DB2 Family: Client/Server Performance Measurement Series
DB2/MVS <> DDCS/2 <> DB2/2 <> DB2/CAE/Windows

Document Number Version 1.1

June 26, 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

IBM

DB2 Performance
Santa Teresa Laboratory
San Jose, California

Trademarks

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

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

All other trademarks are the properties of their respective owners.

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.

Acknowledgements

The following people helped make this paper possible.

- Jerry Heglar
- Serge Limoges
- Mark Ryan
- Bill Wilkins

Contents

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, CCU utilization, network utilization and response times when PowerBuilder clients and OS/2 clients call stored procedures on DB2 for OS/2 and on DB2 for MVS/ESA through DDCS for OS/2 to perform OLTP type transactions. DB2 for OS/2 and DDCS for OS/2 reside on the same hardware for these measurements. Use the results from these measurements to predict performance and ″capacity plan″ in your environment. Particularly noteworthy items are highlighted as **general considerations** in this report.

2. Configuration

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

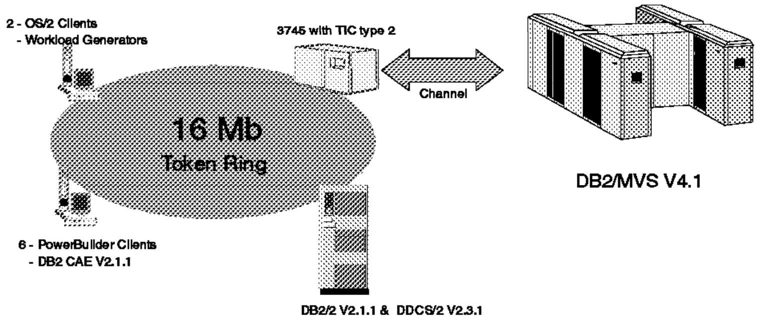

Figure 1. Configuration with DB2 for MVS/ESA, DB2 for OS/2, and DDCS for OS/2

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 to the 16 Mbit LAN via a TIC type 2.
- **DB2 for OS/2** Version 2.1.1. The processor is a PC500 server with a 90 Mhz Pentium processor and 128 MBytes of memory. The PC500 is attached to a 16 Mbit LAN using a LANStreamer adapter card and configured to use early token release.
- **Distributed Database Connection Services for OS/2 (DDCS/2)** Version 2.3.1 running on the same PC500 as DB2 for OS/2.
- **6 - PowerBuilder Version 4.0.2 with DB2 CAE for Windows Version 2.1.1** clients running on 9595 OPT PS/2s with 60 Mhz Pentium Processors and 80 MBytes of memory. The LAN adapter cards 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 OS/2.
- TCP/IP is in use between the OS/2 clients and DDCS for OS/2.
- TCP/IP is in use between the PowerBuilder clients and DB2 for OS/2.
- TCP/IP is in use between the OS/2 clients and DB2 for OS/2.
- LU 6.2 is in use between DDCS for OS/2 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)
- OS/2: System Performance Monitor/2 (SPM/2)
- 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 when the OS/2 server is providing the DDCS function as well as the database server (DB2/2). The full IRWW transaction set is executed against DB2 for MVS/ESA. The ″query only″ subset of the IRWW transaction set is executed against DB2 for OS/2. This is to simulate an environment where DB2 for MVS/ESA performs the OLTP updates and the updates are propagated out to the satellite sites, in this case DB2 for OS/2. These measurements do not include the propagation activities.

A series of measurements are performed, with each subsequent measurement providing an increased workload to the configuration. Since this is a multi-dimensional set of measurements due to the fact that there are two DB2 databases for targets with DB2 for OS/2 sharing the hardware with DDCS for OS/2, the results are presented with the workload constant at one database, while increasing the workload on the other.

All measurements include six PowerBuilder clients running a predefined set of transactions in random order without modification for the entire series of measurements. Three PowerBuilder clients access DB2 for MVS/ESA and the other three access DB2 for OS/2. 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. One of the OS/2 workload generators accesses DB2 for MVS/ESA and the other accesses DB2 for OS/2. Each OS/2 workload generator runs the exact same transaction mix as the three PowerBuilder clients for each respective DB2 they access. 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. The workload on DB2 for OS/2 consists of transactions Tx2, Tx5, and Tx6 which are a query subset of the 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 on DB2 for MVS/ESA as twenty-two percent of the total transaction mix and is the most complicated transaction. Transaction Tx1 does not run on DB2 for OS/2.

- Transaction **Tx2**

Performs various selects and fetches in support of providing the status of an order. Transaction Tx2 runs on DB2 for MVS/ESA as twenty-four percent of the total transaction mix. Transaction Tx2 runs on DB2 for OS/2 as forty-five percent of the total transaction mix.

- Transaction **Tx3**

  Performs various selects, fetches, updates, and inserts in support of received customer payments. Transaction Tx3 runs on DB2 for MVS/ESA as twenty-two percent of the total transaction mix. Transaction Tx3 does not run on DB2 for OS/2.

- Transaction **Tx4**

  Performs an update in support of changing the price of an item. Transaction Tx4 runs on DB2 for MVS/ESA as one percent of the total transaction mix. Transaction Tx4 does not run on DB2 for OS/2.

- Transaction **Tx5**

  Performs various selects in support of providing the price of a set of items. Transaction Tx5 runs on DB2 for MVS/ESA as twenty-five percent of the total transaction mix. Transaction Tx5 runs on DB2 for OS/2 as forty-seven percent of the total transaction mix.

- Transaction **Tx6**

  Performs various selects in support of providing the current stock level of an item. Transaction Tx6 runs on DB2 for MVS/ESA as four percent of the total transaction mix. Transaction Tx6 runs on DB2 for OS/2 as 8 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 on DB2 for MVS/ESA as two percent of the total transaction mix. Transaction Tx7 does not run on DB2 for OS/2.

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

Table Characteristics

Table characteristics on DB2 for MVS/ESA

Before each measurement run, the tables on DB2 for MVS/ESA are put into a consistent state as defined in Table 1 on page 7. The table 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

Table characteristics on DB2 for OS/2

The tables on DB2 for OS/2 are described in Table 2. The table shows the number of rows, the number of columns, the number of indexes, and the buffer pool preload percentage. The amount of preloaded data is greater than available buffer space, so table T9 is loaded first, followed by table T8, then T7, on up to T3. DB2 will remove portions of the earlier preloaded tables as space is needed.

Table 2. Table Characteristics for DB2 for OS/2

Table	Number Rows (length)	Columns	Indexes	PreLoad
T3	80 (Fixed 105)	11	1	100
T5	100,000 (Fixed 86)	4	1	100
T6	240,000 (Fixed 25)	8	2	12
T7	240,000 (Fixed 675)	21	2	12
T8	800,000 (Fixed 306)	17	1	20
T9	2,399,791 (Fixed 60)	10	1	12

Locking Definitions

Locking definitions on DB2 for MVS/ESA

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.

Locking definitions on DB2 for OS/2

The locking is system defined and managed on DB2 for OS/2. All packages are bound with an isolation level of cursor stability.

Physical location

Physical location of data on DB2 for MVS/ESA

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 3. 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.

As a 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 percent preloaded into memory, therefore, I/O contention to these drives is reduced.

Table 3 (Page 1 of 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	9394-1	T8 Table (Partitions 9-16)

Table 3 (Page 2 of 2). Physical Storage Characteristics

Drive	Controller	Drive type	Table/Indexes
DISK33	CNTRL5	3380-E	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)

Physical location of data on DB2 for OS/2

There are four disk drives. One drive contains the DB2 code and log files. The other three disk drives are included in a single database managed tablespace.

As a general consideration, it is best to allow the database to manage the tablespaces.

Buffer Pool Allocation

Buffer pool allocation on DB2 for MVS/ESA

The tables and indexes are also spread out to multiple buffer pools as described in Table 4 on page 10. 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 4. 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	

Buffer pool allocation on DB2 for OS/2

A 48 MByte (BUFFPAGE=12288) buffer pool is allocated for DB2 for OS/2.

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 joint DB2 for OS/2 workload and DDCS for OS/2 workload operating simultaneously on the same platform. There are two dimensions of workload scaling in this configuration. One dimension is the scaling of the workload on DB2 for MVS/ESA which increases the workload for DDCS for OS/2. The other dimension is scaling the workload on DB2 for OS/2. Workload scaling occurs by increasing the number of sessions on the OS/2 workload generators. The results are presented as three measurement sets. Each set holds DB2 for OS/2 workload constant while scaling the workload to DB2 for MVS/ESA. The workload on DB2 for OS/2 scales between each measurement set. The workload combination for each measurement set is summarized in Table 5. The workload, which is represented by the number of sessions, is numerically displayed in the workload columns.

Table 5. Workload Combination

Set #	DB2 for OS/2 Workload	DB2 for MVS/ESA Workload
A	2	5, 10, 15, 20, 25
B	4	5, 10, 15, 20
C	6	5, 10, 15, 20

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 to a database, the throughput increases to that database, while the throughput to the other database in the mix decreases. For example, if DB2 for OS/2 is held constant and the workload to DB2 for MVS/ESA increases, which results in more DDCS activity, the throughput on DB2 for MVS/ESA increases while the throughput on DB2 for OS/2 decreases. The throughput on DB2 for OS/2 decreases because of the PC500 CPU demands to satisfy the DDCS activity on the PC500. The resultant effects on each component is then presented based on the current achieved throughput. For example, when achieving 6.74 transactions per second on DB2 for OS/2 and 28.29 transactions per second on DB2 for MVS/ESA:

- the transfer rate through the 3745 running the NCP is 24,953 bytes per second
- the PC500 running DB2 for OS/2 and DDCS for OS/2 is 95.84 percent utilized
- etc...

For all measurements, the final constraint when scaling the workload is the CPU on the PC500 running DB2 for OS/2 and DDCS for OS/2. Each measurement set scales the DB2 for MVS/ESA workload until the PC500 reaches full capacity.

The scaling of DB2 for OS/2 (approximately 45 percent) degrades the overall throughput on DB2 for MVS/ESA approximately 38 percent. The scaling of DB2 for MVS/ESA (approximately 450 to 500 percent) degrades the throughput on DB2 for OS/2 approximately 64 percent.

As a general consideration, the combination of DDCS and DB2 workloads on the same CPU place large demands on the CPU. If the transactions on DB2 are CPU intensive, you should consider separating DDCS and DB2 on to independent hardware.

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 estimates the number of total PowerBuilder clients based on the throughput obtained by the OS/2 clients plus the throughput of the actual PowerBuilder clients.

For the three measurement sets, Table 6, Table 7 on page 13, and Table 8 on page 13 describe the number of clients activated to achieve the throughput defined in column one. Column two describes the actual number of clients (both OS/2 and PowerBuilder) and estimated number of PowerBuilder clients accessing DB2 for OS/2. Column three describes the actual number of clients (both OS/2 and PowerBuilder) and estimated number of PowerBuilder clients accessing DB2 for MVS/ESA. The equation for evaluating the estimated number of PowerBuilder clients is:

$$\text{Estimated PowerBuilder Clients} = \left[\frac{\text{OS/2 client tps}}{\text{PowerBuilder client tps}} * (\text{\# OS/2 clients}) \right] + (\text{\# PowerBuilder clients})$$

If
OS/2 client tps = 1.86 tps
PowerBuilder client tps = 0.14 tps
OS/2 clients = 15
PowerBuilder clients = 3
then

$$\text{Estimated PowerBuilder Clients} = \left[\frac{1.86 \text{ tps}}{.14 \text{ tps}} * 15 \text{ clients} \right] + (3 \text{ clients})$$

$$= 202 \text{ clients}$$

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

Throughput (TPS) DB2/2 / DB2/MVS	DB2/2 Actual / Est. PB	DB2/MVS Actual / Est. PB
6.47 / 8.96	5 / 38	8 / 63
5.73 / 22.56	5 / 35	13 / 157
5.06 / 34.16	5 / 32	18 / 233
4.14 / 45.83	5 / 27	23 / 327
3.50 / 53.86	5 / 23	28 / 368

Table 7. Number of clients in relation to aggregate throughput (TPS) for measurement set B

Throughput (TPS) DB2/2 / DB2/MVS	**DB2/2 Actual / Est. PB**	**DB2/MVS Actual / Est. PB**
8.91 / 8.59	7 / 54	8 / 61
8.03 / 18.19	7 / 49	13 / 124
6.74 / 28.29	7 / 43	18 / 202
5.68 / 38.70	7 / 38	23 / 276

Table 8. Number of clients in relation to aggregate throughput (TPS) for measurement set C

Throughput (TPS) DB2/2 / DB2/MVS	**DB2/2 Actual / Est. PB**	**DB2/MVS Actual / Est. PB**
9.89 / 7.24	9 / 63	8 / 52
8.73 / 15.55	9 / 54	13 / 111
7.39 / 25.73	9 / 49	18 / 184
6.38 / 33.34	9 / 44	23 / 244

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 accesses DB2 for OS/2 (for this type of transaction mix) should be allowed approximately 0.9 MBytes of RAM and each client that passes through DDCS should be allowed approximately 300 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 PC500 running DB2 for OS/2 and DDCS for OS/2 as the workload increases for each measurement set (A, B, and C). Figure 2 on page 14 graphically displays the effects as the workload increases on DB2 for MVS/ESA, and the graph also shows the relationship between the measurement sets. The solid line is measurement set A where the workload against DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload against DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload against DB2 for OS/2 is held constant at 6 sessions. The squares are the DB2 for OS/2 throughput results and the stars are the DB2 for MVS/ESA throughput results.

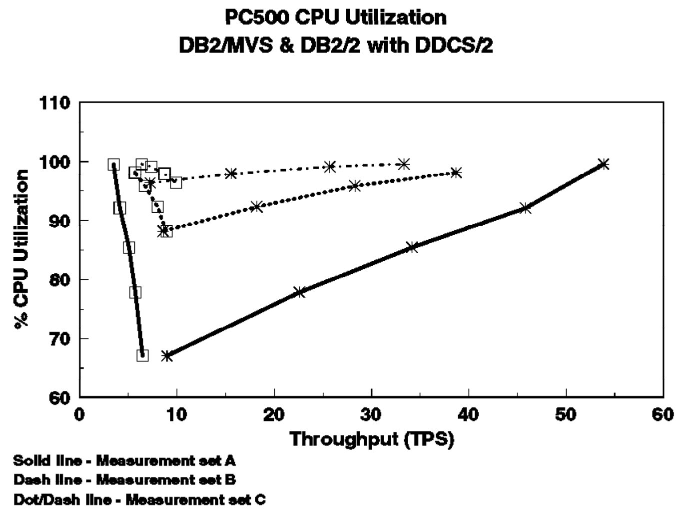

Figure 2. CPU Utilization in Relation to Throughput

At the lower workload (measurement set A), the DB2 for MVS/ESA throughput rate grows faster and higher than measurement set B, and measurement set B grows faster and higher than measurement set C. For each of the measurement sets, the DB2 for OS/2 throughput steadily decreases as the DB2 for MVS/ESA workload increases. This is due to the CPU requirements for DDCS to satisfy the DB2 for MVS/ESA workload that passes through DDCS.

Table 9, Table 10 on page 15, and Table 11 on page 15 provide the numeric values for each measurement set.

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

Throughput (TPS) **DB2/2 / DB2/MVS**	**PC500**
6.47 / 8.96	67.04%
5.73 / 22.56	77.82%
5.06 / 34.16	85.44%
4.14 / 45.83	92.12%
3.50 / 53.86	99.58%

Table 10. CPU utilization in relation to TPS for measurement set B

Throughput (TPS) **DB2/2 / DB2/MVS**	**PC500**
8.91 / 8.59	88.15%
8.03 / 18.19	92.33%
6.74 / 28.29	95.84%
5.68 / 38.70	98.12%

Table 11. CPU utilization in relation to TPS for measurement set C

Throughput (TPS) **DB2/2 / DB2/MVS**	**PC500**
9.89 / 7.24	96.39%
8.73 / 15.55	97.94%
7.39 / 25.73	99.15%
6.38 / 33.34	99.58%

Table 12 on page 16, Table 13 on page 16, and Table 14 on page 16 provide the RAM utilization values for each measurement set. The RAM utilization for the clients accessing DB2 for OS/2 averaged about 250 KBytes per client. The RAM utilization for the DDCS for OS/2 clients averaged about 330 KBytes per client. Other studies in the DB2 Family: Client/Server Performance Measurement Series define a conservative number of 300 KBytes of RAM for clients using DDCS and 900 KBytes of RAM for clients accessing DB2 for OS/2. Unlike the configuration in this report, those reports do not combine workloads against DDCS and DB2 for OS/2. The DDCS client RAM requirements are within 100 KBytes of each other. This is probably due to several things:

- Using a 32 KByte RQRIOBLK size in this measurement vs 4 KByte in the others.
- Using a TCP/IP connection from the OS/2 clients to DDCS in this measurement and LU 6.2 in the others.

The DB2 for OS/2 client requirements are almost 75 percent less for this report. Additional studies are necessary to understand the difference, but the differences might be attributable to:

- Running a query only workload mix in this report vs a mix of query and OLTP transactions in the others.
- Running only three stored procedures in this report vs the full seven stored procedures in the others. Since we are using the hold procedure option, there are memory demands to hold the procedures.
- Effects due to the multi-use (DDCS and DB2 for OS/2)

The base RAM utilization in the tables is based on the number of applications that are running on the server and based on the buffer pool size defined for the database. In this report, the buffer pool size is 48 MBytes. Use the numbers in these tables to estimate the RAM usage for additional clients.

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 12. RAM utilization in relation to TPS for measurement set A

Throughput (TPS) **DB2/2 / DB2/MVS**	**DB2 Clients** **DB2/2 / DB2/MVS**	**RAM on PC500 (MBytes)**
6.47 / 8.96	5 / 8	Not available
5.73 / 22.56	5 / 13	102.80
5.06 / 34.16	5 / 18	104.41
4.14 / 45.83	5 / 23	106.01
3.50 / 53.86	5 / 28	107.06

Table 13. RAM utilization in relation to TPS for measurement set B

Throughput (TPS) **DB2/2 / DB2/MVS**	**DB2 Clients** **DB2/2 / DB2/MVS**	**RAM on PC500 (MBytes)**
8.91 / 8.59	7 / 8	100.57
8.03 / 18.19	7 / 13	103.17
6.74 / 28.29	7 / 18	105.08
5.68 / 38.70	7 / 23	106.48

Table 14. RAM utilization in relation to TPS for measurement set C

Throughput (TPS) **DB2/2 / DB2/MVS**	**DB2 Clients** **DB2/2 / DB2/MVS**	**RAM on PC500 (MBytes)**
9.89 / 7.24	9 / 8	102.25
8.73 / 15.55	9 / 13	103.64
7.39 / 25.73	9 / 18	105.57
6.38 / 33.34	9 / 23	107.06

As a general consideration:

- **For this type of configuration, plan on about 350 KBytes of RAM per client that uses DDCS for OS/2 and plan on about 250 KBytes of RAM per client that accesses DB2 for OS/2.**
- **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 affect 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. 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 following values are set in regards to RQRIOBLK sizes and RUSIZEs:

- 4K RUSIZEs in VTAM
- 32K RQRIOBLK definition in DDCS/2
- 4K RQRIOBLK definitions in the OS/2 clients
- 4K DOS_RQRIOBLK definitions in the DOS clients

The next two sections describe the effects on the NCP and LAN during the measurements.

NCP

Table 15, Table 16 on page 18, and Table 17 on page 18 describe the effects of the workload on the 3745 running the NCP. These measurements do not stress the 3745 beyond 55 percent CCU utilization at a throughput rate of 53.86 transactions per second. The data throughput in the NCP reached 47,050 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 OS/2.

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 DELAY parameter is set to zero. This has the beneficial effect of controlling the CCU utilization with little response time effects at the clients at the higher levels of throughput. The lower levels of throughput experience mildly higher levels of delay than if the GROUP LNCTL=CA DELAY is set to zero.

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 15, Table 16 on page 18, and Table 17 on page 18 provide the numeric values for each measurement set.

Table 15. 3745 utilization in relation to TPS for measurement set A

Throughput (TPS) DB2/2 / DB2/MVS	**Bytes Per Second**	**% 3745 Utilization**
6.47 / 8.96	8,086	11.50%
5.73 / 22.56	19,597	25.00%
5.06 / 34.16	29,984	36.17%
4.14 / 45.83	40,557	48.00%
3.50 / 53.86	47,050	54.50%

Table 16. 3745 utilization in relation to TPS for measurement set B

Throughput (TPS) DB2/2 / DB2/MVS	Bytes Per Second	% 3745 Utilization
8.91 / 8.59	7,670	11.00%
8.03 / 18.19	16,169	21.00%
6.74 / 28.29	24,953	30.86%
5.68 / 38.70	34,161	40.83%

Table 17. 3745 utilization in relation to TPS for measurement set C

Throughput (TPS) DB2/2 / DB2/MVS	Bytes Per Second	% 3745 Utilization
9.89 / 7.24	6,602	10.33%
8.73 / 15.55	13,533	18.00%
7.39 / 25.73	23,119	28.83%
6.38 / 33.34	29,200	36.67%

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 a 4.5 MByte channel, which is under-utilized. If this is replaced by a 56Kbps, 256Kbps, etc..., speed line, the throughput is limited.

As a 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.

As a 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 following throughput rate mix (DB2 for OS/2 / DB2 for MVS/ESA):

6.47 / 8.96 (measurement set A)

- DB2 for OS/2
  - 96 SQL statements per second of which 42 per second are SELECTs
- DB2 for MVS/ESA
  - 30 SELECT statements per second
  - 25 INSERT statements per second
  - 32 UPDATE statements per second
  - 1 DELETE statements per second
  - 39 OPEN statements per second
  - 23 CLOSE statements per second
  - 77 FETCH statements per second

3.50 / 53.86 (measurement set A)

- DB2 for OS/2
  - 54 SQL statements per second of which 24 per second are SELECTs
- DB2 for MVS/ESA
  - 187 SELECT statements per second
  - 153 INSERT statements per second
  - 195 UPDATE statements per second
  - 10 DELETE statements per second
  - 237 OPEN statements per second
  - 138 CLOSE statements per second
  - 468 FETCH statements per second

8.91 / 8.59 (measurement set B)

- DB2 for OS/2
  - 131 SQL statements per second of which 59 per second are SELECTs
- DB2 for MVS/ESA
  - 29 SELECT statements per second
  - 24 INSERT statements per second
  - 30 UPDATE statements per second
  - 1 DELETE statements per second
  - 36 OPEN statements per second
  - 21 CLOSE statements per second
  - 72 FETCH statements per second

5.68 / 38.70 (measurement set B)

- DB2 for OS/2
  - 83 SQL statements per second of which 37 per second are SELECTs
- DB2 for MVS/ESA
  - 136 SELECT statements per second
  - 113 INSERT statements per second
  - 144 UPDATE statements per second
  - 7 DELETE statements per second
  - 173 OPEN statements per second
  - 102 CLOSE statements per second
  - 343 FETCH statements per second

9.89 / 7.24 (measurement set C)

- DB2 for OS/2
  - 147 SQL statements per second of which 66 per second are SELECTs
- DB2 for MVS/ESA
  - 25 SELECT statements per second
  - 20 INSERT statements per second
  - 26 UPDATE statements per second
  - 1 DELETE statements per second
  - 31 OPEN statements per second
  - 18 CLOSE statements per second
  - 62 FETCH statements per second

6.38 / 33.34 (measurement set C)

- DB2 for OS/2
  - 91 SQL statements per second of which 42 per second are SELECTs
- DB2 for MVS/ESA
  - 116 SELECT statements per second
  - 93 INSERT statements per second
  - 120 UPDATE statements per second
  - 6 DELETE statements per second
  - 145 OPEN statements per second
  - 85 CLOSE statements per second
  - 288 FETCH statements per second

As a 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 percent of the transactions complete within the identified maximum response time for the identified throughput. There are three lines in the graphs to represent the three measurement sets. The solid line is measurement set A where the workload against DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload against DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload against DB2 for OS/2 is held constant at 6 sessions.

Figure 3 on page 21 displays the aggregate response time of the transactions on DB2 for MVS/ESA for the identified throughput. The maximum response time for 90 percent of all the transactions is between 1.0 and 2.0 seconds. As an aggregate of all the transactions, the response time did not significantly degrade when scaling the workload on DB2 for OS/2 or on DB2 for MVS/ESA. The primary effects of the scaling is the degradation of the maximum throughput.

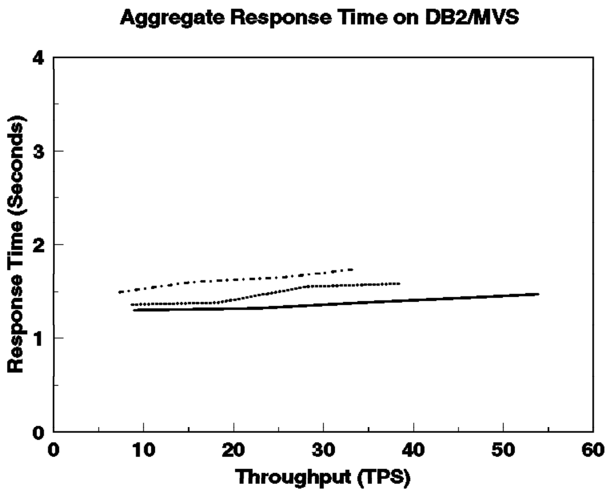

Figure 3. Aggregate Response Time For DB2 for MVS/ESA

Figure 4 on page 22 displays the aggregate response time of the transactions on DB2 for OS/2 for the identified throughput. The maximum response time for 90 percent of all the transactions is between 1.0 and 3.0 seconds. As an aggregate of all the transactions, the response times degrade approximately 0.5 seconds between each measurement set as DB2 for OS/2 scales. The response times degrade approximately 1.0 second within each measurement set as DB2 for MVS/ESA scales to the point where the PC500 is 100 percent utilized. Notice that the horizontal axis decreases throughput when moving from left to right. This is to intuitively orient the graph because the throughput and response times decrease simultaneously as the DB2 for MVS/ESA workload increases.

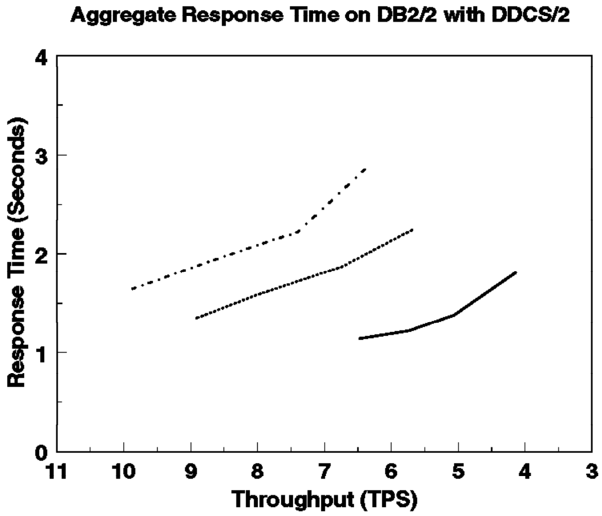

Figure 4. Aggregate Response Time For DB2 for OS/2

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.

Transaction Tx1 only runs on DB2 for MVS/ESA.

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 5 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 percent of the Tx1 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx1 on DB2 for MVS/ESA

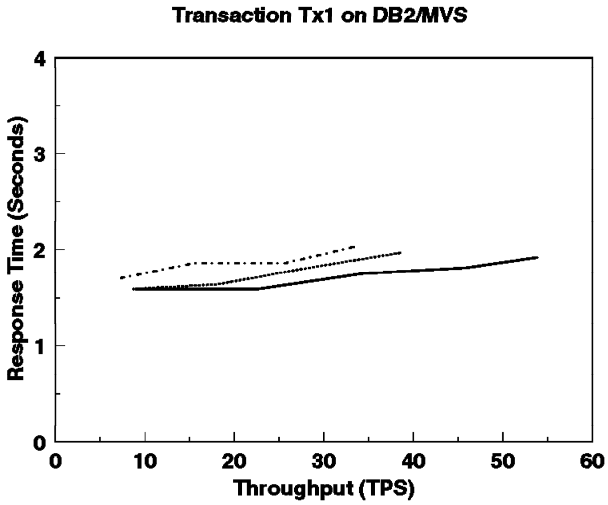

Figure 5. Response Time On DB2 for MVS/ESA For Transaction Tx1

The maximum response time for 90 percent of the Tx1 transactions is between 1.5 and 2.0 seconds. Transaction Tx1 shows a gradual response time degradation for each measurement set which is primarily due to increased locking and latching activity on DB2 for MVS/ESA. Between each measurement set, there is a slight increase in response times. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 18 displays the transactions that have potential lock contention with transaction Tx1 on the defined tables. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

Table 18. 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.

Transaction Tx2 runs on DB2 for MVS/ESA and DB2 for OS/2.

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 6 and Figure 7 on page 26 describe 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 percent of the Tx2 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx2 on DB2 for MVS/ESA

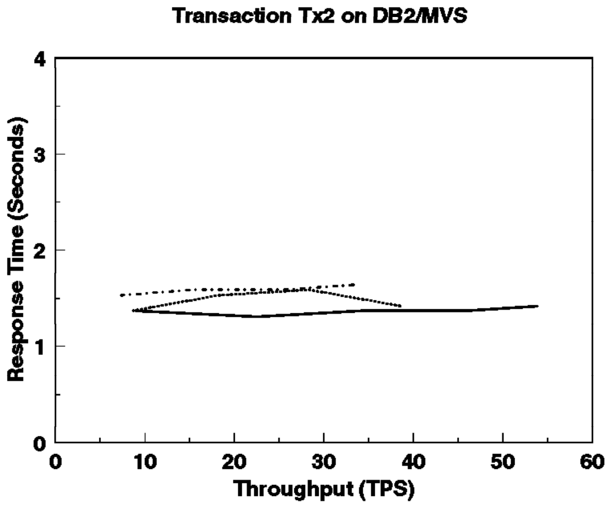

Figure 6. Response Time On DB2 For MVS/ESA For Transaction Tx2

Figure 6 displays the transaction Tx2 response times on DB2 for MVS/ESA for the measurement sets. The maximum response time for 90 percent of the Tx2 transactions is between 1.25 and 1.50 seconds. Transaction Tx2 shows no response time degradation for each measurement set. Between the measurement sets, there is a slight increase in response times. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 19 on page 26 displays the transactions that have potential lock contention with transaction Tx2 on the defined tables. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

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 Tx2 on DB2 for OS/2

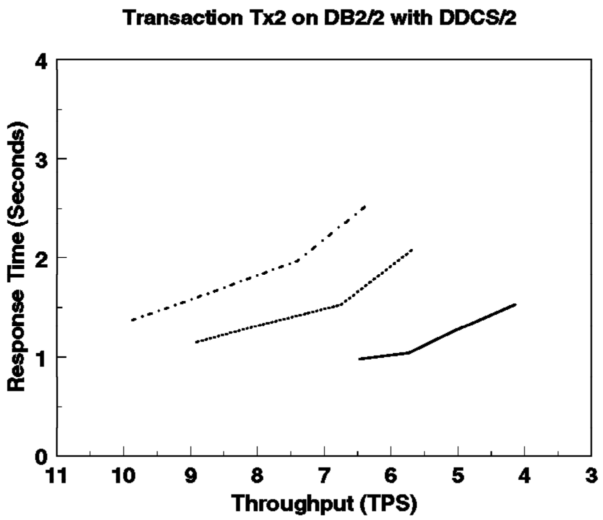

Figure 7. Response Time On DB2 For OS/2 For Transaction Tx2

Figure 7 displays the transaction Tx2 response times on DB2 for OS/2 for the measurement sets. The maximum response time for 90 percent of the Tx2 transactions is between 1.0 and 2.5 seconds. The response times degrade approximately 0.5 seconds between each measurement set. The response times within the measurement sets degrade approximately 0.5 to 1.0 seconds as DB2 for MVS/ESA scales to the point where the PC500 CPU is fully utilized. Notice that the horizontal axis decreases throughput when moving from left to right. This is to intuitively orient the graph because the throughput and response times decrease simultaneously as the DB2 for MVS/ESA workload increases.

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 8 on page 28 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 percent of the Tx3 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx3 on DB2 for MVS/ESA

Figure 8 on page 28 displays the transaction Tx3 response times on DB2 for MVS/ESA for the measurement sets.

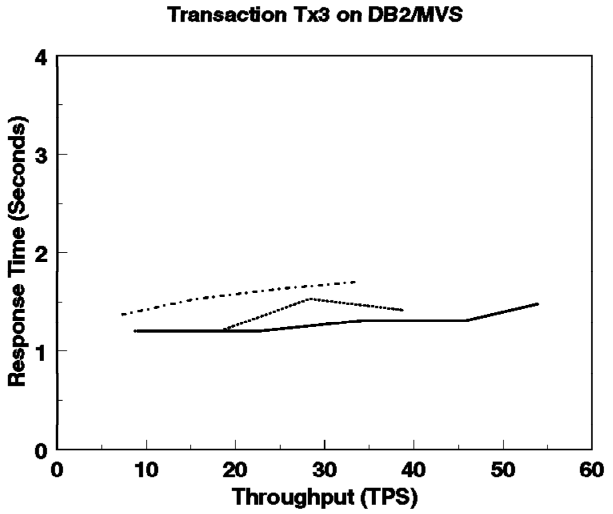

Figure 8. Response Time For Transaction Tx3

The maximum response time for 90 percent of the Tx3 transactions is between 1.0 and 1.5 seconds. Transaction Tx3 shows a gradual response time degradation for each measurement set which can be locking and latching on DB2 for MVS/ESA. Between the measurement sets, there is a slight increase in response times. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 20 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. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

Table 20. 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 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 percent of the Tx4 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx4 on DB2 for MVS/ESA

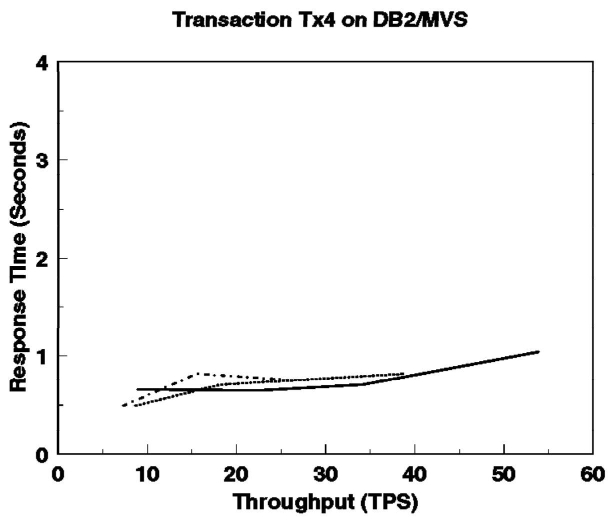

Figure 9. Response Time On DB2 For MVS/ESA For Transaction Tx4

Figure 9 displays the transaction Tx4 response times on DB2 for MVS/ESA for the measurement sets. The maximum response time for 90 percent of the Tx4 transactions is between 0.5 and 1.0 second. Transaction Tx4 shows almost no response time degradation for each measurement set or between the measurement sets. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA. Some degradation begins to occur in measurement set A which can be attributed to some locking and latching activity on table T5 and possibly some increased logging activity.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 21 on page 30 displays the transactions that have potential lock contention with transaction Tx4 on the defined tables. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

Table 21. 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 10 on page 31 and Figure 11 on page 32 describe 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 percent of the Tx5 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx5 on DB2 for MVS/ESA

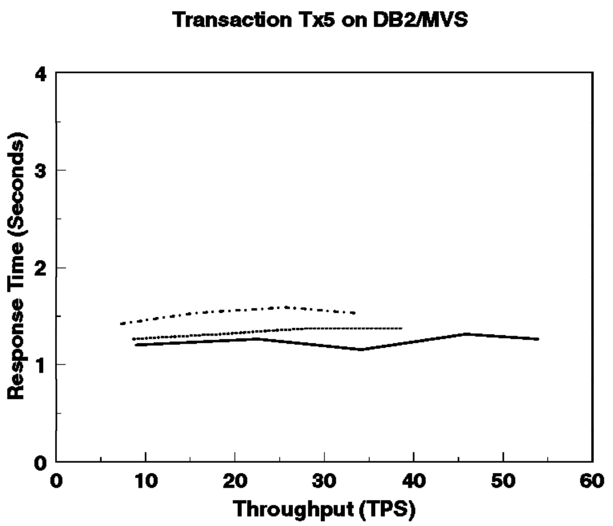

Figure 10. Response Time On DB2 For MVS/ESA For Transaction Tx5

Figure 10 displays the transaction Tx5 response times on DB2 for MVS/ESA for the measurement sets. The maximum response time for 90 percent of the Tx5 transactions is between 1.0 and 1.5 seconds. Transaction Tx2 shows no response time degradation for each measurement set. The T5 table is the only table that Tx5 accesses. There is no lock contention or I/O contention against this table because it is large and it is 100 percent loaded into memory. Between the measurement sets, there is a slight increase in response times. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 22 displays the transactions that have potential lock contention with transaction Tx5 on the defined tables.

Table 22. Potential Lock Contention For Transaction Tx5

	Tx1	Tx2	Tx3	Tx4	Tx5	Tx6	Tx7
T5				yes			

Transaction Tx5 on DB2 for OS/2

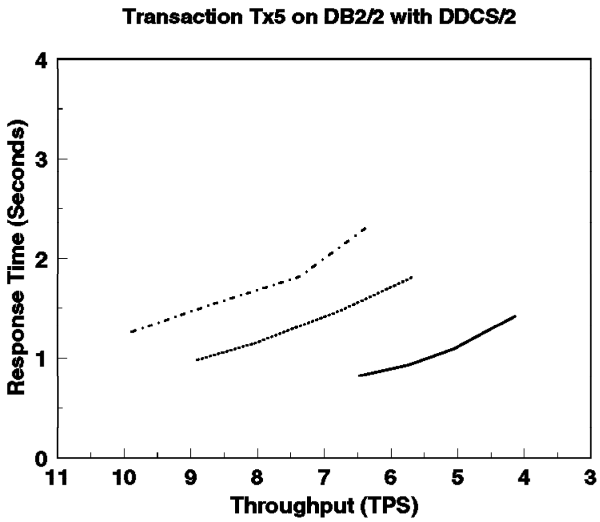

Figure 11. Response Time On DB2 For OS/2 For Transaction Tx5

Figure 11 displays the transaction Tx5 response times on DB2 for OS/2 for the measurement sets. The maximum response time for 90 percent of the Tx5 transactions is between 1.0 and 2.5 seconds. The response times degrade approximately 0.5 seconds between each measurement set. The response times within the measurement sets degrade approximately 0.5 to 1.0 seconds as DB2 for MVS/ESA scales to the point where the PC500 CPU is fully utilized. Notice that the horizontal axis decreases throughput when moving from left to right. This is to intuitively orient the graph because the throughput and response times decrease simultaneously as the DB2 for MVS/ESA workload increases.

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 12 on page 33 and Figure 13 on page 34 describe 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 percent of the Tx6 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx6 on DB2 for MVS/ESA

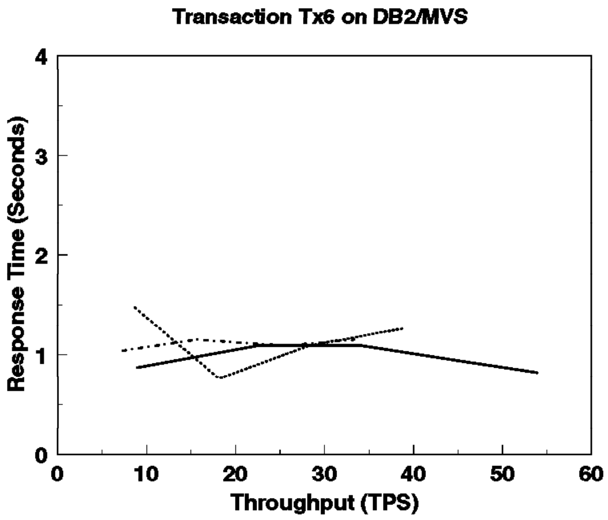

Figure 12. Response Time On DB2 For MVS/ESA For Transaction Tx6

Figure 12 displays the transaction Tx6 response times on DB2 for MVS/ESA for the measurement sets. The maximum response time for 90 percent of the Tx6 transactions is between 0.5 and 1.5 seconds. Transaction Tx6 shows no response time degradation for each measurement set or between the measurement sets. The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 23 displays the transactions that have potential lock contention with transaction Tx6 on the defined tables. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

Table 23. Potential Lock Contention For Transaction Tx6

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

Transaction Tx6 on DB2 for OS/2

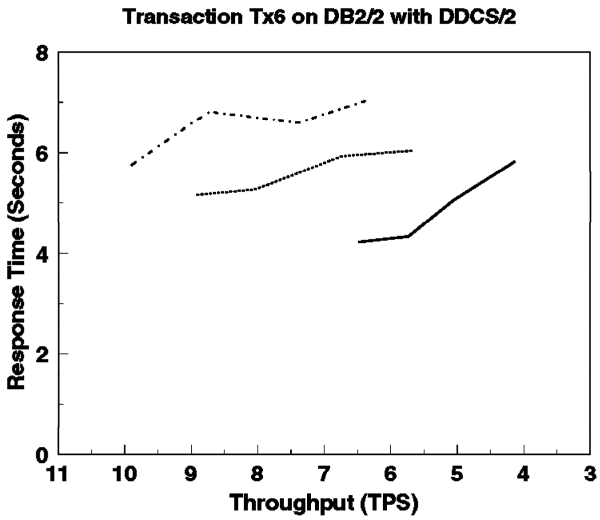

Figure 13. Response Time On DB2 For OS/2 For Transaction Tx6

Figure 13 displays the transaction Tx6 response times on DB2 for OS/2 for the measurement sets. The maximum response time for 90 percent of the Tx6 transactions is between 4.0 and 7.0 second. The response times degrade approximately 1.0 second between each measurement set. The response times within the measurement sets degrade approximately 1.0 to 2.0 seconds as DB2 for MVS/ESA scales to the point where the PC500 CPU is fully utilized. Transaction Tx6 is CPU intensive due to the join between tables T8 and T9. DDCS is also CPU intensive, so between both CPU demands, there are large effects on response times.

Notice that the horizontal axis decreases throughput when moving from left to right. This is to intuitively orient the graph because the throughput and response times decrease simultaneously as the DB2 for MVS/ESA workload increases.

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 14 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 percent of the Tx7 transactions achieve this response time or better. The solid line is measurement set A where the workload for DB2 for OS/2 is held constant at 2 sessions. The dot line is measurement set B where the workload for DB2 for OS/2 is held constant at 4 sessions. The dot/dash line is measurement set C where the workload for DB2 for OS/2 is held constant at 6 sessions.

Transaction Tx7 on DB2 for MVS/ESA

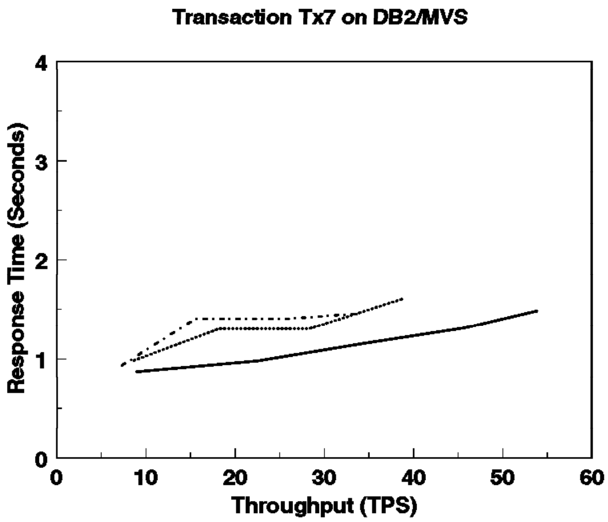

Figure 14. Response Time On DB2 for MVS/ESA For Transaction Tx7

Figure 14 displays the transaction Tx7 response times on DB2 for MVS/ESA for the measurement sets. The maximum response time for 90 percent of the Tx7 transactions is between 0.75 and 1.50 seconds.

The response times degrade less than .25 seconds between each measurement set. The response times within the measurement sets degrade approximately .75 seconds as DB2 for MVS/ESA scales to the point where the PC500 CPU is fully utilized. Transaction Tx7 is CPU intensive due to the join between The primary effects of the DB2 for OS/2 scaling is the degradation of the maximum throughput on DB2 for MVS/ESA.

Although lock and latch contention do not play any significance in these measurements for this transaction, Table 24 displays the transactions that have potential lock contention with transaction Tx7 on the defined tables. Lock contention becomes a factor at throughput rates which are greater than achieved in this measurement.

Table 24. 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 for a client/server relational database configuration that includes DB2 for MVS/ESA, DB2 for OS/2, DDCS for OS/2, and DB2/CAE/Windows with PowerBuilder. DB2 for OS/2 is running concurrently on the same hardware with DDCS for OS/2. 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/2) 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/2?
  - Through the NCP?
- How much processor capacity would I need for my environment considering the throughput and transaction mix? DDCS/2, 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

