

Introduction

- Terminology
 - What is Predict Application Audit?
 - How Predict Application Audit Works
 - Retiring Objects
 - Modifying Deployments
 - Reporting and Customization
 - Security Aspects
 - User Interface
-

Terminology

This section explains some key product-specific terms used in this documentation. Also included are general terms that have a special meaning within the context of PAA. In descriptions of terms, words in italics have a separate entry.

- Deployment
- File Translation Table (FTT)
- Job
- Load
- Location
- Object
- Production Environment

Deployment

A deployment is a list of locations defined to PAA. It defines the final destinations for objects migrated from PAC, that is, it determines where the migrated objects are to be deployed. You must define a deployment in PAA before a PAC migration event can be submitted. A deployment is named after the application name and status of the PAC application that is to use the deployment.

File Translation Table (FTT)

When an object that was developed referencing one set of databases and files is executed referencing a different set, PAC can dynamically recompile the object using a file translation table. The FTT translates the database and file numbers from the development environment to the destination environment. It enables users to execute applications against different databases and files without changing and recompiling the code.

The FTT to be used in a migration to a particular location is specified in the deployment definition link and invoked during the migration event.

Job

A job is the set of PAA objects and control records created by a load.

Load

A load in PAA terms involves the copying of a set of objects migrated from PAC to their appropriate locations. This is a basic function of PAA.

Location

A location in PAA can be any of the type Natural, (Natural libraries), Predict, (Predict files) or foreign (other library types, for example, PDSs). These locations are distinguished as follows:

- A Natural location is defined by the database number and file number of a Natural system file, the name of a Natural library in this file, and the database number and file number of a Predict system file. Note that the database number and file number of the Natural system file must not be (0,0). For the Predict system file, there is no such restriction.
- A Predict location is defined by the database number and file number of a Predict system file.
- A foreign location is defined by the Entire System Server (ESY) node number and the name of a library (for example, PDS name). Optionally, the definition can contain a volume under OS/390. Specification of LMS type under BS2000/OSD is required. A foreign location is further distinguished in PAA by the (user-defined) types of object that can be loaded in it.

Object

Where objects are discussed in this documentation, they are usually qualified further, depending on their location and state in the PAA environment.

- Backed-up Objects
See Unoperational Objects below.
- Current Objects
Current objects are those objects that actually become operational as a result of an activation.
- Domestic objects
Natural and Predict objects are known as domestic objects in PAA. Examples are: Natural sources, modules (cataloged objects), Xrefs.
- Foreign objects
Non-Natural and non-Predict objects are known as foreign objects in PAA. Examples of foreign objects are PDS members on OS/390 systems, LMS elements on Siemens BS2000/OSD, PL/1 source and loadable procedures.
- Operational Objects
See Current Objects above.
- Scheduled Objects
See Unoperational Objects below.
- Unoperational Objects
Unoperational objects are called scheduled objects before they are activated. Unoperational objects after activation and backup are called backed-up objects. A backed-up object always has a current like object in the same location.

PAA does not handle Predict objects. It does, however, handle Xrefs and DDMs. databases, files, etc. are versioned in PAC and can be migrated to their designated locations without involving PAA.

Each object received by PAA from PAC can be categorized as one of four subtypes: NATURAL, ERROR MESSAGE, DDM, or FOREIGN.

NATURAL and foreign objects are further classified into 2nd and 3rd subtypes as illustrated in the table below (subtypes ERROR MESSAGE and DDM have no 2nd or 3rd subtype):

1st subtype	2nd subtype	3rd subtype
NATURAL	a Natural type e.g. LDA, PROGRAM	SOURCE, LOAD, or Xref
ERROR MESSAGE	void	void
DDM	void	void
FOREIGN	user defined	SOURCE or LOAD

Two PAA objects are called "semilike" if and only if

- their names coincide and
- their first subtypes coincide.

Two PAA objects are called "like" if and only if

- they are semilike and
- if their first subtype is Natural and their third subtypes coincide

Production Environment

A production environment is understood as consisting typically of a number of datasets (for example, PDS libraries in OS/390 systems), a number of Predict system files, and Natural libraries scattered over several Natural system files. These Natural system files may additionally contain libraries outside of PAA control. Natural libraries, Predict files, and PDS libraries are locations in PAA.

What is Predict Application Audit?

Predict Application Audit (PAA) is a tool for controlling and auditing a production environment. It is delivered as part of Predict Application Control (PAC), Software AG's system for tracking changes made to applications during their life-cycle. PAA allows you to define and maintain destinations for objects migrated from PAC to a production environment, make the objects operational, and keep track of all object movements and versions under PAA control.

This basic functionality is supported by a number of features such as back-up of objects before they become operational, and retirement of objects from a production library to automatically revert to a previous version.

A PAA system file distinguishes each PAA system. There can be several of these running concurrently at your site. While a single PAC system can serve several environments, a single PAA system serves a single production environment. In a typical configuration, a single PAC (that is, a single (ACF, PCF) pair) collaborates with several PAAs. Each PAA has an administrator of its own, working to an extent independently of the PAC administrator.

A PAA receives sets of PAC objects transferred by PAC migration events, and copies the objects to the appropriate locations in a load operation.

A single PAC object can be copied to several PAA locations in a single load. Several PAA objects can thus be generated from a single PAC object.

How Predict Application Audit Works

In PAC, application production status links only contain references to locations if the application is Predict. In all other cases, the application production status links contain the database number and file number of the PAA system file that controls the production environment to which objects are to be sent.

To allow objects to migrate from a PAC application to a production status, the PAA system file specified in the application - production status link must contain a deployment record named after the original PAC application and the production status.

Defining deployments in PAA is one of your responsibilities as PAA administrator. Note that the individual locations in a deployment are unknown to PAC and need not concern the PAC administrator.

For example, migrating Natural program PROG1 from a PAC application APPL1 to a production status PROD1 controlled by a PAA system file number 111 in database 222 involves copying PROG1 into all Natural libraries listed in PAA deployment "APPL1 PROD1" in file 111 in database 222.

Similarly, migrating a foreign object, for example of subtype JCL, from APPL1 to PROD1, means copying it to all foreign locations listed in the PAA deployment "APPL1 PROD1" as JCL locations.

- Location Considerations
- Criteria for Loading Objects into Locations
- Loading Objects
- Automatic PAA Job Re-scheduling
- Activating Loaded Objects
- Removing Sets of Objects from Locations

Location Considerations

When you add a location to a deployment, the location is automatically defined to the PAA system file. If the location is a Natural or Predict location, it is checked for validity and accessibility.

Natural and Predict locations must be known by Natural Security, otherwise the validity check will fail. However, you, as PAA administrator, need not be authorized to access the location in Natural. Natural and Predict locations are accessible (that is, available for use by PAA), if it is not marked as controlled by another PAA.

PAA marks the Natural and the Predict locations defined to it, thus protecting the location from use by other PAA installations.

When you specify locations in a deployment, the following restrictions apply:

- You cannot specify location ((1,1),LIB1,(1,2)) in a deployment if location ((1,1),LIB1,(1,3)) is used by the same or another PAA (see also Excluding and Forgetting Locations below).
- Locations ((1,1),LIB1,(1,2)), ((1,1),LIB2,(1,2)), and ((1,3),LIB1,(1,2)) can be used each by a different PAA.
- A Predict location (1,2) can be used by at most one PAA.

Note:

PAA does not mark foreign locations. This means other PAAs can use foreign locations. However, Software AG advises against this.

PAC can use a location specified in a PAA non-production deployment. However, Software AG also advises against this.

Excluding and Forgetting Locations

If you wish to include a location in a PAA deployment, but this location is already used by another PAA, you must first exclude or forget the location from the PAA it is currently used by. Excluding a location from the control of a particular PAA entails the following:

- Removing all objects loaded into the location by PAA. In the case of foreign locations, it means removing all scheduled or backed-up foreign objects from the PAA system file.
- Deleting the object records associated with the location.
- Deleting the location from the PAA's deployments.

Locations can also be excluded from only some of the deployments they are listed in. See *Modifying Deployments*.

Forgetting a location from the control of a particular PAA entails the following:

- All related data except current objects will be purged.
- Deleting the object records associated with the location.
- Deleting the location from the PAA deployments.

Criteria for Loading Objects into Locations

Before loading an object in a specified location, PAA runs a check to prevent objects being loaded in unsuitable locations. The following suitability criteria are used:

1. Any Natural location is suitable for any object whose first subtype is NATURAL or ERROR MESSAGE, and whose third subtype is not Xref.
2. Any Natural location whose Predict system file is not (0,0) is suitable for any object whose first subtype is NATURAL and whose third subtype is Xref.
3. Any Predict location is suitable for any object whose first subtype is DDM.
4. A foreign location is suitable for a foreign object if the type of the object coincides with the type attribute of the location in the deployment.

In all other cases, the attempted load will fail.

As you can see from criteria 1 and 2, Xrefs are treated as Natural objects residing in particular libraries. During a load operation, the Natural locations for Xrefs in a PAA deployment are selected accordingly.

For example, consider a Natural module containing cross-references migrating to a production environment as specified in a PAA deployment. Assume the deployment contains the following Natural locations:

```
((1,1),LIB1,(1,2))
((1,1),LIB2,(1,2))
((1,3),LIB1,(1,4))
((1,3),LIB2,(0,0)).
```

The module itself will be copied into all these libraries (see criterion 1).

The Xrefs of the copies of the module in the first two locations will be put into (1,2); the Xref of the copy of the module in the third location will be put into (1,4), but the Xref of the copy of the module in the fourth location will not be copied at all.

Loading Objects

Objects migrated from PAC can be loaded in PAA locations specified in a deployment either directly or indirectly.

In a direct migration, the PAC migration event takes its destination locations from the deployment of the PAA associated with the PAC application. The objects are loaded in their designated locations without requiring further intervention.

Indirect migration means transferring the objects to a workfile. In this case, you must use the MIGLOAD utility to complete the load operation.

Loading objects in locations consists of:

- Storing a job record in the PAA system file.
- Storing unoperational copies of each received domestic object (first subtype NATURAL, ERROR MESSAGE or DDM) in the suitable locations.
- Storing copies of each received foreign object in the PAA system file.
- Storing an object record for each new object copy (that is, for each newly created PAA object) in the PAA system file.

Domestic objects are kept in their designated locations at all times. Whether they are operational or unoperational depends on the changing value of certain descriptors.

Foreign objects are always operational when they are in their designated locations, and unoperational when they are in the PAA system file.

Each object record indicates the number of the job created by the load. To render the objects operational, the job they belong to must be activated.

Unlike other PAC events, those with production destination status can be submitted before their scheduled time. This results in the load taking place before the time the migrated objects can become operational in the production environment. In this case, PAA ensures that the "early arrivals" remain present but do not become operational before the event's scheduled time. When this time comes, you can activate the job, thus making the objects operational.

Automatic PAA Job Re-scheduling

If during a multiple location delivery any of these locations be inaccessible whilst PAA is trying to deliver objects (for various valid reasons such as a NAT3148), PAA will continue with its delivery to those locations that are accessible and mark those locations that where not accessible.

The PAA job will then complete and automatically set the PAA job into a scheduled state so as the user can resolve the initial location accessibility problem and then activate the original PAA job. This will deliver the objects to the now accessible location.

Activating Loaded Objects

If the PAC event, which initiated the load, was scheduled, then so is the resulting PAA job. You must activate the job when its schedule time has come, provided that no jobs associated with the same deployment with earlier schedule times remain inactivated.

If the PAC event was not scheduled, then the resulting PAA job is also unscheduled and is activated automatically as soon as the load is complete.

Activating an object consists of:

- Optionally backing up a like object in its location.
- Assigning a PAA version number to the object being activated.
- Making the object operational.
- Updating the object record in the PAA system file.

A backed-up object becomes unoperational and its record is updated. The record of the backed-up object contains a pointer to the record of the superseding object, and the record of the activated superseding object contains a pointer to the record of the object it has superseded. In this way, PAA maintains an (object version) succession chain consisting of object records.

Apart from the PAC version number of the original PAC object, each PAA object receives a PAA version number in the object record at activation. This PAA version number is 1 plus the highest PAA version number registered in any object record of a semilike object in the same location.

Removing Sets of Objects from Locations

If necessary, you can remove PAA objects from their locations. PAA ensures that only whole sets of objects are removed and that earlier versions of objects become operational where possible. This feature allows you to dispose of obsolete versions and recover older versions of whole products if the new versions prove deficient. At the same time, inadvertent removal of individual objects (possibly leading to faulty applications) is prevented.

PAA offers the following means of removing objects:

- **Job Backout**

When a job is backed out, all its objects are deleted from their locations (or, in the case of backed-up foreign objects, from the PAA system file). Their records are marked "removed" and excluded from the succession chains. If a current object is removed and an object it superseded is backed up, then this superseded object is made current.

Only jobs that have been activated or finalized can be backed out. A job cannot be backed out if any of its objects has been superseded by subsequent activations and the superseding objects still exist.

- **Job Finalization**

When a job is finalized, all objects whose records precede the records of the job's objects in the succession chains are removed.

Only jobs which have been activated but not backed out can be finalized.

- **Job Purge**

Purging a job means backing it out and deleting all its object records and the job record.

Objects can also be removed using a PAC retirement event, see Retiring Objects.

Retiring Objects

You can remove PAA objects generated from specific PAC versioned objects from PAA locations using the PAC retirement function. Retiring a versioned object causes older versions to move in where possible.

With PAC retirement is an event whose origin status is of type PRODUCTION and whose destination status is RETIRE. The event's application name and its origin status determine the PAA system file and the deployment to be addressed by the retirement.

Retiring PAA objects entails the following:

- The PAA objects generated from the PAC versioned objects in the event's migration list (and originally migrated to the same deployment) are deleted from the locations listed in the deployment (or, in the case of scheduled or backed-up foreign objects, from the PAA system file). . .
- The records of the deleted objects are marked "removed" and excluded from the succession chains.
- If a current object is retired and a back-up object it superseded exists, the superseded object becomes current.

If an entry in the PAC event's migration list carries no version specification, then it addresses the PAC versioned object known by PAC to be in the deployment. If the entry has an asterisk in the version position, then it addresses all PAC versioned objects of the specified name and type migrated to the deployment. You can use an asterisk entry only if a PAC versioned object of the given name and type is known by the PAC to be in the deployment.

Like migrations, retirements can be direct or indirect. An indirect retirement must be completed using the MIGLOAD utility.

Modifying Deployments

Once a PAA deployment has been used by a load, it can only be modified by one of the following means:

- Excluding and forgetting a location from a deployment

See Excluding and Forgetting Locations.

- Purging a deployment

The objects are removed and the object records associated with the jobs that have used the deployment are deleted. The job records and the deployment itself are also deleted.

- Refreshing a deployment.

To refresh a deployment is the same as purging it, except that the deployment record remains and can be modified subsequently. A deployment which since its creation or latest refreshment has not been used by any load is called "dormant".

- Renaming a deployment

Use the direct command RENAME DEPLOYMENT. You are recommended to use this command only in exceptional cases because it alters PAA object data records.

Reporting and Customization

This section summarizes some PAA features that add comfort and flexibility to controlling a production environment.

- Reporting Facility
- Adjusting PAA Locations
- Reloading the PAA System File

Reporting Facility

PAA provides online reports about objects, jobs, locations and deployments. The reporting facility lists objects by name, type, location, job, original PAC application, and state.

Adjusting PAA Locations

If an ADABAS file carrying a PAA location is unloaded, deleted, and reloaded with a different database or file number, then the contents of the PAA system file no longer reflects the environment.

You can adjust the database and the file numbers of PAA locations using the direct command ADJUST.

Note that adjusting the definitions in the PAA system file does not change any database or file numbers in Xref stored data.

Reloading the PAA System File

If the system file of a PAA is unloaded, deleted, and reloaded with a different database or file number, then the tags on the locations under PAA's control no longer reflect the true state of affairs, and the PAA cannot use the locations in the normal manner. In this case, you can adjust PAA itself using the direct command ADJUST.

Security Aspects

Unlike PAC, PAA is normally used only by you, its administrator. A PAA administrator is defined as a user linked to library SYSPAAA.

In order to access PAA's reporting facility, you must be authorized to access library SYSPAA. This adds the flexibility of allowing "read-only" access to PAA for some users.

The following users must be members of the group PAA-AUTH:

- The PAC user initiating a PAC migration to a production environment.
- Users starting the MIGLOAD utility to load objects into PAA locations from a workfile.
- Users initiating direct retirements or completing indirect ones.

User Interface

This section provides a starting aid for novices to the PAA user interface. You can work in two modes:

- Direct Command Mode
- Menu Mode

Direct Command Mode

Most functions available in menu mode can also be executed using PAA direct commands. Though the normal mode of work with PAA is menu mode, direct command mode is useful for quick execution of a function if you are an experienced user and wish to circumvent paging through the menus, or if you wish to route output to a file.

Direct commands can be issued from the Natural NEXT prompt with SYSPAA as the current library or in batch.

All direct commands and how to issue them are described in the section Direct Commands.

Menu Mode

Menu mode is the normal working mode with PAA. You enter menu mode by logging on to library SYSPAA (reporting functions) or SYSPAAA (administration functions) and entering MENU. This displays the corresponding main menu.

You can access the reporting facility from an administration screen by entering MENU in the command line. You can access the administration facility from a reporting screen by entering ADMIN in the command line.

The following menu mode operations are described:

Screen contents and available functions in each screen are described in detail in the sections Reporting Functions and Administration Functions respectively.

Using PAA Menus

The following figure illustrates a PAA menu (the Reporting facility main menu):

```

13:53:33          **** PREDICT APPLICATION AUDIT ****          2000-09-20
  User SAGU                      - Reporting -

                                Code  Entities
                                ----  -
                                D  Deployments
                                J  Jobs
                                L  Locations
                                F  File Translation Tables
                                O  Objects
                                ?  Help
                                .  Exit
                                ----  -

                                Code ..... _

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12--
      Help  Menu  Exit                                          Canc

```

You navigate through PAA menus by typing an appropriate one-character code in the Code field and pressing ENTER. Some menus require that you supply additional parameter values in input fields following field labels. In most cases, these are self-explanatory.

Alternatively, you can press a PF key assigned to a particular function, as displayed at the bottom of each PAA screen.

Executing Functions

You execute functions in menu mode by selecting an item with a function code in a PAA menu, display screen or from a list, and pressing ENTER.

For example, the following figure illustrates the Deployment Selection screen:

13:56:01	**** PREDICT APPLICATION AUDIT ****	2000-09-20
User SAGU	- Deployment Selection -	
C	Application	Status
-	-----	-----
-	CLOVER	PROA
-	CLOVER	PROB
-	CLOVER	PROC
-	HEATHER	PROA
-	HEATHER	PROB
-	HEATHER	PROC
		In Use
		Dormant
		Used
		Used
		Dormant
Command ==>		
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12--		
Help Menu Exit Top Canc		

You can execute a function on an item in the list by entering an appropriate function code in the corresponding input field in the column headed C.

You can select more than one item from a list with a function code. The functions are processed top down every time you press ENTER from the selection list display. To interrupt processing, delete the function code(s) or press a PF key.

Generating Selection Lists

If you do not specify a PAA item unambiguously in a menu, PAA automatically generates a selection list of items, depending on the information you have entered in the parameter input fields. You can also force the generation of a selection list by entering selection criteria using wildcard symbols in the parameter input fields. The following selection criteria are available:

Enter this string	to produce this list
<blank> or *	lists all items
SAG*	lists all items whose names start with "SAG"
SAG123<	lists all items whose names precede SAG123 alphabetically and numerically (including "SAG123" itself).
SAG123>	lists all items whose names succeed SAG123 alphabetically and numerically (including "SAG123" itself)

For example, consider the following entries in the Deployment Menu of the reporting facility:

```

13:54:11      **** PREDICT APPLICATION AUDIT ****      2000-09-20
User SAGU          - Deployment Menu -

          Code  Function
          -----
          D    Display Deployment
          S    Select Deployment
          ?    Help
          .    Exit
          -----

Code ..... S
Application .. NAT*
Status ..... PROD

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10---PF11---PF12---
Help Menu Exit                                     Canc

```

This selection lists all deployments whose application name begins with "NAT" and whose status is "PROD".

Using PF Keys

Some standard functions are assigned to PF keys and are available from every PAA screen. Other functions are context-specific and are available by PF key only in some screens.

Press this PF key	for this standard function
PF1	HELP: display online help text.
PF2	MENU: return to the facility's main menu.
PF3	EXIT: save modifications to data and return to previous screen.
PF12	CANCEL: Undo modifications to data and return to previous screen