

ASF Overview

This chapter covers the following topics:

- Introduction
 - ASF Components
-

Introduction

The database monitoring tasks which the Database Administrator (DBA) performs include making regular checks of the current status of the database, e.g. disk and memory utilization, as well as performing long term planning, such as ensuring that the future disk space requirements can be met, based on current trends. The main sources of information for the DBA have been provided in the past by the ADAREP (Database Report) Utility, the Nucleus End Session protocol and ad hoc inquiries with the Adabas Online Services or the Adabas Manager. By using these facilities, the DBA can check the status of individual databases and files. This is often a very time consuming process.

ASF provides an environment which opens up a new dimension in database monitoring:

- ASF automates the process of database monitoring.
- ASF has no impact on the overall database performance.
- ASF is the long-term Adabas monitor for production environments.

The two main components of ASF are:

- A data collection program which collects database status information during an active nucleus session. This program is called the Store Program. Normally the Store Program is scheduled to run at regular intervals (for example once per day) over a period of many weeks or months, in order to collect data which can be statistically evaluated. The Store Program can also be started by the DBA on an ad hoc basis, using commands in the ASF Online Menu System.
- A set of programs to evaluate the statistics gathered by the Store Program and to output summary reports to the screen, the PC or a hardcopy printer. The programs are called Evaluation Programs, and the reports are called Evaluation Reports.

In the evaluation programs, the DBA can

- evaluate the collected data for planning and tuning purpose;
- compare multiple databases in one report;
- report critical database situations (the report can be send as an e-mail);
- extrapolate current trends;
- evaluate extrapolated values; and

- calculate the time when a critical value will be reached.

ASF has approximately 450 criteria called data fields for monitoring Adabas databases and files. The data fields represent many different aspects of an Adabas database, including disk and buffer usage, thread usage, database load, ADARUN parameters, pool usage and frequency of use of particular Adabas commands.

The primary data collection (counting of commands, etc.) is done in Adabas for ADAREP, AOS and the End Nucleus Statistics anyway. There is no special data collection for ASF in Adabas. Therefore ASF has no impact on the overall database performance!

Database information can be collected with the ASF Store Program at the start of a nucleus session, at the end of a nucleus session, and during a nucleus session. The start and end nucleus data, when accumulated over periods of weeks or months, gives an indication of the long term database growth, and permits projections of future database requirements. The nucleus performance data, such as main memory usage and pool usage, permits the DBA to perform tuning analysis of the Adabas nucleus parameters.

ASF Components

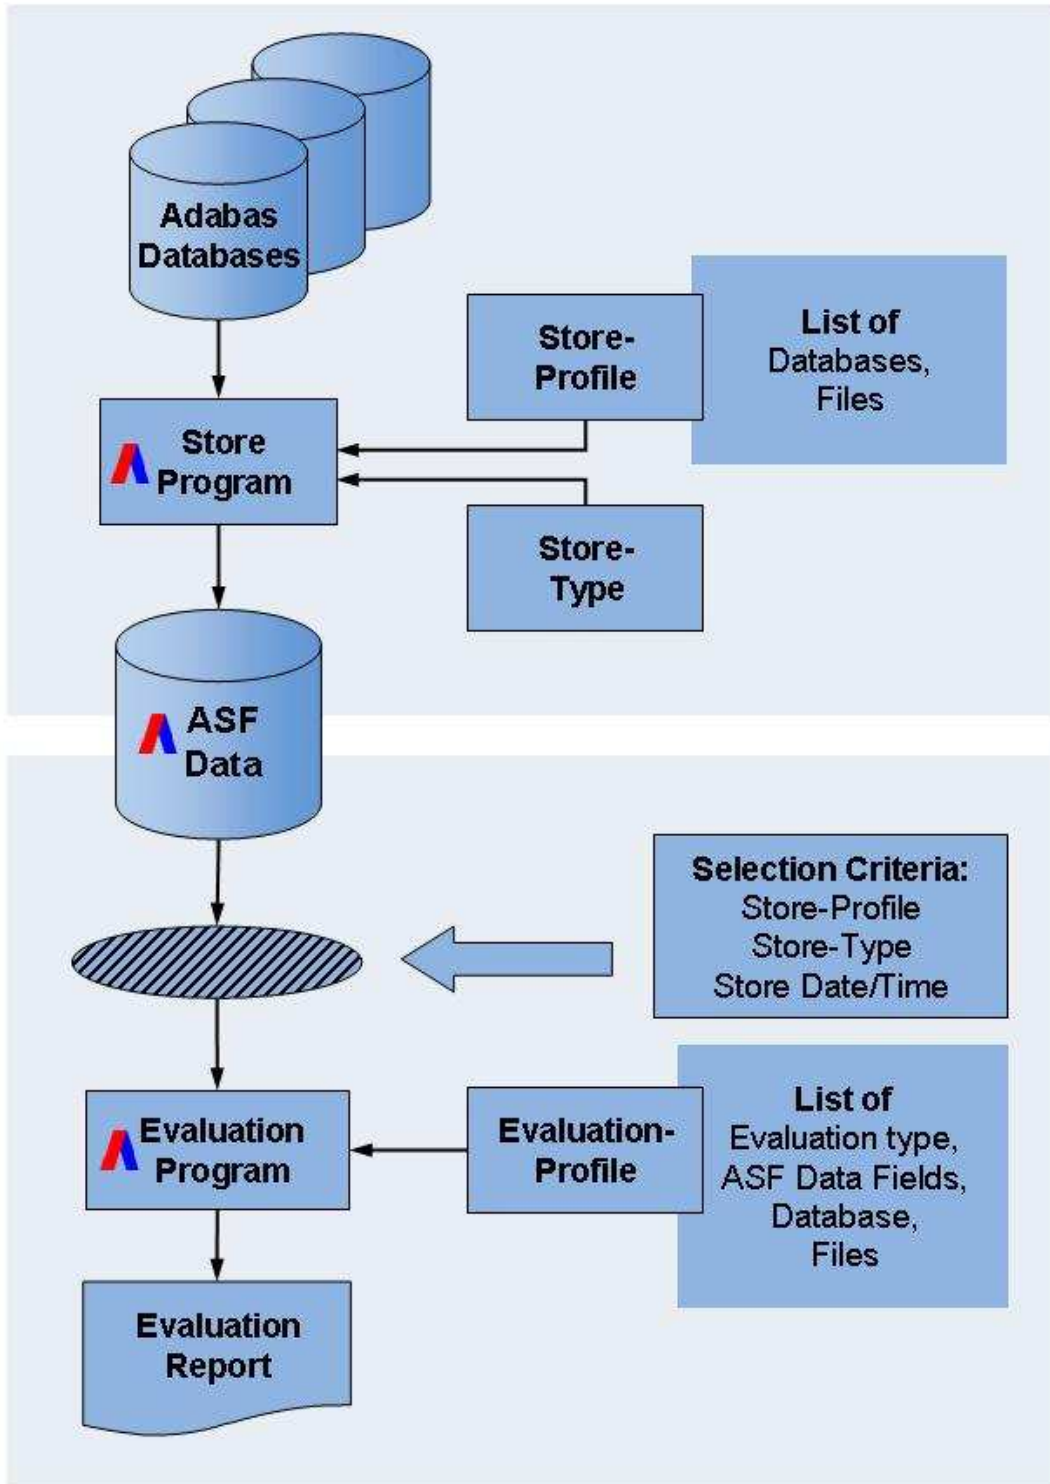


Figure 1-1: The ASF Components

This section main components of ASF are described in the following sections:

The Store Program

The Store Program is the component of ASF which accumulates database and file status information. The Store Program is usually run by the job scheduler as a batch program at regular time intervals, e.g. once per day, and collects information concerning groups of databases and files specified by the DBA. Data which is gathered over a period of several weeks or months in this manner can then be used as a basis for the evaluation reports produced by ASF.

The Store Program collects information concerning the files and databases defined in a Store Profile which is defined by the DBA. The DBA can define many Store Profiles, each specifying a different set of databases and files. When the Store Program runs, it requires as input, amongst others, the name of the Store Profile which contains the IDs of the databases and files to be monitored. The DBA can run several Store Programs concurrently, each with a different Store Profile.

Store Types

The data collected by the Store Program can differ greatly depending on current database activity. For example, the amount of the ASSO database component in use depends on the number of users accessing the database, so this value will be very different at nucleus startup than during normal database operation. For this reason, data collected by the Store Program is classified into several Store Types. When the Store Program is started, the Store Type is supplied as an input parameter. Further information concerning Store Types is given in the description of the "Store Nucleus Records" menu.

Evaluation Report

ASF generates reports which evaluate data collected by the Store Program. The reports summarize current and projected database statistics. The main types of reports are:

- General evaluation: an analysis of the past and present database status. The statistical tables which are generated give an overview of the status of various databases and files. The maximum and minimum values in rows of the output tables can be displayed also, as well as statistical quantities such as the sum and average of the values.
- Trend evaluation: tables of projected statistics, in time steps of days, weeks or months, until a specified end date.
- Critical Report: a report of databases and files for which ASF data fields have reached or exceeded "critical limits", as defined by the DBA.
- Critical Trend Report: a report of databases and files for which ASF data fields will reach or exceed their critical limits within a time frame, based on an extrapolation of current trends.

The programs which generate the Evaluation Reports are run by the DBA, using the appropriate options in the Online Menu System. When one of these programs is run, data which was stored by the Store Program is evaluated and displayed in either screen or hardcopy format, or is stored as a file on a PC.

Store Profiles and Evaluation Profiles

Store Profiles are defined by the DBA using the ASF Online Menu System. Each Store Profile specifies a list of databases and files for which the ASF data fields are to be collected. When the Store Program is run, the name of a Store Profile must be supplied to it as an input parameter. The Store Program collects data only for those databases and files specified in the Store Profile.

Note:

The Store Program stores the values of all ASF data fields for each database and each file specified in the corresponding Store Profile. The selection of ASF data fields to be included in an Evaluation Report is specified in the Evaluation Profile.

Evaluation Profiles are defined by the DBA using the ASF Online Menu System. In an Evaluation Profile, the DBA specifies the following:

- The Evaluation Type, which is a value in the range 1-10. See the section Evaluation Reports for details.
- One or more of the available ASF data fields (database monitoring criteria) which are to be analyzed in the Evaluation Report.
- The databases and files to be evaluated which is normally a subset of the databases and files defined in a Store Profile. If no database is specified in the Evaluation Profile, ASF takes all databases and files for which data is stored by the Store Program.
- The units of measurement of the ASF data fields.
- Statistical operations ("accumulation") to be performed on the report values for General Evaluations.
- Upper and lower values which represent "critical" levels for the data fields for Critical Reports and Critical Trend Reports.

When starting an Evaluation Report, the DBA specifies the name of the Evaluation Profile to be used.






Figure 1-2: Store Profiles, Evaluation Profiles

An example of the relationship between Store Profiles and Evaluation Profiles is shown in Figure 1-2. When the Store Program is run using the Store Profile *s1*, it will collect data for files 1 and 2 in database 25, and files 10 and 80 in database 123. It will also collect summary data for the whole databases 25 and 123. The Evaluation Profile *e1* specifies a subset of the data specified by the Store Profile *s1*, namely file 1 in database 25. The Evaluation Profile *e2* specifies a subset of the data specified by the Store Profile *s1*, namely database 25. The Evaluation Profile *e3* does not specify any database. Therefore all databases and files stored are evaluated, i.e. files 1 and 2 in database 25, and files 10 and 80 in database 123.