

# Bull

## Understanding the Advanced Accounting Subsystem

AIX



**Bull**



# Bull

## Understanding the Advanced Accounting Subsystem

AIX

---

Software

October 2005

**BULL CEDOC  
357 AVENUE PATTON  
B.P.20845  
49008 ANGERS CEDEX 01  
FRANCE**

ORDER REFERENCE  
86 A2 84EM 02

The following copyright notice protects this book under the Copyright laws of the United States of America and other countries which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright © Bull S.A. 1992, 2005

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

### **Trademarks and Acknowledgements**

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX<sup>®</sup> is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

Linux is a registered trademark of Linus Torvalds.

---

# Contents

<b>About this document</b> . . . . .	v
How to use this document . . . . .	v
ISO 9000 . . . . .	v
Related publications . . . . .	v
<b>Advanced Accounting Subsystem</b> . . . . .	1
<b>Data files</b> . . . . .	3
Data file life cycle . . . . .	3
Creating a data file . . . . .	3
Data file management . . . . .	4
Notification messages . . . . .	4
<b>Projects</b> . . . . .	7
Project classification semantics . . . . .	7
Manual project classification . . . . .	8
Project classification through environment variables . . . . .	9
Relative project classification . . . . .	9
Using classification to disable accounting for selected processes . . . . .	9
Adding a project . . . . .	9
Creating a project definition . . . . .	10
Project commands and fast paths . . . . .	10
<b>Policies</b> . . . . .	13
Admin policy . . . . .	13
User and Group policies . . . . .	16
Policy commands and fast paths . . . . .	16
<b>Application Resource Management interfaces</b> . . . . .	19
ARM interface structure . . . . .	19
ARM interface programming model . . . . .	20
Parameters recognized by the ARM implementation . . . . .	21
Internally generated transactional accounting data . . . . .	22
<b>Interval accounting</b> . . . . .	23
System interval accounting . . . . .	23
Process interval . . . . .	23
Interval accounting commands and fast paths . . . . .	23
<b>Hosted accounting policies</b> . . . . .	25
Multi-system accounting policy files . . . . .	26
Configuring an LDAP server to host accounting policies . . . . .	26
Configuring an LDAP client for accounting using the command line . . . . .	27
Uploading the project definition to an LDAP server using the command line . . . . .	28
Uploading the project definition to an LDAP server using SMIT . . . . .	29
Uploading the Admin policy to an LDAP server . . . . .	29
Loading LDAP projects into the kernel using the command line . . . . .	30
Loading LDAP projects into the kernel using SMIT . . . . .	30
Loading LDAP Admin policies into the kernel using the command line . . . . .	30
Loading LDAP Admin policies into the kernel using SMIT . . . . .	30
Downloading LDAP project definitions to a local file using the command line . . . . .	31
Downloading LDAP project definitions to a local file using SMIT . . . . .	31
Downloading LDAP Admin policies to a local file using the command line . . . . .	31

Downloading LDAP Admin policies to a local file using SMIT . . . . .	32
Updating LDAP projects . . . . .	32
Unloading LDAP projects from the kernel. . . . .	33
Unloading LDAP policies from the kernel. . . . .	33
<b>Data aggregation</b> . . . . .	35
System-level data aggregation commands and fast paths. . . . .	35
Project-level data aggregation commands . . . . .	36
<b>Reports and analysis</b> . . . . .	37
Examples of the Process Accounting report . . . . .	38
Example of the Logical Partition Accounting report . . . . .	38
Example of the Transaction Accounting report . . . . .	39
<b>Accounting records</b> . . . . .	41
<b>Index</b> . . . . .	51

---

## About this document

This document provides system administrators with conceptual and procedural information about how to set up, administer, and manage the Advanced Accounting subsystem. Information about projects, policies, transactional accounting, interval accounting, and data aggregation is included. This information is also available on the documentation CD that is shipped with the operating system.

---

## How to use this document

This document is organized to provide overview information and, when needed, conceptual details for major elements of the AIX operating system. Only general descriptions of basic tasks are provided. For task instructions, see the *AIX 5L Version 5.3 System Management Guide: Operating System and Devices*.

## Highlighting

The following highlighting conventions are used in this book:

<b>Bold</b>	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
<i>Italics</i>	Identifies parameters whose actual names or values are to be supplied by the user.
Monospace	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

## Case-sensitivity in AIX

Everything in the AIX<sup>®</sup> operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is not found. Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

---

## Related publications

- *AIX 5L Version 5.3 System Management Guide: Operating System and Devices*
- *AIX 5L Version 5.3 Security Guide*
- *AIX 5L Version 5.3 Installation Guide and Reference*
- *AIX 5L Version 5.3 General Programming Concepts: Writing and Debugging Programs*
- *AIX 5L Version 5.3 Kernel Extensions and Device Support Programming Concepts*
- *AIX 5L Version 5.3 Files Reference*
- *AIX 5L Version 5.3 Performance Management Guide*
- *AIX 5L Version 5.3 Commands Reference, Volume 1*
- *AIX 5L Version 5.3 Commands Reference, Volume 4*
- *AIX 5L Version 5.3 Commands Reference, Volume 5*





---

## Advanced Accounting Subsystem

The Advanced Accounting subsystem, hereafter referred to as Advanced Accounting, is based on mainframe technology and features interval accounting, data aggregation, and dynamic classification of accounting data. You can customize Advanced Accounting for different computing environments. You can configure Advanced Accounting to produce the specific types of records needed for billing applications.

Advanced Accounting provides usage-based information for a wide variety of system resources so that you can develop comprehensive charge-back strategies. You can collect accounting data on resources such as disks, network interfaces, virtual devices, file systems, processors, and memory. Interval accounting gives you the ability to view this data over system administrator-defined time intervals in order to develop chronological views. This has several potential applications, including capacity planning.

Advanced Accounting also provides new statistics from previous accounting tools. For example, the process record provides microsecond-level CPU times, a memory integral based on elapsed time (standard UNIX<sup>®</sup> accounting bases it on CPU times), local and distributed logical file I/O, and local and remote socket I/O.

Interval accounting can be used to periodically record accounting data, enabling the production of more accurate bills. You can configure Advanced Accounting to produce intermediate process records for active processes. These records can be added to the completed process records to produce a bill that reflects the total use of system resources.

Data aggregation is a way to control the amount of data written to a data file. This helps system performance by reducing the overall resource load needed to run Advanced Accounting. Aggregation minimizes a system's I/O requirements, adding accounting records so that fewer records are written to the accounting file. It is transparent to applications and middleware.

Policies are rules that provide for the automatic classification of processes. Classification is done according to users, groups, and applications, categorizing the use of system resources by billable entities. These categories are called *projects*.

APIs are provided so that applications and middleware can describe the transactional nature of their workloads, enabling charge-backs for server processes. These APIs are intended to define and delineate transactions and to identify the end user, if possible. Advanced Accounting measures the resource use of transactions, if possible, and records all of this information in the accounting file.



---

## Data files

Every statistic recorded by Advanced Accounting is written to a data file. When the data file reaches its capacity, the billing application can process the file. After it is processed, that file can be reused, and the cycle repeats.

You must create your accounting data files in order to begin collecting accounting statistics. This involves assessing your company's needs and determining how much data you will be collecting. You can start Advanced Accounting and let it run for a period of time to see how much data is produced, giving you an idea of how much disk space to reserve for accounting, as well as how many files you will need to manage your data.

Although not required, it is a good idea to specify at least two data files so that Advanced Accounting can remain active at all times. Advanced Accounting writes to only one file at a time, but as it does, it requires exclusive access to this file. With two files, Advanced Accounting can write to one file while the other file is processed by the billing application.

---

### Data file life cycle

Each data file goes through a life cycle as it moves through the accounting process.

The following process explains the life cycle of a data file:

1. A pool of pre-allocated data files are created and registered with Advanced Accounting.
2. As accounting statistics are collected from the system, data is written to the active data file.
3. The size of the data file (predetermined when the file is created) limits the amount of data that can be written to the file. Once the data file is full, no more data can be written to it. Advanced Accounting then switches to the next empty data file and begin writing data to it.
4. Once the data file is full, you can then use post-processing tools, such as a billing application, to process the accounting data into statistics for billing purposes.
5. After the data is extracted from the data file and processed, the file is reset so that it can be reused.

---

### Creating a data file

You can create a data file to contain statistics recorded by Advanced Accounting.

This scenario describes how to create a data file using either the System Management Interface Tool (SMIT) or the command line. You can also use Web-based System Manager to create data files.

#### Things to Consider

The information in this how-to was tested using AIX 5.3. If you are using a different version or level of AIX, the results you obtain might vary significantly.
---

1. Change to root user.
2. On a command line, type: `smit create_aacct_file`.
3. In the **Accounting Data File** field, type the name of the file you want to create. You must include the fully qualified path name.
4. In the **Pre-allocated size** field, type the size (in megabytes) of the file you wish to create.
5. Press Enter to save the settings and create the data file.

As an example, to create a 2 MB data file named **testfile** from the command line, type `acctctl fadd /var/aacct/testfile 2`.

---

## Data file management

You can use the command line and SMIT to manage your data files.

The following table shows the commands and SMIT fast paths available for managing accounting data files.

Table 1. Data file management commands

Task	Command	SMIT fast path
Allocate and define an accounting file with specified file name and size. The default size is in megabytes.	<b>acctctl fadd</b> <i>file size</i>	<b>smit create_aacct_file</b>
Remove the specified accounting file from the accounting subsystem. This will not remove the file from the file system.	<b>acctctl frm</b> <i>file</i>	<b>smit release_aacct_file</b>
Indicate that the specified file can be reused by the accounting subsystem.	<b>acctctl freset</b> <i>file</i>	<b>smit reuse_aacct_file</b>
Query the state and current use of the specified file. If no file is specified, all files are queried.	<b>acctctl fquery</b> [ <i>file</i> ]	<b>smit list_aacct_file</b>
Force Advanced Accounting to switch to a new accounting data file. The new file may be optionally specified.	<b>acctctl fswitch</b> [ <i>file</i> ]	<b>smit switch_acct_file</b>

For more information, see *AIX 5L™ Version 5.3 Commands Reference, Volume 1*.

---

## Notification messages

Advanced Accounting continuously writes accounting data to registered accounting data files. It is necessary to monitor the state of these files to ensure that Advanced Accounting always has a place to record data.

Advanced Accounting produces messages to monitor the state of data files and the subsystem status. Use messages related to the data files to trigger actions based on events as they occur. A message is produced when a file is 90% full and when it is 100% full, indicating that administrative action is needed. The **acctctl** command is used to manage accounting data files. For additional information on the **acctctl** command, see **acctctl**.

You can use the **cron** facility to periodically execute a shell script or command that examines and manages accounting data files. By default, AIX records information about the Advanced Accounting subsystem using the **syslog** daemon. For more information about the **syslog** daemon, see *AIX 5L Version 5.3 Commands Reference, Volume 5*.

The **readaacct** command may be used to extract accounting records from accounting data files. It is a sample command, so be aware that it is neither fully tested nor maintained. For more information on extracting accounting records from accounting data files, see “Reports and analysis” on page 37.

## E-mail notification messages

Advanced Accounting sends e-mail messages informing you of activity related to data files.

E-mail notification must be manually configured to function. For more information about configuring e-mail notification, see “Configuring e-mail notification” on page 5.

The following messages are sent through e-mail notification.

Table 2. E-mail notification messages

Subject line	Body
AACCT: File nearly full	1400-330: The accounting file is 90% full.
AACCT: File ready	1400-331: The accounting file is ready for processing.
AACCT: Subsystem out of files	1400-332: The Advanced Accounting subsystem has run out of files for use.
AACCT: Subsystem out of kernel buffers	1400-333: The Advanced Accounting subsystem has run out of kernel buffers.
AACCT: File I/O error	1400-334: The accounting file encountered an I/O error while writing.

Most e-mail programs provide filtering capability so that shell scripts and commands can trigger specific actions when messages are received. Use the string AACCT to identify incoming messages related to Advanced Accounting. To identify the significance of the message, use the message number in the body of the message.

## Configuring e-mail notification

You can configure e-mail notification for Advanced Accounting to send notifications to notify you when a data file is 90% full. Another notification can be sent when the Advanced Accounting subsystem has switched to a new data file.

In the following scenario, you will set up e-mail notification to the address "anywhere@anywhere.com". This scenario describes how to configure e-mail notification using either in SMIT or the command line. You can also use Web-based System administrator to configure e-mail notification.

### Things to Consider

The information in this how-to was tested using AIX 5.3. If you are using a different version or level of AIX, the results you obtain might vary significantly.
---

1. Change to root user.
2. On a command line, type `smit admin_notify`.
3. With the cursor on the **Admin Notification Status** field, press the Tab key to change the e-mail notification status from Off to On.
4. In the **Email-ID for Notification** field, type `anywhere@anywhere.com`.
5. Press Enter to save your settings and turn on e-mail notification.

A valid e-mail address must first be specified using the SMIT interface, and then you can turn notification on and off on a command line. If e-mail notification is activated on a command line, the notifications will be sent to the most recently used e-mail address. To configure e-mail notification from the command line, type:

```
acctctl email anywhere@anywhere.com
acctctl email on
```



---

# Projects

Projects represent billable entities such as users, departments, divisions, companies, or tasks. Each project is composed of a project number, project attribute, and project name, which collectively represent a project definition. Project definitions are entered into the project definition database.

Projects are written to accounting records. Report and analysis commands convert project numbers into project names by looking up entries in the system project definition database. Logically, *projects* are indices into critical business data (for example, customer name, billing address, account number, service level agreement) that are maintained by the billing application.

Project numbers are assigned through project assignment policies, which are composed of project assignment rules. Each rule contains classification criteria that when fully satisfied, yield a classification result. The *classification result* is a project list that is logically assigned to the object, usually a process, being classified. The classification criteria depends on the type of policy.

Project lists enable manual project assignment. If a list of projects is specified, users can change their current project assignment to another project in the list. This provides the capability to launch jobs under different projects, which is useful when you are performing work on behalf of multiple clients. The system administrator can assign multiple projects to any user, and then manually change the project assignment.

---

## Project classification semantics

Project classification semantics are used to classify and assign projects.

For every `exec()`, `initp()`, `setuid()`, and `setgid()` subroutine, the process will be reclassified using the project assignment rules to determine if the current project assignment should be changed. If a project assignment rule is not loaded, or if a rule cannot be successfully applied, then the current project identifier is used.

The default project system identifier is zero (0). It is applied to the base system processes before accounting is enabled and it may be used to signify general system overhead. After it is assigned, a project is inherited from the parent process to the child process using the `fork()` kernel service and `creatp()` kernel service.

The use of the application filter varies between the `initp()` kernel service and the `exec()` subroutine. In the former, the command name of the kernel process that is being started is used to perform the classification. The command name is visible through the `ps` command. In the latter, classification is performed using the FID (inode + device number) of the executable, with the FID of the fully qualified path name that is specified by the policy. Only FIDs that can be computed at policy load time are accommodated.

If a process is classified through a rule that explicitly names an application, then that project identifier should be applied to the process and its children, because the intent is to label a block of work. This implies that subsequent `exec()`, `setgid()`, and `setuid()` subroutines in the process, as well as its children, do not perform project reclassification. Internally, these processes are identified with a sticky bit that indicates they were classified through a rule that explicitly named an application. The sticky bit is visible through the `ps -P` command.

The following is an example of the output after you run the `ps -P` command. The asterisk (\*) preceding the `dev` project indicates it has a sticky bit attached to it:

UID	GID	PID	TTY	TIME	PROJECT	SUBPROJ	CMD
0	0	16922	pts/1	0:00	*dev	0	ps
0	0	19206	pts/1	0:00	*dev	0	acctras
0	0	22286	pts/1	0:00	*dev	0	ksh

When you load a new policy file, all of the processes in the system are reclassified, except those with a sticky bit. Processes with a sticky bit cannot be changed.

When a different project identifier is assigned to a process, a process accounting record is written to the **acct** file, so that the use of resources by a project may be accurately reported. Whenever this occurs, the accounting statistics for the process are reset to zero.

---

## Manual project classification

Users can manually change their current project assignment.

For non-privileged users, authorization is provided through the policy file. The first project listed in the rule is considered to be the default project and is automatically chosen, unless the user changes his current project assignment to another project in the list. Changing the current project assignment applies only to the current session and it is not silently applied to other jobs with the possible exception of the parent process (for example, the shell).

The following rule is an example of a project list.

*Table 3. Example of a project list*

User	Group	Application	Project(s)
User1	-	-	Chemistry, Biology

Because Chemistry is the first project in the list, it is the default project. User1 can change his current project assignment to Biology. The current project assignment is used as a parameter to project classification, which might include Biology as a potential project assignment.

Project lists can also be specified in rules that name applications, allowing administrators to associate projects with specific instances of applications. This is important for server processes, which may be externally configured to manage specific resources that have different cost structures.

This mechanism enables system administrators to charge indirectly for those resources that may not directly support usage-based accounting. For example, a system administrator that wants to charge different rates to access specific databases could start different instances of oracle to manage those databases.

The following rule an example of how applications are specified:

*Table 4. Example of how applications are specified*

User	Group	Application	Project(s)
-	-	/usr/bin/Oracle	Dataset1, Dataset2
User1	-	-	Chemistry, Biology, Dataset2

In the first rule, the application Oracle has two projects associated with it. If Oracle is started and the current project is not Dataset2, then Dataset1 is chosen. Otherwise, Dataset2 is chosen. If User1's current project assignment is Chemistry or Biology and he starts his database application, it will be assigned the project identifier Dataset1. User1 could also choose to change his current project assignment to Dataset2 and invoke a database application again, which would result in the project assignment of Dataset2.

For privileged users (root users or users with Advanced Accounting capability), there is a force option (-f) that ignores the rules.



---

## Project classification through environment variables

You can use environment variables to classify processes.

You can use the environment variables `PROJECTNAME=project name` and `PROJECTID=project id` to classify processes. If you have root authority or Advanced Accounting administrator capabilities, then the assignment is made without consulting project assignment rules. Otherwise, the project is validated against the loaded policy. If both `PROJECTNAME` and `PROJECTID` are specified, `PROJECTNAME` takes precedence over `PROJECTID`.

---

## Relative project classification

You can enable project assignments to be made relative to User IDs and Group IDs with Advanced Accounting.

To relatively classify a project, use the `$UID` and `$GID` keywords in a project list to tell the system that the project code must be computed. A simple expression in the form *keyword* or *keyword + constant* may be used, where *constant* is either a decimal or hexadecimal number (0xffff).

The following rule illustrates the use of relative project classification:

Table 5. Relative project classification

User	Group	Application	Project
*	-	-	\$UID + 1000000

---

## Using classification to disable accounting for selected processes

You can disable accounting for selected processes through the classification process.

You can disable accounting by specifying a "NoAccounting" project list. This attribute is inherited from the parent process to child process.

The following example illustrates how this can be accomplished:

Table 6. Example of the No Accounting specification in a project list

User	Group	Application	Project
Oracle	Oracle	/usr/*/oracle	NoAccounting
Root	Root	kbiod	NoAccounting

Oracle is classified in the first rule by the subroutine `exec()`. The NFS kproc in the second rule is classified by the subroutine `initp()`.

---

## Adding a project

You can create a project in Advanced Accounting.

Projects represent billable entities. Each user, department, division, or company that your business bills must be represented by a project in Advanced Accounting. The `projectl add` command is used to create a project. Once created, the project is stored in the project definition file, located by default in the `projdef` file in the `/etc/project` directory.

**Note:** When you are adding a project definition, make sure aggregation is disabled. To enable aggregation for a project definition, you must first create the project definition with aggregation off, and then manually enable aggregation for that project. For more information about data aggregation, see “Data aggregation” on page 35.

## Creating a project definition

You can create a project definition that includes billable entities.

This procedure describes how to create a project definition. Projects represent billable entities. Each project is composed of a project number, project attribute, and project name, which collectively represent a project definition. The following scenario describes how to create a project definition from both SMIT and a command line. The following describes how to create a project definition called `mktg`.

### Things to Consider

The information in this how-to was tested using AIX 5.3. If you are using a different version or level of AIX, the results you obtain might vary significantly.

1. Change to root user.
2. From a command prompt, type: `smit add_proj`.
3. In the Project Name field, type: `mktg`.
4. Assign a project number to the project. For example, in the Project Number field, type `10`.
5. Verify that Aggregation field is set to **Off**. You cannot create a project in SMIT with the Aggregation field set to **On**.
6. Add an appropriate comments in the **Project Comments** field. For this scenario, in the Project Comments field, type `This is the default project for the Marketing department`.
7. Press Enter to save the values and create the project definition.

To create a project definition called `mktg` from the command line, type `projctl add mktg 10 off "This is the default project for the Marketing department"`.

## Project commands and fast paths

You can manage your projects on the command line or using SMIT fast paths.

The following table shows the commands used to administer projects:

Table 7. Project commands

Task	Command	SMIT fast path
Add a project definition.	<code>projctl add projname projnumber [comment] [{-d projpath   -p [DN]}</code>	<code>smit add_proj</code>
Load or reload a project definition.	<code>projctl ldprojs [{-g [DN]   -p [DN]}] [-d projpath] [-r] [-a]</code>	<code>smit load_proj</code>
Remove the specified project definition.	<code>projctl rm projname [{-d projpath   -p [DN]}</code>	<code>smit rename_proj</code>
Show or change active project definitions.	<code>projctl chg projname [-p pid [, pid] ] [-f]</code>	<code>smit show_chg_proj</code>
Merge two project definitions.	<code>projctl merge sourceprojpath [-d targetprojfile]</code>	None
Start a program with a project assignment.	<code>projctl exec projname [cmd line] [-f]</code>	<code>smit start_proj_prg</code>

Table 7. Project commands (continued)

Task	Command	SMIT fast path
Enable or disable aggregation for a project.	<b>projctl chattr agg</b> <i>projname</i> {-s -u} [{-d <i>projpath</i>   -p [ <i>DN</i> ]}]	None
Show the policies that are loaded on the system.	<b>ojctl qpolicy</b> [-g [ <i>DN</i> ]]	
List all active project definitions.	<b>projctl qprojs</b> [-n]	None
List the specified project definition.	<b>projctl qproj</b> <i>projectname</i>	<b>smit show_chg_proj</b>
Change the project assignment for a process.	<b>projctl chg</b> <i>projname</i> [ -p <i>pid</i> [, <i>pid</i> ] ] [-f]	<b>smit chg_proj_proc</b>
Show project assignments for a program.	<b>projctl qapp</b> <i>appname</i>	<b>smit show_proj_pgm</b>
Manage projects.	None	<b>smit work_project</b>
Remove project definitions.	None	<b>smit remove_admin_proj</b>

For more information, see *AIX 5L Version 5.3 Commands Reference, Volume 4*.



---

## Policies

Policies automate project assignment. A policy is comprised of classification criteria and classification results. The project assignments take place using subroutines and kernel services, such as **exec**, **initp**, **setuid**, and **setgid**.

You can use policies to classify data by user, group, application, or a combination of these attributes. Depending on the type of policy that is created, administrators can specify a user name, group name, application name, and project list in the policy file, although all four components do not have to be present for a policy to function.

Processes can be assigned in the following two ways:

- Using assignment rules when process classification attributes change. This is the most common way that processes are classified.
- Manually assigned to a class by a user with the required authority.

---

## Admin policy

The admin policy uses the user name, group name, and application name process attributes to classify processes. The admin policy is application-based, and provides the ability to gather accounting statistics at the application level.

The admin policy supports the most filters and the use of Korn shell syntax wildcards. Configuring and maintaining Admin policies requires a standalone administration tool, such Web-based System Manager, SMIT, or an editor.

By default, the admin policy is located in the **/etc/project** directory. You can create alternate admin policies to use at various times. For example, you might want to have an admin policy that runs Monday through Friday, and another that runs on Saturday and Sunday. The alternate admin policies are stored in subdirectories of the **/etc/project/alter** root directory. You must specify a name for the subdirectory when you are creating a new admin policy.

## Admin policy assignment rules

The admin policy consists of one or more assignment rules that are placed in the admin policy file.

Admin policy assignment rules must conform to the following syntax:

*user name:group name:application name:Project List::Optional comments go here*

For details on the values that are allowed in each field of the assignment rule, see the following table.

*Table 8. User, group, and application rules*

Type of rule	Description
user	Can contain either a hyphen (-) or at least one valid user name as defined in the <b>/etc/passwd</b> file. An exclamation point (!) can be used before a name to exclude a given user from the class. The list of user names is composed of one or more names, separated by a comma (,). Patterns can be specified to match a set of user names, using full Korn shell pattern-matching syntax. If a nonvalid or an incorrect user name is used, Advanced Accounting ignores the entire rule. If a hyphen (-) is used, Advanced Accounting skips to the next field in the rule.

Table 8. User, group, and application rules (continued)

Type of rule	Description
group	Can contain either a hyphen (-) or at least one valid group name as defined in the <b>/etc/passwd</b> file. An exclamation point (!) can be used before a name to exclude a given user from the class. The list of group names is composed of one or more names, separated by a comma (.). Patterns can be specified to match a set of user names, using full Korn shell pattern-matching syntax. If an incorrect user name is used, Advanced Accounting ignores the entire rule. If a hyphen (-) is used, Advanced Accounting skips to the next field in the rule.
application	Can contain either a hyphen (-), a list of application path names, or the command name of a kernel process. This is the path name of the applications (programs) run by processes included in the class. The application names will be either full path names or Korn shell patterns that match path names. The list of application names is composed of one or more path names, separated by a comma (.). An exclamation mark (!) can be used before a name to exclude a given application. At least one application in the list must be found at load time or the rule is ignored. Rules that are initially ignored for this reason might become effective later on if a file system is mounted that contains one or more applications in the list.

For process assignment policies, the classification criteria is the user name as listed in the **/etc/passwd** file, the group name as listed in the **/etc/groups** file, and fully qualified application name. The classification result is a project list. By default, the first project in the list is used, unless the user changes his current project assignment to a secondary project in the list.

Classification is done whenever an attribute value changes by comparing the value of these process attributes against lists of possible values given in the class assignment rules file, called *rules*. The comparison determines which rule is a match for the current value of the process attributes.

To classify the process, Advanced Accounting examines the top-level Admin policy for the active configuration to determine in which class the process belongs. For each rule in the file, Advanced Accounting checks the current values of the process attributes against the values and lists of values specified in the rule. Advanced Accounting goes through the rules in the order in which they appear in the Admin file, and classifies the process in the project corresponding to the first rule for which the process is a match. Therefore, the order of the rules in the rules file is significant.

The following is a list of the criteria used to determine whether the value of a process attribute matches the values of the same attribute field in the Admin policy file:

- If the field in the rules file has a value of hyphen (-), then any value of the corresponding process attribute is a match.
- If the value in one of the fields is preceded by an exclamation point ("!"), that value is always excluded.
- If the value in one of the fields is followed by an asterisk (\*), then every match with that value will be recognized.

## Examples of Admin policy rules

These examples show how you can use admin policy rules.

The admin policy allows you to specify more than one user, group, or application in their respective fields. For instance, if you wanted Frank, Bob, and John to be given the same attributes, you would specify the following syntax:

```
User1,User2,User3:-::-Project List::Comments
```

The preceding syntax shows that User1, User2, and User3 will be treated the same way for this rule. The dashes in the group name and application fields are wildcards. You can also use an asterisk (\*). As previously mentioned, you can also use wildcards to include all the values of a certain attribute. For example, if you wanted to include every user name beginning with B, you would type B\* in the User Name field. Full Korn shell pattern matching syntax is allowed in all of the fields of the rule.

You can also set up your admin policy is to include certain users, but exclude others:

```
User1,!User2,User3::-:Project List::Comments
```

The preceding syntax shows that User1 and User3 will have the same attributes, but the policy will exclude User2.

**Note:**

1. The kernel only reads numeric values. In the above example, the user names User1, User2, and User3 are converted to numeric values after they are loaded into the kernel.
2. If any changes are made to the policy, it must be reloaded into the kernel using the **projectl** command.

## Aliases for Admin policies

Aliasing provides a shorthand for referring to lists of users and groups, which simplifies admin policies, and makes them more readable and easier to maintain.

You can define an alias name for a list of users or groups so that you don't have to enter the full list of users, groups, or application names in a rule. A separate alias file is used with each admin policy and is placed in the same directory as its associated Admin policy.

To create an alias where User1, User2, and User 3 are grouped into an alias named dev1 would be defined as follows:

```
dev1:User1,User2,User3::Development team 1
```

To use the alias name must be precede the alias name in the admin policy with a dollar sign ('\$'). For example, using the dev1 alias instead of User1, User2, and User3 in the admin policy looks similar to the following:

```
$dev1::-:Project List::This is a rule within a user alias
```

Multiple aliases, separated by a comma, can also be used. Aliases are excluded by placing an exclamation point (!) before the alias name in the admin policy rule.

## Alternate Admin policies

You can create alternate Admin policies that are enforced in different situations that you define.

It is possible to create multiple Admin policies to reflect different billing strategies based on the time of the day or the day of the week. The **projectl** command is used to load alternate Admin policies. Previously loaded Admin policies do not have to be unloaded to load a new policy. The **cron** facility is used to load a policy at the a given time.

Alternate Admin policies are placed in the **/etc/project/alter** directory. For example, an alternate Admin policy with the name "weekend" is placed in the Admin file **/etc/project/alter/weekend/admin**. If this policy used aliases, they would be placed in the **/etc/project/alter/weekend/alias** alias file. SMIT and Web-based System Manager are used to create, change, show, load, unload, and remove Admin policies. Alternate policies are addressed in SMIT by changing the current focus to the name of the policy.

## Creating an Admin policy

The Admin policy is used to automate project assignment.

This policy allows you to gather accounting statistics at the application level. The default Admin policy is named **admin** and is located in the **/etc/project** directory. If you create other admin policies, they are stored in a subdirectory of the **/etc/project/alter** directory. You can either create a subdirectory using the **mkdir** command, or using SMIT.

---

## User and Group policies

The User and Group policies use the process attributes for users and groups. A user or group name and a project list constitute a project assignment rule within the User or Group policies. There is no file associated with a User or Group policy as user policy data is stored in a security database.

The User and Group policies provide project assignment based exclusively on user names or group names, depending on the policy. They are integrated with user administration. Project lists can be specified for users and groups when these users and groups are created.

The following table shows the structure of a user policy:

*Table 9. Structure of a user policy*

User name	Project list
User1	project1,project 2
User2	biology,chemistry

To create a Group policy, complete the following steps:

1. Create a group by typing the following:  
`mkgroup staff`
2. Create a project by typing the following:  
`projct1 add biology_dept 1200 Project Comment`
3. Associate the project `biology_dept` with the newly created group `staff`:  
`chgroup projects=biology_dept staff`

**Note:** To create a user policy, complete Steps 1 through three, substituting **mkuser** and **chuser** for **mkgroup** and **chgroup**.

Once a user or group has been created, a project is associated with it. The users or groups can have several projects associated with them, but only one can be active at a time.

## Creating a User or Group policy

You create User and Group policies by associating projects with a user or group.

User and Group policies are used to automate project assignments for users and groups. The **chuser** command and the **chgroup** command are used to associate projects with users and groups.

---

## Policy commands and fast paths

You can manage your policies on the command line or using SMIT fast paths.

The following table shows the commands used to administer policies.

*Table 10. Policy commands*

Task	Command	SMIT fast path
Show the policies that are loaded on the system.	<b>projctl qpolicy</b>	<b>smit query_policy</b>



Table 10. Policy commands (continued)

Task	Command	SMIT fast path
Load policies.	<b>projectl ldadm</b> [{"-g" [{"name:}DN   name]}]   -p [{"name:}DN   name]} ] [-d admproj] [-r] [-a]  <b>projectl ldusr</b> [-a] [-r]  <b>projectl ldgrp</b> [-a] [-r]  <b>projectl ld</b> [-a] [-r]  <b>projectl ldall</b> [-d localadm] [-r] [-a]	<b>smit load_admin</b>  <b>smit load_users</b>  <b>smit load_groups</b>
Unload policies.	<b>projectl</b> {unldusr unldgrp unldall   {unldrojs   unldadm} [{"-p" [name]   -g}]] [-f]} [-a]	<b>smit unload_admin</b>  <b>smit unload_users</b>  <b>smit unload_groups</b>
Create an Admin policy.	None	<b>smit create_admin</b>
Show or change current focus policy.	None	<b>smit change_show_focus</b>
Remove an Admin policy.	None	<b>smit remove_admin</b>
Add a rule.	None	<b>smit add_admin_rule</b>
Remove a rule.	None	<b>smit remove_admin_rule</b>
Add a user alias.	None	<b>smit add_usr_alias</b>
Add a group alias.	None	<b>smit add_grp_alias</b>
Show or change the specified alias.	None	<b>smit chg_alias</b>
Remove the specified alias.	None	<b>smit remove_alias</b>
Create a project list for a user.	<b>chuser</b> <i>user</i>	<b>smit create_user</b>
Create a project list for a group.	<b>chgroup</b> <i>group</i>	<b>smit create_group</b>
Show or change a specified project list for a specified user.	<b>chuser</b> <b>projects=</b> <i>projectlist</i> <i>user</i>	<b>smit change_show_user_list</b>
Show or change a project list for a group.	<b>chgroup</b> <i>group</i>	<b>smit change_show_group_list</b>
Remove a project list for a specified user.	<b>chuser</b> <b>projects=</b> <i>user</i>	<b>smit remove_user</b>
Remove a project list for a specified group.	<b>chgroup</b> <b>projects=</b> <i>group</i>	<b>smit remove_group</b>
Show project lists for all users.	<b>lsuser</b> -a <b>ALL</b>	None
Show project lists for all groups.	<b>lsgroup</b> -a <b>projects ALL</b>	None

For more information, see *AIX 5L Version 5.3 Commands Reference, Volume 4*.



---

## Application Resource Management interfaces

Applications use the Application Resource Management (ARM) interfaces to describe the transactional nature of their workloads.

Advanced Accounting supports ARM interfaces by recording information that is presented through these interfaces in the accounting data file. To use this information for charge-back purposes, it is necessary to understand the programming model associated with the ARM interfaces and the mechanism for passing critical business information to the Advanced Accounting subsystem, so that this information can be preserved for your billing application.

---

### ARM interface structure

The ARM programming model has a fundamentally hierarchical structure.

Transaction instances are derived from transaction definitions that are defined when registering transactions. Application instances are derived from application definitions that are defined when registering applications. When starting a transaction, a transaction definition that should be used within an application instance is specified, enabling the full set of information to be defined for each transaction instance. For example, all transactions have application and group names.

Advanced Accounting supports the ARM interfaces through a hierarchy of records:

- Application environment record
- Transaction environment record
- Transaction instance record

The application environment record describes a unique combination of application information including the application name, application group name, and properties (both identity and context). Each application environment record is assigned a unique identifier so that it can be referred to symbolically. This identifier is called the *application environment identifier*, and it is included within the transaction instance record.

The transaction environment record describes a unique combination of transaction information, including the transaction name and properties. Each transaction environment record is assigned a unique identifier so it can be referred to symbolically. This identifier is called the *transaction environment identifier*, and it is included with the transaction instance record.

Application environment identifiers and transaction environment identifiers are long lived, but they are not persistent in nature. Each time the system boots, the identifiers are regenerated. Advanced Accounting circumvents this problem by recording the application and transaction environment in each file so report and analysis commands should use the entry in the current file to determine the unique combination of values that should be applied to a transaction.

The transaction instance record is used to describe each individual transaction. It includes the application environment identifier and the transaction environment identifier defined above. These identifiers should be used to look up the corresponding application environment record and the transaction environment records so that the application names, application group names, and transaction names can be associated with the transaction instance. The intent is to minimize the amount of data that needs to be recorded with each transaction instance since a lot of it is static in nature.

Advanced Accounting also supports the aggregation of ARM accounting data. For more information about data aggregation, see "Data aggregation" on page 35.

---

## ARM interface programming model

The ARM interface programming model includes routines that perform basic functions.

The following table describes the basic programming model.

*Table 11. ARM interface programming model*

<b>Interface</b>	<b>Description</b>
arm_register_application	This routine registers the application identity with the ARM implementation. The application identity provides the fundamental instrumented application scope and basis for subsequent ARM 4.0 calls. This routine is called during application initialization.
arm_start_application	This routine establishes the started application instance with the ARM implementation, in preparation for making ARM 4.0 transaction calls. This interface is called during application initialization, after registering the application with ARM.
arm_register_transaction	This routine registers a transaction identity with the ARM implementation. A transaction identity provides a category of transactions, commonly referred to as a transaction type, that will execute under the registered application for subsequent ARM 4.0 calls for transaction monitoring and measurement. This interface is called during application initialization after starting the ARM application instance.
arm_start_transaction	This routine establishes the started transaction instance with the ARM implementation. A started transaction instance provides the foundation for transaction monitoring and measurement. This routine is called during transaction processing.
arm_block_transaction	This routine indicates that the started transaction instance has become blocked behind a particular event. This routine is called during transaction processing.
arm_unblock_transaction	This routine indicates that the blocking event has been relieved for the started transaction instance. This routine is called during transaction processing.
arm_bind_transaction	This routine indicates that the current thread is performing on behalf of a started transaction instance. Binding enables the system to measure the processor utilization of a transaction because it establishes an exclusive processing relationship between the thread and the transaction. This routine is called during transaction processing.
arm_unbind_transaction	This routine indicates that the current thread is no longer performing on behalf of a started transaction instance. This routine is called during transaction processing.
arm_stop_transaction	This routine ends the started transaction instance recognized by the ARM implementation. A call to this interface is the expected way for an instrumented application to end a started transaction instance.
arm_stop_application	This routine ends the started application instance recognized by the ARM implementation. A call to this interface is the expected way for an instrumented application to end a started application instance.

Table 11. ARM interface programming model (continued)

Interface	Description
arm_destroy_application	This routine deregisters the registered application. A call to this interface is the expected way for an instrumented application to deregister an application.

## Parameters recognized by the ARM implementation

The ARM APIs provide a way to delineate application transactions through the use of a named set of parameters that are recognized by the ARM implementation and Advanced Accounting.

The parameters set for ARM enable the operating system to identify them and Advanced Accounting to measure and record them in the accounting data file. The ARM APIs also enable applications to describe their transactions, so that site-specific information can be provided. This is accomplished largely by convention.

Advanced Accounting recognizes the parameters shown in the following table:

Table 12. Parameters recognized by the ARM implementation and Advanced Accounting

Parameter	Description
<i>application name</i>	The application name is specified through the <i>app_name</i> parameter to the <i>arm_register_application()</i> interface. This parameter is set by the application and cannot be overridden by the user. It names the application, so a name similar to "IBM® DB2 Universal Database™" should be used.
<i>application group name</i>	The application group name is specified through the <i>app_group</i> parameter to the <i>arm_start_application()</i> interface. This parameter provides a grouping of the individual application instances that are collectively configured to provide an integrated service, so something identifiable such as "Sample Supply Chain Management" should be used.
<i>transaction name</i>	The transaction name is specified through the <i>arm_register_transaction()</i> interface. The name that is provided should be descriptive so the data file can be analyzed to determine the type of operation that transpired.
<i>identity properties</i>	Identity properties are used to specify properties that never change values. Identity properties can be specified for registered applications and transactions so they can be used at different levels to describe the fixed nature of an application or transaction. Identity properties can be used to identify discount rates or account codes that are known by the Advanced Accounting subsystem. Account codes (similar to project codes) should be specified through the predefined property: EWLM: AIX: Account Class

Table 12. Parameters recognized by the ARM implementation and Advanced Accounting (continued)

Parameter	Description
<i>context properties</i>	Context properties are used for information that changes. Context properties are associated with application and transaction instances, although Advanced Accounting does not capture accounting data for the latter. Only the name of the context property is captured for a transaction instance. Accounting codes (similar to project codes) can also be specified through context properties. In this cases, use the predefined property:  EWM: AIX: Account Class

## Internally generated transactional accounting data

Advanced Accounting also captures data that is generated internally. This information cannot be changed by the user, but is needed by accounting.

The following table describes the internally generated data.

Table 13. Internally generated transactional accounting data

Parameter	Description
<i>User name and identifier</i>	The user name and identifier is derived from the user sub-buffer that is specified when starting a transaction. It is intended to identify the agent that started the transaction. It can be an IP address or a machine name. It is not necessarily a UNIX user name or user ID.
<i>Response and queued times</i>	Response and queued times describe the quality of service that was provided to the transaction. Response and queued times identify the elapsed time of the transaction and the time spent in the application before any real action was taken.
<i>Resource utilization</i>	The resource utilization of a transaction describes the physical use of resources such as the amount of processor time spent on the transaction. This statistic cannot always be computed because an application may work on multiple transactions simultaneously. In fact, this statistic will not be calculated, unless the application uses the <code>arm_bind_thread()</code> interface, which dedicates a thread to a single transaction.

---

## Interval accounting

Interval accounting provides a way to collect accounting data at specified intervals. You can configure it to produce intermediate process records for active processes.

Records from interval accounting can be added to the completed process records to produce a more accurate bill reflecting the total use of the system by a user in a given time period. Interval accounting can also be configured to periodically capture accounting data for system resources like processors, memory, disks, network interfaces, and file systems. This information can be used to generate a bill that reflects the total use of a partition.

Usage information about system resources has other applications beyond chargeback. It can be used to perform capacity planning because it shows the use of physical resources like processors, memory, disks, and network interfaces. Interval accounting provides a time-based view of these resources, so that load can be determined, enabling capacity-based decisions to be made based on the data. For example, it is possible to determine the idle processor capacity in a given interval, which can be used to determine whether more processors are needed.

Interval accounting provides a historical view of resource use, which can be used for performance analysis. For example, the file system records can be examined to determine which file systems were busy when access to the system was slow. The data identifies multiple busy file systems served by the same disk adapter. This information can be used to balance file systems over disk adapters. You do not always know which file sets are being accessed at a given time, so having the ability to replay a log with this information is an asset.

---

## System interval accounting

System interval accounting collects system-related information.

System interval accounting collects resource use information about system resources such as processors, disks, network adapters, file systems, and memory usage. This information is used to profile the use of your system. For most purposes, an interval of once per hour should be sufficient.

The system interval does not collect process accounting data.

---

## Process interval

Process interval accounting collects process-related information.

The process interval is used to capture information about long-running jobs such as data modeling applications that run for months. The standard process completion record is sufficient for most processes. Therefore, the process interval should be relatively large, so that it does not produce unnecessary data. The process interval writes a record for every active process, including newly started jobs. In some situations this can amount to a large number of records, so the process interval should be set to capture information once per day (1 440 minutes). However, if process records are being aggregated automatically by the system, then set the process interval for once per hour.

---

## Interval accounting commands and fast paths

You can manage interval accounting using the command line or SMIT fast paths.

By default, interval accounting is turned off. The following table shows the commands you can use to turn on and turn off system-level and process-level interval accounting:

Table 14. Interval accounting commands

Task	Command	SMIT fast path
Enable system interval accounting, specified in number of minutes by the <i>time</i> parameter, or turn off system interval accounting.	<b>acctctl isystem</b> { <i>timeloff</i> }	<b>smit system_interval</b>
Enable process interval accounting, specified in number of minutes by the <i>time</i> parameter, or turn off process interval accounting.	<b>acctctl iprocess</b> { <i>timeloff</i> }	<b>smit process_interval</b>
Query the state of Advanced Accounting.	<b>acctctl</b>	None

For more information, see *AIX 5L Version 5.3 Commands Reference, Volume 1*.



---

## Hosted accounting policies

An accounting policy is hosted by another system enabling a single point of control to be established for the management of accounting policies.

Using a single point of control simplifies the administration and maintenance of accounting policies. You can use this capability to account for a user who logs into multiple machines across an enterprise. Alternatively, or in conjunction with this capability, you can define server-specific billing strategies. The Advanced Accounting subsystem enables multiple policies to be defined concurrently for a server. You can use this flexibility to implement the right solution for your business.

You implement hosted accounting policies using the Lightweight Directory Access Protocol (LDAP), which defines a standard mechanism for accessing and updating information in a directory (a database) either locally or remotely using a client-server model. The Advanced Accounting subsystem uses this technology to provide a single point of control for the management and distribution of accounting policies and projects to LDAP clients. You must configure each LDAP client system to access accounting data on the LDAP server, but once configured, the client behaves as if the policy were defined locally on that system.

The Advanced Accounting subsystem is fully integrated with user authentication. When creating a new user or group, you can optionally specify a project list for that user or group. This project list is part of the definition of that user or group and is automatically delivered to the client system when the user logs in. The project list is just another user attribute.

In general, you should establish project definitions with the same scope as the user. If a user is defined globally through LDAP, then the user's project definitions should also be defined globally through LDAP. Similarly, if a user is defined locally on a particular system, then the user's project definitions should be defined locally on that system. This simplifies the billing strategy, but is not required by the Advanced Accounting subsystem. Both project repositories can be loaded at the same time to accommodate local and LDAP users. It is recommended that Local and LDAP project definitions are defined in a non-overlapping way, although the Advanced Accounting subsystem does not enforce this. Projects are resolved based on the order that the project repositories are loaded. Local projects may be given precedence over LDAP projects by loading the local project repository before the LDAP project repository, and vice versa. It is also possible to only use local or LDAP projects.

The following advanced-accounting data sets can be used with LDAP:

- Project definitions
- Project lists for LDAP users and groups (for example, LDAP User and Group Policies)
- Admin policy

Admin policies can also be stored on an LDAP server. The Admin policy provides an alternative mechanism for performing classification that in some ways is easier to manage than specifying a project list for each user. The Admin policy is a collection of assignment rules that can be more easily updated because it is not distributed across many individual user definitions. The Admin policy is managed as a single entity. You can define both local- and LDAP-based Admin policies.

You can enable the following policies at the same time. The order of evaluation is:

1. Local Admin policy
2. LDAP Admin policy
3. User accounting policy
4. Group accounting policy

The LDAP server is not Advanced Accounting-aware and can be maintained at a different software level. You can upload a schema that describes Advanced Accounting data to the server, so that it can be used to distribute accounting policies and project definitions without any special knowledge of project and policies.

You use the **mkprojldap** command to configure the connection between the LDAP server and client. Specifically, you use this command to upload the LDAP schema associated with the Advanced Accounting data to the server. This command is also used to define the location on the server where the data is to be stored for a particular client. This enables you to implement a different Admin policy and project repository for each client system, if desired. These items are individually configurable to provide the maximum flexibility. One reason for implementing a different Admin policy could be for administrative purposes, such as making the policy reflect the set of authorized users. Or, you might need to accommodate a different billing strategy for a particular server. For example, the use of server X is always charged to account Y, or server X is intended to be used for projects W and Z only.

You must individually configure each client system for accounting purposes. You must specify the set of policies that should be activated on each system by using the **projectl** command. This command has been extended to provide new functions such as uploading and downloading LDAP-based projects and Admin policies. In general, once the client system has been set up, the location of the policy and project repository is transparent to the end user.

The Advanced Accounting subsystem produces accounting data locally. The **acctctl** command must still be used to define data files and to manage them on an ongoing basis. However, you might want to place these files on shared storage subsystems such as a Storage Area Network (SAN) or in a distributed file system like NFS or General Parallel File System (GPFS) so that a billing application can have access to all of the data.

You can use the **ps** command to show the origin of a project that has been assigned to a process. This information is also recorded in the accounting record for a process, so that report and analysis tools can properly match the assigned project to the proper project repository, assuming that the tool is aware of multiple project repositories. This situation is best avoided by defining non-overlapping ranges of projects for local and LDAP-based projects.

---

## Multi-system accounting policy files

You can upload or download a project definition file and Admin policy files to an LDAP server.

The multi-system accounting policy includes the following data sets:

- Default LDAP projects
- Default LDAP Admin policy
- Alternate LDAP Admin policy

These data sets are stored on the local system in the following directories:

- **/etc/project/ldap/projdef**
- **/etc/project/ldap/admin**
- **/etc/project/ldap/alter/policy name/admin**

You can load multiple Admin policies to the LDAP server. You can use this capability to implement time-based policies, such as policies for peak and off-peak usage.

---

## Configuring an LDAP server to host accounting policies

You must configure the LDAP server to host accounting policies before they can be provided to a client system.

You can perform the setup procedure from any LDAP client that is configured in a general way to access the LDAP server. It is not necessary to perform the setup procedure on the LDAP server. To configure an LDAP server to host accounting policies, you must first upload the Advanced Accounting subsystem schema, which is shipped with AIX. The schema describes the layout of accounting data so that the LDAP server does not need to be Advanced Accounting-aware.

There is no requirement that the LDAP server be at the same software level as the client.

To set up the LDAP server, run the following command for each LDAP server:

```
mkprojldap -u -h hostname -D bindDN -w BindPassword
```

You must then decide where to store accounting data on the LDAP server. Each client system asks for accounting data at a particular location, so it is important to understand the layout of accounting data on the server. You must understand the billing strategy to define the proper layout on the server. If you want to deploy a server-specific billing policy, then you should use Admin policies, since they can be targeted at a specific machine. In this case, you should place the Admin policy and project definitions in a location on the LDAP server that is reserved for that system.

If you want to use an enterprise-level policy that always classifies a user in the same way, then you should use the User or Group policies. In this case, you must define the project repository in a global location on the LDAP server so that it can be accessed by each client. Other strategies are also possible.

To define a base location on the server where you can store accounting data, use the following command:

```
mkprojldap -s -h hostname -D bindDN -w BindPassword -i InstallPoint
```

For example:

```
mkprojlap -s -h ldap.svr.com -D cn=root -w passwd -i -p cn=aixdata,o=ibm -a cn=aixdata,o=ibm
```

This enables Admin policies and project definitions to be stored on the server below the install point. You must run this command as a root user once for each base location.

## Configuring an LDAP server to host accounting policies using SMIT

You can configure an LDAP server to host multi-system accounting policies using SMIT.

To configure an LDAP server to host multi-system accounting policies using SMIT, use the following procedure.

1. Log in as a root user.
2. Access the SMIT menu for the LDAP server setup using the following SMIT path: **smitty aacct** → **Manage Advanced Accounting Subsystem** → **Manage LDAP configuration** → **LDAP server setup**.
3. Enter the server name, Bind DN, Bind password, and the install points for projects and policies in the relevant fields.
4. Press Enter to configure the LDAP server for the Advanced Accounting subsystem.

---

## Configuring an LDAP client for accounting using the command line

You can set up an LDAP client for accounting purposes using the command line.

This topic describes how to set up an LDAP client for the Advanced Accounting subsystem. This section assumes that the client system is configured as an LDAP client.

Use the **mksecldap** command to establish the basic connection between the LDAP client and server.

Use the **mkprojldap** command to provide accounting-specific parameters to establish an accounting-aware LDAP client.

Use the **projctl** command to configure projects and policies as required by the billing strategy. The final step involves specifying a refresh policy for data provided by an LDAP server.

To configure an LDAP client, perform the following steps:

1. Log in as a root user.
2. Run the command `mkprojldap -c -D bindDN -w bindPWD -a default-adminDN -p default-projectDN`, where *default-adminDN* and *default-projectDN* are the base locations on the LDAP server on which the client will look for accounting data. This command adds accounting-specific information to the LDAP configuration file (**ldap.cfg**) and restarts the LDAP client daemon. This is an example of the command: **mkprojldap -c -D cn=testroot -w testpwd -a ou=adminpolicy,ou=aacct,cn=aixdata -p ou=projects,ou=aacct,cn=aixdata**.
3. If you want to upload LDAP projects or Admin policies to the LDAP server, you can do so at this point.
4. If you want to configure the current system to automatically use LDAP projects when a policy is loaded, run the command **projctl ldprojs -g -a**. Projects are resolved in the order that they are loaded. Therefore, if you want local projects to take precedence, run the **projctl ldprojs -a** command first. The **-g** option indicates that data is to be retrieved from the LDAP server. You must configure both sources if you intend to use both sources.
5. If you want to configure the current system to automatically load an LDAP Admin policy, when accounting is started, run the command **projctl ldadm -g -a**. You might also want to configure a local Admin policy. You can do this by running **projctl ldadm -a**. Unlike projects, there is not a precedence issue between Admin policies. The local Admin policy takes precedence over the LDAP Admin policy.
6. Use the **cron** facility to periodically refresh projects and Admin policies that are loaded from the LDAP server. The interval might be once an hour or once a day depending on the site-specific policy for accommodating new users.

You can also perform the above steps through SMIT.

After the client has been configured to use the accounting capabilities provided by an LDAP server, it is not necessary to have LDAP-specific knowledge to administer the Advanced Accounting subsystem, unless you want to add new project definitions to the LDAP project repository or modify an LDAP-based accounting policy or project definition.

---

## Uploading the project definition to an LDAP server using the command line

After you have configured the LDAP server to host accounting data, you can upload the project definitions to the configured DN in the LDAP server from the command line.

By default, the project definitions are uploaded from the default LDAP project definition path **/etc/project/ldap/projdef**. You can specify alternate path names using the **-d** option with the **projctl** command.

To upload the project definition file to an LDAP server from the command line, perform the following steps:

1. Place the file containing the project definitions in the directory **/etc/project/ldap/** with name **projdef**. The directory **/etc/project/ldap/** acts as local repository for the LDAP projects and policies.
2. Run the command **projctl ldprojs -p -d /etc/project/ldap/**. This command uploads the **projdef** file to the default project DN on the LDAP server. To upload **projdef** to a non-default DN, specify the DN parameter with **-p**.

---

## Uploading the project definition to an LDAP server using SMIT

When the LDAP server has been configured to host accounting data, you can upload the project definitions to the configured DN in the LDAP server using SMIT.

To upload the project definition file to an LDAP server through SMIT, use the following procedure:

1. Log in as a root user.
2. Access the SMIT menu **Upload Project Definitions to LDAP Server** using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Upload Project Definitions to LDAP Server**.
3. Enter the local project definition file path and destination location on the LDAP server.
4. Press Enter to upload the project definition file to the LDAP server.

---

## Uploading the Admin policy to an LDAP server

You can upload the Admin policy to an LDAP server.

You can use the command **projectl ldadm -p** to upload an Admin policy to the LDAP server. By default, the source Admin policy is found at the file location **/etc/project/ldap/admin**.

To specify that a different policy be uploaded, use the **-d** command argument. By default, the target location for the policy on the LDAP server is defined by the Admin policy DN that was established when the client was set up. You can upload the policy to a different target location on the LDAP server by specifying a DN with the command argument **-p**. The DN is an optional parameter.

If the Admin policy contains new project definitions, or if this is the first time that the policy is uploaded to the server and the LDAP project definitions have not been uploaded, then you must upload the project definitions so that the policy is resolved correctly when it is downloaded to another client. The Advanced Accounting subsystem downloads project definitions automatically from the LDAP server when a policy is loaded, but it does not automatically upload projects when a policy is uploaded. You must perform this step, because the project repository on the server is a critical resource and the upload operation is a replacement operation.

## Uploading the Admin policy to an LDAP server using the command line

You can upload the Admin policy to an LDAP server from the command line.

To upload the Admin policy file to an LDAP server, use the following procedure.

1. Place the file containing the Admin policies in the directory **/etc/project/ldap/** with name **admin**. The directory **/etc/project/ldap/** acts as local repository for the LDAP projects and policies.
2. As a root user, run the command **projectl ldadm -p -d /etc/project/ldap/**. This command uploads the Admin policy file to the default project DN on the LDAP server. To upload the file to a non-default DN, specify the DN parameter with **-p**.

## Uploading the Admin policy to an LDAP server using SMIT

You can upload the Admin policy to an LDAP server using SMIT.

To upload the Admin policy file to an LDAP server through SMIT, use the following procedure:

1. Log in as a root user.
2. Access the SMIT menu Upload Admin Policy to LDAP Server using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Automatic Project Assignment** → **Work with Admin Policies** → **Upload Admin Policy to LDAP Server**.

3. Enter the local Admin policy file path and destination location on the LDAP server. By default, the policy file will be uploaded with name **default**. A different name can be specified using the **Admin Policy Name** option.
4. Press Enter to upload the Admin policy to the LDAP server.

---

## Loading LDAP projects into the kernel using the command line

Before you use the LDAP projects and policies, you must first load the LDAP projects and policies into the kernel. You can load LDAP projects into the kernel using the command line.

When loading the LDAP projects into the kernel, the projects should be present on the default project's DN on the LDAP server.

To load the LDAP projects into the kernel, perform the following steps:

1. Log in as a root user.
2. Run the **projctl ldprojs -g** command. This command downloads the **projdef** file from the LDAP server to the **/etc/project/ldap/** directory and then loads the project definitions in the file into the kernel.

---

## Loading LDAP projects into the kernel using SMIT

If you want to use the LDAP projects and policies, you must first load the LDAP projects and policies into the kernel. You can load LDAP projects into the kernel using SMIT.

When loading the LDAP projects into the kernel, the projects should be present on the default project's DN on the LDAP server.

To load the LDAP projects into the kernel through SMIT, perform the following steps:

1. Log in as a root user.
2. Access the SMIT menu **Load/Re-load Project Definitions** using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Load/Re-load Project Definitions**.
3. Select LDAP as the project repository.
4. Press Enter to load the LDAP project definitions into the kernel.

---

## Loading LDAP Admin policies into the kernel using the command line

You can load the LDAP Admin policies into the kernel using the command **projctl ldadm -g**

To load the LDAP projects into the kernel, the Admin policy file must be present on the default administrator DN on the LDAP server. To load LDAP projects into the kernel, perform the following steps:

1. Log in as a root user.
2. Run the **projctl ldprojs -g** command. This command downloads the **projdef** file from the LDAP server to the **/etc/project/ldap/** directory, and then loads the project definitions in the file into the kernel.

---

## Loading LDAP Admin policies into the kernel using SMIT

You can load the LDAP Admin policies into the kernel using the SMIT.

To load the LDAP Admin policies into the kernel through SMIT, use the following procedure.

1. Log in as a root user.
2. Access the SMIT menu **Load/Re-load Admin Policy** using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Automatic Project Assignment** → **Work with Admin Policies** → **Load/Re-load Admin Policy**.

3. Select LDAP as the **Current focus is on**.
4. Press Enter to load the LDAP Admin policies into the kernel.

---

## Downloading LDAP project definitions to a local file using the command line

You can download project definitions from the LDAP server to the client to append new project definitions or to delete obsolete ones, and then upload the file. You can edit the project definition file directly, rather than using the **projectl** command to perform these tasks.

You can download an LDAP project definition file to a local directory using the command **projectl**. This command downloads the project definition file from the default project DN on the LDAP server to the specified local directory. To download the project definition file from a specific Distinguished Name (DN), you must specify the DN parameter with the **-g** command argument.

The Advanced Accounting subsystem downloads project definitions automatically from the LDAP server when an accounting policy is loaded (if the client system has been configured to use LDAP as a project repository). The project definitions are downloaded during the client setup phase, when you explicitly load project definitions from an LDAP server and indicate that they should be automatically configured. To do this, use the **projectl ldprojs -g -a** command.

LDAP projects are cached on each client system. You must define a refresh policy for periodically downloading accounting data from the server. There is a SMIT menu item that allows you to specify the frequency of policy refreshes. This delay might impact the ability of the client to properly classify a new user, if the new user is given access to the system before the project is automatically downloaded to that client system. In this event, the user can use the system, but the user is not properly classified. A **SYSLOG** message indicating that the user is not properly classified is generated automatically. You can manually invoke the refresh command on the client systems where users are expected to perform their work. The refresh command is **projectl ldall -r**.

---

## Downloading LDAP project definitions to a local file using SMIT

You can download project definitions from the LDAP server to the client to append new project definitions or to delete obsolete ones, and then upload the file. You may prefer to edit the project definition file directly, instead of using the **projectl** command to perform these tasks.

To download LDAP project definitions through SMIT, use the following procedure.

1. Log in as a root user.
2. Access the SMIT menu **Load/Re-load Project Definitions** using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Download Project Definitions from LDAP server**.
3. Enter the local project definition file path and location on the LDAP server.
4. Press Enter to download the project definition file.

---

## Downloading LDAP Admin policies to a local file using the command line

You can download a local copy of the LDAP Admin policies to the client using the command line. You can make any required updates to the local copy of the LDAP Admin policies file before uploading.

To download the LDAP Admin policy file to a local directory, use the command **projectl ldadm -g -d local-dir**. This command downloads the Admin policy file along with the associated alias and project definition files from the default administrator and project DN on the LDAP server to the specified local

directory. To download the Admin policy file from a specific DN, specify the DN parameter with the `-g` option. When the Admin policy is downloaded from a specific DN, the project definition file will not be downloaded.

---

## Downloading LDAP Admin policies to a local file using SMIT

You can download a local copy of the LDAP Admin policies to the client using SMIT. You can make any required updates to the local copy of the LDAP Admin policies file before uploading.

To download the LDAP Admin policy file through SMIT, perform the following steps:

1. Log in as a root user.
2. Access the SMIT menu Load/Re-load Admin Policy using the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Automatic Project Assignment** → **Work with Admin Policies** → **Download Admin Policy from LDAP Server**.
3. Enter the local Admin policy file path and location in the LDAP server fields. If no Admin policy file name is specified, the default policy file will be downloaded. To download a named Admin policy file, fill in the **Admin Policy Name** field.
4. Press Enter to download the Admin policy.

---

## Updating LDAP projects

You can perform any required updates to LDAP projects without downloading them to the client machine.

To update LDAP projects, use the **projectl** subcommands **add**, **rm** and **chattr**.

To update the project repository on the LDAP server, use the `-p` option.

## Updating LDAP projects using the command line

You can add or remove a new project definitions to the LDAP project definition file from the command line.

To add a new project definition to the LDAP project definition file, use the command **projectl add projname projnumber -p**. This adds the new project definition to the project definition file on LDAP default project DN.

To remove a project definition from the LDAP project definition file, use the command **projectl rm projname -p**. This removes the specified project definition from the project definition file on the default project DN.

To enable the aggregation property of a project definition in an LDAP project definition file, use the command **projectl chattr agg projname -s -p**. This enables the aggregation property of the project definition specified on the LDAP project definition file.

For all of the above operations, if a project definition on a specific LDAP DN needs to be modified, then specify the DN parameter with the `-p` option.

## Updating LDAP projects using SMIT

You can add or remove a new project definitions to the LDAP project definition file using SMIT.

To add a new project definition to the LDAP project definition file through SMIT path, use the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Add Project Definition**.



To remove a project definition from the LDAP project definition file through SMIT, use the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Remove Project Definitions**.

To enable the aggregation property of a project definition in an LDAP project definition file using SMIT, use the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Show/Change Project Definitions**.

For all of the above operations, if a project definition on a specific LDAP DN needs to be modified, then specify the DN parameter with the `-p` option.

---

## Unloading LDAP projects from the kernel

You can unload LDAP projects from the kernel.

To unload the LDAP projects from the kernel, use the command **projctl unldprojs -g**. This command removes only the project definitions that are loaded from the LDAP repository.

To unload the LDAP projects from the kernel through SMIT, use the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Project Definitions** → **Unload Active Project Definitions**.

---

## Unloading LDAP policies from the kernel

You can unload the LDAP Admin policies from the kernel.

To unload the LDAP Admin policy from the kernel, use the command **projctl unldadm -g**. This command only removes the Admin policies loaded from LDAP repository.

To unload the LDAP Admin policy from the kernel through SMIT, use the following SMIT path: **smitty aacct** → **Manage Project Definitions and Assignments** → **Automatic Project Assignment** → **Work with Admin Policies** → **Unload Admin Policy**.



---

## Data aggregation

Data aggregation is a method of accumulating data into an accounting record that is otherwise presented through multiple records.

Aggregated data is written periodically, according to the intervals described in “Interval accounting” on page 23.

**Note:**

1. Interval accounting must be enabled in order to use data aggregation.
2. Set the process interval and system interval to 60 minutes.

Accounting data is totaled inside the kernel with no impact to applications or middleware. The data is made available by interval accounting, which is a kernel function that periodically writes these records to the **acct** file. When the kernel aggregates records, it maps them to a set of aggregated records that are managed internally. These records are pending in the sense that they have been input to the system, but have not been committed to the recording mechanism inside Advanced Accounting. Interval accounting is used to commit the aggregated records so that they are written to the **acct** file.

Because aggregated data is recorded using different data structures, you must verify that the billing application recognizes these structures. Consult the documentation provided with your billing application to determine if your billing application supports data aggregation. Data aggregation can be enabled or disabled at the system level or project level.

---

## System-level data aggregation commands and fast paths

You can manage system-level data aggregation using the command line or SMIT fast paths.

System-level data aggregation commands allow you to turn on and off process data aggregation, kernel extension data aggregation, and ARM transaction data aggregation. The following table lists the commands you can use to turn on and off process data aggregation, kernel extension data aggregation, and ARM transaction data aggregation.

*Table 15. System-level data aggregation commands*

Task	Command	SMIT fast path
Enable or disable system-wide process aggregation.	<b>acctctl agproc {on off}</b>	<b>smit system_paggr</b>
Enable or disable system-wide aggregation for third-party kernel extensions.	<b>acctctl agke {on off}</b>	<b>smit system_kaggr</b>
Enable or disable system-wide aggregation for ARM transactions.	<b>acctctl agarm {on off}</b>	<b>smit system_aaggr</b>
Query the overall accounting state	<b>acctctl</b>	None

For more information, see *AIX 5L Version 5.3 Commands Reference, Volume 1*.

---

## Project-level data aggregation commands

You can manage project-level data aggregation using the command line or SMIT fast paths.

The following table shows the commands to turn on and off aggregation for individual projects.

Table 16. Project-level data aggregation commands

Task	Command	SMIT fast path
Enable or disable aggregation for a project, specified by the <i>projname</i> parameter.	<b>projctl chatr</b> <i>aggprojname</i> <b>{-s -u}</b> <b>[-d]</b> <i>projpath</i>	None
Query the aggregation state for all projects.	<b>projctl qprojs</b> <b>[-n]</b>	None
Query the aggregation state for the specified project.	<b>projctl qproj</b> <i>[projectname]</i>	None

For more information, see *AIX 5L Version 5.3 Commands Reference, Volume 4*.

---

## Reports and analysis

The Advanced Accounting subsystem provides accounting data for a wide variety of resources that are typically included within charge-back mechanisms.

The format of the accounting data file, as well as the format of the individual accounting records are in the **sys/aacct.h** are in the header file. Use the **acctrpt** command to parse the accounting data file and produce accounting reports. You can use this command to produce three types of accounting reports: Process Accounting reports, LPAR Accounting reports, and Transaction Accounting reports.

Along with the **acctrpt** command, AIX provides a sample program that parses the accounting data file. This program is **readaacct**. This program can be used to analyze accounting data and import accounting data into a spreadsheet. This command has the following syntax:

```
/usr/sample/aacct/readaacct [-F file] [-t trid] [-b begin_time] [-e end_time] [-c] [-h]
```

This command has the following options:

- The **-c** flag is used to display information in colon-separated format.
- The **-h** flag is used to display information about the file, such as the hostname and machine model, and serial number of where the data was produced.
- The **-b** and **-e** flags give a time-based view of the information.
- The **-t** flag gives a record-based view of the information.

The following is an example of the output displayed after running the **readaacct -F /tmp/afile -h** command:

```
# readaacct -F /tmp/afile -h
File Name=/tmp/a
Version=0
Flags=0
Offset=3084288
File Size=3145728
State=2
ID=1
First Time=1087266596
Last Time=1087301336
System ID=IBM,01025990A
System Model=IBM,7040-681
Host Name=bigboylp9
Partition Name=bigboylp9
Partition Number=9
-----
Transaction ID=1
Flags=f1
Transaction Project=0
Sub project ID=0
Transaction start time=6-14-2004 21:29:56
UID=0
GID=0
PID=335912
eWLM Service Class=0
Flags=1
Command Name=acctctl
Controlling Terminal's Device Number=23,5
Process Start Time=1087266596
WLM Class key=7770295601810996315
Incrementing Statistics:
Elapsed process time=0.032406 seconds
Elapsed thread time=0.032406 seconds
Process CPU time=0.015238 seconds
Elapsed Page seconds of disk pages=0 seconds
```

Elapsed Page seconds of real pages=5 seconds  
 Elapsed Page seconds of virtual memory=5 seconds  
 Bytes of local file I/O=0  
 Bytes of other file I/O=0  
 Bytes of local sockets=0  
 Bytes of remote sockets=0

---

## Examples of the Process Accounting report

You can generate Process Accounting reports.

### Process Accounting report

The following is an example of the Process Accounting report:

```
# /usr/bin/acctrpt -f /var/aacct/acctdata
```

TIMESTAMP	PROJID	UID	GID	PID	CMD	STARTED	EXITED	(C) PELAPSE (M) VMEM (F) LFILE (S) LSOCKET	TELAPSE DMEM DFILE DSOCKET	CPU PMEM (MB)	(sec) (pg) (MB)
11151936	System	root	system	524386	acctct1	11151936	E	C: 0.3824 M: 25 F: 0.00 S: 0.00	0.3824 27 0.00 0.00	0.2394 0	
11151936	System	root	system	524388	nfssync_kpro	11151936	E	C: 0.0008 M: 0 F: 0.00 S: 0.00	0.0008 0 0.00 0.00	0.0003 0	

### Aggregated Process Accounting report

The following is an example of the Aggregated Process Accounting report. The command line options -P, -U, -G, and -C of this command facilitate preparing aggregated process accounting reports. These arguments correspond respectively to project ID, user ID, group ID, and command name. The administrator can control the order in which the data is presented by changing the order of the command arguments. The following is an example of an Aggregated Process Accounting report based on all user IDs:

```
# /usr/bin/acctrpt -U ALL -f /var/aacct/acctdata
```

PROJID	UID	GID	CMD	CNT	(C) PELAPSE (M) VMEM (F) LFILE (S) LSOCKET	TELAPSE DMEM DFILE RSOCKET	CPU PMEM (MB)	(sec) (pg) (MB)
-	root	-	-	590	C: 54880.8 M: 6 F: 32.3 S: 0.0	281458.2 14100294 1.4 19.8	60.3 14098216	
-	104	-	-	28	C: 3.3 M: 186 F: 8.3 S: 0.0	3.3 299 0.0 0.0	1.0 708	

---

## Example of the Logical Partition Accounting report

You can generate an Logical Partition (LPAR) Accounting report.

## LPAR Accounting Report

The following is an example of the LPAR Accounting report. You use the command line option `-L` of this command to generate accounting reports to specific LPAR resources. The following is an example of file system statistics:

```
# /usr/bin/acctrpt -L FILESYS -f /var/aacct/acctdata
```

File Systems Accounting Report

```
-----  
CNT  DEVNAME          MOUNTPT  FSTYPE RDWR  OPEN  CREATE  LOCKS  XFERS(MBs)  
7    specfs          specfs   16    48853  1066   0        0      1.2  
7    pipefs          pipefs   16    3827   0       0        0      3.3  
7    /export/sp1n1fs2u /farm/cwu  18    4917  4510   750      0      28.0
```

---

## Example of the Transaction Accounting report

You can generate a Transaction Accounting report.

### Transaction Accounting report

The following is an example of the Transaction Accounting report. To generate accounting reports for ARM transactions, use the command line option `-T` of the **acctrpt** command.

```
# /usr/bin/acctrpt -T -f /var/aacct/acctdata
```

```
PROJID  CNT      (A) CLASS      NAME          USER          GROUP          TRANSACTION      (sec)  
-----  ---      (T) RESPONSE   QUEUED        CPU           -----  
System  144      A: -           WebSphere     -             server1        URI  
        T: 0.00      0.00          0.00          0.00  
System  32      A: -           IBM Webserving -             IBM_SERV       Apache/1.3.28(Unix)  
        T: 0.00      0.00          0.00          67.01
```





---

## Accounting records

There are multiple types of records produced by Advanced Accounting.

Accounting records produced by Advanced Accounting are defined in the **sys/aacct.h** file. The following table describes these records.

Table 17. Accounting records

Accounting record	Description
Pad record (type 0)	This record does not provide any meaningful accounting data. Report and analysis tools should skip this record. It is generated for alignment purposes only.
Process record (type 1)	<p>This record is written when a process exits, when a process is reclassified (<b>setUser ID()</b>, <b>chproj()</b>, <b>exec()</b>) and when the system is reclassified. This record is written by the process interval.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"><li>• User ID</li><li>• Group ID</li><li>• Process ID</li><li>• Process flags (exited, core, killed by signal, killed by checkpoint)</li><li>• Base command name</li><li>• WLM class</li><li>• Controlling terminal</li><li>• Process start time (in seconds from the EPOCH)</li><li>• Process elapsed time in micro seconds</li><li>• Combined thread elapsed time in micro seconds</li><li>• Process (combined threads) processor time in micro-seconds</li><li>• Elapsed page seconds of disk pages</li><li>• Elapsed page seconds of real pages</li><li>• Elapsed page seconds of virtual memory</li><li>• Local logical file I/O (JFS, J2) in bytes</li><li>• Other logical file I/O (NFS, DFS™) in bytes</li><li>• Local socket I/O (UNIX domain and loopback) in bytes</li><li>• Remote socket I/O in bytes</li></ul> <p>The process start time and Process ID can be used to correlate interval records for a particular process. The exit flag can be used to distinguish between interval and exit records.</p>

Table 17. Accounting records (continued)

Accounting record	Description
Aggregated process record (type 2)	<p>This record is derived from the process record. A different record is produced for each user by project.</p> <p>This record is produced by the process interval and contains the following information:</p> <ul style="list-style-type: none"> <li>• User ID</li> <li>• Time of first record aggregated (in seconds from the EPOCH)</li> <li>• Number of processes aggregated</li> <li>• Aggregate process elapsed time in micro seconds</li> <li>• Aggregate thread elapsed time in micro seconds</li> <li>• Aggregate process (combined threads) processor time in micro seconds</li> <li>• Aggregate elapsed page seconds of disk pages</li> <li>• Aggregate elapsed page seconds of real pages</li> <li>• Aggregate elapsed page seconds of virtual memory</li> <li>• Aggregate local logical file I/O (JFS, J2) in bytes</li> <li>• Aggregate other logical file I/O (NFS, DFS) in bytes</li> <li>• Aggregate local socket I/O (UNIX domain and loopback) in bytes</li> <li>• Aggregate remote socket I/O in bytes</li> </ul>

Table 17. Accounting records (continued)

Accounting record	Description
Aggregated application record (type 3)	<p>This record is derived from the process record. Records are produced at the user, project, and application level. This record is similar to the aggregated process record, except that the application is named. This record is produced when the process is classified with an application specific rule, which is supported only through the Admin policy.</p> <p>This record is produced by the process interval and contains the following information:</p> <ul style="list-style-type: none"> <li>• User ID</li> <li>• Time of first record aggregated (in seconds from the EPOCH)</li> <li>• Inode of the command that generated the project classification</li> <li>• Device number of the command that generated the project classification</li> <li>• Number of applications aggregated</li> <li>• Aggregate process elapsed time in micro seconds</li> <li>• Aggregate thread elapsed time in micro seconds</li> <li>• Aggregate process (combined threads) processor time in micro seconds</li> <li>• Aggregate elapsed page seconds of disk pages</li> <li>• Aggregate elapsed page seconds of real pages</li> <li>• Aggregate elapsed page seconds of virtual memory</li> <li>• Aggregate local logical file I/O (JFS, J2) in bytes</li> <li>• Aggregate other logical file I/O (NFS, DFS) in bytes</li> <li>• Aggregate local socket I/O (UNIX-domain and loopback) in bytes</li> <li>• Aggregate remote socket I/O in bytes</li> </ul>

Table 17. Accounting records (continued)

Accounting record	Description
Processor and memory use record (type 4)	<p>This record provides information about the use of processors and memory at the system level. It is generated before and after dynamic logical partitioning operations and when the size of the large page pool changes.</p> <p>This record is also generated by the system interval and contains the following information:</p> <ul style="list-style-type: none"> <li>• Reason the record was generated</li> <li>• Number of online logical processors</li> <li>• Entitled physical processor capacity of the partition</li> <li>• Total idle time, in milliseconds</li> <li>• Total I/O wait time, in milliseconds</li> <li>• Total kernel process time, in milliseconds</li> <li>• Total user process time, in milliseconds</li> <li>• Total interrupt time, in milliseconds</li> <li>• Size of physical memory, in megabytes</li> <li>• Size of the large page pool, in megabytes</li> <li>• Large pages in use, in megabytes</li> <li>• Number of page ins from paging space</li> <li>• Number of page outs to paging space</li> <li>• Number of start I/Os</li> <li>• Number of page steals</li> <li>• Elapsed time since start of interval, in milliseconds</li> </ul>
Policy record (type 5)	<p>This record is written when a policy file is loaded or unloaded. It is provided for informational purposes only.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Type of policy: Admin, User, or Group</li> <li>• Load or unload</li> </ul>
File system activity record (type 6)	<p>This record describes the use of file systems at the system level. A separate record is generated for each mounted file system.</p> <p>This record is produced by the system interval and has the following information:</p> <ul style="list-style-type: none"> <li>• Total bytes transferred through read and write</li> <li>• Total number of read and write requests</li> <li>• Total number opens</li> <li>• Total number of creates</li> <li>• Total number of locks</li> <li>• File system type</li> <li>• Device name</li> <li>• Mount point</li> </ul>

Table 17. Accounting records (continued)

Accounting record	Description
Network interface I/O record (type 7)	<p>This record provides information about the use of network interfaces at the system level.</p> <p>This record is produced by the system interval and contains the following information:</p> <ul style="list-style-type: none"> <li>• Logical name of the network interface</li> <li>• Number of I/Os</li> <li>• Total bytes transferred</li> </ul>
Disk I/O record (type 8)	<p>This record provides information about the use of disks at the system level. A separate record is written for each logical disk device.</p> <p>This record is produced by the system interval and contains the following information:</p> <ul style="list-style-type: none"> <li>• Logical name of the disk</li> <li>• Total disk transfers</li> <li>• Total reads</li> <li>• Total writes</li> <li>• Block size of the disk transfer</li> </ul>
Lost data record (type 9)	<p>This record provides information about accounting records that were deleted because Advanced Accounting did not have the ability to record them. This occurs when all of the accounting data files are full. When the ability to write new accounting records is restored, Advanced Accounting produces the lost data record describing the outage.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Number of lost records</li> <li>• Number of microseconds of lost processor time associated with process records</li> <li>• The time that data loss began (in microseconds from EPOCH)</li> <li>• Number of microseconds of lost processor time associated with third party kernel extension records</li> </ul>
Server VIO record (type 10)	<p>This record is produced in hosting partitions. A separate record is produced for each logical device that is shared with a client partition. The system interval can be used to periodically produce this record.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Client partition number</li> <li>• Server unit ID</li> <li>• Device logical unit ID (LUN)</li> <li>• Number of bytes in</li> <li>• Number of bytes out</li> </ul>

Table 17. Accounting records (continued)

Accounting record	Description
Client VIO record (type 11)	<p>This record is produced in client partitions. It describes the use of virtual devices in client partitions. A separate record is recorded for each instance of a virtual device. The system interval may be used to periodically produce this record.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Server partition number</li> <li>• Server unit ID</li> <li>• Device logical unit ID</li> <li>• Number of bytes in</li> <li>• Number of bytes out</li> </ul>
Third-party kernel extension common aggregation record (type 12)	<p>This record provides accounting information for the named accounting record. It is derived from aggregated accounting records that are produced by third-party kernel extensions. This record is written to Advanced Accounting by the system interval.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Command name of the kernel extension (from u-block)</li> <li>• Third-party kernel extension transaction ID, in the range of 129 to 256</li> <li>• Number of accounting records that have been aggregated</li> <li>• Resource use, or accumulated processor time, for this class of transactions</li> <li>• Time of first record aggregated (in seconds from the EPOCH)</li> </ul>
ARM application environment record (type 13)	<p>This record describes an application environment instance. It is created from data that is passed to the operating system through the <b>arm_register_application()</b> system call and the <b>arm_start_application()</b> system call. This record is variable in length. All offsets are calculated relative to the start of the record.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Character set in which the data in this record is recorded</li> <li>• Application environment identifier</li> <li>• Offset to application name</li> <li>• Offset to application group</li> <li>• Offset to application identity properties</li> <li>• Offset to application context properties</li> </ul> <p>The operating system attempts to record the content of the application environment in each accounting data file, so that each accounting data file can be post-processed as a stand-alone item. This is designed to eliminate the dependency between accounting data files.</p>

Table 17. Accounting records (continued)

Accounting record	Description
ARM transaction environment record (type 14)	<p>This record describes a transaction environment instance. It is created from data that is passed to the operating system through the <b>arm_register_transaction()</b> system call. This record is variable in length. All offsets are calculated relative to the start of the record.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Character set in which the data in this record is recorded</li> <li>• Transaction environment identifier</li> <li>• Offset to transaction name</li> <li>• Offset to application identity properties</li> <li>• Offset to application context properties (names only)</li> </ul> <p>The operating system attempts to record the content of the transaction environment in each accounting data file (not guaranteed), so that each accounting data file can be post-processed as a stand-alone item. This is designed to eliminate the dependency between accounting data files.</p>
ARM transaction instance record (type 15)	<p>This record describes an ARM transaction instance. It is created from data that is passed to the operating system through the <b>arm_start_transaction()</b> and the <b>arm_stop_transaction()</b> system calls. It is variable in length. All offsets are calculated relative to the start of the record.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Completion status of the transaction</li> <li>• Application environment identifier</li> <li>• Transaction environment identifier</li> <li>• Offset to user identifier (not User ID)</li> <li>• Offset to user name (not uname)</li> <li>• Offset to accounting code</li> <li>• Response time, in milliseconds</li> <li>• Queued time, in milliseconds</li> <li>• Resource use</li> </ul> <p>The application and transaction environment identifiers are defined respectively in the application and transaction environment records. These records must be used to associate application names, application groups, transaction names, and properties with the transaction instance.</p>

Table 17. Accounting records (continued)

Accounting record	Description
ARM aggregated transaction instance record (type 16)	<p>This record is produced instead of the ARM transaction instance record (type 15), when aggregation is enabled for ARM transactions.</p> <p>This record contains the following information:</p> <ul style="list-style-type: none"> <li>• Completion status of the transaction</li> <li>• Time of first record aggregated (in seconds from EPOCH)</li> <li>• Application environment identifier</li> <li>• Transaction environment identifier</li> <li>• Offset to user identifier (not user ID)</li> <li>• Offset to user name (not unname)</li> <li>• Offset to accounting code</li> <li>• Aggregate response time, in milliseconds</li> <li>• Aggregate queued time, in milliseconds</li> <li>• Aggregate resource use</li> </ul>
Project definition record (type 17)	<p>This record provides a list of project definitions. It is written when the project definition file is loaded. Multiple records may be needed to record all project definitions. This record is used to provide the full set of project information in each data file, so that data files may be treated as stand-alone entities. This may not be required by the billing application, depending on the nature of the billing application. This feature may be disabled by disabling the project definition accounting record.</p> <p>This record is variable in length and contains the following information:</p> <ul style="list-style-type: none"> <li>• Number of projects</li> <li>• Number of bytes in the project definition area</li> <li>• Project definition area</li> </ul>



---

# Index

## A

Advanced Accounting System 1  
usage-based information 1

## C

commands  
acctl 4  
projctl 10, 16  
configuring 5

## D

Data aggregation 35  
enable  
project-level 36  
system-level 35  
Data files 3  
commands 4  
how-to create 3

## E

e-mail notification 5  
E-mail notification  
configuring 4

## F

files  
data 3

## I

Interval accounting 23  
commands 23  
process interval 23  
system interval 23

## P

policies 13  
Admin policy  
aliases 15  
alternate policies 15  
examples 14  
User and Group 16  
Policies  
Admin policy 13  
assignment rules 13  
Policy  
commands 16  
products 49  
projects 7  
accounting records 7  
classification 9

projects (*continued*)  
semantics 7

Projects  
classification  
environment variables 9  
manual 8  
commands 10

## R

records  
aggregated 35  
description 41

## T

terms 50  
Transactional accounting 19  
Application Resource Management 19  
Application Programming Interface 19  
ARM 19  
APIs 19



## Vos remarques sur ce document / Technical publication remark form

**Titre / Title :** Bull Understanding the Advanced Accounting Subsystem

**N° Référence / Reference N° :** 86 A2 84EM 02

**Daté / Dated :** October 2005

### ERREURS DETECTEES / ERRORS IN PUBLICATION

### AMELIORATIONS SUGGEREES / SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

Vos remarques et suggestions seront examinées attentivement.

Si vous désirez une réponse écrite, veuillez indiquer ci-après votre adresse postale complète.

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.

If you require a written reply, please furnish your complete mailing address below.

NOM / NAME : \_\_\_\_\_ Date : \_\_\_\_\_

SOCIETE / COMPANY : \_\_\_\_\_

ADRESSE / ADDRESS : \_\_\_\_\_

Remettez cet imprimé à un responsable BULL ou envoyez-le directement à :

Please give this technical publication remark form to your BULL representative or mail to:

**BULL CEDOC  
357 AVENUE PATTON  
B.P.20845  
49008 ANGERS CEDEX 01  
FRANCE**

# Technical Publications Ordering Form

## Bon de Commande de Documents Techniques

To order additional publications, please fill up a copy of this form and send it via mail to:

Pour commander des documents techniques, remplissez une copie de ce formulaire et envoyez-la à :

**BULL CEDOC**

**ATTN / Mr. L. CHERUBIN**  
**357 AVENUE PATTON**  
**B.P.20845**  
**49008 ANGERS CEDEX 01**  
**FRANCE**

**Phone / Téléphone :** +33 (0) 2 41 73 63 96  
**FAX / Télécopie :** +33 (0) 2 41 73 60 19  
**E-Mail / Courrier Electronique :** [srv.Cedoc@franp.bull.fr](mailto:srv.Cedoc@franp.bull.fr)

Or visit our web sites at: / Ou visitez nos sites web à:

<http://www.logistics.bull.net/cedoc>

<http://www-frec.bull.com> <http://www.bull.com>

CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
____ _ [__]		____ _ [__]		____ _ [__]	
[__] : <b>no revision number means latest revision</b> / pas de numéro de révision signifie révision la plus récente					

NOM / NAME : \_\_\_\_\_ Date : \_\_\_\_\_

SOCIETE / COMPANY : \_\_\_\_\_

ADRESSE / ADDRESS : \_\_\_\_\_

PHONE / TELEPHONE : \_\_\_\_\_ FAX : \_\_\_\_\_

E-MAIL : \_\_\_\_\_

**For Bull Subsidiaries** / Pour les Filiales Bull :

Identification: \_\_\_\_\_

**For Bull Affiliated Customers** / Pour les Clients Affiliés Bull :

**Customer Code** / Code Client : \_\_\_\_\_

**For Bull Internal Customers** / Pour les Clients Internes Bull :

**Budgetary Section** / Section Budgétaire : \_\_\_\_\_

**For Others** / Pour les Autres :

**Please ask your Bull representative.** / Merci de demander à votre contact Bull.



**BULL CEDOC**  
**357 AVENUE PATTON**  
**B.P.20845**  
**49008 ANGERS CEDEX 01**  
**FRANCE**

ORDER REFERENCE  
86 A2 84EM 02