



RUNAS RADIO



<http://www.runasradio.com>



Richard
Campbell

RunAs Radio is a weekly Internet Audio Talk Show for IT Professionals working with Microsoft products. The full range of IT topics is covered from a Microsoft-centric viewpoint.



Greg
Hughes

Text Transcript of Show #139
(Transcription services provided by [PWOP Productions](#))



December , 2009



December , 2009

[Music]

Brandon Wenn: From runasradio.com, you're listening to RunAs Radio, the Internet audio talk show for IT professionals with Richard Campbell and Greg Hughes. This is Brandon Wenn, announcing show #139, with guest Bhargav Shukla, recorded Wednesday, December 2, 2009. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at pwop.com. You can follow the boys on Twitter at twitter.com/runasradio.

Richard Campbell: Thank you, Brandon. This is Richard Campbell. With me as always, my co-host, Greg Hughes.

Greg Hughes: Hey everybody, Richard.

Richard Campbell: Sir.

Greg Hughes: What's happening?

Richard Campbell: Not too much. Are you up for a RunAs?

Greg Hughes: Oh yeah, always up for a RunAs Radio. It's a good thing.

Richard Campbell: Good.

Greg Hughes: More Exchange topics to dive into.

Richard Campbell: You bet. I brought back our favorite Exchange person, Bhargav Shukla. He is a Premier Field Engineer at Microsoft. Always fascinated by computers, he is by education a civil engineer, which explains a lot really. In his 14 years of IT experience, he has had the opportunity to work on a multitude of technologies including messaging, directory services, information security, virtualization and more. He holds industry certifications from Microsoft, VMWare, RSA and Citrix, and when he's not working with enterprise clients, Bhargav contributes to the TechNet forums and blogs about Exchange at www.bhargavs.com. Welcome back, Bhargav.

Bhargav Shukla: Thank you very much, guys.

Richard Campbell: The Premiere Field Engineering, you have an interesting background on a lot of different things but the different areas you've worked in sound like they would give you some really broad and deep experience having to deal with a lot of different types of enterprise systems.

Bhargav Shukla: Yeah and being a premiere field engineer now I focus only on Exchange, that's my area of focus, but as you see before starting as premier field engineer I did work on a lot of other

technologies, and even being PFE working only on Exchange I still, when I work with clients, do see and get exposed to different technologies including working with Active Directory team, I work with customers, Active Directory gurus and talk to them when we have to discuss the Exchange updates, Schema updates and what's all in that. So yes, relatively less now but working on different technologies did give me a good benefit. When I work with customers, I see where they're coming from now, what their thought process could be when they ask a certain question. It definitely helps.

Greg Hughes: Certainly dealing with Exchange you have to touch so many different other things, you know, the big stuff like Active Directory, different types of storage, do you have virtualized environments here, there, maybe not for an Exchange Server per se although I think we've also have crises stuff being done with Exchange.

Bhargav Shukla: Correct.

Greg Hughes: But what a great background to bring to the table.

Bhargav Shukla: Yeah. Like I said, it definitely helps because when you're working with different technologies like Exchange, such as Active Directory, such as communications platform that you're using Microsoft, the OCS, BlackBerry is a favorite one for the enterprises. So it does help.

Richard Campbell: So where are we going to go in Exchange today, guys?

Bhargav Shukla: So I guess we can definitely talk about building upon the last part, guys. We can definitely talk about high availability. That's one of the very important features where Exchange 2010 brings a lot of advances, a lot of new features there.

Richard Campbell: I've been reading on the Exchange site this concept called Mailbox Resiliency.

Bhargav Shukla: Right.

Richard Campbell: That's a great phrase. What does it really mean?

Bhargav Shukla: Does it bring a lot of questions in your mind as well? What does that mean? Why is it mailbox resiliency?

Richard Campbell: What is that, yeah, exactly?

Greg Hughes: It sounds really cool.

Richard Campbell: It sounds good. I like it.

Bhargav Shukla: So if you start from 2003 and 2007; 2003 had this mailbox clustering as you like to



call it, the single copy clusters. It used the Windows clustering and had two servers, one active, one passive. It had the Exchange CNS running on top of it which had the mailboxes on top of it, databases, mailboxes, everything. In 2007, that was improved with using the CCR where now you're not using the shared copy which is SAN. You're not going to the same copy of your database. Now we have two separate copies whether it's on SAN, or whether it's on a direct path storage. Basically you have two copies. Now we have a cluster which provides you the service resiliency or fault tolerance. When service goes down on one server, your service keeps running on the other one, and then with CCR you had the ability to keep it running with a separate copy of database. So if your SAN has an issue on one set of disk and if you're still using the SAN with CCR, you would have the second set of disk working.

Richard Campbell: Right.

Bhargav Shukla: Again that's an assumption. And if use desk, well, obviously you have two separate desks subsystems so that just added a layer of resiliency now. In 2010, what we have done is we have taken it a step further now. You don't have that concept of CMS or CCR. You don't have the CCR, the CCR is gone.

Richard Campbell: Really.

Bhargav Shukla: Yes. So that shouldn't come to you as a surprise if you're an Exchange person and have been catching up on 2010. But if not, hey, CCR is no more and so is SCR, the secondary site copy, that's gone too and what it is replaced by is the Database Availability Groups.

Richard Campbell: Okay.

Bhargav Shukla: So the DAG, what it does is it works as the CCR measure, the one who compares the features. So the DAG basically breaks the components down to the database level so what you can do is you can implement a Database Availability Group across a bunch of mailbox servers and then take a database. So you can have five databases on a server but you can decide to have only a database copied to other servers. So you will have a copy of that given database that you've selected over multiple servers. Now, what that gives you is flexibility of deciding which database will be copied over to each server. You could have 10 servers and you could have one database copied only to three servers and that gives you the ability to have a second database replicated to the other three servers and so on.

Greg Hughes: Ah, okay.

Bhargav Shukla: It's one of the options. It doesn't have to be that way. There are many different ways you can cut it out.

Richard Campbell: Now, when you create this kind of redundancy, are you able to use those machines to improve performance as well? Because it seems like a lot of gear just for redundancy.

Bhargav Shukla: Well, and when you think about it, in CCR you have Server A and Server B. Server A was active primary than the Server B who was sitting there.

Richard Campbell: Right, active/passive.

Bhargav Shukla: With the Database Availability Groups, you have multiple databases. You don't have to have all your databases active on one server and then the other servers are just sitting there doing nothing. What the Database Availability Group now allows you to do is you could have Database One active on Server One, you could have Database Two active on Server 3. You could distribute your load across the servers and that's the beauty of it.

Richard Campbell: In each of those servers is then the redundancy server for one of the other ones.

Bhargav Shukla: Correct.

Richard Campbell: So this is similar to the way we do this in SQL Server as well where...

Greg Hughes: Right, right.

Richard Campbell: Server One has Database A and Server 2 has Database B and they are the failover for each other.

Bhargav Shukla: Right. Now I'm not sure how it works in SQL Server, but in Exchange so what you have is it is similar like you said but you could still have one Database Availability Group so you have this multiple mailbox around servers in the same group, however, you could still have your databases spread across different servers as an active one.

Greg Hughes: I can spread Database A across and have copies of that on multiple other servers, not just on one.

Bhargav Shukla: Right. So your Database One could be active on Server One but it could have a copy of it in Server 3, Server 5, and Server 10 so you have a total of four now.

Greg Hughes: Got you.

Richard Campbell: And those are complete copies?

Bhargav Shukla: They're complete copies.



Richard Campbell: And what's in those databases? Are they just mailboxes that are in that database?

Bhargav Shukla: Yes.

Richard Campbell: Okay. Which is the main thing; people want their mailboxes preserved.

Bhargav Shukla: Right and you have to remember now that these are the copies just like it used to be in C#. So you have a database. It's seeded for the first time, and then you have the transaction logs replayed as they fail. So you will always have one or two, a couple of databases depending on their structure. You would be lagging a little bit behind.

Richard Campbell: And how much is a little bit roughly?

Bhargav Shukla: If it is up to current and if it is given, let's assume that it's in the same datacenter and you're in a high speed links and you don't have any latency on the network side, then usually you would be one, to two databases behind -- sorry, one to two logs behind.

Richard Campbell: Okay. And how often does the log shipped?

Bhargav Shukla: As soon as the log is written. So you have one log that is always current.

Richard Campbell: Right.

Bhargav Shukla: And when that log is filled, it's written on the disk. It's shipped as soon as it's closed. So the service waits for the log to close and once the log is closed you're current logs stays locked, it cannot be shipped. But the one that just got closed would be shipped over to the other server.

Greg Hughes: And in Exchange, we have the ability to define how often those logs are rolled over.

Bhargav Shukla: Yes. Well, if you're talking about the logs itself, when the logs, as the transactions roll in, the log size now in Exchange 2010 and in 2007 is one meg so every one megabyte, it will create a new log.

Richard Campbell: Okay. So now that will totally depend on how busy your Exchange Server is and how fast...

Bhargav Shukla: How big your Exchange Servers are, right.

Richard Campbell: Right.

Greg Hughes: Right.

Richard Campbell: Okay. That makes sense. I think, for me, the distinction here is that adding more servers as part of that cluster doesn't slow down the site because you're not waiting to do synchronization across them before you complete transactions. It's all asynchronous.

Bhargav Shukla: Right. So if you think about it, this is not synchronous, this is not mirror.

Richard Campbell: Right, it's not.

Bhargav Shukla: It doesn't have to wait for the transaction to go. The primary server that has active database, it continues to write its log, it continues to function as it does.

Richard Campbell: Right.

Bhargav Shukla: And when the log is closed, that log gets shipped over to the other server.

Richard Campbell: Right. So this is just like SQL Server log shipping effectively.

Bhargav Shukla: Yes, it is log shipping so the functions are asynchronous.

Greg Hughes: In Exchange 2010, how does a failover work? What is it? What's going on behind the scenes and what does it look like?

Bhargav Shukla: So now here, the improvements help a lot. Let's take the failure at different levels. You have a failure on the server, the server entirely failed and you have two active databases on it. The Database Availability Group looks at it and says okay, the server has failed, where are the copies? The copies are given the preference numbers, so whichever is the next preferred number it would activate that database. So your Mailbox Database 1 has a copy on Server 2, Server 3, and Server 4 but Server 2 is given preference because it has number 2, then that will activate. Activation is usually pretty quick.

Greg Hughes: Well, and what does pretty quick look like?

Bhargav Shukla: Thirty seconds.

Richard Campbell: From the customer's point of view, in Outlook he's going to see lots of connections to Exchange and thirty seconds later regain the connection to Exchange and that's...

Bhargav Shukla: And this is where from the last time we had disclosed the changes in the client access server where now the clients connect to the client access server and not directly to the mailbox server. So what will happen is your connection to the



client access server will still stay alive, and the client access server has a component that basically will check the availability of your mailbox database and it says now that the server 1 is not available, where is the copy activated? Server 2. Okay, so I will now talk to Server 2. So for the user, it would become pretty transparent.

Richard Campbell: Yeah. Now, they're not going to see the lost connection, reestablish connection, because their connection is maintained all the time.

Bhargav Shukla: Right. Now you have to remember if you're using the Outlook web app, you would see an error if you were actively trying to send an email at that particular time. All you have to do is retry the operation.

Richard Campbell: Yeah.

Bhargav Shukla: Refresh your browser or hit the send button again.

Richard Campbell: Well, that's like the classic less than a minute threshold where people will be willing to hit refresh and it will just keep going. It will ignore the error.

Bhargav Shukla: Right. So you would not lose your written email in the Outlook web app that you have to draft again or like you said, just refresh the browser, hit the send again and it will send your mail away.

Richard Campbell: So Bhargav, I'm trying to think through this sort of progression of scaling the system for redundancy or reliability like this. So if I'm starting with a single server running Exchange, do I just duplicate that to start over redundancy or do I move the other roles off first?

Bhargav Shukla: So this is interesting. Now you brought two questions.

Richard Campbell: Yeah.

Bhargav Shukla: Let's address each one separately. So the first one is how do I build my cluster and if you remember previously you have to have all your servers in place, you have to have your Windows clustering configured before you can install your Exchange.

Richard Campbell: Right.

Bhargav Shukla: It then departs from that mentality. So what you can do now is you could have one server installed as your regular server running mailbox role and you could have a database on it now. Tomorrow, if you have finance and you have the ability to buy more servers, you bring in another

server and say, hey, you know what? I have two servers so why don't I just make it cluster?

Richard Campbell: Right.

Bhargav Shukla: Well, what you can do is you can just create a Database Availability Group now and then take your existing database and say I want to make it board off this DAG and I want to replicate it and do so into server. So it allows you to do that later on. You don't have to have it up from day one.

Richard Campbell: Right, because in 2007 I actually had to set up clustering first and then install Exchange and it did a special install of Exchange.

Bhargav Shukla: Right, right and that's what you don't have to do in 2010.

Greg Hughes: So are we saying that in 2010, clustering is sort of the default that's just enabled and you can either just have a single node cluster, or you can just add more as you go. That sounds like what you're saying.

Bhargav Shukla: Well, so what happens is when you install the mailbox on Exchange 2010 it doesn't have clustering enabled. It still relies on the operating system and clustering offered by Windows Serve 2008. But what it does is when you establish DAG, it will go ahead and create the cluster resources for you, and if you don't have the clustering role installed, as in 2008 you have to install different roles, it will install it for you if you are on that mailbox server where you're trying to create the DAG.

Greg Hughes: Got you. So what's happening is Exchange 2010 is making it easier for me and doing some of the work for me in the background.

Bhargav Shukla: Correct. It's doing a lot of that work for you in the background. Now, what you have to remember here is since you're still reliant on that clustering technology, you have to make sure that when you install your operating system it is the correct one which is enterprise and datacenter edition of Windows Server 2008. If you are on 2008 standard edition and then you later on decide to establish Database Availability Group, you won't be able to because it doesn't have the clustering feature set.

Richard Campbell: So at that point I'd probably – let me see if I can figure this out. So I'm on standard with Exchange, I decide I want to use a DAG, I have to stand up a new server with enterprise on it and then set Exchange up on that and migrate to that so I could take down the standard one and rebuild it as an enterprise machine.

Bhargav Shukla: Right, because since its standard, we do not have the upgrade support abilities. So we do not support that you just stay at



December , 2009

your Windows 2008 enterprise CD, stick it in and say upgrade my existing server. That's all it do running Exchange Server. You cannot do that.

Richard Campbell: Okay. Yeah, that makes sense.

Bhargav Shukla: So like you say, you will have to go to the other server, stand it up as an enterprise and then create a database, move your users off because it's the same Exchange, you will be able to do that.

Richard Campbell: Right.

Bhargav Shukla: Now, since you've brought up that point, it's a very good point we can discuss about and that is in 2010 when you move the mailboxes it's Live migration so users don't have to disconnect.

Richard Campbell: And this comes back to your original point that the users aren't connected to the mailbox anymore. They're connected to client access server so you can move mailbox servers around and the client access server will deal with it.

Bhargav Shukla: Right.

Richard Campbell: And the way the mailboxes are moved is different as well.

Bhargav Shukla: So if you're running a 2007 SP2 and then moving the mailboxes off to 2010, that experience is a Live one also for the user, whereas in previous versions, if a member user would get a message that they're either not connected to their mailbox or the administrator is making the changes and they will have to log off and log back on.

Richard Campbell: Well, and we've talked about this a little bit in the previous show where migration from 2003 to 2010 is just fine, it's not a big deal but it's an offline migration. In 2007 SP2, you can do an online migration.

Bhargav Shukla: Right.

Richard Campbell: Okay.

Bhargav Shukla: And when you say offline, relatively usually it is scheduled. So you schedule the users when they're not working, in nonworking hours, so it's still not impactful to users if it's planned properly but it is offline.

Richard Campbell: And it does need planning and it does mean there's a bunch of IT guys working at late hours of the night to make it happen.

Bhargav Shukla: Yup. With this, if you have 2007 SP2 and you're going to 2010, it actually gives the IT guys a relief that, hey, they don't have to work

long hours in the night. Well, that doesn't mean one of the managers won't say I'm not comfortable with this and you still have to work offline, but that is a non-technology issue.

Greg Hughes: Right.

Richard Campbell: Yeah. I need IT guys who actually love the fact that the work at night. They don't have to deal with all the stuff all the time, and then I'll also find folks anytime I could actually do something during the day, I'm happier.

Greg Hughes: Exactly. If it wasn't for the people, and the technology, and the software company, if it wasn't for all that, this job would be great. So I'm kind of curious. I'm going to go way back here and I know this has probably changed overtime, but I'm just curious what it looks like today or if it's even a factor anymore, but thinking way back to the Exchange 5.5, before you started putting 2000 and something on Exchange, and before 5.5, about the fragmentation of the mailbox stores and data stores, is that still an issue today or is that something that we need to think about in terms of the maintenance and does it impact our availability in performance?

Bhargav Shukla: You didn't have that until 2007 and what that means is in 2010 it's not. One of the biggest problems, if you remember, was that your back-ups were running and your fragmentation and online maintenance were running at the same time.

Greg Hughes: Right.

Bhargav Shukla: And that conflicted, the window conflicted, you had to resource issues and what-not. Not anymore. With 2010, your online defragmentation is running 24/7.

Richard Campbell: So the databases are constantly defragging themselves.

Bhargav Shukla: Right and you have to remember why. That's because it is enabled a lot because the storage changes. The way the data is written, we've talked about this briefly last time, it's the cluster I/O, larger byte sizes. A different way of writing the database is sequential instead of random, it gives an opportunity to do that. So what Exchange 2010 does is it runs the online defragmentation in background 24/7, and what it also does is it runs in the background as a lower priority process so it does not affect your primary task of the mailbox server off delivering the mails and doing whatever it is you do with your mails.

Richard Campbell: All right. Now, I'm trying to remember if we really answered the question of the progression of building redundancy. Can you simply have two servers that run all of the services, or does it



December , 2009

make more sense to separate the services into their own machines and then replicate them?

Bhargav Shukla: For that, the answer in the previous version was no. When you have a clustered mailbox server, you have to have other roles on different hardware. Not anymore. In 2010, you can have a planned access, a hub server, and a mailbox server role on the same server and still establish a Database Availability Group and create the database copies. Isn't that beautiful?

Richard Campbell: Yes, this is beautiful but then that back-up machine is only running the mailbox servers effectively.

Bhargav Shukla: You mean in terms of the roles?

Richard Campbell: Yeah.

Bhargav Shukla: You could have the roles on both servers.

Richard Campbell: Okay. So because in theory, if I'm really going to have reliability here, I want to have all of my servers replicated.

Bhargav Shukla: Uh-hmm, and what also happens is the way the hub server has handled the traffic, if you remember from 2007 days, is that when a mail comes in, the mail must go through the hub servers so if I deliver an email to a user who's also on the same database on the same mailbox server, the mail still goes through the hub server.

Richard Campbell: Okay.

Bhargav Shukla: And that's still true in 2010. So what's going to happen is if you're hosting a hub server on the same server as a mailbox server, the database and the transfer resiliency that is destroyed at 2010 will be gone if you don't have another hub server.

Richard Campbell: Yeah.

Greg Hughes: Sure.

Bhargav Shukla: What you will need is you will need another hub server somewhere and remember whether it's co-hosted at your mailbox server on a different hardware or be it a role of its own because what happens now is if you co-host your roles, Exchange 2010 by default will try to look for another hub server and say, hey, do I have another hub server in my site that I can deliver the mail to even if it's local because in that way the transport can keep a copy if your server goes down. The mail is on another server as a copy so when the database fails over it can be redelivered and you don't lose mail.

Richard Campbell: Interesting. Yeah, that's an interesting reflex.

Greg Hughes: So you're saying that part of the logic of 2010 is that it goes out and it looks for availability of services so that they can make its software redundant. Is that what you're saying?

Bhargav Shukla: Right and it does that at multiple layers. In your mailbox resiliency, you have the transport resiliency in form of shadow redundancy. So what it does is when it delivers the mail to the hub transfer, it waits for the next hub, and once it receives the confirmation that I have delivered it to the next hub, then it deletes it from its resiliency or from its shadow resiliency cache.

Richard Campbell: That's very clever.

Greg Hughes: That is cool.

Richard Campbell: And to me it also feels like having two servers running all of the services with the DAG would provide us pretty good redundancy although I mean you're only one machine away from having no redundancy but you're better off. Maybe my question earlier is is that a normal thing to do or do you normally see redundancy being six machines at a minimum?

Bhargav Shukla: Not necessarily. So what 2010 enables you to do is have lower cost storage because of the serial rights. You could easily use a set of drive and host your database on a set of drive. What we also have talked about a lot is that you could now have a RAIDless storage so you can just have a JBOD and host your databases on that. If you want to do that, then you would need more than two copies because what happens is like you say you're one server over. If one fails, then you have just one more copy. So at that point, you know, our support stand is that if you want to do a JBOD, if you want to do a RAIDless storage, make sure that you have three copies of the database. You have the three server DAG.

Richard Campbell: Well, this is new to me because my experience with Exchange has always been that you use pretty redundant reliable machines on this, RAID arrays, redundant power plus and so forth, you really don't want to lose an Exchange box even though you still have redundancy in as well. Now this is sounding more like a web farm where cheap commodity hardware with multiple replicas, and we don't care if we lose one.

Bhargav Shukla: Right and that's because, again go back to where we have come from, the failover perspective, how quick it is and why. So it is quicker in terms of taking over from the failed database whether it's just the database level or the server level



and then your client doesn't see it because they are coming from the guest. So that allows us to do that.

Richard Campbell: Right.

Greg Hughes: So we've addressed the redundancy component of high availability. What does Exchange 2010 have to aid us from the perspective of doing the bigger picture of high availability?

Bhargav Shukla: When we talk about disaster recovery.

Greg Hughes: Yeah and you know datacenters go on available or you have a disaster, that type of thing.

Bhargav Shukla: Right and so for that you used to have SCR in 2007 which I said earlier is gone as well with CCR so the Database Reliability Group now expands around across sites so what you do is you have additional servers on the other side whether it's a separate datacenter or separate datacenters. You could have multiple ones and then you could have different databases go to different places as your disaster recovery plan as a business allows you to do. So think of your DAG and you have two sites, Site A and Site B. So you go from Site A and Site B, and in Site B implement a couple more DAG members, your mailbox servers, and then now you have a single DAG stretching across sites. If your database fails, you could have another server in the same site taken over. If the whole site fails, you could have all your databases activated on the servers on the Site B.

Greg Hughes: Got you. One more question, and this maybe isn't quite high availability, but I'm kind of curious just based on having run into this a lot in the past and if 2010 does anything for us to help make this easier. So you have a lot of mergers and acquisition taking place in business these days where you have multiple companies that are running Exchange or maybe a big company that's constantly acquiring other companies that are running Exchange but they have to migrate all those users into the larger, say the big acquiring companies Exchange environment. Is there anything in 2010 that simplifies that or makes that more robust or takes less time or anything along those lines?

Bhargav Shukla: You mean the process of acquisition like bringing the users over from other form, migrating them across the entire org, is that what your question is?

Greg Hughes: Yeah, that type of thing, that's sort of what I'm talking about.

Bhargav Shukla: Yeah, because usually what you would be doing is you would have two separate farms and depending on the acquisition whether you

have Exchange on another site or not, it could get complicated.

Greg Hughes: Right.

Bhargav Shukla: I'm not sure if 2010 makes it any better in that area.

Greg Hughes: Okay. I'm just curious if there's anything that stood out.

Richard Campbell: Guys, we're getting close to the end of the show here. Bhargav, any stories that you might be able to tell us from your experience from the field of disaster recovery and Exchange? I know they probably wouldn't be able to love the 2010 once, but any experience you can talk about?

Bhargav Shukla: Well, the good thing is I did not have to be a part of a real disaster which is always a good thing, which means I was either a part of environments that were healthy and as a PFE that's one thing that's part of our job. We have to make sure that we have a healthy environment but again, disaster is disaster. No matter how healthy your environment is, disaster can strike anywhere.

Richard Campbell: Sure.

Bhargav Shukla: Luckily though, it didn't strike where we were. I was part of drills, disaster recovery drills, and that was really valuable too to look at your processors overall because when you think about disaster it's not just technology, it's your business decision making, logistics for your back-ups if you have to restore from back-ups, and the cooperation of other teams because you are reliant on other teams and other teams are reliant on you so a lot of disaster recovery actually is processed more so than the technology itself, and this is not my direct experience but one of the people I have talked to. What they had is they had a disaster and the planning before the disaster was pretty good. They had the second datacenter, they had everything in place, they had the processes documented. But where it can go wrong, this is a great example, is that the datacenter they were in where the disaster shock, it was just a one click so what they had done is that you just do this one click and your disaster recovery scenario will drop the rest and you will have your secondary disk interoperation and what-not. The primary disk center had a fire and the smoke was just was in there so they just basically quarantined the building. You cannot go in the building. Well, guess what? Your recovery scenario is gone. You cannot push that button to initiate the recovery sequence.

Richard Campbell: Right.

Bhargav Shukla: So like you say, it takes a lot of planning. It's more planning processes thinking through the processes and obstacles more so than



December , 2009

technology. I mean, if you think about 2010 and let's just say in 2010 your primary datacenter fails, you have the servers on the other site, great. What you did you forget? You forgot your namespace planning. What if your namespace that the user before you do was still pointing from the DNS to your primary datacenter that doesn't exist anymore?

Richard Campbell: Yeah.

Greg Hughes: Right, right.

Bhargav Shukla: So it's more planning, a lot of planning than anything else.

Richard Campbell: And just contending with the fact that what if you can't get access stability? What if there's a smoking hole in the ground where the building used to be.

Bhargav Shukla: Right.

Greg Hughes: Somebody blew up the button I have to push.

Richard Campbell: Yeah. I remember I was working through a scenario like that and realizing, jeez, are disaster recovery manuals only copies in the building? We should keep a copy somewhere else

Bhargav Shukla: We have a very good documentation but it's not here.

Richard Campbell: Yeah, it's burnt up too.

Bhargav Shukla: I just remember one thing about DAG that I would definitely like to mention for our listeners' benefit. It's that you have to keep one thing in mind when you talk about planning. It's that your Database Availability Groups are limited by one factor and that's your maximum number of servers you can have in a group. Your DAG is limited to 15 servers.

Greg Hughes: Fifty, five, zero or 15?

Richard Campbell: One, five.

Bhargav Shukla: So when you plan your DAG whether inside or with different sides, if it's stretched or not, make sure that you don't exceed that number.

Richard Campbell: Wow, and if you do exceed that number, what are you doing? That's a lot of servers.

Bhargav Shukla: Well, you're right. It's not always necessary that you don't have more users than the 15 servers can take, but at that point what you want to do is you want to have a new DAG. You can have multiple DAGs.

Greg Hughes: Ah, okay.

Bhargav Shukla: So just establish a new DAG.

Richard Campbell: Okay, that makes sense.

Greg Hughes: Makes sense, yeah.

Richard Campbell: Well, guys, I think that's the show.

Greg Hughes: Thanks a lot. That's really great information.

Bhargav Shukla: Thank you, guys. I really hope your listeners are enjoying and benefiting from the information we share.

Richard Campbell: Bhargav, thanks so much for coming on the show.

Bhargav Shukla: All right, thank you guys.

Richard Campbell: We'll talk to you next week on RunAs Radio.