

Planning Guide

This chapter covers the following topics:

- General Guidelines for Planning ASF
 - Practical Guidelines for Using ASF
 - Miscellaneous Questions and Answers
-

General Guidelines for Planning ASF

The use of ASF requires thorough planning. With ASF you will be storing statistical database information over a period of weeks or months, and then analyzing the stored data. Therefore it is of great importance to plan which data to store and which types of Evaluation Report are required.

The material presented in this section describes the stages of the planning process for using ASF.

Checklist for using ASF

The main procedures involved in using ASF are presented in the following list, and subsequently described in detail:

- Plan which databases and files are to be stored.
- Plan which Store Type should be used (and create if necessary).
- Create the Store Profile.
- Start the Store Program (Screen/batch).
- Plan which ASF fields are to be evaluated.
- Decide which type of Evaluation Report you wish to generate.
- Create the Evaluation Profile, specifying units and accumulation.
- Run the Evaluation.
- Make a Workplan entry (if required).

Plan which databases and files are to be stored

When an Evaluation Report is generated, it presents an analysis of ASF data which has been stored for particular databases and files over a period of days, weeks or months. It is therefore essential to plan thoroughly the databases and files for which ASF data is to be stored. Storing data for many databases and files has the advantage of providing a greater amount of stored ASF information for later analysis, but has the disadvantage of using larger quantities of disk space to store the data. The ASF Installation Notes show examples of the amount of disk space required in several typical instances. For this reason, the selection of databases and files for storing is normally limited to those which are directly required to produce an Evaluation Report.

Plan which Store Type should be used

When storing ASF data with the Store Program, the data is stored with an associated Store Type. The Store Type is an indication of the circumstances under which the data is stored. ASF has several predefined Store Types, e.g. AH (ad hoc), EN (when the Nucleus is terminated), ST (when the TP monitor is started), ET (when the TP monitor is terminated), CY (cyclic, i.e. based on time of day), DA (daily). You can define additional Store Types if required. The full list of all Store Types supplied is presented in the section Supplied Store Types.

When stored ASF data is retrieved for an Evaluation Report, the Store Type is used as a selection criterion. Thus in an Evaluation Report, it is not possible to combine ASF data records which were stored with different Store Types.

Create the Store Profile

Once you have planned the databases and files for which you wish to store ASF data, and the Store Type to be used, you can create the Store Profile using the "Store Profile Administration" function of the ASF Online Menu System.

Start the Store Program (Online/Batch)

If ASF data is to be stored cyclically, e.g. with Store Type CY (cyclic) or a user-defined Store Type, then the Job Scheduler will have to be updated accordingly to run the Store Program cyclically in batch mode at the appropriate time intervals. Alternatively, the batch job can be started online from the TP monitor (for example, using the COM-LETE "UTIMER" command). In either case, the job which has to run is a Batch Natural job which uses the ASF direct command STORE, specifying the name of a Store Profile and Store Type (see the section Batch Operations in the *ASF User's Guide* documentation for details).

The Store Program can also be started online using the "Store Nucleus Records" feature of the ASF Online Menu System. If the program is started in this manner, the appropriate Store Types are AH (ad hoc) or a user-defined Store Type.

Plan which ASF fields are to be evaluated

The statistical data which the ASF Store Program stores, are the values of ASF data fields. The full list of the ASF data fields is given in the section ASF Group and Field Names. The ASF data fields are subdivided into several groups, each related to some aspect of the database such as physical disk usage, number of Adabas commands issued for the database/file or session parameters (memory usage).

Decide which type of Evaluation Report you wish to generate

When using ASF, you must know which result you are trying to achieve (see the section Examples of Evaluations). The aim of ASF is to produce Evaluation Reports for monitoring databases. An Evaluation Report summarizes database performance data which has been stored over a period of days, weeks or months using the ASF Store Program. There are 10 different types of Evaluation Reports, each representing the stored data in a different manner. The Evaluation types 1-8 summarize stored data in the form of two dimensional or three dimensional tables. Reports generated using these Evaluation Types differ in the units displayed in the various dimensions.

The following table summarizes the combinations of databases, files, times and ASF fields which are used in Evaluation Types 1-8. This is valid when the report is displayed on the screen, printed or downloaded to the PC in Full format.

Screen, Printer, Download "Full"				Accumulation	
Type	Vertical	Horizontal	Third axis	Internal	External
1	Fields	Times	Databases		x
2	Accumulated Fields	Databases	-	x	
3	Fields	Databases	Points in time		x
4	Times	Databases	Fields		x
5	Fields	Times	Files		x
6	Accumulated Fields	Files	-	x	
7	Fields	Files	Points in time		x
8	Fields	Files	Fields		x

When the data is downloaded in standard or CSV format, the result is a two dimensional table. It looks the same for all Evaluation Types. Accumulation values are never calculated, not even for Type 2 or 6.

Download "Standard" or "CSV"				Accumulation	
Type	Vertical	Horizontal	Third axis	Internal	External
1-4	Times, Databases	Fields	-		
5-8	Times, Files	Fields	-		

Evaluation Type 9 displays a list of ASF fields whose last measured values exceeded some user-defined bounds, and Evaluation Type 10 is a list of ASF fields whose values will reach or exceed user-defined limits if current growth rates continue.

A summary of the Evaluation Types is presented in the *ASF User's Guide* in the section Evaluation Reports.

Create the Evaluation Profile, specifying units and accumulation

An Evaluation Profile contains a definition of databases, files and ASF data fields which are to be evaluated. An Evaluation Profile is normally designed to access either the same databases and files as an existing Store Profile, or a subset of the databases and files. In addition, an Evaluation Profile contains the names of ASF data fields which are to be evaluated. See the section Evaluation Profile Administration in the *ASF User's Guide* for details of how to create an Evaluation Profile.

Run the Evaluation.

When the Store Program has run at least once, ASF data is available for analysis and output in an Evaluation Report. The more often the Store Program runs, the more data is available for determining trends in databases and comparison between databases.

Make a Workplan Entry

If you plan to run any given evaluation often, you can save all of the required information for running the evaluation in a Workplan entry. In this way, a frequently used evaluation can be started simply by selecting the appropriate entry from the list of predefined evaluations in the Workplan. See the section Workplan: Predefined Evaluations and Reports in the *ASF User's Guide* for details of how to use the Workplan.

Practical Guidelines for Using ASF

The information presented in this section is based on experience gained in using ASF in production environments.

Data Storage

The most important aspect is to collect data at timed intervals in batch. This guarantees the availability of current and historical data for Evaluations as required.

Storing data once daily is sufficient as a basis for long-term monitoring. The daily storage should be scheduled to take place towards the end of the typical online system usage (for example, about 18:00), and should include all databases and files. For this purpose, ONE Store Profile with dynamic file selection in all databases should be used.

The advantage of this dynamic Store Profile is the automatic detection of active files; you do not have to modify Store Profiles explicitly when new files are created or when existing files are deleted.

The Store Type of this daily store should be "DA" (daily). The Store Type DA is included in the ASF distribution kit.

As described above, the creation of a sufficiently detailed database is catered for by a daily run of the Store Program. Data can, however, also be stored for the purposes of documentation or gap-free database history, or as a basis for certain special evaluations:

- A timer-controlled weekly storage (Friday evening) with its own Store Type "WE" (weekly). The Store Type WE is included in the ASF distribution kit. This data can be used for evaluations on a weekly basis, and can also serve as a basis for evaluations for periods of time for which the daily records have been deleted or transferred to a backup medium.
- A storage at the nucleus termination (Store Type "EN"). This is the basis of gap-free nucleus documentation, and is used for Evaluations of the nucleus termination data of a database (history evaluations 1 and 5). The availability of nucleus termination data is important if you plan to work with delta values. Delta values can be obtained with the ASF User Exit or with ASF Utilities.

For collecting nucleus termination data, you should use the same Store Profile as for the daily runs. The parameter "Reduced on DB-ID" must be set to the ID of the database to be monitored at nucleus termination.

Alternatively, if you are only interested in database information (i.e. no file parameters) at nucleus termination, you should create a new Store Profile which contains databases but no files. This has the advantage that at nucleus termination, only 1 record per database is created. However, on systems offering round-the-clock operation with only a few shutdowns per year, the file information should be stored also.

Summary of guidelines for data storage:

The Store Program should run daily as a timer-controlled batch job. It should be started towards the end of the peak period (e.g. about 18:00) with the Store Type "DA" (daily). In the Store Profile, the "dynamic" feature should be selected for all databases. In some cases it can be advantageous to store data at other times (e.g. at nucleus termination).

Critical Reports

The most important data output by ASF are the Critical Reports. These represent, in compact form, values which are important because their values are "critical". The distribution kit for the ASF product contains predefined Critical Reports called "RED", "YELLOW" and "BLUE" as part of the Workplan. When the Critical Report "RED" displays values, e.g. Extents ≥ 4 (for an Adabas version 7 database) or DATA USED $\geq 95\%$, your immediate intervention is required. Experience has shown that this Critical Report should run immediately after the Store Program runs, so that the any critical values displayed refer to the current database status.

The output can either be sent directly to the printer or be queued by the Operating System. You can send the report as e-mail or to Con-nect cabinets as described in the section Critical Report. In this way you can reduce your daily database monitoring effort to the viewing of a single compact list.

The Critical Report "YELLOW" shows situations which do not require your immediate intervention, e.g. buffer efficiency less than 10, or DATA USED greater than 80 percent.

The Critical Report "BLUE" shows unused resources, e.g. number of Normal Index blocks unused greater than 20000 blocks, or high water marks less than 10 percent.

According to the database setup or to the safety requirements, you can choose to run these reports immediately after running the Critical Report "RED".

Critical Trend Reports

You should plan Trend Evaluations carefully. There is always an element of uncertainty involved, and events which are predicted do not necessarily have to become reality. The value "accuracy", which appears in the Trend Evaluation, acts as an indicator of how reliable the trend values are. Accuracy values of greater than 90% indicate that the trend values are fairly reliable from a statistical point of view.

It is good practice to run a Critical Trend "RED" as part of the "Friday evening" job. The trend should be based on data stored weekly during the last three months (i.e. the relative dates "-90" to "+0"), or data stored daily in the last month (relative dates "-30" to "+0"). The trend limit, i.e. the amount of time for which data should be predicted, should be set to two months, i.e. the relative date "+60". The interval between successive predictions within this period should be set to 7 days.

You must generate trend records if you subsequently want to display the projected development of databases and files as a General Evaluation (e.g. type 1 or 5). You generate trend records using the "Store Trend Records" menu. In the input field "DBs and FIs from" of this menu, you should enter the name of the Evaluation Profile which will later be used to evaluate the trend records. This ensures that trend records are generated for only those databases and files which will appear in a subsequent evaluation. If the Evaluation Profile does not specify databases, use the Store Profile.

General Evaluations

A useful General Evaluation is a database history (Evaluation type 1) every Friday evening, listing the database development over the past few weeks, or the database development over the last 4 weeks with Store Type "WE". For this evaluation, it is best to use the relative dates "-5" to "+0". The evaluations with Evaluation Type 4 and 8 are also very useful. They show for example the development of a parameter over the past 50 days.

You as the DBA should plan which General Evaluations you require and when you require them. You should avoid generating printouts of evaluations which include all databases and files over a long period of time.

Download and Graphical Presentation

If you plan to present the data graphically, choose or create an evaluation which contains all the fields and databases/files to be presented. If you want to download database related data, choose a type 1 – 4, for file related data choose type 5 – 8. Download the data in CSV format (PC-File = "C"), which has the highest download performance. Edit the downloaded data with a spreadsheet tool like MS Excel. If you want to compare various databases or files for one point in time, sort the data by Time/DBID/FILE-ID. If you want to present the history of a specific database/file value, sort the data by DBID/FILE-ID/Time.

Summary of guidelines for Evaluations:

Generate the Critical Report "RED" daily, if possible as part of the job which does the daily run of the Store Program. If required, run the Critical Reports "YELLOW" and "BLUE" also. On a weekly basis (Friday evening), do an evaluation of database development. Download the data in CSV format for graphical presentation.

Maintenance of the ASF Records

ASF records can be moved to an external medium (refer to the EXPORT option in the "Maintain Nucleus Records" menu). Using this feature, all records older than, say, 3 months can be copied to a work file and subsequently deleted from ASF-DATA. To do this, you should create a job with two steps:

1. Copy the records to a work file, using the dates "2000-01-01" to "-90".
2. Delete the records in this period of time, using the Delete function of the "Maintain Nucleus Records" screen.

Summary of guidelines for maintaining ASF Records:

The number of ASF records should be regulated by exporting and deleting records regularly. Exported records can be imported again if required.

Workload and Performance

The following general rules apply:

- Evaluations of type 1, 3, 5 and 7 are relatively fast.

- Evaluations of type 4 and 8 require more time than those of type 1, 3, 5 and 7.
- Evaluations of type 2 and 6 can be lengthy, depending on the number of stored records for the given period of time.
- A Critical Report (Evaluation Type 9) can be lengthy, depending on the number of databases and files in the Evaluation Profile.
- A Critical Trend Report (Evaluation Type 10) can be very lengthy. It is dependent on the number of databases and files in the Evaluation Profile and on the number of stored records for the given period of time.
- Downloading the ASF data in the Standard format is in general faster than in the Full format. Moreover if your target is an Excel file, the CSV format is considerably faster than the Standard format. Therefore it is recommended to use the CSV format if possible. See also the Download Performance test below.

When using ASF, the following points should be noted:

- Use external accumulation only when necessary. If an evaluation which uses external accumulation produces a table of results which is wider than the screen, then when you page to the rightmost screen, ASF must re-evaluate the data from ALL of the horizontal pages in order to calculate the external accumulation values. Thus, external accumulation can cause a lot of database accesses.
- The mathematical quantities DET (coefficient of determination) and DISP (coefficient of dispersion) should only be specified in an Evaluation Profile if they are really required, since ASF needs to read all the relevant data twice to calculate them. Under normal circumstances these two values are not of interest!
- Organize Evaluation Profiles and Evaluations to present information clearly. You could, for example, design "Application Reports" to evaluate files and databases which are required by certain applications.
- Avoid using dynamic Evaluation Profiles in General Evaluations, otherwise the generated reports can be enormous. A database containing 100 files can generate a report 10 or 20 pages wide and thus be completely impractical.

Dynamic Evaluation Profiles can, however, be useful for creating Critical Reports in batch. In this case, all stored databases and files are checked for critical values.

- When generating a set of trend records, you should specify the same Evaluation Profile in the "DBs and FIs from" field as you will later use to evaluate the trend records (provided the Evaluation Profiles specifies databases and files at all). The number of databases and files in this profile should be kept small.
- If the number of records stored in a file is not of interest, the parameter "Get Num.rec.loaded" in the User Profile should be set to "NO". This can speed up the storing of data quite significantly.
- Running ASF Evaluations regularly is an integral part of database management. You should therefore schedule your Evaluations to run automatically as batch jobs.

Download Performance Test

The test processes example data stored with NEW-TEST-DB-ALL. The "Full" and "Standard" download use as target an Excel spreadsheet.

Property	Value
Evaluation profile	SAG-CMDS-3
Evaluation type	3
From time	1992-09-01 , 00:00
To time	1992-09-30 , 23:59
Store profile	NEW-TEST-DB-ALL
Store type	DA
Origin	NU

Download Format	Time
Full (F)	671 sec
Standard (X)	44 sec
CSV (C)	1 sec

A similar run with the ASF 7.1 "Header" (H) format required more than 50 minutes.

Summary of guidelines for Workload and Performance:

- The Evaluations of type 2, 5, 9 and 10 depend on the number of stored records in the period of time which is being evaluated, and also on the number of databases and files in the Evaluation Profile.
- Regularly used Evaluations should run in batch. External accumulation (min, max, sum, avr and especially dis and det) can be time consuming, depending on the volume of data to be processed.
- Consider setting the parameter "Get Num.rec.loaded" in the User Profile to "NO".
- Use the CSV download format.

Miscellaneous Questions and Answers

This section contains hints concerning the usage of ASF. The problems listed in this section are based on customer experience gained in using previous ASF versions.

Question:

How does ASF work in batch?

Answer:**With Natural Security**

Firstly, create a user called BATCH in Natural Security, and link the group ASFGROUP to it. Using the ASF User Maintenance, allocate all privileges to this user. You can also assign a user to a job - in this case, the user name must be the same as the job.

Without Natural Security

The user must have the same name as the job. Batch users should set the "Keep Environment" parameter to "NO" in their profile; otherwise fields can be filled unexpectedly. Forced I/Os are not catered for in batch, but nevertheless the parameters "Limit CPU units" and "Limit ADA-calls" should be set to 9999999.

The printer name specified in the user profile is not used for batch output, but the output medium for batch should be the hardcopy printer. This means that the second dataset contains the printer output. If you select the screen as the output medium, the output is written directly into the job output stream.

If the output medium for batch is the printer, the printer name specified in the user profile is not used for the batch output. The second dataset contains the printer output. If you select the screen as the output medium, the output is written to the printer output as well. If the PC-File has been selected as output medium, the output is written into the Workfile 7 in the same format as it is downloaded online.

Question:

How does the syntax of an ASF batch evaluation looks like?

Answer:

Use as input for CMSYNIN:

```
LOGON SYSASF
MENU
,SELECT nn
.
,FIN
```

where *nn* is the number of the Predefined Evaluation to be performed.

Question:

Why can "Origin" not always be specified as "ALL" ?

Answer:

If tabular evaluations are generated using "TR" (trend records) or "ALL" (trend records as well as nucleus records), there is only a limited selection of ASF fields available in the Evaluations. Normally therefore, only "NU" should be specified for Origin, except in cases where trend records are to be displayed.

Question:

Under which circumstances does ASF use artificial I/Os, and can this lead to problems on some systems?

Answer:

Certain Evaluations (e.g. Evaluations of type 2, 6, 9 and 10) can require a lot of CPU time and thus exceed various TP monitor time limits. To solve this problem without increasing the TP monitor time limits, ASF generates "artificial" I/Os. These are I/Os which cause no change to the screen, but cause the TP monitor to reset its timers.

With IMS DC, however, this leads to problems, since screen updates are managed in a queue and sent asynchronously to the terminals. A bypass for IMS DC is to set the "Limit ADA-calls" and "Limit CPU-units" parameters in the User Profile to 9999999 and run the Store Program in batch only.

Question:

Under which circumstances does the User Maintenance work with Natural Security?

Answer:

User Maintenance works under Natural Security in each of the following circumstances:

1. If the SYSASF library is "people protected" in Natural Security, or
2. If the SYSSEC modules listed in the section Step 5: Setting up User Security in the *ASF Installation* documentation are present, or
3. SYSSEC is defined as a Steplib, or
4. If Adabas Online Services has been installed with INPL and Natural Security has been installed again with INPL, or
5. If your user ID is included in the group ASFGROUP. If the user ID is included in the group, the user ID must not also be linked directly to ASF (the Natural Security command "LINK ... TO ..." links a user to an application). Access to SYSASF may only occur via the ASFGROUP group.

Question:

Why is the cursor position not recognized in BS2000/OSD (e.g. when linking files by positioning the cursor on the database then pressing PF4) ?

Answer:

Before the cursor position can be recognized in BS2000/OSD, you must issue both of the following terminal commands:

```
%T=9756          (activate the terminal driver), and
%KN              (activate the Siemens function key logic)
```

What to do next ...

The steps outlined in this section serve as a checklist for planning and implementing ASF applications. The next section looks in detail at one of the standard Evaluation Profiles, namely SAG-IO-2. This is one of several standard profiles which are distributed with the ASF product. The section examines how SAG-IO-2 was created and shows the format of an Evaluation Report which was generated using this profile. You can copy SAG-IO-2 and modify your copy to fit your requirements, or simply follow the techniques which are described in the section.

The complete list of standard Evaluation Profiles is presented in the section Supplied Evaluation Profiles. If the profile SAG-IO-2 is not directly suitable for your purposes, you might find that one of the profiles listed in this section is more suitable as a starting point. You should not modify any of the standard profiles - instead, you should always make a copy of the profile and modify the copy.

You might also wish to refer to the diagrams shown in the section Evaluation Report Formats. These summarize the formats of the Evaluation Reports which are generated by using different Evaluation Types.