

IBM RS/6000 Enterprise Server S80 and Domino R5 Server Loading Efforts

An IBM Global Services Case Study

January 2002



An International Business Machines Corporation Technical Paper

Introduction	1
<i>Pilot Overview</i>	1
Executive Summary	1
Current Environment in SDC West-Boulder	3
<i>Server Configuration</i>	4
<i>Client Configuration</i>	4
The Server Loading Results	5
<i>Explanation of Performance Statistics</i>	7
<i>Memory</i>	7
<i>CPU</i>	7
<i>Disk Busy</i>	7
Domino Issues Encountered	8
<i>NSF_folderpool</i>	8
<i>Dataseg</i>	8
<i>dataseg Command Parameters</i>	9
<i>TCP Queue Length</i>	10
<i>Domino and AIX Process Memory Usage</i>	10
Client Migrations	11
<i>GNA Large Scale Mail Migration Process</i>	11
Appendix A. AIX and Notes Settings	15
<i>AIX Settings</i>	15
<i>Domino Settings</i>	15
Appendix B. Hints and Tips	17
<i>AIX</i>	17
Appendix C. Reference Information	19
<i>AIX</i>	19
<i>Domino</i>	19
<i>Miscellaneous</i>	19
Acknowledgments	21
Special Notices	23

Introduction

This case study details the second phase of Domino server loading efforts of IBM production users on IBM RS/6000 and pSeries servers. Phase one was documented in an IBM Redpaper entitled: "IBM RS/6000 Enterprise Server M80 and Domino R5 Server Loading Efforts" published in July 2001. The Redpaper, IBM document REDP0201, is available from the following URL:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=REDP0201>

As with the M80 effort, this paper provides a detailed report on the technical activities surrounding the production migration of Domino servers. This document details Domino server loading on a RS/6000 Enterprise Server model S80 server connected to a Storage Area Network (SAN).

This document is intended to be read by system engineers, architects, technical and marketing support personnel, and sales representatives. It presents the logical and physical design of the solution, technical details in terms of configuration changes performed, server statistics and challenges encountered during the server loading of the S80 hardware.

Pilot Overview

The S80 phase of the Domino loading pilot was undertaken when the RS/6000 M80 server was needed to support a commercial account in the IBM Boulder Service Delivery Center (SDC). The M80 was replaced by an 18 CPU, 450 MHz, S80 server with 32GB of RAM. The primary objective of the S80 pilot, similar to that of the M80 pilot, was to architect a solution which would provide the most efficient loading of the much larger S80 server (using number of clients and transactions per server as key metrics) while also maximizing the number of production clients on an individual Domino partition.

In addition, the pilot team collected server statistics in support of marketing efforts to better position the S80 server platform against numbers published by competitive hardware vendors. The server statistics include data from different performance tools that are presented in detail later in this document, such as Server Resource Monitoring Tool (SRM) and Activity Trends.

Executive Summary

Phase two of the pSeries server loading efforts started late in August of 2001 as the S80 server replaced the M80 server. An additional goal of the S80 server loading effort was to stress the number of clients on one Domino server while maximizing the overall number of clients on the S80 itself. The IBM Boulder Service Delivery Center continued to be the host for this effort. The hardware configuration consisted of one RS/6000 S80 running AIX 4.3.3 with eighteen RS64-III 450 MHz processors attached to a SAN fabric.

GNA Domino Loading Pilot RS/6000 S80 Systems Configuration

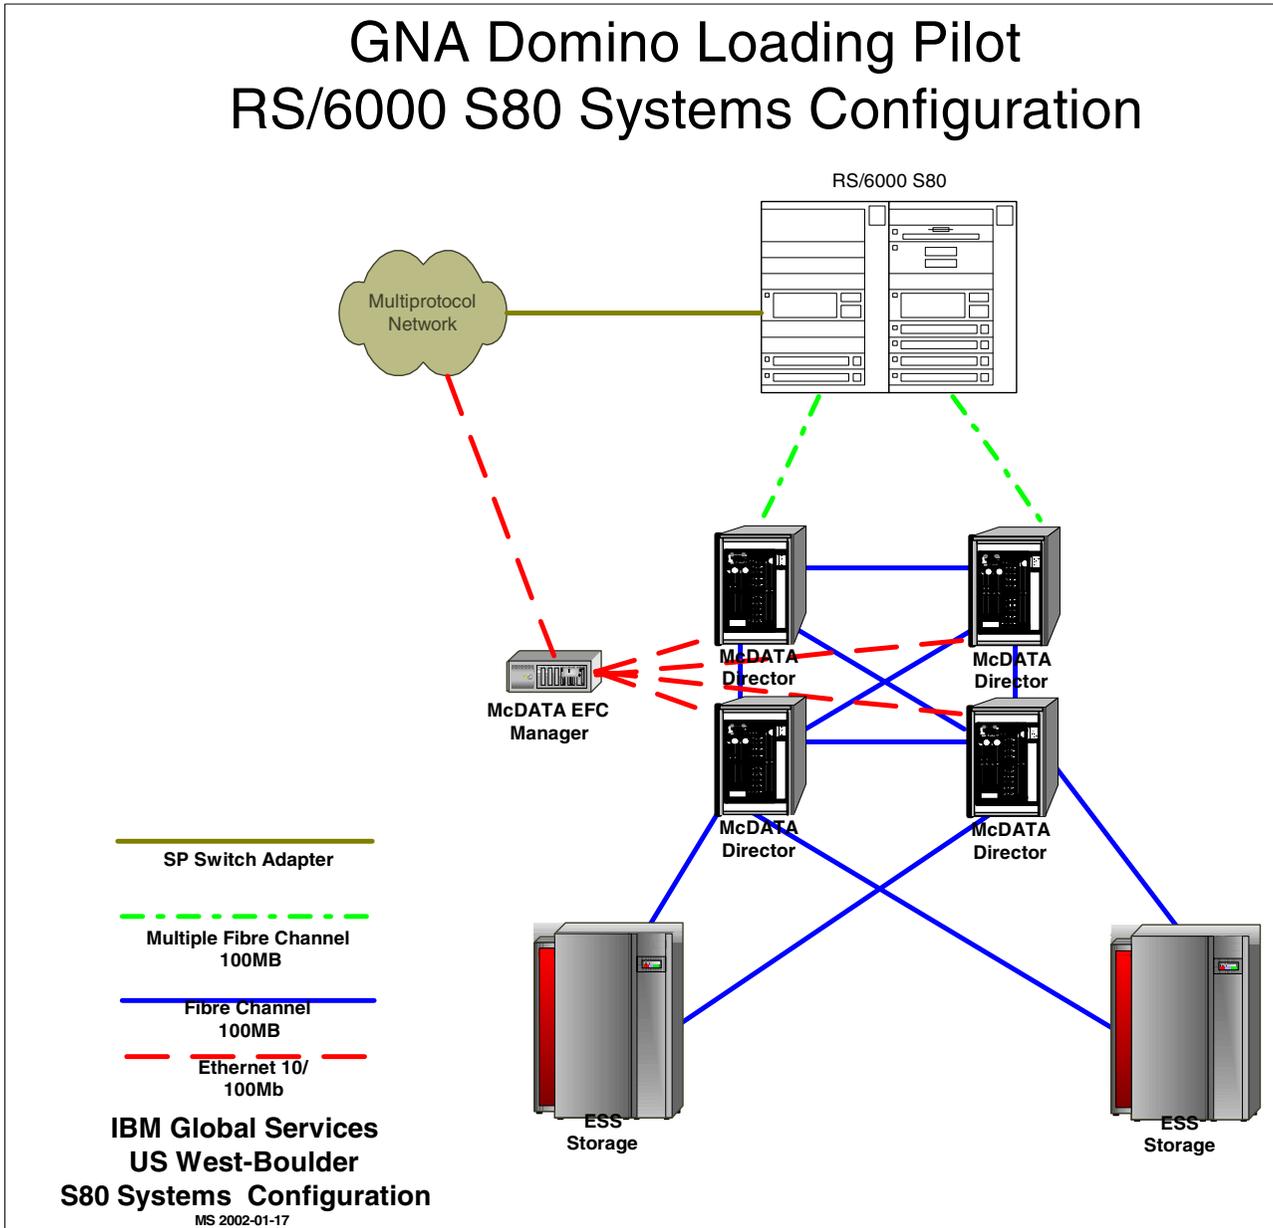


Figure 1.

Migrated from the M80 were two Domino R5.06a mail servers (D03NM800 and D03NM801).

Initial server loading had roughly 2,590 clients on each partition for a total of 5,180 clients on the S80.

The S80 was connected to the Storage Area Network (SAN) with two 32-bit Fibre Channel adapters. The SAN was used for storage of Notes client mail databases as well as for the Domino server transaction logs. The Domino database backup was performed by replicating databases to a dedicated Domino Backup server which was then backed up via the Tivoli Storage Manager (TSM) product to tape drives directly attached to the Domino Backup server.

To load additional Notes clients on to the S80, the Boulder Notes team migrated clients from production SP Silver nodes to each of the production mail servers on the S80. Initial migrations were directed to the D03NM800 server in order to stress Domino while maximizing the number of clients on a single Domino partition. When the number of clients on this server reached 5,800, clients began to encounter the error:

```
Insufficient memory NSF Folder Pool is full
```

To address this and avoid removing clients, a hot fix SPR #SHOR54YPXM provided by Iris, was installed and the server was also upgraded to R5.09. The hot fix increased the size of the NSF folder cache, resolving the problem and allowing for continued client loading on D03NM800. This problem is explained in more detail later in the document.

Loading of clients to the S80 continued until the year end freeze on systems configuration changes took effect in mid-December. The total number of clients at year end on the server reached 11,080 on two Domino partitions (D03NM800: 6,040 and D03NM801: 5,040). At the end of migrations with the 11,080 clients, the S80 prime shift CPU utilization averaged 65% busy with 10.5 to 11.5 million prime shift transactions. This left 35% of the S80's 18 processor system resource available for peak workloads. No servers were clustered with the S80.

During the first three weeks of January the CPU utilization increased to 69% as clients returned from the holidays. A final migration of 1,013 clients to the S80 brought the total number of clients to just over 12,000 (D03NM800: 6,000 and D03NM801: 6,026). With this additional load the prime shift CPU utilization averaged 73% busy with 12 to 13 million prime shift transactions. Note, with the large number of clients the exact user population fluctuated on the server due to attrition, new clients and relocated clients.

As a result of this effort the team was able to consolidate clients from 12 RS/6000 SP Silver nodes and migrate 12,093 Notes clients to the S80. Note the number of servers consolidated can vary based on the number of clients and transactions on the existing servers. Also of note: Hardware Multithreading (HMT), supported on the M80, was not available on the S80 server.

In phase three of the pilot efforts in Boulder, the team plans to migrate Notes Clients from existing SP Silver nodes over to an 8CPU, 800MHz pSeries 690 (aka Regatta H) server. The Regatta H will have one logical partition (LPAR) to be used for mail until additional hardware upgrades provide capacity for multiple LPAR's. Preliminary plans under the single LPAR for the Regatta have the team migrating IBM Notes Clients to Domino R5.09 and Domino 6 beta servers. Once the hardware upgrades (additional processors and memory) are in place, the server will be used for other GNA infrastructure activities (i.e.: database servers, backup servers, etc.). Additional server statistics will be provided for this effort as well.

Current Environment in SDC West-Boulder

Here we discuss the current environment in SDC West-Boulder.

Server Configuration

The environment for our pilot consisted of a RS/6000 S80 (7017-S80) server attached to a SAN. As stated earlier the S80 was not clustered with any SP Silver nodes during the migration efforts. The initial loading consisted of the 5,180 M80 clients and subsequent migrations of 6,846 additional clients for a total of just over 12,000 clients. Initially the S80 was configured with Domino 5.06 on AIX 4.3.3. During the migration efforts to eliminate the "NSF Folder Pool Full" error, the server was upgraded to R5.09 and a hot fix was applied. Table 1 shows a comparison of the S80, M80 and SP Silver node hardware platforms.

Table 1.

	S80	M80	SP Silver node
CPU Speed	450 MHz	500 MHz	332 MHz
Number of CPU's	18	8	4
Memory	32 GB	16 GB	3 GB
DASD adapter	Two 32-bit Fibre Channel	Two 32-bit Fibre Channel	Two Advanced SerialRAID SSA
Network adapter	1 SP switch 1 Gigabit Ethernet Adapter	1 Gigabit Ethernet Adapter	1 SP switch
I/O Drawers	2	N/A	N/A
Domino partitions	2	2	1
Total registered users per machine	12,026	5,180	1,100-1,300 clients on four originally clustered silvers. Fifth server had 800 clients
Concurrent users per machine at peak time	Prime Shift (7am-5pm) Average: 3,600 Prime Shift Peak: 4,200	Prime Shift Average: 1,500 Prime Shift Peak: 1,800	Prime Shift Average: 300 Prime Shift Peak: 400

Client Configuration

At the IBM Boulder site, typical IBM Notes users frequently access their mail files during the workday. Additionally, some clients directly open the server copy of their mail file, while others replicate their mail to their workstation and continue to replicate changes throughout the day.

Furthermore, there is a mix of R5 and R4.5x clients on the two Domino mail servers residing on the S80 server. These IBM clients all use a highly customized IBM mail template, which was derived from the R5 or R4.5x Lotus mail template included with the product. Approximately 45% of the clients on D03NM800 and 43% of the clients on D03NM801 are using the R5 based mail template. The average mail file size for a Boulder IBM client on either of these two mail servers is 78 MB. Along with the use of mail, the IBM clients are also heavy users of Calendar and Scheduling. So, in addition to the resident clients on the server (those defined to the server), the server also plays host to nonresident clients who may open sessions on these servers to either manage or view the calendars of resident clients.

Compared to the industry standard test measurement of NotesBench clients, a production IBM client is two to three times more active in terms of transaction count and mail file size. Additionally, a typical IBM client has larger message sizes (an average mail message size of 70 KB) and sends and receives mail at greater frequencies.

The Server Loading Results

The S80 server loading effort provided significant results. The Boulder SDC team was able to consolidate 12 SP Silver nodes on to one S80 server. This resulted in 12,000 clients on one physical server spread across two Domino mail partitions. During this time daily prime shift CPU utilization averaged 73% .

The data in Figure 2 on page 10 and Figure 3 on page 10 represent the CPU utilization and transaction load on the S80 during the pilot. The CPU utilization statistics were gathered using Server Resource Management (SRM). SRM is an IGS Web application which reports historical trends of key server resources, such as CPU, Memory, and Disk. The transaction rates were collected using the Activity Trends tool (also known as Tivoli's advanced server health monitoring in Domino 6 beta). The primary source of the Activity Trends data collected is the standard Lotus Domino Log database (log.nsf). Activity Trends records and reports statistics which portray the activity of clients against the databases on the Lotus Domino server where this database resides.

The data in Figure 2 on page 10 represents the growth on the S80 on a weekly basis. The CPU utilization climbed and averaged 65% as the 11,080 client mark was reached in mid-December. During the first three weeks of January the CPU utilization increased to 69% as clients returned from the holidays. After the final migration of 1,013 clients to reach the 12,000 client mark, the prime shift CPU utilization averaged 73% with 12 to 13 million prime shift transactions. The weekly transaction rates are a weighted average of all data observations with more recent data having more influence.

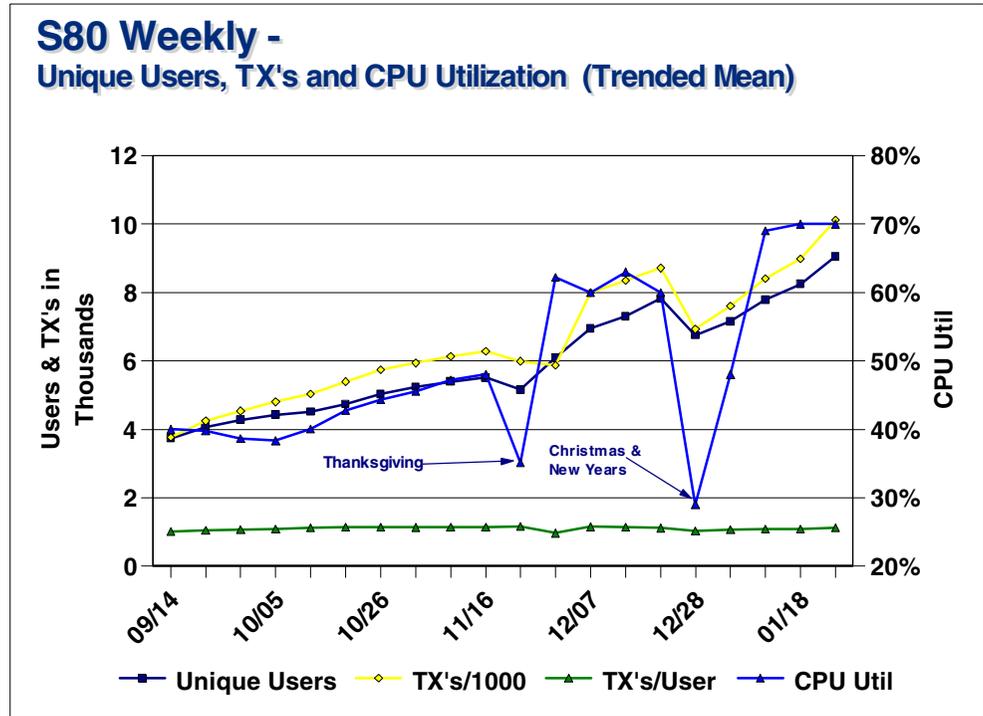


Figure 2.

Daily transactions rates are represented in Figure 3. All data in Figure 2 and Figure 3 reflect prime shift server statistics (7 AM to 5 PM).

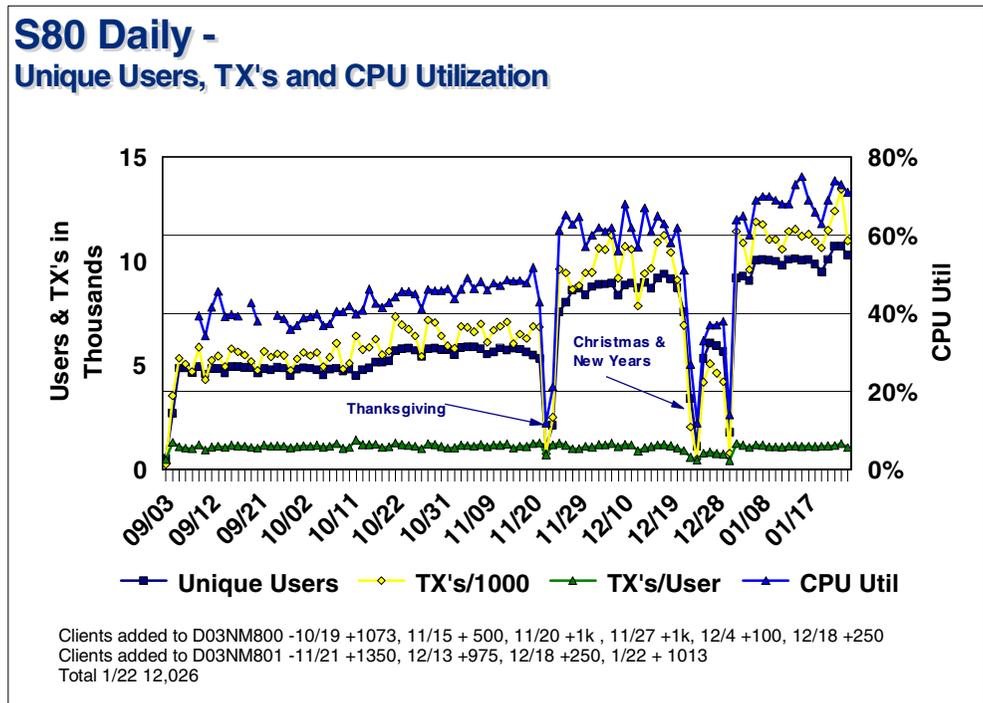


Figure 3.

Figure 4 represents the fifteen minute CPU utilization interval for a one week period on the S80 with the client count at 11,080.

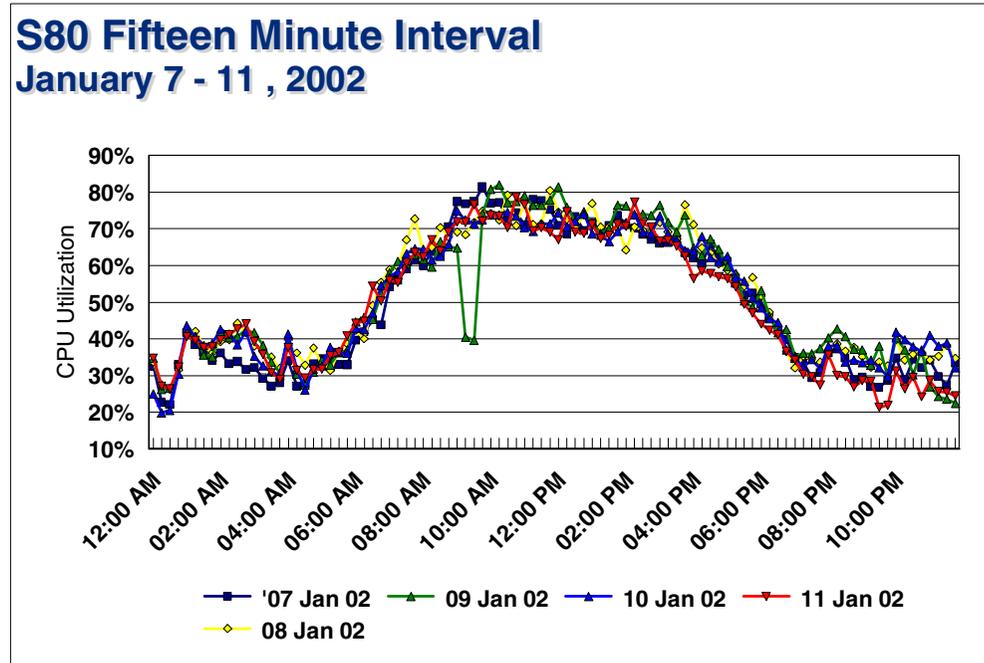


Figure 4.

Explanation of Performance Statistics

Memory

Each of the Domino partitions were allocated 4 GB of memory. There was no paging taking place on the server and no other memory related bottlenecks or issues.

CPU

No form of processor binding was implemented, therefore, all 18 processors were available to any Domino tasks. Utilization was monitored on a per processor basis, but the results were averaged across all processors. The most useful high level processor statistics on the server were the system, user, wait, and idle utilization percentages which are measured by using `vmstat`, `sar`, or `topas`.

Disk Busy

Since our Domino directories are stored internally on SAN servers, our disk statistics such as %busy or cache-hit ratios must be examined with care. The S80 sees only the AIX perspective of disk activity, which is reported as a higher number than that reported by the ESS Expert tool. Each Fibre channel path from the S80 to a LUN on a SAN server appears to AIX as an hdisk. In this environment, there are two paths to each LUN and multiple LUNs for each Domino data directory, so the I/O was spread across two paths (hdisks) per Domino data directory.

Domino Issues Encountered

Here we discuss some of the Domino issues we encountered.

NSF_folderpool

Following the latest migration on the first partition (D03NM800), which raised the registered user count from 5000 to ~6000, users began to encounter the error:

```
Insufficient memory NSF Folder Pool is full
```

when attempting to navigate between folders on their server copy mail database. The cause of this error was traced to an area of shared memory named `BLK_NSF_FOLDERPOOL`, which is the memory pool for caching open folder data. In Domino R5.0-R5.09 this pool is hard coded to a maximum size of 20MB. Due to the high level of concurrent users active (peaks up to 2100+) on this partition, and the large number of folders in their mail files, this maximum value was being reached. Resulting in the NSF Folder Pool full error whenever a user attempted to open a previously un-cached folder on the server copy of their mail file.

Iris supplied a hot fix to R5.09 (fix will be integrated into R5.0.10 and Domino 6), which increases the size of this pool (currently to 50MB), thereby addressing the error occurrence. It is important to note that this error can not be directly correlated to active user count, but is more a function of the number of open folders, which of course has a tendency to increase with user count.

The size of this Folder Pool can be determined by examining a NSD file, or by a memory dump (`server -m`, and then run `memstats` on output file), with a search on the string `FOLDERPOOL`. Since implementing the hot fix the amount of memory utilized by this pool (on partition D03NM800) has grown as high as 30MB.

Dataseg

Domino 5.09 for AIX introduces a new memory management model. AIX 4.3.3 allows up to 11 segments of shared memory to be used by an application, such as Domino. Most instances of the Domino server do not need all 11 segments of shared memory (depends on load) so one or more of these segments can afford to be "taken" from shared memory and "given" to process (i.e. task, such as HTTP, router, etc.) memory.

This is done via the `dataseg` command, which ships as part of R5.09. Please see section below for command parameters. Domino R5.09 ships by default as if this command had been run with an argument of 2, meaning that 2 segments should not be used for shared memory, but instead given to process memory. This leaves 9 segments for shared memory.

This is important to understand for two reasons:

First, in the case of heavily loaded servers (such as D03NM800 in our environment with 6,000 registered users) that need all 11 segments of shared memory, the default R5.09 value for shared memory segments must be changed. This is done by issuing the command:

```
dataseg -f 0 * <path>
```

which allows the Domino binaries to use 11 shared memory segments, as was the case in pre-R5.09 on AIX 4.3.3. This depends on the server, so shared memory segment utilization would need to be monitored (with `ipcs` for example), in order to understand if 9 or 11 is the better setting for a specific environment.

Second, if there are any add-in tasks such as TrendMicro (`tmmscan`, `repscan`, etc.) running on a server with default R5.09 install, they need to be made aware of this new 9 shared memory segment default model. This is done by executing the `dataseg` command:

```
dataseg -f 2 * <path>
```

on the directory where these add-in binary files reside. If this is not done these tasks will expect to be able to see 11 shared memory segments and will generate memory mapping errors when unable to do so.

dataseg Command Parameters

What `dataseg` does is flip a bit on the `bmaxdata` flag post compile time. To use `dataseg`, root authority is required. The Domino server should be down as well. A backup of the desired binaries should be taken should there be a problem with the patching process. The generic command is:

```
dataseg -s *
```

The default for `dataseg` is to add 2 segments to all Domino R5.09 binaries.

Examples:

- Running `dataseg` with only a file argument will tell if the binary has been modified.

```
dataseg http
file = "http"
data/stack seg# = 0 # of shared segs avail. = 11 = 2816 Meg
```

- This is the default change `dataseg` would make with the `-s` flag.

```
dataseg -s http
file = "http"
Old: data seg# = 0 # of shared segs avail. = 11 = 2816 Meg
New: data seg# = 2 # of shared segs avail. = 9 = 2304 Meg
```

- This is an example of the `-f` flag which accepts number between 0-7.

```
dataseg -f 0 http
file = "http"
Old: data seg# = 2 # of shared segs avail. = 9 = 2304 Meg
New: data seg# = 0 # of shared segs avail. = 11 = 2816 Meg
```

Additionally the following command can be used to verify the number of segments of an executable.

```
dump -ov http | tail -2
```

which returns the following:

```
maxSTACK      maxDATA      SNbss        magic        modtype
0x00000000    0x20000000    0x0003       0x010b       1L
```

where `maxData` = `0x20000000` or the hex value associated with 2 additional data segment.

Technote #189972 is available with more detail, see the Miscellaneous reference section in Appendix C, "Reference Information" on page 23.

TCP Queue Length

During the S80 phase of consolidation an unusually high number of "Server not responding" errors were encountered, which were alleviated by addressing the TCP (backlog) queue size. The following is a short explanation of the TCP/IP queue (backlog) length issue dealt with on this S80 server. The main point is how AIX deals differently than other Unix variants with queue length in relation to Domino. This can cause seemingly unrelated symptoms (which are unpredictable and erratic) and therefore can be difficult to diagnose, especially on systems under heavy load.

This is typified by the generic "Server not responding" error during both server-to-server communication (seen in the logs) and also when clients attempt to access the Domino server. While there could be many causes for this error (network latency, server busy due to load, etc.), one area pointed toward in this environment was the size of the TCP queue.

After a TCP connection (destined for an application) has been received by AIX, it still has to be accepted by the application, such as Domino. Once the TCP three-way handshake is complete, and the TCP layer of AIX has accepted the connection, it is then put on the (backlog) queue to be accepted by Domino. The application specifies the size of this (backlog) queue, which can be between zero and the value of `somaxconn` (a `no -a` parameter) set within AIX.

Domino R5.09 sets the backlog value to 5. When an application tells AIX to set the maximum backlog value to 5, it really corresponds to a maximum of this setting (5) or `somaxconn` (1024 for this environment), whichever is less. The Boulder AIX team verified with the AIX utility `crash`, that the queue was not above the value of 5 when this "Server not responding" error was being generated.

The `notes.ini` variable (`TCP_QLENGTH=`) allows the backlog queue size to be specified to AIX by Domino. This setting was adjusted (in increments of 10, from 10 to 30) resulting in a drastic reduction (overall from ~1000 to ~200 instances per day) in "Server not responding" errors on the servers the setting was applied to. This would seem to prove that adjusting this setting was beneficial in this environment. At the same time, however, there are still instances of "Server not responding" (which is a fairly generic error message), even at times when there is very low user count (10% active for example) with almost no load on the server. So there is no conclusive evidence supporting increased queue size as merely a symptom or the actual cause of this error.

It is also important to note that raising this value can effect resource utilization. Including CPU and memory usage, as well as system behavior, such as cluster logic and transaction timing. Therefore it should be raised slowly and with careful monitoring.

Domino and AIX Process Memory Usage

While running the R5.08 code on the S80, AIX crashes occurred because `ulimit` settings in the environment reverted to default settings. This was due to an unintended side effect while using a system utility to migrate files. The `ulimit` parameters were initially set to unlimited on both partitions. If set incorrectly, the

AIX `ulimit` parameter can cause erroneous "insufficient memory" or "out of memory" errors on the Domino server even when there is sufficient physical memory on the system.

For highly scaled configurations such as this, it was important *not* to use the default `ulimit` setting of 128MB. Domino R5.09 does a better job with memory allocations on AIX but the non-optimized `ulimit` setting can still produce erroneous memory errors. The default data value, in particular, limits the data segment size of a process to 128MB in the above example. On a large memory, multiprocessor Domino server such as this S80 it is possible that one or more Domino server tasks such as server, router, amgr or smtp could have data segment size of 128MB or larger. Appropriate S80 `ulimit` settings for Domino R5 are included in Appendix A, "AIX and Notes Settings" on page 19.

Client Migrations

Here we discuss the client migration processes.

GNA Large Scale Mail Migration Process

The mass mail migration process entails a week of activity, beginning with a seven day advance notification sent to the client community and culminating with address book changes that redirect client's mail to their new mail file. Migrations are performed via Notes replication, which means that the customer mail file resides on both the "source" (old) and "target" (new) servers for the entire week. Once address book changes are made to redirect mail to the new mail file, the old mail files are retained for two weeks so that any client who may have been out of the office can still access their old mail file to obtain directions to the new one. The entire mass migration process is managed via the use of a specially developed notes database which automates several of the migration procedures.

Activities performed during the week to prepare for the migration are done during normal business hours. When activities are performed on a server, such as creating new replica mail files on the "target" server or synchronizing mail file replicas between "source" and "target" server, they are done by an administrator familiar with the environment that monitors server performance to ensure there is no impact to clients or response time. This process ensures minimal disruption to clients while maintaining mail integrity.

Generally speaking, the migration process is as follows:

- Seven days prior to the migration:
 - Send notifications to all affected clients, informing them of the migration, their new server, the change record number, and the date the change will take effect.
 - Replica stubs of the client's current mail file are created on the new mail server.
 - Replica stubs of the client's mail file are created on the backup server that correspond with the new mail server, to ensure a viable backup is available on the day of the migration.
 - Replication is initiated between source and target server, as a means of moving mail data between the old and new mail files. Replication is staged to ensure that the impact on the server is minimal. This is accomplished by

limiting the number of replicator tasks running on the target server at any given time.

- Six days prior to the migration:
 - Replication of mail databases between old and new servers proceeds to keep replicas synchronized up to the moment when address book changes take effect. Replication is staged to ensure that the impact on the server is minimal.
 - Document counts are taken using the migration database mentioned above to ensure old and new mail files match and appropriate action is taken as needed which may include ACL modifications, replication setting changes or brand new replicas on the target.
- Day of migration:
 - Required changes are made to the Master copy of the mail domain's address Book prior to 06:00 MT. These changes do not replicate to the mail servers until later in the evening. The affected person documents are modified to indicate the client's new server.
 - Notifications are sent to all affected clients reminding them of the migration and providing detailed information on client changes required to ensure a successful migration with uninterrupted mail delivery. This consists of a series of automated buttons and steps to ensure ease of use and eliminate potential errors. These buttons assist the client in adding the new replica icon to the workspace and changing location documents to access the new server. Notifications are sent prior to 07:00 MT to ensure customers are provided adequate time to prepare as necessary.
 - Replication of mail databases between old and new servers proceeds to keep old and new mail files synchronized. Replication is initiated several times during the day to ensure that data are synchronized, in order to minimize data transfer once address book changes take effect.
 - At 20:00 MT, normal scheduled replication of the mail domain's address book from Atlanta to the Boulder mail hubs occurs. This replication takes 15-20 minutes and is monitored to ensure all necessary address book changes required to complete the migration are replicated.
 - At 21:00 MT, replication between the Boulder mail hubs and all servers involved in the migration (source and target servers) is initiated manually. This is done in advance of normal scheduled replication to the mail servers at 22:00 MT to ensure that all address books on all affected servers contain the same information, which eliminates the possibility of delivery failures occurring. Since normal scheduled replication at 22:00 is staged at different times for different servers, the failure of all servers involved in the migration to receive updates at the same time would lead to mail delivery failures; hence, forced replication prior to 22:00 eliminates this possibility. Once forced replication to the mail servers is complete, mail delivery to the new/migrated mail file begins.
 - A final synchronizing replication is performed to ensure that the new/migrated mail file contains all data from the old mail file that was delivered up to the time of address book replication is available in the new mail file.

- By 22:00 MT - the time advertised for client that mail delivery to the new/migrated mail file begins - delivery is in fact occurring and all data from the old mail file is available as well.
- Two weeks after migration:
 - Mail files remain on the old/source mail server for two weeks following the migration, to allow people out of the office on the day of the migration to access their old mail file which includes the instructions directing them to their new mail file. After two weeks, these files are deleted from the old server and if no other users exist on the old server, it can be decommissioned.
 - ACLs on the new mail files are updated to remove the old server entry.

Appendix A. AIX and Notes Settings

AIX Settings

AIX 4.3.3 maintenance level 9 (ML09)
IBM Subsystem Device Driver: ibmSdd_432.rte.1.2.0.4
Two physical paths to each SAN LUN
Domino transaction logs reside on a SAN server on a separate LUN
Logical volumes setup with "Range on physical volumes" = maximum <code>mklv -e'x'</code>
Maximum number of processes allowed per user = 1024 <code>chdev -l sys0 -a maxuproc='1024'</code>
Ulimits for Domino (from /etc/security/limits): fsize = -1 (unlimited) core = -1 (unlimited) data = 1572864 rss = 589824 nofiles = 4000 stack = 65536
Set vmtune -b 1116 -B 1116 -p10 -P 30 -b = number of file system bufstructs. -B = number of pbufs used by the LVM. -p = point below which file pages are protected from the repage algorithm. -P = point above which the page stealing algorithm steals only file pages.
JFS attributes of Domino data filesystems: frag size: 4096 nbpi: 131072

Domino Settings

Domino R5.09 with hot fix for increasing size of NSF Folder Pool
<code>dataseg -f 0 /opt/lotus/notes/latest/ibmpow/*</code> <code>dataseg -f 0 /opt/lotus/notesapi/*</code>
<code>Server_Max_Concurrent_Trans=2000</code>
<code>Modify Server_Pool_Tasks=200</code>
<code>mail_number_of_mailboxes=4</code>
<code>NSF_DBCache_MaxEntries=3500</code>
<code>SERVER_SESSION_TIMEOUT=30</code>
Anti-Virus scanner enabled for routed mail

Appendix B. Hints and Tips

AIX

Added to the profile running Domino:

<code>export AIXTHREAD_SCOPE=S</code>	# Tells AIX to use 1 kernel thread per user thread, if SCOPE=P, then default is 1:8
<code>export MALLOCMULTIHEAP=1</code>	# Use 32 separate AIX kernel malloc heaps (when MALLOCMULTIHEAP=1), instead of just 1 per system See /usr/lpp/bos/README

Monitor `hd_pendqblked` and `fsbufwaitcnt` with `/usr/samples/kernel/vmtune -a`.

If `hd_pendqblked` increases then increase `-B` or if `fsbufwaitcnt` increases then increase `-b`

If you are going to change either of these parameters they must be added to `/etc/inittab` before `/etc/rc`

If using Gigabit Ethernet on a SMP machine then turn on dedicated kernel threads:

```
ifconfig <interface> thread
```

The following are usually run by default. If not needed, then they should be turned off:

<code>rpc.ttdbserver:</code>	ToolTalk - used by CDE
<code>rusersd:</code>	Responds to queries from the <code>rusers</code> command
<code>rwalld:</code>	daemon handles requests from the <code>rwall</code> command
<code>cmsd:</code>	calendar management functions for CDE
<code>docsearch</code>	<code>/etc/inittab httpdlite</code>
<code>dt:</code> (Common Desktop)	<code>/etc/inittab dt:2:wait:/etc/rc.dt</code>
<code>qdaemon:</code> (print daemon)	<code>/etc/inittab qdaemon:2:wait:/usr/bin/startsrc \ -sqdaemon</code>
<code>writeserv:</code>	<code>/etc/inittab writesrv</code>
<code>muxatmd:</code>	<code>/etc/rc.tcpip</code>
<code>dpid2:</code>	<code>/etc/rc.tcpip</code>
<code>mkatmpvc:</code>	<code>/etc/inittab mkatmpvc:2:once:/usr/sbin/mkatmpvc \ >/dev/console 2>&1</code>

atmsvcd:	/etc/inittab atmsvcd:2:once:/usr/sbin/atmsvcd \ >/dev/console 2>&1
----------	--

If not needed comment out the following in /etc/inetd.conf:

```
chargen daytime discard echo rstatd sprayd talk  
rusersd rwalld dtspc pcnfsd
```

Appendix C. Reference Information

AIX

- AIX fixes and hardware micro code (SCSI and Fibre Channel microcode)
<http://techsupport.services.ibm.com/server/support?view=pSeries>
On left side of web page click on Fixes>Hardware microcode>Download microcode
- PCI adapter placement guide, document SA38-0538 from:
http://www.ibm.com/servers/eserver/pseries/library/hardware_docs
- Search AIX manuals and AIX commands
http://publib.boulder.ibm.com/cgi-bin/ds_form

Domino

- Domino Performance web site
<http://www.lotus.com/home.nsf/welcome/performance>
- Notes.net
<http://notes.net>
- NotesBench Consortium
<http://www.notesbench.org/bench.nsf?OpenDatabase>
- IGS e-business application COC (IBM Intranet only)
<http://w3.ibm.com/services/iga/gad/ebizappcoc/cocweb.nsf>

Miscellaneous

- SRM (IBM Intranet only)
<http://srmweb.raleigh.ibm.com>
- Boulder AMA (IBM Intranet only)
<Http://bldr1dbl.boulder.ibm.com/ama>
- dataseg utility information
<http://www.ibm.com/software/lotus/support>
Technote 189972, see also technote 190356 for related information

Acknowledgments

This paper was written with the input of a number of people who participated in the efforts of the S80 pilot. Specifically:

Rick Andony, IBM Global Services, Boulder, CO

Bill Bocchino (Document Owner), IBM Global Services, Somers, NY

Cameron Hildebran, IBM Global Services, Boulder, CO

Scott Hopper, Lotus, Westford, MA

Steve Mark, Lotus Enterprise CoC, Westford, MA

Mark Smith, IBM Global Services, RTP, NC

It is especially gratifying to have assistance from such a talented core set of people who have made invaluable contributions to the overall S80 project:

James Grigsby, Iris, Raleigh, NC

Wayne Hessler, IBM Global Services, Boulder, CO

Sheila Lavin, IBM Global Services, Boulder, CO

Darrin Saffell, IBM Global Services, Boulder, CO

Special Notices

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, North Castle Drive, Armonk, NY 10504-1785.

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 information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. 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.

The following terms are IBM trademarks:

AIX®
AIX 5L™
Domino™
Enterprise Storage Server™
IBM®
Perform™
pSeries™
RS/6000®
SP™
Tivoli®
Lotus®

Lotus Notes®
Notes®

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