

Level: Intermediate
Works with: Domino 6
Updated: 04-Nov-2002

by Razeyah Stephen

IBM Lotus Domino 6 is packed with new features designed to improve performance. Early performance numbers that we obtained from the pre-release version of Domino 6 show that on the same hardware and system configuration you can increase your scalability and in some cases, even lower your resource usage, compared to R5.0.10.

This article discusses key Domino 6 areas in which we significantly enhanced performance. These include iNotes Web Access, IMAP, transaction logging, streaming replication, centralized directory, network compression, attachment compression, and Domino 6 platform statistics. We include some preliminary performance numbers based on the Domino 6 pre-release. In subsequent articles, we'll present more in-depth performance data, plus actual benchmark comparisons between Domino 6 and R5.0.10.

iNotes Web Access performance improvements

Domino 6 iNotes Web Access (iWA) performance was a top priority for the Domino development team. We focused on optimizing, streamlining, and removing bottlenecks in iWA and in the underlying Domino 6 code. The significant improvements we made include:

- Rewriting and optimizing the Domino 6 HTTP stack for performance (This will benefit *any* Web-based application, not just iWA.)
- Introducing new Notes memory management techniques to make memory allocation faster and to improve memory utilization
- Enhancing NSF database core code for performance
- Optimizing iWA code, especially for Calendar and Scheduling (C&S)
The R6iNotes workload implements C&S; we used this workload during the development cycle to uncover and address performance issues with Domino core C&S code. (For more information on Domino 6 workloads, see the *LDD Today* article "[The new Domino 6 NotesBench workloads: Heavier by request!](#)")
- Minimizing/eliminating semaphore contention in Domino code
- Optimizing the Formula/compute engine code

For information about the Domino iWA server and HTTP persistent connections, see the sidebar "[The Domino Web server supports HTTP 1.1 persistent connections.](#)"

The following table shows results we obtained running the R6iNotes workload on R5.0.10 and Domino 6 (pre-release). We used the same hardware configuration with Web authentication enabled for both tests:

Platform	Increase in users supported on Domino 6 vs R5.0.10	Improvement in CPU cost per user
Solaris 8	50%	50% decrease
AIX 5.1	100%	40% decrease
Windows 2000 SP2	40%	30% decrease

These results show:

- For Domino 6 on Solaris, the number of iWA active users increased by 50% with a CPU cost per user saving of 50%.
- For Domino 6 on AIX, the number of iWA active users increased by 100% with a CPU cost per user saving of 40%.
- For Domino 6 on Windows 2000, the number of iWA active users increased by 40% with a CPU cost per user saving of 30%.

The most important data derived from these results is the substantial CPU cost saving per user. This means a CPU-constrained R5 iWA site can support more users simply by upgrading to Domino 6.

The platform-to-platform variation in performance numbers is primarily due to the hardware used for testing, plus differences in core Domino/operating system code for memory management and semaphore implementation. Full details of these benchmark results (complete with system configurations and graphs for range of users) will be published in an upcoming *LDD Today* article.

Domino 6 IMAP performance improvements

The Domino 6 IMAP server was completely rewritten and re-architected to improve performance and scalability, to simplify administration, and to support additional IMAP extensions. Improved performance and scalability features include:

- An IMAP interface to the Domino database engine (NSF)
This eliminates the translation of IMAP requests into NSF requests, thus making IMAP database elements/structures native to NSF.
- An IMAP worker thread pool composed of LOGIN convert threads, FETCH threads, and FETCH response threads
- Storing folder content and message attributes inside IMAP-enabled mail files
This eliminates folder references and IMAP specific views from mail files.

We ran the R6IMAP workload on R5.0.10 and Domino 6 (pre-release), using the same hardware configuration for both tests. The following table shows the results we obtained:

Platform	Number of users supported on R5.0.10	Number of users supported on Domino 6
Solaris 8	5,000	6,000
AIX 5.1	2,000	10,000
Windows 2000 SP2	3,500	9,000

Notice that:

- For Domino 6 on AIX, the number of IMAP active users increased by 400%.
- For Domino 6 on Windows 2000, the number of IMAP active users increased by 157%.
- For Domino 6 on Solaris, the number of IMAP active users increased by 20%.

Full details on these results will be published in a future *LDD Today* article.

Domino 6 transaction logging performance features

The new Domino 6 lock manager for transaction logging was designed to help performance by only locking down database elements required for a function, rather than locking the entire database. This allows multiple operations on a database at the same time (for example, read access while doing writes). Also, we've improved the flush algorithm to flush the buffers (write data to the databases) more often. This reduces interruptions when the database is locked for write access.

View logging is a new Domino feature that allows transaction logging of view elements. Its purpose is to give faster access to a database with complex views after corruption. It does this by preserving the view information in the transaction log and eliminating the need for a view rebuild. View logging greatly reduces the restart time of large databases. Although this feature comes with some small performance cost, you should enable it on complex views. You can enable this feature per database view with the Domino Designer client.

For more information on Domino 6 transaction logging enhancements, see the *LDD Today* article "[Assessing the](#)

[impacts of new transaction logging features](#)," by Jim Powers. We will provide updated information on these features in an upcoming *LDD Today* article, including details on Domino 6 NRPC (Notes Remote Procedure Call, which is Domino native mail), and performance with and without transaction logging enabled. We'll also present comparison data with R5.0.10.

Streaming Replication

Streaming replication is a new Domino 6 feature enabled by default on servers and clients. To use streaming replication, both parties involved in the replication must be running Notes/Domino 6. Streaming replication is PULL only. This feature uses a single streamed RPC to read a number of documents and attachments. This is an improvement over the non-streaming method of requesting and acknowledging one database note at a time.

Streaming replication saves time and consumes less CPU and network activity to replicate the same number of documents. Network latency is reduced, especially on slower links. Documents are also replicated in ascending size order (smallest first). This allows the client to abort a long replication, but still receive some documents (for example, mail). Partial replication is also supported, allowing an aborted replication to be reinitiated and to continue where it left off. Views in folders are also updated incrementally while a replication is in progress (as opposed to the end of replication). This allows users to begin responding to initial documents (for example, mail messages) while larger documents continue downloading.

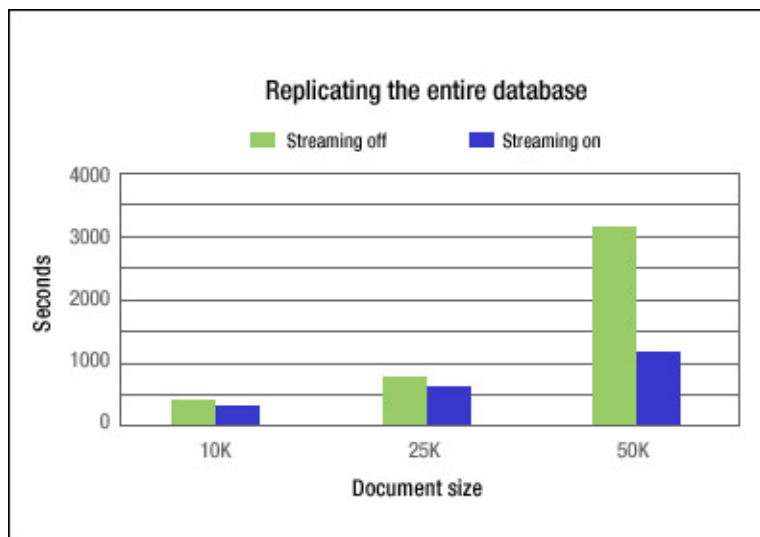
End-user benefits of streaming replication include:

- Shorter replication times
- Ability to start viewing and working on documents before replication is completed
- Replication of smallest documents first
If replication is cancelled, the target system still has documents (especially smaller ones), that were replicated before the replication session ended.
- Incrementally updated folders
You no longer have to wait until the replication is over before seeing replicated documents in folders.

Advantages for IT departments include:

- Faster replication
- Partial replication
- Potentially less network traffic due to single streamed RPC and reduction of ACK (TCP/IP) responses

The following chart compares replicating databases of different sizes with streaming replication on and off.



Databases with 10,000 and 25,000 documents replicated 15% faster with streaming replication on. A database with 50,000 documents replicated 60% faster with streaming replication on. For more information about our streaming replication testing environment, see the sidebar ["Test description and configuration used for streaming replication."](#)

Centralized directory

In Domino 6, you now have the option to move from a distributed directory architecture to a central directory architecture. To do this, you store the complete Domino Directory with all of its person and group information on one central server. The central directory information is available to all users, saving disk space because you no longer need to store the whole directory on each server. You also save time, as you no longer replicate your directory across all the servers in your domain. (You should, however, replicate the central directory to at least one other server for failover.)

On a pre-release Domino 6 server, we compared the performance of a local directory with a central directory. We also compared the name lookup rate of Domino 6 with R5.0.10. We used the IBM directory with about 500,000 names, performing lookups on a subset of 40,000 names. Preliminary results indicate:

- Using a centralized directory does not degrade performance.
- The additional load on the network bandwidth (the increased bytes sent/received when accessing a central directory from a mail server) is about 6 KB per mail server.
- Domino 6 can support *ten times* more name lookups per minute than R5.0.10!

As with other test data mentioned in this article, we will publish more details on Domino 6 centralized directory performance in an upcoming *LDD Today* article.

Platform statistics

Platform statistics is a server-based system performance monitoring tool, available since R5 for the Windows NT and Solaris platforms. Important system level metrics, such as logical disk utilization, system CPU, and memory performance, display on the Domino server console when you issue the command `>sh stat platform`. Monitoring the change in these metrics helps gauge responsiveness and performance of the underlying system (and therefore the Domino server itself).

Domino 6 offers many platform and feature enhancements to the platform statistics tool:

- Platform statistics is now available on AIX pSeries, Windows NT, zSeries (OS390), iSeries (OS400), Sun Solaris, and Windows 2000. In addition to system metrics common to all platforms, metrics specially designed to measure specific operating systems are included. These statistics (along with other Domino statistics) comprise the base data set for the Server Health Monitoring component of IBM Tivoli Analyzer for Lotus Domino.
- New system metrics, such as monitoring individual network and logical disk utilization, process CPU usage, and memory availability, help provide a more comprehensive and in-depth system performance analysis.
- New commands are now available to reset statistics and to help identify issues. For example, the command *Platform Stat Interval Reset Enable* helps you monitor the system on an interval basis.
- Platform statistics are now available via both the Windows and the Web-based Administration client. Statistics can be selected and graphed over time, using the charting mechanism available in the Administration client.
- Wherever possible, we've given statistic names a common naming convention across platforms. This makes it much easier to compare the performance of different systems.

There are many advantages to using platform statistics to monitor system performance compared to using OS-based performance utilities, such as Perfmon (Windows 32) or iostat/vmstat (Unix):

- Platform statistics provides derived and composite statistics that are unavailable through OS performance utilities. Statistics, such as the individual and combined network bandwidth utilization, help you determine performance and plan for network growth.
- Statistic names are homogenized across operating systems. For example, the System Memory Availability statistic translates underlying OS-level statistic names (Avail Kbytes on Win32, Free pages on Unix, and so on) to one common statistic name. This makes it easier for you to monitor disparate systems without advanced OS-level knowledge.
- You don't have to leave the Domino environment to gather these system level statistics. You can measure and chart performance on all servers, regardless of the operating system they're running, with the Administration client. Statistics are also available in the Monitoring database for future analysis.

Network compression

Network compression is an important performance feature offered in Notes/Domino 6. When you enable network compression, data is automatically compressed before it is sent over the network. This improves network performance, especially over slower line speeds.

With network compression enabled on both the sending and the receiving computers, the Notes client and/or Domino server attempts to compress any data, (for example mail or replicated documents) before transmitting it

over the network. Network compression works with mail routing, replication, or any other data sent via NRPC. This includes design elements in new or updated replica databases.

Notes/Domino 6 network compression offers a number of immediate benefits. For example, by reducing the amount of data being transmitted, you can improve the performance of your routing and replicating hub servers, especially if they're currently laboring under heavy workloads. And you can enable it by default, so all your users can take advantage of this functionality without having to select it themselves. Because network compression is a standard, out-of-the-box feature, it doesn't require any additional code, which helps simplify administration. CPU and memory resources used on the system increases by a small amount. But in many environments where it is very expensive or even not an option to upgrade the existing network infrastructure, this is a welcome trade-off.

For a detailed performance analysis of the advantages of network compression, see the *LDD Today* article "[Network compression in Domino 6.](#)"

Attachment compression

We also enhanced attachment compression in Domino 6. Attachment compression can now take advantage of the LZ1 algorithm (the same algorithm used in network compression). In R5, attachment compression used the Huffman algorithm. To use LZ1 compression in Domino 6, you must enable it in the Database Properties box. (By default, databases still use the Huffman algorithm to maintain backward compatibility.) The LZ1 attachment compression typically reduces attachment file size by 5% to 40% over Huffman attachment compression. The percent reduction in size depends upon the type and complexity of the file attachment.

Conclusion

We have discussed all the ways Domino 6 was designed and developed to improve performance. Performance results observed on pre-release Domino 6 show that on existing hardware and system configurations, you can *increase your scalability* and in some cases, *lower your resource usage* compared to R5.0.10. Watch for upcoming *LDD Today* articles that will go into much more detail on Domino 6 performance improvements. You won't be disappointed!

ABOUT THE AUTHOR

Razeyah Stephen is the co-lead on the Domino Performance team. She has worked at Iris since 1998. She came to Lotus from Digital Equipment Corporation, now Compaq, where she worked for five years in their StorageWorks division.



The Domino 6 Web server supports HTTP 1.1 persistent connections

With HTTP 1.1 persistent connections, once a client and server open a connection, it remains open until the client or server specifically requests the connection be closed. While the connection is open, the client can send multiple (but separate) requests, and the server can respond to them in the order in which they are received. Clients are also free to send multiple requests without waiting for the responses, in effect "pipelining" the requests. For example, a client may do this when requesting all graphic images from a particular page. It can also make the requests sequentially, then listen for responses from the server. Persistent connections are on by default in Domino 6.

Persistent connections are especially useful in the following scenario:

1. Browser opens a connection and retrieves a page containing multiple images.
2. Browser opens a second connection and retrieves images over both connections.
3. Browser closes the second connection and maintains the first connection. The server/browser closes this connection after a time-out or when the browser goes to another site.

For clients, persistent connections can make a noticeable difference on slow networks (for example, 56 Kb modems) when rendering a page. On fast networks (especially 100 MB/second), persistent connections aren't as beneficial because the cost of establishing a connection is very low compared to polling persistent connections.

Performance analysis of the Domino 6 iWA server (done on a 100 MB/second network) shows the number of users supported and resources used are about the same with or without persistent connections.



Test description and configuration used for streaming replication

- Domino 6 pre-release
- Server-to-server replication
- Replication of entire database
- Number of documents: 10,000, 25,000, 50,000
- Documents size: 10 KB
- Server A: Four 400 MHz CPUs, 2.5 GB memory, Windows 2000 SP2
- Server B: One 266 MHz CPU, 130 MB memory, Windows 2000 SP2
- Server B performed a pull-only replication from Server A