



Jamie Barnes  
Craig Calhoun  
Mike Ebbers  
Rob Lamb  
Cindy Mestad

## High-Volume Benchmark with Genelco Group+ Proves Capacity of iSeries

What if a traditionally midrange software package could be scaled up to meet the needs of some of the largest IBM @server iSeries customers? If it could be done, the resulting cost savings would be well worth the effort.

Genelco Software Solutions teamed up with the IBM @server iSeries Benchmark Center in Rochester, Minnesota. Genelco wanted to document that the Genelco Group+ group health administration system could be scaled up to meet the needs of a large, high-volume group health insurer. A complicated and comprehensive scenario was created to emulate both the type of data to be moved, as well as the types of transactions and personnel that would be moving it.

This paper documents the resulting successes of two weeks of benchmarking tests and briefly discusses the implications for group health insurers that have not previously considered the reliability and scalability of the Genelco Group+ system and the iSeries server. The lessons learned would be applicable to any medium-to-large iSeries user.

# The performance of the iSeries: A decade of proof

For years, IBM® has impacted the business world with its OS/400®-based servers. Name an industry and you will not have to look far to see how an iSeries is proving itself as one of the most dependable, cost-effective, scalable servers in existence.

However, the iSeries impact has been minimal in the largest organizations, where mainframes are entrenched and the sheer thought of a radical server change would cause panic. Without evidence to the contrary, it is easier for CIOs and IT department heads to believe popular iSeries myths, such as, “The iSeries is a proprietary system,” or “It’s just a green screen legacy system.” The iSeries, though clearly a mainframe by certain definitions, is too small to be considered viable in these companies, because it doesn’t take up an entire room.<sup>1</sup>

At the iSeries Benchmark Center in Rochester, Minnesota, IBM has been proving the myths wrong. The Benchmark Center has spent the last decade demonstrating the stability and scalability of the iSeries in ways that could impact the bottom line of any organization, from startups who need cost-effective, powerful solutions on a low budget to the largest of organizations that need increased reliability with a lower total cost of ownership (TCO).

## Features and benefits of iSeries

iSeries is built on an integrated architecture, one that has proven its versatility in successive generations of computing. iSeries incorporates business technologies from a wide variety of open technology sources and operating environments, offering a flexible, multiplatform solution. It is uniquely positioned to play a leadership role in the on demand era.

iSeries offers the following features and benefits:

- |                    |   |
|--------------------|---|
| <b>Flexibility</b> | iSeries offers a multiplatform operating environment that can simultaneously run OS/400, Microsoft® Windows®, Linux, Lotus® Domino®, ported UNIX® applications, as well as application environments and languages such as WebSphere® and Java™.   |
| <b>Resiliency</b>  | The iSeries dynamic workload manager automatically balances system performance across all operating environments, helping to reduce administration costs and save you time, with up to 32 logical partitions per server, dynamic logical partitioning on a uni-processor, and up to 10 Linux or OS/400 partitions per processor. iSeries self-optimizing storage architecture helps automate storage with balanced application performance. |
| <b>Reliability</b> | A legendary track record of reliability, natural virus resistance, and self-managing, self-healing features means that the iSeries can help increase uptime, optimize performance, and identify and help fix problems before they occur.  |
| <b>Performance</b> | Advanced silicon-on-insulator (SOI) technology makes the POWER4™ chip one of the fastest 64-bit processors in the world, winning an award from MicroDesign Resources as a best workstation/server processor. <sup>2</sup>   |

<sup>1</sup> “What Is an iSeries?,” *iSeries News*, Frank G. Soltis, September 2003

<sup>2</sup> [http://www.mdronline.com/press/releases/jan27a\\_00.html](http://www.mdronline.com/press/releases/jan27a_00.html)

<b>Cost-efficiency</b>	iSeries new capacity on demand capabilities can help increase cost-efficiency. with IBM @server On/Off Capacity on Demand, you can turn on extra processing power when you need it, and turn it off when you don't. And with Capacity Upgrade on Demand, you can activate new permanent processing power as your business grows.
<b>Integration</b>	Integrated storage, database, and e-business tools enable you to manage your entire infrastructure centrally, including Linux and Microsoft Windows workloads. With integrated xSeries® Windows solutions, you can run your Windows servers from your iSeries, plus WebSphere products offered on iSeries allow easy modernization of traditional applications.

## Not just another benchmark

When Genelco Software Solutions decided to take its group life insurance administration solution to the IBM Benchmark Center, it had good reason. Genelco has run up against these objections before. They knew that their Group+ system was scalable and that it would handle high volumes, but they never had a reason to try it, because they were busy servicing their Tier 2 and Tier 3 insurers.

The results of benchmarking a high-volume situation would be of interest to more than just insurance companies. Any Tier 1 or 2 customer who is running an iSeries system, or would like to do so, would be extremely interested in how the iSeries performs in a highly loaded situation. This would move the scalability of the iSeries up past the midrange and into the large account arena, with no disruption to the production environment and support structure.

Genelco clients have enjoyed an unprecedented amount of stability and flexibility within their systems. As word gets around, Genelco has been starting to talk with much larger organizations about the potential use of Genelco Group+ with numbers of policies much greater than those currently being administered. Through corporate acquisitions, IT departments that have been running Genelco Group+, and those that haven't, have been coming into contact with one another and they have begun comparing notes. Questions have been coming up, such as, "Can Genelco Group+ be used as a consolidation solution?" and "If we continue to grow through acquisitions or increased market share, how much more business can Genelco Group+ handle?"

Genelco decided that it was time to find out. Their Director of Client Services, Bruce Johnson, explains: "We were confident that we could demonstrate Tier 1 volume scalability based on our experience in the field working with existing high-volume production clients and further bolstered by our continued attention toward optimized product architecture at both the application software and relational database levels. The timing was right based on market opportunity to substantiate our claims with hard results."

IBM agreed. As Patricia Weber, IBM Segment Executive for the insurance industry in the small/medium sector, put it: "IBM Small and Medium Business (SMB) Insurance Division has developed relationships with Integrated Solution Providers like Genelco to help insurance companies meet the challenges of today by helping them address their business needs more efficiently and more effectively. The IBM iSeries' proven success record is why many business partners have made it the platform of choice to leverage for their insurance solution offerings."

## Growing confidence

The engineers and software developers at Genelco had long thought that they could easily make Genelco Group+ a top-notch Tier 1 solution. There were very few foreseeable barriers to taking the same tools and making them work for a larger number of transactions with a larger number of users and inevitably higher data traffic. Genelco had confidence in the development of the software and in the iSeries. But Genelco would need a benchmark test that would put real data on the framework of possibility.

The primary objective of the benchmark effort would be to demonstrate the scalability of the iSeries-based Genelco Group+ System by emulating the high transaction volume and user load requirements of an organization much larger than any client to date. Processing must complete with very few keying errors and within industry-acceptable time constraints.

Genelco and IBM would then be able to use this data in the industry to prove that the Genelco Group+ solution has the muscle to handle the multiple stresses of massive amounts of data without a meltdown.

The pillars of meeting the primary objective were threefold:

- ▶ To create an accurate test by spending a great deal of time on its high-level approach. Genelco created an “as real as possible” data environment.
- ▶ To produce an interactive test with sub-second average response time and a low failure rate.
- ▶ To complete a batch run in eight hours (thus simulating an overnight run) using an experimental billing multithreaded workload.

## Total cost of ownership

The ability of the iSeries to handle increased workloads while providing consistent response time was a big factor in the expectations of the benchmark participants. Although a Linux or Windows server is often run at a low utilization to deliver acceptable response time even during peak loads, an iSeries server can take on the work of four to five of those servers and still deliver effective throughput. This makes the iSeries total cost of ownership (TCO) very low in practical terms.

IBM ran its own three-in-one benchmark during 2003 to verify this under-publicized iSeries capability. The results were described by Timothy Prickett Morgan in *The Four Hundred* newsletter published by the Midrange Server organization.<sup>3</sup>

## Developing the best possible benchmark scenario

In a formal planning session on March 20-21, 2003, at the iSeries Benchmark Center in Rochester, MN, it was decided to use Genelco Group+ 5.3.6. In addition, features such as Concurrent Single Claim Batch Submission/Lock Management, Expanded Billing Statement Key (to accommodate invoice print stream for up to 999,999,999 print records), and Concurrent Billing Cycle Multithread capabilities were provided.

As part of the benchmark preparation, Genelco needed to define the workload, build a database, determine the database reset procedure, establish a methodology for repeatable benchmark runs, develop success criteria, and put together a realistic benchmark schedule. They were able to complete the benchmark preparation work in six weeks.

<sup>3</sup> See “IBM Three-in-One Benchmark” on page 11 for the articles by IBM and Midrange Server.

Genelco established a benchmark database comprised of 100,000 employer groups, 6 million employees each with 10 coverages (each coverage representing a line of business such as dental or medical) and two dependents, and 2.5 years of accumulative detail premium billing history, premium payment history, and claim payment history. In addition, 230,000 submitted, but not processed, EDI claims (120,000 dental and 110,000 medical) were established to support both batch and interactive claim adjudication testing. The batch process testing focus was on invoice generation, claim check issuance, full data backup, and other production processes. Interactive processes targeted were policy administration (setup, amendment, and inquiry), enrollment, claim processing, and inquiry functions. In all, Genelco defined unique test script business transactions and associated pacing volumes for seven batch and nine interactive processes using the LoadRunner product from Mercury Interactive.

## The environment

The benchmark used the following hardware and software.

### Hardware

The benchmark was conducted using an IBM @server iSeries 890-2488 24-32 way processor configured with 128 GB of memory, 9.5 TB of disk, four 3590E tape devices fiber attached, and two GB Ethernet cards. The 9.5 TB of disk consisted of 270 35 GB 15 k RPM drives, of which 144 were configured in Auxiliary Storage Pool 1 (ASP1), 81 in ASP2, and 45 in ASP3. Only the 144 drives in ASP1 were used for jobs running the simulated production environment. The drives in ASP2 were used for backup purposes, and drives in ASP3 were used to store journal receivers. For the benchmark, the 890 system was configured to only use 12 processors and 48 GB of memory, thus simulating the production of an iSeries model 870.

### Software

The benchmark system was installed with OS/400 Version 5 Release 2 with all Licensed Program Products including WebSphere Development Studio, Backup Recovery Media Services (BRMS), Query/400, Performance Tools, Batch Journal Cache OS Option 42, DB2® Symmetric Multiprocessing OS Option 26, and SQL. The latest V5R2 cumulative PTF package and group PTFs for the database, hipers, and performance tools were also installed.

## The benchmark

Genelco provided a senior technologist project leader and two dedicated technology and application programmers for the effort. The IBM team consisted of a benchmark project leader, a technical/performance analyst, a tools/network analyst, and a hardware analyst.

### Products and tools

The benchmark used the following products and tools:

- |                                    |  |
|------------------------------------|--|
| <b>Management Central monitors</b> | Graphical monitors that show real-time information about how the system is performing such as CPU utilization, memory page faulting, and DASD utilization. |
| <b>Collection Services</b>         | Performance monitor which uses a sampling technique to capture overall performance statistics on the system.   |
| <b>iDoctor</b>                     | Detailed performance analysis tool used to further isolate where resources are being used and in-depth problem analysis.                                   |

## LoadRunner

The LoadRunner product from Mercury Interactive is a load testing product that is designed to predict system behavior and performance. It is intended to exercise an entire enterprise infrastructure by emulating thousands of users to identify and isolate problems. The LoadRunner product's integrated real-time monitors can enable organizations to help minimize test cycles, optimize performance, and accelerate deployment.<sup>4</sup>

Other benchmark assets and tools that we used are described in "Appendix A. List of benchmark assets" on page 12 and "Appendix B. Benchmark data emulation process" on page 14.

## The tests

On May 5, 2003, testing started, beginning with the interactive tests. The nine interactive user processes were assigned a total number of users and a corresponding level of usage (number of transactions per hour). See Table 1.

Table 1 Allocation of usage-to-business transactions

Business transaction	Proposed pacing values/ hour	Total number of users	Total transactions per hour	Script name	Paper script description
New Group Setup	2	20	40	POLADMNC	Policy Admin User - New Case
Amendment	6	80	480	POLADMRC	Policy Admin User - Retroactive Change
Enrollment	40	110	4,400	ENROLL	Enrollment User
Billing	3	75	225	BILL	Sample Bill (daytime reporting)
Claims	48	210	10,080	DNTCLMEP	Dental Claim for Employee
Light User	22	540	11,880	LGTSYSUS	Light System User
Policy Admin Inquiry	20	75	1,500	POLADMIN	Policy Admin User
Pended Claim	15	260	3,900	PENDCLM	Pended Claim Resolution
Single Claim Batch	500	10	5,000	BTCHCLM	Single Claim Batch

The interactive runs included an interactive scaling run. The next tests were batch cycle runs. Additional tests included a client import test, parallel saves and restores using BRMS, and a billing experimental run. The last test was a claims-only batch run that included the multithreaded experimental script.

## Measuring success

Genelco Group+ and the iSeries exceeded the expectations of all who were involved, coming in well within the criteria of the prime objectives. Interactive runs achieved sub-second user response times with an exceptionally low failure rate. Also, batch runs were completed in the allotted time frame. The summary of the individual test results follows, and more detailed data can be found in "Appendix C. Interactive run workload analysis" on page 15 and "Appendix D. Batch cycle workload analysis" on page 17.

<sup>4</sup> For more information, see their Web site: <http://www.mercuryinteractive.com/products/loadrunner/>

### ***The 100% interactive run***

For the 100% interactive run:

- ▶ The total number of users was 1,380.
- ▶ Interactive transaction response time was .08 of a second maximum.
- ▶ Over 14,000 single batch claims were processed in the one-hour run window.
- ▶ Average CPU utilization for the run was less than 70%.
- ▶ Disk IOP utilization average was less than 6%.
- ▶ Disk arm utilization on average for the run was approximately 20%.
- ▶ Machine pool faults were less than 7 per second the entire run.
- ▶ There were no failed LoadRunner transactions out of 1,962,920.
- ▶ This run was completed with journaling turned on.

See Figure 1 on page 8 for a chart of the response times.

### ***The interactive scaling run***

For the interactive scaling run:

- ▶ The total number of users for this run was 2,378.
- ▶ Interactive transaction response time was .18 of a second maximum.
- ▶ Over 10,000 single batch claims were processed in the one-hour run window.
- ▶ Average CPU utilization for the run was less than 90%.
- ▶ Disk IOP utilization average was less than 7%.
- ▶ Disk arm utilization was less than 14%.
- ▶ Machine pool faults were under 7 per second the entire run.
- ▶ There were only 21 failed LoadRunner transactions out of 4,796,513.
- ▶ This run was completed with journaling turned on.

These results far exceeded the objectives. See “Appendix C. Interactive run workload analysis” on page 15 for graphs of the runs.

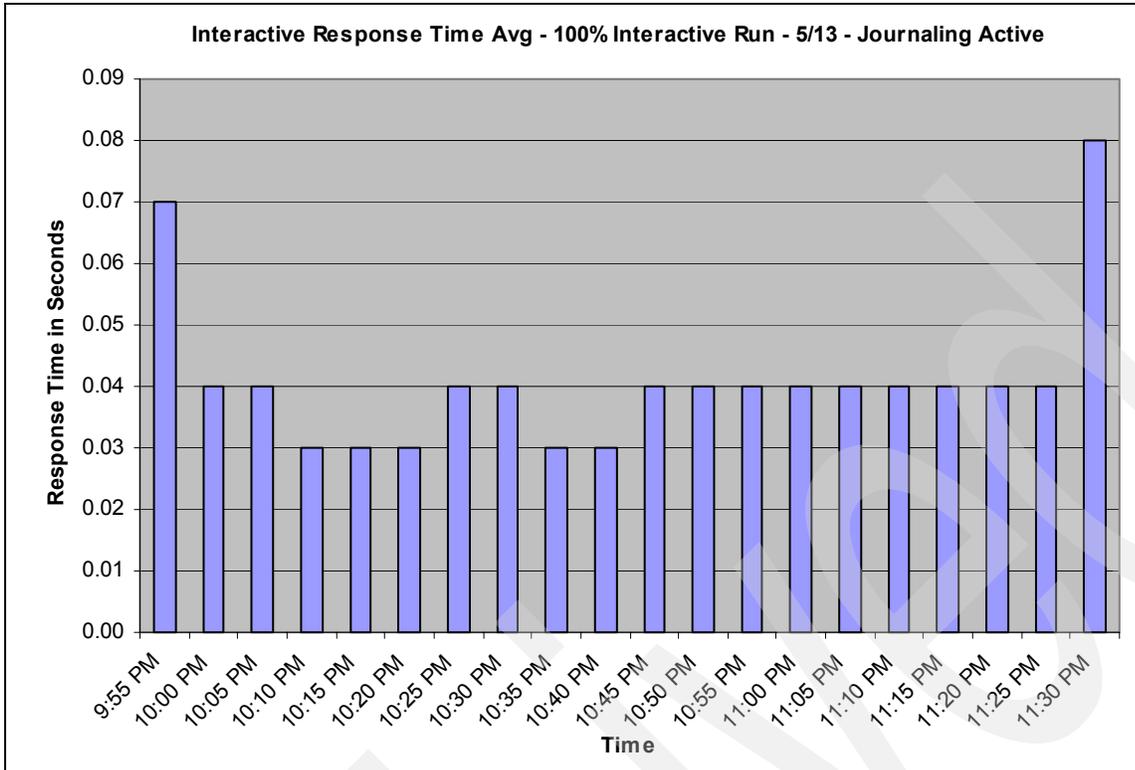


Figure 1 Interactive response time

### Batch runs

Full batch cycles were run without multithreading in order to create a baseline measurement for multithread testing.

**Note:** A single batch job can only run on one processor at a time, unless symmetric multiprocessing (SMP) is enabled. Therefore, in systems with multiple processors, a single-threaded batch job will typically result in low CPU utilization and wasted capacity. The solution is to multithread an application so that it can take full advantage of all the processors.

### The advantage of multithreading

Selecting the number of concurrent jobs is an art and not a science; trial and error is usually required. The goal is to drive CPU utilization as high as possible without over-taxing other system resources (disk drives and memory) and without introducing contention on data (locking on DB records or seizing on indices.)

When choosing to multithread a batch process, one of the challenges is to split the work evenly between multiple jobs. One way is to process the records by relative record, giving each job a unique range of relative record numbers to process. This method guarantees that the jobs are all evenly balanced and will finish at about the same time. In other situations, it is desirable to split the data logically, for example, by group, agency, or date. In this case, there can be many jobs submitted to a batch queue, each one with a parameter indicating its thread (which agency to process, for example). The job queue can be set up to allow a number of these jobs to run concurrently.

Determining the proper number of jobs to process concurrently is where trial and error comes into play. A conservative approach is to start with two jobs per processor. You would increase

from there if the CPU is still under-utilized and if there is enough memory and disk drives to avoid contention. This method of dividing the work among jobs is highly data dependent. It will work well if the data is distributed fairly evenly between agencies, or if there are a lot of agencies. For example, if we have 200 agencies and four processors, we would be running eight jobs at a time. We would hope that the workload would balance out over time so that all jobs complete at nearly the same time. If one of the agencies, for example, has 50% of the total records to process, it might be difficult to achieve this balance.

With that background in mind, here are the results of our multithreaded Genelco Group+ batch runs. See “Appendix D. Batch cycle workload analysis” on page 17 for a table of these runs.

### ***Batch cycle processing: No journaling***

For the batch cycle processing with no journaling:

- ▶ The full batch cycle runs consisted of approximately 70,000 claims and 40,000 bills for each group of 50 employees.
- ▶ Each employee had 10 coverages.
- ▶ Machine pool faults for this run were less than 2 per second.
- ▶ Disk IOP utilization was less than 2% for the entire run.
- ▶ Run time for this batch cycle was approximately 5 hours 41 minutes.

### ***Database backups***

For the database backups:

- ▶ The benchmark database was 2.18 TB with a mixture of a few very large files and several smaller files.
- ▶ The two largest files were 811 GB and 500+ GB.
- ▶ The entire 2.18 TB database was backed up to tape in three hours and four minutes. This was done using a parallel save technique in an iSeries backup software package called BRMS and four optically-attached 3590E tape drives.

### ***Additional batch runs***

Additional batch cycle runs were executed to demonstrate the ability to compress the run times. Enabling concurrent writes on the system cut the billing batch cycle run time almost in half. Additional findings were that the system could handle approximately 88 concurrent billing jobs along with two claims jobs with the concurrent write feature enabled. This pushed the CPU to approximately 80% utilization, and the billing jobs executed in approximately two hours.

**Tip:** The iSeries Benchmark Center can provide customers with remote benchmark access across high-speed Internet or using other protocols. This has proved to be an attractive convenience for customers, especially when running batch workloads.

Also, additional claim batch adjudication jobs were executed remotely (see the previous tip). The jobs were split into multiple concurrent executions and were projected to complete in approximately two hours. With these additional changes, it was projected that the 870 system with 12 processors would be able complete the full batch cycle (billing and claims) in approximately four and a half hours. With backups adding approximately three hours, that would make the entire nightly batch processing run time approximately seven and a half hours. Adding more 3590E optically attached tape drives and using them in the parallel save scenario should improve backup times significantly and further reduce the batch processing window.

Joe Rueda, IBM Business Unit Executive for iSeries ISV sales, observed: “The iSeries platform draws upon its heritage of providing scalable, reliable, and cost-efficient solutions to help businesses thrive in an on demand era. IBM has made a strategic decision to partner with the independent software vendor (ISV) community to provide total solutions that have the flexibility and capacity to meet the needs of today's dynamic businesses. IBM's relationship with Genelco Software Solutions, and this Genelco Group+ benchmark, is another example of partnering that allows customers to grow within the same product line and leverage their existing investments. This Genelco benchmark further strengthens our commitment to continue working with ISVs to deliver scalable on demand solutions to our mutual customers.”

Genelco Client Services Director Bruce Johnson was very pleased with the results: “The IBM Benchmark results provided strong validation of the Genelco Group+ System's ability to fully support Tier 1 client-level transaction volumes. We significantly exceeded success criteria measurements for both batch and interactive tasks during the benchmark effort. Additionally, we nearly doubled peak projected transaction loads during interactive scaling runs, and the application performed admirably with transaction failure rates under .0005% with average response time at .18 seconds per transaction. We can state with conviction that the Genelco Group+ System is fully capable of this level of production usage.”

## In conclusion

Genelco Group+ running on the iSeries 890 exceeded the test criteria, with a record number of transactions and an impressively low failure rate. Genelco clients, and all heavy iSeries users, can have confidence in the scalability of the system and how it will perform in larger environments. The stresses placed upon the system proved it to be very reliable. The results prove conclusively that Genelco Group+ running on an iSeries is a viable long range solution for organizations at any level. It appears to be an especially good solution for those organizations that plan to expand.

Because internal systems consolidation is also a trend, Genelco Group+ has the functionality to operate as the core system for all types of organizations that might be attempting to simplify their environment.<sup>5</sup> In an independent report on iSeries system consolidation, IDC found: “The ability of the IBM @server iSeries platform to support multiple operating system environments on a single server that can be managed by a small number of system administrators has made it an attractive consolidation platform for organizations that wish to lower their total costs while continuing to support the workloads they need to support their business.”<sup>6</sup>

Study after study finds that the IBM @server brand of servers have higher capacity, more effective optimization of the workload, greater reliability, and reduced system management costs. Because of these characteristics, iSeries machines often are chosen during server consolidations.

Sound software solutions operating on reliable hardware configurations are the backbone of the group health insurance industry and most other industries. The characteristics of stable data administration and low maintenance costs combined with a low total cost of ownership gives the iSeries a definite advantage in competitive environments.

<sup>5</sup> “Monitoring Top Technology Concerns,” *META Group Research – Analyst Insights*, Jed Rubin, September 17, 2003

<sup>6</sup> “The ROI of Windows and Linux Server Consolidation on IBM @server iSeries,” an IDC Whitepaper, Randy Perry, Jean S. Bozman, July 2003

## Additional information

For additional information, refer to the following sources:

► Genelco

<http://www.genelco.com>

1-800-983-8114

► IBM Benchmark Center

<http://www.ibm.com/eserver/series/benchmark>

► IBM Three-in-One Benchmark

<http://www.ibm.com/servers/eserver/series/consult/pdf/sageza.pdf>

<http://www.midrangeserver.com/tfh/tfh061603-story02.html>

<http://www.ibm.com/servers/eserver/series/hardware/threeinone/>

# Appendixes

This sections contains the following appendixes:

- ▶ Appendix A. List of benchmark assets
- ▶ Appendix B. Benchmark data emulation process
- ▶ Appendix C. Interactive run workload analysis
- ▶ Appendix D. Batch cycle workload analysis

## Appendix A. List of benchmark assets

The benchmark assets include:

- ▶ Core Data Environment
  - Single Monthly Billed Group
  - Policy Administration Definition
  - Claims Benefit Definition
  - Enrollment
  - Historical Billing and Payment History Accumulation
  - Historical Claim History Accumulation
  - Single Claim Unadjudicated Batches
  - Multiple Claim Unadjudicated Batches
- ▶ Data Replication Utility
  - Replicates Core Group Setup X 100,000
  - Concurrent Multithread Execution
- ▶ Volume of Data
  - 100,000 (100 k) Active Groups
  - 5,000,000 (5M) Active Employees
  - 1 M Terminated Employees
  - 10 Coverages per Employee (for both active and terminated)
  - 10 M Active Dependents
  - 2 M Terminated Dependents
  - 80 k Data Entered and Unadjudicated Dental Claims (single batch claim submissions for interactive test)
  - 80 k Data Entered and Unadjudicated Medical Claims (single batch claim submission for interactive test)
  - 40 k Data Entered and Unadjudicated Dental Claims (accumulated to one batch for batch test)
  - 30 k Data Entered and Unadjudicated Medical Claims (accumulated to one batch for batch test)
  - 2.5 Years of Detail Billing and Payment History (10 coverages per group/month/employee, with historical adjustment accumulation resulting from member termination and coverage changes)
  - 2.5 Years of Accumulative Claim History (300,000,000 dental claims and 300,000,000 medical claims with each claim having three charge lines)

- ▶ Variability of Data
  - Group Number Randomly Generated in Range of 1 – 100,000
  - Employee Number (SSN) Randomly Generated in Range of 1 – 60 per Group
  - Randomized Employee and Dependent Birth Dates
  - Randomized Employee and Dependent Names
  - Claim Numbers Dynamically Generated Based on Date of Service
- ▶ Mercury Interactive LoadRunner Scripts for User Types
  - Policy Admin. New Group Setup – 20 Concurrent Users
  - Policy Admin. Amendment – 80 Concurrent Users
  - Policy Admin. Inquiry – 75 Concurrent Users
  - Manual Enrollment – 110 Concurrent Users
  - Claims Data Entry – 209 Concurrent Users
  - Claims Pend Resolution – 260 Concurrent Users
  - Light System User – 940 Concurrent Users
  - Daytime Reporting/Billing – 75 Concurrent Users
  - Single Batch Claim Submission – 5,000 Claims/Hour
- ▶ Hardware (system, processors, memory, DASD)
  - IBM @server iSeries™ 890-2488 32-way processor configured to only use 12 processors during benchmark runs
  - 128 GB of memory, system configured to use only 48 GB of memory during benchmark runs
  - 9.5 TB disk (270 35 GB 15 k RPM drives attached through 2757 IOPs; only 144 drives used for benchmark runs)
  - Four 3590E tape drives connected by fiber attachment to iSeries
  - Two 1 GB Ethernet cards
- ▶ Group+ 5.3.6 Standard Release
  - Administration Component
  - Claims Component
  - Integrated Enrollment
  - Batch Claim Adjudication
- ▶ Group+ Software Customizations
  - Concurrent Single Claim Batch Submission/Lock Management
  - Expanded Billing Statement Key
  - Concurrent Billing Cycle Multithread
  - Concurrent Client Import Multithread (CARS990)

## Appendix B. Benchmark data emulation process

Genelco designed, tested, and executed a utility to perform the data replication needed to establish the core data environment for the benchmark effort. The utility executed a random number generator from a third-party tool called TAATOOL to randomly assign birth dates to employee and dependent records. Full data replication occurred at the IBM Benchmark Center as one of the preliminary setup tasks for the benchmark effort. The core data environment was saved to file on disk and restored prior to initiation of each of the benchmark tests.

Genelco identified and refined test script business transactions and associated volumes for nine independent interactive processes. Manual (paper) interactive scripts were documented and tested by Genelco. Each script was designed to support the variability of the data in the core environment. Also, the scripts were designed to minimize edit intervention in order to place the full desired interactive load during the benchmark effort.

The nine processes are:

- ▶ Policy Admin. New Group Setup – Create new group setup (billing and claims) from template data.
- ▶ Policy Admin. Amendment – Terminate and revise coverages at the group level with mass update impact to enrollment.
- ▶ Policy Admin. Inquiry – Inquiry about billing and payment status for a group.
- ▶ Manual Enrollment – High-volume oriented new enrollment with two dependents and 10 coverages per enrollment.
- ▶ Claims Data Entry – High-volume oriented online claim entry and adjudication with three charge lines per claim.
- ▶ Claims Pend Resolution – High-volume oriented online claim entry, pend, and subsequent reprocess to paid status.
- ▶ Light System User – High-volume oriented call logging, eligibility inquiry, and claim benefit inquiry.
- ▶ Daytime Reporting/Billing – Daytime submission of invoice generation for trial (projected) invoices.
- ▶ Single Batch Claim Submission – High-volume oriented claim submission using electronic single claim batches.

The LoadRunner product from Mercury Interactive was used to define and execute the test scripts for the benchmark effort. Genelco tested and verified the test scripts against a replicated data environment on the Genelco iSeries prior to the formal benchmark effort. The number of users and pacing values per hour per script were discussed and refined. The next two sections identify the final model used during the benchmark effort.

## Appendix C. Interactive run workload analysis

This appendix describes the 100% interactive run with journaling active.

The total number of users for this run was 1,380. The interactive transaction response time was a maximum of .08 of a second, average CPU utilization for the run was under 70%, disk IOP utilization average was under 6%, and machine pool faults were less than 7 per second the entire run. There were more than 14,000 single batch claims processed in the one-hour run window. There were no failed LoadRunner transactions out of 1,962,920. See Table 2 and the following figures.

Table 2 Interactive run results

Business transaction	Pacing value/hour	Number of users	Total transactions/hour
New Group Setup	2	20	40
Amendment	6	80	480
Enrollment	40	110	4,400
Billing	3	75	225
Claims	48	210	10,080
Light System User	22	540	11,880
Policy Admin Inquiry	20	75	1,500
Pended Claim	15	260	3,900
Single Claim Batch	500	10	5,000

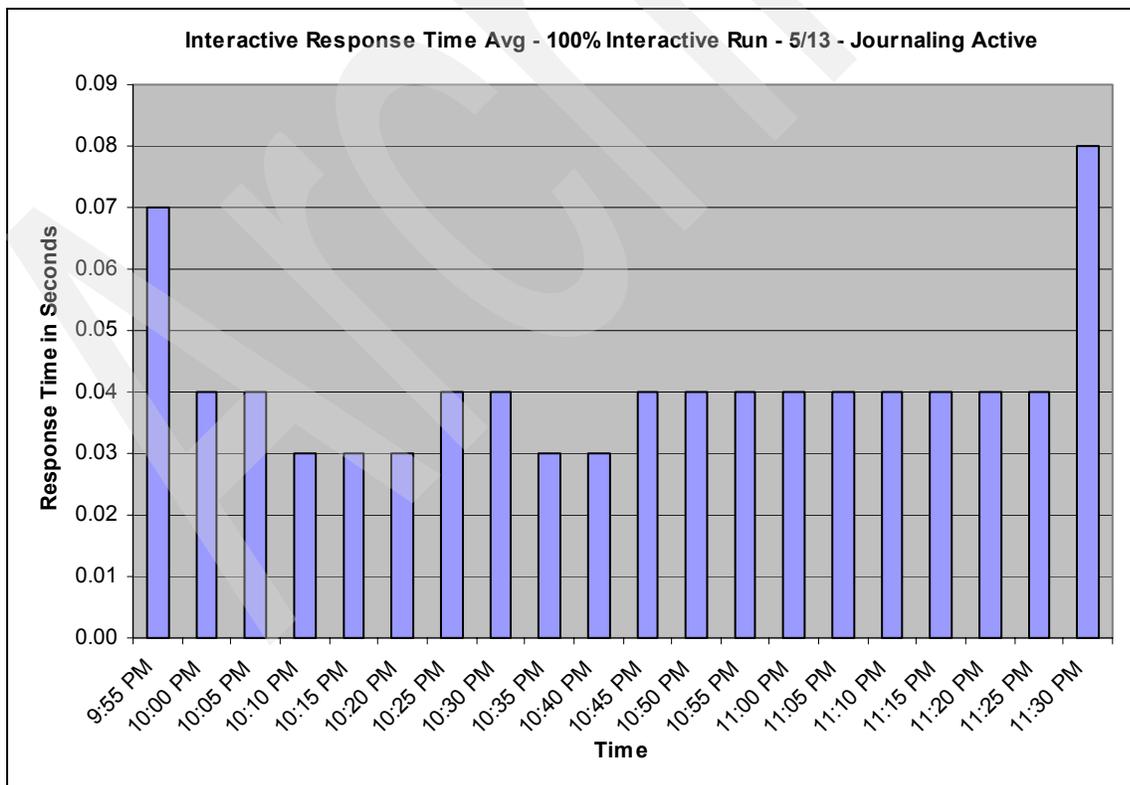


Figure 2 Interactive response time

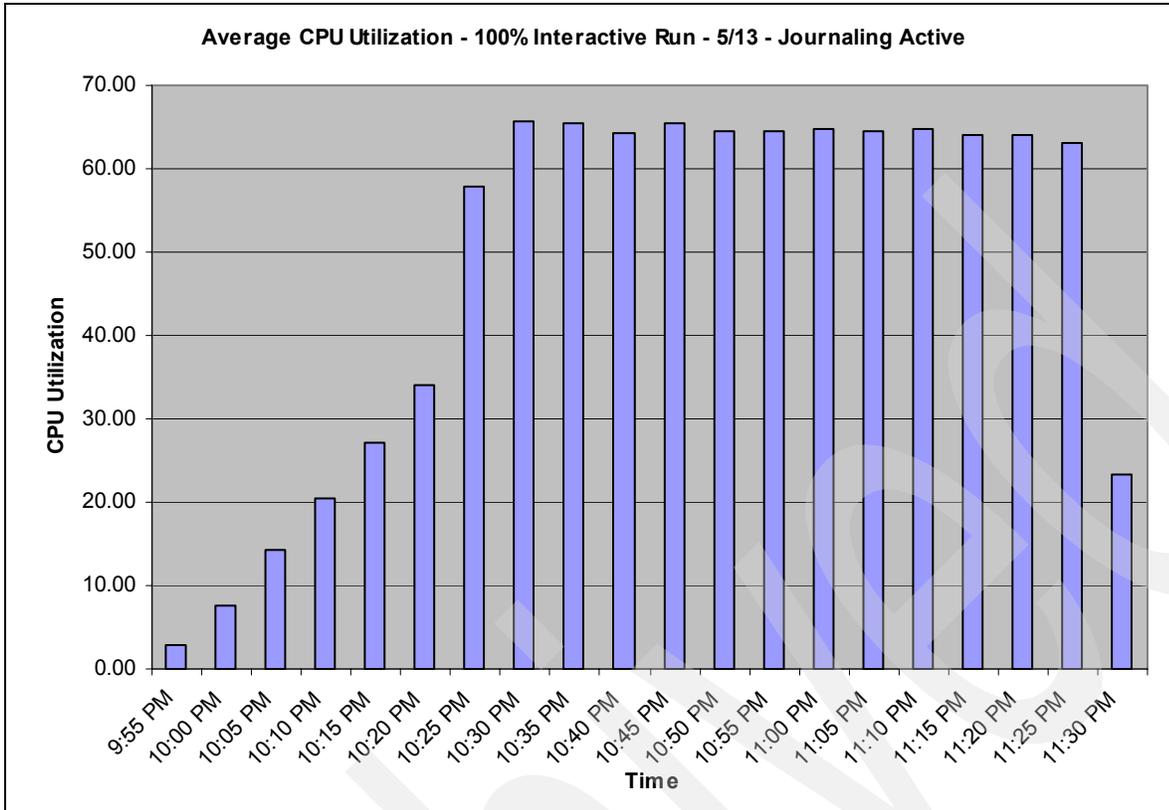


Figure 3 CPU utilization: Interactive run

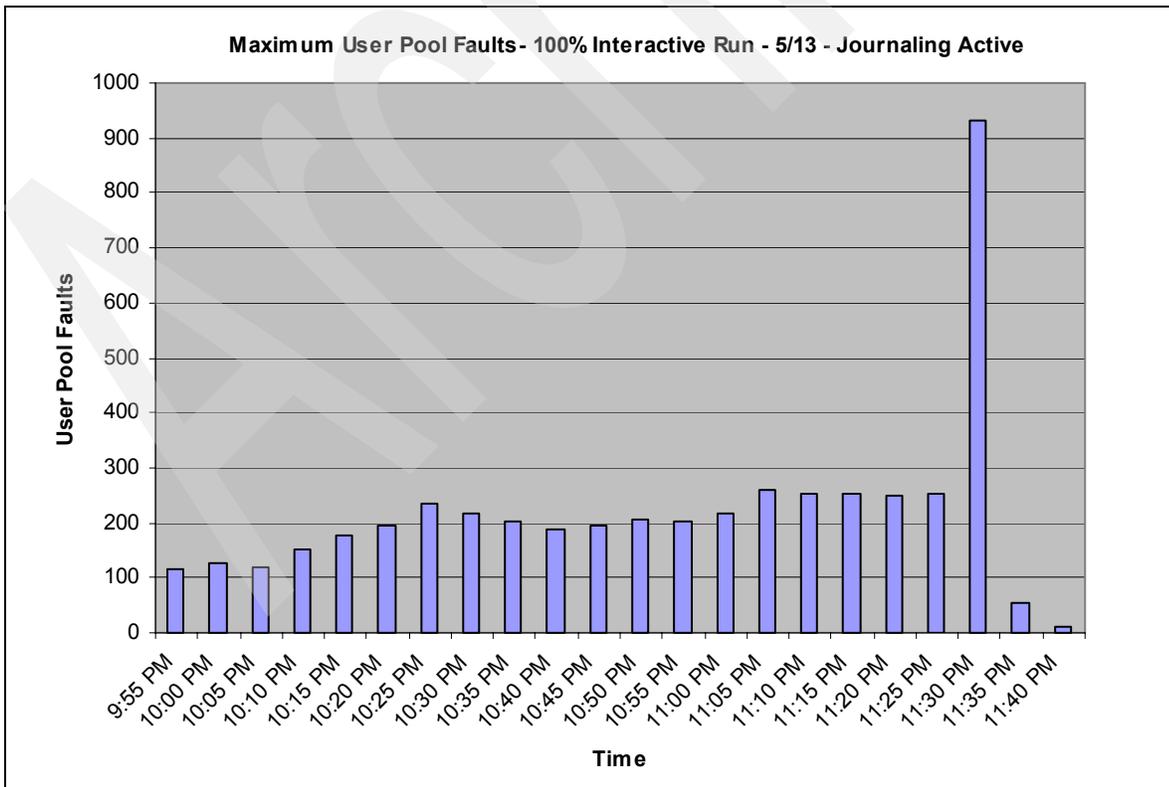


Figure 4 Maximum user pool faults: Interactive run

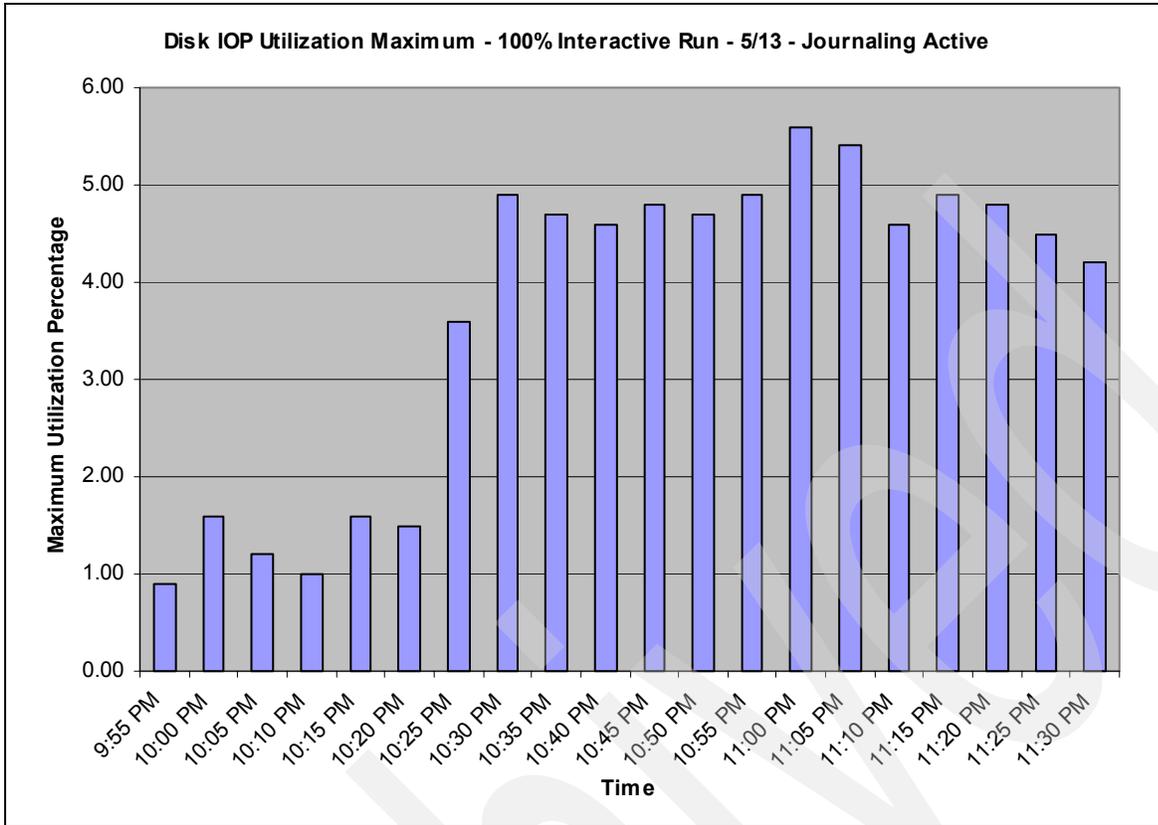


Figure 5 Disk IOP utilization maximum: Interactive run

## Appendix D. Batch cycle workload analysis

This appendix describes the batch cycle processing of May 14 with journaling active.

Journaling was activated for all non-work database files in the Group+ application.

Table 3 Batch run results

Job	Start time	End time	Duration
Dental Batch	0:00:16	4:47:10	4:46:54
Medical Batch	0:00:25	2:09:34	2:09:09
Check Cycle	5:23:30	6:27:20	1:03:50
Update Cycle	6:30:58	6:46:10	0:15:12
<b>Total claims processing</b>	<b>00:00:16</b>	<b>6:46:10</b>	<b>6:46:06</b>
Billing Job 1	0:17:29	4:33:57	4:16:28
Billing Job 2	0:17:29	4:25:08	4:07:39
Billing Job 3	0:17:29	4:37:02	4:19:33
Billing Job 4	0:17:29	4:30:03	4:12:34
Billing Job 5	0:17:29	4:27:49	4:10:20
Billing Job 6	0:17:29	4:37:22	4:19:53

Billing Job 7	0:17:29	4:34:51	4:17:22
Billing Job 8	0:17:29	4:37:12	4:19:43
Billing Job 9	0:17:29	4:35:24	4:17:55
Billing Job 10	0:17:29	4:30:20	4:12:51
Billing Job 11	0:17:29	4:34:00	4:16:31
Billing Job 12	0:17:29	4:23:35	4:06:06
Billing Job 13	0:17:29	4:34:19	4:16:50
Billing Job 14	0:17:29	4:24:36	4:07:07
Billing Job 15	0:17:29	4:36:46	4:19:17
Billing Job 16	0:17:29	4:29:07	4:11:38
Billing Job 17	0:17:29	4:23:36	4:06:07
Billing Job 18	0:17:29	4:36:06	4:18:37
Billing Job 19	0:17:29	4:25:50	4:08:21
Billing Job 20	0:17:29	4:37:26	4:19:57
Billing Job 21	0:17:29	4:37:26	4:19:57
Billing Job 22	0:17:29	4:21:19	4:03:50
<b>Total billing jobs</b>	<b>00:17:29</b>	<b>4:37:26</b>	<b>4:19:03</b>

Journaling only seemed to add about 19% to the run times for long running batch jobs. The High Availability Journal Performance (Journal Caching) was used for journaling. The Journal Caching feature allows batch applications to substantially reduce the number of write operations for which the job must wait, thereby reducing overall elapsed batch execution time.

## The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the IBM @server iSeries Benchmark Center in Rochester, Minnesota.

### Authors

**Jamie Barnes** is a Lead Technology Analyst at Genelco Software Solutions. He has 14 years of experience in the IT field, all with IBM midrange servers. He holds a Bachelor of Science degree in Computer Information Systems from Freed-Hardeman University. His areas of expertise include IBM AS/400® and iSeries system administration with a special emphasis on performance and logical partitioning.

**Craig Calhoun** is an iSeries ISV Business Development Specialist. He is part of the Americas Sales and Distribution organization and has been employed with IBM for eight years. For the past four years, Craig has held various positions that leverage the ISV community in order to help IBM grow their business. He is an IBM @server Certified Specialist for iSeries Solution Sales V5R2.

**Rob Lamb** is a Senior Marketing Coordinator at Genelco Software Solutions. He has 11 years experience in a variety of marketing roles: writer, publicist, brand manager, and direct marketing manager. For the past year, he has been working in Genelco's marketing department writing publicity materials, as well as providing research support to Genelco's sales staff.

**Cindy Mestad** is a Project Manager at the iSeries Benchmark Center in Rochester, MN. She has been with IBM for 23 years, working in various areas of the business such as S/38 Computer Operations, S/38 Vertical Microcode Testing & Development, S/38 & AS/400 System Test, AS/400 Early Programs, AS/400 Systems Performance, and most recently, the iSeries Benchmark Center. She is a certified Project Management Professional.

**Mike Ebbers** is a Consulting IT Specialist at the International Technical Support Organization, Poughkeepsie Center. He has been with IBM for 30 years: 11 in technical marketing, 10 in technical education, and 9 in the ITSO. He writes about benchmarks and leads redbook projects on IBM and Lotus products.

## Reviewers

Bruce Johnson and Susan Wagy, Genelco Software Solutions  
Helen Olson-Williams, IBM Corporation

Archived

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. 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 IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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 Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an Internet note to:  
[redbook@us.ibm.com](mailto:redbook@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, International Technical Support Organization  
Dept. HYJ Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400 U.S.A.



## Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AS/400®  
DB2®  
Domino®  
@server™  
@server™  
eServer™

IBM®  
ibm.com®  
iSeries™  
Lotus®  
OS/400®  
POWER4™

Redbooks(logo) ™  
Redbooks™  
WebSphere®  
xSeries®

The following terms are trademarks of other companies:

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.