LS Retail Office Hours

Overview of replication in SaaS without Data Director

Webinar Q&A January 29, 2025

Presenters: Matthias Matthiasson, Ricardo Moinhos, and Adalbjorg Karlsdottir

Q: **For Azure storage replication, does the customer need to monitor the Scheduler on the POS? Are there any SOP available for checking the Scheduler status on the POS?**

A: When you replicate something with the Data Director or whatever, you run the job, and it creates the request to the Data Director. It's the same thing here. You run the job, and it creates a package on Azure Storage. You don't really have to monitor that, more for Azure storage rather than the Data Director.

So, after creating the data, you would use the Web Replication Monitor to see if the POSs are going in and checking for packages. You could see who has not checked for packages for a certain amount of time and who has checked for packages but has not been completing it, in a centralized way.

Q: **How will the preload actions work?**

A: Preload actions will work exactly like before. You create them by the job, it will create the preload, and then you do a replication, controlled by the preload, same as before. The only thing that is different when you run the replication, you would then have the destination as web replication. You would have the replication job marked as web replication. That's it.

Q: **Will this replication work between LS Central 13 to LS Central 26 during migration?**

A: No. It will not. It will not work that way. It is designed to work within the environment where it is has been created. It has not been designed to go from older versions. This is something we may need to look at: What would we do for the migration phase?

Q: **For preactions, you must add the table you want preactions for to, preaction create and tape creation table. How does this work with time stamps? Will this work for all tables?**

A: Same is true for timestamps, mainly because we have to log delete actions. Therefore, same setup as before. In theory it should work for all tables, but there are some system tables that will not allow you to replicate with web replication (do not support web service update).

Q: Does this apply also for web replication for LS Analytics?

A: LS Analytics does not support Web Replication. There is no internal solution to replace the Data Director, but the product team recommends using this data warehouse extension to export data from LS Central SaaS.

Q: Is there a clear road map of how replication will work for LS Retail? As there are so many options being proposed, from time to time... There are other POS solutions that are really smart at this and simple to manage on a day-to-day basis. Small retailers see this as a pain maintaining and managing. Most of them nowadays don't have IT admins.

A: Regarding maintenance, our recommendation would be to just run the POS online. Of course, you need to update the Hardware Station, through Update Service, but that's maybe once every six months or once a year. So, I would say run online. And we're trying to make this as simple as possible. You have to admit that Azure Storage replication is far simpler than Hybrid Component Server.

Q: Just to be sure, will web replication in SaaS trigger OnBeforeInsert and OnBeforeModify event triggers?

A: Yes, they will run and OnAfter triggers as well, there is no way to compress these events.

Under the hood we're using RecordRefs for the Web Replication feature to be flexible and to allow replication of custom fields without the need to create new API/xmlports to include custom fields.

When inserting (same logic applies to modifying and deletion) records we're using the RecordRef.Insert method, where the RunTrigger parameter is set to false. That means that the OnBeforeInsertEvent will be triggered but the code will not be executed if you're checking if the RunTrigger parameter is set to True to execute the code.

Q: Is store-specific price/promotion replication supported?

A: I assume we are talking about the time stamps here. Or is it about the actions to replicate something specific to a store and not to of the stores? If that's the question: that you first have the preactions and then actions are created according to the location distribution. Yes. It works. So, you run the same job as you always been running, and you just have the task, the set to time stamp. It will use the time stamp to create the action like before, or in that case, it would be a preaction, create an action. In this case, time stamps create an action, and then you run the replication according to actions. So, yes, it will work the same way.

Q: **You mentioned that you can only replicate data from same table to, from table to the same table and field to field. What about replicating a constant?**

A: Yes, it is only field to field when using web replication, no constants or conversion.

Q: **Will I still need to fill records in the *Preaction Creation* table for new tables?**

A: I assume you're talking about time stamps. No. You don't. Preaction table will not be used if you're using timestamp. So, that table would not have any records.

Q: **Will the Monitoring Tool allow us to monitor data coming from store/POS with errors? What type of errors have you encountered with the monitoring tool?**

A: I'm guessing the question is, will the web replication monitor cover the transactional data to head office? And if that is the question, the answer is no. There, you would have to use the client and the server logs. Good input, though. We'll have a look. But if this is something related to the import process, if there's any error when applying the data, then there will be a call to that office stating what went wrong.

If the problem is that the POS is not able to connect to that office because there's no Internet connectivity or some issue with the authentication setup, for example, then, this will not be logged because the POS is not even able to communicate with the head office, meaning that you are not able to understand what's wrong with the POS but you can understand that there's something wrong because the POS is not communicating with head office.

If the POS is able to communicate with head office and the error is rather related to the communication to Azure Storage or the data import process, it should log that the POS has not asked for the packages, or it asked for the packages and fetched them, but an error occurred when applying the data on the POS.

Q: **Will this new approach support object replication?**

A: No. But Update Service covers that.

Q: **Can we do any custom table replication from standard LS table for some specific fields value? (Not the same table it's a custom table)**

A: Web Replication also covers replication of custom tables. But only from that custom table to the same table in the destination

Q: **The time stamp in SQL looks like a random hexadecimal number. Can you go back in time and redo a replication like you could do by setting the replication counter to a specific number?**

A: Yes, timestamps are kept as Big integer (in that sense more like a row version) and can therefore be set backwards.

Q: **Is it understood correctly that you are not able to pull data?**

A: Yes, correct, but it would be interesting for us to better understand in what cases you would need that and why.

Q: **Without preaction records, how do we identify item insert or modify or delete using time stamps?**

A: For the insertion and modifying, the record is in the table, as with the timestamp. But if you delete a record, that record is not there anymore. LS Central will log that into a specific table. And, with timestamps on that table, we'll create a similar result as with the preactions.

Q: **Is there any possibility to replicate old NAV to a new Business Central database by using the new application method and vice versa?**

A: No. We do not support that. But that's something we were interested in; to look at and see how it could be done.

Q: **Do we need to be aware of any limitations of this new replication for large retailers (more than >50 stores)? Some wish to replicate every five to ten minutes.**

A: So, if you're replicating down master data, you must realize you run just 1 replication every 5. Let's say that. And then you run that, and that's just 1 export. So, it creates 1 package in Azure storage. That's it. So, it should be able to handle that. You agree? So, that creates 1 package, and then it's, all the POSs reading from Azure storage, which is independent. So, if you have 500 POSs or 50 POSs, it doesn't really matter regarding head office. It just writes 1 package. And then 500 or 50 passes are all independently going to Azure storage, which should be able to handle that kind of a load. So, it's irrelevant to that.

Q: **What happens if the store is not on an updated version and the head office is updated? What will happen to the updated fields in head office that aren't in the store database? Will the replication work even if the fields are not the same?**

A: Yes, it is included in the process to automatically ignore extra fields sent but not present at destination.

Q: **What will happen if fields are available in the store but not at head office?**

A: What will happen if you're replicating from head office down to the store and you're not part of the package? Of course, it would leave the store fields alone.

Q: **Will other by Action tables still be possible beside pre action and action?**

A: I thought there were only 3 action tables. Right? Preactions, action, and preload. All three work.

Q: Can we have email alerts for success and error for jobs?

A: You can configure to have emails sent if you have packets that do not meet the threshold defined. You can always also rely on other tools on the Azure Stack like Telemetry, Logic Apps and Power Platform to get these kinds of features, to have an email sent when there's an error. You can send some message to Teams and other platforms. It's easier to have a built-in feature, but it would be nice for partners also to explore other kinds of ways to have this information.

Q: For which period will Data Director / Job Monitor be supported alongside Azure application on SaaS? And is it taken out from version 26 onwards or some other version? Or are these versions available when not using SaaS and this new replication?

A: There are no plans to take out anything at the moment, and not on future dates. So, nothing has been decided to take anything out regarding the replication, the Data Director, or change the Scheduler or anything like that. So, if you're using on-premises stuff, please continue. And report if you have problems, we will support it.

Q: Will the Data Director be discontinued?

A: There are no plans to discontinue the Data Director. The Data Director, of course, is a cornerstone of our on-premises implementation, and there's no need to take that away. We are providing a solution that is more SaaS-friendly. There are no plans to stop the Data Director or the Web Monitor that is servicing that.

Q: Is there a retention policy on the Azure web replication?

A: I'm assuming that here they're asking if there's a functionality within LS Central that deletes replicated records or packages from the Azure storage. Currently, we don't have that. And we don't recommend you delete the packages directly from your storage. It needs to be coming from the system. So, currently, packages will be added to the Azure Storage. But we have plans to have this as part of our solution.

Q: The replication counter field is not included in some major tables such as customer sales header and sales line. Some customers need to replicate sales orders from POS to HQ or from HQ the POS. Please consider this a priority.

A: I would like to ask you to create a support ticket about this as a suggestion for a new functionality. And can I put the use cases in so we can address that and we can, give you straight answer on that?

Q: What could be the major bugs and errors in this functionality; that is, web replication? Sometimes web replication stops working. When I check the Schedule setup, I find some packets are stuck in the queue. After deleting the queue, the job works fine, but I am unable to trace what's is causing the issue.

A: Use the Web Replication Monitor to check for errors on individual packages. The scheduler log should provide insight into errors related to the LS Scheduler.

Q: Is the data encrypted?

A: This is probably different based on the replication method used. For Odata4 (direct update), not especially by the web replication itself, but could be by the HttpWrapper we are using. For Azure Storage we are using a function in the BC base application, I am not sure exactly how this is handled in these functions. For Company replication there is no encryption

Q: The setup for data replication involves several steps and can be complex. Is there a plan to simplify it?

A: And then there's the question, which part of it do you want to simplify? And it's very difficult. The scheduled job and the replication are quite simple when you're doing simple things. But once you want flexibility and other possibilities, then it gets very difficult to do it. Just by only going to one direction, push, it is simplifying everything according to that. And, I think, the web replication is simpler, because it has less possibilities. Then there's a question, should we have another schedule that is the simple one, which is, using the same underlying tables? I'm not sure that is the way to go, as it would be used both for web replication and the Data Director. It would be interesting to get your feedback on that.

Q: Web replication doesn't give enough detail when it fails. Is there any plan to make it more user friendly to give a meaningful error message?

A: We always have plans to improve, and it would be great to get suggestions from you on the matter.

Q: Is there monitoring of unprocessed tables within the Azure storage monitor?

A: This must mean unprocessed packages. When you export the data to a package, and you create one package that's going to be used for all POSs. You can define a threshold in the Web Replication Monitor, if location or POS has not checked if they have this package within the threshold defined the package will appear in the monitor. So, you should be able to get a task flagged right away. If there's a package that has not been asked about by a location. And then if they have asked about the package, they know about the package, but they have not processed it. So, it's a twofold threshold, and both can trigger. But there's not, something special that, you've created a package and no one has asked about it. It's all down to each distribution location that is supposed to receive it. So, in a way, you could say, yes.

Q: Is there a limit on how much data can be replicated at once? Sometimes the replication fails to send all data if there is too much.

A: We don't know of a limit, but we can control how big each package is so that the web service will not time out. There might be other reasons why it might have failed. So, I guess the easiest way would be to have a use case or scenario where it failed so that we can have a look and understand what the root cause was. Most of the time an error occurs not because there's a limit but because of other reasons like setup or connectivity issues.

Q: Does Azure storage replication support sub-locations?

A: Currently Sublocations are not supported in Web Replication. We are looking into whether we can include this in future releases.

Q: Can we use Application Insights with the new replication method? The package is treated, but the table within the package is not treated on the POS, and there's no alert. It's not visible in the monitoring.

A: Now this seems to be a specific scenario. If you have a ticket for it, please share it with us so we can address it.