

Notes.net

Iris Today

Home

Download

Iris Today

Iris Cafe

All About Domino

Iris Sandbox

Doc Library



CPU scalability

by
James Grigsby,
Nirmala Venkatraman,
and [Susan Florio Barber](#)

Level: Advanced
Works with: Domino 4.6
Updated: 08/02/99

Inside this article:
[Test methodology and test data](#)

[What did we find out?](#)

Related links:
[Optimizing server performance: Port encryption & Buffer Pool settings](#)
[The top ten ways you can improve server performance](#)
[CPU scalability test results](#)
[NotesBench Web site](#)
[Domino Performance Zone](#)

Get the PDF:



OPT2.pdf 36Kb



Do you want to get the maximum level of performance from your servers? If you are like most administrators, your answer is, "Yes"! You need to know what you can do to get this level of performance, and you probably want concrete test data that backs up the recommendations. If you've already read "[Optimizing server performance: Port encryption & Buffer Pool settings](#)," you might be thinking that the only test data that exists applies to servers running Windows NT. That's not true. Here at Iris/Lotus, we routinely run performance test scenarios with Domino servers running on UNIX. In this second article, we share some of these test results with you.

This article gives you an in-depth look at a performance analysis of CPU scalability. The test shows how changing the number of CPUs running on your Domino server can affect server response time. In a second test, we analyze how changing the disk RAID (redundant array of independent disks) level for the Domino data directory from RAID0 to RAID5 affects response time. We start by defining CPU scalability, then we describe the test methodology and test data, and finally we summarize what the results mean to you. This can help you decide how you want to set up your environment in the future.

For more background information about how we conduct performance analyses here at Lotus/Iris, or an introduction to the tools we use, see "[Optimizing server performance: Port encryption & Buffer Pool settings](#)." To read more recommendations for improving server performance, see "[The top ten ways you can improve server performance](#)."

What is CPU scalability?

CPU scalability refers to the process of adding additional CPUs to a server machine without causing excessive increases in complexity or loss of performance. Ideally, response time should improve with additional CPUs. Most organizations want to know the number of CPUs they should use in order to maximize the performance of their Domino servers. To answer this question, we set up a test scenario to observe how other system metrics increased or decreased when the only system change throughout the test is the number of CPUs running on the system. Most "CPU scalability" tests include several changes. (For example, the tester may change the number of CPUs, the amount of memory, and in some cases, the size of Level Two cache used). Changing multiple components simultaneously makes it difficult to determine if improvements are from additional processors or some other component.

In the second test, we changed the disk RAID level from RAID0 to RAID5. RAID is a data storage method where data, along with information used for error correction, is distributed among two or more hard disk drives in order to improve performance and reliability.

Test methodology and test data

In each test scenario, we used Domino R4.6x to establish a baseline for testing performance improvements in Domino R5. To run the test scenarios, we set up four to five Notes client simulators running our new R5 messaging workload with the following configuration:

- CPUs: Dell Dimension XPS D200, Pentium II processor
- Memory: 256MB RAM
- OS: Windows NT 4.0 Workstation
- Notes: 4.61
- NotesBench: 4.61

We set up a Domino server with the following configuration:

- CPUs: Sun Ultra Enterprise 4000, with 12 Ultrasparc 167MHz processors
- Memory: 1GB RAM
- Hard Drives: four RAID0 drives (total 8GB storage) for 2GB OS SWAP file, six RAID0 drives (total 12GB storage)/six RAID5 drives (total 8GB storage) for the Domino data directory. (For more information about the distinction between the RAID0 and RAID5, configurations visit the [AC&NC Web site](#).)
- OS: Solaris/Sparc 2.6, with Sun Enterprise Volume Manager 2.5
- Domino: 4.61 server for Sparc/Solaris

In particular, we wanted to test the relative impact (the number of users, the response time, and the resource utilization) when transitioning from four, to eight, to 12 CPUs. This test scenario compared response times and the system CPU resource utilization at the same user load, but varied the number of CPUs in the machine. We also wanted to test the relative impact (the number of users, the response time, and the resource utilization) on our Solaris configuration when we transitioned from RAID0 to RAID5. This test compared response times and the system CPU resource utilization at the same user load and with same number of CPUs in the machine, but changed the disk RAID level for the Domino data directory from RAID0 to RAID5.

The workload we used for all the tests is a new R5 workload. This is the R4.0 MailDb, but with the same server message delivery. However, it also includes adding and deleting mail, and the ability to exercise the server's directory for message recipient addressing. In addition, the message size increased by a factor of 10 (to 10,000 bytes).

We ran each test for approximately 90 minutes in a steady state, with a ramp-up period of around one hour (for 1800 users). For all the tests, we set the following shell environment variable:

`Notes_SHARED_DPOOLSIZE=4000000`

As documented in the [Release Notes](#), this variable controls the size of a shared memory segment or mmap files for shared data. We increased this value from the default value of 1MB, so that we didn't reach any limitations on the number of segments or files that the kernel would allow a user application to create.

We ran two of the tests using NotesBench on four to five Notes clients, each launching 300-400 threads (for a total of 1800 users). Using ThreadStagger=2 seconds (which starts each user logon at two seconds apart) on the client helped the server ramp up smoothly, without having connection timeouts during the ramp-up phase. We also configured the Domino Directory so that the Router would deliver all mail messages locally.

You can see the results of this test in the sidebar, "[CPU scalability test results](#)."

What did we find out?

When we tested four, eight, and twelve CPUs with RAID0, even at 1800 users, there was still 50 to 80 percent of the total CPU horsepower left over. Overall, Domino had a good response time at a particular user load when

we increased the number of CPUs from four to eight. Due to current code limitations, we did not see appreciable scalability in terms of response time after eight CPUs. We also did not see appreciable scalability in terms of concurrent users or capacity after four CPUs using RAID0. This laid the foundation (RAID0 provides the best response time but no reliability) for assessing the impact RAID5 would have on a system. We measured the impact by monitoring Domino transactions (NotesMarks), response time, CPU utilization, memory utilization, and disk response time.

When we moved from a RAID0 to a RAID5 disk subsystem for the Domino data directory, we observed a degradation in the user response time (the values increased) at the same user load. For example, when we ran this test with eight CPUs and 1800 users, the response time increased 150 percent (but the response time was still in the acceptable sub-second range), and there was a three percent increase in the amount of the CPU used. Also, the virtual memory page scan rate went up from 75 pages per second to 100 pages per second, and the average disk service times increased from 14 milliseconds to 34 milliseconds. These values show the effect of the disk subsystem on the percentage of the CPU used. You should take these values into account when you decide on the size of your Domino server.

One additional thing we noticed was that using the server's Public Address Book for all address lookups loaded the server heavily, and caused heavy network timeout errors on clients trying to connect to the server. Also, the clients failed to build the Message Recipient List properly because of timeouts and retries. The clients then sent messages to the server without any recipients, and the Router started producing error messages saying, "Unable to deliver message xxxxxx containing no recipients". With the efficient name lookup cache mechanism in R5, we can circumvent this problem and support more users.

In testing CPU scalability, in general, we found that Domino scaled well. Overall scalability for this system would have been higher with Domino partitioning, but we were focusing on a single Domino instance. Based on this information from our tests and with restrictions on the user response time, you can choose a Domino server containing the number of CPUs that best satisfies the load you need to support on your server.

This allowed us to get a baseline measure for CPU scalability on R4.6x servers. We used this information to identify potential performance bottlenecks, and we gave the information back to the developers here at Iris to help them further identify performance improvements in Domino R5. We expect that with major database, NameLookup and Router improvements, R5 will scale significantly better than R4.6x. In addition, when we start using R5 for our tests, we can enable Input Output Completion Ports (IOCP), where persistent worker threads service end-user requests on the Domino server. This will allow us to assess CPU scalability on UNIX servers in terms of the user load that they can support.

ABOUT JAMES

James is the project leader for the Domino Performance team. He came to Iris in 1997 from Lotus, where he worked in Product Management, covering areas such as, competitive analysis, performance, and the Notes server. Previously, he developed IT outsourcing proposals with Computer Sciences Corp. and had a career as an Air Force Officer working with computer systems at bases worldwide.

ABOUT NIRMALA

Nirmala Venkatraman works for Iris as a contractor. She started in April 1998 and primarily works on UNIX performance. She previously worked at Sun Microsystems.



[back to "[Optimizing server performance: CPU scalability](#)"]

CPU scalability test results (sidebar)

The following charts show the results of our tests. The first chart shows CPU scalability using RAID0 and the second chart shows CPU scalability using RAID5. For each test we used the R5.0 maildb as our workload. We ran each test with a user load of 1800 users.

NotesMarks are the Domino throughput metric of transactions-per-minute (tpm). They measure the amount of work the Domino server performs. The values are only comparable when the same workload is being applied to a server (as in our test scenario). The NotesMark values indicate that the server did the same amount of work regardless of the number of CPUs (the minor differences are insignificant).

The Percentage of the CPU Used column shows the percentage of the total CPU utilization, which is the sum of the percentage of user utilization, plus the percentage of system CPU utilization. These percentages come from vmstat, a standard UNIX system utility that provides virtual memory usage statistics. The Scan Rate column shows the pages scanned per second by the virtual memory system. This corresponds to the virtual memory activity as reported by the Scan Rate column in vmstat. The Disk Service Time column shows the average disk service time in milliseconds from the sar disk activity report. This report contains the average disk I/O operation time. Sar, the system activity reported, is a UNIX utility which samples cumulative activity counters in the operating system at the interval specified. This shows you the average time for servicing a read/write disk request.


Table 1: CPU scalability using RAID0

Number of CPUs	NotesMark (Domino Transactions)	Response time (ms)	Percentage of the CPU used	Scan Rate (Virtual Memory)	Disk Service Time (ms)
4	3331	5027	51	75	14
8	3638	100	25	75	14
12	3628	95	14	75	14

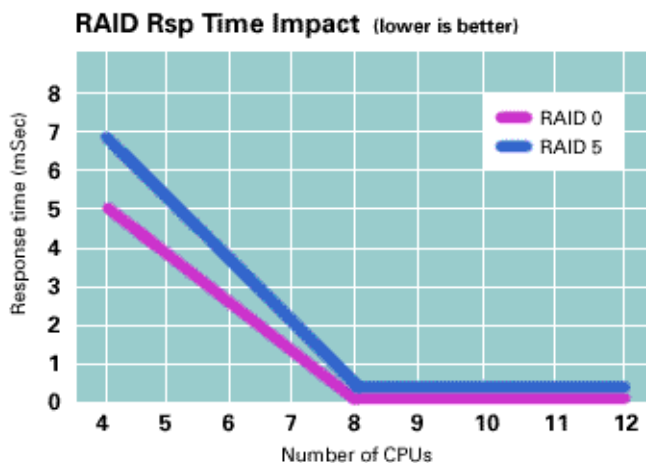
Table 2: CPU scalability using RAID5

Number of CPUs	NotesMark (Domino Transactions)	Response time (ms)	Percentage of the CPU used	Scan Rate (Virtual Memory)	Disk Service Time (ms)
4	3149	6730	60	100	35
8	3596	250	28	100	35
12	3598	246	17.5	100	35

When we compared the information in the two tables, we observed that RAID5 had a 30 percent, to over 200 percent, relative impact on response time. You can see this in the following graph. When we test response time, we often refer to both the absolute and the relative response time of a Domino server. The absolute response time is the measured time reported by Notesnum.exe. Notesnum is a utility that aggregates the response time from each Notes client simulator (we used up to five) to produce an overall response time average for all the simulated users. The relative response time is the difference in time between the series data point (in this case, the relative difference between the response time at four CPUs using RAID0 and four CPUs using RAID5).



However, as you can see in the following graph, the absolute response times were much better at the higher load values. The small-value response times (for example, 100ms versus 250ms) produced large relative response time differences (so it's important to assess both the absolute and relative values when making decisions). You should interpret this data as indicating that RAID5 did have an impact, but not a substantial impact. The increased data reliability provided by RAID5 is worth the increase in resource utilization. End users would not perceive any response time difference between 100ms and 250ms.



**Register
Here!**

About this Site | [Feedback](#)
[Lotus Home](#) | [IBM Home](#) | [Iris Home](#)
Copyright 1999 Iris Associates Inc.

