
Introduction to LS Central Solution Architecture

Webinar Q&A May 29, 2020

Presenter: Bjarni Asmundsson, Regional Director

Q: Can you please share with us some more details on the localization project?

A: If we look at the LS SaaS solutions that are already up and running, then most of those localizations have been done by LS Retail. But there are couple of clusters where done by our partners, where they handled both the localization and testing. Of course, all localizations need to be pure/clean extensions and get thoroughly tested and whitelisted by Microsoft. For further information on this process, please contact our partner operations, pops@lsretail.com

Q: Where will mobile POS fit in this?

A: That is a good question, it fits precisely as the mobile Inventory APP. But potentially the solution to look at here is the LS Central APP, which connects directly to the service tier and can communicate directly with peripherals. Therefore it does not require LS OMNI service.

Q: Can you please clarify what is scale up vs scale out?

A: Scale OUT is adding additional instances or servers.
Scale UP is adding resources to the same server, more RAM, faster CPU etc.

Q: Regarding the KDS, does it need to have a separate server? Or can be installed together with HCCS?

A: The KDS can be installed on the same server as HCCS – but potentially you want to have it installed on separate hardware and in each restaurant if it is a restaurant chain.

Q: We are meaning to run a mobile POS, do we need Omni channel license?

A: The mobile POS APP needs LS Omni service, but there is no license for the LS Omni Server/service. The license needed is a device CAL license for the mobile POS.

Q: Do we have any calculator or tool through which we can calculate the sizing based on the requirement?

A: We do not have the calculator. A few years ago, Microsoft had calculation guidelines but discontinued it. There are so many parameters involved in sizing. We base our recommendations on experiences along with some raw data calculations, when we are involved in this process.

Q: Do you have any customers running an SQL Cluster on their on-prem setups and have they seen any negative performance with it?

A: We do have clients running SQL Cluster but so far have not seen performance issues directly related to SQL Cluster. Any performances issues we have come across, we can normally trace back to a certain piece of code or process, which could be optimized for certain clients. Or, in extreme cases, scale out, so the solution is running a different instance of SQL Server and distributing the load.

Q: Can we run more than 40 POSs directly from HO, client all stores in same geography location. They do not want to buy separate hardware for individual store please suggest!!

A: It is complex question, please email to bjarni@lsretal.com
We do have client running with this setup, but it depends on situation. If it is a supermarket then the answer is no. If it is a low volume retailer then it is an option.
It is not only about number of POSs, also about data volume, concurrent processes etc.
Therefore, there is no plain Yes or No answer.
In APAC, I can think of at least 2 clients with 60 plus clients/POSs connecting to the HO.

Q: Is there any change in the way KDS works in LS Nav 6.3 vs. LS Central?

A: Yes, our development department has made significant enhancements in this area and added a lot of functionality and options.
If you want to bring the new features to older versions it is possible, but you must downgrade the code in LS Nav and use new installers for KDS.

Q: In your setup, you mentioned that logging should be on a SSD drive. Would HPE 12G 10K SAS drives be good enough for this use case? It will be a medium + and an ETA database sizing of 1TB by year 3. Giving our WAN link will only be 50MB so we will only be able to receive 50MB at max

A: We would always recommend using good SSD, and at least Premium SSD on Azure.

Q: Can a KDS be routed to a KOT?

A: Yes, the KDS talks to either display system or KOT printer.

Q: This question may not be relevant, but we just want to know how stable the current version of LS Central is?

A: The current versions and releases of Microsoft Dynamics 365 Business Central and LS Central are good and stable.

Q: How can we minimize locking issues? What maintenance should be done to minimize this? Is this also related to the physical hardware, like the CPU and memory?

A: Again, this might be complicated to answer here, so please follow up in an email if needed.
Many different things can cause locking: undersized hardware, wrongly configured hardware, wrongly configured software on top of that. Start on this, and make sure it is good. After that, look into housekeeping: Index maintenance, free space in the database, VLFs etc.
And now we are first within LS Central – where we can as well see configurations that do not fit the situation, or not optimized planning on integration and/or maintenance jobs. Often, we can make good impact just reviewing those.
But in the end, we might end up reviewing code. Code optimization, which focuses on this clients / projects need, can definitely give us performance enhancements.
Moreover, as mentioned above, Scaling UP and/or Scaling OUT will most often give performance boost. When addressing these, you must already have identified the bottlenecks, so you are addressing this in the right area.

Q: By when can we expect a stable version? Because we are planning for migration to the latest version.

A: The current versions and releases of Microsoft Dynamics 365 Business Central and LS Central are good and stable.

Q: In NAV administration there are some parameters that may increase the performance. Do you have suggestion how to configure it properly for best results? Since, most of the time we just set up as standard. Also, please give us suggestions for the configuration of LS Data Director that can increase the performance of replication.

A: Currently there are not many parameters you can play with in the service tier – potentially Data Cache Size is the one that can give you the biggest impact. But isolating processes and scaling out by adding service tiers and VMs running them has also proven to give good results.

Regarding LS Data Director, there are a few points I would like to make. Potentially they are not directly for the DD, but within the process chain. Make sure you are replicating by actions / by preactions, not Normal. Think about avoiding getting too large packages. In some projects, you might want to put value in “Action Counter Interval” on the Scheduler Subjob, only to avoid that. I also always like to turn off “Move Actions” in the same page. Why send data that you do not need in the destination? Run the Scheduler Jobs frequently enough but do not oversaturate the system. I like to send out every 30-60 mins. as minimum, except potentially setup or configuration tables that should maybe go once a day.

Then in the LS Data Director Configuration tool, I like to review “Max Error Retry” and “Release Job Off Hold”. If the service is not running on its own server, review “Data Thread Priority”, you might want to lower it. Activate “Multi Job Processing” and review “Forward Thread Limit” and potentially “Data Thread Limit”. Always have a separate SDD under the service as it is rather I/O heavy.