RG Solutions®

Comparative Analysis of ITI Processing in Different Platform Environments

Revision 1 (November 2006)

Michael A. Davis

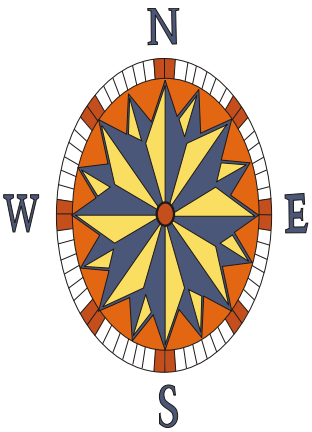

Resource Guidance, Inc.

"Quality Information Since 1987"

5332 W. Pershing Ave
Glendale, Arizona 85304
(602) 843-0023
www.rgsolutions.com

Executive Briefing

RG Solutions®

Background

This is a revision to a study (February 2006) relating to the ITI (Information Technology Inc.) application running on different vendor systems. The initial study was in response to a white paper published by ITI. In that white paper, ITI provided the results of their benchmark tests on the IBM i5 system. The ITI white paper focused on the Primary Update portion of the batch processing. It was felt this focus was too narrow and only looked at a small portion (albeit an important part) of the total user experience. As a result, the initial study was conducted with the intent of measuring all actual processing (batch and online) that is performed by the ITI application in a customer setting.

The initial study included four system types, and there were two each from Unisys and IBM. While these systems did not include the latest offerings available for each vendor, they were the only systems available at the time. Since then additional system information has become available and is included in this revision.

Scope

For this revision, the study added three new systems (two from Unisys and one from IBM), and deleted two other systems (one for each vendor). One system for each vendor was retained to represent some of the installed base and to form a baseline for the systems.

The following systems were included in the revised study:

IBM: AS/400
i5

Unisys: NX6800
FS1300
FS1600

Individual Results

The newer systems were compared to the base system (NX6800 for Unisys and AS/400 for IBM). Each was evaluated as to change in value returned (percent) to the customer. Assuming that all RPM/CPWs (depending on the vendor) cost exactly the same, a positive amount indicates a gain in value and a negative amount indicates a loss in value. For the new systems, the following are the change in values received: Primary Update (the very intense part of batch processing), Total Batch, and Online processing.

IBM i5 from AS/400

Primary Update:	14%
Total Batch:	29%
Online:	- 11%

Unisys FS1300 from NX6800

Primary Update:	61%
Total Batch:	19%
Online:	180%

Unisys FS1600 from NX6800

Primary Update:	657%
Total Batch:	62%
Online:	41%

Resource Guidance, Inc.
5332 W. Pershing Ave
Glendale, Arizona 85304

(602) 843-0023
www.rgsolutions.com
mail@rgsolutions.com

RG Solutions®

Comparative Results

When comparing systems between the two vendors, it would not be fair to compare the older systems with the new. Therefore, the vendor comparison is based on the IBM i5 system and how it compares to the Unisys FS1300 and FS1600 systems.

Each of the following entries will identify the vendor that has the advantage and the magnitude of the advantage. (For example, for i5 vs FS1300, the i5 has a 5.6 to one advantage over the FS1300).

IBM i5 vs Unisys FS1300

Primary Update:	IBM 5.6
Total Batch:	IBM 5.2
Online:	Unisys 4.0

IBM i5 vs Unisys FS1600

Primary Update:	IBM 1.8
Total Batch:	IBM 3.8
Online:	Unisys 2.0

Conclusion

When looking at the comparative results, it is obvious that IBM controls the Batch processing environment while Unisys controls the Online. However, a few points must be made about these results. In the initial white paper, the iSeries system (which was removed from the revision) held a fourteen to one advantage in Primary Update processing. (This type of processing is the basis of the ITI white paper). This advantage has been reduced to just under a six to one advantage on the FS1300, and less than a two to one advantage on the FS1600. This is a substantial improvement (over 200% for just the FS1300 system) for the Unisys systems as compared to the initial white paper.

For the Online processing, the initial white paper gave Unisys about a two to one advantage over IBM. This advantage is maintained for the FS1600 system while the FS1300 system now has a four to one advantage over the i5.

The most interesting result involves the value returned to the customer. For the IBM i5 system, the customer receives increased value for the batch processing while it loses value for the online processing. For the Unisys systems the customer receives increased value for both the batch and online processing, and in some cases, the increased value received by the customer is substantial.

With the introduction of the newer systems from Unisys, the customer no longer has to worry about the Primary Update and Batch Window processing. The newer Unisys systems are capable of handling these requirements within the customer constraints. This type of processing has never been a problem on the IBM systems.

The technical focus of the customer now moves from the static batch processing environment to the dynamic and growing area of Online processing. Meanwhile, the business focus of the customer can now center on the value returned from the vendor solution rather than the issues involving processing constraints.

Notices

Disclaimer

Michael A. Davis of Resource Guidance, Inc. is solely responsible for the content of this document. The analysis contained within represents original work and is protected under Copyright law.

Information used for the analysis contained herein was provided from customers, outside sources, and published documents. All information used was publicly accessible and no collected information was proprietary in nature. An assumption was made that all external information as provided was accurate and correct at the time. The analysis contained within this document is based on a specific set of information and the results of this type of analysis may vary.

Legal Notice

References

This document makes several references to an "ITI White Paper" that was created by Information Technology Inc. The white paper can be downloaded from www.itiwnet.com and is titled:

Performance Benchmarks:

ITI's Premier Banking Software on the IBM eServer i5 Platform

September 13-17, 2004

IBM eServer Benchmark Center
Rochester, Minnesota

Revision Summary

The initial white paper was published in February of 2006. Since that time, the state of the ITI processing in the IBM and Unisys environments has changed. The purpose of issuing this revision to the original white paper is to remove dated information while supplementing the paper with current information. This section outlines the changes made to the original document resulting in this revised white paper.

Unless noted here, the format and content of the document has not been changed. In some cases, the values of items in tables may have changed due to revised methods of system measurement.

Removed

The following systems have been removed from the white paper:

- Unisys NX5800
- IBM iSeries

The initial document contained a section called "IBM Discrepancy Analysis". Since this section pertained to the iSeries processing, it is no longer pertinent and has been removed.

Added

The following systems have been added to the white paper:

- Unisys FS1300
- Unisys FS1600
- IBM i5

An Appendix A has been added as part of the revision. This appendix contains an analysis and discussion of the 3101 update program. This program is a major component of the Primary Update processing.

Major Revisions

There have been no major revisions. All sections present in the initial document (with the exceptions noted here) are present in this revised document, and in the same form. This allows for the comparison of information in the original and revised documents.

Background

Purpose

The intent of the original white paper was to compare the results of the Information Technology Inc. (ITI) white paper to actual user experiences. This revision does not waver from this objective, but the original study did not include some of the newer product offerings provided by IBM and Unisys.

The information scope for the original white paper contained four systems, and included two older Unisys and two older IBM platforms. At the time of the original study, information on the newer systems was not available and could not be included. Since then information about the newer systems has become available and can now be included in this revision.

As before, the intent of this revision is to measure the actual customer value received on different systems and then compare the results in a consistent manner. In this way, all types of processing can be evaluated in a "live" environment that is being experienced by actual customers. This will lead to a more informed buying decision for the customer.

Scope

The scope of the analysis included five systems: Unisys NX6800, Unisys FS1300, Unisys FS1600, IBM AS/400, and IBM i5. While there is a large amount of data with detailed analysis, the results from each system have been averaged to create a composite set of values for the system. The large volume of supporting detail used in the analysis process will not be presented for practical reasons. The content, format, and manner of presentation were deliberately done to protect the identity of the participants in this study.

Workloads

While the objective of the study was to compare the performance characteristics of each system, it was necessary to first establish a baseline of total system processing for each system. This baseline is used to determine the processing characteristics of the system and to determine the similarity or dissimilarity of each system.

For this software application, there are three different and distinct workload types (types of processing). The first is Update processing and is defined as any application program with a number range of 3000 to 5999 (inclusive). For the Update processing, the volume of work is dependent upon the number of accounts and the daily activity to those accounts. Across different system types, the activity per 1,000 accounts remains relatively constant and will be predictable.

The second workload type is Report processing and is defined as any application program with a number range of 6000 to 9999 (inclusive). For the Report processing, the processing is again related to the number of accounts. Since the Report processing generally runs with or following the Update processing, it too will be consistent and relatively predictable for any given 1,000 accounts.

The final workload type is the Online processing and is defined as any application program with a number range of 0000 to 2999 (inclusive). Unlike the other workloads, this processing is dependent upon the marketing and business plans of the bank. These business decisions create a "mix" of online transaction processing that will vary from bank to bank. For this reason, it is reasonable to expect the volume and types of transactions will vary between the systems. The amount and extent need to be established before the systems can be compared.

It will be argued that the criteria used for forming workloads is not correct due to exceptions or other interpretations. This is true. However, the rule set used was applied consistently for all systems and allows for a comparison based on a single standard. If in someone's opinion some processing was categorized incorrectly, then it was done so consistently and for all systems in the study.

As with any system, not all the processing is done by the application itself. There is processing that is performed by the system or by other programs that support the underlying user processing. This indirect processing must be accounted for in the total scheme of things. To ensure this processing is included into the resulting workloads, a proprietary set of algorithms was used. These algorithms not only allocate the indirect processing to the workloads, but ensure it was allocated using a consistent and fair methodology. The result is that the workload types reflect the actual processing associated with it. The sum of all workloads is equal to the sum of all system processing.

During the presentation of the study findings, there will be mention of other groups used for a particular measurement. For example, one such group is called Total Batch and contains all the Update and Report processing for the entire day. As these groups are presented, they will be defined regarding their contents.

Findings

The analysis process was very complex in nature and will be presented in three stages:

- Batch Processing
- Online Processing
- Combined Workload Processing

The findings for the Batch Processing will deal with the specifics of the batch processing. This will include all Update and Report processing, and will be presented in different functional groupings. For this section, there will three groups:

- Primary Update
- Batch Window
- Total Batch

The findings for the Online Processing will follow. In this section, the methods used and the results of the Online Processing analysis will be presented.

The final section is the findings for the Combined Workload Processing. In this section all processing will be combined to form a complete picture of the systems and how they compare with each other. This final section represents the most complete picture of what an ITI customer can expect.

Measurements

The comparison of unlike systems is a complex task so a brief discussion of the comparison methodology is warranted. For Unisys and IBM systems there is no one common value that can be used. Comparison is further complicated by each system in the study having a different configuration and processing power. For this reason a new value had to be used for the comparison and is based on Metering technology.

The first part of the measurement value is based on the power of the corresponding system. For the Unisys systems this will be RPMs, and for the IBM systems this will be CPWs. The theory behind these power values is that regardless of the system's power rating, the amount of value delivered to the user will be the same per unit (RPM or CPW) over time.

The use of the power value alone carries an implied time value. One usually describes their system with only the power rating (such as 2000 RPM), but in reality they mean 2000 RPM Hours (or seconds or some other time duration measure). In this manner, they can describe the delivery of processing benefit to the user over time. For example, if a user has a 2000 RPM system and upgrades to a 4000 RPM system, they would expect to receive twice the benefits of their old system for the same period of time.

You will find in this study that PHrs is used for the basis of measuring work on all systems. The P stands for "Power" and is either RPM or CPW. Since power without time is meaningless, the unit of time used is Hours. (Hours are the most workable time unit for this comparison). By using this method to measure systems, it is not necessary to deal with processor seconds, multiple central processor units, and the varied power ratings of the systems. All these considerations are present in the PHrs value. It also allows for the comparison of systems because the time component is the same (hours) for all measurements.

When the measurement values are presented in the findings, they will be based on a fixed amount of work the system must perform. This is called a Unit of work. Each work unit will be consistent for all systems and the given workload types. For Update and Report, the unit of work is based on the processing required for 1,000 accounts. For Online, the unit of work is based on the processing required for 1,000 transactions.

In all cases except one, the power measurement is based on the underlying processor time used on the study system. The one exception is the comparison of results in the "Update Benchmark Results" section of the ITI White Paper. Since the results are presented in elapsed times, it was necessary to use elapsed hours instead of processor hours for the comparison. For example, for Database "C" and System "D":

System CPW: 3300
Elapsed Time: 1:20:50
Elapsed Seconds: 4850
Accounts: 979,884

This results in:

((4850 * 3300) / 3600) /
(979884 / 1000)

Which results in a value of 4.54 PHrs per Unit (1,000 accounts).

Generally speaking, if there are no external system issues and this value is about the same for all configurations of a particular system type, then it is said that the application "scales" on the system.

Findings: Batch Processing

Introduction

The initial stage of the comparative analysis focused on the Batch Processing of the systems. The batch analysis of the system was broken into three functional parts:

- Primary Update
- Batch Window
- Total Batch

The Primary Update processing is the work that most closely resembles the processing in the "Update Benchmark Results" section of the ITI White Paper. The analysis allows for a comparison of actual system processing with the results found in the ITI White Paper.

The second area of batch analysis deals with the Batch Window. This period of time is called a Batch Window because many times it is the only "window" of time in which the update processing can run. In this situation, the processing must be completed in the "window" of time. Each of the systems had a defined period of time in the evening (and possibly early morning) when the batch processing was performed. The Primary Update is included in this "window" of time, and must be completed along with other batch processing.

The final area of analysis had to do with the Total Batch processing. It was found on all systems that some of the batch processing occurred outside the Batch Window. The reasons for this are not known, but since the processing did occur, it was deemed necessary and worthy of consideration for a complete batch processing analysis.

Primary Update

The "Update Benchmark Results" section of the ITI White Paper provides the elapsed times of various suites of programs running different "mixes" of Update. In order to compare the study findings with the ITI results, it is necessary to first establish a comparable mix of programs that will represent the processing present in the ITI results. While the exact "mix" of the ITI processing is not known, it most likely will never exactly match those found on the study systems. Therefore, it is necessary to find an update set that most closely matches those present in the ITI White Paper.

After some analysis and consultation, the following programs (where x is any number) were included in the study's Primary Update suite:

COD 3xxx
DDA 3xxx
FMS 3xxx
LAS 3xxx
SAV 3xxx

Since the ITI White Paper results were based on elapsed times, it was necessary to create the same results from the study systems. This was done by identifying the programs in the total processing environment, extracting the pertinent information, and then ordering the processing in a time-based manner.

After reviewing the initial results, it was discovered that due to multiple databases and other local scheduling issues, there were "gaps" in the processing streams. These gaps represent times on the systems when there were no

Primary Update programs running. In some cases there was a gap of more than one hour. While it is reasonable to expect gaps in a normal processing environment, it is not expected in a bench mark environment, especially when the benchmark results are using elapsed time for the measurement basis. For this reason, an adjustment was made in the extracted information to eliminate the "gaps."

Table 1 provides the results of the Primary Update analysis. As can be seen (and expected), the actual and adjusted Run Times for the Unisys systems were much longer than on the IBM systems.

In order to compare the Elapsed Times of the systems, it was necessary to use a measurement that is common to all systems. Since all the systems have different processing power (RPM or CPW) and a different number of accounts, the measurement that was used had to take these facts into consideration. The measure used was PHrs per Unit (where P is either RPM or CPW and Unit is 1,000 accounts). As can be seen from the table, the elapsed time for the i5 system ranged from 3.14 (PHrs per Unit) down to 2.07. The corresponding value from the ITI White Paper was 4.36 PHrs per Unit, and this value falls near the range (actual to adjusted) found for the i5 system. Since these values are so close, it would be safe to assume that the selected workload for the Primary Update processing on the study systems is fairly representative of the processing present in the "Update Benchmark Results" section of the ITI White Paper.

When looking at the Ratios for the Elapsed Time (in Table 1), it can be seen that the i5 system is outperforming the Unisys systems by more than a seven to one ratio.

	Unisys			IBM		ITI White Paper
	NX6800	FS1300	FS1600	AS/400	i5	
Run Times						
Actual	07:45:24	04:26:00	02:11:53	01:02:20	00:04:33	
Adjusted	05:12:22	01:55:02	00:36:28	00:33:24	00:03:00	
Elapsed Time						
Actual (PHrs per Unit)	74.37	59.19	80.13	8.03	3.14	4.36
Adjusted (PHrs per Unit)	49.91	25.60	22.16	4.30	2.07	
Actual Ratio	23.685	18.850	25.519	2.558	1.000	
Adjusted Ratio	24.111	12.367	7.057	2.079	1.000	
Processor						
PHrs per Unit	14.02	8.73	2.81	1.78	1.57	
Actual Utilization	18.9%	14.8%	3.5%	22.2%	50.0%	
Adjusted Utilization	28.1%	34.1%	12.7%	41.4%	75.8%	

Notes:
1. Elapsed Times are subject to customer scheduling algorithms.
2. PHrs = RPM or CPW Hours.
3. Unit = 1000 Accounts.
4. Ratios are based on Elapsed Time.
5. Ratios are relative to i5 (1.000).

Table 1: Primary Update

This is a 200% improvement for the Unisys systems as compared in the original white paper. (Note: The comparative processor usage (PHrs) of the systems will be addressed in later sections).

There are some other interesting results from the new systems. When looking at the Adjusted Utilization rates of the processor (in Table 1), the i5 system is at a 75% rate. This suggests the system is near maximum processor utilization and no other user optimization will affect the elapsed time.

This is not the case for the FS1300 and FS1600 systems. For these two systems, there are opportunities for optimization that could significantly reduce the Primary Update elapsed times. For the FS1300 system, the elapsed time could be (theoretically) cut in half to under one hour (as measured here). For the FS1600, the throughput could be (theoretically) improved by almost five times (based on the Adjusted Processor Utilization). This could reduce the FS1600 elapsed time to just over five minutes, which is very close to the i5 value.

One final note about the Primary Processing that is not reflected in Table 1: The main update program for the Primary Update processing is the xxx3101 program (where xxx is COD, DDA, LAS, or SAV). A detailed comparative analysis was done for this program. (Please refer to Appendix A for more information).

It was found that while the Primary Update processing may be functionally equivalent between the Unisys and IBM systems, the actual processing was not. The Unisys systems are required to process a much higher number of writes than the IBM systems. The reason for this is not known, but this requirement adds additional overheads, is highly disruptive to program efficiency, and "hamstrings" the system from achieving good performance. The net result of this difference is that Unisys systems cannot achieve optimal efficiency as compared to the IBM systems. While a functional comparison is valid between the Unisys and IBM systems, it would not be proper (nor fair) to compare the physical capabilities of the systems until the Unisys processing has been changed to coincide with the IBM processing.

While the Primary Update processing has never been a concern with the IBM systems, it was a major concern on the older Unisys systems. With the improvements reflected in the newer Unisys offerings, the Primary Update processing should no longer be of a major concern to the customer even with the processing differences identified (the xxx3101 programs) between the two platforms.

Batch Window

The Batch Window processing is the non-online processing that must complete in a given period of elapsed time. This processing usually occurs in the evening and possibly early morning hours. Due to the nature of the Primary Update processing (which is included in this time period), this time interval had to be tightly managed for the older Unisys systems (due to the long Primary Update times).

For this part of the analysis, it was necessary to analyze all the processing that occurred on the system and determine the time limits of the Batch Window. The start and stop times were based on the major Update and Report activity present on the system and had to include the Primary Update processing.

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Run Times					
Start	18:00:00	18:00:00	18:00:00	21:00:00	21:00:00
Finish	04:29:59	02:17:59	22:32:59	23:25:40	23:37:29
Elapsed	10:30:00	08:18:00	04:32:59	02:25:40	02:37:29
Processor (PHrs per Unit)					
Update	23.29	31.43	12.21	6.33	5.64
Report	30.18	11.61	16.90	2.15	1.43
Total	53.47	43.04	29.11	8.48	7.07
Primary Update					
PHrs per Unit	14.02	8.73	2.81	1.78	1.57
Percent of Batch Window	26.2%	20.3%	9.7%	21.0%	22.2%

Notes: 1. PHrs = RPM or CPW Hours.
2. Unit = 1000 Accounts.

Table 2 Batch Window

Table 2 provides the results of the Batch Window analysis. While the start times were identical for the Unisys systems, they were not for the IBM systems. This was mainly due to the fact that the IBM systems enjoyed a shorter Batch Window than the Unisys systems and this offers more flexibility in batch scheduling. Also, the start times on the Unisys systems are now more traditional rather than due to processing requirements.

Since the Primary Update is part of the Batch Window, it is appropriate to include its processing in the table. It is worth noting that the Primary Update processing of the IBM system now constitutes about the same amount as on the Unisys systems (with the exception of the FS1600). Previously, the i5 was a small portion of the Batch Window processing as compared to the Unisys systems.

Another interesting result is that the Primary Update processing on all systems represents roughly one quarter (or less) of the total Batch Window processor use. While the Primary Update processing was a major concern (especially on a Unisys system), there is obviously more going on during the Batch Window processing (for all systems) than just the Primary Update.

Due to the increased capabilities of the Unisys systems, the user currently has no incentive to reduce the Batch Window processing times. The processing that is performed during this period is much more efficient and currently is not pushing any user constraints. Should the requirements change, the user on the Unisys systems knows they have tuning options available to them that are not available on the IBM systems. Simply stated, there is currently no business case to optimize the processing for the Batch Windows on the Unisys systems.

Total Batch

The final step in the batch analysis was to look at the Total Batch processing for the day. It is not known why batch processing

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Processor (PHrs per Unit)					
Update	27.79	33.19	15.53	9.23	7.91
Report	34.56	19.03	23.03	3.78	2.19
Total	62.35	52.22	38.56	13.01	10.10
Batch Window					
Update (PHrs per Unit)	23.29	31.43	12.21	6.33	5.64
Report (PHrs per Unit)	30.18	11.61	16.90	2.15	1.43
Total (PHrs per Unit)	53.47	43.04	29.11	8.48	7.07
Percent of Day	85.8%	82.4%	75.5%	65.2%	70.0%
Primary Update					
PHrs per Unit	14.02	8.73	2.81	1.78	1.57
Percent of Batch Window	26.2%	20.3%	9.7%	21.0%	22.2%
Percent of Day	22.5%	16.7%	7.3%	13.7%	15.5%

Notes: 1. PHrs = RPM or CPW Hours.
2. Unit = 1000 Accounts.

Table 3: Total Batch

(Update and/or Report) would occur outside the Batch Window, but it did, and so it must be considered as part of the total analysis.

For this phase of the analysis, it was necessary to total all Update and Report processing that occurred for the day. These results are presented in Table 3 along with comparable information for the Batch Window and Primary Update processing.

This type of processing has changed quite a bit since the initial white paper. Previously, the majority of all batch processing was performed in the Batch Window on the Unisys systems. This was mainly due to the very long Batch Windows that occurred on the Unisys systems.

With the advent of the FS1300 and FS1600 systems, the Unisys model of the Batch Window has shifted more toward the IBM model. That is to say, the Unisys customers are running more of the batch processing outside of the Batch Window. The main reason for this is that the Batch Windows are becoming shorter on the Unisys systems, the systems are using less processor for Primary Updates, and the customer is less concerned about the processing activities that occur outside of the Batch Window.

Total Batch Ratios

The final step in the batch analysis was to compare the systems in some form of a meaningful manner. While the Primary Update was previously compared using a ratio based on Elapsed Time, it is not possible to compare the other aspects of the batch processing in the same manner. The reasons are varied, but are mainly due to the scheduling discretion available to the user. What processing is done and when it occurs is under the user control.

The intent of the study is to determine the ability of all systems to perform the required processing. For the comparison, the measure used is based on the weighted

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Total Batch	6.173	5.170	3.818	1.288	1.000
Batch Window	7.563	6.088	4.117	1.199	1.000
Primary Update	8.930	5.561	1.790	1.136	1.000

Notes: 1. Ratios are based on Processor Use.
2. Ratios are relative to i5 (1.000).

Table 4: Batch Ratios

processor usage for 1,000 accounts. This value is noted as PHrs per Unit in the previous tables. The values contained in Table 3 were used to create a new table (Table 4) showing the relative processing capability of the various systems for the three types of processing (Primary Update, Batch Window, and Total Batch). Table 4 provides the resulting system ratios. Since the i5 had the best overall time for the Primary Update processing, it was used as the basis for all other system ratios.

What is interesting from the table are the results for the FS1300 system as compared to the i5. While the i5 system still has an advantage over the FS1300, it is a consistent advantage regardless of the type of batch processing.

The real surprise of the table has to do with the FS1600. When compared with the i5 system for Primary Update processing, the i5 has less than a two to one advantage. This is a quite a change as compared to the initial white paper and previous perceptions about the systems. On the whole, the batch processing on the Unisys systems have improved dramatically as compared to the older systems.

For the batch processing, the i5 customer who migrated from an older IBM system will realize an increase of about 30% in the value returned from the new system. For the same batch processing, a Unisys customer that migrates from an older system will realize from 60 to a 750% improvement depending on the system and type of batch processing.

Findings: Online Processing

Introduction

Online Processing is defined as that system processing having a response time component, and is generally the interaction of a user or customer with the system. This interaction is called a transaction, and generally has an implied response time requirement that ensures timely interaction of the system environment with a human individual. Update and Report processing do not have a response time requirement.

The first step in the online analysis was to understand the similarity and/or dissimilarity of the online processing. The purpose of this is to determine whether the raw system transactions can be used for measurement, or a "standard" transaction must be substituted. To do this, all online processing is allocated into business specific categories and the volume of processing is analyzed. One assumption was made about this processing. Since it is assumed that the ITI application (almost) uses the same code for all platform architectures, then it is assumed that the processing in like categories would be similar across all systems. The only difference would be the rate (number) of transactions in the peak hour.

Analysis found there was too much of a difference in the business activity levels of the systems, so a "standard" transaction had to be used. A standard transaction is based on the e-banking transactions of all systems. For a given system's online processing, the processing was allocated in increments of the system's e-banking transaction. This gives us like transactions (standard) that can be used in comparisons.

Table 5 provides the summary of the Online processing for each of the systems. As can be seen, the

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Standard Trans per 1000 Accounts	0.52	1.32	0.71	0.59	0.45
Percent of Total Online Processor					
ATM	12.6%	17.1%	0.0%	13.3%	11.9%
Electronic Banking	11.2%	24.0%	20.0%	35.5%	25.1%
Exception Item Module	2.5%	7.3%	0.4%	0.3%	0.1%
File Management	1.4%	7.2%	0.7%	4.2%	3.1%
Item Entry System	2.4%	1.6%	0.9%	0.5%	0.4%
Teller	0.0%	1.1%	1.2%	2.7%	2.1%
Transaction Management	0.0%	9.1%	45.7%	4.2%	30.8%
Inquiry Support	28.9%	12.6%	9.0%	12.2%	7.5%
System Control	30.9%	17.4%	18.2%	11.1%	6.4%
Other	10.1%	2.6%	3.6%	16.0%	12.6%

Table 5: Online Processing Summary

Unisys FS1300 system had almost three times the volume of transactions as the i5 system. As stated previously, the "mix" and volumes of these transactions are not determined by the number of accounts on the system, but rather from the business activities of the bank. (Note: The number of transactions used in the analysis represents only those transactions processed by the system. According to the ITI White Paper, these transactions represent about 25% of the total actual user requests.)

What is interesting from the online analysis is the distribution of the work to the various business processes. No two Unisys systems had a similar profile indicating business diversity. It is also interesting to note that the Unisys systems had a higher level of allocations to the Inquiry Support and System Control than the IBM systems.

Comparative Results

Having established the volume of standard transactions and the associated processing for each system, the next task was to measure the processing. This had to be done in such a way so it could be compared across systems. As used in the batch analysis, the use of PHrs was used. This allows for the comparison of systems.

However, using the batch measure based on 1,000 accounts would not be appropriate. The reason for this is simple, each of the systems have a different transaction volume as well as a different distribution of transactions to the business categories. Since all the transactions used in the analysis were standard transactions (all the same on the system), then the standard transactions could be used for the measurement. The measure used was PHrs per 1000 Standard Transactions.

Table 6 provides the results of the online analysis. As can be seen from the table, Unisys exhibits an advantage for this type of processing. Since the FS1300 had the best processing result, it is used for the basis of the comparison. For the online processing, Unisys exhibits almost a four to one advantage over the i5 and a 3.5 to one advantage over the AS/400. In the previous white paper, the best Unisys system exhibited a three to one advantage over the iSeries system and a 1.5 to one advantage over the AS/400 system. The FS1300 system has increased the Unisys advantage in the Online Processing.

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
PHrs per Unit	91.68	32.70	64.83	114.79	129.61
Ratios	2.804	1.000	1.983	3.511	3.964

Notes:
1. PHrs = RPM or CPW Hours.
2. Unit = 1000 Standard Transactions.
3. Ratios are based on Processor Use.
4. Ratios are relative to FS1300 (1.000).

Table 6: Online Comparison

Findings: Combined Workload Processing

Introduction

Having analyzed both the Batch and Online processing for each system, the next task was to analyze the total processing on the systems. This was done with previously established workload types and functional groups that were used in the analysis.

Processing Summary

The first analysis looks at all the workload and functional processing on the system in regard to the total processor used. Table 7 shows the overall allocation of processing to the various workload types and functional groups. As can be seen from the table, the Update and Report (Batch) processing on the Unisys systems constitute more than 42 percent of the total processing while it was less than 16 percent of total processing on the IBM systems. This is not unexpected due to the large processing requirements for the Primary Update on the Unisys systems.

Unlike the previous white paper findings, the results for the IBM systems are consistent. The i5 processing improved for batch activities and caused the allocations for the update and reports to be reduced as compared to the AS/400. As a result, the i5 online allocations was increased 87% of the total system processing.

Likewise, the Unisys systems improved in the Update and Report (Batch) processing. Previously, 55% of the Unisys systems were devoted to batch processing. With the newer systems, this value is around 50% for the FS1300 and 42% for the FS1600 (the latter is mostly due to its improved Primary Update processing).

Since the Unisys systems have a larger allocation to the Update and Report processing, it is reasonable to expect that they would have a smaller contribution to the Online processing. However, it must be noted that the Unisys FS1300 system was handling almost three times the volume of online transactions as the i5 system, and this was done with a smaller allocation of processing resources as compared to almost all the other systems.

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Percent of Total Processor					
Update	26.8%	31.2%	17.0%	11.4%	9.9%
Reports	29.5%	19.7%	25.2%	4.7%	3.0%
Online	43.7%	49.1%	57.8%	83.9%	87.2%
Batch Window	49.1%	45.2%	31.9%	10.5%	9.3%
Primary Update	12.9%	9.2%	6.2%	2.2%	2.1%

Table 7: Processing Summary

Another interesting result from the Processing Summary table has to do with the Batch Window. On the Unisys systems this processing amounts to about 35 percent of the total processing for the day, while on the IBM systems this is 10 percent of all processing.

In order to see what the total system processing would look like if all the systems had the same standard transaction volumes, Table 7 was normalized (based on the FS1300 values). Table 7R provides a normalized view of all systems processing the same relative volume of standard transactions. (It should be noted that the AS/400 system was already at maximum online capacity and could not possibly handle the change in volume.)

An interesting change has occurred in this table. With the normalization applied, the FS1600 online profile moves closer to the i5 profile by having a larger portion of online processing. This is mainly due to the efficiencies released in the batch processing of the FS1600 system.

Workload Usage

Once the workload processing allocations (Update, Report, and Online) were complete, the next task was to determine the system usage amounts for each workload type. This had to be done with an eye toward comparing workloads on each of the systems. This process was complicated by the fact that there is no established methodology for evaluating the effectiveness of a system's processing, let alone across vendors.

The Update and Report processing is relatively constant for a number of accounts, so these two types of workloads are measured based on a group of 1,000 accounts. In this manner we are measuring a unit of Update and Report work (1,000 accounts) and not the volume of work (total number of accounts). This Update and Report unit of work allows for portability across systems.

The Online workload presents a different situation in that it is both content and activity based and therefore required a different measurement basis. The intuitive measurement is based on the transaction activity, but this measurement basis is further complicated by the processing requirements of each type of transaction. This situation was alleviated by using a "normalized" transaction (standard transaction). Therefore, the unit of measurement for the Online processing is based on a group of 1,000 standard transactions.

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Percent of Total Processor					
Update	16.1%	31.2%	11.4%	5.6%	3.7%
Reports	17.7%	19.7%	16.8%	2.3%	1.1%
Online	66.3%	49.1%	71.8%	92.1%	95.2%
Batch Window	29.4%	45.2%	21.3%	5.2%	3.4%
Primary Update	7.7%	9.2%	4.1%	1.1%	0.8%

Table 7R: Normalized Processing Summary

The final requirement was to determine the measurement value that applies to each workload unit. This value has to be of such a nature that it can be compared across the different system types. For the Unisys systems, the RPM is used by the vendor to denote the delivery of completed work to the user. Therefore, the Unisys system measurement value is RPM Hours (delivery of work over time). For the IBM systems, the CPW is used by the vendor to denote the delivery of completed work to the user. Therefore the IBM measurement value is CPW Hours (delivery of work over time).

These two values, RPM Hours and CPW Hours are the means to measure the workload usage of the systems. Since the base unit of time will be equal on all systems, it will be easy to compare the RPM and CPW values for the various workload types.

Table 8 provides the Workload Usage Summary information. When looking at the totals, again there is no surprise regarding the usage allocations (PHrs) for the Update processing. The allocated amounts are expected based on established information and other independent tests.

When looking at the batch processing values (Update and Reports), there are no surprises from the IBM systems. The i5 system has improved in both categories. What is interesting is performance of the new Unisys systems. The FS1600 performs much better than the FS1300 for Update processing while the FS1300 performs much better than the FS1600 for Online processing. The reasons for this would be due mainly to the architectural differences of the systems as well as any program processing that is present in the ITI application.

In the area of Online processing, Unisys has the advantage in all cases. The FS1300 system holds the advantage over all systems, including the FS1600 (by a two to one margin). Relative to the IBM systems, the FS1300 has almost a four to one advantage over the i5 system. It is interesting to note that the i5 system does not perform as well for the Online processing as does the AS/400.

System Comparison

The measurement criteria (PHrs) used previously allows for the comparison of systems. The measurement value used for the

	Unisys			IBM	
	NX6800 (RPM Hrs)	FS1300 (RPM Hrs)	FS1600 (RPM Hrs)	AS/400 (CPW Hrs)	i5 (CPW Hrs)
Update					
Usage per 1000 Accounts	29.23	27.42	13.51	9.23	6.55
Reports					
Usage per 1000 Accounts	32.16	17.26	20.01	3.78	1.97
Online					
Usage per 1000 Transactions	91.68	32.70	64.83	114.79	129.61

Notes: 1. Usage is based on Processor Use.

Table 8: Workload Usage Summary

Unisys systems is RPM Hours, and the measurement value used for the IBM systems is CPW Hours. Both of these measures are based on a single hour and thus allow for a comparison of the delivery of work to the user for each system.

Due to the nature of workloads, comparison of systems should only occur within similar workload categories. The reason for this is simple: all operating systems and hardware platforms are optimized for a specific type of processing. This means that in some cases a specific system might have a processing advantage for one type of workload and be at a disadvantage for another type. While it is possible to formulate a composite single number, it would be situation specific and would not generally apply across the board.

For this reason, system sizing must consider the processing requirements of all system processing. That is to say, the Update, Report, and Online workload processing must all be individually considered in order to properly size a system.

The contents of Table 9A (Unisys) and 9B (IBM) provide the conversion values for each of the system types. For each system, the values are equated to the corresponding values for the other systems. For example, on the i5 system, one Update CPW provides as much processing value to the user as 2.063 RPMs on the FS1600 system. Likewise, on the FS1300 system, one Online RPM provides as much processing value to the user as 3.964 CPWs on the i5 system.

Due to the large portion of Update processing on the Unisys systems, it was possible in the past to size the system based only on the Update processing. In these cases the Online processing volumes were lower and Update allowances coincidently handled the Online

Unisys NX6800 RPM/CPW Conversion

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Update	1.000	0.938	0.462	0.316	0.224
Reports	1.000	0.537	0.622	0.118	0.061
Online	1.000	0.357	0.707	1.252	1.414

Unisys FS1300 RPM/CPW Conversion

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Update	1.066	1.000	0.493	0.337	0.239
Reports	1.863	1.000	1.159	0.219	0.114
Online	2.804	1.000	1.983	3.511	3.964

Unisys FS1600 RPM/CPW Conversion

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Update	2.164	2.030	1.000	0.684	0.485
Reports	1.607	0.863	1.000	0.189	0.098
Online	1.414	0.504	1.000	1.771	1.999

Table 9A: Unisys Conversion Tables

IBM AS/400 CPW/RPM Conversion

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Update	3.166	2.969	1.463	1.000	0.709
Reports	8.504	4.565	5.292	1.000	0.521
Online	0.799	0.285	0.565	1.000	1.129

IBM i5 CPW/RPM Conversion

	Unisys			IBM	
	NX6800	FS1300	FS1600	AS/400	i5
Update	4.465	4.187	2.063	1.410	1.000
Reports	16.330	8.767	10.163	1.920	1.000
Online	0.707	0.252	0.500	0.886	1.000

Table 9B: IBM Conversion Tables

situation. However, this rule no longer applies due to the increase volumes of online transactions (mainly due to e-banking) and reduced Batch Window processing times. For the IBM systems, the Update processing requirement is much smaller and should never be used to size the total system.

One system measurement does not fit all. Sizing the system based on Update alone might have worked on Unisys systems in the past, but it is no longer viable. Any system should be sized for both Online and Batch processing requirements. The greater of the two constitutes the processing requirement.

Conclusions

Introduction

While it is understood that the findings are very technical in nature, it was necessary due to the complex purpose of the study. The findings had to include information that not only provided system comparison, but also support the values used in the comparison.

When entering into the initial study, the only known information was based on the results from the ITI White Paper and other independent sources. All information was focused mainly on the Primary Update processing as this was deemed the critical system processing. Recently this view has begun to change and customers are realizing there is more than one type of processing on their system. While still important, the Primary Update processing contributes less than 15 percent of total system processing and is becoming less critical due to improved system performance.

When the results were initially issued in the ITI White Paper, there was a lot of research and investigation but the results always held. The Primary Update processing was being processed many times faster on the IBM systems as compared to the Unisys systems. What was not known was how the systems handle the other 85 percent of the work.

Batch Processing

As part of the study, the entire batch processing was analyzed at three different functional levels:

- Primary Update
- Batch Window
- Total Batch

The results of the Primary Update processing confirmed the previous known facts. Results indicate the IBM i5 can do the work six times better than the FS1300 systems. However, this value is only two better when compared to the FS1600 system. (Note: this value is based on the processor comparison and not the elapsed time comparison). This is quite a dramatic fact when viewed in isolation.

However, in the previous white paper, the IBM systems enjoyed a sixteen times advantage over the Unisys systems (in Primary Update processing), and this disparity has been reduced to about six (or two) times. This accomplishment is quite remarkable when placed in a context of improved processing on the IBM i5 and the additional processing requirements (write I/Os) imposed by the ITI software on the Unisys systems. One can only speculate what the comparison results would be if all systems were doing exactly the same processing.

As the batch processing was further analyzed, it was found that the results were about the same as the Primary Update. Looking at the Batch Window (which includes the Primary Update) it was found that the IBM i5 has a six (or less) to one advantage as compared to the newer Unisys systems. (This is about the same as before). This comparison can be a bit misleading if taken out of context in that the Unisys Batch Windows contained substantially more processing than

the IBM Batch Windows (additional write activity).

To complete the batch analysis, the Total Batch processing for the entire day was analyzed. The results show that the IBM advantage stayed about the same. The overall batch processing on the newer Unisys systems have become more balanced. While the IBM systems still enjoy an advantage over the newer Unisys systems, the advantage has been greatly reduced in the Primary Update processing while gaining a small amount in the overall batch processing.

As was stated earlier, the Primary Update processing on the Unisys systems is encumbered by additional processing constraints that are not present in the IBM processing. If the constraints were eliminated, the Primary Update processing values would improve for the Unisys systems. Since the Primary Update is part of the Batch Window and Total Batch figures, all categories would be improved. What cannot be determined is the extent of the improvement as compared to the IBM systems.

Online Processing

Along with the batch processing, the study also addressed the area of Online processing. This type of processing was relatively small just a few years ago but has exploded in recent years. The growth is primarily attributed to the increase in e-banking and the demands it makes upon the system.

The online processing presents a unique problem to the traditional sizing model used for the ITI application. In the past, the traditional method of sizing was based on the number of accounts. However, for the online processing, growth occurs with increased accounts and/or increased activity to the accounts. It would be safe to say that increased activity to the accounts currently contributes the largest amount of growth in the online processing. This situation is further exacerbated in that online processing must also ensure a timely response for the human interface.

The comparison of online processing between systems is a difficult and complex process. This is due to the fact that the transaction "mix" and volumes will vary between systems. In the study it was possible to address this situation and to form a basis for comparison between systems. The net result was that for the Unisys systems, the online processing outperformed the IBM i5 system by at least two to one, and as high as almost four to one (depending on the system).

Decision Matrix

In the past it was possible to size a Unisys system by just considering the update processing. This would be based on the number of accounts that needed to be processed in a given period of time. Since the online activity was quite low, the online sizing requirement was handled by default.

With higher transaction volumes (mainly due to e-banking), it is no longer possible to size an IBM or Unisys system based only on the update processing and the number of accounts. The system sizing exercise must now include both batch and online sizing. For the batch sizing it still is the number of accounts to be processed in a given time frame. But for the online sizing, the sizing exercise is based on the number of accounts present as well as the anticipated online activity to those accounts. Once known, the maximum hourly transaction volume

would be put in context to address the response time requirements of the user environment.

The net result is that the acquisition of any system today is a more complex decision which includes many tradeoffs. The first step is to size the system based on the batch processing requirements. This would result in a series of system choices. The second step is to size the system based on the online requirements. This also would result in a series of system choices. Once complete, the results would be put into a decision matrix in order to allow for discussion and analysis. Only those choices that fulfill both the batch and online requirements should be placed in the decision matrix.

Growth

One of the main issues of concern has to do with the growth of activity within the system. In the past, the only growth element was the number of accounts for the bank. This was regular and predictable.

With the advent of e-banking, the increase in the number of accounts is no longer the major concern. Banks must now deal with the increasing volume of external activity to the accounts. Growth on the system comes not only from increasing the number of accounts but the activity to those accounts. In some cases the activity is between the customer's accounts as well as new activities provided for the customer by the bank. The addition of new accounts increases both batch and online activity. However, new marketing strategies by the bank can increase the online activity to higher levels without increasing batch activity. Regardless of how you look at it, online activity is growing faster than batch activity and must be the major focus of any system sizing exercise.

Appendix A: Analysis of the 3101 Program

Introduction

Since the initial publication of the ITI white paper, the emphasis has been on the benchmark results of two vendors, IBM and Unisys. What has always been puzzling was the wide disparity of processing between the two systems. While the exact contents of the ITI benchmark is not known, a good guess can be made and has been used in the preparation of this study.

Due to the time and data constraints in the initial white paper, it was not possible to investigate the disparity in system performance between the two vendors. However, for this revision the time and information was available.

As stated in this white paper, it was assumed that the Primary Update processing consisted of those programs that were in the 3000 to 3999 range. Of these programs, the 3101 update program was the primary update program for the major business functions (COD, DDA, FMS, LAS, and SAV), and was also the program that was most disruptive to the Primary Update processing. For this reason, the 3101 program was chosen for analysis.

Analysis

The user information for the 3101 program was extracted, summarized, and is presented in Table A1. As with other tables, the information is presented based on the processing related to 1,000 accounts. Due to the nature of the information, it is necessary to explain the values in the table and what they represent.

For the Unisys systems, the values represent the number of physical I/Os performed for the related

	I/Os per 1000 Accounts				
	Total	Read	Write	% Reads	% Writes
Cert. of Deposit (COD)					
Unisys	3,282	2,464	818	75.1%	24.9%
IBM	4,268	4,227	41	99.0%	1.0%
Demand Deposit (DDA)					
Unisys	7,832	4,692	3,140	59.9%	40.1%
IBM	12,815	12,644	172	98.7%	1.3%
Loans (LAS)					
Unisys	4,156	2,908	1,248	70.0%	30.0%
IBM	4,024	3,893	131	96.7%	3.3%
Savings (SAV)					
Unisys	3,072	1,939	1,132	63.1%	36.9%
IBM	3,414	3,375	39	98.9%	1.1%

Table A1: Program 3101 I/O Summary (Base)

program processing. These values include both reads and writes, but do not include any reads that may have occurred within the same buffer (logical reads). It has been stated by ITI that the 3101 program processing is highly random and blocking records would have very little effect on the processing. Therefore, the table values for the Unisys systems should closely represent the actual data accesses as seen by the program processing. For example, the COD3101 program will perform three reads for every write for each account during its processing.

For the IBM systems the situation is a bit more complex. First of all, the IBM systems inherently use DB2 for the underlying file system, so the ITI application does not directly control the physical I/Os as on the Unisys systems. Any physical I/Os that occur on the IBM systems are done under the control of DB2. What this means is that the comparison of physical I/Os between the two vendors is out of the question.

However, since the the IBM system uses DB2 for the underlying file structure, it is possible to analyze the activity to the database for the program processing. Since the Unisys processing is using one physical I/O to read or write data, then the IBM system would be doing a similar database I/O to read or write the data. Therefore, the values presented for the IBM systems are the number of database I/Os performed for the related program processing.

As was stated earlier for the Unisys systems, the assumption was little (or no) secondary accesses (logical reads) to the already loaded buffer. Since this is surely not the case, an adjustment should be made for the IBM systems to account for this fact. After all, we are counting all database accesses for the IBM systems and only the physical accesses for the Unisys systems. The adjustment to the table is simply restating the IBM reads to be the same number as the Unisys systems. When these changes are made and applied to Table A1, the resulting Table A2 will contain the new values. (It should be noted that some of the physical writes on the Unisys systems will first require physical reads to retrieve the data, but this consideration is waived for this discussion).

Conclusions

As can be seen in Table A2, the percent of reads and writes are highly disproportionate between the two vendors. The Unisys systems must perform a much higher number of writes as compared to the IBM systems, and this disparity is present across all business functions.

Looking at the table, one would draw the conclusion that the DDA3101 program on the IBM system will perform 172 data updates (writes) for every 1,000 accounts. Since the assumption is that the processing is the same for both vendors, then one would expect that on the Unisys system, the DDA3101 program would do a similar amount of writes. But it doesn't. The Unisys processing must do 3,140 writes. What is the purpose of the other 2,968 writes?

One can speculate as to the purpose of these writes, but they are most likely due to some legacy processing that was not migrated to the IBM system. Something like a recovery scheme for system interruptions comes to mind. Regardless, it is not important. What is important is that the Unisys systems are doing something the IBM systems are not. The requirement for the Unisys processing to perform the additional writes not only inhibits good

	I/Os per 1000 Accounts				
	Total	Read	Write	% Reads	% Writes
Cert. of Deposit (COD)					
Unisys	3,282	2,464	818	75.1%	24.9%
IBM	2,505	2,464	41	98.4%	1.6%
Demand Deposit (DDA)					
Unisys	7,832	4,692	3,140	59.9%	40.1%
IBM	4,864	4,692	172	96.5%	3.5%
Loans (LAS)					
Unisys	4,156	2,908	1,248	70.0%	30.0%
IBM	3,039	2,908	131	95.7%	4.3%
Savings (SAV)					
Unisys	3,072	1,939	1,132	63.1%	36.9%
IBM	1,978	1,939	39	98.0%	2.0%

Table A2: Program 3101 I/O Summary (Normalized)

program performance, but imposes an unwarranted overhead on the system that impacts all program processing.

Data can be accessed in one of three ways: from memory, from an existing file buffer or cache (in memory), or from a disk unit. In the first two cases, the processing is optimized by the system and the overhead incurred by the program processing is almost nil. In the event that a program must do a physical I/O, the processor and system cost is very high (relatively speaking). The I/O control structure must be built, the I/O initiated, the I/O processed by the disk unit, and the I/O must be finished. It is relatively expensive, and this does not include the cost of suspending and resuming a program. If a program does a lot of physical I/Os, then program processing becomes chaotic and the overheads begin to add up quickly.

For the 3101 program processing, the Unisys system is already "hamstrung" because all the data retrieval is done via the existing file mechanisms. This means all data accesses on the Unisys systems are physical I/Os. To contrast this, all data accesses on the IBM systems are DB2 accesses which are done in memory. The IBM systems are further helped by the fact that a lot of the database information is "pre-loaded" into memory and the cache during the online processing day. (Note: What is done for the ITI benchmark is not known). The Unisys systems do not have the advantage of "pre-loading" data, they must perform physical I/Os in lieu of memory access, and they are required to do additional physical writes that the IBM systems do not have to do.

It might be valid to say that the 3101 processing for both vendors has the same net result. However, it is not valid (nor fair) to compare the physical processing for the two vendors. The 3101 program processing is different for each vendor and does not allow for a legitimate "hardware" comparison between the two.

One final point must be made. While it is not possible to circumvent the use of DB2 on the

IBM systems, it does seem possible to make programing changes to the Unisys processing. These changes would result in processing on the Unisys systems that is similar to that found on the IBM systems. These changes do not require the use of DMSII (the Unisys database), and can be implemented within the standard programming environment. These changes would mitigate the processing differences and allow for a more balanced comparison of the physical characteristics of each vendor.