Adabas Review

Concepts and Facilities

Version 4.6

March 2012

Adabas Review

This document applies to Adabas Review Version 4.6.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Document ID: REV-CONCEPTS-46-20120329

Table of Contents

1 Summary of Features

Adabas Review provides a set of tools that allow you to monitor the performance of Adabas environments and the applications executing within them. Information retrieved about Adabas usage helps you tune application programs to achieve maximum performance with minimal resources.

The major Adabas Review features are listed below. Adabas Review:

- Supports *multiple Adabas nuclei clustered under a single database ID,* either in a single operating system image (Adabas Parallel Services) or under z/OS and IBM parallel sysplex, across multiple operating system images (Adabas Cluster Services).
- Can *monitor several databases (DBIDs)* and provide reports for each.
- Can be configured in a straight-forward, uniform, and consistent way under all targeted operating system platforms in a *client/server environment*.
- *Collects, reports, and controls data centrally*. The hub is a true data collection point unencumbered by dependencies on internal operations and structures of the Adabas nucleus.
- Provides unencumbered Adabas operation. The hub *reduces overhead* by eliminating the Adabas Review subtask from the Adabas address space, partition, or region.
- Allows you to *customize command logging*. Adabas Review can physically log selected commands or buffers based on parameters within the Adabas Review report definitions. Each command log corresponds to a particular report definition, and contains only those records that have been selected as meaningful for the end user.
- Allows you to *customize analysis tools*. Adabas Review reports may be customized to perform functions such as cost accounting and application debugging.
- *Collects history data*. Data collected for Adabas Review reports can be written to the Adabas Review repository at specified intervals and stored as history data. The data can be used in trend analysis.
- *Collects client reporting data*. When activated, client reporting data collection will allow you to better determine where command processing time is spent.

- Provides *more than 40 sample report definitions*, which may be used without modification, or edited to suit specific site requirements.
- Provides *interactive screens* via the Edit Report function to help you create a report or edit an existing one.
- Supports a *wide range of data types*. More than 340 data fields are available for specifying the types of data to be collected by Adabas Review.
- Provides *user-defined data fields*. Adabas Review provides a mechanism for defining up to five data fields to meet user specifications.
- *Allows data to be retrieved online*. Adabas Review reports are executed online or in batch mode. Data collected by Adabas Review reports may be viewed online.
- Provides an *online display of Adabas availability*. The Available Availability (AA) function displays a profile of the databases that are monitored by Adabas Review.
- Provides *access to Adabas Online System (AOS)*, an online Adabas maintenance product.

2 Adabas Review Processing and Components

Adabas Review runs in:

- *local mode* in the Adabas address space.
- *hub mode* as a server in its own address space with a client interface in the Adabas address space.
- *batch mode* as a batch job that processes sequential Adabas command log data sets

See the *Release Notes* for this version of Adabas Review for a matrix of supported Adabas versions and other requirements.

Collecting Data in Local Mode

In local mode, the Adabas Review processor is installed as an extension to ADALOG.

The data collection process is partly accomplished by the Adabas Review processor. The data is collected by providing code (RAOSLOCL) that runs as an extension to ADALOG. Adabas calls this module and passes information about resource usage for each command processed by the Adabas nucleus.

This module, in conjunction with the Adabas Review processor and an intermediate Adabas Review buffer, accumulates and tabulates the Adabas information based on various user-defined data requirements.

The Adabas Review data may be:

- displayed in an online environment from the Adabas Review user interface;
- saved automatically in the Adabas Review repository;
- printed automatically when the Adabas nucleus terminates; or
- downloaded directly to a personal computer (PC) using Entire Connection.

The following graphic shows the Adabas Review data collection process for local mode.

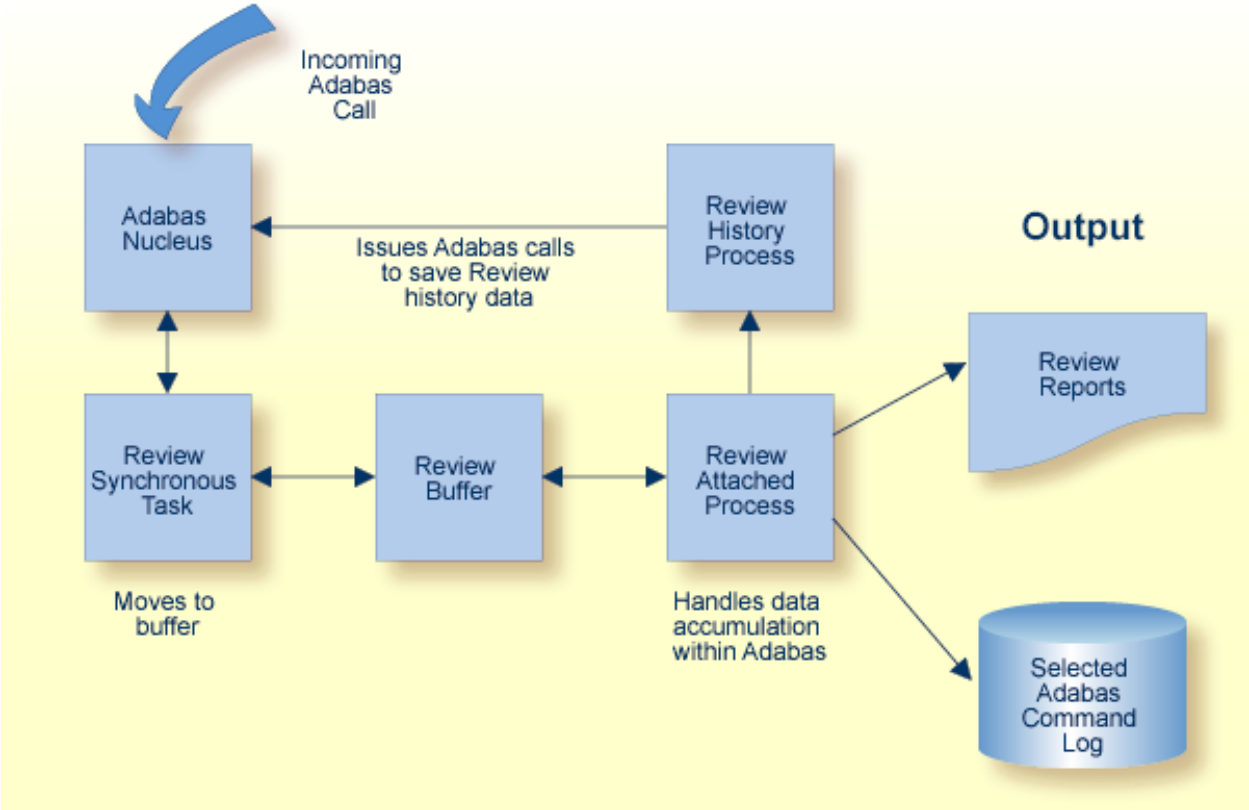

In local mode, Adabas Review job control statements should be added to the Adabas nucleus startup job.

Collecting Data in Hub Mode

In hub mode, Adabas Review uses a client/server approach to collecting data:

- an interface (the client) resides on each Adabas nucleus; and
- the hub (the server) resides in its own address space, partition, or region.

The interface uses the existing Adabas interregion communication process: ADALNK, Trans-port, ADASVC (ADARER), and ADAMPM. This process is consistent across the targeted platforms for Adabas Review. If systems are networked correctly, hub mode supports a multiple platform, multiple operating system, Adabas database environment.

In hub mode, Adabas Review job control statements should be added to the Adabas Review hub startup job.

Note: Trans-port is available only in z/OS environments.

The Server Hub

The Adabas Review hub is a centralized data collector and reporting interface that combines proven components of Adabas and Adabas Review.

It handles the data consolidation and reporting functions for monitoring an Adabas database, including usage information related to applications, commands, command response time, I/O activity, buffer efficiency, and others.

The interactive reporting facility allows you to pinpoint problems quickly, providing detailed and summary data about Adabas activities. Specific information about each database is also available.

The centralized collection server has several advantages:

- A single hub collects information from multiple Adabas nuclei, Adabas Parallel Services clusters, or Adabas Cluster Services (support for IBM's parallel sysplex environment) clusters.
- Because a single hub can support multiple Adabas nuclei, the number of Adabas Review nuclei required to support an enterprise-wide distribution of Adabas nuclei is reduced. This minimizes resource requirements and increases performance.
- Isolating the Adabas Review subtask from the Adabas nucleus enhances the performance of the Adabas main task and minimizes the impact of future Adabas releases on the functioning of Adabas Review.

The hub comprises

- ADAREV, a logic module that manages and supervises the incoming Adabas Review data calls and requests;
- REVHUB, a module to establish and maintain the environmental settings for Adabas Review; and
- the Adabas Review DB nucleus and subsystems including RAOSAUTO, the autostarted report parameter generation routine, and RAOSHIST, the historical data population routine.

The Client Interface

The Adabas Review interface constructs and then transmits the Adabas Review data from the Adabas nucleus to the Adabas Review hub. An Adabas Review interface is integrated with each Adabas nucleus that is monitored.

The interface comprises

- ADALOG, the Adabas command logging module;
- ADACLX, the Adabas command log extension module that is responsible for acquiring additional information not present in the Adabas command log record; and
- ADARVU, which handles the environment conditions for ADACLX and the Adabas API requirements for transmitting the Adabas Review data to the Adabas Review hub.

Interface Calls

To maximize performance, the ADARVU module issues an "optimistic" call from an Adabas nucleus to the Adabas Review hub without waiting for a completion or "post" from the hub; ADARVU assumes that the Adabas Review data was successfully passed to the hub.

However, ADARVU does perform an initialization step to ensure that the hub is active prior to any command processing by the Adabas nucleus. If the hub is not active, ADARVU informs you using WTOs or a user exit. If a user exit is used, you are given the option to wait for the hub to be activated, or continue initialization and call the hub only when it is active.

Command log records are normally buffered to reduce the number of calls from the Adabas nucleus to the Adabas Review hub. In low traffic situations, this might lead to a delay of approximately five seconds.

If the communication to the hub results in a non-zero response code (for example response code 151), ADARVU will call user exit 5 if it is available. If a user exit is used, you are given the option to resend the whole buffer (with or without a wait), although parts of the buffer may already be processed. We recommend that you *not* resend the buffer, but generate a message indicating an error has occurred correct the problem. For more information about user exit 5 read *REVUEX5: Adabas Review Hub Event Handler (Adabas Exit 5)* in the *Adabas Review Reference Guide*.

On the hub side of the call, the elimination of the cross-memory "post" call enhances performance by reducing the overhead of active communication with the Adabas clients. This allows the hub to remain a passive data collector.

Example Client/Server Environment

The following graphic shows the major components of the Adabas Review interface (Adabas nucleus address space) and hub (Adabas Review hub address space) in a client/server architecture
.

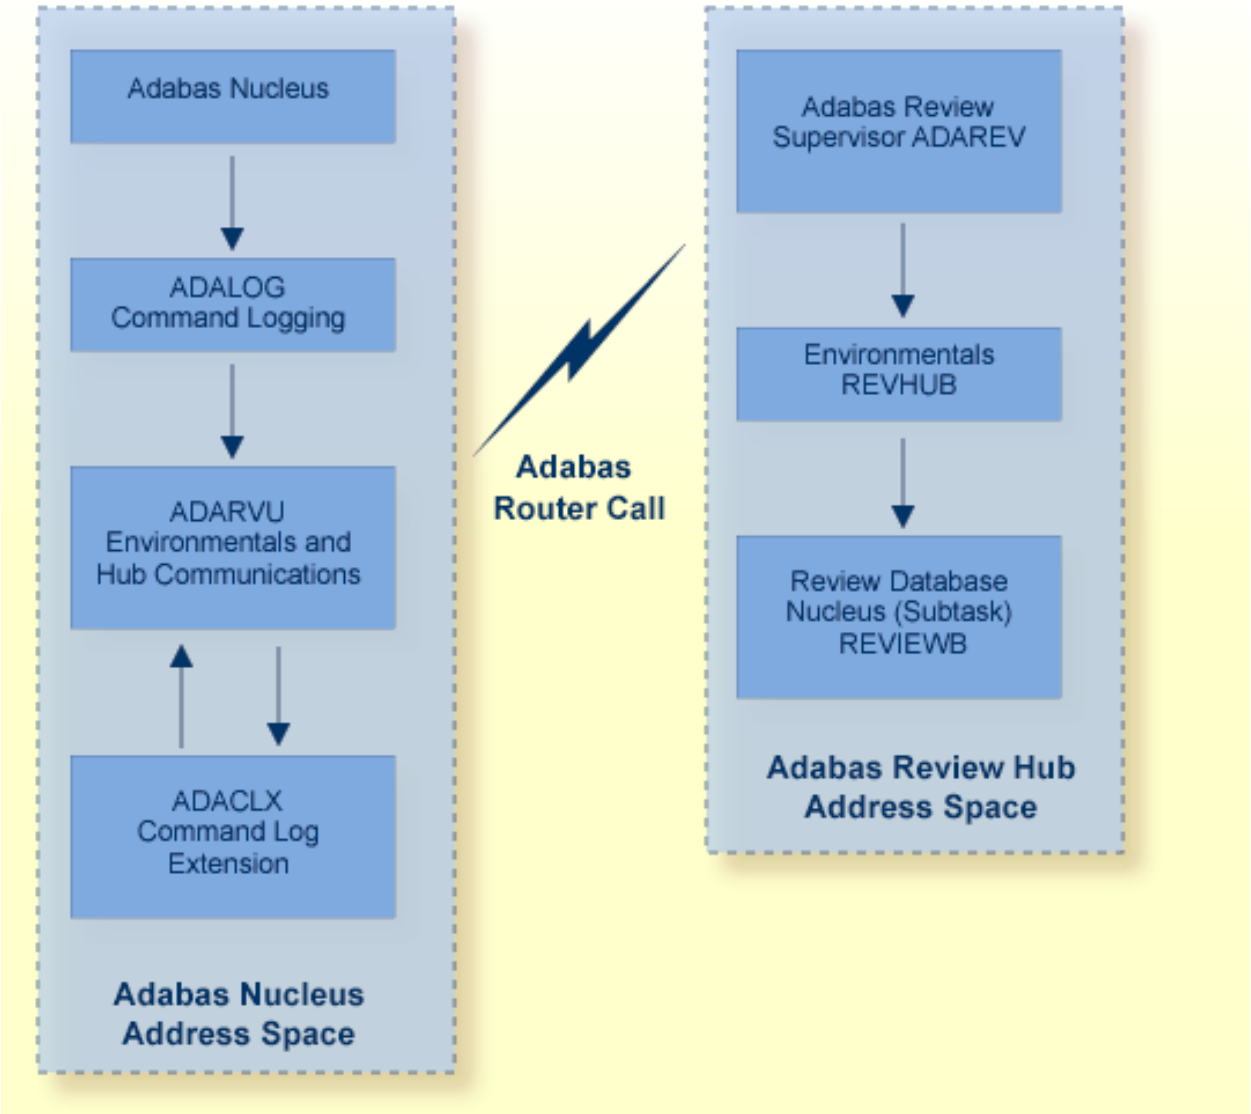

Command Log Processing

The data collection process is partly accomplished by the hub (server) component REVIEWB, the Adabas Review command log processing routine, which runs as an Adabas Review *subtask*; that is, a unit of work that the operating system treats as separately detachable.

At initialization, REVIEWB reads any autostarted report definitions the user has defined and collects data according to the reports' criteria. REVIEWB also processes requests to start, view, and purge reports from the Adabas Review online system.

In hub mode, Adabas responds to requests and calls the interface module ADARVU from ADALOG (Adabas's command logging module) if `REVIEW=dbid` is specified in the Adabas initialization parameters. Adabas passes to ADARVU information about resource usage for each command processed by the Adabas nucleus.

Adabas Review link routine exits are used to pass TP system and Natural information from the user's address space (origin of the Adabas call) to the Adabas address space and, using an extension of the Adabas user buffer, on to ADARVU.

ADARVU queues Adabas command log records received from ADALOG to the Adabas Review hub (server) through an intermediate REVIEW-BUFFER in the Adabas subtask address space. Note that command log records are examined prior to queuing to determine whether record filtering is active. When filtering is active, only command log records that are filtered are sent to the hub.

See **Record Filtering** for more information.

REVIEWB processes the records, accumulating and tabulating various data according to the criteria specified in any user-defined reports that are active.

The resulting nucleus statistics may be

- displayed in an online environment from the Adabas Review user interface;
- saved automatically in an Adabas file called the Adabas Review repository;
- printed automatically when the Adabas nucleus terminates; or
- downloaded directly to a personal computer (PC) using Entire Connection.

Example Hub Mode Data Collection Process

The following graphic shows the Adabas Review data collection process for hub mode. When monitoring multiple databases, Adabas Review allows you to switch from one database to another and provide reports for each.

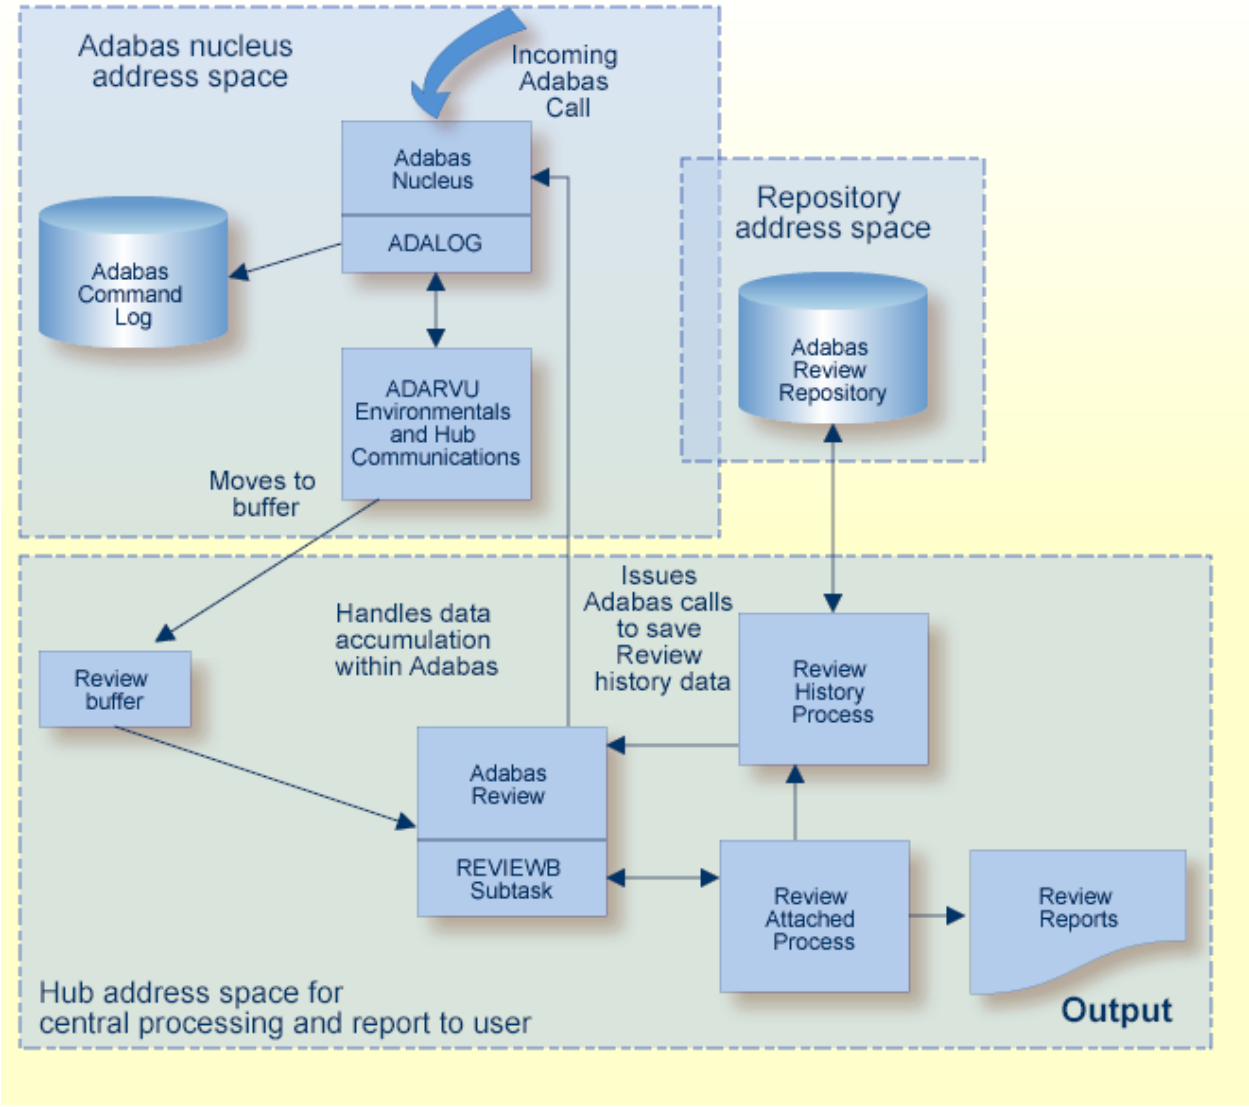

Collecting History Data

History data collection is controlled by RAOSHIST, the Adabas Review historical data population routine.

z/OS and z/VSE, and BS2000

- RAOSHIST executes as a subtask of Adabas Review. At initialization, RAOSHIST reads the RVUALT data set to determine if there are any historical records from the previous Adabas Review session that should be written to the Adabas Review repository. History records are written to RVUALT if the database on which the Adabas Review repository resides is unavailable during Adabas Review termination.
- During normal execution of Adabas Review, REVIEWB is responsible for adding history records generated by detail history reports and by summary history reports that have a history interval.

Processing Data Online Using SYSREVDB

Adabas Review can be accessed online via a standard Natural application that resides in the Natural library SYSREVDB. Using commands and screens provided in SYSREVDB, you can:

- Perform Adabas Review administration (starting, stopping, and defining reports as well as maintaining user profiles)
- Access collected monitoring data online.

For complete information on getting started using SYSREVDB, read *Getting Started*, elsewhere in this guide.

Note: The ADALNK of the TP monitor running SYSREVDB must include the REVEXIT exits. If these are not included, some features might not work correctly in SYSREVDB.

The REVEXIT exits include REVEXIT1 (formerly REVEXITB) and REVEXIT2. REVEXIT should be linked with ADALNK during installation, with the appropriate Adabas link globals table parameters (LGBLSET). The link globals table parameters specified via LGBLSET identify which REVEXIT entry points (REVEXIT1 or REVEXIT2) are called. If LGBLSET parameter REVIEW=YES is set, then REVEXIT is called by ADALNK before the Adabas call as entry point REVEXIT1; if LGBLSET parameters RVCLNT and REVHID are also specified, REVEXIT is also called after the Adabas call as entry point REVEXIT2.

The modules used to perform the linkage between REVEXIT and ADALNK are provided with your Adabas Review installation with names in the format RDBLX*sys*, where *sys* is a three-character code representing the link environment (for example, MVS for z/OS environments, VSE for

z/VSE, COM for Com-plete environments, IMS for IMS environments, and BS2 for BS2000 environments). For more information about performing this linkage, read the appropriate Adabas Review installation documentation.

Note: A return code 4 from the link-edit of the globals table is acceptable because it indicates that the Review exit (REVEXITB or REVEXIT1) has not been resolved by this step. The next link-edit step in the sample jobs should resolve this and finish with return code zero.

Processing Data in Batch Mode

As a batch job, the Adabas Review command log processing routine REVIEWB processes Adabas command log records from a sequential data set. If you use Adabas dual command logging, you must first use the Adabas utility function ADARES CLCOPY to generate a sequential command log data set suitable for input into Adabas Review.

When Adabas Review executes as a batch job, input report parameters that define the data collection criteria selected by the user are read from statements in the RVUPARM data set or the RVUAUT1/RVUAUT2 data sets. These statements can be generated using the GENCARD statement.

The storage allocated for reports is exactly the same as that for Adabas Review executing in interactive (online) mode. However, since REVIEWB is reading the command log records directly from a sequential file, no REVIEW-BUFFER is allocated.

Repositories

The Adabas Review repository is an Adabas file used for storing report definitions, historical data, and target definitions. In hub mode, this file must be located on an Adabas database which will be accessed by the hub using normal Adabas router calls.

Depending on the configuration at your site, more than one Adabas Review repository may be associated with your system. For example, if your site is running Adabas Review against more than one database, you may choose to have an Adabas Review repository for each database.

The Review command `SETFILE` (or SET) may be used to access different Adabas Review repositories and the reports stored on them.

User Profile System

Adabas Review administrators use the user profile system to generate profiles that define access rules for Adabas Review users. Access rules specify the systems or the functions within systems that a particular user is allowed to use.

User profiles may be created for new users, changed for existing users, and purged when no longer required.

A user profile is not required for each user. Adabas Review provides a default profile to allow access for users who do not have a profile defined.

When a user logs on, Adabas Review searches for the user's profile. If one is not found, the default profile is used.

If the default profile is customized so that the access rules meet the needs of the majority of Adabas Review users, the need for individual user profiles can be eliminated.

If a user has access needs that are different from the majority, a user profile can be created to accommodate those needs. Such a profile is generated by customizing a copy of the default profile.

ADARUN Parameters

Adabas ADARUN initialization parameters define and start the Adabas operating environment. Different ADARUN parameters are useful to the Adabas nucleus and the Adabas Review hub in different operating environments. For more information about the ADARUN parameters specific for Adabas Review, read *ADARUN Parameters for Adabas Review*

Storage Requirements

Adabas Review must allocate storage to execute. Storage is required for

- the Adabas Review hub, if used;
- the REVIEW-BUFFER, used as a queueing area for Adabas command log records;
- reports that are executing;
- users accessing the database from the Adabas Review online system; and
- work areas used in Adabas Review's processing.

The type, purpose, and size of these storage areas is discussed in the following sections.

Adabas Review allocates storage *above the line* whenever it is permitted by the architecture of the machine and the operating system on which it is executing.

In *z/OS* environments, Adabas Review allocates all storage from z/OS subpool 5. This allows you to accurately determine the exact amount of storage Adabas Review is using with a z/OS monitoring package.

Storage for the Hub

If you use Adabas Review in hub mode, the hub has a separate storage requirement for its operating queues and working areas. The queues are used to buffer the incoming command log records from the clients until the records can be sent to REVIEWB.

Two queues, both controlled by the database administrator (DBA), are used by the Adabas Review hub: the command queue (sized using the ADARUN parameter `NC`) and the attached buffer (sized using the ADARUN parameter `NAB`). For more information, read about these ADARUN parameters in *ADARUN Parameters*, in the *Adabas Review Concepts Manual*.

Command Queue

The command queue stores information about the client nucleus such as job name, internal ID, etc. Each entry in the command queue represents one command log record from a client.

An entry exists for the time that a command log record is queued and awaiting selection from the hub until the time that the record is sent to REVIEWB. Once the command log record is sent to REVIEWB, the entry is released from the command queue.

This means that the command queue must be large enough to accommodate the backlog of command log records from the client nuclei. If the command queue is too small, it is possible that command log records will be dropped by the hub.

The ADARUN parameter that controls the command queue size is `NC`. The value of this parameter should be set higher for the hub than it is for individual client nuclei.

The NC value should be set to handle the arrival rate based on:

- the number of clients;

- their respective command processing limits;
- the processing power of the CPU(s); and
- the priority settings of the nuclei and hub address spaces, partitions, or regions.

Example

If a client nucleus can process 2000 commands per second, then the expected arrival rate at the hub is a maximum of 2000 command log records per second. When possible, buffering occurs in the client nucleus to buffer several command log records and send them with one call to the hub.

There is no general rule for estimating the NC requirements for a particular hub. However, in this example, you could start with NC=1000 and monitor the results.

The Command Queue and Trans-port (z/OS only)

When using Trans-port in z/OS environments, most communication is handled via a fast buffer mechanism using a fixed buffer length. In such cases, the command queue is not used.

Attached Buffer

The attached buffer is used to store the contents of the command log records and their associated data extensions.

As with the command queue, an element within the attached buffer is allocated to hold the command log record for the duration of time that the record is queued for selection, up to the time the record can be sent to REVIEWB. The element is freed once the record is sent to REVIEWB.

Also like the command queue, the attached buffer must be large enough to hold the queued command log records for the time required to stage the records for REVIEWB. Software AG recommends setting the parameter high to ensure that command log data is not dropped by the hub.

The ADARUN parameter controlling the attached buffer size is NAB. The value of this parameter should also be set higher for the hub than it is for individual client nuclei.

The NAB value must be large enough to buffer the data passed by the client nuclei. The amount of data passed by a client nucleus depends upon the Adabas Review report requirements (for example, whether control buffers are required or whether the I/O list option is being used).

Example

The average size of a command log record and extensions, excluding control buffers, is 2500 bytes.

One approach would be to compute:

```
NAB = (NC * 2500 / 4096)
```

- where 4096 is the size of one NAB segment. If NC=1000 (see the example) , the starting value would be:

```
NAB = (1000 * 2500 / 4096) = 610
```

This computation assumes that there are no control buffers or I/O list elements being passed to the hub.

The Attached Buffer and Trans-port (z/OS only)

When using Trans-port in z/OS environments, most communication is handled via a fast buffer mechanism using a fixed buffer length. The attached buffer area is not used.

Storage for the REVIEW-BUFFER

REVIEW-BUFFER is used to queue Adabas command log records to be sent to REVIEWB. In hub mode, it is located in the hub (server) address space.

The `BUFFER-SEGMENTS` parameter specifies the size of the REVIEW-BUFFER. Each buffer segment is 512 bytes. When a value for BUFFER-SEGMENTS is specified, the total storage indicated by the specification is split in half between 4K buffers and 32K buffers.

In addition, two INPUT statement parameters have been added as replacements for the BUFFER-SEGMENTS parameter. The BUFFER-SEGMENTS parameter will no longer be supported in a future release of Adabas Review. It is still valid in Adabas Review 4.5, but is ignored if the following parameters are specified:

Parameter	Description	Minimum Value	Default Value
BUFFERS-4K	The number of buffer pool entries that have a length of 4096 bytes or less. This parameter is usually specified along with the BUFFERS-32K parameter. If this parameter is not specified, any BUFFERS-32K parameter setting is ignored. If a value below the minimum value is specified for this parameter, the default is used.	124	256
BUFFERS-32K	The number of buffer pool entries that have a length of 4097 bytes or greater. This parameter is usually specified along with the BUFFERS-4K parameter. If this parameter is not specified, any BUFFERS-4K parameter setting is ignored. If a value below the minimum value is specified for this parameter, the default is used.	15	30

For z/OS, z/VSE, and BS2000, it is possible to execute with a REVIEW-BUFFER that is one megabyte.

A larger REVIEW-BUFFER provides a larger queueing area for command log records being sent to REVIEWB and decreases the possibility that Adabas will have to wait for REVIEWB to process these records in the event that REVIEW-BUFFER becomes full.

Storage for Reports

For Control Blocks

When a report is started, either using autostarted report definition parameters or by an online Adabas Review user, storage is allocated for control blocks that define the criteria for the collection of the data.

Typically, the storage allocation for control blocks is two (2) kilobytes, but may be as much as four (4) kilobytes if the report is a history report or the report specifies the collection of many fields.

For Data Collection Areas

In addition to the report control blocks, storage is allocated for the collection of data. The data collection areas are allocated in two (2) kilobyte pieces and a subsequent data collection area is only allocated when the current area is full.

Total Storage Limit

The total storage allocation for a report is limited by the `MAXSTORE` report parameter. When the total storage allocation for a report is equal to the `MAXSTORE` value, the report is marked as inactive and stops accumulating data. When a report is purged, all storage associated with the report is deallocated.

Storage for Online Users

Adabas Review's online system uses Adabas calls to start, view, or purge a report. Each request requires that Adabas Review perform some processing to fulfill the request.

- Each request from the Adabas Review online system results in the allocation of a piece of storage (about 500 bytes or one-half kilobyte) that is deallocated when the request has been satisfied.
- To maintain the integrity of each request, Adabas Review allocates an area for each user requesting Adabas Review to service a request. For example, a request to view a report requires an Adabas call for each record that is to be viewed online.
- In cases where more than one user is viewing the same or different reports, Adabas Review must remember the status of each user between Adabas calls.

Storage for Work Areas

Adabas Review allocates storage for work areas and areas used for reading from and writing to files. These areas are typically small and are kept and used throughout the time that Adabas Review is active.

Size of Adabas Review in z/VSE Environments (Local Mode Only)

In z/VSE environments, the available GETVIS in the Adabas partition may need to be increased to accommodate Adabas Review. An increase of 500K is normal.

3 About Adabas Review Reports

Online or as a stand-alone batch job, Adabas Review processes Adabas command log records and generates reports according to user-defined reporting criteria. The flexible reporting structure of Adabas Review allows you to view the same data in many different ways.

It retrieves performance information about Adabas by collecting data in the form of reports, which capture command log information and provide information about both online and batch transactions. Reports collect data about multiple teleprocessing environments and batch programs that are all using the same database.

Adabas Review provides commands to:

- modify report definitions;
- modify display programs;
- control report data collection; and
- direct report output.

Detailed vs. Summary Reports

Adabas Review can produce detailed and summary reports.

- For detailed reports, data is collected continuously while the database is active. Such reports contain a complete recording of the commands processed while the database is active. Detailed reports cannot be viewed online. However, a display program is generated that may be used to view history data online if the report collects history data. Results of detailed reports are printed at database termination.
- For summary reports, data in the report is summarized by a particular field (for example, by command, user ID, or program). Results can be viewed online. Summary reports can also be printed at the time the summary report is created or at database termination.

Components of a Report

An Adabas Review report has two parts: a definition and a display program.

- A ***report definition*** is a set of parameters that specify the data to be captured and a set of processing rules that identify the conditions under which the data is to be captured.
- A ***Natural display program*** is a Natural program that specifies the appearance of the report output, whether viewed online or downloaded to a personal computer (PC).

Report Definition

Adabas Review uses a set of instructions called a *report definition* to specify the types of data to be collected. Prepared report definitions supplied with Adabas Review may be modified and custom reports may be created.

Report definitions can be created or modified using menu-driven Natural programs. Report options and processing rules allow you to specify the conditions under which the data is to be captured. Report definitions are kept in the Adabas Review repository.

The Edit Report (ER) function is used to create Adabas Review report definitions. In addition, Software AG delivers prepared report definitions called *supplied reports* with Adabas Review. Supplied reports may be modified to meet the requirements at a particular site. Detailed information about supplied reports is provided in *Supplied Report Reference*, in the *Adabas Review Reference Guide*.

An Adabas Review report definition comprises database fields to be monitored, processing rules, and option parameters, all either entered on screens or selected from lists:

- ***Database fields*** from which a report extracts data are listed in the order in which they are to be displayed.
- ***Report processing rules*** identify the conditions under which the data is to be captured by describing how field values are to be extracted for the report.
- **report option parameters**
  Report option parameters specify whether Adabas Review will perform physical command logging, and whether data accumulated by the report will be written to a history file. History data may be used when analyzing database performance trends.

Database Fields

The selection of database fields defines the subject of the Adabas Review report.

A field identifies a type of information to be collected by the report; for example, information about an Adabas command, the Natural program issuing the command, the Adabas nucleus, the teleprocessing monitor, or the operating system being used. The PF11 (Flds) option is used to display a list of all available Adabas Review report fields.

Up to 20 fields may be used in a particular report. However, the number of fields that can be displayed is limited to the total line size (250 bytes) that can be displayed when a report is viewed online.

Fields to be displayed in a report are entered on the Edit Report screen of the Edit Report (ER) function as shown in the following example:

```
20:35:15                A D A B A S  -  R E V I E W               YYYY-MM-DD
                               Edit Report                        LOCL=00009

     Report Name: ____________________________________  DBID to Monitor: _____

      +-------------------------------------------------------------------+
      |  Field      Order   Sum   Min   Max   Avg   Pct   Rate   Round   |
      |-------------------------------------------------------------------|
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      | ________     ___     _     _     _     _     _     _    _______  |
      +---------------------------- Page 1 -----------------------------+


Command: ____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help  Optns Exit        Save  Start  -     +          Rules Flds  Menu
```

Detailed information about individual fields is provided in the *Adabas Review User's Guide*.

Report Processing Rules

Report processing rules are used to set conditions on or limit the data collected by the report.

Processing rules are entered on the Report Processing Rules screen of the Edit Report (ER) function as shown in the following example:

```
20:24:56                  A D A B A S  -  R E V I E W                 YYYY-MM-DD
                             Report Processing Rules                  LOCL=00009

                                    REPORT ONE
+-------------------------------------------------------------------------------+
|  Field    Op        Value                                              And/Or |
|-------------------------------------------------------------------------------|
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
| _________  __  ____________________________________________________________  ___ |
+----------------------------------- Page 1 -------------------------------------



Command: _____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help  Optns Exit        Save  Start  -     +               Flds  Menu
```

An Adabas Review report processing rule is specified by a field, a relational operator (equal to, less than, greater than, not equal to), and a specific value or range of values.

If there are multiple processing rules, the appropriate logical operator (AND or OR) must also be specified.

Fields used in processing rules are not displayed in the report results unless they are listed on the Edit Report screen.

Report Option Parameters

Values for report option parameters can be entered on the Report Options screen of the Edit Report (ER) function as shown in the following example.

Note: One of two different Report Options screens may appear in Adabas Review, depending on whether you have selected a summary report or a detail report. The following sample is for a summary report.

```
19:33:13                 *****   R E V I E W   *****                2009-06-18
                Options for summary report: COMMANDS BY HOUR         HUB=15690

+-- Summary Rpt Options ------- Summary Logging ------------ History ---------+
!  AutoStart .......... Y    Summary Log ........ N    History ............ N !
!  Break .............. Y    Summary File .. RVSUM_    History Int ... _______ !
!  Wrapping ........... N    Num of Logs ....... _2    History DBID ... __123 !
!  Print .............. Y    Log Size ....... 99999    History FNR .... ___45 !
!  Rstrt/Intrvl Y  _______    Log Full Exit _________    History SVC ...... 234 !
!  Max Restarts .. 999999                                                     !
!  Max K ........... ___8                             ---- Summary Exit ----- !
!  ADALimit ...... _____1                              Exit Name ... _________ !
!  Display By ..   SORTED                              Cmd ............... CL !
!  Entries ....... 999999                              Sum ................ Y !
!  Limit ....... 99999999                                                     !
!  Page/Line .. _55 / 133                                                     !
!                                                                             !
!                                                                             !
+-----------------------------------------------------------------------------+


Command: ____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help        Exit             Accpt            Dflts             Menu   ↩
```

Adabas Review is delivered with report options defaults:

- You can modify these defaults for all subsequent reports.
- You can modify the default values for a particular report, either new or existing.

Detailed information about using the Report Options screen is provided in the *Adabas Review User documentation*.

Report User Exits

Adabas Review provides two report user exits: one for detail reports and one for summary reports.

- A *detail report user exit* is driven when a command log record is selected for the report. Only records that pass the processing rules are provided to the user exit. This exit may be used to create SMF records, accounting records, or for any other purpose.
- A *summary report user exit* is driven when a specified Adabas command is selected for the report; and/or a report is summarized. You may control the conditions that trigger the exit.

When creating a report, the user exit is specified on the Report Options screen or in the batch `REPORT` statement. The actual report user exit must be provided in an executable library accessible to Adabas Review.

For more information about these exits, read *REVUXDET: Report Exit for Detailed Reports* and *REVUXSUM: Report Exit for Summary Reports*, in the *Adabas Review Reference Guide*.

Natural Display Program

When a report definition is saved, Adabas Review automatically generates a unique Natural display program, called the Adabas Review report display program. This program determines how the report data will be viewed online or downloaded to a PC. It controls the appearance of the report output for an active database or for historical data. You can run this display program:

- to view data currently being collected by the interactive Adabas Review processor.
- to retrieve and display or download historical data that has been saved and stored in the Adabas Review repository.

Note: No display program is created for buffer pool reports.

The Natural display program generates normal Natural `FIND` and `READ` statements against an Adabas Review DDM to access the data being collected by Adabas Review. By default, the Natural program displays the data at the Adabas Review user's terminal. Options exist, however, to download the data directly to a personal computer (PC).

You may use the Adabas Review-generated display program as it is created, or you may modify it for your individual needs.

The display program is executed when the `VIEW` command is issued to display the report results online.

Note: The output of reports that print at database termination is not controlled by a display program. At database termination, a SYSOUT from Adabas is written by REVIEWB. The format of this output cannot be modified by the Adabas Review user.

Starting Reports

Once the report definitions are edited and saved, the reports can be started. Starting a report tells the Adabas Review data collection process to start accumulating data based on the report definition parameters.

Adabas Review users can display buffer pool information, display active databases, and access the Adabas Online System, if it is available. Adabas Review administrators are also allowed to define and display target objects.

Reports can be run in either interactive or batch mode:

- In *interactive mode*, Adabas passes information to Adabas Review about resource usage for each command processed by the Adabas nucleus. Online reports are started with the `START` command.
- In *batch mode*, Adabas Review runs as a batch job that processes sequential Adabas or Adabas Review command log data sets. Batch parameter statements are generated for the report and submitted in the job stream.

Once a report has been defined, the `SAVE` command is issued and the report definition is written to the Adabas Review repository.

Once data has been accumulated for a report, the results can be:

- saved to a file,
- viewed online,
- downloaded to a personal computer (PC) using Entire Connection, or
- printed when the Adabas nucleus terminates.

As long as a report is not stopped, its data can be viewed online or downloaded to a PC using Software AG's Entire Connection.

Detailed reports do not have any accumulated data (the data is directly saved to a file); therefore, detailed reports cannot be viewed online.

Summary reports do accumulate data. If the report is stopped, the data is processed differently by Adabas Review, depending on what caused the report to be stopped and on the setting of the report options. If the data is kept, it may be stored in RVUPRT, the summary logging file, or the history file.

Reports can be stopped using the SYSREVDB CL (close), PS (purge), RF (refresh), or SU(spend) report commands. Reports stop automatically if one of the following report option limits is set and exceeded:

- Max K

- Entries
- Limit
- Intervl
- History Int

Data is accumulated and written when:

- The report is stopped using the CL command;
- The report stops automatically because of an interval event (INTERVL report option);
- The report has the RESTART=Y option specified and stops automatically because the maximum storage limit (MAX K report option) was exceeded; or
- Adabas Review terminates. Adabas Review terminates in local mode if the nucleus is stopped via the ADAEND or HALT operator commands or the hub is stopped via the ADAEND operator command.

Note: In local mode, data is written with the next Adabas command log record collected by the Adabas Review system.

Data is *not* written when a report is stopped using the PS or RF commands.

When one of the following situations occurs, the data is kept in storage and may or may not be written, depending on the next actions you take:

- The report is stopped using the SU command;
- The report stops automatically because the maximum number of entries that the report can maintain or report has been reached (ENTRIES or LIMIT report option); or
- The report has the RESTART=N option specified and stops automatically because the maximum storage limit (MAX K report option) was exceeded (the report status is set to "I" in this situation).

Autostarted Reports

Adabas Review reports can be set to start automatically whenever Adabas Review initializes.

When you make changes to an autostarted report, delete an autostarted report, or modify the target definition for the database being monitored by the reports, RAOSAUTO (the Adabas Review autostarted report parameter generation routine) automatically regenerates the control statements for all autostarted reports that are defined in the Adabas Review repository file identified by the Natural session (read ***Accessing Adabas Review Repositories***, elsewhere in this section). The parameters for the INPUT statement are generated based on the information of the target definition (read *Listing Target Definitions*, in the *Adabas Review Administration Guide*). Running in hub mode, a target definition of the hub ID should be defined.

RAOSAUTO will then write the parameters to one of two parameter files, RVUAUT1 or RVUAUT2, alternating between them by writing to the older file.

Only when Adabas (local mode) or the hub (hub mode) starts are the files read by Adabas Review (using the RVUAUT1 and RVUAUT2 statements in the job stream) and the reports automatically started.

Note: Under z/OS, the installation procedure defines the statements RVUAUT1 and RVUAUT2 so that they point to members of a PDS. To avoid constant compression of these data sets, the statements may point to sequential data sets.

In exceptional circumstances (for example when the source library becomes too full and requires compressing), you can force regeneration of the control statements for all autostarted reports by issuing the GENAUTO command or by entering the parameters manually using batch parameter statements.

In addition, when you issue the GENCARD command, RAOSAUTO generates report parameter cards for one user-specified report and directs them to a user-specified output file.

In *z/OS* and *z/VSE*, , and *BS2000*, RAOSAUTO executes as a subtask of Adabas Review and is only active when:

- an autostarted report definition is saved in the online system;
- GENAUTO (or GENCARD) is executed in the online system; or
- a database target definition is saved in the online system.

List Report Functions

Adabas Review provides three list functions for reports, which are accessible from the Adabas Review (DB system) menu:

Code	List . . .	This function lists . . .
LR	Report	existing Adabas Review reports by name. Reports may be edited, started, and purged using the commands in this function.
LH	History Reports	reports that have accumulated history data in the Adabas Review repository.
LS	Started Reports	reports that are currently accumulating data.

The lists displayed by the LH and LS functions are subsets of the list displayed by the LR function.

From the resulting lists, certain commands can be issued against a particular report that allow you to maintain report definitions, manage started reports, and manage the output they produce.

The commands are entered on the selection line preceding the name of the report.

Record Filtering

Adabas Review record filtering can decrease the number of command log (CLOG) records passed from the database to Adabas Review. When filtering is active, only those records needed for report processing will be passed to Adabas Review. Record filtering operation is based upon the reports currently active and cannot be directly controlled. However, the Adabas ADARUN parameter REVFILTER may be used to control whether the database allows filtering. See the appropriate installation documentation for a description of the REVFILTER parameter.

Record filtering is limited to and dependent upon select Adabas Review fields (filter fields). Reports with rules based solely on filter fields are eligible to be filtering reports. Any other report is a non-filtering report. Record filtering will be active only when all reports running are filtering reports. That is, record filtering will stop if any non-filter report is started. An example of a filtering report would be the *Exceptional Response Code* report, which has a single rule for the RSP field. If this is the only report running, only those command logs with exceptional response codes will be sent to Adabas Review.

Note: Record filtering is intended for those installations that ordinarily run very few reports (e.g., *Exceptional Response Codes*) at a time. A large number of active filtering reports will negate the performance benefits of record filtering. If your installation runs many rule-based reports, running with the ADARUN REVFILTER=NO setting is recommended. Some analysis of active reports may be necessary to fully realize filtering performance benefits.

Record Filtering Restrictions

The following restrictions apply for Adabas Review record filtering:

- For filtering to be activated, all currently active reports must have processing rules based solely on filtering fields.
- Wildcard values are not permitted. Alphanumeric values containing an asterisk (*) will result in filtering being turned off.
- EXCLUDE statements are not supported. An EXCLUDE statement will result in filtering being turned off.
- Special Adabas Review reports are non-filtering reports.
- CLOGLAYOUT=8 is required for record filtering.
- Filtering will not be active when running with LOGCLEX=YES.
- In hub mode, any active report using the artificial target ID 11111 will result in filtering being turned off.

Adabas Review Fields Eligible for Record Filtering

The following Adabas Review fields can be used in rule processing in order to activate Adabas Review Record Filtering: ACBUSER, ASSOIO, CID, CMD, CMDTYPE, DATAIO, DBID, FILE, ISN, ISNQ, JOBNAME, JOBNAME, NATAPPL, NATPROG, NATSTMT, NATUID, NUCID, RSP, STEPNAME, TPUSERID, TRANSID, and WORKIO.

The use of any other field in rules for active reports will cause filtering to be stopped. For detailed information about particular fields, see the *Field Reference*, provided with this documentation.

4 About Adabas Review Client Reporting

Client reporting allows you to collect client data you can use to determine where command processing time is spent. Client reporting is currently provided for Adabas Review hubs only in z/OS environments using the TSO , CICS and IMS monitors; it is not currently supported in batch or Com-plete. Client reporting is intended as a diagnostic tool and can be activated and deactivated as needed.

Data for client reports is normally collected in the client address space from the client as well as the Adabas data server. Note that data for client reports can be collected when Adabas Review is running either local or hub mode, but can only be transferred to and reported from a hub.

Caution: The process required to collect client data increases the overhead of your system. We therefore recommend that you activate client reporting only for as long as it takes to diagnose a problem. We do not recommend that client reporting be activated all the time.

This chapter covers the following topics:

Installation and Activation of Client Reporting

Client reporting is currently supported for Adabas Review hubs in z/OS environments. For complete information on how to install client reporting, read *Installation Phase 2*, in the *Adabas Review z/OS Installation Guide*. Installation of client reporting support involves performing the following tasks at some point in the Adabas Review installation:

- You will need to set two new LGBLSET macro keywords, RVCLNT (which allows you to indicate that client reporting can be activated for the hub) and REVHID (which specifies the hub ID). New globals tables must be assembled with these keywords.
- The new globals tables must be linked with the appropriate Adabas Review exit modules (which include REVEXIT2).
  1. The Adabas Review exit modules are RDBLXMVS (batch/TSO) and RDBLXCIC (CICS). Support is not yet provided for RDBLXCOM (Com-plete). These modules contain the batch/TSO and CICS components of Adabas Review that will be used with the Adabas Version 8 CICS link routine.
  2. While client reporting is active, the REVEXIT2 link routine will collect data and send it to the preferred Adabas Review hub after each Adabas call.

     Once client reporting is activated, client data collection occurs for each Adabas call from within the ADALNK routine where it has been activated.
- Refer to the *Adabas Review Release Notes* for information about compatibility with Adabas versions.
- If Adabas Review is installed with its CICS TP monitor, SYSREVDB (which is used to activate and deactivate client reporting) must be running in the same CICS and using the same task-related user exit (TRUE) that will process client data.

Status messages are provided during installation indicating whether or not client reporting is supported by the system. After installation, if client reporting is supported by the CICS system, client reporting state messages are produced indicating when client reporting is activated or deactivated.

Once client reporting is installed, it is not automatically activated. Instead, you must manually activate it using SYSREVDB. Activate and deactivate client reporting data collection as needed by your installation. To verify if client reporting is turned on correctly please review any LNKRVX* messages appearing on the console log. For more information, read *Managing Client Reporting*, in the *Adabas Review Administration Guide*. Because of the CICS system overhead that can occur while client reporting statistics are collected, we do not recommend that you leave it activated for very long.

Client Reporting Data Collection Process

The following diagram depicts the client reporting data collection process when client reporting has been activated.

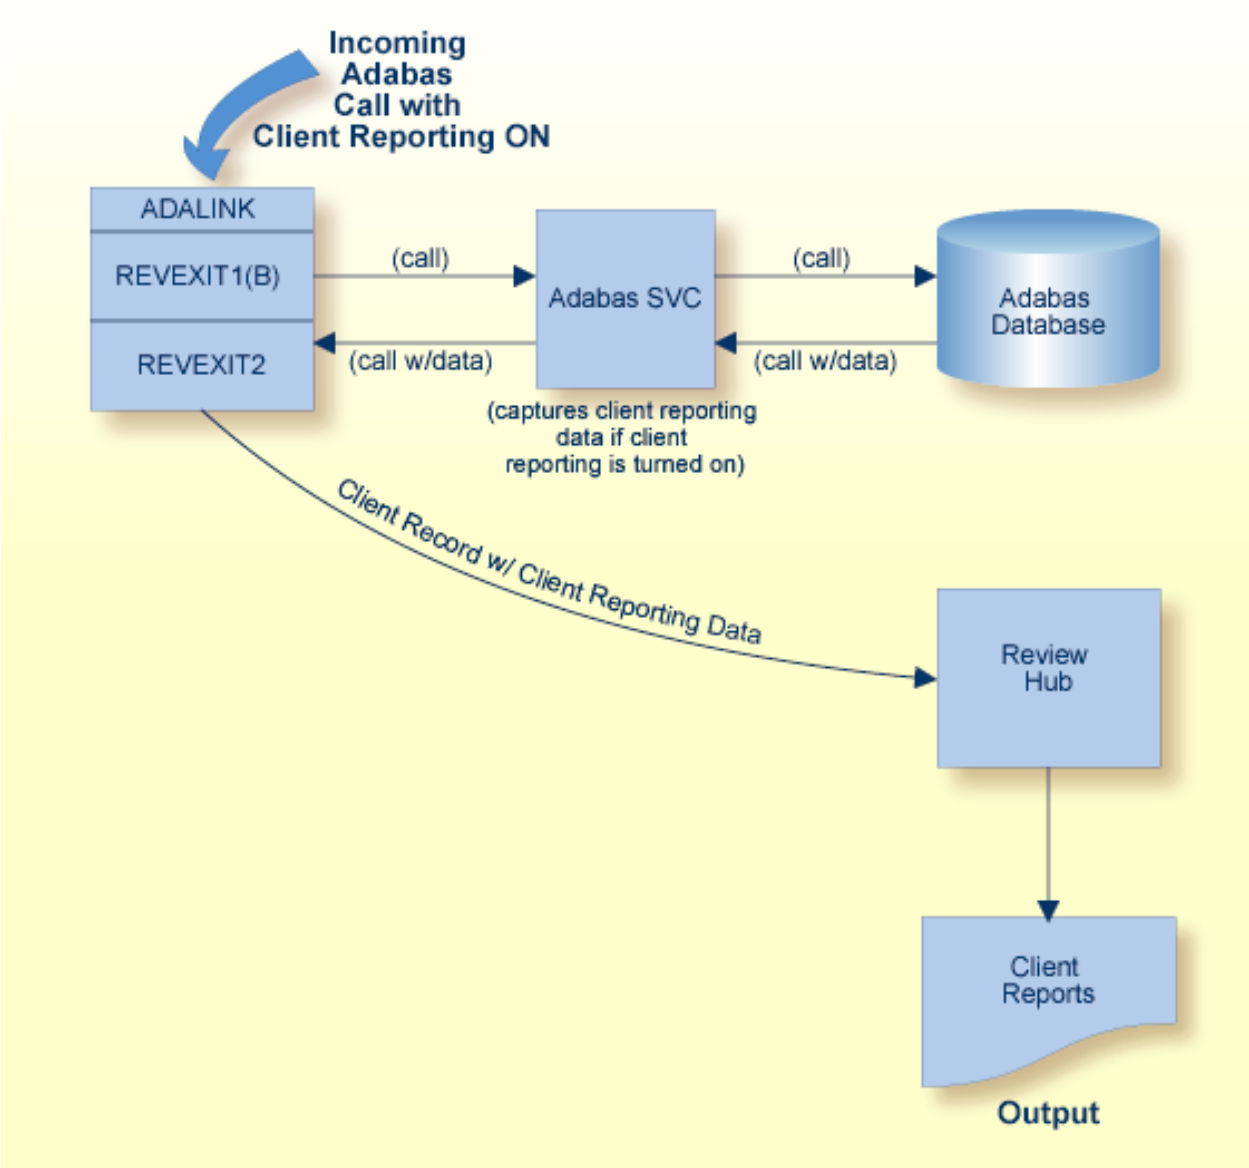

When the Adabas call is processed after client reporting has been activated, the Adabas SVC captures the timing data necessary to produce the data stored in the client reporting fields. Then REVEXIT2 collects and sends that data, via Trans-port, to the preferred Adabas Review hub for the call. The client reporting data is included with a subset of a command log record.

Caution: This process has some impact on the performance of CICS. We therefore recommend that you run client reporting for controlled periods of time as a diagnostic tool only.

Data collection occurs for each Adabas where a globals table has been prepared to support it.

Data Collected for Client Reporting

Data for client reports is normally collected in the client address space from the client as well as the Adabas data server. Note that data for client reports can be collected when Adabas Review is running either local or hub mode, but can only be transferred to and reported from a hub.

When client reporting is activated, three data fields, specific to client reporting, are collected and included with a subset of the command log record:

- The total client duration time is collected and stored in the CDURA field.. This is the total time (in seconds) in which the client waits for the command to be processed by the server and the time it takes the ADALNK portion of the client to retrieve the command results. CDURA is the sum of the CRCVDURA and CWRKDURA fields.
- The client receive time is stored in the CRCVDURA field. This is the time (in seconds) it takes the Adabas link routine to retrieve a processed command from the server.
- The client wait time, or the time in which the server works for the client, is stored in the CWRKDURA field. This is the time (in seconds) in which the client waits for the command to be processed by the server.

Note: For more information on how duration fields are calculated in Adabas Review, read *Adabas Review Duration Field Derivations*, in the *Adabas Review Reference Guide*.

These specific client reporting fields become available for client reporting reports only when client reporting is activated. For complete information about the specific client reporting fields, read *Client Reporting Fields (CMON)*, in the *Adabas Review Reference Guide*.

In addition, a subset of the other field data collected by Adabas Review from the Adabas server is also available for client reporting reports. This includes the data for the Adabas Review fields ADADURA, CQDURA, DBID, JOBNAME, NUCID, ROUTTIME, SEQUENCE, and TOTDURA. If you do not want to collect this data, the appropriate ADARUN parameter of the Adabas server has to be set as described in ***Adabas Nucleus (PROGRAM=ADANUC) ADARUN REVIEW Parameter Setting Considerations for Data Collection Under Client Reporting***, elsewhere in this section.

For a complete list of fields available for client reporting reports, read *Fields Available for Client Reporting*, in the *Adabas Review Reference Guide*.

Reporting on Client Reporting Data

Use the Edit Client Report screen in SYSREVDB to create a client report definition. For more information, read *Maintaining Standard Database and Client Reporting Reports*, in the *Adabas Review User's Guide*.

In addition, you can use the CLASS parameter of the batch REPORT statement to specify whether the batch report to be generated is a client reporting report. Valid values are "STANDARD" (the report is a standard report) and "CLIENT" (the report is a client reporting report). The default is "STANDARD". The CLASS parameter can be specified on both detailed and summary reports. For more information about the CLASS parameter, read *REPORT Statement*, in the *Adabas Review User's Guide*.

Adabas Nucleus (PROGRAM=ADANUC) ADARUN REVIEW Parameter Setting Considerations for Data Collection Under Client Reporting

Data for client reports is normally collected in the client address space from the client as well as the Adabas data server. Note that data for client reports can be collected when Adabas Review is running either local or hub mode, but can only be transferred to and reported from a hub.

Once client reporting is activated (turned on), data is collected for the three specific client reporting data fields: CDURA, CRCVDURA, and CWRKDURA. Data for these fields is collected regardless of the ADARUN REVIEW parameter setting of the databases which are called from within this client program. However, some reporting fields (see the list below) in client reporting reports are affected by the setting of the ADARUN REVIEW parameter:

- To *start, run, or stop* client reporting reports, the ADARUN REVIEW parameter must not be specified in the Adabas nucleus job. You prepare the Adabas Review client (RVCLNT=YES) specify the hub ID (REVHID) in the ADALNK parameter settings that support client reporting reports.
- To *collect* data from the Adabas server in the following Adabas Review fields (while client reporting is turned on), the ADARUN REVIEW parameter specified in the Adabas nucleus job must be set to "LOCAL" or to a hub ID. If it is set to "NO", data is not collected for the following fields:

  ADADURA
  CQDURA
  DBID
  JOBNAME
  NUCID
  ROUTTIME

SEQUENCE
TOTDURA

If you do not want to collect the Adabas server data in these fields for a given database, you can either:

- Set the ADARUN REVIEW parameter to "LOCAL" for the given database, and do not define any reports (manually started or autostarted) for this database; or
- Set the ADARUN REVIEW parameter to an invalid hub ID to avoid any standard reporting for this database.

Data will be collected for these additional fields and the overhead in the Adabas server will be minimal.

For more information about the ADARUN REVIEW parameter, read *REVIEW Parameter: Adabas Review Control*, in the *Adabas Review Reference Guide*.

5 Integration with webMethods Optimize for Infrastructure

Adabas Review key performance indicator (KPI) statistics are provided to webMethods Optimize for Infrastructure. Using the Optimize component of webMethods Optimize for Infrastructure, you can monitor the status of these statistics as well as the overall status of your system. This support is available in version 8.2 (and later versions) of webMethods Optimize for Infrastructure.

The key performance indicator statistics that are passed from Adabas Review to Optimize provide information about the Adabas Review system itself and not about the monitoring data collected by Adabas Review. The KPIs that are passed include statistics such as the number of records sent to Adabas Review, the CPU. time used by Adabas Review, the memory usage of Adabas Review, and number of I/Os performed by Adabas Review.

6 Getting Started

This chapter tells you how to log on to Adabas Review online (SYSREVDB), use the online help system, use function codes, PF keys and commands, and how to end your Adabas Review session.

Accessing Adabas Review Online

Adabas Review can be accessed online via a standard Natural application that resides in the Natural library SYSREVDB.

To access Adabas Review online via SYSREVDB:

1 Access Natural as you do normally.

2 At the NEXT prompt, enter `LOGON SYSREVDB`.

3 At the prompt, enter `MENU`.

If you are running in hub mode and if the default hub ID is specified as "AUTO", the value is determined from the SVC of the current Natural session. If there is a single hub running under the current SVC, that hub ID is used. If more than one hub ID is active, the following pop-up dialog appears. (If no hub is running, the value is determined as if "LFILE" were specified.)

```
NEXT menu                                                        LIB=SYSREVDB
 Logon accepted to library SYSREVDB.


                      Available Review Hubs on SVC 237

                        Please choose a Review Hub
                 ---------------------------------------------
                 _ 00041      _____       _____       _____
                 _ 00558      _____       _____       _____
                   _____      _____       _____       _____
                   _____      _____       _____       _____
                   _____      _____       _____       _____
                   _____      _____       _____       _____
                   _____      _____       _____       _____
                   _____      _____       _____       _____


                 Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7--
                                  Exit
```

Type any character next to the hub you want to use and press Enter.

Your Adabas Review administrator can tell you if an alternative logon procedure has been defined.

When you log on to Adabas Review, the Main Menu screen is displayed.

This section covers the following topics:

- Main Menu
- General Use PF Keys

Main Menu

```
14:29:57                  A D A B A S  -  R E V I E W                 2011-03-01
                                Main Menu                             LOCL=00559

                 Code                Description
                 ----      ---------------------------------
                  AA             Available ADABAS Nuclei
                  AH             Available Review Hubs
                  AO             ADABAS Online System
                  ER             Edit Report Definition
                  ES             Edit Specialty Reports
                  ET             Edit Target Definition
                  LH             List History Reports
                  LR             List Report Definitions
                  LS             List Started Reports
                  LT             List Target Definitions
                  UP             User Profiles
                 ----      ---------------------------------




Command: _____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help        Exit                                                Fin    ↩
```

Notes:

1. For an explanation of the commands in the **Code** column in the screen above, read *Commands* in the *Adabas Review Reference Guide*.
2. LOCL= and HUB= are found in the top right corner of every screen.
   In the screen pictured above, LOCL=00222 means that database 222 is being administered, running in local mode. HUB=201 means that a hub is administered.
   If the message REV00362 - REVIEW IS NOT INSTALLED IN DATABASE 210 in conjunction with display of HUB=210 (in the right top corner) is displayed, then the database 210 is not a HUB and also does not run with Adabas Review in local mode. Changing to a database or HUB is done using the DBID or HUBcommands.

General Use PF Keys

The following PF keys are available on many of the Adabas Review screens:

PF	Command	Description
1	`Help`	Access the help system.
3	`Exit`	Return to the previous screen.
7	`-`	Scroll backward through a list display.
8	`+`	Scroll forward through a list display.
12	`Menu`	Return to the Adabas Review Main Menu.

Other Adabas Review commands and PF keys are local to a particular function or screen. These commands are usually shown on the relevant screen and have PF keys assigned to them.

Commands used in the list functions are not displayed on the screen. These commands may be displayed by using the online help system.

Using Function Codes

Adabas Review is menu-driven. A series of function codes and commands are used to navigate through its functions. Functions may contain several "layers" of screens.

Entering function codes singly accesses one layer at a time; entering a string of function codes enables you to bypass intervening screens and go directly to the screen you want.

For example, you can list the reports that have been started (`LS`) for Adabas Review by entering the following on the command line of the Adabas Review main menu:

```
LS
```

The `LS` code accesses the Started Reports screen.

Using Adabas Review Commands

Commands are used to perform specific tasks within a function.

To issue an Adabas Review command:

- type the command on the command line and press ENTER

  Or:

press the PF key corresponding to the command, if one has been defined.

A command may be included in a string, provided the command is the last element.

Some commands are standard throughout Adabas Review, and may be issued from any screen. These include the following:

Command	PF	Description
COLOR OFF	—	Returns to non-color display.
COLOR ON	—	Displays color attributes, if applicable.
EXIT	3	Terminates the function and returns you to the menu screen that precedes it.
FIN or QUIT	12	Terminates the Adabas Review session. Assigned to PF12 on the Adabas Review Main Menu.
HELP or ?	1	Displays the help screens for a particular screen or field. ? is used to display help for a specific field.
LOGO	—	Displays the Adabas Review logo screen.
MENU	12	Terminates the function, and returns you to the Adabas Review Main Menu.
MSG	—	Displays detailed explanations of Adabas Review messages. May be issued with or without a message number. If no message number is included, Adabas Review displays the explanation for the last message received, unless you have changed the Adabas Review screen or performed a different Adabas Review function since the message was displayed. In this case, specifying the MSG command without a message number produces an error.

Other Adabas Review commands are local to a particular system or function. These commands are usually listed on the relevant screen and have PF keys assigned to them.

Line commands used in the list functions are not displayed on the screen. These commands may be displayed by using the online help system.

Using the Online Help System

Online help is available for Adabas Review functions at any time during the Adabas Review session. You may view general help screens describing functions and commands or, if available, help windows for a specific data entry field:

To access the general help screens:

- Press PF1

  Or:

  Enter the HELP command on the command line.

To access specific help windows:

- Enter a question mark (?) on the data entry field.

  If specific help is not available, the general help for the screen is displayed.

The following PF keys are available within the online help system:

PF	Command	Description
3	Exit	Exits from the display.
7	-	Scrolls one page backwards.
8	+	Scrolls one page forward.
9	-	Scrolls to the top of the text.
10	++	Scrolls to the bottom of the text.
12	Menu	Returns to the Adabas Review Main Menu.

Accessing Adabas Online System (AOS) from SYSREVDB

Adabas Online System (AOS) is a selectable unit of Adabas that enables database administrators to monitor and change aspects of an Adabas database interactively. Refer to the *Adabas DBA Tasks Manual* and the Adabas Online System documentation for more information.

If Adabas Online System is installed on your system and you have access privileges to it, you can access it from Adabas Review.

To access Adabas Online System from Adabas Review:

- Enter the code AO on the command line.

Accessing Adabas Review Repositories

The Adabas Review repository is an Adabas file used for storing report definitions, historical data, and target definitions for Adabas Review.

Depending on the configuration at your site, more than one Adabas Review repository may be associated with your system. For example, if your site has Adabas Review installed on more than one database, there may be an Adabas Review repository for each database.

Using the SETFILE Command

The Review command SETFILE (or SET) and SETALL (or SETA) may be used to access different Adabas Review repositories and the reports stored on them.

The SETALL command differs from the SETFILE command in that it allows you to also change the current DBID/HUBID where Adabas Review is communicating. Changing this value has the same effect as entering DBID=nnnnn or HUB=nnnnn into the command line.

To access different Adabas Review repositories:

1. Enter the command SETFILE or SET on the command line of any Adabas Review screen and press ENTER.
2. Provide the DBID and FNR for the Adabas Review repository you wish to access and press ENTER.

   You are now able to access reports stored on a different Adabas Review repository. The setting remains until you either change it again or log off Adabas Review.

Accessing Technical System Information

The TECH command provides access to information about the Adabas Review system. It invokes the Database System Environment screen, which displays Adabas Review environmental and maintenance information. This function is useful for determining:

- the environment in which Adabas Review is executing; and
- the maintenance (zaps) applied to the current version of Adabas Review.

You can access the Database System Environment screen from:

- the Main Menu
- the List Report Definitions (LR) screen
- the List Started Reports (LS) screen

To access technical system information:

- Type TECH on the command line of one of the appropriate screens and press ENTER.

The TECH command Database System Environment screen is the first to appear:

```
10:00:12                A D A B A S  -  R E V I E W               2011-06-27
                        DATABASE System Environment                HUB=00041


 Review Target                                 TP Monitor
 DBID/HUBID ................. 41               Monitor ................... TSO
 SVC ....................... 237              Vers ...................... 3.12
 Vers .................. 4.6_SP0
                                               REVEXIT1
 Review Repository                             Vers ................... 4.6_SP
 DBID ...................... 559              Status .... Installed correctly
 File ...................... 24               State ................. Unknown
 Status ......... File Available
                                               REVEXIT2
 Natural                                       Vers ...................... N/A
 Natural Vers .......... 4.2_SP6               Status .......... Not Installed
 SYSREVDB Vers ......... 4.6_SP0               State ..................... N/A
 ADATMZ Vers ........... 8.2_SP2               Client .................... N/A


Press PF2 for REVIEW fix list, ENTER to exit
Command: ____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help  Fixes Exit  NatFx                                          Menu   ↩
```

From this screen, you can:

- Press PF1 to invoke help from any screen.
- Press PF2 to obtain information about updates (zaps) to Adabas Review nucleus code. The Review Fixes Information screen appears. The following is an example of Adabas Review nucleus zaps displayed on the Review Fixes Information screen:

```
10:03:23                 Review Fixes Information              2011-06-27
             Zaps applied to the Review nucleus for target 41

**** 1002 1003 **** 1005 1006









Press PF2 to see more zaps
                                                                        ↩
```

Pressing PF2 will cycle through a series of zap screens that provide information on option nucleus zaps, REVEXIT1 zaps and optional zaps, and updates (zaps) to Adabas Review Natural objects.

- Press PF4 to obtain information about updates (zaps) to Adabas Review database Natural objects.
- Press ENTER to exit the Database System Environment screen.

The Database System Environment screen displays information about the following environment variables:

Data Category	Field	Description
REVEXIT1	Vers	The version of the REVEXIT1 link routine exit for the TP monitor currently being used. REVEXIT1 is the new name for REVEXITB; REVEXITB is currently supported, but will no longer be supported in a future release.
	Status	The status of the Adabas Review repository identified by the repository DBID.
	State	The current state of REVEXIT1. Possible values are "Unknown", "N/A", "Active" and "Not Active".
REVEXIT2	Vers	The version of the REVEXIT2 link routine exit for the TP monitor currently being used.
	Status	The status of the REVEXIT2 link routine exit for the TP monitor currently being used.

Data Category	Field	Description
	State	The current state of REVEXIT2. Possible values are "Unknown", "N/A", "Active" and "Not Active".
	Client	The status of client reporting for Adabas Review. Valid values are "On", "Off", or "Not Installed".
Review Target	DBID/HUBID	The database ID in local mode or the hub ID in hub mode.
	SVC	The SVC of this Natural session. This is the SVC where SYSREVDB is running.
	Vers	The version of Adabas Review, either local or for the given Adabas Review hub.
Review Repository	DBID	The database of the Adabas Review repository. The default corresponds to NTFILE 241 in the Natural NATPARM parameter module. This database cannot be in the hub address space.
	File	The file number of the Adabas Review repository.
	Status	The status of the REVEXIT1 link routine exit for the TP monitor currently being used.
Natural	Natural Vers	The version of Natural currently installed at the user's site.
	SYSREVDB Vers	The version of the Natural code of the Adabas Review software.
	ADATMZ Vers	The version of ADATMZ code used and loaded by Natural.
TP Monitor	Monitor	The type of TP system currently being used.
	Vers	The version of the TP monitor currently being used.

To obtain information about updates (zaps) to Adabas Review nucleus code:

- Press PF2 from the Database System Environment screen.

  The Review Technical Information screen appears displaying maintenance (zap) information:

```
+------------------------------------------------------------------------------+
!  11:38:26             REVIEW Technical Information             2011-04-12 ↩
 !
!            Zaps applied to the Review nucleus for target 559                ↩
   !
!                                                                             ↩
   !
!  1001 1002 1003 **** 1005 1006 **** 1008 1009 1000 1011 1012                ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
!                                                                             ↩
   !
! Press PF2 to see optional zaps, ENTER to exit                               ↩
   !
+------------------------------------------------------------------------------+
```

From this screen, you can

- Press PF1 to invoke help from any screen.

- Press PF2 to continue through the maintenance screens.
- Press ENTER to exit the Review Technical Information screen.

Maintenance information screens display, by number, all of the zaps that have been applied to the current version of Adabas Review.

The zap numbers are displayed in rows, from left to right in the order they were applied. If the maintenance screen displays released/test zaps, the test zaps are highlighted. Other maintenance screens display optional zaps.

The types of maintenance screens available vary with the operating environment in which Adabas Review resides. Any or all of the following maintenance information screens may be displayed:

- zaps (released/test, and optional) applied to Adabas Review nucleus for the database or hub;
- zaps (released/test, and optional) applied to Adabas Review link routine exit;
- INPLs (released/test) applied to the Adabas Review Natural code; and
- zaps applied to the Adabas link routine (ADALNK).

Note: For z/VSE systems, the first and the second screen display all zaps that are applied to the Adabas Review nucleus and the Adabas Review link routine exit.

To obtain information about updates (zaps) to Adabas Review database Natural objects:

- Press PF4 from the Database System Environment screen.

  The Review Technical Information screen appears displaying the first set of maintenance (zap) information:

```
+------------------------------------------------------------------------------+
!  11:38:26              REVIEW Technical Information             2011-04-12 ↩
 !
!                      Updates to REVIEW DB NATURAL objects                    ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
!                                                                              ↩
   !
! Press ENTER to exit                                                          ↩
   !
+------------------------------------------------------------------------------+
```

From this screen, you can:

- Press PF1 to invoke HELP from any Technical Information function screen.

- Press ENTER to exit the Review Technical Information screen.

Ending a Session

An Adabas Review session may be ended in any one of the following ways:

To end the current function and return to the menu from which the function was called:

- Press PF3; or enter the `EXIT` command on the command line and press ENTER .

To end the function you are using and return to the Adabas Review Main Menu:

- Press PF12; or enter the `MENU` command on the command line and press ENTER.

To end your Adabas Review session from the Adabas Review Main Menu:

- Press PF3 (Exit); or enter the `QUIT` command on the command line and press ENTER.

  Or:

  Press PF12 (Fin); or enter the `FIN` command on the command line and press ENTER.

To end your Adabas Review session without returning to the Main Menu:

- Enter the `FIN` command on the command line and press ENTER

  Or:

  Enter the `QUIT` command on the command line and press ENTER.

7 Cost Accounting Example

Adabas Review may be used to collect Adabas user resource consumption data that can be manipulated for cost accounting and charge-back purposes. This data can be processed and viewed online or in batch, or it can be downloaded to a personal computer for further processing.

This chapter describes the Adabas Review Cost Accounting Example.

Viewing Data Online

A sample cost accounting report "Cost Accounting Example" is supplied with Adabas Review. It may be used without modification, customized to suit site requirements, or used as a basis for other cost accounting reports. The report definition for the "Cost Accounting Example" report is shown below:

```
21:15:18                  A D A B A S  -  R E V I E W                2009-06-18
                                 Edit Report                          HUB=15690
      Detail/Summary: S
      Report Name: COST ACCOUNTING EXAMPLE_________      DBID to Monitor: _____

      +-------------------------------------------------------------------+
      !  Field      Order   Sum   Min   Max   Avg   Pct   Rate   Round  !
      !-------------------------------------------------------------------!
      ! CQJOB___     _10     _     _     _     _     _     _    _______  !
      ! TPUSERID     _20     _     _     _     _     _     _    _______  !
      ! IOS_____     ___     X     _     _     _     _     _    _______  !
      ! COMMANDS     ___     X     _     _     _     _     _    _______  !
      ! CMDRESP_     ___     X     _     _     _     _     _    _______  !
      ! ADADURA_     ___     X     _     _     _     _     _    _______  !
      ! ________     ___     _     _     _     _     _     _    _______  !
      ! ________     ___     _     _     _     _     _     _    _______  !
      ! ________     ___     _     _     _     _     _     _    _______  !
      ! ________     ___     _     _     _     _     _     _    _______  !
      +----------------------------- Page 1 -------------------------------+


Command: _____________________________________________________________________
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help  Optns Exit        Save  Start  -     +          Rules Flds  Menu   ↩
```

In this report, the generated NATURAL display program was modified to attach a cost value to the amounts of CPU time and Adabas time used, and to the number of Adabas calls and I/Os performed by each user.

By issuing the VW command, the report may be used to examine data that has been stored in the Adabas Review repository.

The following screen shows the type of information collected by the report:

```
21:15:59                    COST ACCOUNTING EXAMPLE                    2009-06-18
                 20:15:31 2009-06-18 Thru 21:15:17 2009-06-18          HUB=15690
                                                                       Page:    1
                      Total         Total          Total            Total
 CQ-Job  TPUserid  Num-of-IOs     Commands       Cmd-Resp         ADA-Dur
-------- -------- ------------ ------------ -------------- ----------------

 ALSCCOMP                  210          151          1.1202       5.152688
 ******** SUB-T==>         210          151          1.1202       5.152688
 COMPLETE                    6           12          0.0702       0.251056
          USER1            318          652          1.3057       6.523424
          USER2             57          186          0.2409       0.903504
          USER3           1241         2775          5.5287      26.014992
          USER4             97          219          0.3831       2.337776
          USER5            271          831          1.2486       7.082896
          USER6            169          147          0.4203       2.965808
          USER7            117          129          0.4519       2.433072
          USER8            136          386          0.5802       3.022160
          USER9           5573         8030          7.7051     107.307856
          USER10           323          707          1.3217       6.409696

 Command: ____________________________________________________________________
 Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
       Help  Sort  Exit                                                  Menu
```

Additionally, the displayed resource consumption data may be downloaded to a personal computer using the `DL` command. To use this feature, the Software AG product Entire Connection must be installed.

For more information, read *Downloading Report Output* in the *Adabas Review User's Guide*.

Processing Data in Batch Mode

Cost accounting values can be defined for batch command log processing by using the COST batch parameter statement.

The member REVCOST, shown below, is a sample report for cost accounting that uses the COST statement. REVCOST is located in the Adabas Review source data set.

```
*
* This sample report allocates a cost value to three summary
* fields based on the JOBNAME and Adabas user's TPUSERID.
* In this example, each I/O is assigned a cost of $0.02,
* each command response second is valued at $0.25 and each second of
* Adabas duration costs $0.10.
*
 REPORT TYPE=SUMMARY,
    TITLE='COST ACCOUNTING EXAMPLE',
```

```
    PROGRAM=SR-00028,
    RESTART=Y,
    MAXSTORE=8
 DISPLAY JOBNAME,TPUSERID
 SUM     IOS,CMDRESP,ADADURA
 COST    IOS*0.02,CMDRESP*0.25,ADADURA*0.10
*
```

For additional information about Adabas Review batch processor parameter statements, read *Using Batch Facilities,* in the *Adabas Review User's Guide*.

A sample of the data produced by the batch report is shown on the next section.

Adabas Review Command Response Time Calculation

The Adabas Review command response time calculation is based on the following information.

Adabas Review stores in the command table the minimum Adabas duration for each command type returning a zero response code.

The command table is updated whenever a lower duration value is encountered.

Index

Symbols

A

B

C

D

E

F

G

H

I

L

M

O

P

Q

R

S

T

U

V

W

Z