

SPAZIO MFT/S ver. 2.30 for Distributed Platforms

Benchmarks

January 20th, 2007

Introduction

The goal of the test was to perform a benchmark between the SPAZIO MFT/S 2.3.0 release and the following ones:

- SPAZIO 2.2.7 RP03
- SPAZIO 2.2.7 RP06 (currently in production at Customer sites)
- SPAZIO 2.2.5 SP94 (last available 2.2.5)

The Involved Systems

The test was made using the following systems that were totally dedicated to the benchmark activities.

Machine	Function	OS	Description	Notes
ROBIN	Initiator	Sun Solaris 8	Spazio FTOM (WMQ Server) Spazio Orchestration Suite	Sun Fire 280r CPU: 2 UltraSparc III 750 MHz Mem: 2 GB;Disk: 40-MB/sec. Ultra SCSI (SCSI-2) (synchronous)
MISSMARPLE (MM)	Responder	Linux	Spazio FTOM (WMQ Client)	CPU: 2 Intel Xeon 2.4 GHz Mem: 4 GB
RACF (RF)	Responder	Linux	Spazio FTOM (WMQ Client)	CPU: 2 Intel Xeon 2.8 GHz Mem: 2 GB

The Load

The test was performed exchanging the following quantity of data:

Num Of Files	Total File Size (MB)	Avarage File Size (KB)	Needed Bandwidth considering 1 hour time span	From	To
6.300	4.000	635	9 Mbit/sec	ROBIN	MM
6.300	400	63	900 Kbit/sec	ROBIN	MM
6.300	100	16	222 Kbit/sec	ROBIN	MM
4	8.000	2.000.000	18 Mbit/sec	ROBIN	RF
6.300	400	63	900 Kbit/sec	MM	ROBIN
25.204	12.900		About 30 Mbit/sec		

Shared Memory Segments

The number of Shared Memory Segments have been checked before and after each test and NO LOSS have been found.

Benchmark Results

SPAZIO Release	2.3.0	2.3.0 SPOSNOFS YNC=ON	2.2.7 RP03*	2.2.7 RP06	2.2.5 SP94**
Elapsed time (minutes):	68	278	245	101	451
Throughput: (files/minute)	371	91	103	249,5	56
CPU consumption minutes (Note 1)	78	87	334	83	28
Number of system calls:	163,960,471	216,870,976	723,896,860	174,948,430	550,226,080
Average CPU usage % (Note 2) user/system/wait/idle	75/ 20/ 3/ 2	66/ 23/ 9/ 2	60/ 18/ 20/ 2	70/ 25/ 3/ 2	15/ 15/ 65/ 5

Quick comparison: SPAZIO MFT/S 2.3.0 as base reference

SPAZIO Release	2.3.0	2.3.0 SPOSNOFSY NC=ON	2.2.7 RP03*	2.2.7 RP06	2.2.5 SP94**
Elapsed time	100	408	360	148	663
Throughput	100	24	27	67	15
CPU consumption	100	111	491	106	035
System calls	100	132	441	106	335

Note 1: CPU consumption was obtained only for the resident modules. No measurements could be done for SPROUPR4 since these processes last only the session time and then they close.

Note 2: The CPU usage distribution has the following meaning:

- user: SPAZIO kernel processes
- system: Operating System on behalf of the process's system calls
- I/O wait: the potential CPU that could be used but that cannot be utilized since the processes are waiting for the disc I/O to end
- idle: the unused CPU, due to that fact that there are no processes in READY queue and/or all are waiting for network or terminal I/O to complete.



* **Comments to SPAZIO 2.2.7 RP03 results:** after the file transfers were over str_smon process has continued using CPU for about 68 minutes deleting the files “xxx.rms” present under the “rt” directory. This CPU time has been included in the total CPU time reported above.

** **Comments to SPAZIO 2.2.5 SP94 results:** the 2.2.5 SP94 release could not sustain the load and thus it was necessary, after all the dispatching processes were over, to shut it down and restart it immediately after. This was due to a failure in str_smon whose SPOS queue got filled with 512 queued messages, thus generating errors in the sending processes with subsequent hang situation.

This release suffers of a very high I/O usage and in fact the % CPU in wait state is very high (averagely more than 1 CPU in wait state for I/O).

Conversely it uses very little CPU and this is good for systems using very fast I/O devices. Our benchmark system has not a very fast I/O and thus the impact on the performance is extremely high.

Maximum value of CPU usage

The maximum CPU usage has been detected in str_smon and sd_iscl modules.

module	2.3.0	2.2.7 RP03	2.2.7 RP06	2.2.5 SP94
str_smon	4 min 12 sec	192 min 15 sec	9 min 15 sec	10min 18 sec
sd_iscl	11 min 19 sec	10 min 33 sec	10 min 40 sec	9 min 58 sec

Final considerations

This very high difference between 2.3.0, 2.2.7 RP03 and 2.2.5 SP94 is due to the following facts:

Better I/O utilization. The 2.3.0 release as well as the 2.2.7 RP06 allow the use of the SPOSNOFSYNC=ON option in SPOS. This option reduces considerably the overhead paid whenever a SYNC operation is requested.

The Transport Monitor has been greatly modified in order to reduce its overhead and in fact it consumes (in 2.3.0) 48 times less than the str_smon of the 2.2.7 RP03 release!

The CPU time usage, for SPSA and STR_FTRA processes has decreased, because of the better efficiency with which the Transport Monitor Restart File is managed.