

STEPHEN Good afternoon, everybody. Thank you for taking the time to sign up for this class. Thank you
PORATH: to those of you that responded to my inquiry about a week and a half ago. I greatly appreciated that.

This class is designed for those who are either responsible for administering their SQL Server databases, or influential for those who are not doing direct administration. So just in case you were curious if you're in the right spot or not.

So we'll start with letting Traci introduce herself. Go ahead, Traci.

TRACI Hi, everybody. My name's Traci Peurasaari. I have been a SQL/DBA for over 25 years, and an
PEURASAARI: Oracle DBA for approximately 15. So SQL is my world. I've been with Autodesk for a year in Enterprise Priority Support-- mostly with Vault-- but really supporting databases in general.

STEPHEN All right. And I am Stephen Porath. Call me Steve as well. I've been with Autodesk just shy of
PORATH: three years. I've got about 15 years doing enterprise software support-- Various platforms, specifically SQL Server-based, Microsoft .NET applications that run typically with IIS as well. Traci and I worked together previously at another company, and they were going through a transition between Oracle and SQL as their database platform. So that was where I was introduced to maintenance plans that we'll be covering here, as Oracle database maintenance is very, very different from SQL database maintenance.

So some considerations for this. This is applicable to all recent versions of SQL Server. It is not necessarily for Vault only. It will be kind of spoken to towards Vault users, but you could be using any of the Autodesk applications that run on SQL Server, or anything that runs on SQL Server. And this can be beneficial. This uses Ola Hallengren's base maintenance plan. This is pretty much like the industry standard maintenance plan. Towards the end here there's some links showing you where to get the maintenance plan, and then some supporting information as well-- Microsoft and some others.

These slides are available online. I have a link to that at the end as well. And then we'll be putting together some of the supplemental information depending on the questions that come out of here. If you do have a question on something worth speaking to directly, feel free to let us know. We'll try to answer it. If it's something that's going to take a bit longer, we'll try to follow back up at the end of class. So what we'll be covering here is what a database

maintenance plan is, where you can find one, why it is important, and then how to implement one as well.

So we'll start with the "what." This is Microsoft's definition for a maintenance plan. Think of it as a list of tasks-- or a workflow-- that are required to be executed against a database environment to keep it stable, backed up, and free of any kind of inconsistencies. And that's done through integrity checking. There's a note towards the end of their definition. You can set up a maintenance plan as a part of the SQL Server Maintenance Plan Wizard. However, they make note of the fact that creating plans manually gives you a lot more flexibility, and that's what we'll be doing here.

So to think of this a bit differently. I would guess a lot of people here have probably purchased a car or gone through the process of purchasing a car. If you buy a car, there are certain maintenance items that you're going to want to consider such as an oil change, tire rotations, and tire replacement. You also are going to want to make sure that car is properly insured. Just like having insurance for a car is not optional, maintenance plans are not optional.

So here's a brief look at our agenda for today. We will go through an overview of the maintenance plan, and we'll implement it. Then we will break into each of the components of that maintenance plan-- backup, integrity check, re-index or index, as well as the statistics update. And then we'll address any final considerations as well. Any questions at this point?

All right. So here's some of the setup tasks we'll go through. We will take a look at the SQL Server Agent service. By default, that installs with SQL Server to start manually. It's something that, if you're implementing this plan, it is utilized pretty heavily. So we're going to switch that over to start automatically, which we'll address. If the server is rebooted, it'll start back up and the task that it runs will also execute appropriately.

There is some configuration for the backup location. That is an important piece of this because out of everything in this maintenance plan, that is the one thing that could cause you problems if you don't configure it appropriately. Because by default, it is set to the C Drive, and obviously you're sending database backups to your C Drive. If you have a rather large database, you could fill up that C Drive quickly, and that's a problem.

A note on network location-- and we'll swing back to that-- but network location is also something you may want to consider in this maintenance plan. And we'll explain why later. And then something else that we'll be filling in for the maintenance plan is the backup hours, which

is basically the retention of your backups that this maintenance plan is executing-- how many hours you want to keep them around. So we'll set it for 168 hours, which is seven days.

So I'm going to switch over to Traci. Traci, you're going to want to switch the slide. Traci's going to basically bring up another machine we have here. I have a virtual machine running on there that has SQL Server Dev edition installed. But I wanted to show everybody the maintenance plan solution website, and then there is a link right there on the front page to download the maintenance plan. That is basically all you need to do to get the maintenance plan. I have a link to that site at the end of this as well.

So Traci has switched over to the virtual environment, and she's going to bring up the maintenance plan. There's only a very few areas that need to be addressed. So we touched on the backup location and the cleanup time both. Those are the only things that she's going to be changing today. So, because this is a virtual environment, I only set it up with a C Drive. We're just going to leave that one as-is for now. But cleanup time, she can put in 168. And that is basically the extent to modifications on this maintenance plan.

Before you execute, Traci, you're going to need to modify the service as well. So again, the service that runs this is going to be your Agent service and, specifically, this one is a Vault environment. So Vault runs as an Autodesk Vault instance, installed on whatever server it has been installed on. So that is SQL Server Agent, and then you can see that it is specifying Autodesk Vault. Any questions on that? Is this clear?

All right. So now she's executed the maintenance plan. And we can see the message indicates that it's completed successfully. If she expands the Agent section within SQL Server Management Studio, you're going to see that there are several jobs listed there. The one at the very bottom is the only one that would have existed if we looked at this prior to implementing. But all of the other jobs were created by this maintenance plan. And these are what we'll dig into next.

TRACI

PEURASAARI:

So, if you haven't figured out what we're trying to get across here, again, maintenance plans are not optional. They really are important to keep your database working well.

So a little bit about the core components of a maintenance plan. There're decisions that you need to make when you're putting your maintenance plan into place. And the things that you want to think about are how much time to get back up and running, and how much data loss you can afford. And that's what you use to determine what kind of backup you want to do. The

types of backup are transactional backups, differential backups, and full backups. Now, one thing you might notice is, by default, Vault is usually set up in simple mode, which means transactional backups really don't come into play. But if you've changed it, and it's not running in simple mode, you would really want to implement a transactional log backup to keep track of that.

Differential backups. One really important note to pay attention to is if you're using differential, when you go to do a restore, you have to have the full backup and every single differential that's happened after it. And if you're missing anything, you will not be able to fully restore your database. So those are the kind of decisions that you need to make when you're trying to determine what your retention and what your restore period is-- what your backup period is.

Virtualization. A lot of people run on virtual servers nowadays, and snapshots are the easiest thing in the world to do on a virtual server. A virtual snapshot will not take care of backing up a SQL Server unless you've brought SQL down, because any transaction that is in the middle of processing will not be covered. So again, you want to make sure that you're doing a SQL backup or a third party backup of some sort-- in this case we're talking maintenance plans and SQL backups. If you do a virtual snapshot to grab your backups, that's fine. But you cannot use virtual snapshots to get SQL backups.

DR recommendation. It's really important from a DR standpoint that if you are backing up to your SQL Server-- as we showed in the demo here, we were going to the C Drive, or if you're going to the D Drive-- that you do copy them offline. However, in the maintenance plan we use, you can use a UNC location and you can directly backup to an offline storage as well.

And then the SQL Server Service and the SQL Server Agent. You want to run them-- you saw me go into Services and start them. But you want to run them with a login account-- domain account-- that has access to the folder that you're backing up to.

STEPHEN

So, just a quick show of hands. How many of you are running in a virtualized environment

PORATH:

today? I know a lot of-- yeah. Just about everybody it seems like. So, that's a very important consideration to make. You can do a backup of a virtualized environment, but guaranteeing it's stable and aware of SQL activity is only really possible by doing backups through SQL.

TRACI

Integrity checks are one of the pieces of the maintenance plan it goes through. There is a

PEURASAARI:

command that SQL will use, and it checks the status of your database and the integrity of your

indexes. One of the reasons this is kind of a critical point is, if you get errors on this, you want to make sure that you run through and do a correction and then redo your backup. If you have an integrity error and then you do a backup, and then you have to restore from that backup, you are not going to get a solid restore. You're going to be trying to restore bad backup. So following the integrity check and checking the output that you get from this is really important to know that you are restoring from a good backup.

**STEPHEN
PORATH:**

And one note with anybody running replication-- Replication increases the transactional complexity of your database quite a bit because replication is being tracked against all of those databases. So every time you add a server, every single transaction that has occurred against that database is theoretically going up one more time for every additional server. So this becomes even more important in those environments.

**TRACI
PEURASAARI:**

So now, for most of you for Vault, this probably becomes really the thing that a lot of people pay attention to, which is performance. Maintaining your indexes is one of the key points of having good performance. Your indexes are what tells SQL how to read your data. And keeping your indexes clean, and nice and organized, is what keeps SQL performing nice and neat. There's two options for defragmenting your indexes-- a rebuild and a reorganize.

A rebuild can be done only offline. A reorganize can be done online. The reorganize is lighter weight and quicker. So those are some of the considerations, again, when you're making your decisions in the maintenance plan. One of the options, if you drill down into the maintenance plan, is what fragmentation level you're going to use to determine when you're going to do which option. So again, it's when you can have downtime, when you can afford to have downtime.

Statistics are the other thing that highly affect your performance. Statistics tell SQL what indexes to use. As your statistics get stale, SQL will ignore the indexes that have been created specifically for that purpose. So rebuilding your statistics are another critical point. As an example, we were working with a customer who was not building statistics. We went in and we put a regular statistics job. So a query that was reading one million reads took 90 milliseconds. After rebuilding their statistics, that same one million reads took three milliseconds. If you multiply that by how many users, how many queries, how many check-ins, how many checkouts, it just kind of gives you an example of the performance difference that you can see in that single piece of maintenance.

STEPHEN Yep. That's a very good example. The last job that is created by the maintenance plan is a
PORATH: purge job. And this is basically something that really just turns this maintenance plan into something that, ideally, is going to sustain itself. It's going to clean up its own backups. It's going to clean up all of the logging that it does at the specified intervals. So from this example, we did seven days, 168 hours.

If we move on to the next important piece of this which is scheduling-- We've got some minimum recommendations for the scheduling of this backup. Most places are going to want to back-up at least once a day-- full backup. And then transactional log backups can be done hourly or in-between that at whatever interval you desire. Do you want to fill in anything on the simple part of it? The simple mode?

TRACI No.
PEURASAARI:

STEPHEN OK. And then integrity check daily but before the backup. As Traci said, this is going to identify
PORATH: errors. So running it before the backup allows you to identify an error and then hopefully address it before the backup is done. So if you're doing a backup and you have an integrity issue, you're backing up that integrity issue. So fix it first, then backup.

Index and statistics-- very much dependent on volume. We've got kind of a recommendation of a daily rebuild and a weekly re-index. So one is less aggressive than the other, and Traci kind of covered those previously already. But basically one is fixing the index that's there. The other is rebuilding it. And then for the purge and cleanup, once a day is probably sufficient.

So Traci's going to do a quick demo on the scheduling. But the scheduling screen is right there. There are all of the typical expected scheduling options that you have available for, that such as daily, hourly-- so on and so forth.

TRACI It actually is very simple. Once you've got these jobs created, you determine when you want to
PEURASAARI: run them-- what time for each job. And it is simply a matter of-- I'll use the cleanup just because it's a small job-- going into Properties and selecting the schedule, and adding a new schedule. And determining when you're going to run it.

STEPHEN Naming's important only for you going after yourself, and others possibly picking this up. So
PORATH: you can name it whatever you want. But I would try to be somewhat consistent so either somebody else can follow in your own footsteps, or you can follow yourself.

TRACI So you can see it's created the job, and then now you can see right here it's in this schedule.

PEURASAARI: It's set to run. Next time it's got to run.

STEPHEN And maybe touch on this screen, too. This is a good way to see how your jobs have been

PORATH: running, whether they've been running successfully, or they're scheduled to run the next time. Those are very important tasks after the fact.

TRACI I got there just by double-clicking the job activity monitor, if anybody didn't see that.

PEURASAARI:

STEPHEN So the cleanup is looking at what is in your backup folder. And if it is older-- this example of

PORATH: 168 hours-- if it's older than 168 hours, it'll delete it. So it cleans up the backup so that you don't have them being dumped into a single folder forever. It also does clean up its own logs, too. So there's some logging that is created here, and it backs those up and just-- Sorry, I should have followed up with the question. What does the purge job do, is the question.

AUDIENCE: Under your demo that you have, you have separate jobs [INAUDIBLE].

STEPHEN Oh, sure. So the question-- as long as I'm getting it right-- he was asking about the difference

PORATH: between user and system databases. I guess I'll give the non-SQL DBA answer, which is system databases are basically the Microsoft-created databases that are going to exist on any SQL Server anywhere. They are basically the base that creates all other databases. It's the structure that Microsoft uses. The user databases, for Vault users-- that would be all of the Vault-related databases on SQL Server. For any other applications or whatnot, those are basically the created databases that store your vital data. Both can be very important. Obviously, most people here are probably mostly concerned with the user databases, as that's where all of your design data and whatnot-- for Vault, at least-- is stored.

TRACI It would include your Content Center, your KBM, as well as your Vaults.

PEURASAARI:

STEPHEN So something that's obviously very important here is plan execution. You want to-- at least in

PORATH: the very beginning-- be monitoring this plan, making sure that it's executing when it's supposed to be. If it's not, there are error links in anything SQL Server-related that will start pointing you into the right direction as to why something is failing. And then once you've confirmed that things have been running on the appropriate schedule, and successfully, you probably can start monitoring a little less frequent and only when there's errors.

And that kind of leads into the last demo that we have here, which is notifications. So backup and integrity checks are really kind of the ones that you want to be concerned with any errors. So if an integrity check results in an error, or a backup results in an error, a notification of that is something that you may want to configure.

TRACI PEURASAARI: So assuming that you have email configured on your SQL Server-- or that you would request whoever is running your SQL Server to configure email-- you can have any of these send you a notification. Steve mentioned like using the backup job. But again, it's just properties from the job, and there is a notification option. And you can have it send you an email on failure or success, or whatever option you want. There are also options for alerts-- to add alerts. If you don't have email configured, you can have it do alerts and it will write alerts to a log file. And then you can look in the log file to see what's going on. So there are options if you don't have email configured.

STEPHEN Thank you.

PORATH:

AUDIENCE: [INAUDIBLE].

Where can I find the alerts?

STEPHEN The alerts? She wants to know where she would be able to find the alerts.

PORATH:

TRACI PEURASAARI: The log file for the alerts are created in the maintenance job. You can put a location for those just like you put the location for the backup folder. Otherwise, they will go in the same folder where you've written the backups, unless you've specified otherwise. It's part of the job that we ran in the beginning. One of the options is to tell it where you want to write, or where you want to log the alerts. One option is to log it to a table. Another is to log it to a file.

AUDIENCE: Yeah, I have it logged to a file. And I'd like to [AUDIO OUT].

STEPHEN PORATH: Yeah, that's what she was actually-- So the question, assuming I got it right, is you want to get an email when there is a problem. So the last thing that she demonstrated is the area that you would figure that. It does go with a caveat of needing to have SQL Server Mail configured as well. So assuming that is configured, which is more of a server-level task, you can then set up

notifications based around each one of the jobs, whether it's executed successfully or not.
That good?

AUDIENCE: [INAUDIBLE].

STEPHEN No. So it's pretty simple as a part of the job. But configuring SQL Mail-- You want to show it
PORATH: again? Yeah. Go ahead.

TRACI If you're looking at a job, it a property of the job. You just go to the job that you're configuring.
PEURASAARI:

AUDIENCE: [INAUDIBLE].

STEPHEN Oh, SQL Server Mail?

PORATH:

TRACI Yeah. Well first, under properties of the job itself, under Notifications, this is where you tell it
PEURASAARI: you want it to send you an email to that effect. To set up SQL Mail, there is a query that you
can run that turns the SQL Mail on on SQL. Or there's a job that you can configure SQL Mail.
You would have to work with-- like for Mail administrator-- to have it use email, computer
email.

AUDIENCE: Thank you.

STEPHEN You're welcome. Any other questions so far? Hopefully making sense.

PORATH:

AUDIENCE: Maybe. When you use Vault, you can configure for the application, I'm sure, to do the backup
and backups to the filestores as well.

STEPHEN Yes.

PORATH:

AUDIENCE: Now we're talking just backing up the index.

STEPHEN Yep.

PORATH:

AUDIENCE: It can't conflict in the wrong way.

STEPHEN Yes. Yeah, there is a snippet on that that we'll cover. So you want to let me know after we

PORATH: cover that if it hasn't addressed that question. Hopefully it does. But yeah, that is a different set of backups and whatnot.

TRACI There also is a class tomorrow that talks about doing Vault backups that covers using the

PEURASAARI: backup that you're talking about-- the Vault backup. Yeah.

AUDIENCE: Just we use the one-- Should be very careful with the other, especially with incremental [INAUDIBLE].

STEPHEN Yes. Yeah. You want to make sure that both are roughly identical time. So under those

PORATH: additional considerations, I'll start from the bottom. So the Vault-- (LAUGHING)

TRACI There was a good lead-in.

PEURASAARI:

STEPHEN So the Vault filestore is not being addressed by this backup. This backup is very much SQL

PORATH: Server. It is not necessarily dependent on you running in a Vault environment. You could be running in any SQL Server environment. It is only backing up the database. And this solution is backing up all databases, system and user. So that is something to make sure you keep in mind. You can actually execute some of that. Sorry about the audio issues in here.

So something to be aware of. You could execute the filestore backup piece as a part of SQL as well. I don't know if you have any comments to add to that as well.

TRACI Yeah. Basically an answer to your question-- Well first, in the class tomorrow, we could cover

PEURASAARI: using third party versus using [? EDMS ?] backup. But yes, anybody who does chose to use SQL as backup, you want to make sure that it's possible at the same time. We've just been, for the purpose of this class, we're looking mostly at maintaining your SQL Server and [? above ?] specific just Vault. So it was a little bit broader. But if you are using this as your only method of backing up Vault, you want to make sure that your parameter files work at the same exact moment. Or after you use backup. Yes?

AUDIENCE: Do you run your backups from Vault? When you do this, I see it's grabbing files towards this [INAUDIBLE] other than-- Is it getting the SQL database as well? So if you do it from Vault do you get both ends?

STEPHEN Yes.

PORATH:

AUDIENCE: I'm not trying to say you shouldn't do what you're speculating.

STEPHEN Well-- and, and, and-- well, if you want, I'll address and then you can add to it.

PORATH:

TRACI Yeah, I'm addressing if it is the SQL databases?

PEURASAARI:

AUDIENCE: Yes.

TRACI So, when it in the Vault, it does get the Vault databases and it gets the filestores. It does not
PEURASAARI: get the system databases. That was stored in SQL, which-- if you have to do a manual restore, especially in a replicated environment-- those come into play. So, yeah. In the folder you would see all your Vault databases had been backed up. Your filestore had been backed up. It would wind up in Vault.

STEPHEN And to add to that, too-- The other maintenance tasks are not being done when you do a Vault
PORATH: backup. So those tasks are definitely highly important, especially for the performance implications and whatnot. The Vault backup does cover the user-specific databases that Vault is aware of. This you could do against the server running various vendors' databases, or even internal databases, and have it back everything up, too.

TRACI I mean, if you're using EDMS to back up your databases-- and that's probably what you're
PEURASAARI: running right now-- I've set up a customer where I've run this. And then enabled the jobs to do the [INAUDIBLE] and to do this integrity. And to do indexes but not enable backup jobs. Or just enable the system database backup tab, because the backups are already covered. And there's no reason why you can't pick and choose which piece of this you enable, you feel like you're covered comfortably the way you are. And that's fine. That's why they were individual jobs.

STEPHEN Yep. And to add one more thing to that I wanted to add to the discussion was the backups
PORATH: here. That's only taking the databases into account. But you could-- we've had experience with this between her and I-- where the Vault backups, because of the size of a filestore and everything, were taking more than two days to complete. This solution, plus a manual backup of the filestore, was able to be executed in less than a day.

So there are other things that the Vault backup does that are very good to do. But if you're going through-- in this scenario we were doing a backup of the database and filestore so that an upgrade could be performed. Doing that, we want to do the upgrade over a weekend. You want to make sure your backup's not holding up the whole process. So, something to consider.

AUDIENCE: Will there be class tomorrow?

**STEPHEN
PORATH:** The class tomorrow. I've got that title at the end here. And then he's the main presenter and she's the co-speaker. So I'll let them-- We can address it at the end if you want. I don't think we have too much more. So one message we're just trying to impose upon you-- maintenance plans are not optional.

So something I'd like you guys to do. When you return to your office, I'd like you to download this plan, modify the necessary areas of it, plan it out-- think about the recovery time that you need and all of that to suit your day-to-day needs with the Vault, or whatever databases are in this equation-- and plan that out. Implement it, and then review its progress.

Some maintenance thoughts from Benjamin Franklin. And here's the links we've been referring to. The URL to this is located at that link up there. If you grab that, you'll get all the rest of them. But Ola Hallengren-- that is the maintenance plan we're utilizing here. I've got a link to Microsoft discussing the different backup options available within SQL Server.

And then Brent Ozar is a SQL Certified Master, which is a certification that's not available any longer. He's probably one of a handful in this country. He has a blog that I would highly recommend, and then he also-- I put a link in here-- he discusses modifying some of the options for this maintenance plan specifically. It's all related to fragmentation levels, so doing it at basically different levels in the defaults.

**TRACI
PEURASAARI:** Yeah, I was going to say the lead for the modification is really good because industry standards are going to say that if your fragmentation is less than 30%, you should do a re-org. And greater than a certain percent, you should do a rebuild. And actually if you do use the index maintenance within Vault, I asked him to put that [INAUDIBLE]. Make sense? But in this he actually changes it and recommends a 15% threshold because the backup rebuild is done offline. You take it offline while it's rebuilding the index. So from a performance standpoint, especially for those of you that might have multiple locations that are all [INAUDIBLE], it's hard

to find the times when you can support your table going down long enough to rebuild your index. Or to have the best slowdown. So this doesn't really reach out to talking about what is a good way to manage that.

STEPHEN

There are some caveats, too, that I'm thinking of that I wanted to speak to a little bit earlier.

PORATH:

But the SQL Server Express versus Standard versus Enterprise. Enterprise-- as she has mentioned-- can do the re-index or index rebuild online. The other editions cannot. And then the SQL Express edition is limited further. I think you do not have the ability to do scheduling, and you do not have obviously the online re-index or index rebuild. You'd have to take the database down.

TRACI

So you create when those [INAUDIBLE] you have for each job [INAUDIBLE].

PEURASAARI:

STEPHEN

Yeah. You have to go a slightly different route. And then these are some of the accolades on the maintenance plan we're using here. It's voted best free tool by *SQL Server Magazine*. It was voted an MVP award by Microsoft on the Data Platform. It is something that Brent Ozar, the previously-mentioned blogger, has said, you know, there is no better maintenance solution for a database. And if you want to see some of the largest companies in the world using this, there is a list. Every industry you could think of, including some that are known for massive databases, such as insurance, health care, so on. So that's on that site.

PORATH:

And then Chris is going to be doing the class tomorrow with Traci, and that's a list of that class. So I would highly recommend everybody look that up and attend if you're still here at 1:00 in the afternoon tomorrow. Did anybody have any questions that we had not addressed? Anyone? We're a little bit early.

AUDIENCE:

Previously when talking with my reseller and Autodesk support, the stance has always been you use the EDMS console and you don't touch SQL. So I'm kind of confused how we're sitting here being told yes, we should do these things, and we should maybe do a separate backup plan.

STEPHEN

I think there was. So question for everybody. EDMS console backup versus a SQL backup. He was recommended by a reseller not to do SQL backups-- which I think there was a time where that was a recommendation. I don't know that that's necessarily what you would hear now or that it would be accurate if you heard it now. We do have a white paper published on doing backups, and I think he's going to address the Vault side of it. But the white paper talks about

PORATH:

third party backups, and using their solutions and what you need to do to be more or less supported using them.

CHRIS BLUHM: My comment on that would be to-- Is if you were asked, depending especially on who's asking, and you ask us, hey should I take my own homegrown solution or use your baked in, our answer is going to be do the baked in unless you can go home and say, I'm using this maintenance plan and I'm also backing up files for the same time there's a plan. They're doing essentially the same things. But what would frequently happen is if someone was told, oh you can do it on your phone and then parts of it fails, they will come back in here-- you guys told me it was OK, now everything's ruined.

So if there's a plan-- you're watching it, you're following up on it-- then that's the same as the EDMS backup. But EDMS is going to be a lot simpler for the lay person to use. The assumption is everybody in here is a SQL admin or has access to SQL Administrator. You're more than just, I got put in charge of EDMS and I'm told I'm supposed to back it up somehow.

TRACI PEURASAARI: And again, as Steve mentioned, in some of the situations the volume is dictating this [INAUDIBLE]. The volume just cannot. Because the EDMS is a linear process, the [INAUDIBLE] the filestore, your backup of filestore, and your backup database, customers could have high, high volume are getting in situations where three to four days to get a backup done by the time it's finished. So an alternate situation needs to be explored.

STEPHEN PORATH: Those are usually--

AUDIENCE: My third party backup does my Vault in a tenth of the time.

STEPHEN PORATH: Yeah.

AUDIENCE: And that's full SQL and local filestore.

STEPHEN PORATH: Yeah. And the real big differentiator that I can just instantly say that Vault has, is that it does a filestore validation. So it's validating your filestore before it does the backup. And that takes a significant amount of time. I would guess it probably takes, you know, 70% of it. If I had to guess. And you can do that independently, too.

So you can do all of the pieces of the Vault backup solution independently-- databased with

this-- and back-up a filestore separately. And you can validate it through the Vault command line. So if you hit the size where that makes sense, that's an option.

CHRIS BLUHM: And the only last point. I don't know that we're telling you you shouldn't be only doing SQL backups. This is more just an option. The only thing that Traci could say is, you should look at the indexing and the rebuild statistics. Make sure you've got those and then if it fits in your plan using SQL backups, that would work. But make sure you have backups that are working in verifying--

STEPHEN That's really the ultimate goal.

PORATH:

AUDIENCE: I have a third party backup solution that does SQL and the filestore. And I've done pretty much everything you guys have touched on, but I figured it out on my own. I haven't been able to find any support from Autodesk in this process. And now I'm hearing a different story that you [INAUDIBLE].

[INTERPOSING VOICES]

TRACI --was willing to using a third party backup solution. It's very much part of this.

PEURASAARI:

AUDIENCE: Yeah they go hand-in-hand. You can't do one without the other. You've got to do both.

STEPHEN Yeah. Yes. They are both useless without the other, basically.

PORATH:

TRACI Or you might just [INAUDIBLE].

PEURASAARI:

AUDIENCE: I have a question.

STEPHEN Sure.

PORATH:

AUDIENCE: We've got EDMS console backups scheduled full incremental throughout the week. You guys got me to thinking about getting an email notification on the SQL backups. I'm assuming you'd do the same thing through Task Scheduler. It's the first thing I want to look into when I get

back. Can you set that up? I mean, I've got something where I go check make sure my backups are going through successfully. But all that's going to get emailed every time there is a failure. Does that make sense?

STEPHEN I'm sure--

PORATH:

TRACI I would tell you that if you talk to your reseller, or whoever you have support for CSM, there are triggers that probably can be set up to send SQL Mail when a task happens. And AD [? before ?] EDMS is using a task within SQL to do that backup somehow. So you could reach out to whoever is your support. They could probably help you with that, to get it if you're not doing it from within SQL. If you're doing it from within SQL, then your SQL Mail can be set up.

CHRIS BLUHM: And if you're calling the EDMS console backup, if you have a scheduled task that's running the command, you'd be able to set that up in Windows task schedule respected. The difficult part will be working with your internal IT to get the server to be allowed to send emails.

AUDIENCE: Yeah, I've already been running that through.

STEPHEN Whereas the SQL Mail may already be established. You might have that in place. So they--

PORATH:

AUDIENCE: How to determine it [INAUDIBLE]?

STEPHEN Yeah that would be also some logic that you'd have to have built into it. Yeah.

PORATH:

CHRIS BLUHM: Normally, we write down the [INAUDIBLE] log file, and then send an email to a log file being sent that day. You're right. If the job fails, and your email step-up alert these job fails, it's never going to get to that step. So really what you need is two jobs-- one to run it, and the other to do a check to verify that job is sending email if it failed. Or use the daily verification. Then you're overwhelmed and you stop paying attention because you're getting a daily email eventually sent to the folder you're never looking at.

TRACI But yeah. You can do-- for example-- you can go up to [INAUDIBLE] for a solution with a job scheduler. Or you can get as fancy as you can get with other options. Or, using SQL [INAUDIBLE].

AUDIENCE: And again, [? power shell ?] will send an email if a test fails, then you have to correct an [INAUDIBLE].

STEPHEN It becomes a more complicated affair, I think, if you were trying to do that. Anybody else?
PORATH: Anything?

AUDIENCE: What about the big group Plant in the Vault?

STEPHEN If you keep-- Oh, Plant? So he's kind of the guy to answer that, but this would cover the Plant
PORATH: database on SQL Server. Absolutely.

AUDIENCE: If it's a different server?

STEPHEN Well, yes. You can implement this there as well. And all of the tasks can be--
PORATH:

AUDIENCE: Aligned--

STEPHEN What's that?
PORATH:

AUDIENCE: To align.

STEPHEN Sure. That's an important consideration. That's where I'll kick over to him.
PORATH:

CHRIS BLUHM: I can tell you want Vault database backs up first, then your Plant database. And then have your filestore anytime after Vault, because you can always do a Plant/Vault reconcile to pull down from the Vault to your Plant database. But if your Plant database is bigger than your Vault, you'll never recover. I'm happy with-- well not happy-- but I'm completely [INAUDIBLE].

[LAUGHTER]

STEPHEN Anybody else? Well, thank you, everybody, for attending. There are surveys so please feel
PORATH: free to fill those out. I appreciate it. Thank you.