



IBM InfoSphere DataStage Performance and Scalability Benchmark Whitepaper

Data Warehousing Scenario

Contents

5	Overview of InfoSphere DataStage
7	Benchmark Scenario
	Main Workload Description
	Data Characteristics
11	Benchmark System
13	Performance Results
15	Discussion
18	Conclusion

Executive Summary

Information available to business leaders continues to explode, creating the opportunity for a new kind of intelligence to drive insightful, informed business decisions. On a Smarter Planet, information is the life blood of an organization, offering a transformative power for businesses that are able to apply it strategically.

In order for business and IT leaders to unlock the power of this information, they need to work together to pursue an information-led transformation. This relies on three areas of focus:

- **Planning an information agenda**, a strategic plan to align information with business objectives, reflecting aspects that may be unique to the industry or organization.
- **Applying business analytics to optimize decisions** and identify patterns in vast amounts of data and extract actionable insights through planning, monitoring, reporting and predictive analysis.
- **Establishing a flexible information platform** that will provide the necessary agility across the technology platform, infrastructure and common software services at an optimal cost.

Creating a flexible information platform relies heavily upon an information integration architecture that can scale as the volume of information continues to rise, as well as perform optimally to deliver trusted information throughout the organization.

IBM® InfoSphere Information Server is an extensive data integration software platform that helps organizations derive more value from the complex, heterogeneous information spread across their systems. It includes a suite of components that provides end-to-end information integration that offers discovery, understanding, transformation and delivery capabilities.

IBM InfoSphere DataStage is the information integration component of InfoSphere Information Server. Recognized as an industry-leading integration product by analysts and customers alike, InfoSphere DataStage delivers the performance and scalability required by customers who have growing, complex information integration needs.

In the whitepaper entitled, “The seven essential elements to achieve highest performance & scalability in information integration,” IBM laid out requirements companies must consider when developing an information integration architecture equipped to grow with their needs. The seven requirements are:

1. A dataflow architecture supporting data pipelining
2. Dynamic data partitioning and in-flight repartitioning of data
3. Adapting to scalable hardware without requiring modifications of the data flow design
4. Support and take advantage of available parallelism of leading databases
5. High performance and scalability for bulk and real-time data processing
6. Extensive tooling to support resource estimation, performance analysis and optimization
7. Extensible framework to incorporate in-house and third-party software

InfoSphere DataStage is designed to support these seven elements and provide the infrastructure customers require to achieve high performance and scalability in their information integration environment.

This whitepaper provides results of a benchmark test performed on InfoSphere DataStage 8.1 to illustrate its capabilities to address customers' high performance and scalability needs. The scenario focuses on a typical – that is, robust – set of requirements customers face when loading a data warehouse from source systems. In particular, the benchmark is designed to use the profiled situation to provide insight about how InfoSphere DataStage addresses three fundamental questions customers ask when designing their information integration architecture to achieve the performance and scalability characteristics previously described:

- How will InfoSphere DataStage scale with realistic workloads?
- What are its typical performance characteristics for a customer's integration tasks?
- What is the predictability of the performance in support of capacity planning?

In order to fully address these questions, the benchmark was designed to consider a mix of operations and scalability elements including filters, transformations, lookup and enrichment operations, joins, grouping, sorts and custom business logic. The workload reflects complex business logic and multi-step and multi-stage operations. Further, the benchmark demonstrates performance of integration jobs using data volumes that scale up to one terabyte.

The results of this benchmark highlight InfoSphere Data Stage's ability to deliver high performance and scalability consistent with what many customers seek for their integration environments.

- **Scalability** in this benchmark demonstrates that InfoSphere DataStage's execution time remains steady as data volumes and the number of processing nodes increases proportionally. This underscores InfoSphere DataStage's capability to use its parallel framework to effectively address a complex job.
- **Performance characteristics** are demonstrated by the speedup and throughput tests that were conducted. These tests show InfoSphere DataStage's capability to manage the same workload proportionally faster with additional resources. Further, the benchmark demonstrates InfoSphere DataStage's capability to leverage parallelism of its operators when the number of InfoSphere DataStage configuration nodes and hardware cores scale.

- **Capacity requirements** can be estimated through the predictability demonstrated by the benchmark results. Therefore, using a combination of a customer's workload expectations with the throughput tests, it is possible to determine a relative complexity factor for a given environment. Using this factor and the known throughput rates, customers can estimate the number of processing cores needed for their workload.

InfoSphere DataStage is a high performance, scalable information integration architecture. This benchmark represents a sample of a common integration scenario of loading a data warehouse. While specific customer requirements and objectives will vary, the benchmark can be used as an example to estimate how InfoSphere DataStage may address a customer's questions regarding scalability, performance characteristics and capacity requirements.

Overview of InfoSphere DataStage

InfoSphere DataStage provides a designer tool that allows developers to visually create integration jobs. *Job* is used within InfoSphere DataStage to describe extract, transform and load (ETL) tasks. Jobs are composed from a rich palette of operators called stages. These stages include:

- Source and target access for databases, applications and files
- General processing stages such as filter, sort, join, union, lookup and aggregations
- Built-in and custom transformations
- Copy, move, FTP and other data movement stages
- Real-time, XML, SOA and Message queue processing.

Additionally, InfoSphere DataStage allows pre- and post-conditions to be applied to all these stages. Multiple jobs can be controlled and linked by a sequencer. The sequencer provides the control logic that can be used to process the appropriate data integration jobs. InfoSphere DataStage also supports a rich administration capability for deploying, scheduling and monitoring jobs.

One of the great strengths of InfoSphere DataStage is that when designing jobs, very little consideration to the underlying structure of the system is required and does not typically need to change. If the system changes, is upgraded or improved, or if a job is developed on one platform and implemented on another, the job design does not necessarily have to change. InfoSphere DataStage has the capability to learn about the shape and size of the system from the InfoSphere DataStage configuration file. Further, it has the capability to organize the resources needed for a job according to what is defined in the configuration file. When a system changes, the file is changed, not the jobs. A configuration file defines one or more processing nodes with which the job will run. The processing nodes are logical rather than physical. The number of processing nodes does not necessarily correspond to the number of cores in the system.

Following are additional factors that affect the optimal degree of parallelism:

- CPU-intensive applications, which typically perform multiple CPU-demanding operations on each record, benefit from the greatest possible parallelism up to the capacity supported by a given system.
- Jobs with large memory requirements can benefit from parallelism if they act on data that has been partitioned and if the required memory is also divided among partitions.
- Applications that are disk- or I/O-intensive, such as those that extract data from and load data into databases, benefit from configurations in which the number of logical nodes equals the number of I/O paths being accessed. For example, if a table is fragmented 16 ways inside a database or if a data set is spread across 16 disk drives, one should set up a node pool consisting of 16 processing nodes.

Another great strength of InfoSphere DataStage is that it does not rely on the functions and processes of a database to perform transformations: while InfoSphere DataStage can generate complex SQL and leverages databases, InfoSphere DataStage is designed from the ground up as a multipath data integration engine equally at home with files, streams, databases, and internal caching in single-machine, cluster, and grid implementations. As a result, customers in many circumstances find they do not also need to invest in staging databases to support InfoSphere DataStage.

Benchmark Scenario

The benchmark is a typical ETL scenario which uses a popular data integration pattern of loading and maintaining a data warehouse from a number of sources. The main sources are securities trading systems from where details about trades, cash balances and accounts are extracted. The securities trading system source data sets are generated from the industry standard TPC-E¹ benchmark. Flat files are created for the source data using the TPC-E data generator and represent real-world situations where data arrives via extract files. The integration jobs then load the enriched and transformed data into a target database. See Figure 1.

While this data warehouse integration scenario reflects the financial services industry, it is applicable across many industries, especially those in which there is a need to analyze transaction data using a variety of dimensional attributes. For instance, call detail records (CDRs) could replace trades to obtain a telecommunications data warehouse scenario. Likewise, point-of-sale (POS) transaction data could be substituted for trades to create a retail industry data warehouse scenario.




Figure 1: Data Warehousing ETL Scenario

Main Workload Description

The benchmark consists of a variety of jobs that populate the warehouse dimension and fact tables from the source data sets. Each job contains a number of stages and each stage represents processing logic that is applied to the data. In InfoSphere DataStage, the stage processing logic is applied to data in memory and the data flows to the memory areas of subsequent stages using pipelining and partitioning strategies.

There are two main jobs that load the *Trades* and *CashBalances* tables from their respective source data sets. These jobs perform the initial load of the respective data into the warehouse. These two jobs are described below.

In the *Load_Trades* job (Figure 2), input data for each trade is retrieved from a file and various fields are checked against other dimension tables (through lookup operations) to ensure a trade is valid. The trade data then flows to a transform stage where additional information is added for valid trades while invalid trades are redirected to a separate *Rejects* table. All valid trades are then bulk loaded into the target database table. The *Load_Trades* job processes roughly half of the of the input data volume. If the input data set is 1 TB, it processes 500 GB of data.

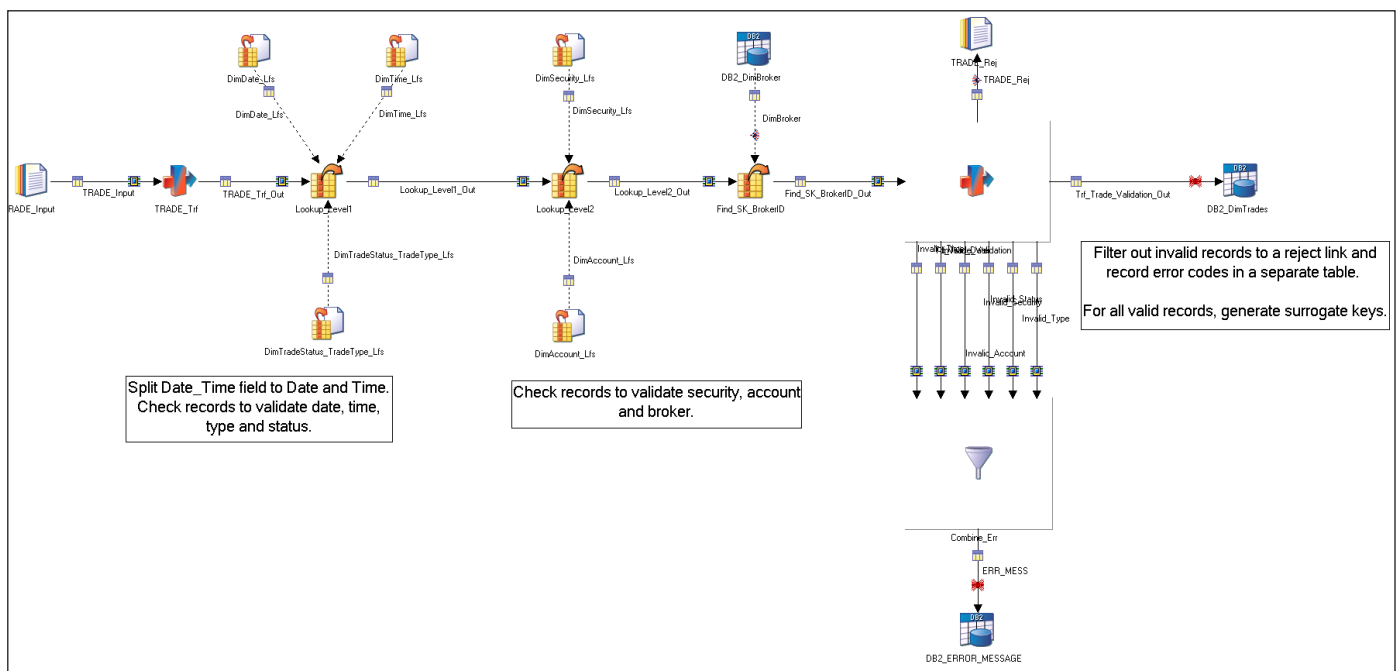


Figure 2. Load_Trades Job

In the *Load_CashBalances* job (Figure 3), input data for each cash transaction is retrieved from a file and various fields are checked against other dimension tables. The transaction data is joined with *Trades and Accounts* (which are loaded in earlier steps) to retrieve the balance for each customer account and establish relationships. A transform stage is used to compute the balance after each transaction for every individual account and results are bulk loaded into the target database table.

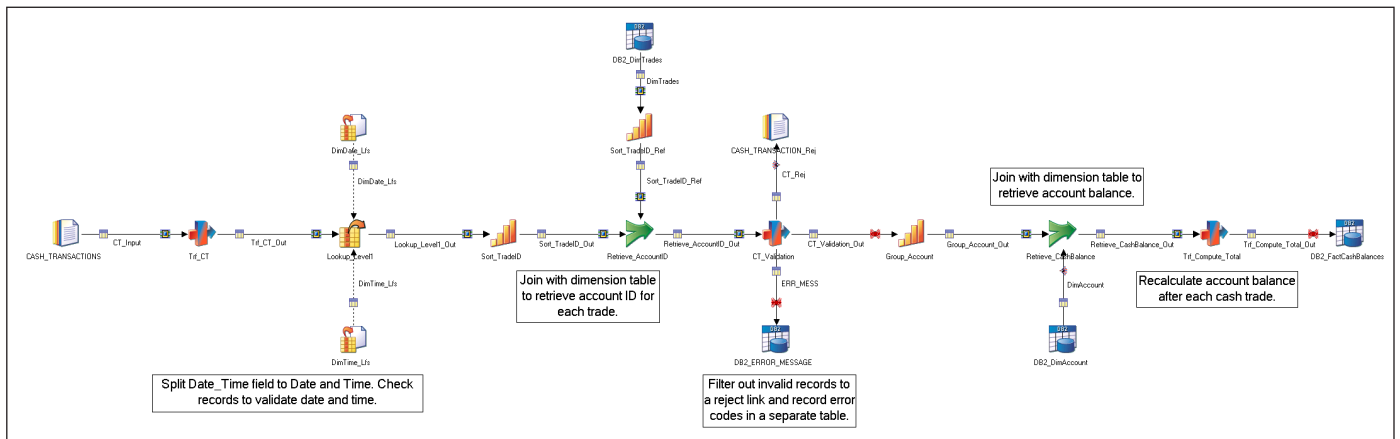


Figure 3. *Load_CashBalances* Job

These two sources account for a significant amount of processing in this scenario. They are representative of typical jobs in data integration since they include many standard processing operations such as filters, joins, lookups, transforms, sorts, grouping operations and multiple outputs. The *Load_Cashbalances* job processes almost 95% of the input data volume. If the input data set is 1 TB, it processes 950 GB of data.

Data Characteristics

The main source of data for this benchmark is generated using the TPC-E data generator tool. The tool generates source data following established rules for all the attributes and can generate data of different scale factors.

One important goal of the benchmark is to test input data sets of different scale factors. For this benchmark, data sizes of 250 GB, 500 GB and 1 TB have been chosen to study the scalability and speedup performance of the main workload tasks.

The main source data files are Trade, CustomerAccount and CashTransaction. Data types in the records include numerics (integers, decimals and floating point), string and character data, identifiers and flags. Table 1 summarizes the characteristics of the main input source files in the benchmark.

Table 1. Source Data Files Characteristics			
Source File Name	Number of columns	Average Row size (bytes)	Data types
Trade	15	179	Int, timestamp, string, decimal
CustomerAccount	6	91	Int, string
CashTransaction	4	141	Int, timestamp, string, decimal

Benchmark System

The benchmark system is a 64-Core IBM P6 595 running the AIX 6.1 operating system.

This system was divided into eight logical partitions (LPARs) allocated equally between InfoSphere DataStage and IBM DB2. These LPARs are allocated in order to test and demonstrate the scalability of InfoSphere DataStage in a Symmetric MultiProcessor (SMP) configuration using one LPAR and in a mixed SMP and cluster configuration using multiple LPARs. See Figure 4.




Figure 4. Server Configuration

Storage was provided by six DS4800 SAN storage servers. Two storage servers were dedicated to InfoSphere DataStage and the other four were dedicated to IBM DB2.

The table below shows the characteristics of each LPAR:

Table 2. Benchmark LPAR Characteristics	
Processor	<ul style="list-style-type: none"> • 8 physical cores, dedicated • SMT (Simultaneous Multithreading) enabled • 5GHz
Memory	<ul style="list-style-type: none"> • 32GB
Network	<ul style="list-style-type: none"> • 1Gbps Ethernet, MTU 1500 • 10Gbps Virtual Ethernet, MTU 9000
Storage	<ul style="list-style-type: none"> • Internal SCSI disks • 2 Fiber Channel Adapters connected to DS4800 SANs

InfoSphere DataStage Storage Configuration

Four of the eight LPARs were assigned to InfoSphere DataStage. The operating system and InfoSphere DataStage binaries were installed on the internal SCSI disks. The staging area, data set and scratch area for InfoSphere DataStage were configured to use the DS4800 SAN storage. There are nine logical units numbers (LUNs) configured as RAID-0. A total of 1,394GB of storage was made available for each InfoSphere DataStage LPAR and used for source / reference data files, scratch and other working storage. It should be noted that this storage was not exclusive to this benchmark. Only a portion of this storage was used in the individual benchmark experiments. InfoSphere DataStage can be configured to make a server appear as many logical nodes since it uses a shared-nothing processing model. In this benchmark, InfoSphere DataStage was configured to use from one to eight logical nodes per LPAR.

IBM DB2 Server Storage Configuration

The other four LPARs were assigned to IBM DB2. The operating system and IBM DB2 were installed on the internal SCSI disks. IBM DB2 tablespaces and transaction logs were located on the DS4800 SAN storage. There were four LUNs for each of the tablespaces and eight tablespaces per LPAR. These are DMS tablespaces on JFS2 filesystem in AIX. Each of the LUNs was set up as RAID-5 (3+Parity). IBM DB2 is configured to use eight logical partitions in an LPAR. This allows the IBM DB2 shared-nothing capability to take maximum advantage of the eight cores in each LPAR.

Table 3. Server Software	
Operating System	AIX 6.1 TL2 SP2, 64-bit
DataStage Server	IBM InfoSphere Information Server 8.1 • 1 to 8 logical nodes per LPAR
Database Server	IBM DB2 9.5 FixPak 3 • 8 Partitions per LPAR.

A Windows client (IBM System x3550) loaded with the InfoSphere DataStage client was used for development and for executing benchmark tests.

Performance Results

The benchmark included scalability and speedup performance tests using the *Load_Trades* and *Load_CashBalances* jobs on the 64-Core IBM P6 595 AIX server. Three data volumes, as described above, were considered: 250 GB, 500 GB, and 1 TB. The number of nodes for InfoSphere DataStage processing varied from eight nodes to 32 nodes scaling out across the one to four AIX LPARs.

Scalability

Figure 5 shows the scalability performance of the *Load_Trades* workload for different data volumes and nodes. The figure shows that the execution time of this job remains steady as the data volumes and the number of processing nodes increase proportionally. This result shows that InfoSphere DataStage's processing capability is able to scale linearly for very large data volumes. Doubling the volume and available resources resulted in approximately the same execution time.




Figure 5. Scalability of *Load_Trades* Job




Figure 6. Scalability of *Load_CashBalances* Job

Figure 6 shows the scalability performance of the *Load_CashBalances* job. For this job, only the 500 GB and 1 TB data volumes and corresponding number of nodes were considered. This is a complex job performing three sorts, two join operations and a grouping operation incurring significant I/O and CPU utilization during these different stages. This job also scales linearly showing that InfoSphere DataStage is able to partition and pipeline the data and operations within this complex job effectively. Doubling the volume and available resources resulted in approximately the same execution time.

Speedup

In the speedup tests, the jobs were executed against a fixed data volume as the number of processors was increased linearly. Figure 7 shows the speedup of the *Load_Trades* job for 250 GB of data using 8, 16 and 32 Power6 cores with 8, 16 and 32 InfoSphere DataStage configuration Nodes. The x-axis shows the number of nodes while the y-axis shows the execution time of the job. The y-axis shows execution time increasing downward in order to show that processing speedup is increasing with more processing nodes. The figure shows linear speedup of this workload which indicates that InfoSphere DataStage is able to process the same workload proportionally faster with more resources.




Figure 7. Speedup of *Load_Trades* Job. The figure shows linear improvement in the processing speed of the job as more processing nodes are added.

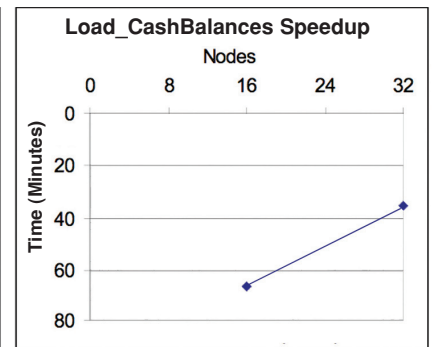


Figure 8. Speedup of *Load_CashBalances* Job. The figure shows linear improvement in the the processing speed of the job as more processing nodes are added.

Figure 8 shows the speedup of the *Load_CashBalances* job for 250 GB across 16, 32 Power6 cores with 16, 32 InfoSphere DataStage configuration Nodes. Linear speedup for this job is also observed. Note that speedup is obtained without having to retune any of the parameters of the job or make any changes to the existing logic. InfoSphere DataStage is able to take advantage of the natural parallelism of the operators once they have been specified even as the number of CPUs is scaled up.

Throughput

The overall throughput of InfoSphere DataStage can be computed for these jobs based on the above experiments. The overall *throughput* is defined as *the aggregate data volume that is processed per unit time*. As expected, throughput increases with an increase in the degree of parallelism of the jobs. In these experiments, due to linear scalability and speedup, the largest throughput rates are measured for the maximum number of nodes in the system.




Figure 9. Throughput Scalability of Load_Trades Job

Figure 9 shows the aggregate throughput of the *Load_Trades* job when varying the degree of parallelism. The throughput increases linearly with the number of nodes and reaches 500 GB/hr using 32 Power6 cores with a 32-node InfoSphere DataStage configuration. Note, adding more CPU capacity will increase the throughput since the limit of scalability has not been reached for the job. For instance, it can be projected that throughput will approach 1 TB/hr if 64 cores available for use by InfoSphere DataStage and that other resources such as I/O scale in a similar manner to InfoSphere DataStage. Similar throughput results are observed for the *Load_CashBalances* job.

Discussion

The Executive Summary highlighted the seven essential elements to achieve performance and scalability in information integration. These elements are architectural and functional in nature and they enable an information integration platform to perform optimally. This benchmark demonstrated a number of these elements in the following ways:

- Dataflow and pipelining architecture across all stages and operations to minimize disk I/O
- Data parallelism and partitioning architecture for scalability
- Using the same job design across multiple infrastructure / server configurations

- Support for integration with leading parallel databases
- Extensive tooling for designing jobs and for monitoring the behavior of InfoSphere DataStage engine

The results from the benchmark demonstrate the ability of InfoSphere DataStage to provide a high performance and highly scalable platform for information integration.

In addition, the introduction posed the following questions regarding the customer's expectations on performance and planning.

- How will InfoSphere DataStage scale with realistic workloads?
- What are its typical performance characteristics for a customer's integration tasks?
- What is the predictability of the performance in support of capacity planning?

These questions are addressed below using the benchmark results that have been described in the previous section.

Scalability

The benchmark demonstrates that InfoSphere DataStage is capable of linear scalability for complex jobs with high CPU and I/O requirements. Linear scalability up to 500 GB/hr was observed on 32 Power6 cores for these jobs and further scalability is possible with the addition of processing nodes. If a customer's job is less intensive, higher throughput rates should be possible. Many customers have challenging data integration requirements needing to process hundreds of gigabytes of data within short daily windows of three to four hours. This benchmark indicates that InfoSphere DataStage has the capability of processing significant amounts of data within short windows of time. Also, customer data volumes continue to grow. The extensibility of InfoSphere DataStage permits the addition of more processing nodes in the system over time to address monthly/quarterly/yearly workloads as well as normal business growth. Given the scalability results of this benchmark,

customers can gain insight into how InfoSphere DataStage might address their particular scalability and growth requirements, depending upon their expected data volume expansion.

Performance Characteristics

Most typical customer workloads will use a similar mix of operations and transformations as in this benchmark. Hence, the CPU and I/O requirements and characteristics of a customer's workload could be similar, depending on the customer's specific job design. Therefore, it is possible to project that typical customer workload processing rates will be similar to the rates that were observed in this benchmark. If the customer is able to provide an overview of the job descriptions and input sizes, one can project the behavior of the jobs using the benchmark throughput rates. For instance, if the customer's job is CPU intensive, the profile may be similar to the *Load_Trades* job. The performance of the customer's job can be projected accordingly. Suppose a new job needs to process 500 GB of data and it includes a mix of operations similar to the *Load_CashBalances* job. From the benchmark results, it can be estimated that this particular job can be processed in 140 minutes using 16 Power6 Cores in a 16-node InfoSphere DataStage configuration as described in this study.

Capacity Planning

Assume a workload is specified in terms of input data sizes, a coarse description of integration operations and the desired execution window. One can compare the benchmark workload's operations to this new workload and estimate a *complexity factor*. The complexity factor indicates whether the workload is *similar, less complex or more complex* than this benchmark. Using this factor and the known throughput rates, customers can estimate the number of cores required for this new workload by leveraging the proven predictability. For example, suppose a workload needs to process 300 GB of data in two hours and it has similar complexity to the benchmark jobs. Using the throughput scalability graph (Figure 9), it can be estimated that roughly ten cores could be used to process the workload in two hours.

It should be noted that the benchmark guidance on capacity is one important aspect of a rigorous capacity planning activity and is not meant to be a replacement. Customers should perform a thorough capacity planning activity by identifying all key requirements, of which performance is just one important aspect. Other characteristics such as availability, development and test requirements as well as platform characteristics should be considered. IBM has tools and resources to assist customers in such a capacity planning activity.

Conclusion

Organizations embarking on an information-led transformation recognize the need for a flexible information platform. This information platform needs to provide high levels of performance and scalability in order to process terabyte and petabyte data volumes in appropriate time windows. An information integration platform designed according to the seven essential elements of a high performance and scalable architecture is pivotal for satisfying such volumes and velocity of information.

This benchmark illustrates the capabilities of IBM's InfoSphere DataStage within a common scenario that will help many customers estimate how InfoSphere DataStage might bring performance and scalability to their integration needs. Linear scalability and very high data processing rates were obtained for a typical information integration benchmark using data volumes that scaled to one terabyte. Scalability and performance such as this is increasingly sought by customers investing in a solution that can grow with their data needs and deliver trusted information throughout their organization.



© Copyright IBM Corporation 2010

IBM Software Group
Route 100
Somers NY 10589
U.S.A.

Produced in the United States of America
February 2010
All Rights Reserved

IBM, the IBM logo, ibm.com, DB2, IMS, Informix, InfoSphere, Optim and VSAM are trademarks or registered trademarks of the IBM Corporation in the United States, other countries or both.

Microsoft and SQL Server are trademarks of Microsoft Corporation in the United States, other countries or both.

All other company or product names are trademarks or registered trademarks of their respective owners.

References in this publication to IBM products, programs or services do not imply that IBM intends to make them available in all countries in which IBM operates or does business.

Endnote

- 1 <http://www.tpc.org/tpce/default.asp>