WHITEPAPER / MARCH 13, 2018 / QL052976

Comparing IBM Db2 11.1 and Microsoft SQL Server 2017 for Transaction Processing

Sponsored by: IBM

Market Situation

Databases remain the cornerstone of business IT infrastructure, supporting both analytics workloads to facilitate decision making, as well as critical transaction processing workloads to drive revenue and growth. Due to the current rate of unprecedented technological advancement, organizations must continuously modernize in order to maintain competitive advantage. This can be achieved through digitization of legacy processes and implementation of next generation applications, improving the effectiveness and efficiency of business operations. In doing so, organizations should still prioritize the ubiquitous concerns of improving reliability, availability, and security.

In service industries such as healthcare, telecommunications, and financial services, data assets tend to be large, complex, and sensitive. Customers value their privacy, yet require rapid, reliable access to personal accounts and records. The challenge for organizations becomes how to adopt new technologies to satisfy customer demands without compromising transaction processing systems.

Industries such as ecommerce, manufacturing, and distribution are subject to the constraints of supply chain management. Organizations in these industries must improve efficiency, even if only marginally, to meet the demands of customers who value time and convenience above all else. System downtime can cause inefficiencies that reverberate across the entire supply chain, potentially resulting in lost revenue, negative press, and/or regulatory penalties.

Database platforms supporting core online transaction processing (OLTP) may directly affect an organization's productivity. When selecting a new database solution, in addition to standard considerations such as costs, businesses should pay attention to aspects such as risk tolerance, ease of integration, and effects on personnel productivity. Ideally, to minimize downtime, a database platform should have high levels of reliability, availability, and serviceability (RAS), along with robust security features.

In addition, autonomic features and capabilities that facilitate ease of system management will streamline database administrator (DBA) tasks, and interoperability with modern, intuitive applications can increase overall employee productivity.

This paper compares the downtime impacts experienced by users of two industry leading databases, IBM Db2 11.1 and Microsoft SQL Server 2017 for organizations in the discrete manufacturing, telecommunications, and payment processing industries. Financial impacts were calculated based on industry risks as well as average planned and unplanned downtime reported by users in each industry. Db2 sample installations include Power Systems with Linux and/or AIX whereas SQL Server 2017 installations used Windows Server and x86 hardware.

Although many organizations view planned downtime as a necessary evil, installations, migrations, and/or upgrades can still be disruptive and should be minimized to improve operational efficiency. Planned downtime can be reduced by techniques such as online updates, autonomic system management and configuration, and preventative maintenance. Unplanned downtime, typically caused by unforeseen circumstances such as natural disasters or security breaches, are more catastrophic and more difficult to recover from. Organizations need to maintain high levels of fault tolerance and employ excellent disaster recovery strategies to mitigate the cost of unplanned downtime.

Organizations using Microsoft SQL Server, on average, experience 4.5 times higher planned downtime and 6.6 times higher unplanned downtime than Db2 organizations. Estimated costs of downtime for Db2 11.1 average 78 percent less for planned outages, and 82 percent for unplanned outages, compared to SQL Server. Three-year planned and unplanned downtime costs for the sample industries are shown in Figures 1 and 2 respectively.

SQL Server's higher costs of downtime can be largely attributed to less efficient recovery processes resulting from complex architectures that require constant DBA intervention. Db2 pureScale environments are more autonomic, with redundant architecture to provide continuous availability. Costs of downtime are further reduced by Db2 features such as rolling maintenance, automatic backup, and automatic failover detection and recovery.

FIGURE 1: Three-year Planned Downtime Costs

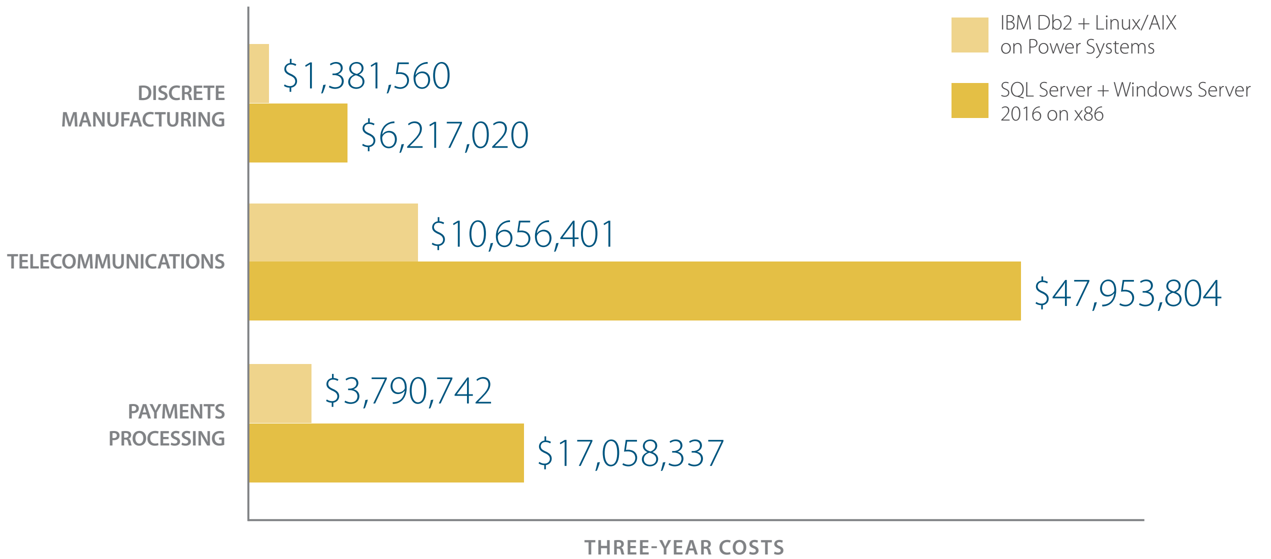

SOURCE: Quark + Lepton (March 2018)

FIGURE 2: Three-year Unplanned Downtime Costs

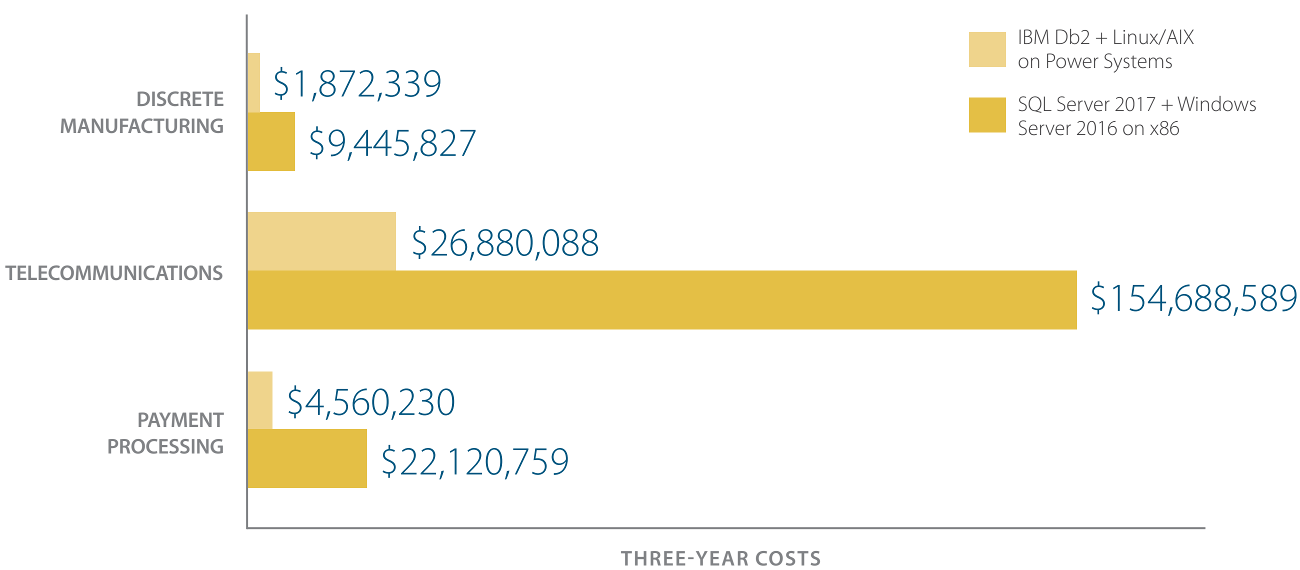

SOURCE: Quark + Lepton (March 2018)

OLTP Technology Differentiators

IBM Db2 and Microsoft SQL Server compete in the enterprise software market as popular database solutions. Although Microsoft has offered users a cost-efficient database platform with satisfactory reliability, there was the limitation of being locked into the company's Windows operating system ecosystem prior to the availability of SQL Server 2017 on Linux. Despite the advancements made to garner more users, the Linux version of SQL Server is not fully featured. Table 1 lists various critical capabilities SQL Server on Linux lacked in the General Availability (GA) release, but featured in SQL Server for Windows.

Db2, on the other hand, has historically been available for a wide range of operating systems and hardware architectures, such as x86-based servers, Power Systems, and System z. Supported on Linux, Unix, and Windows operating systems, Db2 11.1 excels in the cross-platform environments found in most organizations.

For business-critical OLTP workloads, Db2 harnesses pureScale clustering capabilities to provide enterprise-grade continuous availability and processing performance. Using Db2 pureScale, organizations can grow their database and OLTP workloads organically and seamlessly to meet business requirements. For even greater fault tolerance as well as flexibility, Db2 on Power Systems with Linux or AIX can leverage PowerVM virtualization capabilities.

Organizations using SQL Server for core OLTP workloads tend to experience lower levels of reliability, flexibility, and scalability. High availability (HA) features for SQL Server require use of Always On availability groups along with Windows Server Failover Clustering (WSFC), and have limited compatibility with third-party software and hardware. Users who wish to leverage the full functionality of SQL Server are also limited to using Windows Server as their operating system.

TABLE 1: 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)

SCALABILITY AND PERFORMANCE

Designed to run even the most demanding enterprise workloads, Db2 pureScale's processing capabilities scale efficiently by optimizing workload balance between multiple database members. Db2 pureScale supports up to 128 members that share storage, but process requests independently. Figure 3 depicts the continuous availability architecture of Db2 pureScale, where cluster caching facilities (CFs) and cluster services (CS) facilitate workload management and failover between members.

Extreme scalability in pureScale is a result of CFs and the high speed remote direct memory access (RDMA) InfiniBand network linking members, directly improving performance for OLTP workloads. For all but the most demanding workloads, pureScale environments can also be run on RDMA over Ethernet or standard TCP/IP networking. As members are added, the capability to process transactions also scales in a near-linear manner due to the centralized locking and caching mechanisms found in CFs. CFs provide a shared communication area that members can use as clusterwide shared memory, significantly reducing the communications overhead between members that generally increases when systems scale out.

Databases without these features, such as SQL Server, are much less efficient at scaling out. In addition, scaling out for SQL Server tends to be a highly manual process, commonly involving third-party software, whereas scaling out in a Db2 pureScale environment simply involves adding additional members due to application transparency. Members can be added without the need to reconfigure applications or storage, mitigating the planned downtime normally required for other databases. There are no SQL Server features equivalent to Db2 pureScale.

FIGURE 3: High Availability Architecture of Db2 HADR with pureScale

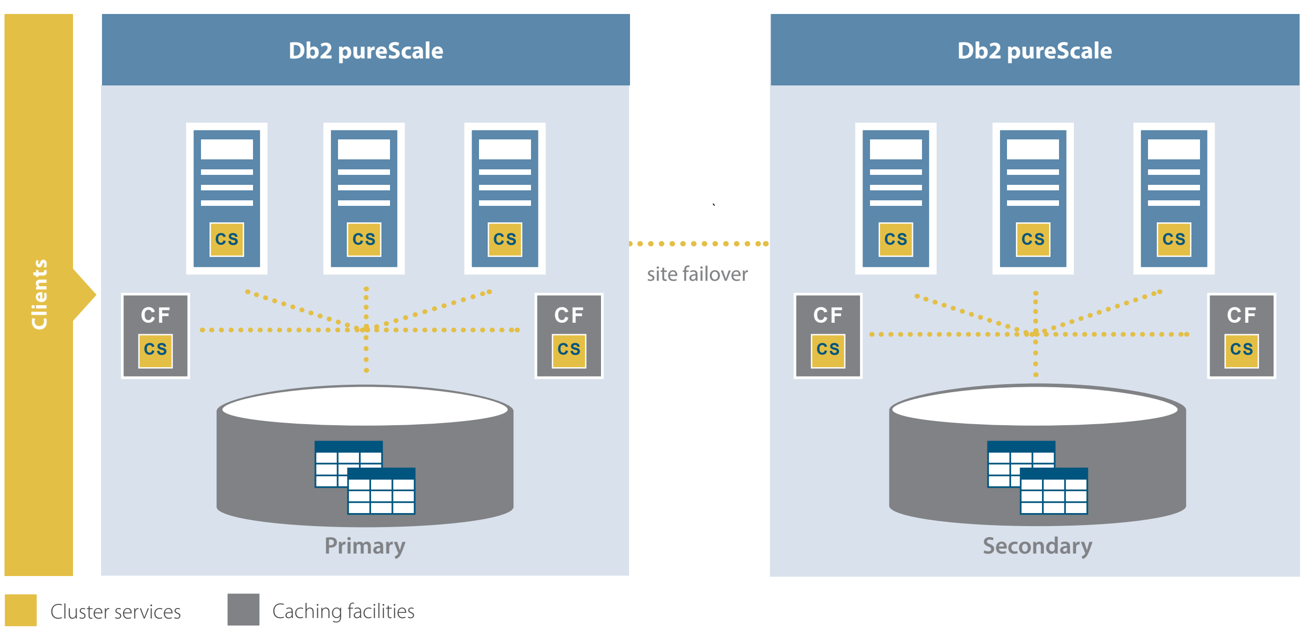

SOURCE: Quark + Lepton (March 2018)

OPTIMIZATION AND AUTOMATION

As database products evolve to meet the needs of growing businesses, such as increased availability and security features, few do so with the goal of increasing DBA efficiency, and many tend to implement tools requiring high levels of manual configuration. IBM Db2 pureScale is positioned to meet demanding reliability requirements without sacrificing performance or personnel efficiency.

For analytical workloads influencing decision makers, intuitive applications and modern dashboards with high levels of customization allow users to effectively extract business intelligence (BI). On the other hand, systems with high levels of automated self-recovery and tuning functionality, as oppose to DBA intervention, contribute to superior transaction processing efficiency.

Db2 has been designed with emphases on simplification and automation to improve personnel efficiency. Db2's autonomic environment minimizes personnel requirements compared to SQL Server, resulting in lower costs of ownership.

For example, DBA tasks can be streamlined using Data Server Manager (DSM), a tool that consolidates monitoring, tuning, configuration, and administration of multiple Db2 databases across multiple systems through one browser-based graphical user interface (GUI).

In terms of resource management, capabilities such as self-tuning memory, automatic storage, and adaptive compression help reduce DBA workloads. A list of Db2 11.1's performance optimization features, some of which are enabled by default and others manually, may be found in Table 2.

Db2's self-tuning memory manager (STMM) automatically and iteratively adjusts configuration settings to allocate memory resources optimally, contributing to improved 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).

Db2 automatic storage, on by default, simplifies storage management by only requiring specification of storage paths and shifting the responsibility of creating, extending, and adding containers to the Db2 database manager. Furthermore, table and index compression are automatic once specified at table creation.

DSM also configures databases automatically to optimize performance at database creation and 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 that performance is not negatively impacted. Db2 can also collect database performance statistics in real-time, and optimize for workloads using these statistics.

Conversely, because SQL Server does not offer the breadth and sophistication of autonomic tools found in Db2, DBAs are responsible for more database configuration and management tasks, especially in complex enterprise environments. Automatic tuning is the primary feature used to boost query performance in SQL Server. The SQL Server database engine continuously monitors query performance of current query execution plans and regresses to previous plans, either manually or automatically, if current access plans are determined to be less efficient based on statistics collected by the database engine. These corrective actions often involve manual intervention from DBAs; as advised by Microsoft, “whenever you notice the plan choice regression, you should find some previous good plan and force it instead of the current one.”

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

Storage optimization technology in Db2 also tends to be more efficient than that of SQL Server. For organizations looking to maximize storage efficiency, Adaptive Compression can achieve higher compression rates than classic row compression techniques.

Adaptive Compression uses a dual approach, combining table- as well as page-level compression for row-organized tables, allowing organizations to benefit from both compression techniques simultaneously. SQL Server offers comparable capabilities in row- and page-level compression.

TABLE 2: Performance Optimization Features of Db2 11.1 pureScale

PERFORMANCE OPTIMIZATION FEATURE	DESCRIPTION
Self-tuning Memory for Single Partition Databases	Automatically & iteratively adjusts memory & buffer pool configurations without DBA intervention.
Automatic Storage	Enables database manager to govern container & space allocation for table spaces.
Data Compression	When specified at table creation, table & index compression is fully automatic. Temporary tables & indexes are compressed automatically by default.
Automatic Backups	Automatically backs up database fully to various types of media.
Automatic Reorganization	Periodically checks to see if tables & indexes are fragmented & automatically schedules defragmentation.
Automatic Statistics Collection	Collects table statistics, asynchronously or synchronously, to optimize query execution plans.
Health Monitor	Continuously monitors database environment, without performance overhead, to inform & advise DBAs on system health.
Utility Throttling	Regulates maintenance utilities such as statistics collection, back-up operations, rebalancing operations, & asynchronous index cleanup to minimize performance impacts.

SOURCE: Quark + Lepton (March 2018)

Business Continuity

Threats to core business systems, such as security breaches, remain omnipresent due to the evolving sophistication of cyberthreats and increasing reliance of users on online business processing.

For some industries, such as ecommerce and banking, downtime can directly impact the bottom line. Downtime can manifest from a variety of sources, and risks of downtime increases with the diversity of components in an organization's infrastructure. Integration and security are essential considerations when making any changes to core IT infrastructure.

For most enterprises, especially those handling sensitive customer information or confidential data, exposure to downtime is increased by the ever-expanding channels of digital access. In the complex environment of electronic transactions using hardware and software from a multitude of vendors, careful implementation and integration of new technologies is required to prevent system failures. Legacy technology requires updates and patches to maintain stability and avoid downtime. Additionally, it has become critical for organizations dealing with sensitive information to implement robust security policies to minimize the risks of intrusion.

Security breaches have increased in recent years as cybercriminals develop more refined and sophisticated methods of intrusion (Figure 4). The Identity Theft Resource Center (ITRC) reported a total of 1,339 security breaches affecting over 174 million records in 2017, an 18.5 percent increase over the 1,091 breaches tracked in 2016. Breaches tracked by the ITRC originate from media sources and/or notifications from government agencies and typically include exposure of sensitive information such as

FIGURE 4: Data Breaches Tracked by the Identity Theft Resource Center

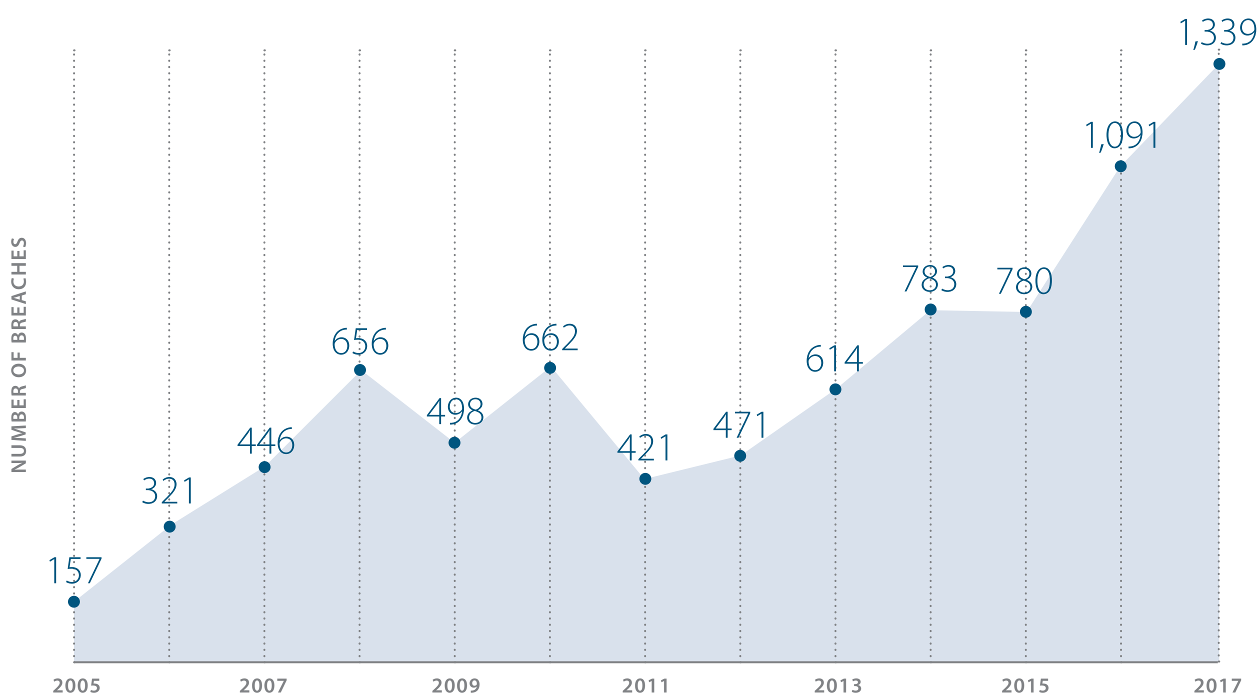

SOURCE: Quark + Lepton (March 2018), Identity Theft Resource Center (ITRC) Breach Statistics individual names, social security numbers, driver license numbers, medical records, and/or financial records. As cybersecurity threats become more sophisticated, breaches in any part of an organization's IT infrastructure can impact other IT assets.

For example, differences in the environments in which these databases run can have implications on core infrastructure security. Statistics from the National Vulnerability Database for selected operating systems used by SQL Server and/or Db2 are detailed in Table 3. Windows Server, the only platform SQL Server has been supported on until the most recent edition, has reported significantly more vulnerabilities than enterprise Linux and IBM AIX.

For additional security and workload management capabilities, Db2 on Power Systems running Linux or AIX can also leverage the advantages of PowerVM virtualization, which enables highly granular allocation of computing resources. PowerVM's hardware-based logical partitions (LPARs) provide added security through robust isolation and enable organizations to increase server utilization by running mixed workloads.

Power Systems with PowerVM can support various operating systems, such as AIX and/or Linux, concurrently on the same server. In addition, Active Memory Sharing and Virtual I/O Server (VIOS) virtualization features increase workload efficiency by flexibly allocating memory and I/O resources between virtual machines (VMs). VIOS also allows VMs to share a pool of virtualized storage. Microsoft Hyper-V and other x86-based hypervisors provide only limited support for these capabilities.

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 Release date	*2012 Oct '12*	*2016 Oct '16*	*6 Nov '10*	*7 Apr '14*	*11 Mar '09*	*12 Oct '14*	*7.1 Sept '10*	*7.2 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

AVAILABILITY AND STABILITY

For most organizations, making changes to critical OLTP systems presents many more challenges than just meeting budget and time constraints. Modernizing core IT infrastructure such as database platforms may disrupt personnel productivity and increase exposure to business disruption. Introducing components that have low interoperability with an existing infrastructure can increase risk of unexpected downtime and may prolong the recovery process.

While SQL Server 2017 port to Linux is intended to reach a wider spectrum of users, it is still a first-generation technology with limited features. If customers want to access all features in SQL Server, they are limited to deploying SQL Server on Windows Server, generally regarded as a less reliable operating system compared to Linux or Unix-based environments.

In contrast, Db2 remains highly secure, reliable, and available. Db2 integrates tightly with IBM software and hardware, and also takes advantage of a wide ecosystem of innovative applications offered by multiple vendors. Many Db2 features were implemented by IBM to satisfy requests from a committed customer base. In addition to support, IBM continues to invest in Db2 by adding compatibility with new technologies as well as maintains support for a wide range of operating systems.

Db2 offers additional fault tolerance, flexibility, and scalability capabilities with the pureScale shared disk clustering solution for on premises and hybrid deployments. Db2 pureScale is available for Linux and AIX on Power Systems and was designed to provide scalability and continuous availability for OLTP workloads across densely populated data center and enterprise environments.

To reduce business disruptions, Db2 also includes the High Availability Disaster Recovery (HADR) feature. HADR protects against data loss during site failures through remote replication of the primary database to the standby database(s). HADR minimizes downtime through rerouting client requests to standby databases during primary database failure and recovery.

HADR supports reads on standby, allowing organizations to shift read-only workloads to HADR standby members to streamline the performance of the primary database. This enables organizations to reduce their costs of ownership by improving the utilization of database resources while maintaining industry-leading HA.

Microsoft SQL Server offers comparable HA and DR functionality through Always On availability groups that consist of a set primary databases, which failover together, and up to eight sets of active standby replicas. Similarly, only primary databases are read/write capable, and secondary replicas do not support transactions from end users or applications.

When creating an availability group, SQL Server on Windows Server will use WSFC by default, but SQL Server 2017 on Linux requires use of external cluster managers from different vendors, such as Pacemaker, introducing more complexity to the stack. Each replicated set of the availability databases must reside on a separate node in the failover cluster. Due to this, nodes supporting secondary replicas do not contribute to processing performance and may be cost inefficient compared to Db2 pureScale.

SQL Server allows replicas to support read-only workloads, but it is a different configuration and does not support HA. Administrators must prioritize whether read-scaling or HA is more important, as SQL Server availability groups cannot use both at the same time, unlike Db2 HADR.

To maintain continuous availability, organizations may leverage Db2 pureScale's capabilities such as built-in failure detection, recovery automation, and data-sharing architecture to complement HADR's capabilities in protecting against site failures. HADR supports up to three remote standby databases to takeover when the systems supporting the primary database fail.

Standby databases are synchronized via TCP/IP, and separating them geographically from primary databases can help mitigate downtime during complete site failures. Db2 pureScale can distribute members across geographic locations using geographically dispersed Db2 pureScale cluster (GDPC) technology, increasing the level of disaster recovery protection. HADR can reduce recovery time from minutes to seconds, eliminating the need to restart the failed server and reinitiate the primary database.

For critical transaction processing workloads, Db2 pureScale minimizes downtime by maintaining availability during failure. During member failures, Db2 cluster services will automatically restart any failed resources or members in the Db2 pureScale instance. Members can be restarted either on their original host or a different host, without affecting the remainder of the instance. During recovery, rollback of in-flight transactions is prioritized to free locks and maintain consistency.

Db2 cluster services also provide automatic, heartbeat failure detection and recovery. Failed members are automatically isolated from the cluster system. During recovery of a failed member, the database instance remains available and accessible to all other cluster members. Applications and clients, whose connections to the failed member are momentarily disrupted, are automatically rerouted to another member in the cluster. During site failures, pureScale with HADR provides disaster recovery by failing over to a secondary cluster at a different site.

Cluster Facilities (CF) contribute to rapid failover and recovery by providing a shared buffer pool that coordinates caching across cluster members. CFs contain data from local buffer pools and can act as a shared communication area, which server members may use as cluster-wide shared memory. In addition to 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.

SQL Server Always On with Windows Server Failover Clusters do not have the advanced centralized cache and lock features of Db2 cluster facilities, nor the automated system management and workload balancing features found in pureScale environments. During failures, WSFC is less agile than pureScale because metadata is distributed to each node and changes to a node or cluster resource's status must be repropagated to all other nodes.

Figure 5 depicts the architecture of SQL Server Always On availability groups with Windows Server Failover Cluster.

SECURITY AND RISK REDUCTION

To minimize risks of downtime, organizations must implement effective security measures as well as robust high availability and disaster recovery strategies. Db2 encryption capabilities are integrated into core processing and apply to data at rest 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 supports integration with third-party key management, which are transparent to applications and schemas, and protect both the database and any backup images created. This symmetric encryption scheme automatically detects and exploits hardware acceleration for cryptographic operations, such as Intel AESNI.

Db2 utilizes the Federal Information Processing Standard (FIPS) 1402 certified cryptographic libraries and cryptographic algorithms that meet the requirements of National Institute of Standards and Technology (NIST) SP 800131a. Db2 on AIX can further leverage the AIX encrypted file system to protect all files in the system.

DBAs can also use row and column access control (RCAC) to manage user access to tables at the row or column-level, enabling an additional level of security and customization. Label-based access control (LBAC) allows administrators to further tailor access to meet an organization's needs by enabling read and/or write access based on specific criteria.

FIGURE 5: SQL Server Always On Availability Groups on Windows Server Failover Cluster

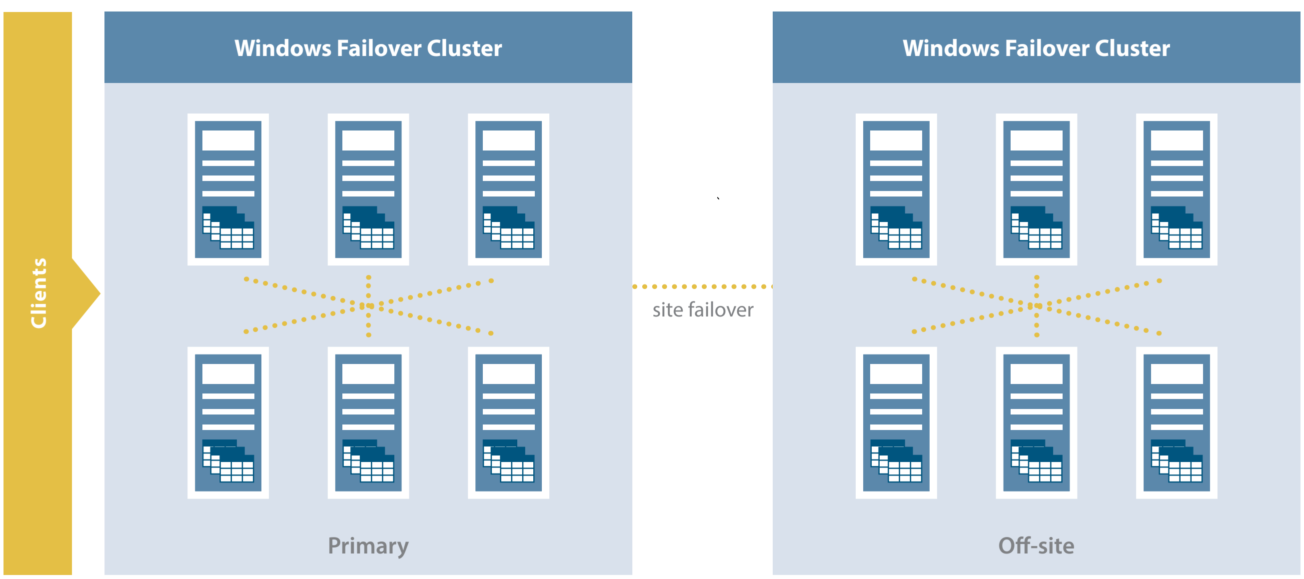

SOURCE: Quark + Lepton (March 2018)

Downtime Cost Factors

Cost of downtime comparisons presented were based on sample installations in the discrete manufacturing, telecommunications, and payment processing industries (Table 4).

Discrete Manufacturing. Costs included lost sales; idle and underutilized capacity; handling of delivery delays; additional inventory carrying costs; costs of change for procurement, production and logistics processes; customer billing and payments delays; late delivery and imperfect order penalties; and costs of remedial actions such as rebates and discounts required to win back customer business.

Telecommunications. Costs included regulatory fees, breach of contracts with third-party vendors and suppliers, loss of revenue due to service level agreements with customers, damaged reputation, decreased personnel productivity, increased customer attrition, as well as equipment and infrastructure replacement costs.

Payment Processing. Costs included customer attrition (lost customer income); lost transaction fees (including ATM/debit fees, and fees for transactions conducted online and through call centers); and other costs, (including lost interest, lost customer acquisition expenditure, as well as productivity loss by branch call center, and other customer-facing staff during outages). Costs of mobile banking outages were included in customer attrition and lost fee income.

TABLE 4: Profile Installations for Cost of Downtime Comparisons

DISCRETE MANUFACTURING COMPANY	TELECOMMUNICATIONS COMPANY	PAYMENT PROCESSING COMPANY
BUSINESS PROFILE		
Employees: 3,500 Revenue: $1 billion Locations: 15 production facilities	Employees: 15,000 Revenue: $15+ billion Subscribers: 120 million	Employees: 5,000+ Revenue: $6+ billion POS locations: 15+ million
APPLICATIONS		
Inventory Management, Production Planning, CRM, Financials, Billing & Payments, Supply Chain Management	CRM, Billing & Payments, Enterprise Asset Management, Facilities Management	CRM, Billing & Payments, Risk Management, Analytics
SQL SERVER ANNUAL PLANNED & UNPLANNED DOWNTIME		
21.3 hours	45.3 hours	18.2 hours
DB2 SERVER ANNUAL PLANNED & UNPLANNED DOWNTIME		
4.2 hours	7.9 hours	3.9 hour

SOURCE: Quark + Lepton (March 2018)

Conclusions

Core information technology systems become increasingly complex with the integration of next-generation software and hardware. As businesses transform their operations to meet the needs of modern customers, they may also introduce new risk factors to their infrastructure. The need to continuously implement innovative functionality to maintain competitive advantage, in conjunction with the rising sophistication of cybercrime, increase the risks of system downtime and exposure of sensitive company and/or customer data.

Enterprises rely more than ever on a transaction processing database platform that integrates with a variety of next generation applications and offers web and mobile access to users and customers. Databases must not only be high performing, but must also maintain high levels of availability and security. While many databases offer tools and add-ons to meet these requirements, not all databases implement these features in an integrated, efficient manner.

Db2 11.1 pureScale clustering technology minimizes disruptions for critical business operations, through its advanced security, high availability and disaster recovery, and autonomic systems management and configuration features. Because of these features, businesses not only experience industry-leading low levels of planned and unplanned downtime, but increased productivity as well.

Organizations using SQL Server with Windows Server experience up to 4.5 times higher planned downtime and 6.6 times higher unplanned downtime than organizations using Db2 11.1. Three-year planned and unplanned downtime costs averaged 78 percent and 82 percent less for use of Db2 11.1 compared to SQL Server respectively.

Downtime costs for different industries are affected by a variety of factors, but all tend to result in loss of revenue through customer attrition or decreased productivity. As IT systems become more complex and interconnected, and reliability requirements constantly increase, organizations considering new transaction processing database solutions should look for highly secure and reliable platforms offering ease of management as well as integration into existing and future infrastructure.

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/

Identity Theft Resource Center (ITRC) found at idtheftcenter.org/images/breach/Overview2005to2016Finalv2.pdf

Quark + Lepton®

CORPORATE OFFICE
Boulder, Colorado USA

www.quarkandlepton.com
info@quarkandlepton.com