4.10 Maintain Database

4.10.1 Manage Record Data Aging Status

Please see ***Section 4.5.15*** for details on this section.

4.10.2 Update DIS Node Configuration

by James F. Cornwall

The ND_EDIT program is used to update the Node Configuration information in the NWIS database. This information is used to specify node hardware descriptions and configurations. This section briefly discusses how to update Node File information.

Introduction

The Node Configuration tables contain local node information. This information specifies the printers, tape readers, and batch queues that are available to the users. Additional information to describe the characteristics of the local installation, network node identification, and maintenance control parameters may also be set using the ND_EDIT program. The Node Configuration information must be set up by the Database Administrator (DBA) before users can access any programs or data.

Program Operation

The ND_EDIT program allows the DBA to display, list, update, add, and delete data in the Node Configuration file. The program's submenus provide the ability to list, update, add, or delete relevant items in each function. Once the Node Configuration information has been entered, very little updating of the file is necessary. An example of an update requirement would be to create new entries for printers, or to modify a parameter to meet new processing requirements.

The screen shot below is the main program menu. Sub-menus and their options are discussed in following sections:

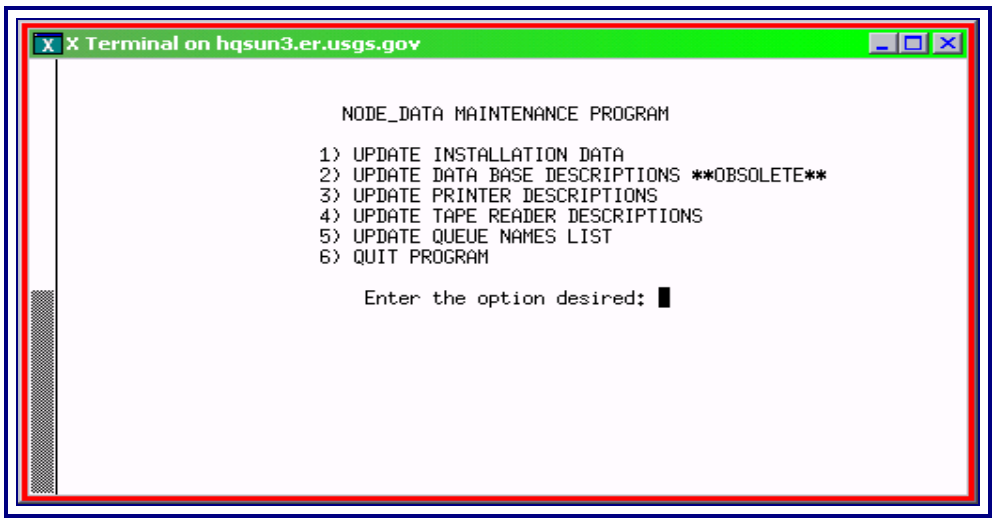

Selection of a menu option will present the user with one of the following screens:

Update Installation Data

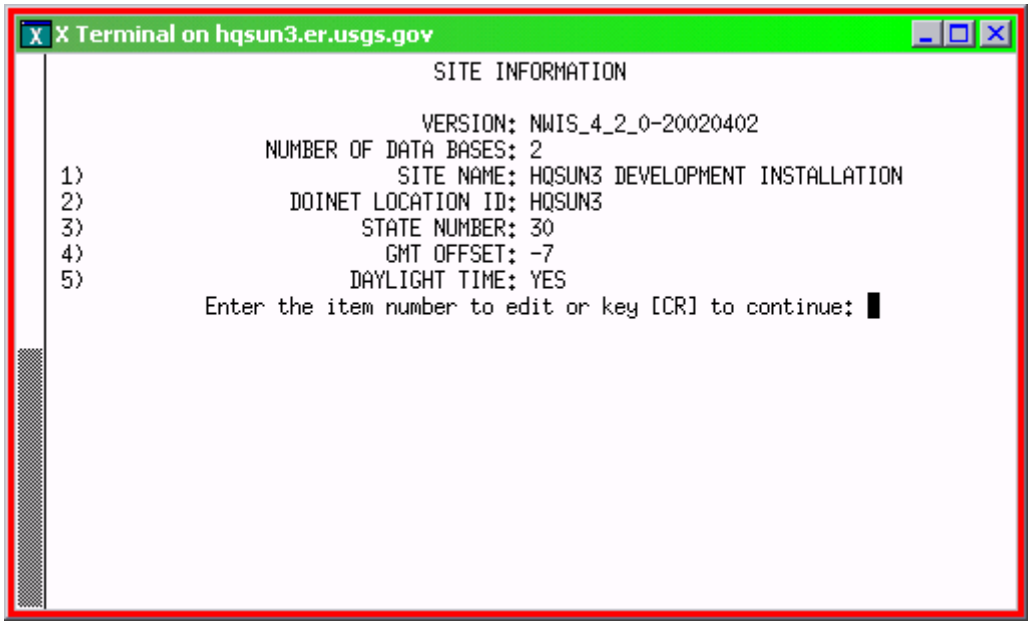

The Installation data consists of general information that describes the ADAPS installation. Any of the fields displayed may be edited, except for the ADAPS Release Number and the number of databases that are established on the system.

Update Database Descriptions **Obsolete**

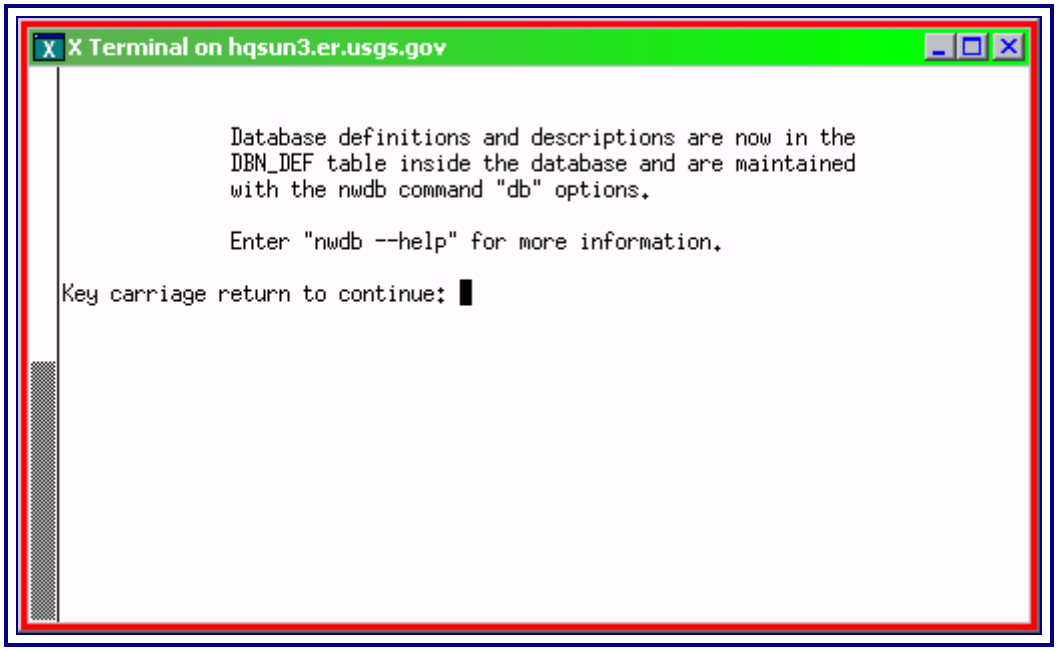

This menu option is obsolete with the release of NWIS 4.2. The user is presented with a message on the new program used for defining database information.

Update Printer Descriptions

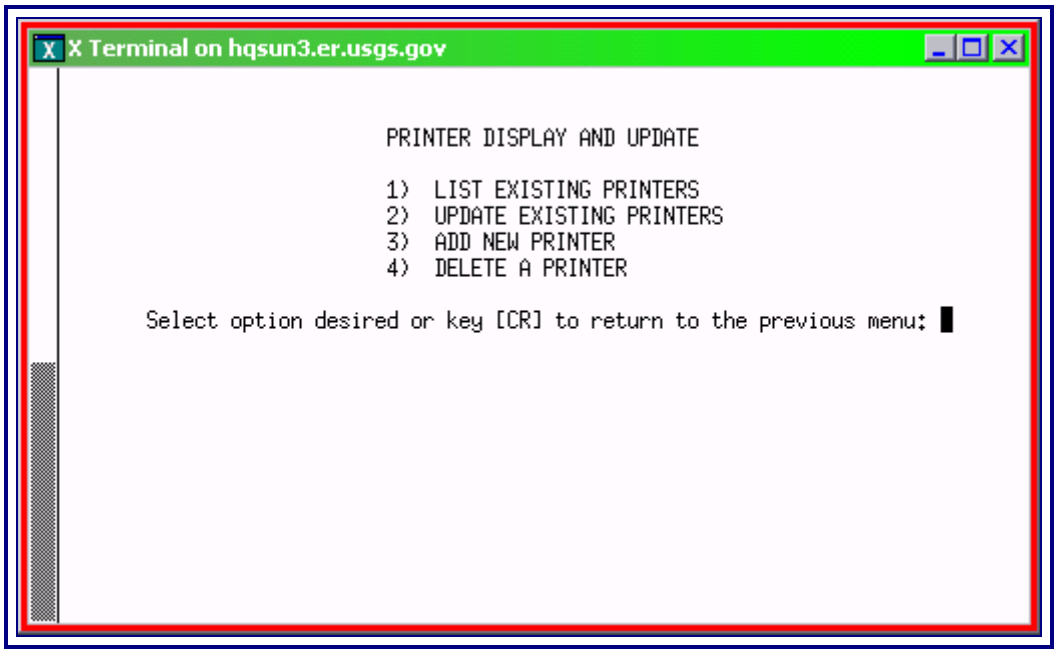

ADAPS requires that all printers have the following information stored: 1) a printer description, 2) a printing width, 3) whether or not there is to be header suppression, and 4) how communication is carried out. This information allows the users to select their output destination. Communications with each printer is specified by a system destination name, form name, or user number, and thus allows for a wide range of printers to be available including system printers, remote printers, and printers tied to user terminals. The available submenu options are as follows:

List Existing Printers

This option provides a list of all printers stored in the Node file.

Update An Existing Printer

This option allows the DBA to change any or all of the printer information in order to meet new needs. First, the user is given a list of available printers defined in the system, and when one is selected, the screen below is shown with the editable parameters:

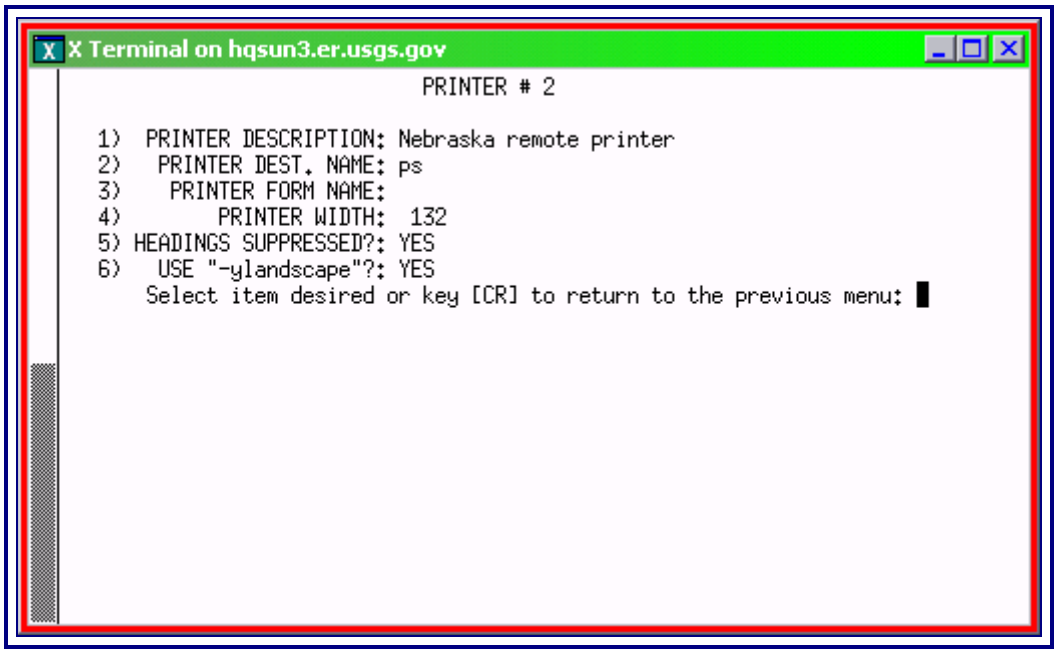

The user may change any of these parameters as needed.

Add New Printer

When adding a new printer to the Node file, the DBA is defining how the printer is described in printer lists, what its characteristics are, and how ADAPS communicates with the printer. The printer name or the form name, if supplied, defines where to send the output. The ND_EDIT program will prompt the user to enter all required information.

Delete An Existing Printer

This option is used by the DBA to delete a printer no longer in use.

Update Tape Reader Descriptions

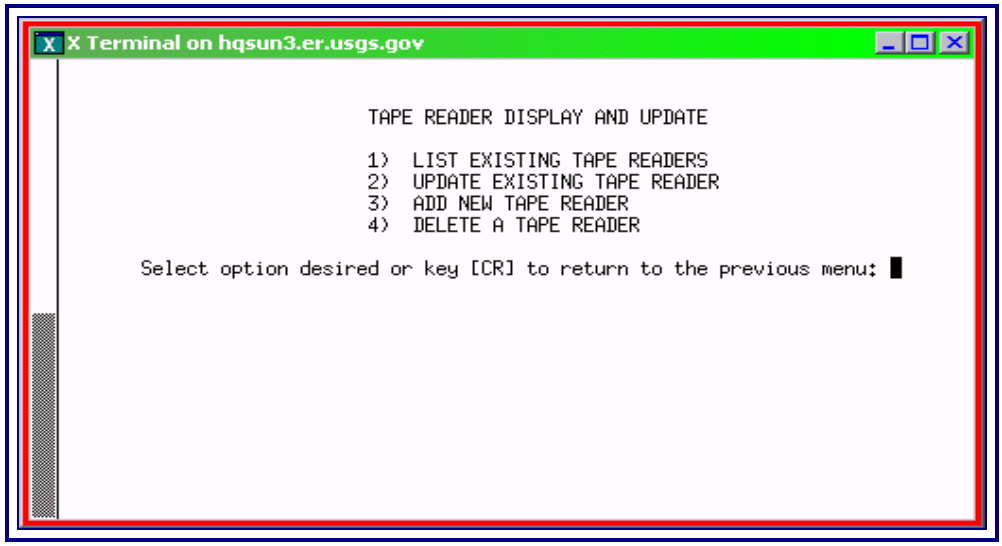

In order to translate ADR 16-channel paper tape data, ADAPS requires that certain configuration information be available. The information needed is a description of the tape reader, the reader type, the path to the assigned port, and communications protocol information (parity, baud rate, stop bits, and XON/XOFF usage). The available sub-menu options are as follows:

List Existing Tape Readers

This option provides a list of all ADR paper tape translators defined for the system. If there are no tape readers defined on the system, the user is given a warning message telling him so.

UPDATE EXISTING TAPE READER

This option allows the DBA to change any or all of the selected tape translator information in order to meet new needs.

ADD NEW TAPE READER

This option allows the DBA to enter the information that defines a new ADR paper tape translator. The program will prompt for all items required to define the tape reader (description, reader type, host name, path to port, etc.)

DELETE A TAPE READER

This option allows the DBA to delete an ADR paper tape translator that is no longer in use.

UPDATE QUEUE NAMES LIST

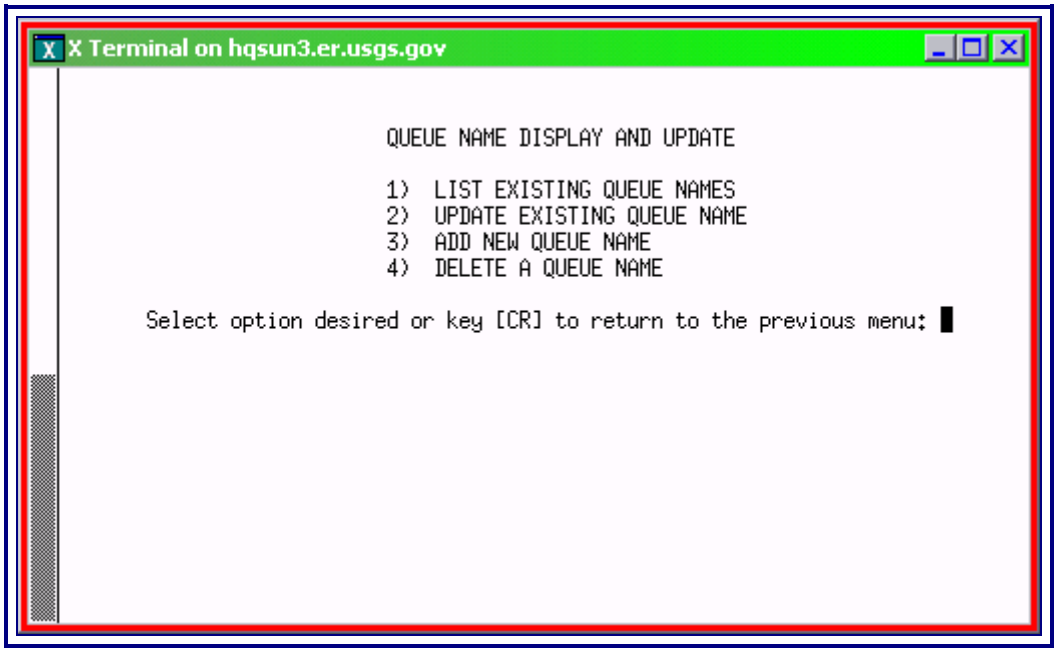

ADAPS uses the system's batch queues to run a variety of programs in batch mode. This is done to avoid overloading the system. The DBA must establish which batch queues are to be available to the users. These batch queue names must be entered into the Node configuration tables. The available submenu options are as follows:

LIST EXISTING QUEUE NAMES

This option provides a list of the available defined batch queue names.

UPDATE EXISTING QUEUE NAME

This option allows the DBA to change any of the existing queue names.

ADD NEW QUEUE NAME

This option allows the DBA to enter a new queue name.

DELETE A QUEUE NAME

This option allows the DBA to delete queue names that are no longer in use.

4.10.3 Review/Delete User File Entries

The US_EDIT program is used to add, delete, and update user information for ADAPS. The program allows the system or Database Administrator to review the settings for a user (useful when attempting to resolve problems), and to delete non-users of ADAPS from the database. It also allows the administrator to add new users, although this usually happens automatically the first time a new user runs the ADAPS software. The user's group files may also be optionally deleted by use of this program.

Program Operation

The US_EDIT program displays a list of all the users defined in the ADAPS database. The user is queried for a user name to review, or <CR> to quit. If the user name is in the list of existing entries, the program will display the user's settings and query what action to take (No action, Delete User, or Delete User and his Groups). If the user's name is not in the list, the program will query if it should be added. Answering "yes" will add the user to the ADAPS database with default settings. The list of user names is then repeated and the user queried for the desired action again.

4.10.4 Archive Unit-Values to Disk for Tape Storage

The UV_ARCHIVE program is used to dump unit-values (UV) data from the NWIS database to disk files suitable for archival to tape or CD-ROM for long-term offline storage. The data may optionally be deleted from the database after archival in order to reclaim and conserve disk space used by the database.

Introduction

Due to the large number of unit values which are accumulated and stored in the NWIS database, it may be necessary from time to time to dump portions of the data out to disk files so they can be archived to offline storage media (tapes, CD-ROM, etc.) to conserve disk space. The UV_ARCHIVE program allows an administrator to do this when required.

The UV_ARCHIVE program operates on a complete water year, writing the unit-values data for the specified year to a series of disk files, one per station. The program allows the user to delete the data from the database to reclaim disk space, or the data may be left intact and just copied to the disk files. The disk files can then be backed up to tape or CD-ROM and deleted from the disk.

Depending on the number and size of the UV sub-files, the program can either be run for all unit-values sub-files at once, or it can be with only selected unit-values sub-files. This mode allows the user to make repeated sets of archival files, which can be copied to offline storage media and deleted from disk after each run if there is insufficient storage space to dump all the UV sub-files at once.

Users (including sentry) can continue to use ADAPS during the archival process as long as they do not attempt to process data in the water year being archived. There is, however, nothing in the program to enforce this restriction. The user needs to know what is taking place in the database and decide accordingly whether to allow access while the UV_ARCHIVE is running.

Archival Procedures

When archiving unit values, follow these steps:

- Select (or create) a directory to hold the archive files. Ensure that this directory is on a partition with a large amount of free space.
- Exclude all other users from NWIS using the "**locknwis**" command.
- Stop SENTRY for all database numbers.
- Manually run a checkpoint of the NWIS database.
- Warn users that in the event the UV_ARCHIVE process fails, the database may have to be rolled forward to the checkpoint just taken.
- Start SENTRY for all database numbers.
- Allow users back in with the "**unlocknwis**" command.
- Run the UV_ARCHIVE program, as described in the next section.
- After the UV_ARCHIVE processing is finished, run the Ingres **schema_editor** utility to compress the unit-values tables and recover the disk space freed by the archival.

Archive Program Operation

The UV_ARCHIVE program is an interactive preprocessor program for the actual processing, which is done in batch mode. The program receives program options from the user, builds a control file, and submits a batch job to do the actual work. The program uses the standard ADAPS startup routines, which allow the user to select a database, water year, and work directory for program operation and control files.

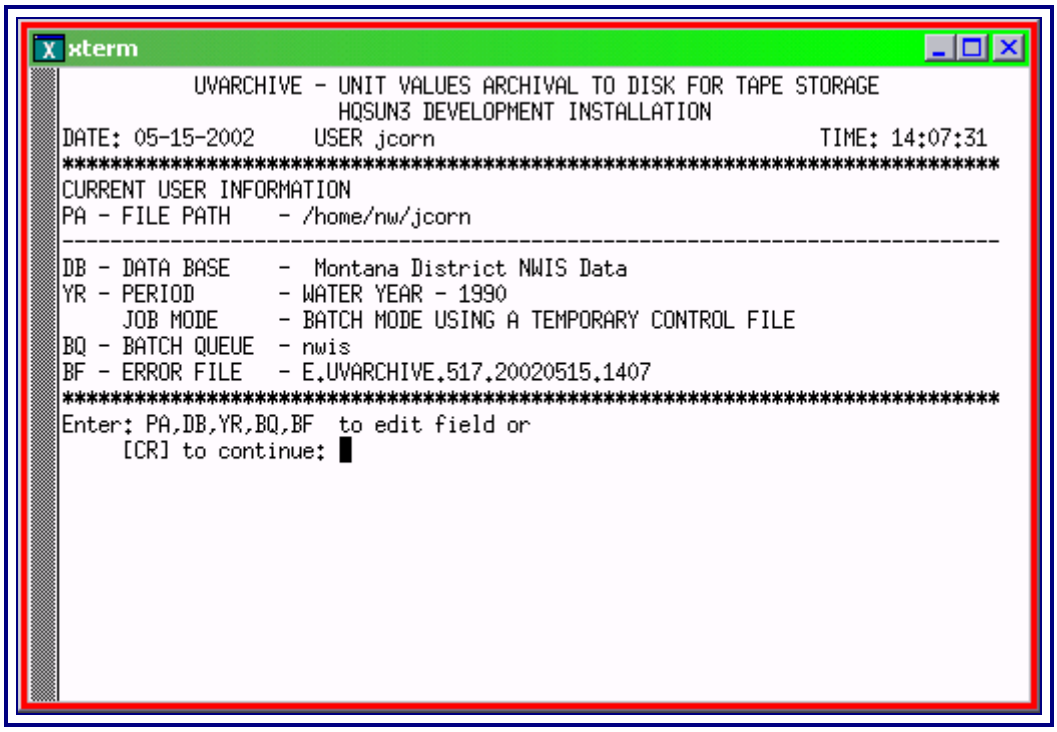

The user is then presented with a warning screen and a list of items that should have been performed.

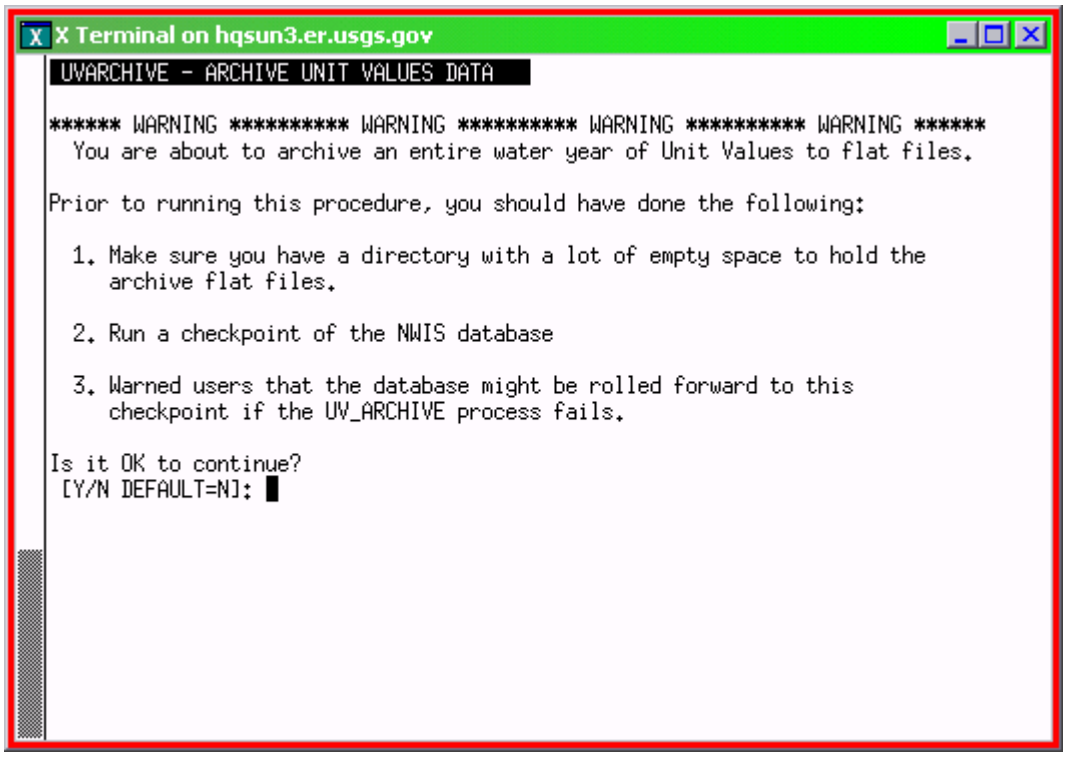

Finally, the program will query the user for the options desired for the run.

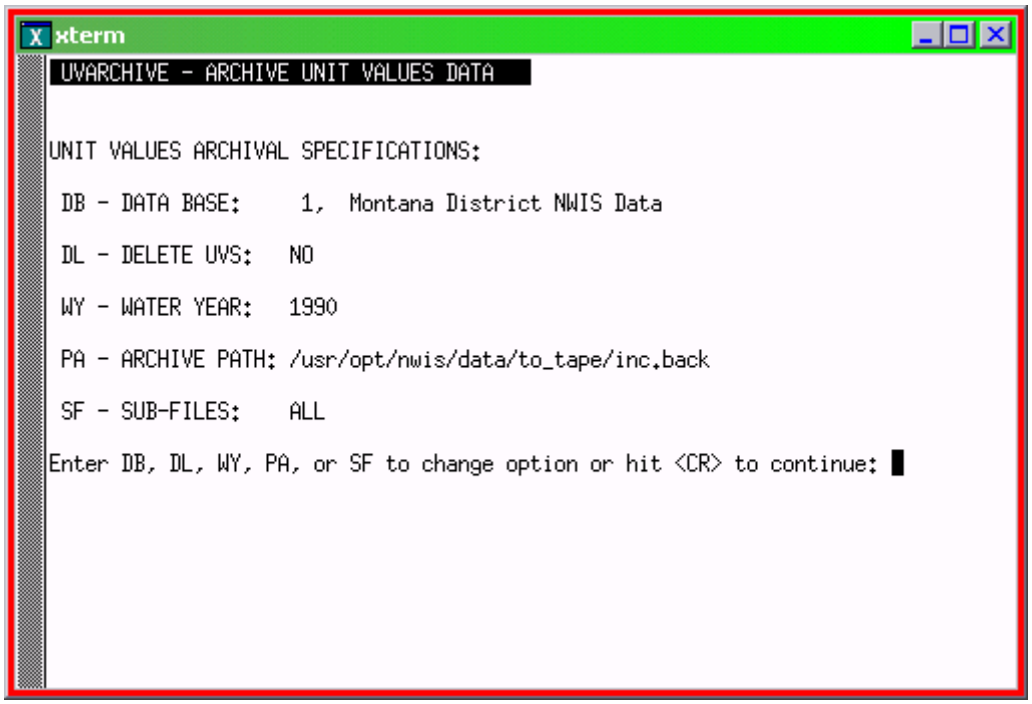

At this point, the user may choose a new database or water year, select specific UV sub-files to archive, or specify the path to another directory where the archival disk files will be written. All the archive files will be placed in the specified directory. The "DL" option allows the user to specify whether the UV data will be deleted from the database after the archival disk files are written. The default choice is to retain the data rather than deleting it. If the data is to be deleted, the user must run an Ingres utility **schema_editor** in order to reclaim the space in the database. Entering a <CR> at this screen will submit the batch job to do the archiving, and the user is returned to his previous menu screens.

The batch postprocessor will read every record in the selected UV sub-files and will write the records to a series of disk files in the target directory. If the user opted to delete the UV data, the unit values are then deleted from the database after the file has been written. Each station which is being archived, will have the UV data written to a separate file. File names are generated in the format "**UDBnn.agncy.station.yyyy**" where:

```
nn      = the database number,
agncy   = the agency code (up to 5 characters),
station = the 8-15 character station number, and
yyyy    = the water year.
```

All file names are uppercase. This may be a cause for errors if the files are transferred via FTP or another mechanism onto systems and the file names change case. The ADAPS program used to restore the data into the NWIS database, UV_RESTORE**,** expects the file names to be in uppercase.

The files are written using the WRD Standard Format as described in ***Section 4.4.8*** of this manual. Although the UV restoration program, UV_RESTORE, would normally be used to retrieve data

from archival files, the use of this standard record format allows data from these files to be recovered with the STD_STOR utility program as well.

4.10.5 Restore Unit-Values from Archive Files

The UV_RESTORE program is used to restore unit-values (UV) data from disk files created by the unit values archival program UV_ARCHIVE. The data are read from these disk files and stored into the NWIS database.

Introduction

The ADAPS unit-values archive files are produced using the UV_ARCHIVE program, which retrieves all unit values for a specified water year from the NWIS database and stores them into disk files (one station per file) using a WRD standard archive format. The UV_ARCHIVE program then (optionally) deletes those unit values from the database. When these unit values are needed in the database for some purpose, the UV_RESTORE program allows the user to read the data from these disk files and store them back into the NWIS database. The UV_RESTORE program allows the user to either restore the data from all archive files in a specified directory, or to specify certain files (sites) to be restored from the archival files. The WRD standard unit-values archive format is described in ***Section 6.3*** of this manual.

Note: Since file names are case-sensitive on UNIX systems, the user must ensure that the file names specified for restoring of unit values data exactly match the file names specified when the archives were created. The UV_ARCHIVE program creates all file names in uppercase; however files may have been renamed or changed case during FTP or other transfer operations. If a restore operation fails, check to see that file names are in uppercase as expected within the UV_STORE program. The uppercase can be done with the following Korn-shell script:

```
#!/bin/ksh
for afil in udb01.*
do
mv $afil UDB01.${afil##udb01.}
done
```

Program Operation

The UV_RESTORE program uses the standard ADAPS startup routine to specify the database, mode of operation (Batch/Interactive) and the user's ADAPS work directory. A menu of default options is then displayed:

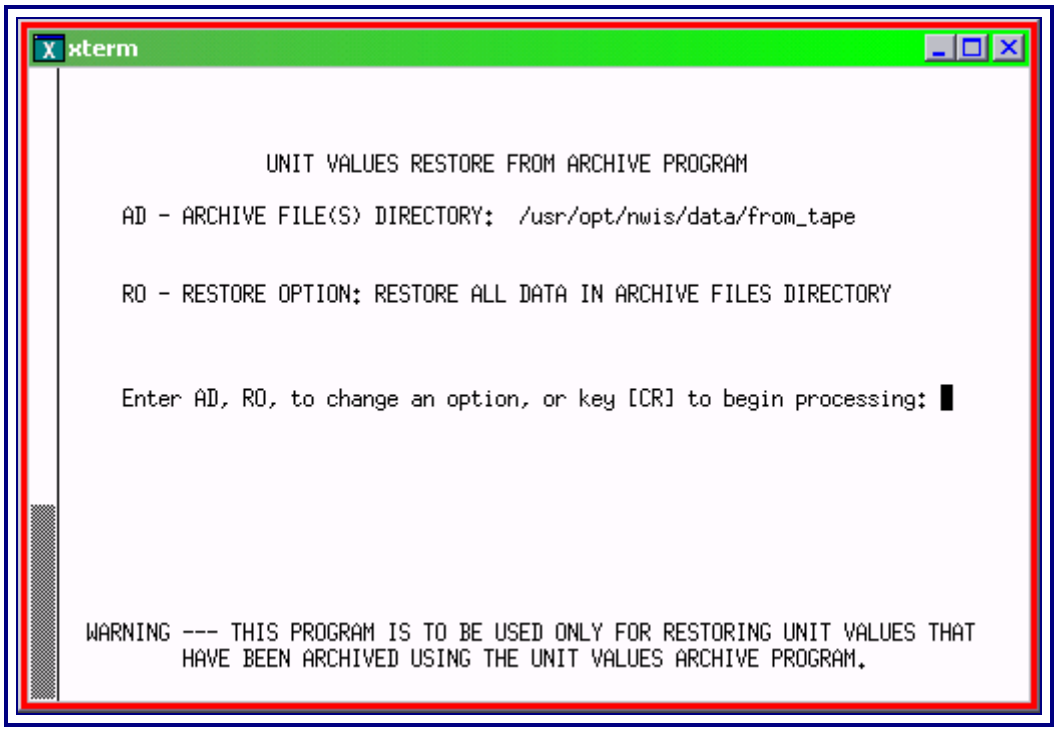

The first option, "AD," allows the user to select the directory where the archive files reside.

The second option, "RO," allows the user to change the restore option. The default specification is to restore all standard archive files in the specified directory. To restore only certain files in the specified directory, select the RO option. This menu option is not displayed if the user selected Batch mode operation. In Batch mode, all files in the specified directory will be processed. If satisfied that these two options are correct, enter a <CR>.

If the program is in Interactive mode, and the option to specify selected files (sites) is selected, the program:

- searches the specified directory for all archive files, then
- displays each site found, queries if the user wishes to reject the site or load the data into the database, and if any existing data is found queries the user if it is OK to overwrite the data with the archived values, then
- Loads the UV data from the file into the database (or skips the load if not OK to overwrite existing data), and displays the next site found in the specified directory.

If the program is in Batch mode, then all files in the specified directory will be processed. Before submitting the batch job, the program will query the user if any existing data should be overwritten. If the user answers "Yes," then the archived data will be stored into the database over the existing data for all input files. If the user answered "No," the occurrence will be logged and the next input file processed.

If the option to load all archive files in the directory is selected, the user is not queried for other options. The program then processes the unit-values archive files as specified. The progress of the

program is written to a log file, along with any generated error messages. For each file processed, site and parameter that has data processed, a summary including the site name, the parameter name, the type of data, and number of data values is printed.

4.10.6 Split a UV Sub-file in two

Prior to the release of NWIS version 4.2, the ADAPS Unit-Values data (UVs) were stored in a collection of disk files which consisted of a UV Index file, an active file, and several sub-files. Though these disk files were replaced by attributed database tables in the NWIS 4.2 release, the term "sub-file" will still be used within ADAPS. When too many values are stored in a sub-file, the database cannot maintain acceptable data access speed, and efficiency and response times will suffer. The UV_SPLIT program allows the Site Administrator to select a UV sub-file and divide the data into two sub-files, reducing the data size roughly by half and increasing efficiency. This is done by copying every other site from an existing sub-file into a new, empty sub-file, and deleting the copied sites from the old sub-file.

Preliminary Steps required before running UV_SPLIT

To run the UV_SPLIT program, the user (a Site Administrator) must have a database access level of "SYST." Before running UV_SPLIT**,** the Administrator must complete these preliminary steps:

(a) Log in as user "ingres" and determine Ingres data areas and UV sub-file sizes.

- Use the program ACCESSDB to determine the Ingres data areas in use (see accessdb documentation).
- Use the command line function ***nwis nw_list_ingres_file_sizes*** (for detailed instructions see the NWIS Database Administrators Guide) to determine the sizes of the Unit Values sub-files. To sort the resulting report by the database table name, use:

  nwis nw_list_ingres_file_sizes > output.file.name

  To sort the report by file size instead, use the "-s" option:

  nwis nw_list_ingres_file_sizes -s > output.file.name

(b) Use the program NWDB to create the new UV sub-file tables in the database, and to grant all necessary rights and permissions for the new sub-file tables (***see nwdb documentation***.). The specific commands are listed below (***db_no*** is the database number {01, 02, etc.}, ***uv_no*** is the new UV sub-file number {01, 12, etc.}, and ***data_area*** is the Ingres data area where the database is located).

- **nwdb -v create tables AD *db_no* UV *uv_no* --data *data_area***
- **nwdb -v modify tables AD *db_no* UV *uv_no***
- **nwdb -v create procs AD *db_no* UV *uv_no***
- **nwdb -v create rias AD *db_no* UV *uv_no***
- **nwdb -v grant all AD *db_no* UV *uv_no***

(c) Exclude ALL other users from ADAPS, including SENTRY.
(d) Take a checkpoint of the database for possible recovery if there are problems with the UV split.
(e) Run **uv_split**.

Program Operation

The UV_SPLIT program first displays the standard ADAPS startup menu, allowing the user to select the database. Next, a screen is displayed that reminds the user of the preceding steps that must already have been done. If the preceding steps have not been done, the user may now halt processing.

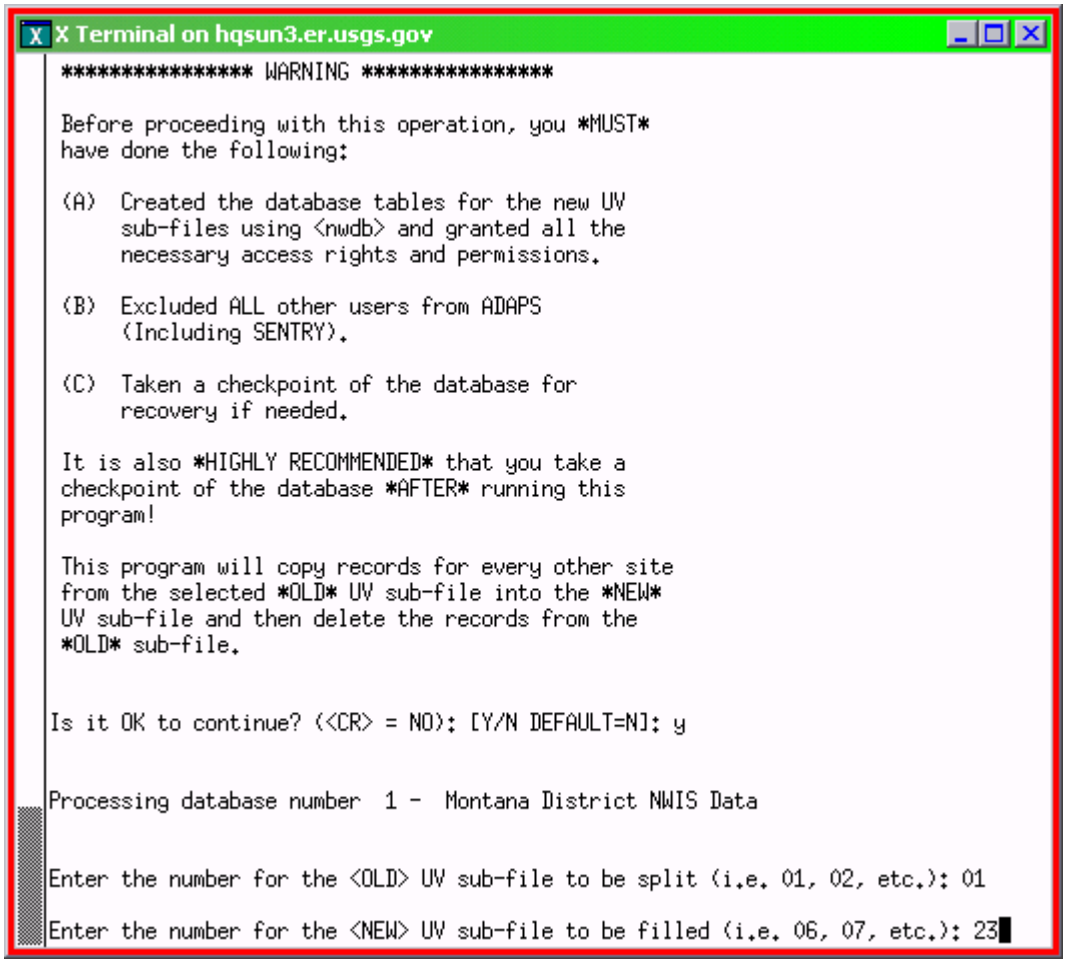

The user is then prompted for the numeric identifiers for the "FROM" and "TO" sub-files (i.e. – "01" for Sub-file 1). The program will validate the information to ensure that table names for both the old and new sub-files exist, that the user has write-access to the tables, and that there is no data already in the "TO" data tables. Error messages will be displayed to the screen if the new sub-file tables do not exist, or if the new sub-file is found to contain data already. When these errors occur, the program execution is aborted with an error message telling the user what happened. Once the user's access level and the status of the old and new sub-files have been validated, the split process

begins. Every other site is read out from the "FROM" sub-file and written to the "TO" sub-file. Once the writes have been verified as successful, the index record for each moved site is updated to indicate the correct sub-file for that site, and the site is deleted from the original "FROM" sub-file. At completion of the process, the count of site records copied is displayed at the user's terminal.

4.10.7 Delete Measured Unit-Values

The UV_DELMEAS program allows the user to delete the "Measured" Unit-Values data for a specified Data Descriptor (DD) for a user-specified time period. Access level of "ADBA" or "SYST" is required. This program deletes the actual sensor data from the database, and should be used with caution.

Introduction

This program is designed to delete from the database the raw measurements from a sensor. This may be required when there are transmission errors from an automated data input connection (DCP satellite link, EDL upload, etc.) or other problems that require reloading the raw data.

Unless the measured data is being deleted in preparation for reloading the raw inputs from the original source (satellite reprocessing, EDL, chart, etc), the sensor data will be obliterated from the database. For this reason, access to the program is restricted, and a warning message is presented to the user (see below) to make certain that the user understands the impacts of what the program will do.

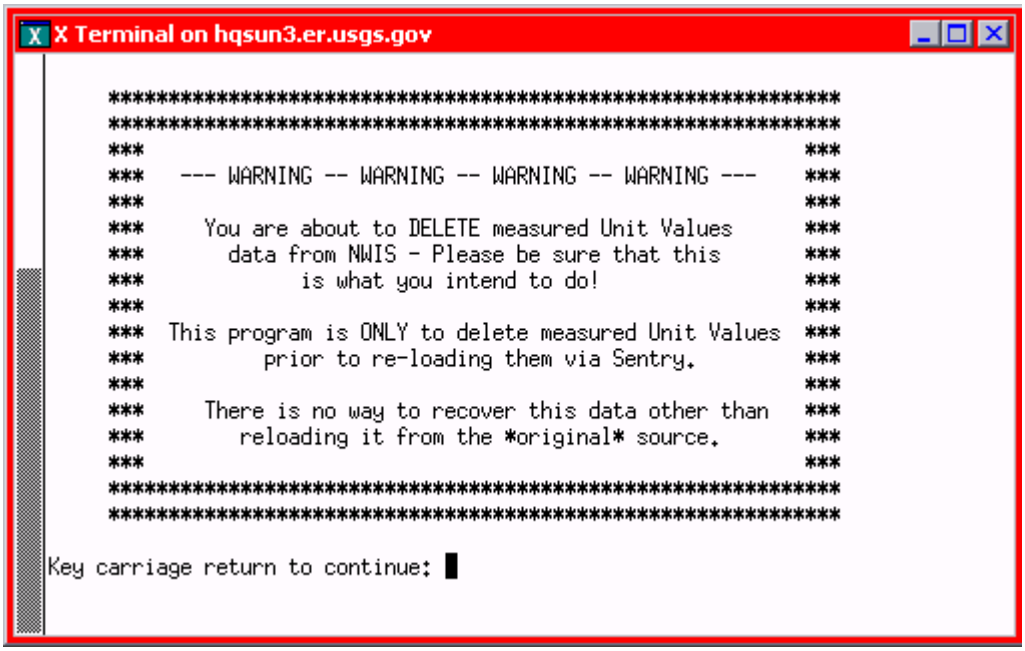

Program Operation

After the warning screen is displayed and acknowledged, the UV_DELMEAS program uses the standard ADAPS startup screens to let the user choose the station, DD, dates, and output file destination, as shown below:

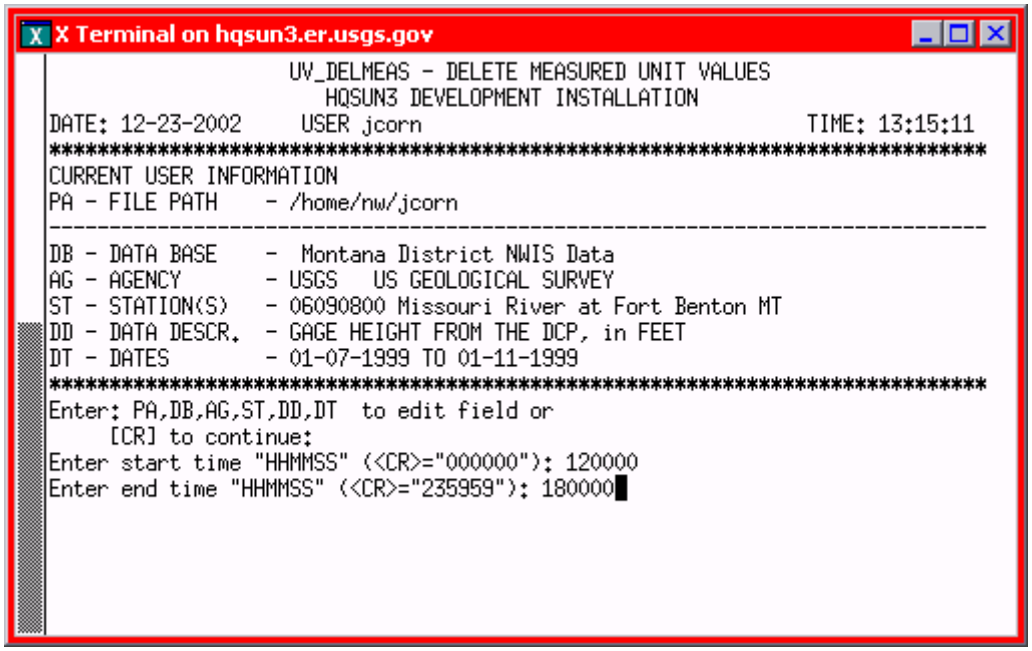

When the user has chosen the correct parameters for the data to be deleted, and entered a <CR> to continue, the program will ask for a starting and ending time, then scan the database to determine what "Measured" UV data sources are present. If there is more than one transport method found for the DD (for example, the Preferred Input is DCP, with an EDL defined for backup data), the program will display a list of the available transport methods and query the user to select one. Once the transport method has been chosen, or if there is only a single transport method for the DD, the program will display the following screen:

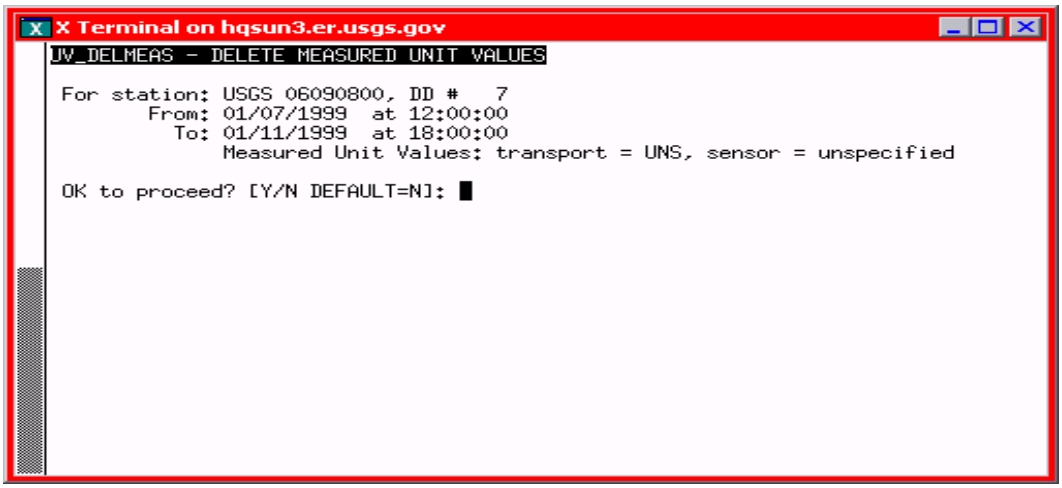

To delete the "Measured" UV data, answer this query with <Y>, or enter <CR> to cancel the operation. After the data has been deleted from the database, the program will automatically perform a recomputation of the period deleted to maintain data consistency. ***This recomputation will wipe out all "edited," "shift," "corr," and "computed" Unit-Values and any Daily-Values in the database, unless the DV data is "in Review" or "Approved" status.*** After the recompute is completed, the program returns to the ADAPS startup screens where another station, DD, or date range may be selected. Once the "measured" UV data has been deleted, the raw data must be reloaded via Sentry or manual downloads and reprocessed to generate the "edited" and "computed" UVs and the DVs for the period.