



Level: All
Works with: Domino 6
Updated: 04-Jun-2002

Jim Rouleau on Domino 6 server availability

Interview by
[Laura Rutherford](#)

Keeping servers up and running so people can get their work done. It's a basic—and critical—task, that is also the main goal of our server availability effort. This month we talk with Jim Rouleau about what's new for Domino 6 server availability and how the automatic fault recovery features really set Domino 6 apart when it comes to server availability. He also talks about how these features came to be included in Domino 6 and what he thinks is in store for future releases.

Tell us about server availability—what does that term mean? How does it impact users? How does it impact a corporate network?

Availability simply means the amount of time that the product is available to be used. A user might think, "Well, if I can get to my email, that means the server is available." If users can't get to what they need in order to do their work, they are frustrated and their work stops. One down server can potentially affect many people across a company. Suddenly you have a hundred or a thousand people who are affected. That's a big deal. The administrators have to rush to get the problem fixed, and so it tends to disrupt a company.

Availability is actually the A in a bigger term called RAS, which stands for Reliability, Availability, and Serviceability. IBM has started a big effort to try to improve these areas across all of its products. I'm the project leader for the team that is going to be focusing on these areas for Domino.

When you say users can't get what they need, can you give some examples of what that means?

For instance, I've had the situation where I've spent time typing an email and I try to send it and the server goes down and then I can't send it. Or I'm expecting something important and the servers are down. You feel very cut off when that happens. We're selling a product that people can depend on and build their business on. People rely on our products. It's important that we do our part to keep our part of the chain up and running.

Give us a summary of Domino 6 features that improve or enhance server availability.

There's automatic fault recovery. There's view logging, which is in the transaction logs. We do view logging for special views, such as the \$ServerAccess views and \$Users views. It logs those views so that if you have to restart after a crash, you don't have to rebuild those major views, which helps the server to restart faster after a crash. We have a new database cache for a faster restart—so you don't have to reread all the databases on large servers after you restart. These are the big new fault recovery features that we added. We also tried to enhance the basic features we had in R4 and R5. For instance, we enhanced the fixup options, which allow you to do a fixup on the database while the server is running. You used to have to take the server down when you fixed up the open database.

We also have a new Java-based Domino Console, which is indirectly helpful because it allows you to get to servers. Even when the server is down—you can get to it and get it back up and running more quickly. And it's available on all of our operating system platforms. [See the *Iris Today* interview, [Mallareddy Karra on the Domino Console](#), for more about the this new feature.]



Tell us about the automatic fault recovery features. Are these features new to Domino or were they available in R5?

These features are not exactly new to Domino. Our recovery features actually started in R4 but they were only for the UNIX platforms. They were never really documented or played-up, and only a few people even knew about them. IBM was our biggest customer back then. They were big on availability and they wanted it—they were using AIX servers. Basically this stuff [fault recovery] was information that support people knew about and it would trickle out, but it was never really documented.

When I came back to work on UNIX late in R5, after having worked on mobile products for a couple of years, I suggested that we should make fault recovery available for all the platforms, including NT. I started that work about a year and a half ago—to make sure it was documented and was a real, solid feature.

For Domino 6, have you enhanced the fault recovery features that you offered in UNIX?

Yes. At the same time when I was doing the work for Windows, we filled in the gaps that needed to be filled. That was a great time to do it all. It used to be that the server would come down and restart. Then IBM said, "Well, we want to be able to run our own script or program when this happens." I think they used fault recovery to page their administrators and inform them of the server crash. So we enhanced it to say, we'll give you a NOTES.INI variable and we'll launch a script for you. We added that. We recently added a feature that will automatically send an email to the administrator that says a server crashed and was automatically restarted.

When I was down at DevCon last year, I was talking with some developers there. I mentioned how we were doing fault recovery for NT. They said they could not use that feature because they have seen problems when sometimes Domino will crash while it is initializing—it comes up halfway and crashes, comes up and crashes, and keeps doing that. As a result of that conversation, I added a feature that enables administrators to set some limits. For example, you can specify that if the server crashes three times within five minutes, then don't restart it again—let it just crash and stay down.

Can you tell us why these automatic fault recovery features are particularly important to users and administrators?

If the server crashes with fault recovery turned off, you get some kind of indication that the server crashed, but it is basically sitting there dead or simply terminated. The administrator then has to go over and do something to take action, to tear down all the processes and completely take the server down and restart it. Automatic fault recovery is going to do this automatically. Instead of the administrator having to walk over, tear everything down and restart

it, this process automatically happens. We skip the human step.

The benefit is speed. Instead of waiting for the administrator to figure out that something actually went wrong and having, for example, all the customers calling him or her saying they can't get their email, the server automatically detects when there has been a problem and takes everything down, runs your script if you choose to have one, and brings things back up. The time can be as little as seconds and your server can be back up again. The funny thing is that your users may not even realize that the server went down; the connection just has to be re-established.

Another benefit to users came from an internal request for client crashes. When the client crashes, users sometimes have to run utilities that take all the processes down; otherwise they have to restart their system. The client people said, "Can we use the fault recovery features to eliminate this problem?" So we fixed that issue in Notes 6. If your client crashes, you should just be able to restart it; you don't have to do any other tricks. This feature will be included automatically, you don't have to configure anything.

The server availability features are standard in the product, but do administrators need to do any configuration to use them?

Fault recovery is a section in the Server document that you enable and disable. We hope to have this in the setup so you can choose at setup time whether you want fault recovery on or not. In any case, it's going to be disabled by default because we don't want to affect the way that people are used to the server currently working.

How did you come to include fault recovery in Domino 6? Did customers request it?

Customers didn't really request it because they did not really know about it. I saw how well fault recovery worked in UNIX and how the bigger customers were using it and getting benefit out of it. I knew everybody would benefit from it. It's just that it was not designed for NT, it was designed for UNIX, so nobody wanted to tackle the job of enhancing it for NT. I stepped up and said, "Let's do it."

What challenges did you face when designing fault recovery for NT?

When fault recovery was designed for UNIX, it was designed to take advantage of operating system features unique to UNIX. There were system facilities that UNIX had that NT didn't and that's originally why we said we could do fault recovery with UNIX. We had to come up with equivalents for NT.

Think about things like, what if you have a password on your server? Normally the server console just sits there and prompts you for your password. How is that going to work with fault recovery? That would require an administrator to physically type in the password, which would defeat the entire purpose of fault recovery. So what we had to do is to store the password away so that it is available on the restart, but it has to be done in a very secure way. On UNIX it was encoded and stored in kernel memory. We had to come up with a similar method for NT; we had to encode it and put it in shared memory. But in NT, when all the processes go away, all the shared memory automatically goes away. We had to do something tricky; when everything shuts down, we had to start up a special process that does nothing but holds onto that password memory until the server comes back up and grabs it. We had to do a lot of that passing-the-baton-type-of-thing with fault recovery for NT since it was designed for UNIX.

Was fault recovery easier to do in UNIX?

It was easier to do in UNIX because of the way some of the system facilities work on UNIX. It was not impossible to do on NT, but it did require more designing. On the other hand, some issues were easier. We have to do a lot of work on UNIX to track and remove shared resources, but on NT we get that for free!

Once you made Domino 6 fault recovery available to customers, what was the feedback?

The customers who gave feedback from the early betas thought it was a great feature. The customers also asked for additional things, such as being able to email the administrator in this process. We put that in. I want customers to use this and I think this feature could be one of the real sleeper features for Domino 6. Customers want availability even though they don't always say it; they shouldn't have to say it. When the servers are available, they are happy. That will be my on-going challenge.

Who are your customers? Administrators?

Anybody who is depending on the product. Not just administrators, though they are the ones who will see it more. If a server crashes, people want it fixed and back up. But I expect that most of the feedback will come from administrators.

What are your ultimate goals when thinking about server availability?

If you want to define an ultimate goal, I'd say my ultimate goal is that no Notes user should ever see the dialog box that says Server Not Responding. We should be able to eliminate that. We know that we are never going to be

able to get every product perfect and that is really not our goal, though we put a lot of effort toward it. We want to be able to handle those times when we run into problems by getting the server back up as quickly as possible. We also want to be able to analyze what went wrong—and that includes maybe eventually making it easier for customers to get us that information, so we can do more data-analysis about what is going wrong and what do we need to fix that. The goal is to improve the product.

We test the product a lot here before we ship it, but our customers are savvy. They do things that we do not anticipate. They find problems that we sometimes can't. You want to make it a positive experience for them. In other words, we want the customer to tell us "Hey we found this problem. Here's the information." Then we analyze the data and say "Here's the fix." That would be an ideal goal.

What would you say sets Domino apart in terms of server availability options?

Maybe I am a bit naive, but I don't know of any product out there that has the ability to take itself down and start itself automatically. Most products when they crash, they crash. When it comes to servers, availability is the big story. If it is not available, it's not good. In addition to Domino's cross-platform availability options, we also work with each of the platform vendors to incorporate their availability features. Availability is so important, because it reduces TCO (which is total cost of ownership).

What do you see on the horizon for server availability post Domino 6?

If we do get fault recovery enhancement requests from customers, we will put them in—but the serviceability area is where we really want to put a lot of effort. We want to be able to analyze why a server crashed. If we can get the server back up quickly, we've done our job in terms of availability, but we need to understand why it has crashed so that we can fix the problem and prevent it from happening again. We've had customers repeat a crash scenario a number of times before we can get enough information to tell what went wrong. We want to reduce and hopefully eventually eliminate that.

Another goal is to be able to improve our first failure data capture story so that when the server crashes the first time, we get the information we need to analyze and fix the problem. That's an ongoing goal that we are going to be focusing on throughout the releases of Domino 6. That challenge will keep me very busy for a long time, but I'm looking forward to it because I've had the good fortune of making customers happy in the past and it's very satisfying.

About Jim Rouleau

Jim was born and raised in Fitchburg, MA, about 20 miles from Westford, where he works now. He spent 6 years in the Air Force working as a Russian translator. He's also conversant in French and enjoys spending time in Quebec. After the Air Force, Jim completed his degree. He received degrees in Computer Science and Mathematics and a Masters degree in Computer Science, all at Fitchburg State College. He's worked in the computer field for 13 years. Jim joined Lotus (and started working on Domino) in 1994. He lives in Gardner, MA with his wife Deb and cat Oreo. He keeps busy with his family, home improvements, and guitar. He and his wife also love to travel.