

**S/390 Division**



---

**S/390 Multiprise 3000 Integrated LAN Adapter Feature  
Performance Report**

November 1999

---

# CONTENTS

|   |   |
|---|---|
| <b>INTRODUCTION</b> .....                   | 2 |
| <b>INTEGRATED LAN ADAPTER FEATURE</b> ..... | 2 |
| Function .....                              | 2 |
| Connectivity .....                          | 2 |
| Theoretical Raw Bandwidth .....             | 2 |
| I/O Subsystem Processor .....               | 2 |
| PCI Bus .....                               | 3 |
| <b>PERFORMANCE SUMMARY</b> .....            | 3 |
| <b>WORKLOAD DRIVER</b> .....                | 3 |
| Primitive Traffic .....                     | 3 |
| Streams .....                               | 4 |
| Request/Response .....                      | 4 |
| Connection Request/Response .....           | 4 |
| Sessions .....                              | 4 |
| <b>MEASUREMENT RESULTS</b> .....            | 4 |
| Environment .....                           | 4 |
| Measurement Coverage .....                  | 5 |
| Tables .....                                | 6 |
| Fast Ethernet TCP/IP Results .....          | 6 |
| Bulk Data Streams .....                     | 6 |
| Token Ring TCP/IP Results .....             | 7 |
| Bulk Data Streams .....                     | 7 |
| Request/Response .....                      | 7 |
| Connection Request/Response .....           | 7 |
| <b>I/O PROCESSOR CONTENTION</b> .....       | 7 |
| <b>ABBREVIATIONS</b> .....                  | 8 |

---

## INTRODUCTION

Multiprise<sup>®</sup> 3000 is an extension to the IBM<sup>®</sup> original Multiprise family, which itself offered innovative technology and new levels of performance for smaller-scale server environments. Each of the new Multiprise server models provides the capacity necessary to handle existing mission-critical applications while providing room for growth; each server also provides additional capacity for handling new e-business applications within the same environment, greatly simplifying integration and management.

Multiprise 3000 offers standard ESCON<sup>®</sup> and parallel channels as well as an Integrated Local Area Network (LAN) Adapter feature and a Wide Area Network (WAN) Adapter feature. Network connections can be made through standard ESCON or parallel attached controllers, or through the LAN or WAN subsystems.

The Open Systems Adapter found in Multiprise 2000 and other S/390<sup>®</sup> models is not supported on Multiprise 3000.

This paper provides performance data on the Integrated LAN Adapter subsystem to assist you in choosing the right connections for your e-business and other networking applications.

---

## INTEGRATED LAN ADAPTER FEATURE

The Integrated LAN Adapter feature offers greater configuration options for the S/390 environments with the ability to connect to Fast Ethernet and Token-Ring. The Integrated LAN Adapter feature is plugged directly into a standard PCI slot, becoming an integral component of the system and enabling convenient LAN attachment.

The Integrated LAN Adapter feature has four PCI slots for LAN adapters. Each Integrated LAN Adapter feature adapter provides a single LAN port.

## Function

The Integrated LAN adapters provide several functions:

1. HMC connectivity
2. 3172 emulation for VTAM<sup>®</sup>
3. 3172 emulation for TCP/IP
4. Non-SNA, DFT 3270. The 3174 emulation function takes the place of an external 3174; it translates Non-SNA, DFT protocol into TN3270E protocol, providing 3270 sessions either on the local console, or over the LAN. This function is especially useful for OS/390<sup>®</sup> consoles.

## Connectivity

The Multiprise 3000 Integrated LAN Adapter feature supports four LAN adapter slots. If you have an HMC console, you might want to dedicate one LAN for that function. If you have multiple TCP/IP stacks, you will need an adapter for each stack. If you connect to SNA and TCP/IP you might want an adapter for each protocol. If you are running SNA and TCP/IP over Fast Ethernet you need an adapter for each if one protocol is to run over an EtherNet IEEE 802.3 LAN and the other is to run over an EtherNet DIX LAN. If you have more than 256 SNA PU connections you will need an adapter for each set of 256 PUs.

## Theoretical Raw Bandwidth

The Integrated LAN Adapter feature supports full-duplex (bidirectional) communication with devices on dedicated segments such as IBM's 8271 EtherStreamer<sup>®</sup> Switch and 8272 LANstreamer<sup>®</sup> switch. Full-duplex support for the Fast Ethernet and Token Ring environments permits simultaneous transmit and receive on the LAN.

The raw full-duplex bandwidth of Fast Ethernet is effectively 200 Mbps: 100 Mbps transmit and 100 Mbps receive. That is 25 MBps total. The raw full-duplex bandwidth of Token-Ring is effectively 32 Mbps: 16 Mbps transmit and 16 Mbps receive. That is 4 MBps total.

In half-duplex mode the adapter may not transmit and receive simultaneously, hence, the raw bandwidth is half that of full-duplex mode.

The amount of total raw bandwidth that can be put to effective use depends on a number of factors outside and inside the Integrated LAN Adapter subsystem.

## **I/O Subsystem Processor**

The I/O Subsystem Processor resides on the PCI bus and governs all LAN traffic through the subsystem.

In many cases the I/O Subsystem Processor will be the bottleneck.

The I/O Subsystem Processor also governs the WAN traffic for the RS-232, X.21, V.35 and MultiPort Model 2 Wide Area Connection (WAC) adapters attached to the ISA-bus based Integrated WAN Adapter feature. Thus, Integrated WAN feature traffic processing may degrade LAN performance.

## **PCI Bus**

Traffic from Internal Disk does not use the I/O Subsystem Processor but does use the PCI bus.

Internal Disk traffic on the PCI bus may also impact LAN performance.

---

## **PERFORMANCE SUMMARY**

A single Integrated Fast Ethernet adapter is capable of delivering up to 6.3 MBps of TCP/IP bulk data streams throughput from S/390 to the network.

The I/O Subsystem Processor is the limiting factor, therefore adding adapters will not increase performance.

Applications requiring more than 6.3 MBps should use ESCON or parallel channel connections in addition to Integrated Fast Ethernet.

A single Integrated Fast Ethernet adapter is limited to 2.5 MBps when transferring TCP/IP bulk data streams from the network to S/390.

As in the S/390-to-Network case, adding adapters will not increase throughput.

A single Integrated Token- Ring adapter is limited to 1 MBps in the S/390-to-Network direction and 1.7 MBps in the Network-to-S/390 direction.

Performance data is only available for the measurements in this report.

If your applications require more performance than the Integrated I/O LAN Adapter feature delivers, consider using an ESCON or parallel channel attached controller,

such as the 2216 by IBM or any controller from Bus-Tech<sup>®</sup> or Cisco<sup>®</sup>.

---

## **WORKLOAD DRIVER**

The performance measurement workloads used in this report can be categorized as bulk data streams or as interactive transactions. Bulk data streams traffic represents the continuous flow of large amounts of data between source and destination. Interactive transaction traffic represents the transfer of small request-response message pairs between client and server.

The workload tool used in this study drives traffic at the application layer of the TCP/IP protocol stack. This means that not only is the Integrated LAN Adapter subsystem hardware measured, but the entire networking infrastructure is measured as well.

All data movement occurs between client-side memory and server-side memory, therefore all outside bottlenecks, such as DASD or tape, are eliminated.

The workload is driven as hard as possible so as to make the Integrated LAN Adapter subsystem the bottleneck and to achieve maximum performance. Measuring performance in this way, with a number of workloads and protocols at the application layer, without outside constraints, demonstrates the performance available to a wide variety of applications.

Whether or not an application can make full use of the Integrated LAN Adapter subsystem depends on a number of factors. One such factor might be a bottleneck at a peripheral device. For example, in an FTP transfer, the source or destination DASD might limit throughput to only 3 MBps even though the networking infrastructure offers twice that throughput.

## **Primitive Traffic**

The workload driver generates three primitive types of traffic: Streams, Request/Response (RR), and Connection Request/Response (CRR).

Streams is a bulk data streams workload. RR and CRR are interactive transaction workloads.

The primitive-driver moves data between application memory and the protocol stack. The driver code is very efficient and is not intended to represent any particular customer application code.

The primitive-driver allows control of workload parameters such as stream size and request/response message size and also TCP socket parameters such as the application program interface (API) data buffer size.

Values for the most important network parameters are given with the measurement data.

### **Streams**

This type of traffic models bulk transfer workloads such as FTP and ADSM™ in the TCP/IP environment. In this model each primitive application client-server pair continuously pushes large streams of data through the network.

Multiple instances of the primitive-driver are run to provide multiple concurrent sessions.

The key measure is throughput, in bytes per second.

### **Request/Response**

This type of traffic models TCP/IP fixed-connection transactional applications such as TN3270 and Enterprise Resource Planning (ERP) applications such as SAP. In this model, the client creates a TCP connection with the server and then sends a request-message to the server. In response the server sends a response-message back to the client. Once the response-message is received the client sends another request-message to the server, and so on.

The request-response message pair constitutes a transaction.

Note that each TCP connection is created once and is maintained throughout the run.

The size of the request and the response is fixed at the start of the run.

Multiple instances of the primitive-driver are run to provide multiple concurrent transaction sessions.

The key measures are throughput, in transactions per second, and response time, in seconds per transaction.

### **Connection Request/Response**

This type of traffic models connection-intensive workloads such as those generated by Web based

applications and popular Web-workload drivers. In this model, the client creates a TCP connection with the server, sends in a request-message to the server and in response receives a response response-message from the server. The client then breaks the connection and starts the process all over again.

The connection setup and takedown and request-response message pair constitute a transaction.

Note that TCP sets up and tears down the TCP connection for each transaction.

The size of the request and the response is fixed at the start of the run.

Multiple instances of the primitive-driver are run to provide multiple concurrent transaction sessions.

As with request/response, the key measures are throughput, in transactions per second, and response time, in seconds per transaction.

### **Sessions**

In any particular session, each stream of data, or each transaction, must complete before the next can begin.

When running multiple streams together, each stream operates independently, except of course for the contention that often occurs at a shared resource.

---

## **MEASUREMENT RESULTS**

### **Environment**

All measurements were taken on a model 7060-H70 (2-way) server running eNetwork Communication Server OS/390 V2R7.

The Integrated LAN Adapter feature was running GA-level driver code during the Token Ring RR and CRR measurements. During all other measurements, the machine was running post-GA driver level D24Q.

All resources in the machine were dedicated to the measurements.

In the TCP/IP environment, the clients were IBM RS/6000 F50 (604e 332 MHz) workstations running AIX 4.3.2.

In the TCP/IP environment the Multiprise 3000 and workstation Token Ring LAN adapters were connected to an IBM 8228 Multistation Access Unit (MAU).

In the TCP/IP environment the Multiprise 3000 and workstation Fast Ethernet LAN adapters were connected to an Alteon Networks ACEswitch™ 180, 10/100/1000 Mbps Ethernet Server Switch. Each S/390-workstation pair was a separate LAN in the switch.

The TCP/IP send and the receive buffer size was set at 262,144 bytes in the TCP/IP profile on S/390 and 65,535 bytes in the network options (the 'no' command) on the workstation.

TCP/IP message-wait-all and TCP/IP no-delay were set as follows:

|                           | <b>S/390 To Network Streams Traffic</b> | <b>Network To S/390 Streams Traffic</b> | <b>Interactive Traffic</b> |
|---------------------------|---|---|----------------------------|
| <b>S/390 Server</b>       |   |   |                            |
| Message-wait-all          | No                                      | Yes                                     | No                         |
| TCP/IP no-delay           | No                                      | Yes                                     | Yes                        |
| <b>Client Workstation</b> |   |   |                            |
| Message-wait-all          | Yes                                     | No                                      | No                         |
| TCP/IP no-delay           | Yes                                     | No                                      | Yes                        |

When Message-wait-all is set to Yes, TCP fills the application programming interface receive buffer before handing off the data buffer to the application. When TCP/IP no-delay is set to Yes small packets such as acknowledgments are sent to the network without delay.

These parameters are often set by the application.

In all cases TCP/IP RFC1323 was not used. Without RFC1323, the TCP/IP window size is limited to 64 kilobytes.

## Measurement Coverage

This report contains TCP streams performance measurement data for single and multiple S/390-to-Network and Network-to-S/390 streams on a single Fast Ethernet adapter and on two Fast Ethernet adapters. A single Fast Ethernet adapter was also measured with a mix of S/390-to-Network and Network-to-S/390 traffic.

This report also contains TCP streams performance measurement data for single and multiple S/390-to-Network and Network-to-S/390 streams on a

single Token-Ring adapter. A single Token Ring adapter was also measured with 200-byte/1000-byte RRs and 200-byte/15,000-byte CRRs. (The number before the slash is the request-message size and the number after the slash is the response-message size.)

This report does not show the effects of Integrated WAN Adapter feature traffic or Internal Disk traffic on the LAN subsystem.

## Tables

The column headings in the tables following are defined as such:

**MTU Bytes:** Maximum Transmission Unit (packet size), in bytes

**Num LAN:** Number of Integrated LAN Adapters

**Number Sessions:** Number of primitive-driver sessions

**Total MBps:** Total Megabytes per second over all LAN adapters and all sessions.

**ISP %:** Estimate of I/O Support Processor utilization, in percent.

**CP %:** SAD estimate of S/390 CPU utilization, in percent. Total across both CPUs.

**SAP %:** SAD estimate of S/390 System Assist Processor utilization, in percent. The SAP is often referred to as the S/390 I/O support processor, however, the SAP is separate from the I/O Support Processor on the PCI bus.

The accuracy of most results was high in terms of the percent of the target bandwidth achieved (95% or better) and in terms of the precision (very low variability in throughput during the measurement). Those measurements that had medium or low precision are marked.

The following markers are used in the tables:

\* TCP retransmissions; M: Medium precision; L: Low Precision; n/a: not available

## Fast Ethernet TCP/IP Results

All Fast Ethernet measurements were taken in full-duplex mode.

### Bulk Data Streams

| MTU Bytes               | Num LAN | Number Sessions | Total MBps             | ISP % | CP % | SAP % |
|-------------------------|---------|-----------------|------------------------|-------|------|-------|
| <b>S/390-to-Network</b> |         |                 |                        |       |      |       |
| 1500                    | 1       | 1               | 6.3                    | 100   | 14   | 9     |
| 1500                    | 1       | 2               | 6.2                    | 100   | 15   | 9     |
| 1500                    | 1       | 3               | 6.1                    | 100   | 15   | 9     |
| 1500                    | 2       | 1 per LAN       | 6.0<br><i>2.7+3.3</i>  | 100   | --   | --    |
| 1500                    | 2       | 2 per LAN       | 5.8M<br><i>2.5+3.2</i> | 100   | --   | --    |
| 576                     | 1       | 1               | 4.6                    | 100   | --   | --    |
| 576                     | 1       | 2               | 4.6                    | 100   | --   | --    |
| 576                     | 1       | 3               | 4.4L                   | n/a   | --   | --    |
| 576                     | 2       | 1 per LAN       | 4.4M<br><i>1.7+2.7</i> | 100   | 12   | 6     |
| 576                     | 2       | 2 per LAN       | 4.2L<br><i>1.4+2.8</i> | 100   | 13   | 6     |
| <b>Network-to-S/390</b> |         |                 |                        |       |      |       |

| MTU Bytes  | Num LAN | Number Sessions | Total MBps                               | ISP % | CP % | SAP % |
|--|---------|-----------------|--|-------|------|-------|
| 1500   | 1       | 1               | 2.5*                                     | 100   | 12   | 3     |
| 1500   | 1       | 2               | 2.3*                                     | 100   | 11   | 4     |
| 1500   | 1       | 3               | 2.3*M                                    | 100   | 13   | 4     |
| 1500   | 2       | 1 per LAN       | 2.8<br><i>1.3+1.6</i>                    | 90    | 13   | 4     |
| 1500   | 2       | 2 per LAN       | 2.7*M<br><i>1.2+1.5</i>                  | 95    | --   | --    |
| 576  | 1       | 1               | 1.7*                                     | 100   | --   | --    |
| 576  | 1       | 2               | 2.5*                                     | 100   | 17   | 4     |
| 576  | 1       | 3               | 2.5*M                                    | 100   | 17   | 5     |
| 576  | 2       | 1 per LAN       | 2.8*M<br><i>1.2+1.5</i>                  | 95    | 17   | 5     |
| 576  | 2       | 2 per LAN       | 2.6*M<br><i>1.2+1.4</i>                  | 95    | --   | --    |
| <b>Mixed S/390-to-Network and Network-to-S/390</b> |         |                 |  |       |      |       |
| 1500   | 1       | 2               | 6.0*L<br><i>In: 0.9*L<br/>Out: 5.1</i>   | 100   | 15   | 9     |
| 1500   | 1       | 4               | 5.6*L<br><i>In: 0.7*L<br/>Out: 4.91M</i> | 100   | 15   | 8     |
| 576  | 1       | 2               | 4.3*L<br><i>In: 0.8*L<br/>Out: 3.5L</i>  | 100   | 14   | 6     |
| 576  | 1       | 4               | 3.7*L<br><i>In: 1.2*L<br/>Out: 2.5L</i>  | 100   | 15   | 7     |

Throughput through a single Fast Ethernet adapter at 1500-byte packets is 6.3 MBps in the S/390-to-Network direction and 2.5 MBps in the Network-to-S/390 direction. At 576-byte packets S/390-to-Network throughput is 4.6 MBps and Network-to-S/390 throughput is 2.8 MBps.

In general, adding adapters does not significantly increase throughput through the Integrated LAN Adapter subsystem, because in many cases the I/O Subsystem Processor is the bottleneck.

For the same reasons, adding streams, in general, does not significantly increase throughput. In fact, in many cases we see a slight decrease in throughput with multiple streams sessions on a single adapter and with multiple adapters.

In the S/390-to-Network cases, the I/O Subsystem Processor is fully utilized and is therefore thought to be the bottleneck. In most Network-to-S/390 cases the TCP retransmissions noted in the tables are a side-effect of the bottleneck.

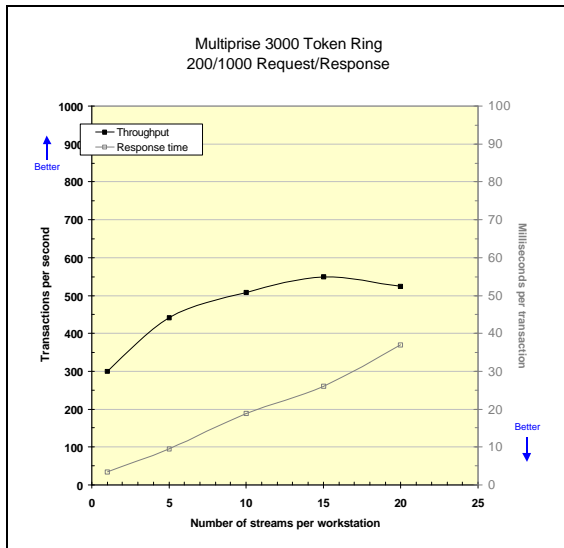
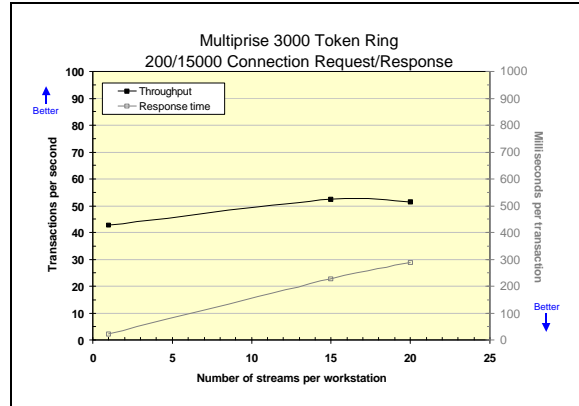
## Token-Ring TCP/IP Results

All Token-Ring measurements were taken in half-duplex mode.

**Bulk Data Streams**

| MTU Bytes               | Num LAN | Number Sessions | Total MBps | ISP % | CP % | SAP % |
|-------------------------|---------|-----------------|------------|-------|------|-------|
| <b>S/390-to-Network</b> |         |                 |            |       |      |       |
| 4000                    | 1       | 1               | 0.9        | 35    | 4    | 3     |
| 4000                    | 1       | 2               | 0.9L       | 35    | 6    | 3     |
| 4000                    | 1       | 2               | --         | --    | --   | --    |
| <b>Network-to-S/390</b> |         |                 |            |       |      |       |
| 4000                    | 1       | 1               | 1.7        | --    | --   | --    |
| 4000                    | 1       | 2               | --         | --    | --   | --    |
| 1500                    | 1       | 1               | 0.8M       | --    | --   | --    |
| 1500                    | 1       | 2               | 1.0L       | --    | --   | --    |
| 1500                    | 1       | 3               | 1.0L       | --    | --   | --    |
| 1500                    | 1       | 4               | 1.0L       | --    | --   | --    |

One workstation client was used.



**Request/Response**

The RR workload was set to issue a 200-byte request-message from the client and return a 1000-byte response-message from the server. One workstation client was used.

**Connection Request/Response**

The CRR workload was set to issue a 200-byte request-message from the client and return a 15,000-byte response-message from the server.

**I/O PROCESSOR CONTENTION**

As described in I/O Subsystem Processor on page 2, contention for the I/O Support Processor may degrade the performance of the Integrated LAN Adapter feature.

For this report we measured the contention from the SAD display and a custom OS/2<sup>®</sup> performance monitor that we used to monitor I/O Support Processor utilization. In our test, Fast Ethernet throughput dropped from 6.406 MBps to 6.301 MBps when the monitors were in use. This insignificant drop suggests that our performance tools had little impact on throughput.

These tools and the Warp Center System Activity Monitor on the I/O Subsystem Processor were active in all the tests contained in this report.

**ABBREVIATIONS**

|            |                             |
|------------|-----------------------------|
| <b>CRR</b> | Connection Request/Response |
|------------|-----------------------------|

|                    |   |
|--------------------|---|
| <b>ERP</b>         | Enterprise Resource Planning  |
| <b>LAN</b>         | Local Area Network  |
| <b>MBps</b>        | Megabytes per second, that is, 1000000 bytes per second   |
| <b>Mbps</b>        | Megabits per second, that is, 1000000, bits per second  |
| <b>MTU</b>         | Maximum Transmission Unit (TCP packet size)   |
| <b>NW-to-S/390</b> | Network to S/390  |
| <b>S/390-to-NW</b> | S/390 to network  |
| <b>RR</b>          | Request/Response  |
| <b>SAP</b>         | S/390 System Assist Processor: the S/390 I/O support processor, not to be confused with the Multiprise 3000 Integrated LAN Adapter feature I/O Support Processor. |
| <b>SNA</b>         | Systems Network Architecture (IBM)  |
| <b>TCP/IP</b>      | Transmission Control Protocol/Internet Protocol   |
| <b>TPNS</b>        | Teleprocessing Network Simulator (IBM program product)  |
| <b>VTAM</b>        | Virtual Telecommunications Access Method (IBM)  |
| <b>WAC</b>         | Wide Area Connection  |
| <b>WAN</b>         | Wide Area Network   |



(C) International Business Machines Corporation 1999

IBM Corporation  
Marketing Communications Servers  
Route 100  
Somers, N.Y. 10589

Printed in the United States of America, November 1999  
All Rights Reserved

This publication was produced in the United States.  
IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM's World Wide Web server at <http://www.ibm.com>.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

® IBM, IBM Logo, eNetwork, ESCON, Ether Streamer, Lan Streamer, Multiprise, OS/2, OS/390, RS/6000, S/390 and VTAM are trademarks or registered trademarks of the International Business Machines Corporation.

Lotus Notes is a trademark owned by Lotus Development Corporation.

Tivoli ADSM is a trademark of Tivoli Systems, Inc.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Windows NT is a registered trademark of Microsoft Corporation.

All other registered trademarks and trademarks are the properties of their respective companies.

# RELIABLE COMPUTER

Reliable Computer Sales, LLC  
1627 Colonial Parkway  
Inverness, IL 60067  
Phone (847) 705-0700 Fax (847) 705-0861  
Email [webmaster@reliablecomputer.com](mailto:webmaster@reliablecomputer.com)  
<http://www.reliablecomputer.com>  
click in this box to visit our web site

®

IBM, ACF/VTAM, AD/Cycle, AIX/ESA, BookManager, CICS, CICS/ESA, CICS/VSE, Database 2, DB2, DFSMS/VM, Enterprise Systems Architecture/390, Enterprise Systems Connection Architecture, ES/9000, ESCON, EtherStreamer, GDDM, Hipersorting, Infoprint, Language Environment, LANStreamer, Multiprise, NetView, OpenEdition, OS/390, Parallel Sysplex, RACF, RAMAC, SAA, SecureWay, S/390, VisualAge, VM/ESA and 3090 are registered trademarks of International Business Machines Corporation in the United States and/or other countries.

TM

C/370, DFSMS, e-business logo, eNetwork, ESA/390, Hiperbatch, IMS, MVS, MVS/ESA, NetRexx, Network Station, Process Resource/Systems Manager, PR/SM, RMF, S/370, S/390 Parallel Enterprise Server, SQL/DS, Ultrastar, VSE/ESA, WebSphere and 3090 are trademarks of International Business Machines Corporation in the United States and/or other countries.

Lotus is a registered trademark and Domino is a trademark of Lotus Development Corp.

Tivoli is a registered trademark ADSM is a trademark of Tivoli Systems, Inc.

Java is a trademark of Sun Microsystems, Inc.

Windows is a registered trademark of Microsoft Corp.

SET is a trademark of Secure Electronic Transaction, LLC

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other products and services are registered trademarks or trademarks of their respective companies.