0.1 Standard Operation - Technical Customizing Settings

It is possible to set parameters that apply only to a **specific archiving object**.

The technical customizing settings will be set via SAP transaction SARA -> Button: Customizing -> Archiving-object specific customizing -> Technical Settings.

Archive File Size	**Recommended settings**
Max. Size in MB	100 – 300
Max. No. of Data Objects	Blank
Settings for Delete Program	**Recommended settings**
Not scheduled/Start automatically/After Event/	Not scheduled
Build Index	Not activated
Place File in Storage System	**Recommended settings**
Content Repository	to be maintained by the customer
Start Automatically	Activated
Sequence	**Recommended settings**
Sequence	Store Before Deleting

The column 'Recommended settings' gives information about the recommended settings by SAP. Dependent on the archiving objects other settings can be recommendable e.g. regarding performance aspects. You will find the specialized recommendations in the single archiving chapters of the analyzed areas.

1. **Archive File Size**

- **Max. Size in MB**

It is recommended to define the archive file size with 100 to 300 MB as the interface to Archive Link could cause problems if the size is higher than 1 GB. We recommend choosing a file size of maximum 300 MB, as 300 MB can still be handled well, are easy to transfer to an external content server. On the other hand, a larger file size of 300 MB will reduce the number of archive files being created and therefore will reduce the number of delete jobs that have to be scheduled, which may simplify the scheduling but which may limit the options for parallelization of delete jobs.
Note that the maximum size of an archiving file is limited by the operating system and also by the external storage system if one is connected.

- **Max. No. of Data Objects.**

It is recommended to leave the field 'Number of Objects' blank. You only should operate with the parameter 'Size of Archive File [MB].

Please note: The value that is reached first triggers the creation of a new archive file. If you leave both fields blank, only one archive file is created.

2. **Settings for delete program**

- **Start automatically**

This radio button is of interest primarily if you want to save the archive files before the delete phase.
Please note:

a. If there is a long time period between the write phase and the delete phase, there is a danger that the data could be changed before it is deleted, which means that the data in the archive and in the database are no longer the same. The delete phase should therefore be carried out as soon as possible after the write phase.
b. If you activate the Start Automatically radio button, the delete program is run automatically immediately after archiving. However, if it was specified in Customizing that the file storage is to take place before the delete phase, the delete program is not started until the file has been stored.
c. Since the deletion program is automatically started for each archive file, you can carry out delete and write jobs in parallel if relatively small archive files are created. This parallel activity can have a positive effect on the runtime of an archiving session, as the database is used more efficiently. If a file is too small, the number of processes rises and the system is negatively affected.

3. **Place File in Storage System**

- **Content Repository**
  The name of the content repository is maintained here. You can find more information about content repositories in the SAP Library under Content Repositories.

- **Start automatically**
  Indicates that archive files, after successful processing, are automatically transferred to a connected storage system.
  If you mark this check box, an archive file is automatically stored in the Content Repository that was specified in Customizing.

- **Sequence**

  - **Delete phase before storage:** The file is stored in the content repository after it has been processed by a live deletion program. If the deletion program is run in test mode, the file is not then automatically stored.
  - **Storage before delete phase:** The file is stored in the storage system after the write program has created an archive file but before the deletion program has been run for this file.

0.2 House-Keeping

In SAP S/4HANA, the technical job repository carries out the scheduling of periodic technical jobs in the ADMIN and BUSINESS clients. This mechanism is carried out automatically by the system and does not require any user interaction, unlike the earlier "Standard Jobs" function in SM36.

Among other things, the technical job repository contains job definitions for the standard jobs mentioned in SAP Note 16083.

Transaction SM36 in SAP S/4HANA systems has the function "Job Repository", which branches to transaction SJOBREPO, instead of the function "Standard Jobs".

Transaction SJOBREPO displays an overview of all job definitions delivered by SAP in the technical job repository.

You can use SJOBREPO to change a job definition (and the background job generated from it) within certain limits or to completely deactivate it.

Please see SAP Note 2190119 “https://launchpad.support.sap.com/#/notes/2190119 for further details.

0.3 Sources of further information

Book from SAP Press: “Archiving SAP Data – Practical Guide”. ISBN: 978-1-4932-1279-8 (English).

SAP Trainings
- BIT615 SAP Archive Link Document Management
- BIT660 Data Archiving
  (overview about Data Archiving and its tools)
- BIT665 Information Lifecycle Management (SAP NetWeaver ILM)
- BIT670 Data Archiving – Programming
  (especially applicable to integrate an archive-access in Z-Reports)
- WDE680 Data Retention Tool DART (only for Germany)

For details see www.sap.com/education

0.4 Data Volume Management Service Portfolio

The data volume management service portfolio of SAP (DVM) helps you to set up a data volume management strategy that defines how to manage and reduce future data growth and the existing database size by following a holistic approach that considers and integrates the following options: data avoidance, data aging, data summarization, data deletion, and data archiving.

While implementing data avoidance and summarization are one-off activities affecting future data growth, data deletion, data archiving and data aging are recurring operations aimed at reducing the existing data volume. For HANA databases the data aging concept offers you the option of moving large amounts of data within a database to reduce the memory footprint.

SAP's service and support portfolio for Data Volume Management (DVM) consists of:

- DVM Analysis & Monitoring - cloud application for S/4HANA

- DVM Guided Self-Service

- DVM Continuous Quality Check (CQC)

- DVM Enterprise Support Value Map

- DVM Workcenter and Fiori Apps in SAP Solution Manager

- DVM Engineering Service (ESRV)

Here are some details for each of these building blocks:

DVM Analysis & Monitoring:
The DVM Analysis & Monitoring application is a cloud-based application that allows you to monitor you HANA-based system to get insights into the DVM topic such as:
- An evaluation of the reduction potential for Memory and Disk.
- Applicable Reduction measures with a reduction potential based on SAP’s recommended residence times.
- Simulation tool to see the impact of setting your own residence times.
- Time-based and Accurate reduction potential calculations for Archiving, Aging or both.
  - “Accurate” reflects the results after the business checks imposed by archiving or aging programs have been applied.
- HANA Memory Usage
- HANA Disk Usage
- Saving Potential on Disk
- Saving Potential in Memory
- Time Based data distribution overview for all analyzed tables
- A suite of HANA specific analyses including…
  - Tables with more than 1 Billion Records
  - Trop Growing Tables
  - Top Shrinking Tables

- Largest Tables
- LoB Overview
- Saving Potential using Inverted Hash Keys for Indexing

The cloud application is available through the SAP One Support Launchpad at the following address:
https://supportshell-supportportal.dispatcher.int.sap.hana.ondemand.com/#/dataoverview/

DVM Guided Self-Service:
SAP supports your implementation of a data management and data archiving strategy with a Self-Service for Data Volume Management (DVM), which is an SAP tool-based approach powered by SAP Solution Manager. The self-service generates a best practice document that describes how to handle your largest data objects using the methodologies of data avoidance, summarization, archiving, and deletion. The best practice document shows the type and amount of data that can be archived or deleted as well as the corresponding archiving objects or deletion reports. SAP's recommended residence times are used to calculate possible savings.

To assist you in using the self-service and to reduce the learning curve, SAP offers an 'Expert Guided Session' in the form of remote training. For a detailed schedule and registration, see the 'Expert Guided Implementation Calendar'. https://service.sap.com/Expert-Guided-Implementation

DVM Continuous Quality Check (CQC):
This service uses a similar technical infrastructure as the Guided Self-Service (see above) to generate a service report but offers more flexibility as it is performed by an SAP service consultant together with you. For more information, see: https://support.sap.com/content/dam/library/SAP%20Support%20Portal/support-programs-services/supportprograms/enterprise-support/academy/delivery-format/cqcis/cqcdvm.pdf

DVM Enterprise Support Value Map:
The SAP Enterprise Support Value Map for Data Volume Management is a social collaboration platform offered by SAP Enterprise Support. The value map provides information on each step involved in SAP Data Volume Management from initial assessment, through implementation, to improvement. It provides details of the Enterprise Support services that can assist you and also a forum where you can ask questions and create discussions. SAP Focus Advisors who have experience with DVM are available to assist you with your queries.

Other customers also participate in the value map, so they may be in a position to share their experiences and areas of shared interest with you.

If you are interested in joining the SAP Enterprise Support value map, either register at https://support.sap.com/valuemaps or contact your local SAP Enterprise Support Advisory Center.

DVM Work Center and Fiori Apps in SAP Solution Manager:
The Data Volume Management Work Center in SAP Solution Manager offers capabilities to gain insights into the source of data volume movements in single and especially in multisystem landscape environments. The solution is based on SAP NetWeaver BW and provides a holistic landscape overview of your data. This analytical and reporting infrastructure from SAP provides you with a rich array of functions and features that allow you to:
- Get transparency of system landscape data consumption at different levels
- Reveal potential for creating/optimizing a data volume management strategy
- Leverage SAP Best Practices to drive your data volume management strategy
- Simulate different data volume scenarios (for example, moderate versus aggressive archiving)

- Provide a compliance check of a corporate data volume management strategy

The workcenter is accompanied by new Fiori Applications with Solution Manager 7.2 SP05. These new Apps provide insights and transparency on the following:
- Data Aging App (only for HANA systems)
- Unused Data App
- BW NLS and Unused Infoprovider
- BW Housekeeping and Administration

SAP offers an 'Expert Guided Session' for setting up the Data Volume Management Work Center.
For a detailed schedule and registration, see the 'Expert Guided Implementation Calendar'
(https://service.sap.com/Expert-Guided-Implementation).

Enterprise Support Value Map Data Volume Management:
The SAP Enterprise Support Value Map for DVM is a social collaboration platform offered by SAP Enterprise Support. The value map provides information on each step involved in SAP Data Volume Management from initial assessment, through implementation, to improvement. It provides details of the Enterprise Support services that can assist you and also a forum where you can ask questions and create discussions. SAP Focus Advisors who have experience with DVM are available to assist you with your queries.

Other customers also participate in the value map, so they may be in a position to share their experiences and areas of shared interest with you.

If you are interested in joining the SAP Enterprise Support value map, either register using the following menu path:
SAP Service Market Place SAP Support Portal -> alias ESACADEMY (http://support.sap.com/esacademy) -> Value Maps -> Join Now (register for DVM) or contact your local SAP Enterprise Support Advisory Center.

0.5 Background information on Data Aging

The SAP HANA database uses the concept of partitions. For the purposes of this document we will refer to two partitions i.e. the Current partition and the Historical partition. The Current partition resides in the “in-memory” area of the HANA Database and is intended for “operationally relevant” data. The Historical partition resides on disk and is intended for “non-operationally relevant” data. The data in the Historical partition is still accessible via SQL on request.

The main goal of aging is to reduce the main memory footprint and improve the speed of database queries. This is achieved by keeping only operationally relevant data in the current partition. The historical data is placed on (less expensive but usually slower) secondary storage.

In other documentation you may find reference to “Current data / HOT area” – this is the in-memory area. Also, where you read “Historical data /COLD area” this refers to the historical partition. HOT and COLD are old terminology but for a transition period they will co-exist.

Data Aging

Data Aging offers you the option of moving large amounts of data within a database to reduce the memory footprint.

Data Aging requires application knowledge on how to move data from the Current area to the Historical area The move influences the visibility of the data and its accessibility. Typically, only the in-memory area is accessed. This results in much shorter run times when you run queries of large amounts of data. To be able to apply Data Aging to your data, you need to fulfill certain requirements regarding the database and the application.

Note: When designing the Historical partitions for storing the aged data, it is really important to know the volume of data that can be expected by timeframe/partition. Avoid creating partitions that will later have to store more than one billion records.

The following information should provide a general overview and basic knowledge on Data Aging.

Technical structure of Data Aging

Data Aging is a Suite-tailored concept for reducing the SAP HANA* memory footprint = A new, cost-saving volume management capability for Big Data making sure that:

• Only operationally relevant (Current data) data is loaded into the SAP HANA main memory

• Other (Historical data) data remains primarily stored on disk, not affecting Current data performance, yet Historical data remains accessible via SQL on request.

To achieve a substantial memory reduction, data residing in the Current data area of SAP HANA should be limited to “open” items only, e.g., only a couple of months of actual data.

There might be exceptions when “old” data is needed to be loaded in the Current area of SAP HANA. However, such situations should be minimized as much as possible.

Audits and reporting requires that Historical data in SAP HANA must be included in analytical SQL queries.

To achieve sufficient performance for OLAP (Online Analytical Processing), the Historical storage needs to be optimized for OLAP processing (disk-based columnar storage).

Moving data from Current data to Historical data requires application knowledge, it is not pure DB processing. This is related to the complexity of identifying a document as “complete”.

The completeness of such a document (and so the move from Current data to Historical data) depends on several criteria including the status of predecessor and successor documents. The lifecycle of successor or

predecessor documents is determined by several criteria including dependent documents (contracts, preliminary documents) and, e.g., quantities completely received.

Transition from Current data – Historical data

Prerequisites for Data Aging

To be able to apply Data Aging on your data, you need to fulfill various requirements:

- You are using a database that supports Data Aging.
- You have switched on the Data Aging functionality.

You can switch on the business function only if the database requirements for Data Aging have been fulfilled and the application server ABAP (AS ABAP) supports additional system requirements. Note that even before you activate the business function, the system checks for these requirements.

- The SAP application provides the necessary enhancements per business objects with Data Aging object functionality.

Data Aging Procedure

As mentioned earlier, Data Aging mainly contains functions that support moving business data from the database to the disk i.e. from Current area to Historical area.

The Data Aging procedure can be divided schematically in two different functional areas, which are as follows:

Design Time

Before you can start to execute Data Aging, you need to have an overview of your tables and decide how you want to partition them.

- First, you need to create Data Aging objects, and enhance them if applicable. To do so, the tables participating in Data Aging are determined by the SAP application and assigned to a Data Aging object or an enhancement. You will find a list of Data Aging objects and enhancements provided by the individual SAP applications in the transaction Data Aging Objects (DAGOBJ).
- Prior to moving the tables intended for Data Aging into partitions, these tables must be assigned to a partitioning object. Usually, SAP applications suggest partitioning objects that you can apply. For more information about the available partitioning objects, see transaction Customizing for Partitioning (Transaction DAGPTC) under Edit Partitioning Objects.

- If the assignment of a table to a partitioning object does not fulfill your requirements, you can create customer-specific partitioning groups without making modifications. Partitioning Groups override the assignment of tables to partitioning objects if a partitioning object as well as a partitioning group contains the same table. To decide whether the partitioning objects delivered by the SAP application are sufficient, you first need an overview of the tables and then you can decide how you want to partition them. It is critical when defining the historical partitions, to be aware of the expected volume of records by timeframe/partition. Avoid creating partitions than need to store more than 1 billion records.

Runtime

After you have defined how you want to partition the tables during design time, you can start to run Data Aging.

The following information is relevant for running Data Aging:

- You use the application-specific runtime class to determine the data for which Data Aging is intended. The SAP application assigns these runtime classes to the relevant Data Aging object so that the runtime class can be called and processed in a Data Aging run.
- To be able to move the data from the Current area of the database to the Historical area according to the specified partitioning objects and partitioning groups, all of the participating tables must be partitioned for Data Aging. All of the tables in a Data Aging object of the SAP application must be enhanced with the _DATAAGING column for this. The added column is the basis for a RANGE partition with which the data field is represented:
  - 00.00.0000 or NULL stands for the HOT partition
  - Time period between the intervals 01.01.0001 to 31.12.9999, represent the COLD partitions.

The following activities are relevant for running Data Aging:

- For each system, you need to define the partitions for the corresponding tables of a Data Aging object. If the conditions are not fulfilled, the Data Aging run is not started. You use transaction Manage Partitions (DAGPTM) to partition tables.
- After the partitions have been defined, you can activate the Data Aging object. The system runs through various checks so that the Data Aging object can be used for the run. Choose transaction Data Aging Objects (DAGOBJ) to activate the Data Aging object.
- Finally, schedule the Data Aging run for your tables. Create Data Aging groups for this purpose. Select transaction Data Aging Runs (DAGRUN) to schedule the runs.

Data Aging Transactions

An overview of all transactions available for Data Aging is listed in the table below.

Transaction Code	Name of Transaction
DAGOBJ	Data Aging Objects
DAGPTC	Customizing for Partitioning
DAGPTM	Manage Partitioning
DAGRUN	Overview of Data Aging Runs
DAGADM	Managing Data Aging Objects
DAGLOG	Data Aging Logs

To see the list of tables included in an aging object, use the transaction DAGOBJ and display the details of the corresponding aging objects.

0.6 Additional Information

Examination of tables
During this data volume strategy session, all the tables are examined for:

- Data avoidance
- Data aging
- Data summarization
- Data deletion
- Data archiving

Only those sections for which action is necessary or possible are listed in the report.

Data Volume Management prevents your database from growing unnecessarily. If archiving or deletion is seldom scheduled, or not at all, this will increase the data volume and result in a performance loss and longer database maintenance times (backup, restore/recovery).

If you do not have a data volume management strategy, set up a project to develop one. Since this takes some time, you should start as soon as possible.

All table entries that have exceeded the residence time should be archived or deleted during the cleanup phase. During the operational phase, the archiving runs should be executed on a regular basis.
Make a backup copy of all tax-relevant data before any data is archived. Select a storage format that the tax authorities can access. Take into account the country-specific legal requirements of your tax authority.

The SAP Data Retention Tool (DART) is available in the USA and Germany. In Brazil, transaction IN68 provides comparable functionality. DART gives access to tax-relevant data, which meets the requirements of these tax authorities. For more information, see the alias "DART" on SAP Service Marketplace.

For analyzing tables prior to archiving, see SAP Note 317219 (transaction TAANA).

Interfaces to other systems:
Running your system in a solution landscape with systems such as ERP, CRM, BW, or other systems requires you to check the possible impact on these systems when following the recommendations in the DVM report. Since the applicability of DVM recommendations depends on the customer scenario, this DVM report cannot consider potential interface problems.

SAP Notes:
Before starting any data volume management activities (data avoidance, aging, summarization, deletion, and archiving), ensure that you have applied the **latest relevant SAP Notes for your current release, or have applied the latest Support Package.**

Use the name of the archiving or aging object to search for SAP Notes.
Use the table name and the appropriate term, such as 'Avoidance', ‘Aging’ to search for SAP Notes relating to avoidance, data aging, summarization, and deletion.