Download your recording (external link)

SPEAKER_14: already discussed, but maybe to take a step back and to try to understand what are the exact status and plan if there is a global agreement on it, if we can find to have a way to summarize them in a way that are clear and agreeable by all of us, including the experiments and the sites. This is the goal of today's session, focusing on FTS. And on Thursday, we will have the session similar to the FTS one, but focused on XRudy. And from XRudy, we have been working with the team. The team has been working by themselves also, but also in terms of exposing and explaining the milestones and the path towards ILUMi, we have been working a bit with the team to try to have something that is also understandable at the level of not the super experts. And also, what we will have on Thursday, a more detailed discussion on the XRudy monitoring, which was one of the main points that has been discussed and raised in the last year. Um, that's all for the introduction. We don't have anything else. So we would leave the floor to the FTS experts. I think is Maria. If you don't have any questions. Thanks, Maria.

SPEAKER_03: Hello, Ale. Thank you very much for the introduction. And where is my presentation, please? Could you help me? Because now I can't see. I don't want to touch. So, hello also everybody from remote. Thanks for attending. Actually, we wanted to do this hybrid, so then all the plans and ideas for FTS in the Touch Challenge 27 are reaching everyone. So, thanks a lot that Ale was organizing this to make it possible. And yeah, Mija is preparing me the slides. I hope you can see them.

SPEAKER_15: OK.

SPEAKER_03: This is PowerPoint, but you have. Yeah. And in order to pass the slides here now. OK, so let's go. So in this presentation, I will focus on what we are preparing for FTS4 for the data challenge, what we are going to deliver, but also what comes next. Also, you were hearing in the previous presentations that we were talking about these numbers, that we will go to 48 concurrent transfers, or that we will even get 20 million queue requests. And this is trying to align what Data Challenge is preparing. So in Data Challenge 24, we were going up to 20K. And this was actually 25% of the round four. So in Data Challenge 27, the aim is to go to 50%. So that is why we were simply multiplying our numbers, OK? However, no one was requesting these numbers, but we wanted to. try to align and that is why we expect to go up to 40k as maximum in concurrent transfers and 20 millions in queue requests and this is the dates for the data challenge. So then maybe just to summarize why FTS-4. Actually in FTS-3 we were having four main structural problems. One, and the first one was the ground scheduling decisions. We were seeing that the links were starved. Priorities were not respected under heavy load. And also the storage endpoints were, you know, reaching the capacity limits so we were We were not protecting the configurations that the storage endpoints were putting up front. This we were having many SNOW tickets even nowadays with FTS3 and this is one of the main things that we wanted to do for FTS4. Then as well we were not having deterministic behavior from the scheduler. And this is because it was totally stochastic. So under heavy load, we couldn't explain to you, and this was a bit embarrassing, why there were these delays. Then another thing that we were tackling in FTS4 is the complexity. Actually, not just the scheduler, but all the FTS3, it

was kind of monolithic. And of course, future development was blocked. So all requests from experiments were very difficult to accept or implement. It was almost impossible. There were so many patches around that it was impossible to proceed. And then the fourth one is the database. It was a hard selling. So actually, scheduling and monitoring were competing for the same DV RAM. And of course, this was not helping in the scalability that even we want to go today with its 20 millions of queue requests. So. I don't want to scare you with all this but what I wanted to do is to collect all your requests, classify them in the different components that actually today you were getting some presentations about it, then also who was requesting what, and tagging as well the ones that were already presented in the previous FTES workshop. But even before going through them, I just want to say that FTES4 will be available at CERN, And this is because, this is not because FTS will not be ready for, it is not going to be ready for data challenge with D7. It's more that we would like to provide the proper support to our tier ones and to the rest of non-WLCG sites. And at this point, we have to prioritize the development of FTS4 and the focus and the objectives of the Data Challenge 27. But that doesn't mean that we are going to provide a proper deployment to the tier ones with the proper support on how to use it and so on. And this we will, again, we are going to provide a momentum to also do the support, mini data challenge, and so on that we are going to follow up with the rest of tier 1s. So saying that, because I think it's important that we include them there, these are the different items that we were that we were prioritizing for Data Challenge 27. As you can see, the scheduler and the flow controller were the top ones. And you can see the different tickets regarding no starvation, activity search that were mentioned by CMS, the DC27 mode optimizing the throughput, the tape factor, you know, to give priority as well to tapes. Then the flow controller that as well was, there was a lot of discussion of protecting the storage limits that were configured per endpoint and so on. All these by the way, thanks a lot to Alessandra Forti, to Katie and even Chris Haim for the last moment because indeed they were validating that I was actually collecting all properly, so thank you very much. But as you can see, we were focused on the ones that were heavily requested. And this is actually the work that actually Nicola was doing and presenting. So if you want to see more details, please go to his presentation from this morning. Then, as well, we have a heavy part here that is the workers' scalability, and I will say performance. But this is the part where will help us to go to this 40K of concurrent transfers. And this is the work that was carried out by Louis. You will have more details in his presentation. I mean, he was already presenting it this morning. Yes, if you want to have a look. But this will unblock the scalability of FTS4. Then for the REST API, we didn't change too much, honestly, because we just focused on security patches. So this was the main thing, plus one request from CMS from the workshop that was end users can change activity of queue transfers to jump on it. Then, even if it was not requested heavily, but it's quite important for DC27, that is the test with tapes, we were doing a heavy effort there to provide a proper tape scheduler and tape demon with FTS4. And this is the work that was very well presented by Joao this morning, so if you want to have more details, you can have a look. Then also maybe you can see here is Minor, the database and deployment part, with actually all the database deadlocks that my colleague Nicola was explaining. But actually, Steve was the one that was following up in all of this. And thanks a lot to Atlas Test that actually

were helping us to discover all these deadlocks. Actually, these deadlocks were inherited, many of them, from FTS3. that we tackle and fix all of them, plus another ones that we're introducing by FTS4, to be honest. But we, of course, there could be more, but nowadays we couldn't find more, neither with Atlas, neither with CMS, and we are quite happy about that. Then regarding the web UI, as you can see, there is not much that we are doing because actually we were thinking more, I don't, actually the web UI was for the security patches because the REST API, sorry for the error, the REST API actually was just introducing the tokens and this activity. I don't know why I say that. But yeah, for the web UI, it was more focused on the security part. We have a lot of plans for the web UI, overall for the tier ones. For WLCG per se, of course, it's important for the challenge, but not that much, because we have the WLCG monitoring with the active MQ messaging. So that is why we didn't thought that this will have a heavy weight. for Data Challenge 27. However, you will see that we have a lot of plans as well for the web UI. And then we cannot forget that a part of all this development and implementation, and even internal testing, because even that we are asking and we are getting help from Atlas, from CMS, and so on, we are doing a lot of heavy internal testing. Of course, we are doing continuous testing with experiments, that this is an effort as well from the team to validate, to debug, to provide a proper communication with them. And overall, even now, even we had an incident and we had a post-mortem, and you know, it implies a lot of a lot of work as well for the team. So these are the plans that what FTS4 will deliver for Data Challenge 27. But as I was saying, FTS4 is not a project. So it's not a project that ends in DC27. Actually, FTS4 is like another new version that FTS will have, because it should be a continuous evolution. Now that we are unblocking the monolithical model that we were having in FTS3, now FTS4 is opening the doors for accepting requests, for evolving, and so on, as any other service, as EOS, from 4 to 5, X2D from 5 to 6. So FTS will be in continuous evolution. So all your requests will be considered, of course, we will decide if we will implement or not together in coordination with experiments and so on. But now is a momentum that we could accept and implement, but of course organized by priorities, okay? So, saying that, this is what we were collecting for post-DC27 because the work is not ending there. The first thing and our priority, the official public release for the tier ones. And this will be in Alma 10. So we are not going to do the effort on providing in Alma 9, then in Alma 10. So we go directly to Alma 10. And the idea is that we will give you the proper support, how to use the scheduler, and so on, so you get the most out of it. Then we have more things for the REST API. In this sense, they remove the M2Crypto. Also, you can see that here we were having many Important item from the experiments that was the configuration import-export for the different sites and so on. We didn't forget about it. Actually, it was in our program of work. But due to a problem with the resources, we had to postpone it. And this is something that we know that is very important for all the experiments. So we will do it a posteriori. And this is a high priority after post-DC27. Then, okay, the permanent banning and what I was saying, the web UI. The idea of the web UI even before was to do it even in February, when actually we were having the plan to have the technical student and so on. So actually, the technical student at the end comes next month. So we will do this. We will even start this next month. And here you can see all the things that experiments were requesting. and that I know that is very interesting and so on, but cross-checking with

some of them, indeed you realize that it can be after this C27. Then of course more things regarding tape, more things regarding the workers, for the timeouts, then in the database also we want to optimize as well some parts in the database for the final transfers and then in the flow controller we have this item as well based on the throughput. But that doesn't mean that these are the only items. We are expecting, indeed, that after Data Challenge 27, we are going to have more things to do. And then with that, we will reprioritize according to what is needed. OK? So yeah. And now I'm going to talk about the rollout. So actually, Ale was presenting somehow the rollout. Actually, I'm not putting exactly your screenshot, but I'm putting actually even the same items. Maybe the numbers are changed because, indeed, I will add something else as well. First, okay, we were saying the testing with experiments. Actually, the scale was more in internals with EOS, as you could see in the presentations. And then for the second one, you see that actually we advanced this one because the idea was to have FTS4 in production for Atlas Transfer in Q4. Actually, we did it a bit in advance in Q3. And unfortunately, we had this incident. But as we were mentioned, we already did a post-mortem in this workshop. So you can even have a look from the people that is in remote, because with this closer, there was more than in this workshop. So then everybody knew up front what happened. But anyway, don't worry about that, because as we were discussing, it's not blocking the continuous testing and putting even FTS4 in production in Atlas to continue with that. Then as we were saying in October we want all the TAPE token test October is next month so we have to be careful about that we want that all the TAPE sites start testing TAPE tokens And as I was promising, after DC27, we will want to focus on the public release in Alma 10, and then, of course, to check with any one of them when to do the mini data challenge. Saying that, I'm going to say some things that we were doing at the same time was that, indeed, we did internal functional test tokens with CTA. And actually, it succeeded. But it was internal. So the next step, and for the For the normal workflow, probably there are corner cases that we have to test and improve there. But we would like that, indeed, we do this. Or maybe we could start with FTS in production for Atlas to test with CDA type tokens with Ruthio. So then, like this, we can really validate and sign it. Then, of course, Atlas wanted to be the first one for getting the risk, not to be the first one in production, but CMS also is very interested of starting with production transfers in FTS4. And we will do it after the workshop. Then another test that we would like to do, and we will see with whom, if Atlas or CMS, is to validate the first version of the flow controller. And this is quite important because this is one of the key elements that you were looking for, for as well Data Challenge 27. And yeah, we want to evaluate the first version with some adjustments. And this will be, you know, beginning of October. Then also, I was just discussing to Chris Hayen that it's very important to have LHCb on board. And he agreed, more or less, about the dates that I was proposing. But I think it's very important to have the three experiments and the LHCb doesn't have surprises really at the end. So I think it's important. And anyway, I hope they will be comfortable when Atlas and CMS were suffering already before. But I think it's crucial because then I would love actually to do some mini data challenge and I think Katie will love it as well to have the three experiments using FTS4 and mini data challenge that could help us to see how far we are for the data challenge 27 or if there are things that we have to adapt. And that's all. And then, of

course, then we will have the data challenge, the Tier 1 support. As you can see, we have a heavy timetable. That is why for the Tier 1s, we think that it's much better that we postpone, and then we give the proper support to them. Then going to the DC27 expectations. So actually, in FTS4, we were already up from providing some numbers, thinking about the scalability. So one thing was the 40K active transfers, the 20 millions of files. So we were doubling the numbers from DC24. Then for the transfer protocol, we say that it's better to do it as we were doing in Data Challenge 24, with HTTP only, with no streaming. And then the day token support, as we were saying, with Petra. Thanks a lot, Petra. He was testing already the cache. However, we would like, with a complete workflow, if possible, with, I mean, and the different D-cash that are going to participate in the data challenge. The same with CTA and the same with STORM. We would like that this is done by end of October so that we know who will be there for TAPE tokens. Okay, because it's a way to stop doing much more implementation even in FTS or even in the different software sites. Then, this is what we were saying to Katie, Alessandra and so on, and I think it's quite acceptable when I say, but then here I have some assumptions. So actually it's true that FTS was assuming that We will have all the links during Data Challenge 27, that we will have around 100 participant sites, but maybe it's not true. This is a number that we would like to know, but OK. Anyway, we are assuming all. But it's true that we are missing maybe to know how many sites are going to participate in Atlas, in Data Challenge 27, which are the sites, actually, and even the type of sites, you know, the cache, storm, and so on, that are going to participate in CMS, in LICB. I think we don't have yet that. Or maybe I missed it. But this is important. And then another thing that already I say to Katie upfront is that it will be essential that there will be an operator representative per site and service in our follow-ups meetings during the data challenge. So this will facilitate the debugging and the discussions and really the experts are responsible for that because then like this, yeah, I think we can make it more productive. And then an open question, but I think Katie was already following up, is that we don't know the files per second per link, but I think maybe I can kind of calculate. I don't know. We will see. And what throughput, then, can each link sustain? But it's true that, OK, I can get more as a lower number, because I don't know the sites that are going to participate. So actually, this will be, as well, very interesting. And we need these two to correlate with the size. So I think that's all. And maybe I just wanted just to show this map, because you know, always there is this kind of maps in Ruzio and so on. And the Ruzio is FTS. So actually, also, to let you know that we are, as well, doing this extra effort to help the different sites, not for Data Challenge, but for other cases. So just to let you know that our small team is doing many things apart of just preparing for Data Challenge 27. And I don't have anything else, but thank you very much for the experiments, for your passion, and also for this nice collaboration that we have. So thanks. Thanks. Thanks.

SPEAKER_14: Thanks a lot, Maria and the FTS team. You can hear me, right?

SPEAKER_03: Yes. Very well.

SPEAKER_14: There is a bit of echo. Maybe you might need to mute while I speak. But basically, thanks a lot for this super overview. And thanks a lot for all the work. I have written down a few

things that we might want to ask, but maybe there are other colleagues that want to start with their own questions. I don't see the room, so maybe you handle the room.

SPEAKER_03: Yes, I can handle the room.

SPEAKER_15: Yes, we have one from Dimitrios in the room.

SPEAKER_02: Thank you, Maria. So the immediate question for Atlas, and I guess for CMS as well, is for DC27, this means two unfavorable alternatives. Either we use FTS3 at the kind of other FTS instance that we have, like we run with this asymmetry, or we put all the traffic at the FTS4 instance at CERN, and I wonder, would that affect the quality of the data challenge? Katie, I don't know if you agree with this.

SPEAKER_05: My thoughts was we were going to run everything from the CERN instance. What do you think, Maria?

SPEAKER_03: I think there will be more pressure, but

SPEAKER_02: I say that this might be potentially alternative because once again we're changing the parameters compared to last data challenge and I don't know if it's too much to kind of expect for experiments as Atlas and CMS to have one instance handling like all this.

SPEAKER_05: Yes, it will be a lot and I think maybe the one thing we can't do with FTS3 would be if we were doing tape with tokens transfers. Is that true? Is that the only functionality we'd be missing?

SPEAKER_03: The only functionality? I mean, the scalability is in FTS4, okay? And of course, the DAPE tokens is in FTS4. So these are the two things that will be missing. But I will say that why not maybe to do FTS4 at 10 as you were having in plan and having the other with FTS3 and then maybe we can start. pushing in FTS4 and see how this goes.

SPEAKER_13: So, can I comment? Can I comment on this? No, there was Alessandra, maybe they wanted to comment? Yeah, I want to comment. But take the mic from Ale. Ale, please take the mic.

SPEAKER_06: I think we should... It would be preferable to have only FTS4.

SPEAKER_03: Sorry.

SPEAKER_06: Instead of mixing, it would be preferable to have an FTS for itself.

SPEAKER_05: Just one FTS running everything? Yes.

SPEAKER_06: Yes. We cannot count it. It is a challenge anyway.

SPEAKER_14: Ale, you are a bit low as a volume, but we got your point. We got your point. You said it would be better to have FTS4 only at CERN for comparison reason, but I heard Dimitri's points, which is basically saying that it needs to change the configuration for all the production. And then, I mean, I heard you. That's why we have this meeting to actually try to sort out an agreement.

SPEAKER_06: I mean, isn't that what we have done already in production before the incident?

SPEAKER_05: Sorry, I didn't understand. She said what you had in production already before the incident.

SPEAKER_02: Right, but this is the current average daily production, so the question is, all of DC27, if a single FTS4 instance can handle all of it? Like, from your side, Mihai?

SPEAKER_15: So we are definitely preparing the CERN deployment to do all the transfers that are sent to CERN. So at CERN, we always imagine we're only going to run FTS4, not run the two in parallel. If your question is, can we put every single transfer on FTS4 at CERN, that's maybe ambitious.

SPEAKER_03: Exactly. I will start with a normal setup and then maybe we can, if needed, to put more pressure in FPF.

SPEAKER_06: Sorry, during the challenge we will not have time to reconfigure.

SPEAKER_00: Yes.

SPEAKER_06: So we have to decide in advance, not now, but in advance.

SPEAKER_14: I personally think that the situation Dimitrios highlighted is more a matter of comparison with respect to DC24 than a configuration issue. We could keep on running the production like we are doing, and we could think about injecting all the DC27 additional load into the FTS4. Would this be something reasonable from your point of view, Dimitrios?

SPEAKER_02: Yes, this part. So on the configuration side, I don't think there's any concerns for Atlas CMS. We can do the same configuration. It's mostly whether that would distort the results, whether, let's say, in North America, things would appear a bit worse than they actually were just because some of the normal production traffic was flowing through FTS3 at the time.

SPEAKER_03: I think maybe I will need to do it comparable, no? Of what we were doing in Data Challenge 24, but with FTS4 at CERN in that side. I don't know.

SPEAKER_14: Why do you think that it should be in North America different? I mean, if we have a scalability issue of FTS4 in terms of handling the links and so on, I mean, there are completely different instances, the FTS4 at CERN and the FTS3 at CERN, right, Maria and Mihai? completely different.

SPEAKER_02: Yeah. Right. But DC27 will replace the FTS3 instance at CERN.

SPEAKER_03: Yes.

SPEAKER_02: And for Atlas, we'll have the FTS instance still at BNL. So what if- Yes, it's the BNL what he- Not reactable to the flow, the traffic, and we end up in a situation where it does not put the existing transfer well enough.

SPEAKER_16: We're coming on downtime. So the US side, the Canadian side, are performing worse.

SPEAKER_03: Yeah, this is-

SPEAKER_16: just because of that affinity. That's the question that we have.

SPEAKER_06: I cannot hear anything.

SPEAKER_03: Yeah, because he was not having the cue. Sorry, Alessandra.

SPEAKER_02: I mean, quick summary. Yeah, the concern is that the results might be skewed for the sites still handled by the FTS3 instance at the other side of the Atlantic.

SPEAKER_06: I mean, particularly the US sites which have a larger bandwidth, they would benefit from FTS4.

SPEAKER_14: Exactly.

SPEAKER_06: At least this was my point.

SPEAKER_14: So maybe, Dimitris, I would like, so first of all, this one is an open point, and it's very good that we are raising it. So this might require a bit more of discussion and agreement. We might not finish it today, as Cathy was implying. I have a related question. That is, the monitoring, the web UI, I think this one is more for Cathy than for Maria, but it's together. Basically, the new monitoring is the old monitoring, in the sense that FTS4 will publish in the ActiveMQ, so in theory, all the dashboards that we have, they should keep on working transparently. Is this correct? Of course, yes. Perfect. If this is the case, then Dimitrios, I hear you in terms of skewing up, maybe we could have sites that we are not able to completely understand if we are pushing enough. But I think that in terms of throughput, we do have a measure still.

SPEAKER_03: No, I mean, I just want to go.

SPEAKER_15: Yeah, I have to admit the meteors. This is a very good question. And yeah, if the same thing that happened in DC 24 overall in FTS3, now we could still see it in the US sites. But I don't think we have a question for this on the spot. One thing, we've asked for measurements of how many transfers we expect each link to do. It was also in Maria's presentation, and we also raised it a bit earlier in the DC talks. One idea would be if that number is reasonable, we're already preparing FTS4 for the 40,000 in concurrent transfers. If it fits in that number, we could think about having everything at CERN and then having another instance as a backup. But for the moment, our plan was FTS4 at CERN, FTS3 outside. But this might skew the results. This is true.

SPEAKER_14: There is Andrew Melo with the RACER band. Please, Andrew. No, you, you had your end result.

SPEAKER_09: Yeah, go ahead. Yeah, it's working. Okay, sorry. I'm not used to this app. Yeah, I was just gonna say that we did some, this is Andrew Mello from Vanderbilt. From the US CMS side, we did with our eight tier twos and the tier one, the tier twos were doing something on the order of over 100 gigabits a second. So I think it's difficult to backtrack and say how many concurrent transfers need to happen, because that's not something that we have readily available. But if you were to sum over the US facilities, you're in excess of a terabit a second of what the storage in the network can provide, if that helps set a bound on things.

SPEAKER_03: Take the queue one second. There is a reply from here.

SPEAKER_08: In the last mini-challenge, there was like between 200 and 300 per tier 2, if that's the number you're looking for, concurrent transfers to Fermilab. That raised to 700, 250 gigabits per second to Fermilab. That number helps?

SPEAKER_14: So I have a related question on this. Mihai, you and Maria, in your slide one or two, you made the diagram of the four boxes that you worked to improve. And if I'm not mistaken, the wrong scheduling decision on the top left. Um, it basically, you mentioned that we had problem with the links to start, but then the, um, capacity limit ignored. Uh, how do we measure that FTS4 is doing better with respect to FTS3? Do you already have a plan on it? The question is related because if you know how to measure things, maybe then this would actually allow us to renormalize the skewed data if it happens in US.

SPEAKER_15: Well, we actually had a presentation this morning on exactly this. We were running tests using a stable EOS from Alice P2 to EOS Pilot. So this was kind of our sandbox. And there we could run about 40,000 concurrent transfers. However, if we need to double that, I'm not sure. There's eventually hit a ceiling with the database, which we haven't pushed that far to see where that one would be. But at 40,000, we were OK.

SPEAKER_03: Yeah, and this is what we were expecting to do with FTES for Atlas at CERN, to do with this number. If we add more, I mean, that is the thing.

SPEAKER_06: I mean, the alternative is that you release FTS 4 to BNL.

SPEAKER_03: Sorry, could you repeat, please?

SPEAKER_11: I think she said that you could release FTS 4 to BNL.

SPEAKER_14: Apologies. Alessandra, can you please repeat, Andrew? Let's try with Ale. Ale, please go ahead.

SPEAKER_06: The alternative is that you release FTS4 to BNL and Fermilab.

SPEAKER_14: That's not what the FTS team mentioned.

SPEAKER_06: I know, I know.

SPEAKER_11: What about a different alternative? In the past, I mean for the last data challenge, we put also a big amount of load to pilot. So you could run an extra instance FTS4 pilot that will behave as a BNL and

SPEAKER_03: But this is an extra that we have just in game, no? Yeah. But if it will be used, then maybe we need another extra if we want to debug other things and so on.

SPEAKER_11: But during the challenge, you will follow the challenge, you will not... Yeah, no, but okay, we have other bots as well. You can deploy extra...

SPEAKER_14: Can we take a step back on this? I think it's important that Dimitri's question is relevant. Maybe Dimitri, we need to try to understand how impactful could be this problem with respect to the results that we are trying to get. Because if this one is not perfectly understanding the results, then it's one thing. But if it's still allow us to understand if the scheduling of FTS4 is going the right direction, then maybe this basically is the results are not the exact results we are

expecting. but it's good enough from our point of view in terms of understanding the gaps and the problems that the FTS4 team needs to work in the next years. Do you think this is really a showstopper in terms of understanding the test? This to me is not, because we have 100 here by the second with FTS3, moreover, whatever it is. I mean, like I said, grosso modo, we can discuss what's the exact number. And I think that overall, you want to go at another factor three of that, more or less. So we are still putting a lot of load on FTS4. So to me, we are still good in terms of understanding what's happening there. But expand me a bit more on your thoughts there.

SPEAKER_02: So I'm repeating concerns by others in the Atlas community, and the fear is that What if it's something that we'll only see the effect afterwards and it will, as I said, reflect poorly on the sites that were not running on FTS4 if that is the decision we're going with. So it felt appropriate to raise that we have to make a choice somewhat soon, I should say, about what the deployment, like what the configuration would look like. And yeah, ensure that this would be sufficient, like in terms of resources. Okay.

SPEAKER_14: Thanks, Dimitrios. Again, my understanding, Maria, correct me if I'm wrong, is that we will have FTS-3 in the configuration for production like you have in Atlas, in CMS, wherever you want to have it. And the FTS-4 will get injected from Rusio all the additional load. Is this correct?

SPEAKER_03: No, no, no, no, no. It's not correct. So actually, Atlas at CERN will use FTS4 for all their transfers, Atlas at CERN. CMS at CERN will use FTS4. And LHCbE at CERN will use FTS4 completely for all the transfers that they used to do at CERN. This I got it. This I got it.

SPEAKER_14: And the rest, FTS3. Sorry, we inject in FTS-3 for the US side? This would be the idea?

SPEAKER_03: No, the US side, they don't have other option than using FTS-3 because we are not providing FTS-4 for the tier 1s. So FTS-4 will be only at CERN.

SPEAKER_14: But sorry, Dimitris, we can have two things in common. We could push to Midwest tier 2 even with FTS-4 at CERN. This is just a matter of how we configure FTS-4.

SPEAKER_02: Right, hence the question, do we move all the load, all the traffic to the single FTS4 instance at CERN or do we run with an asymmetric configuration where some transfers, it could be not the data challenge transfers but the production transfers, that they still go to the existing FTS3 instance at BNL. Like those are what I see the two scenarios. Now Luca mentioned what if we have a second FTS4 instance at CERN like the pilot that replaces the BNL instance during the data challenge. So these are our options. Maybe one of them is sufficient.

SPEAKER_01: And can I make a real quick comment on configuration? So currently FTS server is configured per destination, which means that in Russia you set it on a destination site and then this FTS instance is used for transfers with this destination. So. Currently in CMS we have the Fermilab and the CERN one. The CERN one is handling all the EU traffic, Eurasian traffic, and the Fermilab one is handling all the US traffic. So I think this is what Dimitris is describing.

SPEAKER_02: Wow, so many hands. But again, just comment, this we can provide you with a patch that

SPEAKER_03: unlocks this like you can do like activity based yeah something like this so if necessary we can share the same what i suggested maybe can i say but just just maybe to dimitrios just just one second and dimitrios for you which is the main concern for you of using FTS-4 at CERN and FTS-3, you know, which is the main one. Sorry, passing the box to Dimitrios. Which is really the main, main concern. Because you will see in FTS-4 at CERN how it is scaling compared with Data Challenge 24. Now what will be in Data Challenge 27.

SPEAKER_02: Now, one concern is that FTS3 is not, for whatever reason, pushing enough traffic. And then when we look at the end results, at what was achieved, we see that from the FTS4 site, everything was handled well. And from the FTS3 site, it didn't. And that reflects poorly on the affected sites. And then the sites complain to us that this is a challenge, in a way. one of the concerns.

SPEAKER_14: I would like to, maybe we need to pause this discussion. You are in this workshop, and I think we are not going to reach an agreement right now. I think that mixing the two instances is bad. I would really love to see only FTS4. That's my personal preference, but maybe it's not possible. Basically, what I would suggest is that you keep on discussing these things during your dinner tomorrow, and then maybe We take it over on Thursday. We can try to sneak out a bit more time to see if we have an agreement in the next session in two days from now. That would be really ideal. So. Again, I don't think that we can solve it right now, but it's a very valid point, and this needs to be discussed. So maybe the ball is back to Cathy that is there, and Johannes that I guess is remote like us. So Cathy, most probably you need to drive this in such a way that for the C27, you have a clear understanding of the configuration. You agree, Cathy? OK, she agreed. Great. I have another couple of points, but people could have others too. Other people with other questions. Yeah, we still have quite a couple. Yeah, go ahead then.

SPEAKER_15: Please say your name, say your name. So, before we continue, a small summit, then we do a final question on this topic and we'll move to the next. But yeah, I think question noted. It's a very valid concern, but I don't think we're going to give an answer right at this minute, so we have to think about it. Now, we do the final round on this topic. No, no, no, let's stop it.

SPEAKER_14: No, sorry. Stop this question now.

SPEAKER_12: You discuss it later on there.

SPEAKER_03: Yeah, yeah, yeah, there are more questions.

SPEAKER_12: I just want to point out, so this here, yeah? Anyway, so we do have a development instance of FTS. There's nothing preventing me to install it as long as developers are willing to let me install it, yeah? So that's the only thing, okay? So just throw that out, okay?

SPEAKER_14: Discuss it with them and come out to us with a plan and where we can agree. I mean, this is something you need to discuss with the team there. They are there. Maybe you

need to offer them enough beers or whatever, and then you sort it out. They don't want it. And they explain why. It's reasonable why they don't want. They want to be able to offer support. So it's not up to me to say anything here, right?

SPEAKER_04: So that it would be more a question or like a follow-up since I saw a lot of information related to the token authentication for the tape. For example, in the case of Firmula for Dune experiment, who is going to be one of the participants for the data challenge, we are using CILogon as a token user for American transfers. And the problem that we have is that it only works at the transfer with unmanaged tokens. And I know that this is a long discussion from a long time ago. I just wanted to follow up on this because I know that in the past I was discussing this with the CILO developers and with the FTS3 team. and they said that they have to put some parts, but... I just want to say, because this is related to your expertise with the CILogon token history that you were following. If you remember about the CILogon, you know, managed token, that it was not working, and since DUNE is going to be part of the data challenge as well as a participant, For American transfer, we can only use the CILogon tokens with unmanaged tokens. So, I remember that we had this discussion before with the CILogon developers that it has to be some kind of patch from the FTS3 side. So, I was wondering the status of that, because I know that talking with Steve Murray, they say that you guys were going to follow up, but we didn't have more news.

SPEAKER_16: Yeah, I can.

SPEAKER_15: I think you're the last one on that. Yeah. Thank you. So back to me, Mihai, for the token part with the CI logon. Yeah, there's a discussion, and it has to do with FTS doing the token implementation and the token exchange request response based on IM interaction. But that's actually not compliant with the RFC book. And CILogon does it exactly by the RFC. So it's something we'll have to change, but we envisioned it more for FTS4, to be honest. You'd probably need it if you want it for the data challenge, depending on how these things will plan out. This might be needed in FTS3. So we'll get back to you on that. And there's a second part with CILogon. The fact that CILogon uses multiple VOs, but I think this question has never been answered at the entire WCG level. So I say this second topic, we park it after the data challenge. But the first one, I think that we're planning for FTS4 to have it. We could also, depends, we'll see if we... Yeah, which are your plans for D7?

SPEAKER_03: Yes.

SPEAKER_14: Cuba, please.

SPEAKER_04: Sorry. Yeah, I would have to take a look with the June liaison, which would be the plan for that. And I know that the CMS people and the June people, they were asking which are our plans to move forward FTS 3.4. But I think that we have to sit down all together in order to, you know, agree in a timeline when it's the best for everyone.

SPEAKER_03: Exactly, so then just keep us in the loop when it's the best moment. If it's after the challenge, there is no issue, then before we would like to understand actually which sites we are

planning and then which is going to be your strategy for tokens and even if you are considering TAFs or not.

SPEAKER_04: I can definitely follow up on that.

SPEAKER_14: Okay, thanks. As an organizational point, I do not know, I hope you do not have a strict 5 p.m. end of the meeting, right? I hope. Because we are late with our discussion, we have a few points. Is it okay if we stay a bit more?

SPEAKER_15: Ale, this is a tough one to answer for any of us organizers, because at the same time, I think everyone here also deserves a good break after an intense day, but... Okay, so let me try to be fast on two other points, and then... Yes, please, I would like to listen to you.

SPEAKER_14: Food for thoughts, and then we see when we can discuss more. The first point is we just realized that there is more discussion needed, we discussed it, now Cathy will help us, we are there. I have two points that I still are quite hot for me. One is the tapes and in your presentation you are showing correctly that you need to have Tier 1 states that are going to test that the tokens will work. And here the question is, how do we follow up that each of them, or at least a good representative part of them, has done it? That's the first question. And the second is completely different and is on IIM. Do we have any way of measuring if IAM is doing fine enough? Is it scaling up as we expect? So these are the two topics I still would like to make sure they are covered for DC27. The third is a bonus, but we might discuss later on, is that we focus on DC27 right now. There will be follow-up. And then we'll discuss further plan later on. But let's focus on these two, basically. The tape, how do we make sure that each tape is answering to us, to you, in this case?

SPEAKER_03: Dimitrios will answer. Go, Dimitrios.

SPEAKER_02: Thank you, Dimitrios. I suppose Atlas and CMS will handle that part, chasing the sites, ensuring that everything has been set up properly and validating it. I don't know if that's sufficient to answer your question. I do have a note. There's a snag with the interface between Ruchio and FTS that I'll describe tomorrow. But when it comes to the sites, I think the two main experiments, LHCbI assume the same, should be the ones that chase the sites to deploy the support. I'm hoping it will not be that difficult, at least for the decache sites.

SPEAKER_14: I hope, too, if there is any problem and you need something else from the WCG, some push at the level of the management board when we could have the responsible tier ones there, let us know. We can do it. But it's very urgent that the timeline is important and very urgent also from what Maria presented. So thanks, Dimitris. This makes perfect sense to me.

SPEAKER_03: Yeah, I tell you, for me, sorry, was the natural way, since Atlas and CMS are going to do transfers in FTS4, then I'm doing transfers with tapes, doing with tape tokens, no? I think is the next step. So then you can bring your Storm, Dcash and CTA instance to the loop, and then we validate all at once.

SPEAKER_14: I think this would be the most... We are ambitious. We could have other problems there, but it makes sense.

SPEAKER_16: Of course.

SPEAKER_14: And then we see. And the second one was IIM. I don't know if you want to answer on it.

SPEAKER_02: Well, I repeat what we said earlier this meeting, that maybe someone from IIM should be present in this workshop, as we said last year as well. I think Atlas is the most expensive user on IIM and currently with the latest version where they don't store the access tokens. And we have not used it ourselves to verify, but based on what they presented with their own tests, it looks very promising. So what I do hope, again, Atlas that we're doing the file-specific destination tokens, LACB also doing the file-specific tokens. Yeah, so I think between the two of us, we've put enough load on EIM to assess in advance whether that, or by that time, we'll know if it can survive those token rates. Mihai. Yes.

SPEAKER_14: Erk, thank you very much. I see you have already raised your hand, Bert. Can you hear me?

SPEAKER_10: Yes, please go ahead. So this is Berk from IAM side. Currently with the numbers of the tokens that Atlas is creating, IAM has no issues. And also with the last performance test we did, which includes all the last of the includes not storing the tokens in the database. The performance are way, way better than before. And if we just double the numbers from the DC24, which is the expected throughput, IAM can easily handle this.

SPEAKER_14: Very good. Thank you very much, Berk. Is there anything that you might want to test in addition for IIM for the C27? I know it's not fully related to FTS, but I mean, it's very much related. So is there anything you think you might be useful if FTS team does differently or?

SPEAKER_15: Maybe I can actually, I have a request and it's mostly to do with the managed token workflow. For the unmanaged, we saw even in the incident presented, 42% of ATAS transfers or 45% were using unmanaged tokens. But we never really tested that production scale, the managed workflow, which is, I think that's what CMS plans to do also in the data challenge. Unless if you go with unmanaged and then there's no problem.

SPEAKER_10: Just to clarify, when you say manage, do you mean using the token exchange and then refreshing those tokens?

SPEAKER_15: Exactly. Exchange and refresh.

SPEAKER_10: I mean, from our side, we would be happy to test this workflow with higher loads, for example, in the D-team instance, which is usually used for testing operations ahead of the data challenge so that we would see if there's any issues. Again, there also the performance has been increased much more compared to the last data challenge.

SPEAKER_15: But this is more of a question of scale. With Atlas, we saw half of production data using tokens and that was fine. But with CMS and FTS4, we never got to a quite high production number.

SPEAKER_03: Yeah, we have it.

SPEAKER_15: Yeah, but I want it in FTS4 also.

SPEAKER_14: What I would love to have is an understanding of these couple of things, Mihai. It is not clear to me exactly what you are suggesting and what the IAM team would like to test or not. I think it's very important that also this, maybe you can try to write it down in the notes here that we have in the live notes or somewhere else. But Berk, you got exactly what you want, what Mihai wants?

SPEAKER_10: I understand what Mihai wants but I don't know if there's a way to currently test it without just using this flow in the production or running it on a test instance but I'm happy to follow up with Mihai after to find a way to test it to not have a surprise later but I'm not expecting any surprise with the current performance-wise, but Mihai is right that it's not tested with high loads in production.

SPEAKER_02: But Atlas can do a test where we switch to managed tokens instead for a day or so. So this should get you, I think, representative statistics and see how that works. We'd be happy to do this.

SPEAKER_15: I think that's exactly what I want to see as well. FTS4 running production data, but with a managed workflow. So we're a bit unmanaged, but okay, yeah.

SPEAKER_14: I think it can be arranged. I think it's great because we are coming up with a lot of other important bits that needs to be tested as a part of being ready for the C27. James, am I forgetting anything else? Are there other questions that people have in the room or remote before the people in the room want to escape?

SPEAKER_05: Just a quick comment on the sites that will be involved. So we have obviously CERN and we have 16 tier ones that are associated to at least one of the LHC VOs. In terms of the tier twos, I've got a list here of over 90, of which I know CMS are planning to test all, however small the rate would be. For Atlas, I think they have not quite decided yet, but they're going to make some kind of cutoff based on size. Alessandra, do you want to comment on anything on like that at the moment?

SPEAKER_06: yes so we are at the moment what i'm thinking is to cut everything below two gigabit per second for the past year

SPEAKER_05: Okay, so Maria, how would that affect what you were asking to do with the number of links, was it, that you were interested in?

SPEAKER_03: We wanted to know the sites participating in Atlas and then the files per second per link.

SPEAKER_05: Okay, so you want to know how many per VO are involved?

SPEAKER_03: How many sites per view? Yes. And if possible, the files per second, per link that you plan to do.

SPEAKER_06: So just to be clear, I thought these numbers were going to be presented at Doma General.

SPEAKER_05: Okay, perfect. End of the month, we have Doma General on the 30th of September and the LHC, the CERN VOs will give their numbers there. Whether we have the files per link, that sounds like incredibly detailed information. Yeah, indeed.

SPEAKER_06: He may need another iteration after the 30th.

SPEAKER_03: Okay, anyway, it will be at least enough that we know which sites we are going to use. So that number will be at least something.

SPEAKER_13: Okay, this seems... Thank you very much, Katie.

SPEAKER_14: James and I have been trying to write down the notes as usual. I see that there is also a transcription, so if the owners then later on want to share with me and James, we can see if we did some major error. But also, please, you all can try to go there and see and improve our minutes there. I see that there are a few follow-ups that we need to do. With Cathy, we will discuss some other most, and you will tell us whatever you are able to come up later on in the next days. Berk and Miei, you agree that you will discuss between you and then most probably at this point, it seems to be the natural place where to discuss more details is the Doma General on the 30th of September, right, Cathy? This seems to be the natural place where we want to basically have another checkpoint.

SPEAKER_05: Yes, we may need to extend the meeting beyond an hour so we can have more time for discussion.

SPEAKER_14: And then there is the workshop one month after where we might want to have a discussion. So I don't think that everything should be settled there, but the ballpark should be settled.

SPEAKER_05: So the idea is that we'll cover the rates in the Doma General and then we won't have to discuss that too much in the WLCG workshop. There we will discuss different tests we want to do, the schedule and so on.

SPEAKER_14: There is plenty of space in the W3 workshop also to discuss more details. So again, all the FTES experts and X3D experts also, they are welcome to join and to register to come in Bologna. We are not going to have remote participation in the workshop for technical reasons, but you will attend in person.

SPEAKER_00: Yeah, and we should probably also address this questions of the FTS3 versus FTS4 configurations in the DOMA general as well. Even if it's resolved here in side discussions, it would be good to sort of address that in the DOMA general.

SPEAKER_14: Okay, so thank you very much for having hosted us for this hour and nine minutes. Apologies if we went a bit long. Thanks a lot for the FTS team for the great work that you presented to us and to all the people here that participated in the discussion. And me and James will join again for a similar one on XRudy on Thursday at 4 p.m. Thank you very much from our side is all. Thank you very much. This will work perfectly in a T-shirt, like Zoom. Zero problem. Thank you very much. We disconnect. Thanks.

SPEAKER_15: Thank you very much, everyone. I think it's time to end the second day of the FT