

WHITE PAPER

Date September 1995

Prepared By NetWare
Systems Development

Compaq Computer
Corporation

CONTENTS


History of SMP at Compaq	3
NetWare and MP	4
NetWare 4.1 SMP.....	4
Why Do I Need SMP?	6
Installing NetWare SMP ..	7
MP Driver	8
SMP Kernel	9
Dedicated Mode.....	9
Shared Mode.....	10
Priority Mode.....	10
Tips	11
Other NLMs	11
MONITOR.NLM Statistics	12
Summary.....	12
Appendix: Installing CD Drive Support.....	13

FIGURES

Figure 1: Processor Cache Architecture	3
Figure 2: SMP Architecture.....	5
Figure 3: Database Performance	6
Figure 4: File/Print Performance	7
Figure 5: Dedicated Mode.....	9
Figure 6: Shared Mode .	10
Figure 7: Priority Mode	10

NetWare Symmetric Multiprocessing

NetWare Symmetric Multiprocessing (NetWare SMP) supports the use of multiple processors in a NetWare environment. Each processor has equal access to the system's resources, including system interrupts, memory, I/O bus and peripheral support. SMP provides more processing power, often at a lower cost than multiple single processor systems. The most common need for more processing power is in processor-intensive database environments. NetWare SMP allows you to add additional processors to a system without adding the overhead of an entire machine with its own hardware (disk drives and network cards) and software (operating system and disk drivers). This paper looks at the evolution of Compaq multi-processing systems, focuses on the development of NetWare SMP, provides basic installation instructions and hints for using NetWare SMP on your Compaq server.



COMPAQ

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, registered United States Patent and Trademark Office.

Aero, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, SilentCool, LTE Elite, Presario, SmartStation, MiniStation, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, TriFlex Architecture and UltraView, CompaqCare and the Innovate logo are trademarks and/or service marks of Compaq Computer Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1995 Compaq Computer Corporation. Printed in the U.S.A.

Microsoft, Windows, Windows NT, Windows NT Advanced Server, and SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare Symmetric Multiprocessing

First Edition (September 1995)

Document Number 173A/0795

HISTORY OF SMP AT COMPAQ

Multiprocessing (MP) in the PC industry began with the Compaq Systempro. While this design enhanced performance and supported up to two processors, the support was asymmetrical: only one processor could service interrupts. In the case of the Systempro, all interrupts were processed by the boot processor. The second processor could run applications but could not service interrupts from peripheral boards, timers, or other components.

The Systempro XL and ProLiant family were the first Compaq SMP servers and used the 8259 programmable interrupt controller. This architecture is symmetric: any processor can service any interrupt and there are separate Direct Memory Access (DMA) channels, timers and Programmable Interrupt Controllers (PIC) for each processor. Because of expanding customer needs, Intel has introduced the Multiprocessor Specification (MPS) which is becoming a standard for MP.

Intel's MPS assumes that an Advanced Programmable Interrupt Controller (APIC) is available in the system. The APIC was developed specifically for multiprocessing environments. Later versions of Pentium and P6 chips have the APIC embedded in the processor to simplify design requirements and to guarantee that the APIC functionality will be available for use by operating systems. For SMP environments, the APIC design provides better interrupt handling than previous designs. The largest benefit is that the APIC scheduler dispatches interrupts to the processor executing the least important task.

An additional performance increase is realized through the operating system scheduler. The scheduler recognizes high priority tasks and avoids interrupting these tasks when possible, thus increasing the overall efficiency of the system. The scheduler also attempts to use the same processor that was last used by a particular application. Because each processor has its own cache, if a processor is using information that has already been called from system memory and placed in its cache, performance is improved (see Figure 1).




Figure 1: Processor Cache Architecture

NETWARE AND MP

Standard NetWare is not multiprocessing capable nor is it preemptive. Novell's first multiprocessing development was the NetWare SFT III architecture which uses more than one processor and is asymmetric. One processor is dedicated to the I/OEngine, which responds to hardware events. The second processor is dedicated to the MSEngine which provides network and file services to clients. This design is asymmetric because only the I/OEngine handles interrupts.

NetWare SFT III, like standard NetWare, is non-preemptive. Each NetWare Loadable Module (NLM) must yield control of system resources before another NLM can begin a task. NLM code must be written so that no single NLM controls the system for too long. In a preemptive environment, the operating system controls how long each process is allowed to run before handing control to a different one. Preemptive environments are more useful when servicing priority tasks because the operating system can preempt other tasks and does not have to wait until the NLM yields control.

NetWare 4.1 SMP

The most recent development from Novell is NetWare 4.1 SMP. The standard NetWare operating system still exists and runs strictly on the boot processor as though there were only a single processor available. The SMP functionality, called the NetWare SMP kernel, is an addition to the traditional NetWare architecture and runs on the boot processor and secondary processors.

The SMP kernel design has been added to the NetWare OS architecture to allow tasks to be handled symmetrically and preemptively. The SMP kernel design includes an algorithm for scheduling multiple incoming tasks to the processors available in a given system. In the past, this algorithm was not required because new tasks were added to the queue and processed in sequence by the available processor.

The SMP kernel is preemptive so that database applications can service priority tasks more efficiently. Applications that are written to take advantage of SMP, referred to as SMP aware applications, use the SMP kernel to run preemptively on whichever processors are available in a system. Oracle is an example of an SMP aware application.

Just above the processor level are the MP driver and the Platform Specific Module (PSM). The MP driver activates the secondary processors (processors beyond the boot processor). The PSM is provided by the hardware vendor in order for the SMP kernel to communicate with the processors. The PSM that Compaq supplies is CPQPSM.NLM. This hardware support provides the same functionality as any device driver. For example, a network interface card requires a device driver to communicate between the network card and the operating system. The PSM is the same type of essential communication layer between the OS and the multiple processors.

Note: NetWare SFT-III servers do not currently support SMP.

Figure 2 shows the relationship of the different pieces of SMP architecture and uses Oracle as an example of an SMP aware application.

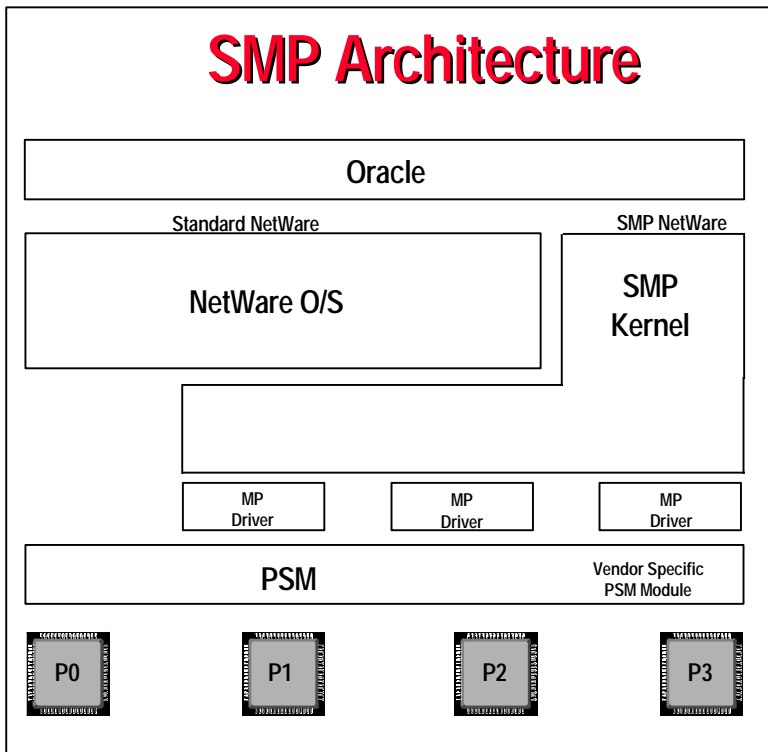


Figure 2: SMP Architecture

NetWare 4.1 SMP is scalable; the more processors in a system, the better the potential performance. Keep in mind that standard NetWare NLMs are not actually SMP aware; they run strictly on the boot processor (P0). NetWare does not schedule standard NLMs on secondary processors. The extent to which the SMP aware applications can use the boot processor is configurable by setting the SMP mode.

WHY DO I NEED SMP?

NetWare SMP offers the greatest potential returns in environments that require intensive processing power. The most common example of this can be found in application serving environments where large enterprise applications stress the capacity of the typical NetWare file server.

Often in the past, the only option users could find to address the capacity requirements of these environments was to implement a new multiprocessing operating system. In most organizations, introducing a new operating system represented a costly decision requiring a significant financial commitment to develop resources to manage the new environment. However, with NetWare SMP, organizations have the option of leveraging the existing skills of their support staffs to implement a high-capacity enterprise application serving environment.

Benchmark tests run in Compaq Integration Labs suggest that NetWare SMP can provide a highly scalable platform for application serving environments. Figure 3 illustrates the results of database tests run to assess the scalability of NetWare SMP.

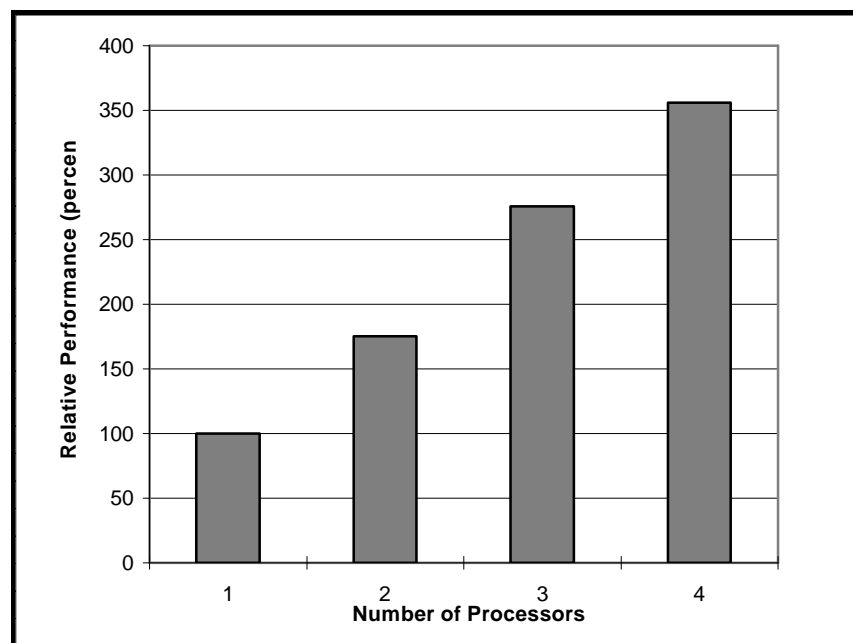


Figure 3: Relative Performance in a Database Environment

To obtain this benefit in an application serving environment, your network applications must be designed to take advantage of NetWare 4.1 SMP. Contact your application vendor to find out about current support for NetWare SMP. Novell is also updating applications, such as GroupWise and MultiProtocol Router v3.0, to support the SMP architecture.

NetWare SMP also offers some opportunities to organizations running NetWare to provide basic network file and print services. Large file serving environments that stress the processor capacity of existing NetWare servers could see significant benefit from NetWare SMP. In NetWare 4.1 SMP, the LAN driver interrupts are moved to the secondary processors, reducing the load on the boot processor and making more capacity available for other functions. Server environments that support multiple LAN segments and maintain high processor utilization are most likely to benefit from this capability. Figure 4 illustrates the results of file server benchmark tests run to assess the scalability of NetWare SMP in this environment.




Figure 4: Relative Performance in a File/Print Serving Environment

INSTALLING NETWARE SMP

Before beginning the SMP installation process, NetWare 4.1 must be installed and the server must be running. You must also have an operational CD drive. If you do not have support loaded for your CD drive, see the appendix for instructions.

NOTE: Using Compaq Systems Configuration Utility, make sure that interrupts are not being shared between LAN and non-LAN devices. For example, the Compaq NetFlex-3 and the Compaq 32-Bit Fast-Wide SCSI Controller cannot share an interrupt. Sharing of interrupts between LAN and non-LAN devices will lockup the server.

1. To start the server, type `SERVER` at the DOS prompt.
2. Insert the CD containing the SMP files into the CD drive.
3. Type `LOAD INSTALL` at the system console prompt.
4. Choose "Product Options," and then "Install a Product Not Listed" from the menu.

5. The screen that appears assumes the install will be from the A drive. Press F3 to enter the path of the CD-ROM drive. Press F10 to continue.
6. When prompted, check the option to install SMP NetWare. You may also check the option to install documentation. Press F10 to continue. Several files will be copied from the CD-ROM to the directory where SERVER.EXE resides on the C drive and to the SYS:\SYSTEM directory.
7. When prompted, insert the diskette containing the Compaq PSM, CPQPSM.NLM.
Note: CPQPSM.NLM will be renamed PSM.NLM and the LOAD PSM.NLM command will be placed in the STARTUP.NCF file.
8. When prompted for the drive and/or path for booting NetWare, indicate the drive and path where SERVER.EXE resides, for example: C:\SERVER.4X.
9. When prompted for the license diskette, insert the NetWare license diskette in the A drive and press <ENTER>.
10. Exit INSTALL.
IMPORTANT: Check the AUTOEXEC.NCF file to ensure that the MPDRIVER is loaded before the LAN drivers. If it is not loaded first, LAN interrupts will not be distributed properly across all available processors. Also, the latest MSM and TSM files are only copied to the SYS directory during installation. If, in the AUTOEXEC.NCF file, you are loading these files from a local drive, such as C:, make sure you copy these files to that directory. SMP will not load without the correct MSM and TSM versions.
11. Type DOWN to down the server, and then EXIT to return to DOS.
12. Type SERVER at the DOS prompt to restart the server as an SMP server. You can use MONITOR.NLM to inspect the state of your server.

MP DRIVER

During installation, the following command is automatically added to AUTOEXEC.NCF. This command loads the MP driver, MPDRIVER.NLM, once for each secondary processor (any processor beyond the boot processor).

```
LOAD MPDRIVER ALL
```

If you wish to use only selected secondary processors, you must unload the MP drivers for all of the secondary processors, and then load the MP driver for the selected processors. For example, if there are four processors in a system, the MP driver may be loaded for any of the three secondary processors using the following commands:

```
UNLOAD MPDRIVER (unloads SMP support for all secondary processors)
LOAD MPDRIVER 1 (brings up the second processor in the system)
LOAD MPDRIVER 2 (brings up the third processor in the system)
LOAD MPDRIVER 3 (brings up the fourth processor in the system)
```

To see which processors are on line, use MONITOR.NLM.

SMP KERNEL

The SMP kernel support, SMP.NLM, is added to STARTUP.NCF during the installation process. There are three SMP modes: dedicated, shared, and priority. These modes control how SMP aware applications use the boot processor, thus controlling the extent to which NetWare NLMs, which run only on the boot processor, are disrupted.

- Dedicated** Boot processor is dedicated to NetWare. NetWare request from SMP aware applications are queued with requests from standard NetWare NLMs.
- Shared** NetWare and SMP aware applications share the boot processor. (The default mode.)
- Priority** Boot processor is dedicated to NetWare, but NetWare requests from SMP aware applications have priority on the boot processor over requests from standard NetWare NLMs.

To change the SMP mode, use one of the set parameters described below. (Note: When the server is started it will be in shared mode by default unless the parameter is changed manually.)

```
SET SMP NETWARE KERNEL MODE = DEDICATED
```

```
SET SMP NETWARE KERNEL MODE = SHARED
```

```
SET SMP NETWARE KERNEL MODE = PRIORITY
```

Dedicated Mode

Dedicated Mode should be used when you want file server response (NetWare NLMs) to be minimally affected by the SMP aware applications. The boot processor is dedicated to NetWare. Even when extra CPU time is available, SMP aware applications cannot use the boot processor to service SMP application calls. When an SMP aware application makes a NetWare call, the call will be added to the tail of the queue and serviced by the boot processor when the request reaches the front of the queue. Figure 5 illustrates dedicated mode.

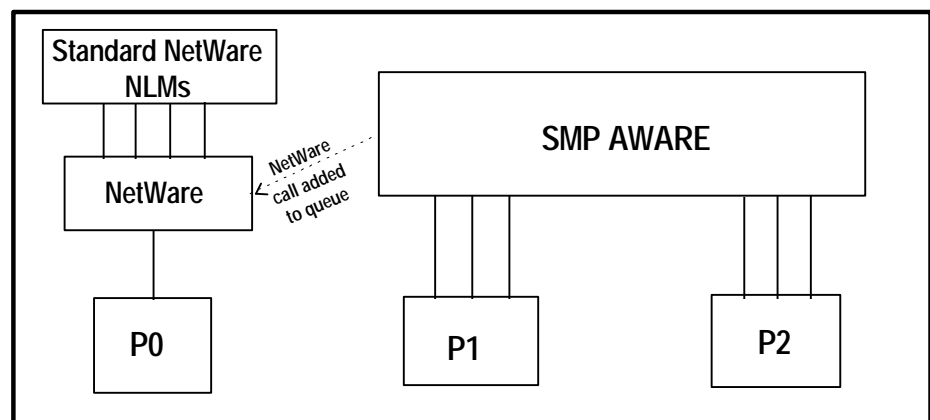


Figure 5: Dedicated Mode

Shared Mode

In shared mode (the default mode), both NetWare and SMP aware applications use the boot processor. This mode should be used when the file server activity does not need priority service and the SMP aware application could benefit from more processing power. CPU time is distributed on a first come, first served basis. In shared mode, NetWare NLMs running on P0 will be more affected than when using dedicated mode. Figure 6 illustrates shared mode.

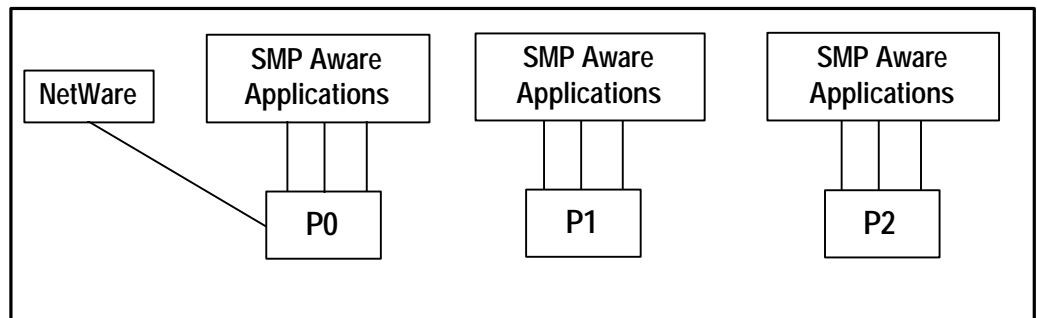


Figure 6: Shared Mode

Priority Mode

Priority mode is designed for a system containing eight or more processors. The boot processor is reserved for NetWare, but if an SMP aware application makes a NetWare call, that call is placed at the front of the queue and will be the next request serviced. SMP aware applications do not interrupt a NetWare task being serviced, but numerous NetWare requests from SMP aware applications queued ahead of other outstanding NetWare requests will slow down their processing. Priority mode, more than the other two modes, affects the servicing of NetWare NLMs running on P0 the most. Figure 7 illustrates priority mode.

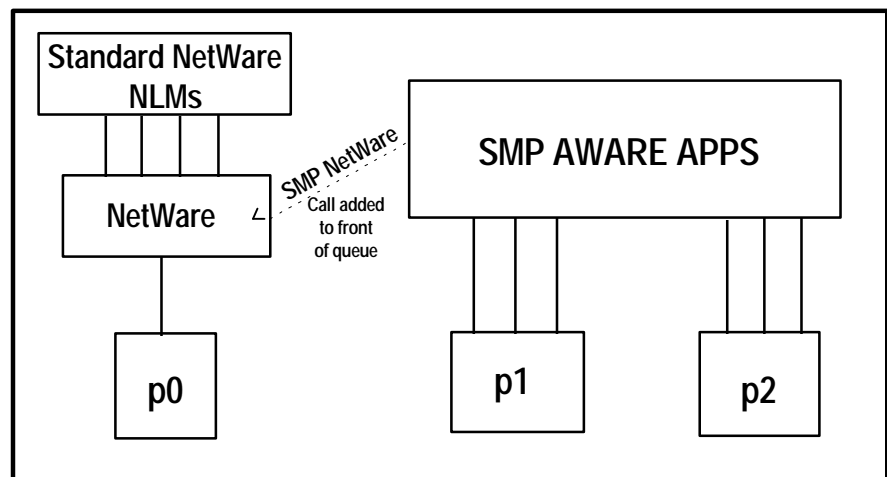


Figure 7: Priority Mode

TIPS

- You must use shared mode to run SMP NLMs with only one processor (the boot processor) loaded.
- If you are trying to run an SMP aware application, but the SMP NLMs are not loaded, the application may not load. However, some SMP aware applications, such as MSM, will load and adapt themselves to the single processor environment.
- At least 64 megabytes of memory is recommended. The minimum requirement is 32 megabytes.
- All NetWare 4.1 patches should be applied before loading SMP.
- You should use the latest NLMs as they become available. CPQPSM.NLM (renamed PSM.NLM during installation), MSM, ETHERTSM and TOKENTSM ship on the NetWare Programs from Compaq disk. Other NLMs can be obtained from Novell.
- Once you load the SMP NLMs, you must down the server to unload them. First, to prevent SMP from automatically loading when you restart the server, remove the following lines:

```

from STARTUP.NCF    LOAD PSM.NLM and LOAD SMP.NLM
from AUTOEXEC.NCF  LOAD MPDRIVER ALL
    
```

Second, down the server. When you restart the server, it will no longer be an SMP server. You may use MONITOR.NLM to verify that SMP is not loaded.

OTHER NLMs

Several NLMs necessary for the SMP server to operate properly are added or upgraded automatically during the installation process. As new versions of the NLMs become available, they should be upgraded. The following NLMs are upgraded for you automatically during the installation process:

CLIB.NLM You must use this latest CLIB.NLM when running SMP aware applications. CLIB.NLM is backwards compatible for use with non-SMP servers and legacy NLMs.

STREAMS.NLM
 ETHERTSM.NLM or
 TOKENTSM.NLM
 MSM.NLM

The Compaq LAN drivers and disk drivers do not require an upgrade. Compaq's current drivers provide full SMP support. However, you should always use the most current drivers available. Compaq supports several forums for obtaining the latest drivers:

Electronic Media Addresses	
BBS (Compaq support bulletin board)	Compaq Download Facility at (713) 378-1418
CompuServe	GO COMPAQ
Internet email	Support@compaq.com
World Wide Web	http://www.compaq.com
FTP	FTP://FTP.COMPAQ.COM

MONITOR.NLM STATISTICS

MONITOR.NLM has also been upgraded and now includes SMP statistics. When you load MONITOR.NLM, you will see several new statistics that will help you inspect the state of your SMP server. In addition to seeing the number of processors in a system and the percentage utilization of each processors, you can tell if the processors are on-line or off-line.

You can also monitor the amount of activity on each processor. This statistic can help you determine whether bottlenecks exist within the system. For example, if your system is configured in dedicated mode and you see that very little activity is occurring on the boot processor, you may want to change the mode to shared mode.

In particular, you may want to watch the behavior of your SMP aware applications by checking the Mutex Information. This statistic indicates that a thread is locking a resource. Any resource, such as buffers, registers, memory, and devices can be locked by an SMP thread (multiple threads make up an application). Upon completing an operation, an SMP thread unlocks the resource it used; however, during the time a resource is locked, no other thread can use it. Resources that are shared by multiple SMP aware threads need locks to ensure the integrity of an operation being serviced. For example, the mutex list may be locked temporarily by an SMP aware thread. This prevents any other SMP aware thread from changing or adding to that list while the first thread is completing its operation. The longer and more frequently resources are locked, the more likely performance will be affected. If you are concerned about performance, you may want to consider removing applications that lock resources.

MONITOR.NLM can also be used to determine whether your system is running SMP aware applications. Because NetWare always runs on the boot processor, you can check MONITOR.NLM for any tasks that are being run on the secondary processors. If the secondary processors are being utilized, your system is running SMP aware applications.

SUMMARY

Adding processors to your system can be useful when you are running SMP aware applications or using Compaq's ODI compliant NIC drivers. You can then configure your system so that the SMP aware applications use the additional processors, leaving the boot processor free to run NetWare. Early tests indicate that systems with heavy LAN traffic or 100 Mb/s LAN traffic directed towards SMP aware applications can show a performance improvement when run in an SMP environment. As Compaq learns more about performance benefits and configuration parameters, the information will be published. Check the World Wide Web (<http://www.compaq.com>) or Compaq's other electronic media addresses for future updates.

APPENDIX: INSTALLING CD DRIVE SUPPORT

You will need the following set of Compaq DOS CD-ROM drivers which provide support for all Compaq 32-Bit SCSI-2 Controllers. The drivers are located on the NPFC diskette v2.57 or greater. For more information see the INSTALL.RDM file on the NPFC diskette.

```
FWS2ASPI.SYS
FWS2CD.SYS
```

Complete the following steps to install the drivers.

1. Copy the FWS2ASPI.SYS and FWS2CD.SYS files from the /SCSI/DOS directory of the NPFC diskette to the root directory of the C: partition:

```
COPY A:\SCSI\DOS\FWS2ASPI.SYS C:\
COPY A:\SCSI\DOS\FWS2CD.SYS C:\
```

2. To allow the driver to be loaded the next time your system is restarted, edit the CONFIG.SYS file and add the following lines:

```
DEVICE = FWS2ASPI.SYS
DEVICE = FWS2CD.SYS /D:name
```

NAME: Name is an eight-character name you assign to the device. For example to assign the name CPQCDR01:

```
DEVICE = FWS2CD.SYS /D:CPQCDR01
```

Note: The name parameter must not match possible filenames that are on a CD-ROM disk. For example, if you assigned the name parameter as "CDROM," files on the CD, such as CDROM.NLM, could not be loaded or accessed.

3. Delete or comment out the following statements if they exist in the CONFIG.SYS file.

```
HIMEM.SYS
SETVER.EXE
```

4. If you are using DOS 5.0 or earlier, copy MSCDEX.EXE from the diskette that came with your CD-ROM drive to your DOS directory. If you have DOS 6.0 or greater, do not copy this file because it is already present in your DOS directory.
5. Edit the AUTOEXEC.BAT file and add the following line:

```
C:\DOS\MSCDEX.EXE /D:name
```

Remember, the name parameter can be any name of up to eight characters, but it must match the name assigned in the CONFIG.SYS file. For example to assign the name CPQCDR01:

```
C:\DOS\MSCDEX.EXE /D:CPQCDR01
```

6. Reboot the machine to load the DOS CD-ROM support.