/sql-server-linux-release-notes

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

**Quark + Lepton®**

**CORPORATE OFFICE**
Boulder, Colorado USA

www.quarkandlepton.com
info@quarkandlepton.com