



APRM[®] Version V6R1M0 Release Notes (and) User Memorandum

Introduction

This is a maintenance release of APRM which includes several improvements to the product. With the exception of V4R5M0 of OS/400 (i5/OS), this release will operate wherever the prior versions of APRM would. The improvements in this release over V5R4M3 and earlier releases are:

- Problems with non-USA language identifiers and/or job CCSIDs (see [APRM Technical Alert # 006](#) for a discussion of this problem and for the temporary APRM patch that is built-into this release). **(note)**
- Automatic resource tuning for all partitions could have been turned off without user action.
- HMC Version 7.3.1 and later can respond before they complete a Dynamic LPAR action.
- This version of APRM has all the "creation data" required to pass IBM's ANZOBJCVN process to assess readiness for V6R1M0 and later releases of i5/OS. **(note)**
- There are now 2 license levels supported.

(note) These improvements bring user considerations with them - see below.

Problems with non-USA language identifiers and/or job CCSIDs.

Versions of APRM prior to V5R4M4 were created with default CCSID values for alphanumeric fields in most of the data files. When APRM is installed on a typical USA-based system, all these fields get assigned a CCSID of 37 (IBM's standard EBCDIC code) including some fields that are used to hold non-character (binary) data. If all users of the CFGAPRM command to supply user names and passwords also operated with this default code page, any binary data that was stored was later retrieved with no translation and was therefore unchanged. If, however, the system standard language identifier was not ENU or code page was not 37 or a user with job values other than these supplied data that was stored in these files, the values retrieved by APRM's batch code was not the same as it appeared to the person entering the data.

This caused several problems, most notably that the encrypted password for HMC access was not acceptable to the HMC and, therefore, APRM would not correctly start on some systems. This deficiency is cured in the V5R4M4 and later APRM releases by forcing the non-numeric data fields in the files to CCSID 37 for real character fields and to 65535 for character fields used for binary data.

There is a side effect of this change however as mentioned above as "**(note)**". If you are operating in an HMC-controlled environment, it will be necessary, once this CCSID change is first done, to re-enter all HMC passwords in each partition before APRM can start in that partition. To do this, Use command CFGAPRM. Select option 10 (HMC Connection Entries) and then, for each entry, use option 2=Edit. Retype the password (and correct the user name, if needed) and press ENTER to perform the update. After that is done, APRM may be (re)started in the partition by using command STRAPRM.

Subsequent APRM release upgrades after this has been done will not require this manual step to be repeated.

Automatic resource tuning for all partitions could be turned off without user action.

There is a very infrequent occurrence of a request for system partition information from the HMC that returns no data at all. In prior releases of APRM, this caused the configuration retrieval code to believe that no partitions existed (clearly impossible since that code runs IN a partition). The result was that all existing partition records in file APRMLPARC were deleted. This actually caused no permanent trouble by itself because the data was quickly repaired by the next configuration retrieval cycle a short time later.

However, the lack of any partition records caused an undesirable side effect. That is, tuning was turned off for all partitions. The logic behind this action is that, if a partition was deleted and then recreated with a significantly different hardware configuration such as processor power and memory and their lower/upper limits, outdated tuning rules could cause significant performance problems. Thus it was for safety that tuning rules that would likely be inappropriate should not be used without manual review. Moreover, no update was done to the timestamp of the tuning rule records nor was there any indication what agent was responsible for the change.

Release V5R4M4 and later releases avoid this undesirable effect by treating a missing response from the HMC simply as a condition to be ignored and no action is taken to update APRM's configuration files.

HMC Version 7.3.1 and later can respond before they complete a Dynamic LPAR action.

Unlike prior versions of the HMC's software, versions 7.3.1 and later accept a Dynamic LPAR (DLPAR) action and return quickly to handle the next user request in some cases. Prior versions "blocked" and would not respond to commands referencing the server's status or configuration until the DLPAR action was complete. This can be viewed as a welcome improvement but, since the HMC does not yet know that the action was completed, the HMC still reports the prior configuration until a subsequent automatic refresh of its data is done. What is more difficult to deal with is that if a series of commands is given to the HMC in rapid succession where the acceptability of the second command depends upon the successful completion of the first, the HMC may reject the second command. This is particularly true for APRM's tuning of processor power and, where applicable, interactive CPW. This is because APRM typically will remove these resources from one or more partitions and then redistribute them to one or more others in various measures. This can also apply to switching ownership of I/O slots in order to switch an I/O Adapter that controls a some hardware between partitions such as a tape device for backup use by more than one partition.

The problem was evidenced on one of our systems by frequent HMC error messages such as "HSCL0616 Cannot use the 5250 CPW setting. Check whether the setting is valid." and/or occasional failures to switch an IOP-less IOA between partitions.

The cure for this problem, as recommended by IBM, was to introduce a short delay between the removal of resources from one or more partitions and their subsequent addition to another partition. The recommended value of this delay is 15 seconds at HMC Releases 7.3.1 and 7.3.2 and will likely be less than this value in some subsequent HMC release. In order for APRM to accommodate this requirement, a new datum was added to the CFGAPRM processing application on the page for option 7 (Other APRM Parameters). This new datum is titled "How long to delay between removal and addition of CPU or I/O resources (HMC delay)". The parameter must be entered as a number of seconds between 0 (zero) and 120. As with other configuration items, an entry made in any active APRM partition will be shared with other active APRM partitions so that your entry only needs to be made once per partitioned system.

An APRM "Technical Alert" on the [APRM Service page](#) will provide recommended values for later HMC code releases as they are determined.

This version of APRM has "creation data".

One of the preparation steps recommended by IBM before installing version V6R1M0 of OS/400 (i5/OS) is to assure that all of your programs, modules, etc. have the ability to be retranslated. This requires them all to have "creation data" - whether or not it is observable. Many software vendors remove as much observability from their distributed products as possible both for their asset protection and for size considerations. APRM has always done this for both of those reasons. Version V5R4M4 and later versions of APRM do, however, retain enough creation data, though it is unobservable, to allow retranslation. Therefore the preparatory step of using IBM's Analyze Object Conversion (ANZOBJCNV) command will not flag any component of APRM as a problem.

There is a side effect of this change however as mentioned above as "**(note)**". The minimum release level of OS/400 (i5/OS) that version V6R1M0 can be installed on is V5R1M0 rather than V4R5M0 as prior releases supported.

There are now 2 license levels supported.

Prior versions of APRM provided a 70-day grace period before any license was required and, thereafter a license was required. This is still true at V6R1M0, but there are now 2 levels of licenses. The first level will support all of the features of APRM except tuning of CPU, memory and interactive performance percentage; this is for users who only wish to use APRM for switching I/O slots between partitions to share devices such as tape and optical drives. The second level provides all the functions of APRM. During the "grace" period, all of the functions of APRM are available.

Tuning of resources, when licensed or allowed by a "grace" period, is still restricted to hardware that does not support multiple virtual processor pools. This will be supported by APRM version V6R1M1 and later releases. These multiple virtual processor pools can exist for the purpose of limiting the maximum CPU power that will ever be assigned to one or more partitions even with uncapped processors configured. This concept was created to allow compliance with licensing by number of processors.