



SAP NetWeaver on IBM i
Recommendations and Review

Version 3.1
August 2020

IBM Systems
IBM Corporation

Please direct questions or comments to:
Ron Schmerbauch
rschmerb@us.ibm.com



TABLE OF CONTENTS

- I Introduction
- II Highly Recommended Actions
- III Further IBM i OS and DB items
- IV Further SAP items
- V Infrastructure Data

I INTRODUCTION

The purpose of this document is to provide guidance regarding SAP environments running on IBM i. This material represents a collection of review steps based on SAP/IBM documentation, SAP Notes, IBM Redbooks, and experiences of the IBM i Development team.

This document should be helpful to improve SAP on IBM i performance and/or implementation but is not intended to cover all issues or situations.

The document is relevant across the SAP on IBM i Database and SAP Application Servers, which can be running on IBM i or Windows. The following releases are applicable:

- SAP NetWeaver with kernels 7.21, 7.22, 7.49 and 7.53
- IBM i 7.1 or after
- SAP components dependent on these releases (e.g. ECC 6.0, BW 7.0)

The commands described in this document will require the appropriate authorizations in order to execute successfully, such as the appropriate SAP authorization profile for SAP transactions, and appropriate privileges for IBM i and Windows commands.

SAP and all SAP product and service names mentioned herein are trademarks or registered trademarks of SAP SE in Germany and several other countries around the world.

IBM®, POWER®, and Db2®, are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

Other products may be trademarks or registered trademarks of their respective companies.

II HIGHLY RECOMMENDED ACTIONS

Take the following actions on all SAP on i partitions and systems to make sure currently recommended default settings for function and performance are all in place.

1) Load and apply the latest recommended PTFs known to help in SAP environments.

- Recommended [PTFs](#) for IBM i 7.4
- Recommended [PTFs](#) for IBM i 7.3
- Recommended [PTFs](#) for IBM i 7.2
- Recommended [PTFs](#) for IBM i 7.1

2) Take the appropriate SAP actions for any known issues with your specific IBM i OS release level as described in the SAP Note for that release. Some of these actions require that the SAP system is down.

- SAP Note [2786037](#) for IBM i 7.4
- SAP Note [2299407](#) for IBM i 7.3
- SAP Note [2011710](#) for IBM i 7.2
- SAP Note [1432783](#) for IBM i 7.1

When upgrading your IBM i release level, consult SAP Note [68440](#).

3) Check OS level settings as documented in SAP Notes [428855](#) and [49201](#). Some settings are dynamic and take effect immediately. Some will take effect at the next IPL. Set “Max Active” for the main storage pools used by your SAP instances high enough to keep “Wait->Inel” or “Active->Inel” at 0.0 when viewed from WRKSYSSTS. To see these settings, press F21 from WRKSYSSTS, set the assistance level to intermediate or advanced, then press F11.

Set IBM i system value QDBFSTCCOL with the CL command CHGSYSVAL QDBFSTCCOL *ALL. This is a dynamic setting and can be applied at any time. If resources are available, a QDBFSTCCOL task will collect DB statistics that are useful to the optimizer. If the system is busy to the point that resources are not available, the task will sleep and not disrupt operations.

4) Use WRKJRNA to display the current attributes of each SAP system. If the attributes do not match those of the CHGJRN statement below, use the statement to update them. Only change the settings when the SAP



system is down. CHGJRN JRN(R3<sid>DATA/QSQJRN) JRNRCV(*GEN) RCVSIZOPT(*RMVINTENT *MAXOPT3 *MINFIXLEN) MINENTDTA(*FILE) JRNOBJLMT(*MAX10M) **This statement represents the combined effect of SAP Notes [424866](#), [448552](#), [1449715](#) and [1366799](#). Specific PTF levels are required as documented in the Notes. Journal caching (option 42 of the operating system and the JRNCACHE(*YES) journal attribute) has little effect in SAP environments because Db2 commitment control forces journal entries to disk at each SQL COMMIT operation regardless of how many entries are in the journal cache.**

- 5) SAP JVM is the only JVM supported by SAP as of October 2012. See SAP Note [1907153](#) for details and relevant notes. JDBC connection options for SAP JVM are found in SAP Note [654800](#).
- 6) Review SAP transaction ST02 in each SAP ABAP instance. Hit ratios should be high (usually >97%) and swaps should be low (usually <100 per day). Poor ST02 settings can cause performance issues regardless of your HW resources, OS or DB platform. The following SAP Notes may be applicable depending on your release level and which buffers are swapping: [16244](#), [649613](#), [808607](#) and [2098347](#). Changes to SAP buffer settings take effect at the next SAP instance restart.
- 7) A performance benefit from reduced CPU consumption may be realized with kernel release 7.20 and after, especially on systems with larger numbers of processors. This is achieved by changing the ES/TABLE and related instance profile parameters for shared memory handling as described in SAP Notes [808607](#) and [2098347](#).
- 8) SAP kernel release 7.4x requires profile parameter updates when coming from prior kernel releases. Certain profile parameters no longer provide optimal performance and others become necessary to obtain optimal performance. SAP Note [2098347](#) summarizes necessary actions. A standalone enqueue server also becomes highly recommended as described in SAP Note [2013043](#).
- 9) Check SAP BW performance using the BW Sanity Checks tool to make sure critical DB indexes are in place for your SAP BW InfoCubes. Missing EVI indexes will result in suboptimal performance. Make sure you have the version described in SAP Note [2342139](#), then run the checks as described in SAP Note [541508](#), taking action as needed.

- 10) Consider using the flat InfoCube for BW infrastructure that is available for all SAP database platforms as of SAP BW 7.40 SP15 or SAP BW 7.50 SP4. Previously reserved for the SAP HANA database only, this simple internal data structure change can significantly improve BW load times and may also help query performance without requiring any changes to the BW application design. See SAP Notes [2073519](#), [2634508](#) & [2446889](#).
- 11) SAP NetWeaver 7.5 introduces significant new requirements when coming from prior NetWeaver releases.
- Unicode conversion is required prior to 7.5 as described in SAP Notes [1051576](#), [800791](#) and [928729](#).
 - Split dual stack systems are described in SAP Note [1797362](#).
 - Change central instances as described in SAP Note [2146940](#).
 - Convert pool and cluster tables as needed – follow SAP Notes [1892354](#), [2236867](#) and [1784377](#) for planning, estimating disk capacity needs and conversion requirements.

III FURTHER IBM I OS AND DB ITEMS

These checks and recommendations may apply to IBM i and/or Db2.

- 1) If you often import transports to running SAP systems which change the structure of tables, implement SAP Note [1952271](#) to enable an option that will make it far more likely that necessary SQL table locks are successfully obtained without encountering lock time out limits.
- 2) If multiple SAP systems on a partition are using the /QFileSvr.400 directory to access the same server in parallel, they may see a benefit by allowing multiple connections from the client to the server when the client is using IBM i release 7.2 or later. To enable parallel access to /QFileSvr.400, enter the following command on the client side:
ADDENVVAR ENVVAR(QIBM_RFS_CONNECTION_POLICY) VALUE('1') LEVEL(*SYS). The value will take effect once hostname directories in /QFileSvr.400 are recreated and jobs using those directories are restarted. See SAP Note [2202005](#).



- 3) Periodically check for any severe HW or OS issues in each IBM i partition with the CL commands DSPMSG QSYSOPR and WRKPRB.
- 4) Use IBM i CL command WRKDSKSTS. Check for unbalanced capacity or individual disk units that have far higher utilization rates than all others over a sustained period of time. Remedy with STRASPBAL, either with the *CAPACITY option or after running TRCASPBAL to collect information. It may be helpful to use STRDSKRGZ on occasion to consolidate disk space when you have reached a percentage used within an ASP.
- 5) To avoid performance problems by disk arm contention, we recommend keeping the journal receivers (e.g. in library R3<SID>JRN) in the same ASP as the database library (e.g. R3<SID>DATA) in most cases. Only for production systems without data replication we suggest a separate ASP for journal receivers, so that data recovery up to the point of failure is possible in case of disk failures beyond the RAID protection level of your system (see SAP Note [654801](#)).
- 6) To automatically monitor detachment of Journal receivers so they can be saved and deleted, see SAP Note [82079](#).
- 7) Periodically (quarterly or yearly) use SAP transaction DB02 to identify tables that have significant numbers of deleted rows and might be candidates for reorganization. Usually, the space occupied by deleted rows is being reused when new rows are inserted to a table, so in general it is not necessary to reorganize tables. However, after a significant archiving project, or if you have deleted SAP clients with no intention of replacing them it may be helpful to reorganize tables with a large number of deleted rows. Sizes of deleted data reported by DB02 can be overestimated because the calculations are based on the maximum row length size instead of on the actual size – but variable length fields may not be using the maximum length in all rows. Offline reorganizations are preferred whenever possible as online reorganizations have limitations and raise potential issues for heavily used files. See SAP Note [84081](#).
- 8) Periodically examine the list generated by SAP transaction DBACOCKPIT, Detailed Analysis, 50 Slowest Queries. Generally, the SAP indexing strategy is very good. Customer-written queries over SAP tables may use fields that do not have appropriate indexes. Most often customer created tables (they start with the letter Z) are not indexed properly, if at all. Check for these same Z tables in the System i Navigator Database Index Advisor and consider implementing an



appropriate index that includes fields from the WHERE, GROUP BY and ORDER BY clauses of the query. More than one index may be useful for a given table to provide benefits to multiple queries.

With kernel 7.20 or later, SAP logs additional information (program, user, workstation, transaction) with each statement that is sent to DB2. These can be used to quickly identify the program or user that is responsible for a long running SQL statement. See SAP Note [2196060](#) for more details.

- 9) Periodically check for objects and jobs that may be approaching database and operating system limits. This may be done in the DBA Cockpit when using SAP NetWeaver 7.0 Enhancement Package 2, SAP NetWeaver 7.3, and higher. Refer to SAP Note [2244176](#). If deleting or archiving data is not an option to reduce object sizes, table partitioning can be a solution as described for non-BW systems in SAP Note [2187681](#) and for BW systems in SAP Note [815186](#).
- 10) If you are using Independent ASP (IASP) or Cross Site Mirroring (XSM), review SAP Note [568820](#). SAP integration when using IBM PowerHA is described in SAP Notes [1635602](#) and [2390472](#). When using IASP, make sure there is available capacity in the System ASP for temporary storage such as Db2 temporary objects, the Db2 plan cache, and shared memory for each SAP instance. It may be helpful to increase the Plan Cache size manually or to allocate more storage to the system ASP.
- 11) If your system is using 3-tier application servers on IBM i, see SAP Note [834218](#) and if using a 3-tier application server on Windows see SAP Notes [705886](#) and [1644051](#).
- 12) If you are using IBM Db2 Mirror for i, follow the requirements and configuration options described in SAP Note [2808345](#).
- 13) Use the free [IBM Insight tool](#) to periodically capture throughput information specific to your workload. Make sure to record “normal” workload periods and “peak” workload periods such as month-end or year-end processing. This data can be valuable input to your next sizing exercise and can provide a historical record of past performance for future comparisons, such as after a software or hardware upgrade.

IV FURTHER SAP ITEMS

The following checks and recommendations may be of benefit within the SAP application layer.

- 1) Keep up to date with SAP Host Agent patches on all partitions using SAP applications. See SAP Note [1031096](#).
- 2) Regularly examine the SAP Alert Monitor through SAP transaction RZ20 for any reported alert conditions.
- 3) Review the SAP Table Buffer configuration and performance through SAP transaction ST10:
 - Are tables assigned to wrong buffer? This can occur when a table is assigned to the SAP Single Record Buffer and is accessed primarily through an ABAP SELECT statement. The SELECT does not use the Single Table Buffer. This will result in a high number of calls to the database server.
 - Tables with very high update counts may not be good candidates for buffering.
 - Non-buffered tables with mostly read access may be good candidates for assignment to the Single or Generic Table Buffer.
- 4) Check that SAP collector and cleanup jobs are running. Use transaction SE36 and choose “Standard Jobs” as described in SAP Note [16083](#).
- 5) Check that the SAP ABAP performance improvement to RSDB4UPD described in SAP Note [977676](#) has been implemented.
- 6) Review the work process number and configuration through SAP transactions SM50, SM51, and SM66 to determine if there are too few or too many allocated work processes.
- 7) Use transaction DB02 to check that the SAP Data Dictionary is consistent with the SQL objects that are in the database. See SAP Note [117219](#).
- 8) Exclude libraries SAP<sid>IND and SAP<sid>Inn from backup operations. If you do attempt to save these libraries, timestamp



updates may lead SAP believe it needs to rebuild some directory information when in fact it does not need to do so. These libraries will be recreated from IFS directories at an SAP instance restart if they do not exist, so they do not need to be saved. See SAP Note [825473](#). Also exclude the library SAP<sid>IND and SAP<sid>Inn from your HA data replication software processing to make sure backups complete successfully. See SAP Note [1637588](#).

V INFRASTRUCTURE DATA

Document the architecture implemented for the complete SAP on IBM i solution. The following information should be considered. Record changes as they are made so any effect can be traced back to the change. This data may be useful when explaining your environment to consultants or service representatives.

- 1) IBM Power Systems Hardware Configuration. Power Systems models, number of active and inactive cores.
- 2) IBM i LPAR Configuration. For each server, the number of LPARs and the processor, memory and storage resources associated with each of them. Also include virtualization environments: Shared processor pools, capacity and virtual processors assigned to each partition, and any VIOS partitions.
- 3) If applicable, External storage configuration.
- 4) If applicable, additional Application Server Hardware Configuration:
 - Windows: server model, number of processors, amount of memory. Windows version.
- 5) Host names for all LPARs and application servers
- 6) SAP System IDs (SIDs) and instance numbers
- 7) ABAP/JAVA dual stack or separate stacks
- 8) SAP BASIS Release
- 9) SAP Kernel Release
- 10) IBM i OS Version and Group PTF levels