Configuration and Tuning of Sybase Adaptive Server Enterprise 11.5.1 for Microsoft Windows NT on Compaq Servers

White Paper

Prepared By
Database Engineering
Compaq Computer Corporation

July 1998

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.

The software described in this guide is furnished under a license agreement or non-disclosure agreement. The software may be used or copied only in accordance with the terms of the agreement.

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.

Configuration and Tuning of Sybase Adaptive Server Enterprise 11.5.1 for Microsoft Windows NT on Compaq Servers

First Edition (July 1998)

Document Number ECG037/0798

Compaq Computer Corporation

Table of Contents

Introduction 1

Tuning Goals 1

Planning and Installation 1
- Primary Domain Controller, Backup Domain Controller or Windows NT Server 1
- System Processor Planning 2
  - UniProcessor and MultiProcessor Support from Compaq 2
  - Processor Scalability 3
  - Adding Engines to Sybase 5
  - Sybase Engine-to-CPU Affinity 5
  - Engine Groups and Execution Classes 6
  - Resource Limits 6
- Memory Planning 6
  - Initial Memory 6
    - Reducing Memory Access Overhead 7
    - User Stack Size 7
    - Greater than 2GB User Memory Support 7
    - Modifying Sybase ASE 11.5 for 3GB Support 8
  - Virtual Memory 9
    - Windows NT Memory Management Setting 10
- Disk Subsystem Planning 10
  - RAID Performance Comparison 11
  - Array Accelerator: Its Function and Benefit in a Sybase Environment 12
    - Checkpoints, Housekeeper and Transaction Log Writes 12
  - Fault Tolerance Considerations 13
    - 1. RAID-based Disk Controllers 13
    - 2. Windows NT-based RAID 14
    - 3. Sybase Mirroring/Duplexing 14
  - Support for Devices Over 2GB in Size 15
  - FAT vs. NTFS vs. Direct Disk Access 15
  - NT Disk Device Drivers 17
  - Large I/O Sizes 17
  - Tempdb Placement 17
  - Asynchronous Prefetch 18
  - I/O Pacing 19
  - Compaq System Configuration Utility Considerations 20
- Network Planning 21
  - Network Characteristics of a Sybase Environment 21
  - Network Protocols: Packet Sizes 21
  - Reducing Network Overhead 22
  - Multiple Network Protocols 23

Parallel Processing Features 23
- Parallel BCP and Table Partitions 24
- Parallel Backup 25
  - Parallel Dump 26
  - Parallel Load 27
- Parallel Processing 28
  - Parallel Query 28
    - Parallel Query Index and Table Scans 28
  - Parallel Index Creation 29
  - Parallel Sorting 30
  - Engine Groups and Execution Classes 30
- Parallel DBCC commands 31

Dbcc checkstorage 32
Parallel Recovery 32

Creating a Database Using RAID and Parallel Processing 33

Microsoft Windows NT 4.0 vs. SCO UnixWare 7.0 34

Performance Monitoring and Tuning 36
Performance Monitoring Tools 36
Windows NT Performance Monitor 36
Microsoft Network Monitor 37
Sybase Monitor Services 37
Sybase SQL PerfMon Integration 37
Sybase sp_sysmon Command 38
Compaq Insight Manager 38
System Processor Utilization, Monitoring and Optimization 38
Optimizing Windows NT 4.0 & ASE Operations 38
Windows NT Thread Priority 39
Sybase –P Boost Priority Option 39
Windows NT Maximize Throughput for Network Applications 39
Sybase Runnable Process Search Count 39
Sybase Max Online Engines 40
Unnecessary Services 40
CPU Performance Monitor Objects and Their Meaning 40
Memory Utilization, Monitoring and Optimization 41
Window NT Paging: Detection and Avoidance 41
Sybase Tuning Parameters 41
Memory 41
Dbcc Memusage 42
Procedure Cache 42
Disk Subsystem Utilization, Monitoring and Optimization 43
Recovery Interval 43
Diskperf 43
Disk Performance Monitor Objects and Their Meaning 44
Network Utilization, Monitoring and Optimization 45
Sybase sp_sysmon command 45
Network Performance Monitor 45

Sybase Runserver Command File and the Windows NT Registry 45

Configuring Sybase ASE for High Availability 46
Compaq On-Line Recovery Server 46
Microsoft Cluster Server 46

Configuring Compaq Tape Devices for Sybase Backup Server 47
Verifying Compaq SCSI Device Drivers 47
Verifying Compaq Tape Device Drivers 47
Defining Sybase Tape Device Capacity 47
Performing a Sybase Backup 47

Configuring Windows NT and Sybase ASE for High Security 50
C2 Security Rating 50

Conclusion 52

Sybase SQL PerfMon Integration Counters 53

Configuration and Tuning of Sybase ASE 11.5.1 for Microsoft Windows NT on Compaq Servers

Introduction

The purpose of this document is to share the knowledge acquired by Compaq Systems Engineers in the area of configuration and tuning of Sybase Adaptive Server Enterprise (ASE) 11.5.1 for Microsoft Windows NT on Compaq Servers. It is our desire to deliver the best technical information possible on a specific topic in a timely manner and in a highly usable format. Any comments, suggestions and feedback are appreciated.

The information presented in this document is applicable to Sybase ASE 11.5.1 running on Microsoft Windows NT 4.0 Server Enterprise Edition with Service Pack 3. Because some tuning techniques might differ between releases, we strongly encourage you to upgrade to the latest available versions of Microsoft Windows NT, Microsoft Windows NT Software Support (NT SSD) from Compaq, the Compaq Resource Paq for Microsoft Windows NT and Sybase ASE.

The reader should already be familiar with the installation and reconfiguring of Windows NT Server, Compaq SSD for Windows NT, Compaq Resource Paq for Windows NT and Sybase ASE for Windows NT.

Tuning Goals

In order to achieve the best performing system possible there are several factors that must be reviewed. These include optimization of the hardware, the Sybase dataserver, the operating system and the application software. This paper will focus on the hardware, Sybase and the operating system. It is also important to tune the Sybase application to take advantage of the system. Due to the diversity of database applications, they are beyond the scope of this paper. This paper will discuss processor scalability and tuning, disk controller optimization and I/O tuning, memory tuning, and network tuning. Also, specific Sybase and Windows NT configuration and tuning issues will be presented throughout the paper.

Planning and Installation

This section discusses the planning and installation of Microsoft Windows NT on your network server. These planning considerations cover the areas of the system processor, initial and virtual memory configurations, disk subsystem, and network planning.

Primary Domain Controller, Backup Domain Controller or Windows NT Server

During the installation of Windows NT, you can designate the ProLiant system to be a Primary Domain Controller , Backup Domain Controller, or a Windows NT Server.

The Primary Domain Controller provides a centralized management location for the network and performs the validation of all user logins and permissions. To provide the best performance, avoid running Sybase on a Primary Domain Controller server. The performance impact can vary depending upon the number of user logins, level of security, and the available system bandwidth.

The Backup Domain Controller[1], does not incur as much overhead in the network management and administration as the Primary Domain Controller. The Backup Domain Controller server is part of a domain and incurs overhead in maintaining a replicate copy of the Primary Domain Controller server. One could use this system to run Sybase on but if the Primary Domain Controller goes off-line; the Backup Domain Controller running Sybase is elected (automatically) and promoted to Primary Domain Controller. The impact on the server can vary depending upon the number of user logins, level of security, and the available system bandwidth. For this reason it probably is not the best candidate to run Sybase on.

A Windows NT Server can be member of a Windows NT Domain or a Windows NT Workgroup. As a member of a Windows NT Domain, the Primary Domain Controller system handles all of the administrative tasks for that domain. This leaves all of the Windows NT Server system resources available for running Sybase. This is the optimal system to install and run Sybase. As a member of a Windows NT Workgroup, there is no Windows NT Domain controller and the management of the entire network, users, and shared resources is on a per system basis. The throughput of the Windows NT Server running Sybase depends on the level of requests from other users on the network in the same workgroup.

To receive optimal performance of Sybase on your Windows NT network, use four ProLiant systems. The first ProLiant system would be configured as the Primary Domain Controller with Windows Internet Name Service (WINS), Dynamic Host Configuration Protocol (DHCP), and Gateway modules. This system would incur all of the overhead of user logins, network gateway functions and user permissions. The second ProLiant server would be configured as the Backup Domain Controller. The third ProLiant server, configured as a Windows NT Server participating in a Windows NT Domain, would be dedicated as the Sybase database server. The fourth ProLiant server, also configured as a Windows NT Server participating in a Windows NT Domain, would provide shared file and print services.

For additional information on Domain Controllers, refer to the *Windows NT Server, Concepts and Planning Guide*.

System Processor Planning

The objective of this section is to outline the configurations of the Compaq ProLiant Family of Servers and to provide some performance information necessary to determine the best configuration for your environment. As is the case throughout the document, a certain level of experimentation and internal performance testing in your own environment is expected.

UniProcessor and MultiProcessor Support from Compaq

Windows NT utilizes different versions of the following files for single processor and multiple processor systems:

- *HAL.DLL*
- *KERNEL32.DLL*
- *NTDLL.DLL*
- *NTOSKRNL.EXE*
- *WINSRV.DLL*

The Windows NT installation program automatically selects the correct versions of these files according to the number of processors installed. However, if you upgrade your system from one

[1] You must have a Domain Controller installed and active on the network before you install a Backup Domain Controller.

system processor to two or more, you must install the multi-processor support files[2] by running the *SETUP.CMD* file on the Windows NT SSD from Compaq[3].

Processor Scalability

With the rapid growth of processor technology the systems of today are performing tasks once only thought possible in a mini- or mainframe computing environment. As the systems have become more efficient the scalability of the systems has increased along with them. Applications that a few years ago required the power of multi-processor systems are now capable of running on single processor systems.

Chart 1 shows the performance increases when adding additional PentiumPro processors to a ProLiant 7000 system. The test was based upon an OLTP test of a parts inventory control database.

Chart 1: PentiumPro Processor Scalability of Compaq ProLiant 7000

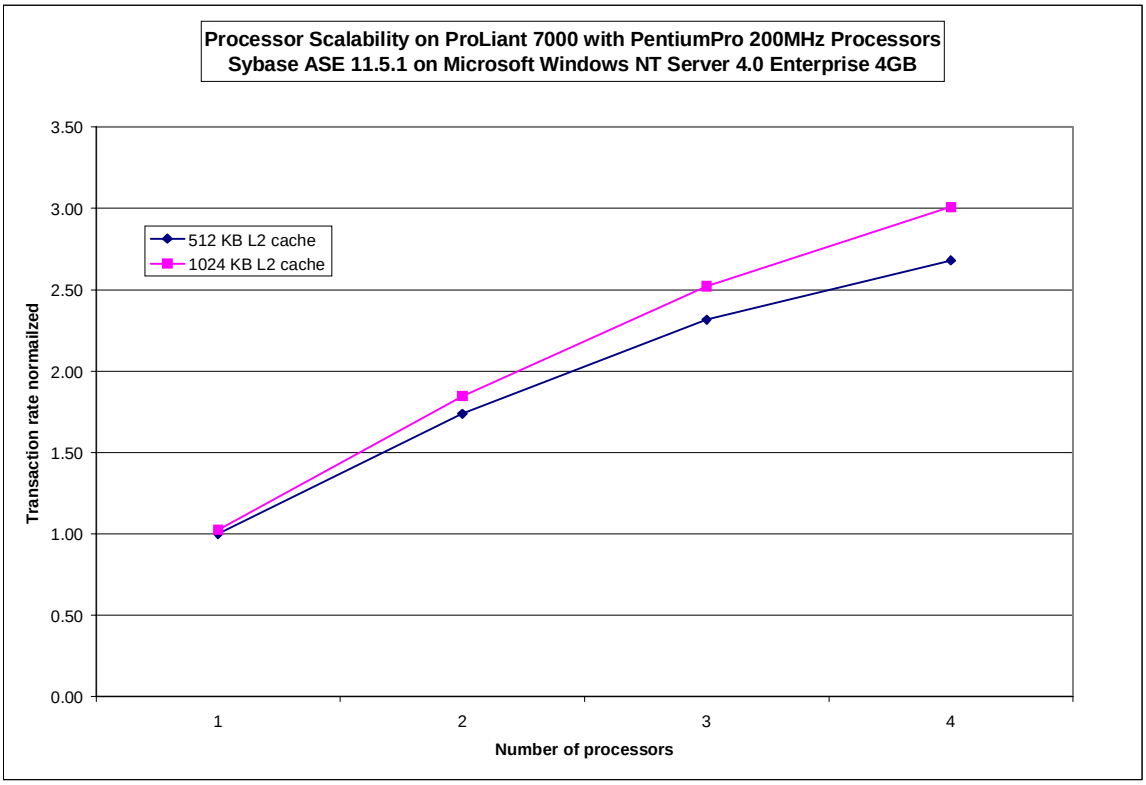

Chart 2 shows the effect of upgrading a ProLiant 5000 with PentiumPro 200MHz 256 KB level 2 cache processors to 512 KB or 1024 KB level 2 cache processors.

[2]This procedure is not necessary if you already have a multiprocessor configuration and you add another processor.

[3] The Compaq Support Software Diskette for Windows NT (NT SSD) contains the latest drivers and utilities that enable you to take advantage of specific capabilities offered on Compaq products. Always make sure you have applied the latest available NT SSD. You can verify the currently installed versions of these files via the File Manager, Properties option.

Chart 2: PentiumPro Processor Scalability of Compaq ProLiant 5000

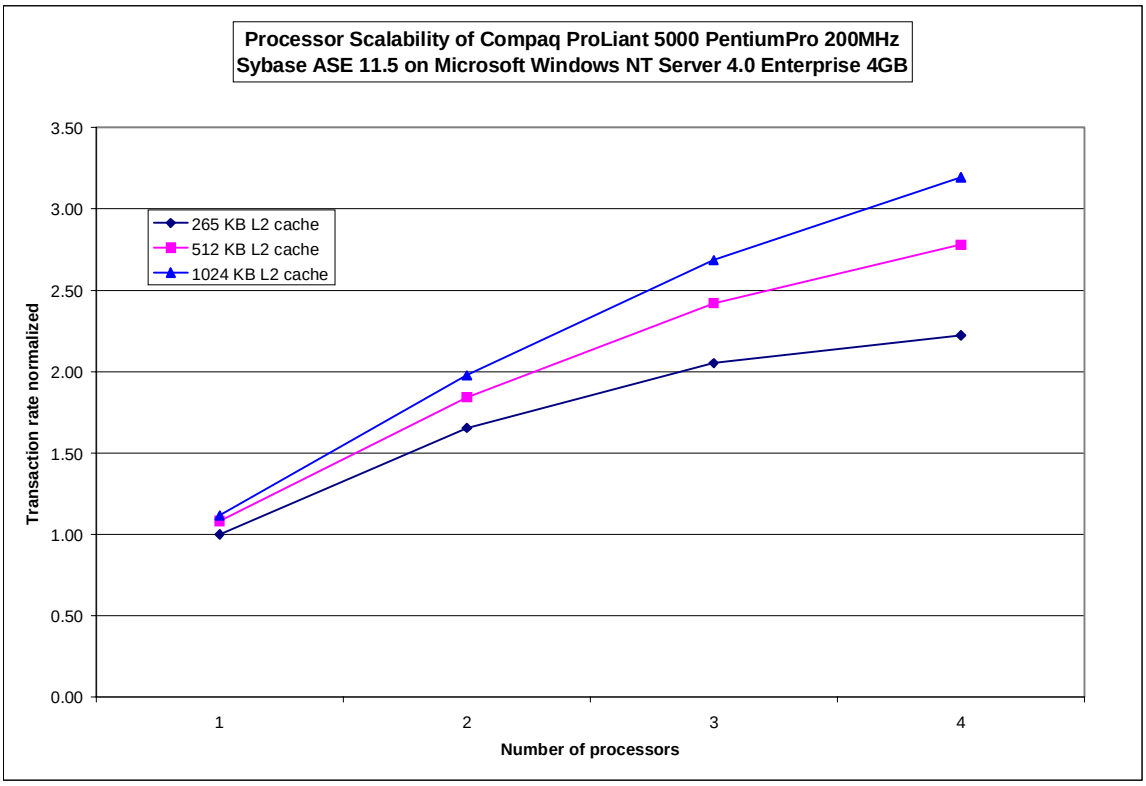

At a first glance it would appear that the Compaq ProLiant 5000 is experiencing better scalability than the Compaq ProLiant 7000 but this is not the case. Both systems were near the maximum configuration for the number of SMART-2 Array disk controllers and drives. The Compaq ProLiant 7000 was configured with 8 SMART-2DH Array controllers connected to 98 @ 9.1GB RAID0 data drives and 8 @ 9.1GB RAID1 log drives. The Compaq ProLiant 5000 was configured with 6 SMART-2/P Array controllers connected to 72 @ 4.3GB RAID0 data drives and 12 @ 4.3GB RAID1 log drives. The Compaq ProLiant 7000 was configured with roughly 68 GB of raw ASCII data placed into a database allocation of 136 GB. The Compaq ProLiant 5000 was configured with roughly 52 GB of raw ASCII data placed into a database allocation of 104 GB. The ProLiant 7000 was running with approximately 9000 simultaneous OLTP users where as the Compaq ProLiant 5000 was running with approximately 6400 simultaneous OLTP users.

Chart 3 shows the effect of upgrading the ProLiant 5000 PentiumPro 166MHz processors with 200MHz processors. In this OLTP test the processor throughput was the limiting factor. As the processor clock frequency increased roughly 21% from 166 MHz to 200 MHz the overall throughput of the system increased only 6%. This can be attributed to the fact that as the processor clock frequency increased the performance bottleneck appears to be either disk or memory bound. An increase in the amount of system memory or number of disk drives and controllers maybe necessary to see the full benefit of upgrading the processors. You may experience different results based upon the type of workload in your environment.

Chart 3: 166MHz to 200MHz Processor Upgrade on Compaq ProLiant 5000

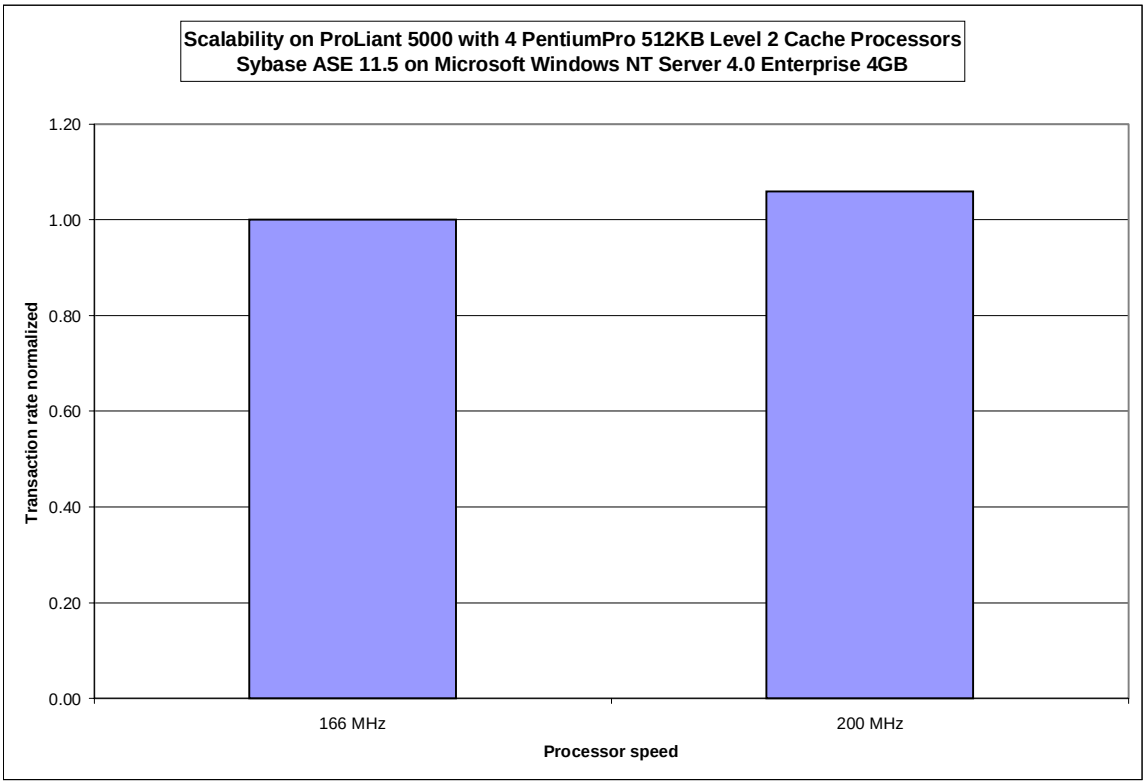

Adding Engines to Sybase

If you add an additional processor or processors to a server, you may need to reconfigure the number of Sybase engines. The parameter that sets the maximum number of Sybase engines that can be started when Sybase ASE boots is called *max online engines*. If this parameter is set to a number greater than the number of processors in the machine, then Sybase automatically starts with one engine per processor. The Sybase server *errorlog* will have a message indicating the number of configured engines that failed to start.

For a single processor environment, *max online engines* should be set to one. For SMP environments, one engine per processor is common. For example, if you have 4 processors in the server and most or all processes being performed on the system are Sybase processes, then you should set *max online engines* to 4. If there is a need to execute non-Sybase applications on the system, then it may be appropriate to set the parameter to 3, using only 3 of the 4 processors for Sybase engines. This will leave one processor available at all times to perform the non-Sybase application processes.

Sybase Engine-to-CPU Affinity

Sybase ASE for Windows NT supports CPU affinity. When you enable engine-to-CPU affinity, the probability that a process that was executing on a particular CPU and is suspended awaiting the completion of an event will resume execution on the same CPU is increased. This can potentially increase the throughput of a high throughput system, like an OLTP system. To enable or disable CPU affinity use the *dbcc tune cpuaffinity* option. To enable CPU affinity for all Sybase engines and CPUs issue the following command:

dbcc tune (cpuaffinity, 0, on)

Please refer to the *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide* for a more detailed discussion of the *dbcc tune cpuaffinity* option and how to monitor its effect in your environment using the *sp_sysmon* utility.

Engine Groups and Execution Classes

The Logical Process Manager (LPM) provides the ability to distribute the engine resources between different tasks. This feature provides the methodology necessary to run both OLTP transactions and DSS queries on the same system with minimal performance impact on either class of users. For example on a 4 processor system you could define three of the Sybase engines for OLTP activity and the fourth engine for DSS activity.

Sybase provides a wide range of flexibility in the definition of and assignment of engine groups and execution classes. Please refer to the following documentation for a more in-depth discussion of this topic:

- *"Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide"*, Document Id# 32645-01-1150-02
- *"Configuring Sybase® Adaptive Server™ Enterprise for Windows NT"*, Document Id# 32645-01-1150
- *"Great OLTP Performance in a Mixed OLTP/DSS Environment"*, Sybase Technical White Paper, March 1998

Resource Limits

Sybase ASE supports the definition of resource limits to restrict I/O costs, elapsed time and resultant set size of a transaction or query. These restrictions can be defined and associated with an access time range, login id, or application. By properly defining the system resource limits along with engine groups and execution classes the System Administrator can balance the workload among the various OLTP, DSS and other processes executing on the dataserver. Please refer to the section on "*Limiting Access to Server Resources*" in the *Sybase Adaptive Server Enterprise System Administration Guide* for a more through discussion of implementing and managing resource limits.

Memory Planning

The objective of this section is to provide you with a starting point for system base memory and to introduce you to some important considerations regarding the system memory on Compaq servers.

Initial Memory

The *Sybase Installation Guide for Microsoft Windows NT* states the minimum system memory requirement for running Sybase is 32 megabytes[4], which includes the memory requirements for Windows NT.

Some general guidelines to use to calculate starting memory values are shown in the following formula. Refer to the *Sybase System Administration Guide* for additional memory requirements and calculations.

Sybase memory =

5 megabytes for database kernel and data structures
\+ (2% * total data and index space)
\+ (51 kilobytes * number of users)

[4] For satisfactory performance, do not run Sybase on a system with less than 32 megabytes of memory.

The 51 kilobytes for each user includes 10 kilobytes for procedures. Make appropriate calculations to determine the percentage value of procedure cache to total memory using the result for the "procedure cache" value of sp_configure.

Reducing Memory Access Overhead

Sybase ASE for NT upon completion of a task runs an additional task to cleanup the memory area and check for possible memory errors before releasing the memory for use by the next Sybase task. Use the *dbcc tune cleanup* option to disable this task and reduce the processing overhead of cleaning the memory before releasing it. Refer to the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for additional information.

User Stack Size

Sybase ASE 11.5.1 for Windows NT has changed the manner in which the user stack area is allocated, managed and configured. For Sybase ASE 11.5.1 the *sp_configure stack size* and *stack guard size* settings have no effect on the system under Windows NT. For Sybase versions prior to ASE 11.5.1 the *sp_configure* configuration settings for *stack size* and *stack guard size* do effect the system under Windows NT.

The Sybase user stack is now allocated by and managed by Windows NT. To reconfigure the Sybase ASE binary to increase or decrease the user stack size, use the Microsoft Utility program *imagecfg.exe* with the *–k* option to alter the c:\sybase\bin\sqlsrvr.exe file. The default stack size is 0x20000 or 131072 bytes. To decrease it to 0x10000 or 65536 bytes execute the following command.

Imagecfg –k 0x10000 sqlsrvr.exe
sqlsrvr.exe contains the following configuration information:
Subsystem Version of 4.0
Image can handle large (>2GB) addresses
Stack Reserve Size: 0x10000
Stack Commit Size: 0x4

The Windows NT editbin utility can also be utilized to alter the Sybase ASE NT user stack size. The syntax for this command would be:

Editbin /stack 65536,4 sqlsrvr.exe

Greater than 2GB User Memory Support

By default, Microsoft Windows NT Server 4.0 and Server 4.0 Enterprise Edition permit only 2GB of the installed memory to be allocated as user memory. With the advent of Microsoft Windows NT Server 4.0 Enterprise, user applications like Sybase ASE can utilize up to 3GB of memory. The remaining 1GB of a 4GB system is reserved as system memory for use by the Windows NT system and is unavailable as user memory.

After installing Microsoft Windows NT Server 4.0 Enterprise Edition you will see the message as shown in Figure 1. To save you the time and trouble of looking up the procedure for modifying your system to utilize the 4GB of memory, we have outlined it below.

Figure 1: Microsoft Windows NT Server, Enterprise Edition 4GB Support

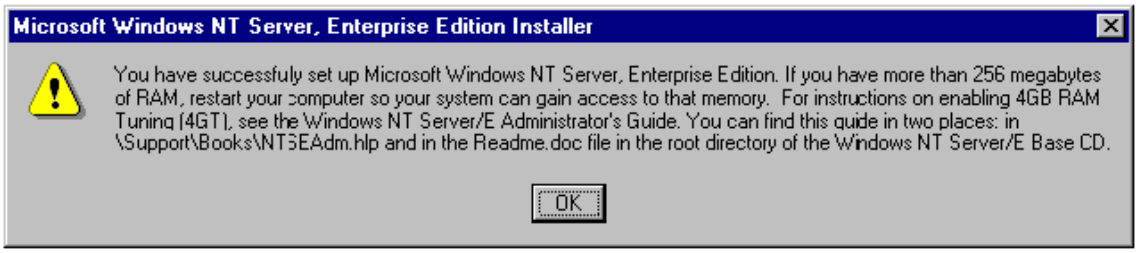

To enable the 3GB user memory feature you will need to edit the system hidden file c:\boot.ini and add the command line option /3GB to the boot up string and reboot your system.

[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(2)\WINNT
[operating systems]
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server, Enterprise Edition Version 4.00 [3GB user]" /3GB
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server, Enterprise Edition Version 4.00"
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT Server, Enterprise Edition Version 4.00 [VGA mode]" /basevideo /sos

Modifying Sybase ASE 11.5 for 3GB Support

Sybase ASE 11.5.1 is shipped pre-configured, out of the box, to support the full 3GB of user memory available under Windows NT 4.0 Enterprise Edition. Therefore Sybase ASE 11.5.1 does not need this modification to utilize memory above the 2GB user memory limit.

If you are running with Sybase ASE 11.5 on Windows NT Server 4.0 Enterprise Edition and want to utilize the memory above the default 2GB user memory limit, you will need to reconfigure the Sybase ASE 11.5 binary. Use the Microsoft Utility program *imagecfg.exe* with the *–l* option to alter the c:\sybase\bin\sqlsrvr.exe file. The Sybase TechNote *Addressable Memory Limits in Adaptive Server Enterprise 11.5*, Document ID: 20101, March, 25, 1998 correctly states that Windows NT 3.51 and Windows NT 4.0 have an addressable memory limit of 2GB for user based applications like Sybase ASE 11.5. Modifying the Sybase ASE 11.5 binary with this procedure will permit it to address memory above the 2GB limit only when running under Windows NT Server 4.0 Enterprise Edition.

Imagecfg –l sqlsrvr.exe
sqlsrvr.exe contains the following configuration information:
Subsystem Version of 4.0
Image can handle large (>2GB) addresses
Stack Reserve Size: 0x10000
Stack Commit Size: 0x4

Chart 4 shows the increase in throughput when increasing the memory available for Sybase on a ProLiant 5000 4GB system from 1GB to 2GB and upwards to 3GB.

Chart 4: PentiumPro Memory Scalability Under Windows NT Server 4.0 Enterprise

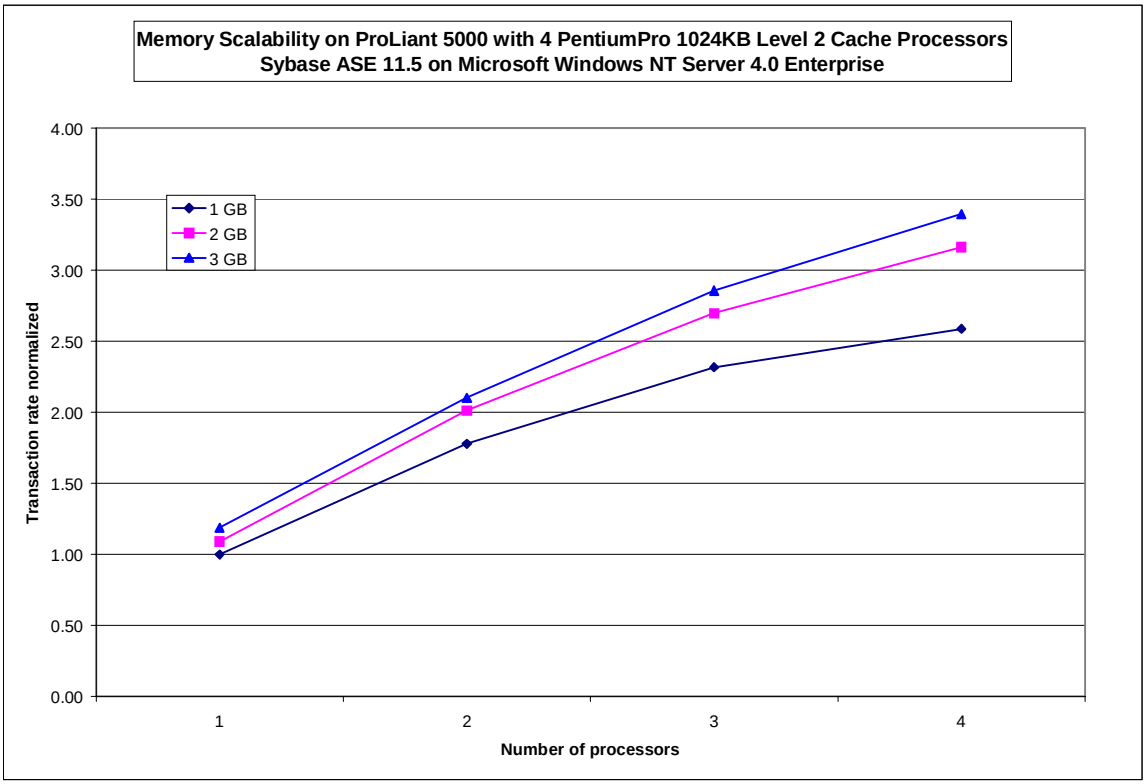

Memory requirements are environment-specific. In the "Memory Utilization, Monitoring, and Optimization" section found later in this document, we provide suggestions on how to detect and correct low memory situations.

Virtual Memory

During the initial installation of Windows NT, a paging file is setup automatically on the default boot drive with a size of physical memory + approximately 64 megabytes. In a properly configured and tuned Sybase environment, no paging should occur after the initial startup of the dataserver. The size of the paging file will need to be altered if you increase or decrease the physical memory in the system.

Sybase ASE uses the Windows NT paging file for generation of a memory mapped file during startup of the dataserver. After the database is online there should be little or no I/O activity to the Windows NT paging file. The size of the NT paging file must be greater than the Sybase total memory setting.

If you have over committed the virtual memory allocated on your system the following warning message will appear. When this happens the performance of your system is severely impaired.

Figure 2: Out of Virtual Memory

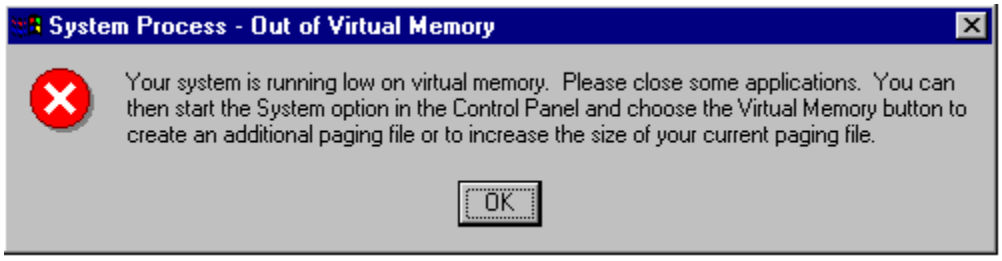

Windows NT Memory Management Setting

To help reduce excessive paging, check the Windows NT Registry setting for *LargeSystemCache* under *HKEY_LOCAL_MACHINE, System, CurrentControlSet, Control, Session Manager, Memory Management.* This value should be set to zero to instruct Windows NT to favor the process working set over the system cache.

Refer to the *Microsoft Windows NT Server, Concepts and Planning Guide* for a more in-depth discussion of the paging file.

Disk Subsystem Planning

The objective of this section is to provide information on the benefits of the Compaq SMART and SMART-2 Array Controllers. The pros and cons of various fault tolerance methods. The selection of the file system. How disk space is consumed by the combination of Sybase and Windows NT system level files.

Additional information on disk subsystem configuration and Compaq drive array technology can be found:

- *Configuring Compaq RAID Technology for Database Servers*, White Paper, 2nd edition, May 1998, Document Number ECG011/0598
- *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide*, Document Id# 32645-01-1150-02
- *Configuring Sybase® Adaptive Server™ Enterprise for Windows NT*, Document Id# 32645-01-1150 Part Number AA0539
- *Sybase® Adaptive Server™ Enterprise System Administration Guide*, Document Id# 32645-01-1150-02
- *Sybase® Adaptive Server™ Enterprise Security Administration Guide*, Document Id# 32645-01-1150-02
- *Microsoft Windows NT Server Resource Kit*
- *Microsoft Windows NT Server, System Guide*
- *Microsoft Windows NT Server, Concepts and Planning Guide*

RAID Performance Comparison

The throughput of your system does not have to be limited by the level of fault tolerance you choose. Chart 5 shows the measured throughput of Sybase ASE 11.5 on Microsoft Windows NT Server 4.0 Enterprise running an OLTP warehouse and order entry application on a Compaq ProLiant 5000, 4 PentiumPro/200MHz processor, 4GB memory system. The number of SMART-2/P Array controllers and 4.3GB hard drives remained constant. The log drives were left configured as a single RAID1 logical drive. The Array Accelerator for the Log drives was disabled (off) and for the database data drives either disabled (off) or enabled (on) using cache settings of 100% write, 50% read / 50% write or 100% read.

Chart 5: RAID Levels and Array Accelerator Benefits

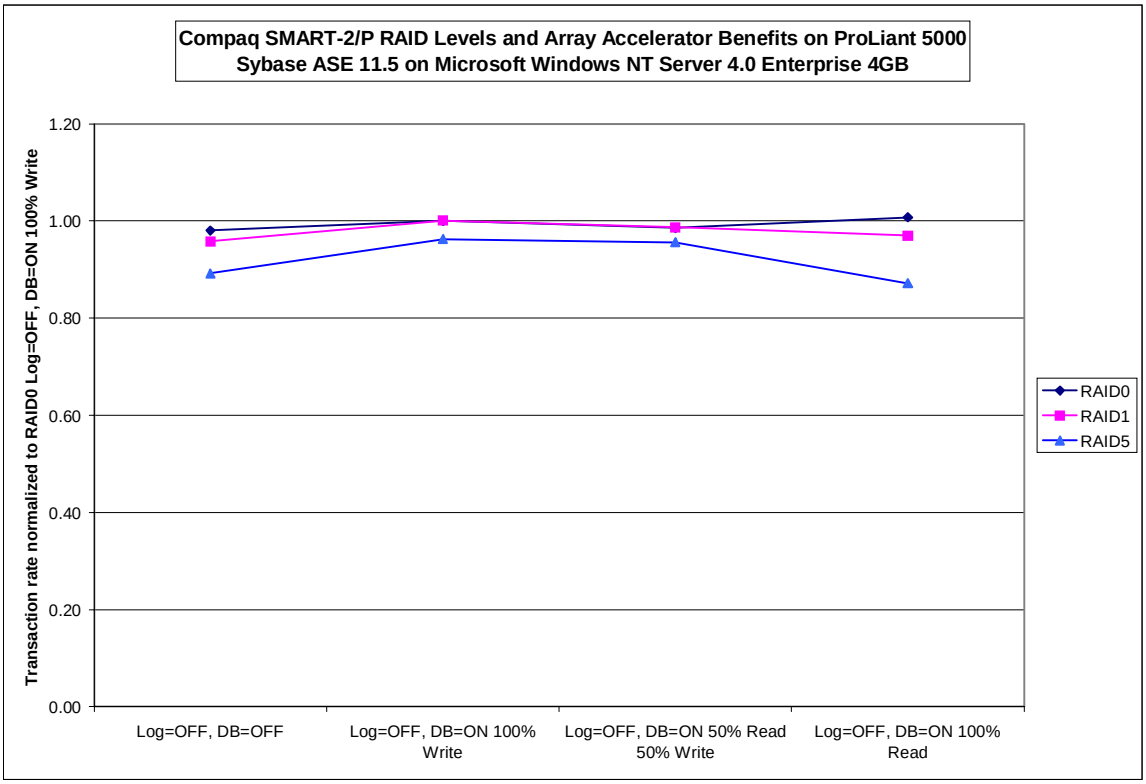

The OLTP application profile has a read to write ratio of approximately 3:1. The measured number of I/Os per second per drive, at the application level, is between 30 and 35. The Compaq SMART-2/P Array controller may perform additional I/Os depending upon the RAID level settings. The total number of I/Os performed by the SMART-2/P Array controller for the different RAID levels are given in Table 1.

Table 1: Application I/O Requests Mapped to Controller RAID Level I/Os

	Controller Physical Reads	Controller Physical Writes
RAID0 – application physical read	1	0
RAID0 – application physical write	0	1
RAID1 – application physical read	1	0
RAID1 – application physical write	0	1 primary drive, 1 mirror drive
RAID5 – application physical read	1	0
RAID5 – application physical write	1 data, 1 parity	1 data, 1 parity

The measured application I/Os per drive is between 30 and 35 I/Os per second. For the purposes of the following calculations we will use 32 I/Os per second per drive.

RAID0 – 24 reads 8 writes = 32 physical I/Os per second per drive

RAID1 – 24 reads and 16 writes = 40 physical I/Os per second per drive

RAID5 – 40 reads and 16 writes = 56 physical I/Os per second per drive

From Table 4 listing drive I/O capacities, the 4.3GB drives support 55-60 random I/Os per second. For all 3 configurations the drives are operating within specifications. The RAID0 and RAID1 configurations have ample additional I/O bandwidth to handle future increased workloads. The RAID5 configuration is already at its peak performance level as evidenced from the data in the chart and calculations. To increase the bandwidth in the RAID5 configuration you will need to either upgrade the drives to higher throughput drives or add an additional SMART-2/P Array controller and 4.3GB drives.

The choice of configuring a system to use RAID0, RAID1 or RAID5 can effect the overall performance of the system if you exceed the stated I/O capacities listed in Table 4. There is little or no degradation of system throughput when performing a read from a data table, clustered index or non-clustered index on a RAID1 or RAID5 drive volume when compared to a RAID0 drive volume. The potential for sever system throughput degradation can occur when performing writes to a RAID5 drive volume as in the case of a BCP load, the creation of a clustered or non-clustered indexes, or a system that performs a lot of data updates and inserts to tables having multiple indexes. By careful evaluation of the database scheme and proper configuration of the Sybase ASE table partitioning and parallel features one can utilize RAID 1 and RAID5 drive volumes and maintain I/O throughput levels comparable to a system configured using RAID0. The test database used to collect the results presented Table 4 just happens to be a good example of this.

Array Accelerator: Its Function and Benefit in a Sybase Environment

The main function of the Array Accelerator is to increase the throughput of all I/O operations by storing data in cache memory on the controller. In the case of a write, the drive controller informs the host immediately that the write operation has completed, and writes the data to the hard drives at a more convenient time. In the case of sequential reads, the controller performs the read of the requested sector and pre-reads the next couple of sectors into the Array Accelerator cache in anticipation that the data for the next read request will be available in the cache thus reducing the likelihood of physically performing another read. The onboard rechargeable battery pack and fault tolerant memory banks guarantee data integrity at all times.

Without the Array Accelerator, the application must wait until each write request is written out to the disk. Writing to a disk device can be slower than posting the write request in the Array Accelerator, thus resulting in very different performance characteristics.

Checkpoints, Housekeeper and Transaction Log Writes

There are three main write-intensive operations Sybase performs: checkpoints, housekeeper and transaction log writes.

- During **checkpoints**, Sybase generates a large number of write requests in a short time interval. The main objective of the checkpoint is to write **all** dirty pages from the data cache to the disk in the shortest time possible by simultaneously writing as many pages as possible. This technique requires careful tuning considerations for the Compaq SMART and SMART-2 Array Controllers.

  In some environments, the amount of write activity that the checkpoint generates can saturate the Array Accelerator, thus interfering with read requests pending at the controller. Proper tuning of both the checkpoint and housekeeper processes can alleviate this problem. To tune the checkpoint process use the *dbcc tune maxwritedes* option. Refer to the I/O Pacing section later in this document for more details. Also refer to the *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide* for a more information on I/O Pacing.

- The **housekeeper**, if enabled, waits until it detects idle time in the Sybase dataserver then begins to write dirty pages from the data cache to the disk at a lower priority than the

checkpoint process. The main objective of the housekeeper is to write as many dirty pages from the data cache to the disk in the shortest time possible with minimal impact on the throughput of the database application. Unlike a checkpoint process, which must write all dirty pages from the data cache to disk before terminating, the housekeeper writes only what it can during idle times. If the system is idle for a long enough period of time the housekeeper may actually have written all dirty pages from the data cache to disk. When this occurs the housekeeper notifies the checkpoint process and requests that a checkpoint be performed on the database.

To tune the housekeeper process use the *housekeeper free write percent* and *dbcc tune deviochar* option. Refer to the *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide* for more detailed information.

- The **transaction log** activity is composed exclusively of sequential writes and does not saturate the Array Accelerator. However, the difference between writing to the Array Accelerator and writing to the disk drives can be substantial. For optimal performance the Array Accelerator should be configured for 100% write and enabled.

The Array Accelerator improves performance of all write activity, including the transaction log writes, the housekeeper writes and the checkpoint writes.

Fault Tolerance Considerations

Your transaction log needs to always be protected against a disk failure. Protecting your database devices is at your own discretion. Most mission critical sites protect both the transaction log and the database devices and they usually choose hardware-based fault tolerance. Non-mission critical sites are often satisfied with the protection of the transaction log only and do frequent backups/dumps.

You have three choices of protecting your data:

- Use a RAID-based disk controller such as the Compaq SMART and SMART-2 Array Controllers
- Use Windows NT-based RAID
- Use Sybase-based mirroring/duplexing.

Keep in mind that in an OLTP environment the transaction log activity is characterized by sequential writes, whereas the database activity is composed of random reads and writes. In a DSS environment the transaction log activity is characterized by sequential writes, whereas the database activity is composed mostly of sequential read with few random reads and writes.

Below are some key points to be aware of when choosing the appropriate fault tolerant method. The performance differences between fault tolerance implementations can vary from insignificant to very significant, depending on your configuration and environment. We recommend that you evaluate the various fault tolerance methods using your own applications.

1. RAID-based Disk Controllers

- Hardware RAID is configurable on a logical volume basis[5]. Therefore, the whole logical volume is protected by the appropriate fault tolerance. The capacity reduction depends on the size of the logical volume and the type of fault tolerance configured.
- Hardware RAID offers the best reliability and automatic recovery. When a drive fails, the system continues operating. Once the failed drive is replaced, the controller automatically rebuilds the new drive and restores the system to its full fault tolerant capabilities.

[5]Logical volumes for the Compaq SMART SCSI Array controller are defined using Compaq System Configuration Utility. Logical volumes for the Compaq SMART-2 Array controller are defined using Compaq Array Configuration Utility

- Since the data protection occurs at the hardware (controller) level, there is no overhead on the system processor. This type of RAID is totally transparent to the operating system and the applications. With a CPU-intensive application/environment, such as Sybase, hardware-based fault tolerance can provide the best performance. Please refer to *Configuring Compaq RAID Technology for Database Servers, 2nd Edition*, May 1998, Document Number ECG011/0598 published by Compaq for a more complete discussion.

2. Windows NT-based RAID

- To take advantage of Windows NT-based fault tolerance, you need to have Windows NT Server installed. The current versions of Windows NT Workstation do not provide software based RAID.
- You have an option of controller duplexing, if you place the mirrored partition on a different disk controller. Controller duplexing adds redundancy at the controller level, but you must make sure that your mirrored partitions are on physically different drive controllers.
- Windows NT fault tolerance is based on logical drives. The capacity reduction depends on the size of the logical drive and the type of fault tolerance installed.
- Windows NT fault tolerance offers reliability in a sense that you will not lose your data when your disk crashes. Your system will continue running after a drive failure. However, the recovery procedure is not automatic, and you will need to bring your system down.
- The data protection occurs at the operating system level, which induces an additional overhead on the system processor. Consequently, in a CPU-intensive environment, Windows NT fault tolerance will impact performance. Similar to hardware-based fault tolerance, mirroring is preferred to striping with parity.

3. Sybase Mirroring/Duplexing

- Sybase mirroring is based on Sybase devices. This offers you the finest granularity and lowest capacity reduction due to duplicated data. Refer to *Sybase, System Administration Guide* for guidelines on choosing which devices to mirror.
- You have an option of controller duplexing, if you place the mirrored device on a different disk controller.
- Sybase supports both serial and noserial write mode of mirroring. When utilizing the default, serial write mode, writes to the first device must finish before writes to the second device begin. Changing from serial to noserial write mode allows the writes to both devices to be queued immediately, one to each mirrored device. The noserial write mode will incur less overhead than serial writes that results in a slight performance improvement.
- Reliability and recovery procedures are very similar to those of Windows NT fault tolerance. Your system remains running, but you will have to take the system off-line and upon a failed drive replacement you will have to manually re-mirror the devices.
- Again, mirroring through Sybase induces an additional processing overhead on the system, resulting in a lower performance.

For further information on Compaq RAID technology with Sybase, refer to the following document:

Configuring Compaq RAID Technology for Database Servers, 2nd Edition, May 1998, Document Number ECG011/0598

Support for Devices Over 2GB in Size

Sybase ASE for Windows NT supports database and backup devices up to 32GB in size. Be forewarned that the disk init command will accept values larger than 32GB without generating an error message. The operation of and the data integrity of your database will be compromised if you attempt to use devices larger than 32GB. Sybase ASE can support 127 database devices of 32GB each bringing the total database size to just under 4 terabytes. The Sybase Backup Server supports a maximum of 32 striped backup devices of 32GB each for a total of 1024GB or 1 terabyte per volume. When the first volume of 32 backup devices are full, the Backup Server will request a volume change to mount the next set of 32 backup devices. Please refer to the Sybase TechNote *Device Size Limits for Backup Server, SQL Server11.0.x and Adaptive Server Enterprise 11.5*, Document ID: 800, May 7, 1998 for more details.

FAT vs. NTFS vs. Direct Disk Access

Direct disk access, a.k.a. raw I/O for Windows NT, does have a slight performance edge over NTFS and FAT file systems. Sybase ASE 11.5.1 greatly improved the I/O throughput to NTFS disk volumes, so the performance advantage of utilizing raw I/O is probably only necessary where you need to maintain high I/O throughput levels. Testing on Compaq servers showed only a slight performance differences between FAT and NTFS file systems in a Sybase environment. While requiring a little bit more time in disk planning and management, direct disk access under NT will give you the best performance. If ease of disk management and data recoverability are a priority then, it is recommended that NTFS be used for both transaction log and database files.

Direct disk access is available only for user defined devices under control of Sybase SQL Server, Backup Server and ASE 11.5. The master database device must remain on either a FAT or NTFS file system. The tempdb database can be extracted from the master device onto its own direct disk access device for improved performance. Refer to the section on *Tempdb Placement* in this document for more details. Other Sybase products such as SQL Remote, Replication Server and *OmniConnect™* can communicate with a Sybase ASE dataserver configured with direct disk access devices via the Open Server/Open Client interface but at this time do not support direct disk access devices for themselves.

Be aware that the Windows NT C2-level security features are only available on NTFS volumes. A complete list of advantages and disadvantages of FAT and NTFS file systems is available in the *Microsoft Windows NT, Concepts and Planning Guide*.

Third party tools such as MKS Toolkit[6] for Windows NT provide UNIX-style commands for the Windows NT environment. The MKS Toolkit for Windows NT does support access to the direct disk access devices with commands like dd. The dd command can be useful if you want to do perform an image copy of the complete database installed on direct disk access devices from a controlled test environment system to a production environment system or vica versa. It also can be used to perform an image level backup of the entire database.

Figure 3 and the following description describe “how-to” use Windows NT direct disk access with Sybase ASE. The first time Disk Administrator is run permit it to write the NT disk signature on the disk devices. Leave the disk devices that are to used by Sybase ASE as unformatted disk devices. From the Disk Administrator note the actual disk numbers for the devices attached to your system. In this example the boot OS drive is located on a FAT formatted partition and the NT Paging File is located on a NTFS formatted partition both of which are on Disk 23 and not on Disk 0 as one would have assumed. Sybase ASE requires that the master database device be on a FAT or NTFS formatted partition therefore Disk 24 is formatted as a NTFS partition with master.dat located on it.

[6] The MKS Toolkit for Windows NT is available from Mortice Kern Systems, Inc., please refer to the Web Site http://www.mks.com for more detailed information on this product.

Figure 3: Windows NT Disk Administrator

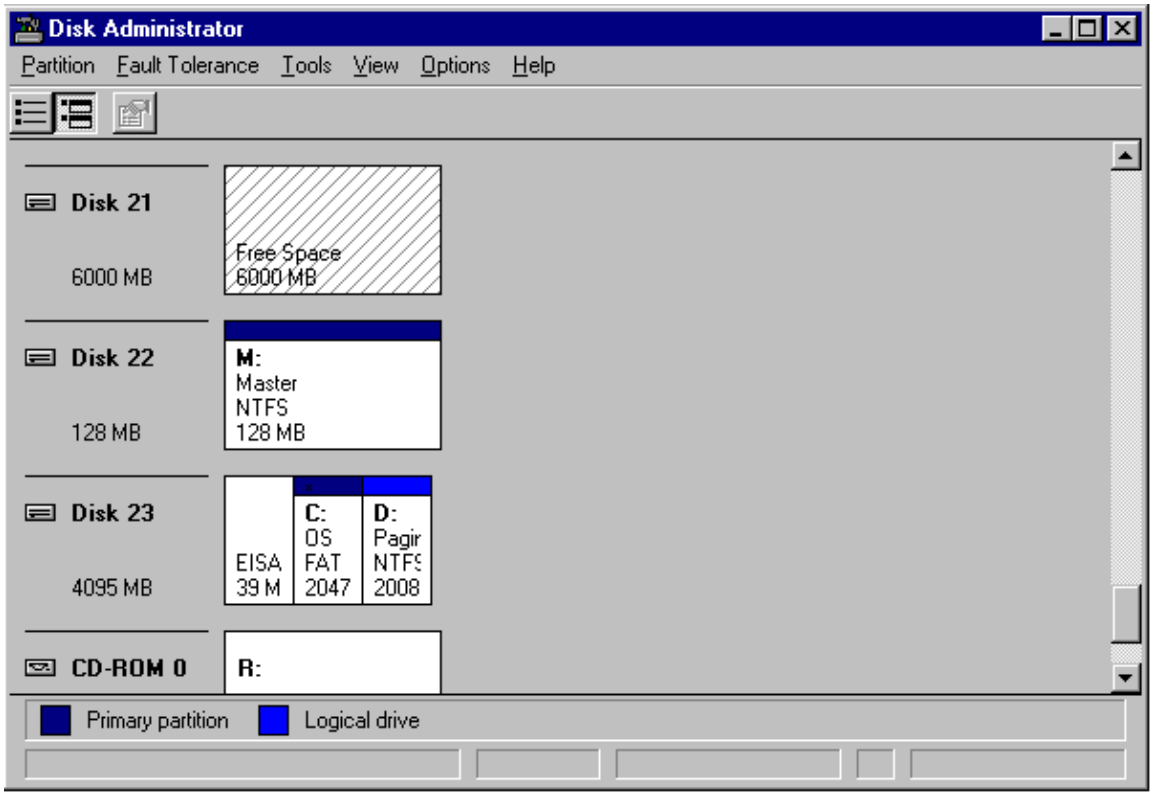

To have Sybase ASE create a virtual device of 6000MB on Disk 21 issue the following SQL commands:

```
disk init
  name = "customer_01",
  physname = "\\.\PhysicalDrive21",
  vdevno = 15,
  vstart = 1,
  size =3072000
```

The *vstart=1* command is important, this instructs the Sybase disk initialization routine to skip over the first page, 2KB, of the disk device. Microsoft Windows NT Disk Administrator places its own disk signature in the first 512 bytes. Without the *vstart* command the potential exists for one or more of the Windows NT or 3rd party utilities to accidentally overwrite the first page of your database device. One item to consider that makes system management of physical drives more difficult is that under Windows NT the physical disk numbering of the drive volumes will change if you add or remove hard disk drives and/or disk controllers. Please refer to the following documents for more information:

- *"Sybase® Adaptive Server™Enterprise 11.5 for Microsoft Windows NT"*, Sybase Server Journal, Second Quarter 1997 issue
- "*Physical Device I/O on Sybase SQL Server and Adaptive Server for Microsoft Windows NT*", Sybase White Paper, February 1998, Document Number 20172

- *"Device Size Limits for Backup Server, SQL Server 11.0.x and Adaptive Server Enterprise 11.5"*, Sybase TechNote[7], Document ID: 800, May 7, 1998

NT Disk Device Drivers

Compaq provides two different implementations of the NT disk device driver that interfaces to the SMART-2 Array controllers. The standard *cpqarray.sys* driver that ships with Microsoft Windows NT and the *cpqsmrt2.sys* driver. The latest version of the *cpqarray.sys* driver is available from the Compaq NT SSD. The latest version of the *cpqsmrt2.sys* driver is available from the Compaq Web site as a SoftPaq downloadable file.

The *cpqsmrt2.sys* device driver is a performance optimized device driver using the enhanced command interface (ECI) to communicate between Windows NT and the SMART-2 Array controller. The ECI driver reduces the processor overhead required to perform the I/O operations. This results in increased throughput of the system, especially when both the processor and disk sub-systems are heavily utilized. A more in-depth discussion of the differences between the two device drivers is described in the *readme.txt* file on the SoftPaq diskette. The Compaq SMART-2 Array Controller ECI Driver v1.04 was used for the testing performed in this document and is downloadable as SoftPaq number SP4005.exe from the Compaq Web site.

Large I/O Sizes

Configuring Sybase ASE to use large I/O sizes can provide you with a reduction in the number of physical I/O operations performed resulting in increased system throughput. Sybase ASE supports I/O sizes of 2K, 4K, 8K and 16K. To have a database, table or index use both 2K and 16K I/O perform the following tasks, in sequence. Configure a dedicated named cache with both a 2K and 16K I/O Buffer Pool. Bind the cache to the database object. Then use the *dbcc tune iosize* option to enable the large I/O. The setting needs to be reset after each reboot of the database server. The Sybase optimizer will now automatically determine if and when the large I/O is the most efficient method of performing the I/O. Use the Sybase Monitor Server and *sp_sysmon* utility to monitor the effects of enabling large I/O. Refer to the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for more details on using large I/O and its effect on your system.

Tempdb Placement

The size and placement of tempdb database may drastically affect the overall performance of your system. Your choices are leave the tempdb database in master device, place it into a dedicated named cache or move it out of master device to its own device. Obviously the size of the tempdb database, the type of I/O, and the I/O activity rate will determine which choice is best for your environment.

For most OLTP systems the tempdb database is relatively small, less than 100 MB in size. The I/O activity is relatively high but is mostly composed of writes. The first choice for improving the throughput of these systems it is to assign a dedicated named cache of a size equal to or slightly larger than the tempdb database itself. Then configure and enable both 2K and 16K I/O buffers to it. Monitor the throughput of your system using the *sp_sysmon* utility. If the reallocation of Sybase data buffers from other named caches results in a loss of performance then reduce the size of the tempdb dedicated named cache to between 2 and 5 % of the actual size of the tempdb database. Redistribute the released tempdb data buffers to other named caches. Move the tempdb database on a separate database device using NT direct disk access, configured on a SMART-2 Array controller using RAID1 for fault tolerance. Refer to the *"Sybase® Adaptive Server™Enterprise 11.5 for Microsoft Windows NT"*, Sybase Server Journal, Second Quarter 1997

[7] The Sybase TechNote *Device Size Limits for Backup Server, SQL Server 11.0.x and Adaptive Server Enterprise 11.5*, Document ID: 800, May 7, 1998 incorrectly states that Windows NT does not support raw I/O. Sybase has supported direct disk access, raw I/O, for Windows NT since the release of Sybase SQL Server 11.0.2 for Windows NT.

issue for detailed information on "how-to" move tempdb out of the master device and onto a separate Sybase database device. Additional information maybe found in the document *Maximizing Use of tempdb on a Non-Master Device,* Sybase TechNote, October 30, 1996, Document ID: 506.

On the other hand DSS systems may have a tempdb database of 512 MB or larger in size. Random reads and writes characterize the I/O activity to it. Assign a dedicated named cache of approximately 5% of the tempdb database size. Configure and enable both 2K and 16K I/O buffers pools for the tempdb named cache. Place the tempdb database on a separate database device using NT direct disk access, configured on a SMART-2 Array controller using RAID1 for fault tolerance.

In either configuration use the *sp_sysmon* utility to monitor the overall effectiveness of these changes and modify them to provide the best overall throughput for your system. Refer to the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for a more detailed discussion of these features and there impact on your system.

Asynchronous Prefetch

The Sybase Logical Memory Manager (LMM) allows the configuration of dedicated named data caches for various database objects. Asynchronous prefetch (large I/O) is enabled when defining a dedicated named cache for a table, index, text or image object. When a read command to the database is issued, the optimizer can automatically request up to the number of pages specified by the *async prefetch limit* value. These pages are then read from the database and placed into the dedicated named cache. This will greatly improve the performance of OLTP transactions and DSS queries that perform table and index scans. Prior to Sybase ASE 11.5.1, the number of pages the could be pre-fetched was limited by the implicit prefetch value of 8 pages.

The *sp_configure global async prefetch limit* command determines the percentage of the dedicated named cache pages to use for asynchronous prefetch reads. This setting is for all dedicated named caches. To adjust this on an individual dedicated named cache basis use the *sp_poolconfig* command to configure to each I/O size buffer pool within the dedicated named cache.

Use the *sp_cachestrategy* command to verify the current setting for asynchronous prefetch (large I/O) and also to enable or disable it.

Sp_cachestrategy db_contacts, t_zipcode

Object name	index name	large I/O	MRU
Dbo.t_zipcode	idx_zip	ON	ON

Do not confuse the asynchronous prefetch definition of large I/O with that of configuring large I/O sizes for the 2K, 4K, 8K and 16K dedicated named cache buffer pools. Refer to the section on *Large I/O* Sizes in this document for information on using I/O sizes greater than 2K.

In the case of a relatively small table or index that is read-only or can fit entirely in a dedicated named cache, the throughput of the system maybe increased by disabling asynchronous prefetch. Use *sp_cachestrategy … prefetch, "off"* command to disable asynchronous prefetch for a particular database object. Asynchronous prefetch can be disabled on a global scale by using the *set prefetch "off"* command.

Refer to the *Sybase® Adaptive Server™Enterprise Reference Manual, Sybase® Adaptive Server™Enterprise System Administration Guide* and the *Sybase Adaptive Serve Enterprise Performance and Tuning Guide* for a more detail discussion on configuring and tuning of asynchronous prefetch and buffer pools.

I/O Pacing

Sybase ASE 11.5 introduces the ability to configure the checkpoint and housekeeper processes to limit the number of additional or outstanding I/Os generated. The checkpoint process can be configured by using the *dbcc tune maxwritedes* option. The housekeeper process can be configured by using the sp_configure *housekeeper free write percent* option and the *dbcc tune deviochar* option.

Some of the reasons why you would want to adjust the checkpoint and housekeeper processes are:

- To increase or decrease the average elapsed time for the checkpoint process to complete
- To increase or decrease the number of I/Os to a device to better match its I/O capabilities
- To reduce the semaphore contention during I/O for a particular device
- To minimize the effects of increased user response times during the checkpoint process

The System Administrator and Database Administrator should carefully consider the impact of altering these values on the overall throughput of the system. If you configure the checkpoint process, recovery interval and housekeeper process incorrectly you may end up with a system that is in a continual state of performing checkpoints. For example, on a heavily utilized OLTP system, if the *recovery interval* is set to 5 minutes, the *housekeeper write percent free* is set to 1, *dbcc tune deviochar* is set to 1 and *dbcc tune maxwritedes* is set to 1, the first checkpoint may not have completed when a second one is queued waiting to run.

Chart 6 shows the effects of setting the *recovery interval* to 1440 minutes (24 hours), disabling the housekeeper process and just varying the value for the *dbcc tune maxwritedes* option. The system was configured with sufficient I/O capacity so that the average user response time, system throughput and semaphore contention were not effected.

Chart 6: Dbcc Tune Maxwritedes Effect on Duration of Checkpoint

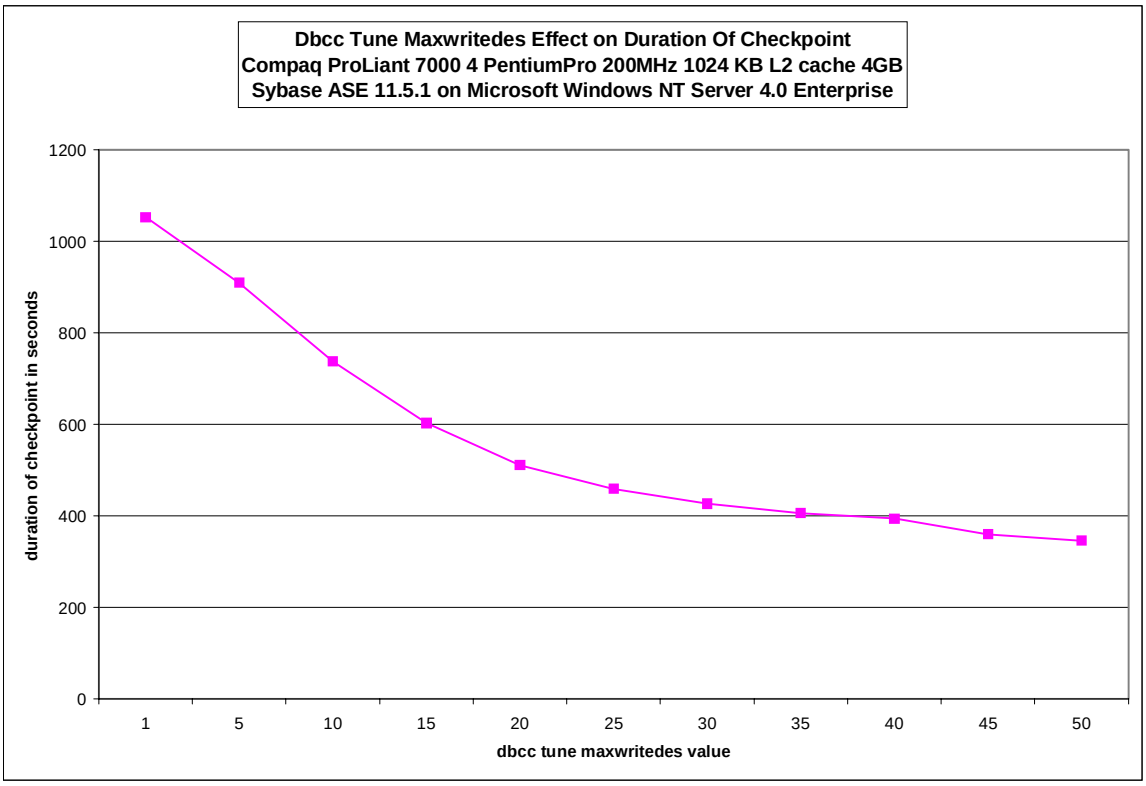

Compaq System Configuration Utility Considerations

Compaq System Configuration[8] utility allows you to select the operating system for each Compaq SMART and SMART-2 Array controller. This selection allows the configuration utility to choose the best physical characteristics, such as striping block size, number of sectors per track, etc., for a given environment.

Once you choose the *MS Windows NT* option as your selection for the operating system, you will be presented with the following options:

- MS Windows NT 4.0
- MS Windows NT 3.51
- MS Windows NT 3.5
- MS Windows NT 3.1

WARNING: Changing controller characteristics such as the operating system through the Compaq System Configuration utility causes the controller to apply different physical characteristics to the drive array. As a result, you may lose data on such a drive array. If you have existing data on such a drive array, and need to change these controller characteristics, we strongly suggest you back up your data completely before doing so.

Current partition limitations are as follows:

- Windows NT limits the boot partition to 4GB, regardless of the file system.

[8] Make sure you obtain the latest available version of Compaq System Configuration Utility from Compaq.

- Non-bootable partitions to be formatted with FAT file system are limited to 4GB. Non-bootable partitions to be formatted with NTFS file system are typically limited by the physical space.

Network Planning

The objective of the *Network Planning* section is to provide information on how the performance, optimization and interpretability of the network can influence performance of Sybase. The information in this section is being provided to give the reader information necessary for the evaluation of implementation alternatives. The results presented here are based upon network tests using both Windows NT and SCO UnixWare workstations running Sybase Open Client applications against Sybase ASE 11.5.1 on a Windows NT Server system. Therefore the information presented may not be indicative of other workstation/server combinations.

Network Characteristics of a Sybase Environment

On-line transaction processing (OLTP) environments typically generate a large number of small packets, and the amount of data being transmitted between the clients and the server is relatively small. On the other hand decision support systems (DSS) environments typically generate a large number of large packets, and the amount of data being transmitted between the clients and the server is relatively large. Many applications heavily utilize stored procedures, which not only decrease processing at the server, but also reduce network traffic. Stored procedures are invoked by passing the name of the stored procedure along with its parameters to the server. Even though stored procedures can return any amount of data, a typical stored procedure returns only a very small result set.

However, using stored procedures or not, the bulk of OLTP and DSS processing still occurs at the server, with the physical network layer (network interface controller or cables) seldom having a significant influence on the overall performance.

From the hardware perspective, you should have a 16- or 32-bit bus-master network card installed in the server to minimize the processing overhead associated with non-bus master network cards.

Check to insure that the Sybase system is not the server responsible for being the Default Gateway, Primary WINS Server, Secondary WINS Server, or Domain Name Service Server (DNS). Having these services enabled on the Sybase system will impact the overall performance of the server and increase the amount of network traffic to this server. When possible, move these services to another Windows NT server on the network.

Network Protocols: Packet Sizes

Sybase communicates with clients through *Tabular Data Stream* (TDS). TDS packets have a default size of 512 bytes. Applications can manipulate the packet size using the DB-Library dbsetlpacket() call; with BCP and ISQL utilities you can change the packet size using the *[-a packetsize]* parameter. The TDS packet size also needs to be incremented at the dataserver for the clients to connect properly to the server. From sp_configure alter the setting for *maximum network packet size* to match or exceed the maximum application packet size.

Utilizing larger TDS packet sizes also requires that you increase the memory requirements for the Sybase kernel. The sp_configure value for *additional netmem* controls the size of this kernel memory allocation and will need to be increased. Please refer to the *Sybase, Systems Administration Guide* for more guidelines on computing the amount of *additional netmem* required.

Chart 7 shows the for larger data transfers between the clients and the server, such as BCP loads and unloads, it may be beneficial to increase this packet size to a higher value (2048 bytes and higher). In this case the elapsed time for BCP to load 100,000,000 rows, with an average row length of 106 bytes, was reduced from over 21 minutes for the default packet size of 512 to approximately 15 minutes simply by increasing the *max network packet size*. Chart 7 also shows

that after a certain point there is no additional benefit to increasing the packet size, it only wastes memory due the corresponding increase of the *additional netmem* value. A packet size of 4096 or 8192 would be the best choice to provide a balance of increased throughput, reduced BCP load time, minimizing number of network packets transferred and efficient utilization of memory resources.

Chart 7: BCP Load of 10,000,000 Rows

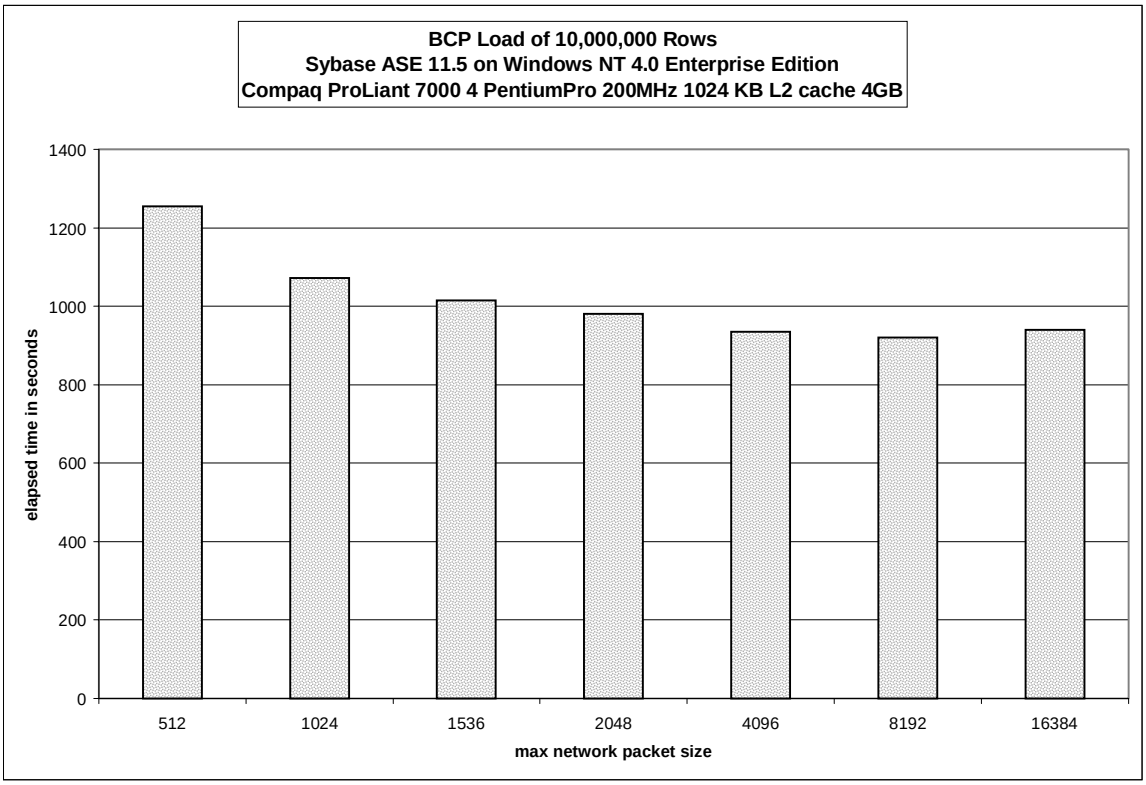

For the typical daily processing performed on an OLTP parts order entry system a TDS packet size of 4096 bytes provided the best throughput. In a DSS environment a TDS packet size of 16384 bytes provided the best throughput. Experimentation is necessary to determine the best packet size for your environment.

Reducing Network Overhead

Whenever a transaction or query completes Sybase ASE returns to the user the resultant data set plus some additional information. This additional information typically consists of the number of rows message. In systems with heavy network loads disabling the return of the additional information may increase the system throughput. To disable[9] the returned information use the *dbcc tune doneinproc* option. Refer to the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for more information.

[9] Disabling the dbcc tune doneinproc option should be performed for DB-Lib based applications ONLY. Some CT-Library based applications and ad hoc query generators rely on the return of the additional information to determine if the query was a success or failure, additional processing is needed or if an error has occurred. Otherwise the time the returned information is ignored.

Multiple Network Protocols

Windows NT and Sybase support multiple network protocols. To view the protocols utilized in your environment, use Control Panel, Network, Bindings, Server to list the protocol bindings that are configured on this Server. Take note of the protocol bindings listed. The Server will accept whatever protocol the Workstation/Client logs in with if it is listed here. Remove or disable any unused protocols to reduce the amount of system memory used by the network handlers. Adjusting the protocol binding order of the Server, Network Bindings has no effect on the server performance.

Within Sybase adjust the *sp_configure* parameter value for *max number network listeners* to match the number of network protocols used by the server.

If the same physical hardware platform that is running Sybase is also acting as a Workstation/Client running Sybase SQL Monitor and/or Sybase Server Manager check the Workstation Network Bindings. If more than one protocol binding is listed, highlight and promote the protocol that is being utilized by Sybase SQL Monitor and/or Sybase Server Manager to the top of the listing. This will reduce the chances for failed logins and network protocol time-outs on a heavily loaded system.

Refer to the following documents for more information on configuring, monitoring and optimizing the network for Sybase ASE.

- *Microsoft Windows NT Server, System Guide*
- *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide*, Document Id# 32645-01-1150-02
- *Configuring Sybase® Adaptive Server™ Enterprise for Windows NT*, Document Id# 32645-01-1150-02 Part Number AA0539
- *Sybase® Adaptive Server™ Enterprise System Administration Guide*, Document Id# 32645-01-1150-02
- *Installing Sybase® Adaptive Server™ Enterprise and OmniConnect™ on Windows NT* , Document Id# 34454-01-0200-02

Parallel Processing Features

The parallel processing features of Sybase ASE include BCP, index creation, backup, query, sorting, dbcc commands, some database maintenance utilities and recovery. In the following sections we will demonstrate the performance benefits when using the Sybase ASE parallel processing features. In each working environment the set of raw data, table structures and indexing requirements are different therefore the system throughput may or may not show similar benefits when using parallel processing.

The Sybase ASE parallel processing features can provide increased levels of throughput in single processor environments. By enabling the parallel query features you can have multiple parallel processes running, performing separate I/O tasks, thereby making more efficient use of the single processor and the disk system. Without parallel query capabilities the single processor system must execute the query in a serial fashion and maybe idle while awaiting for an I/O task to complete.

In the test environment the database had a space allocation of 3125248 KB for the table, indexes and log. After loading approximately 566 MB of raw data into the database table and generating the indexes the breakdown was a follows: 896028 KB data, 12 KB clustered index, 352832 KB non-clustered index and 42 MB used by the log. The data consisted of 10737418 rows with a fixed row length of 106 bytes. The clustered index consisted of a smallint field. The non-clustered index consisted of a char(25) field. The log consisted of 683884 rows and was approximately 42 MB in size after the initial data load and index creation was performed.

For a more in-depth discussion of the parallel processing features of Sybase ASE please refer to the section on *Parallel Query Concepts and Tuning, Chapters 13, 14 and 15* in the *Sybase Adaptive server Enterprise Performance and Tuning Guide*. For a discussion on the configuration of parallel features please refer to the section on *Parallel Queries* in *Chapter 11* of the *Sybase Adaptive Server Enterprise System Administration Guide*.

Parallel BCP and Table Partitions

To use parallel BCP, after creating your database tables and before loading any data, use the *alter table mytable partition* command to create the desired number of partitions to optimize the parallel loading of the data. A "rule of thumb" is to create one table partition for each disk drive in the table. In the case of the SMART-2 Array controller configure one table partition for each physical drive in the logical drive definition.

```
Partition table mytable 10
```

Then start multiple copies of the BCP utility, 1 BCP per table partition.

```
BCP parallel_db..mytable in mydataset1:1 –A 8192
…
BCP parallel_db..mytable in mydataset10:10 –A 8192
```

The BCP processes will load the data into the table in parallel. Each BCP process will use a different table partition to insert the new records. This eliminates the I/O contention between the competing BCP processes as they insert the new data records into the table.

Sybase has additional information about using Parallel BCP and Table Partitions in the following documents:

- *Using Parallel BCP to Reduce Partition Skew*, Sybase TechNote, March 24, 1998, Document ID: 20092
- *Tuning Sybase SQL Server and Adaptive Server Enterprise for Bulk Loads*, Sybase White Paper, March 24, 1998, Document ID: 1376
- *Improving BCP Performance by Increasing Extent Allocation*, Sybase TechNote, February 10, 1997, Document ID: 715
- *FAQs about Table Partitioning*, Sybase FAQ, November 17, 1997, Document ID: 1271

Chart 8 shows that the elapsed time to load 10737418 rows of data into a table with 14 partitions and 14 BCP processes is half the time it takes to the same data using 1 partition and 1 BCP process. The data to being loaded was read from one SMART-2/P controller RAID0 volume and written to another SMART-2/P controller RAID0 volume. Each RAID0 volume was striped at the hardware level across an array of 14 drives at 9.1GB each for a total volume size of over 127 GB. Each volume was formatted using NTFS. The Array Accelerator was configured as 100% Read cache on the source controller and configured as 100% Write cache on the target controller. Because each working environment, set of raw data, and table structures are different the throughput of using Parallel BCP vs. Serial BCP may vary.

Chart 8: Serial BCP vs. Parallel BCP

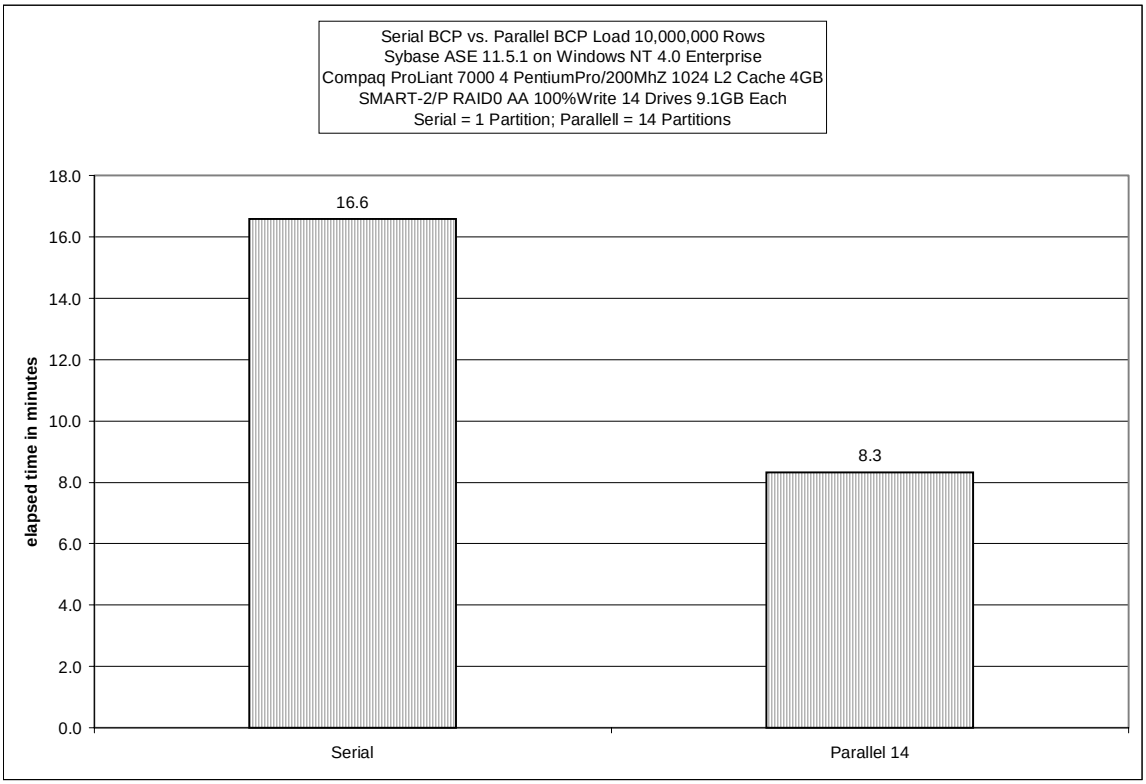

For an in-depth discussion on parallel BCP and table partitioning refer to the *Sybase Adaptive Server Enterprise Performance and Tuning Guide,* the *Sybase Adaptive Server Enterprise System Administration Guide* and the *Sybase Adaptive Server Utility Programs for Windows and Windows NT* document.

Parallel Backup

The Sybase Backup Server supports dumping of the database from disk to disk, disk to tape or disk to another Sybase Backup Server dump device over the network. The quickest and easiest method to perform throughput testing of the Backup Server is to use the disk to disk dump. The database is dumped to a local disk, then the output dump files can transferred to tape or over the network to another server for long term storage.

Sybase provides parallel backup features in the Backup Server when striping the dump or load across 2 or more backup dump devices. Each dump device causes the Backup Server to start another backup process. These backup processes operate in parallel to quickly and efficiently dump the database or transaction log to the dump devices. Each backup process can be viewed in the Performance Monitor as the process *sybmbuf*. The *dump database blocksize* for disk dump devices is fixed at 2048 bytes and can not be changed. During the dump of disk to disk the Backup Server may dynamically increase the *blocksize* from 2048 up to 16384 bytes. The *dump database blocksize* for tape devices can be configured from 2048 up to 65536 bytes for optimal data transfers. The *blocksize* used for the dump read and write activity can be viewed via the Performance Monitor, object Logical Disk, counters Avg. Disk Bytes/Read and Avg. Disk Bytes/Read. The number of I/Os performed per second can be viewed from the Performance Monitor, object Logical Disk, counters Disk Reads/sec and Disk Writes/sec. Remember to divide the Disk Reads/sec or Disk Writes/sec values by the number of drives in the array to get the number of reads or writes per second per disk drive.

For an in-depth discussion on parallel backup refer to the *Backup and Recovery* section of Chapter 23 *Maintenance Activities and Performance* in the *Sybase Adaptive Server Enterprise Performance and Tuning Guide* and the *Backup and Recovery* chapters in the *Sybase Adaptive Server Enterprise System Administration Guide.*

Parallel Dump

Chart 9 demonstrates the increase in throughput when increasing the number of dump devices[10] from 1 to 32 for the dump database command. The throughput increased from 7.6 MB per seconds for one dump device to a peak of 9.3 MB per second for 7 dump devices. The throughput at 32 devices was approximately 8.0 MB per second the same as with 2 dump devices. Overall the throughput averaged between 26 to 32.5 GB per hour for the Backup Server to dump the database.

The Backup was performed while dumping from one SMART-2 controller to another SMART-2 controller. Each controller was configured as one logical drive consisting of 14 physical drives striped using RAID0. No other processes or database activities were performed during the throughput measurement intervals on the system. The SMART-2 controller Array Accelerator was enabled for 100% Read cache on the source controller and enabled for 100% Write cache on the target controller.

The throughput of the Backup Server can be greatly influenced by the I/O rates of the database data and dump devices. During the dump process, typically the speed at which one can read the data from the database data device is not the limiting factor. Rather it is the speed at which the data can be written to the dump device. If the dump device has a greater write I/O capacity than the read I/O capacity of the database data device then the speed at which one can read the data from the database becomes the limiting factor. The I/O rates for reading data from the SMART-2 controller are not influenced by the RAID settings of the controller.

[10] In the test case, the 32 database dump devices were disk based NTFS files located on a single SMART-2DH Array controller. If they had been located on tape devices the throughput maybe lower than stated in this document.

Chart 9: Sybase Backup Server 11.5.1 Throughput

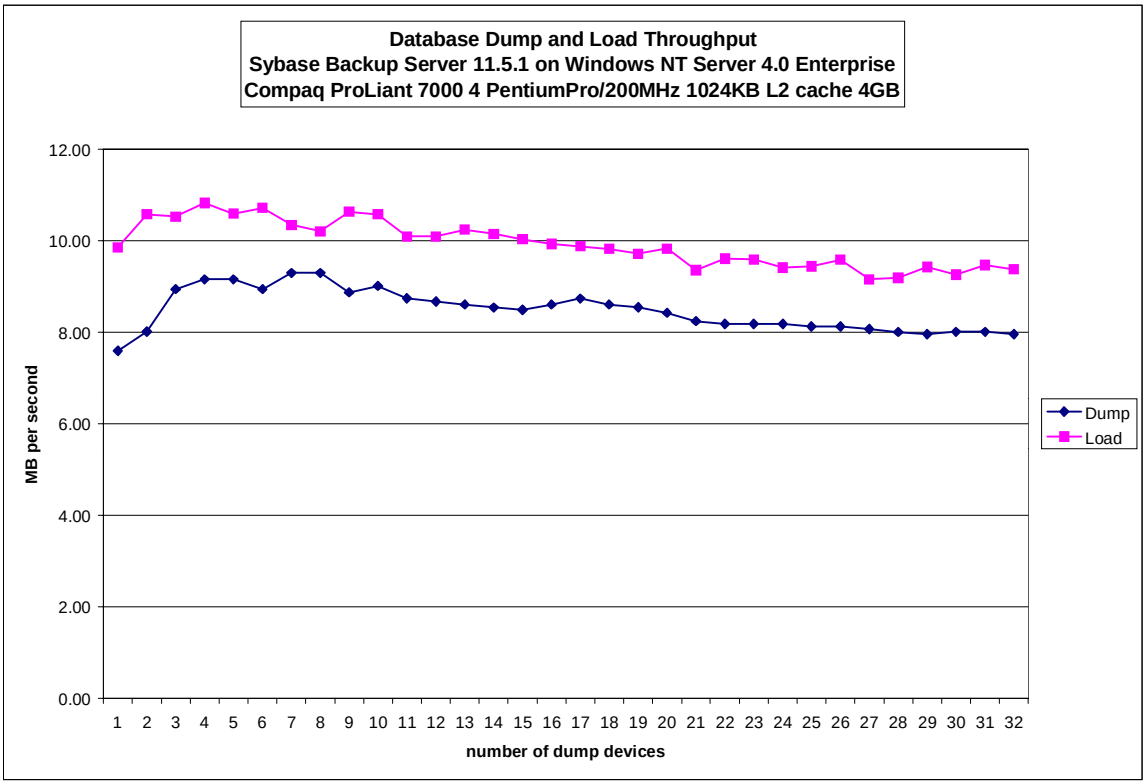

For an in-depth discussion on parallel backup refer to the *Backup and Recovery* section of Chapter 23 *Maintenance Activities and Performance* in the *Sybase Adaptive Server Enterprise Performance and Tuning Guide* and the *Backup and Recovery* chapters in the *Sybase Adaptive Server Enterprise System Administration Guide.*

Parallel Load

Chart 9 demonstrates the increase in throughput when increasing the number of dump devices from 1 to 32 for the load database command. The throughput increased from 9.8 MB per seconds for one dump device to a peak of 10.8 MB per second for 4 dump devices. The throughput at 32 devices was approximately 9.3 MB per second. Overall the throughput averaged between 34 to 37.9 GB per hour for the Backup Server to load the database.

The Backup was performed while loading the data from one SMART-2 controller to another SMART-2 controller. Each controller was configured as one logical drive consisting of 14 physical drives striped using RAID0. No other processes or database activities were performed during the throughput measurement intervals on the system. The SMART-2 controller Array Accelerator was enabled for 100% Read cache on the source controller and enabled for 100% Write cache on the target controller. Just opposite of the setting they used for the dump process.

The throughput of the Backup Server can be greatly influenced by the I/O rates of the database data and dump devices. During the load process, typically the speed at which one can read the data from the database dump device is the limiting factor. If the dump device has a greater read I/O capacity than the write I/O capacity of the database data device then the speed at which one can write the data to the database becomes the limiting factor. This can occur when loading the data into a database located on a SMART-2 controller configured using RAID5 for fault tolerance.

For an in-depth discussion on parallel backup refer to the *Backup and Recovery* section of Chapter 23 *Maintenance Activities and Performance* in the *Sybase Adaptive Server Enterprise*

Performance and Tuning Guide and the *Backup and Recovery* chapters in the *Sybase Adaptive Server Enterprise System Administration Guide*.

Parallel Processing

When the *sp_configure* settings for *number of worker processes*, *max parallel degree* and *max scan parallel degree* are enabled, the optimizer may choose to perform parallel processing of the query. You can force parallel processing of the query by placing a hint in the SQL statement.

For a discussion on the configuration of parallel features please refer to following references:

- *Parallel Queries* section of the *Sybase Adaptive Server Enterprise System Administration Guide*
- *Quick Start Guide to Parallel Query Processing*, Sybase TechNote, June 30, 1997, Document ID: 20019
- *FAQs about Parallel Query Processing*, Sybase FAQ, February 2, 1998, Document ID: 20043

Parallel Query

Parallel query processing can be performed on either an index, a table or both. The optimizer makes the decision to perform parallel queries based upon the type of query, size of index or table being scanned, I/O buffer size, number of table partitions, current sp_configure "Parallel Query" values and the query hint. The optimizer may choose to perform serial or parallel query processing based and at times will choose to perform parallel queries even though there is more I/O being performed than during serial query processing. The 'extra' I/Os being performed by the parallel query processing still results in a faster total execution time of the query over the serial query processing because 2 or more processes in parallel can complete the reads faster than a single serialized process.

For a more in-depth discussion of the parallel processing features of Sybase ASE please refer to the *Chapter23* in the *Sybase Adaptive Server Enterprise Performance and Tuning Guide*.

Parallel Query Index and Table Scans

Chart 10 shows the elapsed time necessary to perform a table scan on a table partitioned with 14 partitions. When 2K I/O buffers were used the time to complete the parallel table scan is less than half of the time for performing the serial table scan. When 16K I/O buffers were configured and utilized, both the serial and parallel table scans executed quicker than when using the 2K I/O buffers. The serial table scan when using 16K I/O is still slower than parallel table scan using 2K I/O buffers. Similar performance gains can be expected when performing parallel index scans on clustered and/or non-clustered indexes.

Chart 10: Parallel Query Table Scan

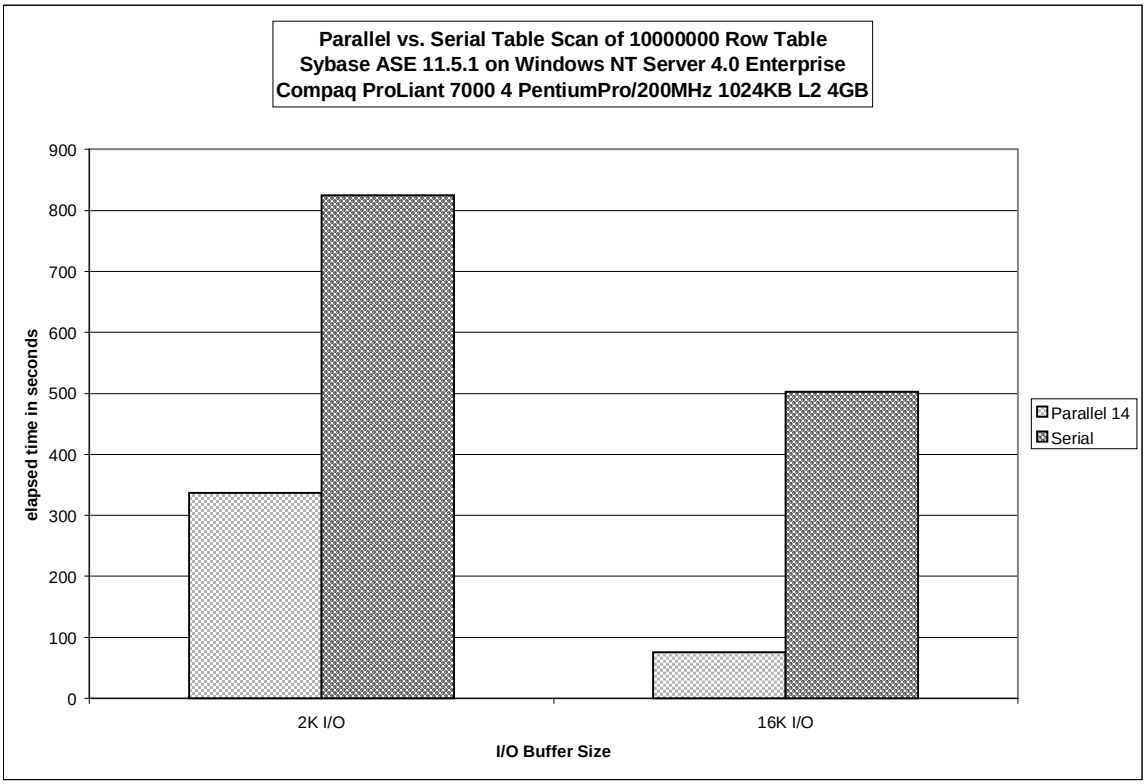

Parallel Index Creation

The index creation process will be performed using parallel processing techniques when the table that is being indexed has been partitioned. The number of parallel processes executing are governed by the *sp_configure* "*Parallel Query"* options for *number of worker processes* and *max parallel degree*. It is recommended, for best throughput, that you configure one worker process for each partition in the table. Other ways to improve the performance of the parallel index creation process are:

- have enough free space available for the worktables during the execution of the sort
- for clustered indexes, increase the number of table partitions
- for non-clustered indexes, increase the number of database devices on the target segment
- use *sp_sysmon* output to view the resources utilized by the create index command and reconfigure the system to reduce the total execution time

Chart 11 shows the execution times for the creation of both a clustered index and a non-clustered index using serial and parallel processing techniques. The time required to create the indexes in parallel on a 10000000 row table was less than half the time required to create the indexes in a serial manner.

Chart 11: Parallel Index Creation

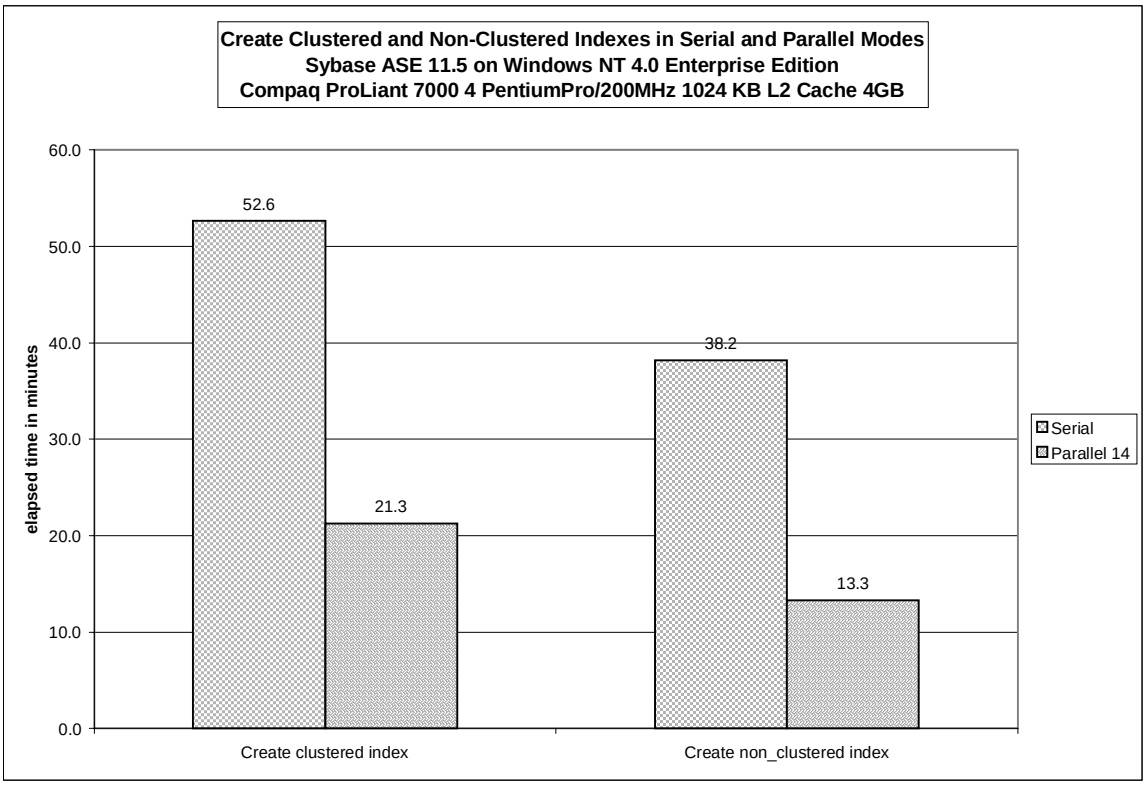

The table was partitioned to use 14 partitions, one for each of the 14 physical hard drive configured in the RAID0 drive array. The *sp_configure* value for *number of worker processes* was set to 2 times the number of table partitions yielding 28 worker processes. The *sp_configure* value for *max parallel degree* was configured to be equal to the number of table partitions + 1, in this case 15.

For a more in-depth discussion of the parallel index creation process, please refer to the *Sybase Adaptive Server Enterprise Performance and Tuning Guide* and the *Sybase Adaptive Server Enterprise System Administration Guide*.

Parallel Sorting

The parallel sort process can occur when creating an index, when reformatting the data, or executing a select statement with either an order by, distinct, or union clause. The ways to improve the performance of the parallel sort process are

- have enough free space available for the worktables during the execution of the sort
- increase the number of database devices used by the tempdb system segment
- use the *set sort_resource on* command to view the sort processing plan and reconfigure your system to optimize the sort operations

For a more in-depth discussion of the parallel processing features of Sybase ASE please refer to the *Chapter15* in the *Sybase Adaptive Server Enterprise Performance and Tuning Guide*.

Engine Groups and Execution Classes

By implementing engine groups and execution classes one can partition the database engines to execute only a particular type of transaction or query. One could limit the OLTP transactions to use only 3 out of 4 database engines with a medium priority and limit the DSS queries to use the remaining database engine at a high priority.


```
Sp_addengine 0, OLTP_GROUP
Sp_addengine 1, OLTP_GROUP
Sp_addengine 2, OLTP_GROUP

Sp_addengine 3, DSS_GROUP

Sp_addexeclass OLTP_CLASS, MED, 0, OLTP_GROUP

Sp_addexeclass DSS_CLASS, HIGH, 0, DSS_GROUP

Sp_bindexeclass OLTP_APP, AP, NULL, OLTP_CLASS
Sp_bindexeclass OLTP_proc01, PR, NULL, OLTP_CLASS
…
Sp_bindexeclass OLTP_proc99, PR, NULL, OLTP_CLASS
Sp_bindexeclass OLTP_user001, LG, NULL, OLTP_CLASS
…
Sp_bindexeclass OLTP_user999, LG, NULL, OLTP_CLASS

Sp_bindexeclass DSS_APP, AP, NULL, DSS_CLASS
Sp_bindexeclass DSS_user01, LG, NULL, DSS_CLASS
…
Sp_bindexeclass DSS_user09, LG, NULL, DSS_CLASS
```

One can dedicate a database engine just for use by the upper management. By implementing execution classes you can assign the execution class to a particle engine group. Simple rules like all CEO and CIO logins are assigned to an execution class "Keep_Them_Happy" and an engine group of "Go_Higher_Go_Faster" with a high priority will go a long way to improving the System Administrator and DBA relationship with the higher echelons.

```
Sp_addengine 0, Go_Higher_Go_Faster
Sp_addexeclass Keep_Them_Happy, HIGH, 0, Go_Higher_Go_Faster
Sp_bindexeclass CEO, LG, NULL, Keep_Them_Happy
Sp_bindexeclass CIO, LG, NULL, Keep_Them_Happy
```

For a more in-depth discussion of engine groups, execution classes and parallel processing features of Sybase ASE please refer to the following documentation:

- *"Sybase® Adaptive Server™Enterprise Performance and Tuning Guide"*, Document Id# 32645-01-1150-02
- *Great OLTP performance in a mixed OLTP/DSS environment*, Compaq Database Technology and Solution Center, Technical White Paper, Database Performance and Optimization Series, March 1998
- *"Configuring Sybase® Adaptive Server™Enterprise for Windows NT"*, Document Id# 32645-01-1150
- *"Great OLTP Performance in a Mixed OLTP/DSS Environment"*, Sybase Technical White Paper, March 1998

Parallel DBCC commands

The some of the *dbcc* commands and maintenance tasks such as *update statistics* are capable of parallel processing. You do not need to partition the database tables for these commands to take advantage of parallelism. If you do have tables partitioned, large I/O buffer pools configured and global asynchronous prefetch enabled they will execute even less time. If you *sp_configure global async prefetch limit* to 10 percent and have a 16K I/O Buffer Pool configured you will need to make certain that the pool size is at a minimum of 4960 KB. The 2K Buffer Pool needs to be configured with at least 140 KB. If you increase the *sp_configure global async prefetch limit* to 20 percent you will need to increase the 16K I/O Buffer Pool size to at least minimum of 9920 KB. The 2K I/O Buffer size needs to be increased to at least 280 KB. For a more in-depth

discussion of the *dbcc* parallel processing features of Sybase ASE please refer to the *Chapter18* in the *Sybase Adaptive Server Enterprise Performance and Tuning Guide*.

Dbcc checkstorage

The *dbcc checkstorage* command uses the dbccdb database to store information about each of the databases in the system. The *dbcc checkstorage* command takes advantage of the dedicated named cache, large I/O and parallel processing features of Sybase ASE 11.5 when properly configured. After executing the *dbcc checkstorage* command, execute the *sp_dbcc* commands to retrieve detailed reports on the status of the database.

To use *dbcc checkstorage* command you will need to setup the dbccdb database. It is important that the dbccdb database be placed on a fault tolerant device, the information stored in the dbccdb database can be used to repair a damaged user database. Follow the steps outlined in the section entitled *Preparing to Use dbcc checkstorage* in Chapter 18 of the *Sybase Adaptive Server Enterprise System Administration Guide*. For a discussion on the reports that can be extracted from the dbccdb database refer to section entitled *dbcc Stored Procedures* in the *Sybase Adaptive Server Enterprise Reference Manual*. For a list of recovery actions to perform when you obtain a dbcc checkstorage fault please refer to either of these documents:

- *Analyzing dbcc checkstorage Faults*, Sybase TechNote, June 6, 1998, Document ID: 2856
- *Sybase Adaptive Server Enterprise Troubleshooting and Error Messages Guide*, Document ID: 49604-01-1150-01[11]

Parallel Recovery

Upon startup of the ASE dataserver, parallel recovery will occur if the servername.cfg has the Parallel Query features enabled. The ASE 11.5 dataserver recovery process runs multiple parallel recovery processes but uses only 1 CPU processor resource. The parallel recovery feature can use only the default data cache but can take advantage of large I/O buffer sizes and asynchronous prefetch if enabled. On a 4 processor system the following servername.cfg settings can provide a dramatic reduction in the recovery time of a downed database.

```
[Backup/Recovery]
        recovery interval in minutes = 5
[Cache Manager]
        global async prefetch limit = DEFAULT
[Named Cache:default data cache]
        cache size = 256 M
        cache status = default data cache
        cache status = HK ignore cache
        cache replacement policy = DEFAULT
[2K I/O Buffer Pool]
        pool size = 64 M
        wash size = 512 K
        local async prefetch limit = DEFAULT
[16K I/O Buffer Pool]
        pool size = 192 M
        wash size = 512 K
        local async prefetch limit = DEFAULT
[O/S Resources]
```

[11] The *Sybase Adaptive Server Enterprise Troubleshooting and Error Messages Guide* may not be in the documentation kit that shipped with your copy of Sybase ASE 11.5. Please refer to the Sybase TechNote "*How Customers Can Order Latest Troubleshooting & Error Messages Guide*", Document ID: 968 dated May 15, 1998 for instructions how to acquire this document.


```
        max async i/os per engine = 512
        max async i/os per server = 512
[Parallel Query]
        number of worker processes = 8
        memory per worker process = 4096
        max parallel degree = 8
        max scan parallel degree = 8
[Processors]
        max online engines = 1
        min online engines = DEFAULT
[SQL Server Administration]
        housekeeper free write percent = 5
```

On a system with 2GB of memory configure the default data cache size to 1536 MB, 2K I/O Buffer Pool to 256 MB and the 16K I/O Buffer Pool to 1280 MB. The settings for *recovery interval in minutes* and *housekeeper free write percent* will not speed up the recovery of the database. These settings, during normal processing, cause the dataserver to perform automatic checkpoints and buffer pool management tasks that would reduce the time necessary to recovery the database when an unexpected system crash were to occur.

For a more in-depth discussion of the parallel processing features of Sybase ASE please refer to the *Sybase Adaptive Server Enterprise Performance and Tuning Guide*. For an in-depth discussion of the *sp_configure* commands please refer to the *Sybase Adaptive Server Enterprise System Administration Guide*.

Creating a Database Using RAID and Parallel Processing

The Sybase ASE 11.5.1 features of table partitioning, parallel BCP load and parallel processing can be used to minimize the elapsed time to initialize, create, load, index and update statistics of a database. Chart 12 shows the elapsed times to perform these tasks on the initial loading of 30GB of raw data into a database configured with 90GB space for tables and indexes. The tasks were performed on the database using RAID0, RAID1 and RAID5 fault tolerant disk configurations. As evidenced in the chart the major time differences were measured when creating the clustered and non-clustered indexes. The time difference to create the database on a RAID5 fault tolerant system took approximately 12 hours longer than the creation of the database on a RAID0 fault tolerant system.

Chart 12: Effects of Using RAID and Parallel Processing to Build a 30GB Database

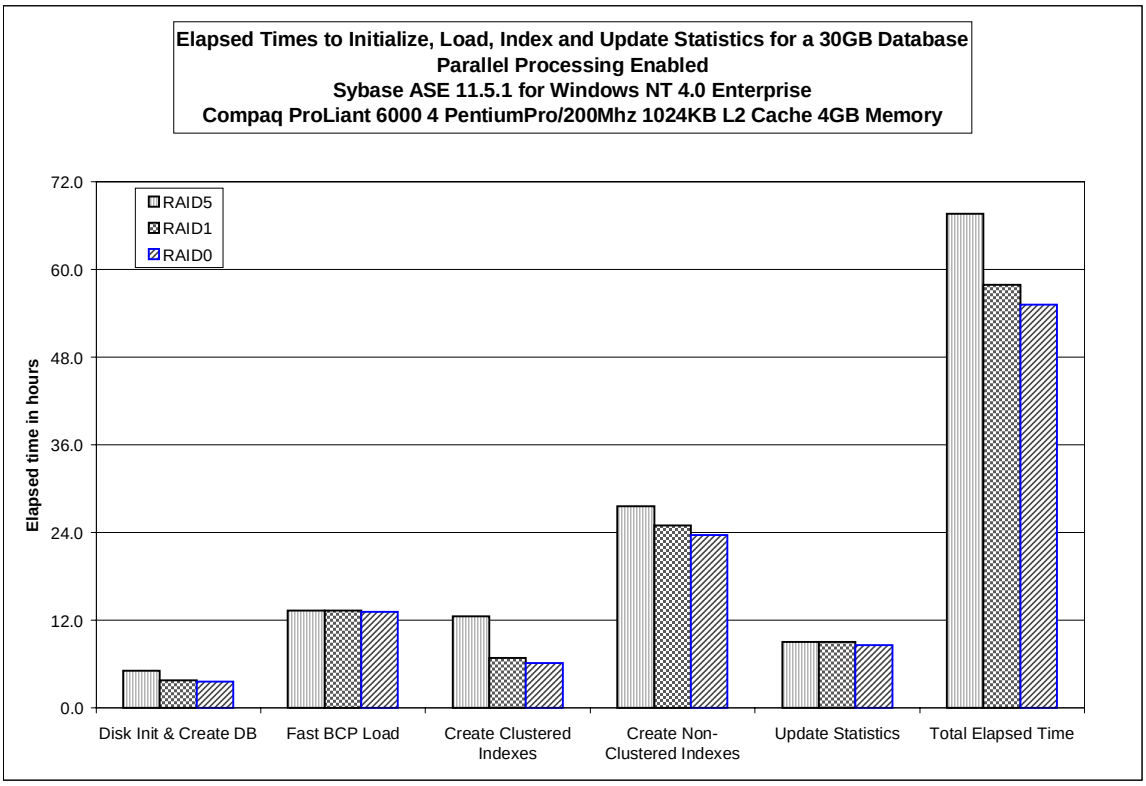

The ProLiant 6000 system was configured with 3 SMART-2DH controllers each with 14 drives of 4.3GB each, for a total of 42 drives. The elapsed time to create this database could have been reduced if the I/O capacity of the disk system was increased by adding additional controllers and disk drives or replacing the existing 4.3GB disk drives with higher throughput drives. Please refer to Table 3 Maximum I/O Operations per Second per Disk Drive later in this document obtain the I/O rates for the drives.

The total number of table partitions was 60. The sp_configure values for Parallel Query were:

Number of worker processes = 120
Memory per worker process = DEFAULT
Max parallel degree = 60
Max scan parallel degree = 1

For loading the data into the database 60 separate parallel BCP processes were simultaneously started.

Microsoft Windows NT 4.0 vs. SCO UnixWare 7.0

One of the most often asked questions from System Administrators and Database Administrators is which operating system to choose for best system throughput, scalability, integration with existing systems, manageability, security and ease of use. A proper discussion of this topic would consume the time remaining in this century and therefore is beyond the scope of this document.

Chart 13 demonstrates the throughput and processor scalability of Sybase ASE 11.5.1 for each operating system on the Compaq ProLiant 7000 PentiumPro 200MHz system. For each operating system, the ProLiant 7000 was configured with 4GB of memory, 7 SMART-2DH Array controllers connected to 98 9.1GB drives configured using RAID0 for the data, 1 SMART-2DH Array controller connected to 8 9.1GB drives configured using RAID1 for the log and one NetFlex-3 network controller configured to use TCP/IP as the network protocol. Sybase ASE for

NT was configured with Sybase total memory of approximately 2.9GB. Sybase ASE for UnixWare was configured with Sybase total memory of approximately 3.8GB. It is interesting to note that Sybase ASE 11.5.1 scales almost identically regardless of the underlying operating system.

Chart 13: Throughput and Scalability Comparison of Windows NT and UnixWare

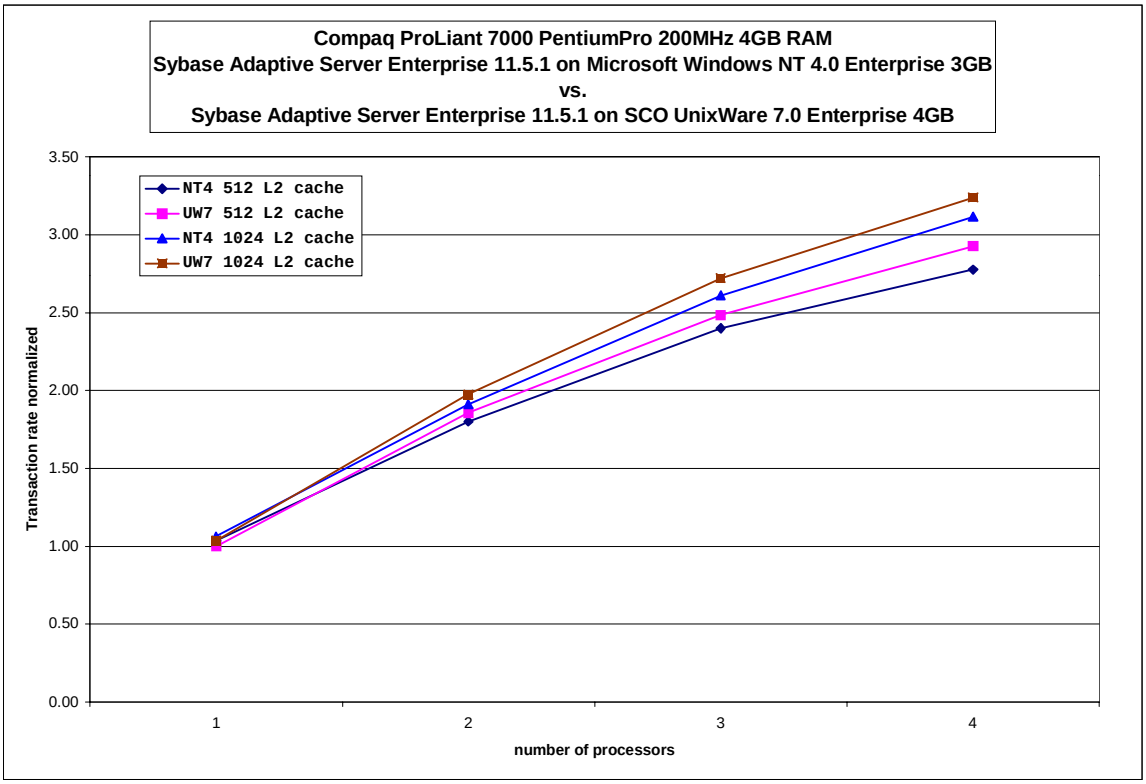

Chart 14 demonstrates the throughput and scalability when the Sybase total memory value is set identically for both operating systems. The *sp_configure total memory* value of 1524880 Sybase pages or approximately 2.9GB was used. It is interesting to note that when both Sybase ASE 11.5.1 for Microsoft Windows NT and Sybase ASE 11.5.1 for SCO UnixWare are configured identically and running the identical OLTP application workload there is no difference in the throughput of the systems.

Chart 14: Throughput and Scalability Comparison of Windows NT and UnixWare

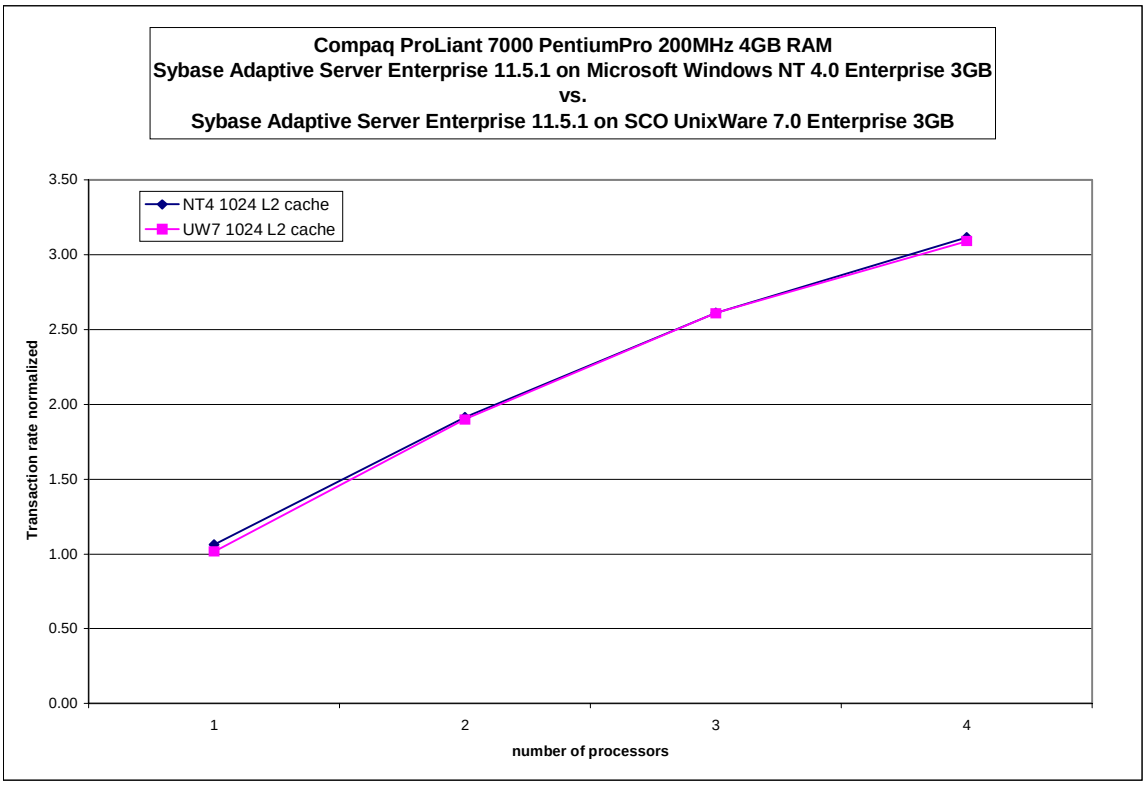

Performance Monitoring and Tuning

Performance monitoring of Sybase ASE can be performed in several different methods. You may want to use all, some or none of them depending upon the requirements of your environment. Most System Managers and Database Administrators want to know the current workloads to resolve daily problems and the long term trends to determine future hardware and software requirements.

A concise method exists for altering and verifying the many tuning parameters that affect the performance of Sybase ASE 11.5 for Windows NT. All configurable items in Sybase are stored external to the database in an ASCII format data file of the name, servername.cfg. This makes it possible to quickly fix any configuration errors with a text editor like Notepad without the hassles of previous versions via the buildmaster utility.

Performance Monitoring Tools

One can use the Windows NT Performance Monitor, Microsoft Network Monitor, Sybase Monitor Services, Monitor Viewer for Sybase Central, Sybase Historical Server, Sybase SQL PerfMon Integration for Windows NT Performance Monitor, Sybase sp_sysmon command, Compaq Insight Manager or a Sybase Monitor Client Library application to determine the throughput of the system. Each of these tools can be used independently or in a coordinated effort to monitor the system. The results from the monitoring can be used to determine system bottlenecks, plan for future hardware expansion or simply to help understand the complex interactions between the Windows NT operating system, Sybase ASE and the hardware.

Windows NT Performance Monitor

The Windows NT Performance Monitor can collect the performance counters and display them in real time or store them in a data file for later review. It can be used to determine system level and

application level bottlenecks or determine long term trends in system wide resource utilization. Please refer to the *Microsoft Windows NT Workstation 4.0 Resource Kit* for a concise and in-depth discussion of the performance counters and there meaning. The *Microsoft Windows NT Server 4.0 Resource Kit* does provide some information on the performance counters but does not have the detailed information on how to interpret them to detect bottlenecks that the *Microsoft Windows NT Workstation 4.0 Resource Kit* contains. If you need the information on performance monitoring of the network please refer to the section entitled *Using Performance Monitor with TCP/IP Services* in the *Microsoft Windows NT Resource Kit for Windows NT Workstation and Windows NT Server Version 3.5*. This section appears not to be in the Microsoft Windows NT Server 4.0 or Workstation 4.0 Resource Kits.

Microsoft Network Monitor

The Microsoft Network Monitor Tools and Agent provide Windows NT Server, Windows NT Workstation and Windows 95 systems the ability to perform monitoring, capture and tracing of local networks as well as RAS connections. It also can provide real-time statistical information like percentage of network utilization, frames per second, bytes per second and so on. These features make the Network Monitor useful in detecting and resolving problems that require low level analysis of the transport mechanisms.

The Microsoft Network Monitor Tools and Agent are a part of the Microsoft Systems Management Server product. These utilities are provided with Windows NT Server and Windows NT Workstation but are not installed during a regular install of the Windows NT operating system they can be added later via the *Control Panel, Network Services, Add* button. Please refer to the *Microsoft Windows NT Server Networking Guide* of the *Microsoft Windows NT 4.0 Server Resource Kit* for more detail on configuring and using the Microsoft Network Monitor Tools.

Sybase Monitor Services

The Sybase Monitor Server is a Sybase Open Server application that provides real time information on the performance of your system. This information is made available to Sybase Central Monitor Viewer, Sybase Historical Server and applications that use the Monitor Client Library. It is very useful for resolving performance problems and bottlenecks of the Sybase database.

For proper operation of the Sybase Monitor Server you may have to reconfigure the following sp_configure items.

Event buffers per engine = 2000
Max SQL text monitored = 1024
Shared memory starting address = default

Please refer to *Sybase® Adaptive Server™Enterprise Monitor Server User's Guide, Sybase® Adaptive Server™Enterprise Historical Server User's Guide* and *Sybase® Adaptive Server™Enterprise Monitor Client Library Programmer's Guide* for information on configuring your system to use the Sybase Monitor Services.

Sybase SQL PerfMon Integration

Sybase ASE 11.5.1 allows the monitoring of the data server via the Microsoft Performance Monitor for NT when *SQL PerfMon Integration* is enabled. The counters will appear for each instance of Sybase ASE configured and running. In other words you can monitor multiple instances of Sybase ASE whether on they are running on the same system or different systems all from one central location and correlate this activity with standard Windows NT performance counters.

In Table 4, located at the end of this document, the SQL PerfMon Integration counters are listed along with the Sybase parameter that can be tuned to alter its performance. For some of the counters the information shown is just that information, no Sybase sp_configure parameter can be

altered to change its value. Please refer to the *Sybase® Adaptive Server™Enterprise System Administration Guide* and the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for more information.

Sybase sp_sysmon Command

The Sybase *sp_sysmon* command provides a wealth of information on the internals of the database server. The reports generated from it can be several KB in size. This tool is the most useful of all the monitoring tools for database tuning in a controlled environment. Please refer to the *Sybase® Adaptive Server™Enterprise Performance and Tuning Guide* for very detailed instructions on interpreting the results and tuning suggestions.

Compaq Insight Manager

The Compaq Insight Manager (CIM) can be used to remotely monitor the performance and health of the Compaq server. The Compaq Insight Manager Agents report on hardware level performance characteristics of the processors, EISA bus, PCI bus, network controllers, SCSI controllers, SCSI devices, SMART Array controllers and drives arrays. CIM can be configured to send an alert message via e-mail, pager or phone to the System Administrator in case of a catastrophic hardware failure such as a power outage or failed hard drive. CIM also provides pre-failure warnings of potential hardware problems such as a disk drive that is experiencing a high number of retries or a memory module that has exceeded the maximum threshold for correctable memory errors. This allows the System Administrator to schedule preventative maintenance during non-peak hours of operation.

For the Compaq Insight Manager console to receive information from a Compaq server the following products must be installed and properly configured:

- Compaq supplied device drivers from the Compaq SSD for NT
- SNMP Service from the Windows NT distribution CD
- Compaq Insight Manager Agents for NT from the Compaq Systems Management CD

The Compaq device drivers can be installed via the setup program of the Compaq SSD for Windows NT. The SNMP Service is installed from the Windows NT distribution CD via the *Control Panel, Network, Services, Add* selection. The Compaq Insight Manager Agents for NT are installed via the setup program on the Compaq Systems Management CD.

The Insight Manager Agents can be configured via the *Control Panel, Insight Agents, SNMP Settings* tab to perform data collection at intervals from as frequent as 5 seconds and as infrequently as 60 minutes. To maintain high OLTP transaction rates and DSS query throughput set the *Data Collection Interval* to 60 minutes.

Please refer to the Compaq Insight Manager online Help file and the *"Compaq Insight Manager"*, White Paper, March 1998, Document Number ECG061/0398 for more information on the configuration and use of Compaq Insight Manager.

System Processor Utilization, Monitoring and Optimization

This section provides information on parameters that you can set to influence performance from the system processor perspective and tools you have available to monitor the behavior of the system.

Optimizing Windows NT 4.0 & ASE Operations

Both Windows NT and Sybase ASE are designed with many different tunable parameters and settings that can potentially increase the throughput of your system. The problem is knowing which settings are beneficial for your environment.

Windows NT Thread Priority

By default, all Sybase threads fall into the ***variable priority class***, priorities ranging from 1 to 15. Most threads in the system, including applications, fall into this class, and the Windows NT dispatcher can vary the priority of these threads to optimize system response time. Another priority class exists, called ***real-time priority class***, with priorities ranging from 16 to 31, and is used by time-critical applications requiring immediate processor attention. Once a thread runs in the variable priority class, its priority cannot be dynamically raised to the real-time class, and vice versa.

The *Control Panel, System, System Properties, Performance tab, Application Performance, Boost* option setting determines how Windows NT schedules base thread priorities. The *Boost maximum* slide setting informs Windows NT to schedule the initial foreground base thread priority at level 9 and the background base thread priority at level 7. If you choose the *Boost* middle slide setting, the initial foreground base thread priority is at level 8, background base thread priority remains at 7. If you choose the *Boost None* slide setting the initial foreground and background base thread priority will be at level 7.

If the server is dedicated to Sybase, change the Boost setting to *Maximum*. If the server is not dedicated to Sybase and is being used to run other application programs or printer services then change the *Boost* setting to *None*.

Refer to the *Windows NT Workstation and Windows NT Server Resource Kits* for more detailed information on this topic.

Sybase –P Boost Priority Option

By default, Sybase threads[12] are started as variable priority class threads having a base priority of 7-9. You can increase the base priority of the Sybase threads before starting Sybase with the -P option. This starts the Sybase threads with a base priority of 13-15. We recommend using the -P option to boost Sybase priority on dedicated Sybase multiprocessor systems. On dedicated Sybase single processor systems do not use the -P option.

The Sybase -P parameter will have a greater effect on the performance of your server than the *Application Performance Boost* option setting.

Windows NT Maximize Throughput for Network Applications

The *Control Panel, Networking, Services, Server, Properties, Optimization* should be set to *Maximize Throughput for Network Applications*, [13] this tells the Windows NT memory manager to "back off" and allow Sybase to manage its' own memory area(s). What you really are doing is altering the way Windows NT controls its' memory optimization routine from disk cache management to process working set management.

Refer to *Microsoft Windows NT Server Resource Kit* and *Microsoft Windows NT Server, System Guide* for a more detailed discussion.

Sybase Runnable Process Search Count

The tunable parameter *schedspins* is now renamed as *runnable process search count*. The default value of 2000 is satisfactory for most environments and no longer needs to be tuned except in rare cases. If you would like to experiment with this tunable the following settings are recommended:

runnable process search count

[12] You can check the priority of Sybase ASE threads via the Performance Monitor, Thread, Priority Base and Priority Current for the Instance of sqlsrvr -> 0 through 15.

[13] Some software packages may alter this setting so after installing any new software or changing any other Windows NT settings check to insure that the proper setting of "*Maximize Throughput for Network Applications*" is enabled.

UniProcessor systems
    I/O intensive systems use value of 1
    CPU intensive systems use value of 5
Multiprocessor systems
    I/O intensive systems range 50-200
    CPU intensive systems range 750-1250
Systems using Binary Large Objects (BLOBS)
    leave at default of 2000

Sybase Max Online Engines

The sp_configure *max online engines* setting, determines how many CPU resources should be dedicated exclusively to Sybase. It is recommended on dedicated Sybase multiprocessor systems that *max online engines* value be set equal to the number of CPUs in the system. This option greatly increases performance, but since it can degrade performance of non-Sybase applications running on the same machine, it needs to be used only on dedicated Sybase multiprocessor systems. Experimentation will be necessary to obtain the best performance for your non-dedicated database server environment. If you set the value for *max online engines* higher than the number of CPUs installed in your server, you will still be able to run Sybase but your overall system performance will be severely impacted.

max online engines
    dedicated systems
        number of Sybase engines = number of processor boards
    non-dedicated systems
        number of Sybase engines = number of processor boards - 1

Unnecessary Services

You may turn off or stop all unnecessary services running on the server, since they add a processing overhead on the system processor(s). Even though you can run with as few as two services (*Server* and *Sybase*), we recommend also running *EventLog*. You can disable services such as *ClipBook Server, Computer Browser, Messenger, Network DDE, Network DDE DSDM, Schedule* and others if you do not need them. You can turn these services off through Control Panel/Services. Always monitor the behavior of your system, and reverse any changes that adversely impact the performance or impair the stability of your system.

CPU Performance Monitor Objects and Their Meaning

This section provides you with some important Performance Monitor objects that you can monitor to gain insight into the system processor utilization.

Object: Processor

Counter: % Processor Time

The *% Processor Time* counter monitors the percent of CPU utilization of all threads (both Privileged Mode and User Mode threads) running on the system. This counter is very useful in determining whether or not you have a CPU bottleneck. Consistent CPU utilization of over 95% on all processors is a safe indicator of a CPU bottleneck. In such case, add another system processor or upgrade the existing processor(s) to the next generation processor(s). You may also review your logical database design to make sure it is designed as efficiently as possible. As stated elsewhere, hardware tuning and upgrades can only partially compensate for an inefficiently designed database.

Object: Processor

Counter: % Privileged Time, % User Time

In a dedicated Sybase environment, *% Privileged Time* is typically the percentage of time the system processor is spending executing NT kernel commands, much of this time is associated with

processing Sybase I/O requests. *% User Time* is the percentage of time the system processor is spending executing user applications, such as Sybase.

If Sybase finds all (or most of the frequently used) objects in the data cache, and does heavy processing with these objects (such as table joins), very little I/O is generated. As a result, *% Privileged Time* can be low, may be in the 5-15% range, whereas *% User Time* can be as high as 85-95%.

On the contrary, if Sybase generates a large amount of I/O, *% Privileged Time* will be substantially higher (30-40%) and *% User Time* will be substantially lower (60-70%).

Both of these counters are very useful in determining how different types of operations are utilizing the system processor(s). If your system is spending too much time doing I/O, you may need to further investigate the disk subsystem and how to relieve it of some I/O. You may also need to add more memory. If your system is spending most time doing Sybase computing, you may want to investigate denormalization, reduction of the number of joins, horizontal partitioning, or adding/upgrading your system processor.

Memory Utilization, Monitoring and Optimization

This section provides information on parameters that you can set to influence performance from the system memory perspective and tools you have available to monitor the behavior of the system.

Window NT Paging: Detection and Avoidance

In a dedicated Sybase environment, Windows NT Paging can occur when you allocate too much memory to Sybase and the remaining memory does not satisfy Windows NT. In a non-dedicated Sybase environment, memory allocated to other programs also plays a significant role.

Keep in mind that to reduce or eliminate paging, you can either add more physical memory or reduce the amount of memory allocated to your processes, such as Sybase. We recommend starting with allocating a minimum of 32 megabytes to Windows NT and the rest to Sybase and carefully monitoring the behavior of the system.

Object: Memory

Counter: Page Faults/sec

This counter monitors total page faults by all running processes. A page fault occurs when a process makes a reference to a virtual memory page that is not in its working set in the memory. You can also monitor Page Faults/sec per individual process, such as the number of page faults per second Sybase generates. On a dedicated Sybase system, Sybase-generated page faults and total page faults per second will often be identical.

You should have very little or no paging once the system reaches a steady state; occasional paging prior to steady state is often acceptable.

Sybase Tuning Parameters

Sybase provides a wide variety of methods of configuring the memory in your system via the *sp_configure* options for *total memory*, *additional network memory*, and *procedure cache percent* and the Logical Memory Manager (LMM). The LMM permits the definition of dedicated named caches that can be assigned to a database, log, table, index, text object or image object. The dedicated named cache can be configured in size from a minimum of 512KB up to several GB. A dedicated named cache can use either strict or relaxed LRU cache replacement policies.

Memory

The correct amount of memory allocated to Sybase varies from one environment to another. As a rule of a thumb, the more objects (tables, indexes, etc.) you can cache, the better performance you

will have. If you can't cache them all, attempt to cache the more frequently accessed objects in the following priority: tempdb, clustered indexes, non-clustered indexes and data pages.

The LMM feature in Sybase ASE permits the refinement of which database objects you want to cache, how much space to allocate for the cache, and how to manage the cache. To determine the initial size of a dedicated named cache, use the following procedure. In the following example we assign the *tempdb* database to the *c_tempdb* dedicated named cache.

1. Use *sp_spaceused* to determine the size of the database object
2. Edit the server.cfg file to establish the dedicated named cache object for the database object

```
[Named Cache:c_tempdb]
    cache size = 2M
    cache status = mixed cache
    cache status = HK ignore cache
[2K I/O Buffer Pool]
    pool size = 2M
    wash size = 512K
```

3. Use *sp_bindcache* to associate the database object to the dedicated named cache
4. Use *sp_helpcache* to verify the cache binding

Please refer to the *Sybase System Administration Guide* for more in depth information on the tuning of the LMM.

Dbcc Memusage

Dbcc memusage[14] is a very useful tool in monitoring what objects are in the Sybase cache. Remember that out of the total memory allocated to Sybase, the Sybase engine itself will allocate enough memory for itself and its required data structures, and then will divide the remaining memory between the procedure cache and data cache according to the *procedure cache percent* parameter. Also remember that Sybase allocates memory in 2 Kbyte pages.

Dbcc memusage reports the following information:

- Sybase memory allocation at startup
- Memory used by 20 largest objects in the data cache
- Memory used by 12 largest stored procedures, triggers, views, rules and defaults in the procedure cache

Run d*bcc memusage* occasionally to determine which objects, data and procedures are in the Sybase memory areas. Refer to the *Sybase System Administration Guide* for more detailed information on how to use this feature.

Procedure Cache

The amount of procedure cache you need for your stored procedures, triggers, views, rules and defaults depends on their number and size. Keep in mind that multiple users accessing the same stored procedure, for example, will cause Sybase to store one copy of the same stored procedure for each user.

Since reading stored procedures from disk to the procedure cache is costly, you want to keep all needed stored procedures in the procedure cache. In other words, you want to configure the procedure cache large enough to hold all of your frequently executed stored procedures.

Use d*bcc memusage* to monitor the 12 largest procedures in the procedure cache.

[14] For accurate results, set the database to single user mode before issuing this command. Also it maybe necessary to issue the *dbcc traceon (3604)* command to redirect the output to your command window instead of the Sybase SQL Server errorlog.

If you have more than 12 stored procedures in the procedure cache, *dbcc memusage* only shows the 12 largest ones, and you will have to get a little creative to find out if you have enough procedure cache. Below is an example.

1. Make a copy of one of your largest stored procedures, call it a different name, and execute it. This step will place this stored procedure in the procedure cache. Use d*bcc memusage* to verify this.
2. Execute the rest of your stored procedures, except for the one created in step 1, several times.
3. Use *dbcc memusage* to determine which 12 largest stored procedures remained in the procedure cache. If you have enough procedure cache, your dummy stored procedure should have remained in the list of top 12. If you don't have sufficient procedure cache, your dummy stored procedure should have aged out and has been pushed out of the procedure cache by other stored procedures.

The above procedure works if you have a relatively few number of stored procedures and works just fine in a test environment. But a simpler method that works in production environments without is available. To find the number of 2K pages that a stored procedure requires in memory, you can run this isql statement:

select (count(*)/8) +1 from sysprocedures where id=object_id("procedure_name")

Execute this isql statement for each of the most commonly executed stored procedures in your environment, sum the results and reconfigure the setting of *procedure cache percent*.

Disk Subsystem Utilization, Monitoring and Optimization

This section provides information on parameters that you can set to influence performance from the disk subsystem perspective and tools you have available to monitor the behavior of the system.

Recovery Interval

Sybase uses the *recovery interval* parameter to calculate how often to perform a checkpoint. Sybase internally computes the frequency of the checkpoint based on the *recovery interval* value and the transaction rate to guarantee that in case of a system crash the database(s) would be recovered within the time specified by the *recovery interval* parameter.

For example, in a heavy transaction processing environment, a *recovery interval* value of 5 (minutes) means that at the present transaction rate Sybase would recover the database within 5 minutes after restarting from a crash. Taking into consideration the present transaction rate, Sybase would perform the checkpoint approximately every 2 minutes. Also if *trunc log on chkpt* is enabled then a checkpoint will occur every minute.

Refer to the *Disk Subsystem Planning* section for more detailed information on how checkpointing can effect disk throughput.

Diskperf

The *diskperf* option[15] allows you to monitor the disk subsystem activity. It is very useful when monitoring performance of the drive subsystem. If this option is disabled, performance monitor will not be able to monitor low level disk-related activity, such as *LogicalDisk* and *PhysicalDisk* counters.

Having this option enabled slightly degrades performance. Enable *diskperf* only when needed, then disable it to get maximum performance.

[15]The *diskperf* option is described in a greater detail in the *Microsoft Windows NT Workstation Resource Kit*.

You can enable or disable *diskperf* using Control Panel/Devices or by issuing the *diskperf -y|n* command from the system prompt. When using the Control Panel, set *diskperf* to automatically start at boot time to enable. You must restart your system for the *diskperf* option to become effective.

Disk Performance Monitor Objects and Their Meaning

You can monitor *LogicalDisk* and *PhysicalDisk* related objects, such as *Avg. Disk sec/Transfer* or *Disk Transfers/sec*. Monitoring *LogicalDisk* and *PhysicalDisk* related objects requires you to have the *diskperf* option enabled, which slightly reduces performance.

Note that resetting the *diskperf* option requires you to reboot your machine.

Object: Logical Disk

Counter: Avg. Disk sec/Read, Avg. Disk sec/Write

The *Avg. Disk sec/Read* and *Avg. Disk sec/Write* counters monitor the average number of seconds for read or write operations from or to a disk, respectively.

Use the following guidelines when optimizing the disk subsystem. Do not have more average disk seconds per read or write operation than the values in Table 2 below.

Table 2: Average Disk Access Times

	Avg. Disk sec/Read	**Avg. Disk sec/Write**
Transaction Log Access (with the Array Accelerator enabled)	N/A	<=10msec
Database Access (with the Array Accelerator enabled)	20-40msec	<=10msec

If your values are much higher than those suggested in the above table, you may want to increase the speed of your disk subsystem by adding more drives or using faster drives.

Object: Logical Disk

Counter: Disk Transfers/sec

The *Disk Transfers/sec* counter monitors the rate of read and write operations on the disk. It is important to realize that if you have a Compaq SMART or SMART-2 Array controller and several drives allocated to a logical volume, this counter monitors the total number of disk transfers per logical volume. To calculate number of disk transfers per second per drive, you must divide the *Disk Transfers/sec* value by the number of drives in the logical volume.

Use the following guidelines when optimizing the drive subsystem. You should not have more I/O requests (disk transfers) per second per drive than the values in Table 3 below.

Table 3: Maximum I/O Operations per Second per Disk Drive

2 KB I/O requests	**1.05 GB**	**2.1GB**	**4.3 GB**	**4.3 GB 10,000 rpm**	**9.1 GB**	**9.1 GB 10,000 rpm**
Sequential Write	≈150	≈160	≈180	≈290	≈210	≈290
Random Read/Write	≈30-40	≈50	≈55-60	≈75-80	≈70-75	≈75-80

NOTE: With the Array Accelerator enabled, you may actually see substantially higher I/O per second per drive rates than suggested above, especially during checkpoint. This increase is due to the Array Accelerator write posting some of these I/Os. In the Compaq Database Performance

labs, during the checkpoint process, we have actually measured rates of up to 90 random I/Os per second per drive on the 4.3 GB drives. This increase in throughput can be attributed to some of the I/Os being serviced by the SMART-2DH Array Accelerator.

If your values are much higher than those suggested in Table 4, you should increase the throughput of your disk system by adding more drives, using faster drives and/or adding additional disk controllers.

Network Utilization, Monitoring and Optimization

The network performance characteristics of your system can be determined by using several different tools. The Windows NT Performance Monitor, the Microsoft Network Monitoring Tool, Sybase *sp_sysmon* command and 3rd party hardware and software based network sniffers. The discussion of using 3rd party hardware and software based network sniffer is beyond the scope of this document. Typically one can properly configure the network without the use of a network sniffer.

Sybase sp_sysmon command

The Sybase *sp_sysmon* output provides you with information on the network workload across all Sybase database network engines. If the workload appears unbalanced the system is typically processing a large number of small network packets or BCP activities.

The counter for *total network I/O requests* in conjunction with a knowledge of the your networks' bandwidth capacity can be used to determine if you are encroaching upon the maximum capacity of the network.

If the counters for *average bytes received per packet* and/or *average bytes sent per packet* are close to the same value as the *sp_configure default network packet size* you should try increasing the *sp_configure max network packet size* and reconnect the users with the larger network packet size.

Please refer to the *Sybase® Adaptive Server™ Enterprise Performance and Tuning Guide* for an in-depth discussion on interpreting the *sp_sysmon* output.

Network Performance Monitor

The Microsoft Windows NT Performance Monitor can provide you counters on most all network activities. The SNMP Service must be installed on TCP/IP based systems for the Performance Monitor to gather and report the statistics. If it is not installed use the *Control Panel, Services, Add* selection to install it.

Please refer to the *Microsoft Windows NT 4.0 Server Resource Kit, Networking Guide*, the *Microsoft Windows NT 4.0 Workstation Resource Kit* and the chapter entitled *Detecting Network Bottlenecks* in the *Microsoft Windows NT Resource Kit for Windows NT Workstation and Windows NT Server Version 3.5*. For some odd reason the chapter on *Detecting Network Bottlenecks* did not survive the migration from version 3.5 to 4.0 of the Resource Kit.

Sybase Runserver Command File and the Windows NT Registry

The Sybase startup parameters are written in both the *runserver* file and the Windows NT Registry. The recommended method for altering the Sybase startup parameters is through the Sybase Configuration Server Utility. There are other methods of altering the contents of the *runserver* file and the Windows NT Registry with utilities such as Edit and Regedt32 but without careful attention to details an inconsistency may exist between the two files. When this happens the manner in which Sybase runs, may not be as expected. The order in which Sybase and Windows NT read the startup parameters is defined so you can determine what parameters are being utilized. If you start Sybase via the Services Manager or the Control Panel, Services menu, the Registry values are utilized. If you manually start Sybase at the DOS prompt or use the

servername.bat file then the parameters specified in the *servername.bat* file will take effect. If the Registry has an entry for a parameter not specified at the DOS prompt or in the *servername.bat* file then the Registry parameter value(s) will be used as the default.

The Sybase Server Configuration Utility does not verify or resolve any inconsistencies between the *runserver* file and the Registry. You will need to manually check that the command line parameters specified in the *C:\Sybase\install\servername.bat* file match the Registry entries for arg0-argx under the heading of HKEY_LOCAL_MACHINE, SOFTWARE, Sybase, Server, Servername, Parameters. In the *servername.bat* file, if multiple entries are listed for the same parameter, the last parameter is the one utilized. In the Registry, if multiple entries are listed for the same parameter, the arg with the highest number associated with it is the one utilized.

This also applies to the entries for Sybase Central, Backup Server, Monitor Server and Historical Server.

Configuring Sybase ASE for High Availability

With the advent of such convince systems as 24 hour ATMs, global communications, the Internet and world wide commerce systems the demand for database servers with high availability has increased tremendously in the past few years. To provide high availability of your Compaq ProLiant servers, Sybase ASE supports both the Compaq On-Line Recovery Server and the Microsoft Cluster Server (MSCS).

Compaq On-Line Recovery Server

The Compaq On-Line Recovery Server is a hardware based solution that provides high availability for the Compaq ProLiant servers. The Compaq Remote Monitor Service and the Compaq Online Storage Recovery Utility from the Compaq SSD for Windows NT provide the tools necessary to monitor and manage the Compaq On-Line Recovery Server systems. For more details on the Compaq Remote Monitor Service please refer to *"Implementing Online Storage Recovery Option under Windows NT"*, Compaq White Paper, Document Number 066A/0797.

The following documents provide detailed information on "how-to" install and configure your ProLiant Server, Windows NT and Sybase ASE for use with the Compaq On-Line Recovery Server option:

- *"Configuring Sybase SQL Server for High Availability: Compaq On-Line Recovery Server"*, Sybase TechNote, May 1997, Document ID: 2854
- *"Compaq On-Line Recovery Server"*, second edition, White Paper, November 1996, Document Number 286A/1196
- *"Compaq Recovery Server Option User Guide"*, Part Number 213818-003

Microsoft Cluster Server

The Microsoft Cluster Server (MSCS) is a software based solution for high availability under Microsoft Windows NT 4.0 Server, Enterprise Edition. It is recommended that Service Pack 3 or higher be installed on the server for sane operation of MSCS. The Compaq ProLiant servers and Sybase ASE are supported running MSCS. The Compaq Cluster Verification Utility located on the current Compaq SSD for Windows NT can be used to determine if your Compaq ProLiant server is setup for use in a MSCS environment. This utility will run in the Windows NT 4.0 Server and Windows NT 4.0 Server Enterprise Environment.

The following documents provide detailed information on "how-to" install and configure your ProLiant Server, Windows NT and Sybase ASE for use with MSCS:

- *"Configuring Sybase® Adaptive Server™Enterprise for High Availability with Microsoft Cluster Server"*, Sybase White Paper, March 1998, Document ID: 20205

- *"Microsoft Windows NT Cluster Server Administrator's Guide"*
- *"Installing Sybase® Adaptive Server™Enterprise and OmniConnect™ on Windows NT"*

Configuring Compaq Tape Devices for Sybase Backup Server

Before proceeding with this section make certain that you have applied the latest Compaq SCSI device drivers and Compaq tape device drivers from the Compaq SSD for NT. This will insure the proper operation of Compaq Tape Devices under Windows NT. Failure to install the proper SCSI device and tape device drivers will result in possible data corruption on the backup tape set. This data corruption is usually not detected during the backup phase but rather during the verification of the backup or worse yet during the restoration of the database.

Verifying Compaq SCSI Device Drivers

Use the Control Panel, SCSI Adapters, Devices tab, select the Compaq 32-Bit SCSI-2 Controllers, Properties, Driver tab to verify that the installed device file is the Compaq cpq32fs2 device driver file. One can also use the Windows NT Utility program regedt32 to view this information. On some installations the Windows NT Install process will install a either a Symbios Logic or NCR device driver instead of the Compaq device driver. Replace the Symbios Logic or NCR device driver with the correct Compaq device driver from the Compaq NT SSD diskettes. Failure to install the proper SCSI device drivers will result in data corruption on the backup device.

Verifying Compaq Tape Device Drivers

Use the Control Panel, Tape Devices, Devices tab, select one of the listed device, Properties, Settings tab to verify that the SCSI Adapter Information Name is the Compaq cpq32fs2 device driver file. Also check the General tab and record the Device Map information. The Device Map information along with the Settings information will allow you to determine which physical tape drive unit is mapped to the Windows NT tape device names \\.\Tape0, \\.\Tape1 and so on. One can also use the Windows NT Utility program regedt32 to view this information.

Defining Sybase Tape Device Capacity

Most tape devices state the capacity of the tape device using both the uncompressed and compressed capacities. The compressed capacity assumes a hardware data compression ratio of 2:1. Therefore theoretically you can place 20GB of data from disk onto a 10GB tape cartridge. For proper operation of the Sybase Backup Server do not use the compressed capacity value.

For example when defining the Sybase tape dump device information for the Compaq DLT 10/20 tape device. Use only the uncompressed tape capacity value of 10240MB (10GB). Do not define the dump device with the compressed capacity of 20480MB (20GB). You can not be guaranteed that the hardware compression algorithm in the tape device will yield a 2:1 compression ratio when backing up your database.

Performing a Sybase Backup

In the following example we demonstrate "how to" perform a Sybase backup of the master database using a Compaq DLT 10/20 tape drive on a Compaq ProLiant 4500.

1) Use the Windows NT Utility Program regedt32 to view the Registry and verify that the installed SCSI device driver is the Compaq **cpq32fs2** device driver. Please refer to Figure 4. On some installations the Windows NT Install process will install a either a Symbios Logic or NCR device driver instead of the Compaq device driver. Replace the Symbios Logic or NCR device driver with the correct Compaq device driver from the Compaq NT SSD diskettes. Failure to install the proper SCSI device drivers will result in data corruption on the backup device.

Figure 4: Regedt32 verification of Compaq SCSI Device Driver

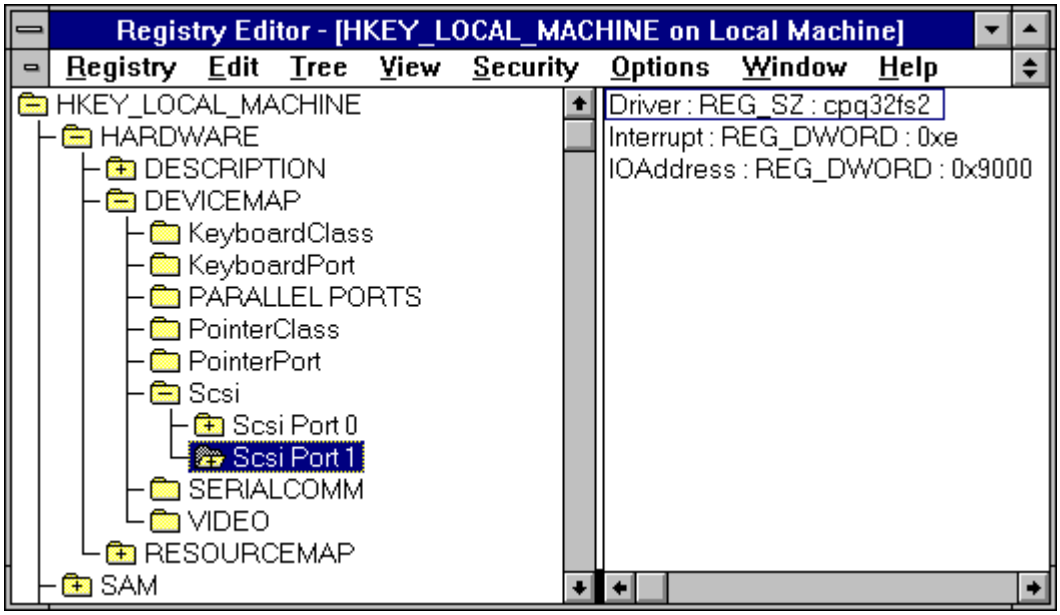

2) Please refer to Figure 5. Check the Registry to verify the tape device settings for your system. On this system the DeviceName is Tape0 , Identifier is DLT2000 and Type is TapePeripherial. The DeviceName maybe different on your system.

Figure 5: Regedt32 Verification of Compaq Tape DeviceName

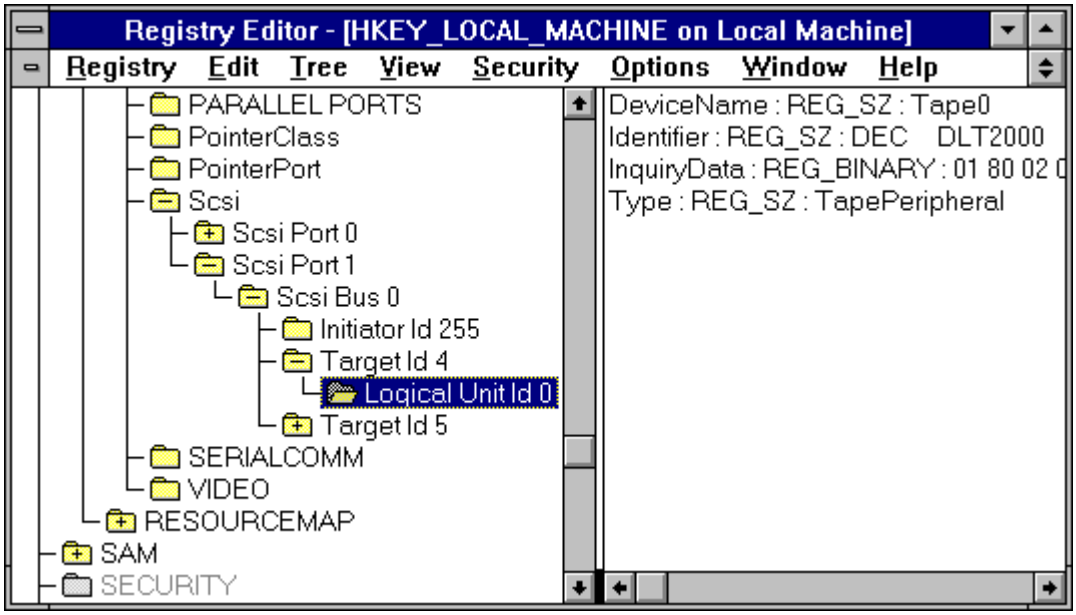

3) Use isql to verify that the Sybase Backup Server SYB_BACKUP exists

```
C:\>isql -Usa -P -SROLLING_ROCK
1> select * from sysservers
2> go
 srvid  srvstatus srvname              srvnetname
 ------ --------- -----------          --------------
```



```
      1          1 SYB_BACKUP          ROLLING_ROCK_BS

(1 row affected)
```

4) Check for the existence of the Sybase pre-defined tape dump devices

```
1> sp_helpdevice
2> go
 device_name    physical_name        description
------------------ ------------------------------------------------------------------------------
master              master.dat                  special, default disk, physical disk, 21.00 MB
sysprocsdev      C:\SYBASE\DATA\SYBPROCS.DAT       special, physical disk, 16.00 MB
tapedump1              \\.\TAPE0       tape,      625 MB, dump device
tapedump2              \\.\TAPE1       tape,      625 MB, dump device
 (4 rows affected, return status = 0)
```

5) Remove the Sybase pre-defined tapedump1 device because it incorrectly defines our tape device '\\.\Tape0' with a capacity of 625 MB. Also remove the Sybase pre-defined tapedump2 device to avoid misinforming the Backup Operator of the dump devices available on this system.

```
1> sp_dropdevice tapedump1
2> go
Device dropped.
(return status = 0)
1> sp_dropdevice tapedump2
2> go
Device dropped.
(return status = 0)
```

6) Add the tape dump device information for the Compaq DLT 10/20 tape device. Use only the uncompressed tape capacity value of 10240MB (10GB). Do not define the dump device with a capacity of 20480MB (20GB). You can not be guaranteed that the hardware compression algorithm in the tape device will yield a 2:1 compression ratio when backup your data.

```
1> sp_addumpdevice "tape", "DLT_2000", "\\.\Tape0", 10240
2> go
'Tape' device added.
(return status = 0)
```

Verify that it installed properly

```
1> sp_helpdevice
2> go
 device_name         physical_name        description
------------------------------------------------------------------------------------------------------------
DLT_2000             \\.\Tape0             tape,      10240 MB, dump device
master                master.dat           special, default disk, physical disk, 21.00 MB
sysprocsdev          C:\SYBASE\DATA\SYBPROCS.DAT       special, physical disk, 16.00 MB
 (3 rows affected, return status = 0)
```

7) Start the database dump from an isql session. The dump database command doesn't require the capacity option because the capacity was specified in the sp_addumpdevice command.

```
1> dump database master to "DLT_2000"
2> go
```


WARNING: In order to LOAD the master database, the SQL Server must run in single-user mode. If the master database dump uses multiple volumes, you must execute sp_volchanged on another SQL Server at LOAD time in order to signal volume changes.
Backup Server session id is: 5. Use this value when executing the 'sp_volchanged' system stored procedure after fulfilling any volume change request from the Backup Server.
Backup Server: 6.52.1.1: OPERATOR: Volume to be overwritten on '\\.\Tape0' has unrecognized label data.
Backup Server: 6.78.1.1: EXECUTE sp_volchangednt@session_id = 5,@devname ='\\.\Tape0',@action = { 'PROCEED' | 'RETRY' | 'ABORT' }

8) From a second isql session[16] execute the sp_volchanged command.

1> *sp_volchanged @session_id = 5, @devname = '\\.\Tape0', @action = PROCEED*
2> *go*
(return status = 0)

9) From the first isql session you can view the Backup Server progress and verify that the dump completed without error.

Backup Server: 6.28.1.1: Dumpfile name 'master962360EDF0 ' section number 0001 mounted on device 'tape drive'
Backup Server: 4.58.1.1: Database master: 944 allocated kilobytes DUMPed.
Backup Server: 4.58.1.1: Database master: 2066 allocated kilobytes DUMPed.
Backup Server: 3.43.1.1: Dump phase number 1 completed.
Backup Server: 3.43.1.1: Dump phase number 2 completed.
Backup Server: 4.58.1.1: Database master: 2070 allocated kilobytes DUMPed.
Backup Server: 3.43.1.1: Dump phase number 3 completed.
Backup Server: 4.58.1.1: Database master: 2074 allocated kilobytes DUMPed.
Backup Server: 6.45.1.1: Be sure to remove the tape/floppy from drive \\.\Tape0 (server: , session id: 5).
Backup Server: 3.42.1.1: DUMP is complete (database master).

Configuring Windows NT and Sybase ASE for High Security

Microsoft Windows NT 4.0 and Sybase ASE 11.5.1 have different security and auditing features that can be enabled to provide high levels of security. Among these are integrated login features, limited hours of availability based upon time and date, resource limitations to avoid a single user or application from devouring all of the system resources and features to assist with audit trails. For a more in depth discussion of these features please refer to the section *"Limiting Access to Server Resources"* and the *sp_configure Security Related* section in the *Sybase Adaptive Server Enterprise System Administration Guide*.

C2 Security Rating

At the time of this publication, Sybase Adaptive Server Enterprise 11.5.1 for the Microsoft Windows NT 4.0 Server Enterprise Edition with Service Pack 3 executing on a Compaq ProLiant server has not been evaluated to meet the Class C2 Security Level requirements of the Department of Defense (DOD). The C2 security level criteria are outlined in the DOD publication DOD 52.00.28-STD, *Department of Defense Trusted Computer System Evaluation Criteria* (TCSEC), affectionately known as the "*Orange Book*".

Configuring the hardware and software components to meet the C2 security level criteria may result in a decrease in throughput when compared to an unsecured system. The reduction in

[16] The second ISQL session must be logged into the same ASE dataserver as the dump database command was issued from. A common mistake here is to login to the Backup Server instead of the ASE dataserver.

throughput can be attributed to the extra workload necessary for maintaining the secure login environment, audit trail capabilities and increased overhead in performing I/O to NTFS volumes instead of raw disk partitions for the database devices.

Microsoft Windows NT version 3.5 with Service Pack 3 was evaluated by the National Computer Security Center (NCSC) and meet the Class C2 Rating security level in July 1995. It also received the NCSC recognition for B2 Trusted Path and B2 Trusted Facility Management. There have been no significant changes in Windows NT 3.51 and 4.0 that effect the security design features. Therefore it is theoretically possible to configure Microsoft Windows NT Server 4.0 Enterprise Edition with Service Pack 3 to meet the NCSC C2 Rating. Please refer to the *Microsoft Windows NT 3.5 Guidelines for Security, Audit and Control* and the *Microsoft Windows NT Server 4.0 Resource Kit* both published by Microsoft Press for more details. The *Microsoft Windows NT Enterprise Planning Guide: Implementing NT Security*, Part No. 098-54661 available from Microsoft also is a wealth of information on planning and deploying the security model for your environment.

Sybase SQL Server 11.0.6 running on the HP 9000 HP-UX BLS, 9.09+ platform was evaluated by the National Security Agency (NSA) and meet the Class C2 security level in 1996. For a complete discussion of the evaluated configuration please refer to *Appendix A* in the *SQL Server Installation and Configuration Guide for HP 9000 HP-UX BLS, 9.09+*.

Sybase Adaptive Server Enterprise 11.5.1 for Windows NT contains all of the security features of Sybase SQL Server 11.0.6 for HP-UX. Therefore it is theoretically possible to configure Sybase Adaptive Server Enterprise 11.5.1 running on Microsoft Windows NT Server 4.0 Enterprise Edition with Service Pack 3 to meet the Class C2 Security Level Rating. Please refer to the *Sybase® Adaptive Server™ Enterprise Security Features User's Guide*, the *Sybase® Adaptive Server™ Enterprise Security Administration Guide* and the *Sybase Adaptive Server Enterprise System Administration Guide* for more details.

Conclusion

We would welcome feedback from your configurations and experiences to improve our information products in the future. Please send us any comments or suggestions on the attached form, attach addition sheets if necessary. This will help us tailor the future information products to your needs, and will enable us to make future revisions of this document and related new information products available to you.

Sybase SQL PerfMon Integration Counters

Table 4 list the various Sybase PerfMon Integration Counters that appear in the Microsoft Windows NT Performance Monitor when the *sp_configure parameter SQL PerfMon Integration* is enabled.

Table 4: SQL PerfMon Integration Counters for the NT Performance Monitor

Object	Counter	Instance	Counter Definition	Tunable
Access	Endxact log record count		Number of end xact records in the transaction log.	
	Endxact log record count per second		Number of end xact records in the transaction log per second.	
Buffer	% of buffer hits in wash area	@average, @total and one for each named cache	Not that big of a performance hit, but extra IO is performed when value is high. Wash area is too big.	Sp_poolconfig
	% turnover rate (2k pages)	@average, @total and one for each named cache	If this value is high, it means transactions have to wait for IO to complete before a buffer becomes available. User should increase the 2k buffer cache.	Sp_cacheconfig, sp_poolconfig, sp_cachestrategy
	% turnover rate (4k pages)	@average, @total and one for each named cache	If this value is high, it means transactions have to wait for IO to complete before a buffer becomes available. User should increase the 4k buffer cache.	Sp_cacheconfig, sp_poolconfig, sp_cachestrategy
	% turnover rate (8k pages)	@average, @total and one for each named cache	If this value is high, it means transactions have to wait for IO to complete before a buffer becomes available. User should increase the 8k buffer cache.	Sp_cacheconfig, sp_poolconfig, sp_cachestrategy
	% turnover rate (16k pages)	@average, @total and one for each named cache	If this value is high, it means transactions have to wait for IO to complete before a buffer becomes available. User should increase the 16k buffer cache.	Sp_cacheconfig, sp_poolconfig, sp_cachestrategy
	HK disk overload	@average, @total and one for	HouseKeeper terminated a batch	Dbcc tune (deviochar, vdevno,

		each named cache	due to overloading a disk in a batch.	batch_size)
	House keeper activity	@average, @total and one for each named cache	House Keeper % of activity is scale indicator that says the house keeper is actively working.	Sp_configure housekeeper free write percent
	House keeper backoff count	@average, @total and one for each named cache	Number of times house keeper process has backoffed because we hit the configuration limit on redirty count.	
	Housekeeper extra writes	@average, @total and one for each named cache	% of extra writes done by housekeeper due to redirty pages.	
Database	Transaction log free pages remaining	@average, @total and one for each database	Transaction Log free pages remaining (statistic only available on currently active databases).	
	Transaction log usage (%).	@average, @total and one for each database	% of Transaction Log used(statistic only available on currently active databases).	
Disk	% device contention	@average, @total and one for each virtual device	If this value is high, that means we have IO contention on this device.	Partition table or distribute data over more devices
	% IO reads	@average, @total and one for each virtual device	Percentage of IO reads to the device.	
	% IO writes	@average, @total and one for each virtual device	Percentage of IO writes to the device.	
Engine	% bytes read of all network io	@average, @total and one for each engine	This is the % of network bytes read by this engine of the entire number of network bytes read.	
	% bytes sent of all network io	@average, @total and one for each engine	This is the % of network bytes sent by this engine of the entire number of network bytes sent.	
	% engine CPU busy	@average, @total and one for each engine	This is the % of CPU used for an instance of an SQL Server Engine	

Lock	Exclusive intent		number of exclusive intent locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Exclusive page		number of exclusive page locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Exclusive table		number of exclusive table locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Shared intent		number of shared intent locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Shared page		number of shared page locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Shared table		number of shared table locks	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “table”,

				null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Total exclusive		number of exclusive locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Total intent		number of exclusive and shared intent locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Total locks		total number of locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Total page		number of exclusive, shared and update page locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Total shared		total number of shared locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock

				promotion PCT
	Total table		number of exclusive and shared table locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
	Update page		number of update page locks.	Fillfactor; max_rows_per_page; set transaction isolation level; sp_setpglockpromote “server | database | table”, null, lock promotion LWM, lock promotion HWM, lock promotion PCT
Memory manager	Memory usage (pages)		Number of memory pages allocated.	Sp_configure total memory
Network	Ave Packet Size (read)		Average packet size of network packets read.	Sp_configure default network packet size, sp_configure max network packet size
	Ave Packet Size (send)		Average packet size of network packets sent.	Sp_configure default network packet size, sp_configure max network packet size
	Bytes Read (per interval)		number of bytes read from network in sampling interval.	
	Bytes Read per second		number of bytes read from network per second.	
	Bytes Sent (per interval)		number of bytes sent to network in sampling interval.	
	Bytes Sent per second		number of bytes sent to network per second..	
	Packets Read (per interval)		number of packets read from network in sampling interval.	
	Packets Read per second		number of packets read from network per second.	
	Packets Sent (per interval)		number of packets sent from network in sampling interval.	

	Packets Sent per second		number of packets sent from network per second.	
Procedure manager	% sproc cache hits		Store Procedure hit percentage in procedure cache	Sp_configure “total memory”; sp_configure “procedure cache percent”
Process	Process CPU Usage Ratio (%)	@average, @total and one for each process	% ratio of the increase of the CPU Usage Counter for this Process versus the sum of the increases of the CPU Usage Counters for all Processes in this interval.	
	Process CPU Usage (accumulated)	@average, @total and one for each process	Accumulated total of the CPU Usage Counter for this Process.	
	Process CPU Usage (delta)	@average, @total and one for each process	Increase of the CPU Usage Counter for this Process in this interval.	
	Process IO Issued Ratio (%)	@average, @total and one for each process	% ratio of the increase of the IO Issued counter for this Process versus the sum of the increases of the IO Issued Counters for all Processes in this interval.	
	Process IO Issued (accumulated)	@average, @total and one for each process	Accumulated total of the IO Issued Counter for this process.	
	Process IO Issued (delta)	@average, @total and one for each process	Increase of the IO Issued Counter for this process in this interval.	
Sqlsrvr	% of all engines busy		Total CPU busy for all SQL Server Engines.	Sp_configure “max online engines”
Transaction log subsystem	% of Log Page Flushes due to log page full		When this parameter is low, it means we are writing the log pages too often. Sometimes this may be cause by disk controller caching being turn on for the device.	Sp_configure “user log cache size”
	Ave xact records per log page		Average log records per log pages written to transaction log	

			per sampling interval.	
	Ave xact records per xact		Average log records per transaction for committed transactions in the sampling interval.	
	Log page writes		Number of log page writes for sampling interval.	
	Log page writes per Xact		Number of log page writes for sampling interval per Transaction.	

User Registration/Evaluation Form

Please fill out and return to us this registration/evaluation form to help us keep you up to date with future revisions of this document and related new information products. Your effort will help us improve the quality of the future information products.

Name: _____
Title: _____
Company: _____
Address: _____
_____
_____
Phone Number: _____

Please evaluate the quality of this document:

	Excellent				Poor
Technical Accuracy	1	2	3	4	5
Organization	1	2	3	4	5
Clarity	1	2	3	4	5
Completeness	1	2	3	4	5
Timeliness	1	2	3	4	5

Please indicate the type of environment you have at your site:

Operating Systems	RDBMS	Processing Type
❑ Microsoft Windows NT	❑ Sybase ASE 11.5	❑ On-Line Transaction Processing
❑ SCO UnixWare 7	❑ Microsoft SQL Server 6.5	❑ Decision Support
❑ Novell InternetWare	❑ Oracle 8	❑ Batch Processing
❑ SCO Open Server	❑ Informix	❑ Other:
❑ IBM OS/2	❑ DB2	
❑ Other:	❑ Other:	

Please indicate the type of information you would like us to provide in the future:

Topic	Operating Systems	RDBMS
❑ Configuration and Tuning	❑ Microsoft Windows NT	❑ Sybase ASE 11.5
❑ Capacity Planning	❑ SCO UnixWare 7	❑ Microsoft SQL Server 6.5
❑ Integration Information	❑ Novell InternetWare	❑ Oracle 8
❑ Competitive Analysis	❑ SCO Open Server	❑ Informix
❑ Systems Management	❑ IBM OS/2	❑ DB2
❑ Other:	❑ Other:	❑ Other:

Additional Comments:

Return to:
Database Engineering
Compaq Computer Corporation
MailCode 090803
20555 SH 249
Houston, Texas 77070

