



Exploit Storage Resources for SAP on IBM i

A summary of options to boost performance and reduce the storage footprint of SAP systems running on IBM i

Ron Schmerbauch

*IBM Systems and Technology Group ISV Enablement
May 2016*



Table of Contents

Document Change History	1
Abstract	2
About the Authors	2
Prerequisites	2
Introduction	3
Recommendations	5
Software Options	6
SAP housekeeping to remove unnecessary data.....	6
SAP Basis Tables Clean-Up	6
Deleting Client Copies	6
Batch Input Processing Clean-Up	6
Change and Transport System Clean-Up	7
SAP Data Management Guidelines and Archiving	7
Other options to clean up SAP data	7
IBM i: Journal settings to reduce journal receiver sizes	8
DB2 for IBM i: Automatic reuse of deleted rows	8
DB2 for IBM i: Automatic index balancing	8
Hardware Options	9
Hardware Tape Compression	9
Parallel Tapes	9
IBM System Storage Copy Services.....	9
Mirroring	9
FlashCopy	10
Real Time Disk Compression	10
Summary	11
Resources	11
Trademarks and special notices	12



Document Change History

Version	Changes
October 2012	Initial version
May 2014	Correct transaction name to SM36.
June 2016	Adjust authors list, note SVC compression support allows for more HW options beyond V7000

Abstract

This paper provides a summary of the various options that are available to reduce the size of an SAP system running on the IBM i operating system and DB2 database. Performance of SAP applications and system maintenance tasks will improve due to the reduced data volume. Also described are hardware based solutions to improve the performance of such tasks.

About the Authors

Ron Schmerbauch is the technical leader of the SAP on IBM i team, part of IBM i Development in Rochester. He has many years of experience with SAP on IBM i, beginning with the original porting effort of SAP R/3 to OS/400® in 1994. Before joining the SAP team, Ron worked in Database development. He holds a degree in Computer Science from the University of Southern California. He can be reached at rschmerb@us.ibm.com.

Prerequisites

You are running SAP systems on the IBM i operating system and DB2 for i database platform.

You are looking for ways to reduce the storage footprint of your SAP systems.

You want to backup and restore your systems faster and use fewer tapes.



Introduction

Over time, SAP applications store ever larger amounts of data in file and database systems and can eventually present challenges to any server platform. For example, the database size can have an impact to the overall system performance each day. Specific operations such as building indexes over large tables can take more time when more data is present. Not only the performance of maintenance tasks like backup and restore is affected by the file system and database size, also SAP release upgrades and enhancement package installations will run faster with less data in the system.

With SAP NetWeaver 7.5, certain files may grow in size as a result of manual “declustering” processes, raising the priority of efficient handling of large tables. More information about this may be found in [SAP Note 1892354](#).

There are several opportunities available to address these types of challenges. In general, one can either reduce the amount of data that is handled or one can more efficiently handle the data. The following sections of the document summarize those options. They are grouped by software and hardware provided options. If you are looking for a solution to a specific problem, the following table quickly guides you to the appropriate sections.



Page	Section	Ways to... reduce disk requirements	speed up application performance in general	speed up save operations	speed up restore operations	reduce number of tapes required for backups	speed up system copy/clone	avoid database reorganizations
6	SAP Basis Tables Clean-Up	x	x	x	x	x	x	x
6	Deleting Client Copies	x	x	x	x	x	x	x
6	Batch Input Processing Clean-Up	x	x	x	x	x	x	x
7	Change and Transport System Clean-Up	x		x	x	x		
7	SAP Data Management Guidelines and Archiving	x	x	x	x	x	x	x
8	IBM i: Journal settings to reduce journal receiver sizes	x	x	x	x	x		
8	DB2 for IBM i: Automatic index balancing	x	x	x	x	x	x	x
8	DB2 for IBM i: Automatic reuse of deleted rows	x	x	x	x	x	x	x
9	Mirroring			x			x	
9	Hardware Tape Compression					x		
9	Parallel Tapes			x	x		x	
10	Real Time Disk Compression	x					x	
10	FlashCopy	x		x			x	

Recommendations

The options available and summarized in this paper differ in the amount of storage saved, effort required and performance gained. Some options may not apply to your SAP system. If you don't want to do everything listed in this paper, we recommend that you perform the options in this order:

- 1.
2. [DB2 for IBM i: Automatic reuse of deleted rows](#) and [DB2 for IBM i: Automatic index balancing](#) are taken care of by the DB2 for i database.. You still want to read the sections to know what it is about.
3. Check your [IBM i: Journal settings to reduce journal receiver sizes](#).
4. Do [SAP Basis Tables Clean-Up](#) and [Change and Transport System Clean-Up](#). If you do batch inputs to your SAP systems, also do [Batch Input Processing Clean-Up](#).
5. Consider [Deleting Client Copies](#), if you have any unused clients.
6. Not enough? Follow the instructions for [SAP Data Management Guidelines and Archiving](#). This will probably generate the most in storage savings and a noticeable performance boost.
7. If you are still looking for fewer tapes and/or faster backup times, consider the hardware provided [Hardware Tape Compression](#), [Parallel Tapes](#), [Mirroring](#) and [FlashCopy](#).
8. Consider [Real Time Disk Compression](#) for minimizing the amount of disk storage required for the remaining data.

Software Options

This section summarizes SAP, IBM i operating system and DB2 for i database provided options to reduce the volume of data that is being used by the SAP systems.

SAP housekeeping to remove unnecessary data

There are a number of background jobs pre-defined in an SAP system that will do the basic housekeeping in your SAP system. These jobs need to be scheduled to prevent certain files and database tables from growing unnecessarily. Starting with SAP release R/3 4.6C, the standard jobs can be scheduled by calling transaction SM36 and clicking on the 'Standard Jobs' button. Then click on 'Default Scheduling' to schedule all standard clean-up jobs. SAP note [16083 - Standard jobs, reorganization jobs](#) contains more details about the standard jobs. Also for SAP releases prior to R/3 4.6C this note contains instructions how to schedule the standard jobs manually. Check note [1411877 - New standard jobs](#) to verify that you've scheduled all necessary jobs.

Besides the standard jobs, there are more areas in an SAP system that need clean-up from time to time. Those areas are listed in the following sections.

SAP Basis Tables Clean-Up

SAP systems are using a number of database tables to store administrative data. Those called basis tables may grow in size over time and affect system performance. There are more than 50 different categories of basis tables like application log tables, IDoc and RFC related tables. Each table category has its own clean-up method, ranging from simple deletion of data to specific archiving procedures. SAP note [706478 - Preventing Basis tables from increasing considerably](#) contains the full list of basis table categories and the according clean-up methods. Most of the clean-up methods are described in their own SAP note which is listed in note 706478.

Deleting Client Copies

Client copies occupy space in the database. If there are client copies that are not needed, consider deleting those to save space. Existing clients in your system are listed in transaction SCC4. Clients can be deleted with transaction SCC5. SAP note [70643 - CC-TOPIC: Client Deletion \(SCC5\)](#) describes known issues when deleting a client copy.

Batch Input Processing Clean-Up

Batch input processing can be used to import data from non-SAP systems into SAP. When this is done periodically, the according log files and administrative table entries can occupy a large amount of disk space.

Batch input log files can be cleaned-up by running report RSBDC_REORG, which is a newer version of report RSBDCREO. The report has an input screen that lets you choose batch sessions by name and time window to reorganize batch input files. SAP note [147354 - Batch input: Reorg. and delete sessions and logs](#) describes how the report works. SAP note [18307 - Batch input logs and reorganization](#) describes more technical details of the batch input log reorganization report RSBDCREO.

The APQD table is used to store batch input session information. This table can grow significantly over time if you have many incomplete batch input sessions or sessions marked as “keep”. SAP note [36781 - Table APQD is very large](#) describes preventive measures how to keep table APQD small and actions how to reduce its size when the table is large already.

Change and Transport System Clean-Up

The change and transport system which is used to distribute changes between various SAP systems is using files in the file system and database tables to store administrative data and the transport requests itself. Without regular clean-up, the transport file system directories and database tables can become very large.

When you import many transport requests individually instead of using the mass transport feature, those individual transports request remain on the import queue as already imported. The import queue grows in size over time, and among other effects displaying the queue takes longer. SAP note [565574 - Deleting imported entries from TMS import queue](#) explains how to remove entries from the import queue that are already imported.

Transport data in the file system that is no longer needed, like log files and old transport requests, can be deleted by calling the TP program with the ‘clearold’ option. SAP note [41732 - Deletion of data in transport directory](#) contains detailed information about the procedure.

SAP Data Management Guidelines and Archiving

SAP systems often contain many more years of history than are required to satisfy the current needs of the business. The SAP Data Management Guidelines document provides comprehensive information about how to handle application data growth. It describes archiving strategies for the most common SAP Business Suite applications ERP and CRM and database tables that are common to all SAP NetWeaver based applications. The document can be downloaded from the Information Lifecycle Management section of the SAP Service Marketplace (service.sap.com/ilm) under the Data Archiving topic.

A special topic is reducing the data volume of profit center accounting. The following SAP notes cover size reduction options for profit center accounting:

- [203545 - Archiv. profit center: Analysis of table GLPCA](#)
- [178919 - Table GLPCA: How to reduce the data volume?](#)
- [91615 - Archiving/profit center: new archiving programs](#)

Other options to clean up SAP data

The following SAP notes contain information about further clean-up options:

- [1401395 - Report to clean up residual data in user tables](#)
- [16513 - File system is full - what do I do?](#)

IBM i: Journal settings to reduce journal receiver sizes

Database journals are used to keep track of changes to database table entries. The journal entries are stored in journal receiver objects. Although the journal receiver objects are usually saved and then deleted from the system periodically, they still may occupy larger storage space between the save and delete points. The journal settings RCVSIZOPT and MINENTDTA influence the size of the journal receiver objects. To minimize the journal receiver sizes in one easy step that never needs to be repeated, use the IBM i CHGJRN command to set RCVSIZOPT(*MINFIXLEN) and MINENTDTA(*FILE). For further information about the journal settings see SAP note [1449715 - IBM i: Current recommendations for journal configuration](#).

DB2 for IBM i: Automatic reuse of deleted rows

When deleting records from a database table, the free space will be reused by the DB2 for i database automatically, unless the “Reuse deleted records (REUSEDLT)” option is set to *NO for that table. The default setting for all SAP tables is *YES for REUSEDLT. Therefore no reorganization of database tables is necessary during normal system operation.

If you have deleted many rows from the database, for example during archiving or client deletion, you may want to immediately regain the free storage space. Then you can use the Reorganize Physical File Member (RGZPFM) command. SAP note [84081- iSeries: Reorganization of database tables \(RGZPFM\)](#) contains important information about how to use the RGZPFM command in an SAP environment.

DB2 for IBM i: Automatic index balancing

As data in a database is changed and deleted, the associated indexes are updated to reflect the data changes. Depending on the order that data is added and changed, the underlying data structure (typically based on a binary tree) can become unbalanced. On other systems, an administrator would manually re-balance the underlying data structure by reorganizing the system.

Although there is no action necessary to enable it, you should be aware that on the IBM i platform, the database engine is continually monitoring the underlying index structure and rebalancing the index structure internally when necessary. This allows your applications to continue running without having to take the system down for index reorganization.

Automatic rebalancing also makes sure that each index is no larger than necessary.

Hardware Options

This section summarizes options that are provided by hardware solutions to address the challenges posed by SAP systems with large data volumes.

Hardware Tape Compression

If reducing the number of tapes that must be managed or used for each save operation is desired, hardware tape compression can be an answer.

Many tape devices support data compression, which increases the apparent capacity of your media by encoding the data to use less space. The data is compressed and decompressed by the hardware each time it is read or written on your tape device and is not apparent to applications.

Depending on the device, compression ratios are often in the 2:1 or 3:1 range. Therefore it can be possible to reduce the number of required tapes by half or more by taking advantage of devices featuring hardware tape compression.

Parallel Tapes

If reducing the time it takes to perform a save or restore operation is desired, one option is to use more than one tape device in parallel. It is possible to use the tape compression techniques described in the previous section together with parallel tape strategies.

One way this can be done is manually by using multiple save or restore operations, each operating on specific objects concurrently.

IBM i Backup, Recovery and Media Services (BRMS) may be used to control automated methods of saving and restoring multiple objects. Two strategies exist, parallel-parallel or parallel-serial. Parallel-parallel will split individual objects across tapes, while parallel-serial is keeping objects intact on a tape, but saving multiple objects in parallel.

All three methods are described at: [www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/IBM Backup, Recovery and Media Services \(BRMS\) for i/page/IBM i parallel save and restore?lang=en](http://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/IBM_Backup,_Recovery_and_Media_Services_(BRMS)_for_i/page/IBM_i_parallel_save_and_restore?lang=en)

IBM System Storage Copy Services

IBM System Storage copy services may be used to address more than one aspect of large SAP data volumes.

Mirroring

Mirroring can be used to create and maintain an exact duplicate of a set of disk volumes, even for large data sets. Once a copy is created, it can be taken offline so that operations one might want to perform on it that can take a significant amount of time such as a save operation can be performed without impacting the performance on the original copy. Mirroring is provided by many IBM System Storage solutions. Refer to www.ibm.com/systems/storage for most current information.

FlashCopy

FlashCopy can be used to quickly create an exact duplicate of a set of disk volumes, even for large data sets. When creating a copy of an existing system, a flash copy will be much faster than performing a save then a restore operation. FlashCopy functionality for IBM i is provided by IBM System Storage solutions Storwize V7000, SAN Volume Controller, XIV and DS8000. Further storage systems may support FlashCopy in the future. Refer to www.ibm.com/systems/storage for most current support information.

Real Time Disk Compression

If you are interested in reducing the amount of disk capacity required by SAP components on an IBM i LPAR, you might consider compressing the disk volumes associated with that LPAR. Real time disk compression reduces the amount of required disk space without significant impact to the I/O response time for most workloads. Compressing an entire LPAR is easy to do and, best of all, requires no ongoing administration effort.

Some database platforms offer file based compression. This requires choosing which files in the database to compress, followed by an ongoing monitoring effort to make sure the efficiencies are not lost over time. Administrators on other platforms must also manage the frequency of table and tablespace reorganizations, which generally require downtime.

The disk compression rate for a typical SAP ERP system is between 50 and 70 percent. It should be noted that compressed disk volumes remain compressed when flash copied. However, when data is delivered to the server for processing it is decompressed. In this case hardware tape compression could be used to recompress the data as it is saved.

The real time disk compression feature is currently available for IBM Storwize Family members, most prominently V7000, and storage devices behind an SAN Volume Controller(SVC). The IBM Redpiece "Real-time Compression in SAN Volume Controller and Storwize V7000", available at www.redbooks.ibm.com/redbooks.nsf/RedpieceAbstracts/redp4859.html?Open, contains more detailed information about how to use the compression feature.

Summary

On any platform, SAP solutions can eventually result in very large databases and/or file systems.

The IBM i platform automatically works to minimize the amount of data capacity necessary, most notably with automated index balancing and the reuse of deleted records.

SAP administrators have opportunities from within the SAP application to reduce the size of the database and file system by either archiving or deleting data that is no longer necessary to have in the active system. They can also easily minimize the size required by the IBM i journal receivers with little effort by running a IBM i command.

Hardware options such as disk compression, tape compression and parallel tape drives can also be employed to handle large data volumes as efficiently as possible. FlashCopy and Mirroring technologies provide options for working with snapshots of large datasets in an offline manner.

Resources

These websites provide useful references to supplement the information contained in this paper:

- The SAP developer community network (SAP SCN) offers a SAP on IBM i space with discussion forum, blog articles and other featured content for SAP on IBM i: scn.sap.com/community/ibm-i
- The IBM information page for SAP on IBM i contains latest news and links to the current info APAR for SAP on IBM i: www.ibm.com/developerworks/ibmi/sap
- An IBM Redbook providing a technical deep dive into all aspects of a SAP on IBM i installation is available from the IBM Redbooks website at www.redbooks.ibm.com/abstracts/sg247166.html



Trademarks and special notices

© Copyright IBM Corporation 2012. All rights Reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is



presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.