

CHRIS BLUHM: So hi, everybody.

AUDIENCE: [INAUDIBLE]

CHRIS BLUHM: Hi. See, we got some engagement. I like that. All right. Welcome to restoring your vault, real-world tools for restoring your data. I'm Chris Bluhm. Traci Peurasaari will introduce herself more in just a moment. So thanks, everybody for coming. I know, speaking as one of the people who was here on Sunday disrupting all of your travel, we're almost done. AUs just about the end of this class. I'm going to have to go over here to get us to work. Or it just won't work at all. I promise I tested this.

Here we go. We're almost done. This class is the last mile. You all look a little bit better than I did on Sunday. And don't mind the watermark, because I wasn't paying \$60 for a photo that ugly. But we're in the last mile. Irvin's class next is going to be the finish line. But let's power through. Will all stay awake. We'll ignore the hot room and the lunch. And we'll learn some great things today. Let's keep going.

So this class is going to cover the strategies, possible pitfalls, and best practices of restoring your vault. So hopefully that's why you're here. Traci is my co-speaker. She's going to be going through a video later. She has 25 years of being a SQL DBA. She's the brains behind the operation. She works with me in enterprise priority support. And she is our SQL DBA. So any especially in-depth taking technical questions she's going to be able to help out a lot.

So as I said, again, my name is Chris Bluhm. I have been with Autodesk about two and 1/2 years. And before that, I had 10 years as an IT admin. That's where my focus really has been during my career. And I had a lot of experience doing DR and restore plans for large multimillion dollar firms. So I mean I guess not large, multibillion would be large. But it was large enough.

So at the end of this class you're going to be able to understand issues that need to be addressed in your recovery plans, what you're restoring and backing up, and what you need to do to plan to restore for your own servers. And then also I sent out an email to everybody asking what you wanted to learn. And everybody said replication, replication, and then more replication. So hopefully that's why you're here. And if not, just stay anyway. It's going to be a good time. I don't like this clicker.

So the first things first, before you're going to restore, you need to back up. We can't do any restore. So we're going to quickly go over the backup, and what you're going to have to have in order to successfully do that. So what do you need to back up before you can restore? There is going to be a couple of things. You can just go through the ADMS, that hopefully everybody's touched. But what are you doing? So as part of vault, there is for every vault installed there is your Knowledge Vault Master, the KVM database. So that's part of vault that keeps track of all your users, and that's the master settings for your vaults.

And then there's your vault database or databases. So I'm not going to point out somebody who's in charge of the product, but it's a little confusing, because we call the product Vault. We call the things inside the Vault vaults. But for every Vault that you create, whether it's going to be named after your offices, based on location, so Tulsa, Delaware, China, Berlin, or if you name them after your internal groups, engineering, or CAD, or anything else, each one of those is going to have its own separate database.

There's also your content center databases, and then the SQL Server system databases. So we recommend you back those up outside of Vault. If you went to Steve's class yesterday, if not you can look that up. It was maintaining your Vault and setting up maintenance plans. There were comments on that on why you want to back up your SQL databases. But it's going to assist in any restore you want to do.

There's also the Web.config file. Some of you may not know what that is. It's what we use to control all the settings for how your clients interact through IIS server to Vault. There may be times when support will edit things in there. So what size files you can upload, how many you can download at the same time, and that isn't saved inside of the Vault, because it's more of an IIS setting. So that's something you want to grab. The location of that is in the handout that I've put up for this class. And you'll be able to see it now or in two weeks. I believe it's also in this presentation. And again, you back that up, because if we ever changed it, you want to save that setting so you don't open a new support case and have to fix it again after a restore. And of course, the filestore-- I hope nobody is missing or forgetting that you need to back up the filestore.

So you want to back it up? Well, you can use are built in tool in the ADMS and just go through tools, backup and restore, and choose to do a backup there. You can use a third party. In the handout that's attached to this class I do have a white paper on what we suggest for third

party backups. And then if you do a third party backup, you want to back up all your SQL databases and your filestore. And again, as we said in the previous slide, make sure you get your SQL system databases and your web.config file. And last, schedule frequent full backups, weekly, bi-weekly, or daily. But if nothing else, make sure you have them.

And then, of course, it's not a backup unless you've tested it. So this is just my personal-- we get tons of calls of, oh, we checked the backup and the job says that it was running, and then you go look at the files and they're not there. I'm sure I'm preaching to the choir here, but just want to make sure we're covering that.

And then now it's time to restore. So what do you need to think about? What should you consider? Well, if this is a DR, or if you're doing it in a test environment, how powerful do your servers need to be? How many people are going to access this? Think about that. If it's a DR, what do you have ready to go? You don't want to be putting your production level servers that can handle 1,000 people accessing it on an old Pentium. You want to make sure that all the servers that you're restoring to are the same OS level, and service pack as before. You don't want to have your production on 2012 and a test or DR on 2008. You want to make sure you're maintaining consistent levels so that if there are any problems or issues, we're replicating them in the same environment.

Also, for upgrades or a disaster recovery scenario, it's really best if you can name the new server the same as your old server. It's not a requirement, but it's going to make your life a lot easier. In terms of worrying about changing the SQL server name, and the databases, and where it keeps track of, and the KVM that we mentioned before, it keeps track of the local server it's installed to. So it's a helpful thing that'll save you time, but not a requirement.

When it comes to SQL, just the same as your OS. You want to make sure that you're restoring to the same level. If you're doing a production in a test environment, or if you're in a DR and you're restoring to a new server, you always want to make sure your SQL that you're restoring to was a higher level than what you backed up from. An earlier version cannot restore from a newer version. Microsoft won't let you do that. You have to have something that knows about the past. It won't let you restore from the future.

Also, a tip that we do get relatively frequently, or that I've run into also, make sure that if you use a third party tool, and you're restoring your databases, the database owner has to be vaultSys, otherwise you're going to run into problems. If you use a third party tool, or anything

else that restores it, it may automatically set it to SA, or your user account, or something else. It has to be vaultSys. So the ADMS sets that, and you'll have if you do it separately.

Also, the order of your store is important. So I know everybody here is an expert, especially if it comes to building appliances, toys for kids for Christmas, or anything else. But before you start the restore go grab our read-mes. Look through them. Make sure that there's no gotchas, anything else you need to worry about. We do-- well I don't. But other people spend time on them. We try and make them accurate, and they're going to help you out before you go through a restore.

And then, it's not all about the Benjamins, it's all about the KVM. So without the KVM, you're going to have a bad time. The KVM, as I mentioned before, controls your users, and all the settings about the Vault. And without your users, you're not going to know who owns the files and who last access them within the Vault. And you wouldn't be using Vault if you didn't care about data management control. So you have to make sure that your KVM is at the same time level as all your other things that you're restoring.

If you're restoring a KVM-- if the only backup you have is your KVM from six months ago, that's when you need to restore everything else from. You can't say, I'm going to grab my KVM from six months ago, but my Vault database and my filestore from yesterday. While technically you might be able to shoehorn it in, you're going to have the problems, because the KVM is only going to know about users you had six months ago. And anybody who edited since then who own files, they don't exist. So you have files who have some mystery owner that the KVM doesn't know about. It's going to be a really, really bad time. So if I can say again, just pound you over the head with it, because we've had this happen before, where people have lost months of data, because that's the last KVM backup they have, if they're doing a third party tool or anything else. So make sure you're getting those.

And then of course, make sure that the Vault restore location is empty, and you have enough space to restore your filestore. And then just my own personal pet peeve, or at least my OCD, try and set your Vault filestore location something other than name of-- if you only have a single vault, the name of your vault, because when we do the restore, your filestore is going to be named the name your Vault. That sounds a little confusing and I know that.

But if you look at the screenshot I have you'll see under This PC, DATAPART, and my D drive, and so I set my Vault filestore location as vaultFS. I'm very imaginative, very creative. And we

have the Vault name as Your_Vault. So if I had set my filestore location as the name of my vault, it would be D, Your_Vault, and then another folder under that Your-Vault. So it's not the end of the world, but for anybody else who's a little too focused on things like that, and will get annoyed, it's going to be hard to change it, so just something to think about. You guys have got to deal with my OCD, sorry.

Another thing-- if you haven't had to do a restore, if you create multiple vaults, you cannot restore just one of them. So if you have Tulsa, Delaware, and Florida, and you decide that you only want to restore Florida's, we don't do that. They're all connected. They all have entries into the KVM. If you want to do a restore, you're replacing everything. So it's a question we get frequently of, oh, our one office screwed up a whole bunch of stuff, can we just fix only them, and restore them back to last week. No, sorry. I know it would be really great. I don't think it's ever going to change with the way they're all incorporated in KVM. But that's just something you need to be aware of and think about if you're creating numerous vaults and telling people how to use them. And we do warn you when you go to restore that you're going to blow everything.

So things to think about for your environment. It's almost like I've said this, but make sure they return to the same version or newer. Vault isn't going to let you restore a 2016 restore to 2014. We've covered that. And then one of the questions we get to is, how do I know what version? My server died, I don't remember what version my Vault was. What do I need to do? How can I find that out? Well, we're going to go over a little tip that you can do. So when you create your backup set through the ADMS, you're going to get four folders.

I'll let you take a picture. This PowerPoint will be available. I think it's usually about two weeks turn around, maybe a little bit more. Ha ha. Yeah, sorry. The classes will be available. The PowerPoint's available immediately after this class. But Irvin's here to help me.

So when you create your backup it's going to have a databases folder, a filestores folder, backup contents, and backup history. So that little XML on file, if you crack it open, it's going to list what product the back it is from. So here it's going to tell you, you backed this up on version 21.1.4.0 You say, that's great. Where do I download version 21.4.0, because I don't know what the heck that is. Well, we can help you too. If you then go and take that version number and go to the knowledge base, we have a list of all the build numbers, and what you need to get. So granted, we do make you hop through six different places. But I'm sure that could be something that you guys could submit to the idea station, which everybody uses all

the time. Yes? If not, we're going to tell you about it at the end.

And so here you'd see that 21.1.4 2016 SP1. So that's what you'll know you have to set up. And you'll download 2016 and grab the service pack, and then you'll be able to restore it without getting a error. If you don't do all that, all we do is tell you this version is incompatible with your restore. We don't tell you necessarily why or what versions. I'm sure there is very good logic for that. But it can be confusing. So that's the tip you can use. You can also go look, I think, if it the restore fails, it will tell you in the log that it wasn't compatible there also.

So the reason why, at least, I'm sure half of you came-- what to do in replicated environments. So if you just have a AVFS environment-- pretty easy right? You've got your primary central ADMS server, and a whole bunch of AVFS servers. This is really popular since we introduced. It cuts down on the license costs for SQL. And the main thing you have to do there is that all your remote sites, if you want to make your life easier, back at the filestore. So you want to make sure you have a backup locally at that site of the filestore, so if the server crashes and burns, you have the filestore ready to go.

And then you just reinstall the AVFS software. And when you install it, it'll validate the filestore. So you don't have to copy everything from the ADMS down, you can say, all right, I have my ADMS, I have an AVFS that I just installed, and I have the filestore from yesterday. And then as soon as you tell them to sync, and enable it, the ADMS will go through and compare each file that the AVFS has locally to what it should be, and validate that each one is the correct version. Any new versions will be replaced, but you don't have to go through and wait for it to sync over the WAN.

Now an important note about that, and this is something that you do need to think about for upgrades or anything else, is if you have, let's say, 10 AVFS servers, but only one ADMS, only one AVFS can validate at a time. So if you're doing an upgrade, and they all have to re-validate, it shouldn't be a bad upgrade. But if somehow five sites die, you're going to have to wait, and each one takes an hour, you're going to have to wait five hours. It is not a concurrent process. It is a single process at a time. They're going to queue up and each AVFS will go in. So just a little thing to think about in terms of any recovery plans or upgrades that you're going to plan for yourself, remember this, and build in that extra time.

To pre-answer answer any questions, no, we can't tell you how long it takes to validate. I'm sorry. You can test it yourselves. But we don't have any tool to go through and warn you about

that, to give you an estimate of time.

A multisite SQL environment-- the biggest complaint and biggest issue we hear from people, is it takes too long for me to set up a remote site over the WAN. How can I make that faster? Well, there's two things, your SQL database replication, and your filestore replication, that we know are being affected by that WAN speed. So for the filestore replication, there's not a whole lot we can do to help you. You can limit what folders are synced. I hope everybody knows that. When you go to set up the remote site, by default, everything is checked. You can go through and you can say, you know what, this site really only needs designs for the widgets. That doesn't mean they can't get everything else, but that's just what can be pre-seeded at that site.

So if they need to get other files, when they go through the Vault client, and say, I want to work on this, it's going to warn them, your file has been copied to the local site. Do you want a copy now? They say, yes. It gets brought down, and then it exists at a local site in perpetuity until it gets deleted. The other way, if you need to set something up, and we've had plenty of customers do this, is they copy all the files to the hard drive, and fly it out there. And then again, just as we talked about before, if you get all the files there, once you start a replicated ADMS or AVFS, it's going to validate those files and make sure everything's OK. So it's perfectly fine to throw that filestore on a hard drive and ship it out.

And then also, what you can do is create a-- this is a SQL issue. And it's in the handout. It's called an alternate snapshot method of restoring or setting up a subscriber. So if any of you have set up a replicated workgroup within the ADMS and Vault, you're going to know at one point it says, what is the replication folder you want to use for SQL. It looks like this. So it's going to say specify the location of your replication shared folder. You shouldn't name it your vault, because then you'll be really confused. But I just threw that in there for the purpose of this.

So during the group replication set up, it asks you for this folder. And that's important, because in that folder that's actually where-- again this is all a SQL process-- SQL stores the snapshot of the database that it's going to use to seed any subscribers. It stores the snapshot of the publisher to seed all the subscribers. And there's a job that refreshes that every few days. I think it's by default every three days. So if you set up five subscribers, they're all getting that same seed. And that's how they're going to build up and then sync any deltas from the last time that seed was created.

So the important part is you go through and you set up workgroup replication. It creates that replicated folder and puts the snapshot in there. And then you'll set up your subscriber server that you want to set and you want it to be copy too. But let's say it's in Timbuktu. Why not? Really far away-- and you've got a 1 mg connection to. That's going to be terrible. So how can we make your life a little easier? But you go, you set the publisher, you set the subscriber, and it's created the snapshot.

And now you need to go actually disable the replication jobs. You need to stop them, so it stops trying to replicate. So you're going to do that. You're going to disable the snapshot creation on the publisher, because you don't want a new snapshot being created. SQL will only respect the most recent version of the snapshot. If you get a snapshot from three weeks ago, it won't accept it, because the delta is too big. You want the most recent copy that snapshot. And then you want to disable the replication job on the subscriber. And the reason you need to disable that job on the subscriber is-- sorry, I skipped ahead of myself. I apologize.

So the snapshot is created in that replication folder. And it's called-- it's very helpful. It's just unc. And then it'll have the name of the folder underneath that and the snapshot. So you want to grab that unc folder. And hear the big benefit is you can compress that. So you can zip it down. You can FTP that. So SQL is really good. It's a transactional database relational. It wants to make sure all the data it communicates between the two are perfect.

But I'd also say, when it comes to talking to each other, it's sort of like trying to play whisper down the lane, and describe entire chapters of *War and Peace* to each other through whisper it on the lane. If anything gets missed, if one sentence gets messed up, SQL is going to throw it all away, and then bring it back. For anyone here who's actually a network head and knows about topologies, and everything else, I know you're cringing, but this is a good analogy.

So what you can do if you do a Windows file copy, or FTP, or anything else, it's going to allow those little errors and instead of saying, you know what, you just told me an entire paragraph, retell it to me, it'll say, I heard the beginning and the end, but I missed one sentence in the middle, so FTP windows file copy is even going to be a lot faster than how SQL replicates. So you can zip those snapshots up, you can send them through any other method that's going to be a lot faster, or again, and we've had customers do this, throw them on the hard drive, overnight it, ship it, do you anything you need to get it over there.

And then you've disabled the job on the replication folder. And you'll go back into the SQL Agent jobs, which if everybody has touched SQL, hopefully you know. If not, it's going to show your databases, and then it'll say SQL Agent, and they're going to be jobs underneath there. And you've got that alternate snapshot, and you'll go to your job properties, your replication merge job, and at the end you're going to edit what it's running. And you're going to say, alternate snapshot folder.

So this alternate snapshot, again, is Microsoft's term, a lot of people get confused by it, but that's the reason they call it that. And you'll say, what folder do I want to point to? Where have you copied that unc folder that we pulled from where the publisher has created, and you'll put it in here. And then you re-enable the job, and you can see the replication status. And it's going to sync through like orders of magnitude faster than going through from the publisher. Maybe it'll take 5% or 10% of the time to completely see the database, and go through. And then after, it's finished seeding and after the replication status says that it's finished, you're going to go back into the job and switch it to continuous. And the continuous command is what by default it is, but that means always be checking between your publisher and your subscriber to make sure that you have the most recent copy of the data.

So now Tracy, I know we've put a live demo-- sorry, we lied to you a little bit, because it was hard to fit that in. So we've got a video that Tracy is going to cue up, and I will switch over, that she's going to walk you through on how you could also restore a subscriber through SQL. Tell me when you're ready.

TRACI Go. I'm ready.

PEURASAARI:

CHRIS BLUHM: After you log in.

TRACI I guess I do have to do have to do that. A couple of notes-- this is specific to the subscriber.

PEURASAARI: However, if you are doing SQL for all of your backups, the beginning would also hold true if you were going to restore a publisher.

CHRIS BLUHM: You ready?

TRACI Yeah. So when you are restoring a subscriber, if you are not trying to do it from a snapshot, or alternate snapshot, or CD, and you've got backups of your subscriber, Chris talked about using ADMS to do a backup and restore. If any of you've noticed, you don't have that option

on a subscriber. You can't do an ADMS backup.

So the first thing you want to do when you need to restore subscriber, is as Chris mentioned, stop the jobs on the publisher, because you can't publish to a subscriber, or a partially existing subscriber. Then you're going to go over to your subscriber, and you would have needed to have installed SQL. Now Chris said that you can restore to an earlier version of SQL, but in this case you cannot, because you need to restore master, and it has to be the same version. Master from 2008 forward, will only let you restore to the exact same version of SQL Server.

So in order to restore the master database, you need to do it in single user mode. So if you notice, I'm opening a command prompt, and I am going to stop the service, although you could do that in services. And then I'm going to restart SQL Server in single user mode, and also tell it that I'm going to run SQL in command mode. So the single user mode switch is a slash m. And then SQL command in quotes, and it is case sensitive-- is what tells it that I'm going to run it in command line mode. The reason that you need to restore master when you're restoring the subscriber, if you're restoring it from SQL Server, is that is what contains your users, and the database structure.

And as Chris mentioned earlier in his, you want vaultSys to exist. You want vaultSys to be the owner of the database, and you want your SQL users all to be restored. And all that's held in master. So you start the database in single user mode telling it you're going to run SQL command. I'm running SQL command here. And now I'm going to tell it to restore master. And the command is just restore database master from disk equal and then in quotes the location of wherever you have your master backup. So you're going to do that from here.

And then we're also, after that, going to restore the MSDB database from the command line. The reason we need to restore the MSDB database from the command line is because as soon as you start up Enterprise Management Studio, MSDB starts being used, and it'll tell you the database is in use. So once master is restored, I have to restart the services. Again, you can do it from here. I did it from here, because I'm running in a command line. You can also, however you're comfortable, you can go into services, you can start your SQL Server service. And it's just a preference. So you restart SQL, and then it's going to be the same command for MSDB that it was for master, restore database MSDB from disk equal, and then wherever you had the backup.

Once you've gotten your MSDB and your master restored, from there it's a preference. You

can continue to restore all your databases from here if you're comfortable. A lot of people don't mind using SQL command. However, at this point, the next step for me is to go into SQL Enterprise Management Studio and to continue restoring. As Chris mentioned, you need to restore everything together. So from here you need to restore all your vaults, your KVM-- a note, and it is in the white paper. If you are using SQL to back up your publisher and your subscriber, you need to make sure that the master and the MSDB are backed up at the same time on the subscriber.

And on the publisher, you need to make sure that the master, the MSDB, and in the distribution database, which is called Autodesk Distribution, I think-- and its under System databases. And those need to be backed up at the same time, because time is critical. Just like the KVM, they care that they're all together. So here is Management Studio. The first time I'm going to go in and I'm just going to say, OK, I'm going to use the GUI. And if you have one Vault or a KVM, this is fine. Work's easy. I'm going to pick a database. I'm going to say, hey, I'm restoring AOTC admin. And if you want to fast forward past this a little bit, this is just using the GUI.

CHRIS BLUHM: And they're in recovery pending because they don't exist right now?

TRACI If you see up there, see how it says recovery pending-- I'm sorry, I did need to mention that.

PEURASAARI: That's because the master has put them out there and said, hey, the structure exists, but there's no data in them. So thank you, Chris. So it's just going to drill down, find the backup, and tell it to restore-- under Options, Overwrite, existing, and then restore. Now for my preference, I actually have this all scripted out. You can use Transact SQL. You can use the exact same commands that I used in the command line.

So I have all the rest of the databases scripted out. It's restore database xyz from disk, equal, with replace, and I'm commenting out the one that I just did manually. So I mean if you want to pre-script this, you can type it ahead of time. It's a lot easier than just doing one at a time, obviously, manually Go through and restore the rest of the remaining databases from the KVM and through all the vaults. It will tell you on the bottom that they all restore. This one goes quick. I didn't--

But from there you are back at the same point you were when your system went down. And if you're doing regular backups, you're good to go. If you, at this point, start your SQL services, because you restored MSDB, all your replication jobs are now restored. And they will be

running and ready to go. So replication will just kick right back in where it was when your system went down.

So the subscriber will be pulling, it'll be grabbing from the publisher. And then you need to go back to your publisher, and restart the jobs. However, some of the things to consider are has the subscriber been down a little bit longer? If it's been down for a length of time, you may not want replication to just kick right back in. You may want to grab a new snapshot from the publisher. So in this case, you wouldn't start the jobs up right away. You'd want to tell the publisher to re-initialize the snapshot, so that when the subscriber comes up it refreshes all the databases on the subscriber.

And again, you can use the same method that Chris referred to with creating an alternate snapshot, so that you're not trying to do this brand new snapshot all the way across your WAN for every single database at the location of the subscriber. And again, that really is going to be dependent on when did your subscriber go down, when was your last backup, how long before you tried to restore it. So those are decisions that you make from there.

And then as Chris said, your filestore has to be backed up or restored the same time as your databases. So in this case I just had a local copy copied someplace else on my system. I'm restoring my filestore. It has to be restored in to the same location that it was originally from. So I'm putting the filestore back in place here under filestore, copying the filestore back. If you don't, it will also try to replicate the entire filestore from the publisher. And again, that's not a short process.

From here, I'm just going to show you the jobs that are running. You can go in and check for yourself. You can take a look at the job. You can look at the history. You can verify that replication is running. You can see where they're at in the process, if they're done, if replication is happening from this window here. So if you're going to View history, you can see the errors where I broke it in the beginning. And you can see the most recent job where it's running.

And then on the bottom here you can see what the status is. You can see where it's replicating from, who it's talking to, and it also makes reference to the replication monitor. If you want to look at the replication monitor, that exists on the publisher. So if you go to the publisher, and you right mouse click from within SQL on replication, it gives you an option to look at the replication monitor. You can see all the jobs happening. So here I'm restarting the jobs on the

publisher, so that everything starts getting in sync.

One other thing, while your jobs are down, just a couple other options-- here's marking for re-initialization if you want to. Tell it, hey, restart all the snapshots. Once you do that in SQL it tells SQL that you can tell it to create a new snapshot, use the existing snapshot, tell SQL that you don't want to just do an update from the publisher. It also keeps the subscriber from pushing any data that the subscriber had up to the publisher.

And again, depending on when your system went down, you may not want that to happen, or you may. It really is totally dependent on when your system had issues. And then here I've reopened Vault. You can now see both the subscriber and the publisher. They both exist. I'm telling it to replicate the site now. And it's good to go. With This is the subscriber here.

When you go into Management Studio from this point, you can expand SQL again, and you will see the list of subscribers in SQL. So now you see they're all there. They're all now being replicated. If you want to look at any given one, you can see. And this is interesting. Chris was talking about how long it takes across the network as opposed to an alternate snapshot. See this, you can see how it's saying populating table xyz. If this is happening across the network, you can literally see this, and you can open it's 10 minutes later, and it'll still say populating network.

Right now when I was doing this my computers were side by side. But if you're doing this across the network you can open this, and you can see it's on the same table for quite a while. And that's back to what Chris was referring to, as opposed to shipping it if you're going across a slow connection or a long distance. Here you can just see where it's now propagating all the tables and syncing them up. This one shows interestingly enough, the KVM-- 60 seconds. It's already done. The KVM has already been synced. It's all ready to go. At this point, you can either install Vault, if you haven't yet. Or if you've already installed Vault, you can open Vault on the subscriber, and you'll be able to see the subscriber and the publisher on both systems from within Vault.

Important thing to note, once you've done this, your SQL, if you're doing backup scheduled, MSDB should have also restored your SQL backup jobs. But you do want to make sure that they are back, enabled and running again, if you don't want to stop your backups. You want to make sure your backups are happening again, and all that is back in place. So beyond that, I think I do pop in showing the Vault, but I don't think we need to go into that. Mr. Chris, you

can't stop it. You can actually show your ADMS.

CHRIS BLUHM: Do you want me to stop it.

TRACI Yeah. Go ahead.

PEURASAARI:

CHRIS BLUHM: But I will say just a note-- not on this one. But for everybody playing at home, you do know that the subscription is working when it says waiting 60 seconds for the next update. So I tried to leave a fair amount of time for questions. I figure we've covered a lot. I thought we'd probably get a few. So any questions, comments? Yes sir.

AUDIENCE: [INAUDIBLE]

TRACI Your subscriber is a pull.

PEURASAARI:

AUDIENCE: [INAUDIBLE]

TRACI The subscriber will update the publisher as well. But it doesn't update your other subscribers.

PEURASAARI: So subscriber to publisher, subscriber to publisher. It is to a subscriber to publisher.

AUDIENCE: [INAUDIBLE]

TRACI Right. Not subscriber to other subscriber.

PEURASAARI:

CHRIS BLUHM: The question was is does the traffic go between publishers and subscribers? Is it just one way? Or is it two way? And then also if you have a hub and spoke, does one subscriber talk to all the other portions of the wheel? No, so it's two way where the databases talk back to each other and update. So if you check in a file on your subscriber, and you save metadata there, it's going to update that, and send it up, and share it, and you're going to see those two sharing. But no subscribers is ever going to update another subscriber. That's what the publisher is going to be.

You still always have the ownership. Only one site is going to have ownership of a file at any given time, and that gets updated through the metadata through the vaults. And importantly, that's only for an ADMS replicated workgroup. If you have a single ADMS, and many AVFSes, when you're checking in the file at your AVFS site, it's actually sending that metadata to the

ADMS. That's why if you ever lose the connection in an AVFS site, you practically can't do anything.

TRACI
PEURASAARI: But the next time the publisher updates the other subscribers, that metadata should go to the other subscribers. As Chris said, they don't go.

CHRIS BLUHM: So if you have Washington publisher, Florida subscriber, and Texas subscriber, Florida updates Washington, and then Washington will update Texas. So if it's happening every two minutes, you're going to have a three minute window where Texas wouldn't know about it, up to a three minute. And then you'll have to wait for the metadata to copy down. And then Texas will have that data locally for it to view. But again, you don't have to worry about anybody messing anything up, because while Florida has that, it owns the file, so Texas can't touch it. And if Texas tried to get it, then it's going to get warned, Florida owns that file. Anything else?

AUDIENCE: [INAUDIBLE] the example was if a subscriber crashes and you are restoring this subscriber. So you said that it could work either way, this example could work [INAUDIBLE].

TRACI
PEURASAARI: You would use the same method if you had SQL backups on publisher. It's a little bit more difficult to update the data. If you've done a recent sync, you can get as much data as exists on the subscriber to the publisher. But you do have the opportunity to lose more data.

AUDIENCE: That's what I was getting at. When you restore the publisher, then what happens to that data that's sitting on that subscriber? Is there just a delta [INAUDIBLE]?

TRACI
PEURASAARI: You can do a sync. It will look and see if there's any differences that it can update. But the publisher is still the boss per se. So if he can't reconcile it--

AUDIENCE: [INAUDIBLE]

TRACI
PEURASAARI: Well, SQL has a bunch of rules to determine the newness, what happened first, what happened last. And when you do backups and restores to the publisher, it actually recommends that you run a sync first.

CHRIS BLUHM: So there's the retention period within SQL, that should be 14 days that it's going to keep that data on the subscribers. So that's what you have.

AUDIENCE: So I'm just trying to figure out whether or not when you do the restore if it's a delta like sync or if you have to totally throw that whole [INAUDIBLE]?

TRACI You'd have whatever was already in the backup anyway. I mean, you're going to have the majority of it.

CHRIS BLUHM: So it'd be the delta.

TRACI And as Chris mentioned, if it's within the retention period, 15 days, that data's still valid on the subscriber.

CHRIS BLUHM: Yes, sir.

AUDIENCE: [INAUDIBLE]

TRACI I'm not touching that one.

PEURASAARI:

CHRIS BLUHM: Especially with the man next to me, I would say you need to treat your publisher as the golden main source. So that should be your primary site. And that's why we threw the ADMS-- So the question was, is it OK to plan to use your subscriber to restore the publisher if it dies.

[LAUGHTER]

So the answer is, there's a reason within ADMS that we only let you back up and restore at the publisher level. That's what is recommended. That's what you should do. While technically possible, what Tracy is describing is going through the publisher. I wouldn't say anybody should put that into their plan.

TRACI I said, publisher is the boss.

PEURASAARI:

CHRIS BLUHM: Yeah, publisher is the boss. And that's what should be the golden source. That's why we focused on it. If you didn't want to use the alternate snapshot method to restore subscriber, you could do this if you had all the SQL backups including master, MSDM, et cetera. Yeah, sorry.

AUDIENCE: [INAUDIBLE]

CHRIS BLUHM: Yeah, so the question was, if you're restoring a publisher, does it matter if it's the same computer name? It doesn't matter. But as I said, it's more steps. So if you restore SQL databases, if you're using a third party tool, you're going to have to go into the SQL database

to the master database and rename what it thinks what server it's on. And then from Vault you'd have to do a restore or a detach your Vault and reattach them to point them to the different server name.

And so there are settings in the web.config, and during the install that go to the knowledge vault master databases, and other steps, that have that server name. So you're going to have to do a re-install of Vault, and give it the new server name. So that keeps it locally. Or if you haven't installed Vault, detach the vaults inside of it, or reattach them so that they know the new SQL Server name.

TRACI
PEURASAARI: Because SQL will install, but the underlying metadata inside master will still think that it's the old server name until you run the stored procedure telling it that you want to rename it and give it a new internal server name. So it's the same thing.

AUDIENCE: In addition to that, if you're [INAUDIBLE]. Your subscribers know the old name. [INAUDIBLE]

TRACI
PEURASAARI: I also wanted to mention something. When you're restoring the subscriber or just in general with replication to the subscriber, for those of you who might be SQL people, and you are comfortable restoring replicated environments, there is something that you have to remember. SQL tracks replication. So does the KVM. So if you tried to restore this, just like you would, oh, this is a replicated environment, I'm going to restore a replicated environment. It won't work, because the KVM doesn't get it.

On the flip side, if you just try to restore it from the Vault side, and then try to restore it from the SQL side, SQL won't get it. You have to do it together, in the right order, because the KVM keeps track inside Vault that it's replicated, and it uses SQL inside SQL keep track that it's replicated. They both have to be in sync together or it won't work. You'll try to start up your Vault, and it'll say, hey, wait a minute, I'm supposed to be replicated. What's going on? Even if you set it up in SQL, or vice versa. SQL will start up and say I'm supposed to be replicated, what's going on? And KVM is like, huh, oh, huh. So you literally do have to make sure that you get both pieces setup back in sync.

CHRIS BLUHM: Follow up.

AUDIENCE: [INAUDIBLE]

TRACI
PEURASAARI: Test, test, test.

AUDIENCE: [INAUDIBLE]

CHRIS BLUHM: I think he deserves some swag.

TRACI I think so too. But actually, I mean from not only from Autodesk, just from a DR perspective,
PEURASAARI: that's just a nice--

CHRIS BLUHM: Can I interest you in a Death Star, a PokeBall? You can have the working one. It'll open.

[LAUGHTER]

I just caught you. That was the whole thing. Yeah. If it was your class, you wouldn't have that.

So then the only other thing I'll say in terms of questions, it was a question that was emailed, just to stomp on everybody's dreams, especially if you've gone through all this. There was a question, is there any way you can take a production subscriber, and in the same location switch that to a test subscriber? No. You could use the alternate snapshot method, and use the same filestore to set that up. But if you have a production environment, as Irvin said, because of the SQL Server, and the names, they know what names they're syncing to. The most you can do is really take your test environment, and take your test publisher, do an alternate snapshot, and seed your tests of subscribers that way.

AUDIENCE: [INAUDIBLE]

TRACI Build a new publisher?

PEURASAARI:

AUDIENCE: Take the publisher, back it up to a development box, and then let it be stand alone to play in that environment without replication?

TRACI Yeah, you just have to make sure you turn it off in Vault, because when you first restored it, it
PEURASAARI: would have the workgroups. Yeah, you would just have to break the workgroups off.

AUDIENCE: If you're the subscriber in it then you can't do that.

CHRIS BLUHM: Yeah, you're not going to copy a subscriber-- it's good to be the king. All right, any other questions? OK. You're not allowed to leave yet. You can pack up though.

So this is my tip, please fill out your surveys. Or as I like to say, it's not a survey, it's a quiz.

The answer is 5.

[LAUGHTER]

But the surveys are really great. We get the feedback. You're going to tell us what we can improve on. It's going to improve other classes. It's going to improve future AUs. It's a big thing for us. We really appreciate it if you can do the feedback.

Resources after today-- you can go read *Under the Hood* blog. Irvin and a whole lot other people publish there. We put on tips and tricks for Vault. It's right here, so I can trip over it. And then if you have an idea, go to the Idea Station. Irvin and the development group constantly watch the Idea Station. And quarterly they're going to look at all the ideas, and talk about what we can bring into the product.

You're also going to know once you submitted an idea, it gets a status, if it's under review, if it's under consideration, if it's going to be implemented in a further product. This is the number one way that we're going to take your ideas or your suggestions for what needs to be put into Vault, and get it into the product. So the more upvotes that you give out, or the more upvotes that your idea gets, the more likely it is that it's going to show up in the future.

So if you come to these things for free toys, I'm not going to say bribe, free toys, but maybe a gentle reminder to also fill out your survey. After this, feel free to come up grab dinosaurs, Death Stars, giant PokeBalls.

TRACI Anybody who passes the survey gets the giant-- pass the quiz.

PEURASAARI:

CHRIS BLUHM: The most favorite will be our Autodesk egg. So I 3D printed all of these using our really cool MakerBot-- so future of making things, yeah, yeah, sure, sure. Thank for coming. I hope you guys had a great time and learned something, and that's all we got. So grab a toy if you want.

[APPLAUSE]

TRACI When the stuff from the class gets published, if anybody's interested, I will publish commands
PEURASAARI: that were used to do the restore in the single. I'm not going to publish the recording, because it's too hard to follow.

AUDIENCE: [INAUDIBLE]

TRACI

Yeah. But I will publish the commands that you use to do the single server mode and the

PEURASAARI:

restore. I'll just put it in a text file. It should be out there as soon as you walk out of here.