



HIRO Database Benchmark

Sybase vs. Informix

July 18, 1994

This document and its contents are the property of Holiday Inn Worldwide, Inc.; it shall not be reproduced, copied, or disseminated without the written consent of Holiday Inn Worldwide, Inc.

Table of Contents

1.0	Background	1
1.1	Benchmark Criteria	1
1.2	Benchmark Test	1
2.0	Architecture	2
2.1	Parallel Processing/Scalability	2
2.2	Concurrent Access	3
2.3	OnLine Tuning	3
2.4	Transaction Logs	4
2.5	Archive Restore	4
2.6	Architecture Conclusion	4
3.0	Vendor Support	5
4.0	Performance	6
5.0	Recommendation	6
6.0	Database Creation	7
6.1	Database Creation Results	7
6.2	Database Creation Conclusion	8
7.0	Property Loads	9
7.1	Property Loads Results	9
7.2	Property Loads Conclusion	10
8.0	Drop and Create Index	11
8.1	Drop and Create Index Results	11
8.2	Drop and Create Index Conclusion	12
9.0	Database Archive and Restore	13
9.1	Database Archive and Restore Results	13
9.2	Database Archive and Restore Conclusion	14
10.0	Miscellaneous Queries (49 Property Database)	15
10.1	Miscellaneous Queries Results (49 Property Database)	15
10.2	Miscellaneous Queries Conclusion (49 Property Database)	16
11.0	Miscellaneous Queries (392 Property Database)	17
11.1	Miscellaneous Queries Results (392 Property Database)	17
11.2	Miscellaneous Queries Conclusions (392 Property Database)	18
12.0	Concurrent Property Loads	19
12.1	Concurrent Property Loads Results	19
12.2	Concurrent Property Loads Conclusion	20
13.0	Multi-User Queries	21
13.1	Multi-User Queries Results	21
13.2	Multi-User Queries Conclusions	22
14.0	Altered Multi-User Queries	23
14.1	Altered Multi-User Queries Results	23
14.2	Altered Multi-User Queries Conclusion	24

Executive Overview

1.0 Background

Due to significant performance problems with the HIRO application, a benchmark was performed between Informix 6.0 and Sybase 10.01 Relational Database Management Systems (RDBMS). In the current Sybase implementation, the HIRO application has an upper limit of approximately 50 hotel properties per database instance. Above this limit, performance suffers to unacceptable levels.

In order to support the proposed roll out of 2,000 HIRO properties, the current level of performance would require 40 database instances and a considerable investment in additional hardware and database administration expense. It is imperative that HIW determine what corrective action can be taken to reduce the number of database instances and establish an acceptable level of performance. Our goal is to have at a minimum 2 and at a maximum 4 database instances.

1.1 Benchmark Criteria

The focus of the benchmark was the demonstration of scalable performance in a Symmetrical Multi-Processing (SMP) environment. This included a detailed evaluation of both the current and long term strategy of the DBMS vendor architecture.

1.2 Benchmark Test

A benchmark test was constructed which stressed all aspects of a production database. Database creation, index drop and rebuilds, database loads and single and multi-user queries were developed to evaluate each vendors performance on several data sets varying in size and number of concurrent users.

The tests were run on a Sun SparcCenter 1000 with eight processors and 1 GB of memory running the Solaris 2.2 OS with patches. The database was stored on 5 disk drives: three 2.9 GB drives on single controller and two 2.1 GB drives on another separate controller. Sybase and Informix were allowed to modify any Solaris parameters and any of their own database tuning parameters. The client/server tests were performed across a dedicated 10Mb ethernet segment.

Tests were conducted by representatives of each participant and validated by HIW staff. The results were tabulated and presented in this document.

2.0 Architecture

The platform for HIRO is a Symmetrical Multi-Processor (SMP) Sun SparcCenter configuration. Each vendors ability to scale and utilize the resources of this SMP technology is crucial to preserve the investment of HIW. Each vendors architecture was evaluated in the following areas:

- * Parallel Processing/Scalability
- * Concurrent Access Strategy
- * Online Database Administration Support
- * Transaction Logging Strategy
- * Archive/Restore Strategy

2.1 Parallel Processing/Scalability

Informix Symmetric Multi-Processor (SMP) tuned product available today. Included as part of base product.

Sybase No SMP tuned products. Promises Massively Parallel Processor (MPP) optimizer that will work on SMP machine. Separate added cost of \$400,000 for entry level Navigation Server and scale into the millions of dollars as processor amounts grow.

The META Group, as well as other analysts, believe that commercial adoption of MPP technology is at least 3 years away. Since Sybase has chosen to implement the Navigation Server, an optimizer purchased from NCR and designed for MPP machines, this leaves Sybase users without a product that is optimized for the cost-effective, as well as scalable computing environment already available with SMP. The HIRO project has already made over a \$1,000,000 commitment to the SMP architecture.

The problem today is there is no Sybase product available to provide the parallelism required for HIRO. In actual experience, the HIRO project has found that when a database exceeds 1 gigabyte, performance drops dramatically. When multiple processors are added, there is a severe system drain even when there is little database activity. Sybase does not appear to use system resources effectively.

2.1 Parallel Processing/Scalability (continued)

According to the Gartner Group, "the base [Sybase] engine has problems with scaling over four processors and we have found, even at four, a disappointing 50% scalability". In the July 18, 1994 issue of ComputerWorld, it is reported that Sybase has opted to delay support of Navigation Server and System 10. A ComputerWorld quote follows: "Sybase has been touting System 10 as the database for client/server architectures, but now a key product - Navigation Server - will not support it".

The result of this lack of parallelism in Sybase is poor scalability. This can be seen in the benchmark results. When a 392 property database was created, Sybase could not run 100 concurrent users when presented a standard set of insert, delete and update statements. Also, Informix could load 8 properties concurrently in less time than Sybase could load one property. For a more detailed explanation, see the benchmark results section of this report.

2.2 Concurrent Access

Informix Row level locking

Sybase Page level locking

Providing users with concurrent access to data is extremely critical in an OLTP environment. When considering the effects upon a system having hundreds of concurrent users and gigabytes of data, "lock granularity" can substantially affect productivity.

This is a fundamental flaw in the Sybase product. This means with Sybase in an OLTP environment where hundreds of users are manipulating data, users will wait for data because entire pages of data may be locked. Several different HIRO hotels data could be contained on a single page, forcing a deadlock as jobs wait for time critical optimizations. In order to eliminate current deadlocking in the current Sybase production environment, records needed to be padded to a page size to "simulate" row level locking, wasting disk space and increasing administrative costs. With Informix's support for row level locking, concurrency is increased since only those records requested are locked, not entire pages.

2.3 OnLine Tuning

Informix Tune engine while on-line.

Sybase Must take down engine to tune.

The ability to perform database administration and tuning activities without taking down the database is vital for a mission-critical environment. Informix allows on-line tuning of

2.3 OnLine Tuning (continued)

the database engine including adding shared memory and modifying the number of CPU virtual processors. In the HIRO project, this will be very useful during the Nightly processing. With fewer users and optimizations, more system resources could be added to the database dynamically to speed nightly processing. Sybase experiences a degradation in performance with multiple CPU's during "idle periods" of access. Currently with Sybase, performance tradeoffs are required to balance interactive and batch process requirements.

2.4 Transaction Logs

Informix Multiple logs

Sybase Single log

A current problem with Sybase and the HIRO system has been the transactions logs filling and processing halting until the log files have been dumped. Many man hours have been dedicated to work around the Sybase limitation of one log. With Informix, multiple logs could be implemented eliminating the fear of a single log filling. Older logs can also be backed up to tape off-line and cleaned before reuse. The number of logs can also be adjusted dynamically with Informix.

2.5 Archive/Restore

Informix If media failure occurs, only the failed partition needs to be recovered. Application remains online during recovery.

Sybase If media failure occurs, full database recovery is required. The application must be taken offline during recovery.

In the event of a disastrous media failure, this is an important point to consider. Based on the projected size of 50GB of RAID 5 storage for the HIRO database and the results of these tests, this could mean a difference of hours to days to restore HIRO.

2.6 Architecture Conclusion

Informix has demonstrated their architecture features available today. Industry experts along with Sybase executives admit delays and problems with Sybase's architecture. With the given facts, the HIRO Information Technology Management is more comfortable and confident with the current and future direction of Informix's architecture and Informix's ability to meet the HIW technical and business requirements of our strategic mission objectives.

2.6 Architecture Conclusion (continued)

Additionally, a beta version Informix 7.0 was tested to evaluate the direction of the Informix architecture. Version 7.0 contains the Parallel Data Query (PDQ) capability which maximizes database throughput by breaking the queries into parallel tasks, thus leveraging to the fullest extent the SMP technology. For the queries which were run with version 7.0, a dramatic improvement of 300% to almost 900% was observed. Clearly, Informix has a better strategy to leverage SMP technology today. Moreover, Informix's PDQ capability is built into the core or kernel of the database engine, providing a true parallel processing architecture. Sybase plans to utilize a bolt-on approach by adding an additional layer, Navigation Server, in order to obtain parallel processing capability.

3.0 Vendor Support

Since late last year, Sybase has been aware of our performance issues. HIW has paid for several Sybase engineers to review our physical and logical designs and to observe the HIRO system under user loads. Each time we were told to buy more hardware or rewrite the application. For \$1,500 per day per person, this was not the advice we were seeking. We were also told that the problem was with Solaris not supporting asynchronous I/O. Never did Sybase confirm that there may be a problem with their system even though there are many industry articles articulating this fact as well as our experience with their product.

Informix supplied three people full time for six weeks to participate in the benchmark effort. Because the application was written for Sybase, Informix had to port all benchmark tests from Sybase to Informix. Informix completed all tests within the original time requested. All Informix time, travel and expenses were performed free of charge.

Sybase supplied one part time SE until a letter of concern was sent by HIW to Sybase two weeks before the benchmark was to be completed. Sybase then added another part time SE. Although Sybase was the incumbent and started the benchmark one week before Informix, Sybase did not complete over half the tests. Even though all of Sybase's time, travel and expenses were performed free of charge, Chris Johnston, Regional Technical Manager, demanded that HIW sign a consulting agreement for \$1,500 per day per person for a team of Sybase consultants to rewrite the HIRO application to perform within the limitations of Sybase.

4.0 Performance

Due to Informix's inherent multi-threaded architecture, significant performance gains were achieved over Sybase. An average of 130% to 38,571% improvement in performance was obtained during the benchmark when using Informix. Additionally, Informix's database engine is continually selected by hardware vendors to showcase their hardware capabilities. The Transaction Processing Council (TPC) benchmark for OLTP and Decision Support Systems, TPC-C, has 18 independently published results with Informix. Sybase has 1 TPC-C benchmark for a uniprocessor box.

5.0 Recommendation


In each of the identified criteria: performance, architecture and support, Informix was clearly superior to Sybase. Benchmark results in the subsequent sections document Sybase's substantial problems performing this benchmark. Sybase management also questioned the motives of the HIRO staff and the benchmark tests themselves. The goal was simply to determine the most cost-effective solution with guaranteed benefits for HIW and this important component of the WwHS strategic initiative. Based on the results of this benchmark test, Informix is recommended as the preferred vendor for HIRO. The justification is contained in the following sections.

6.0 Database Creation

Create a database that is large enough to contain 400 hotel properties.

6.1 Database Creation Results

<i>Timings in Minutes</i>	
	DB Create
Informix 6.0	0.12
Sybase 10.01	45.00
Informix 6.0 Faster	38571%



6.2 Database Creation Conclusion

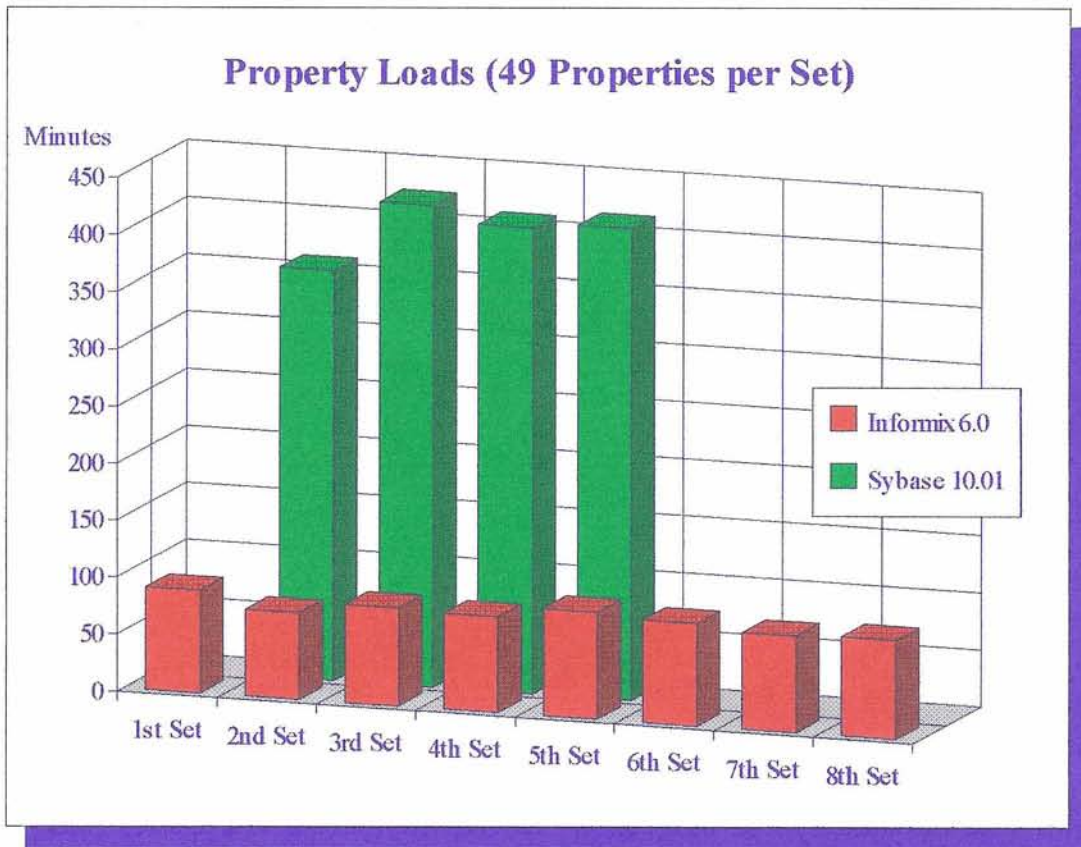
Informix was 38,571 percent faster. After talking with Sybase, the speed difference comes from the need of Sybase to initialize all database space. Informix does not need to initialize the space. In the event of a major failure, database creation speed is important for recovery.

7.0 Property Loads

Load a set of 49 properties into the database. Continue to add 49 properties for a total of 392 properties. A single property consists of about 20 MB of data; a set of 49 properties is about 1 GB.

7.1 Property Loads Results

<i>Timings in Minutes</i>								
	1st Set	2nd Set	3rd Set	4th Set	5th Set	6th Set	7th Set	8th Set
Informix 6.0	90	76	86	84	94	90	85	87
Sybase 10.01		359	422	408	413			
Informix 6.0 Faster		472%	491%	486%	439%			



7.2 Property Loads Conclusion

Informix was over 400 percent faster in sets two through five. When loading several properties in the database, speed is important to minimize downtime for a property. During the loading process, a property is unable to use the HIRO system.


For the first set of 49 properties, no time was recorded for Sybase because it was run before a Sybase SE tuned the database. Sets 6 through 8 times were rejected because Sybase loaded the database in an unsupported mode. This mode is not in the documentation and could jeopardize the consistency of the data and violate our support agreement.

8.0 Drop and Create Index

Drop and create an index for a 49 and 392 property database. The table selected was the HIRO pral table which contained over 3,000,000 and 26,000,000 rows respectively. The index contained five fields in the key.

8.1 Drop and Create Index Results

<i>Timings in Minutes</i>		
	49 Properties	392 Properties
Informix 6.0	7	58
Sybase 10.01		115
Informix 6.0 Faster		198%



8.2 Drop and Create Index Conclusion


Informix was almost 200 percent faster on the 392 property database. From 49 to 392 properties, Informix scaled almost linearly. Index creation speed is important in a production environment if an index should get corrupted and need to be recreated. Sybase did not attempt the 49 property test due to time.

9.0 Database Archive and Restore

Perform an archive and restore on a 49 property database. The test was performed using a single tape drive.

9.1 Database Archive and Restore Results

<i>Timings in Minutes</i>			
	Archive (No Logging)	Archive (Logging)	Restore
Informix 6.0	29	68	57
Sybase 10.01			



9.2 Database Archive and Restore Conclusion

Informix completed the test with and without database logging. Being able to restore a large database in under an hour is very fast compared with the restoring times experienced today. Today using Sybase it takes 45 minutes just to create the database without data.


Sybase did not complete this test due to time even though Sybase insisted that this test be added to the benchmark.

10.0 Miscellaneous Queries (49 Property Database)

Seven complicated queries were selected from the HIRO system. The first test was to run the queries on a 49 property database.

10.1 Miscellaneous Queries Results (49 Property Database)

<i>Timings in Minutes</i>							
	Query 1	Query 2	Query 3	Query 4	Query 5	Query 6	Query 7
Informix 6.0	0.3	0.1	6.2	12.9	109.2	18.5	0.0
Informix 7.0	0.1		0.7	9.9	30.1	1.1	
Sybase 10.01							
Informix 7.0 Faster	375%		886%	130%	363%	1760%	



10.2 Miscellaneous Queries Conclusion (49 Property Database)

Performance of Informix 6.0 was within expectations; performance of Informix 7.0 was outstanding. With the performance gains achieved in Informix 7.0 much less hardware will be required to run the HIRO system.


Sybase did not run this test due to time. Informix was not required to run the queries using Informix 7.0. Informix flew two people to their benchmark center in California to perform the tests on Informix 7.0.

11.0 Miscellaneous Queries (392 Property Database)

The same seven complicated queries that were run on the 49 property database were also run on the 392 property database.

11.1 Miscellaneous Queries Results (392 Property Database)

<i>Timings in Minutes</i>							
	Query 1	Query 2	Query 3	Query 4	Query 5	Query 6	Query 7
Informix 6.0	6.4	8.7	51.5	109.8		72.2	16.2
Informix 7.0	0.2		5.8			7.7	
Sybase 10.01	0.6					30.0	
Informix 7.0 Faster	271%		894%			389%	



11.2 Miscellaneous Queries Conclusions (392 Property Database)

On five of the seven queries, Informix beat Sybase. On two of the seven queries, Sybase beat Informix 6.0 but did not beat Informix 7.0. This validates the fact that Sybase works well for some queries. However, other run times were not captured because the queries took too long to run. Decision support systems, like HIRO, run many complicated queries. If complicated queries are not completed quickly, other processes will suffer a degradation in performance.

Sybase only posted times for queries one and six. Sybase could not post times on the other queries due to run times. The only way Sybase could run queries two, three, four, five and seven was to add indexes or join conditions thus violating the benchmark criteria. Holiday Inn did not accept these times because the benchmark required the queries to be run without modifications.

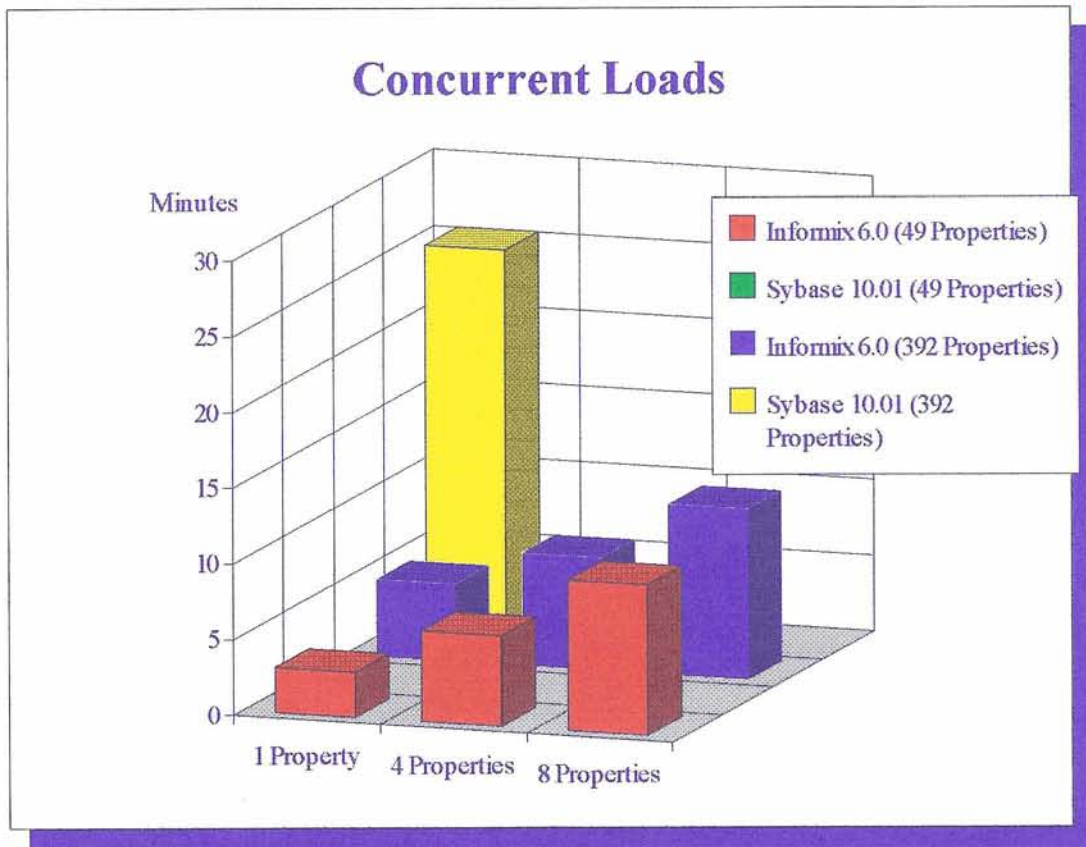
Informix was not required to run the queries using Informix 7.0. Informix flew two people to their benchmark center in California to perform the tests on Informix 7.0.

12.0 Concurrent Property Loads

Load one, four and eight properties concurrently on a 49 and 392 property database. A property consists of about 20 megabytes of data that is contained in several flat files; one flat file per table.

12.1 Concurrent Property Loads Results

<i>Timings in Minutes</i>			
	1 Property	4 Properties	8 Properties
Informix 6.0 (49 Properties)	3	6	10
Sybase 10.01 (49 Properties)			
Informix 6.0 (392 Properties)	5.1	7.5	11.3
Sybase 10.01 (392 Properties)	25.4		
Informix 6.0 Faster	494%		



12.2 Concurrent Property Loads Conclusions

The loading of one property in Sybase took more than twice as long as Informix loading 8 at a time. Informix scaling was excellent. With Informix, eight simultaneous property loads took less than twice as long as loading one property in Sybase.

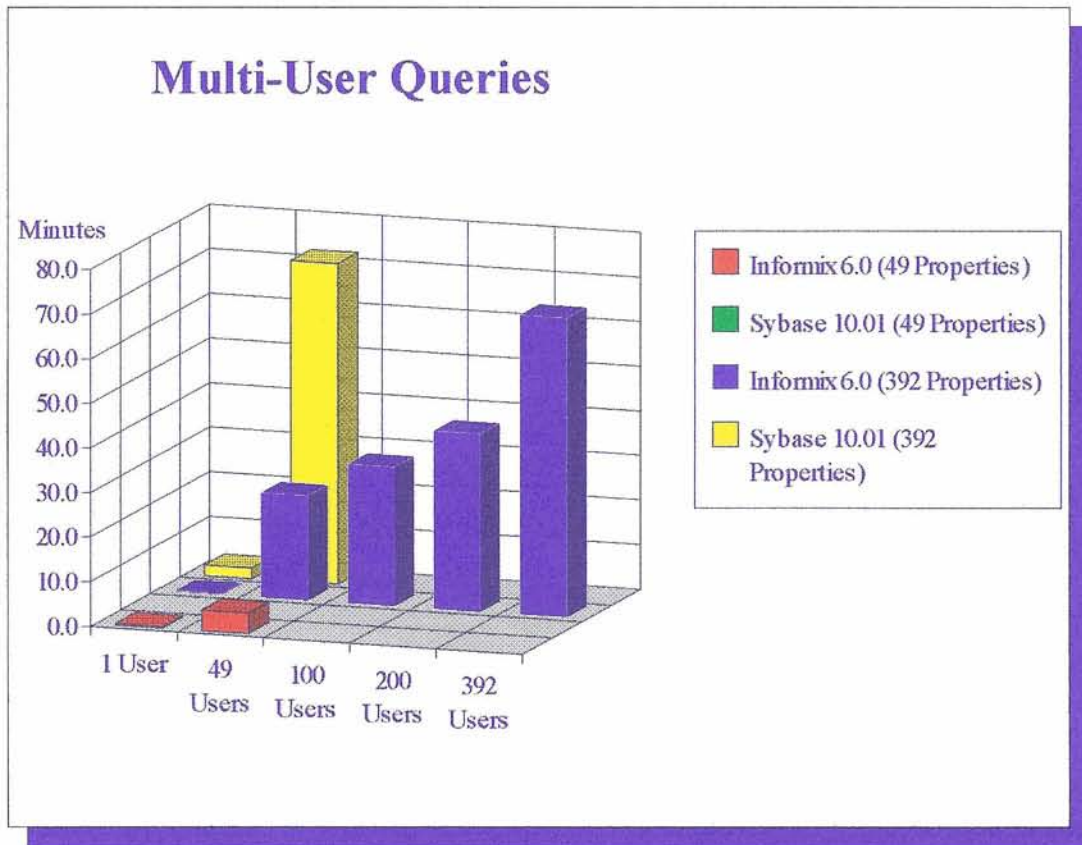
Sybase did not run the 49 property test due to time. On the 392 property test, Holiday Inn told Sybase not to run the four and eight concurrent load test because the time for one was more than double the time Informix could run eight.

13.0 Multi-User Queries

Run a set of inserts, updates, deletes, table creates and table drop statements. Each set of queries applied to only one property. To help reduce contention, the users were spaced by five second pauses. The SQL statements were actual statements from the HIRO nightly processing.

13.1 Multi-User Queries Results

<i>Timings in Minutes</i>					
	1 User	49 Users	100 Users	200 Users	392 Users
Informix 6.0 (49 Properties)	0.4	4.7			
Sybase 10.01 (49 Properties)					
Informix 6.0 (392 Properties)	0.4	23.5	31.2	40.0	67.2
Sybase 10.01 (392 Properties)	2.5	72.0			
Informix 6.0 Faster	639%	306%			



13.2 Multi-User Queries Conclusions

This was the most important test in the benchmark. Although Informix had an unscalable hop from 1 to 49 users on the 392 property run, from 49 to 392 users, Informix was extremely scalable. For 392 properties, Sybase was only able to run 49 users. Sybase's 49 users took longer than Informix's 400 users. This validates the problem that HIRO is currently having with Sybase; once the numbers of users increases, the Sybase database becomes too slow and unusable. Due to the scalability of Informix, HIRO management is estimating that Informix can handle a 1,000 property database. This meets our goal of a minimum of 2 database instances for the 2,000 property HIRO roll out. Today, Sybase can only handle a 60 property database.


Sybase did not run the 49 property test due to time. On the 392 property test, Sybase did not post times for more than 49 users because of exponentially increasing run times.

14.0 Altered Multi-User Queries

The vendors took the multi-user script and rewrote the script to be tuned with their engine. The only restrictions placed on the vendors was that they not violate any business rules with the queries. This was confirmed by the HIW benchmark team. The script was only run on the 392 property database.

14.1 Altered Multi-User Queries Results

<i>Timings in Minutes</i>					
	1 User	50 Users	100 Users	200 Users	400 Users
Informix 6.0	0.3	5.3	9.3	18.5	38.5
Sybase 10.01	0.4	13.5	29.4	56.1	110.4
Informix 6.0 Faster	137%	255%	318%	303%	287%



14.2 Altered Multi-User Queries Conclusions

Both vendors demonstrated performance improvements when tuning the queries for their database engine. It must be noted, however, that Sybase was unable to perform several of the HIRO queries without tuning the queries specifically for their database engine.

Informix was almost 300% faster on the 400 user test than Sybase. With this test, the HIRO management predicts that the combination of tuning the SQL and using Informix, the HIRO application would experience dramatic speed increases over the existing Sybase application and no additional computing resources will be required.