High Performance Parallel Database Solutions

White Paper

Data

Management

Solutions

IBM®

First Edition (January, 1996)

Abstract

This document describes the high performance parallel database market place and how IBM participates in this emerging technology as the way to implement high performance oriented data processing applications. This paper will also address current trends and directions and how IBM is addressing customer high performance requirements with parallel database and parallel hardware strategies, operating systems, other IBM solution technologies and products to provide customer solutions.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to state or imply that only that IBM product, program or service may be used. Any functionally equivalent product, program or service that does not infringe any of the intellectual property rights of IBM may be used instead of the IBM product, program or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Corporation
IBM Director of Licensing
208 Harbor Drive
Stamford, Connecticut 06904
U.S.A.

Trademarks

The following terms, denoted by an asterisk (*) in this publication, are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both:

AIX	IMS
CICS	IMS/ESA
CICS/ESA	Information Warehouse
DataGuide	MQSeries
DataHub	MVS/ESA
DataPropagator	OS/2
DataRefresher	POWERparallel
DB2	RISC System/6000
DB2/6000	Scalable POWERparallel Systems
ESCON	SP
ES/9000	Sysplex Timer
HACMP for AIX	System/390
IBM	S/390

The following terms, denoted by a double asterisk (**) in this publication, are trademarks or registered trademarks of other companies:

Encina is a trademark of Transarc Corporation
HP-UX is a trademark of Hewlett-Packard Corporation
Intel is a registered trademark of Intel Corporation
Solaris is a trademark of Sun MicroSystems, Inc.
TPC-A, TPC-C and TPC-D are trademarks of the Transaction Processing Performance Council
Tuxedo is a trademark of Information Management Company
UNIX is a registered trademark of UNIX System Laboratories, Inc.
Windows is a trademark of Microsoft Corporation
Windows NT is a trademark of Microsoft Corporation

Contents

Introduction

This paper describes how IBM* is delivering very high performance to common customer application requirements that demand the fastest response times, highest systems throughput, large transaction volumes and big data stores. This high performance is delivered via the synergy of IBM database management products - **IMS* DB, DB2* for MVS/ESA* and DB2 Parallel Edition for AIX***, with IBM parallel execution platforms - **System/390* and AIX**.

This paper will also look at the emerging market place for parallel solutions as customers look for advanced technologies to handle high performance applications in the ′90s and beyond. This is especially important as the corporate world is experiencing data overload, yet thirst for timely business information. In today′s global and highly competitive economy, most enterprises acknowledge the value of historical data. Inhibitors to implementing high performance solutions have been the cost of storing massive amounts of data, response times to access and analyze data and tools and procedures to manage data.

Executive Summary

In past years, large and complex application requirements have challenged the resource capacities of large computer systems even when high-end systems were configured to address huge transactional or query applications. It has become increasingly difficult to justify these large applications, given the very expensive ratio of cost versus the business value these applications deliver. This is for processing many concurrent transactions and queries with rapid response times. Add to this, massive amounts of data to be processed by transactions or queries and the need for data to be shared by concurrent users.

You may have a demand for a high performance solution for high speed execution and large resource capacity for such application types as Online Transaction Processing (OLTP), Online Analytical Processing (OLAP), decision support or client/server with huge number of users doing simultaneous work. You may also have requirements for continuous availability applications that are vital to your enterprise. For example, applications that need to be available 7 days a week, 24 hours a day to your business users or customers. This is called ″**high availability**″ or ″7**x24**″.

The addition of parallel processing to existing database technology offers the very highest performance solution available in the computer industry. This solution meets the following requirements:

- High transaction or query throughput with rapid response times
- Uninterrupted operations with outstanding system reliability
- A way to grow incrementally
- Ability to process very large databases
- Analysis of massive amounts of information

All these advantages come with better price/performance than previous technologies, plus you can choose the platform best suited to meet your requirements and complement your existing environment. An added value of the IBM high performance solution and its product set is that it fits in and coexists with other solution technologies you may have implemented or are evaluating, such as a data warehouse, data mining, client/server, data replication and systems management. These solutions leverage development and operational environments, management and personnel cultures. An IBM high performance solution will not disrupt existing production applications nor current data processing operations.

IBM High Performance Database Heritage

IBM has been an innovator in high performance database technology and is the computer industry leader in providing database products with ″industrial strength″ features to deliver high performance power and facilities for ″mission-critical″ applications. IBM database management products providing these capabilities are IMS DB, DB2 for MVS/ESA and DB2 Parallel Edition for AIX.

- IBM introduced Information Manager System Database Manager (IMS DB) in 1968 and it is based on a hierarchical data model. It was the first commercially available database management system and became the standard for high performance, data and system integrity, reliability and recoverability. The companion front-end product is Information Management System Transaction Manager (IMS TM) and it handles teleprocessing and communications interfaces to IMS DB and DB2 for MVS/ESA. A facility called **Fast Path** provides superior performance and availability. IMS DB, Fast Path and IMS TM are used to develop OLTP applications that run most firm′s core business applications on mainframe processors.

  Most corporate data is managed by IMS DB and it is the most successful database manager with more installations than any other MVS database manager. About ninety percent of Fortune 1000 companies use IMS DB to run their business. We estimate that IMS DB manages over 12 million gigabytes of production data with about 8 billion transactions per day servicing close to 15 million end-users. If IMS were a software company, it would rank as the 9th largest revenue producing software company.

  TPC-A** performance measurements were done in December, 1993 for IMS DB Version 4 and IMS TM Version 4 running on an ES/9000* 9121-742. Resultant throughput was 1,427.07 transactions per second (tpsA) with a price/performance of $13,586 per tpsA. An additional internal non-TPC measurement was done in a lab performance study. It was a banking debit/credit application with Fast Path databases on an ES/9000 9021-942. Resultant throughput was 2,478 transactions per second.

  IMS DB remains a strategic IBM database management product and is part of other IBM solutions: client/server; open database; object technology; data replication and systems management. IMS Client Server for Windows**, IMS Client Server Object Manager for OS/2* (IMS CSobject), VisualGen and VisualAge can be used to develop IMS/ESA* applications. IMS TM can be used with MVS Open Edition, Application Server/IMS and MQSeries* for MVS/ESA for other application development. An optional IMS/ESA feature is **Remote Site Recovery (RSR)**. It extends IMS/ESA recovery to remote locations to eliminate single points of failure that can disrupt end-user service. This provides improved continuous availability. For data replication, the **DataRefresher*** and **DataPropagator* NonRelational** products can be used to subset, unload and replicate IMS DB data to and from other data types, such as relational.

- IBM relational technology research began in the late ′70s and after prototype systems showed promise, the first IBM relational database product was available in the early '80s. This also introduced the **Structured Query Language (SQL)** to do data manipulation on relational tables. In short time, IBM delivered DB2 products on other operating systems and recently IBM has introduced DB2 products on non-IBM platforms, e.g., HP-UX**, Solaris** and Windows NT** operating systems. The DB2 family has seen tremendous growth and acceptance as relational technology has become the standard for application development, providing user productivity and high performance execution. DB2 family products span diverse client and server platforms from workstations and LANs to large mainframe and parallel processors.

DB2 for MVS/ESA has been available for more than 11 years. There are over 7,000 licenses installed worldwide and it is by far the most popular MVS relational database, installed on over fifty percent of all MVS installations. Nine of ten MVS customers with a relational database management system choose DB2 for MVS/ESA. The amount of production data managed by this product is over 2 million gigabytes. There are many business partner and third party vendor applications and tools. DB2 for MVS/ESA is in production in many shops doing transactional, decision support and batch work. Performance has greatly improved from initial releases. Benchmarks (using the DB2 Data Systems Workload) went from about 30 transactions per second in Version 1 to over 950 transactions per second in Version 3. The bottom line - DB2 for MVS/ESA has been accepted as the foundation for MVS relational database application development and has proven itself viable to execute diverse production applications with varying performance attributes. DB2 for MVS/ESA is highly integrated into other IBM solutions: client/server and distributed database; decision support; data warehouse; data replication and systems management.

- As UNIX** operating systems matured and became commercially viable to support traditional relational database applications, opportunities arose for this platform to deliver good price/performance for development and execution of customer database workloads. IBM joined the UNIX industry to deliver a relational database product leveraged to the IBM AIX operating system. The first offering on this platform was **DB2/6000* Version 1**. Many new facilities and functions were added, based on new technology and prototyping from IBM research (the Starburst project), and the second version, **DB2 for AIX Version 2**, was made available. This product and DB2 for OS/2 Version 2 were the initial members of the **DB2 Version 2 Common Server**. This is a common code base for DB2 Version 2 family products that run on Intel** (workstation) and RISC (AIX and UNIX) platforms, and that provide equivalent function while using operating system services and optimization to fully leverage each platform's environment. DB2/6000 and DB2 for AIX have a high performance emphasis toward decision support and client/server processing. DB2 for AIX has additional Symmetric Multiprocessor (SMP) enhancements.

  TPC-C** measurements were made for DB2 for AIX Version 2 on several RISC System/6000* configurations. The results (as of June 19, 1995):

  - 3119.16 tpmC at $349/tpmC on a RISC System/6000 Model J30 (8-way)
  - 3119.16 tpmC at $355/tpmC on a RISC System/6000 Model R30 (8-way)
  - 1470.06 tpmC at $666/tpmC on a RISC System/6000 Model R24 (POWER2 Uni)
  - 1438.04 tpmC at $492/tpmC on a RISC System/6000 Model J30 (4-way)
  - 1438.04 tpmC at $510/tpmC on a RISC System/6000 Model R30 (4-way)

The IBM high performance massively parallel processor (MPP) relational database offering on AIX is DB2 Parallel Edition for AIX (DB2 PE). Its design point is scalability for complex queries, including OLAP, against large databases. DB2 PE runs on IBM′s **Scalable POWERparallel Systems* (SP*)** hardware. DB2 PE is functionally equivalent to DB2/6000 and adds parallel technology from IBM Research to deliver even higher performance to process much larger databases. Applications automatically take advantage of the parallel processing enhancements provided by DB2 PE, because parallelism is implemented in DB2 PE, not in the user application. With the addition of DB2 PE, IBM now offers two high performance databases on the AIX platform, DB2 for AIX Version 2 (SMP) and DB2 Parallel Edition for AIX (MPP), that can be used to implement customer applications.

IBM Database Family Advantage

The key to delivering high performance execution is a set of robust database managers, rich in performance oriented facilities, resources and options that use IBM hardware, system and subsystem software. Our relational database managers have facilities that include: very intelligent and cost-based optimizers; multiple data indexing and locking choices; deadlock detection; data partitioning; integrated system catalog; flexible configuration options; hardware implementation of sort and data compression; governor function to prevent ″run-away″ queries and performance tracing and monitoring facilities.

Our hierarchical database manager - IMS DB, along with Fast Path, are the products of choice for OLTP operational applications to support high availability, large capacities and very rapid response times. IMS DB allows complex data relationships and predefined structures to be handled with efficient I/O operations, un-paralleled data integrity and recovery, functions to insulate end-users from planned and unplanned system outages and the synergy with IMS TM for definition and management of terminal and network interfaces.

Additional IBM products span the S/390 and AIX platforms for hierarchical and relational database managers. They are part of the other IBM solutions of data replication and systems management and aid data base and systems administrators, developers, operations and end-users. Products in these areas are:

- **DataHub*** (centralized control point for database administration)
- **DataGuide*** (business information catalog/directory)
- **DataJoiner** (transparent SQL access to distributed relational and non-relational data sources and locations)
- **DataPropagator Relational** (data replication within the DB2 family from source location to target location)
- **DataPropagator NonRelational** (replication of IMS DB data to and from DB2 for MVS/ESA data)
- **DataRefresher** (batch data movement from one data source/location to another data target/location)

Business Value of High Performance

High Performance Dilemma

We have traditionally thought of large mainframe processors as the only way to deliver high performance execution characteristics required by many of today's business applications.

As a definition - **high performance is very high transaction or query processing rates with rapid throughput and excellent response times**.

The development process of prior high performance applications was usually time consuming because it required extensive user requirements gathering, precise project analysis, detailed application and database design and traditional coding by skilled programmers with a fixed set of development tools. Although this was appropriate for operational mission-critical applications, you were unable to apply high performance techniques, tools and skills to non-operational, non-mainframe and/or un-traditional applications such as decision support. It may not have been feasible to implement new and emerging technologies, ones that could help you make better business decisions. They were too costly to create, hard to justify and non-mainframe platforms could not provide the required high performance throughput, quick responsiveness and system reliability.

Expanding High Performance Opportunities

IBM high performance solutions of IBM database management systems running on IBM parallel processors gives you the opportunity to have all the advantages of a high performance solution system applied to a wide array of new application opportunities. This is in addition to its use to address the operational world of OLTP applications working with hierarchical and relational databases.

Why is this important to you? How can the IBM high performance parallel database solution benefit your enterprise? Let's answer these questions by examining how parallel processing works to help deliver a wider array of valuable business applications. With new hardware technology on multiprocessors coupled with software advances, you can now have parallel solutions that allow individual tasks to be run simultaneously. One way is for large tasks to be divided into smaller, independent tasks that can execute concurrently. By doing this, it is possible to complete tasks (like a large query) fast enough to meet end-user's response time needs. A second way is to route "units of work" (like a transaction) to one of several available execution resources (like a processor). Transactions can then operate in parallel, each doing autonomous work. This high degree of parallelism is integrated into the IBM high performance solutions, in IBM database management systems (hierarchical and relational) and on IBM parallel hardware platforms (S/390 and AIX).

Justification for High Performance

For many application types, parallel processing can handle transactions or queries in a shorter time period, leading to increased earnings, like analyzing retail sales patterns. When a company can achieve this enhanced price versus performance throughput ratio, there is increased revenue potential. On today′s parallel platforms this is accomplished by using lower costs RISC-based (reduced instruction set computing) and CMOS-based (complementary metal oxide semiconductor) processors.

By reducing the time it takes to do a query, applications not previously feasible can be performed in a cost effective manner. These queries can be of great value to an enterprise to provide in-depth knowledge about their data, i.e., valuable information on which to make business decisions. Queries can execute to retrieve important information to gain a competitive advantage in your industry, resulting in substantial increases in revenue and cost reductions. A business can now access and profitably use large amounts of historical data. This was not practical before because the required hardware was too expensive or the response times might have been poor. With the ability to access valuable historical corporate information, a business can get a competitive edge and possibly improve their bottom line.

Another advantage of a high performance parallel database solution is the ability to expand a system incrementally. This means you can have a startup system with smaller, lower-cost configurations to match initial size requirements. Then, as applications grow (database size increases, additional users, more transactions to process, etc.), you can easily add appropriately sized resources to accommodate growth. This is as opposed to having to ″buy a bigger processor″ as in the past. This incremental growth capability is non-disruptive to your user community and operations environment.

High availability is addressed by this solution. Because there can be several processors, a system can be set up so that if a processor becomes unavailable, the system can be reconfigured for the remaining processors to handle work introduced to the system. This can occur without affecting existing work. There is no system outage when individual processors fail or have to come down for a planned outage, e.g., to do scheduled maintenance. The system continues to process transactions and queries with no end-user down time.

New Application Enablement

Let′s explore exciting new data processing opportunities a high performance parallel database solution makes possible and cost justified. Certain applications may not have been economically feasible on serial (non-parallel) processors. Now you can efficiently and quickly process vast amounts of data for tactical (operational or OLTP) and strategic (informational or decision support) business purposes.

- *OLTP:* Transactional systems to handle daily operational needs continue to grow. On the OLTP high-end, large systems have approached or exceeded the limits of large bipolar processors and you can no longer buy a bigger processor to handle these growing workloads. This problem needs to be addressed by new technology that maintains mainframe reliability to which you are accustomed, with flexible, non-disruptive growth options. Also required is the need to use existing database and teleprocessing products (like IMS DB, DB2 for MVS/ESA, IMS TM and CICS/ESA*) to be able to run current OLTP workloads within this framework without having to make changes to applications or databases.

An IBM high performance parallel database solution provides the framework to enable execution of larger OLTP workloads. It maintains a single systems image, has the reliability of previous systems, with the added benefits of data sharing, workload balancing and insulation from system outages. An advantage of our OLTP solution is its synergy with other solutions such as client/server, open systems, data replication and systems management.

- *Data Warehouse:* Also called an ″**Information Warehouse***″, is a solution for decision support. It provides an infrastructure for transforming detailed corporate data into more useful business information to support the strategic decision making processes. By being able to provide for and report on this valuable corporate information, you can make more reliable decisions that can lead to significant business advantages. An example is being able to react more quickly to changing market dynamics. A departmental or workgroup sized data warehouse is called a ″**datamart**″. The major components of a typical data warehouse or datamart implementation are:

  - ″Data Store″ - housing of informational data on a relational database. The source of this data is typically from transactional or ″legacy″ systems and data can be reformatted, e.g., summarized, for query usage
  - Tools and techniques to populate and update or refresh the informational database from external data sources at desired time intervals
  - Report and query tools used by business users to access and analyze the informational data in the data warehouse or datamart
  - Administration tools to do process and workflow management, central database management and user information catalog maintenance
  - Data cleansing tools to insure the quality of the information

  The challenges a data warehouse presents to traditional technology are large databases and adequate response times for queries that can be complex and require searching an entire database. S/390 and AIX are ideal platforms to base a data warehouse implementation. Both have parallel hardware running relational database software, capable of handling large data volumes with reduced elapsed times.

- *Data Mining:* This exciting new solution is enabled by high performance decision support applications. Our fulfillment of this capability is called **IBM Data Mining Technology**. It will help identify previously hidden associations of data and sequential behavior patterns embedded in large relational databases. Such implicit and useful patterns can represent a competitive advantage to your business. IBM will provide unique industry-oriented algorithms to help identify customer behavior. IBM Data Mining Technology is planned for S/390 and AIX to access DB2 databases in both environments.

High Performance Database Products and Directions

Each of the three IBM database management systems (IMS DB, DB2 for MVS/ESA and DB2 Parallel Edition for AIX) can be used to develop high performance applications. This section will describe the functions and facilities of these products that will deliver high performance. In the hardware platform section, we will examine how the databases work on IBM parallel hardware to deliver even higher performance for transactions and queries.

IMS/ESA

IMS DB is the original high performance database manager. IMS TM continues to set the standard for mainframe OLTP. Together, they form the basis to deliver rapid operational application systems. The Fast Path option provides the highest transactions per second volumes for database applications. IMS TM can also be the front-end to DB2 for MVS/ESA to provide excellent response time transactions for MVS relational applications.

To achieve greater transaction volumes than IMS DB or Fast Path on bipolar processors can deliver, IBM has paired IMS/ESA with a parallel hardware medium called the **S/390 Parallel Sysplex**. The synergy of software and hardware in this environment forms an architecture to easily handle many concurrent IMS/ESA transactions with accuracy to accommodate large data volumes, complex data structures and multiple update requests. A benefit of this parallel architecture is its ability to execute existing IMS/ESA applications. No redesign or reformatting of data, database unloading/reloading nor new application development is required. For example, your current COBOL programs with DL/I calls will still work, no modification is required.

IMS DB was the first database manager to exploit the processing power of the S/390 Parallel Sysplex for read and update processing and introduced **IMS DB data sharing**. This is the ability of up to 32 IMS/ESA subsystems to perform transaction work and concurrently access and share IMS DB data. IMS/ESA transactions (units of work) can execute on the processor which the submitting terminal is attached to or another IMS/ESA subsystem which the transaction is forwarded to on another processor within the S/390 Parallel Sysplex. This ″N-Way″ data sharing support gives increased capacity, higher availability and a way to grow incrementally. Specialized hardware, called **coupling technology**, along with the enhanced IMS Resource Lock Manager (IRLM), facilitates data sharing for update among multiple processors (handling locks, buffer invalidation, etc.). This is done with minimum impact to system performance.

IMS/ESA Directions

IMS/ESA continues to address customer demands as new versions are built and introduced to the field:

- **Investment protection.** Applications currently in production should continue to run, unaltered, as new IMS/ESA versions are installed
- **New application development and end user facilities.** Provide rapid application development tools with additional end-user functions
- **Maintain ″industrial strength″ facilities.** High performance, availability and systems management functions continue to be improved

- **Synergy with and exploitation of new and emerging technologies.** IMS/ESA being enabled to distributed, open systems and object solutions while keeping the IMS/ESA system aspects in these new ventures

The performance of IMS DB, Fast Path and IMS TM continues to improve to provide for higher transaction rates, more throughput and larger capacities in a S/390 Parallel Sysplex environment. A future direction in this area would be to enhance dynamic workload balancing. This could be accomplished with better queue management (also called ″shared message queues″) and would help insure that IMS/ESA transactions submitted from any processor in the S/390 Parallel Sysplex could be more evenly spread across parallel resources to minimize performance ″hot spots″ and bottlenecks in system execution.

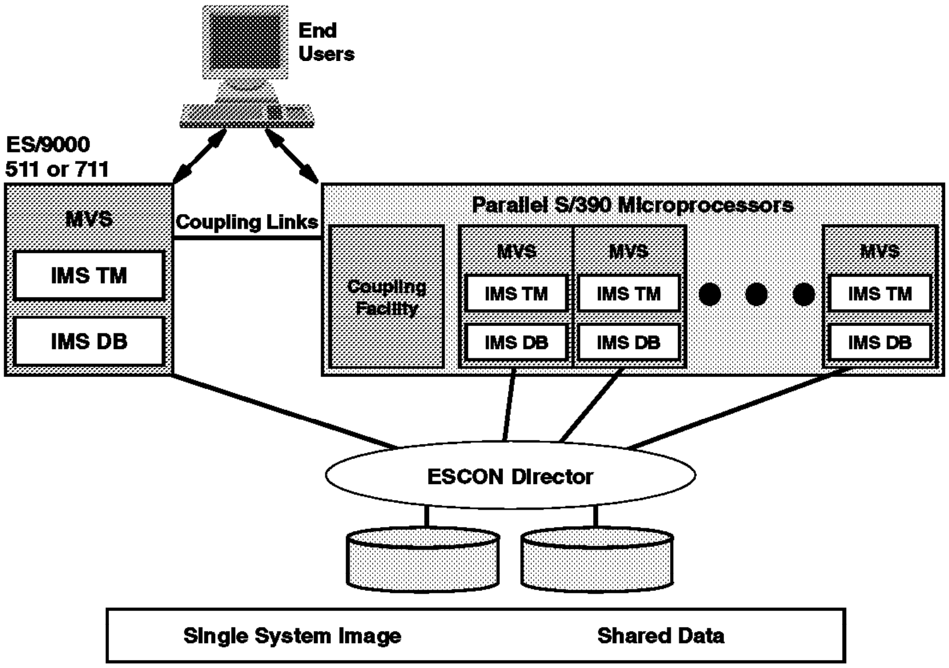

Figure 1. IMS/ESA Within a S/390 Parallel Sysplex Environment

DB2 for MVS/ESA

DB2 on the MVS platform was introduced in 1984. Since that initial release, there has been immense growth in customer usage and acceptance. Performance has tremendously improved, making this product viable to handle transaction workloads and query applications. DB2 for MVS/ESA has an outstanding reputation as a high performance relational database, due in part to it being leveraged for the **S/390 Parallel Enterprise Server** (a CMOS processor). This is either stand-alone or within a shared server environment - the S/390 Parallel Sysplex. The parallel database capability is called **DB2 data sharing**. Now DB2 for MVS/ESA is in line with the definition of high performance - very high transaction or query processing rates with rapid throughput and excellent response times. Added benefits of DB2 data sharing are high availability, reliability, scalable growth options and the capability to process large volumes of shared data. The DB2 for MVS/ESA high performance solution is important as customers look to activate new application technologies with characteristics of faster CPU execution against massive amounts of data - the most common being the implementation of data warehousing, data mining, decision support and client/server applications.

You can now have high performance relational database applications, delivered by DB2 data sharing and it has four key benefits:

- **Higher throughput with better price/performance.** Information systems can now process more work with higher performance at a much lower overall cost. This is primarily achieved by using lower-cost CMOS processors.
- **Reduced elapsed time and application expense with parallel processing.** DB2 for MVS/ESA′s parallel solution makes it feasible, practical and efficient to develop and execute query intensive applications such as data warehousing and data mining.
- **High system availability.** This parallel implementation allows for multiple DB2 for MVS/ESA subsystems (up to 32) to access shared data. This is key to 7x24, mission-critical applications because it insulates systems operations from scheduled and unscheduled outages. This continuous availability capability does not require changes to existing relational applications.
- **Flexible scalability and capacity.** You can install the appropriate amount of parallel resources (software and hardware) to match response time and capacity needs. As numbers of users grow, workloads become larger or data volumes increase, you can non-disruptively introduce additional parallel resources to match future size and capacity requirements.

A DB2 data sharing system can support up to 25,000 defined users per subsystem with 2,000 active at one time. With up to 32 DB2 for MVS/ESA subsystems in a S/390 Parallel Sysplex, the number of supported users goes up to 800,000 defined and 64,000 active.

DB2 for MVS/ESA has high performance features within the product. These are in addition to the parallel functions provided by DB2 data sharing.

- **Parallel I/O.** DB2 for MVS/ESA can perform parallel I/O operations on table partitions. This will greatly improve performance of I/O intensive queries because multiple data partitions are scanned in parallel (not serially) to reduce transaction elapsed time.
- **Parallel query processing.** A single query can be broken up into multiple tasks, each one operating on multiple processors in parallel. This facility improves overall performance of process intensive queries.
- **Parallel utility operations.** DB2 for MVS/ESA utilities, such as LOAD and REORG, can process table partitions concurrently instead of sequentially to speed up total elapsed time.
- **New index manager.** The index manager was enhanced to provide row level locking. This is in addition to page, table and tablespace locking. Now the appropriate lock type can be selected to increase data concurrency to minimize lock outs and data contention.
- **Improved client/server features.** DB2 for MVS/ESA is better positioned as a ″super server″ in an open client/server world. This was done with the addition of stored procedures to reduce network traffic and improve elapsed times of client applications.

The product - **DB2 Performance Monitor for MVS** has kept pace with new DB2 for MVS/ESA versions by immediately supporting its new functions. This product can help you understand the value of high performance functions, e.g., for parallel processing, it can analyze all parallel tasks of a split query and provide relevant performance values to the originating task. This will give a clear picture of how much the performance has improved using parallelism.

DB2 for MVS/ESA Directions

Expect future functions to address customer requirements in the areas of higher performance, client/server implementations, larger table sizes and capacities, increased availability and programmer/end-user productivity. DB2 for MVS/ESA seeks to improve with each new release in the areas of: better optimizer technology; more conformance with the DB2 Version 2 Common Server; more compliance with SQL and other standards; additional synergy with IBM hardware and operating system software; improved and additional distributed database functions and faster utilities.

For DB2 family futures, IBM is looking to port the DB2 engine to additional non-IBM platforms, e.g., other UNIX systems. Recent activity in this area has produced DB2 for HP-UX, DB2 for Solaris and DB2 for Windows NT. Another direction would be to expand client access opportunities to DB2 family products. This could be done by additional native interfaces to other network protocols.

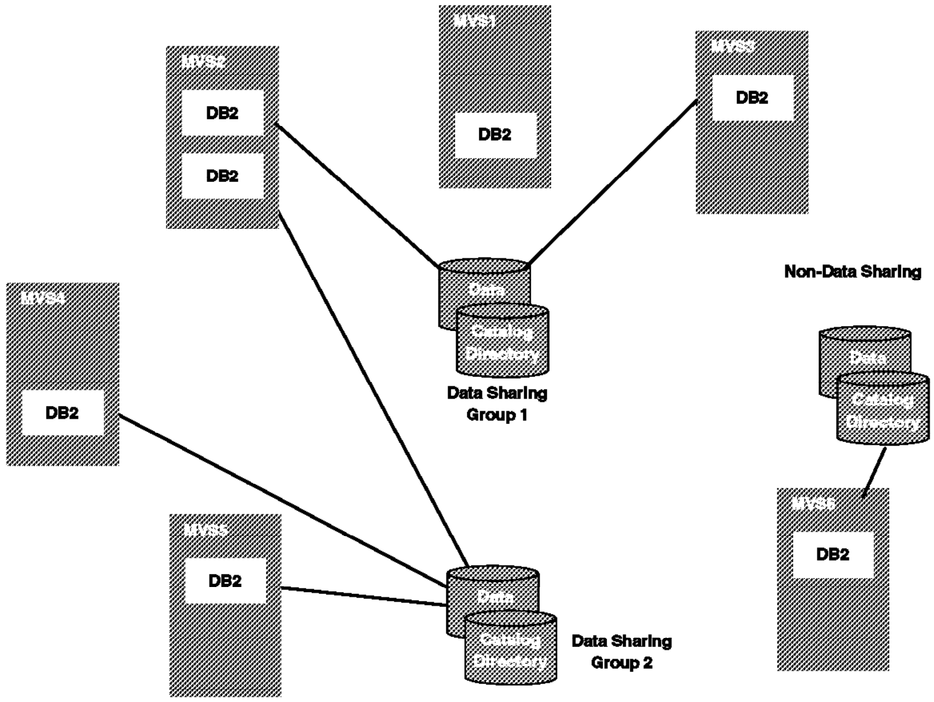

Figure 2. DB2 Data Sharing

DB2 Parallel Edition for AIX

DB2 PE is the newest IBM parallel database offering and its design point is for users who want an AIX high performance parallel database solution to address requirements beyond the capabilities of DB2 for AIX. DB2 PE is appropriate for informational (decision support) and smaller operational (OLTP) applications. The former allows for implementation of high performance data warehouse and data mining applications. DB2 PE executes on the IBM Scalable POWERparallel Systems (SP) and provides a flexible, integrated growth path capable of running on up to 128 processors (nodes) - all in parallel. Because DB2 PE and SP have the same architecture, this combined solution can offer near linear scalability. This means very little additional overhead when new processors are added to the total solution. No processor resources are shared (disk, memory, logs, etc.), resulting in minimum system overhead because processors are not trying to access each other′s resources.

DB2 PE features that aid the implementation of high performance query and decision support applications include:

- All database functions are performed in parallel, e.g., data manipulation, data definition and utility operations.
- A parallel, cost-based optimizer providing optimal use of parallel resources that results in maximum query speed-up with minimum data administration. All access plans are created for parallel operation.
- Flexible configuration options that exploit the advantages of the SP hardware platform. This allows for incremental growth to increase capacity and accommodate large data volumes.
- Large table support. Maximum table size on 128 nodes is 8 terabytes.
- The ability to rebalance data across nodes when new nodes are added to the SP complex and analysis and correction of distribution of data by a rebalancing function.
- **IBM High Availability Cluster Multi-Processing for AIX (HACMP for AIX*)** support. This provides increased data and system availability and for concurrent database operations over the processors in an SP complex.

Because DB2 PE was designed with the SP architecture in mind, a major benefit is achieved for parallel processing. As the number of processors increases, and potentially each one trying to access the same data row, it becomes important to minimize data contention. This potential bottleneck is eliminated with the use of DB2 PE. It allows each data row to be accessed only by the processor to which it is attached. Other processors will efficiently send messages to a processor if access to data is needed. Because of this ″shared nothing″ architecture, DB2 PE requires no distributed lock manager, thus avoiding potential performance degradation (additional processing to perform lock, buffer and cache management for updating data from multiple processors).

Additional synergy of DB2 PE and the SP architecture provides efficient processing of parallel functions for communications across processing nodes. To minimize data traffic between nodes, DB2 PE uses a technique called ″function shipping″. The parallel optimizer splits an SQL statement into subtasks. These subtasks are then sent to processors to operate on data they own and return only the answer set to the originating processor. This is done for all SQL data manipulation operations of SELECT, INSERT, UPDATE and DELETE. The concept of function shipping has been proven effective for decision support applications. Earlier non-IBM parallel products used this technique. DB2 PE on the SP enhances this implementation by correcting problems of prior deliverables, e.g., better pipelining between processors and data balancing correction.

DB2 PE is a viable product to base AIX high performance applications. Running in an SP environment gives all the advantages of parallel processing to provide cost effective solutions. To summarize, DB2 PE benefits on the AIX platform:

- Support of very large databases
- Ideal platform for decision support, data warehouses and datamarts
- Improved response times for complex queries
- Incremental growth with linear scaling
- Improved database availability

DB2 Parallel Edition for AIX Directions

In the future, DB2 PE would look for ways to incorporate new technology and provide for implementation of the latest and emerging industry and parallel applications. Another direction would be to merge DB2 for AIX and DB2 PE technologies. As new features are added to DB2 for AIX and DB2 PE, you could expect those features to also be added to the other product. This could include new data base functions and additional synergy with AIX hardware and software. DB2 PE would like to be able to exploit other RISC platforms. This could include focus on parallel OLTP and could be ported to other UNIX platforms.

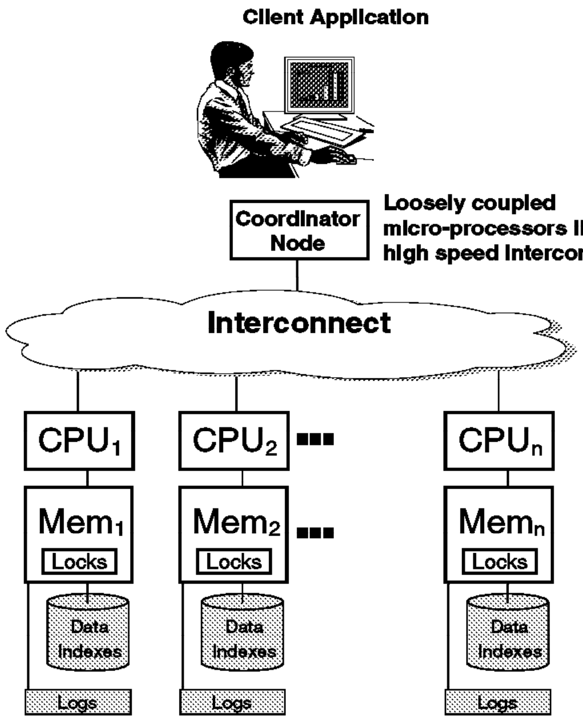

Figure 3. DB2 PE - Shared Nothing Architecture

High Performance Parallel Platforms

IBM has two high performance parallel operating opportunities:

- **System/390 platform:** IBM S/390 Parallel Sysplex with IBM S/390 Parallel Enterprise Server and IBM S/390 Parallel Query Server
- **AIX platform:** IBM Scalable POWERparallel Systems

Based on preference, specific applications or performance needs, you can select the platform to build your enterprise′s operational and informational systems for current and future high performance application requirements. Let′s examine the two high performance parallel platforms in more detail.

S/390 Parallel Sysplex with S/390 Parallel Enterprise Server

This solution evolved from the need for high performance access to transactional data, where detailed business information is processed with a high speed front-end teleprocessing system, such as IMS TM or CICS/ESA, with access to database management systems, such as IMS DB or DB2 for MVS/ESA. These OLTP applications have been high performance oriented in the past and are the basis of most company′s operational applications. However, customer demands are challenging even the largest S/390 mainframe processors and their capacities. OLTP workloads have the characteristics of higher transaction volumes, large databases, multiple concurrent users, massive update activity, big terminal networks - all with the need to maintain rapid response times. To address these emerging very large transactional requirements, IBM has a strategic S/390 platform solution - the S/390 Parallel Sysplex. This system can also be used to handle decision support and batch workloads. IBM services are available to help install and operate a S/390 Parallel Sysplex system.

This multiple processor environment is based on the MVS/ESA operating system and performs like a single system. It is not one product, but a combination of IBM processors - S/390 Parallel Enterprise Server (9672), 9121 (511-based) and 9021 (711-based) processors, along with IBM software, all cooperating to process work in parallel and without interruption. The S/390 Parallel Enterprise Server is a CMOS processor that can also execute as a stand-alone processor to perform general purpose workloads. The processors in the S/390 Parallel Sysplex are connected via coupling technology, like the **S/390 Coupling Facility (9674)**. A component called the **Sysplex Timer*** is the centralized time source and maintains time-of-day synchronization among processors.

The S/390 Parallel Sysplex was designed for continuous availability. Through a combination of coupling (hardware) and data sharing (software) technologies, a S/390 Parallel Sysplex can manage up to 32 IMS DB or DB2 for MVS/ESA subsystems. Data sharing is the ability of multiple subsystems to concurrently perform read and update operations on shared data with full integrity. Should one processor or subsystem experience an outage or be brought down for scheduled maintenance, a facility called the **MVS Workload Manager** will help shift incoming work to available resources. This product, in cooperation with IMS/ESA or CICS/ESA, also allows you to prioritize transactions to ensure the most important jobs are executed first.

With the themes of high performance and continuous availability, the S/390 Parallel Sysplex is the logical platform for critical OLTP workloads. You can design and configure a system made up of bipolar and/or CMOS processors. Transaction volume and complexity can be matched with appropriate parallel resources and capacities. This prevents bottlenecks and maintains high performance. IMS TM and CICS/ESA can be used to front-end IMS DB and DB2 for MVS/ESA applications. High-speed fiber links connect the coupling facility to each processor and the disk subsystem is accessible to all processors.

The concept of scalability takes on new meaning with this solution. It offers the largest single systems image, built in small or large increments depending on your needs. Before this solution, if you wanted to add processing power to boost performance execution, you had to install a larger processor. Now you can build a S/390 Parallel Sysplex out of lower cost CMOS processors or couple them with existing ES/9000 9121 air cooled or ES/9000 9021 water cooled processors. There is no impact to existing applications, that can run without program or data changes. This will protect application development investments in data processing resources and personnel skills.

The same features that make this solution attractive for OLTP make it appealing for decision support. Flexibility and scalability are key reasons to create decision support workloads on this platform. Flexibility to have a high performance, scalable, highly available solution to support data warehousing, data mining and OLTP requirements with a single, shared database. Or, the ability to have separate systems, one to handle operational tasks and another to process informational workloads. The choice is based on performance demands - either way, the hardware investment is one.

Batch execution remains a significant portion of many computer workloads. Most batch jobs can run on CMOS processors. Larger batch jobs can execute on ES/9000 9121 or ES/9000 9021 processors within a S/390 Parallel Sysplex. These are the same large bipolar engines that efficiently run today′s massive batch jobs. This is another example of application investment protection.

Systems management is significantly simplified by the single systems image. This is important to administrative and operational staffs because it makes daily operations easier and virtually the same as the operation of traditional S/390 systems. Change management can be performed on one processor and then replicated (or ″cloned″) across other processors. Other tasks, like adding capacity or doing maintenance, can be done to individual processors without affecting systems availability.

The target audience for a S/390 Parallel Sysplex has these attributes:

- Desire very high system and application availability
- Superior scalability solution with flexible growth options
- Need for lower cost solution
- High transaction volumes with rapid response times
- Database investment in IMS DB and/or DB2 for MVS/ESA
- Have existing IBM TP networks of IMS TM and/or CICS/ESA
- Wants single control point for operations and maintenance

Look for enhancements through improved CMOS price/performance delivery and subsystem functions. New functions would not impact current functions and could deliver even higher performance OLTP, decision support and batch solutions. Future directions that would provide enhancements are VSAM record level sharing and the ability to have IMS TM shared message queues on the coupling facility. This could allow for dynamic workload balancing of IMS TM transactions introduced to the S/390 Parallel Sysplex.

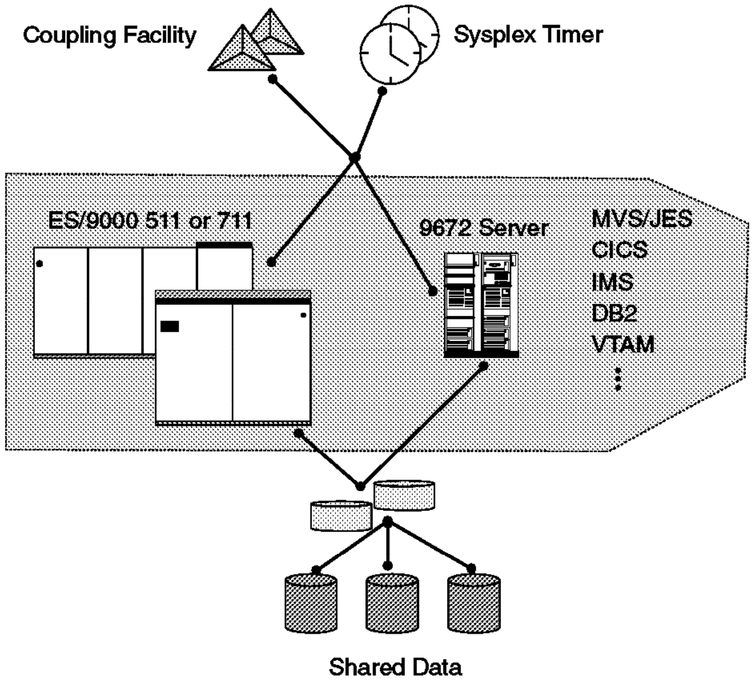

Figure 4. System/390 Parallel Sysplex

S/390 Parallel Query Server

For high performance decision support requirements, IBM has a specialized and dedicated solution - the S/390 Parallel Query Server (9673). This special purpose CMOS multiprocessor is the key component of a complete ″turn key″ solution that includes hardware, software and services. An SQL SELECT statement is issued from a ″front-end″ processor and routed to the attached ″back-end″ S/390 Parallel Query Server where the statement is split into subqueries, each executing in parallel on up to 48 processors. The subqueries operate on shared DB2 for MVS/ESA data and offer a high degree of parallelism. This reduces elapsed time for long running queries, making them more economically feasible.

This solution offers lower cost of computing, while assuring proven S/390 reliability and integrity along with CMOS technology. ESCON-based channels provide fast disk storage and external communication access. The front-end S/390 system provides physical network connections, security checking and database updates. Applications can run under CICS*, TSO or Call Attach Facility (CAF). This solution′s hardware and software are completely customized when delivered. Base services (e.g., database design, installation support and performance monitoring) minimize skills required for installation and operation. Optional services (e.g., training and database maintenance) can be tailored to specific needs. This solution complements client/server environments. Popular client tools running on DOS, OS/2 or UNIX can issue dynamic SELECT calls to obtain fast access to large amounts of business data.

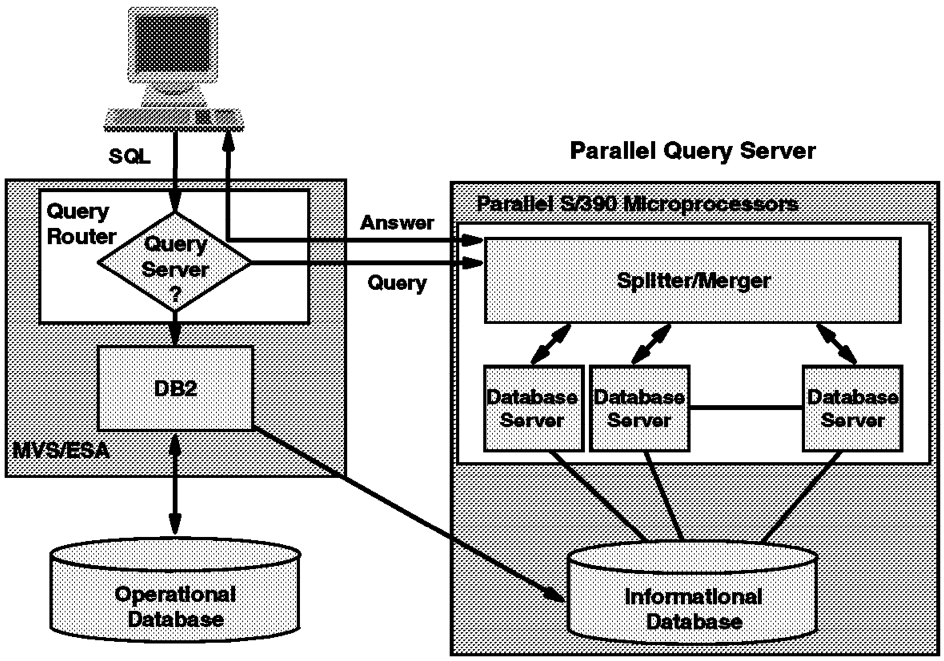

Figure 5. System/390 Parallel Query Server

Scalable POWERparallel Systems

Scalable parallel systems are being used to address AIX application areas requiring performance levels significantly higher than what predecessor symmetric multiprocessors could provide. These areas include traditional technical (engineering and scientific) and commercial (decision support and OLTP) computing applications and newer emerging areas like data mining, data warehouses, datamarts, digital libraries and video production and distribution. These applications have challenged the resources of other solutions and previously were not economically feasible.

The SP is a general purpose open scalable parallel system designed to address a wide range of application types. It is bigger and more powerful than prior systems and uses the latest POWER2 chip (77MHz). The SP supports ESCON channel attach with a throughput rate of almost 10 MBs per second. Memory and disk capacity per node has been increased. Total disk storage capacity ranges from 4 gigabytes to over a terabyte. The SP supports the **IBM 7133 Serial Storage Architecture (SSA) Disk Subsystem**. Total memory capacity ranges from 256 megabytes to 256 gigabytes. By doubling the maximum number of nodes and introducing a ″wide″ node and ″thin″ node mix capability, the SP is more scalable. SP customers can now select from a range of 7 nodes (2 wide and 5 thin) with different power, up to 128 nodes in a single system. The High Performance Switch has been improved and can provide increased bandwidth and decreased latency to allow for more transaction throughput with faster response times as this feature boosts the possible number of nodes to 512.

Each node has its own copy of the AIX operating system and other standard RISC System/6000 system software. New software products designed for the SP allows its parallel capabilities to be exploited (e.g., DB2 PE). OLTP support is provided by leading UNIX TP monitors: CICS for AIX*, Encina** for AIX and Tuxedo**. The SP enables one systems administrator to manage an entire system using a single control point (RISC/6000 workstation). The SP system includes systems management, job scheduling and systems monitoring software.

Like DB2 PE, SP′s design point is a ″shared nothing″ architecture, where each processor has its own private memory and disk storage. Message passing is the method the processors use to communicate with each other. SP supports HACMP for AIX. This additional subsystem software allows up to eight SP nodes to be configured in a highly available cluster to provide a continuous availability option for the SP system. Because of the substantial improvement of overall capacity and management of the SP system, it is more high performance oriented. Combined with DB2 PE, the SP is an excellent choice for running data-intensive workloads such as those found in data warehouses and data mining, in addition to traditional technical and commercial applications.

POWERquery Solution

POWERquery is an IBM packaged solution offering of highly scalable and flexible ″building blocks″. The components are hardware, software and IBM services specifically tailored to an installation′s needs. This solution addresses the implementation of large-scale, cost-effective decision support applications. By using POWERquery you can achieve competitive and strategic benefits like successful marketing campaigns, efficient resource allocation and a way to better analyze your customers′ buying patterns. Industries that can benefit from this solution are: retail; consumer goods manufacturing; banking; insurance; airlines and telecommunications.

A POWERquery solution includes:

- The two AIX parallel components - the Scalable POWERparallel Systems and DB2 Parallel Edition for AIX
- IBM′s most advanced storage offering for RISC-based systems - the IBM 7135 RAIDiant Array
- IBM services such as planning, implementation and training assistance

Optional POWERquery solution components are:

- POWERquery Data Replication Tool: Software tools to populate and maintain decision support databases from operational data sources
- POWERquery Database Tool: Performance monitoring, database administration and capacity planning tools to manage AIX operations

Parallel Performance Measurements

A common question asked about IBM high performance solutions is - "What are the performance numbers"? Typical requests are for number of transactions per second or minute, TPC benchmark results and cost per transaction. Also requested is performance data for specific environments. IBM constantly conducts performance studies and benchmarks in our product development labs. Test are made to insure new product versions give better performance, throughput and response times than previous versions and that nothing in the execution environment creates bottlenecks or performance degradation. For example, extensive tests were made and are continuing to insure that our high performance databases run efficiently on parallel hardware platforms.

Our performance tests may not reflect your exact situation. Thus, the answer to a request for performance numbers may be - "it depends". This means many factors need to be taken into consideration to measure an exact customer application workload in a specific execution environment. IBM's goal of doing performance studies and measurements are three-fold:

- To provide general guidelines for hardware sizes and software configurations required to adequately perform anticipated workloads
- To insure that when installing new IBM software product versions (operating system, database, etc.) or upgrading hardware (processor, network, DASD, etc.), the results will be better, not degraded, performance
- That various configuration combinations have been tested to insure there are no bottlenecks caused by individual components or product combinations

Much information exists on performance of IBM high performance solutions, but is too voluminous and detailed to be included in this paper. Because performance measurements are constantly being performed with various workload types on different hardware platforms with several operating system, subsystem and database product versions, any numbers quoted here would soon be obsolete or would not entirely apply to your conditions. Your IBM representative has access to performance information if you need it.

A common concern is the overhead of additional resources when scaling up the number of processors in a system complex. As was previously mentioned, such performance related issues are of concern to IBM and are factors as we implement new functions. For example, DB2 for MVS/ESA studies show that as new processors are added to a S/390 Parallel Sysplex (each additional increment from 2 to 32 CMOS processors), a 99% to 99.5% linear performance increase is achieved for DB2 data sharing applications. Only a 0.5% increase in total system overhead is seen. This is an extremely minimal performance impact, especially in light of the near perfect linear performance boost. CMOS scaling is lower in cost than scaling up with non-CMOS processors.

IBM is moving more toward industry standard performance measurements. We hope to continue to provide TPC-C numbers for DB2 Family products running in AIX and UNIX environments. Recently, a non-IBM hardware vendor (Sun Microsystems, Inc.) announced outstanding TPC-C benchmark numbers for a DB2 product (DB2 for Solaris) running on their operating system. IBM was first to publish audited results on the new TPC-D** industry standard benchmark (oriented to decision support and query workload measurements).

High Performance Platform Decision - Solution Differentiation

An advantage of an IBM high performance parallel database solution is the choice of IBM platforms for implementation. Selecting one platform may be obvious, an easy determination process or may be difficult. You may want help in choosing a platform for current and future requirements.

The following are general guidelines to help select a high performance platform. This process will focus on what is different and what is common between the S/390 and AIX solutions. A distinction will be made between OLTP and decision support. These are general guidelines and suggestions. In a real platform selection process, much care should be taken to fully study, examine and compare both solutions, their costs, impact to existing development and operational cultures and desired benefits based on solution justification. Consider the strengths of each platform and determine if they can be leveraged to a favorable position for you.

An AIX advantage is its heavy investment in the implementation of and ongoing participation in emerging open standards. A S/390 advantage is its heavy investment in systems reliability, data integrity and systems management. For overall availability, the two solutions match up fairly evenly with a slight edge to S/390. For transaction capacities, AIX can handle about 1,000 concurrent users doing hundreds of transactions per second, while S/390 can process about 10,000 concurrent users doing thousands of transactions per second. Each platform is making advancements in each other′s areas of strength, but clearly there will not be a full merger of functionally of these two cultures in the near future.

For decision support applications, the S/390 and AIX platforms are similar and can support equivalent workloads. Differentiation can be in a customer′s special demands, platform preference, existing people skills, system′s heritage and overall costs. These will be factors for helping you decide which platform to select. Before continuing, you should answer the following questions that influence a selection process methodology. These questions will determine your current processing situation and preferences. The two choices (answers) are S/390 and AIX:

- Where is major hardware investment for existing applications?
- Which platform is preferred by management, systems administrators, database administrators, development staffs, end-users?
- Where do most personnel technical skills exist?
- What database products are installed? If multiple, which one has most mainstream applications? Which one houses most corporate data?
- Do you currently have a data warehouse or datamart? If so, on what database(s) and on what hardware platform(s)?
- Which database product is preferred for production operational data? For decision support informational data?
- What is the anticipated database size of a proposed high performance decision support or OLTP application? How fast will the database grow?
- Does the proposed high performance application need to share data with existing applications on a specific platform? Will it be a single application or integrated with existing S/390 host or AIX system applications?

- What are system and application availability requirements?
- What are performance and response time expectations?
- On which platform are most end-users?
- Which platform is used the most by application developers?
- On which platform do most application development tools reside?
- Which platform is used by the majoring of application packages?

Considering your company′s future and trends in your type of business, look at the above questions again, this time asking where do you want to be in the next 5 to 10 years with a production S/390 or AIX system? The results could indicate a combination of both.

Here are several statements that are platform recommendations to implement certain application types with specific requirements. Remember, a S/390 or AIX high performance parallel database implementation could have one of two characteristics:

- A stand-alone solution - where a dedicated system will be installed to perform specific OLTP or decision support workloads
- An OLTP or decision support solution to coexist with other general application workloads

Because of the diverse requirements and different aspects of OLTP and decision support applications, we rarely see customers with one system performing both application types with high performance characteristics.

1. There is a clear delineator for OLTP implementation. Since IMS/ESA runs on MVS, an OLTP IMS/ESA high performance solution can only execute on a S/390 Parallel Sysplex. If you are an IMS DB shop (with a IMS TM and/or CICS/ESA front-end), your growth path is a S/390 Parallel Sysplex. DB2 for MVS/ESA could be added for a relational OLTP implementation. For non-IMS OLTP applications requiring very high availability, the better choice is S/390. For other non-critical OLTP relational applications with fewer number of users issuing lesser number of transactions, the better AIX performer is DB2 for AIX on an SMP. If the workload requires more than an SMP, you can design the application as a distributed application and use multiple DB2 for AIX databases, where each could be running on an SP node.

2. On the opposite side of the high-end S/390 spectrum, if you only have AIX (no MVS systems installed) and are looking to implement a data warehouse, datamart or decision support application, the SP is the logical choice. You also have a entry level OLTP option using DB2 for AIX or DB2 PE with a choice of AIX TP monitors.

3. For decision support, if your answers showed a clear bias toward one platform or if you only have one platform installed, then a new high performance implementation could be developed and executed on that platform. If you are an MVS shop with no AIX systems installed, it would make sense to install either a S/390 Parallel Query Server for read-only relational workloads or the S/390 Parallel Sysplex (or S/390 Parallel Enterprise Server) for more diverse demands. However, if AIX only or you envision AIX as the growth platform for new decision support applications or packages, then an AIX choice is logical.

As stated earlier, heritage plays a very important role in a selection process. One solution might appear less expensive or more appealing to decision makers. However, installing and operating a new platform requires additional skills, new management concerns and more resources (hardware, software and people). In other words - additional and hidden costs. Take this into consideration as you evaluate platforms. If your heritage is mostly S/390 or largely AIX, then a similar platform decision will not cause a major disruption to your data processing culture. You would be familiar with a platform′s application development environment, systems management issues, backup/recovery methods, operational procedures, configuration options, etc.

4. A clear distinction between the two platforms is system size combined with usage and performance expectations. The more transactions to be processed, the higher the throughput expectations, the larger number of end-users - all point to the S/390 platform.

5. AIX supports more end-user device types and has more native network support. AIX supports PCs, UNIX workstations, ASCII terminals, 3270 terminals, etc., while S/390 (non-gateways) natively supports mostly 3270 terminals. Network support for AIX includes TCP/IP, Ethernet, Token-Ring, FDDI, etc., while network support for S/390 is IMS TM and CICS/ESA.

6. If you have an existing decision support application under a specific operating system and database and you want to port it to a high performance platform, then the logical migration is to a similar platform. This would avoid application rewrites and migration rework. For example, a DB2 for AIX application would easily port to DB2 PE on an SP and a DB2 for MVS/ESA application would easily port to a S/390 Parallel Enterprise Server running stand-alone or in a S/390 Parallel Sysplex. However, if an organization has made a strategic decision to develop new high performance decision support applications with AIX, then the SP solution is viable and most of the MVS investment is protected because of similarities within DB2 family products.

7. If you want to migrate off a non-IBM parallel solution, the AIX platform is probably closely suited to this requirement. Most non-IBM parallel solutions are similarly architected to a ″shared nothing″ approach, where the nodes or other technology, distribute the functions (as opposed to data sharing). Thus, DB2 PE on a SP is an ideal solution for this type of conversion.

Now let′s get into the ″it depends″ considerations for platform selection.

A. Although S/390 seems capable of handling large system ″objects″, this may not be the case for database size considerations. S/390 can accommodate large amounts of total DASD, but the current DB2 for MVS/ESA table limit size is 64GBs. Many data warehouses are likely to be more than 200GBs with individual tables greater than 64GBs. The advantage for table size goes to AIX. Coupled with heavy insert activity for a large decision support application, DB2 PE is a good solution for large tables processing.

B. Considerations that can direct you toward one platform versus another is a combination of application development environment and/or end-user preferences. For example, if end-users are on many TCP/IP networks or your programmers are using AD tools unique to UNIX, then this might be the most important criteria. This assumes the two platforms match up fairly evenly in other criteria areas.

C. If you have both platforms and intend to keep them for application development, then a decision is more complicated. Factors to consider are your corporate direction toward a particular platform. Is there any dependency of a new application to existing applications on one platform or another? It can be tempting to view solution costs as the major decision making factor. There may be additional and hidden cost associated with a decision. Make sure you understand all platform selection issues of hardware, software and people skills, and what other costs there might be for future system growth. Subsequent application enhancements and increased usage may force you onto another platform and the conversion and migration costs are never budgeted for nor anticipated.

D. If a decision for a high performance solution is being driven by the growth of or anticipated use of an application package, then this may be the major decision criteria. Some vendor packages may only run on one platform, while others are available in S/390 and AIX versions. If this is the case, insure both package versions have the same functions and find out if the package is leveraged for each platform (i.e., can it take advantage of specific platform functions).

The S/390 Parallel Query Server is a niche solution with a very specific marketplace. To a lesser degree, DB2 PE is a specialized solution but with a much wider range of potential uses. They both excel at processing large queries and analysis against huge amounts of data, returning answers in seconds. OLTP is IMS/ESA′s niche. DB2 for MVS/ESA and DB2 for AIX are all-purpose database products that can support some OLTP and most decision support requirements. DB2 for MVS/ESA has more systems management functions, however, DB2 for AIX is more ″open″. An important thing to remember is the DB2 family and that over time you have the flexibility to change platforms should your business or data processing requirements change. You can do this without having to redesign applications or rewrite queries. The bottom line is investment protection.

Lastly, consider total costs of ownership. These are additional and usually unanticipated costs like staff salaries, education expenses, software fees, systems management utilities and environmental factors (floor space, air conditioning, electrical power demands, etc.). Similarly configured high performance solutions could be very close for hardware and software price and cost of operation. An equally important decision criteria should be - will the S/390 or AIX high performance parallel database solution meet your needs for now and the future.

Summary

IBM high performance parallel database solutions combine unprecedented price/performance value with complete scalability. Applications you now develop can grow, economically and non-disruptively, as your business grows. Whatever your mission or platform, IBM′s unsurpassed operating systems, database software and parallel hardware deliver the right mix of speed, throughput, high availability and capacity. You can leverage existing investments in systems, applications and skills, resulting in power over your corporate data you never had previously. This is the power to process and analyze information in ways not possible or cost-effective with traditional computer processing methods - to implement an OLTP system to process massive transaction volumes, to create a data warehouse or datamart for decision support and/or, to do data mining to find patterns of data and to spot business trends. The bottom line - to transform your corporate data into strategic information that will give your business a competitive edge.

Additional Information

Related White Papers

This white paper is one of a suite of Data Management papers available from IBM. The complete set consists of:

- *Client Server Solutions for IMS*
- *Data Management: IBM Strategies and Solutions*
- *Data Replication: Data on the Go for Businesses on the Move*
- *Database Systems Management: IBM's Solution*
- *DataJoiner - A Multidatabase Server*
- *DB2 Relational Extenders*
- *High Performance Parallel Database Solutions*
- *IBM Information Warehouse Solution: A Data Warehouse - PLUS!*
- *Information Technology and IMS in the 21st Century*
- *Multi-dimensional Analysis: Extending the Information Warehouse Framework*
- *Multimedia and the DB2 family*
- *Object Support and the DB2 family*
- *Open IMS*

Solutions Directory

For a list of complementary tools and applications available to work with IBM database products, please refer to the *DB2 Client/Server Solutions Directory*; this directory describes over 1300 applications and is available on CompuServe** in the IBMDB2 forum - please look for the DB2VEND file in the DB2/2 section.

Also available is the AS/400 Client/Server Application Directory which includes over 1500 DB2/400 client/server applications organized by industry/market segment. The IBM publication number is G325-6202

Technical Report

In addition the following is available as IBM Santa Teresa Laboratory Technical Report: *Object and Relational Databases: IBM's Strategy*.(03.457)

Ordering Information

To obtain free copies of the *DB2 Client/Server Solutions Directory*, the *AS/400 Client/Server Application Directory*, the IBM white papers listed above, or the technical report on *Object and Relational Databases: IBM's Strategy*, readers in the U.S. and Canada may call 1-800-IBM-3333 and ask for extension STAR 706.

To send a fax requesting these materials, please fax your name, mailing address, and the titles of specific materials that you want to 1-408-463-2750. Please address your fax to "Data Management Information".

IBM employees may obtain these papers online through the DB2INFO conference disk, as described below.

- To request soft copy of *Client Server Solutions for IMS*, issue this VM/CMS command:

  request imsspcs from db2info at torolab2

- To request soft copy of *Data Management: IBM Strategies and Solutions*, issue this VM/CMS command:

  request datasolu from db2info at torolab2

- To request soft copy of *Data Replication: Data on the Go for Businesses on the Move* issue this VM/CMS command:

  request drswp from db2info at torolab2

- To request soft copy of *Database Systems Management: IBM's Solution*, issue this VM/CMS command:

  request dsm_wp from db2info at torolab2

- To request soft copy of *DataJoiner - A Multidatabase Server*, issue this VM/CMS command:

  request dj1wp from db2info at torolab2

- To request soft copy of *DB2 Relational Extenders*, issue this VM/CMS command:

  request db2relex from db2info at torolab2

- To request soft copy of *High Performance Parallel Database Solutions*, issue this VM/CMS command:

  request hppardb from db2info at torolab2

- To request soft copy of *IBM Information Warehouse Solution: A Data Warehouse - PLUS!*, issue this VM/CMS command:

  request iwsolu from db2info at torolab2

- To request soft copy of *Information Technology and IMS in the 21st Century*, issue this VM/CMS command:

  request imsit from db2info at torolab2

- To request soft copy of *Multi-dimensional Analysis: Extending the Information Warehouse Framework*, issue this VM/CMS command:

  request mdaandiw from db2info at torolab2

- To request soft copy of *Multimedia and the DB2 Family*, issue this VM/CMS command:

  request db2mm from db2info at torolab2

- To request soft copy of *Object Support and the DB2 Family*, issue this VM/CMS command:

  request db2obj from db2info at torolab2

- To request soft copy of *Open IMS*, issue this VM/CMS command:

  request imsop from db2info at torolab2

IBM employees may also obtain these papers online through the MKTTOOLS conference disk, as described below.

- To request soft copy of *Client Server Solutions for IMS*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get imsspcs package

- To request soft copy of *Data Management: IBM Strategies and Solutions*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get datasolu *

- To request soft copy of *Data Replication: Data on the Go for Businesses on the Move*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get drswp *

- To request soft copy of *Database Systems Management: IBM's Solution*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get dsm_wp *

- To request soft copy of *DataJoiner - A Multidatabase Server*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get dj1wp *

- To request soft copy of *DB2 Client/Server Solutions Directory*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get db2vend package

- To request soft copy of *DB2 Relational Extenders*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get db2relex *

- To request soft copy of *High Performance Parallel Database Solutions*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get hpdbsolu *

- To request soft copy of *IBM Information Warehouse Solution: A Data Warehouse - PLUS!*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get iwsolu *

- To request soft copy of *Information Technology and IMS in the 21st Century*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get imsit package

- To request soft copy of *Multi-dimensional Analysis: Extending the Information Warehouse Framework*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get mdaandiw *

- To request soft copy of *Multimedia and the DB2 Family*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get db2mm *

- To request soft copy of *Object Support and the DB2 Family*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get db2obj *

- To request soft copy of *Object and Relational Databases: IBM's Strategy*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get oostrat *

- To request soft copy of *Open IMS*, issue this VM/CMS command:

  tools sendto usdist mkttools mkttools get imsop package

Printed in U.S.A.