

Notes.net Exposed: Using Domino clusters for your Web site

by Barbara Burch

[Editor's note: This article resides in "Iris Today", the technical Webzine located on the <http://www.notes.net> Web site produced by Iris Associates, the developers of Domino and Notes. First published in Oct. 1997, this article marked the beginning of our regular columns on exposing how the Notes.net site is run. The article explains how to set up a Domino server cluster for your Web site. This update also now includes information on how we use Cisco Systems' LocalDirector for load balancing, and details on the new clustering features in Domino R5.]

So, you've got a great Web site with an edge -- everyone's talking about it, and everyone wants to see it. A Webmaster's dream? Maybe -- if you've got the tools to handle it. In this article, you'll learn how you can use Domino server clusters to make sure that your great Web site can actually be appreciated.

A cluster is a way to group up to six Domino servers together to ensure high availability of your site, and to balance the workload between the servers. Domino includes the real-time Cluster Replicator, which ensures that all data changes are continuously replicated around the servers in the cluster. No matter how popular your Web site is right now, a cluster is a good investment, because it allows you to grow your system as the number of hits to your site increases.

In the rest of this article, we'll introduce you further to Domino clusters. You'll learn why clusters are a good idea for your Web site -- so good, that we use them here at Notes.net. Then, we'll walk through the steps for setting up a basic cluster.

An introduction to clusters

Clusters were first available for R4 servers as part of the Notes for Public Networks, an offering for telecommunication and Internet service providers. Then, Lotus released clusters for general customer usage as part of the Domino Advanced Services in R4.5. One of the big selling points of clusters is that they provide high availability to data, because you can access the same data from any of the servers in the cluster. For more information on this high availability, please see the article "[Lotus Domino Advanced Services: High Availability Powered by Notes](#)."

In addition to high availability, a Domino cluster for a Web site also offers the following advantages:

- *Fault tolerance:* You can now have a site that responds gracefully to an unexpected hardware or software failure. Clustering means more servers (up to six) are capable of responding to client requests to your Web servers.
- *Scalability:* You can easily add servers to the cluster as your Web site becomes more popular.
- *Workload balancing:* You can distribute user accounts across the cluster and balance workloads to optimize performance. The load balancing is almost instantaneous because of a single virtual address for all the servers.
- *Ease of multi-server administration:* You can use the real-time Cluster Replicator to keep site design and site content synchronized across the servers in the cluster. Because the Cluster Replicator is event driven, any changes to a database or to the cluster membership itself, are immediately passed to other databases or servers in the cluster. You can have multiple replicators on a clustered site.
- *Analysis tools:* You can use the Cluster Analysis tools to determine what is happening in the cluster and what you should do to optimize the performance of the cluster.

Prior to Domino R5, clusters provided load balancing and automatic failover only for Notes clients. To do this for Web clients, you needed to use an alternative method, such as [round-robin DNS](#) or another software or hardware-based solution. In the DNS method, a linear rotation distributes the Web user requests across the server cluster. You simply assign multiple IP addresses to a single DNS name. The


Domain Name Server then responds to Web user requests by rotating through the IP addresses of the physical servers in the list. Domino's real-time cluster replication ensures that the servers in the cluster are identical to one another. A drawback to this approach is that you rotate to a server regardless of its state. Plus, other site DNS servers and/or Proxy servers tend to cache URL references, which hinders your ability to turn off DNS rotation for a particular server. In some cases, it may take one to two days before this information is properly updated by other sites.

To overcome the drawbacks of the DNS method, you can use another software or hardware-based solution. These solutions work like a traffic cop. They "sit" in front of your server cluster and appear as the final destination Web server for all Web client requests. But really, they're a "virtual" server with a virtual IP address assigned via DNS. When a request comes in, they distribute the request to one of the "real" servers -- servers with a real IP address that reside in the cluster. The benefit is that now you can add and remove servers from the cluster, and not worry about cached references out on the Internet. You can do this with a proxy-software package, such as Netscape's or Microsoft's Proxy Server, or a hardware solution, such as Cisco Systems' LocalDirector. For more information on your different load balancing options, see the excellent article in Web Techniques magazine on "[Load balancing your Web site](#)."

Now, with Domino R5, you don't need to use these alternatives -- clusters provide load balancing for *both* Notes and Web clients. Domino R5 includes a new cluster component called the Internet Cluster Manager (ICM), which is like the software solutions we just described above but better, because it's Domino-aware. The ICM maintains information on the availability of the servers in a Domino cluster, and also maintains information about the distribution of databases on the servers. Then, when a request comes in from a Web client, the ICM determines the best server to receive that request. For more information on the ICM, see the article "[Domino R5: The Domino Internet Cluster Manager](#)."

Putting clusters to work at Notes.net

Now that you understand more about Domino clusters, let's look at how we use them at Notes.net. The following diagram is an overview of how we have a Domino cluster set up for the Notes.net site.



Because we're still using Domino R4.x servers, we use [Cisco Systems' LocalDirector](#) to provide load balancing for the site. All Web requests first go to the LocalDirector, which has the virtual IP address of `www.notes.net`. Then, the LocalDirector distributes the request to one of our "real" servers with real IP

addresses: www1.notes.net, www2.notes.net, or www3.notes.net. Of course, as our site migrates to R5, we'll investigate using the Domino Internet Cluster Manager for load balancing.

Each server contains replicas of all the databases in the cluster. Net1 is the administration server for the cluster and the replication hub, so it handles all replication with the rest of the servers in the domain and all agent execution.

Although our scheduled replication is set for every five minutes, Domino's real-time cluster replication is crucial for keeping our Iris Cafe discussions and Public Address Book synchronized. For example, it's possible that a user could create a topic, decide to edit it, and return to a server that has not replicated yet with the original server. If the discussion database replicates in real time, this time lapse is prevented. Similarly, the registration process and creating users in the Public Address Book is processed more efficiently. If an edit to an entry is required, the real-time replicator forces the change immediately. To maintain performance for the site, cluster replication is enabled only for the Notes.net Public Address Book and discussion databases.

Components of a Domino cluster

Before you begin setting up your Domino Web cluster, it might be helpful to get some background on the specific components that make up a Domino cluster. As with most things in Domino, a server cluster has its roots in the Public Address Book (NAMES.NSF). The Public Address Book is where you create a cluster, and then add or delete servers from that cluster. When you add a server to a cluster, the Administration Process (ADMINP) adds the Cluster name to the Server document.

For the actual day-to-day operations of the cluster, Domino uses the following components:

- **Cluster Administration Process (CLADMIN)** -- This server task is responsible for the correct operation of all cluster components. On clustered servers, the process runs automatically at server startup and whenever the cluster membership changes.
- **Cluster Analysis** -- Available from the administration panel, this tool enables you to analyze the configuration of a cluster and determine whether it has been set up correctly. By default, it stores the results in the **Cluster Analysis database (CLUSTA4.NSF)**.
- **Cluster Database Directory Manager (CLDBDIR)** -- This server task tracks all the databases on each server in a cluster, and stores the database information in the **Cluster Database Directory (CLDBDIR.NSF)**.
- **Cluster Manager** -- The Cluster Manager tracks all members in a cluster. It keeps a list of which servers in the cluster are currently available and maintains information about the workload on each server.
- **Cluster Replicator (CLREPL)** -- This server task is responsible for synchronizing databases and their replicas in a cluster. The Cluster Replicator uses the Cluster Database Directory to determine which local databases have replicas on other cluster members. Then, whenever a change occurs to a database, the Cluster Replicator pushes the change to the other replicas in the cluster.

Note: The Cluster Replicator only handles replications within the cluster. The standard replicator task (REPLICA) is still responsible for replicating changes to and from servers outside of the cluster.

As mentioned earlier, the Internet Cluster Manager (ICM) is a new Domino R5 cluster component that provides failover and workload balancing for Web clients. For more information on the ICM, see the article "[Domino R5: The Domino Internet Cluster Manager](#)."

Setting up your Web cluster

Each server that you want to include in your Domino Web cluster, must have the following:

- TCP/IP support
- Domino 4.5 or later (any supported platform)
- [The Domino Advanced Services License](#)
- A hierarchical server ID (required by the Administration Process)

In the following steps, we assume that you have at least two existing Domino servers that share a Public Address Book -- that is, the servers are in the same domain. The servers should also be able to successfully replicate with each other. You can set up database replicas before or after you create the cluster.

Part I. Planning for the Web cluster

Before you begin to create your cluster, you must first complete the following steps.

1. Verify that each server for the cluster has a hierarchical server ID. To do this, choose File - Tools - User ID. Select the Basics icon, and check that the ID type is hierarchical.

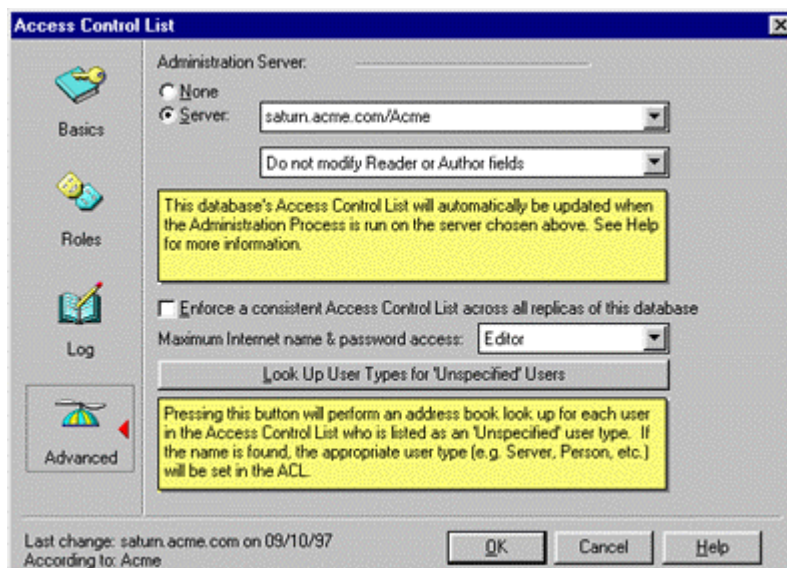
If you are using an existing server that has a flat name ID, you must convert to a hierarchical ID. For the conversion steps, please see the Administration Help.

2. Make sure that the Administration Process (ADMINP) is set up for your domain. To do this, you must first define an administration server for the Public Address Book.

Note: With Release 4.6, when you set up your first Domino server, it automatically becomes the administration server for the Public Address Book. However, you must complete these steps if you're *upgrading* to Release 4.6.

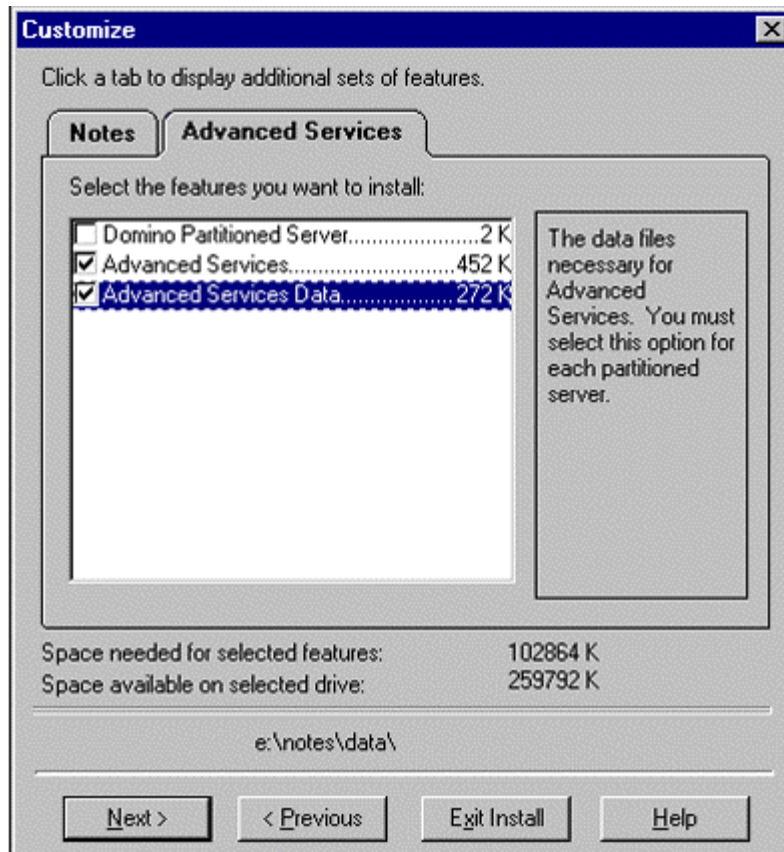
-- From the workstation of the server, open the local copy of the Public Address Book.

-- Choose File - Database - Access Control, and select the Advanced icon. At the top of the dialog box, click Server and select the name of the server that will be the administration server. Don't enable "Modify All Reader and Author fields."



-- Then, make sure that ADMINP is running correctly on each server. For more information on the Administration Process, please see the article "[Setting up the Administration Process in your Domain](#)."

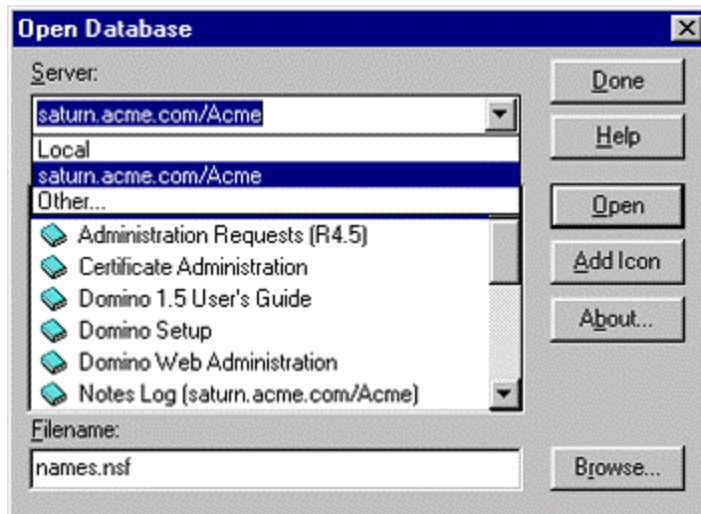
3. Install or upgrade to the Lotus Domino Advanced Services license on each server.
 - From the Install Options dialog box, select "Customize features - Manual install" and select the same directories as for your original server install.
 - On the Advanced Services tab, select Advanced Services and Advanced Services Data. Click Yes when prompted whether you purchased the Advanced Services license. (You do not need to install the Domino Partitioned Server.)



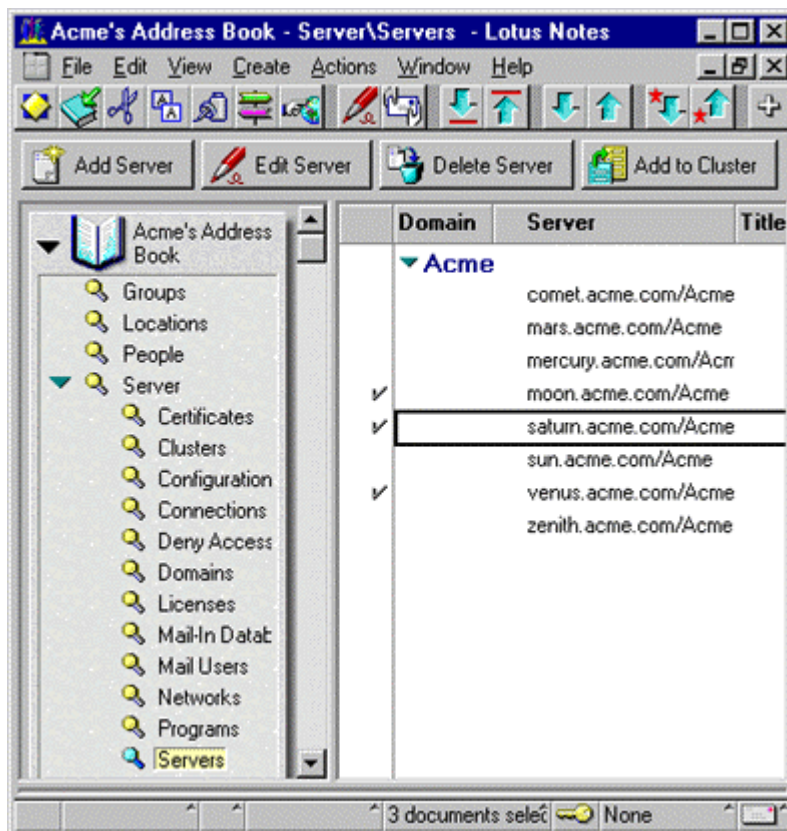
Part II. Creating the cluster

If possible, you should use the administration server for the Public Address Book to create the cluster, and add servers to the cluster. This administration server does not have to actually be part of the cluster.

1. Open the server copy (rather than local copy) of the Public Address Book. Choose File - Database - Open, select the server, and type the filename names.nsf. (If you are denied access, try switching to your administrator's ID.)



2. From the Server - Servers view, select the servers you want to add to the cluster, and click "Add to Cluster." Click Yes when prompted if you want to add the selected servers to a cluster. (You can have only up to six servers in a cluster.)



3. Select Create New Cluster, then type the new cluster name. Click OK. You are then prompted that your request to add the servers to the cluster has been successfully submitted to the Administration Requests database (ADMIN4.NSF).

Congratulations -- You've created your cluster! The rest of the steps in this section are optional, because they simply help you verify that the components of your cluster are now set up.

4. Run the Administration Process (ADMINP).

-- If you are using the administration server of the Public Address Book to create the cluster, ADMINP immediately runs and adds the cluster name to the Server document(s).

-- If you are not using the administration server, your request to create the cluster is queued in the Administration Requests database. When this request is replicated to the Administration Requests database on the administration server, ADMINP executes the request to update the appropriate Server document(s).

After ADMINP makes changes to the Server document in the administration server's Address Book, the changes replicate to the servers. (You can also force replication between the administration server and the added servers.)

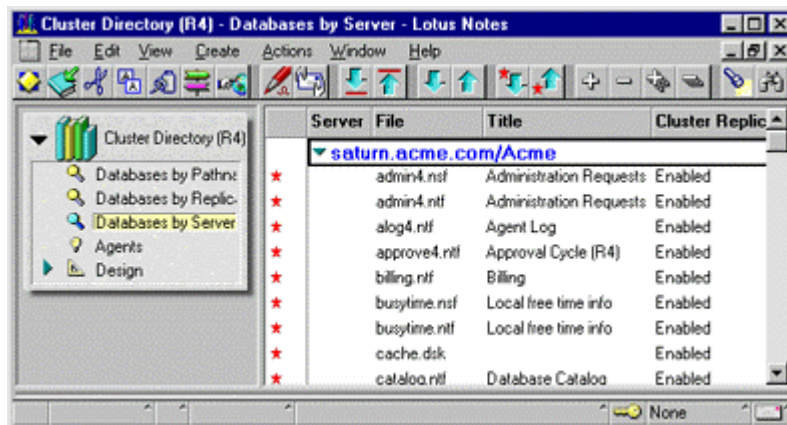
5. On the server you just added to the cluster, open the server copy of the Public Address Book. Verify that ADMINP added the cluster name to the server's Server document.

SERVER: saturn.acme.com/

Basics	
Server name:	saturn.acme.com/Acme
Server title:	
Domain name:	Acme
Cluster name:	NotesNet
Master address book name:	

The Server - Clusters view should also show the cluster, and the servers you just added to the cluster.

6. Again, from the added server, verify that the Cluster Administration Process (CLADMIN) created the Cluster Database Directory (cldbdir.nsf). Choose File - Database - Open, select the server, and type the filename cldbdir.nsf. All databases on the server should be listed in this directory.

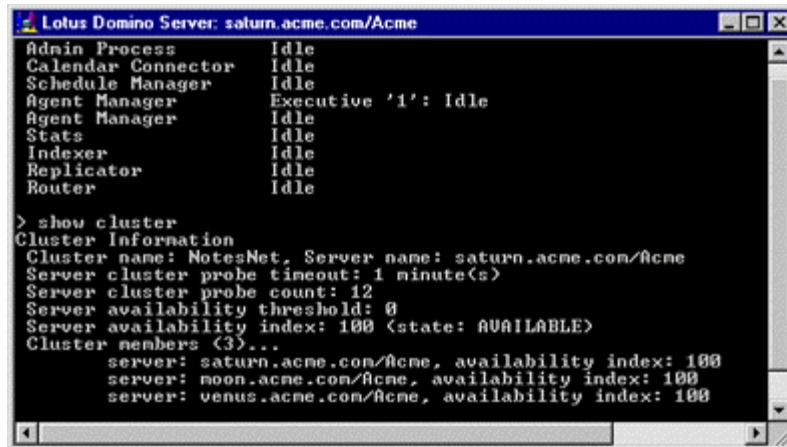


7. Now, check that the Cluster Administration Process (CLADMIN) loaded the Cluster Replicator task (CLREPL).

- From the server console, enter "Show tasks". The Cluster Replicator should be listed as Idle.
- If the Cluster Replicator is not listed, enter "Load CLREPL" to load the task.

When you add a server to a cluster, Domino automatically adds the server tasks CLDBDIR and CLREPL to the ServerTasks entry in the NOTES.INI file. This entry means that these tasks will start automatically whenever you re-start the server.

8. From the server console, enter "Show cluster" to show the status of the cluster. For example, you can now see the members of the cluster, and if they're available.



```

Lotus Domino Server: saturn.acme.com/Acme
Admin Process      Idle
Calendar Connector Idle
Schedule Manager   Idle
Agent Manager       Executive '1': Idle
Agent Manager       Idle
Stats              Idle
Indexer            Idle
Replicator          Idle
Router             Idle

> show cluster
Cluster Information
Cluster name: NotesNet, Server name: saturn.acme.com/Acme
Server cluster probe timeout: 1 minute(s)
Server cluster probe count: 12
Server availability threshold: 0
Server availability index: 100 (state: AVAILABLE)
Cluster members (3)...
  server: saturn.acme.com/Acme, availability index: 100
  server: moon.acme.com/Acme, availability index: 100
  server: venus.acme.com/Acme, availability index: 100

```

(Optional) Part III. Selecting databases for cluster replication

When you create a cluster, each database in the cluster is set to replicate in real time (that is, the "Cluster replication" field for each database is "Enabled" in the Cluster Database Directory). However, real-time cluster replication is only necessary for databases that you expect to change often. For example, the administration server for Notes.net has more than 100 databases, but only the Public Address Book and discussion databases are enabled to replicate in real time.

Note: Disabling cluster replication in this way only prevents event-driven replication changes -- that is, replication changes made by the Cluster Replicator. Any standard, scheduled replication (REPLICA) will continue to be performed both from and to the databases.

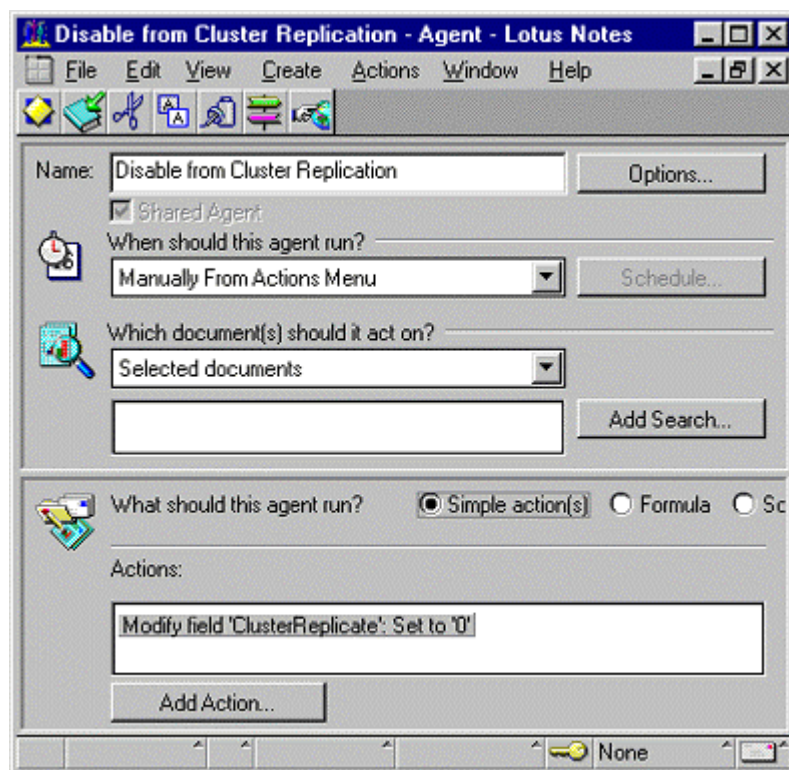
1. From the administration server for the Public Address Book, open the server copy of the Cluster Database Directory (CLDBDIR.nsf). Open any of the database documents, and notice that the Cluster replication field is enabled.

DATABASE:

BUSYTIME.NSF on server
saturn.acme.com/Acme

Server:	saturn.acme.com/Acme
File:	busytime.nsf
Title:	Local free time info
Replica ID:	8525650E:006F460B
Cluster replication:	Enabled
Database out of service:	No
Database marked for delete:	No

2. Choose Create - Agent. Then, name the agent, and select for it to run "Manually from the Actions Menu" and on "Selected documents."
3. Select for the agent to run "Simple actions," and click Add Action. Select to modify the Cluster Replicate field, and replace the value with 0. Save the action.



4. Select the databases that you want to disable, and run your new action from the Actions menu. To easily see which databases have replication enabled, sort the Databases by Pathname view by the Cluster Replication column.

Conclusion

Domino clusters are an excellent way to ensure that data on your Web site is current and dynamic -- so you don't lose your edge. Creating your cluster is just the beginning. You should check out the following sources for details about cluster tasks, how to set up workload balancing, and how to use the Cluster Analysis database for better performance for your Web servers:

- The "Iris Today" articles "[Fine points of configuring a cluster](#)," "[Workload balancing with Domino clusters](#)," and "[Domino R5: The Domino Internet Cluster Manager](#)"
- The [Domino Administration Help](#)
- The [Domino Cluster Zone](#) on the Lotus Web site
- [IBM Redbooks](#) on Domino clustering

Copyright 1998 Iris Associates, Inc. all rights reserved.