DB2 Family: Client/Server Performance Measurement Series
Variable Line Speeds

Document Number Version 1.1

July 31, 1996

Robert Gilles
gilles@vnet.ibm.com
GILLES at STLVM14

Todd Munk
munk@stlps.vnet.ibm.com
MUNK at STLVM14

Hugh Smith
hugh@vnet.ibm.com
SMITHHJ at STLVM14

DB2 Performance
Santa Teresa Laboratory
San Jose, California

© **Copyright International Business Machines Corporation 1996. All rights reserved.**
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Trademarks

The following are trademarks of the IBM Corporation. An asterisk indicates the trademark is registered:

- DB2 *
- Distributed Database Connection Services/6000
- IBM *
- LANStreamer
- MVS/ESA
- Netview *
- OS/2 *
- PS/2 *
- Resource Measurement Facility
- VTAM *

All other trademarks are the properties of their respective owners.

© Copyright IBM Corp. 1996

Disclaimer

References in this technical report to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program, or service.

The information contained in this technical report has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this technical report was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this technical report should verify the applicable data for their specific environment.

© Copyright IBM Corp. 1996

Acknowledgements

The following people helped make this paper possible.

- Jerry Heglar
- Mark Ryan
- Bill Wilkins

© Copyright IBM Corp. 1996

Contents

Trademarks . . . iii

Disclaimer . . . iv

Acknowledgements . . . v

1. Introduction . . . 2

2. Configuration . . . 3
Specifications . . . 3

3. Workload Description . . . 5
Transaction Characteristics . . . 5
Table Characteristics . . . 6
Locking Definitions . . . 7
Physical location . . . 7
Buffer Pool Allocation . . . 8

4. Results . . . 9
Client Equivalence . . . 9
CPU and RAM Utilization . . . 11
Network . . . 13
  NCP . . . 13
  LAN . . . 14
SQL Rates . . . 15
Transactions . . . 15
  Transaction Tx1 . . . 16
  Transaction Tx2 . . . 19
  Transaction Tx3 . . . 21
  Transaction Tx4 . . . 23
  Transaction Tx5 . . . 25
  Transaction Tx6 . . . 27
  Transaction Tx7 . . . 29

5. Conclusions . . . 32

6. Other sources of information . . . 34

© Copyright IBM Corp. 1996

1. Introduction

This technical report is one of a series of technical reports dealing with performance prediction and capacity planning issues for the DB2 family of products operating in a client/server configuration. The goal of these technical reports is to help the customer:

- Assess proposed configurations for acceptable performance
- Assess performance of current configurations
- Choose proper hardware capacity for a proposed configuration
- Choose the proper software programs for a proposed configuration
- Tune a new or current configuration

This report is unique because it describes the impacts on the numerous components involved in a client/server environment as the workload increases in that environment. Use this report as an example configuration that could be analyzed and potentially applied to other configurations. This report describes the throughput, CPU utilization, RAM utilization, network utilization and response times when PowerBuilder clients and OS/2 clients call stored procedures on DB2 for AIX and on DB2 for MVS/ESA through DDCS for AIX to perform OLTP type transactions. The network line speeds are varied to illustrate the effects of various line speeds on the transactions. Use the results from these measurements to predict performance and plan for capacity in your environment. Particularly noteworthy items are highlighted as **general considerations** in this report.

© Copyright IBM Corp. 1996

2. Configuration

This technical report provides information in regards to the following configuration:

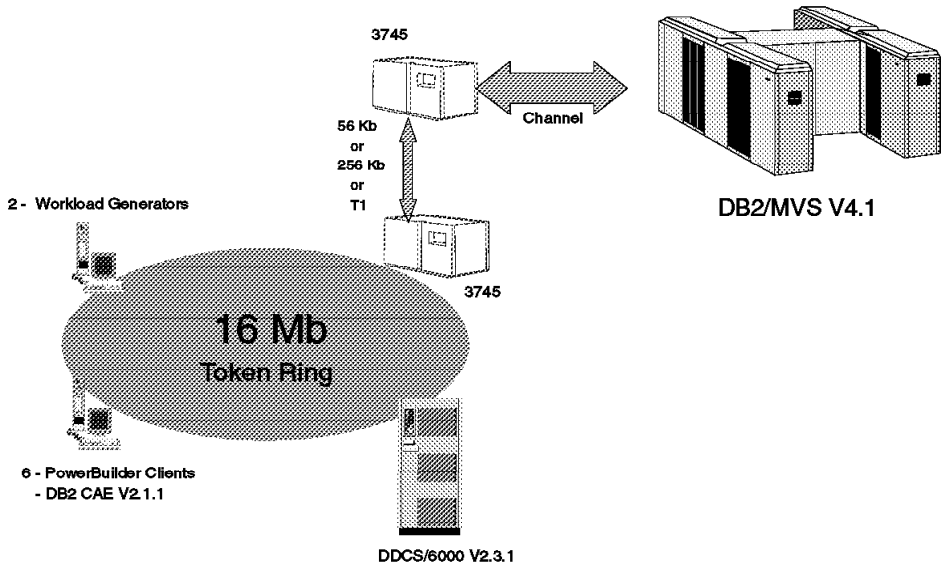

Figure 1. Configuration with DB2 for MVS/ESA, DB2 for AIX, and DDCS for AIX

Specifications

- **DB2 for MVS/ESA** Version 4.1 running on MVS Version 5.2 with VTAM Version 4.2. The processor is a 9121-742 (4-way) with 1 GByte of main storage and 1 GByte of expanded storage. The processor is partitioned into four logical partitions, with any one partition having the ability to utilize the other processors, if needed. The partition running DB2 has 576 MBytes of storage available to it. The processor is attached to a 3745 model 210 with 8 MBytes of storage on a 4.5 MByte block multiplex channel. The 3745 is attached via a NCP internally clocked SDLC link to another 3745 model 210 with 8 Mbytes of storage. This 3745 is attached to the 16 Mbit LAN via a TIC type 2. The SDLC link between the 3745s is used to measure the effects of various line speeds.
- **Distributed Database Connection Services for AIX (DDCS/6000)** Version 2.3.1 running on a C10 PowerPC 601 80 Mhz processor and 256 MBytes of memory.

© Copyright IBM Corp. 1996

- **6 - PowerBuilder Version 4.0.2 with DB2 CAE for Windows Version 2.1.1** clients running on PC350′s with 133 Mhz Pentium Processors and 32 MBytes of memory. The machines contain LAN Streamer adapter cards which are configured to use early token release.
- **2 - 9595 OPT PS/2s** with 60 Mhz Pentium Processors running OS/2 Version 3.0 (WARP). These PS/2s are used to increase the workload in the configuration.
- **16 Mbit Token Ring LAN**

Network protocols in use between the components are as follows:

- TCP/IP is in use between the PowerBuilder clients and DDCS for AIX.
- TCP/IP is in use between the OS/2 clients and DDCS for AIX.
- APPC is in use between DDCS for AIX and DB2 for MVS/ESA

The components in this configuration are monitored using the following programs:

- MVS: Resource Measurement Facility (RMF)
- DB2 for MVS/ESA: DB2 Accounting and Statistics traces
- 3745: Netview Performance Monitor (NPM)
- AIX: VMSTAT
- Clients: Internal tools to measure elapsed time and transactions per second

3. Workload Description

The IBM Relational Warehouse Workload (IRWW) is used for the measurements. IRWW is a simulated product warehouse, customer order, and order delivery system.

These measurements focus on the effects of varying the line speeds on the connection between DDCS for AIX and DB2 for MVS/ESA.

All measurements include six PowerBuilder clients running a predefined set of transactions in random order without modification for the entire series of measurements. The PowerBuilder clients provide the ″end user″ perspective of the effects of increasing the workload in the configuration. The workload is increased each run by increasing the number of sessions on the OS/2 workload generators. Each OS/2 workload generator runs the exact same transaction mix as the six PowerBuilder clients. The PowerBuilder clients simulate a real application by performing screen I/O, where the OS/2 workload generators do not perform screen I/O.

Transaction Characteristics

The workload on DB2 for MVS/ESA consists of seven transactions.

Each transaction consists of a stored procedure call followed by a commit or rollback. This results in two ″round trip″ network flows between the client and the server. The stored procedures are written in ″C″ and they perform seven distinct functions in a predefined mix. The seven transactions, a brief description, and the percentage of the mix are:

- Transaction **Tx1**

  Performs various selects, fetches, updates, and inserts in support of the receipt of new customer orders. Transaction Tx1 runs as twenty-two percent of the total transaction mix and is the most complicated transaction.

- Transaction **Tx2**

  Performs various selects and fetches in support of providing the status of an order. Transaction Tx2 runs as twenty-four percent of the total transaction mix.

- Transaction **Tx3**

  Performs various selects, fetches, updates, and inserts in support of received customer payments. Transaction Tx3 runs as twenty-two percent of the total transaction mix.

- Transaction **Tx4**

  Performs an update in support of changing the price of an item. Transaction Tx4 runs as one percent of the total transaction mix.

- Transaction **Tx5**

  Performs various selects in support of providing the price of a set of items. Transaction Tx5 runs as twenty-five percent of the total transaction mix.

- Transaction **Tx6**

© Copyright IBM Corp. 1996

Performs various selects in support of providing the current stock level of an item. Transaction Tx6 runs as four percent of the total transaction mix.

- Transaction **Tx7**

  Performs a join and various selects, updates, and deletes in support of the delivery of a group of orders. Transaction Tx7 runs as two percent of the total transaction mix.

These transactions are described in greater detail later in this report.

Table Characteristics

Before each measurement run, the tables on DB2 for MVS/ESA are put into a consistent state as defined in Table 1. Table 1 shows the initial number of rows at the start of a measurement run, whether the number of rows is fixed or will grow throughout the measurement run, the number of columns, the number of indexes, the preload percentage, and whether the table is partitioned into multiple parts so the table can be spread out to multiple physical disk drives. The preload value is the percentage of the table that is preloaded into memory prior to the measurement run. The transactions manipulate these tables in various ways and they are designed to have a high degree of contention on some of the tables, such as the T2 table.

Table 1. Table Characteristics for DB2 for MVS/ESA

Table	Initial Rows (length)	Row Count	Columns	Indexes	PreLoad	Partitioned
T1	0 (Fixed 62)	Variable	8	0	0	No
T2	8 (Fixed 103)	Fixed	9	1	100	No
T3	80 (Fixed 110)	Fixed	11	1	100	No
T4	72,000 (Fixed 22)	Variable	3	1	12	No
T5	100,000 (Fixed 91)	Fixed	4	1	100	No
T6	240,000 (Fixed 42)	Variable	8	2	12	No
T7	240,000 (Variable, 679 max)	Fixed	21	2	12	Yes
T8	800,000 (Fixed 322)	Fixed	17	1	20	Yes
T9	2,399,791 (Fixed 75)	Variable	10	1	12	Yes

Locking Definitions

Type 2 indexes are in use for all indexes on DB2 for MVS/ESA. Row level locking is in use for the table T4 and page level locking for all the rest. T2 and T3 tables have one row per page, with the net equivalence of row level locking for those two tables. All packages are bound with an isolation level of cursor stability. Tablespace locks are set for use/deallocate.

Physical location

The configuration on DB2 for MVS/ESA is tuned by spreading out the data and indexes to multiple unique disk drives. Also, T7, T8, and T9 tables and indexes are additionally spread out by partitioning to unique drives. This is summarized in Table 2. Fictitious labels are used to illustrate the relationships between the physical locations of the tables and indexes. For example, if two tables are on DISK1, they share the same disk drive.

For general consideration, it is best to spread the tables and indexes out to unique drives and controllers to reduce I/O contention to the drives.

For the configuration in this report, there is some overlap of tables and indexes. In some cases, such as the T2 and T3 tables, the tables are 100% preloaded into memory, therefore, I/O contention to these drives is reduced.

Table 2. Physical Storage Characteristics

Drive	Controller	Drive type	Table/Indexes
DISK1-8	CNTRL1	3390-3	T7 Table (Partitions 1-8) T9 Table (Partitions 1-8)
DISK9-16	CNTRL2	3390-2	All indexes
DISK17-24	CNTRL3	9394-1	T8 Table (Partitions 1-8)
DISK25-32	CNTRL4	3390-3	T8 Table (Partitions 9-16)
DISK33	CNTRL5	3390-3	T1 Table T2 Table T3 Table
DISK34	CNTRL5	3380-E	T6 Table
DISK35	CNTRL5	3380-E	T4 Table
DISK36	CNTRL6	3380-E	T5 Table
DISK37-44	CNTRL6	3380-E	T7 Table (Partitions 9-16) T9 Table (Partitions 9-16)

Buffer Pool Allocation

The tables and indexes are also spread out to multiple buffer pools as described in Table 3 on page 8. Spreading the tables and indexes out to multiple buffer pools allows for buffer pool tuning based on the access characteristics of the tables and indexes. See the DB2 Administration Guide for a discussion of buffer pools.

Table 3. Buffer Pool Characteristics

Buffer Pool	Size	Table/Index
BP0	2000	System defaults
BP1	2000	T1 Table
BP2	3000	T2 Table and Index T3 Table and Index T5 Table and Index
BP3	6000	T7 Index 1 T7 Index 2 T8 Index
BP4	7000	T4 Index T9 Index T6 Index 1 T6 Index 2
BP5	4000	T4 Table T6 Table
BP6	20000	T9 Table
BP7	12500	T7 Table
BP8	20000	T8 Table
Total	76500	

4. Results

The next sections describe the results from the measurements. The effects on each component in the configuration are described as well as the performance of each transaction. The component results can help you understand hardware capacity issues. The transaction results can help you understand the perceived performance prediction issues.

The results are primarily focussed on the effects of varying the line speeds while accessing DB2 for MVS/ESA. The results are presented as three measurement sets. Each set holds the line speed constant and scales the workload on DB2 for MVS/ESA. The measurement sets and line speeds are:

- Measurement set A - 56 Kbps via external modem clocking SDLC link (56 Kbps)
- Measurement set B - 245.760 Kbps via 3745 internal clocking SDLC link (256 Kbps)
- Measurement set C - 1,474.560 Kbps via 3745 internal clocking SDLC link (T1)

The results are presented in relation to the overall throughput achieved at each workload level. This is described as transactions per second (TPS). As the workload increases, the throughput to the database increases. The resultant effects on each component is then presented based on the current achieved throughput. For example, when achieving 65.06 transactions per second on DB2 for MVS/ESA:

- the transfer rate through the 3745 running the NCP is 62,685 bytes per second
- the C10 running DDCS for AIX is 73 percent utilized with 20 percent user and 53 percent system
- etc...

For two of the measurement sets, the constraint is the line speed capabilities. For one of the measurement sets, the line speed has capacity, but lock and latch contention on DB2 for MVS/ESA is the constraint.

Client Equivalence

The configurations include six PowerBuilder clients and a variable number of OS/2 clients. All clients run the same transaction mix, but the PowerBuilder clients present the output data on the computer monitor. This adds to the response time, which decreases the throughput for the PowerBuilder clients. This section describes the formula used to estimate the number of PowerBuilder clients based on overall throughput.

The equation for estimating the number of PowerBuilder clients is:

© Copyright IBM Corp. 1996

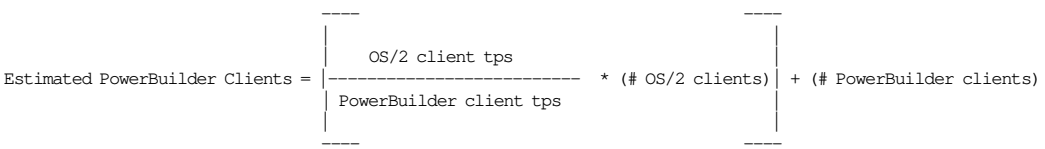

If
OS/2 client tps = 1.282 tps
PowerBuilder client tps = 0.157 tps
OS/2 clients = 50
PowerBuilder clients = 6
then

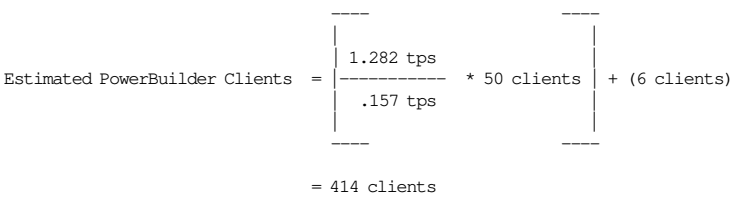

Table 4, Table 5, and Table 6 display the estimated PowerBuilder client equivalence for the three measurement sets. These estimates are based on throughput of the PowerBuilder clients and the OS/2 clients.

Table 4. Number of clients in relation to aggregate throughput (TPS) for measurement A (56 Kbps)

Throughput (TPS)	Actual (PB + OS/2)	Est. PB
3.25	8	21
4.90	10	32
6.24	12	42
7.00	14	48

Table 5. Number of clients in relation to aggregate throughput (TPS) for measurement set B (256 Kbps)

Throughput (TPS)	Actual (PB + OS/2)	Est. PB
16.60	16	103
26.88	26	172
34.76	36	227
38.00	46	253

Table 6. Number of clients in relation to aggregate throughput (TPS) for measurement set C (T1)

Throughput (TPS)	Actual (PB + OS/2)	Est. PB
28.98	16	174
45.88	26	281
57.71	36	357
63.12	46	390
65.06	56	414

The values obtained in these measurements are obtained based on the actual number of clients. When evaluating the estimated number of PowerBuilder clients, other capacity planning factors need to

be addressed. For example, each client that passes through DDCS should be allowed approximately 450 KBytes of RAM. These factors need to be taken into consideration when designing a configuration. One last point on the number of clients. These numbers are based on the clients running without think time or keystroke time. So if these aspects are factored in, the actual number of clients could be increased to fill the capacity left available due to the think time and keystroke time.

CPU and RAM Utilization

This section describes CPU utilization and RAM utilization for the C10 running DDCS for AIX for the three measurement sets.

CPU and RAM capacity on the C10 is not a factor in these measurement sets. CPU reaches a maximum of 73 percent for the greatest workload.

Table 7, Table 8, and Table 9 provide the numeric values for the three measurement sets.

Table 7. CPU utilization in relation to TPS for measurement set A

Throughput (TPS)	C10 (User)	C10 (System)	C10 (Total)
3.25	1%	3%	4%
4.90	1%	4%	5%
6.24	2%	5%	7%
7.00	2%	5%	7%

Table 8. CPU utilization in relation to TPS for measurement set B

Throughput (TPS)	C10 (User)	C10 (System)	C10 (Total)
16.60	5%	13%	18%
26.88	8%	20%	28%
34.76	10%	26%	36%
38.00	11%	30%	41%

Table 9. CPU utilization in relation to TPS for measurement set C

DB2 for AIX throughput (TPS)	C10 (User)	C10 (System)	C10 (Total)
28.98	8%	21%	29%
45.88	14%	35%	49%
57.71	17%	45%	62%
63.12	19%	50%	69%
65.06	20%	53%	73%

Table 10 on page 12, Table 11 on page 12, and Table 12 on page 12 provide the RAM utilization values for the three measurement sets. The RAM utilization can be compared within a measurement

set, but not across the measurement sets. This is because the base RAM level at the start of each measurement set is different. The RAM utilization for the DDCS for AIX clients is determined to be between 400 and 500 KBytes per client.

The base RAM utilization in the tables is based on the number of applications that are running on the C10.

Additional studies with RAM utilization indicate that a client that performs more transactions per second does not place a significantly larger demand on RAM requirements. For example, if a client is running the same transaction mix, but at 50 percent the rate as another client, the RAM demands are still about the same.

Table 10. RAM utilization in relation to TPS for measurement set A

Throughput (TPS)	Number DDCS for AIX Clients	RAM on C10 (MBytes)
3.25	8	68.16
4.90	10	68.94
6.24	12	68.96
7.00	14	69.75

Table 11. RAM utilization in relation to TPS for measurement set B

Throughput (TPS)	Number DDCS for AIX Clients	RAM on C10 (MBytes)
16.60	16	71.43
26.88	26	76.02
34.76	36	80.94
38.00	46	85.12

Table 12. RAM utilization in relation to TPS for measurement set C

Throughput (TPS)	Number DDCS for AIX Clients	RAM on C10 (MBytes)
28.98	16	73.37
45.88	26	77.40
57.71	36	81.54
63.12	46	85.58
65.06	56	89.78

For general consideration:

- **For this type of configuration, plan on about 500 KBytes of RAM per client that uses DDCS for AIX.**
- **The actual transaction rate per client does not noticeably increase the RAM demands.**
- **Avoid RAM demands greater than the RAM available. This results in paging which is a serious performance impact.**
- **Avoid over utilizing the CPU that runs DDCS if the clients are holding locks on the target database. Over utilization of the CPU will create delays in freeing the locks, which can effect client access to the same tables from other sources.**

Network

The network plays a big part in how well the client/server application will perform, and how big an effect the client/server applications will have on the performance of the database. The line speed has a significant effect on the throughput. Over-utilization of a line will introduce delays. Delay is the biggest enemy. How big a role delay plays, is dependent on the application. Every time a trip across the network is required, if locks are held, they are held for the time it takes for the messages to move between the client and the database. For these measurements, stored procedures are used, which has the benefit of decreasing the number of SQL calls that must flow across the network as messages. Locks are only held for the time to process the stored procedure and for one network round trip to issue a commit or rollback. So although delay does play a role in these measurements, it could play a larger role in other environments.

The next two sections describe the effects on the NCP and LAN during the measurements. The following values are set in regards to RQRIOBLK sizes and RUSIZEs:

- 4K RUSIZEs in VTAM
- 4K RQRIOBLK definition in DDCS for AIX
- 4K RQRIOBLK definitions in the OS/2 clients
- 4K DOS_RQRIOBLK definitions in the DOS (PowerBuilder) clients

NCP

Table 13 on page 14 and Table 14 on page 14 describe the effects of the workload on the 3745 running the NCP. These measurements do not stress the 3745 beyond 77.40% CCU utilization at a throughput rate of 64.80 transactions per second. The data throughput in the NCP reached 62,685 bytes per second. Use the numbers in the tables to understand the amount of byte activity that is occurring while achieving the defined throughput rates on DB2 for MVS/ESA and DB2 for AIX.

The DELAY parameters in the NCP play an important role in the performance of the NCP as well as the response time observed at the clients. The greater the value on the PCCU, and/or GROUP LNCTL=CA parameters, the lower the CCU utilization, but the greater the response time at the clients, especially at lower throughputs.

For these measurements, the GROUP LNCTL=CA DELAY parameter is set to 0.1 and the PCCU parameter is set to zero. This had the beneficial effect of controlling the CCU utilization with little response time effect at the clients at the higher levels of throughput. The lower levels of throughput experienced mildly higher response times than if the GROUP LNCTL=CA DELAY parameter equals 0.

The NCP for the measurements are dedicated to the IRWW workload activity. Typically, an NCP is doing other work and this work should be factored into the capacity requirements for the NCP.

Table 13 on page 14, Table 14 on page 14, and Table 15 on page 14 provide the numeric values for the three measurement sets.

Table 13. 3745 utilization in relation to TPS for measurement set A (56 Kbps)

DB2 for MVS/ESA throughput (TPS)	Bytes Per Second
3.25	3,800
4.90	5,291
6.24	6,472
6.99	7,142

Table 14. 3745 utilization in relation to TPS for measurement set B (256 Kbps)

DB2 for MVS/ESA throughput (TPS)	Bytes Per Second
16.60	16,644
26.88	25,575
34.76	32,724
38.00	35,758

Table 15. 3745 utilization in relation to TPS for measurement set C (T1)

DB2 for MVS/ESA throughput (TPS)	Bytes Per Second
28.98	28,616
45.88	44,204
57.71	55,623
63.12	61,424
65.06	62,685

These results can be used to analyze the line capacity requirements between the LAN and the DB2/MVS host system. For these measurements, the connection is through varying line speeds (56 Kbps, 256 Kbps, and T1).

For general consideration:

- **Assess the line speeds between the LAN and DB2 for MVS/ESA. The throughput is only as good as the slowest line. Over-utilizing any line will result in delays.**
- **Introducing multiple communications devices (e.g. 3745s, Bridges, Routers) on the communication path between the LAN and the database introduces delays.**
- **Tune the NCP definition DELAY parameters. If response time at the client is most important set all DELAY values to zero. If CCU constrained, set the GROUP LNCTL=CA to 0.1 or 0.2. 0.2 will create a greater response time delay than 0.1. The following NCP statements have DELAY parameters:**
  PCCU
  GROUP LNCTL=CA

LAN

For these measurements, the LAN capacity does not factor in. The 16 Mbit LAN is lightly utilized.

For general consideration, monitor your LAN to make sure it is not being over utilized. Over utilization of the LAN could introduce delays with resultant bad effects on database performance.

SQL Rates

The following rate and mix of SQL statements occur when achieving the maximum throughput rate for each measurement set:

7.00 (measurement set 56 Kbps)

- 25 SELECT statements per second
- 19 INSERT statements per second
- 25 UPDATE statements per second
- 1 DELETE statements per second
- 31 OPEN statements per second
- 17 CLOSE statements per second
- 61 FETCH statements per second

38.00 (measurement set 256 Kbps)

- 136 SELECT statements per second
- 105 INSERT statements per second
- 136 UPDATE statements per second
- 7 DELETE statements per second
- 165 OPEN statements per second
- 95 CLOSE statements per second
- 328 FETCH statements per second

65.06 (measurement set T1)

- 230 SELECT statements per second
- 185 INSERT statements per second
- 238 UPDATE statements per second
- 13 DELETE statements per second
- 288 OPEN statements per second
- 168 CLOSE statements per second
- 573 FETCH statements per second

For general consideration:

- **When using stored procedures in a client/server environment, tune the database as if the access is local.**

Transactions

The next section describes the transactions and the achieved response time at the PowerBuilder clients in relation to the aggregate throughput. The response time duration starts when the end user submits the transaction request. This is after the network connection is established and the end user has entered the input data. The duration ends after the commit or rollback has been issued and all the response data is available to the end user (visible on the computer monitor).

In this section, graphs depict the point where 90% of the transactions complete within the maximum response time for the identified throughput. There are three lines in the graphs to represent the three measurement sets. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

Figure 2 displays the aggregate response time of the transactions on DB2 for MVS/ESA for the identified throughput for the three measurement sets. The largest impact of decreasing the line speed is the decrease in throughput. Response times degrade more quickly when adding additional workload on the slower line speeds. Full utilization of the line is the primary factor limiting additional throughput for the 56 Kbps and 256 Kbps line speed measurement sets. Lock and latch contention is the primary factor limiting additional throughput for the T1 line speed measurement set.

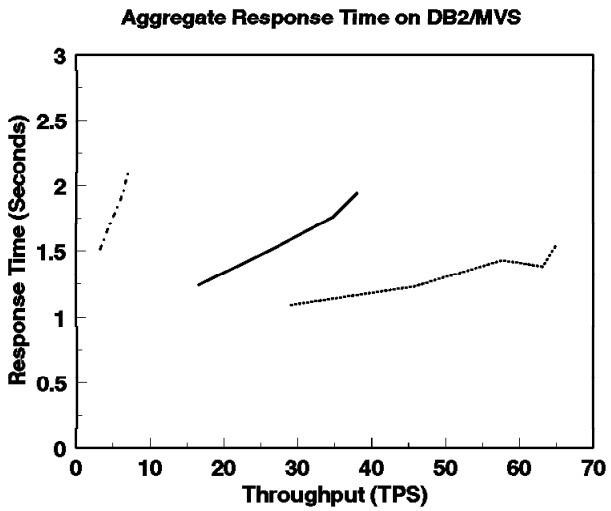

Figure 2. Aggregate Response Time

Each procedure described in the following sections, perform different operations which creates various degrees of contention on the tables. The SQL calls in the order they are called, are described for each procedure.

Transaction Tx1

Transaction Tx1 is the most involved of the seven transactions. The transaction calls a stored procedure which performs write operations against tables T3, T4, T6, T8, and T9. The stored procedure performs read operations against tables T2, T3, T5, T7, and T8. Index matching predicates in the WHERE clauses optimize access to all tables. The stored procedure has two input parameters for a byte count between 22 and 190. There are nine output parameters with a maximum byte count of 713.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
Select
   T7
   T2
Fetch and Update
   T3
Insert
   T4
   T6
Fetch
   T5
Loop up to 15 times
   Fetch and Update
      T8
   Insert
      T9
End loop
```

Figure 3 describes the achieved response time for transaction Tx1. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx1 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

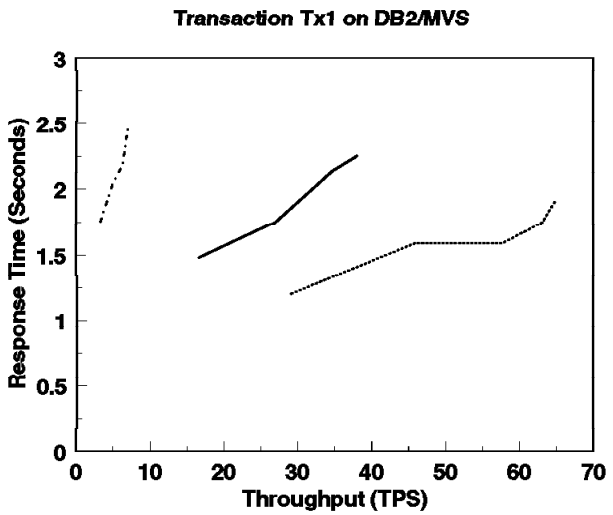

Figure 3. Response Time For Transaction Tx1

The 56 Kbps measurement set ranges from 1.75 seconds to 2.5. The 256 Kbps measurement set ranges from 1.5 seconds to 2.25 seconds. The T1 measurement set ranges from 1.0 and 2.0 seconds.

Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention on DB2 for MVS/ESA.

Table 16 displays the mean and standard deviation for the three measurement sets.

Table 16. Response Time Mean And Standard Deviation For Transaction Tx1

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.42	0.28
A	4.90	1.63	0.31
A	6.24	1.80	0.42
A	7.00	1.90	0.41
B	16.60	1.12	0.30
B	26.88	1.31	0.40
B	34.76	1.60	0.44
B	38.00	1.79	0.46
C	28.98	0.94	0.28
C	45.88	1.10	0.44
C	57.71	1.15	0.42
C	63.12	1.30	0.47
C	65.06	1.46	0.58

Table 17 displays the transactions that have potential lock contention with transaction Tx1 on the defined tables.

Table 17. Potential Lock Contention For Transaction Tx1

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T2			yes				
T3	yes		yes				
T4	yes						yes
T5				yes			
T6	yes						yes
T7			yes				yes
T8	yes						
T9	yes						yes

Transaction Tx2

Transaction Tx2 calls a stored procedure which performs read operations against tables T6, T7, and T9. Index matching predicates in the WHERE clause optimize access to these tables. The stored procedure has one input parameter for 26 bytes and four output parameters with a maximum byte count of 516.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
For 60% of the time:
    Select
       T7
    Loop up to 4 times
       Fetch from
          T7
    End loop
End 60% of the time
    Select
       T7
       T6
    Loop up to 15 times
       Fetch
          T9
    End loop
```

Figure 4 on page 20 describes the achieved response time for transaction Tx2. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx2 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

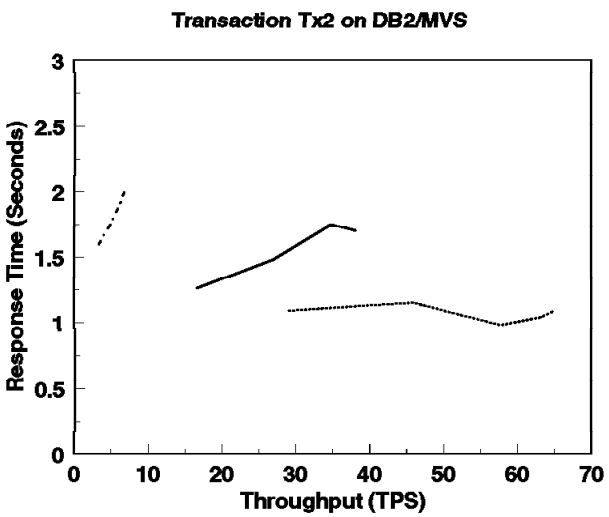

Figure 4. Response Time For Transaction Tx2

The 56 Kbps measurement set ranges from 1.60 seconds to 2.1. The 256 Kbps measurement set ranges from 1.5 seconds to 1.75 seconds. The T1 measurement set ranges around 1.0 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention on DB2 for MVS/ESA.

Table 18 displays the mean and standard deviation for the three measurement sets.

Table 18. Response Time Mean And Standard Deviation For Transaction Tx2

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.26	0.25
A	4.90	1.42	0.28
A	6.24	1.50	0.31
A	7.00	1.66	0.32
B	16.60	0.92	0.26
B	26.88	1.03	0.30
B	34.76	1.25	0.36
B	38.00	1.39	0.34
C	28.98	0.72	0.29
C	45.88	0.78	0.39
C	57.71	0.87	0.44
C	63.12	0.89	0.26
C	65.06	0.94	0.37

Table 19 on page 21 displays the transactions that have potential lock contention with transaction Tx2 on the defined tables.

Table 19. Potential Lock Contention For Transaction Tx2

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T6	yes						yes
T7			yes				yes
T9	yes						yes

Transaction Tx3

Transaction Tx3 calls a stored procedure which performs write operations against tables T1, T2, T3, and T7. The stored procedure performs read operations against tables T2, T3, and T7. Access to all tables, except T1, use index matching predicates in the WHERE clauses to optimize access to the tables. The stored procedure has one input parameter for 39 bytes and six output parameters with a maximum byte count of 551.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
For 60% of the time
      Select
         T7
      Loop up to 4 times
         Fetch from
            T7
      End loop
End 60% of the time
      Fetch and Update
         T7
      Conditionally Fetch and Update
         T2
         T3
      Insert
         T1
```

Figure 5 on page 22 describes the achieved response time for transaction Tx3. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx3 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

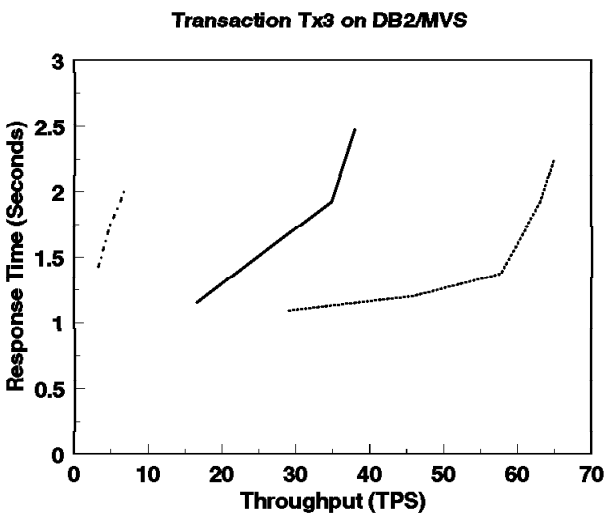

Figure 5. Response Time For Transaction Tx3

The 56 Kbps measurement set ranges from 1.4 seconds to 2.0. The 256 Kbps measurement set ranges from 1.2 seconds to 2.5 seconds. The T1 measurement set ranges from 1.0 and 2.2 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention.

Table 20 displays the mean and standard deviation for the three measurement sets.

Table 20. Response Time Mean And Standard Deviation For Transaction Tx3

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.16	0.24
A	4.90	1.33	0.32
A	6.24	1.50	0.34
A	7.00	1.63	0.34
B	16.60	0.82	0.25
B	26.88	1.11	0.35
B	34.76	1.32	0.48
B	38.00	1.70	0.56
C	28.98	0.65	0.25
C	45.88	0.77	0.34
C	57.71	0.89	0.37
C	63.12	1.14	0.69
C	65.06	1.41	0.80

Table 21 on page 23 displays the transactions that have potential lock contention with transaction Tx3 on the defined tables. Both tables T2 and T3 tables are small, so the contention potential is great.

Table 21. Potential Lock Contention For Transaction Tx3

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T1			yes				
T2			yes				
T3	yes		yes				
T7			yes				yes

Transaction Tx4

Transaction Tx4 calls a stored procedure which performs a write operation against T5. Index matching predicates in the WHERE clause optimize access to this table. The stored procedure has one input parameter for 8 bytes and three output parameters with a maximum byte count of 99.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
Update
  T5
```

Figure 6 on page 24 describes the achieved response time for this transaction. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx4 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

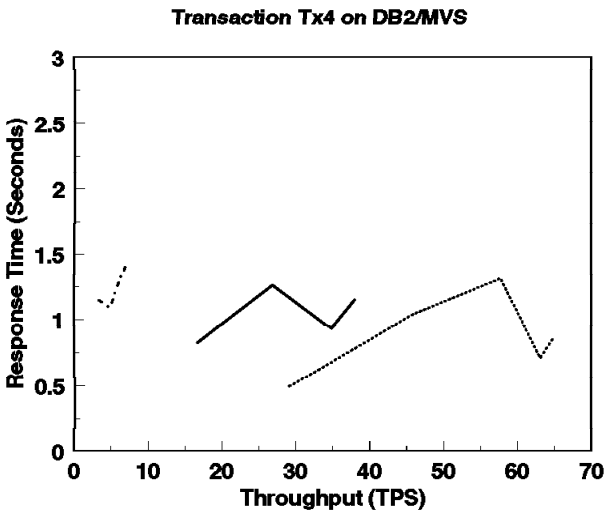

Figure 6. Response Time For Transaction Tx4

The 56 Kbps measurement set ranges from 1.1 seconds to 1.5. The 256 Kbps measurement set ranges from 1.8 seconds to 1.2 seconds. The T1 measurement set ranges from 0.5 and 1.2 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention.

Transaction Tx4 runs only one percent of the transaction mix, so low sampling anomalies account for the graph characteristics.

Table 22 on page 25 displays the mean and standard deviation for the three measurement sets.

Table 22. Response Time Mean And Standard Deviation For Transaction Tx4

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	0.64	0.27
A	4.90	0.89	0.15
A	6.24	1.05	0.19
A	7.00	1.22	0.18
B	16.60	0.64	0.16
B	26.88	0.86	0.24
B	34.76	0.76	0.11
B	38.00	1.05	0.17
C	28.98	0.35	0.12
C	45.88	0.55	0.26
C	57.71	0.64	0.29
C	63.12	0.59	0.10
C	65.06	0.76	0.08

Table 23 displays the transactions that have potential lock contention with transaction Tx4 on the defined tables.

Table 23. Potential Lock Contention For Transaction Tx4

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T5				yes			

Transaction Tx5

Transaction Tx5 calls a stored procedure which performs read operations against the T5 table. Index matching predicates in the WHERE clause optimize access to this table. The stored procedure has one input parameter with a byte count between 8 and 120. It has seven output parameters with a maximum byte count of 590.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
Loop up to 15 times
   Select
      T5
End loop
```

Figure 7 on page 26 describes the achieved response time for transaction Tx5. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx5 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

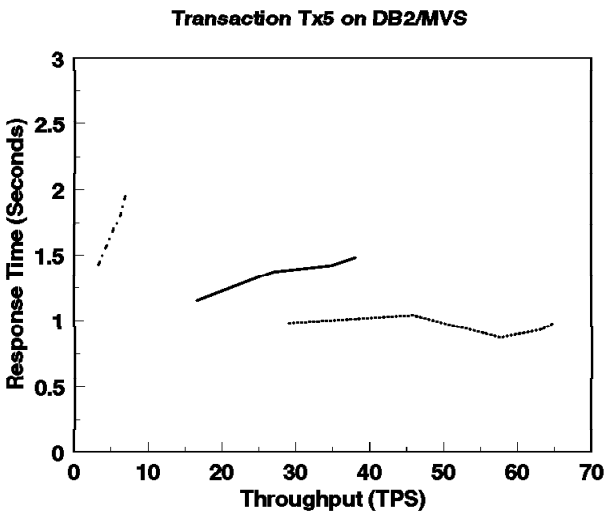

Figure 7. Response Time For Transaction Tx5

The 56 Kbps measurement set ranges from 1.4 seconds to 1.8. The 256 Kbps measurement set ranges from 1.1 seconds to 1.4 seconds. The T1 measurement set ranges from 0.7 and 0.85 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention.

Table 24 on page 27 displays the mean and standard deviation for the three measurement sets.

Table 24. Response Time Mean And Standard Deviation For Transaction Tx5

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.16	0.24
A	4.90	1.33	0.28
A	6.24	1.41	0.30
A	7.00	1.53	0.33
B	16.60	0.79	0.24
B	26.88	0.95	0.25
B	34.76	1.09	0.26
B	38.00	1.24	0.23
C	28.98	0.65	0.22
C	45.88	0.64	0.22
C	57.71	0.68	0.18
C	63.12	0.75	0.15
C	65.06	0.82	0.16

Table 25 displays the transactions that have potential lock contention with transaction Tx4 on the defined tables.

Table 25. Potential Lock Contention For Transaction Tx5

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T5				yes			

Transaction Tx6

Transaction Tx6 calls a stored procedure which performs read operations against tables T3, T8, and T9. Index matching predicates in the WHERE clause optimize access to the rows. The stored procedure has one input parameter with a byte count of 8 and three output parameters with a maximum byte count of 78.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
Select
   T3
Select
   T8 joined with T9
```

Figure 8 on page 28 describes the achieved response time for transaction Tx6. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx6 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

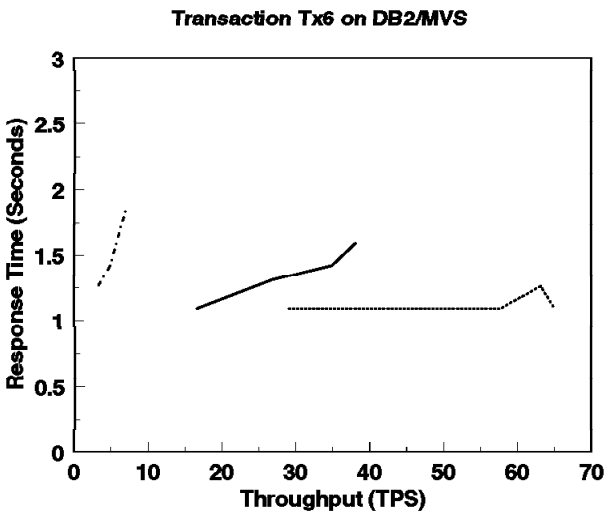

Figure 8. Response Time For Transaction Tx6

The 56 Kbps measurement set ranges from 1.25 seconds to 1.75. The 256 Kbps measurement set ranges from 1.1 seconds to 1.6 seconds. The T1 measurement set ranges from 1.1 and 1.2 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention.

Table 26 on page 29 displays the mean and standard deviation for the three measurement sets.

Table 26. Response Time Mean And Standard Deviation For Transaction Tx6

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.07	0.24
A	4.90	1.12	0.26
A	6.24	1.28	0.30
A	7.00	1.38	0.36
B	16.60	0.90	0.28
B	26.88	0.94	0.31
B	34.76	1.05	0.26
B	38.00	1.21	0.31
C	28.98	0.78	0.33
C	45.88	0.89	0.36
C	57.71	0.85	0.28
C	63.12	0.91	0.36
C	65.06	0.97	0.39

Table 27 displays the transactions that have potential lock contention with transaction Tx6 on the defined tables.

Table 27. Potential Lock Contention For Transaction Tx6

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T3	yes		yes				
T8	yes						
T9	yes						yes

Transaction Tx7

Transaction Tx7 calls a stored procedure which performs write operations against tables T4, T6, T7, and T9. The stored procedure performs read operations against tables T4, T6, and T9. Index matching predicates in the WHERE clauses optimize access to these large tables. The stored procedure has one input parameter for 6 bytes and three output parameters with a maximum byte count of 78.

A summary of the SQL calls in the stored procedure and the order in which they occur are as follows:

```
Loop up to 10 times
   Fetch and Delete
      T4
   Fetch and Update
      T6
   Update
      T9
   Select
      T9
   Update
      T7
End loop
```

Figure 9 describes the achieved response time for this transaction. The vertical axis represents the response time as seen by the PowerBuilder clients. The horizontal axis represents the overall throughput rate for the aggregate of transactions.

For each measurement set, a line represents the point where 90% of the Tx7 transactions achieve this response time or better. The dot/dash line is measurement set A with a line speed of 56 Kbps. The solid line is measurement set B with a line speed of 256 Kbps. The dash line is the measurement set C with a T1 line speed.

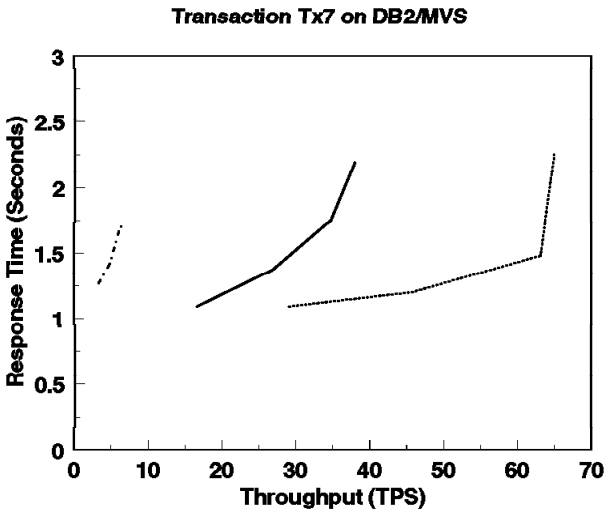

Figure 9. Response Time On DB2 for MVS/ESA For Transaction Tx7

The 56 Kbps measurement set ranges from 1.25 seconds to 1.75. The 256 Kbps measurement set ranges from 1.1 seconds to 2.25 seconds. The T1 measurement set ranges from 1.1 and 2.25 seconds. Line utilization is the limiting factor for the 56 Kbps and 256 Kbps measurement sets and the T1 measurement set is limited by lock and latch contention.

Table 28 on page 31 displays the mean and standard deviation for the three measurement sets.

Table 28. Response Time Mean And Standard Deviation For Transaction Tx7

Measurement set	Throughput (TPS)	Mean (seconds)	Standard Deviation (seconds)
A	3.25	1.06	0.19
A	4.90	1.11	0.25
A	6.24	1.31	0.31
A	7.00	1.29	0.28
B	16.60	1.08	0.60
B	26.88	1.08	0.21
B	34.76	1.19	0.31
B	38.00	1.71	1.51
C	28.98	0.92	0.26
C	45.88	1.33	1.46
C	57.71	1.92	2.67
C	65.06	1.77	2.49

Table 29 displays the transactions that have potential lock contention with transaction Tx7 on the defined tables.

Table 29. Potential Lock Contention For Transaction Tx7

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T4	yes						yes
T6	yes						yes
T7			yes				yes
T9	yes						yes

5. Conclusions

This technical report shows the results and analysis of some performance measurements over varying communications line speeds for a client/server relational database configuration that includes DB2 for MVS/ESA, DDCS for AIX, and DB2/CAE/Windows with PowerBuilder. The goal of these measurements is to provide an example environment and analysis that will help a customer:

- Assess proposed configurations for acceptable performance
- Assess performance of current configurations
- Choose proper hardware capacity for a proposed configuration
- Choose the proper software programs for a proposed configuration
- Tune a new or current configuration

Throughout the report, ″**General Considerations**″ are identified. This information is specially identified as information to consider for tuning and capacity planning analysis. This does not imply that all information is summarized in these general considerations, because the information provided in this report can provide additional information dependent on the analyzer′s reference point.

This report focuses on the use of stored procedures for online transaction processing types of environments. Network delays play an important role in how well the environment performs. Through the use of stored procedures, network delay plays a lesser role, but is not eliminated. For example, each transaction consists of a stored procedure call and a commit (or rollback). This results in two ″round trip″ message flows from the client to the database. After the stored procedure completes, any locks created will remain held until the commit or rollback flows across the wire. Without stored procedures, locks might be held much longer because there are two ″round trip″ messages for every SQL statement.

The burden of providing good throughput in a client/server environment falls on several different people:

- Network designers to make sure the network has capacity and an efficient path exists between the client and target database.
- Database administrators to make sure the target database and gateway (DDCS/6000) is operating within capacity and tuned for client/server access.
- Application programmers to make sure the program is efficiently manipulating the database, for example, use of indexes.
- Any other administrators of software programs that might exist between the client and the target database (i.e. gateways).

The importance of this report is that it provides a reasonably tuned client/server database configuration. It achieved a certain level of throughput using a combination of components. Some questions to ask when analyzing this report to apply towards your environment are:

- How does the data throughput compare to my environment or proposed environment?
  - On the LAN?
  - Through DDCS/6000?

© Copyright IBM Corp. 1996

  - Through the NCP?
- How much processor capacity would I need for my environment considering the throughput and transaction mix? DDCS/6000, DB2, etc...
- Is the response time at the clients reasonable, or can I trade worse response time for additional clients?
- How much network line capacity would I need considering throughput?
- How many network devices are between the clients and the database?

6. Other sources of information

- **DRDA Connectivity Guide** SC26-4783
- **DB2/MVS Administration Guide** SC26-3265
- **DB2 Administration Guide For Common Server V2** S20H-4580

© Copyright IBM Corp. 1996