

IBM* LAN SERVER 4.0 WHITE PAPER
PERFORMANCE, CAPACITY ENHANCEMENTS, & TUNING TIPS

IBM LAN SYSTEMS PERFORMANCE ANALYSIS
DEPARTMENT 55LS
AUSTIN, TEXAS
MARCH 1995
(Revised)

Contents

	Page #
Introduction	1
LAN Server 4.0 Performance Tuning Assistant	2
Introduction to the Tuning Assistant	2
HPFS 386 Cache Size Calculation	3
Examples of Key Parameter Calculations	5
Using Tuning Assistant in "What if" Mode	8
LAN Server 4.0 Configuration Defaults	10
DOS LAN Services Client Performance Considerations	12
LAN Server 4.0 Capacity Enhancements	13
LAN Server 4.0 Support of SMP	14
NetBIOS over TCP/IP	15
Design Considerations	15
Enhancements	15
Performance Characteristics	17
Tuning TCP/IP	18
Recommendation: Dual Protocol Stacks	18
Additional Useful Information	19
Reducing NetBios Broadcast Frames	19
DCDB Replication Performance	19
Upgrading from LAN Server 3.0	19
Considerations when RAW SMBs are disabled	20
DOS TCP/IP	20
Configuring DLS with Windows for Workgroups	20
Additional Tips for LAN Server 4.0 Performance	22
Entry vs. Advanced Server	22
Fixed Disk Utilization	22
CPU Utilization	23
Network Interface Cards	23
Network Media Utilization	24
Performance Benchmark Comparison	25

Introduction

LAN Server 4.0 includes features that allow increased capacity and performance over LAN Server 3.0. Architectural limitations of LS 3.0 have been addressed in LS 4.0. Parameter defaults have been increased so that a newly installed Advanced Server supports 100 users without modifications. Client-side caching has been added to the DOS client resulting in improved performance. A new protocol driver that runs at ring 0 privilege along with OS/2* NetBIOS over TCP/IP provides greater performance than the OS/2 NetBIOS for TCP/IP used by LS 3.0. In addition, LS 4.0 provides a tool (LAN Server 4.0 Tuning Assistant) to help users tune their specific configurations for optimum performance.

The features described above are documented in the LS 4.0 publications. It is the intent of this paper to provide additional information, such as design considerations and performance analysis results, from the LAN Server Performance Analysis group.

LAN Server 4.0 Performance Tuning Assistant

Introduction to the Tuning Assistant

The Tuning Assistant was designed to satisfy the following usability and performance objectives:

- Provide an easy way for users to tune LS 4.0 to their configurations
- Provide performance tuning based on users' unique situation
- Optimize performance parameters but leave safety margin
- Provide a tool to allow "what if" calculations

The following are general rules implemented by the Tuning Assistant during calculations and modifications

- Never reduce parameters below their default value
- Never add or delete lines from any configuration file
- Never exceed maximums such that system will not boot
- NetBios resource requirements will be spread equally over adapters
- LAN Server has priority if NetBios resources are overcommitted

In most environments the important elements in tuning LAN Server for best performance (in priority order) are the following:

- Configure the largest HPFS386 cache possible
- Provide a sufficient number of NUMREQBUF
- Provide a sufficient number of Commands ('x2' parameter in the 'netx' line of the IBMLAN.INI)
- Enough adapters (and NETBIOS resources) for number of users
- Provide a sufficient number of NUMBIGBUF(Entry Server or print spooling only)
- Reserve sufficient memory for GUI if it will be used frequently and delay for swapping is undesirable

HPFS386 Cache Size Calculation

The HPFS386 cache size calculation involves determining all the uses of memory in the system and assigning the remainder to the cache. Listed below are the factors used by the Tuning Assistant in the calculation of the HPFS386 cache size of a system with a memory size of 32.0 MB.

Memory Allocation (MB):

- OS/2 base	2.8
- Spooler	0.7
- MPTS base	0.6
- Memory per adapter	0.2
- LAN Server base	3.5
- IBMLAN.INI additional	1.5
- Heap reserve	1.0
- Cache mgmt (64 bytes/1K bytes for memory in excess of 12M)	1.3
- Reserve for other apps	1.0
- Safety margin (5 percent)	1.6
- Memory assigned to HPFS386 cache	17.8
TOTAL	32.0 MB

The entries for OS/2, MPTS and LAN Server base are the same as found in the Memory Estimating Worksheets in the Network Administrators References Volume 1, Appendix A.

The "IBMLAN.INI additional" entry is for increases made to certain parameters by the Tuning Assistant.

The Heap reserve memory is set to 1 MB; this memory is used by LAN Server for internal file system control needs such as open file handle tables, search handles, filename parsing, etc. This memory is not assigned to the Heap parameter but merely set aside for availability.

The cache management entry is necessary when dealing with large cache sizes as it becomes a significant amount of memory. The formula is applied after subtracting 12 MB from the system memory size since at least 12 MB of memory is always needed and is never available for cache.

By default, 1 MB is always reserved for other user applications to be run on the server. This is a parameter used by the Tuning Assistant to provide the administrator with a significant input to the cache size calculation.

Note: The administrator should determine the memory requirements of any application that is to run concurrently with LAN Server and provide that value in the 'Application Reserve Memory' entry field of Tuning Assistant.

An important example of an additional user application on the server is the new LS 4.0 Graphical user Interface (GUI). If the administrator will regularly use the GUI administration feature, at least 5 MB should be entered for this parameter to provide good performance of the GUI. If this amount of memory is not available, significant swapping will occur when the GUI is started. If only occasional use of the GUI on the server is expected, the recommendation is to leave this parameter at 1 MB and use the additional 4 MB system memory for HPFS386 cache.

Examples of Key Parameter Calculations

NUMREQBUF (IBMLAN.INI)

Optimum number is 2 to 3 per "active" user. Since NUMREQBUF locks the memory from other processes, we want to be efficient. Also, since most uses of NUMREQBUF require a corresponding command it is wasteful to allocate more NUMREQBUF than commands. Only 250 commands per adapter are configured by Tuning Assistant, thus only 250 NUMREQBUFs per adapter will be configured.

Calculation: 2.2 times MAXUSERS with maximum of 250 per adapter.

Special Considerations: Memory used by NUMREQBUF is calculated in Tuning Assistant using a hard coded value of 4096 bytes for each request buffer. If the user wants to change SIZREQBUF from 4096 to 2048, then the calculated HPFS386 cache size can be increased by $(\text{NUMREQBUF} * \text{SIZREQBUF}) / 2$.

A parameter related to NUMREQBUF is the USEALLMEM parameter in the Requester section of the IBMLAN.INI. This parameter allows request buffers to be defined in the memory above 16 MB. If no network interface cards (NICs) are limited to 24-bit direct memory access (DMA), and more than 16 MB of RAM is installed in the machine, set this parameter to 'YES'.

NUMBIGBUF (IBMLAN.INI)

NUMBIGBUF are used only by the ring 3 (Entry) Server when files are accessed on a FAT or HPFS file system or the Printer Spooler is accessed. Because better performance can be obtained by using all available memory for HPFS386 cache, NUMBIGBUF will not be increased if the LAN Server Advanced package is installed.

Calculation: If Advanced Server, NUMBIGBUF = 12 (default). If Entry Server, NUMBIGBUF increases to a maximum of 80 as MAXUSERS increases.

Commands (IBMLAN.INI)

For optimum performance commands also needs to be 2 to 3 per "active" user because of its close relationship with NUMREQBUF. Obviously, if 250 users are logged on through one adapter, each user will not have 2 to 3 commands always available and performance will be less than optimum. The IBMLAN.INI parameter that specifies commands is 'x2' in the 'netx' statement.

Calculation: 2.2 times MAXUSERS with maximum of 250 per adapter

Special Considerations: Commands (NCBS in Protocol.ini) are NetBios resources and must be shared with other NetBios applications like DB2*, Lotus Notes**, etc. If a user specifies a MAXUSERS >=114, commands of 250 will be set for LAN Servers' net1 line, leaving only 4 commands for other NetBios applications. Users should manually reduce commands ('x2') in the net1 line to allow the other applications more NCBS resources if required.

Maxusers (IBMLAN.INI)

Calculation: # DOS/Windows** users + # OS/2 users + # additional servers (if Domain Controller)

Maxshares (IBMLAN.INI)

Calculation: Number home dirs + # aliases + (3 * number shared apps).

Maxconnections (IBMLAN.INI)

Calculation: (MAXUSERS + # additional servers) * 4

Special considerations: Advanced Server maintains its own set of connection resources; this parameter pertains only to resources shared by the ring 3 (Entry) Server such as print aliases. This is also true for the following parameters which are not changed for Advanced Server (if HPFS only): MAXLOCKS, MAXOPENS.

Maxsearches (IBMLAN.INI)

Calculation: The Tuning Assistant sets MAXSEARCHES = 700

Special considerations: Advanced Server maintains its own set of search resources; this parameter pertains only to searches done by the ring 3 (Entry) Server. This value was chosen to provide ample search memory for the ring 3 (Entry) Server.

Sessions (PROTOCOL.INI)

Calculation: DOS/WIN requesters + OS/2 requesters + # additional servers
(if DC) + Lotus Notes requesters + DB2 requesters + User logged
on server + other NetBios requirements

NCBS Calculation (PROTOCOL.INI)

For optimum performance NCBS also need to be 2 to 3 per "active" LAN Server user plus the NetBios commands needed by other NetBios applications.

Calculation: 2.2 times MAXSESSIONS + other NetBios reqmts up to a maximum of 254 per adapter

Special considerations: NCBS in PROTOCOL.INI are shared with other NetBios applications like DB2, Lotus Notes, etc. LAN Server will use a maximum of 250 of the 254.

Using Tuning Assistant in "What if" Mode

In response to feedback from Beta users the "What if" mode was added, although it is not described in the product documentation. This is the capability to run Tuning Assistant calculations on a machine other than where LAN server 4.0 is installed. This will allow a user to provide system configuration information to Tuning Assistant and create tuned configuration files for use on other machines. IBM is interested in your experience with Tuning Assistant. Please post your comments and any operational concerns to the LS40 CFORUM and someone will respond.

The Tuning Assistant's filename is LS40TUNE.EXE; it is located in the IBMLAN directory. The additional parameters which can be used in the command line launch of Tuning Assistant are as follows:

/D:DOMAIN1 - Domain name (has no effect on calculations)
/S:SERVER1 - Server name (has no effect on calculations)
/T:DC(or AS) - Type: Domain controller or Additional server
/P:Entry(or Advanced) Package - Entry or Advanced version
/M:XX - System Memory in MB
/A:N - Number of network interface cards(adapters)
/U - User supplied files (CONFIG.SYS, IBMLAN.INI, PROTOCOL.INI, and HPFS386.INI)

NOTE: /T is always required whenever /U is specified if running on a system with no server installed.

Example 1

```
LS40TUNE /D:DOMAIN1 /S:SERVER1 /T:DC /P:ADVANCED /M:32  
/A:2 /U
```

To run Tuning Assistant this way, all four of the Advanced version configuration files must be located in the current subdirectory with LS40TUNE. This will run on a machine with or without a server installed. The command line values will take precedence over any actual system version of these parameters. If the "Apply" pushbutton is chosen the user supplied files will be changed and no backup files will be made.

Example 2

```
LS40TUNE /M:32 /A:2
```

This will run only on a machine with a server installed. The command line values will take precedence over actual system versions of these parameters. This example could be useful to look at the effects of more system memory or network interface cards on the tuning calculations. When the "Apply" pushbutton is chosen the system configuration files will be changed and backup files will be copied into the \IBMLAN\BACKUP subdirectory with names like IBMLAN.001, PROTOCOL.001, etc. All files that are updated when the Tuning Assistant calculation is 'applied' will have the same suffix.

Warning: The "What if" feature is useful in examining the logic of the Tuning Assistant, but you should be careful when creating actual configuration files for use on systems other than the one on which the tool was executed.

LAN Server 4.0 Configuration Defaults

The Advanced version of LS 4.0 may be used in larger configurations than previous versions. Therefore, default values of a number of parameters have been increased. The objective is to allow many users to run LS 4.0 out-of-box with little or no customized tuning. The Advanced Server will support 100 users in typical environments, however, running the Tuning Assistant may provide an additional performance improvement for some customers. Some changes to the Entry Server and Peer Services were also made. A summary follows:

IBMLAN.INI PARAMETERS	ADVANCED SERVER		ENTRY SERVER		PEER SERVICES	
	LS3.0	LS4.0	LS3.0	LS4.0	LS3.0	LS4.0
maxopens	576	256	576	160	576	128
maxsearches	50	350	50	150	50	50
numbigbuf	12	12	12	6	12	4
numreqbuf	36	250	36	48	36	10
maxshares	16	192	16	64	16	16
maxusers	32	101	32	32	5	5
maxconnections	128	300	128	128	26	26
x1(in net1)	32	102	32	34	32	34
x2 (in net1)	50	175	50	70	50	70

PROTOCOL.INI PARAMETERS	SAME FOR ALL SERVER VERSIONS	
	LS3.0	LS4.0
sessions	40	130
ncbs	95	225

HPFS386 Cache Defaults

The HPFS386 cachesize was specified in the IFS line in CONFIG.SYS in LS 3.0. For LS 4.0 it is specified in the \IBM386FS\HPFS386.INI file with a line reading "cachesize = xxxx" in the FILESYSTEM section.

The algorithm for determining the default HPFS386 cache size has also changed. Previously the cache size was set at 20 percent of the remaining memory after OS/2 was started. This gave a cache size of 2.9 MB on a 16 MB system. This formula will still be used as long as there is less than 20 MB of memory in the system. If the system has at least 20 MB of memory and the user has indicated that the server can use memory above 16 MB for cache, the default cache size will be 60 percent of remaining memory after OS/2 has started. This will yield a cache size of around 18 MB on a 32MB system. This will enable LS 4.0 to provide excellent performance on most systems without any tuning.

As with earlier releases of LAN Server, the USEALLMEM parameter defaults to 'NO'. This restricts access to memory above 16 MB. If no network interface cards (NICs) or disk adapters are limited to 24-bit direct memory access (DMA), and more than 16 MB of RAM is installed, this parameter should be set to 'YES'. This parameter used to be in the CONFIG.SYS file but is now in the FILESYSTEM section of the HPFS386.INI file.

DOS LAN Services Client Performance Considerations

OS/2 Lan Server 4.0 comes with DOS Lan Services (DLS) clients. DLS clients offer substantial performance improvements over the DOS Lan Requester (DLR) clients provided with LAN Server 3.0. Significant performance benefit is realized through the implementation of client-side caching algorithms. In brief, client-side caching offers local caching, reducing requests to the server, thereby increasing overall system performance.

Client-side caching is enabled by default with DLS clients. This means that the AUTOCACHE parameter is set to YES in the NETWORK.INI file (since this a default, it is not specifically written into NETWORK.INI during installation). With AUTOCACHE=YES, the DLS client will allocate big buffers in extended memory(XMS). Each big buffer is 8K in size. The number of big buffers is calculated by the system and is dependent on the amount of XMS available (up to a maximum of 30 big buffers). If a machine does NOT have any XMS the AUTOCACHE parameter is effectively ignored. If you want to configure big buffers on a DLS client that has no XMS, set the following parameters in the NETWORK.INI file:

1. AUTOCACHE=NO
2. SIZBIGBUF=xxxx (in bytes)
3. NUMBIGBUF=xx (integer)

This will allocate big buffers on the client that can be used for large data transfers. However, these buffers will be put in Upper and/or Conventional Memory, reducing available memory for applications.

Another parameter of importance with DLS clients is the WORK BUFFERS. These are the buffers that are used on the requester to process an application's request for data. The default values for WORK BUFFERS on the DLS client are as follows (also set in NETWORK.INI).

1. SIZWORKBUF=1024
2. NUMWORKBUF=2

The above default values are the generally recommended values for the best system performance. If you are unable to use the AUTOCACHE option, you may want to experiment with these two parameters for possible improvements in your environment.

LAN Server 4.0 Capacity Enhancements

As the number of workstations connected to LAN Server 3.0 grew into the hundreds in some installations, an architectural limitation was discovered which has been addressed in LAN Server 4.0. Specifically, a data structure design was limiting the number of request buffers(NUMREQBUF) which could be configured to a maximum of around 350. In large installations this could cause a performance degradation. LAN Server's new design provides future extensibility by allowing the value of NUMREQBUF to be as large as 2000.

The current recommended value for NUMREQBUF is 2.2 per user up to a maximum of 250 for each adapter, or 1000 if four adapters are in the system.

LAN Server 4.0 Support of SMP

LAN Server 4.0 has been tested with and shown to support symmetric multi-processor(SMP) machines running under OS/2 for SMP. LAN Server 4.0 Advanced does not gain additional performance benefits from SMP machines. Its architecture has been optimized to the point where most requests are processed "on interrupt" when received from the network component. The queuing time for a request to be processed is usually extremely short since there are rarely instances when a file/print server's CPU approaches 100 percent utilization. Under these conditions, it would not be expected that an additional CPU would improve response time to the requester. This design provides industry leading performance as evidenced by the LANQuest** report of October 1994. See page 25 for more information.

There are some situations in which LAN Server 4.0 support of SMP does lead to an improvement in total system throughput performance. Since OS/2 is a multi-tasking operating system, other applications can run in the same machine as LAN Server. For other applications which make extensive use of the CPU(e.g. Lotus Notes, etc.), additional processors may make sense. Whenever the CPU workload approaches 100 percent, the additional processor can make a significant difference in the system throughput. LAN Server 4.0 Advanced accommodates the use of the additional processor unless its own workload is unusually high in which case it takes precedence over other applications. LAN Server 4.0 Entry runs with the same privilege as other OS/2 applications and does not take precedence in an SMP environment.

NETBIOS Over TCP/IP

Design Considerations

NetBIOS over TCP/IP is an implementation of NetBIOS that has been specifically designed to operate with IBM TCP/IP. It enables a workstation to be geographically isolated from its domain yet communicate with it transparently.

NetBIOS over TCP/IP is an implementation of the Request for Comments (RFCs) 1001/1002 standards which describe how to enable NetBIOS applications over TCP/IP. It is a B-node, or Broadcast node implementation with routing extensions. A broadcast node uses broadcasting to exchange information between hosts. The routing extensions allow nodes to span subnets through IP routers. These extensions plus the remote name cache discussed below simplify the configuration of RFC 1001/1002 NetBIOS nodes into TCP/IP environments.

NetBIOS over TCP/IP uses an expanded syntax for NetBIOS names that is transparent to NetBIOS applications. The Local NetBIOS Name Scope String is appended to the NetBIOS name creating an expanded name that has the effect of limiting the scope of a NetBIOS name. Two RFC-compliant NetBIOS nodes can communicate only if they have the same Local NetBIOS Name Scope. The Local NetBIOS Name Scope string is defined by the LOCALSCOPE parameter in the TCPBEUI section of the PROTOCOL.INI.

NetBIOS over TCP/IP supports only 1 logical netbios adapter and should therefore be added to only 1 network interface card during the installation/configuration process. However, if TCP/IP is installed on multiple adapters, NetBIOS over TCP/IP will make use of those adapters.

TCPBEUI is IBM's high performance, ring zero protocol driver which maps NetBIOS API calls into the TCP/IP protocol. NetBIOS over TCP/IP contains enhancements over the RFC 1001/1002 standards which improve system performance by decreasing broadcast storms, and expanding communications over routers and bridges. These enhancements, described in the next section, are transparent to NetBIOS applications and do not interfere with other B-node implementations that lack similar functions.

Enhancements

Three of the enhancements to NetBIOS over TCP/IP are in the form of routing extensions. These extensions allow communication between networks and over IP routers and bridges. These extensions are:

1. The broadcast file. A broadcast file contains a list of host names, host addresses, or directed broadcast addresses. It is read at startup and each valid address is added to the set of destination addresses for broadcast packets. Remote nodes included in the broadcast file are then treated as if they were on the local network. Use of a broadcast file has the effect of extending a node's broadcast domain to its own subnet plus any other subnets listed in the broadcast file. A maximum of 32 broadcast file entries are supported, each of which could include additional subnets, thus extending the node's broadcast domain.
2. The names file. A names file consists of NetBIOS name and IP address pairs. NetBIOS over TCP/IP will conduct a prefix search of the names file before broadcasting on the network. The prefix match succeeds if the entry in the names file matches the given name, up to the length of the entry. The first match is used, therefore, the order in which NetBIOS names are listed in the names file is important.

To enable this routing extension, set the NAMESFILE parameter in the TCPBEUI section of the PROTOCOL.INI to a nonzero integer that represents the number of names file entries.

3. The Domain Name Server (DNS). A network administrator can maintain NetBIOS name and IP address pairs in a DNS. If a name query fails NetBIOS over TCP/IP can append the NetBIOS Domain Scope String to the encoded NetBIOS name and issue a request to the DNS to look up an IP address for that NetBIOS name.

The Domain Scope String is defined by the PROTOCOL.INI parameter DOMAINSCOPE.

Another enhancement NetBIOS over TCP/IP provides is a cache for storing remote names that have been discovered. This cache is enabled by setting the NAMECACHE parameter in the TCPBEUI section of the PROTOCOL.INI to a nonzero integer that represents the number of names stored in the directory (NAMECACHE=xx).

The information in the remote names cache (or directory) is also stored on disk and periodically updated. When the system is restarted, this information can be preloaded into the cache at bootup time. Preloading can reduce the amount of broadcast frames on the network since NetBIOS will not have to rediscover names for remote names. To preload the remote names cache, set PRELOADCACHE=YES in the TCPBEUI section of the PROTOCOL.INI.

NOTE: When NetBIOS over TCP/IP is searching for a name, the name cache is checked first, then the names file, the broadcast file, and finally the Domain Name Server.

Recommendation: When running NetBIOS over TCP/IP in a Wide Area Network (WAN), turn name caching on at the server (e.g. NAMECACHE=100).

Performance Characteristics of NetBIOS over TCP/IP

The performance difference between NetBIOS over TCP/IP and NetBEUI can range widely depending on the environment. Some environmental factors that can affect performance are the type of client (OS/2 or DOS), the server CPU workload, the type of network operations being performed, the network media, network congestion, and communication line speeds. We've observed the performance of NetBIOS over TCP/IP being anywhere from 10% slower to as much as 4 times slower than NetBEUI.

One of the environments in which we conducted performance tests was a medium-sized Local Area Network on 16Mbps Token Ring with no WAN connections. We ran a set of industry standard business applications on OS/2 NetBIOS over TCP/IP clients and again on OS/2 NetBEUI clients. In this environment, NetBIOS over TCP/IP was 20% slower than NetBEUI. The performance of DOS NetBIOS over TCP/IP clients was significantly less than that of the OS/2 clients.

Database applications generally use small records when accessing shared databases residing on the server. Often these small records are retrieved from the file system cache with no physical disk access being required. The performance of this type of application on NetBIOS over TCP/IP may be noticeably slower than if the application were run using NetBEUI. However, if the number of database accesses of this type in performing a typical operation is in the order of hundreds, not thousands, the user may not notice a difference in performance in the two protocols.

It may be necessary to periodically update client applications or other files by copying them from the server disk. DCDB replication from a domain controller to a remote additional server also generates I/O operations sometimes known as file transfers. This type of file I/O activity over a network will show little or no performance difference between NetBEUI and NetBIOS over TCP/IP due to protocol characteristics. One should be aware, however, that most WAN connections today are made over relatively low speed communication lines when compared with a LAN speed of 4 to 16 Mbps. File transfer operations over WAN communication lines will probably be slower than over LANs but most likely not due to the network protocol.

Tuning TCP/IP

If you're using NetBIOS over TCP/IP in a Local Area Network environment, file transfer performance might be improved by increasing the maximum transmissible unit (MTU) size. We have seen up to a 20 percent increase in performance of large file transfers by using an 8KB packet instead of the default 1500 bytes. The default of 1500 was chosen because of ethernet's packet size limitation and prevalence in TCP/IP environments.

The MTU size can be changed with the IFCONFIG command in TCP/IP's SETUP.CMD. Set the MTU size to the desired packet size plus 40 bytes, the maximum TCP/IP header size. The desired packet size should be a multiple of 2048.

Your network adapter must be configured to support transmission of buffers that are at least the size specified for the MTU. On an IBM 16/4 Token Ring adapter, this would be accomplished by setting the XMITBUFSIZE parameter in the Token Ring section of the PROTOCOL.INI file. Check your network interface card documentation for information on configuring your adapter.

Recommendation: Dual Protocol Stack

Because there may be a performance difference in a particular environment, it is recommended to configure and use NetBEUI in the Local Area Network (LAN) environment, and NetBIOS over TCP/IP in the Wide Area Network (WAN) environment. The Multi-Protocol Transport Services (MPTS) shipped with LAN Server 4.0 provides the capability of configuring your LAN workstation or server with both NetBEUI and NetBIOS over TCP/IP on the same network interface card.

The dual protocol stack can be configured through the LAN Server installation/configuration program. When selecting protocols, install logical adapter 0 with NetBEUI and logical adapter 1 with TCP/IP and NetBIOS over TCP/IP. This dual protocol stack configuration allows local sessions to continue running with NetBEUI performance while also providing Wide Area Network connectivity with NetBIOS for TCP/IP.

Additional Useful Information

Reducing NetBIOS Broadcast Frames

A key concern with many NetBIOS users is the amount of broadcast traffic that occurs on the network. Broadcasts are used to communicate between nodes. Broadcast storms can slow network performance and overwhelm routers. Use of the Remote Name Directory (RND) function can help to minimize this broadcasting by sending frames to specific nodes when possible.

When using RND, the local station caches the node addresses of remote names that it has located. Any messages sent to that remote name after the node address has been saved is sent directly to that node rather than broadcast to all nodes.

The RND function in LAN Server 4.0 has been extended to include datagrams. RND stores only unique names and no group names, so if an application uses mostly group names for sending datagrams, RND should not be used. Another enhancement to the RND function is that the maximum number of directory entries has been increased from 255 to 2000 when running on OS/2 2.0 or greater.

The parameter RNDOPTION in the NETBEUI section of the PROTOCOL.INI specifies whether RND is turned on or off. Set this parameter to 1 to enable use of the RND function. If RNDOPTION is chosen, make sure that DATAGRAMPACKETS in the NETBEUI section is greater than 2. A related parameter, also found in the NETBEUI section, is NAMECACHE. This parameter specifies the size of the remote name directory. This parameter defaults to 1000 entries.

DCDB Replication Performance

Changes to the DCDB Replicator service for LS 4.0 have yielded substantial performance improvements. In some configurations, users may see up to an 80 percent increase in performance over the LS 3.0 DCDB Replicator service.

Upgrading from LAN Server 3.0

Upgrading from LAN Server 3.0 to LAN Server 4.0 will cause parameters in the PROTOCOL.INI file to be set to the LS 4.0 default values. This may cause performance problems in previously tuned servers. Users who have fine-tuned their PROTOCOL.INI for LS 3.0 should be aware that they may need to make the same changes for LS 4.0.

Considerations when RAW SMBs are disabled

The multiplex read and write SMB protocols are used if the RAW SMB protocol is disabled. These protocols divide data transfers into buffer-size chunks (sizeworkbuf) and chain them together to satisfy large read or write requests.

A parameter that affects performance when working in multiplex mode is PIGGYBACKACK in the NetBEUI section of the PROTOCOL.INI file. This parameter specifies whether NetBIOS sends and receives acknowledgements piggybacked with incoming data. When used with RAW SMBs, piggybackacks improve performance. However, users that attempt to use piggybackacks with multiplex SMBs may see performance degrade by up to 3 times for large file transfers.

Note: The RAW SMB protocol is disabled on a server when srvheuristic 19 in the IBMLAN.INI file is set to 0 (default=1). The RAW SMB protocol on an OS/2 client is disabled when IBMLAN.INI wrkheuristic 11 is set to 0 (default=1) and wrkheuristics 14 and 15 are set to 1 (default=1).

DOS TCP/IP

The LAN Server Performance Team has tested a number of vendor TCP/IP products for DOS. These include Network Telesystems, Wollongong, and FTP TCP/IP offerings. In many cases, these performed considerably better than the IBM TCP/IP protocol stack shipped with LAN Server 4.0. The Network Telesystems product, in particular, showed significant throughput improvement.

While IBM continues to refine their DOS TCP/IP offering, the performance content of each of the OEM products reviewed may provide a near-term solution for running DOS clients in a TCP/IP environment.

In addition to the TCP/IP protocol stack, each of the vendor products includes the normal TCP/IP applications such as FTP, mail, SNMP, etc.

Configuring DOS LAN Services with Windows for Workgroups

You can install both Windows for Workgroups and DOS LAN Services on the same workstation. However, you cannot use the network function of Windows for Workgroups with this configuration. To run DOS LAN Services and Windows for Workgroups on the same workstation, use the following procedure:

1. Install Windows for Workgroups
2. Install DOS LAN Service

3. In the WINDOWS\SYSTEM directory, rename the following files:

From:	To:
VNETSUP.386	VNETSUP.WFW
VREDIR.386	VREDIR.WFW
NETAPI.DLL	NETAPI.WFW
PMSPL.DLL	PMSPL.WFW

4. In the CONFIG.SYS file, REM out the following line:

'DEVICE=C:\WINDOWS\IFSHELP.SYS'

5. In the Windows SYSTEM.INI file, under the '386enh' section, change the line that contains the 'network=' statement to the following:

'network=vnetbios.386,vnetsup.386,vredir.386'

The fix for APAR IC08963 makes the same changes, so you can use the APAR fix if you do not want to change the CONFIG.SYS and SYSTEM.INI files manually.

Additional Tips for LAN Server 4.0 Performance

A number of the major factors affecting performance of LAN Server 4.0 are reviewed in the following sections. Although a few parameters are discussed, most of the tips are aimed at getting you to think about your particular environment in relation to LAN Server's system resources.

Because there will always be a bottleneck in any computer system, the objective of performance tuning is to remove the current bottleneck. Hopefully the resulting system performance has its new bottleneck at an operating point outside normal operating conditions.

Entry vs. Advanced Server

If your LAN Server is to share file, applications, or printers for less than 80 users, the Entry Server will fit your needs with very good performance. The LANQuest report described on page 25 contains a comparison of Entry vs. Advanced Server performance. A subsequent upgrade to Advanced Server is available with minimum impact to your business.

If your immediate requirements are for high performance and high capacity, you will want the Advanced Server. To gain the performance advantage of the Advanced version, your applications and data files must reside on an HPFS386 partition, not on a FAT partition. Neither OS/2 2.1 nor LS 4.0 must be installed on an HPFS386 partition because accesses to system software are infrequent after initial loading.

Fixed Disk Utilization

The disk subsystem, an electromechanical device, can often be the system bottleneck even though the system provides a lot of memory for caching files. If you have observed that your fixed disk activity indicator (the little light that flashes when the hard disk is in use) is on more than it is off for long periods of time, you probably have a disk bottleneck. Your options for improving performance include:

- Distribute the disk-intensive workload from a single physical disk drive to multiple disk drives, enabling concurrent disk seeks and read/writes.
- Off-load some users, files, or applications to another server.
- Install the Fault Tolerance feature of LAN Server to enable disk mirroring. This not only protects your data by backing up your disk but also

improves performance since the additional disk drive will also be used to read data (split reads).

- Adding fixed disks and striping data across them (RAID architecture) will sometimes improve performance as well as enhance data integrity in an environment where data is predominantly looked up (read) without a subsequent update (write), for example, databases used for price lookup, part number information, etc.

CPU Utilization

Server performance can degrade when the computer (CPU)'s ability to process incoming instructions is overtaxed. If there are many users (usually hundreds) with high interaction rates to the server, a CPU performance bottleneck may occur (the Advanced server CPU efficiency is several times greater than the Entry server). You may see a lot of fixed-disk activity and suspect the disk subsystem, but this may be lazy-write activity, which is not necessarily the system bottleneck. To check CPU utilization you can use System Performance Monitor/2 or LAN NetView* Monitor for a detailed analysis.

To get a rough idea of how your server uses the CPU, start the Pulse applet from the OS/2 Workplace Shell* Productivity folder and observe its display during a heavy server workload period. If the CPU utilization level is 80 percent or greater for much of the time, performance is being impacted by the CPU's ability to satisfy its workload demands. Replacing standard network interface cards (NICs) with busmaster NICs will provide additional CPU power and usually improve server performance. Another remedy is to offload some of the users, files, applications, or functions (e.g., domain controller or print server) to another server or to upgrade to a more powerful hardware system.

Network Interface Cards (NICs)

Let's assume that your fixed disk activity is not excessive and that your CPU utilization is generally less than 30 to 40 percent, but you still feel that your server could respond more quickly. Your network interface card (NIC) is analogous to a nozzle which physically limits the amount of traffic flowing to/from the server. Depending on the number of users, speed of the client machines, type of data transactions, etc., server performance can be NIC-limited. NICs come in 8-bit, 16-bit and 32-bit bus widths. Some 32-bit NICs are busmasters, which means they can handle most data transfers with their built-in processors, relieving the server CPU of this task.

You can improve an NIC-limited condition by changing to a faster NIC and/or adding additional NICs to your server. As you add additional NICs, your server

CPU utilization will increase as the server will be busier than before servicing the additional traffic coming through the NICs (nozzles). If you add busmaster NICs, the increase in server CPU utilization will be less significant, as you might expect. LS 4.0 will automatically load-balance sessions across all NICs when you initiate a session. When using standard 16/4 token-ring NICs, we recommend that you use a 16 KB shared RAM size for best performance and memory utilization.

Both LAN Server versions 3.0 and 4.0 now support more OEM NICs than the initial release of LAN Server 3.0. You may obtain the current lists of supported NICs from CompuServe** with the following selections:

1. GO IBM
2. Technical Service and Support
3. IBM OS/2 Forums
4. OS/2 Developer 2 Forum (Browse)
5. LAN Server Library

Network Media Utilization

The physical media over which network traffic flows has a finite capacity. The Ethernet bandwidth limit today is usually 10 megabits per sec (Mbps); token rings today are running at 4 Mbps or 16 Mbps. It is quite possible that with powerful servers and hundreds of clients, LANs can almost saturate the physical media providing interconnection. This is much more likely to occur in Ethernet networks due to the broadcast/collision detection/re-broadcast nature of that architecture.

In large networks interconnecting many clients and servers, the level of network traffic on the wire can impact token- ring network performance. A simple (but not always viable) remedy is to change your network topology. You could add NICs to your server, and separate and isolate clients into LAN segments so that all network traffic is not passing through all machines. The net effect is that the server with two Ethernet NICs now has a greater potential bandwidth (20 Mbps) plus a lower collision level on each of the two segments than on a single Ethernet segment.

This solution is not viable if the machines on the two isolated segments must communicate, since LS 4.0 does not internally route the NetBIOS protocol. More sophisticated ways to reduce network utilization include using the traditional backbone rings and bridges plus the new intelligent switches, hubs, and routers now becoming available.

Performance Benchmark Comparison

In October 1994 LANQuest Labs published a Performance Benchmark Comparison Report assessing the performance of LAN Server 4.0 Advanced and Entry, Windows NT** Server 3.5, and NetWare** 4.02. The results of this benchmarking showed that LAN Server 4.0 Advanced was 38% faster than Windows NT Server and 11% faster than NetWare. For copies of this report, call 1-800-IBM-4FAX and request document 2014.

Trademarks denoted by an asterisk (*) are IBM trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries: IBM, OS/2, DB2, NetView, Workplace Shell

Trademarks denoted by a double asterisk (**) are registered trademarks of their respective companies.