



RODIN Data Asset Management

High Performance Extract/Transform/Load Benchmark

Performed at the IBM iSeries Teraplex Center

Rochester, MN.

June 2002

Detailed Results



Table of Contents

BENCHMARK OVERVIEW	3
EXECUTIVE SUMMARY	3
WHY IS THIS BENCHMARK IMPORTANT ?	4
BENCHMARK HARDWARE	5
PLATFORM	5
SERVER CONFIGURATION	5
TUNING.....	5
BENCHMARK SOFTWARE	6
MODIFICATIONS TO STANDARD SOFTWARE.....	6
RODIN ETL ARCHITECTURE.....	6
BENCHMARK SCENARIOS	8
SOURCE DATA	8
TEST SCENARIO #1. COMPLEX LOAD OF DETAIL LEVEL TABLE (ALL INSERTS).....	9
TEST SCENARIO #2. COMPLEX LOAD OF SUMMARY LEVEL TABLE (BOTH INSERTS AND UPDATES).....	9
TEST SCENARIO #3. COMPLEX LOAD OF BOTH DETAIL AND SUMMARY LEVEL TABLES CONCURRENTLY	10
RESULTS	10
TEST # 1: LOAD 200 MILLION ROWS - ALL INSERTS	10
TEST # 2: LOAD 200 MILLION ROWS INTO SUMMARY TABLE – INSERTS AND UPDATES	11
TEST # 3: LOAD 200 MILLION ROWS INTO BOTH DETAIL AND SUMMARY TABLES	12
MEASUREMENT METHODOLOGY	12
PREVIOUS BENCHMARK RESULTS	13
DISTRIBUTED DATABASE ENVIRONMENTS	14
CONCLUSIONS	15
ABOUT COGLIN MILL AND RODIN	16
ABOUT THE TERAPLEX CENTERS	16
CONTACT INFORMATION	17
FOR MORE INFORMATION REGARDING RODIN:.....	17
FOR MORE INFORMATION REGARDING THE IBM I SERIES TERAPLEX CENTER:.....	17

Benchmark Overview

This benchmark is the sixth in a series of similar benchmarks that have been performed by Coglin Mill at the iSeries Teraplex Center in Rochester Minnesota, since 1997. Coglin Mill was the first IBM business partner to utilize the resources of the Teraplex Center, just weeks after it was established. It is no coincidence that we have also been the first to be invited in to the Center to test and benchmark our software on each new generation of hardware released since that time.

This latest benchmark was performed on hardware, which had been announced, but was not (at the time of the benchmark) generally available - the IBM **@server** iSeries i890. Similarly, the latest version of OS/400, V5R2, including the latest version of DB2 UDB for iSeries, was used in the testing.

The purpose of the benchmark is twofold:

- To demonstrate the scalability and performance of the iSeries platform for high end Data Warehousing and data integration applications.
- To demonstrate how the RODIN Data Asset Management software takes advantage of the iSeries architecture and it's unique parallel load capabilities to perform very large ETL (Extract/Transform/Load) processes at industry leading speeds.

Executive Summary

The major conclusions that can be drawn from the results of this benchmark include:

- The IBM **@server** iSeries is a highly scalable platform that can easily handle very large scale data warehousing and data integration applications,
- The RODIN Data Asset Management software fully leverages and enhances the advanced technologies of this hardware and database platform, and
- Together this hardware and software platform delivers performance levels, in a real world environment, not achieved to date by any other published benchmark on any other hardware / software platform.

Existing iSeries customers can be certain their chosen platform is highly suited to Data Warehousing, Business Intelligence and Customer Relationship Management applications, even at the very high end of business volumes.

Furthermore, customers with a mixed technology environment should consider using the iSeries and RODIN if they are looking for scalability, ease of use and a low cost of ownership (leading to improved ROI) when deciding on their DW / BI / CRM platform.

Why is this Benchmark Important ?

Data volumes are growing at an exponential rate. Several years ago very few organisations considered they might need to manage terabytes of data in their data warehouses, but today this is quite common. Visionaries are now talking about **petabytes** of data (1 petabyte = 1,024 terabytes).

For most organisations, data volumes this size are still a long way off, however the trend is indisputable. As data warehouses grow to tens or hundreds of terabytes, it is clear that both hardware and software need to scale similarly.

While the hardware, database, and disk subsystems need to manage these huge amounts of data once loaded, the ETL (Extract, Transformation, and Load) process is just as important. It must be capable of loading many gigabytes of data on a daily basis, and to do this must take advantage of all hardware processing resources – which is not an easy task.

Speed of processing and scalability are just as important for smaller companies as they are for very large ones. Knowing the software is fully optimized to the available hardware resources ensures throughput is maximized and processing times minimized on all iSeries models, whatever the size. This delivers the best possible TCO and ROI figures at all levels within the entire iSeries range.

Real world benchmarks such as this one are the proving ground for very large-scale iSeries data warehouse applications as well as for the small to medium size applications currently being implemented by the majority of organizations today. No matter what its size today, its very important and comforting to know your data warehousing implementation can grow as and when the need arises without any costly and time consuming hardware platform or software technology replacements.

And now, for the first time, this benchmark proves the IBM **@server** iSeries i890 and RODIN Data Asset Management software are in the same performance class as high-end Unix, mainframe and Teradata data warehousing implementations. With the added benefits of greatly improved ease of use and proven low TCO, this solution provides a real alternative to the currently more prevalent data warehousing platforms.

Benchmark Hardware

Platform

RODIN runs natively on the IBM **@server** iSeries platform. This benchmark was performed on the latest addition to the iSeries line – model i890.

The i890 system features the eighth generation 64 bit PowerPC processor, which utilizes IBM's copper and silicon-on-insulator technologies. These processors are regarded as the fastest 64 bit microprocessors available today by a clear margin.

Server Configuration

Model:	IBM iSeries i890, feature code 2488.
CPU:	32 x POWER4 1.3GHz 64-bit RISC Microprocessors.
Memory:	256GB total: 240GB available in storage pool for benchmark.
Disk:	15.9 TB (10% utilized prior to benchmark) 704 17GB drives. 172 36GB drives. Total 876 drives/disk arms. RAID 5 protected.
Operating System:	OS/400 V5R2
Database:	DB2 UDB for iSeries (integrated database).

Tuning

Very little tuning was performed (or necessary) on the system prior to the benchmarks.

Unlike most other servers, database tuning (e.g. playing with partitions, containers and tablespaces) is not necessary. The integrated nature of the DB2 UDB for iSeries database and the unique single level storage concept completely removes the need for time-consuming database tuning.

SMAPP (System Managed Access Path Protection) was turned off prior to the benchmarks to prevent the system from journaling access path changes.

Automatic Performance Adjustment (System Value QPFRADJ) was switched off, to ensure a fixed size to the main storage (memory) pool for the benchmarks.

No other adjustments were made. All normal operating system tasks (e.g. host server jobs, etc.) remained active during the benchmarks.

Benchmark Software

RODIN Data Asset Management Version 3 Release 2 Modification 0, PTF level 6 was used for the benchmark. This is the latest commercially available release of RODIN.

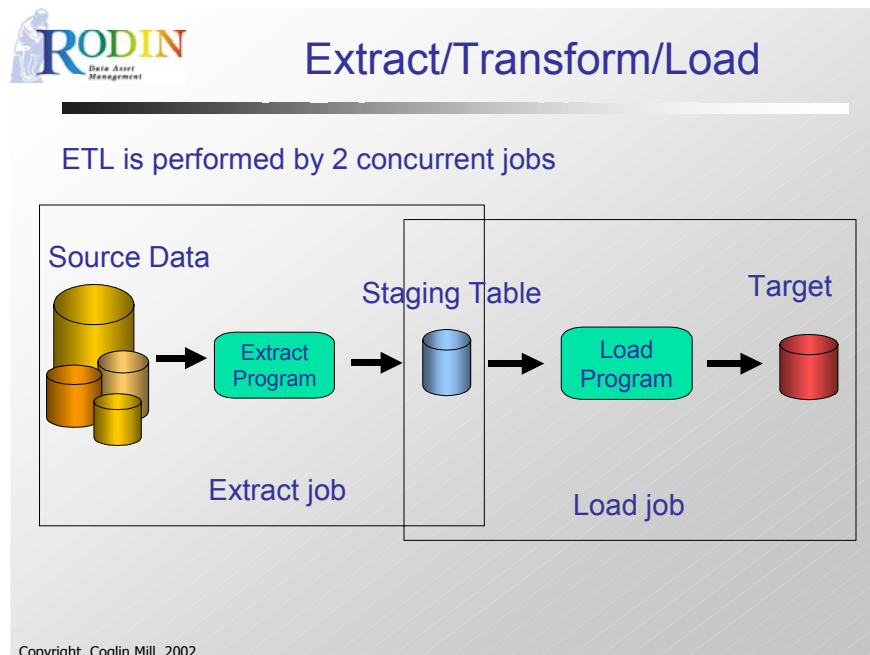
Modifications to standard software

None. We included this section only to highlight that NO modifications, adjustments or steroids were used to enhance performance. The RODIN software typically requires no set-up or tuning to take advantage of the configuration of the server hardware on which it is installed. Unlike many benchmark activities that are conducted on unrealistic system configurations to optimize to a particular measurement, this benchmark represents a very realistic hardware and software configuration that would be installed in many customer situations.

RODIN ETL Architecture

RODIN is designed to take advantage of all applicable functionality in the OS/400 operating system and integrated DB2 database. Automatically generated ILE RPG programs perform both the extract from the source tables and the load into the target table(s).

Figure 1 demonstrates the 2-stage design of a RODIN ETL process: 2 separate batch jobs concurrently perform the extract and load. This greatly enhances throughput in a multi-processor environment.



Copyright, Coglin Mill, 2002

Figure 1. RODIN ETL Architecture

For large loads on multiple CPU (n -way) systems, RODIN's unique parallel processing technology can also easily split the source data into n job streams to fully utilise the resources of all CPUs, as per figure 2.

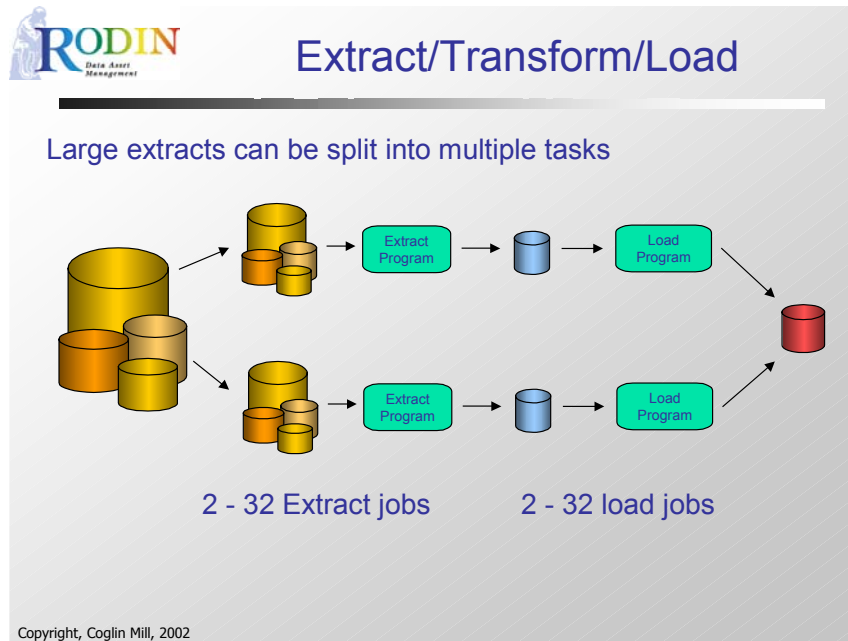


Figure 2. RODIN Parallel load

RODIN also has the capability to load multiple target tables at the same time – even in parallel. Figure 3 shows a 2-way parallel load of 4 different target tables.

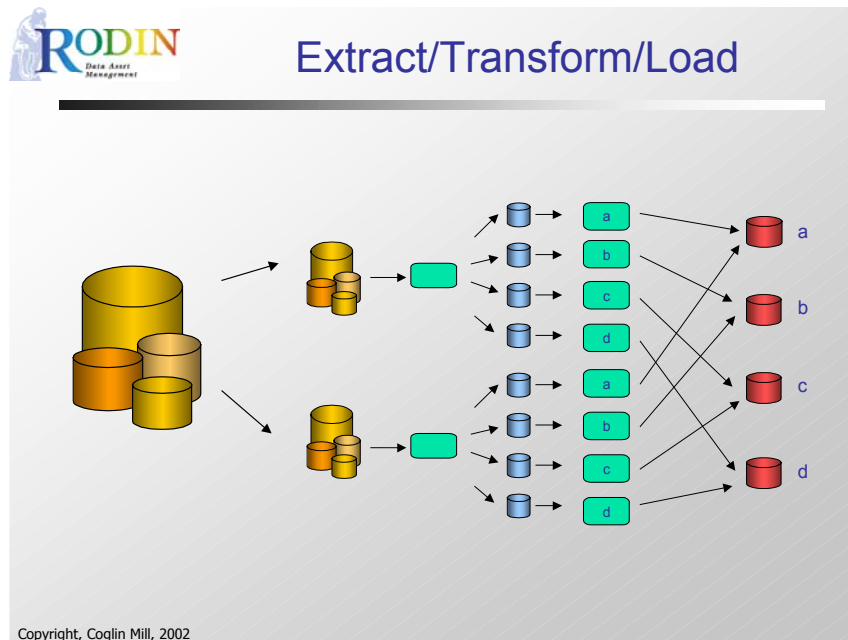


Figure 3. RODIN Parallel load of 4 tables

Benchmark Scenarios

Source Data

The source data for all benchmarks was a set of related tables, containing sales transaction information:

Primary table:

Shipments Table, containing 200,000,000 rows. Record length 105 bytes

Associated Tables:

7 reference tables, containing Store, Geography, Customer and Product information were accessed.

These tables contained between 500 and 1.2 million rows, and varied from 27 bytes to 202 bytes in record length. Figure 4 shows the table relationships (actual key joins involve multiple columns in all cases). The total number of bytes of data associated with each source 'transaction' was 875 bytes. Extrapolated out this is logically equivalent to $200,000,000 * 875 = 175\text{GB}$ of source data.

The shipments table did contain a primary key, however it was accessed sequentially. All associated tables contained a suitable primary index for the required joins.

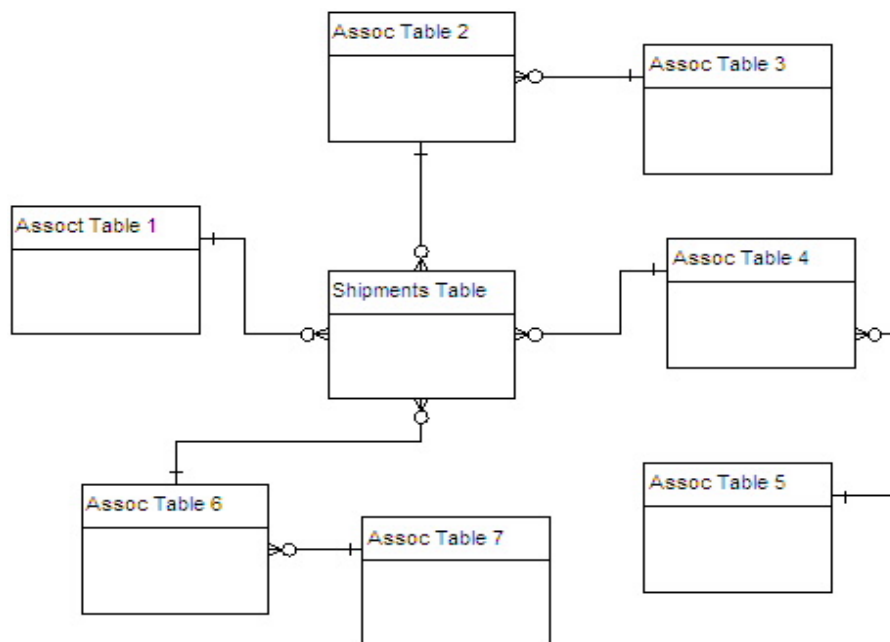


Figure 4. Source Table relationships

Test Scenario #1. Complex load of detail level table (all inserts).

In this benchmark scenario, two similar tests were conducted. The same source data was used in each test:

1. This test represents a scenario of a small fact table in a star schema data warehouse. It inserts 200 million rows into a target table with no index, and a record length of **100** bytes. The table contained 12 columns of 3 different data types:
 - 1 Date column
 - 5 Character columns
 - 6 Packed Decimal columns
2. This test is representative of the load of a typical large table in a relational data warehouse. It insert 200 million rows into a target table with no index, and a record length of **500** bytes. The table contained 28 columns of 3 different data types:
 - 2 Date columns
 - 16 Character columns
 - 10 Packed Decimal columns

Various business rules and transformations were performed:

- Referential integrity to ensure a matching row was found in each associated table with associated error handling and reporting.
- A number of user defined data validation rules to test the integrity of the data, with associated error handling and reporting.
- Arithmetic operations.
- Date conversion from numeric CYYMMDD format to *ISO database format.
- Substring

Test Scenario #2. Complex load of summary level table (both inserts and updates)

In this benchmark scenario, a 100-byte target table, identical to the table from test #1 was used, however this time the 200 million source rows were aggregated to 48 million rows in the target table. A unique primary index existed on the target table, and this index was maintained during the load. The same referential integrity, business rules and transformations were applied.

Test Scenario #3. Complex load of BOTH detail and summary level tables concurrently

In this benchmark scenario, the 100-byte detail (non-indexed) target table and the 100 byte indexed summary table were loaded concurrently.

Results

Test # 1: Load 200 million rows - all inserts

a) Load of 100-byte table

Number of Job streams	1	12	24	32
Elapsed Time (seconds)	14746	1534	1385	1425
Rows/Hour (in millions)	48.8	469.4	519.9	505.3
GB/Hour	4.2	40.7	45.1	43.8

Table 1 – Load of 100-byte table

b) Load of 500-byte table

Number of Job streams	1	12	24	32
Elapsed Time (seconds)	14746.0	2494.0	1970.0	2239.0
Rows/Hour (in millions)	48.8	288.7	365.5	321.6
GB/Hour	23.1	136.3	172.6	151.8

Table 2 – Load of 500-byte table

Notes

GB/hour measurement is based on the target table size.

Analysis

These results are slightly better than expected, based on projections from previous benchmarks on earlier systems.

It is interesting to note that the maximum throughput was achieved using a 24-way parallel load, whereas earlier benchmarks saw maximum throughput where the degree of parallelism matched the number of CPUs. It was therefore expected that a 32-way split would achieve the best results.

This new behaviour is attributed to the imbalance in capacity / performance between the 32 processors and the IO subsystem on this particular server configuration. The IO

subsystem becomes saturated when servicing the 48 concurrent jobs, each performing extremely high IO velocity. Adding additional parallel jobs simply increases the requests to the IO subsystem resulting in a detrimental affect on overall performance. This is not seen as a negative, in fact it indicates that the system is making more efficient use of the available resources, negating the need to split the job into more parallel tasks.

The 100-byte load achieved the highest rows/hour, whereas the 500-byte table load achieved a significantly higher GB/hour rate, at the expense of rows/hour.

Test # 2: Load 200 million rows into summary table – inserts and updates

Number of Job streams	1	12	24	32
Elapsed Time (seconds)	14853.0	1589.0	1469.0	1503.0
Rows/Hour (in millions)	48.5	453.1	490.1	479.0

Table 3 – Load of 100-byte table summary table

Notes

GB/hour is not measured in this test, as it would be a confusing metric. The detail loads demonstrate the GB/hour performance.

Rows/hour measurement is based on the number of source rows processed. The 200 million rows were aggregated to 48 million target rows.

Aggregation is achieved by updating existing rows in the target, rather than performing aggregation in memory and writing the final result to the table. This approach allows full re-start recovery in the event of a system failure (unlike memory based processes), as well as other unique RODIN capabilities.

Analysis

These results were noticeably better than expectations: being within 4% of the equivalent detailed table load with no index. This is attributed to the same factors that were noted in the detail loads in test #1. We are at the physical IO limit of this server configuration and the system is easily managing to keep the access path updated with minimal overhead.

Test # 3: Load 200 million rows into both detail and summary tables

Number of Job streams	1	12	24	32
Elapsed Time (seconds)	23188.0	2575.0	2304.0	2551.0
Source Rows/Hour	31.1	279.6	312.5	282.2
Target Rows/Hour	62.2	559.2	625.0	564.5

Table 4 – Concurrent Load of 100-byte detail *and* summary tables

Notes

GB/hour is not measured in this test, as it would be a confusing metric. The detail loads demonstrate the GB/hour performance.

Since two inserts/updates are occurring for each source row, the rows/hour metric is calculated at both source and target level.

Analysis

This test demonstrates the value of RODIN's ability to load multiple target tables concurrently, with support for both inserts and updates.

Measurement Methodology

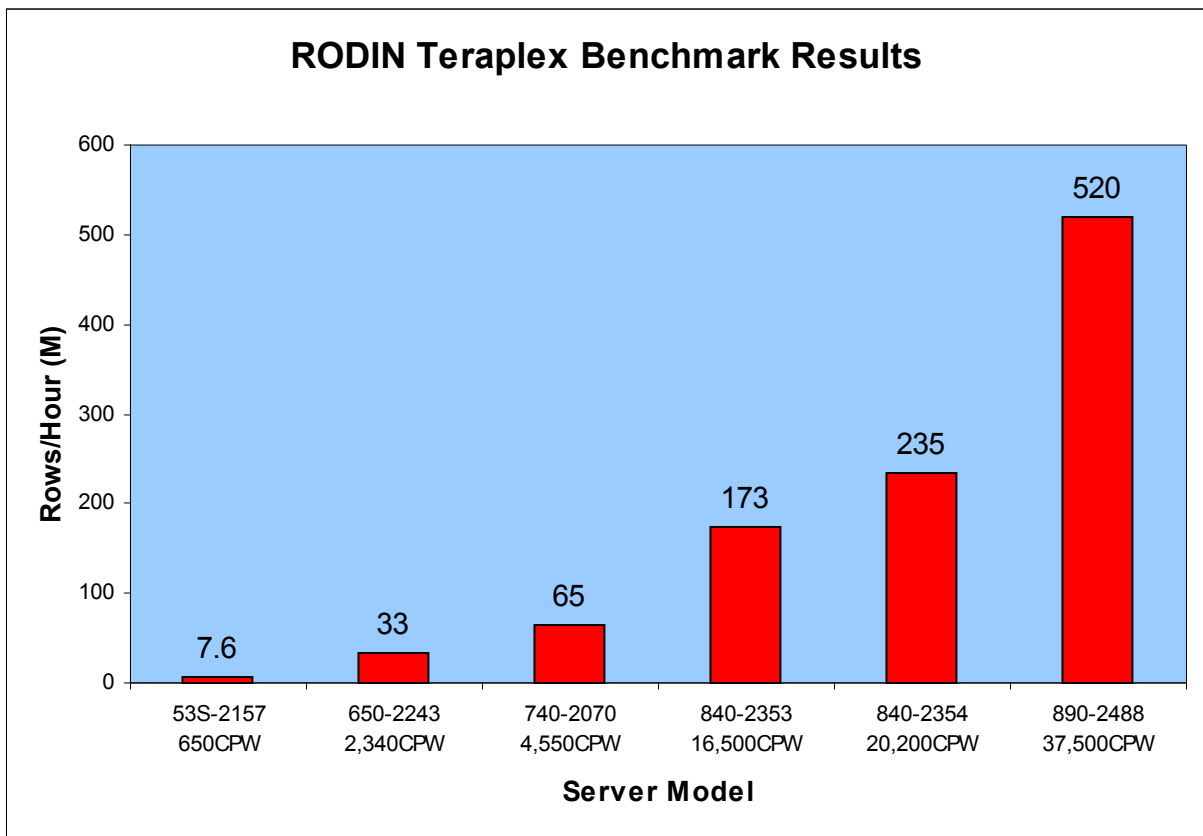
The test data sets used in these benchmarks is smaller than ideal, chosen because of time constraints. RODIN provides significant value add in the ETL process by automatically providing audit statistics and error reporting. Producing these audit and error reports at the end of the extract, as well as the similar overhead in initiating the ETL jobs, are a fixed overhead and these smaller test sets are negatively skewed by this overhead.

To compensate for this, we omitted the initialisation and termination stages from the times. This is not a significant factor in real time – for example these stages took less than 2 minutes on a 24-way parallel load, however because the load rates are so high and the data set not that large, the impact on the overall rates were noticeable.

As a comparative example of this, consider how Ferrari might demonstrate the top speed of its sports cars. If it did a timed run over a quarter of a mile from a standing start to a standing finish and calculated the average speed, the result would be unspectacular. Of course they would time the same Ferrari over a quarter mile stretch *after* it has reached top speed, and the result is the true maximum speed of the car. This is how we calculated these results.

Previous Benchmark Results

The following chart demonstrates prior benchmark results for comparison. The same benchmark scenarios were used in each of these tests (load of a 100 byte target table). Each benchmark used the latest available version of the RODIN as well as the latest IBM hardware.



This chart dramatically emphasises the huge increase in server performance in just a few years. The very first benchmark, performed in 1997 was conducted using a model 53S-2157 4-way server, which was at the time the largest AS/400 available.

Since then we have seen incredible growth – today’s i890 model is over 57 times as powerful, while RODIN’s benchmark results outperform this hardware growth, demonstrating a 68 fold increase in performance.

We expect this growth to continue well into the future.

Distributed Database Environments

Although the high end iSeries servers offer tremendous processing power and can manage up to 72TB of DASD, it is possible that very large organisations could need even greater processing power or storage capacity. In this event, up to 32 iSeries servers (each with up to 32 processors) can be linked to create a distributed database in an MPP (Massively Parallel Processing) environment. With the 32-way i890 servers, this allows for 1,024 processors and over 2 *petabytes* of disk to act as a single distributed database image.

In previous benchmarks at the Teraplex Center, RODIN was used to load a distributed target table in a 3-node MPP scenario. The results of that test indicated that this configuration performed within 90% of a single server implementation. Using these metrics, we can extrapolate these latest results to suggest load rates close to **15 billion rows per hour**, or over **5 terabytes per hour** are possible.

Realistically, it's unlikely that any organization will ever need to create distributed databases using more than 2 or 3 nodes. RODIN has been proven in a 3-node environment, and there's no technical reason why a 32-node environment would not be just as successful.

Conclusions

These benchmark results demonstrate that RODIN and the iSeries continue to hold an industry leading position in high performance, scalable ETL solutions.

Advances in technology by both Coglin Mill and IBM continue to deliver real benefit to organisations needing to implement and manage large-scale data asset management environments. We continue to stay ahead of the curve by providing proven scalability that allows customers to be confident that their business intelligence infrastructures will not be constrained by hardware or software limitations.

While other architectures force you into complex clustering technologies that require significant DBA resources to monitor database performance, the iSeries unique architecture of single level store, I/O Processors, and the latest Power PC Processors, provide a foundation without the complexity, and RODIN takes full advantage of this.

RODIN's unique parallel processing architecture, combined with its native implementation on the iSeries, superior error management capabilities and extensive functionality combine to make it the premier tool for building and managing iSeries Data Warehouse and Data Mart infrastructures.

It is important to recognize that this benchmark was not a straight 'data load' exercise where source columns are simply mapped to target columns. A complex 8 way join was done for each row, validation rules were applied to many of the source columns, numeric date columns were verified and converted to date format etc. These activities are always necessary in ETL situations and were therefore included in these benchmark scenarios, to ensure they are representative of real-life situations.

Organisations who are planning to implement a large data warehouse, or who have already implemented and are facing incredible data growth can be assured that the iSeries continues to raise the bar, and that with RODIN there is an industrial strength native ETL tool that can utilise all of the processing capacity of the hardware to load and manage enterprise data warehouse and data mart environments.

Load rates of over 500 million rows/hour, or 172GB per hour on a single server are unparalleled on any other hardware platform. These metrics exceed the requirements of most organizations, but for those with even greater need *today*, RODIN and the iSeries distributed database technology can provide the solution. If your need is not quite so immediate, just wait for the next generation of iSeries servers - they're sure to raise the bar even further.



About Coglin Mill and RODIN

Coglin Mill is a privately held software development company, established in 1986. The company has been developing System 38, AS/400 and now iSeries software since its inception, and has been a long standing Advanced Level member of IBM's Partnerworld program. The development team enjoys the sunny climes of Adelaide, South Australia, while the Northern hemisphere sales, marketing, support and advanced technology groups are located in Rochester, MN, close to IBM.

RODIN Data Asset management is the leading solution for building and deploying successful and sustainable data warehouse and data mart environments on the iSeries. Recommended by IBM's business intelligence specialists, RODIN provides superior functionality and ease of use, as well as proven ROI.

RODIN's customer base includes household names such as HSBC, Harley-Davidson, Viking Office products, Scholastic Publishing and many others.

About the Teraplex Centers

IBM's Teraplex Integration Centers are in place to facilitate the implementation of Business Intelligence and related solutions on IBM hardware and software platforms. The primary service offered by the Teraplex Centers is large-scale proofs of concept conducted at IBM facilities, either during the platform selection phase or the architecture design and validation phase.

Located in Rochester, Minnesota, the iSeries Teraplex Center is designed to investigate, research and demonstrate very large database Business Intelligence implementations and related technologies running on the IBM eServer iSeries server.

Contact Information

For more information regarding RODIN:

Coglin Mill

421 First Avenue SW
Rochester, MN 55902
USA

<http://www.coglinmill.com/>

Sales and Marketing inquiries:

Pete Wangen
Sales Director
(507) 282-4151 ext 45
pwangen@coglinmill.com

Technical inquiries (including this benchmark document):

Alan Jordan
Vice President of Development
(507) 282-4151 ext 42
ajordan@coglinmill.com

For more information regarding the IBM iSeries Teraplex Center:

Mike Cain
Team Leader – iSeries Teraplex Center
(507) 253-0256
mcain@us.ibm.com

For more information regarding iSeries:

iSeries Servers : <http://www.ibm.com/servers/eserver/series/>

DB2 UDB for iSeries : <http://www.ibm.com/servers/eserver/series/db2/>