Redpaper

Mike Ebbers
Gene Neyer
Tangy Ward

A Banking Application: Fundtech Global PAYplus Test Results on IBM POWER5+

A group of United Kingdom banks made plans to reduce payment times of transactions (involving telephone or Internet instructions from customers) from days to seconds. The Faster Payments Initiative is a proposal to update their payment systems to enable connection with the new central infrastructure. They plan to implement this very shortly. Fundtech ran a series of benchmark tests to demonstrate that their Global PAYplus application can handle the high transaction volumes required.

Introduction and management summary

In May 2005, the Payment Systems Task Force, chaired by the Office of Fair Trading, announced that an agreement had been reached to reduce clearing times on electronic payments among UK banks from days to seconds for telephone or Internet payment instructions from customers, and for customer Standing Orders.

Ten banking groups have committed to implement the scheme. Among them, they represent 95% of the payments made in the UK. The other banks and building societies will be making decisions about how, when, and if to join the scheme over the next year or so.

Driving high transaction volumes: The UK Faster Payments Initiative

These ten banking groups need to update their payment systems to enable connection with the new central Faster Payments infrastructure and complete testing within a very compressed time frame. Their “go live” date has been set for November of 2007.

The Faster Payments scheme presents unique challenges to the existing payment systems in a bank. It simultaneously needs to support:

- ACH-type batch processing with volumes of several million transactions an hour during peak periods. Although processed as a batch, these transactions need to be examined individually to isolate and to process internally those transactions that are between the account holders of the bank or between the account holders of the bank and its sponsored agencies.
- ATM-type processing with volumes of several million transactions per hour during peak periods.
- Near Real Time (NRT) processing with volumes of around one hundred thousand transactions per hour for Single Immediate transactions (SIMs). The guaranteed round-trip time will be measured in seconds.

At the request of a large financial institution, Fundtech Corporation conducted a benchmark test to demonstrate that Global PAYplus, already a best of breed platform for processing individual payments, meets the requirements of this bank and the industry for batch processing performance under the prescribed service level agreement (SLA) for Faster Payments. The performance testing was conducted in the IBM Labs in Poughkeepsie, New York.

Parties to the test

Fundtech is a leading global provider of electronic payments, cash management, and settlement solutions. Fundtech’s flagship product, Global PAYplus, is a rules-based payment system implemented on Open technology that provides a global financial institution with distributed access via browser to payment processing and centralized control of their payment operations worldwide.

IBM installed and configured the hardware and much of the software on eight 8-way p575 systems, as described in “Configuration details ” on page 4. In addition, the DS4800 storage system with four EXP810 drawers was installed and configured.

Windows® XP was installed on the five x335 systems that were used to generate the test load. IBM provided ongoing system administration support across the test series and also MQ consultation. Oracle® RAC support was also provided to support the benchmark center’s need to use eight separate systems rather than a single system. IBM responded promptly to all problems or changes that were required to ensure the success of this benchmark.

Summary of test objectives

At the request of a large bank, Fundtech Corporation conducted a performance test to demonstrate that Global PAYplus (GPP) could meet the bank's requirements for performance, as described in the service level agreement (SLA) for the Faster Payments Initiative.

The bank defined the objectives and success criteria that Fundtech must demonstrate to ensure that GPP was a workable solution:

- Demonstrate a throughput of 300 TPS on any one of the High Volume (STO) business flows that model FPS. This is 10% more than the expected operating requirement in year 2007. Fundtech extended this objective to 600 TPS, which corresponds to the operating requirement in the year 2012, assuming CAGR of 15%.
- Demonstrate the ability to scale.
- Demonstrate the performance on the IBM System p™ hardware and software platform.

To give the process focus and direction, the objectives were prioritized into Gold, Silver, and Bronze. The objectives defined by this bank were classified as "Gold". There were other objectives defined by Fundtech that were classified as "Silver" and "Bronze". Figure 1 summarizes these levels.

Standard	Standard Definition
Gold	Fundtech and GPP MUST meet the success criteria for this type of objective.
Silver	Fundtech and GPP SHOULD meet the success criteria for this type of objective after the gold objectives are achieved. In the event of failure, Fundtech must demonstrate to the bank's satisfaction a reasonable explanation of the failure and a resolution path.
Bronze	As part of the test, Fundtech and GPP SHOULD meet the success criteria for this type of objective, time permitting. Bronze standard objectives were deemed "nice to have."

Figure 1 Explanation of levels of objectives

Summary of test results

The objectives and the outcome of each test are summarized in Figure 2.

Objective ID	Objective	Standard	Venue	Outcome	Supporting test runs
O1	Demonstrate 600 TPS on one high volume business flow	Gold	IBM Labs	Passed	All
O2	Demonstrate FPS style processing	Gold	IBM Labs	Passed	All
O3	Demonstrate ability to scale one aspect of processing	Gold	IBM Labs	Passed	W1-W3
O4	Demonstrate 600 TPS on both inbound and outbound flows	Silver	IBM Labs	Passed	All, except for W4, W8
O5	Collect the metrics for HW sizing appropriate to bank volume	Bronze	IBM Labs	Passed	W1-W8

Figure 2 The overall test objectives

Summaries produced by Fundtech for each test undertaken can be found in “Test runs: Summary” on page 6. In addition, Fundtech provided reports capturing test results, system configuration, and where appropriate, follow-up actions.

Conclusion

We are able to report complete success in terms of the criteria set down by a large financial institution prior to the commencement of the test (the gold standard objectives). We believe that Global PAYplus has proven to be an exceptional platform on which the bank can underpin its Faster Payments solution as well as the future strategic payments solution.

Situation and environment

This section describes our test environment, including hardware, software, and testing schedule and guidelines.

Configuration details

This section lists the hardware and the major software subsystems that we used.

Server type: IBM eServer p575

The hardware included:

- 8 CPUs
- 32 GB
- 1.9 GHz

Partitions

The partitions were:

- Application server (LPAR 1-5, r85n01 - r85n05):
  - Host name r85n01.pbm.ihost.com- r85n05.pbm.ihost.com (129.40.17.211-215)
  - Number of processors 8 (SMT-enabled)
  - Processor speed 1.90 GHz
  - Amount of memory 32 GB
  - Operating system information:
    - AIX® 5300-04-03
    - 64-bit kernel
    - Mirrored rootvg
    - Paging space 25 GB
  - Additional software installed:
    - WebSphere® 6.1.0, installed on r85n05
    - MQSeries® 5.3 CSD9
    - Java™ 1.4.2 (32-bit and 64-bit)
    - Oracle 10g client
    - BEA Tuxedo 8.1
    - GPP 3.0

- GPP simulators
- Network adapter information:
  - r85n01 - r85n05 129.40.17.211 - 129.40.17.215 en0 1 GB
  - r85g01 - r85g05 129.40.33.211 - 129.40.33.215 en1 1 GB
  - r85p01 - r85p05 129.40.47.211 - 129.40.47.215 en2 1 GB
- Additional filesystems:
  - /export/home 15 GB
  - /MQRaw1 20 GB (128k striped across hdisk0, hdisk1)
  - /var/mqm 5 GB (128K striped across hdisk2, hdisk3, hdisk4, hdisk5)
  - /u01 6 GB (256k striped across hdisk2, hdisk3)

- Database Server (LPAR 6-7, r85n06 - r85n08):
  - Host name r85n06.pbm.ihost.com- r85n08.pbm.ihost.com (129.40.17.216-218)
  - Number of processors 8 (SMT-enabled)
  - Amount of memory 32 GB
  - Operating system information:
    - AIX 5300-04-03
    - 64-bit kernel
    - Mirrored rootvg
    - Paging space 30 GB
  - Additional software installed:
    - Oracle 10g
    - Java 1.4.2 (32-bit and 64-bit)
  - Network adapter information:
    - r85n06 - r85n08 129.40.17.216 - 129.40.17.218 en0 1 GB
    - r85g06 - r85g08 129.40.33.216 - 129.40.33.218 en1 1 GB
    - r85p06 - r85p08 129.40.47.216 - 129.40.47.218 en2 1 GB
  - Additional filesystems:
    - /export/home 15 GB
    - /u01 15 GB hdisk2 mirrored hdisk3
    - Additional filesystems
  - External disk: One DS4800 Storage Subsystem with 4 enclosures each with 16 73 GB 15k RPM disks:
    - 13 RAID 10 Arrays (4 disks in each array):

      - 24 20 GB LUNS for Oracle Data use

      - 5 300 MB LUNS for Oracle ASM Administrative use

Suggested production hardware layout

Based on the results of the tests, no changes are recommended to the hardware configuration used. This is detailed in “Appendix B: Physical configuration” on page 18.

Test description

The test cycle was conducted in three phases:

1. Internally to Fundtech, where the majority of the changes to the code base and initial tuning were performed. This phase lasted from September 4th to September 20th.
2. Internal tests at the IBM labs in Poughkeepsie, NY, where the initial emphasis was on scaling the system to a 4 Node (32 CPU) configuration and then to an 8 node (64 CPU) configuration. This phase started on September 27th and concluded on October 20th.
3. Witnessed tests at the IBM labs in Poughkeepsie, NY, conducted on October 16th and 17th. A large UK financial institution provided a team of engineers and architects to review and witness the technical and business aspects of the tests.

IBM provided all of the equipment for the test and assigned their staff to review the results.

The static data in the test system is configured as shown in “Appendix A: Test background” on page 17. The test data was staged and injected into the test system at the maximum rate (simulating a “flood” condition).

The tests are cumulative; that is, the database is not reset and transactions continue to accumulate from test to test. At the conclusion of the test, the database contained approximately 14 million transaction records.

Individual tests follow the same basic execution pattern:

1. Process the inbound transactions as shown in Appendix C.
2. Process the outbound transactions as shown in Appendix C.

Only the number of processing nodes was varied and tuning parameters were changed between the runs.

The test was considered started at the time of the earliest time stamp and was considered concluded at the time of latest time stamp.

Test runs: Summary

The results met and exceeded the objectives and are listed in the tables below. In summary:

- Outbound transactions per second were generally above the 1,000 range, with a peak of 1,783. Inbound transactions per second were well above 1,000, with a peak of 2,707. There were some exception cases, which would need attention, and the entire configuration would require some adjustments and tuning when the test system from the FPI scheme is made available. These results significantly exceed bank objectives.
- Outbound processing scaled nearly linearly in the number of the application server nodes. The variations around 575 TPS/node have more to do with the interaction between outbound and inbound processing than with the increased resource utilization and locking having to do with adding nodes. This result matches bank objectives.
- Inbound processing scalability has not been conclusively demonstrated due to the interaction of inbound and outbound processing. The quanta are the maximum number of records that will be read and processed at one time. Although no effect was observed at quanta of 20,000; the scaling is better than linear at quanta of 2,000. This provides sufficient evidence that both processing paths are scalable in the number of nodes.

Pre-witness test: October 13

Here are the results of the preliminary tests that Fundtech ran to assess the situation. The information is separated into two tables for readability. Table 1 summarizes test conditions; Table 2 summarizes the test results.

Table 1 Test conditions

	#Nodes							
Test Code	Out	In	Quanta – in	Start Time	End Time	Duration	# msg - out	# msg - in
A	2	2	20,000	11:29:32	12:29:35	60:03 min	3,420,000	3,680,000

Table 2 Test results

	Outbound TPS			Inbound TPS			App server CPU util.		Database CPU util.	
Test Code	Avg	Peak	Trough	Avg	Peak	Trough	Average	Peak	Avg	Peak
A	1,675	5,089	428	1,326	3,000	-	85%		65%	

Witnessed tests: October 16-17

These were the test runs where everyone was watching. The information is separated into two tables for readability. Table 3 summarizes test conditions; Table 4 on page 8 summarizes the test results.

Table 3 Test conditions

	#Nodes							
Test Code	Out	In	Quanta - in	Start Time	End Time	Duration	# msg - out	# msg - in
W1	2	2	20,000	14:38:30	14:48:30	10:00 min	570,000	530,000
W2	2	3	20,000	15:03:45	15:13:00	09:15 min	570,000	525,000
W3	1	3	20,000	15:30:15	15:39:00	08:45 min	285,000	510,000
W4	3	1	20,000	19:06:30	19:17:30	11:00 min	855,000	445,000
W5	2	2	2,000	9:25:30	9:48:15	22:45 min	570,000	630,000
W6	2	3	2,000	12:40:45	12:51:15	10:45 min	570,000	627,531
W7	2	3	2,000	13:21:00	13:35:00	14:00 min	570,000	630,000
W8	2	2	200	14:14:30	14:35:00	20:30 min	570,000	630,000

Table 4 Test results

	Outbound tps			Inbound tps			App server CPU utilization		Database CPU util.	
Test Code	Avg	Peak	Trough	Avg	Peak	Trough	Avg	Peak	Avg	Peak
W1	1,056	1,273	294	1,536	2,048	66	43.7	97.9	n/a	n/a
W2	1,205	1,278	945	978	2,193	53	34.6	97.7	44.9	87.5
W3	612	652	296	1,013	2,178	19	35.5	95.8	41.3	81.7
W4	1,693	1,831	999	711	845	490	25.6	97.4	n/a	n/a
W5	1,146	1,260	121	462	572	226	36.4	100	41.1	93.8
W6	1,037	1,262	40	980	2,017	327	56.5	99.5	35	79.7
W7	1,178	1,258	385	754	1,267	332	51.3	99.9	43.5	93.8
W8	1,086	1,289	18	467	561	289	51.1	99.9	40.9	85.1

Scalability extracted from the above tables

In Table 5, the O/I TPS or OTPS column lists the total outbound/inbound TPS for the solution. The O/I TPS/Node or OTPS/Node column is the number of outbound/inbound TPS processed by an individual node.

Table 5 Outbound processing

#Nodes Out	OTPS	OTPS/Node	% decrease
1	612	612	
2	1118	559	9%
3	1693	564	-1%

The variations around 575 OTPS/node have more to do with the interaction between outbound and inbound processing than with the increased resource utilization and locking having to do with adding nodes.

See Table 6 on page 9 for inbound processing scalability.

Table 6 Inbound processing

#Nodes In	ITPS	ITPS/Node	%	Quanta
1	711	711		
2	978	489	31%	
3	996	332	32%	20000
1				
2	422	211		
3	867	289	-37%	2000
1				
2	401	201		
3				200

Inbound processing scalability has not been conclusively demonstrated due to the interaction of inbound and outbound processing. Although no effect was observed at quanta of 20,000, the scaling is apparent at quanta of 2,000 and is actually better than linear. This provides sufficient evidence that both processing paths are scalable in the number of nodes.

Impact of quanta on TPS and resource utilization

These tests were done in the four application server configuration: two inbound processors and two outbound processors.

The quanta are the maximum number of records that will be read and processed at one time. Thus, if the quanta are set at 200, if the MQ depth is 200 or more, 200 records will be processed in a single commit. However, if the queue depth is less, then the read will not block. The available number of transactions will be processed. See Table 7.

Table 7 Impact of quanta on TPS and resource utilization

		Inbound App Server - CPU utilization - averaged over two servers		Inbound App Server - I/O rates - averaged over 2 servers		Database - CPU utilization		Inbound TPS
Test Code	Quanta - in	Avg	Peak	Avg	Peak	Avg	Peak	Avg over the period
W1	20,000	85%	98%	900	1,500	n/a	n/a	1,536
W5	2,000	82%	100%	350	625	41%	94%	462
W8	200	80%	100%	250	775	41%	85%	467

The results of these test runs show that decreasing the quanta by the factor of 10 decreased the TPS and I/O by 50%. Decreasing the quanta by a factor of 10 again had negligible impact on the throughput of the inbound processing (in fact, it increased slightly) and decreased the local I/O by another 30%.

As you can readily see from the graphs below, at a quanta of 2,000, there is still a fair amount of variability in the processing rate. That is not an issue for the batch processing. We have used the opportunity to vary the quanta to determine the appropriate setting for processing a high volume of Single Individual Messages. With the quanta of 200, the processing is very stable around mean, which will result in predictable response times within the SLA (see Figure 3).

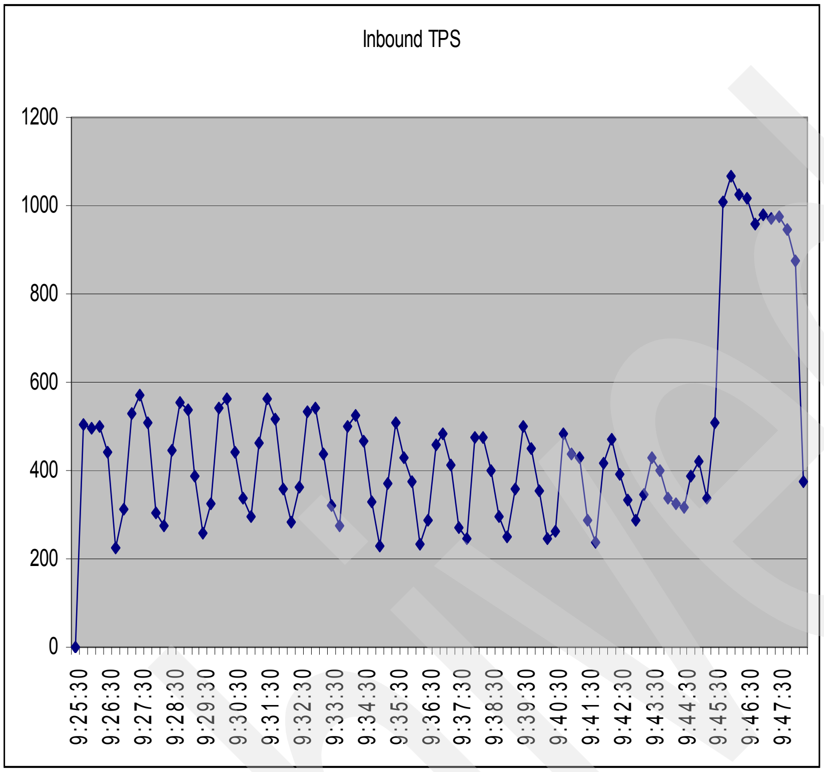

Figure 3 Inbound TPS for a quanta of 2,000

See Figure 4 on page 11 for Inbound TPS for a quanta of 200.

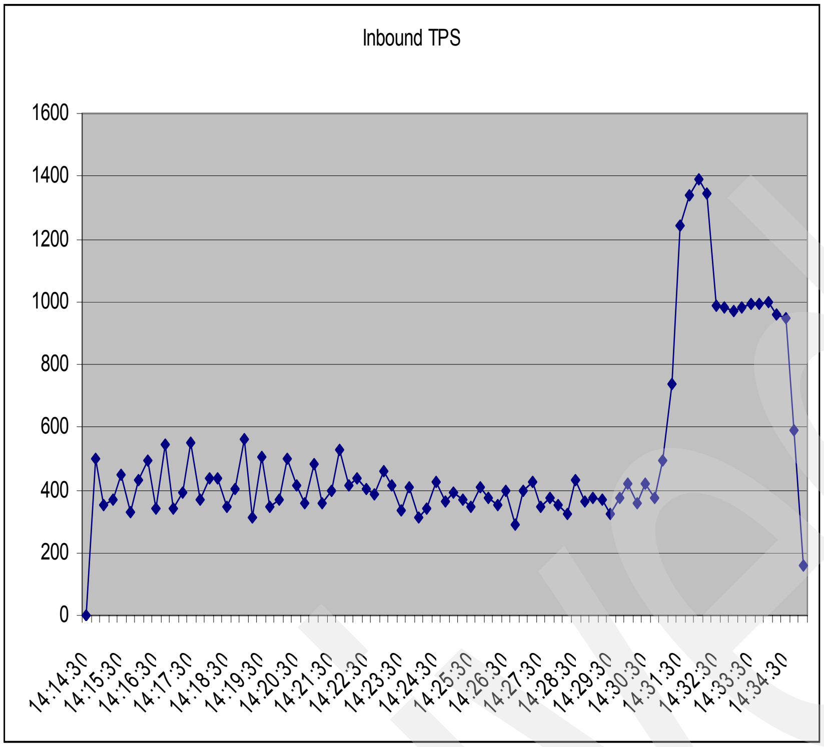

Figure 4 Inbound TPS for a quanta of 200

In the last two to four minutes of the test, the processing spikes and then tails off as the outbound processing and then inbound processing winds down. Presented here for completeness, this period should be disregarded during analysis and planning.

Findings and recommendations

The findings and recommendations are:

1. The key finding is that a combination of distributed (cluster) architecture and GPP software is very well suited for the demands of the Faster Payments Scheme, which combines the elements of ACH, High Value, and ATM processing.
2. The Fundtech application utilizes BEA Tuxedo software for transaction management, a non-IBM-standard software stack. Fundamentally, the Fundtech application uses Tuxedo as a “black box”, and it is transparent to the user. Fundtech did say it might be possible to use WebSphere Application Server for transaction management, but such alterations, as with any changes, would take time, effort, and expense.

   It should be noted that GPP already uses WebSphere Application Server to provide GUI services and thus is already familiar with this IBM product.
3. The reduction in quanta from 20,000 to 2,000 caused a doubling in process time. Yet when we reduced the quanta from 2,000 to 200, this had a negligible effect. Fundtech will do some more research into possible causes.
4. Although Oracle 10g RAC is recommended for the database, we all agreed that the RAC cluster should be set up in Active/Passive mode, because when it runs in Active/Active mode, the traffic over the RAC interconnect has been seen by Fundtech to impair the performance. The RAC cluster can still be used for availability in Active/Passive mode and

will still give instantaneous failover, it will just not involve the complexity of high volumes of traffic going over the interconnect to keep an Active/Active database in sync.

Fundtech did say that GPP has been designed to work with any XA compliant database (including DB2®). However, qualifying GPP operation with DB2, as with any changes, would take time, effort, and expense.

5. Because the benchmark was performed on IBM storage arrays, we recommend that the end solution is built on an IBM storage array (DS8000™). In the event of a need for performance tuning, an end-to-end IBM solution would be easier to compare with the benchmark infrastructure than anything else.

Representative test run: W2

Because all individual test runs follow the same pattern, only one run (W2) is detailed here (Table 8 and Table 9). However, this level of detail is available for all of the tests described in the tables above; contact the authors.

Table 8 Test conditions

	#Nodes							
Test Code	Out	In	Quanta - in	Start Time	End Time	Duration	# msg - out	# msg - in
W2	2	3	20,000	15:03:45	15:13:00	09:15 min	570,000	525,000

Table 9 Test results

	Outbound TPS			Inbound TPS			App server CPU utilization		Database CPU util.	
Test Code	Avg	Peak	Trough	Avg	Peak	Trough	Avg	Peak	Avg	Peak
W2	1,205	1,278	945	978	2,193	53	34.6	97.7	44.9	87.5

Transaction processing rates

Figure 5 on page 13 illustrates the transaction processing rates of the overall solution. Because 50% more capacity was devoted to the inbound processing, it completes significantly before the outbound processing completes. It should be noted that inbound processing only minimally impacts the outbound processing.

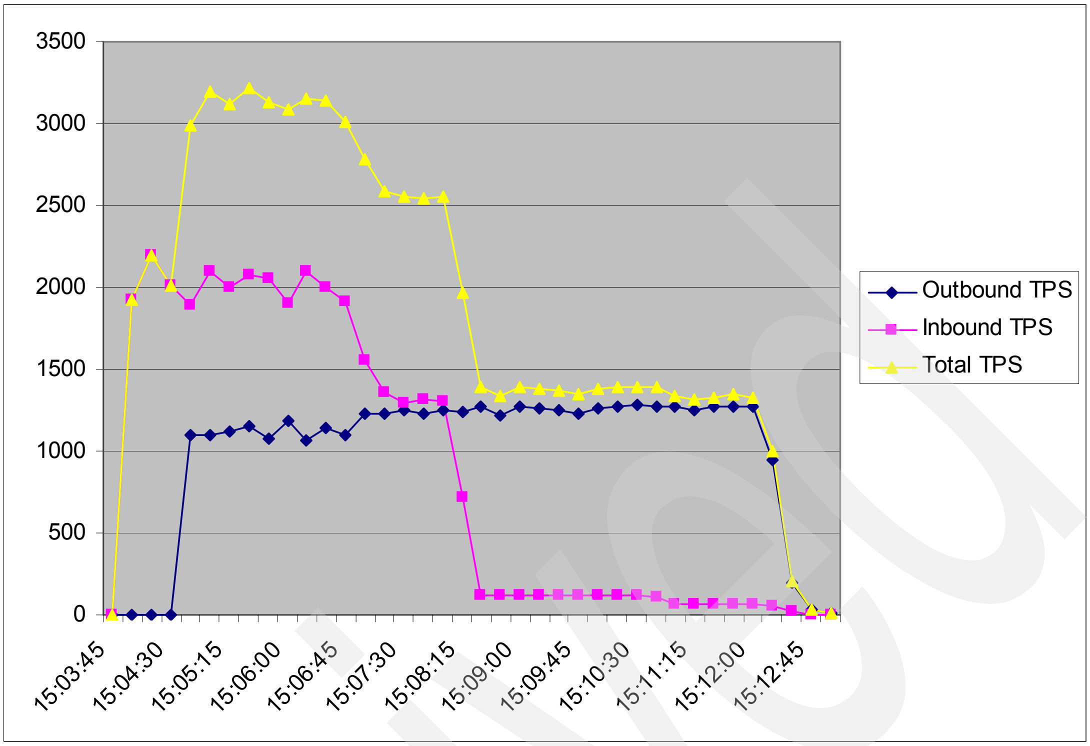

Figure 5 Transaction processing rates of the overall solution

Resource utilization rates

The following graphs (Figure 6 through Figure 11 on page 15) illustrate the CPU utilization and I/O rates of the solution. The CPU utilization and I/O are correlated to the TPS rates. No degradation of processing was observed even when CPU utilization exceeded 90%, demonstrating that GPP makes efficient and complete use of the IBM platform.

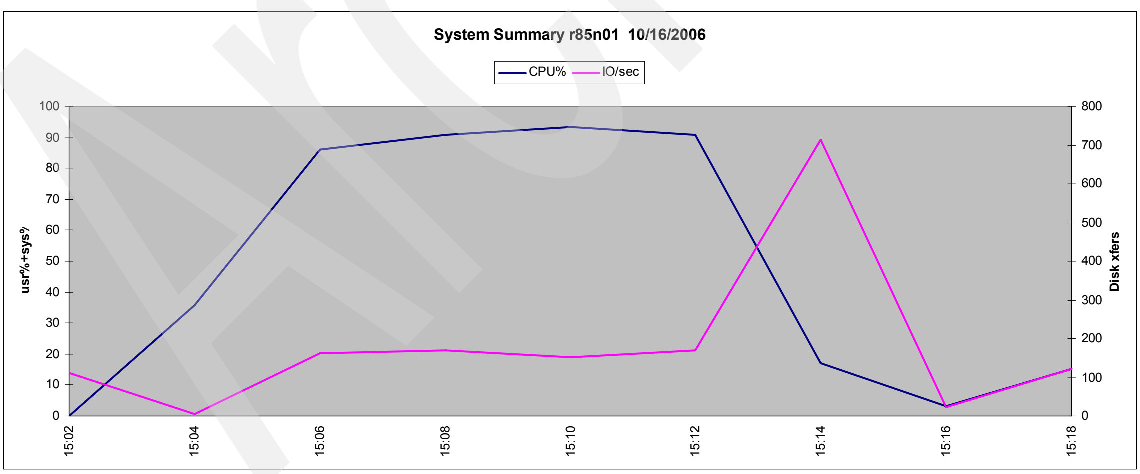

Figure 6 Graph: Outbound 1

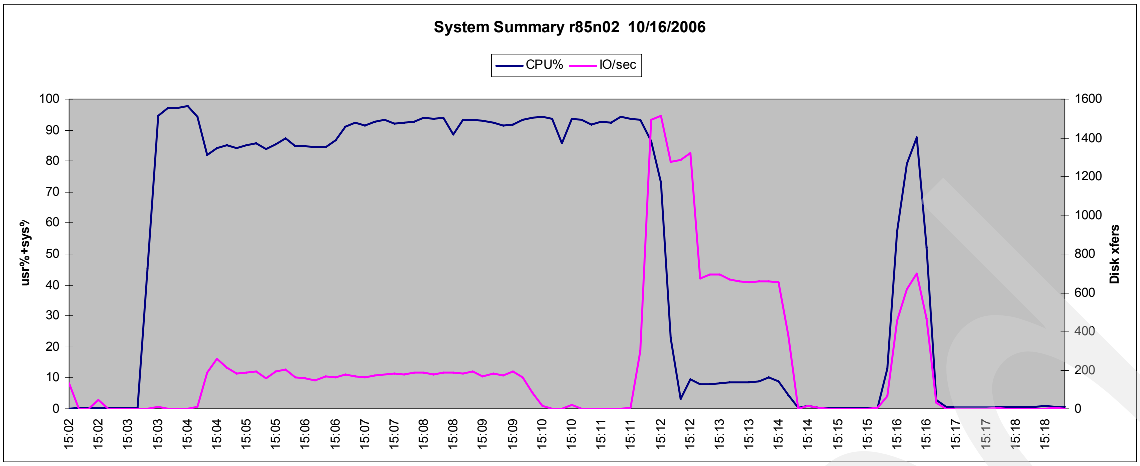

Figure 7 Graph: Outbound 2

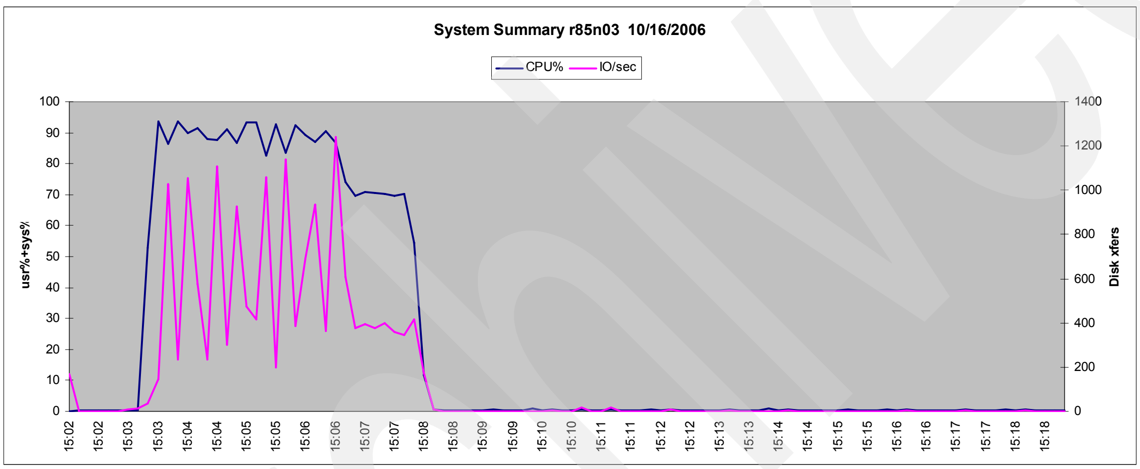

Figure 8 Graph: Inbound 1

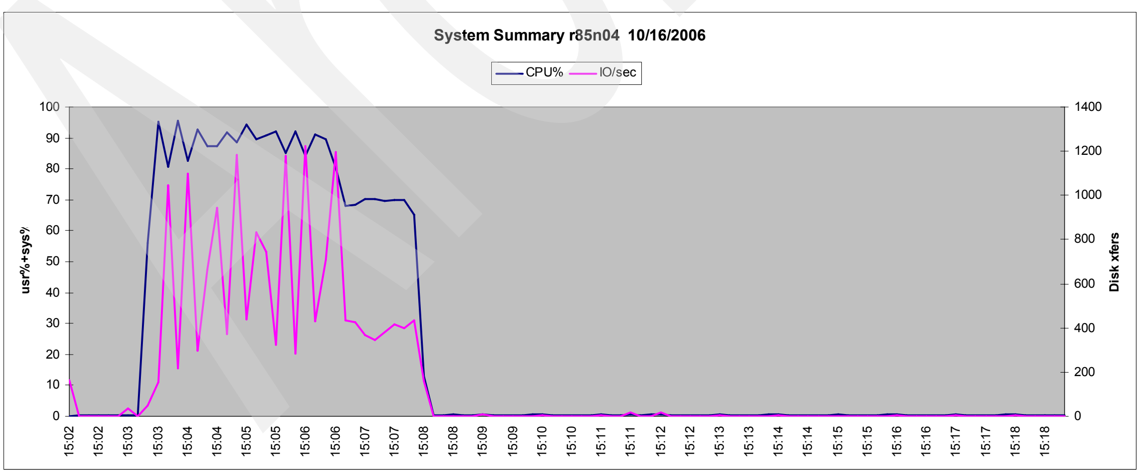

Figure 9 Graph: Inbound 2

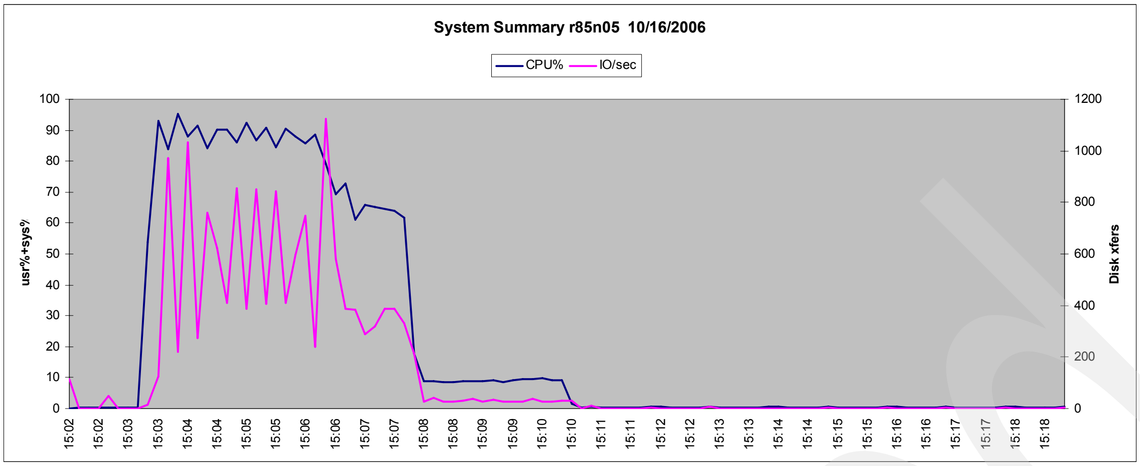

Figure 10 Graph: Inbound 3

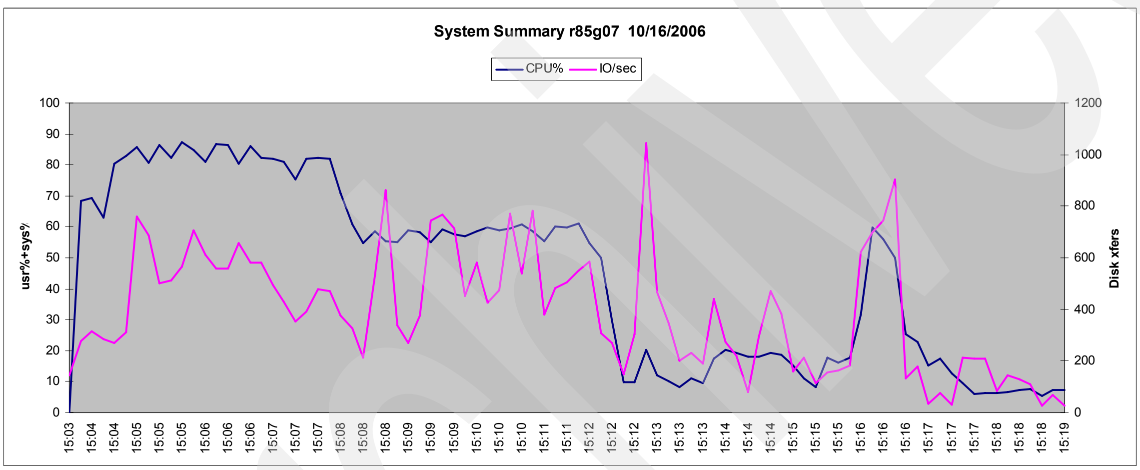

Figure 11 Graph: Database

The resource utilization profile is consistent with the application configuration and test execution:

- The objective of the test was to process transactions fully utilizing the available resources. The CPU processing is the first bottleneck that has been encountered and is addressed by horizontal scaling (adding application server nodes).
- The I/O subsystems in the application server nodes are not seen as a bottleneck. The I/O rates for outbound processing are significantly lower than those for the incoming nodes. This is consistent with the configuration: the larger the quanta (30,000 for outbound as opposed to 20,000 for inbound), the lower the expected I/O load.
- As evidenced by the resource (CPU and I/O) utilization, a single database node is sufficient to support five application server nodes and RAC replication with approximately 10 to 20% spare capacity even at peak transaction loads. This validated the database interface design. Utilizing a small number (fewer than 256) of application processes that deliver medium size transactions into the database is an effective design pattern for Oracle 10G in the high volume transaction processing scenario.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Mike Ebbers is a Consulting IT Specialist and Project Leader at the International Technical Support Organization, Poughkeepsie Center. He has worked for IBM for 33 years.

Gene Neyer is a Senior Vice President and Group Product Manager for Global Payments at Fundtech in Jersey City, NJ. He has over 20 years of experience in the Financial Industry as a product manager, management consultant, and technologist.

Tangy Ward is an Accredited IT Specialist and Project Manager at the Industry Solutions and Proof of Concept Centers Organization in Poughkeepsie, NY. She has worked for IBM for 24 years.

Thanks to the following people for their contributions to this project:

Eli Biton: VP of global R&D at Fundtech in Ramat Gan, Israel.
Robert Coakley: SVP, Head of Technology at Fundtech in Jersey City, NJ.
Avi Landau: VP Engineering Support Group at Fundtech in Jersey City, NJ.
Andriy Khrystoforov: Senior DBA at Fundtech in Jersey City, NJ.

Appendix A: Test background

Table 10 Test information

Number of customer accounts	200,000
Number of accounts exercised (by payment messages)	100,000
Number of sender parties	3
Outbound batch - Message Types	MT103
Number of unique messages	2,225,000
Responses for outbound - Message Types	Proprietary
Number of unique messages	One for each payment request
Inbound payment requests - Message Types	MT103
Number of unique messages	2,339,445
Responses for inbound - Message Types	Proprietary
Number of unique messages	One for each payment request
Clearing links: simulating queues from Back end and to VPS	25 queued per node (total of 75) 5 – Inbound payment requests 5 – ACKs generated from the inbound requests 5 – Files of outbound requests 5 – Outbound payment requests 5 – ACKs for outbound requests
Modelled FPS/"On Us" ratio	90:10
Modelled local/remote accounts ratio	100:0 [95:5 in subsequent tests]
Modelled valid/invalid accounts ratio	99:1

Appendix B: Physical configuration

Full eight-node configuration

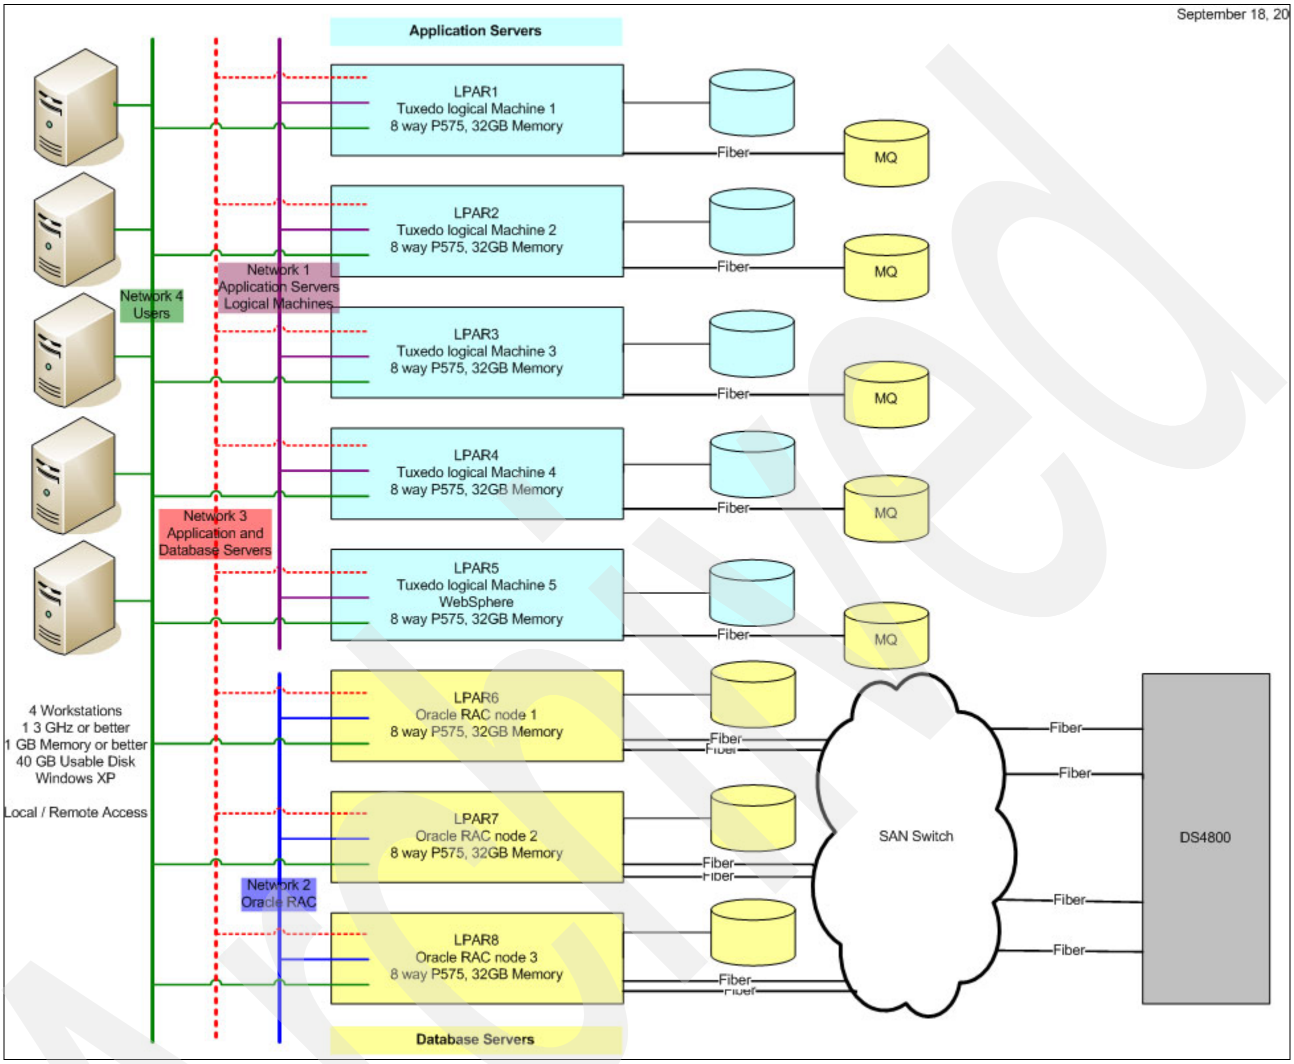

Figure 12 Full eight-node configuration

Actual use

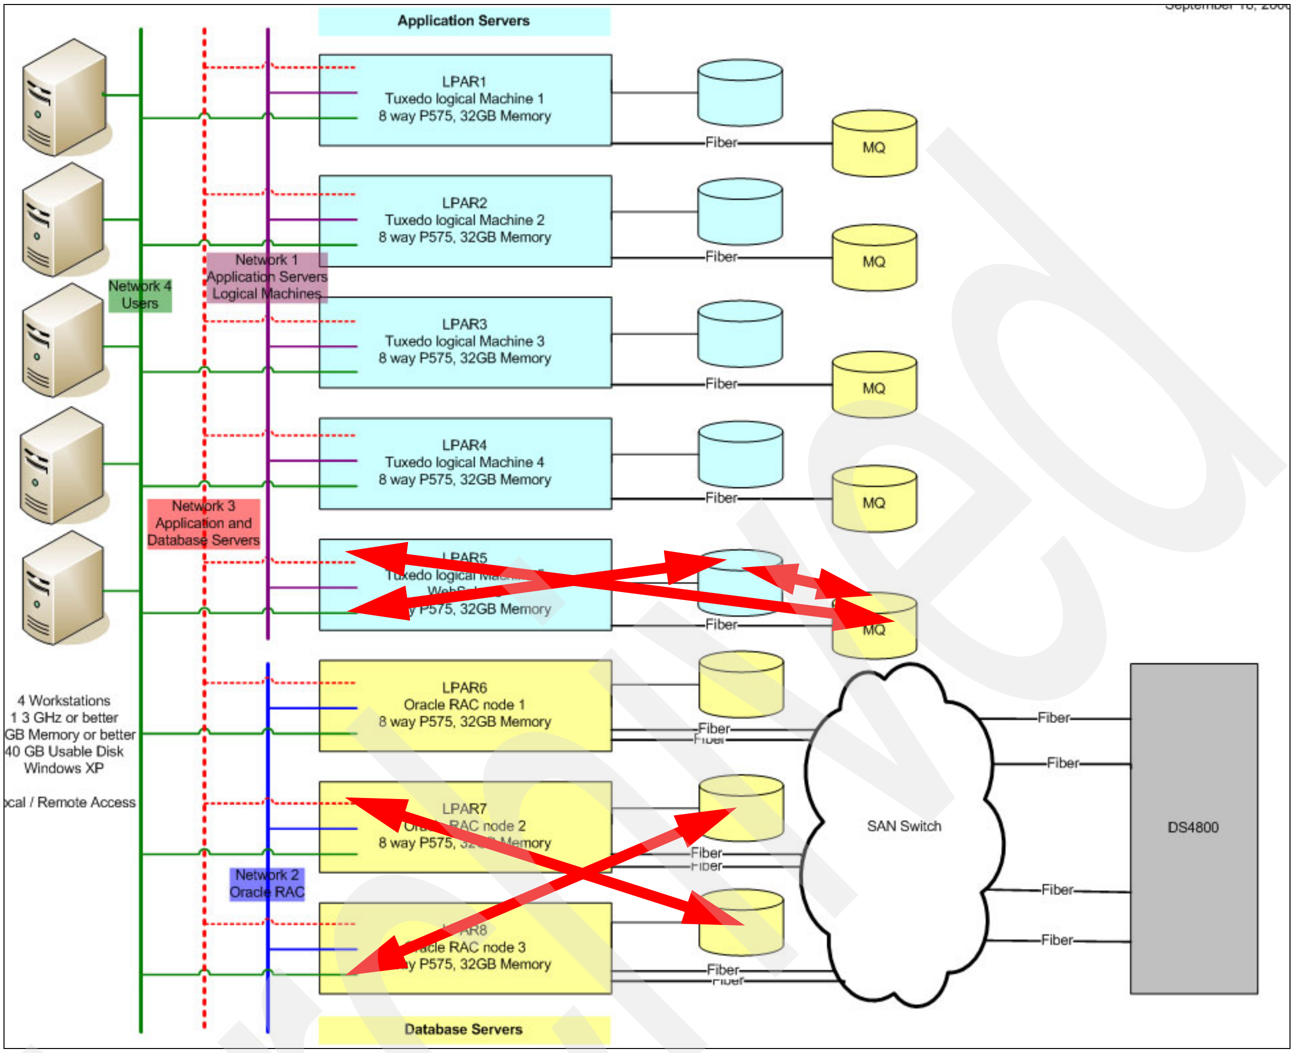

Figure 13   Diagram showing components actually used

Appendix C: Processing steps

Table 11 Outbound and inbound processes

	Outbound: Payments to be sent to CI	Inbound: Payments from CI
Process	1. Receive multiple files from RPS into GPP. De-bulk them, separate "on us" items, process FPS bound items through liquidity checks and send to the gateway. 2. Receive and apply the responses. 3. On reject responses, reverse liquidity. If reject reason is transient, requeue for processing; otherwise, reverse accounting.	1. Receive individual credits into VPS and forward them to GPP. 2. GPP will validate the beneficiary account by doing a local lookup or going to an external system, generate the appropriate response to VPS. If valid, GPP will update liquidity and post funds into account.
Simulation	1. Set up 2,250+ files of 1,000 SWIFT messages across 3 receivers with 10% to the bank's sort codes ("on us" transactions). 2.The processing will simulate: a. Channel assignment b. Validation of sort code (through the look up in the membership table) c. Checking of limits d. Checking liquidity (100% will pass) 3. The simulator sends back a positive response with FPID, qualifier code, settlement, and so forth. [Reject processing does not meaningfully affect the performance profile.]	1. Set up MQ simulator to generate 2.25+MM credits, debiting one of the sort codes, crediting one of the bank's customer accounts. 2. Look up the credit account: a. Evaluate posting restrictions and either complete (SUCCESS) or reject (POSTREST) the transaction. 3. Generate response to the MQ, which will be read and discarded by the simulator. 4. Add funds to liquidity and position.

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 illustrate 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.

Send us your comments in one of the following ways:

- Use the online **Contact us** review redbook form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD 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:

Redbooks (logo) ™	DS8000™	System p™
eServer™	IBM®	WebSphere®
AIX®	MQSeries®	
DB2®	POWER5+™	

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates.

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

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

FUNDtech© is a trademark of Fundtech LTD.

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