



## How I Lost Half My Domino Mail Server Capacity Without Users Noticing

by Danny Lawrence

"Hey Dan, Bill Thomas can't open his mail file. He's getting the 'This file must be on a Domino server' error."

That message from the Desktop Support staff dropped me into a rabbit hole of hardware failures, Domino server design, and the meaning of "high availability." Along the way, I learned some lessons that I want to share with you.

When we built the company's two Domino mail servers (Mail1 and Mail2), we ensured that each server was powerful enough to handle all the company's mail. (The company is a financial services firm, and many of the users "live" in their e-mail.) We then put half of the users on each server, and clustered the servers to provide high availability. Because the company has a lot of "pack rat" users, and it didn't want to impose restrictions on them, we built the servers with lots of drive space configured into a RAID array. At the time, I said to myself, "Given that the servers will be clustered, the RAID array is overkill, but if the company's willing to fund it, it's worth having."

Now, back to Bill's problem. I try to access Bill's mail file from my workstation and get the same error. Assuming that the backup software hasn't let his file go, I check the backup logs and discover that not only has the backup job finished, but it lists the mail file as successfully backed up. Checking Mail.BOX, I see several messages waiting to be delivered to Bill. As Alice would say, "Curiouser and curiouser."

I go into the server room to see what I can do from the server console. I'm unable to directly manipulate the file; Windows gives me some errors that make no sense at the time. I figure that the first thing to do is get Bill his e-mail again, so I change his home mail server to Mail2 and watch as all the pending mail routes into his mail file. I tell Desktop Support to tell Bill to restart his Notes client, and he should then be able to access his mail. Bill can now access his mail — but my problems are just beginning.

In the server room, I notice a red light blinking on one of the drives in the Mail1 RAID array. I call Dee, the hardware guru who built the boxes. Dee looks at the red light and says, "Oh, one of the drives has failed. I'll replace it before I go home tonight. That's probably the source of the problem." Another problem solved, I think to myself as I go back to my workstation and leave myself a reminder to move Bill back to Mail1 tomorrow after Dee has rebuilt the RAID array.

The next day, I run down Dee to ask how the RAID array is and get this unsettling reply: "It's strange. I put in the new drive last night, and it started to rebuild, but when I came in this morning, one of the other drives was dead, and we don't have another spare. In theory, the array can carry on until we get the spare, but the equipment is all less than three months old, and this kind of thing shouldn't be happening. I'm going to call the vendor and see if they can shed some light on this."

I inform the Desktop Support team that Mail1 is having some problems and to let me know if anyone calls in. There aren't any user calls, but late in the afternoon, Dee gets back to me with some more bad news. "I

called the vendor, and they say we're going to have to upgrade the firmware on the RAID controller. They're sending us a new controller, which will arrive tomorrow. I'll install it immediately. This means that Mail1 is going to have to be taken down, and the entire disk array is going to have to be rebuilt. This is going to take at least a day. Do you think Mail2 can handle it?"

My response is, "I hadn't wanted to do a failover test just now, but I guess we're going to find out whether Mail2 can handle the load!"

At the end of the day, we kill the Domino server task on Mail1, but Dee leaves the box running some diagnostics overnight to see if anything else shakes out. I leave a message for Desktop Support to inform them that mail access might be a little slow because we're having hardware problems and are working to remedy them.

The next morning, Dee tells me, "I've got bad news and worse news. The bad news is that the diagnostics revealed yet another bad disk in the RAID array, and the vendor says that we're going to need two additional new disks because the disks that failed and some of the ones that are currently in the array are part of 'a bad bunch of disks' that the vendor is recalling and sending replacements for. The worse news is that they won't arrive until tomorrow, and although I can install the new controller, I can't start rebuilding the machine without the new disks."

I ask Desktop Support how the users are faring, and Desktop Support says that, except for a few minor problems (e.g., a couple of users had set their inbox as their homepage, so they had to be told how to point to Mail2 to get to their mail), everything is fine. It's now past noon, and we've gotten through the "9:00 a.m. e-mail rush hour," so I'm feeling confident that Mail2 can handle the load. We monitor Mail2 closely through the rest of the day but it requires no attention.

The next day, the new drives arrive, and Dee begins the process of rebuilding Mail1. Mail2 continues to handle the company's entire e-mail load. User calls to Desktop Support are on a par with a normal day. When Dee asks me for the Domino CD-ROM to install the server, I say, "We built those boxes with Domino 5.0.4. As long as we're rebuilding the machine, let's put 5.0.8 on it. We can upgrade Mail2 later."

Dee asks, "Can you do that? Run different versions on two machines in a cluster?"

I reply, "Dee, I've clustered Domino servers running different OSs. Trust me, the point release version isn't going to make a difference."

We rebuild Mail1, restore the mail files, and set Mail1 to be unavailable for the rest of the day to give Mail1 time to "burn in" and to let the replication catch up the restored files with the current ones. We also figure that switching the users back to Mail1 during off hours will be less disruptive than doing it during business hours.

At the end of the business day, we make Mail1 available. The next morning, everything is back to normal. I even remember to change Bill Thomas's Person document to move him back to Mail1.

## **Clustering Saves the Day**

So in a nutshell, we had a catastrophic failure of half of our mail server capacity, and the impact on users was virtually zero. Because in Domino clustering, servers share no resources, even the failure of the RAID array didn't bring down the entire mail system. Because we were able to keep the mail system running, we could wait for the new hardware to be delivered and rebuild the dead server properly. Because we didn't have

management or users screaming, "When is e-mail going to be back?" we didn't have to implement a stopgap solution while we were rebuilding the server and then try to figure out a way to recombine the interim solution with the existing solution (think of the disruption to the users that would have caused). Taking our time on the rebuild led to a better server, as well. Because of Domino replication, we were able to restore users' day-old mail files and let the server bring the files up-to-date with virtually no intervention from the administrative staff. Because of Domino's version interoperability, we were able to do a point-release upgrade of the server while we were rebuilding.

Although the RAID array kept the Domino server running as the "bad disks" were failing, relying on one system for high availability still leaves you vulnerable to the failures of that system. That's why I recommend that all Domino administrators take a look at Domino clustering, especially because clustering is part of the enterprise server in Notes and Domino 6. In my case, clustering kept a hardware disaster from turning into a company-wide crisis. Clustering can make your infrastructure more scalable and reliable as well. In this case (because the company had the money for a deluxe solution), we were able to use RAID and clustering as a "belt and suspenders" approach. In other circumstances, where the choice is "belt or suspenders," I believe that a Domino cluster will give you more "bang for the buck" than a RAID array.

*Danny Lawrence is president of Tiassa Technologies Inc. You can reach him at [Tiassa@TiassaTech.com](mailto:Tiassa@TiassaTech.com).*

---

Copyright © 2003 - Penton Technology Media