The Role of IMS in Today's Enterprise

Colin White, BI Research
May 2014
Sponsored by IBM

TABLE OF CONTENTS

INTRODUCTION

Since its first release in 1968, IBM's Information Management System (IMS) has enjoyed nearly five decades as an industry-leading system for high-volume transaction processing. It has a strong reputation for its reliability and security and solid performance, and is still used today by many of the Fortune 1000 to run core business processes such as order processing, inventory control, supply chain management, shipping, invoicing, and so forth. Its high performance is demonstrated clearly by the latest release of the product, IMS 13, achieving a staggering 117,000 transactions per second on IBM's IMS performance benchmark.

Over the last forty or so years since IBM initially developed IMS in conjunction with Rockwell and Caterpillar for the Apollo space program, both transaction and data volumes have grown dramatically and show no signs of leveling out. Leading hardware and software vendors are therefore constantly improving and optimizing their products to support these volumes and help their customers maintain required service levels. This is why IBM continues to invest significant resources in enhancing IMS to increase its performance and keep ahead of customer requirements.

Many new enterprise technologies have been introduced since the introduction of IMS and another key area of IBM investment in IMS has been to enhance the product to coexist and interoperate with these new technologies. Examples of technologies here include relational database systems, cloud computing, mobile platforms and rapidly evolving big data technologies such as Hadoop. This coexistence is important because **no single technology today can support all customer needs**, and ease of integration across different technologies is becoming a key success factor for both vendors and customers.

The objectives of this paper are to look at the role of IMS in today's highly complex IT environment and to examine how organizations can integrate IMS into an open and flexible enterprise IT infrastructure that can evolve as technology changes.

ENTERPRISE PERFORMANCE: HISTORICAL PERSPECTIVE

Organizations have spent over fifty years automating their business processes and during that time there have been some key technology advances that have had a major positive impact on IT, which in turn has led to significant benefits to the business. Four key technologies of importance here are online transaction processing (OLTP) in the early 1960s, relational database management systems (DBMSs) in the early 1980s, business intelligence (BI) and data warehousing (DW) and in the early 1990s, and lastly big data, which is the latest technology innovation that promises significant benefits to organizations.

OLTP Systems

OLTP systems dramatically improved the ability of companies to automate business processes and react rapidly to changing operational business requirements. One of the first OLTP systems built was the Sabre airline reservation system, which was custom developed by American Airlines in conjunction with IBM. The initial Sabre system handled some 84,000 transactions per day, which at the time was considered to be a very *high* transaction volume, since it stretched the limits of what was possible using available hardware and software at that time. High availability and fault tolerance were also critical to the success of the system.

The concept of what is considered to be a *high* or *big* volume has changed over the years as the introduction of point-of-sale terminals, bank tellers machines, mobile devices, the Internet, and sensor networks has led to more automation of business processes and correspondingly higher transaction and data volumes. Sabre, for example, now processes some 60,000 transactions per second. Many IMS systems today handle similar transaction volumes.

Relational Database Management Systems

The introduction of relational DBMSs gave IT developers and business users for the first time a standard data access language (SQL) and logical data structures (relational tables and views) that were independent of the way the data was physically stored, accessed and managed. This technology not only improved usability, but also enhanced portability and interoperability. Today, relational DBMSs dominate the database marketplace for both transaction and business intelligence processing. Nevertheless, there are many companies today in business domains such as banking and telecommunications that need to support transaction volumes and service levels that relational technology cannot handle, and these companies therefore continue to use systems such as IMS for that purpose.

Data Warehousing and BI

As the number of data sources and the volume of transaction data grew so too did the problem of accessing this data for reporting and business decision making and planning. This issue led to the birth of the data warehouse where historical transaction data was consolidated into a single logical data store for use by BI reporting and analysis applications. This development radically improved the ability of users to analyze and optimize business operations and make more informed operational and strategic decisions.

Constantly growing transaction and data volumes has inevitably led to a corresponding increase in the size of data warehouses. As in the transaction processing environment, hardware and software have been improved and optimized to support these large data warehouses. Today, multi-terabyte data warehouses are no longer the exception as they once were, and some companies are now managing several petabytes or more of data in their data warehousing systems.

Big Data

The concept of big data evolved from solutions developed initially by web companies such as Google and Yahoo. These companies needed to manage and index huge volumes of web data, and existing technologies were unable to support this in a timely or cost effective manner. To solve this problem, these web companies developed their own non-relational data management technologies to provide the required performance.

Several of these systems have been donated to the open source community, which has led to the growing use of them in a wide range of industries. These systems enhance and extend the existing data warehousing environment by enabling many new types of data to be managed and analyzed, and offering improved price/performance for certain types of workloads.

The four disruptive technologies outlined above have led to a rich set of capabilities for automating business processes and deploying data warehousing and BI solutions. All four have needed to support new types of application processing and new types of data, while also providing good performance and open interfaces. No single technology can satisfy all needs, and organizations today need to employ different technologies depending on application needs. The challenge of course is *how* to determine which technology to use for any given project. As always in IT, satisfying performance requirements is a key selection criterion.

UNDERSTANDING APPLICATION PERFORMANCE NEEDS

Performance has many dimensions (See Figure 1). In terms of data management, the dimensions are the amount of data to be managed (data volume), its rate of generation or change (data velocity), the types of data to be managed (data variety), and the number of data sources, structures and relationships involved (data complexity). In the case of workloads, the dimensions are the types and complexity of the application processing (workload complexity), data currency and response time requirements (workload agility) and the makeup of the overall workload (workload mix).

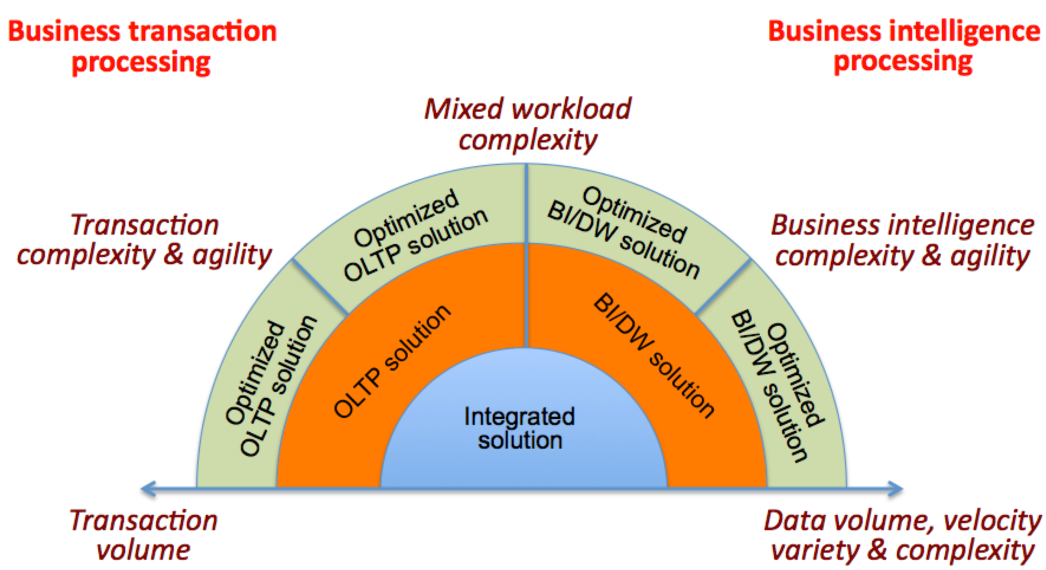

FIGURE 1. THE APPLICATION PERFORMANCE SPECTRUM

The product and technology needed to support any given project will be determined by the actual performance requirements in any of these data and workload dimensions. The higher the requirement in any dimension, the more important it is to choose a system that is optimized to provide good performance in that dimension. In Figure 1, the blue (integrated solution) and orange (OLTP and BI/DW solutions) parts of the application spectrum are dominated today by relational DBMS technology. As performance needs move into the green parts of the application spectrum, i.e., towards the extremes of the spectrum, then careful technology evaluation is required and the need for optimized solutions grows. For example, in the case of high-volume transaction processing, this need is satisfied by IMS, whereas for handling large amounts of multi-structured data this is where Hadoop is starting to gain ground.

THE IBM IMS PERFORMANCE BENCHMARK

The current release of IMS is Version 13 (IMS 13). As with every release, the IMS development group continues to focus on improving performance and scalability, usability, and availability. The performance improvements in IMS over the years can be seen when examining and comparing IBM IMS performance benchmark results.

The IBM IMS benchmark simulates a credit card processing environment. The application workload represents a high-volume transaction workload with four credit card transactions – credit card validation, credit limit check, debit/credit account processing and lost/stolen card reporting. The transactions provide an online update capability with full integrity and recovery facilities for both transaction and database processing. A Teleprocessing Network Simulator (TPNS) system simulates 8,000 workstations entering transactions into the IMS system.

In 1978, 10 years after IMS contributed to the successful Apollo moon shot, the IBM IMS benchmark achieved approximately 88 transactions per second (TPS). In 1987, at the IBM Santa Teresa Lab in San Jose, California, IMS accomplished a benchmark transaction rate of 1,012 TPS on a single IMS system. Some 27 years later, IMS 13 achieved a sustained average transaction rate of over 117,000 TPS using the most current IBM hardware. The IMS 13 benchmark results showed a significant performance improvement over the IMS 12 result of 46,000 TPS.

It is interesting to compare the IBM hardware used for the 1987 and the IMS 13 benchmarks. The 1987 benchmark was run on a mainframe that required, in total, at least 600 square feet of floor space. By contrast, the IMS 13 benchmark used a moderately sized IBM zEnterprise EC12 mainframe processor that used less than one-tenth the amount of floor space (the size of a standard cubicle), using a fraction of the power, and still on a single IMS system.[1] If required, the largest single zEnterprise EC12 processor could drive an additional 10 times as much work.

It is useful at this point in the discussion to compare and contrast IMS with Hadoop,

[1] A detailed report on the IMS 13 benchmark can be downloaded using this link: www-01.ibm.com/software/data/ims/benchmark.html

which is a non-relational system that is enjoying significant attention in the marketplace as a system for handling very large volumes of big data. Hadoop was initially designed for building and processing large Internet search indexes, but it is now also being used for loading, transforming and analyzing large volumes of data, especially multi-structured data such as Web pages and logs, text documents, sensor event data and so forth.

The attraction of Hadoop is that it is open source software that runs on low-cost hardware. Hadoop achieves its performance by exploiting the power of parallel computing on a large number of servers. The number of server nodes involved may run into hundreds and even thousands. The theory of this divide and conquer approach is that using hardware to achieve performance is simpler and easier than employing costly sophisticated commercial software to solve the problem. The downside of the Hadoop approach is that more hardware nodes means higher power and floor space requirements and an increasing number of failure points. The Hadoop approach is therefore very different from the IMS development approach where the focus is on chipping away at removing bottlenecks and instruction path lengths. This is why the IMS 13 results were achieved on a single IMS server.

The comparison between IMS and Hadoop is somewhat "tongue in cheek" because the two products support very different workload and availability requirements. The comparison does demonstrate, however, that the purchase cost of the hardware is not the only factor to consider. The total cost of ownership of the complete system is an equally, if not more, important metric. The comparison also shows that there is no single solution that can address all customer needs, especially at the extremes of the application performance spectrum. Instead, to move forward, companies will need to employ multiple systems to satisfy their IT needs and there will be a growing requirement for these systems to interoperate with each other.

INTEGRATING IMS WITH OTHER SYSTEMS

In addition to on-going needs to provide ever higher performance, the other driving forces behind IMS development are to improve ease of installation and maintenance, and to enhance the ability of IMS to coexist with other systems. Coexistence with other systems is focused in two main areas of development (see Figure 2):

1. Allow non-IMS systems to interact with IMS at a business transaction level and also at a database level (items 1 and 2 in Figure 2). Examples of non-IMS systems here include, for example, web services, cloud-based environments, mobile platforms, and SQL-based BI applications and tools. IMS provides two interfaces to enable this interaction: Open Transaction Manager Access (OTMA) and Open Database Manager Access (ODBA).
2. Enable new IMS transaction applications to be developed using a modern programming language such as Java and to give these applications access to data managed by both IMS and non-IMS systems (item 3 in Figure 2). Access to non-IMS systems includes, for example, DB2 for z/OS relational databases and z/OS files.

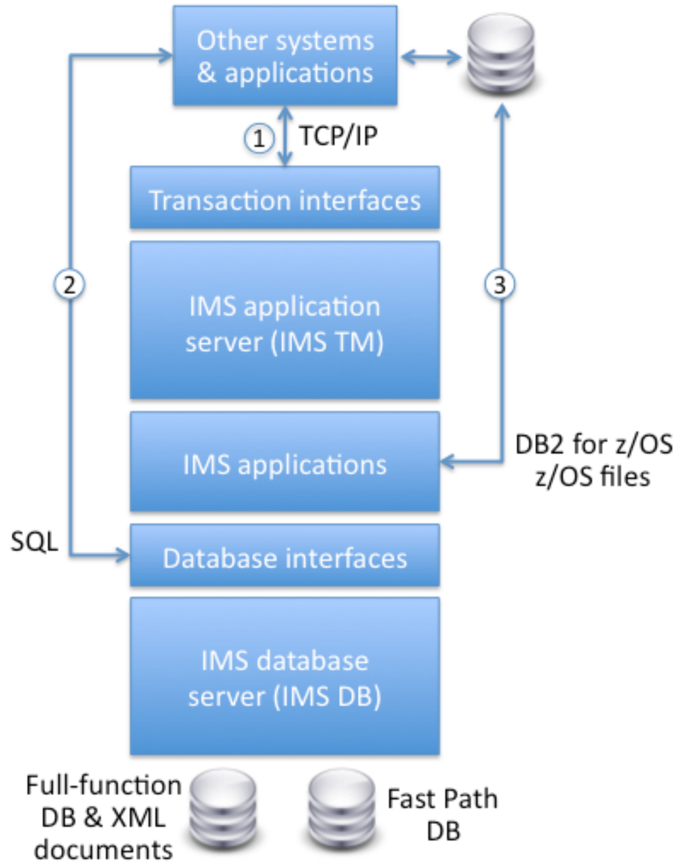

FIGURE 1. INTEGRATING IMS WITH OTHER SYSTEMS

Interacting with IMS at a Business Transaction Level

Early IMS applications were driven from so-called "green screen" terminals that interacted with IMS using a SNA/VTAM networking connection. With the advent of client/server computing, the Internet and mobile computing, IMS customers needed to move beyond terminal-driven and proprietary SNA networking approaches to exploit these new technology developments. To meet this need IBM developed the Open Transaction Manager Access (OTMA).

OTMA allows other IBM software and customer applications to send business transactions to IMS and receive the responses. IBM provides a number of OTMA connectors and software that provide a wide range of different approaches and technologies for entering IMS transactions. Two key focus areas here are support for web services (using technologies such as JCA, SOAP/XML and REST/JSON) and messaging (using technologies such as JMS and IBM WebSphere MQ). These interfaces enable IMS to be integrated, not only with existing enterprise and partner systems, but also systems in rapidly growing environments such as mobile and cloud computing.

To provide a secure and governed interface to IMS from the mobile and cloud environments, IBM offers the IMS Mobile Feature Pack (IMS Mobile), which enables developers to discover, model, enable, and publish IMS transaction as services. These services can be consumed by both mobile and cloud applications.

IMS Mobile leverages the IBM WebSphere Application Server, which provides support for the lightweight REST web protocol and the JSON data-interchange format. IMS Mobile can be used to supplement existing roll-your-own mobile solutions as well as mobile solutions that use either the IBM Worklight mobile platform, or the IBM WebSphere DataPower appliance (discussed later in this paper), or both.

Interacting with IMS at a Database Level

IMS provides two main types of databases: full-function databases and Fast Path databases. IMS full-function databases can be partitioned using IMS's High Availability Large Database (HALDB) capability. To support large database volumes, HALDB databases can split into up to 1001 partitions, each of which can be up to 40 gigabytes in size. One advantage of HALDB partitions is that they can be reorganized online for high availability, and can be taken offline independently of each other, thus easing the management of large databases.

In addition to the classic DL/I interface to IMS data, IMS also provides SQL access to the data from applications written in COBOL and Java. The Microsoft .Net development environment is also supported for programming languages such as C# and Visual Basic. In an SQL environment, each segment in an IMS database hierarchical structure is seen as a single relational table. Each field in the segment then becomes a column of the table. The SQL support is based on the SQL-92 standard, but obviously there are some restrictions when building an SQL interface to a non-relational system.[2] IBM has extended its Cognos BI toolset to support SQL access to IMS data.

IMS full-function databases can also be used to manage XML documents. An XML document can be stored intact (together with indexes into those intact documents) or can be "shredded" across multiple IMS DB segments. The XML data can then be accessed using a standard SQL SELECT statement.

IMS allows external systems and applications to access IMS DB data through the Open Database Manager Access (ODBA) interface. Both local and remote access using TCP/IP are supported. This interface provides broad support for applications written in Java, which can use both DL/I calls and SQL statements to access IMS data. SQL access is enabled through a supplied JDBC driver.

[2] A detailed report on IMS SQL support and other IMS DB features can be downloaded using this link: publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.ims13.doc.apr/dfsaprk0.pdf?noframes=true

An Appliance Approach to Integration: WebSphere DataPower

Many IMS customers are using the features outlined above to exploit new technologies and to integrate both new and existing enterprise systems with IMS. To simplify the integration work and modernize their IMS environments a number of them are using the WebSphere DataPower appliance. This purpose-built hardware and software appliance supports many of the security and integration technologies provided by IBM that support IMS. For example, the appliance provides the security, control, integration and optimization needed for mobile workloads. This is done using an integration gateway to mobile traffic coming through the IBM Worklight server. WebSphere DataPower provides similar features for connecting to a cloud-based computing environment. The end point for both of these examples can be an IMS system. A leading Canadian bank and financial service provider, for example, is using a DataPower Appliance to modernize its IMS system, and support Web and mobile interfaces to IMS using TCP/IP, HTTP, FTP protocols and SOAP and REST Web services.

Conclusion

After nearly 50 years of development, IMS still remains a key component of mission-critical applications across a wide range of organizations throughout the world. During this time both performance requirements and technologies have changed dramatically. Even the IBM mainframe today bears little resemblance to the systems that were used a few years ago. As has hopefully been demonstrated in this paper, IMS has managed to keep up with these changes while also at the same time protecting customer investment in older technologies. IMS programs written in 1969, for example, still work on IMS 13 today.

In the area of performance, IMS is actually ahead of customer needs. When IMS first achieved 1,000 TPS in 1987, no customer required those volumes, but now these rates are commonplace. Today, few customers need to support 100,000 TPS, but given ever-growing transaction volumes it is only a matter of time before they do and IMS is already well positioned to support this volume. Given this high performance and IMS's increasing focus on supporting and integrating with ever changing enterprise systems, the product will be part of the enterprise environment for the foreseeable future.

About Colin White

Colin was an IBM employee from 1970 until 1983 where he worked on both IMS and DB2. After leaving IBM, he formed DataBase Associates and BI Research and worked closely with the Relational Institute founded by Dr. E.F. Codd and C.J. Date. He is now an industry analyst and educator specializing in business intelligence, data warehousing, database management and big data.

