DB2 Family: Client/Server Performance Measurement Series
DB2/MVS <> DDCS/6000 <> DB2/CAE/Windows

Document Number Version 1.3

June 24, 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:

- DatagLANce
- AIX *
- DB2 *
- Distributed Database Connection Services/6000
- IBM *
- MVS/ESA
- Netview *
- OS/2 *
- PowerPC
- PowerPC 601 *
- 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
- Serge Limoges
- Mark Ryan
- Akira Shibamiya
- 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 . . . 6
Physical location . . . 7
Buffer Pool Allocation . . . 7

4. Results . . . 9
Client Equivalence . . . 9
CPU Utilization . . . 10
Network . . . 12
  NCP . . . 12
  LAN . . . 14
DB2 Utilization . . . 15
Transactions . . . 15
  Transaction Tx1 . . . 16
  Transaction Tx2 . . . 18
  Transaction Tx3 . . . 19
  Transaction Tx4 . . . 21
  Transaction Tx5 . . . 22
  Transaction Tx6 . . . 23
  Transaction Tx7 . . . 25

5. Conclusions . . . 27

6. Other sources of information . . . 29

© 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 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 to perform OLTP type transactions. 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.

© Copyright IBM Corp. 1996

2. Configuration

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

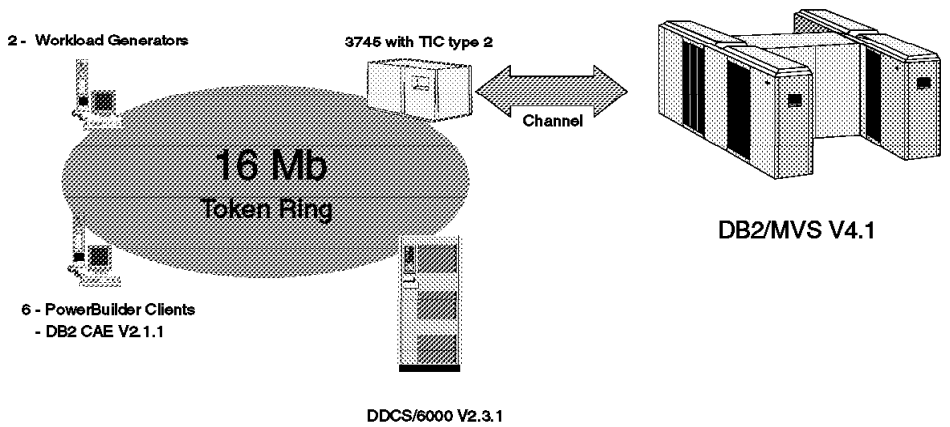

Figure 1. Configuration with DB2 for MVS/ESA

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.
- **Distributed Database Connection Services for AIX (DDCS/6000)** Version 2.3.1 running on a C10 PowerPC 601 with a 80 Mhz processor and 256 MBytes of memory.
- **6 - PowerBuilder Version 4.0.2 with DB2 CAE for Windows Version 2.1** 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.

© Copyright IBM Corp. 1996

- **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/6000
- LU 6.2 is in use between the OS/2 clients and DDCS/6000
- LU 6.2 is in use between DDCS/6000 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
- LAN: DatagLANce Network Analyzer for Ethernet and Token Ring for OS/2

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.

A series of measurements are performed, with each subsequent measurement providing an increased workload to the configuration. 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 two OS/2 workload generators. The OS/2 workload generators run the exact same transaction mix as the six PowerBuilder clients. The only difference is the management of screen I/O at the client. The PowerBuilder clients simulate a real application, where the OS/2 workload generators do not perform screen I/O.

Transaction Characteristics

The workload 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**

  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**

© Copyright IBM Corp. 1996

Performs 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. 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

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 used for all indexes. 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 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 and spreading the partitions out 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 percent 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	9394-1	T8 Table (Partitions 9-16)
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)

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 presented in relation to the overall throughput achieved at each workload level. This is described as transactions per second (TPS). As the workload increased, the throughput increased. The resultant effects on each component is then presented based on the current achieved throughput. For example, when achieving 57.1 transactions per second:

- the 3745 running the NCP is 74 percent utilized
- the transfer rate through the 3745 running the NCP is 50,922 bytes per second
- the C10 running DDCS/6000 is 97 percent utilized
- etc...

The process of tuning the application and database is ongoing. These measurements achieved a throughput rate of 57.1 transactions per second. At this point, several factors are beginning to show signs of stress. Primarily, there are CPU contraints on the C10 running DDCS/6000. The CPU utilization on the C10 reaches 97 percent. Secondarily there are large levels of locking and latching occurring. Also the device activity rate and the average response time accessing disks 37-44 is increasing. These disks are 3380-E and they contain 8 of the 16 partitions for the T7 Table and the T9 Table. The other 8 partitions are on 3390-3 disks, which are faster and the device rate and average response time reflects this.

Multiple iterations with tuning occurred prior to achieving these results. These results are probably not the absolute best, but to continue tuning and pushing for the absolute limits is not the goal of this project. Other benchmarking projects deal with those goals. The goal of this project is to provide a reasonably tuned environment. An important point to note is that the tuning of DB2 for MVS/ESA in this client/server configuration should be the same as tuning DB2 for MVS/ESA in a local configuration. This is true because of the use of stored procedures.

Client Equivalence

The configuration includes 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 screen. This adds to the response time, which decreases the throughput for the PowerBuilder clients. This section estimates a rough number of total PowerBuilder clients based on the throughput obtained by the OS/2 clients plus the throughput of the actual PowerBuilder clients.

Table 4 on page 10 describes the number of clients activated to achieve the throughput defined in column one. Column two describes the number of PowerBuilder clients. This is always 6. Column three describes the number of OS/2 clients. This is increased by ten for each measurement run. Column four is the actual number of clients, which is the sum of columns two and three. Column five

© Copyright IBM Corp. 1996

is the estimated number of PowerBuilder clients. This column is derived by evaluating the number of equivalent PowerBuilder clients each OS/2 client represents in terms of throughput. The equation for evaluating this is:

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

If
 OS/2 client tps = 1.41 tps
 PowerBuilder client tps = 0.128 tps
 # OS/2 clients = 40
 # PowerBuilder clients = 6
then

$$\text{Estimated PowerBuilder Clients} = \left|\frac{1.41\ \text{tps}}{0.128\ \text{tps}}\right| * 40\ \text{clients} + (6\ \text{clients})$$

$$= 447\ \text{clients}$$

Table 4. Number of clients in relation to Aggregate Throughput (TPS)

Throughput (TPS)	PowerBuilder Clients	OS/2 Clients	Actual Clients	Est. PowerBuilder Clients
24.53	6	10	16	179
43.67	6	20	26	340
53.68	6	30	36	402
57.1	6	40	46	447

The values obtained in these measurements are obtained based on the actual number of clients as reflected in column four. When evaluating the estimated number of PowerBuilder clients, other capacity planning factors need to be addressed. For example, DB2 for MVS/ESA V4.1 can handle up to 25,000 connections, with up to 1,999 active threads, and each client that passes through DDCS/6000 should be allowed approximately 500KB of RAM on DDCS/6000. These factors need to be taken in 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 Utilization

This section describes CPU utilization and RAM utilization as the workload increases for each measurement. Table 5 on page 11 and Figure 2 on page 11 describe the CPU utilization and Table 6 on page 12 describes the RAM utilization for the C10 running DDCS/6000.

At a throughput rate of 57.1 transactions per second, the 9121-742 processor reaches 48.57 CPU utilization. The 9121-742 is capable of scaling up to 98 percent CPU utilization, therefore, there is excess processor power available.

The C10 running DDCS/6000 reaches CPU utilization of 97 percent at a throughput rate of 57.1 transactions per second. The C10 is identified as the primary bottleneck that prevents additional throughput.

Table 5. CPU utilization in relation to TPS

Throughput (TPS)	C10 Running DDCS/6000 - User	C10 Running DDCS/6000 - System	9121-742 Running DB2/MVS
24.53	8%	37%	22.15%
43.68	14%	57%	38.07%
53.68	17%	72%	45.72%
57.1	17%	80%	48.57%

In Figure 2 the dashed line represents the system CPU utilization of the C10, the dotted line represents the user CPU of the C10, and the solid line represents the CPU utilization of the 9121-742.

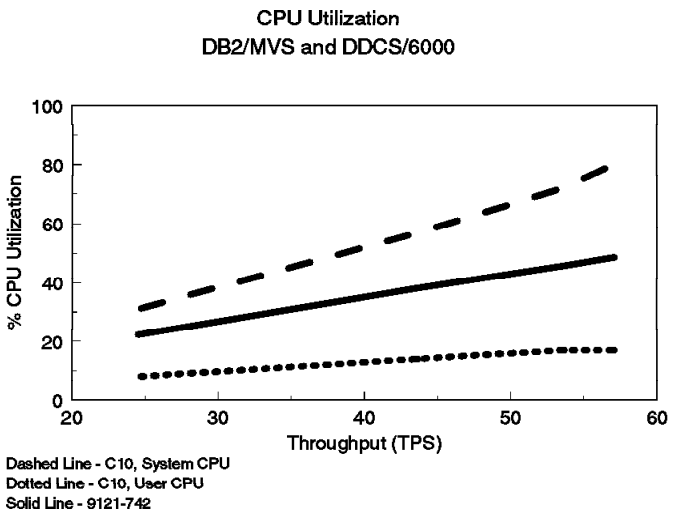

Figure 2. CPU Utilization in Relation to Throughput

Table 6 on page 12 shows the increase in RAM utilization for the C10 as the number of clients increases. The initial value of 75.54 MBytes for 16 clients should not be used as an indication of RAM usage. This number is a factor of the types and numbers of applications activated on the C10. Treat 75.54 as the base number which increases as clients are added. The average amount of RAM used per client is approximately 450 KBytes. Conservatively, plan for 500 KBytes for additional clients.

Table 6. RAM utilization on C10

Number of Clients	MBytes of RAM
16	75.54
26	80.33
36	84.78
46	89.31

For general consideration:

- **DDCS/6000 should run on a fast processor with enough memory to handle the number of expected clients. Plan for 500 KBytes for each client passing through DDCS/6000. The base level of RAM for the C10 should be based on the number of applications and RAM requirements for those applications.**
- **Hard disk capacity is not an issue for DDCS/6000 because the processing requirements do not require large amounts of disk capacity.**

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 next two sections describe the effects on the NCP and LAN during the measurements. For this mix of transactions, the message lengths are never greater than 1,200 bytes, so the following values are set:

- 4K RUSIZEs in VTAM
- 4K RQRIOBLK definition in DDCS/6000
- 4K RQRIOBLK definitions in the OS/2 clients
- 4K DOS_RQRIOBLK definitions in the DOS clients

NCP

Table 7 on page 13 describes the effects of the workload on the 3745 running the NCP. The 3745 reached 74 percent utilization at a throughput rate of 57.1 transactions per second. The data throughput in the NCP reached 50,922 bytes per second.

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 DELAY parameter values 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 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 levels of delay than if the GROUP LNCTL=CA DELAY parameter equals zero.

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

Table 7. 3745 utilization in relation to TPS

Throughput (TPS)	Bytes Per Second	% 3745 Utilization
24.53	24,105	35
43.67	39,379	58
53.68	47,935	70
57.1	50,922	74

These results can be used to analyze the line capacity requirements between the LAN and the DB2 for MVS/ESA 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.

Additional measurements using a 3174-61R attached to the LAN with a 56Kbps connection to the 3745 achieved a data throughput rate of approximately 4847 bytes per second (See Figure 3 on page 14). This is a 69 percent utilization of the line. At this utilization, the throughput rate reached approximately 5 transactions per second. Although the line has more capacity, the increase in transaction per second rate begins to level off. For this report, there is not enough information to identify why 69 percent is the threshold. Is it due to message sizes, hardware constraints, protocol constraints, etc...? Additional research or measurements are needed to identify the cause.

Note: The 3174 measurements use a DDCS/2 on a PC/500. Although this is different than using DDCS/6000 on a C10, the results are the same regardless which type of component is driving the throughput.

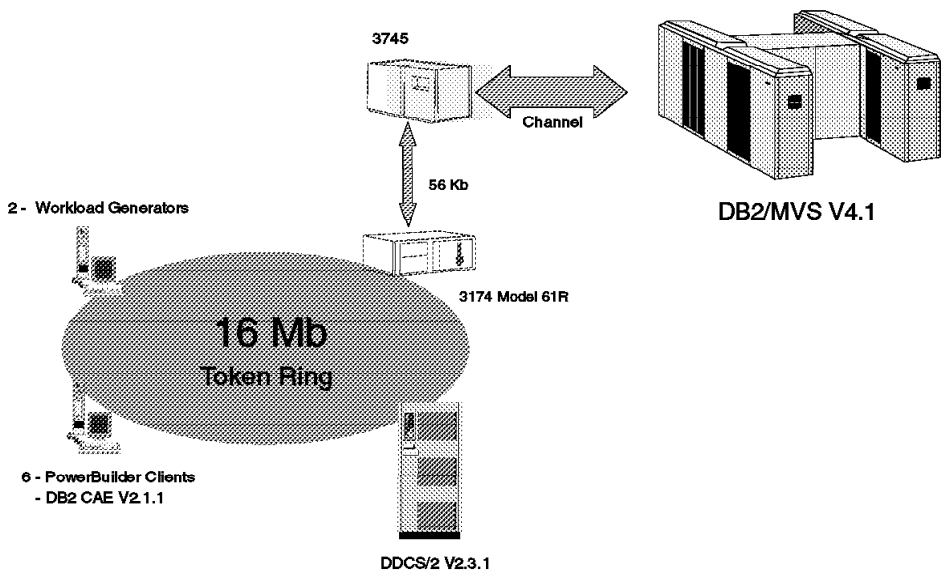

Figure 3. Configuration with 3174

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

Table 8 on page 15 shows the percentage of LAN utilization and byte rate as the workload is increased. For these measurements, the LAN capacity does not factor in. The 16 Mbit LAN is lightly utilized. Typically, a production environment will have more activity on the LAN. Use these numbers to analyze the additional load introduced on the LAN due to the mix of transactions used in this configuration.

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.

Table 8. LAN utilization in relation to TPS

Throughput (TPS)	Bytes Per Second	% LAN Utilization
24.53	86,985	4.35
43.67	148,656	7.43
53.68	178,345	8.92
57.1	190,265	9.51

DB2 Utilization

The following rate and mix of SQL statements occur when achieving a transaction throughput rate of 57.1 transactions per second:

- 203 SELECT statements per second
- 159 INSERT statements per second
- 205 UPDATE statements per second
- 11 DELETE statements per second
- 248 OPEN statements per second
- 143 CLOSE statements per second
- 495 FETCH statements per second

There are two network flows involved in each transaction. The first flow calls the stored procedure. The second flow commits or rolls back the transaction. Locks acquired during the stored procedure processing are held until the commit or rollback is received. These locks are held for the amount of time it takes to send the stored procedure response back to the client, the client processes the response and issues the commit or rollback. DDCS/6000 running at maximum CPU utilization increases this delay. Delays result in locking contention at DB2 for MVS/ESA.

Secondarily, the average response time accessing 8 of the 16 partitions for the T7 table and T9 table is reaching levels that have a delay effect due to those partitions residing on slower speed disk drives (3380-E).

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 percent of the transactions complete within the identified maximum response time for the identified throughput. Since this is the 90th percentile, this is an indication of the greatest response time when sampling the fastest 90 percent of the transactions. Most transactions achieve a better response time. Figure 4 on page 16 displays the aggregate response time of the transactions for the identified throughput.

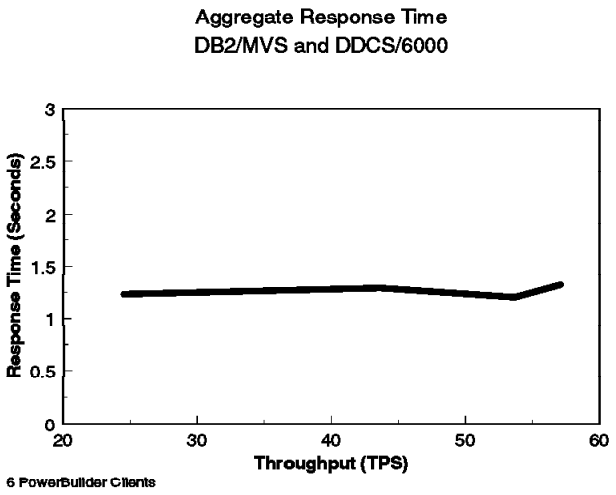

Figure 4. Aggregate Response Time

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. Locking contention is the primary factor for the response time degradation with access to slower speed disk drives also becoming a factor.

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 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. The solid line represents the point where 90 percent of the Tx1 transactions achieved this response time or better.

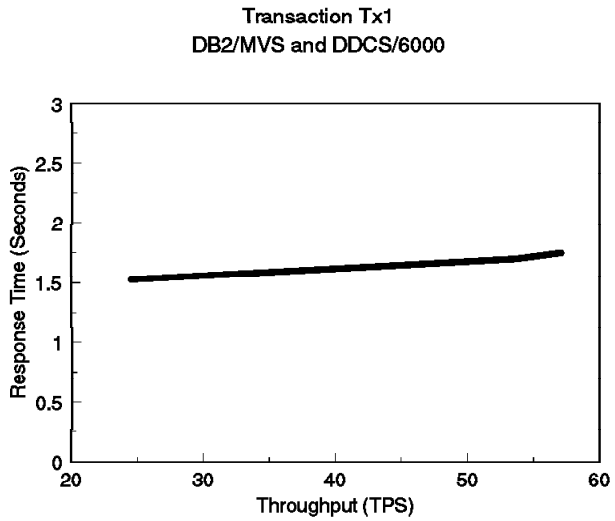

Figure 5. Response Time For Transaction Tx1

The maximum response time for 90 percent of the Tx1 transactions is between 1.4 and 2.0 seconds. Transaction Tx1 shows a gradual response time degradation as the workload increases. This can be attributed to lock contention on the T3 table. Both the Tx3 and Tx6 transactions update the T3 table. Transaction Tx1 also inserts the T9 table and half of those rows are resident on slower disk drives. A greater reponse time degradation begins at 54 transactions per second and this is attributed to CPU contraints on the C10 running DDCS/6000.

Table 9 on page 18 displays the transactions that have potential lock contention with transaction Tx1 on the defined tables.

Table 9. 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 and 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 on page 19 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. The solid line represents the point where 90 percent of the Tx2 transactions achieved this response time or better.

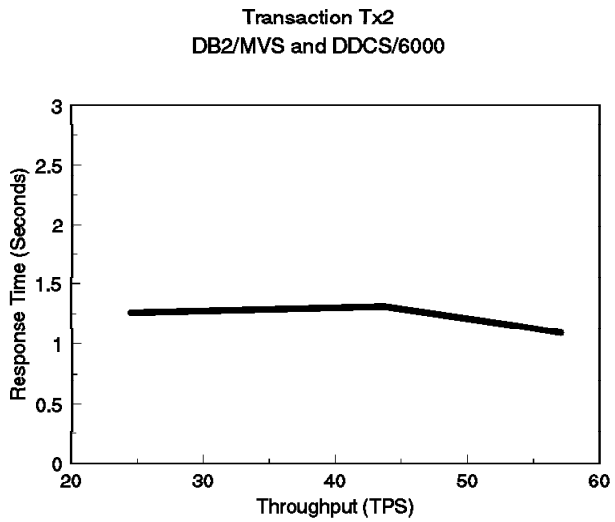

Figure 6. Response Time For Transaction Tx2

The maximum response time for 90 percent of the Tx2 transactions is between 1.0 and 1.5 seconds. The response time improvement can be attributed to the data becoming available in the buffers due to the increase workload.

Table 10 displays the transactions that have potential lock contention with transaction Tx2 on the defined tables.

Table 10. 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 7 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. The solid line represents the point where 90 percent of the Tx3 transactions achieved this response time or better.

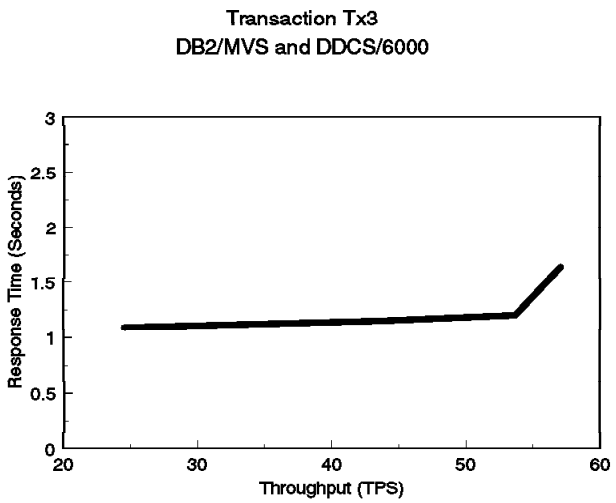

Figure 7. Response Time For Transaction Tx3

The maximum response time for 90 percent of the Tx3 transactions is between 1 and 1.75 seconds. Transaction Tx3 shows a response time degradation at 54 transactions per second. This is attributed to delays introduced by CPU contraint on the C10 running DDCS/6000 which causes delays acquiring exclusive locks on tables T2, T3, and T7. Both the T2 and T3 tables are small. The Tx7 transaction updates the T7 table. The Tx1 transaction updates the T3 table and reads tables T2 and T7. The Tx2 transaction reads the T7 table.

Table 11 on page 21 displays the transactions that have potential lock contention with transaction Tx3 on the defined tables.

Table 11. 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 procecdure 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 8 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. The solid line represents the point where 90 percent of the Tx4 transactions achieved this response time or better.

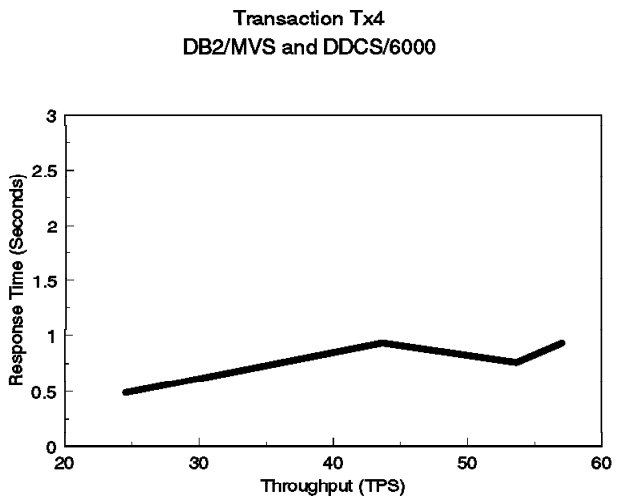

Figure 8. Response Time for Transaction Tx4

The maximum response time for 90 percent of the Tx4 transactions is less than 1 second. The response time degradation is gradual for the scope of the measurements. Tx4 updates the T5 table, which is 100% preloaded into memory. Transaction Tx4 runs 1% of the time, therefore, the sampling is low which can result in sampling anomalies.

Table 12 displays the transactions that have potential lock contention with transaction Tx4 on the defined tables.

Table 12. 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 9 on page 23 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. The solid line represents the point where 90 percent of the Tx5 transactions achieved this response time or better.

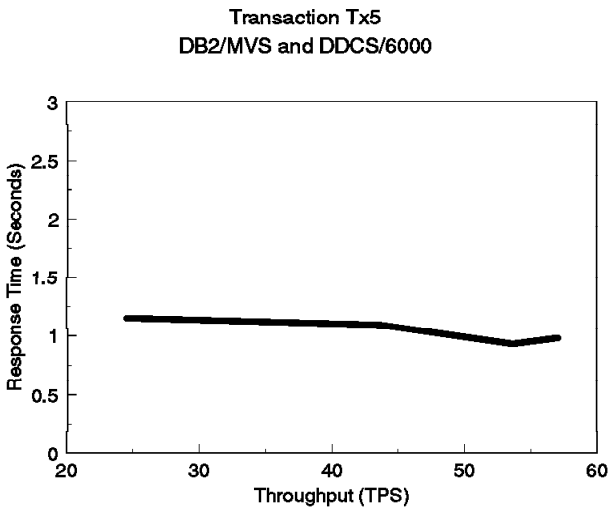

Figure 9. Response Time of Transaction Tx5

The maximum response time for 90 percent of the Tx5 transactions is less than 1.25 seconds. Tx5 shows a little change in response time. 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% loaded into memory.

Table 13 displays the transactions that have potential lock contention with transaction Tx5 on the defined tables.

Table 13. 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 10 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. The solid line represents the point where 90 percent of the Tx6 transactions achieved this response time or better.

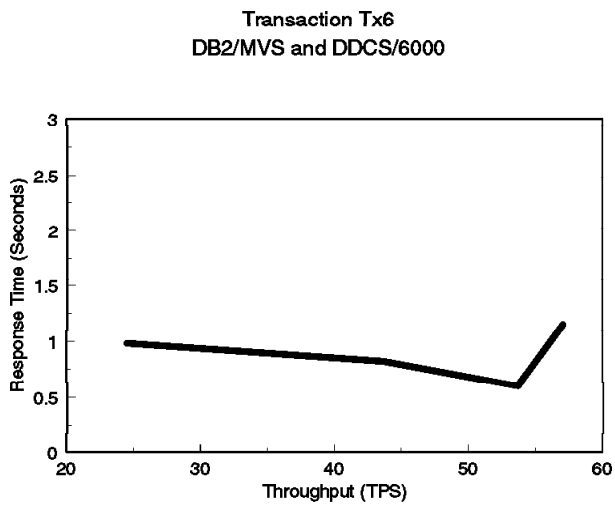

Figure 10. Response Time for Transaction Tx6

The maximum response time for 90 percent of the Tx6 transactions is between 0.5 and 1.0 seconds. Tx6 show a gradual improvement in response time until 54 transactions per second. This is attributed to data becoming available in the buffers as the workload increases. At 54 transactions per second the response time degrades. This can be attributed to CPU contraint on the C10 running DDCS/6000 which causes locking contention on the T3 table. The Tx1 and Tx7 transactions perform updates on the T3 table.

Table 14 displays the transactions that have potential lock contention with transaction Tx6 on the defined tables.

Table 14. 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 11 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. The solid line represents the point where 90 percent of the Tx7 transactions achieved this response time or better.

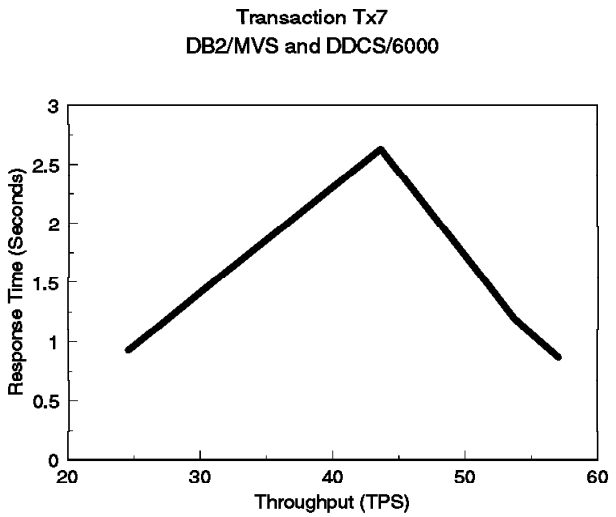

Figure 11. Response Time For Transaction Tx7

The maximum response time for 90 percent of the Tx7 transactions is between 0.75 and 2.5 seconds. The spike in response time at 40 transactions can be attributed to the random sampling that created a sample with lock contention. Since Tx7 runs only 2% of the time, a single sample with contention can have drastic effects on the final numbers.

Note: Other measures not included in this report show this transactions response time gradually degrades as the workload increases. Spikes such as these are not uncommon for small transaction samples such as found with Tx7 and Tx4.

Table 15 displays the transactions that have potential lock contention with transaction Tx7 on the defined tables.

Table 15. 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, DDCS/6000, 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