Lotus

IBM

Lotus Domino for S/390
Performance Tuning and Capacity Planning

B. Macfaden T. Draper I. Karge M. Lang D. Raisch

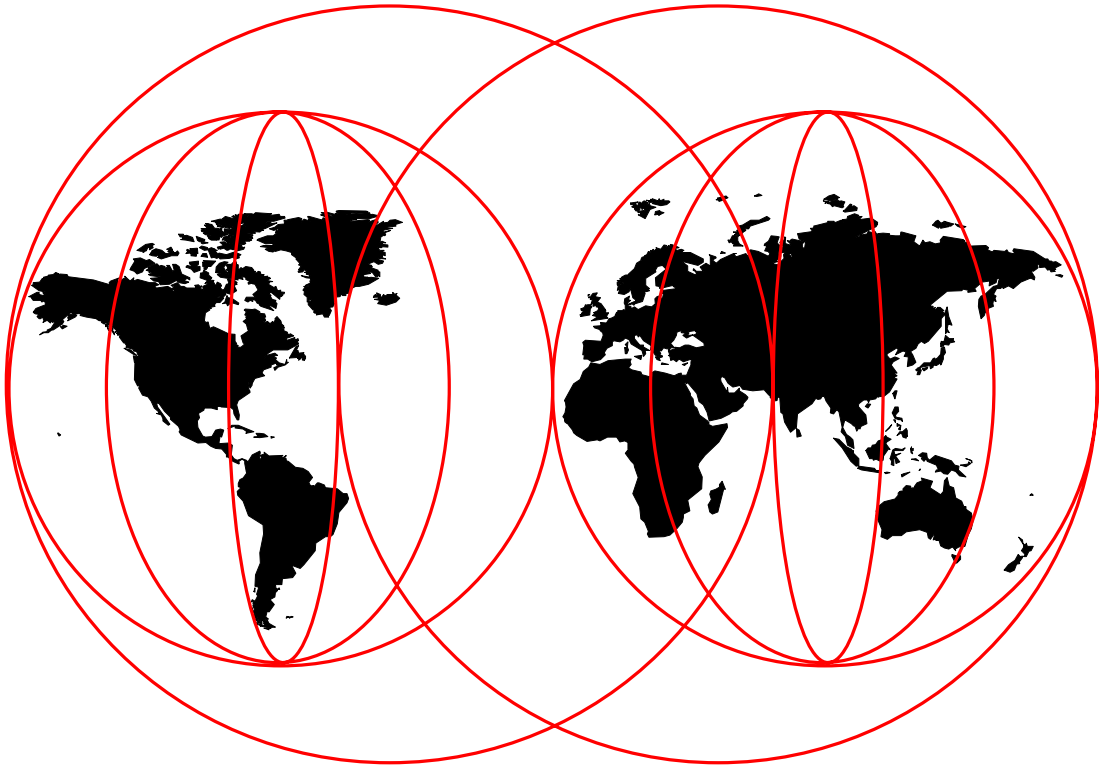

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5149-00

International Technical Support Organization

SG24-5149-00

Lotus Domino for S/390
Performance Tuning and Capacity Planning

May 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, “Special Notices” on page 85.

First Edition (May 1998)

This edition applies to Lotus Domino for S/390 Release 4.51 for use with OS/390 Version 1 Release 3 and later releases. Many of the techniques will apply to later releases of Lotus Domino for S/390, although some details may change.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HYJ Mail Station P099
522 South Road
Poughkeepsie, New York 12601-5400

Contents

Figures

Tables

Preface

This redbook will help OS/390 system programmers and Lotus Notes administrators monitor and tune Domino for S/390.

It identifies the factors that affect Domino performance on S/390 and makes recommendations for the server configuration and parameter settings to tune the Domino server for optimal performance.

It also provides information on how to monitor Domino performance, and how to do capacity planning for your installation.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization Poughkeepsie Center.

Brian Macfaden is a program manager at the International Technical Support Organization, Poughkeepsie Center. He writes extensively and teaches IBM classes worldwide on Lotus Domino for S/390 and on the value of the IBM System/390 platform. Brian has twenty years experience in the information technology industry, specializing in large systems, distributed systems and networking.

Terry Draper is a Large Systems Specialist in the United Kingdom. He has 25 years of experience with the S/390 platform. For the past 13 years he has worked for IBM as a specialist in S/390 hardware and software. His areas of expertise include performance of the operating system, subsystems and batch processing. He also specializes in competitive processor marketing.

Ingo Karge is a Host Office Migration Specialist and an Information Technology Architect for the host office customers in IBM Germany. He worked as an account SE for six years for large customers, and now works in the IBM Lotus integration group, where he is responsible for migration-related questions.

Mike Lang is a Consulting Market Support Representative in the IBM Lotus Integration Center, a part of the Dallas System Center, in Dallas, Texas, USA. He has provided support for OS/390 office systems for 17 years and is currently supporting Domino for S/390. He specializes in capacity planning and performance tuning. Prior to joining the Dallas System Center Mike was a Systems Engineer for 13 years in Cincinnati, Ohio supporting CICS and IMS application development.

Daniel Raisch is a Certified Senior Information Technology Specialist in Brazil. He has 17 years of experience with the S/390 platform. For the last 10 years he has been working on performance analysis and capacity planning. Daniel's areas of expertise include client/server architecture, OS/390 UNIX Services and Web implementation.

Thanks to the following people for their invaluable contributions to this project:

Joe Bostian	IBM S/390 Division
John Brooks	IBM S/390 Division
Pamela Dewey	IBM S/390 Division

Cathy Eilert	IBM S/390 Division
Jim Ray	IBM S/390 Division
Bill Schoen	IBM S/390 Division
Bob St John	IBM S/390 Division
Joe Sweeney	IBM S/390 Division
Marianne Hammer	IBM Global Services
Elaine Murphy	IBM Global Services
Rich Conway	ITSO, Poughkeepsie Center
Mike Ebbers	ITSO, Poughkeepsie Center
Bob Haimowitz	ITSO, Poughkeepsie Center
Paul Rogers	ITSO, Poughkeepsie Center

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 99 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

  For Internet users `http://www.redbooks.ibm.com/`
  For IBM Intranet users `http://w3.itso.ibm.com/`

- Send us a note at the following address:

  `redbook@us.ibm.com`

Chapter 1. Introduction

This book provides information to help you tune the performance of Lotus Domino for S/390, and to assist you in capacity planning a Domino workload on S/390.

As you approach this topic it is useful to remember:

- Domino was originally written for the PC platform. Therefore it was developed with different requirements than those for traditional S/390 applications such as CICS, IMS and DB2. The characteristics of the application are different, and one of our objectives is to explain the characteristics of Domino so that you are better able to monitor and tune the system.
- Domino was written to be platform independent. This means that it does not exploit all OS/390 functions directly. Contrast that to CICS, IMS and DB2, which were all originally written for OS/390 and designed specifically to exploit OS/390's capabilities.
- Domino was ported to the OS/390 UNIX Services (also called OpenEdition) interfaces of OS/390. You may not be familiar with these interfaces. This book will guide you in performance considerations for OS/390 UNIX Services.

This book assumes knowledge of Lotus Domino for S/390 at least to the level of the redbook, *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083. In particular, chapter 13 of that book describes how Domino interfaces to OS/390, and the settings for many OS/390 parameters. We do not repeat that information here, except where necessary for the topic being discussed.

1.1 Recommendations

Our key recommendations for performance tuning, performance monitoring and capacity planning of Lotus Domino for S/390 are:

- Monitor OS/390 closely as you introduce Domino.

  Domino does not run the same way as other S/390 workloads that you may be used to. Make sure that you are not artificially constraining Domino with parameter settings designed for a different type of application.

- Use a team approach.

  Domino tuning on S/390 is a combination of:

  – OS/390 tuning, including UNIX Services tuning
  – Domino server tuning
  – DASD tuning
  – Network tuning

  Therefore, the Notes administrator, the OS/390 performance specialist, the DASD administrator and the network administrator must work together to optimize Domino performance.

- See the more detailed list of tuning recommendations at the start of each chapter.

- Use valid measurements for capacity planning.

We discuss some valid capacity planning approaches in Chapter 6, “Capacity Planning” on page 65. Avoid using simple single-function tests as an indicator of the capacity of S/390 as a Domino server. These tests are easy to run, but they can give very misleading results.

1.2 A First Step

We do not consider this book to be the final word on Domino for S/390 performance tuning and capacity planning. There are performance and capacity improvements in plan for:

- Lotus Domino
- OS/390 and its UNIX Services interface
- OS/390 Hierarchical File System
- S/390 hardware

However, we hope that this initial work will assist you to better tune and manage your Domino servers on S/390.

Chapter 2. OS/390 Tuning

This chapter describes how Domino interfaces to OS/390, and in particular its address space and task structure. We discuss the tuning considerations for OS/390 when you are running one or more Domino servers.

2.1 Recommendations

These are the key recommendations for tuning OS/390 to run Domino:

- Make sure you have enough memory for Domino.

  Domino creates many address spaces, some of which have large working sets. OS/390 UNIX Services also makes extensive use of storage for performance, in the OS/390 UNIX Services kernel address space and DFSMS buffer manager address space. You need enough real memory for these or you will page to DASD, which will impact Domino response times.

- Check that OS/390 is set up to allow for the extra address spaces.

  We recommend that you use Workload Manager goal mode, since it automatically allocates resources so that the workload achieves its goals. If you use compatibility mode, make sure that OS/390 will allow for the large increase in address spaces and real memory. See 2.6.2, “Compatibility Mode” on page 10. Failure to do this has sometimes meant that the Domino server will not even start.

- Check the BPXPRMxx member in SYS1.PARMLIB for your OS/390 UNIX Services settings and make them large enough.

  There are several OS/390 UNIX Services parameters in BPXPRMxx which, if too low, will cause problems in the Domino server. It will not be able to create new tasks or get the virtual memory it requires. See *Lotus Domino for S/390 Install Guide for Servers* for the recommended values and follow them. See *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083 for further information.

- Make the Domino server a high priority OS/390 workload.

  The Domino server is an online system. It needs the same level of service that you would give to your CICS or IMS production systems. If the priority within OS/390 is not set high enough, then response times to the Notes client will suffer and the client request will time out.

- Implement PLO Instruction.

  A new instruction, Perform Locked Operation (PLO), has been introduced on the IBM CMOS processors to improve the performance of OS/390 UNIX Services locking. We have seen a significant reduction in the CPU time used by Domino servers when using this instruction. To exploit the PLO instruction, you need:

  – The PLO instruction on the processor. This instruction is supported on the 9672 Generation 3 (Model Rn4) and Generation 4 (Model Rn5) processors and later 9672 processors. PLO requires microcode level 98 which was available in November 1997 for the Generation 4 processors and in January 1998 for the Generation 3 processors.

  - OS/390 Version 2 Release 4 or later and the PTF to APAR OW32071.
  - Domino for S/390 Release 4.53.
- Tune and monitor the entire OS/390 system.

  If you add a Domino server to a badly tuned OS/390 system, it will suffer. Do all the things you would normally do for performance. Get IBM to help set up the system before you install the Domino server if you need the expertise. Then as you start to use Domino, monitor the system closely to see how Domino performs.
- Check out the OS/390 UNIX Services performance tuning tips on the Web.

  The URL is `http://www.s390.ibm.com/products/oe/bpxpftgt.html`.

2.2 OS/390 Address Spaces

Domino uses multiple OS/390 address spaces, as shown in Figure 1.

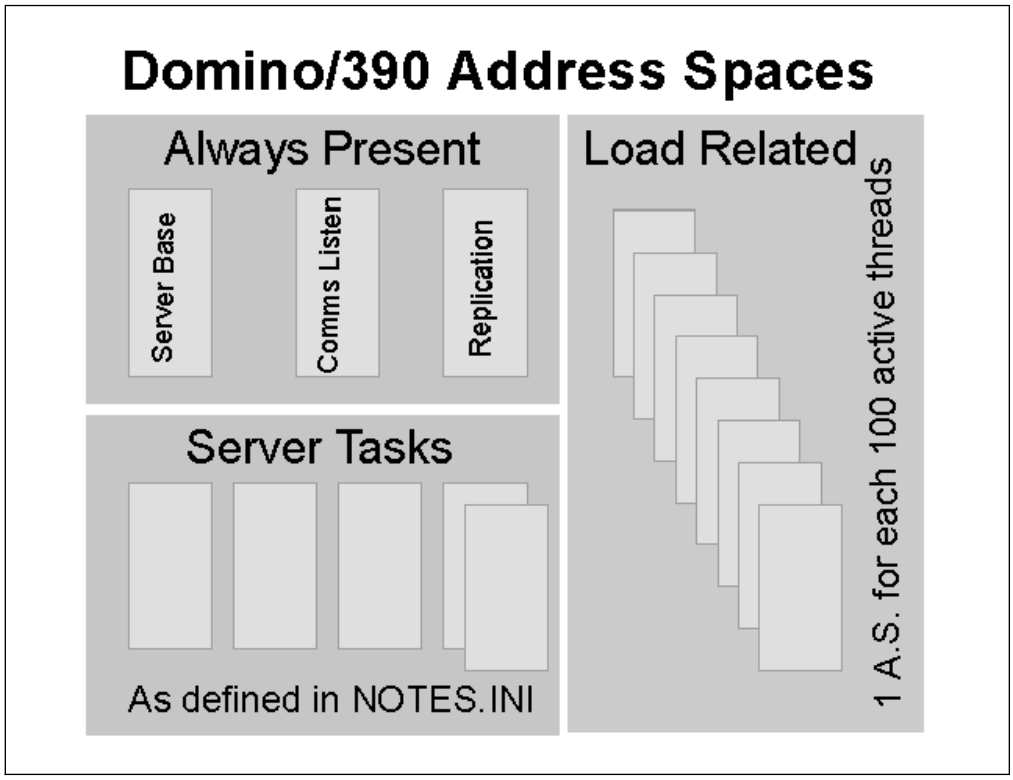

Figure 1. Domino Tasks and OS/390 Address Spaces

At initialization three address spaces are always started. These are:

- Server Base
- Comms Listen
- Replication

Additional server address spaces are created at startup depending on the functions requested in the initialization control file `notes.ini`. These server address spaces perform background tasks or functions requested by the user address spaces. Server address spaces can also be started at the request of the Domino

administrator while the Domino server is running. Domino also dynamically starts and stops server address spaces to perform some long running background tasks.

2.2.1 Address Spaces for User Connections

In addition to the above address spaces, Domino creates address spaces to process work for the user, that is, the Notes client. As each client connects to the Domino server it is assigned a thread. This thread has its own subtask within an address space and thus its own task control block (TCB). A user address space has 100 threads. Thus the first 100 users are in one address space. The 101st user will cause a new address space to be created, which will have the next 100 users. As a user connection is terminated, the thread becomes available for reuse.

For example, if, at a point in time, you have 520 connected users, you would have at least six (520/100) user address spaces. The reason for saying “at least” is that after a user address space is created it is never deleted, even when there are no threads in use within that address space. The number of user address spaces reflects the high water mark of the number of concurrent user connections and pending connections.

A thread is used during the setup of the connection and if the attempted connection fails the thread is freed for future use. This means that you can have more threads than connected users.

2.2.2 Address Space Management

In OS/390 Version 1 Release 3 APPC is used to manage the Domino address spaces.

In OS/390 Version 2 Release 4 and later releases the Domino address spaces are created and managed by the Workload Manager. After a Domino process has completed, the Workload Manager keeps the address space for 30 minutes to allow reuse of the address space by another Domino process. Address spaces are retained in both Workload Manager goal mode and compatibility mode. Domino address spaces retained by the Workload Manager keep their Domino job name but have a step name of BPXAS.

2.2.3 Address Space Naming

The Domino server creates all its address spaces with a job name of the user ID under which the server is started, followed by a number. This is shown in Figure 2 on page 6. This system is running a Domino server. The display is produced using the MVS command D A,L.

```
D A,L
IEE114I 12.55.05 1998.069 ACTIVITY 163
 JOBS     M/S    TS USERS    SYSAS    INITS   ACTIVE/MAX VTAM     OAS
00019    00023    00003      00028    00067    00003/00030       00024
 LLA      LLA      LLA      NSW  S  JES2     JES2     IEFPROC  NSW  S
 VTAM44   VTAM44   NET      NSW  S  OSASF    OSASF    OSASF    IN   S
 TCPIP    TCPIP    TCPIP    NSW  S  DFSMSHSM HSMSC52  DFSMSHSM NSW  S
 VLF      VLF      VLF      NSW  S  RMF      RMF      IEFPROC  NSW  S
 APPC     APPC     APPC     NSW  S  ASCH     ASCH     ASCH     NSW  S
 IHV      IHV      PSTEP01  OWT  S  SDSF     SDSF     SDSF     NSW  S
 DFRMM    DFRMM    IEFPROC  NSW  S  OPTSO    OPTSO    OPTSO    OWT  S
 RACF     RACF     RACF     NSW  S  EZAFTSRV EZAFTSRV EZAFTSRV NSW  S
 PORTMAP  PORTMAP  PMAP     OWT  S  RXPROC   RXPROC   RXSERVE  OWT  S
 IOASRV   IOASRV   SERVER   OWT  S  FTPD1    STEP1    STC      OWT  AO
 TCPIPOE  TCPIPOE  TCPIP    NSW  SO TSO      TSO      STEP1    OWT  S
 INETD1   STEP1    OMVSKERN OWT  AO PORTMAPO STEP1    STC      OWT  AO
 MVSNFS   MVSNFS   GFSAMAIN NSW  SO RUNCF    RUNCF    IEFPROC  NSW  S
 MACFAD62 STEP1    MACFAD6  OWT  AO MACFAD67 STEP1    MACFAD6  NSW  AO
 MACFAD68 STEP1    MACFAD6  IN   AO MACFAD69 STEP1    MACFAD6  IN   AO
 MACFAD61 STEP1    MACFAD6  NSW  AO MACFAD62 STEP1    MACFAD6  IN   AO
 MACFAD63 STEP1    MACFAD6  IN   AO MACFAD64 STEP1    MACFAD6  IN   AO
 MACFAD65 STEP1    MACFAD6  IN   AO MACFAD65 STEP1    MACFAD6  IN   AO
 MACFAD66 STEP1    MACFAD6  IN   AO MACFAD67 STEP1    MACFAD6  IN   AO
 MACFAD68 STEP1    MACFAD6  IN   AO MACFAD69 STEP1    MACFAD6  OWT  AO
 MACFAD61 STEP1    MACFAD6  OWT  AO MACFAD62 STEP1    MACFAD6  OWT  AO
```

Figure 2. Domino Address Space Names - D A,L Command

The Domino server was started from the user ID MACFAD6 and therefore the server address space names are MACFAD6x. As you see, you can have duplicate job names. You cannot predict what the job name will be for a specific Domino function. Therefore, you cannot differentiate between the various address spaces created by a Domino server.

2.2.4 Address Space and Program Name

If you wish to obtain the function names of address spaces on a running Domino server you should issue the OS/390 command D OMVS,A=ALL.

If you wish to display just one address space, then use the DA panel on SDSF to find the address space identifier (ASIDX) of the required address space, as shown in Figure 3 on page 7.

```
SDSF DA SC52  SC52     PAG    0 SIO   387 CPU  23/ 23  LINE 53-74 (74)
COMMAND INPUT ===>                                      SCROLL ===>
NP   JOBNAME  JOBID    OWNER    C POS DP REAL PAGING    SIO   CPU% ASID ASIDX
     MVSNFS   STC02386 STC        NS  FE 2273   0.00   0.00   0.00  105  0065
     MACFAD61 STC10573 MACFAD6    LO  FF  686   0.00   0.00   0.01  106  0080
     RMFGAT   STC10126 STC        NS  FE 6209   0.00   0.06   0.70  109  0081
     MACFAD64 STC10550 MACFAD6    IN  C1  734   0.00   0.00   0.06  110  0082
     MACFAD67 STC10577 MACFAD6    IN  C0 2208   0.00  93.30  12.25  111  0083
     MACFAD61 STC10586 MACFAD6    LO  FF  273   0.00   0.00   0.00  112  0084
     MACFAD68 STC10576 MACFAD6    IN  C2  649   0.00   0.00   0.01  115  0085
     MACFAD66 STC10575 MACFAD6    IN  C1  644   0.00   0.00   0.08  117  0086
     MACFAD73 STC10549 MACFAD7    LO  FF  188   0.00   0.00   0.00  118  0087
     MACFAD67 STC10570 MACFAD6    IN  C3  661   0.00   0.00   0.03  119  0088
     MACFAD6  TSU09810 MACFAD6    LO  FF  730   0.00   0.00   0.00  120  0089
     MACFAD5  TSU10539 MACFAD5    IN  C2 1395   0.00   0.06   0.04  121  008A
     MACFAD63 STC10579 MACFAD6    IN  C2  734   0.00   0.00   0.16  124  008B
     MACFAD62 STC10572 MACFAD6    IN  C1  682   0.00  77.27   0.33  127  008C
     MACFAD69 STC10449 MACFAD6    NS  C5  843   0.00   0.00   0.03  129  008D
     MACFAD64 STC10571 MACFAD6    IN  C2  745   0.00   1.19   0.24  130  008E
     MACFAD66 STC10569 MACFAD6    NS  C4  872   0.00   0.00   0.15  132  008F
```

Figure 3. SDSF Display of Domino Server Address Spaces

Then issue the OS/390 command `D OMVS,A=xxxx`, where `xxxx` is the ASIDX of that address space. Figure 4 shows the results of the command `D OMVS,A=84` for the job name MACFAD61 in the previous screen.

```
D OMVS,A=84

BPXO040I 17.36.15 DISPLAY OMVS 288
OMVS     000F ACTIVE          OMVS=(4D)
USER     JOBNAME  ASID        PID       PPID STATE   START     CT_SECS
MACFAD6  MACFAD66 0084 1073741841 1040187426 HS    17.32.00      3.951
  LATCHWAITPID=         0 CMD=/usr/lpp/lotus/notes/latest/os390/server
    MACFAD61 STC10586 MACFAD6    LO  FF  273   0.00   0.00   0.00  112
```

Figure 4. Program Name for a Domino Address Space

The function being performed by this Domino address space can be determined from the program name, which is the last qualifier in the command displayed (CMD in the D OMVS results). Here the address space is executing the server program.

Note: If the path name is too long, the CMD field may not be large enough to display the full program name.

A list of the program names is in 5.3, “Domino Server Tasks” on page 45. For a full list refer to the *Lotus Notes Administrator's Guide*. Note that the program names may appear in upper or lower case, but they are the same program. For example, the server program appears as `SERVER` in the SMF Type 30 record, but as `server` in the CMD field in the D OMVS command. An uppercase name indicates that the named program is running from LPA. See 2.7.3.1, “Domino Program Modules in LPA” on page 14.

In a future release it is planned to change the Domino server to create predictable job names. This will allow you to identify the function performed by an address space more easily.

2.3 Multiple Domino Servers

Multiple Domino servers can run on a single OS/390 image. An instance of a Domino server is called a Domino partition. Each Domino partition has the set of address spaces just described. Each Domino partition is independent of other Domino partitions, with its own address spaces and HFS files. Two Domino partitions use TCP/IP to communicate and transfer data. This means they are not dependent on being in the same image or processor.

This gives you a high degree of flexibility on where to run a Domino server. It can easily be moved from one OS/390 image to another. The files corresponding to the Domino partition, both Domino system files and user files, can be unmounted from the HFS directory on one OS/390 image and mounted in the directory of another OS/390 image.

As the number of users increases, it will be necessary to create more Domino partitions. The recommendations for the number of users in each Domino partition are in 6.4, “The Use of Domino Partitioned Servers” on page 69. By splitting the users across several Domino partitions you reduce the level of locking and the size of queues managed within a single Domino server. We have found this to be the most efficient way of running Domino servers for a large user population.

You can run multiple Domino partitions in the same OS/390 image, or in different PR/SM Logical Partitions (LPARs) on a single processor, or spread across more than one processor. You can configure the Domino servers to meet the capacity required to support your Notes clients.

2.4 Domino on Symmetric Multiprocessors

Domino is ideally suited to drive multiple engines on processors such as the 10-way CMOS processors. Domino uses a number of address spaces to perform server functions, and multiple dispatchable units within some of its address spaces. You can specify that more than one copy of a server function is started in the `notes.ini` initialization parameters, if required.

The busiest address spaces are usually the Domino user address spaces. Since the connection for each Notes client runs under its own TCB, the workload for the connected users can be spread across the engines of the processor.

TCP/IP also uses multiple TCBs to process its workload and can thus exploit multiple processor engines.

2.5 Setting OS/390 Parameters

It is important to set up OS/390 and OS/390 UNIX Services to allow for the increased number of address spaces and the large number of tasks (threads) created by a Domino server. Since Domino also makes extensive use of storage, it is important that the OS/390 UNIX Services storage parameters in SYS1.PARMLIB(BPXPRMxx) have the recommended values.

The recommended parameter values are listed in *Lotus Domino for S/390 Install Guide for Servers*. They are discussed in detail in *Lotus Domino for S/390 Release*

4.5: Installation, Customization and Administration, SG24-2083. Table 1 on page 9 lists the parameters that you should check.

Table 1. OS/390 Parameters

Parameter type	SYS1.PARMLIB member	Parameters relating to address spaces and tasks	Parameters relating to storage
OS/390	IEASYSxx	MAXUSER	SQA
OS/390 UNIX Services	BPXPRMxx	MAXPROCSYS MAXPROCUSER MAXTHREADTASKS MAXTHREAD IPCSEMNOPS MAXCPUTIME	MAXSHAREPAGES IPCSEMNIDS IPCSHMNIDS IPCSHMMPAGES IPCSHMNMSEGS IPCSHMSPAGES FORKCOPY MAXASSIZE
APPC (OS/390 Version 1 Release 3 only)	ASCHPMxx	CLASSADD MAX CLASSADD MIN TPDEFAULT TIME	

2.6 Managing the Domino Workload

Your Domino servers should get the same consideration that you would give a production CICS or IMS system. It should have the resources required to give good response times to the user.

We recommend that you use the Workload Manager in goal mode to manage your OS/390 system when running a Domino server. The Workload Manager can make sure that your Domino server receives the resources it requires.

It is not possible, using Workload Manager goal mode or compatibility mode, to single out individual Domino functions for specific treatment. In fact, it is recommended that you do *not* attempt to single out individual Domino address spaces for higher or lower priority. Domino uses extensive locking between tasks, so all tasks in a Domino server should run at the same priority.

You can manage and report Domino server partitions separately, if each Domino server is started under a different user ID.

If you plan to run the Domino server and other high priority workloads in the same OS/390 image, ensure that you have enough processor capacity to meet the needs of all workloads without impacting the responsiveness of each.

It is important to maintain the response time of the Domino server to the Notes client. If the Domino server does not respond within a specified period, as set on the Notes client, then the client treats this as though the server is not available. This TCP/IP session timeout is specified in the Notes Ports dialog box on the client. (see the *Lotus Notes Network Configuration Guide*). The default value is 5 seconds. The user will then have to decide whether to retry the request, switch to another server or work locally on the client. This situation should be avoided.

2.6.1 Workload Manager Goal Mode

OS/390 Workload Manager goal mode is the recommended way to manage Domino servers. Even if you are only running a single Domino server on OS/390, goal mode simplifies the setup of your system to meet its performance objectives and avoids some potential errors if you have to set parameters in compatibility mode. Within IBM we use Workload Manager to control our Domino servers for a large user population.

The Domino server should be placed in a service class with a high velocity. We recommend at least 60 for the velocity. The Domino server should also have a high importance relative to other workloads.

When setting up Workload Manager for a Domino server, you need to:

- Define a workload for OS/390 UNIX Services work.
- Define service classes for OS/390 UNIX Services work:
  - OS/390 UNIX Services startup processes which run under the USERID OMVSKERN, in the OMVS subsystem.
  - Forked address spaces, such as Domino server address spaces, which run under the user ID used for the Domino server, in the OMVS subsystem.
- Define classification rules for the OMVS subsystem:
  - Start up processes under USERID OMVSKERN
  - Forked address spaces, such as the Domino servers

  Note: OMVS (OS/390 UNIX Services Kernel address space), SYSBMAS (DFSMS buffer manager) and BPXOINIT (OS/390 UNIX Services initialization process) should take the started task default priorities.

An example of the Workload Manager definitions for a set of production Domino servers can be found in Appendix A, “Workload Manager Example” on page 75.

2.6.2 Compatibility Mode

If you are not running in Workload Manager goal mode, you have to take extra care with your compatibility mode parameters. The Domino server can create a number of address spaces, both at initialization and dynamically during the running of the server. There can be over 30 address spaces for a Domino partition supporting many users. You need to make sure that your IPS specifications in SYS1.PARMLIB take these additional address spaces into account.

2.6.2.1 Multiprogramming Level (MPL)

A Domino address space can be logically swapped because of detected wait. This means that the address space's storage is retained in memory but the address space is removed from the OS/390 in queue. When the address space is ready to run, the working set should still be in memory unless the logical swap is converted to a physical swap because of memory constraints. However, it is important to make sure that a Domino address space is not physically swapped to the page data sets on disk. This delays the execution of the address space during the swap in and also as it builds up its working set in central storage.

In compatibility mode the multiprogramming levels are controlled by the MPL parameter in IEAIPSxx in SYS1.PARMLIB. You must make sure that the maximum

MPL for the domain is much larger than the number of address spaces the server will create. You should allow for growth in the number of address spaces as the use of the server increases. If you have more address spaces than the maximum MPL, the SRM will have to keep swapping Domino address spaces to keep the number below this maximum. If a Domino address space is swapped, response time to the user will be severely affected and the Notes clients may time out.

Important: Make sure that you avoid the situation where the number of Domino address spaces exceeds the maximum MPL for the domain. In those circumstances the Domino server will not perform well, or it may not even start successfully.

2.6.2.2 Priority

Make sure that the Domino server is not waiting for CPU because of higher priority tasks. The Domino server address spaces should be above any batch processing, TSO (other than maybe a short period 1) and other online systems that consume large amounts of CPU.

Only workloads that act as servers to Domino should be at a higher priority. This includes:

- TCP/IP
- OMVS (OS/390 UNIX Services Kernel address space)
- SYSBMAS (buffer management address space)
- BPXOINIT (OS/390 UNIX Services initialization process)
- OS/390 system address spaces
- Monitors such as RMF

If you have two online production workloads, such as CICS and a Domino server, you will have to make a business decision on their relative importance.

The priority of address spaces in compatibility mode is controlled by the DP parameter in IEAIPSxx in SYS1.PARMLIB. The absolute priority of an address space is not important. What is important is where the address space sits in the order of priorities of all the address spaces in the OS/390 image.

2.6.2.3 Storage Isolation

The Workload Manager in goal mode manages the working set of central and expanded storage for you. In compatibility mode, when there is not enough free storage available OS/390 can take storage from Domino servers if they are not protected. The way to protect the Domino server address spaces is to use storage isolation.

Storage isolation attempts to protect the real memory working set of selected workloads. If you use storage isolation, then OS/390 prefers to take real memory from address spaces that do not have storage isolation.

For storage isolation you specify the minimum and maximum working set size. We recommend that you specify a minimum value and not control the maximum, by setting it to “*.” You can specify that OS/390 controls the working set size depending on the paging rate of the address space, and this is recommended. You should try to achieve zero paging for the Domino address spaces. The working set controls for an address space are PWSS and PPGRTR in IEAIPSxx in SYS1.PARMLIB.

Storage isolation is important if you are short of central or expanded storage. It is also required if you use storage isolation for other workloads, otherwise OS/390 steals pages from the Domino server and gives them to the other workloads.

2.6.2.4 Setting the Definitions

To add the definitions for running a Domino server in compatibility mode, you need to update the IEAICSxx and IEAIPSxx members of SYS1.PARMLIB as follows:

- Add a subsystem definition, SUBSYS=OMVS, for OS/390 UNIX Services address spaces.
- In the OMVS subsystem, define a USERID=OMVSKERN for OS/390 UNIX Services initialization processes and give it a high priority.
- Also define the USERIDs you will run the Domino servers under and give them sufficiently high priority, as discussed.

  Note: OMVS (OS/390 UNIX Services Kernel address space), SYSBMAS (DFSMS buffer manager) and BPXOINIT (OS/390 UNIX Services initialization process) should take the started task default priorities.

2.7 Domino's Use of Memory

In 2.2, “OS/390 Address Spaces” on page 4 we saw that a Domino server creates a large number of address spaces. It can have more than 30 address spaces when supporting a large number of users. Some of these address spaces are large. They use a large amount of virtual memory, particularly in extended private (above 16MB addressing) memory.

Domino also uses large amounts of common storage, also above 16MB, in the Extended System Queue Area (ESQA). You need to make sure that you have defined sufficient ESQA in the IEASYSxx member of SYS1.PARMLIB. A guideline for the additional ESQA required for the Domino server can be found in 6.3, “Estimating Processor Storage” on page 68.

2.7.1 Real Memory

With the large virtual memory requirement for a Domino server it is not surprising that it also requires a large amount of real memory. This real memory can be a mixture of central and expanded storage. Make sure you have sufficient central storage to avoid excessive page movement between central and expanded storage.

You should not attempt to add Domino servers to an OS/390 image that does not have sufficient real memory to support the combined workloads. It would cause paging to the page data sets and affect the Domino servers and the other workloads on the OS/390 image.

When you have several large Domino partitions in a single OS/390 image, or where a Domino server is added to an OS/390 image that is already using large amounts of central storage, you may reach the limit of 2GB of central storage for an OS/390 image. You should then use expanded storage or split the workloads over more than one LPAR. A processor can be configured with more than 2GB of central storage, but a single LPAR can only have up to 2GB of central storage.

The real memory required to support a number of Domino users is given in 6.3, “Estimating Processor Storage” on page 68. The real memory required is

influenced by the number of address spaces and the working set size of these address spaces. The working set size is the number of 4K-byte pages that the program frequently references. These pages include instructions within the programs and the data. This is smaller than the virtual memory size, because there are parts of programs and data that are not used during normal running of the server. These include parts used only at Domino startup, infrequently used functions and error routines.

The unreferenced pages will be stolen by OS/390 for use by other address spaces. If the OS/390 image has far more storage than is required to run its workloads, then these pages are not stolen and the working set remains larger than is actually required.

When we ran two Domino server partitions, with a large number of users, on a system that had more real memory than required, we saw the memory use shown in Table 2. However, note that this system was not constrained. If additional workloads were added to this OS/390 image, then the working sets of these address spaces would become smaller.

Table 2. Memory Use on a Production System

Address spaces	**Working set size (MB)**
OMVS	308
SYSBMAS	308
TCP/IP	182
Domino server 1	173
Domino server 2	183
ESQA	133

2.7.2 Page Data Sets

Because of the large virtual memory requirement for Domino, there will be a significant increase in page space needed. Therefore you should evaluate the size and number of page data sets on the system. The page data sets must be able to hold the virtual memory requirements of the address spaces in the OS/390 image. The size of the page data sets and the number used can be found from the PAGESP report in RMF.

Always have far more page data set space than you require for the normal running of your OS/390 system. If you have an increased demand for real memory, such as when OS/390 takes a system dump of a large address space or part of the real memory goes offline, then you must have enough page data set space to handle these situations.

2.7.3 Place Program Modules in LPA

We recommend that you place certain Domino and C/C++ Runtime Library programs in the Link Pack Area (LPA). (Nearly all the modules go above 16MB in the Extended Link Pack Area (ELPA).) This reduces the virtual and real memory used. When multiple address spaces use the same program, it is more efficient to place a single copy in common storage than a copy in each address space.

2.7.3.1 Domino Program Modules in LPA

The Domino server installation process provides a PUTINLPA job, which places certain modules in LPA. It is important that you run this job. If the Domino server modules are not in LPA, approximately 20MB of program modules will be loaded into each address space. As we have seen, Domino uses many address spaces, so this could add up to a significant storage requirement that would have a negative impact on system performance.

The procedure to place these modules in LPA is described in *Lotus Domino for S/390 Install Guide for Servers* and also in *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083, chapter 5 “Install the Domino Server Code.”

You should also check that the sticky bit is set on for those programs in the Hierarchical File System (HFS), as described in these books. With the sticky bit on, the module is accessed from LPA rather than loaded from the HFS.

To make absolutely sure that you do not load the modules from the HFS into the Domino address spaces, you can replace the modules in the HFS with dummy modules. Just take a backup copy of the real modules, then replace them with dummy files with the same permission bit settings. Make sure that these dummy files have the sticky bit on. If everything is working correctly, the original programs are loaded from LPA. However, if there is an error, such as if the modules are removed from the LPA data sets or if the HFS sticky bits are set off, then the dummy modules will be loaded. Since these are not valid modules, this will cause an error in the Domino server. This may be preferable to having the modules loaded into many address spaces, which will degrade system performance but will give no indication of the problem.

2.7.3.2 C/C++ Runtime Library Program Modules

For a list of the C/C++ Runtime Library program modules to place in LPA, and how to do it, see *OS/390: OpenEdition Planning*, SC28-1890, chapter 16 “Tuning OpenEdition Performance.”

C/C++ Runtime Library programs that you do not put into LPA should be cached in the Virtual Lookaside Facility (VLF).

2.8 Monitoring OS/390

An essential part of running any application is to be able to monitor the resources consumed. This provides input into both performance tuning and capacity planning. In this section we discuss the major sources of information about a Domino server from the OS/390 system on which it runs. We discuss Domino monitoring techniques in 5.11, “Monitoring Domino” on page 58.

Traditional OS/390 information sources provide a lot of information for capacity planning, resource accounting and performance tuning. These have been enhanced for applications such as Domino that use the OS/390 UNIX Services interfaces. First we discuss RMF, which provides system level information that is useful for capacity planning and tuning the system. Then we look at the data available in SMF, which is useful for monitoring Domino resource use, both CPU and I/O.

2.8.1 Monitoring OS/390 with RMF

Resource Management Facility (RMF) is used to monitor the activity of the OS/390 system. It is useful for capacity planning, performance tuning and problem diagnosis. RMF gives information about system resources, workload groups, DASD controllers and DASD volumes.

There is a lot of detailed information available with RMF, but we focus here on the OS/390 UNIX Services OMVS report, some key system indicators, and useful data about a Domino server workload.

Details on producing RMF reports and interpreting them are found in:

- *OS/390 Resource Measurement Facility User's Guide*, SC28-1949
- *OS/390 Resource Measurement Facility Report Analysis*, SC28-1950

RMF has two online monitors, which you can access from TSO:

- The RMFMON (Monitor II) command shows data about processor CPU, storage, address spaces and DASD volumes.
- The RMFWDM (Monitor III) shows how an address space is being delayed and what is causing the delay. RMFWDM keeps information in storage or on DASD so you can look at delays after the event.

2.8.1.1 RMF OMVS Report - OS/390 UNIX Parameters

A new OMVS Kernel Activity report is included for OS/390 UNIX Services. RMF Monitor III must be active to collect the information for the OMVS report.

This report is useful to see if you have specified values in the OS/390 UNIX Services SYS1.PARMLIB(BPXPRMxx) member that are too low. It shows the maximum values you have set and also minimum, average and maximum values observed during the RMF reporting interval. It also shows the number of overruns, that is, the number of times processes could not be created or storage was not available.

Figure 5 on page 16 is an example of an OMVS Kernel Activity report. Note that this report was run on our test system and therefore shows very little activity.

```
                                        O M V S   K E R N E L   A C T I V I T Y

           OS/390                  SYSTEM ID SC52            START 03/11/1998-15.00.00  INTERVAL 001.00.00
           REL. 02.04.00           RPT VERSION 2.4.0         END   03/11/1998-16.00.00  CYCLE 1.000 SECONDS

TOTAL SAMPLES =  3580

                                                OMVS SYSTEM CALL ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------
                  MINIMUM  AVERAGE  MAXIMUM
-------------------------------------------------------------------------------------------------------------------------------
 SYSCALLS (N/S)   0.000    67.73    759.0
 CPU TIME (H/S)   0.000    0.081    1.000

                                               OMVS PROCESS ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------
                          PROCESSES                        USERS                    PROCESSES PER USERS
 MAXIMUM  (TOT)             300                              50                       10125
-------------------------------------------------------------------------------------------------------------------------------
                  MINIMUM  AVERAGE  MAXIMUM      MINIMUM  AVERAGE  MAXIMUM      MINIMUM  AVERAGE  MAXIMUM
-------------------------------------------------------------------------------------------------------------------------------
 CURRENT  (TOT)       51    51.24       53            0    0.000        0
 OVERRUNS (N/S)    0.000    0.000    0.000        0.000    0.000    0.000        0.000    0.000    0.000

                                            OMVS INTER-PROCESS COMMUNICATION
-------------------------------------------------------------------------------------------------------------------------------
                    MESSAGE QUEUE IDS             SEMAPHORE IDS              SHARED MEMORY IDS            SHARED MEMORY PAGES
 MAXIMUM  (TOT)             20000                        20000                      20000                        2621K
-------------------------------------------------------------------------------------------------------------------------------
                  MINIMUM  AVERAGE  MAXIMUM    MINIMUM  AVERAGE  MAXIMUM    MINIMUM  AVERAGE  MAXIMUM    MINIMUM  AVERAGE  MAXIMUM
-------------------------------------------------------------------------------------------------------------------------------
 CURRENT  (TOT)    5.000    4.972    5.000      121.0    120.9    122.0      8.000    8.116    9.000       1920    8.116     2165
 OVERRUNS (N/S)    0.000    0.000    0.000      0.000    0.000    0.000      0.000    0.000    0.000      0.000    0.000    0.000

                                                  OMVS MEMORY MAP
-------------------------------------------------------------------------------------------------------------------------------
                  MEMORY MAP STORAGE PAGES        SHARED STORAGE PAGES
 MAXIMUM  (TOT)             4096                         32.8M
-------------------------------------------------------------------------------------------------------------------------------
                  MINIMUM  AVERAGE  MAXIMUM    MINIMUM  AVERAGE  MAXIMUM
-------------------------------------------------------------------------------------------------------------------------------
 CURRENT  (TOT)    435.0    432.8    436.0      24928    25273    29167
 OVERRUNS (N/S)    0.000    0.000    0.000      0.000    0.000    0.000

 Units:   (TOT) = Total Value, (N/S) = Number per Second, (H/S) = Hundredths of seconds per Second
```

Figure 5. RMF Kernel Activity Report

The RMF reporting interval is typically set at 15 minutes. You are only interested in whether you are exceeding the values in BPXPRMxx and not in how the values change during the day. We suggest that you run an RMF report just for the OMVS Kernel Activity specifying a complete day, or more, in the DINTV parameter. The DINTV has a maximum value of 100 hours.

You should check periodically that the maximum used values (CURRENT (TOT) in the MAXIMUM column) are well below (MAXIMUM (TOT)) in all the process and interprocess communication and memory map sections. The MAXIMUM (TOT) fields are the values from the active BPXPRMxx in SYS1.PARMLIB. You should not have any overruns (OVERRUNS (N/S) in the MAXIMUM column) when running a Domino server, since this would cause problems with Domino unable to start processes or obtain storage.

2.8.1.2 RMF Summary Report - Overall System Performance

You should monitor the same things for Domino servers as for other workloads that require good response times. Areas to consider are CPU, DASD response times, paging and storage.

The RMF Summary Report provides a useful overview of the system. An example is shown in Figure 6 on page 17.

```
                                    R M F   S U M M A R Y   R E P O R T

           OS/390                  SYSTEM ID SYS1           START 02/26/1998-09.00.00  INTERVAL 00.14.59
           REL. 02.04.00           RPT VERSION 2.4.0        END   02/26/1998-17.00.00  CYCLE 1.000 SECONDS

DATE   TIME     INT   CPU   DASD  DASD  TAPE   JOB   JOB   TSO   TSO   STC   STC  ASCH  ASCH  OMVS  OMVS SWAP DEMAND
MM/DD HH.MM.SS MM.SS  BUSY  RESP  RATE  RATE   MAX   AVE   MAX   AVE   MAX   AVE   MAX   AVE   MAX   AVE RATE PAGING

02/26 09.00.00 14.59  62.2     8 252.5   7.0     2     0     7     6   101   101     0     0    49    49 0.00   0.00
02/26 09.15.00 15.00  59.4     9 278.4   6.6     3     1     8     7   102    99     0     0    49    49 0.00   0.07
02/26 09.30.00 15.00  66.0     9 241.8   6.1     0     0     8     8   100   100     0     0    49    49 0.00   0.00
02/26 09.45.00 15.00  66.2     8 238.5   5.1     0     0     9     9   101   101     0     0    50    49 0.00   0.00

02/26 10.00.00 14.59  69.0     7 212.3   4.9     0     0     9     9   101   101     0     0    50    49 0.00   0.00
02/26 10.15.00 14.59  66.9     7 204.5   8.1     1     0     9     9   103   100     1     0    52    50 0.00   0.00
02/26 10.30.00 15.00  69.6     7 219.7   8.5     1     0     9     9   104   103     0     0    53    50 0.00   0.00
02/26 10.45.00 14.59  67.7     6 222.2   4.9     1     0     9     9   103   103     0     0    59    54 0.00   0.01

02/26 11.00.00 15.00  67.7     8 294.0  10.4     2     0     8     8   107   106     0     0    60    56 0.00   0.02
02/26 11.15.00 15.00  66.0     5 242.5   7.1     1     0     9     9   107   107     0     0    54    54 0.00   0.00
02/26 11.30.00 15.00  61.7     4 260.6   9.5     1     0     9     9   106   104     0     0    54    54 0.00   0.01
02/26 11.45.00 14.59  61.1     5 238.8   8.0     1     0    10    10   101   101     0     0    55    54 0.00   0.02

02/26 12.00.00 14.59  58.3     7 221.5   7.8     0     0    11    10   103   103     0     0    54    52 0.00   0.01
02/26 12.15.00 14.59  62.7     5 219.4   8.1     0     0    10    10   103   103     0     0    52    52 0.00   0.00
02/26 12.30.00 15.00  62.6     5 182.1   7.8     0     0    11    10   103   101     0     0    52    52 0.00   0.00
02/26 12.45.00 14.59  62.7     4 239.4   8.7     0     0    10    10   101   101     0     0    53    52 0.00   0.00
```

Figure 6. RMF Summary Report

You should produce the summary report for the whole day and check that:

- CPU BUSY is not too high.

  If the utilization is high, over 80% during peak periods, you need to make sure that the Domino server is set to a high enough priority so that it is not impacted. If the utilization reaches 100%, then you are running more work than the OS/390 image can support and workloads are suffering. Make sure this is not the Domino server, by making it a high priority workload.

- DASD average response time is not high, preferably less than ten milliseconds.

- DEMAND PAGING is zero.

  If there is paging, make sure that the Domino address spaces and the common storage areas are not impacted by using storage isolation in compatibility mode.

2.8.1.3 RMF LPAR Report

When running in logical partition (LPAR) mode, make sure the processor is not at high utilization. Also make sure that the LPAR has a sufficient share of the processor so that it is not being constrained by other LPARs. You can see this on the Partition Data Report, shown in Figure 7 on page 18.

```
                                    P A R T I T I O N  D A T A  R E P O R T

           OS/390                  SYSTEM ID PROD           START 02/26/1998-10.00.00  INTERVAL 000.14.59
           REL. 02.04.00           RPT VERSION 2.4.0        END   02/26/1998-10.15.00  CYCLE 1.000 SECONDS


MVS PARTITION NAME                PROD

NUMBER OF CONFIGURED PARTITIONS          4

NUMBER OF PHYSICAL PROCESSORS           10

WAIT COMPLETION                         NO

DISPATCH INTERVAL                  DYNAMIC

--------- PARTITION DATA ---------    -- LOGICAL PARTITION PROCESSOR DATA --   ---- AVERAGE PROCESSOR UTILIZATION PERCENTAGES ----

                                     NUMBER OF  ----DISPATCH TIME DATA-----   LOGICAL PROCESSORS  ----- PHYSICAL PROCESSORS -----
  NAME    STATUS  WEIGHTS  CAPPING   LOG PRCRS   EFFECTIVE        TOTAL      EFFECTIVE    TOTAL  LPAR MGMT   EFFECTIVE     TOTAL

TEST         A       10         NO          3   00.01.42.824   00.01.51.793       3.81     4.14      0.10        1.14      1.24

PROD         A       35         NO          7   01.11.35.588   01.12.26.860      68.20    69.02      0.57       47.74     48.31

SYSP         A       15         NO          2   00.03.37.460   00.03.47.129      12.08    12.62      0.11        2.42      2.52

DEVT         A       35         NO          7   01.08.24.722   01.09.08.900      65.17    65.87      0.49       45.62     46.11

*PHYSICAL*                                      00.01.10.252                         0.78                  0.78

                                            ------------   ------------                      ------      ------    ------

  TOTAL                                     02.25.20.596   02.28.24.937                      2.05       96.92     98.97
```

Figure 7. RMF Partition Data Report

Create the report at times of peak loading for the processor and check the following:

- PHYSICAL PROCESSORS TOTAL (TOTAL line)

  This shows how busy the processor is. If it is at or near 100%, then the LPARs will receive resources according to their weights. If the OS/390 image is in an LPAR with dedicated engines, then you should consider only the LPAR utilization.

- WEIGHTS

  The weights are used to allocate CPU resources between LPARs that are sharing engines. These weights come into effect when the engines assigned to shared LPARs are 100% busy. An LPAR receives its weight divided by the sum of the weights of the active LPARs. In Figure 7 the PROD LPAR would receive 35/(10+35+15+35), which is 36.8% of the shared engines.

- NUMBER OF LOG PRCRS

  Make sure that you are not constraining an LPAR by giving it too few Logical Processors (engines). In Figure 7 we have 10 engines on the processor, but the PROD LPAR has only 7 engines specified and so can never use more than 70% of the processor, even if the other LPARs are idle.

- CAPPING

  It is not recommended to cap an LPAR running a Domino server. You could affect the server response time when CPU resource is available, by restricting the LPAR to its specified weight.

You can find more information about LPARs in *ES/9000 and ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123.

2.8.1.4 RMF Processor Storage Report

The key indicators of central and expanded storage are Unreferenced Interval Count (UIC) and Migration Age. Figure 8 shows the RMF Paging Activity Report which reports these indicators.

```
                                        P A G I N G   A C T I V I T Y

           OS/390                  SYSTEM ID PROD            START 02/26/1998-10.00.00  INTERVAL 000.14.59
           REL. 02.04.00           RPT VERSION 2.4.0         END   02/26/1998-10.15.00  CYCLE 1.000 SECONDS

OPT = IEAOPT00                          EXPANDED STORAGE MOVEMENT RATES - IN PAGES PER SECOND
-------------------------------------------------------------------------------------------------------------------------------

     ESF CONFIGURATION                                                                                          HIGH UIC  MIGR AGE
   INSTALLED      ONLINE                                                                                      MIN   254   186628
   ---------      ------                                                                                      MAX   254   187944
    327680       327680                                                                                      AVG   254.0 187285.1
                                                 MIGRATED       FREED
                     WRITTEN TO     READ FROM       FROM     WITHOUT      *----- EXPANDED STORAGE FRAME COUNTS -----*
                       EXP STOR      EXP STOR    EXP STOR   MIGRATION              MIN           MAX           AVG

  TOTAL         RT         36.74          3.64       0.00        0.00           80,670        90,359        83,834
   PAGES         %        100.0%        100.0%      100.0%

  HIPERSPACE    RT          0.00          0.00       0.00                          30            30            30
   PAGES         %          0.0%          0.0%       0.0%

  VIO           RT          0.00          0.00       0.00                           0             0             0
   PAGES         %          0.0%          0.0%       0.0%

  SHARED        RT          0.00          0.00
   PAGES         %          0.0%          0.0%

                                                FRAME AND SLOT COUNTS
-------------------------------------------------------------------------------------------------------------------------------

                     CENTRAL STORAGE              EXPANDED STORAGE              LOCAL PAGE DATA SET SLOT COUNTS
                    ---------------------------   ---------------------------   -----------------------------------------------
   (91 SAMPLES)          MIN       MAX       AVG       MIN       MAX       AVG                          MIN         MAX         AVG

   AVAILABLE             159     2,758       544   237,321   247,010   243,846   AVAILABLE SLOTS    864,000     864,000     864,000

   SQA                30,707    33,419    31,404         0         0         0   VIO SLOTS                0           0           0
   LPA                 6,966     6,967     6,966        29        29        29
   CSA                 1,993     2,118     1,996         0        17        16   NON-VIO SLOTS            0           0           0
   LSQA               21,448    21,966    21,585       689       700       693
   REGIONS+SWA       453,696   457,261   456,080    80,583    90,079    83,641   BAD SLOTS                0           0           0

   TOTAL FRAMES      524,288   524,288   524,288   327,680   327,680   327,680   TOTAL SLOTS        864,000     864,000     864,000

                        FIXED FRAMES                                                  SHARED FRAMES AND SLOTS
                    ---------------------------                                   -----------------------------------------------

   NUCLEUS             5,709     5,709     5,709                                   CENTRAL STORAGE    141,669     143,172     142,076
   SQA                30,200    32,912    30,897                                   EXPANDED STORAGE       122         122         122
   LPA                    85        85        85                                   FIXED TOTAL              0           0           0
   CSA                    39        39        39                                   FIXED BELOW 16 M         0           0           0
   LSQA               17,572    17,987    17,679                                   AUXILIARY SLOTS          0           0           0
   REGIONS+SWA         1,175     1,347     1,236
   BELOW 16 MEG          573       581       573                                   TOTAL              160,855     160,856     160,856

   TOTAL FRAMES       54,811    58,062    55,645
```

Figure 8. RMF Paging Activity Report

You should regularly monitor the following:

- HIGH UIC AVG

  The UIC is a measure of how central storage pages are referenced. This value should be at or near 254. The higher the number, the better. If it is in single

figures, then the OS/390 image requires more central storage. If it less than 50 you should consider increasing central storage.

- MIGR AGE AVG

  The migration age is the length of time, in seconds, a page of storage remains in expanded storage before migrating to the page data sets. This number should be in the thousands of seconds. A low number means expanded storage is overcommitted and a very high number means it is underutilized.

In Workload Manager goal mode the storage used by address spaces is managed to achieve the velocity or response time goals requested. In compatibility mode, if you are storage constrained, you will need to protect the Domino server address spaces by using storage isolation. This is discussed in 2.6.2.3, “Storage Isolation” on page 11.

2.8.1.5 RMF Workload Reports

Whether running in goal mode or compatibility mode, you should define your workloads to report each Domino partition separately. You can run all the Domino partitions under the same priority level (Service Class or Performance Group), but should place each Domino partition in its own reporting group (Report Class or Report Performance Group). This allows you to monitor the resources used by each Domino partition.

Appendix A, “Workload Manager Example” on page 75 has an example of the Workload Manager goal mode definitions using Report Classes.

In Workload Manager goal mode, the SYSRPTS(WLMGL(SCLASS,RCLASS)) RMF report gives the CPU used by service classes and report service classes. Figure 9 shows an example of a report class.

```
      OS/390                  SYSPLEX SYSPRIM           START 02/26/1998-10.15.00 INTERVAL 000.14.59   MODE = GOAL
      REL. 02.04.00           RPT VERSION 2.4.0         END   02/26/1998-10.29.59

                                      POLICY ACTIVATION DATE/TIME 02/19/1998 17.35.42


REPORT BY: POLICY=PRODPRM                        REPORT CLASS=DOMP01
                                                        Domino server

TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----   --SERVICE RATES--   PAGE-IN RATES    ----STORAGE----
AVG     18.39   ACTUAL                162   SSCHRT   0.0   IOC     43569    ABSRPTN       247   SINGLE     0.0   AVG     3200.80
MPL     18.39   EXECUTION             158   RESP     1.2   CPU      3912K   TRX SERV      247   BLOCK      0.0   TOTAL   58870.1
ENDED       9   QUEUED                  3   CONN     0.9   MSO         0    TCB        1553.5   SHARED     0.0   CENTRAL 58073.0
END/SEC  0.01   R/S AFFINITY            0   DISC     0.0   SRB    125545    SRB          49.9   HSP        0.0   EXPAND   797.14
#SWAPS    475   INELIGIBLE              0   Q+PEND   0.3   TOT      4081K   RCT           0.8   HSP MISS   0.0
EXECUTD     0   CONVERSION              0   IOSQ     0.0   /SEC     4534    IIT           0.0   EXP SNGL   0.0   SHARED   903164
                STD DEV               115                                  HST           0.0   EXP BLK    0.0
                                                                           APPL %      178.3   EXP SHR    0.0

VELOCITY MIGRATION:   I/O MGMT  78.7%    INIT MGMT 78.7%

        ---RESPONSE TIME---  EX   PERF   AVG   --USING%--  --------------- EXECUTION DELAYS % ----------------  ---DLY%--    %
        HH.MM.SS.TTT         VEL  INDX  ADRSP   CPU   I/O  TOTAL   CPU                                          UNKN IDLE QUIE
GOAL
ACTUALS
PROD                        78.7%  N/A   22.0  11.3   0.0    3.1   3.1                                          60.0 25.7  0.0
```

Figure 9. RMF Workload Manager Report Class

Under the Service Rates column you see the TCB and SRB (which are the CPU seconds used by this workload during the interval), as well as the APPL% (which is the percentage of the interval when the workload was processing). For a Domino server you may see APPL% above 100%, because the server has multiple address

spaces with many tasks that can be executing on more than one engine concurrently. This report can be used to obtain the CPU time used by a Domino server.

In compatibility mode, the REPORTS(WKLD(GROUP)) report provides similar information (excluding the velocity migration and execution delays). You need the performance group number or report performance group number used by the Domino partition to identify the correct set of data.

You can also obtain the central and expanded storage used by the address spaces in this workload group from the same report. However, these are weighted numbers based on the time during the interval that the address spaces were resident (not swapped out). The actual storage required by the address spaces may be higher when all the address spaces are resident.

2.8.2 Monitoring OS/390 with SMF

System Management Facility (SMF) is an important source of information on Domino resource use. SMF collects data about processor, storage and I/O consumption for address spaces in OS/390.

Details of the SMF records and how to manage them are found in *OS/390 MVS System Management Facilities (SMF)*, GC28-1783.

The RMF monitor also writes its records to the SMF data set as record types 70 through 79. The RMF record types are discussed in 2.8.1, “Monitoring OS/390 with RMF” on page 15.

RMF reports can be created using the RMF post processor, but no utility is provided to report on the other SMF records. Many installations have products that can report on SMF data. You can also write a program to produce a desired report, although the format of the SMF records is complicated.

We discuss some useful SMF record types here.

2.8.2.1 Interval Recording of SMF Records

You specify interval recording in the SMFPRMxx member of SYS1.PARMLIB, where xx is the suffix of the currently active member. You can have interval accounting for all workload types, or you can select specific workload types. You will need interval recording enabled for the OMVS workload type.

Figure 10 on page 22 shows an example of the SMFPRMxx parameters. You should use this as a guide only, since you may require your own values for other parameters. The important parameters are INTVAL, SYNCVAL and INTERVAL. In the example, the interval is set to 15 minutes, but you can choose the frequency of reporting you require. If a job step starts at 10:50, interval records are written at 11:00, 11:15, 11:30 and so on until the step ends, when a final interval record is written. Each interval record shows the resources consumed since the previous interval record was created.

```
ACTIVE                        /* ACTIVE SMF RECORDING             */
DSNAME(SYS1.&SYSNAME..MAN1,   /* SMF DATA SET NAMES               */
       SYS1.&SYSNAME..MAN2,   /* SMF DATA SET NAMES               */
       SYS1.&SYSNAME..MAN3)   /* SMF DATA SET NAMES               */
NOPROMPT                      /* DON'T PROMPT THE OPERATOR        */
REC(PERM)                     /* TYPE 17 PERM RECORDS ONLY        */
INTVAL(15)                    /* INTERVAL RECORDING EVERY 15 MIN  */
SYNCVAL(00)                   /* SYNCHRONIZE TO THE HOUR          */
MAXDORM(3000)                 /* WRITE AN IDLE BUFFER AFTER 30 MIN */
STATUS(010000)                /* WRITE SMF STATS AFTER 1 HOUR     */
JWT(0800)                     /* 522 AFTER 2 HOURS                */
SID(&SYSNAME(1:4))            /* SYSTEM ID IS &SYSNAME            */
LISTDSN                       /* LIST DATA SET STATUS AT IPL      */
LASTDS(MSG)                   /* DEFAULT TO MESSAGE               */
NOBUFFS(MSG)                  /* DEFAULT TO MESSAGE               */

SYS(TYPE(0:255),EXITS(IEFU83,IEFU84,IEFU85,IEFACTRT,IEFUJV,IEFUSI,
               IEFUJP,IEFUSO,IEFUJI,IEFUTL,IEFU29),
               NOTYPE(99),INTERVAL(SMF,SYNC),DETAIL)
```

Figure 10. SMFPRMxx Example for SMF Interval Recording

2.8.2.2 Common Address Space Work (Record Type 30)

The type 30 record contains a lot of information about CPU and storage use by address space. These record types can be produced at specified intervals during an address space execution, which is very useful for the long running Domino address spaces. We recommend that you enable interval recording for SMF Type 30 records.

The Type 30 records can be quite large. Therefore, if you are not already collecting them, you should plan for your SMF data sets filling more rapidly and allow for this in your operational procedures.

There are many fields with similar types of data, for example CPU data, in the Type 30 record. The fields that we found useful are:

SMF30JBN JOB name

SMF30EXN Program name of the OS/390 UNIX Services process. Do not use SMF30PGM, which is too short to hold the OS/390 UNIX Services name. The program name in OS/390 UNIX Services can be up to 16 characters and so you must use SMF30EXN. The OpenMVS Accounting chapter in *OS/390 MVS System Management Facilities (SMF)*, GC28-1783 describes the format of the program name.

SMF30CPT CPU time under a TCB

SMF30CPS CPU time under an SRB

SMF30OFR and SMF30OFW
HFS reads and writes to regular files. There are also counts of reads and writes to HFS pipes and special files in the fields SMF30OPR, SMF30OPW, SMF30OSR and SMF30OSW.

SMF30TEP Blocks Transferred. All I/O to the HFS is reported in the OMVS OS/390 UNIX Services Kernel address space. The I/O is performed in cross memory mode under the TCB of the Domino

address space. The block count in the Domino address spaces is an indication of the blocks transferred for the Domino address space. The OMVS address space includes the blocks count for all OS/390 UNIX Services users on the OS/390 image. There is thus double counting of the block count for OS/390 UNIX Services I/O to DASD.

SMF30AIS Start Subchannel (SSCH). This field is zero for the OS/390 UNIX Services Domino address spaces. The OMVS address space record contains the number of I/Os performed for all OS/390 UNIX Services users. Also in the record are the connect, disconnect and pending times for these I/Os.

2.8.2.3 DFSMS Statistics and Configuration (Record Type 42)

The data set statistics, subtype 6, record provides useful information about OS/390 data sets. The OS/390 UNIX Services HFS file system is stored in OS/390 HFS data sets. This record therefore shows the activity to all of the UNIX files in each HFS data set.

By default, the records are written at data set close time. For an HFS data set this means from the time the data set was mounted until it is unmounted. That is not useful for performance purposes. To get significant statistics for the Domino address spaces, you should have interval recording active. The type 42 records are written at the same time as the SMF type 30 records.

These records contain useful information about I/O performance, including:

- Number of I/Os
- Average connect, disconnect, pending and control unit queue time, as well as response time
- Cache performance, including number of cache candidates and cache hit rates

You will find these records useful when you have an I/O performance problem and wish to find greater detail than the DASD volume information provided by RMF.

2.8.2.4 Open/MVS File System Activity (Record Type 92)

The type 92 record does not provide interval records as types 30 and 42 do. Information is reported for the duration of the mount for HFS data sets, and for the duration of the open for UNIX files.

Record subtype 5 contains information about the HFS file system; it is written when the file system is unmounted. This should not be a frequent occurrence with a Domino server, so there may be no useful information from this record because the activity it reflects covers a long time.

The UNIX file close record, subtype 11, provides useful information on the requests made to a UNIX file by the OS/390 UNIX Services Domino address spaces. It contains:

- Read and write calls
- Directory I/O blocks
- Blocks read and written
- Bytes read and written

Unfortunately, the record does not contain the UNIX file name, but there is a way to obtain which UNIX file the data is for. It is documented in *OS/390: OpenEdition Planning*, SC28-1890, chapter 15 “Monitoring the OpenEdition Environment” in the section “Using SMF Record Type 92.” IBM intends to add the UNIX file name to this record in a future release of OS/390, subject to other business priorities.

Collecting Record Type 92: Collecting type 92 records, particularly subtypes 10 and 11, generates a large number of SMF records. You may want to suppress these subtypes, in SMFPRMxx in SYS1.PARMLIB, until you need this detailed level of information. However you should note that to create these records:

- Subtype 5 must be collecting at the time the HFS file system is mounted and unmounted.
- Subtype 11 must be collecting at the time the UNIX file is opened and closed.

Chapter 3. DASD (Disk) Subsystem Tuning

In this chapter we discuss the performance of the DASD (or disk) subsystem, and how you can monitor and tune it.

I/O response time can be a significant part of application response time because, while processor activities take times in the order of nanoseconds, I/O activities take times in the order of milliseconds. Therefore, good disk response time is an important part of Domino performance.

3.1 Recommendations

These are the key DASD tuning recommendations:

- Monitor the system.

  Look for ESCON channels at more than 45% utilization, or parallel channels at more than 30% utilization. Watch for DASD devices with high activity that are not performing well.

- Use cached DASD devices.

  The OS/390 Hierarchical File System (HFS) writes all data to disk synchronously. Therefore, use DASD Fast Write and Dynamic Cache Management Enhanced (DCME) to give the fastest I/O response time.

- Avoid multiple extents in HFS data sets.

  HFS data sets can expand to multiple extents. As for all MVS data sets, consolidate them to one extent using DFSMSdss. This reduces the time to access the data set.

- Spread high-activity HFS data sets and volumes across multiple DASD control units.

  Monitor the DASD subsystem and make sure that the activity is spread across the available DASD volumes and control units. Keep user HFS data sets separate from system HFS data sets.

- Place UNIX files as close to root as possible.

  When accessing a file system, HFS locks all paths to the file. By placing files close to the root, you speed up HFS access.

- Avoid collecting unnecessary data.

  Domino can collect a lot of statistics and log data. This can drive a lot of I/O activity. Make sure you collect only the data you are interested in.

3.2 Cached DASD

The S/390 platform supports a broad range of DASD subsystems, from the 3380 models to the latest RAMAC family products. The later products generally provide better performance characteristics, and, therefore, you should consider these for online workloads such as Domino.

A key element of the DASD subsystem is cache storage. Volumes that contain HFS data sets should be placed on cached DASD with DASD Fast Write for good performance.

The latest 3990/9390 control units are provided with Dynamic Cache Management Enhanced (DCME). DCME provides more comprehensive and responsive cache management techniques through the use of improved algorithms. This function determines which data sets are using cache effectively and dynamically turns cache off for data sets that do not benefit from caching. DCME also provides additional cache activity and I/O activity statistics which DFSMS/MVS stores in the SMF data set to be used for customer information reports.

We strongly recommend that you allocate Lotus Domino HFS data sets on cached DASD with the DASD Fast Write (DFW) feature enabled and under DCME capable controllers.

For more details about cache implementation, see *IBM 3990/9390 Storage Control Planning and Installation*, GA32-0100.

3.3 Hierarchical File System (HFS)

OS/390 provides a standard UNIX hierarchical file system (HFS) for applications that use the UNIX interfaces. The HFS on S/390 is made up of a number of OS/390 data sets, called HFS data sets, that are mounted into the HFS.

An HFS data set is a standard OS/390 PDSE data set with DSNTYPE=HFS file organization. The PDSE access method is used only when allocating an HFS data set. After that the PDSE access method is no longer involved in I/O operations. Many cached storage control units provide enhanced internal processing for PDSE searches, which can result in improved performance. However HFS data sets do not benefit from this feature because they do not have a conventional PDSE file structure.

For more information on the HFS, see *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083, chapter 2.12 “Notes Directory Structure.”

3.3.1 HFS Data Set Allocation and DFSMS/MVS

HFS data sets must be allocated as system managed data sets on volumes managed by Data Facility System Managed Storage (DFSMS). When allocating these data sets, DFSMS uses four constructs, as shown in Table 3.

Table 3. DFSMS Constructs

DFSMS Construct	**Definition**
Storage Class	Specify data set performance and availability
Management Class	Specify data set backup and migration criteria
Data Class	Specify data set allocation parameters
Storage Group	Specify volume name and free space thresholds

You must define at least one Storage Group and one Data Class. The other constructs are optional. However, in order to specify performance requirements,

you should define and use the Storage Class. The Storage Class is the only way to manage performance in a DFSMS environment.

You also need to check your SMS ACS routines to ensure that appropriate constructs are assigned to HFS data sets for Domino use.

3.3.1.1 DFSMS Storage Class

A sample Storage Class for Lotus Domino data sets is shown in Figure 11. This uses a “may cache” policy for those files, resulting in better performance and better use of the available hardware.

```
  Panel  Utilities  Scroll  Help
------------------------------------------------------------------------------
                         STORAGE CLASS DISPLAY               Page 1 of 2
Command ===>

CDS Name  . . . . . : ACTIVE
Storage Class Name  : DOMINOSG
Description  : STORAGE CLASS FOR LOTUS-DOMINO HFS FILES - DR18FEV98

Performance Objectives
  Direct Millisecond Response  . . . : 5
  Direct Bias  . . . . . :
  Sequential Millisecond Response  . : 5
  Sequential Bias  . . . . . :
  Initial Access Response Seconds  . : 2
  Sustained Data Rate (MB/sec) . . . :
Availability . . . . . : PREF
Accessibility  . . . . . : NO
Guaranteed Space . . . . . : NO
Guaranteed Synchronous Write . . . . : NO


Use DOWN Command to View the next Page;
Use HELP Command for Help; Use END Command to Exit.
```

Figure 11. DFSMS Storage Class

The performance objectives parameters tell the I/O subsystem what response time you require for data sets that are assigned this storage class. A data set may have different attributes for sequential access and for direct access. DFSMS will attempt to meet those requirements by allocating the data set on a suitable device. If you have DCME-capable control units, DCME will also manage the control unit cache to try to meet the objectives of all of the data sets that it is managing.

DFSMS uses hiperspaces to automatically stage PDSE data sets to expanded storage when the MSR (millisecond response) of the SMS Storage Class indicates a “must cache” data set. However, because HFS data sets are not conventional PDSE files, it does not use this feature.

See *Storage Management Library Version 1 Release 2*, SC26-3124 for more details.

3.3.2 HFS Data Set Management

This section provides some management recommendations for the HFS data sets.

3.3.2.1 Placement of HFS Data Sets on DASD Volumes

As with other large subsystems, one of the keys to good I/O performance is to spread the I/O load across multiple DASD volumes, control units and channels. You therefore need to identify the high-activity HFS data sets. Since each HFS data set contains one or more UNIX directories and multiple files, you also need to understand which are the high-activity UNIX files.

The likely high-activity files are:

- The Notes data directory (`/notesdata` by default). This directory, by default, contains the Public Address Book, help files, and the billing, log, journal, and statistics files.
- Notes databases that are shared by many users.

You should monitor your system regularly to check for bottlenecks in I/O activity.

3.3.2.2 Separating High-Activity Files

If you have an HFS data set with a lot of activity and you wish to split it you need to:

- Move the high-activity UNIX files into different directories. Domino can still find databases that you move by the use of Domino links. See 3.3.2.3, “Domino Links.”
- Make sure that those directories are in different HFS data sets. A single HFS data set contains all of the files in a single directory and may or may not contain files in subdirectories.
- Make sure that the high-activity HFS data sets are placed on different volumes. You can move an HFS data set using DFSMSdss.

3.3.2.3 Domino Links

By default, Domino expects databases to be in the Notes data directory (`/notesdata` by default) or in a subdirectory of it. However, you can place a database in any directory on the system and point to it using a Notes directory or database link. Thus you can spread your databases across multiple directories and HFS data sets, while logically keeping them in a single directory for ease of management. In addition, this allows you to spread out those system databases that Domino expects to find in the Notes data directory.

For information on how to set up Notes directory and database links, see the *Lotus Notes Administrator's Guide*.

3.3.2.4 Consolidate Extents

HFS data sets can expand to use multiple extents. As with other OS/390 data set types, accessing a data set with multiple extents is less efficient. Therefore, you should monitor the number of extents in your HFS data sets and consolidate data sets with more than one extent using DFSMSdss.

3.3.2.5 Minimize Directory Levels

When OS/390 UNIX Services opens a UNIX file, it locks all of the path it uses to access the file. Therefore by minimizing the number of directory levels, you minimize the number of directories that need to be locked, and thus reduce contention.

3.3.2.6 RAMAC Virtual Array (RVA)

The IBM RAMAC VIRTUAL ARRAY (RVA) DASD models are worth a special mention. RVA provides a new concept in DASD architecture. It allows you to define up to 256 virtual volumes over a storage pool, emulating any mixture of 3380 and 3390 models.

This allows you to place a high-activity data set on a smaller dedicated volume. OS/390 will only start a single I/O at a time to a volume. With only one data set on the volume you will eliminate any queue time (called IOSQ time) in OS/390.

RVA also supports SNAPSHOT, a powerful backup/dump utility.

3.3.2.7 Domino Server Partitioning

Domino server partitioning is the capability to run multiple Domino servers on the same OS/390 image. Partitioned servers share the same Domino code but each partition has its own data files, including the Notes data directory and user databases. Files cannot be shared between Domino servers.

From a performance perspective we recommend that you separate the files for each server into separate UNIX directories, HFS data sets and DASD volumes. This will help you avoid data access contention and also provides the simplest management environment.

3.4 Monitoring DASD I/O Activity

You should monitor the DASD I/O activity for Domino files. You can do that at a number of levels.

3.4.1 RMF

Use RMF to monitor the DASD response time for volumes that contain Domino HFS data sets. 2.8.1.2, “RMF Summary Report - Overall System Performance” on page 16 shows the RMF summary report that lists the average DASD response time and DASD I/O rate.

Other reports provide you with information on:

- Channel path activity
- I/O queueing activity
- Cache subsystem activity (see 3.4.1.1, “Monitoring DASD Cache” on page 30)
- DASD activity.

These reports are useful for identifying performance issues such as:

- Poor response time from a DASD volume (aim for at most six to ten milliseconds, although it will depend on the activity rate and the DASD device type)
- Too much activity to a single DASD volume

- Channels that are too busy (more that 45% busy for ESCON channels with four paths, more that 30% busy for parallel channels)
- Insufficient paths - we recommend at least four paths to every control unit

3.4.1.1 Monitoring DASD Cache

The RMF Cache Subsystem Activity report and the IDCAMS LISTDATA STATUS command are useful tools to monitor cache use and effectiveness. Figure 12 shows an example.

```
                                        C A C H E   S U B S Y S T E M   A C T I V I T Y
                                                                                                                              PAGE  1
            OS/390                  SYSTEM ID ML99             START 03/12/1998-11.30.00  INTERVAL 000.14.59
            REL. 02.04.00           RPT VERSION 2.4.0          END   03/12/1998-11.45.00
SUBSYSTEM  3990-06     CU-ID  6300     SSID 6107     CDATE  03/12/1998     CTIME 11.30.01      CINT  00.14.59
------------------------------------------------------------------------------------------------------------------------------------
                                                    CACHE SUBSYSTEM STATUS
------------------------------------------------------------------------------------------------------------------------------------
SUBSYSTEM STORAGE           NON-VOLATILE STORAGE        STATUS
CONFIGURED     1024.0M      CONFIGURED        64.0M     CACHING                  - ACTIVE
AVAILABLE      1019.9M      PINNED             0.0      NON-VOLATILE STORAGE     - ACTIVE
PINNED            0.0                                   CACHE FAST WRITE         - ACTIVE
OFFLINE           0.0                                   IML DEVICE AVAILABLE     - YES
------------------------------------------------------------------------------------------------------------------------------------
                                                   CACHE SUBSYSTEM OVERVIEW
------------------------------------------------------------------------------------------------------------------------------------
TOTAL I/O   192802   CACHE I/O   192802   CACHE OFFLINE        0
TOTAL H/R    0.800   CACHE H/R    0.800
CACHE I/O    -------------READ I/O REQUESTS-------------    ----------------------WRITE I/O REQUESTS----------------------      %
REQUESTS       COUNT     RATE      HITS     RATE     H/R      COUNT     RATE     FAST     RATE      HITS     RATE     H/R    READ
NORMAL        116082    129.1     82736     92.0   0.713      67105     74.6    67105     74.6     61964     68.9   0.923    63.4
SEQUENTIAL      6564      7.3      6473      7.2   0.986       3051      3.4     3051      3.4      3051      3.4   1.000    68.3
CFW DATA           0      0.0         0      0.0     N/A          0      0.0        0      0.0         0      0.0     N/A     N/A
TOTAL         122646    136.4     89209     99.2   0.727      70156     78.0    70156     78.0     65015     72.3   0.927    63.6
-----------------------CACHE MISSES-----------------------           ------------MISC------------          ------NON-CACHE I/O---
REQUESTS      READ    RATE   WRITE    RATE  TRACKS    RATE                          COUNT   RATE                      COUNT RATE
                                                                     DFW BYPASS       502    0.6         ICL              0  0.0
NORMAL       33346    37.1    5141     5.7   37767    42.0           CFW BYPASS         0    0.0         BYPASS           0  0.0
SEQUENTIAL      91     0.1       0     0.0    6801     7.6           DFW INHIBIT        0    0.0         TOTAL            0  0.0
CFW DATA         0     0.0       0     0.0                           ASYNC (TRKS)   28448   31.6
TOTAL        38578    RATE    42.9
---CKD STATISTICS---     ---RECORD CACHING---
WRITE           1582     READ MISSES    25968
WRITE HITS       107     WRITE PROM      2581
                                        C A C H E   S U B S Y S T E M   A C T I V I T Y
                                                                                                                              PAGE  2
            OS/390                  SYSTEM ID ML99             START 03/12/1998-11.30.00  INTERVAL 000.14.59
            REL. 02.04.00           RPT VERSION 2.4.0          END   03/12/1998-11.45.00
SUBSYSTEM  3990-06     CU-ID  6300     SSID 6107     CDATE  03/12/1998     CTIME 11.30.01      CINT  00.14.59
------------------------------------------------------------------------------------------------------------------------------------
                                                 CACHE SUBSYSTEM DEVICE OVERVIEW
------------------------------------------------------------------------------------------------------------------------------------
VOLUME     DEV  DUAL      %     I/O    ---CACHE HIT RATE--   ----------DASD I/O RATE----------     ASYNC   TOTAL    READ   WRITE     %
SERIAL     NUM  COPY    I/O    RATE     READ     DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER     RATE     H/R     H/R     H/R  READ
*ALL                  100.0   214.5     99.2    72.3    0.0    42.4    0.6    0.0    0.0    0.0     31.6   0.800   0.727   0.927  63.6
*CACHE-OFF              0.0     0.0
*CACHE                100.0   214.5     99.2    72.3    0.0    42.4    0.6    0.0    0.0    0.0     31.6   0.800   0.727   0.927  63.6
0E0031    6300          6.5    14.0      7.5     4.0    0.0     2.4    0.0    0.0    0.0    0.0      1.1   0.828   0.759   0.999  71.0
0E0032    6301          3.3     7.0      3.5     1.3    0.0     2.1    0.0    0.0    0.0    0.0      0.7   0.693   0.621   0.999  80.9
0E0033    6302          2.7     5.9      2.6     1.5    0.0     1.7    0.0    0.0    0.0    0.0      0.7   0.699   0.592   1.000  73.7
0E0034    6303          2.7     5.7      2.8     1.4    0.0     1.6    0.0    0.0    0.0    0.0      0.8   0.724   0.636   1.000  75.8
0E0035    6304          3.2     6.9      3.3     2.0    0.0     1.5    0.0    0.0    0.0    0.0      0.8   0.773   0.682   0.999  71.3
```

Figure 12. RMF Cache Report

On the Cache subsystem activity report, observe the following fields:

CACHE SUBSYSTEM STATUS Make sure all the fields under the STATUS column are active.

OFFLINE Shows the amount of cache disabled for use. This can indicate a hardware problem.

DEVICE OVERVIEW Take a look at the device overview section to see whether one of the volumes exhibits unusual behavior compared to the others.

3.4.2 SMF

SMF provides you with more detailed information on I/O activity. See 2.8.2, “Monitoring OS/390 with SMF” on page 21 for details of the SMF records.

3.4.2.1 Record Type 42

SMF record type 42 subtype 6 provides information about HFS data set performance. I/O statistics are provided for each HFS data set.

3.4.2.2 Record Type 30

SMF record type 30 reports activity on a job and job step basis. The OS/390 UNIX Services process section also provides information on file system lookups, which can consume significant resources on systems with hierarchical files. A lookup is an access to the directory path in order to access a file. Monitor the number of lookups on your system. To reduce the number of lookups, keep frequently used files in the same directory and avoid too many levels on the file structure. If key jobs appear to be doing many lookups, your installation may be able to reduce this overhead by reorganizing the file system and mounting these key files closer to the root of the file system.

3.4.2.3 Record Type 92

SMF record type 92 subtype 5 collects information at the time the HFS is mounted and unmounted.

SMF record type 92 subtype 11 collects information at the time the HFS is opened and closed.

3.4.3 UNIX Statistics

You also can obtain some online statistics at the UNIX level by issuing the UNIX command `/usr/sbin/kerninfo stats`. The output is shown in Figure 13 on page 32.

```
> /usr/sbin/kerninfo stats
KERNINFO: 4 Mar 1998 09:23:14

Total lookups: 7653997
Look-aside hits: 93%
Age-out misses: 0%
Latch set: BPXPRTB1.PPRA.LATCH.SET
  Latches: 4097  HWM: 64  Fast obtains: 8729544  Slow obtains: 2605
Latch set: BPXFSLIT.FILESYS.LATCH.SET
  Latches: 16384  HWM: 64  Fast obtains: 7534803  Slow obtains: 134
Misc file system slow obtains: 0
NFS
   ** no vfs **
CINET
   AF_INET
      dev=17(11x)  cache count=9  vnode count=40
      VFS slow obtains=0  Vnode slow obtains=0
      Reads=357902  Writes=138954  Dir I/O=0
      Read Blks=0  Write Blks=0
UDS
   AF_UNIX
      dev=16(10x)  cache count=0  vnode count=3
      VFS slow obtains=0  Vnode slow obtains=0
      Reads=189  Writes=28141  Dir I/O=0
      Read Blks=0  Write Blks=0
HFS
   OMVS.DOM453.SERVER2.MAIL.HFS
      dev=22(16x)  cache count=0  vnode count=1
      VFS slow obtains=0  Vnode slow obtains=0
      Reads=1422  Writes=2018  Dir I/O=418
      Read Blks=1667  Write Blks=2834
   OMVS.DOM453.SERVER2.DATA.HFS
      dev=21(15x)  cache count=0  vnode count=2
      VFS slow obtains=0  Vnode slow obtains=0
      Reads=28782  Writes=8941  Dir I/O=181635
      Read Blks=1354875  Write Blks=112834
```

Figure 13. Kerninfo Stats Command

This display reports the number of reads, writes, blocks, cache count and other variables for each HFS data set.

3.4.4 Domino I/O Statistics

Domino provides statistics on database use in the Notes log database `log.nsf`. See the views “Database by sizes or usage.” This gives you information on documents read and written, data transferred and number of transactions. These statistics are helpful for monitoring which are the most accessed databases, how full they are, the access mode pattern, and distribution across time. Based on these statistics, you may consider repositioning a database. See 5.11.1, “Notes Log” on page 58.

Chapter 4. Network Tuning

Notes clients and Web browsers are connected to a Domino server on S/390 through a TCP/IP network. In this chapter we discuss the performance of the TCP/IP network, and how you can monitor and tune it.

Note: When deploying OS/390 as a Domino server, TCP/IP is the only network protocol supported.

A client/server application such as Lotus Notes is designed so that parts of the application run on the client, parts run on the server, and these parts must communicate with each other. The connection is made through a set of resources called a network. The network usually consists of multiple hardware and software resources. The challenge is to ensure good performance along the entire path and consistent response time to the user. Network performance is an important part of Domino performance.

4.1 Recommendations

These are the key network tuning recommendations:

- Use OS/390 TCP/IP for UNIX Services or a later release.

  OS/390 TCP/IP for UNIX Services, also called OS/390 TCP/IP OpenEdition (or TCP/IP Version 3 Release 3 in some early documentation), was especially developed for applications that use the OS/390 UNIX Services sockets interfaces, such as Domino. It provides improved performance over previous releases. Make sure that all required service and performance PTFs are applied.

- Consider the whole network.

  The network consists of many hardware boxes and links. Monitor all the components in the network and address bottlenecks.

- Define REGION=0M for the TCP/IP region size.

  TCP/IP buffers are dynamically allocated in the private area. REGION=0M allows TCP/IP to allocate the maximum virtual storage above and below the 16MB line. If you have implemented an IEFUSI exit check that it will allow the specification of 0M for the region size.

- Estimate TCP/IP buffers accurately.

  TCP/IP's basic task is buffer management. The correct buffer sizes will optimize TCP/IP performance. Use the TCP/IP shutdown statistics to estimate and tune the number and size of buffers.

4.2 TCP/IP

Lotus Domino for S/390 uses the TCP/IP network protocol to communicate with Notes clients and Web browsers. This section discusses TCP/IP considerations on OS/390.

4.2.1 TCP/IP Release Level

TCP/IP Version 3 Release 2 is the minimum level required to run Domino on S/390. Domino does not require the TCP/IP Version 3 Release 2 OpenEdition Application Feature, except that you will need it if you wish to use Telnet or to FTP directly into the HFS. Domino uses TCP/IP port number 1352.

However, we *strongly recommend* that you implement at least OS/390 TCP/IP for UNIX Services to benefit from better performance. Between TCP/IP Version 3 Release 1 and TCP/IP Version 3 Release 4 we have measured improvements of more than 30% in total throughput and up to 60% in CPU reduction in some tests.

4.2.1.1 TCP/IP Version 3 Release 2

TCP/IP Version 3 Release 2 was first released with OS/390 Version 1 Release 3. It is a mixed version of TCP/IP. It includes support for both MVS socket applications and OS/390 UNIX Services socket applications. You can open a Telnet or FTP session into the regular MVS interfaces or directly with OS/390 UNIX Services, depending on your port definitions. We recommend that you do not use this release.

To access OS/390 UNIX Services directly you need the optional no-charge TCP/IP OpenEdition Application Feature.

4.2.1.2 OS/390 TCP/IP for UNIX Services

OS/390 Version 2 Release 4 comes with a new TCP/IP called TCP/IP for UNIX Services (FMID JTCP329). It is also known as TCP/IP OpenEdition. This supercedes the TCP/IP OpenEdition Application Feature (FMID JTCP327). This is a separate stack release because it only supports OS/390 UNIX Services socket applications; it does not support MVS socket applications. TCP/IP for UNIX Services does not replace the entire TCP/IP Version 3 Release 2 product stack.

Therefore, if you want support for all MVS and OS/390 UNIX Services functions you must run two stacks: TCP/IP Version 3 Release 2 for regular MVS and TCP/IP for UNIX Services for OS/390 UNIX Services support.

OS/390 UNIX Services applications benefit from enhanced performance and enhanced reliability in TCP/IP for UNIX Services. The enhanced performance comes from the newly designed stack, the Communications Storage Management (CSM) facility and the Multi-Path Channel (MPC) I/O process. CSM and MPC are supplied with ACF/VTAM Version 4 Release 4 and are used by both VTAM and the TCP/IP for OS/390 UNIX Services stack. CSM reduces data moves and manages subsystem data pools. MPC handles multiple I/O dispatchable units of work.

4.2.1.3 Communications Server

The TCP/IP shipped with the Communications Server in OS/390 Version 2 Release 5 can be thought of as TCP/IP Version 3 Release 4. In this release, the OS/390 UNIX Services component is merged back into the TCP/IP basic product. This version of TCP/IP supports both OS/390 UNIX Services socket applications and MVS socket applications.

4.2.2 Monitoring TCP/IP

TCP/IP provides several commands that are useful for monitoring the system and highlighting problems. In particular we discuss:

- PING
- TRACERTE
- NETSTAT
- TCPIPSTATISTICS

You should also use System Display Service Facility (SDSF) and Resource Management Facility (RMF) to monitor TCP/IP at the OS/390 level. Monitor CPU usage by the TCP/IP address space, paging and I/O rates for communication devices for example. See 2.8.1, “Monitoring OS/390 with RMF” on page 15.

Beyond these tools you should use a network monitor product for more detailed monitoring, tuning and troubleshooting.

4.2.2.1 PING

The PING command sends an echo to an IP host to determine whether the host is accessible. The response shows how long the round trip message took. This is the first tool to use to identify connectivity problems, and it can also show performance problems.

The PING command can be issued from any system that supports TCP/IP, including a Notes workstation client or a Web browser, where you would issue it from a command prompt. PING can also be issued from OS/390 TSO, as shown in Figure 14.

Note: The PING command is valid in OS/390 Version 2 Release 5. In earlier versions use OPING.

```
                                ISPF Command Shell
Enter TSO or Workstation commands below:

===> tso ping 9.12.2.17 (length 481 count 5


PING: Ping #1 response took 0.022 seconds. Successes so far 1.
PING: Ping #2 response took 0.022 seconds. Successes so far 2.
PING: Ping #3 response took 0.022 seconds. Successes so far 3.
PING: Ping #4 response took 0.022 seconds. Successes so far 4.
PING: Ping #5 response took 0.023 seconds. Successes so far 5.
***
```

Figure 14. TCP/IP PING Command

In this example, the PING command is sending 481 bytes of data five times. The time for the data to be sent to the remote host with IP address 9.12.2.17 and then returned to this system was 0.022 seconds. High values for the response time can show performance problems somewhere between the issuing system and the target system.

4.2.2.2 TRACERTE

The TRACERTE command is similar to the PING command except that it reports the round trip time by intermediate hosts. It is helpful to check how each connection between the issuing system and the target system is performing and to determine bottlenecks along the route. You must be defined in the OBEY statement of the TCP/IP profile data set to issue this command. Figure 15 shows an example of this command.

```
                              ISPF Command Shell
Enter TSO or Workstation commands below:

===> tso tracerte 9.12.2.17


Trace route to 9.12.2.17 (9.12.2.17)
1     (9.12.14.75) 15 ms 7 ms 7 ms
2     (9.12.14.7) 5 ms 6 ms 6 ms
***
```

Figure 15. TCP/IP TRACERTE Command

Note: The three times reported for each line show that there were three attempts for each connection.

For more information on the TRACERTE command, see *IBM TCP/IP for MVS: User's Guide*, SC31-7136.

4.2.2.3 NETSTAT

The NETSTAT command can be used to monitor TCP/IP buffer usage. It has many parameters; use NETSTAT HELP to display all of them. Figure 16 on page 37 shows an example of the NETSTAT command.

Note: The NETSTAT command is valid in OS/390 Version 2 Release 5. In earlier versions use ONETSTAT.

```
                                ISPF Command Shell
Enter TSO or Workstation commands below:

===> tso netstat poolsize


MVS TCP/IP Netstat Version 3 Release 2

TCPIP Free pool status:
Object       # alloc     # free      Lo-water    Permit size
======       =======     ======      ========    ===========
ACB           12500       12493       12488        1250
CCB             150          89          88          15
Dat buf        8096        8090        8087         809
Sm dat buf     1200        1199        1193         120
Tiny dat buf   7600        7598        7598         760
Env             750         750         746          75
Lrg env         100          99          99          10
RCB              50          50          48           5
SCB             256         217         215          25
SKCB           7600        7593        7593         760
TCB           14000       13994       13992        1400
UCB             100          98          97          10
Add Xlate      1500        1496        1495           5
IP Route        300         295         294           6
***
```

Figure 16. TCP/IP NETSTAT POOLSIZE Command

4.2.2.4 TCPIPSTATISTICS

TCPIPSTATISTICS provides TCP/IP statistics when TCP/IP is shut down. To get these statistics, specify this parameter on the ASSORTEDPARMS statement in the TCP/IP profile data set. Figure 17 shows an example of these statistics.

```
ipInReceives      =         622   vmcfSendsOK       =         188
ipOutRequests     =         371   vmcfSendsAbnormal =           0
ipForwDatagrams   =           0   vmcfSendsFatal    =           5
ipReasmReqds      =           0   iucvReplies       =           0
ipReasmFails      =           0   iucvReceives      =           0
ipFragCreates     =           0   iucvSends         =           0
ipFragFails       =           0   iucvRejects       =           0
ipInHdrErrors     =           0   ioReads           =         320
ipInAddrErrors    =           9   ioWrites          =         224
icmpInMsgs        =           1   ioInOctets        =      234860
icmpOutMsgs       =           4   ioOutOctets       =      134482
arpInRequests     =          58
arpOutReplies     =           0
arpOutRequests    =           3
tcpInSegs         =         449
tcpOutSegs        =         360
tcpRetransSegs    =          72
udpInDatagrams    =           0
udpOutDatagrams   =           0

ShutDown at 287.143 seconds
```

Figure 17. TCP/IP Shutdown Statistics

See the *IBM TCP/IP Performance Tuning Guide*, SC31-7188 and the *MVS TCP/IP User's Guide*, SC31-7136 for more details.

4.2.3 Tuning TCP/IP

This section contains some key actions that you should take to tune TCP/IP. However, TCP/IP tuning is a specialist skill and you should refer to the *IBM TCP/IP Performance Tuning Guide*, SC31-7188 for more detailed tuning information.

4.2.3.1 TCP/IP Priority in OS/390

We recommend that you set the MVS dispatching priority of TCP/IP to be approximately the same as VTAM. Other TCP/IP servers, such as FTP and Telnet, should be set at a slightly lower priority than TCP/IP.

Also make sure that TCP/IP is getting the OS/390 resources that it needs, by monitoring it using RMF. If TCP/IP is constrained it will not perform well.

4.2.3.2 Specify REGION=0M for TCP/IP

TCP/IP buffers are dynamically allocated in the private area. REGION=0M allows TCP/IP to allocate the maximum virtual storage above and below the 16MB line, and thus does not constrain the number of buffers it can use.

4.2.3.3 Make Buffer Pools Large Enough

TCP/IP uses a lot of buffers to store data being sent through the network. It is important that these buffers are large enough. Otherwise TCP/IP is unable to accept more data until some buffers are cleared. Therefore, you should monitor buffer pool use carefully.

The buffer pool parameters are stored in the TCP/IP profile data set. Recommendations for the initial settings for Domino are given in the *Lotus Domino for S/390 Install Guide for Servers* and you should follow those. The buffer pool parameters are documented in *TCP/IP Users Guide*, SC31-7136.

Once you have Domino running you should monitor the use of the buffer pools, as described in 4.2.2.3, “NETSTAT” on page 36. From the NETSTAT POOLSIZE output calculate the maximum number of buffers used for each pool, using the formula:

```
Maximum used = # alloc -  Lo-water
```

See Figure 16 on page 37.

Consider the following when changing the number of buffers:

- Collect this information for several days.
- Reduce the overallocated buffers and increase the underallocated.
- The number of buffers is limited only by the amount of virtual storage available. If you are not constrained by virtual storage, define more buffers.
- CCBs are required for OBEY and INFORM statements in the TCP/IP profile data set, one for each user and one for each FTP server.
- Add one UCB for each open UDP port.

4.2.3.4 OS/390 UNIX Services Network Parameters

There are two parameters related to the network in the OS/390 UNIX Services startup parameters in SYS1.PARMLIB(BPXPRMxx):

- NETWORK DOMAINAME(AF_INET) MAXSOCKETS

  TCP/IP supports AF_INET sockets for communication with OS/390 UNIX Services applications. Notes clients use AF_INET sockets to communicate with the Domino server. This parameter should be estimated based on the maximum number of clients connecting to the server.

- NETWORK DOMAINAME(AF_UNIX) MAXSOCKETS

  AF_UNIX sockets are used when two OS/390 UNIX Services applications establish a connection. Use the default value of 100 for Domino.

For more information about these parameters, see *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083.

4.3 Network Connection

The S/390 server can connect to the IP network using:

- S/390 Open System Adapter (OSA)
- 3174 Networking Server
- 3745/3746 Communications Controller
- 3172-3 Interconnect Controller

You should monitor the connection that you use to ensure good performance.

4.3.1 Monitoring Connection Activity

You can monitor the connection from an OS/390 perspective, using RMF postprocessor or MONITOR-II. To use RMF MONITOR-II:

1. Start Mon-II data collection.
2. From TSO, enter RMFMONITOR.
3. Enter DEV COMM in the input area of the primary menu. This creates an I/O activity report for communication devices.
4. Enter D ON in the input area. This includes only the data since the last report, rather than a cumulative total.
5. Press Enter to begin.
6. Look at a couple of sequential intervals and see if there are any constraints (see the DELAY columns).

Figure 18 on page 40 shows a sample report.

```
                                  CPU= 13 UIC=254 PFR=   0 DEV      T
13:15:17        DEV      ACTV RESP IOSQ ---DELAY--- PEND DISC CONN %DEV %D
STG GRP  VOLSER NUM LCU  RATE TIME TIME DPB CUB DB  TIME TIME TIME UTIL RV I%
                00F --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
                102 --- 15.21   42    0     0.0 0.0  0.2 63.5  0.1 96.8  0 61
                158 --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
                159 --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
                9CE --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
                9CF --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
                DAA --- 0.000    0    0     0.0 0.0  0.0  0.0  0.0 0.00  0 61
```

Figure 18. Sample RMF Report to Monitor Communications Device

In this example address 102 is an OSA device. The address number matches the DEVICE statement in the TCP/IP profile data set.

For more information on OSA, see *Planning for System/390 Open Systems Adapter Feature*, GC23-3870.

4.4 Network Performance

As well as considering TCP/IP and the network connection on the S/390 server, you also need to consider all of the other network components. That is beyond the scope of this book. However, we discuss the option of splitting the network.

4.4.1 Splitting the Network

Under certain circumstances you may consider splitting the network in order to improve performance. You might do this to separate different types of workloads. For example, you could have a separate network for Domino traffic. This is likely to be more costly than running all traffic over a single network, but it can offer these benefits:

- It allows Domino work to be isolated from other workloads. Therefore, Domino will not be impacted if the other workloads generate a lot of network traffic.
- It allows you to tune the separate networks for their specific workloads. You could, for example, customize the TCP/IP parameters to a specific application pattern usage, such as Domino.

Note: Lotus recommends that you run a separate network for interserver traffic if you use Domino clustering. That ensures good performance for that traffic.

Some of the ways that you could split the network traffic are:

- Add more network interfaces to the host.

  By adding more network interfaces you can balance the network traffic arriving at the S/390 Server. If the interface, such as an OSA card, is shared between several MVS images, you may consider adding additional ones and dedicating them to an MVS image.

- Run multiple copies of TCP/IP.

  This allows you to tune each TCP/IP for the specific applications that it supports.

- Define additional PORT addresses.

Many LAN gateway interfaces support more than one port per gateway, such as OSA-2 and 3172-3.

- Use MVS Virtual IP Addressing (VIPA).

  VIPA allows you to define more than one IP address on the same TCP/IP stack. This is very useful if you want to drive the network traffic through different IP addresses for workload balancing or security reasons.

4.5 Domino Network Statistics

Domino provides statistics on network use in the Notes log database `log.nsf`. (see 5.11.1, “Notes Log” on page 58). See the view “Usage By Date or User.” This gives you information on the number of Kbytes of data transferred across the network. It is helpful to analyze and forecast network traffic. You can monitor the network usage throughout the day to determine peak times and bottlenecks.

Chapter 5. Domino Tuning

This chapter discusses the Domino parameters you can set that affect the performance of your system. It shows how you specify which tasks the Domino server will run, and provides information on the relative resource use of those tasks. We also discuss useful Domino information that you should monitor.

5.1 Recommendations

These are the key recommendations for tuning Domino:

- Set Notes_SHARED_DPOOLSIZE to 32000000.

  The Notes_SHARED_DPOOLSIZE variable sets the amount of storage that the Domino server will ask for in one request. By setting it high Domino will get more storage in each request and will therefore make fewer requests to the system. However, see the important notes on this parameter in 5.2.2, “The Notes_SHARED_DPOOLSIZE Parameter” on page 45 before you change it.

- Do NOT change the following parameters:

  – Server_Max_Concurrent_Trans
  – Server_MaxSessions
  – Server_MaxInitialThreads
  – Server_Secondary_Threads

  Leave the values to default. Otherwise you may cause problems with the Domino server on S/390.

- Do not run unnecessary Domino tasks.

  Check the `notes.ini` file and remove any server tasks that you do not need. For example, you may not need tasks such as logging and billing. You can run multiple versions of some tasks, but check whether you need them.

- Schedule tasks for non-peak times.

  Some tasks run on a scheduled basis. Schedule them to run at non-peak times. Examples are compact, index (to create a new full text index), update, updall, and fixup. The schedule is set by the Notes administrator but should be agreed with the OS/390 operations staff to make sure that it does not interfere with other peak workload or administration tasks on the server.

- Balance work between the server and the client.

  Lotus Notes is a client/server application. It runs on a combination of the Domino server and the Notes client. (If you are using a Web browser as the client, the work that can be done on the client is limited.) Consider running some functions on the client rather than on the server, to minimize the server load. For example, you could tell your users to:

  – Put all users that they send mail to in their local Public Address Book to avoid searching the Public Address Book on the server when sending mail.

  – Replicate their mail and other databases to their Notes client to offload processing from the server.

  As you look at these ideas, remember how much server offload you might get. For example, you may not be doing your mail on the server, but you do need to

replicate with the server. Also consider the management aspects of users having database replicas on their workstations.

- Educate users on replication.

  If users replicate databases to their Notes clients, educate them on the replication interval. We recommend a value of 15 to 30 minutes or more. As each replication uses server resource, even if there is no data to transfer, a longer interval will reduce the server load.

  Also educate users to use scheduled replication only on databases for which that makes sense. Within IBM we recommend that some applications be run locally by replicating them to the client. However, for applications used infrequently (for example, to submit expense claims), we recommend that those databases do not replicate automatically. Instead the application reminds the user to replicate the database manually whenever they update information in the local replica.

- Build applications to replicate from the server to the client.

  In our tests with Domino for S/390, creating a new replica on the server (replicate to server) was slower than creating a new replica on the client (replicate to client). Consider this when you design applications.

5.2 Domino Initialization Parameters

The `notes.ini` file provides the parameters for the initialization of the Domino server. During server start this file is read to set up the right Notes environment and load the predefined tasks and programs.

`notes.ini` contains parameters to define:

- Which server tasks are started
- Tuning parameters
- Other parameters

In the following sections we discuss the Domino server tasks and some of the performance parameters. We include a list of the Domino tuning parameters that can be set in `notes.ini` in Appendix B, “Domino Parameters That Affect Performance” on page 79. For further information on the `notes.ini` file, see the *Lotus Notes Administrator's Guide*.

5.2.1 Changing Parameters in `notes.ini`

To modify parameters in `notes.ini` you can:

- Use a server configuration document
- Use the Set Configuration server command at the console
- Edit the server document
- Edit the file directly with an editor

These are all described in the *Lotus Notes Administrator's Guide* and *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083.

5.2.2 The Notes_SHARED_DPOOLSIZE Parameter

The Notes_SHARED_DPOOLSIZE parameter is a UNIX environment variable that you can set before you start the Domino server. This variable tells the Domino server how much UNIX shared memory to ask for in each request to OS/390. The default value is 1000000. However, we recommend that you set this value to 32000000.

With the current implementation of shared memory, there is a significant amount of work involved in getting a new piece of shared memory and in making that accessible to all of the Domino address spaces. To minimize this overhead, set Notes_SHARED_DPOOLSIZE to 32000000 so that Domino needs to make fewer memory requests to OS/390.

Important

This parameter is closely related to the OS/390 SYS1.PARMLIB(BPXPRMxx) parameter IPCSHMMPAGES. Do *not* change Notes_SHARED_DPOOLSIZE without first checking the value of this parameter; otherwise your Domino server may not start.

See the redbook *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083 for more information on Notes_SHARED_DPOOLSIZE and the IPCSHMMPAGES parameter in SYS1.PARMLIB(BPXPRMxx).

5.3 Domino Server Tasks

Domino consists of a number of server tasks. These are shown in Table 4.

Table 4 (Page 1 of 2). Description of Domino Server Tasks

Task Name	Description	See Note
Adminp	Administration process	3
Amgr	Agent manager	
Billing	Collects billing information	3
Calconn	Calendar connector	
Catalog	Cataloger	2
Chronos	Updates indexes	
Cladmin	Cluster administration	
Cldbdir	Cluster database directory	
Collector	Collects statistics for multiple servers	
Compact	Database compactor	2
Delcnflt	Delete replication conflicts	
Event	Monitors events on a server	
Fixup	Locates and fixes corrupted databases	2
HTTP	Web server	
Object	Object store manager	

Table 4 (Page 2 of 2). Description of Domino Server Tasks		
Task Name	**Description**	**See Note**
POP3	Maildrop for POP3 clients	
Report	Reports statistics for a server	
Replica	Replicates databases with other servers	
Router	Routes mail to other servers	
Sched	Free time manager for booking meetings	
Server	Used for the Domino user address spaces as well as comms listen and replicate control address spaces	
SH	Used for the Server Base address space	
Stats	Collects Domino statistics	3
Updall	Refresh for index and database views - temporary task	2
Update	Refresh - permanent task	3
Web	Web retriever	

Notes:

1. Task names may appear in various cases. If they are in upper case, they are being loaded from the LPA. See 2.7.3.1, “Domino Program Modules in LPA” on page 14.
2. Some of these tasks are required, while others are optional. Some, such as Fixup, Updall, Compact and Catalog can be scheduled to run at non-peak times.
3. Look at the tasks that your server is running and see if there are some that you do not need or that you can schedule at non-peak times. Stats, Update, AdminP and Billing are some possible candidates that you may not need to run at peak times.

For a detailed description of the various server tasks, see the *Lotus Notes Administrator's Guide*.

5.3.1 Domino Tasks Started

When the Domino server starts, it will start a number of tasks as specified in `notes.ini`. If you issue the server command `show tasks` from the Domino server console, you will see the output in Figure 19 on page 47. This is the set of tasks for the default `notes.ini` file produced by the Domino server setup on S/390.

```
Server name:            Domino1/Poktest

Server directory:       /notesdata
Partition number:       3
Elapsed time:           07:28:02
Transactions/minute:    Last minute: 0; Last hour: 0; Peak: 42
Peak # of sessions:     2 at 03/09/98 03:29:09 PM
Transactions:           133
Shared mail:            Not enabled
Pending mail:  0        Dead mail:  0

      Task                Description

 Database Server      Perform console commands
 Database Server      Listen for connect requests on TCPIP
 Database Server      Idle task
 Database Server      Idle task
 Calendar Connector   Idle
 Schedule Manager     Idle
 Admin Process        Idle
 Agent Manager        Executive '1': Idle
 Agent Manager        Idle
 Stats                Idle
 Indexer              Idle
 Router               Idle
 Replicator           Idle
 Billing              Idle
```

Figure 19. Output From the show tasks Command

Each of these tasks runs in its own OS/390 address space, so you will have multiple address spaces as described in 2.2, “OS/390 Address Spaces” on page 4.

5.3.2 Resource Use by Domino Server Tasks

The resources used by the various Domino server tasks will vary depending on what type of application is being run. As an example, Table 5 shows the resource use by the server tasks for a production Domino server running mail for several hundred users. The table shows the CPU time consumed in a 15 minute interval, and the I/O activity to the Hierarchical File System (HFS).

Table 5 (Page 1 of 2). Resources Used by Domino Tasks in a 15 Minute Interval

UNIX process name	Approximate CPU time used	HFS activity	See note
Adminp	8 sec	read/write - low activity	
Amgr	3 sec	read/write - low activity	
Calconn	7 min	read only	
Chronos	15 sec	read/write	1
Compact	20 sec	read/write	1
Delcnflt	4 min	read only	1
Event	15 sec	read/write	
Replica	1 sec		
Report	4 sec		
Router	3 sec	read/write	

Table 5 (Page 2 of 2). Resources Used by Domino Tasks in a 15 Minute Interval			
UNIX process name	**Approximate CPU time used**	**HFS activity**	**See note**
Sched	50 sec	read/write	
Server	20 sec	read/write - low activity	2
Server	5 min	read/write	2
Server	5 min	read/write	2
Server	5 min	read/write	2
Server	5 min	read/write	2
Update	5 min	read/write	

Notes:

1. This is a temporary process.
2. One of the server processes is part of the base Domino server. The other server processes run the user connections. See 2.2.1, “Address Spaces for User Connections” on page 5.

5.4 Replication

Replication is one of Domino's key functions. Replication ensures that all users on a global network can access the latest version of a database without having access to one centrally located database. By accessing a local copy users can get better response times.

During replication, the replica copies of a database exchange modifications in one or both directions. There are two types of replication:

- Replication between servers
- Replication between a client and a server

5.4.1 Replication between Servers

Replication between servers is managed by the Notes administrator. An administrator can set up the system so that databases will be automatically replicated at intervals.

In a large Domino network, replication is a important factor in planning network traffic and system load. You should carefully plan the times, frequency and type of database replication. To optimize this area, the Notes administrator should work closely with the OS/390 system programmer and the network specialist. There are tools available to customize and optimize the process of replication for your system.

Server-to-server replication is controlled by the following parameters in the `notes.ini` file:

Allow_Access Lists groups or servers that can access the server
Deny_Access Lists groups or servers that cannot access the server
Domain Determines the domain for the server
KeyFilename Location of server or user ID file
Log_Replication Log start and stop of replication and display on console

Repl_Error_Tolerance	Number of allowed errors of the same type
ReplicationTimeLimit	Time limit for replication
Replicators	Number of concurrently running replication processes
Repl_Push_Retries	Number of times attempting to replicate to destination server
ServerKeyFileName	Specifies server ID file
ServerName	Full hierarchical name of the server
ServerNoReplRequests	Forces the server to refuse replication requests
ServerPullReplication	Specifies that scheduled replication initiated from this server be pull-push
ServerPushReplication	Specifies that scheduled replication initiated from this server be push-pull
Servertasks	Tasks to start automatically when the server starts

5.4.1.1 Number of Replicator Tasks

On a processor with a single engine, run a single replicator task. The task will handle replication requests one at a time.

On a processor with multiple engines we recommend that you start with a single replicator task. If you have a replication backlog you can start more replicator tasks, but the maximum number should be one less than the number of engines. This will ensure that one engine is always available for other, more time critical, work.

For more information see the *Lotus Notes Administrator's Guide*.

5.4.2 Replication between a Client and a Server

Databases can also be replicated between a client and a server. Creating a local copy of a database on your workstation will significantly increase performance and availability. It will reduce the amount of time it takes to open and access the database (such as your mail), as well as let you continue to process your mail when the network or server is temporarily unavailable.

Client replication is managed from the Notes client. The Notes administrator cannot control this type of replication, other than by restricting databases from being replicated and by providing guidelines to the users.

To tune client-to-server replication, consider the following:

Replication task	Client replication runs as a user task, and does not use the replica task. During the creation of a new replica, a task with a program name of `server` is created for building the new replica.
Large database	A single replication can take a large amount of CPU time and do many I/Os to the HFS, depending on the size of the database.
Direction of replication	We noted that creating a replica from the server to the client was up to 20% faster in elapsed time than creating a replica from the client to the server. You should consider this in your application design.
Resource used	Each replication uses resources even when there are no changes to the databases.

5.4.2.1 Replication Frequency

We did some tests to identify the resources used by replication when there are no changes to the database. We set up three Notes users and got them to do frequent replication of their mail databases. There were no changes to the database during the test. For *each* replication we found the following:

- There were 9 Domino transactions recorded.
- There was 0.2 seconds of CPU time consumed on a 9672-RX4 processor. Note that this is 0.2 seconds of a single engine. This time also includes other background Domino tasks that ran at the same time, although we did our best to minimize those.
- There was 2.1 KB of data transferred, as reported by Domino.
- There were 33 reads and 10 writes to the HFS.

Therefore, if 1000 Notes clients automatically replicate one database every five minutes, then the resources used in each hour would be approximately:

- 2400 seconds of CPU time on a 9672-RX4 processor, which is 67% of a single engine on that processor
- 25 MB of data transferred across the network

You could use a significant amount of resources replicating databases between the clients and the server even when there are no changes to the data. Therefore, you should encourage your users not to do unnecessary replication. The replication frequency is specified in the `repeat every` keyword in the Notes client's location document. This value should not be set below 15 minutes. Thirty minutes or higher is recommended.

5.4.2.2 Monitoring Replication Activity

There is no administration tool to limit the frequency of replicating databases to a client (the minimum value is one minute). However, the Notes administrator can see the replication activity in the log.nsf database. You can request that client session events are logged at server setup time, or with the `notes.ini` parameter `Log_Sessions`, or with the `Log_Replication` parameter.

We recommend using the Log_Replication parameter and setting it to 1. This will reduce the amount of message traffic at the server console (compared to Log_Sessions), since there won't be messages logged every time a user opens a database on the server for non-replication purposes. Log_Replication will show more details about replication, depending on the value you select.

You will see messages in the Notes log for each client to server replication. Figure 20 on page 51 shows that documents for the user Ingo ik Karge are checked for replication every minute.

```
03/10/98 06:01:02 PM  Opened session for Ingo ik Karge/Poktest2 (Release
03/10/98 06:01:04 PM  Closed session for Ingo ik Karge/Poktest2
Databases accessed:     3   Documents read:     0   Documents written:
03/10/98 06:01:32 PM  Database mail/draisch.nsf created by Ingo ik Karge/
03/10/98 06:02:05 PM  Opened session for Ingo ik Karge/Poktest2 (Release
03/10/98 06:02:06 PM  Closed session for Ingo ik Karge/Poktest2
Databases accessed:     3   Documents read:     0   Documents written:
03/10/98 06:02:48 PM  Database mail/bmacfade.nsf created by Ingo ik Karge
03/10/98 06:03:08 PM  Opened session for Ingo ik Karge/Poktest2 (Release
03/10/98 06:03:09 PM  Closed session for Ingo ik Karge/Poktest2
Databases accessed:     3   Documents read:     0   Documents written:
```

Figure 20. Client Replication Events in Notes Log

Thus you can see how often your users have set up replication to run on their Notes clients.

5.4.2.3 Access Database on the Server or Replicate and Access Locally?

The decision on whether it is more efficient to access a database on the server or replicate it to the client and access it locally depends on many factors such as:

- Whether the user will access all of the data in the database, or just a subset of it
- How often the user needs to replicate to have up-to-date information
- Whether users all access the server at the same time, resulting in a large workload peak on the server which could be eliminated if replication is used

5.5 Indexing

Indexing server tasks are used to update and repair the views and the index of a database after the database has been changed. Update and Updall are the two tasks responsible for this process.

Only a user with designer or manager access to the database can create a new index for a database. You also need this access level to create a new database on the server. Therefore, the Notes administrator can select people who can create new indexes and full text indexes. That is important because indexing can use a significant amount of system resource.

Chronos is another server task that Domino runs hourly to update the index.

5.5.1 Update and Updall Tasks

Update is a standard server task and can be defined in `notes.ini`. This task is usually automatically started when the server is started. You can run more than one update task. However, each update task will use additional resources, so plan this carefully.

Updall updates all views that have been accessed once and all full text indexes for all databases on the server. It runs as a temporary server task. After finishing its work the task is shut down. Updall is also defined in `notes.ini`. Check whether you need this function in your environment.

Both Update and Updall can be scheduled using the `ServertasksATx= Taskname` parameter in `notes.ini`.

Note: The Updall server task has a number of arguments. We recommend that you run Updall without any arguments specified. In particular beware of the `-R` argument. It rebuilds from scratch all database views that have been accessed at least once and updates all full text indexes. This can be time-consuming and affect server performance, so use it only as a last resort to solve corruption problems.

5.5.2 Full Text Index

A *full text index* is a collection of files that lets users search database information. Only a user with designer or manager rights can create a full text index.

The creation of a full text index takes system resources:

- CPU cycles are used to create the index.
- I/Os are generated to build the index.
- DASD space is used to hold the index.

Depending on the size of the database, and the amount of text, indexing can take a significant amount of resource. Therefore we recommend:

- Run indexing at times of low system load.
- Limit which users can do indexing.

5.5.3 Number of Indexing Tasks

On a processor with a single engine, run a single indexing task. The task will handle indexing requests one at a time.

On a processor with multiple engines we recommend that you start with a single indexing task. If you have an indexing backlog you can start more indexing tasks, but the maximum number should be one less than the number of engines. This will ensure that one engine is always available for other, more time critical, work.

5.5.4 Indexing Parameters

We recommend the following parameter settings for indexing:

Server_Name_Lookup_Noupdate=1

You should set this value to 1 in `notes.ini`. This allows users to access the name lookup views in the Public Address Book while the Domino server is running the view index process.

If this parameter is not set, during the period when name lookup views are being indexed they are locked and not accessible to the users. This process may take several minutes, and the user would not be able to look up names in the Public Address Book during that time.

Update_Suppression_Time

This parameter specifies the delay time between full text index and view updates. Increase the value to index less often.

Update_Suppression_Limit

This parameter overrides `Update_Suppression_Time` and forces an update of the views when a specified number of updates have been received.

5.6 Compact

When documents and attachments are deleted from a Notes database there are blocks of unused space left within the database. Compact is the Domino task to get back this unused space. Databases should be compacted periodically to reclaim this space. Statistics reports can be used to monitor databases sizes.

Compact can be run:

- Against a specific database.
- As a server task called Compact. You can start it with the `load compact` command from the Domino console. All databases on the server are compacted. You can use parameters when running this task.

Compact is not automatically started with the default `notes.ini`. It runs as a temporary task.

Compaction can use a lot of processor time and create a large amount of I/O activity, depending on the amount of data being compacted. Therefore, you should run it at a quiet period on the system.

We recommend that you run compact at least monthly to recover unused space, and more often if there is a lot of deletes in your databases. You should also run Compact for a specific database after a lot of delete activity.

Notes:

1. Compacting a database creates a temporary copy of it in the same directory as the original. Make sure there is enough disk space available for the temporary copy in the HFS data set, especially if the database is large.
2. If you use Domino Release 4 to compact a database created with Notes Release 3, the database is converted into Domino Release 4 format.

See the *Lotus Notes Administrator's Guide* for more information on Compact.

5.7 Partitioning

Partitioning is one of the Domino Advanced Services functions. It allows you to run multiple Domino servers on a single OS/390 system. There are several benefits of partitioning, including:

- Isolation of workloads from each other
- Consolidation of multiple Domino servers quickly and easily
- Minimizing costs of running multiple servers
- Scalability, to be able to run many users on a single server
- Higher availability, by isolating workloads from each other

See *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083, for a longer discussion of partitioning. In the next section, we look at the performance aspects of partitioning.

5.7.1 Costs of Partitioning

Running partitioned servers uses increased resources in some areas:

- Disk space to hold the server files and multiple copies of databases. Each server must have its own set of files and databases. There is no sharing between servers.
- Processor memory to run the server. There is a base memory requirement for each server.
- Processor cycles for communication between servers. This would include mail routing and database replication. In a single server, these tasks would not be needed.

5.7.2 Efficiency of Partitioning

Apart from the benefits of partitioning already discussed, there is also a potential performance benefit. Domino is still a relative newcomer to very large systems. Lotus and IBM are aware of improvements that can be made to improve performance on large servers, and design changes are planned.

Tests on IBM's S/390 and AS/400 servers, run with Domino Release 4.51 and many thousands of users, have shown that you can support more users on a server by running multiple Domino servers rather than a single server. We believe this will apply to other Domino server platforms also.

We do not know the exact reason for this, but it is possible that the way in which a single Domino server uses memory, locks or queues could become less efficient with a high workload. We certainly saw this in the past with OS/390 (and MVS) as we grew the size of the servers, and we made design changes over the years to address these areas.

Therefore, as discussed in 6.4, “The Use of Domino Partitioned Servers” on page 69, we recommend that currently (with the Domino server Release 4) you should plan for each server to support approximately 1,000 registered users. This is provided as a general guideline only; if your users are less active, then you can plan on supporting more of them on each server. Note that this is a tradeoff, since having more servers increases resource use, as previously discussed.

5.7.3 Parameters That Affect Partitioning

These parameters in `notes.ini` affect partitioning performance:

- Server_MaxUsers
- Server_Max_Concurrent_Trans

Note: On a S/390 system, be sure to set Server_Max_Concurrent_Trans to -1, which is the disabled setting. If you set it to other values (as recommended in the *Lotus Notes Administrator's Guide*) you will impede performance on the S/390 platform.

5.8 Clustering

Clustering is another of the Domino Advanced Services functions. A cluster is a group of Domino servers that work together. The servers pass database updates between each other, using real-time replication, so that the data on the various servers is kept in synchronization. Up to six Domino servers can be in a single cluster. These servers can be running on different platforms.

The benefits of clustering are:

- Higher availability, since, if one server fails, users can automatically switch to another server in the cluster
- Workload balancing across multiple servers, resulting in better performance
- Scalability, by allowing multiple servers to act as one

See *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083, for a longer discussion of clustering. We look at the performance aspects of clustering next.

5.8.1 Cluster Components

The following additional tasks run on each server in a cluster:

- The Cluster Administration Process (Cladmin) runs temporarily during creation of the cluster.
- The Cluster Database Directory Manager (Cldbdir) tracks the state of all the members in the cluster. You can get status information with the `show cluster` command.
- One or more Cluster Replicator (Clrepl) tasks run the real-time replication.

You will also see additional databases created:

- CLUSTA4.NSF contains analysis data on the configuration of the cluster.
- CLDBDIR.NSF is the Cluster Database Directory that contains cluster status information.

When a cluster is first defined on a server, the two cluster tasks Cldbdir and Clrepl are started. They are also added to the `ServerTasks` parameter in `notes.ini` so that they are started automatically when the server is started in the future.

You can also start additional cluster replicator tasks. The *Lotus Notes Administrator's Guide* recommends that on each server in a cluster you should run one fewer cluster replicator tasks than there are servers in the cluster. In this way there is one cluster replicator task to handle replication to each other server in the cluster.

5.8.2 Parameters That Affect Clustering

These parameters in `notes.ini` affect clustering:

- Cluster_Admin_On
- Server_Cluster_On
- MailCluster_Failover
- Server_Restricted

See the *Lotus Notes Administrator's Guide* for more details.

5.8.3 Costs of Clustering

Running clustered servers uses increased resources in the following areas:

- Processor cycles, for sending database updates to other servers and applying them to all copies of the database.
- I/O activity, since updates are done multiple times, once on each server.
- Disk space, to hold the multiple copies of each database. Each server must have its own set of files and databases. There is no sharing between servers.
- Network traffic, to pass the updates to all servers in the cluster. Lotus recommends providing a separate network for the server-to-server traffic in a cluster, to ensure good performance.

For each database you can decide which servers in the cluster contain a replica, from only one server to all servers. Make that decision carefully for each database in order to limit the number of updates being passed through the cluster and therefore minimize the resource use.

Note: If you are also running partitioned servers, you must take into account the increased resource for that also.

5.8.4 Monitoring Clustering

You can get statistics on clustering from the log database `log.nsf` and from the output of the Domino server commands:

- `show tasks`
- `show statistic replica`
- `show statistic <servername>`

To look for replication backlogs issue the server command `show statistic replica`. Look at the `Replica.Cluster.WorkQueueDepth` value, which shows the current number of modified databases waiting for replication. If this value is consistently greater than zero you should consider running additional cluster replicator tasks.

5.9 Billing

Billing is the third Domino Advanced Services function. It collects information that is useful for understanding who uses Domino resources, and is therefore useful for billing purposes. The information is placed in a billing message queue. The billing server task periodically moves the data from the queue and writes records into the billing database `billing.nsf`.

On OS/390, the setup process automatically adds the billing task to `notes.ini` if you install the Domino Advanced Services files. You can stop and start this task using the Domino server commands `tell billing quit` and `load billing`.

You specify which types of information you wish billing to collect using the BillingClass parameter in `notes.ini`.

Since collecting billing information will use system resources, we recommend that you review the information being produced. Collect only those billing classes that

you will use. If you do not need any of the billing information, remove the billing task from the `ServerTasks` parameter in `notes.ini`.

5.9.1 Parameters That Affect Billing

These parameters in `notes.ini` affect billing:

- BillingClass
- BillingSuppressTime
- BillingAddinOutput
- BillingAddinWakeup
- BillingAddinRunTime

See the *Lotus Notes Administrator's Guide* for more details. Also see *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083 for a longer description of billing.

5.10 Domino Web Server

The Lotus Domino Server Release 4.5 and later releases provides Web server capability. Web browsers can connect to the Domino server and access Notes databases. Most functions available to a Notes client are available to a Web browser, although some, such as database replication, are not supported due to technology differences between the two user interfaces. The Domino server provides Web server function, and also provides a conversion between the Web HTTP and HTML protocol standards and the Notes internal standards. This is shown in Figure 21.

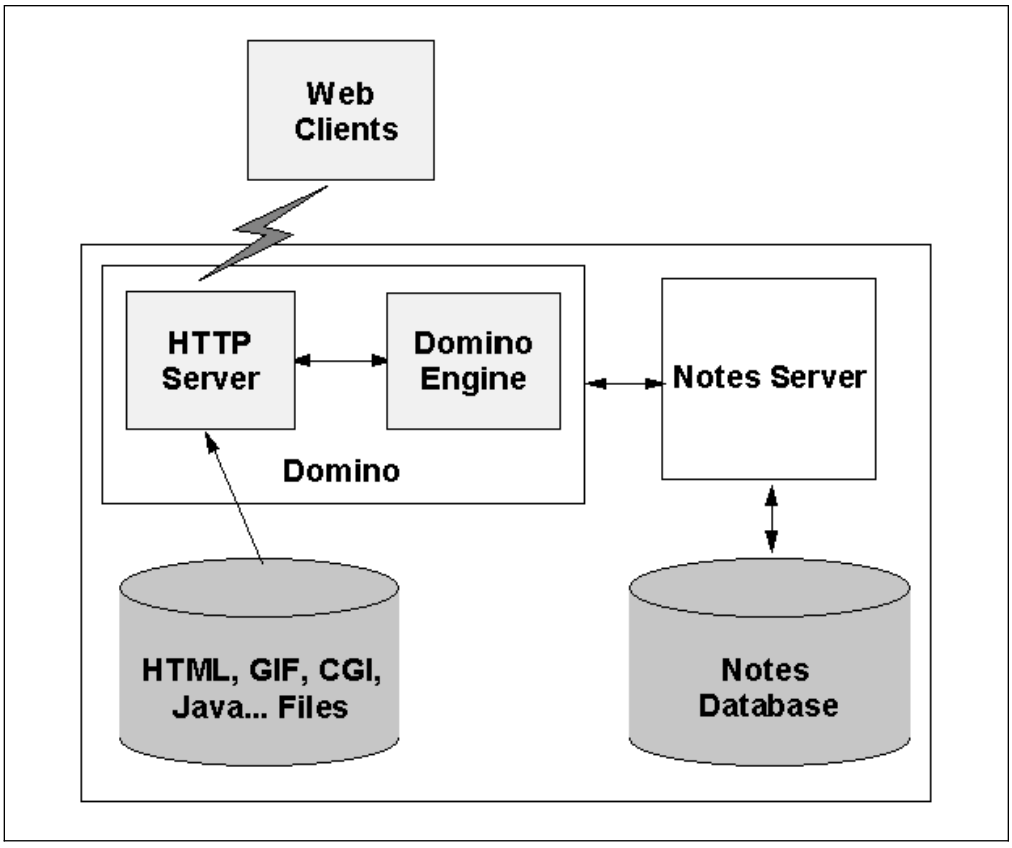

Figure 21. How the Domino Web Server Works

The web server task is called `http`. You can start it by adding it to the `ServerTasks` parameter in `notes.ini`, or by issuing the Domino server command `load http`. OS/390 creates an address space for the http task.

5.10.1 Web Server Parameters

The http task reads the parameters in the http server section of the server document in the Public Address Book. There you can specify the minimum and maximum number of threads that you want the http task to run. When a request is received from a Web browser, it is run on one of these threads. When the response is sent to the browser, the thread is released and is available for the next request. There is no long-running user connection as there is for a Notes client.

If there are no available threads when a request is received, the request is queued until a thread is released. Therefore, you should ensure that you have the maximum number of threads set sufficiently high. You can see the number of threads used by issuing the command `show stat domino` on the Domino server console.

You should set the minimum number of threads high enough to handle a typical workload. This will reduce the need to create threads, and thus minimize the system load.

For more information on the Web server task and related tuning parameters, see *Working with Lotus Notes and the Internet*.

Also see the tuning tips for the Lotus Domino Go Webserver on S/390 on the Web at the URL: `http://www.ics.raleigh.ibm.com`.

5.11 Monitoring Domino

Domino provides a set of tools for monitoring the server. Some of the tools run automatically. Others have to be started using a server command or started task. We discuss the various sources of monitoring information in the following sections. These tools and databases will be familiar to a Notes administrator. Administrators may use additional tools to analyze the data. You can also get information about the Domino server by using the `show` command.

Note: Do not collect information that you do not need. Data collection takes system resources.

See the *Lotus Notes Administrator's Guide* chapter 11 for more details about monitoring a Domino server.

5.11.1 Notes Log

Every Domino server logs information about server activities in the Notes log database log.nsf. The information includes:

- Database usage by user
- Database size
- Mail routing, replication and other events
- Usage of the system by user, including
  - User name
  - Connect time
  - Number of documents read and written, by database

- Amount of data transferred across the network
- Number of transactions run

You can choose to collect replication and client session event records when you initially set up the server. You can change these settings later with the `Log_Replication` and `Log_Sessions` parameters in `notes.ini`.

5.11.2 Statistics and Events Log

The statistics and events database events4.nsf is used to configure Notes server event handling, statistic monitoring, ACL monitoring, and replication monitoring. The database contains the names of all statistics monitored by the server, and thresholds for producing event records. It also contains error and status messages from the server.

This database can be used to:

- Configure Notes server event handling
- Look up information about a specific statistic or event message
- Set statistic thresholds
- Assign types and severities to server events

You run events by running the server task `Event`.

5.11.3 Statistics and Reporting Database

If you run the Reporter task on the server, a set of statistic reports will be created at scheduled intervals. These are put in the statistics and reporting database statrep.nsf. You can specify the interval between records. For information on setting up the Reporter task see the *Lotus Notes Administrator's Guide*.

The information contained in the statistics and reporting database is shown in Figure 22 on page 60. This is an exported version of the data for one collection interval. See 5.12.1, “Exporting the Contents of a Database” on page 63.

```
Agent.Daily.AccessDenials:  0
Agent.Daily.ScheduledRuns:  557
Agent.Daily.TriggeredRuns:  145
Agent.Daily.UnsuccessfulRuns:  0
Agent.Daily.UsedRunTime:  2258 Seconds
Agent.Hourly.AccessDenials:  0
Agent.Hourly.ScheduledRuns:  58
Agent.Hourly.TriggeredRuns:  1
Agent.Hourly.UnsuccessfulRuns:  0
Agent.Hourly.UsedRunTime:  17 Seconds
Calendar.Total.All.Appts.Reservations:  52794
Calendar.Total.All.Users.Resources:  480
Calendar.Total.Appts:  52794
Calendar.Total.Reservations:  0
Calendar.Total.Resources:  0
Calendar.Total.Users:  480
Database.BufferControlPool.Peak:  3270300
Database.BufferControlPool.Used:  1143584
Database.BufferPool.Maximum:  200000000
Database.BufferPool.Peak:  163872000
Database.BufferPool.PerCentReadsInBuffer:  99
Database.BufferPool.Reads:  200013
Database.BufferPool.Used:  163171718
Database.BufferPool.Writes:  84578
Database.DbCache.CurrentEntries:  47
Database.DbCache.HighWaterMark:  587
Database.DbCache.Hits:  11764
Database.DbCache.InitialDbOpens:  34449
Database.DbCache.Lookups:  52602
Database.DbCache.MaxEntries:  508
Database.DbCache.OvercrowdingRejections:  0
Database.NIFPool.Peak:  4120578
Database.NIFPool.Used:  1280960
Database.NSFPool.Peak:  2681646
Database.NSFPool.Used:  729963
Disk./.Free:  17117184
MAIL.AverageDeliverTime:  157
MAIL.AverageServerHops:  2
MAIL.AverageSizeDelivered:  32
MAIL.Dead:  0
MAIL.Delivered:  5614
Mail.Domain:  IBMUS
MAIL.MaximumDeliverTime:  10400
MAIL.MaximumServerHops:  16
MAIL.MaximumSizeDelivered:  9111
MAIL.MinimumDeliverTime:  1
MAIL.MinimumServerHops:  1
MAIL.MinimumSizeDelivered:  1
MAIL.PeakByteTransferRate:  169176
MAIL.PeakMessagesTransferred:  95
MAIL.PeakMessageTransferRate:  1
Mail.PeakMessageTransferTime:  03/02/98 03:30:48 PM
Mail.PeakTotalBytesTransferred:  10150597
MAIL.TotalFailures:  42
MAIL.TotalKBTransferred:  51203
MAIL.TotalRouted:  8954
MAIL.Transferred:  2333
MAIL.Waiting:  0
MAIL.WaitingRecipients:  0
Mem.Allocated:  209267848
Mem.Allocated.Process:  24003886
Mem.Allocated.Shared:  185263962
Mem.Availability:  Plentiful
Mem.PhysicalRam:  67108864
NET.Log.X99XX999/01/M/IBM.UnwrittenEntries:  2
```

Figure 22. Exported Data from the Statistics and Reporting Database (Part 1 of 2)

```
NET.TCPIP0A.BytesReceived:  180011959
NET.TCPIP0A.BytesSent:  812272784
NET.TCPIP0A.Sessions.Established.Incoming:  19925
NET.TCPIP0A.Sessions.Established.Outgoing:  2343
NET.TCPIP0A.Sessions.Limit:  65535
NET.TCPIP0A.Sessions.LimitMax:  65535
NET.TCPIP0A.Sessions.LimitMin:  10
NET.TCPIP0A.Sessions.Peak:  306
NET.TCPIP0A.Sessions.Recycled:  0
NET.TCPIP0A.Sessions.Recycling:  0
Reporter.Time.Analysis:  03/02/98 01:00:01 PM
Reporter.Time.Collected:  03/02/98 11:57:51 PM
Reporter.Time.Elapsed:  10:57:51
Server.Administrator:  ******** this has been removed
Server.BootID:  6367873
Server.BusyTimeQuery.ReceivedCount:  1180
Server.Location:  Poughkeepsie
Server.Name:  CN=X99XX999/OU=01/OU=M/O=IBM
Server.OpenRequest.MaxUsers:  0
Server.OpenRequest.PreV4Client:  0
Server.OpenRequest.Restricted:  0
Server.OpenRequest.V4Client:  33839
Server.Path.Data:  /x99xx999
Server.Ports:  TCPIP0A
Server.PoweredBy:  Notes
Server.Sessions.Dropped:  6
Server.Task: *********** all the tasks for connected users and
             *********** server tasks are listed here
             *********** these have been removed
Server.Tasks:  44
Server.Time.Start:  03/02/98 12:56:46 PM
Server.Title:  Poughkeepsie S390 Mail Server
Server.Trans.PerMinute:  61
Server.Trans.PerMinute.Peak:  1676
Server.Trans.PerMinute.Peak.Time:  03/02/98 01:18:28 PM
Server.Trans.Total:  229114
Server.Users:  18
Server.Users.1MinPeak:  89
Server.Users.1MinPeakTime:  03/02/98 01:17:32 PM
Server.Users.5MinPeak:  145
Server.Users.5MinPeakTime:  03/02/98 01:17:32 PM
Server.Users.Peak:  273
Server.Users.Peak.Time:  03/02/98 02:20:01 PM
Server.Version.Notes:  Release 4.5.3
STATS.Time.Current:  03/02/98 11:57:51 PM
STATS.Time.Start:  03/02/98 12:41:18 PM
$UpdatedBy:  CN=X99XX999/OU=01/OU=M/O=IBM
```

Figure 23. Exported Data from the Statistics and Reporting Database (Part 2 of 2)

5.11.3.1 Reporting Period

Be careful about the reporting period for the various statistics. There are several types of statistics in the interval record:

- A sample at the end of the interval
- Accumulation for the duration of the interval
- Maximum value since the server started

Unfortunately there are not many values accumulated for the duration of the interval, which is the most useful period for performance monitoring.

5.11.3.2 Useful Domino Statistics

Statistics that you may find useful for performance monitoring are:

Server.Trans.Total

The transaction total is the most useful statistic collected. This is the transaction count since the server started, as sampled at the end of the interval. By looking at this value for adjacent interval records you can determine the number of transactions processed in an interval.

The transaction rate should be an indication of the use of the server. This can then be related to information from SMF and RMF records to determine the CPU use per transaction. The transaction rate also allows you to check whether the load across several servers is balanced.

Server.Trans.PerMinute.Peak

This is the peak number of transactions during a one-minute interval since the server started. It is *not* the peak during this interval. You are given the date and time this occurred.

Server.Users

This is the number of users with connections to the server at the end of the interval when statistics were collected.

Server.Users.1MinPeak

This is the peak number of users during a one-minute interval since the server started. It is *not* the peak during this interval. The peak number of users in a five-minute interval is also reported. You are given the date and time these occurred.

There are also counters for the number of calendar and mail requests, and also agent processes. These may be of use if you are investigating the average profile of your mail users.

5.11.3.3 Coordinating SMF, RMF and Domino Data Collection

If you wish to correlate the Domino statistics, such as transaction count, with the SMF and RMF data, you should try to coincide the start of Domino statistics collection with the start of the SMF and RMF intervals. SMF and RMF have controls to synchronize the time of interval collection to the time of day, for example at zero and 30 minutes after each hour. Domino has no such control. It records statistics every so many minutes after the statistics task is started. Therefore, to get Domino to record its records at the same time as SMF and RMF, you must start the statistics task at the desired time.

5.11.3.4 Statistics Commands

You can also use:

- The `load stats <servername>` command to create statistics on demand for a remote server
- The `show statistics` server command to display the complete list of statistics for the server

5.11.4 Billing Database

If you run the Domino billing task the output will be written to the database billing.nsf. The information recorded can include:

Session The session class tracks when Notes sessions start and end, and when any ongoing activities, document editing or replication occur during the session. It records bytes read and written for each user name.

Database The database class tracks when a user or server opens and closes a database, and the duration of use.

Document The document class tracks read and write activity for specified documents.

Replication The replication class tracks when a billing server initiates replication with another server or a workstation.

Mail The mail class tracks when a mail router transfers messages to another Domino server.

Agent The agent class tracks when users or servers run a scheduled agent on a server. It also tracks the elapsed run time of those agents.

5.12 Processing Domino Statistics

To process Domino statistics you can:

- View them in the database.
- Use a tool to analyze the data and produce reports.
- Extract the data into another form for further processing.

5.12.1 Exporting the Contents of a Database

If you wish to combine Domino statistics with SMF data, for example, you will probably need to extract the Domino statistics into a sequential OS/390 file. You can then use current analysis tools to analyze the data.

To export the contents of a database into a sequential file on your workstation:

1. Open the database on the server.
2. Select the view that you want.
3. Select the documents that you want to extract based on the collection time.
4. Select **File/Export**.
5. Set the file name to a file on your workstation.
6. Save as type Structured Text.
7. Press the Export Button.
8. Choose:
   - **Selected documents**
   - **Char code 12** as the interdocument delimiter
   - Wrap words at 75 characters per line
9. Then choose OK.

The file will be written on your workstation. You can then upload it to OS/390 and input it into your analysis program.

A sample of an exported file for one collection interval is shown in Figure 22 on page 60.

Chapter 6. Capacity Planning

This chapter provides capacity planning guidelines in the areas of:

- Processor capacity
- Processor storage (or memory)
- The use of Domino partitioned servers
- The use of S/390 logical partitioning (LPAR)
- DASD space
- Network capacity

We also provide information on a large IBM production Domino server on S/390, and show how you can get assistance with your capacity planning.

6.1 Recommendations

These are our recommendations for capacity planning:

- Estimate carefully.

  Initial estimates of the resources required to support a Domino server are based on estimates of:

  – Number of registered users
  – Number of connected users
  – Transaction rate generated by the users
  – Average CPU time used per transaction
  – Amount of data you will have, including the average amount of data a user will have in their mail database

  The more accurately you estimate these values, the more accurate your estimate of resource needs will be.

- Base your capacity planning on your own measurements as soon as possible.

  Every Notes environment is different. Therefore, you should monitor your system as soon as you have a significant Domino workload, and then revise your plans based on that. If you have other Domino servers, then you can collect information on your workload profile from them and use that to help validate your assumptions.

- Make use of the capacity planning assistance available.

  See 6.9, “Capacity Planning Assistance” on page 72.

- If you will have more than 1,000 users, consider using Domino server partitioning.

  See 6.4, “The Use of Domino Partitioned Servers” on page 69.

- Remember that these guidelines will change.

  These guidelines were put together in March 1998 for Domino Release 4.5, OS/390 Version 2 Release 4, and 9672 generation 4 processors. We expect significant Domino capacity improvements from changes in:

- The Lotus Domino Server, which has been making major improvements in its scalability since Lotus Notes Server Release 3
- OS/390 and the UNIX Services interface
- OS/390 Hierarchical File System
- System/390 processors
- TCP/IP releases

We have already seen a significant increase in capacity with Release 4.5.3 of Domino for S/390, and the new Perform Locked Operation (PLO) instruction on the 9672 processors.

Allow for improvements in your planning.

- Beware of running your own performance tests and misinterpreting the results.

  It is possible to set up some very simple performance tests. However, we have seen very misleading results when trying to use the run time of a single Domino function to estimate the performance of a production workload. Unless you understand exactly which system resources the function uses, and can accurately relate that to the resources your production workload will use, you can get very erroneous results. In addition, unless you understand how Domino performs the functions, it is possible that you are measuring something quite different from what you think you are measuring. You could, for example, be measuring network performance instead of server performance.

  Realistic tests are, unfortunately, complex and expensive to run. Use the information in this document, and the other information sources listed, to get better capacity estimates for production workloads.

6.2 Estimating Processor Capacity

The processor capacity required for Domino depends on:

- The number of users
- The number of server transactions generated by the users
- The CPU time used per transaction

We discuss each of these items, and then show how to take that information to produce a processor capacity estimate.

6.2.1 Number of Users

You can probably estimate the number of registered users on your server relatively easily: it is the number of users in your company or organization.

6.2.2 Transaction Rate

The rate at which transactions arrive at the Domino server depends on the number of active users and the rate at which they are working. For example, if an installation with 1,000 registered users has 200 active users working at a rate of 90 server transactions per hour, then the server would need to process 18,000 transactions per hour.

Domino does not provide any statistics on active users (a connected user may, or may not, be active), so base your estimates on the number of registered users.

On a production system, you can easily measure the transaction rate. The number of transactions run in an interval can be calculated by subtracting the Domino transaction count at the ends of two consecutive intervals. The Domino transaction count is recorded in the `Server.Trans.Total` statistic, as discussed in 5.11.3.2, “Useful Domino Statistics” on page 62.

You can calculate the transaction rate per registered user by dividing the number of transactions by the number of registered users.

Note that the transaction rate per user could be higher during a pilot when users are “experimenting” with this new tool. Alternatively the transaction rate per user could be lower until you have a critical mass of users on the server.

6.2.3 CPU Time Used Per Transaction

There are many variables that affect the CPU time used by a server transaction. The size of the address book used for address resolution, the size of the items being sent, and the number of items in the user's mail database are just a few of the variables that affect the CPU time per transaction.

You can measure the CPU time used per transaction on your system. For a time interval you select you can get the total CPU time used by Domino by adding up the CPU time used by each of Domino's address spaces. You can get that information from the SMF records - see 2.8.2, “Monitoring OS/390 with SMF” on page 21. You should include a proportion of the CPU time used by OS/390 and TCP/IP, depending on the workload mix on the system. Divide the CPU time used by the number of Domino transactions recorded for the same interval, as described above.

6.2.4 Calculating the Processor Capacity Needed

Based on the number of registered users, transaction rate per user and CPU time used per transaction, you can now calculate the processor capacity needed. It is the number of registered users x transaction rate per user per hour x CPU time used per transaction. This is the number of CPU seconds per hour needed on a specific processor. The percentage utilization of the processor is: (CPU seconds per hour / number of engines / 3600) x 100%.

To estimate the capacity on any other S/390 processor, you use the relative throughputs of the two processors. The Internal Throughput Ratios (ITRs) of many S/390 processors are published in the Large Systems Performance Report (LSPR) on the Web page `http://www.s390.ibm.com/sections/lspr.html`

6.2.5 Consider the Peak Times

There are times during the day when the capacity needed by Domino is higher than the average and there are times when it is lower.

If Domino will represent less than 50% of the total workload on the S/390 server, and where the other workload is of lower priority, you can probably use the average workload for planning. The OS/390 Workload Manager will manage the overall resources and make them available to Domino at the peak times if you define the Domino workload to have a high priority.

For dedicated Domino servers, or for S/390 systems where Domino is not the highest priority workload, a peak load factor should be considered to ensure consistently good response times.

Determining the peak load factor to use requires some understanding of the work schedule of the users. If the users are teachers that can only use the system between classes, or at the beginning or end of the day, then the peak load factor could be much higher than two. If the majority of the users are using a local copy of their mail database on their workstation, then the peak load factor could be much less than two.

The use of a local copy by a high percentage of users results in a more uniform workload throughout the business day, but could also extend it. Users may start their replication earlier so that it finishes before they arrive at work.

If Domino will share a S/390 server with other high priority workloads, consider the Domino peak in relation to the other workloads. If the users of your Domino system are also the users of your CICS or IMS system, they are unlikely to be using both systems heavily at the same time of day. Since the peaks of the two workloads occur at different times of the day, the total capacity you need is less than the sum of the peaks for the two workloads. By recognizing this, many S/390 users have less capacity installed than if they ran each workload on a separate server.

6.3 Estimating Processor Storage

The amount of storage that is needed on the S/390 server depends on the number of Domino connected users. A Domino user is connected to the server from the time that they initially access it until the server ends the connection based on the `Server_Session_Timeout` parameter in `notes.ini`. A connected user is known to the server and therefore uses processor storage.

Table 6 shows the estimated amount of processor storage and the amount of ESQA that is needed to support Domino on a S/390 server, based on the number of connected users. Note that processor storage is the sum of central storage and paging storage.

The maximum number of connected users would be less than the total number of registered users. A conservative estimate is that 70% of the registered users would be connected during peak activity periods.

Table 6 (Page 1 of 2). Estimated Central and Paging Storage in Megabytes

Connected Users	Storage (MB)	ESQA (MB)
100	70	0.3
200	90	0.7
500	200	2.1
1,000	300	5.5
1,500	500	10.3
2,000	600	16.5
2,500	700	24.0

Table 6 (Page 2 of 2). Estimated Central and Paging Storage in Megabytes		
Connected Users	**Storage (MB)**	**ESQA (MB)**
3,000	900	33.0
3,500	1,000	43.3
4,000	1,100	54.9
4,500	1,300	68.0
5,000	1,400	82.4
6,000	1,700	115.4
7,000	2,000	153.8
8,000	2,300	197.8
9,000	2,600	247.2
10,000	2,900	302.1

Notes:

1. The storage estimates are the additional amount of storage required to add a single Domino server to an existing OS/390 system that has the UNIX Services running. For a new OS/390 system you need to add the base OS/390 storage requirement of approximately 128 MB.
2. The estimates are based on measurements of active users doing mail.
3. The estimates are based on Domino Release 4.5 running on OS/390 Version 2 Release 4. Storage requirements may change for later releases.
4. These storage estimates are comparable to the storage estimates for other Domino platforms.

6.4 The Use of Domino Partitioned Servers

Domino partitioned servers provide the ability to run up to 16 Domino servers on one OS/390. If you plan to have more than 1,000 registered users on your OS/390 system we recommend that you consider running multiple Domino servers to support them.

With Domino Release 4.51 on OS/390 Version 2 Release 4, we have observed that optimum performance is achieved when the number of concurrent active users per server is kept below 550. For an installation where 40% of the registered users are active, that would equate to 1,375 registered users per server.

Since the number of concurrent active users will vary we believe that 1,000 registered users per Domino server is a sensible planning number based on current experience with current products. You may be able to achieve much more than that with your workload, but it would be prudent to keep the possibility of multiple servers in mind. For future releases we would expect to increase this planning number.

Some of the reasons behind this recommendation are given in 5.7.2, “Efficiency of Partitioning” on page 54.

6.5 The Use of S/390 Logical Partitioning (LPAR)

The limitations on a single OS/390 are 2 GB of central storage and a 2 GB virtual storage size for address spaces. Given the storage estimates in Table 6 on page 68, if you plan to have more than 10,000 registered users on your OS/390 system we recommend that you consider running multiple OS/390 systems in S/390 logical partitions. This is to minimize OS/390 paging and give good response times to your Domino servers.

Storage use will vary between installations, but we believe that 10,000 registered users per OS/390 system (LPAR) is a sensible planning number based on current experience with current products. You may be able to achieve much more than that with your workload, but it would be prudent to keep the possibility of multiple LPARs in mind. For future releases we would expect to increase this planning number.

You can track the storage use on your system using the high UIC value - see 2.8.1.4, “RMF Processor Storage Report” on page 19.

You can run up to 15 OS/390 systems on a single S/390 server using logical partitioning (LPAR).

6.6 Estimating DASD Space

The amount of DASD space that is needed for Domino depends on the amount of data that you will store. It is the same as for other Domino platforms.

The amount of DASD space required for mail databases depends on the amount of data that each user will store. This is normally in the range of 40 to 100 megabytes per user and will typically grow as the users become more advanced in their use of Domino. Installations with thousands of users will therefore require hundreds of gigabytes, if not terabytes, of DASD space.

OS/390 provides a standard UNIX hierarchical file system (HFS) for applications such as Domino that use OS/390 UNIX Services. You should plan your use of the HFS carefully to avoid the need to move data around later. See section 2.12, “Notes Directory Structure” in *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083 for more information on the HFS.

6.7 Network Considerations

Proper planning of the network will minimize delays due to an overloaded network. As users begin to exploit the advanced functions of Lotus Notes, such as attaching documents, spread sheets and presentations to notes, the amount of data flowing through the network will grow.

One difference with a large Domino server, such as S/390, is that you will have a lot of network traffic into one point. S/390 users already understand how to do that, but it is different from an implementation where there are many small servers. Make sure that you configure at least *two* paths into the S/390 server, for both performance and availability.

A large server also results in less network traffic overall. It eliminates the server-to-server network traffic for mail routing, database replication and calendar tasks.

6.8 The Production Domino for S/390 Server Running in Poughkeepsie

Many people have shown an interest in the production Domino server on S/390 that IBM has been running in Poughkeepsie, New York, USA since November 1996. This is a true production Domino server and is run as part of IBM's rollout of Lotus Notes. It has also been used as a production proving ground for Domino since 1996. We provide this information as at March 1998 to show you a large Domino for S/390 system.

6.8.1 Hardware and Software Environment

The Poughkeepsie production system is running on an IBM 9672-RY5 with four logical partitions (LPARs). This is shown in Figure 24.

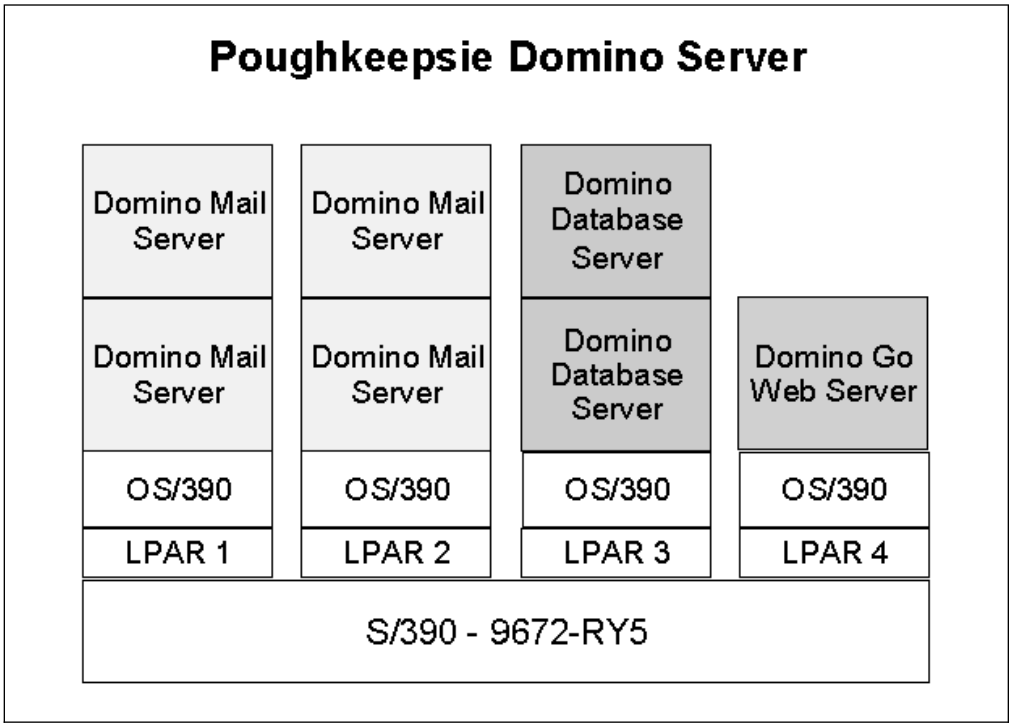

Figure 24. Poughkeepsie Production Domino Server

The system is not fully used. We saw a significant reduction in CPU capacity used when we migrated from Lotus Domino Release 4.5.1 to Release 4.5.3 and enabled the Perform Locked Operation (PLO) instruction.

Three of the LPARs are used for Domino servers. Each LPAR runs two Domino servers, using Domino server partitioning. Two of the LPARs run two mail and calendar servers each. The third LPAR runs two Domino servers for database applications.

People often ask why we run two LPARs for the four mail servers, rather than all four servers in one LPAR. The reason is historical, since when we first split the users across two Domino servers in the summer of 1997, Domino partitioning was

not available. Therefore, we had to use LPAR. We believe that we could run all four servers in one LPAR. However this is a production server and reducing the number of LPARs has not been a high priority.

Each LPAR has 2 GB of central storage and 1.2 GB of expanded storage. The two mail server LPARs have 7 logical processors defined and an LPAR weight of 35. The database LPAR has 2 logical processors and has a weight of 15. A total of 790 GB of DASD has been allocated to the three LPARs. Another 136 GB of DASD will be added to accommodate another 1200 registered users. Each registered user is allocated 100MB of DASD space. Some users have already exceeded this.

6.8.2 Workload

The S/390 server currently supports 3,072 registered users. A typical day's workload consists of:

- More than 40,000 mail items delivered.
- The average size of a mail item is 50 KB.
- The maximum size of a mail item of 23 MB.
- More than 65% of the registered users are active during the peak hour.
- More than 70% of the registered users use the calendar daily.
- More than 600 server transactions per registered user are run per day.
- More than 50 server transactions per registered user per hour are run during the nine prime time hours.

6.9 Capacity Planning Assistance

An initial capacity planning estimate for a Domino workload is available through your IBM or Lotus Marketing Representative. A more detailed analysis is available through various chargeable IBM offerings.

6.9.1 CALLS390

Your IBM or Lotus Marketing Representative can request an initial sizing for Domino on S/390 by requesting a form from CALLS390, completing this form with your help, and submitting the completed form to CALLS390. If you can estimate the number of registered, connected, and active users you will have and the type of work you expect them to do, CALLS390 will provide you with an estimate of the processor size you will need to support this workload. IBM and Lotus employees can contact CALLS390 at the Lotus Notes address `CALLS390 CALLS390/Poughkeepsie/IBM`.

6.9.2 SNAP/SHOT

Host SNAP/SHOT Services System Capacity Planning supports many applications, including Domino for S/390. SNAP/SHOT can show you the effect of adding Domino to your existing S/390 workload and can help you determine what changes, if any, should be made to your S/390 and network configuration to accommodate this additional workload. For more information, visit their Web site or call them toll-free as follows:

Web Site `http://www.as.ibm.com/asus/performance.html`

Telephone (USA) 1-800-426-4682
Telephone (outside) (919) 301-4141

6.9.3 IBM Testing Services

IBM Testing Services provides the hardware, software, drivers, network, and expertise to perform comprehensive stress testing. Whether for proof of concept testing, capacity and performance testing, or full stress testing, there are times when only a “real” test will do. For more information, visit their Web site or call them toll-free as follows:

Web Site `http://www.as.ibm.com/asus/svs.html`
Telephone (USA) 1-888-474-TEST
Telephone (outside) (301) 240-8535

Appendix A. Workload Manager Example

This appendix contains an example of the workload manager goal mode definitions for a Lotus Domino production environment. There are six Domino servers called ML96NP, ML97NP, ML98NP, ML99NP, DBL90NP and DB99NP. These servers run in three LPARs of a sysplex and are the prime workloads for these LPARs.

A.1 Workload and Service Class Descriptions

Table 7 shows the workloads that the workload manager will use. A workload is a named collection of work to be reported as a unit.

Table 7. Workload Descriptions

Workload	Description
PRIMEAPP	APPC workload
PRIMEBAT	Batch workload
PRIMESTC	STC Workload
PRIMETSO	TSO workload
PRIMOMVS	OpenEdition workload

Table 8 shows the service classes that the different workloads will be assigned to. A service class is a named group of work within a workload with similar performance characteristics.

Table 8. Service Class Descriptions

Service Class	Description
BATCHLOW	Low priority batch
NOTESSRV	Lotus Notes Server
OMVSLOW	Low Priority OMVS Work
OMVSTASK	OMVS Support Tasks
OMVSUSER	OpenEdition users
STCLOW	Low Priority STC
TSOPRIME	Normal TSO users

A.2 Report Classes

Table 9 shows the report classes set up to provide more granular reporting. Report classes allow this reporting for individual Domino servers while managing all the Domino servers the same in a single service class (NOTESSRV).

Table 9 (Page 1 of 2). Report Classes

Report Class	Description
ADSM	ADSM Report Class
DBL90NP	DBL90 DB server

Table 9 (Page 2 of 2). Report Classes	
Report Class	**Description**
DB99NP	DB99 DB server
DFHSM	DFHSM Report class
JES2	JES2 Report class
ML96NP	ML96 Notes servers
ML97NP	ML97 Notes server
ML98NP	ML98 Notes server
ML99NP	ML99 Notes servers
MVSNFS	MVSNFS Report Class
OMVS	OMVS Kernal Report Class
RMFIII	RMF Monitor III Report Class
SYSBMAS	SYSBMAS Report Class
TCPML	TCP/IP Report Class

A.3 Classification Rules

Classification rules specify how a workload, such as a Domino address space, is assigned to a service class. It can also be assigned to a report class if required. In this example, the Domino server all have a service class of NOTESSRV but have different report classes.

The Qualifier Type shows on what input the classes are selected. In our example (Table 10) they are JOBNAME (TN) and USERID (UI).

Table 10 (Page 1 of 2). Classification Rules				
Subsystem	**Qualifier Type**	**Qualifier Name**	**Service Class**	**Report Class**
JES		Default	BATCHLOW	
OMVS		Default	OMVSUSER	
OMVS	TN	ADSM*	OMVSLOW	ADSM
OMVS	UI	OMVS	OMVSTASK	
OMVS	UI	SERVER	NOTESSRV	
OMVS	UI	ML99NP	NOTESSRV	ML99NP
OMVS	UI	ML98NP	NOTESSRV	ML98NP
OMVS	UI	ML97NP	NOTESSRV	ML97NP
OMVS	UI	ML96NP	NOTESSRV	ML96NP
OMVS	UI	DB99NP	NOTESSRV	DB99NP
OMVS	UI	DBL90NP	NOTESSRV	DBL90NP
STC		Default		
STC	TN	MVSNFS*		MVSNFS

Table 10 (Page 2 of 2). Classification Rules				
Subsystem	**Qualifier Type**	**Qualifier Name**	**Service Class**	**Report Class**
STC	TN	OMVS		OMVS
STC	TN	SYSBMAS		SYSBMAS
STC	TN	TCPML		TCPML
STC	TN	RMFGAT		RMFIII
STC	TN	DFHSM		DFHSM
STC	TN	JES2		JES2
STC	TN	ADSM*	STCLOW	ADSM
TSO		Default	TSOPRIME	

Note: Blank in the Service Class column means “started task default” (also blank) to get into SYSTEM and SYSSTC classes.

A.4 Service Definition Service Class Goals

Having put workloads in the defined service classes, you now need to define the performance goals required for these service classes. These goals may be velocity (the rate at which you expect work to be processed) or response times (a percentage completed within a specified time). You also assign a number to specify how important it is that the workload achieves its goal (with 1 being the highest in importance). Table 11 shows our example.

Table 11. Service Definition Goals					
Service Class	**Workload**	**Period**	**Duration**	**Importance**	**Goal**
BATCHLOW	PRIMEBAT	1			Discretionary
NOTESSRV	PRIMOMVS	1		1	Velocity 60
OMVSLOW	PRIMOMVS	1			Discretionary
OMVSTASK	PRIMOMVS	1		2	Velocity 30
OMVSUSER	PRIMOMVS	1	400	2	80% 00:00:05.000
		2	2000	2	Velocity 30
		3			Discretionary
STCLOW	PRIMESTC	1			Discretionary
TSOPRIME	PRIMETSO	1	200	1	90% 00:00:00.500
		2	400	1	90% 00:00:02.000
		3		1	Velocity 30

Appendix B. Domino Parameters That Affect Performance

This appendix lists the parameters in the `notes.ini` file that affect Domino performance. They are sorted by function. For detailed information see appendix A in the *Lotus Notes Administrator's Guide*.

Caution

Do NOT change the following parameters:

- Server_Max_Concurrent_Trans
- Server_MaxSessions
- Server_MaxInitialThreads
- Server_Secondary_Threads

Leave the values to default. Otherwise you may cause problems with the Domino server on S/390.

B.1 System Parameters

These parameters affect the Domino server overall:

- Fixup_Tasks
- NSF_Buffer_Pool_Size
- Server_MaxSessions (do not change the default)
- ServerTasks
- XPC_Pool_Size

B.2 Monitoring Parameters

These parameters affect monitoring:

- Console_Loglevel
- Log
- Log_MailRouting
- Log_Replication
- Log_Sessions
- Log_Tasks
- Log_ViewEvents
- Log_MailCompactDisabled
- RTL_Logging

B.3 Administration Automation Parameters

These parameters affect automation of administration activities:

- AdminPInterval
- AMgr_DocUpdateAgentMinInterval
- Amgr_DocUpdateEventDelay
- Amgr_NewMailAgentMinInterval
- Amgr_NewMailEventDeley

B.4 Replication Parameters

For the parameters that affect replication see 5.4.1, “Replication between Servers” on page 48.

B.5 Indexing Parameters

These parameters affect indexing:

- FT_Max_Instances
- FT_Intl_Setting
- Update_No_Fulltext
- Updaters
- Default_Index_Lifetime_Days
- FTV_Max_Fields
- FTV_Fields_database
- Update_Suppression_Time
- Update_Suppression_Limit

B.6 Mailing Parameters

These parameters affect mailing:

- MailCompactDisabled
- MailDisablePriority
- MailDynamicCostReset
- MailEncryptIncoming
- MailLowPriorityTime
- MailMaxThreads
- MailTimeout
- MailTimeoutMinutes
- SecureMail
- SharedMail

B.7 Domino Advanced Services Parameters

For the parameters that affect the Domino Advanced Services functions of partitioning, clustering and billing see:

- 5.7.3, “Parameters That Affect Partitioning” on page 54
- 5.8.2, “Parameters That Affect Clustering” on page 55
- 5.9.1, “Parameters That Affect Billing” on page 57

Appendix C. Sample notes.ini File

This is an example of a `notes.ini` file from a production system. Some names have been changed.

```
[Notes]
KitType=2

; Local machine configuration
ExistingServerName=CN=D01HUBM1/OU=01/OU=H/O=IBM
Domain=DOM01
Directory=/d01ml001
IconPath=/d01ml001
FileDlgDirectory=/d01ml001
MailSystem=0
MailServer=CN=D01ML001/OU=01/OU=M/O=IBM
CertificateExpChecked=/d01ml001/d01ml001.id 11/01/98
ServerKeyFileName=d01ml001.id
KeyFilename=d01ml001.id
Admin=Jane Smith,David S Jones

; OS/390 configuration variables
OS390SPAWN=1
SERVER_CONSOLE_CODEPAGE=1047
SERVER_MAXINITIALTHREADS=3    (do not change this parameter)
SERVER_SECONDARY_THREADS=100  (do not change this parameter)

; Debug variables
MAILLOGTOEVENTSONLY=0
LOG_MAILROUTING=40
SERVER_CLOCK=0
LOG_SESSIONS=1

; Trace, performance, and logging
TRACE_SHOW_TIME=1
SERVER_SHOW_PERFORMANCE=1
LOG_AUTHENTICATION=0
LOG_REPLICATION=2
LOG_SESSIONS=1
PhoneLog=2
Log=log.nsf, 1, 0, 7, 40000
Passthru_LogLevel=0
Console_LogLevel=2

; Security
SECURE_LOG=0
SECURE_LOG_ENCRYPTION=0

; Started tasks
ServerTasks=Replica,Router,Update,Stats,AMgr,Adminp,Sched,CalConn,Report,Event
ServerTasksAt1=Updall
ServerTasksAt2=
ServerTasksAt5=

; Network stuff
TCPIP0A=TCP, 0, 15, 0,,12201
Ports=TCPIP0A
TCPIP0A_TCPIPADDRESS=0,9.999.99.99:1352
SPX=SPX, 0, 15, 0
Serial1=XPC,1,15,0,
Serial2=XPC,2,15,0,
DisabledPorts=SPX,Serial1,Serial2

; Other stuff
$$HasLANPort=1
Preferences=2148011121
VIEWIMP1=Lotus 1-2-3 Worksheet,0,_IWKSV,,.WKS,.WK1,.WR1,.WRK,.WK3,.WK4,
VIEWIMP3=Structured Text,0,_ISTR,,.LTR,.CGN,.STR,
VIEWEXP1=Lotus 1-2-3 Worksheet,0,_XWKS,,.WKS,.WK1,.WR1,.WRK,
```

```
VIEWEXP3=Structured Text,0,_XSTR,,.LTR,.CGN,.STR,
EDITIMP1=ASCII Text,0,_ITEXT,,.TXT,.PRN,.C,.H,.RIP,
EDITIMP2=MicrosoftWord RTF,0,_IRTF,,.DOC,.RTF,
EDITIMP3=Lotus 1-2-3 Worksheet,0,_IWKSE,,.WKS,.WK1,.WR1,.WRK,.WK3,.WK4,
EDITIMP4=Lotus PIC,0,_IPIC,,.PIC,
EDITIMP5=CGM Image,0,_IFL,,.GMF,.CGM,
EDITIMP6=TIFF 5.0 Image,0,_ITIFF,,.TIF,
EDITIMP22=PCX Image,0,_IPCX,,.PCX,
EDITEXP1=ASCII Text,2,_XTEXT,,.TXT,.PRN,.C,.H,.RIP,
EDITEXP2=MicrosoftWord RTF,2,_XRTF,,.DOC,.RTF,
EDITEXP3=CGM Image,2,_XCGM,,.CGM,.GMF,
EDITEXP4=TIFF 5.0 Image,2,_XTIFF,,.TIF,
EDITEXP5=Ami Pro,2,_XW4W,W4W33T/V0,.SAM,
EDITEXP14=WordPerfect 5.1,2,_XW4W,W4W07T/V1,.DOC,
DDETimeout=10
NAMEDSTYLE0=020042617369630000000000000000000000000000000000000000000000000000000000000000
00000000001010100000A0000000000000100A0050000A0050000000000
00000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
0000000000000000
NAMEDSTYLE1=020042756C6C6574000000000000000000000000000000000000000000000000000000000000000
00000000001010100000A0000000000000000807000008070000000000000000000000000000000000
00000000000000000000000000000000000000000000000000000000000040000000000000000
NAMEDSTYLE2=02004865616
46C696E650000000000000000000000000000000000000000000000000000000000000010101010B0C00
00000000000100A0050000A00500000000000000000000000000000000000000000000000000000000000000
00000000000000000000000000000000000000000000000000000000
$$$OpenSpecial=NotesNIC
$$$NotesNIC=CN=Home/OU=Notes/O=NET, welcome.nsf, Notes NIC Welcome, Notes
Network Information Center on the Internet
TemplateSetup=2
Setup=47
ServerSetup=7
ZONE_SET=1
Timezone=5
DST=1
SERVER_SESSION_TIMEOUT=30
MAILTIMEOUT=3
ADMINPINTERVAL=15
ADMINPMODIFYPERSONDOCUMENTSAT=0
LOG_AGENTMANAGER=1
SHOW_TASK_DETAIL=1
LOG_VIEW_EVENTS=1
NSF_BUFFER_POOL_SIZE=200000000
MAILMAXTHREADS=10
LOG_MAILROUTING=40
Fixup_Tasks=75
SERVER_NAME_LOOKUP_NOUPDATE=1
DateOrder=DMY
<parameter=value>
Update_No_Fulltext=0
NIF_POOLSIZE=25000000
NIF_POOL_SIZE=25000000


NPN=1
VIEWIMP4=Tabular Text,0,_ITAB,,.PRN,.RPT,.TXT,.TAB,
VIEWEXP4=Tabular Text,1,_XTAB,,.LTR,.RPT,.CGN,.TAB,
EDITIMP7=BMP Image,0,_IBMP,,.BMP,
EDITIMP8=Ami Pro,0,_IW4W,W4W33F/V0,.SAM,
EDITIMP17=WordPerfect 5.x,0,_IW4W,W4W07F/V1,.DOC,
EDITIMP28=Binary with Text,0,_ISTRNGS,,.*,
EDITIMP29=WordPerfect 6.0/6.1,0,_IW4W,W4W48F/V0,.WPD,.WPT,.DOC,
EDITIMP30=Excel 4.0/5.0,0,_IW4W,W4W21F/V4C,.XLS,
EDITIMP31=Word for Windows 6.0,0,_IW4W,W4W49F/V0,.DOC,
EDITIMP32=GIF Image,0,_IGIF,,.GIF,
EDITIMP33=JPEG Image,0,_IJPEG,,.JPG,
EDITEXP21=WordPerfect 6.0,2,_XW4W,W4W48T/V0,.DOC,
EDITEXP22=WordPerfect 6.1,2,_XW4W,W4W48T/V1,.WPD,.WPT,.DOC,
EDITEXP23=Word for Windows 6.0,2,_XW4W,W4W49T/V0,.DOC,
BillingClass=Agent,Database,Document,Mail,Replication,Session
BillingAddinOutput=1
MailClusterFailover=1
```

```
KillProcess=1
server_MaxUsers=0
NPN=1
Debug_enable_core=1
```

Appendix D. Special Notices

This redbook will help OS/390 system programmers and Lotus Notes administrators monitor and tune Domino for S/390. See the product manuals for documentation and programming interfaces for *Lotus Domino for S/390*.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ACF/VTAM	AIX
AS/400	BookManager
CICS	DB2
DFSMS	DFSMS/MVS
DFSMSdss	ES/3090
ES/9000	ESA/390
ESCON	IBM
IMS	MVS
OfficeVision	OpenEdition
OS/390	PR/SM
Processor Resource/Systems Manager	PROFS
RACF	RAMAC
Resource Measurement Facility	RMF
RS/6000	S/390
SNAP/SHOT	System/390
VTAM	

The following terms are trademarks of other companies:

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

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

Other company, product, and service names may be trademarks or service marks of others.

Appendix E. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

E.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 91.

- *Lotus Domino for S/390 Release 4.5: Installation, Customization and Administration*, SG24-2083
- *Deploying Domino in a S/390 Environment*, SG24-2182
- *Porting C Applications to Lotus Domino on S/390*, SG24-2092
- *OS/390 Release 4 Implementation*, SG24-2089
- *OS/390 Release 3 OpenEdition Installation and Customization*, SG24-2087 (available after first quarter 1998)
- *Lotus Notes Release 4.5: A Developer's Handbook*, SG24-4876
- *Using ADSM to Back Up Lotus Notes*, SG24-4534
- *Lotus Solutions for The Enterprise, Volume 1 Lotus Notes: An Enterprise Application Platform*, SG24-4837
- *Lotus Solutions for the Enterprise, Volume 3 Using the IBM CICS Gateway for Lotus Notes*, SG24-4512
- *Lotus Notes Release 4 In a Multiplatform Environment*, SG24-4649
- *Enterprise Calendaring with Lotus Notes: The Notes to OfficeVision Connection*, SG24-4811
- *Secrets to Running Lotus Notes: The Decisions No One Tells You How to Make*, SG24-4875
- *Accessing OS/390 OpenEdition MVS from the Internet*, SG24-4721
- *TCP/IP Tutorial and Technical Overview*, published by Prentice Hall, ISBN 0-13-460858-5 or GG24-3376

E.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

E.3 Lotus Documentation

These books are shipped with the Lotus Domino for S/390 server code CD-ROM:

- *Lotus Domino for S/390 Install Guide for Servers*, Lotus part number AB0999
- *Lotus Notes Administrator's Guide*, Lotus part number 12755
- *Working with Lotus Notes and the Internet*, Lotus part number 12763
- *Lotus Notes Deployment Guide*, Lotus part number 12757
- *Lotus Notes Network Configuration Guide*, Lotus part number 12756

E.3.1 Lotus Documentation Online

Lotus documentation is available online after the product has been installed. The database name is listed after the title of the documentation.

- *Domino 1.5 User's Guide*, domguide.nsf
- *Notes Help*, help4.nsf
- *Notes Help Lite*, helplt4.nsf
- *Release Notes: Domino/Notes 4.5*, readme.nsf
- *Install Guide for Servers*, doc/srvinst.nsf
- *Install Guide for S/390 Server*, doc/srvs390.nsf
- *Install Guide for Workstations*, doc/wksinst.nsf
- *Lotus Notes and the Internet*, doc/internet.nsf
- *Notes Administration Help*, doc/helpadmn.nsf
- *Notes Migration Guide*, doc/migrate.nsf

E.4 Other Publications

These publications are also relevant as further information sources:

- *OS/390 OpenEdition Planning*, SC28-1890
- *OS/390 OpenEdition User's Guide*, SC28-1891
- *OS/390 OpenEdition Command Reference*, SC28-1892
- *OS/390 OpenEdition Messages and Codes*, SC28-1908
- *OS/390 MVS System Messages, Vols 1-5*, GC28-1784 to GC28-1788
- *OS/390 MVS Planning: Workload Management*, GC28-1761
- *OS/390 MVS Initialization and Tuning Reference*, SC28-1752
- *OS/390 Information Roadmap*, GC28-1727
- *OS/390 Resource Measurement Facility User's Guide*, SC28-1949
- *OS/390 Resource Measurement Facility Report Analysis*, SC28-1950
- *OS/390 MVS System Management Facilities (SMF)*, GC28-1783
- *ES/9000 and ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123
- *IBM 3990/9390 Storage Control Planning and Installation*, GA32-0100.

- *Storage Management Library Version 1 Release 2*, SC26-3124
- *MVS TCP/IP User's Guide*, SC31-7136
- *IBM TCP/IP Performance Tuning Guide*, SC31-7188
- *TCP/IP for MVS: Planning and Migration Guide*, SC31-7189
- *IBM TCP/IP for MVS:Customization and Administration Guide* , SC31-7134
- *Planning for System/390 Open Systems Adapter Feature*, GC23-3870

E.4.1 Web Sites

See also the following Web sites:

IBM site for Lotus Domino for S/390	`http://www.s390.ibm.com/products/domino`
OS/390 UNIX System Services	`http://www.s390.ibm.com/oe/`
S/390 Partners in Development	`http://www.s390.ibm.com/products/s390da/`
IBM Home Page	`http://www.ibm.com/`
Lotus Home Page	`http://www.lotus.com/`

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at `http://www.redbooks.ibm.com/`.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type `GOPHER.WTSCPOK.ITSO.IBM.COM`
- **Tools disks**

  To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks, type the following command:

  ```
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  ```

  To register for information on workshops, residencies, and redbooks, type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
  ```

  For a list of product area specialists in the ITSO: type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Web Site on the World Wide Web**

  `http://w3.itso.ibm.com/redbooks/`

- **IBM Direct Publications Catalog on the World Wide Web**

  `http://www.elink.ibmlink.ibm.com/pbl/pbl`

  IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

  Index # 4421 Abstracts of new redbooks
  Index # 4422 IBM redbooks
  Index # 4420 Redbooks for last six months

- **Direct Services** - send note to `softwareshop@vnet.ibm.com`
- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com/
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

  With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to `announce@webster.ibmlink.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank).

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Glossary

address space. In ESA/390, the range of virtual storage addresses that provide each user with a unique address space and which maintains the distinction between programs and data within that space.

APAR. Authorized program analysis report. A report of a problem caused by a suspected defect in a current unaltered release of a program.

central storage. Storage that is an integral part of the processor unit. Central storage includes both main storage and the hardware system area.

channel path. Is the physical medium by which a channel subsystem exchanges data with an I/O device in ESA/390 mode.

CICS. Customer Information Control System. An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

DASD. Direct access storage device. A device in which access time is effectively independent of the location of the data.

ESCON. Enterprise Systems Connection.

expanded storage. Optional integrated high-speed storage that transfers 4K-byte pages to and from central storage.

FTP. File transfer program

GB. Gigabyte. 1,073,741,824 bytes, or 1,024 megabytes.

HFS. Hierarchical File System (DFSMS/MVS, POSIX 1003.1(a) compliant file system)

HTML. Hypertext Markup Language

HTTP. Hypertext Transfer Protocol

IBM. International Business Machines Corporation

I/O. Input/output

ITSO. International Technical Support Organization

JES (Job Entry Subsystem). A system facility for spooling, job queuing, and managing I/O.

KB. Kilobyte. 1,024 bytes.

logical partition. In LPAR mode, a subset of the processor unit resources that is defined to support the operation of a system control program (SCP).

logically partitioned mode (LPAR). A mode that allows the operator to allocate hardware resources of the processor unit among several logical partitions.

LPAR. See *logically partitioned mode.*

MB. Megabyte. 1,048,576 bytes or 1,024 kilobytes.

Processor Resource/Systems Manager (PR/SM). A function that allows the processor unit to operate several system control programs (SCPs) simultaneously in LPAR mode. It provides for logical partitioning of the real machine and support of multiple preferred guests.

PTF. Program Temporary Fix. A temporary solution or by-pass of a problem diagnosed by IBM as resulting from an error in a current unaltered release of the program.

S/390. System/390. Any ES/9000 system including its associated I/O devices and operating system(s).

TCP/IP. Transmission Control Protocol/Internet Protocol. A public domain networking protocol with standards maintained by US Department of Defense to allow unlike vendor systems to communicate.

TSO. Time Sharing Option

UNIX. An operating system developed at Bell Laboratories (trademark of UNIX System Laboratories, licensed exclusively by X/Open Company, Ltd.)

URL. Uniform Resource Locator

virtual storage. The storage space that can be regarded as addressable main storage by the user of a computer system in which virtual addresses are mapped into real addresses. The size of virtual storage is limited by the addressing scheme of the computer system and by the amount of auxiliary storage available, not by the actual number of main storage locations.

Index

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

ITSO Redbook Evaluation

Lotus Domino for S/390 Performance Tuning and Capacity Planning
SG24-5149-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

_____

_____

_____

_____

_____

SG24-5149-00
Printed in the U.S.A.

Lotus Domino for S/390 Performance Tuning and Capacity Planning

SG24-5149-00