

[Home](#)[Download](#)[Iris Today](#)[Iris Cafe](#)[All About Domino](#)[Iris Sandbox](#)[About This Site](#)

John Banks-Binici: R5 Messaging & Directories

Interview by

[Betsy Kosheff](#)

Level: Beginner
Works with: All
Updated: 05/03/99

Inside this article:

Related links:
[Domino R5 Technical Overview](#)

[Art Thomas: Domino R5 Administration](#)

[Russ Holden: Domino R5 Database Improvements](#)

[Out of the Inbox: Mail processing with the new R5 mail agents](#)

Get the PDF:



JBinici.PDF_{62Kb}



[Editor's Note: To learn more about R5 messaging and directory features, check out the discussion with John in the [Developer Spotlight](#).]

As manager for the R5 messaging and directories teams, John Banks-Binici is responsible for a significant portion of the product's number-one stated goal: Deliver the best enterprise-messaging product that money can buy for companies of all sizes. But support for Internet standards, such as native MIME, SMTP, POP3, IMAP4, HTML, and LDAP, are just part of the equation. Being the best is about a lot more than interoperability. Here we see how his team's laser focus on ease-of-use, reliable security, and transparent interoperability help Iris meet the challenge.

Let's start with messaging. What problems are you trying to solve?

The whole R5 messaging effort is really about being both a world-class Internet mail system and interoperating with the other mail systems and products that our customers and *their* customers use. At the same time, we've paid lots of attention to cost of ownership. So our administration team has worked to develop a new task-oriented UI that makes the product easier to set up and administer. We've also spent a significant amount of energy addressing the needs of cc:Mail customers that want to migrate. The administration team has built migration tools into the core product, and we've added the router controls that the cc:Mail administrators expect to see.

All along, the underlying theme has been to help our customers reach out to their customers through the Internet. Until recently, companies had been concentrating mostly on how to solve their internal communications problems. Some of the more innovative ones have gone further to formalize communications standards with the rest of the companies and business partners in their value chain. They've been using Notes very effectively to do that. Now, the Internet is enabling them to turn and point their systems directly at consumers for business advantage, and we're right there helping to make that scale in a cost effective way.



What if customers don't do this -- what if they really just want to send e-mail?

That's fine; there's nothing wrong with that. But for certain businesses,

failure to recognize the fundamental shift in their customers' expectations can have a dramatic impact. For example, Fidelity had determined that they were not going to worry about the small investor that wanted to trade stocks online at discount, but were going to focus instead on the high end of the market. And they ended up losing a significant portion of their customer base as a result. You'll hear John Landry, the former CTO of Lotus, make this point in his speeches. Cases like this have made many of our customers focused on interoperability issues. Our job is to provide the solutions, and something the market has clearly been asking for is newer and better ways to point their systems at the Internet.

What are the big messaging enhancements that you're working on for R5?

We're working on four major areas across the whole product: *native* RFC 822 messaging and associated MIME and HTML support, *native* SMTP routing, *native* Internet addressing, and alternate name support for international users that want to display their names one way for those recipients who share their language, and another way for those who don't. In addition, the client has added native RFC 822 messaging, MIME, and HTML support; IMAP4 support for access to IMAP4 mail accounts; LDAP3 support for directory queries; and mail rules for synchronous processing on delivery.

Related to messaging, the database team has added transactional storage capability, online and in-place compaction, as well as API's for online backup. We've done a lot of work on the Domino Directory too, but we can talk about that later. (Editor's Note: For more information on the database team's work, see our [interview with Russ Holden](#), project leader of the database team.)

By the way, we're really a consumer of the native MIME work, which has been done by another group. The native MIME team has had one of the toughest jobs in R5 and my hat goes off to them for the great job they've done.

Let's talk about MIME -- what was the technical challenge?

Well, it's well known that Notes had a proprietary format for storing messages in rich text because we've been doing this for more than a decade, long before there were any standards. So in the past, we had to convert Notes rich text (called CD records) to MIME plain text format when we sent messages to the Internet and we had to do the opposite when messages arrived from the Internet. This process had the potential to cause bottlenecks at the Lotus MTA (message transfer agent), where we did those conversions. We realized that what we needed to do was to make RFC 822 messages and their MIME and plain text, or HTML bodies a native format of Notes, so that no conversion would be necessary. This also gives the best level of interoperability since there's a chance of fidelity loss on conversion. Along the way, we've had a number of integration points: the client, mailer, mail template, server, router, and directory. These are places where we've had to code logic to process the new formats and protocols in ways that are compatible with prior versions of Notes. It's been challenging for all of us. I've had the pleasure of supporting a team of truly brilliant mail and directory people in doing this work, so at the same time, it's been exhilarating.

Does MIME actually do as good a job as the rich text environment of Notes?

First, I need to clear up a little terminology. Technically, what we've done in R5 is to store RFC 822 messages from the Internet natively as Notes items and objects. We call this "native MIME" support, but in some ways, that's a misnomer since an RFC 822 message doesn't need to contain MIME. However, most messages that get created these days are

structured as MIME. The MIME message very often contains an HTML root and some, but not necessarily all, of the objects that the HTML references. This is referred to as an MHTML message. So, it's the HTML that is the functional equivalent of CD in this context.

Now, I can answer your question, which is really about how HTML compares to CD. I would say HTML covers probably 98% of what Notes rich text can do. There are still constructs you can express in a Notes CD format -- for example, "collapsible sections" -- that can't be done in HTML. This doesn't present a big issue for customers who want to interoperate with other mail systems since there are reasonable mappings, but it does put the onus on Lotus to try to standardize the CD technology that doesn't overlap. That's something many people think we need to look at right after we ship R5.

How does the MIME support in R5 improve upon 4.6?

We're really eliminating the need for the MTA to provide conversions. This is going to improve interoperability as well as remove a potential bottleneck, because you no longer have the MTA when you're administering Internet messaging for your company. The role of providing MIME-related services is distributed among R5 servers instead of being concentrated at one or two MTAs. Also, the way we store the MIME internally is more efficient than in R4.6.

So you no longer need to do conversions? Will that improve performance?

Yes, this does improve performance quite a bit. But it's not true that we never do conversions. We still need to provide the service to support down-level clients and servers. R5 doesn't need conversions. In a mixed R5 and pre-R5 mail system, we've worked very hard to reduce the probability that a conversion will be required, and to improve the performance of the conversions when they do happen.

How did you accomplish that?

The server provides, in place of the MTA, those conversions that are needed to support pre-R5 clients. We also built the smarts into our SMTP support, so that when an RFC 822 message arrives from the Internet at an R5 system, it can be stored in a format that doesn't require conversion, and can be easily served up in pieces to clients as they need them. In Release 4.6, we stored these messages, but they were stored as what we've termed an RFC 822 or MIME "blob," which is just an attachment. That's okay, but when clients connect and try to read the message, the server actually has to deal with the entire blob. One of the whole points of our "native MIME" work is that it breaks up the RFC 822 messages into parts so the server doesn't have to read the entire message in order to process it. Being able to work this way improves performance quite a bit, especially for SMTP, IMAP, and POP3.

Do you see your job as shielding people from these kinds of complexities?

Absolutely. Customers deploying R5, especially those who have a previous version of Notes already installed, will have all of this content managed for them transparently. We built a lot of the smarts in, so we will automatically do the right thing to the content. So if someone sends MIME to a R4.x user, we make sure it gets converted to CD for them. If the message is going to a R5 client, there's no need to convert at all.

How are the different types of user preferences specified?

Every user in the Domino Directory can have one of three settings: they can be Internet only, Notes and Internet, or just Notes, and this is determined in the Person document. This setting specifies what types of content the user's client prefers to handle. So, for example, the default

setting for the R5 client is "Notes and Internet" because the R5 client can send and receive either format. By knowing a recipient's setting, we do a lot of the work to make sure that the data gets optimally stored and delivered to that user. So, if "Internet only" is the setting for you, say, because you're using an IMAP client, then R5 clients automatically know that you want MIME because they'll look you up in the Domino Directory or Catalog (more on that later), and they'll send you the content you prefer from the start, relieving the need for conversion in the first place. If someone sends your IMAP client CD records, we'll make sure that we convert the message when it arrives, and store it in this new and optimally accessible structure we call "native MIME."

Is there an added benefit of using the R5 client here?

Sure. If it's the R5 client that's doing the sending and you're sending to ten or so people, it will look each one up in the directory, and determine who wants MIME, who needs CD records, or who doesn't care. The Notes editor, working with the new "two-pass mailer," will split the message and send MIME to those who need it, and CD records to those who need that. This process is not a conversion. Conversion is bad because once something is already in HTML, converting to CD records or vice versa doesn't always map perfectly. The only thing that really does a good job of generating HTML properly is our editor, since it creates the content in the first place and understands its structure above and beyond what tagged text can convey. So, the most optimal place to generate HTML (or CD) is at the editor. Furthermore, at the source client, we can know what format the recipient wants to receive, so the message can also be encrypted in the appropriate manner. This is important because you can't encrypt or sign something and later decide to convert it. Encryption and signing need to happen only once and at the source.

So the R5 client knows what type of format it should be sending to an individual recipient, and what algorithm to use to sign or encrypt that particular format, which allows us to seamlessly incorporate S/MIME signing and encryption side-by-side with Notes signing and encryption. It's simple really. You just select the people you want to mail to, and then we figure out what should be sent. If the recipients aren't listed in the directory, they'll get the default format designated by your administrator. Furthermore, we've built in the smarts to recognize RFC 821 and 822 addresses, so we can default to sending MIME to users on the Internet without requiring them to be looked up. This two-pass mailer approach is just a huge win, especially for extranets, which have lots of different kinds of mail clients.

You've talked about the existing Notes customer; what's the win for people who have never used Notes before?

It's really about the richness of our features and the stability of our products, which is what people have come to expect from Iris, Lotus, and IBM. Our security infrastructure is renowned for its robustness. Our high availability clustering technology is state of the art. The time it takes for us to load balance or failover a mail client to another server is measured in milliseconds. We scale from the small business to enterprise, from the aircraft carrier to the fleet. Platform availability is also a big plus. All of the areas mentioned are ones where we have a significant technology lead. But we've gone beyond all that and added plenty of new features for R5. For example, we've done a lot to enhance the execution of agents on mail delivery, and that's a benefit of the Notes mail system overall, one that you gain beyond complying with the SMTP standard.

And in-market experience means everything when it comes to providing a stable mail environment. Our router has been out there for over ten years and has a number of great features. We've also added a number of new

ones. For example, we've added router controls so you can filter out spam and unwanted junk mail from the Internet or your intranet. We've also listened to our cc:Mail customers and added controls that let you bounce a message or change its priority so that it gets routed at off-peak hours, depending on its size. These are all settings in the Domain Configuration document.

What else is new in the router?

Overall, we've enhanced our router so it can route over both SMTP and Notes Remote Procedure Calls (RPC). You get a lot of benefits from RPC routing. Internet standards like SMTP are really just designed to specify how you get mail from one domain to another, but they don't do anything to determine rules for how it gets delivered end-to-end once the mail arrives in a domain. So, if you have a lot of different kinds of protocols, like SPX, and your mail has to travel over those, then SMTP isn't going to work for you. But Notes RPC does support SPX, as well as X.25 and NETBIOS. So our router can still get your MIME message to you without conversions, which is a benefit you will not get from other mail systems. This sort of flexibility is something that IS people love since they never know which company they're going to be merging with next, and what sort of infrastructure they'll have to work with. It's a cost of ownership issue for them. Another example of where we've worked hard to decrease cost of ownership is in SMTP configuration. We've simplified it down to one form.

What about router performance?

We spent a lot of time working on the performance of our router. We've increased the degree of parallelism we're capable of. For example, we use multiple instances of mail.box databases (where mail being routed gets deposited). Contention for access to a single mail.box is one of the largest constraints on router performance, so supporting a number of mail.box databases and randomizing access to them eliminates a major bottleneck.

Our router was always multi-threaded, but transfer to any given server was always handled by a single thread. Now, we also have multiple transfer threads between the same two endpoints, so if you are pumping one long message from Server A to Server B, it will not hold up other messages that need to get through just as quickly. This especially increases the throughput if we need to convert messages for down-level clients.

And we have multiple delivery threads. In the past, if you had New Mail agents, they could only be run asynchronously. One of the problems with this was that, depending on what they were doing, your agents could take a long time to run to completion, in which case your mail delivery performance would be affected. That's why we only supported asynchronous new mail agents. Now, with multiple delivery threads, you can set up a synchronous New Mail agent and it won't adversely affect the router's performance for everyone else. (Editor's Note: For more information on the new mail agents in R5, see "[Out of the Inbox: Mail processing with the new R5 mail agents.](#)")

Any other new router features?

One I like is that now you can execute canned mail processing rules on delivery. These are written in C so they're really fast. There are four different kinds of rules and you can set up a selection criteria to determine which messages the rules apply to, for example, based on who the messages are from or whatever criteria you prefer. The rule can be set up so that you can decide to delete the mail, or perform a folder operation on it, or change the priority. So for example, if you subscribe to a mailing list, you can have those messages put in the appropriate folder automatically.

There's also a Pull router feature that allows us to pull mail. In the past, we

only supported this over the XPC protocol and it was referred to as "opportunistic routing." But now, we've got it working over all the supported Notes protocols including SMTP over TCP/IP, so you can periodically connect to your service provider and get your mail.

We've also done a lot on IMAP performance enhancements. Native MIME support alone is a big win -- we've already more than doubled the performance and we're certainly not done yet.

What does the new support for native Internet addressing bring to users?

Native Internet addressing means that every user can have both a Notes e-mail address and an Internet address. Depending on how you deploy messaging in your organization, you could depend on one or the other, or both. Plus, we've provided tools that allow the administrator to verify the uniqueness of those new Internet addresses and to generate them programmatically.

So for example, you can apply a pattern in which everyone in your organization has the first letter of their first name, followed by the first ten letters of their last name as the local part of their RFC 821 address to which we'll append an @ sign followed by their domain. The administrator can use the tools to register users and also continue to maintain existing users to make sure no one has changed their address or forced it to something that is already owned by someone else.

Is this new addressing going to affect current users?

We spent a lot of time on backward compatibility, so that when you deploy native Internet addresses, your existing applications continue to work. Each message that we create has both a Notes address and your Internet address -- depending on whether or not it's going to go out over SMTP, we'll pick the right address to send over that protocol. If it goes out over SMTP with your native Internet address and is received at a Notes server, we know how to take the embedded Notes address and promote it on the receiving server so your workflow applications can continue to work.

Is there a benefit besides backward compatibility to having two different addressing approaches?

There is. Notes addressing works by specifying at the sender, what route you want the message to take, or an individual administrator can set up domain documents that know how to get the message to the right domain. On the Internet, it is always assumed that every domain is connected. If you want mail to go from Domain A to Domain C, you connect to Domain C, which is specified in the mail address, and just give them the message. With Notes mail, you don't have to do that -- you can have Domain B in between and have it pass on the message to Domain C. The benefit is that you don't have to have a connected network, which is often desirable. What we did is allow you to still send mail from one Notes user to another and not have to sacrifice either of those approaches. Everyone has two addresses and you can use either one. If you use an Internet address -- great, we'll connect to Domain C and hand them the message. If you use a Notes address, we know what to do with that too. It's really up to the administrator to decide what sort of mail routing infrastructure they want to deploy. We're flexible.

What about LDAP and directory interoperability?

There are several levels of directory interoperability and LDAP plays a key role in making it happen. Ultimately, we want to go beyond the level of interoperability we have today to what is sometimes described as interchangeability -- the ability to swap out one directory system for another. For example, a customer wants to leverage the security infrastructure of our directory while working with other LDAP directories.

That's their first concern, and they want to know how to pass LDAP referrals and mail addressing requests around their directory environment. We're doing that for them with Domino Directory. But over time, they may decide that they want to use Domino as their primary enterprise directory; or they may not want to have to use our directory at all. We want to do a great job at fulfilling both of these requirements.

What steps have you taken?

We've built on our previous work by adding LDAP v3 support so you can read and write information, not just read. Our overall goal is to make the Domino Directory the best general-purpose directory for any application to use -- not just Notes applications. To do that, you need to let non-Notes applications not only read and write, but store their own application information in our directory. You have to add extensible schema support so that these applications can add and administer their own objects. And that's what we've done with R5.

Can you give an example of extensible schema directory support?

IBM has a suite of Windows NT BackOffice applications, for example DB2, and these applications work with LDAP directories. So, IBM will use the Domino directory to store their application information using LDAP. The applications can store special information about the user, and add that information to the user's Person document when the application is installed for use. We've written an API so a Person document can be automatically extended with subforms. You use Domino Designer to create the subform and then ship it with your application.

It sounds like the R5 directory can be the main LDAP source for the enterprise.

Yes, that's what a general-purpose directory has to provide. We've also added the ability for the Domino HTTP server to authenticate clients using a different directory. So you can set up an LDAP directory that contains information like passwords and X.509 certificates for a set of people who may not be Notes users. You can point our Directory Assistance database to reference that particular "foreign" LDAP directory for lookups. This solves the problem of people having to put names in the Domino directory when they're already elsewhere. We've also done the opposite; we've successfully enabled Netscape Enterprise Server to authenticate users with our directory through LDAP.

What is the R5 Directory Catalog?

People are very excited about this. You can store an optimized local copy of your enterprise-wide Domino Directories and have type-down addressing on very large sets of names, right from your laptop. Even if you're on an airplane! And it takes up a minimum amount of space. For example, the IBM worldwide directory, including Iris and Lotus -- about 350,000 entries -- can be stored in well under 50MB on your hard disk. You can look up someone's e-mail address (and whatever else you've configured the catalog to store) without ever connecting to your directory server. Not only have we enabled you to look up addresses for your entire enterprise at 37,000 feet, or from your back yard, but we've taken that one step further: you can even send them mail encrypted with their public key! You don't have to fill up your hard disk with 350,000 public keys at 1 to 4KB each just to do disconnected, secure mail. We do it using "just-in-time encryption." We look up the keys when you're transferring the messages to your server. I'd say this approach sets a new benchmark for scalability. I really think this means you can bank on Notes continuing to dominate disconnected mail client technology.

Which LDAP directories are you currently working with?

We test with Microsoft, Netscape, IBM, and Novell. Generally, if the directory supports LDAP, we'll be able to use it to authenticate users. You

also have the ability to put your LDAP server into our Directory Assistance database (formerly known as the Master Address Book) and have it be one of the servers that we will consult when trying to resolve queries. You put rules in the Directory Assistance database, and it will search for a user in the appropriate directory.

What are your plans for Microsoft's Windows NT 5.0 Active Directory?

LDAP is our integration point with Active Directory. Our strategy is to be able to use any directory, and we're already well on that path with this release by being able to do external authentication with foreign directories. But LDAP is just the first step towards the holy grail of a universal federated directory. The next steps are really to address the fact that all directories have different schemas, name spaces, and access control mechanisms, and in order to federate those directories, they need to be able to synchronize with each other. The whole industry acknowledges that it's going to be a while before those standards are nailed down. We'll play a big role in that. For example, LDAP doesn't specify a standard replication mechanism, so we are working with IBM and the rest of the industry through the IETF to help solve that problem by leveraging our years of experience in multi-master directory replication. From there, the next step is to de-couple applications and mail from caring what directory they are using so that customers can deploy the directory of their choice. Of course, we hope their choice will be the Domino Directory. That's what we mean by providing interchangeability, not just interoperability.

So, do you feel confident that you are going beyond supporting the Internet standards to value-added messaging, as promised?

Definitely. We're providing a unified, secure and integrated architecture to exploit the protocols. We obviously have a lot of customers using our products and we've done a lot of thinking and work to seamlessly bring their applications to the Internet for them. With R5, we've delivered on the promise of a world-class Internet mail system and we've continued to lead in this space with features like the Directory Catalog, just-in-time encryption, router controls/filters, synchronous mail agents, and mail rules. These are things that not only round out our interoperability story, but continue to show our commitment to being a leader in the industry.

BIOGRAPHY

John Banks-Binici is a software architect at Iris, and has worked on much of the design and implementation for the NTI networking layer in Notes. His background is in Email-enabled Calendar and Scheduling applications, and he's currently working on mail and directory issues for R5. He also plays guitar in Iris's rock band the "Iris Notes."

What do you
think about
this article?

Register
Here!

[Lotus Home](#) | [IBM Home](#) | [Iris Home](#) | [Feedback](#)
Copyright 1999 Iris Associates Inc.

