

# Replication Tuning for Mobile Users

by Doug Conmy

*[Editor's note: This article resides in "Notes Today", the technical Webzine located on the <http://www.notes.net> Web site produced by Iris Associates, the developers of Domino/Notes.]*

As I sat staring at my blinking cursor (a Notes editor cursor, of course), I tried to determine how to focus my thoughts and direct this article. I knew I wanted to provide a framework for the mobile Notes user to better understand how to manage their time most effectively.

Of course, this is a very large subject, and there are numerous sub-currents that could divert or overtake us along the way. I will try to avoid digressions, but please feel free to skip the introductory overviews and read on if you are a sophisticated Notes user. There will be enough content to offer something new for all audiences in these areas:

- Setting up your local databases to take advantage of replication options
- Defining and editing locations on- and off- the network
- Viewing and activating replication from the Replicator page

## Stacked Databases and Your Notes Workspace

When you start-up Notes, a number of folders appear with database icons. This is your Notes workspace. It serves as a starting point for performing your daily work. One of the first databases that you will probably encounter as a Notes user is your mail file. This is represented with a description, such as "Doug's Mail File on Local" and the graphic image of an envelope.

When you work remotely, at least two replica copies of your mail database will exist, one on the home server and one on your laptop machine. Two (or more) databases are "replicas" of each other if their replica IDs match. These IDs are displayed on the database properties page. Notes analyses this identifier to determine if two databases are actually the same.

The database replicas are represented visually on your workspace by a single "stacked" icon. Stacking them is a visual short-hand that simplifies your workspace management.

The icon on top of the workspace stack defaults to the replica database that makes the most sense for the current location or mode of use. Initially, a server copy of a database will be on the top of the stack (if you have network connectivity), while the local copy will be on the top for travel or home use (where you will generally have dial-up connectivity).

If you don't care for the default, you can switch replicas by choosing the stacked icon menu widget (affectionately referred to as a "thing-a-ma-jig") in the upper right corner of the icon. All you keyboard users out there can press the space bar while the database is selected to bring up the server picker menu. The last copy used will be remembered for the current location.

The mail database is the notable exception to the stacked icon default. To keep you oriented in your current context, the icon on the top of the mail database stack will always default to the replica specified in the current location record. Of course, you can change the mail database to another replica in the location document.

## Replication Is Your Friend

Yes, replication is indeed your friend, especially on the road. So what is replication, precisely? Although I have at times used the word "database copy" above, a database replica is more than just a copy.

The key is remembering that two databases are replicas if their IDs are identical. It is quite possible and often valuable to make a separate copy of a database (using File, Database, New Copy). This copy is assigned a different ID than its original. Such a copy is not a replica. Essentially, a database replica is exactly the same database as its original database, except that we can slice and dice its documents and data to meet our departmental or individual needs. Then, when we are ready, we can ask Domino to synchronize the replica with the original.

Replication provides the basis for convenient and resource-savvy management of distributed databases. For instance, if you are a remote worker, your database documents probably reside on your laptop as well as in regional offices in New York, Boston and Los Angeles. Over time, via replication, changes made in one database replica will be propagated to all other replicas.

The user with manager access to the database determines which users can read documents, create new documents and modify existing documents. A mobile user is usually granted the same level of access to corporate data while traveling as while seated in their office. Granted, the data is only as recent as the last replication, but this is often sufficient for many applications.

### **Replicas To Go**

Most Notes users have a small set of databases, besides their mail file, that are mission-critical for access and modification on a daily basis. When you are out of the office, you can always dial into your office using your modem and interactively access needed databases from the server while online. Of course, this assumes ready availability of a phone line as well as a healthy expense account, since interactive access generally forces heavy data transmission over the phone line as you browse through databases.

All things being equal (here is where the art of specifying your replica comes into play), a better way to use these databases is to create a local replica of the database on your laptop. The best time to do this is before you leave the office, since network access is faster and cheaper, but you can always make a replica over the modem as well.

### **Right-Sizing A Replica Database**

If the databases stored on the server are large, Notes enables you to specify the size of your replica according to several key parameters under your own control.

Although you can keep an entire copy of a database on your laptop, you may want to create a subset of a database to conserve disk space, reduce replication time or both. Always remember that you can change replica parameters later, reconnect to the server and reload the needed database subset.

The simplest user-defined parameter answers this question, "how recent are the documents you need?"

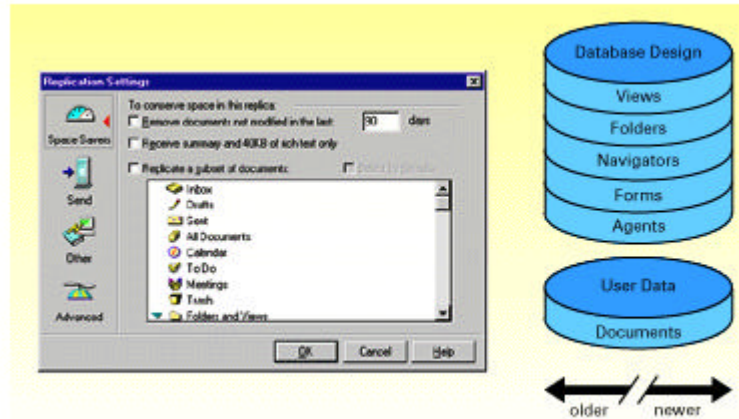
This setting will limit the database to documents that have been modified in the last "n" days (where "n" defaults to 90 days but can take user-specified values). This works well for mail databases or others that are strongly chronological, such as a discussion database. To set this or the other available parameters, let's actually create a database replica.

1. Select a database icon in your workspace for transfer to your laptop and choose "File, Replica, New Replica." Since laptops are not the world's most secure environment, you may want to encrypt the database so that you (more precisely, the ID file that is currently in use) and only you can open the database.

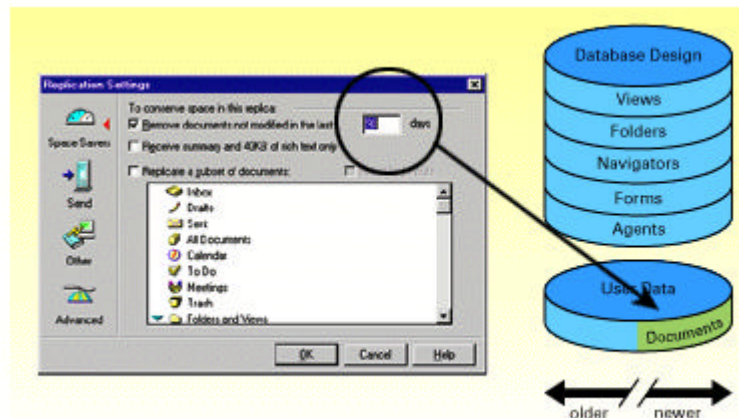
The encryption button dialog explains the pros and cons of each encryption level as it is selected. The default recommendation is medium encryption. However, note that medium encryption is not compressible. The database subset may actually be larger on machines that use compression software such as Stacker or DoubleSpace. The mysteries of computing are indeed great. You can

always encrypt a database later by choosing "File, Database Properties" and then "Encryption."

2. To minimize the size of the database, we bring up the replication settings dialog which allows us to specify what parts of the database we want. The first panel below, Space Savers, contains most of the settings a remote user would specify.



3. The first setting removes documents that have not been modified in the last 90 days (or some other number of days, if entered), and then replicates the most recent documents from the server database to your laptop computer. Note that the documents are not actually deleted from the server. You could consider the document subset a local (although detached) view of the document superset still stored in the server database.



4. The next setting, "receive summary and 40KB or summary data", is also referred to as "truncation" in the on-line documentation. Truncating documents brings only a small portion of a large document to the local database, namely, the document summary and the first 40K of text.. Bitmaps and any attachments are excluded. Of course, the full document can be obtained either immediately while online or by request at the next replication of the database.

Now for the truncation "if's":

*If this setting is chosen by the database designer when the database is created, truncation will occur by default for replication at all locations. A database designer can choose to visually indicate truncated documents from the view level. The mail database uses a truncated document icon in place*

of the attachment icon (it looks like a torn or partial document) so that a user can quickly determine which documents are missing some data. To get the rest of the document, the user can choose the menu command "Action, Retrieve Entire Document" after selecting one or more documents from a view.

*If this option is chosen to reduce replication while remote, a user may turn off document truncation while replicating in the office by changing the truncation option for the database on the replicator page.*

*If this option is chosen to save both time as well as disk space, a user will typically choose to truncate documents for this database at all locations.*

If you have several databases which contain truncated documents, you may want to retrieve one or more entire documents at the next replication after viewing their summaries. To do this, select the documents you would like to retrieve across the appropriate databases. Then, initiate replication from the replicator page of your workspace (that page, which cannot be deleted, is always to the farthest right of your visible workspace).

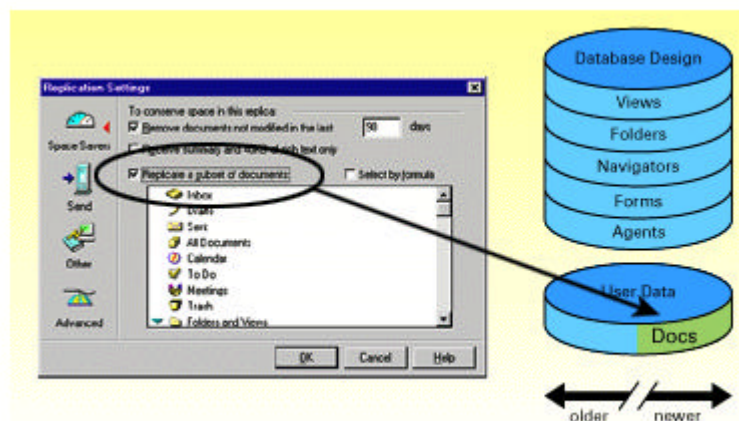
5. The final way to reduce the size of a local replica is to replicate a subset of documents. You can choose one or more folders, views, or combinations of both from the displayed dialog.

For example, a sales database may define different views that show the stored documents by region (often, North, South, East and West). You can select the view, 'Sales in the North', to replicate your region's documents only. If the views and folders available don't represent the documents you want, and you are comfortable with the Notes formula language, you can write a formula to specify precisely the documents you would like.

Note that the conditions you specify are additive. Should you have already specified retrieval of documents modified during the last 90 days while also replicating a document subset, both conditions will apply.

For the example above, the replica will retrieve documents from the view, 'Sales in the North' that were modified within the last 90 days.

6. Selecting "OK on the Space Savers" dialog will bring you back to the new replica dialog.



Select "immediately" to replicate the database right away. Replication proceeds in the background, allowing you to create additional local replicas or perform other work. To check on progress of the new replica, click the replicator tab from the workspace. The replication progress-to-date for the

currently active database will appear at the bottom of the page with an estimated time-to-completion for replication.

### Some Real-World Space Savings

To give you a feeling for space savings associated with the above options, here is some real-life data from three databases that I use frequently:

Database Name/Type	Local Replica Size/Docs	Server Replica Size/Docs	Subset Information
My Mail File	12MB 800 documents	50MB 2,220 documents Truncated	3 months of documents
Iris Office Notes	14MB 2,000 documents	160 - 180MB Truncated	3 months of documents
Large Company Database	68MB 7,560 documents	1,150MB 44,000 documents Truncates all except a select few	My group's documents

As you can see, the space savings can be considerable. Still, you may ask yourself, "why should I go to the bother of creating local replicas of databases when I can simply dial into a server and browse the most current documents while connected?"

This is a fair question. Beyond the often critical issue of free space on your laptop, there are equally critical performance factors that may come into play.

Recall that a local database replica has either obtained the data quickly over the network before going remote or has obtained the data once, in batch mode, over the modem. A local database will offer the best response time to your documents by a large factor since it is far faster to retrieve and display the document from your hard drive than when obtained over a 14.4 or 28.8 modem.

But there are equally important performance factors to consider on behalf of your team members as well. Server resources are always finite. By limiting the number of interactive users, those who must be connected across the network are given the access they require for their work, which often entails design, development and testing of Notes applications to serve your remote needs.

Nevertheless, if you use only a small number of documents in a database or cannot quite determine which documents you may need at a given time, the wise course is to obtain documents interactively through direct modem connection to the server. Likewise, if you do not use a particular database often enough to justify periodic replication of its data or the surrender of precious local disk space, by all means use the server database.

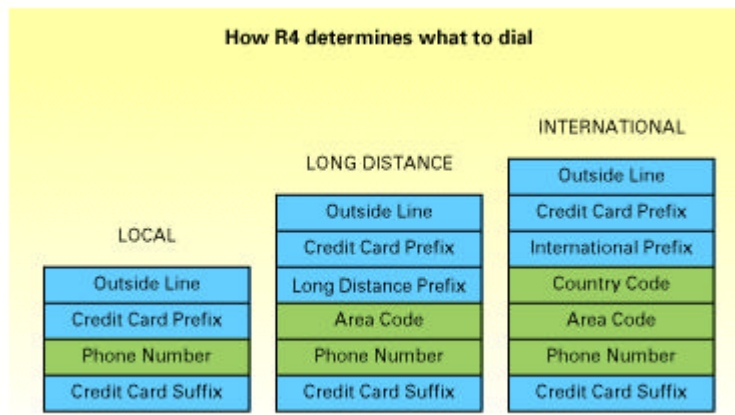
In other words, remote replication is an option, not a requirement. Deciding when to replicate comes with experience. The basic rule-of-thumb is this: the gain of replication must exceed the pain of interactive access.

### Modifying Your Connection Document

New Notes users are confronted with a lot of new information, concepts and decisions to make. Your administrator may have simplified life by setting up a personal profile when you were initially registered as a new user. The profile may contain links to databases of interest to you, as well as phone numbers to servers available to you while traveling.

Notes refers to the personal profile during initial setup to add icons to your workspace automatically, as well as to create connection documents in your name and address books. Connection documents provide Notes the information needed for successful server dial-up.

You can view the phone numbers of available servers, as well as create new connection documents for other servers by choosing the menu choice, "File, Mobile, Server Phone Numbers." The information in the connection document below (light blue for the Country Code, Area Code and Phone Number for a server) is used along with the information in your current location (dark blue below) to determine what number to dial to reach a particular server while you are at a particular location.



### Managing Locations

One concept that I have alluded to, but not yet explained, is "Locations." A location can be thought of as a physical place, such as your corporate "Office", or as a mode of operation, such as "Travel." A number of locations are created automatically to get you started. They are:

- *Office.* This location typically supplies direct network connectivity. When connected, you use server-based mail and databases. Mail is transmitted and received immediately.
- *Travel/Home.* This location often supports dial-up connectivity. You use a local mail database until actually connected to the server. A "Home" location is similar to a "Travel" location, except that a travel location prompts for phone information such as the number to dial to get an outside line, as well as the area code and country code at a particular location. If you travel regularly to several locations, you may want to create a unique travel location document by filling in the needed connection information, and by giving the location a meaningful name (for instance, "California Travel" or "Hilton Hotels."
- *Island.* This location is similar to the home location, except for the assumption that no ports are available. You cannot reach a server from this location, but must work exclusively in a remote, local mode.

You can easily access the active location setting from the status bar for editing. Simply click on the location name and select "Edit Current." The changes you make will take effect immediately after you close and save the location document.

### Working With the Replicator Page

At last! You have set up your local databases, verified your server phone numbers, and defined the locations where you work. Now, how do you get Notes to replicate what you want, when you want and show you what happened?

Recall that the replicator page is a visual window into your local databases and their replication status. This page can be thought of as a simple script that specifies the replication order for local databases. Replication is executed from the top down. Consequently, make certain that your most important databases are towards the top of the list.

A database can be selected and dragged up to the top of the replicator so that it is replicated before the other databases. Conversely, a database can be dragged to a lower position so that it is replicated later.

The clock icon at the top of the replicator page shows the next time that any scheduled replication activities will occur. To enable or disable scheduled replication, click the check box next to the entry. If you want to change the schedule, select the setting button to the right. This navigates to the current location where this information is stored.

By default, the replication schedule for the replicator page is disabled. You will probably want to enable it to keep your databases up-to-date. You will need to decide how often this needs to be done while in the office, as well as if a previously scheduled replication should occur while at home or traveling. To avoid scheduled replications, initiate replication of a single database or a list of databases manually.

If you know you are going to be traveling and wish to ensure that your databases are up to date with respect to a particular server (such as the one you might dial into), you can go to the replicator page, choose "Action, Replicate with Server" and specify the server targeted for replication. This will replicate each selected database with one on the server if a matching replica exists there. Again, this is best done in the office, ensuring quick and speedy connections and replications.

### **Calling and Hanging-Up Server Connections on the Replicator Page**

Please note, as shown below, that a "call server" entry is added to the top of the page, and a hang-up entry is added to the bottom of the page for use when dialing in from remote locations. This assumes that the server databases activated by the call entry can be connected to directly through the call, or by passing through the server that is called. For other network situations, a call server entry may initiate a remote LAN connection (such as RAS). In these cases, the other servers can be accessed because the user is attached remotely to the network.

Users who are authorized to reach multiple servers can create additional call entries. If one server is inaccessible, the next one will be tried. Once a connection is made, subsequent call entries will be ignored until Notes encounters a database to replicate or some other task to perform.

A call entry ensures that Notes performs the fewest possible phone calls. Instead of being forced to call multiple servers and replicate databases with each, we can call one server and replicate with all the other servers that are accessible from that server. If the user initiated replication, and a called server is busy, the server will be retried up to ten times.

A hang-up entry, typically placed at the bottom of the replicator page, is a convenient way to inform Notes as to whether the phone line should be hung-up. If the hang-up entry is checked, if the replicator initiated the call, and if the user is not browsing databases interactively from the workstation, then the dialed connection will be terminated.

On the other hand, if the entry is not checked, the connection will be retained until normal time-out periods cause eventual termination.

Finally, if the user deletes the hang-up entry or if it is not present, then the phone line is always terminated after replication has been completed.

### **Editing Database Settings On the Replicator Page**

Since all settings on the replicator page are location-dependent, any



database changes made to the page (for instance, their order, checked state and other settings) are for that location only. However, the databases themselves are not location-dependent. A database deleted from the replicator page is deleted for all locations. Likewise, a database added to the page is added for all locations.

If a database is removed from the replicator page, Notes assumes the user wants to turn off database replication for all locations. To disable replication of a database at a particular location, uncheck its entry. To prevent replication from all locations without removing the database itself, change its replication settings to disable replication.

This can be done by selecting the database on the replicator page, right mouse clicking and choosing replication settings. There are a couple of other options available from the right mouse menu.

You can conveniently change other settings for a database for a particular location beyond replication order. These are controlled through the settings button positioned next to a database. These include specifying what server or servers a database should replicate with, whether or not documents should be sent (--->), received (<---) or both (<--->), and whether or not replication should obtain an entire document, or just the beginning of the document if the document is too large.

To the right of the settings button, the replicator page displays summary information about the last replication time, and the number of documents sent or received. When these numbers are displayed in red, it indicates that new documents have been received since the last replication. The counts are cumulative for scheduled replications. If you replicate frequently, but only look at the page occasionally, you will see that you sent and received documents. However, if you initiate replication of a database manually, the count is reset to 0, and the count reflects the documents sent and/or received during the current replication.

As a convenience, Notes enables you to double click on a database entry on the replicator page to shortcut directly to the corresponding icon for the local copy of that database, select its workspace page and highlight the database. This is handy if you see some documents replicated to a database and want view them immediately.

### **Go Forth and Replicate**

. . . Or something like that. The most important point to remember above all is: don't sweat the details. Unless you have been given special permission by your administrator to delete documents on the server, there is nothing you can do with replication that will harm important corporate or personal documents.

In fact, we are working constantly at Lotus to simplify replication. Its complexities are not invented but real. They stem from the unavoidable complexity of synchronizing tens of thousands of documents created by hundreds of authors across dozens of locations worldwide . . . in real time. Replication is the racing engine under the hood of Domino/Notes that differentiates it from its competitors. A few moments devoted regularly to database replication will be repaid many times over by your increased productivity as a mobile Notes user.

### **ABOUT THE AUTHOR**

Doug has been with Iris since September of 1991. For Notes Release 4, he helped refine Notes for remote use, the focus of this article. Recently, he has focused his energies on Calendaring and Scheduling in the latest release of Notes. Before joining Iris, Doug worked at Wang Labs which he joined straight out of college. In his off-hours, Doug enjoys skiing, traveling and spending time with his 3 year old daughter, Kaitlyn.