

DCE Performance Study

By Bob Russell

This article presents DCE performance characteristics observed in a complex heterogeneous enterprise environment. It addresses some of the frequently asked performance questions, focusing on DCE topology performance, while noting the behavior of the DCE administrator, OS/2 Database Manager, and Network File System (NFS) services contending for the system resources.

Distributed Computing Environment (DCE) application behaviors were studied on a variety of OS/2-, AIX-, UNIX-, and Windows™-based systems. Over 1,500 clients were supported by 12 to 16 DCE and application servers. We used off-the-shelf hardware and software including IBM and non-IBM solutions for the Distributed Computing Environment (DCE) performance study. The IBM Point-of-Sale (POS) DCE application benchmark was used to drive the workload for this evaluation.

Study Requirements

The needs of IBM customers who are implementing applications on DCE set the direction for this performance study. During the first half of 1994, we surveyed and met with many of these customers. Figure 1 shows the results that led to the following requirements for the study:

- ◆ **More than 1,500 clients.** The surveys indicate the highest interest is between 100 and 500 clients. We studied workloads up to 2,000 clients.
- ◆ **Heterogeneous client and server environment.** The following clients and application servers were studied:

–**Clients:** Windows, OS/2, OS/2 WARP, AIX/6000™, Sun® and Hewlett-Packard (HP) workstations

–**DCE and application servers:** Sun, HP, and PS/2® and RISC System/6000 (RS/6000™) systems

- ◆ **Performance comparison between a large single cell and two smaller cells.**
- ◆ **Performance comparison between single and multiple cells connected by a Wide Area Network (WAN).** The survey feedback indicates the greatest interest in system geographies beyond a single contiguous location.

Methodology


The philosophy for these tests was to view the system behavior from a high level. We did not use high-resolution timers to obtain microsecond-level measurements. Generally, we monitored end-user response time and system throughput, and used the following guidelines for interpreting the results:

- ◆ Response times less than one second are considered instantaneous.
- ◆ Response times greater than one second may be perceived by the user as a long time.
- ◆ The *maximum throughput* is the highest arrival rate applied to the system without causing an objectionable degradation in end-user response time.

In many of the test cases, higher throughput could have been achieved. However, if we had pushed the arrival rate any higher, the response times would have begun to grow above our one-second guideline.

Standard Configuration

Figure 2 shows the standard system configuration for these studies. The performance characteristics of the standard system were studied before introducing variations. The standard configuration



Bob Russell

| Platform | Windows | OS/2 | AIX | Midrange | Mainframe |
|---------------------|-----------|---------------|-------------|-----------|-----------|
| Client Workstation | 5 | | 10 | | 6 |
| Application Server | | 10 | 9 | 4 | 3 |
| DCE Server | | 9 | 10 | 3 | 2 |
| Number of Clients | 1-99 | 100-499 | 500-999 | 1000-4999 | 5000+ |
| | 2 | 10 | 7 | 5 | 1 |
| System Geography | Building | City | State | Country | Global |
| | 8 | 10 | 1 | 6 | 2 |
| Type of Application | Database | Data Transfer | Graphical | | |
| | 9 | 10 | 3 | | |
| Data Transfer Size | 1 KB-4 KB | 32 KB-64 KB | 512 KB-1 MB | | |
| | 10 | 6 | 4 | | |

Relative importance (normalized): 1 = least important, 10 = most important

Figure 1. Customer survey results

consists of three separate networks connected by WANs.

Headquarters is the primary network and provides most of the system support. In the standard configuration, all DCE servers and POS application servers are in the headquarters network. The baseline throughput for the standard system at the Headquarters level (shown in Figure 3) is 1,640 Customer Sales Per Minute (CSPM), and the 16 Mbps Token Ring is 40% to 45% utilized. The POS application servers in the standard configuration can support about 1,640 POS clients. An OS/2 Database Manager Online Transaction Processing (OLTP) benchmark, also running at 15 Transactions Per Second (TPS) on a separate server, is not affected by the POS traffic on the network.

The **Region Office** is a smaller network connected to Headquarters by a high-speed WAN. We simulated the high-speed WAN using two IBM Token-Ring Network Bridges linked together by a 4 Mbps Token-Ring network. This is faster than the 1 Mbps links generally used, but the network analyzer indicated that traffic on the 4 Mbps link remained below 0.8 Mbps during these tests. The maximum workload exerted by the clients in the Region Office is 190 CSPM.

The **Branch Office** has a small 16 Mbps Token-Ring network with two PS/2 clients: one running DCE for OS/2 and one running DCE for Windows. The Branch Office is linked to Region Office using IBM LAN Distance™. The link used in the study between the Branch and Region was a null modem set at 38,400 bps. The maximum throughput of the Branch Office was limited to 7.5 CSPM by the line speed of the null modem.

Applications

The following benchmark applications provided the system load:

- ◆ The Point-of-Sale DCE application benchmark is patterned after a retail order-entry system. Since our objective was to look at DCE performance, the POS workload is the predominate application in the study.
- ◆ A suite of rgy_edit, acl_edit, rpccp, and cdscp updates is run from one workstation. This simulates the ongoing activity of a DCE system administrator.
- ◆ An OS/2 Extended Services 1.0 Database Manager OLTP benchmark using the NetBIOS transport for the Database Application Remote

Enterprise Server Topology




Figure 2. Server topology for the enterprise

Interface (DARI) stored procedures is run on a small number of PS/2 clients on the Headquarters' Token-Ring network.

- ◆ IBM LAN Server/Requester 3.0 provided file sharing for the benchmark applications and tools. The LAN Server performance was not measured in the study.

Results

Due to the complexity of the system topology shown in Figure 2, this article discusses the sub-

topologies (A-H) and the performance of each sub-topology individually.

The DCE Servers

In the standard configuration, the DCE servers (Figure 2, Sub-topology A) are all PS/2s:

- ◆ The Primary DCE Cell Directory Service (CDS) server was a PS/2 9585-50 MHz. The CPU utilization remained under 40% during all tests. We found that the type of DCE server has a minimal impact on the overall cell throughput; the total difference between a 33 MHz and 50

POS Throughput in the Standard Configuration




Figure 3. POS throughput

MHz PS/2, RS/6000 (Models 220, 520, and 580), and the HP 9000 Model 827S was less than 5%.

- ◆ The Secondary CDS server, a 50/25 MHz PS/2, also remained less than 40% utilized.
- ◆ The DCE Security and Distributed Local Time servers ran on a 33 MHz PS/2. The CPU of the DCE Security Server remained under 10% utilized except when there was a flurry of DCELOGIN activity. There was no measurable impact when a faster RS/6000 or HP 9000 security server was used.

Security Registry

The following effects of increasing the size of the Security Registry were noted.

- ◆ The disk space required for the Security Registry grows by nearly 1 KB per principal and account, which is about 1 MB per 1,000 principals and accounts.

- ◆ The memory working set for the DCE Security server grows by about 1 KB for each unique principal currently logged into DCE—about 1 MB per 1,000 unique principals.
- ◆ The cell throughput and response time of DCELOGIN are not measurably affected after adding the 1,000 principals and accounts.
- ◆ The only negative effects were that both the startup time of the Security Server Daemon (SECD) when the DCE Security Server is restarted, and the initial startup time of each client's Security Client Daemon (SCLIENTD) are about two times longer.

CDS Namespace

The memory and disk requirements for CDS objects are somewhat higher than for Security Registry entities. The effects of increasing the size of the CDS Namespace follow.

Distributed POS Schematic

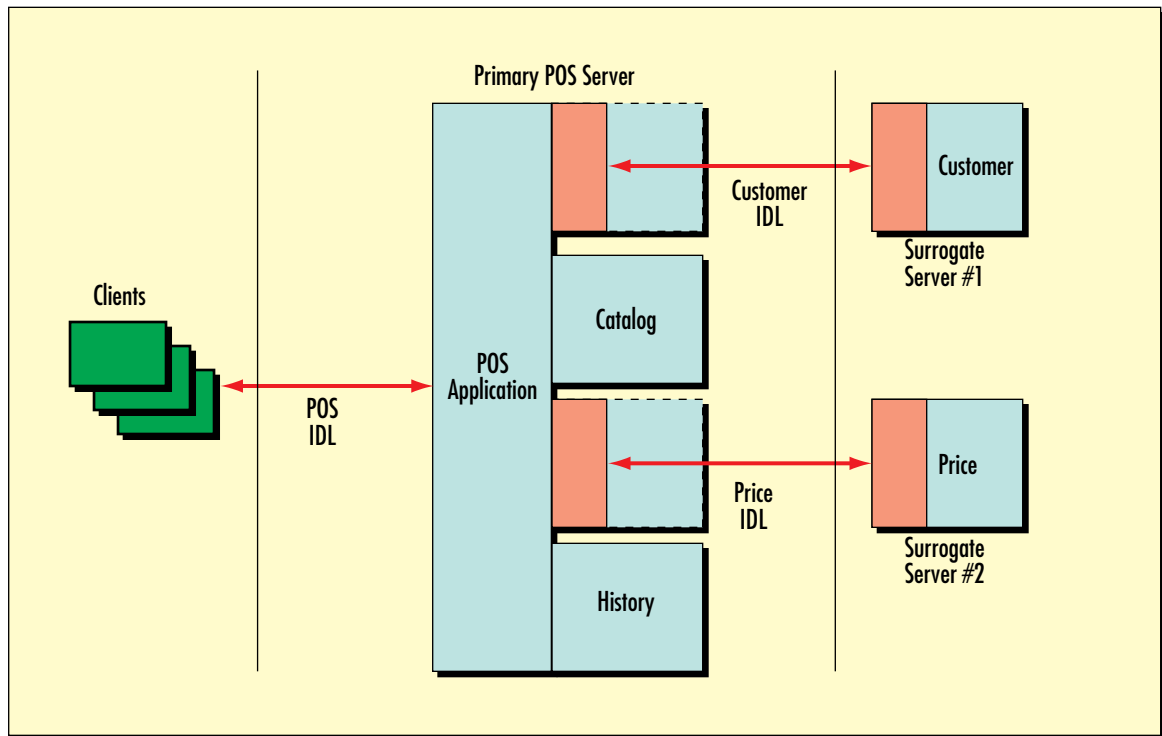


Figure 4. Distributed POS schematic

- ◆ Since CDS directory objects require 14.2 KB of CDS server memory and disk space for each object, the 1,000 directory objects in our tests required 14.2 MB. The base OSF® code allocates a maximum size of 14.2 KB per object, regardless of the actual space needed. Generally, only about 11% of the allocated object space is actually used. There is currently no committed plan by the OSF technology provider to address this memory/disk concern.
- ◆ CDS program objects each required about 1 KB of CDS server memory and disk space, or 160 KB for the 160 program objects in the CDS Namespace of our test cell.
- ◆ There was no measurable effect on cell throughput or client response time.
- ◆ The only negative effect was that the CDS Server Daemon (CSDS) startup time increased up to 10 times.

POS Application Servers

The POS Application Servers are shown in Figure 2, Sub-topology B. The POS application ran in a distributed configuration with five primary and two surrogate POS servers. Figure 4 illustrates the

distribution of the four POS server Remote Procedure Call (RPC) routines. Each client imports a new RPC binding from the CDS server for one of the five primary POS servers before each customer sale. The `rpc_ns_binding_import_next` Application Programming Interface (API) randomly selects one of the five available POS primary servers, providing a random and somewhat even distribution of the workload across the five primary POS servers.

The two surrogate POS servers were shared by the five primary POS servers, allowing all clients to access a single copy of the POS Price and POS Customer databases. The five primary POS servers call the two surrogate servers through nested RPC calls within the POS primary server application.

In the standard configuration, the following hardware supports the POS application:

- ◆ The POS Price surrogate server is on an RS/6000 Model 580. The POS Price RPC routine was called most often: 2.5 times for every customer sale. The CPU and disk utilization of this POS Price server was about 15% when the system throughput was 1,640 CPSM.

- ◆ The POS Customer surrogate server was on an RS/6000 Model 220. The CPU and disk utilization was about 65% at 1,640 CSPM. When a Model 520 was substituted, the CPU was about 40% utilized.
- ◆ All five of the primary POS servers are 33 MHz PS/2 Model 90s or 95s. At 1,640 CSPM, the CPU utilization on all five primary servers was near 100%—definitely the bottleneck in the standard configuration. To achieve higher throughput, the primary POS servers must be upgraded to faster hardware. A single RS/6000 Model 580 primary POS server was fully utilized at about 1,100 CSPM, suggesting five 580s could support about 5,000 POS clients.

It might seem reasonable to substitute a faster machine for one or two of the primary POS servers. However, this will not help since the `rpc_ns_binding_import_next` DCE API call randomly distributes binding handles for multiple instances of the same application. The CPU capacity of the slower machines would continue to gate the total throughput.

A general rule-of-thumb is that the total throughput is equal to n times the throughput of the slowest machine where n is the number of instances of the application server. In effect, the faster machine would respond faster, but would be idle much of the time while most of the requests wait in the queues of the slower machines. Therefore, to improve the total throughput, all instances of an application server must be upgraded to faster machines. This characteristic of DCE currently has no workaround.

Normally, the POS clients import a new binding from CDS for every customer sale. If the clients remain bound to the application server for longer periods, throughput will be much higher. In the standard configuration, the throughput is 1,640 CSPM with rebinding for every customer sale. If rebinding is not done, the throughput

increases to 2,080 CSPM. If there is no specific need to rebind frequently, the POS application realizes a 22% improvement in total throughput.

Elimination of frequent rebinding drops utilization of the CDS server CPU from about 40% to near 0%. The CPU utilization of the application server also drops as a result of eliminating handshakes between the client and application server to negotiate the RPC authentication and reestablish the binding. The client's aggregate response time for a complete customer sale is reduced by about 0.5 seconds when rebinding is eliminated. Higher throughput can be achieved while maintaining acceptable end-user response time.

The LAN and OS/2 Database Manager Servers

The LAN and OS/2 Database Manager servers, shown in Figure 2 (Sub-topology C), include the following:

- ◆ **Two IBM LAN Server 3.0 domain controllers:** An Advanced LAN Server and an Entry LAN Server support benchmark program sharing, performance data collection, and various network install services. All OS/2 systems are also LAN Requester 3.0 clients.
- ◆ **OS/2 Extended Services 1.0 Database Manager Server:** This server supports the database OLTP benchmark clients and runs continuously to provide network loading.
- ◆ **NFS for OS/2 Servers:** The two LAN servers, also NFS servers, are used to compare the NFS I/O performance among OS/2, AIX/6000, SunOS/Solaris®, and HP/UX servers.

The presence of the LAN and database load have no measurable effect on our system performance other than increased traffic on the Headquarters Token-Ring LAN.

The Token-Ring/Ethernet Gateway

The Headquarters' network includes an Ethernet LAN segment to support the HP and Sun

The memory and disk requirements for CDS objects are somewhat higher than for Security Registry Entities.

machines (Figure 2, Sub-topology D). The 16 Mbps Token Ring and 10 Mbps Ethernet are linked together by an RS/6000 Model 220 acting as a gateway router.

The Ethernet begins as a thick-wire connection to the RS/6000 gateway, then to thin-wire, and finally to 10BaseT terminating at the Sun SPARC-server 1000. The HP 9000 827S server and 730 client are attached to the thin wire; the Sun SPARCstation® LX is attached to the thick wire.

Two simple tests were run to evaluate the efficiency of this particular Token-Ring-to-Ethernet LAN gateway. We did not intend to compare the performance of specific hardware and operating system platforms, so only general observations about this LAN segment will be discussed.

- ◆ The first Token-Ring-to-Ethernet gateway test measured a 4,096 byte ping from a single OS/2 client in each LAN segment. (The LAN segments are Headquarters' Token-Ring segment, Headquarters' Ethernet segment, Region Token-Ring segment, and the Branch Token-Ring segment.) The OS/2 TCP/IP ping function was used for measurement since it reports its response time with a timer resolution of 31 milliseconds. Each of the selected OS/2 workstations ping each of the 15 PS/2, RS/6000, HP, and Sun servers.

Within each subnetwork (Headquarters, Region, and Branch), the 4 KB ping responded in 31-34 milliseconds (one timer tick). A ping between the Region and Branch (null modem) took about 2,580 milliseconds, between the Region and Headquarters (4 Mbps link) about 62-64 milliseconds, and crossing the Token-Ring/Ethernet gateway added an additional 10-30 milliseconds.

- ◆ The second gateway test was to indirectly access the four POS data repositories using the NFS mount command. For this test, a single instance of the POS primary server was running.

Although applications usually employ a Relational Database Management System (RDBMS) to manage their data, the POS data was stored in binary files and accessed using portable ANSI®/POSIX C language calls. This was a result of our objective of portability and comparability across hardware and software platforms. Using a specific RDBMS would have precluded this.

The results of these NFS-redirectioned I/O tests are normalized. The results of independent exter-

nal performance tests are used to establish the basis for normalization of these servers. Accessing the POS data via NFS mount on an RS/6000 on the Headquarters' 16 Mbps Token Ring was assigned an efficiency factor of 1.0 (100%). When accessing data on the HP and Sun servers on the Ethernet, two levels of degradation are encountered:

- ◆ When the NFS mount was from a Token-Ring-attached POS Primary server through the gateway to the Ethernet-attached servers, the efficiency was 0.75—a 25% loss in maximum throughput.
- ◆ When the POS Primary server is moved to the gateway server and the NFS mount is performed on this gateway server, communication is directly to the Ethernet adapter. The efficiency improved to 0.88—only a 12% loss in maximum throughput.

The following additional observations were made during the gateway tests:

- ◆ Although the line speed of Ethernet is 10 Mbps, the Sniffer Network Analyzer indicated that the Ethernet was less than 20% utilized. The CPU and disk of the HP and Sun NFS servers were less than 30% utilized. This suggests that the bottleneck is the communication adapter, not the capacity of the HP or Sun NFS servers.
- ◆ When POS Primary servers are switched among the RS/6000 Models 580, 520, and 220, the efficiency ratio is consistent. This suggests that server congestion and arrival rate are not factors in this test.
- ◆ When the results of the independent performance tests on the RS/6000, Sun, and HP servers are compared to our NFS throughput results, the inconsistencies further suggest a network hardware bottleneck.

Remote Offices and WANs

Sixty percent of the customers surveyed indicate that their systems extend beyond one contiguous location. It is therefore interesting to look at options for linking remote locations.

Region Office (Figure 2, Sub-topology F): Two servers in the Region Office that are used only for the “one versus two cell” testing are discussed later in this article. These servers are not used in the standard configuration.

Sixty percent of the customers surveyed indicate that their systems extend beyond one contiguous location.

Response Time Cost for WANs




Figure 5. WAN response time cost

The DCE clients in the Region Office are OS/2, OS/2 WARP, and Windows. In the standard configuration, these clients can exert the workload of 190 clients, or 190 CSPM. The 190 CSPM includes the 7.5 CSPM workload of the Branch Office.

Branch Office (Figure 2, Sub-topology H):

The Branch Office represents a remote location or work-at-home user connected to the Region Office using IBM LAN Distance and a modem. For our tests, we installed two PS/2 clients on the Branch Token Ring: one OS/2 client and one Windows client. The maximum throughput of the Branch Office clients is gated by the null modem line speed.

High-Speed Link (Figure 2, Sub-topology E):

A typical high-speed link is 1 Mbps. We did not have access to a link of this speed, so we implemented a 4 Mbps link. Two PS/2s with two Token-Ring adapters were installed in the local office (16 Mbps) rings of both the Headquarters and the Region. These two bridges were connected by an isolated 4 Mbps Token Ring attached to

their second adapter. The IBM Token-Ring Network Bridge Program ran on the bridges.

When running the entire system at 1,640 CSPM, the network analyzer shows that less than 0.8 Mbps flows across the 4 Mbps link. The throughput exerted by the Region Office is 190 CSPM. Therefore, if we assume a 1 Mbps link, the 0.8 Mbps traffic is nearly full utilization of the link. Thus, 190 clients would be about the maximum that a single 1 Mbps link could support.

Low-Speed Link: The Branch Office is connected to the Region Office using an IBM LAN Distance low-speed link (Figure 2, Sub-topology G). The physical link is a null modem set at 38,400 bps.

We did encounter a problem configuring the Windows client in the Branch Office. The increased response time across the null modem caused an unidentified time-out situation, which terminated the Windows 3.1 desktop. When the Novell® LAN Workplace 4.1 was replaced with the IBM TCP transport for DOS, the problem was

eliminated. The IBM TCP transport for DOS is slower than the Novell transport or the other two non-IBM transports that we tested, and it would not normally have been our first choice.

Several tests were performed to evaluate the performance impact of traversing the links between Headquarters and the Region Office, and between the Region and Branch Offices. Figure 5 shows the portion of response times attributable to each link for the following tests. A normalized value of 1.0 is assigned to each test when performed locally on the Headquarters' Token-Ring network. These measurements are from a single OS/2 client in an otherwise idle system.

- ◆ **TCP/IP 4 KB ping:** Probably the closest application to the TCP/IP transport level is ping. Since it is also the most subject to the line speed of each LAN segment, it should be expected to have the greatest degradation across the WANs.

The Region-to-Headquarters ping is two times slower than the local ping. This is good considering the 4:1 difference between the 16 and 4 Mbps line speeds.

Between the Branch and Region (38,400 bps), the ratio is 75:1, whereas the line-speed ratio is 426:1.

The ratio is 77:1 between the Branch and Headquarters (across both WANs).

- ◆ **Authenticated RPC 4 KB Data Transfer:** This test used an application program to transfer 4 KB of data from a client to a server in an authenticated RPC call. This test should perform better than the ping test, but since the application is still data transfer-intensive, the degradation remains high because of the effect of the line speeds.

Using DCE Security call-level authentication, the ratio was 36:1 across the 38,400 bps link and 1.4:1 across the 4 Mbps link.

Using DCE Security packet-privacy (encryption) authentication level, the ratio was 8:1 across the 38,400 bps link and 1.1:1 across the 4 Mbps link. The ratio improved because the larger portion of the total response time was expended by the server and client encrypting and decoding the data packets.

- ◆ **POS Customer Sale:** In real life, the time spent within the client and server application is much longer than the time spent actually moving data across the network. Therefore, we should expect the data transfer ratio to improve to a more acceptable level since the RPC portion of the POS application is less than 5% of the aggregate response time when run

Major AIX Events in Europe

AIX Systems Europe announces two major events in Europe to share IBM's strategy for AIX and the RISC System/6000 with Customers, Business Partners, and other professionals interested in working with AIX and the RISC System/6000.

Each event will consist of main tent and technical electives. The main tent sessions will be presentations by key IBM executives discussing IBM's strategies. The technical electives will be presented by IBM's technical experts and will cover a broad range of topics.

AIX in Focus

Date: February 6-10, 1995
Audience: Technical AIX Professionals, Software Developers, Business Partners
Topics: IBM's Technology Principles
Middleware (Databases, Networking, and so on)
"From Laptop to Supercomputer"
Object Technologies
Competition

The Customer Colloquium

Date: February 15-17, 1995
Audience: Customers
Topics: Re-engineering Your Business
The IT Imperatives
Solutions – Not Just Technology

For more information on either event, contact either of the following:

John Childs
(UK): 44-256 344875

Richard Sandback
(UK): 44-256 344661

locally in Headquarters (the average data size was 3.2 KB, with a range of 1 to 16,384 bytes).

The POS ratio between the Branch and Region on the 38,400 bps link was 9.2:1, and it was 1.5:1 between the Region and Headquarters.

- ◆ **DCE Calls:** Although many DCE functions and API calls make some RPC calls under the covers, most response time is spent in the client DCE programs.
 - **DCELOGIN:** The ratio was 1.8:1 across the 38,400 bps link and 1.4:1 across the 4 Mbps link.
 - **rpc_ns_binding_import_next:** The ratio was 2.3:1 across the 38,400 bps link and 1.5:1 across the 4 Mbps link.
 - **DCESTART:** The ratio was 1.4:1 across the 38,400 bps link and 1.3:1 across the 4 Mbps link.

End-user response time across the WAN is closely tied to the frequency and data size of data transfer operations. The POS application makes an average of 5.5 RPC calls and 7 DCE API calls for each customer sale. The overall response time ratios for a complete POS customer sale (null modem 9.2:1 and 4 Mbps link 1.5:1) are probably more representative of many interactive end-user applications across these specific WANs.

The two clients in the Branch Office can saturate the 38,400 bps link at 7.5 CSPM. So under the best conditions, this LAN Distance link can handle up to 7.5 clients. In reality, LAN Distance is intended for a single client, and extrapolating from our results, a minimum 5,120 bps link would be sufficient for one client.

When we initially installed the LAN Distance link, some LAN Server 3.0 Client time-out problems were encountered due to the longer response times across the null modem. The following changes were recommended by our LAN Distance performance analyst:

- ◆ In `IBMLAN.INI`
 - `wrkheuristics` bits 0, 6, 8, 9, and 11 were changed to 0
 - `sizeworkbuf` was changed from 512 to 4096
- ◆ In `PROTOCOL.INI` (`NETBEUI.NIF` section)
 - `dlcretries` was changed to 20
 - `TI` was changed from 30000 to 60000
 - `T1` was changed from 500 to 10000
 - `T2` was changed from 200 to 2000

One Versus Two DCE Cells

We tried adding a second DCE cell server and a POS application server in the Region Office to handle both the Region and Branch clients' DCE and POS activity, while the Headquarters' cell handled only the clients on the Headquarters' network. LAN Server activity was the only traffic across the 4 Mbps link between Headquarters and the Region Office. There was no DCE Global Directory Server (GDS) for this test. Results were as follows:

- ◆ The throughput in Headquarters drops from 1,640 CSPM to 1,570 CSPM because of the reduced client horsepower. When we started more clients to replace those in the other cell, the throughput returned to 1,640 CSPM; therefore, the capacity of the Headquarters' DCE cell was not changed by splitting the Region and Branch into a new cell.
- ◆ The throughput in the Region Office remained at 190 CSPM with a PS/2 DCE server and an RS/6000 Model 520 POS server in the Region DCE cell. Since the RS/6000 Model 520 can support about 360 POS clients in this configuration, the 190 CSPM is clearly a limitation of the client's ability to exert workload rather than a limitation of either the WAN or POS servers when configured in a single cell.

Implementing two DCE cells to eliminate traffic across the 4 Mbps WAN improved the aggregate throughput. However, a similar improvement could be achieved by adding the additional POS application server in the original single-cell configuration. The only performance gain from this two-cell experiment was a small improvement in the response time in the Region and Branch by eliminating the hop across the 4 Mbps WAN.

This two-cell configuration could result in a performance improvement if there were a resource bottleneck in the original single DCE cell. In our test configuration, we could expect a resource bottleneck if the arrival rate was increased to 4,000 or 5,000 CSPM. Both the 16 Mbps Token Ring and the PS/2 CDS server could become saturated if we had sufficient client hardware to drive the POS workload.

Summary

The following summarizes the findings from the DCE performance study:

- ◆ IBM DCE 1.2 for OS/2 and AIX provide good performance and interoperability in our

testing. IBM DCE for AIX 1.3 (OSF DCE 1.0.3) is now available. When we ran the new version on our test system, we found the new version to be compatible with OS/2 DCE 1.2 clients, IBM DCE for Windows 1.0 clients, and our current HP DCE/9000 1.2 installation. There were no significant differences in the system performance.

- ◆ Performance problems have been encountered while integrating the HP 9000-827S system into our test environment stemming from our specific Token-Ring/Ethernet gateway and network hardware configuration.
- ◆ Using WANs provided satisfactory performance, given the line speeds. No DCE interoperability problems were encountered using IBM LAN Distance or IBM Token-Ring Bridge software.
- ◆ We encountered no interoperability or performance concerns while integrating OS/2, Windows, and AIX DCE clients.
- ◆ We encountered no network contention problems with DCE (TCP/IP), LAN Server (NetBIOS), and OS/2 Database Manager running concurrently.
- ◆ The overall performance of our test system is well-behaved, and we did not encounter any bottlenecks that would preclude expanding to much higher workloads.
 - The OS/2-based DCE configuration could be expanded to 4,000 or 5,000 clients in a single DCE cell. The limiting factor would be the single 16 Mbps Token Ring.

– Using an AIX-based DCE configuration, the estimated capacity of our single DCE cell would be even higher. The capacity of the single 16 Mbps network would probably be the limiting factor.

References

Corn, Cindy; Li, Tim; Pekowski, Ray; and Sante-ford, Bob. "Performance of Key Functions in DCE for OS/2." *IBM Personal Systems Technical Solutions*. January/February 1994.

Perry, Benetta and Russell, Bob. "IBM DCE for OS/2 Multi-user Application Performance." *IBM Personal Systems Technical Solutions*. January/February 1994.

Perry, Benetta and Russell, Bob. "Distributed Performance Characteristics of IBM DCE for OS/2." *IBM Personal Systems Technical Solutions*. July/August 1994.

Pekowski, Ray. "Performance of Key Services in IBM DCE Client for Windows." IBM White Paper. Available from MKTT00LS in the DCE Performance package from your IBM representative.



Bob Russell, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Russell is team leader for the LAN Systems Performance Workgroup Environment team. His most recent positions include performance analysis for OS/2 Database Manager and DCE for OS/2.



RISC System/6000 Developer's Discount Program

As a developer, you need the latest equipment for development and testing. Your ability to create quality applications and solutions depends on having the latest technology available. And at prices you can afford.

IBM offers special purchase programs for qualified hardware and software developers. These programs include the following:

- ◆ Hardware
- ◆ Software
- ◆ Peripherals

The program includes systems ranging from the P40 through the Symmetric Multiprocessors (SMP). With discounts as great as 50% and up to 10 machines in one year, you can't afford not to evaluate the program.

Programs may vary outside the United States. Contact your local representative for more information.

For information, call **1-800-222-2363** and ask for extension 63. ■

IBM reserves the right to change terms and conditions of this program without notice.