



Using the ryaacct Toolkit

July, 3rd, 2008

Table of contents

Table of contents	2
1. Introduction	3
2. Planning	3
3. Installation	4
4. Configuration and usage	4
5. Command Reference	5
6. File Format	6
7. Known Problems	7
8. Changes	7

1. Introduction

This document describes the usage of the ryaacct collection of scripts to be used with the AIX Advanced Accounting Subsystem as introduced in AIX 5L Version 5.3. The document applies to ryaacct Version 0.6 and AIX 5.3 or AIX 6.1

1.1 Purpose of ryaacct

This collection of tools is meant to provide clients and IBM Service to implement a fast, easy-to-understood solution to measure utilization of virtualized environments. Its main focus is on simplicity so it provides only a small subset of all the functions available with AIX Advanced Accounting. For users seeking full functionality IBM Tivoli Usage and Accounting Manager or native AACCT is the required solution

1.2 Packaging

The ryaacct collection divides into two parts, a client and a server fileset. The client fileset (rtools.ryaacct.rte) is meant to be installed on all nodes that need to be accounted. The optional server fileset (rtools.ryaacct-server.rte) is meant to be installed on a management server (i.e. a CSM Management Server, a NIM Server or an IBM Director Server running AIX) and provides additional capabilities for simplified collection of distributed data

1.3 Limitations

This package only provides accounting data for the following metrics, although enhancements are planned:

- CPU seconds used by a micropartition
- Utilization data of the central electronic complex (CEC)

In addition to that, it provides informational data about

- Memory assigned to a micropartition
- Entitled capacity of a micropartiton

Currently full functionality is only obtained for partitions connected to a Hardware Management Console. IVM managed systems provide all data except CEC utilisation

2. Planning

Effective use of the tools requires some planning. On the machine some space in the /usr filesystem is needed. Furthermore /var is expanded by 16 MB to hold the accounting files for machines having the client fileset installed.

2.1 Planning for the client fileset

The client fileset can be installed on every machine running the AIX Version 5.3 or higher. Due to enhancements in the AIX Advanced Accounting system it is recommended to use at least AIX 5.3 TL05 or higher. The rtools.ryaacct.rte fileset can be installed on every machine where bos.acct is installed – which is the only requirement.

2.2 Planning for the server filset

The server fileset is meant to be installed on a single machine acting as a central management point. It is recommended to use a CSM Management Server because some of the manual setup process needed can be skipped.

The server fileset alone is not of use, it needs at least one client (or the same machine), where the rtools.ryaacct.rte fileset is installed and configured. Furthermore a working ssh is required, where the central focal point can connect to all nodes without requiring a password.

The csm.dsh fileset is required as prerequisite on this server – it is usually installed as part of AIX per default. Furthermore, to obtain utilization data access to a HMC is required without prompting for a password. This can be obtained using CSM and the HMC as a hardware device or mkauthkeys on the HMC.

Using the ryaacct collection of tools

3. Installation

Both packages were delivered in standard AIX install format. They can be installed using any kind of fashion, i.e. NIM Server, /usr/sys/inst.images or CSM software distribution. They do not require acceptance of a license. During installation no services were started and no files were modified. The tools install in the /usr/rtools directory, where the following subdirectories were created:

- /usr/rtools/bin - All scriptes reside here
- /usr/rtools/samples - Examples of files to use reside here
- /usr/rtools/docs - Basic documentation

In case of uninstalling the fileset it is required to unconfigure it first, in case it was configured before. Then follow the usual procedure. Not doing this lead to unneeded waste of space in the /var/aaacct/ directory as well as some cron jobs not longer working.

4. Configuration and usage

4.1 Configuring the client fileset

Once installed the tool can be configured using the /usr/rtools/bin/acctsetup command. This command does the following steps

- Create an entry in the inittab for aaacct subsystem to start at boot time
- Start the aaacct subsystem
- Configure aaacct mail send to root (default) or a user defined with the -m flag
- Create 2 files in /var/aaacct each size 5MB to store native accounting data. Default name is accounting1.dat and accounting2.dat. Can be changed using the -f flag
- Defines track record ids and aggregation to AACCT
- Creates an empty reporting file (default /var/aaacct/report.csv), can be changed using the -o flag
- Creates a cron entry to run acctcron daily after midnight

After configuration the acctcron command runs daily shortly after midnight to extract the collected data from AACCT to the file. As default a 24 hour reporting is activated. In case you want to have an hourly measurement of the CPU usage only you need to manually modify the crontab entry.

4.2 Using the client filest

Usually, the tools act automatically due to the cron job.

The /usr/rtools/bin/acctcron command runs usually from crontab and performs the following operations:

- Switches accounting files, so the data from the last day can be used to be analyzed. (Note that running acctcron manually switches the files too, so you might miss information on that day when the cron job runs. In this case use the /usr/rtools/bin/acctgenreport command
- Collects used CPU seconds, entitled capacity and memory settings of the lpar and write it as “;” separated data in the /var/aaacct/report.csv file. (Or any other filename defined with the -f flag).
- It is possible to write a record for each hour, using the -i flag and to write only a pair of values for time and cpu seconds using the -f flag.
- It is possible to limit the amount of data collected to a certain number of entries using the -c flag (0.6.2.0)

In cases you want to manually see entries in your accounting files, the tool acctgenreport can be used. It extracts values from the original accounting tools but does not alter the state of such a file. Note that running this script against a file that is not in the state “processing” does not provide accurate results, because aaacct buffers some data internally for a while.

The produces file (report.csv) can be postprocessed by any user-written script or can be imported into a Spreadsheet, like Excel.

In case you want to get rid of ryaacct, the acctremove script cleans up the system, especially

- Removing the cron entries

Using the ryaacct collection of tools

- Removing the /etc/inittab entry for aacct to start automatically
- Stopping the aacctl subsystem
- Removing the accounting files

Everything you might have altered should be removed manually. To fully uninstall the package follow the known AIX deinstallation procedures.

4.3 Configuring the server filset

The server filset consists of a collection of tools optionally to be used to gather cross-lpar information and to collect data from different lpars. There are two separate tasks

1. Collecting cross-lpar usage from an HMC
2. Collecting files generated from ryaacct from different nodes and extract some values.

For obtaining the first task a setup script called acctsetup-hmc is provided. It has to be run for every CEC and does the following

- Connects to the HMC to enable collection of utilization data of that CEC
- Create a crontab entry to collect utilization data shortly after midnight for the previous day

For obtaining the second task a list of nodes has to be provided where all hosts are listed – one per line. An example of such a file resides in /usr/rtools/samples/. The hostnames in that file need to be resolved and ssh/scp to the nodes should work without prompting for a password

4.4 Using the server fileset

There are the following utilities included in the fileset that can be used to centralize data.

- The acctcec command is usually run from cron and connects to a HMC to collect utilization data of that machine for a day. It captures online CPU and Memory and appends this value to a “;” separated file with the name of the CEC which resides per default in /var/aacct/
- The acctcollect command is used to copy remote reports (the report.csv) from nodes listed in a hostfile to the local node. Once copied the files reside in /var/aacct/remote_files/ for later processing by user generated post-processing. It could be implemented as a cron job to automatically collect all values from all nodes once per day or once per month – depending on the reporting structures
- The acctreport command is used to extract exactly one month of reporting data from the files collected by acctcollect. This can be used to feed a database , a self written application or can be imported in a spreadsheet.

5. Command Reference

Each command as a -h option that prints out a usage function.

acctcron

This program is usually run from cron and is used to extract accounting data from a period of time to a flat ASCII File.

Options: [-h] [-o MODE] [-f OUTFILE] [-i TIME] [-c ENTRIES]

-o : Specify the operation that acctcron uses. This defines the output format as described later. There are two options possible. CSV specifies a detailed view of CPU data together with informations on entitled capacity and memory settings, FILE specifies only timestamp and consumed CPU seconds and is suited for shorter sample intervals.

-f : Usually acctcron uses a default file /var/aacct/report.csv. Using this flag allows setting of other names and subdirectories

-i: Specify the time interval for gathering accounting data. The default is 0 for once per day. It is also possible to set 1 for once per hour. Please note that acctcron in any case gathers information on a per day base, only the level of detailed extracted later is different.

Using the ryaacct collection of tools

-c: Specify a number of entries to keep in the report file. The default is to unlimited grow the data so the administrator must take care on its own to not exceed filesystem size. With the -c flag introduced in 0.6.2 it is possible to limit this by skipping the oldest entries.

acctsetup

This program is usually run from cron and is used to extract accounting data from a period of time to a flat ASCII File.

Options: [-h] [-f FILENAME] [-m MAIL] [-o OUTFILE] [-c ENTRIES]

-o : Specify the outputfile. Should be the same as the one used by acctcron. Default is /var/aacct/report.csv

-f : Specify the names of the binary accounting files. Usually /var/aacct/accounting1.dat and /var/aacct/accounting2.dat. You only need to specify the beginning of the filename, i.e. accounting. The 1.dat or 2.dat will be appended automatically.

-m: Specify the user that should receive accounting related mail. Default is root.

-c: If given the entry for crcron uses the -c flag of acctcron.

acctgenreport

This program is a viewer to the accounting data and it acts a lot like acctcron beside that it does not switch files and the user must manually supply begin and enddate.

Options: [-h] -S STARTTIME -s STOPTIME -f FILENAME

-S: Starttime of accounting data to be gathered (MMDDHHMMYY)

-s : Stoptime of accounting data to be gathered (MMDDHHMMYY)

-f: Filename of binary accounting file to process. It is a design of acct that a file that is not in the state "processing" does not necessarily show all informations due to internal bufferings.

acctremove

This program is used to unconfigure the client fileset from the machine. It does not deinstall the package.

Options: [-h] [-f FILENAME]

-f : Specify the names of the binary accounting files. Usually /var/aacct/accounting1.dat and /var/aacct/accounting2.dat. You only need to specify the beginning of the filename, i.e. accounting. The 1.dat or 2.dat will be appended automatically.

6. File Format

There are usual two types of output generated by the tools that are either kept locally on the machines or are held in a central repository. The types of output differs on how the acctcron command was invoked.

The format of the acctcron program depends on the mode defined by -o. The field separators are always a semicolon ";". For own scripting it is guaranteed that the existing columns will be there, additional columns might appear afterwards.

For the -o FILE it is:

Date;CPU seconds

The following list gives an example

1203100007;3690.7

1203110007;3692.9

1203120007;3717.1

1203130007;3746.8

1203140007;3747.1

1203150007;3736.7

1203160007;3766.2

Using the ryaacct collection of tools

Please note, that the Date is written in MMDDHHMMYY format, meaning two digits for the month followed by two digits of the day followed by two digits of the hour followed by two digits of the minute followed by two digits of the year.

For the -o CSV option it is

Date;CPU seconds entitled capacity, CPU seconds; Consumed CPU seconds, Memory (MB)

The following gives an example:

```
1128160407;25920;116.5;5.2;2.4;1943.1;4.9;2072.1;1024
1128160407;25920;2941.2;54.7;48.6;47994.9;101.2;51140.6;1024
1129000107;25920;1689.9;6.2;25.9;27094.8;45.5;28862.3;1024
1130000107;25920;5297.7;56.2;112.4;82751.7;227.1;88445.1;1024
```

So after the time (column 1) follows the entitled capacity in CPU seconds for one day (column 2), it is calculated from the entitled capacity of an LPAR (i.e 0.8) x 24 x 60 x 60. In this example the entitlement was 0.3 physical CPU. The next 5 columns (3-7) list the CPU seconds used for hypervisor, usr, sys, idle and wait and the next column (column 8) is a sum of them. This is the column to be accounted and this is the same value as given by the -o FILE command.

Column 9 lists the memory the lpar had at the time acctcron was running. It is informational and not the aggregated value. It however gives a good representation as memory add and remove operations do usually not occur that often.

7. Known Problems

When a server is rebooted more than once a day, ryaacct can not fully account this server. This is due to the fact that on each reboot, acct switches files from Active to Processing. As ryaacct uses exactly two files this leads to acct internally buffering data which can not be written if another reboot occurs. As a workaround manually create more accounting files with the same naming conventions. Those can be handled by ryaacct.

Housekeeping is not performed by ryaacct per default. Most files get data appended. It is in the responsibility of the administrator to periodically clean up data. It is implemented this way as different clients might want to hold data for different periods of time. However the report file now can be limited using the -c flag of acctcron

8. Changes

rtools.ryaacct-server

0.6.0.0: Initial release of the rtools.ryaacct collection

Rtools.ryaacct.client

0.6.0.0: Initial release of the rtools.ryaacct collection

0.6.1.0: Allow to have entries per hour, not only per day; robustness enhancements

0.6.2.0: Allow to limit the report file; robustness enhancements, problems fixed with acctsetup



© IBM Corporation 2005
IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
December 2007
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the AIX 5L, Micro-Partitioning, POWER, POWER5, pSeries, Virtualization Engine are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

Other company, product, and service names may be trademarks or service marks of others.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

Many of the IBM @server p5 and pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html.

The IBM home page on the Internet can be found at: <http://www.ibm.com>.

The IBM @server p5 and pSeries home page on the Internet can be found at: <http://www.ibm.com/servers/eserver/pseries>.

The IBM @server OpenPower home page on the Internet can be found at <http://www.ibm.com/eserver/openpower>.