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.

This chapter covers the following topics:

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.

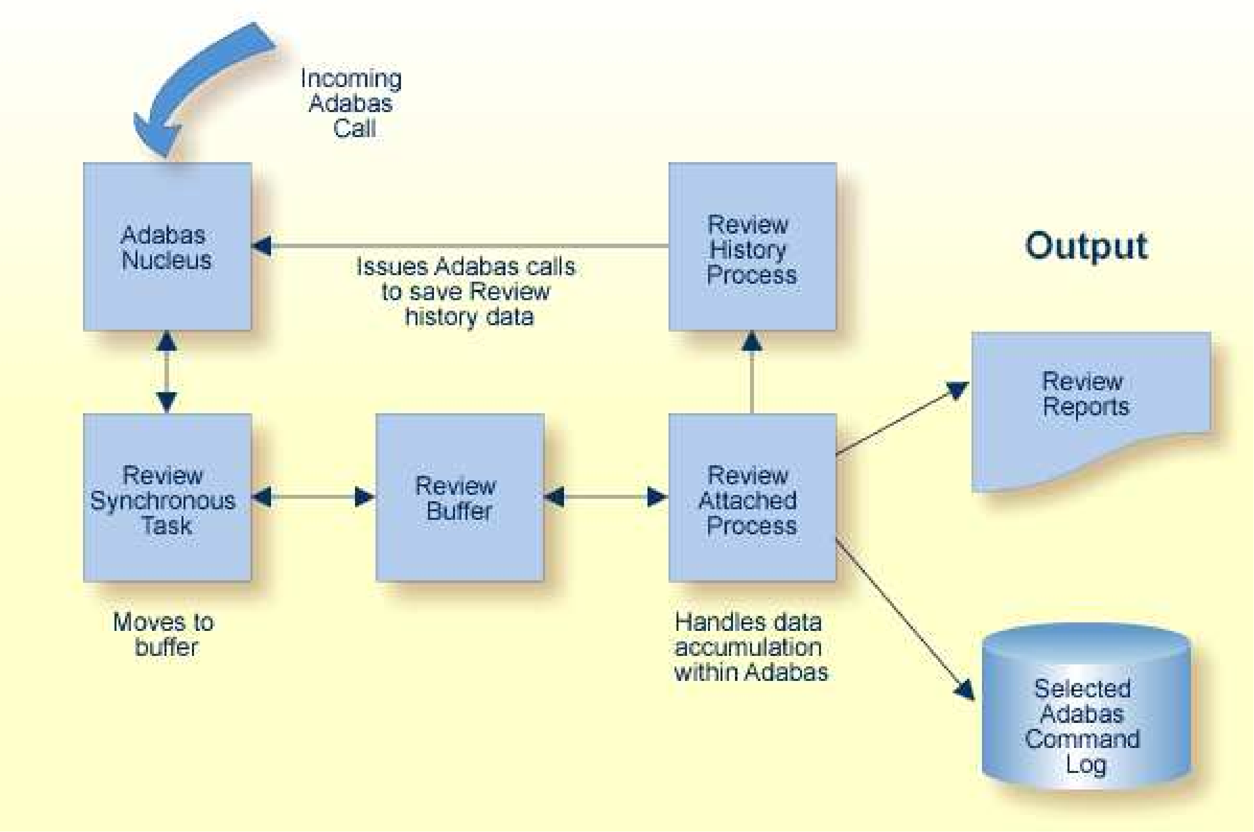

In local mode, Adabas Review job control statements should be added to the Adabas nucleus startup job.

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 Client Interface
- Interface Calls
- Example Client/Server Environment
- Command Log Processing
- Example Hub Mode Data Collection Process

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)*.

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 .

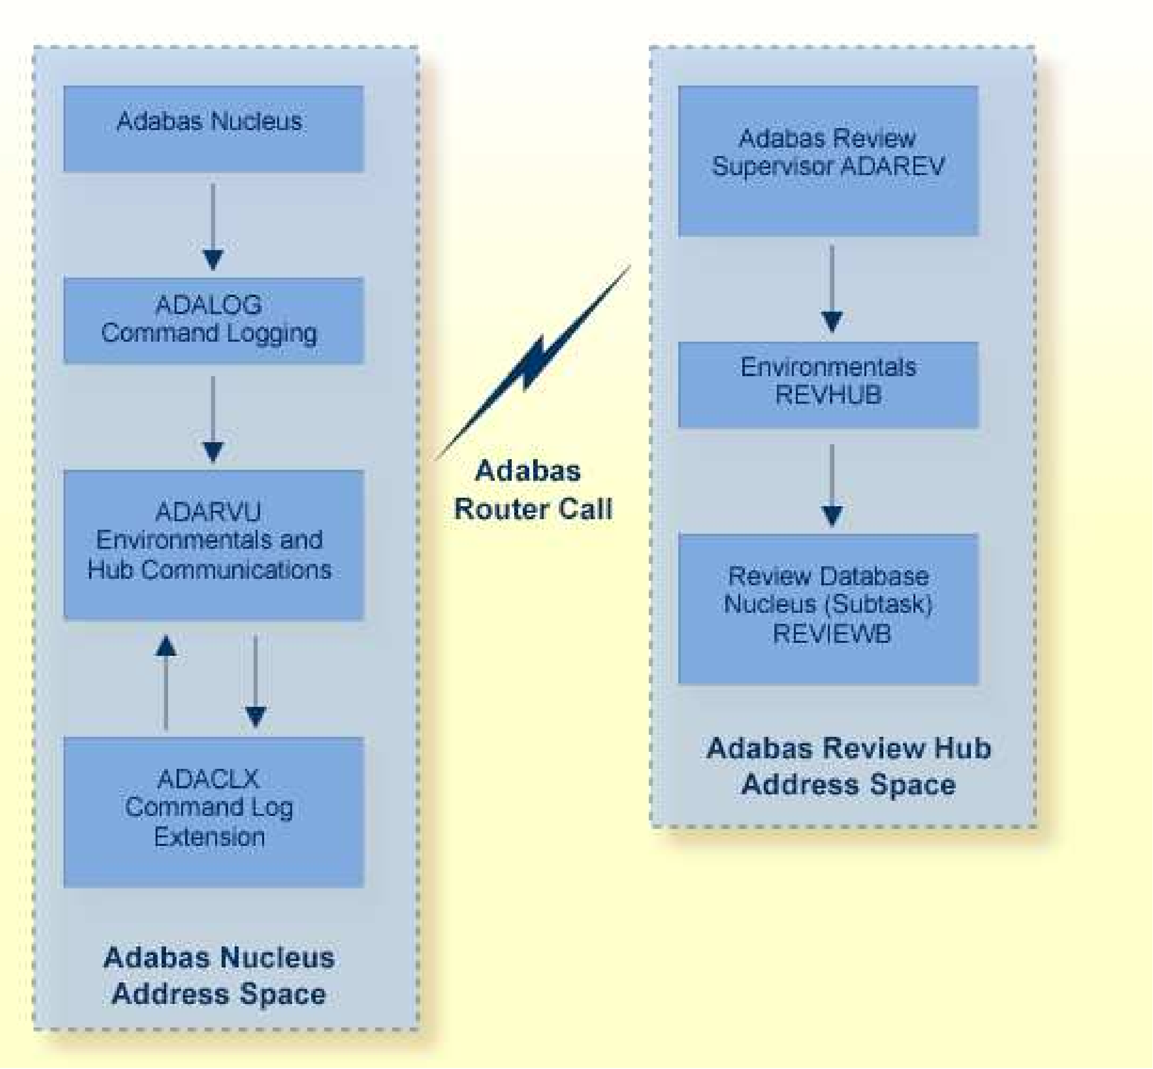
Adabas Nucleus
ADALOG Command Logging
ADARVU Environmentals and Hub Communications
ADACLX Command Log Extension
Adabas Nucleus Address Space
Adabas Router Call
Adabas Review Supervisor ADAREV
Environmentals REVHUB
Review Database Nucleus (Subtask) REVIEWB
Adabas Review Hub Address Space

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.

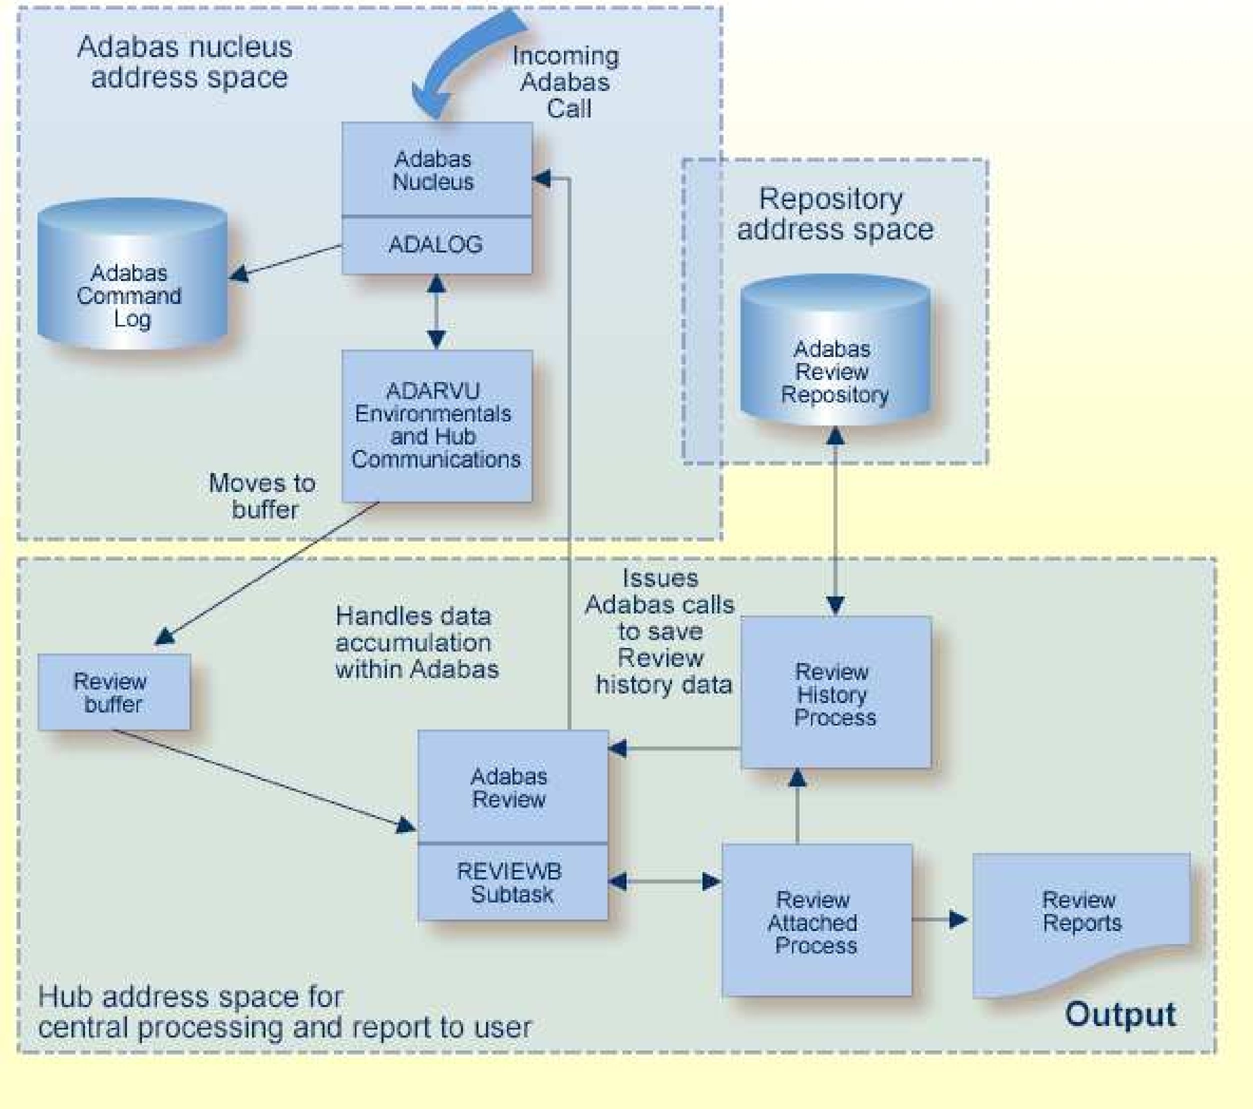

Collecting History Data

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*.

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
- Storage for the REVIEW-BUFFER
- Storage for Reports
- Storage for Online Users
- Storage for Work Areas
- Size of Adabas Review in z/VSE Environments (Local Mode Only)

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*.

- Command Queue
- Attached Buffer

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.