

09 October 2026

MHS Information Platform (MIP) MHS Data Repository (MDR) and SAS Computing Environment (SCE) Users Guide

Version 2.0

Revision History

Version	Date	Description of Changes
1.01	December 2022	• Major updates to restructure document for MIP SCE environment considerations will be published as a working draft for use.
1.1	January 2023	• Includes updates identified during the MIP MDR Limited User Testing event. Removed the "DRAFT" watermark to finalize the version.
1.2	January 2023	• Made document 508 compliant. • Enhance modifiedDate example.
1.3	February 10, 2023	• Edited the filesGet syntax for Command Line Interface. • Edited smart quotes.
1.4	May 2, 2023	• Added contact information for MDR Operations and Processing. • Added clarification on filesGet,Remove functions. • Updated user training and resource links.
1.5	September 13, 2023	• Updated saswork workspace and job submission information.
2.0	October 9, 2026	• Updated for RHEL8 server changes and processes that have been modified since last version of this guide.

Table of Contents

I. Introduction and MDR Overview 5
II. System Overview and Technical Architecture 6
2.1 Connections to Internal and External Resources 6
2.2 Data Refresh and Availability 6
2.3 System Training 7
2.4 File Versioning 7
III. Space Allocation and Related Issues 7
IV. Storing and Securing Privacy Act and HIPAA Data 8
2.5 Restricted Files 9
2.6 Keeping Logs 9
V. Getting Started 11
2.7 Connecting to the SCE 11
2.8 AVHE Login 11
2.9 Getting Help 11
2.10 Obtain "Username" 12
2.11 MIP AWS- WinSCP Application 12
2.12 MIP AWS- PuTTY Application Guide 12
2.13 MIP AWS-SAS EG (Enterprise Guide) Application 12
2.14 Working with AWS S3 12
2.15 Submitting SAS Jobs in the SCE via PuTTY -- SAS Job Queue 13
2.16 Submitting SAS Jobs in the SCE via SAS Enterprise Guide 18
2.17 Scheduling SAS Jobs 19
2.18 Moving Files to the SCE from External Systems: Data Imports and Extractions 22
2.19 Commonly Used UNIX Commands 22
2.20 Linux File Naming Conventions 28
2.21 The Linux Directory Tree 29
2.22 Getting Around Directories 30
Appendix A: The NANO Editor (For editing in Linux, a clone of PICO) and VI editor 31
Appendix B: MIP SCE Quick Reference Guide 39
Appendix C: AVHE MIP AWS -- WinSCP Login to MIPSAPSASP02 40
Appendix D: Launch PuTTY SSH Session to MIPSAPSASP02 40
Appendix E: SAS Create Enterprise Guide Profile to MIPSAPSASP01 41

Appendix F: Connecting to SQL Databases 43
2.1 F.1: Pulling Data from Redshift with Explicit Pass-Through SQL 44
2.2 F.2: Pulling Data from Redshift with Implicit Pass-Through SQL 45
2.3 F.3: Pushing a Temp Table from SAS to Redshift Using Bulkload 46
2.4 F.4: Establishing a Kerberos Ticket to Enable the Connection from SAS to SQL Farm 47
2.5 F.5: Pulling Data from SQL Farm with Explicit Pass-Through SQL 48
2.6 F.6: Pulling Data from SQL Farm with Implicit Pass-Through SQL 49
2.7 F.7: Pushing Data from SAS to SQL Farm 50

I. Introduction and MDR Overview

The SAS Computing Environment (SCE) provides authorized individuals with access and query capability to specified Military Health System (MHS) Data Repository (MDR) data sets and other data sources within the MHS Information Platform (MIP), including Redshift and SQL Farm. This User Guide addresses the processes for accessing and utilizing the MDR/SCE. The SCE is for experienced healthcare analysts, who are enterprise-level power users possessing a thorough knowledge of SAS and existing MHS data sets.

Data sets contained within the MDR and MIP include workload, cost, population, enrollment, demographic, and reference data for the MHS. The data sets are arrayed in a structured manner. They are available post processing and are represented as a subset of the MDR file catalog:

- According to a defined periodicity or update cycle;
- Processed using known, tested software - with identified business rules;
- Quality controlled, documented, and reviewed; and
- Released for specified use.

MDR Information Resources

Additional information regarding MHS MDR Data Dictionary, MDR Functional Guide, and data import/export transmittal can be found at the following link below.

https://health.mil/Military-Health-Topics/Technology/Support-Areas/MDR-M2-ICD-Functional-References-and-Specification-Documents

The MDR Users Group

This group meets every second Wednesday of each month at 1300 ET. Topics covered are relevant to all MDR users and MDR data partners. Historical documentation is available at the link below.

https://info.health.mil/staff/cdao/MDRUsersGroup/Forms/AllItems.aspx

To join the MDR Users Group, contact:

Mr. James Huber: james.t.huber4.civ@health.mil
Mr. Keith Hofmann: keith.d.hofmann.ctr@health.mil

Information Alerts (Blasters)

Periodically, “blaster” messages will be sent via the Enterprise Intelligence and Data Solutions (EIDS) MIP Users Group (MUG) Alert system to communicate significant milestones and/or events (e.g. outages, system wide changes). To sign up for the MUG alerts, follow the steps in “How to Subscribe to EIDS Alert Calendar”.

How to Subscribe to EIDS Alert Calend

For questions about Calendar Alerts and import/export transmittals use the following email.
dha.jbcharleston.j-6.list.eids-mip-mdr-ops@health.mil

For questions about user profiles, data quality or currency users should submit a ticket. See the Getting Help section for information on how to submit a ticket.

II. System Overview and Technical Architecture

The MDR/SCE resides in the MHS Information Platform (MIP), which is in an AWS GovCloud environment specifically managed by the Enterprise Intelligence and Data Solutions (EIDS) Program Management Office (PMO). The data in the MIP are a combination of SAS data sets, ASCII flat files, and Redshift schemas. Most processed data exist in SAS data set form and takes advantage of SAS procedures and library functions.

Available MIP Tools and Software

The MIP offers a variety of tools for analysts to utilize (e.g. Python, WinSCP). See the link below to find the current versions in Application Virtual Hosting Environment (AVHE).

https://carepoint.health.mil/sites/MIP/mugteam/SitePages/Tools%20Main%20Page.aspx

2.1 Connections to Internal and External Resources

Authorized users can access the following external databases using their credentials: TED-Bench, TED-Query, PCDW-Q, PEPR. Libname statements to include in your SAS code will be provided once you are authorized access.

Authorized users can access MIP resources such as Redshift and SQL Farm. Users will need to submit an update to their eSAARs to obtain access.

Please see the following link for assistance.

Resources link for MIP users of Redshift and SQL Farm

2.2 Data Refresh and Availability

The MDR data sets in this environment will be refreshed with current week, month, quarter, or yearly data immediately after MDR processing is completed. For MDR data sets that are

processed, the update will be loaded upon receipt and completion of Quality Control (QC). For more information on data sets you can consult the MDR data dictionary located at

https://health.mil/Military-Health-Topics/Technology/Support-Areas/MDR-M2-ICD-Functional-References-and-Specification-Documents

2.3 System Training

Users are encouraged to join the MDR Users Group as described above.

The MIP Users Group (MUG) offers a multitude of training content to assist in learning how to leverage MIP capabilities. See the site navigation panel on the left side of the page for more information on Training.

https://carepoint.health.mil/sites/MIP/mugteam/SitePages/Home.aspx

SAS specific training can be found at the link below:

Link where SAS specific trainings can be found.

2.4 File Versioning

Backup of files that exist in the MDR/SCE are handled by Amazon Web Services (AWS) S3 versioning which is enabled for the SCE. Any output/work product removed from AWS S3 requires MIP system administration staff to restore. Versioning in Amazon S3 is a means of keeping multiple variants of an object in the same bucket. S3 Versioning is active in the SCE and allows users to preserve, retrieve, and restore every version of every object stored in S3 buckets. Versioning allows for recovery more easily from both unintended user actions and application failures.

All requests for file restoration are to be submitted via EIDS One Front Door or GSC ticket. See section: Getting Help

III. Space Allocation and Related Issues

A directory is a method of storing and organizing data files, such as SAS programs, output files, SAS data sets, and text files. It also allows for manipulation and retrieval of data within and across directories. Every directory on the SCE has rules that set permissions for user access that are determined based on sponsoring organization and project designation. Permissions are controlled by Role Based Access Controls (RBACs). Users can modify the RBACs to which they belong by submitting an electronic System Authorization Access Request (eSAAR):

https://carepoint.health.mil/sites/esaar/SitePages/eSAARHome.aspx

As these directories contain private rather than MDR files, the government and/or functional owner of the directory is responsible for:

- Ensuring the directory contains only data authorized by law and regulation;
- Ensuring Privacy Act and Health Insurance Portability and Accountability Act (HIPAA) laws and regulations are followed; and
- Notifying EIDS of both who is empowered to provide access authorizations for that directory, and through that agent, the list of users allowed into those project or sponsor- designated files.

The public MDR data (/mdrfile) is currently allocated 48 TB. This includes data for MDR PUB, MDR APUB, MDR REF, and MDR Restricted.

The private SCE locations (/scefile) is currently allocated 24 TB. This space allocation is shared by all users of the SCE. Users can monitor the amount of space they are using with the "du" command in their directory. Users can also monitor if the SCE is reaching its limit using the "df" command. Automated processes with S3 and Lustre will also help control space utilization in the SCE.

EIDS is implementing a tiered storage solution in the MDR/SCE. Older/less commonly used data will be archived, moved to different S3 buckets, and eventually transitioned to Glacier. This data is still available, but may require a GSC ticket (see Getting Help) to recover the files.

- Archived data from the MDR (/mdrfile) will be visible (though not necessarily accessible) in /mdrfile/mip_archive. After that path, the path will mimic the normal MDR directory structure. For example, the file that was originally stored here: /mdr/apub/genesis3/admission/fy24/d231129/fy24.sas7bdat is now archived here: /mdrfile/mip_archive**/mdr/apub/genesis3/admission/fy24/d231129/fy24.sas7bdat**
- Archived data from the SCE (/scefile) will be visible (though not necessarily accessible) in /scefile/mip_archive. After that path, the path will mimic the original directory structure.
- Note that "/mip_archive" is aliased to "/scefile/mip_archive", but to access the MDR archive, you need to use the "/mdrfile" prefix.
- In most cases, the prefix (/mdrfile vs. /scefile) is not necessary.

IV. Storing and Securing Privacy Act and HIPAA Data

Health Insurance Portability and Accountability Act (HIPAA) requires healthcare systems to record any use or disclosure of an individuals' healthcare data.

Users are responsible for full compliance with HIPAA and the Privacy Act for data created or placed into all files systems on the SCE. EIDS does not track data stored in user workspace within these file systems. Private Health Information (PHI) data created as a result of programs extracting data from the MDR datasets for download or those originating outside of the SCE and imported/uploaded/keyed into these directories must be reported via e-mail to the SCE Functional Proponent (james.t.huber4.civ@health.mil) with the subject line HIPAA download Request. This applies to the movement of data outside of the SCE architecture.

The email should be sent **before the file is downloaded or uploaded**. Multiple files can be listed in the same email. The e-mail needs to include a brief description of the file, approximate size, source, and destination directory, along with your identity and organization. This information may be used to determine if an Import or Export Transmittal form is required (e.g., large amounts of data).

Downloads of PHI data that are 10,000 rows or greater require an Export Transmittal request. Users should not download this data themselves via WinSCP or any other FTP software.

EXAMPLE OF DOWNLOADING
EMAIL:
25-March 2022: Downloaded Direct Care SIDR and SADR records from /xyz/abc to hard drive; ~2,500 records total, with PHI and PII data. Delivered to OTSG Decision Support Center via secure FTP (HIPAA compliant). Data deleted from hard drive after delivery. John Doe, ABC Company (Decision Support).

EXAMPLE OF UPLOADING
EMAIL:
25-March- 2022: Uploaded DMDC-provided file containing 2,500 activated guard/reserve records with SSNs and minimal activation information, saved into the "ABC Company" file system. John Doe, ABC Company (Decision Support).

2.5 Restricted Files

Some data in the MDR is considered sensitive and requires special justification. To obtain access to files in /mdr/restricted, users must complete a special justification form signed by their Government Sponsor/supervisor. The form can be found at the following link and must be submitted to Mr. James Huber.

https://www.health.mil/Military-Health-Topics/Technology/Support-Areas/MDR-M2-ICD-Functional-References-and-Specification-Documents

2.6 Keeping Logs

In healthcare systems such as the SCE, EIDS can identify data access or disclosure to the point of knowing which SCE user extracted data from an MDR data set on a specific date. To fulfill the remainder of the requirement for records of access, **each SCE user must retain logs of program executions on any files containing Privacy Act data**. These logs can be the normal electronic logs produced by SAS and stored by the user on the SCE. Upon request, these logs must be made available to EIDS, DHA Privacy Office, and/or OWHA.

One example for automatically saving the SAS logs to a dedicated log directory includes the following SAS macro which is coded at the beginning of every SAS program:

```
%let dir1=%str(/main/secondary/task/log/);
%let pgmname=%str(xxxpgm_name);
%let date=&sysdate;
%let time=&systime;
%let ext1=%str(log);
%let break=%str(-);
%let final1=&dir1&pgmname&date&break&time..&ext1;
```

```
%put &final1;
proc printto log="&final1";
run; quit;
```

/main = the main directory of your dedicated workspace
/secondary = secondary directory, if applicable
/task = specific task name
/log = log directory set up for each task
xxx = users initials
pgm_name = program name

Before running this, make sure you have created the log directory in the file path. This is where all the logs, every time the program is run, will be stored for future reference, if necessary.

Notes:
1) create task areas for easier reference
2) use program names that make sense for the task
3) user's initials allows for identification of who ran the program

By default, the system uses the --altlog option, which also puts a *.log file in the directory you submitted the program from and is where the programmer can review the results of the program. So, essentially two duplicate logs are created, one for the programmer to review, the other to be HIPAA compliant. The HIPAA compliant log never gets overwritten. The reviewer log gets overwritten every time the program is rerun.

V. Getting Started

2.7 Connecting to the SCE

To obtain access to the AVHE, MIP, and MDR/SCE, users need to follow the process listed at the following link on the MUG CarePoint site. MIP applications (including MDR) are only available via AVHE.

https://carepoint.health.mil/sites/eSAAR/SitePages/eSAARHome.aspx

2.8 AVHE Login

Use this link https://avhe.health.mil/Citrix/AVHEWeb/ to access AVHE.

2.9 Getting Help

To get assistance, you can use the EIDS One Front Door on the MUG CarePoint site:

https://carepoint.health.mil/sites/mip/Pages/MUGHome.aspx

Once logged onto the site, select "Get Help" on the menu on the left side of the screen. Select EIDS Front Door and follow the instructions on that page.

You can also submit Global Service Center (GSC) Ticket for issues, by sending an email in the following format. GSC tickets are sometimes called Service Now or SNOW tickets. You may be asked to provide your DHACS or DHAINC ticket number (e.g. DHACS11398873).

The Subject MUST contain "MHS Information Platform (MIP)" and "MIP MDR" in the SUBJECT in order to get the ticket to the correct team.

TO: dha.jbsa.j-6.mbx.MHS-Service-Desk@health.mil

CC: **DHA JB Charleston J-6 List PS** dha.jbcharleston.j-6.list.ps@health.mil

SUBJECT: **GSC Ticket: MHS Information Platform (MIP) Ticket for MIP MDR**

BODY:

Route To: MHS Information Platform (MIP)
- **Name:** [REQUIRED]
- **Phone Number:** [REQUIRED]
- **E-mail Address:** [REQUIRED]
- **Site: (optional)**
- **Organization:** [REQUIRED]

Department: (opt)
- **Building Number: (opt)**
- **Room Number: (opt)**

Full Computer Name (opt)

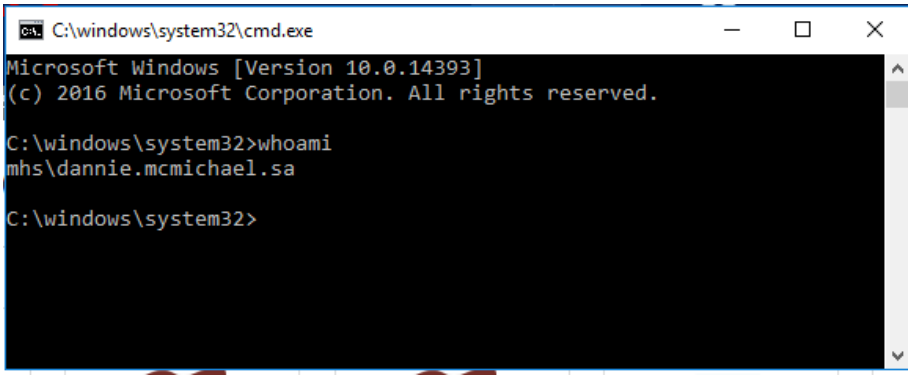
A detailed description of the issue/request: (include as much details as possible such as URLs, screenshots, Tableau bitab folders, application names, etc.) [REQUIRED]

SIGNATURE BLOCK

2.10 Obtain "Username"

Background: Your "Username" is required when accessing the MDR/SCE through Putty or WinSCP. SAS Enterprise Guide does not require a user to know their username. Usernames do not follow a standard naming convention. To determine what your "Username" is:

1. Open AVHE 'MIP AWS -- Command Prompt'
2. Enter 'whoami' command
3. Username does not include the domain 'mhs\' (e.g. dannie.mcmichael.sa)



Once the account is approved, permissions are established, and the user obtains their username, the user will be able to access the SCE.

2.11 MIP AWS- WinSCP Application

See Appendix C: AVHE MIP AWS -- WinSCP Login to MIPSAPSASP02

2.12 MIP AWS- PuTTY Application Guide

See Appendix D: Launch PuTTY SSH Session to MIPSAPSASP02

2.13 MIP AWS-SAS EG (Enterprise Guide) Application

See Appendix E: SAS Create Enterprise Guide Profile to MIPSAPSASP01

2.14 Working with AWS S3

The MIP SAS servers are now using Amazon FSx for Lustre, an AWS service that integrates high-performance local storage with Amazon S3.

- All files are always stored in S3. Only recently accessed files are cached locally on Lustre for faster processing.
- There are no more .S3 sync files (in use from 2023-2026). Files are automatically backed up to S3, eliminating the need for manual S3 synchronization.
- Deprecated macros. The filesGet, filesPut, and filesRemove macros (in use from 2023-2026) are no longer required and have been retired.

To help control local storage costs:

- SAS Administrators will run a nightly script to release files from local Lustre storage if they have not been accessed in the last 10 days.
- Releasing a file only removes it from local Lustre storage — the file remains safely stored in S3.
- If you access a released file, it is automatically retrieved from S3 and made available locally again.

S3 permissions are dependent upon the user's access. If a user has read access, they will be able to read/obtain files from S3. If a user has write access, they will be able to overwrite a file within S3.

Note that if a user uses the "rm" command in PuTTY (or delete in WinSCP) this will delete both the local copy in Lustre and the backup copy in S3. This is different than from the era when filesGet/filesRemove was in use (2023-2026), when using "rm" would delete the local copy only. If a file is accidentally removed (or overwritten) from S3, it can often be recovered. File a help desk ticket within 60 days that takes advantage of the versioning in S3 described above (see Getting Help).

The following shows an example of the interaction between S3 and the local SAS server:

The metadata for all files is visible using the "ls" command, but these files are not cached locally on Lustre (are not available on the SAS server). They are in S3 only. We can tell by using the "du" command, because they take up no space on the SAS server.

```
[khofmann@MIPSAPSASP02 jan21]$ ls -lrt
total 6
-rw-rw----. 1 root sas_scehcret  2461134 Oct 31  2022 colo.txt
-rw-rw----. 1 root sas_scehcret 18718273 Oct 31  2022 colon.sas7bdat.gz

[khofmann@MIPSAPSASP02 jan21]$ du -h *
512     colon.sas7bdat.gz
512     colo.txt
```

Once the user interacts with "colo.txt" (in this case by viewing it with the vi editor), the file is implicitly/automatically downloaded from S3 and made available on the SAS server. Once this interaction occurs, the file takes up space on the SAS server.

```
[khofmann@MIPSAPSASP02 jan21]$ vi colo.txt

[khofmann@MIPSAPSASP02 jan21]$ du -h *
512     colon.sas7bdat.gz
2.4M    colo.txt
```

2.15 Submitting SAS Jobs in the SCE via PuTTY – SAS Job Queue

SAS jobs run in the background (by default) on the SCE (via PuTTY) to prevent running jobs from failing if the SSH session times out. You may create or edit SAS programs on the SCE using the NANO (PICO) or vi editor. If you prefer to work on your PC, you may edit your programs in Windows and then use a secure file transfer to copy them to the server. Please note that Windows may add hidden characters that cannot be interpreted by the SAS engine.

Due to the high usage of the SCE, there is a limit of 10 programs per user running at a time. If you already have ten programs running, you'll need to wait until one job finishes before setting off another SAS program. If more than ten are submitted, the most recent submission will be 'in queue' until one of the prior programs finish and release resources.

It is recommended that you name your programs with a .sas extension (e.g. program_name.sas). To submit a SAS job:

1. Log on to the server using an SSH client and change to the directory where the SAS program is located in your workspace.

   Example:
   [mipsapsasp02}:/home/med.ds.osd.mil/$USER > **cd** /hpae2/

   Note: Do not store code or data in your home directory, specifically /home/med.ds.osd.mil/$USER. Files in this area may be deleted without warning.

2. Use the "sas" command to submit your program to the jobQueue:

   sas <program_name.sas>

Your program will be submitted to the jobQueue. The SCE can have up to 48 jobs running at a time. Any jobs after that enter a queue to be executed sequentially. Keep this in mind on busy days and try to only run jobs you need when you see jobs waiting.
You can see the queue by typing **jobQueue –list**.

Using Environment Variables

Any environment variable defined in the shell when the jobQueue is invoked will be passed to the SAS program. The jobQueue has key arguments, which can be found by executing jobQueue –help. These key arguments are specifically jobQueue options, all other arguments will be passed the SAS program. Passing a SAS argument can be done using the following:

sas <program_name> -<variable_name> <value>

SAS Temporary Workspace

There is one SAS temporary workspace area on the SCE, located here: /saswork. When submitting a SAS job, it will default to this location.

Commands for Submitting SAS Jobs

- To start a job running in the default workspace, the command is **sas *jobname.sas***

- To determine what jobs are running in the SAS Temporary Workspaces simply type **'jobQueue –list'** at the command prompt and look under "Queue" to determine the jobs running or queued to run in the SAS Temporary saswork area.

Previously, there was a "sasbig" command and a –workspace option (until Sept 2023), but those are no longer recommended.

When you submit a SAS job, a temporary work folder will be created like this:
/saswork/$USER/<jobPID>/
(see below where you can find jobPID)

For example: /sas_data/saswork/khofmann/68971/

SAS Arguments Configuration

Memory

Users can control the amount of memory available to their SAS program. The default is 2GB and should be sufficient for most users. Using more than 2GB will impact performance for other users.

If necessary, use the -MEMSIZE option to make adjustments for a single use program run:

sas *jobname.sas -MEMSIZE 4G*

Checking Job Status

To find the current status of running jobs, type **'jobQueue –list'** at the command prompt.

=

```
-bash-4.2$ sas caper.sas
-bash-4.2$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue        Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/05 18:03 khofmann   15434   15725   sas caper.sas

Run: R=Running, Q=Queue, S=Scheduled
```

This job is running because the "Run" column has R and there is a progPID:

```
jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue        Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/05 18:03 khofmann   15434   15725   sas caper.sas

Run: R=Running, Q=Queue, S=Scheduled
```

Queued (Q) or Scheduled (S) jobs will have a jobPID, but not a progPID

To find the current status of running jobs only for the user, type **'jobQueue –list –user $USER'** at the command prompt.

```
$ jobQueue --list --user khofmann

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue        Run Run-Date    User       jobPID  progPID Program
------------ --- ----------- ---------- ------- ------- -------------------------
saswork       R  09/04 15:19 khofmann   580489  579643  sas caper.sas

Run: R=Running, Q=Queue, S=Scheduled
```

All submitted jobs have completed when the output is empty:

```
[dannie.mcmichael.sa@mipsapsasusrp01 ~]$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue        Run Run-Date    User       jobPID  progPID Program
------------ --- ----------- ---------- ------- ------- -------------------------

[dannie.mcmichael.sa@mipsapsasusrp01 ~]$
```

Check the .log files for results.

Canceling a submitted Job – When job is in queue or scheduled status, but before job is 'Running'

To cancel a submitted job, use the **jobQueue –remove <jobPID>** command. It will remove the job from the jobQueue.

1. List your jobs within the jobQueue: **jobQueue –list –user $USER**
2. Take the 'jobPID' of the job you wish to remove from the Queue: **jobQueue –remove <jobPID>**

```
-bash-4.2$ date
Wed Dec  7 11:59:48 EST 2022
-bash-4.2$ sas test61.sas --datetimeExec '2022-12-07 13:00:00'
-bash-4.2$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue       Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/07 11:54 djuckett   119322  119599  sas ...Freq_FY12+_v1.sas
saswork500G  S  12/07 13:00 khofmann   15238           sas test61.sas

Run: R=Running, Q=Queue, S=Scheduled

-bash-4.2$ jobQueue --remove 15238
jobPID: 15238 has removed
-bash-4.2$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue       Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/07 11:54 djuckett   119322  119599  sas ...Freq_FY12+_v1.sas

Run: R=Running, Q=Queue, S=Scheduled
```

This job was not yet running. It was scheduled (S) and there was no progPID. Therefore, we used jobQueue --remove (and the jobPID).

Note: jobQueue --remove <jobPID> will only allow a user to remove their own submitted jobs.

Note: Sometimes SAS jobs which terminate abnormally leave residual data within the temporary SAS workspace. This data will be removed by the system administrator on a routine basis.

Remember, do not store programs, scripts, data, formats, etc. under your user home directory: /home/med.ds.osd.mil/$USER/. There is very limited space available for the home directories. If this space becomes full, all users will be locked out of the SCE. If the system administration staff sees files in the home directories other than those placed there by the staff, they will be removed after due diligence.

Canceling a submitted Job – When job is 'Running'

To cancel a submitted job, use the **jobQueue –stop <progPID>** command. It will remove the job from the jobQueue.

1. List your jobs within the jobQueue: **jobQueue –list**
2. Take the 'progPID' of the job you wish to remove from the Queue: **jobQueue –stop <progPID>**

```
-bash-4.2$ sas test61.sas
-bash-4.2$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue       Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/07 11:54 djuckett   119322  119599  sas ...Freq_FY12+_v1.sas
saswork500G  R  12/07 11:56 khofmann   127359  127652  sas test61.sas

Run: R=Running, Q=Queue, S=Scheduled
-bash-4.2$ jobQueue --stop 127652
progPID: 127652 has stopped
-bash-4.2$ jobQueue --list

jobQueue - List Jobs in Queue
--------------------------------------------------------------------------------
Queue       Run Run-Date    User       jobPID  progPID Program
----------- --- ----------- ---------- ------- ------- -------------------------
saswork500G  R  12/07 11:54 djuckett   119322  119599  sas ...Freq_FY12+_v1.sas

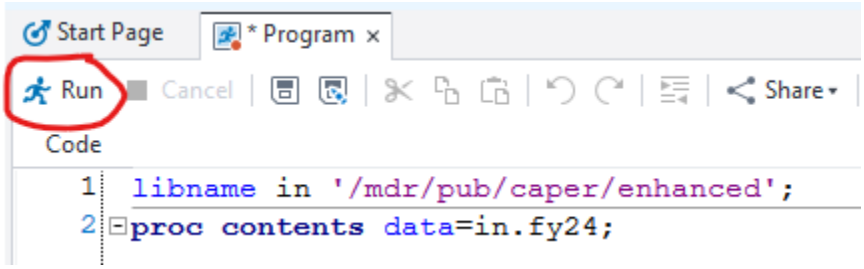
Run: R=Running, Q=Queue, S=Scheduled
```

This job was running (R) and there was a progPID, therefore, we used jobQueue --stop (with the progPID)

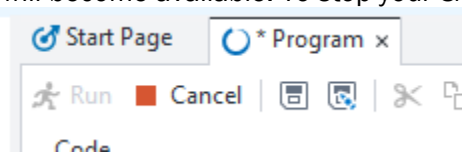
Note: jobQueue --stop <progPID> will only allow a user to stop their own submitted jobs

2.16 Submitting SAS Jobs in the SCE via SAS Enterprise Guide

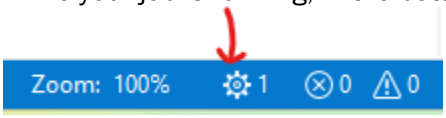
Unlike submitting jobs via PuTTY, SAS programs submitted via SAS Enterprise Guide (EG) run in the foreground. You must stay connected to SAS EG (and AVHE) while your program runs. Jobs submitted in SAS EG do not enter the jobQueue. To submit a SAS program in SAS EG, simply click "Run".



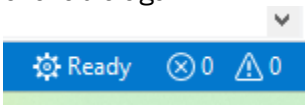
While your job is running, the "Run" button will be greyed out and the "Cancel" button (red square) will become available. To stop your SAS program, press "Cancel".



While your job is running, in the bottom right corner of SAS EG, there will be a spinning gear.



Once your job finishes, the gear will stop spinning and will say “Ready”. The “Run” button will become available again.



2.17 Scheduling SAS Jobs

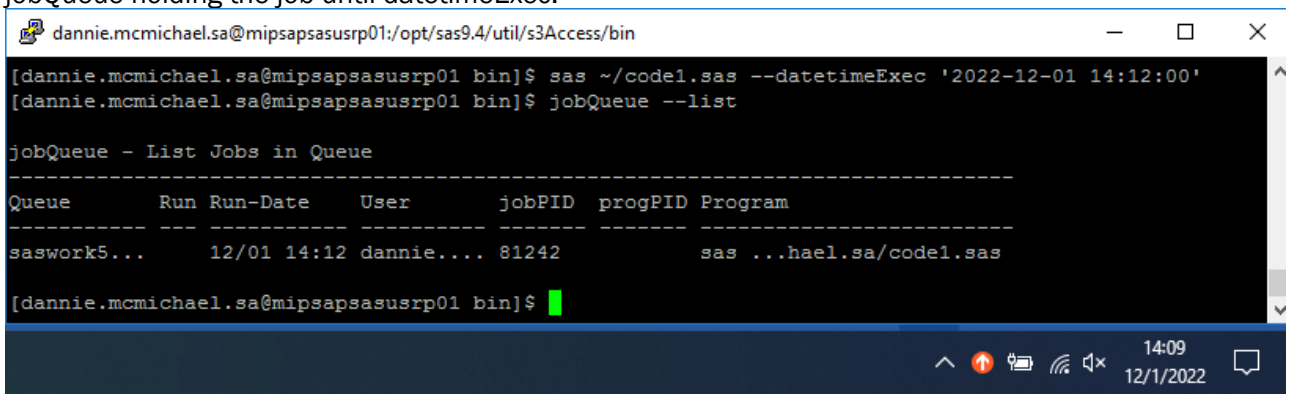
Background: There are two different types of scheduled jobs, one-time execution and recurring. Depending on execution requirements, refer to the appropriate scheduling method. Jobs can only be scheduled via PuTTY, not SAS EG nor WinSCP.

One-time execution:

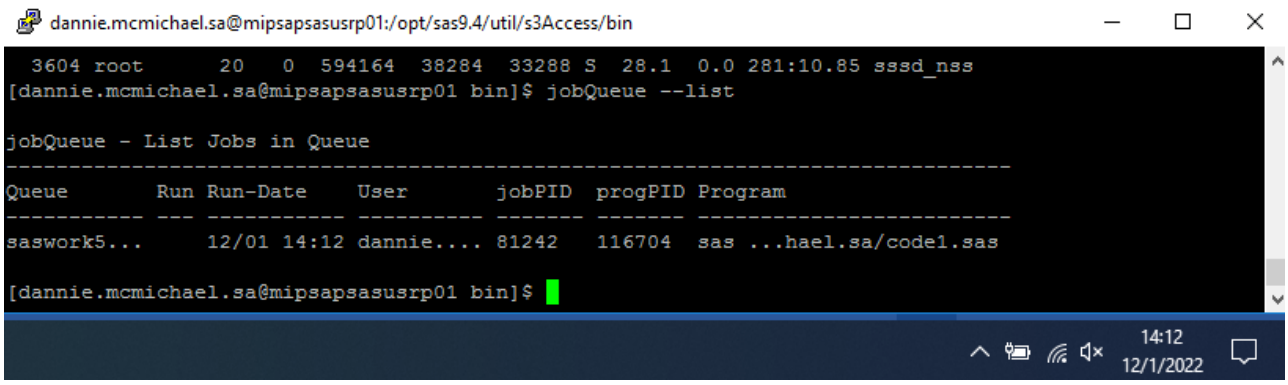
A one-time execution utilizes the jobQueue to schedule the job at the set time. The jobQueue will hold the job until execution time has been met, then releases the job into the queue. If there are other jobs exceeding the queue limits, the job will execute once resources are available.

Syntax: sas /dir/file.sas --datetimeExec '2023-01-01 23:59:59'

jobQueue holding the job until datetimeExec:



JobQueue release job at datetimeExec:



Recurring execution:

Crontab is a Linux service that allows users to execute commands at specific times.
Crontab Man: https://linux.die.net/man/5/crontab

A recurring execution requires the use of cron. A user can edit their cron by using the following command: **crontab -e** (which will default to the vi editor). If users prefer to use the nano/pico editor, they can submit this command: **EDITOR=nano crontab -e**. Users of cron tab need to consider that most MDR data is refreshed weekly or monthly. Please do not run daily cron jobs unnecessarily. Consider scheduling cron jobs to run at off-peak times like overnight or weekends.

A cron job cannot submit a SAS program directly. To schedule a SAS program, users can use cron to schedule a script which calls a SAS program (cron -> script -> SAS).

First, build a reusable "wrapper" shell script (assume it is saved as /path/sas_wrapper.sh) that can submit a SAS program:

```
#!/bin/bash

# source the SAS environment
source /opt/sas9.4/util/startup/sas-initial.sh

# Before running SAS, change to the directory where the SAS program lives
# This will allow relative directory references to work properly
cd "$(dirname "$1")" || { echo "ERROR: Could not cd to directory of $1"; exit 1; }

# Run SAS, passing the .sas program in as the first argument
/opt/sas9.4/util/jobQueue/sas "$1"
```

In your crontab, set the schedule you want, refer to the wrapper script, and feed it the SAS program you want to run as the first argument:

```
$ crontab -l
#Run my test71 program on the 6th of each month at 17:03 eastern time
3 17 6 * * /path/sas_wrapper.sh /hpae2/kennell/keith/random/test71.sas
```

Color-coded version:

```
3 17 6 * * /path/sas_wrapper.sh /hpae2/kennell/keith/random/test71.sas
```

With this cron, at 17:03 (eastern time) on the 6th of each month (regardless of day of the week), cron will call the "sas_wrapper.sh" script. That script will set up the SAS environment and then call the **sas** command on the program /hpae2/kennell/keith/random/test71.sas.

Reference the crontab manual for help creating a scheduled execution. When done creating the entry, if using the default vi editor, **:w** will write the entry to crontab and **:q** will quit crontab.

Note: Your processes are dependent on Kerberos authentication. Kerberos operates by validating users and issuing tickets. These tickets expire and requires to re-authenticate. Periodically logging into the SCE server will refresh your Kerberos ticket. See section F.4 of Appendix F for more details.

2.18 Moving Files to the SCE from External Systems: Data Imports and Extractions

SCE users may periodically require files be made available on the server that only exist on external systems or export data from the MIP. Please use the following link to find the data import and export transmittal forms.

https://health.mil/Military-Health-Topics/Technology/Support-Areas/MDR-M2-ICD-Functional-References-and-Specification-Documents

2.19 Commonly Used UNIX Commands

Following are some of the basic commands that you will need in the UNIX environment.

Command Reference Table

ls

The “ls” command lists files in a directory. There are many options you can use with the "ls" command. For example "ls -a" will list ALL of the files in a directory (files and directories whose name begin with a period are considered hidden files). "ls -al" will list ALL files in LONG format. Here is an example directory listing:

```
www:~> ls -laF
total 34
drwxr-xr-x 4 root  other   2048 Sep 24 11:25 ./
drwxr-xr-x 102 root  root   2048 Sep 24 11:24 ../
-rw-r--r-- 1 root  other   384 Jan 6 1998 .bashrc
-rw-r--r-- 1 root  other   383 Jan 6 1998 .cshrc
-rw-r--r-- 1 root  other   706 Jan 6 1998 .defaults
-rw-r--r-- 1 root  other   1858 Jan 6 1998 .emacs
-rw------- 1 root  other   2797 Jan 6 1998 .history
-rw-r--r-- 1 root  other   137 Jan 6 1998 .login
-rw-r--r-- 1 root  other   138 Jan 6 1998 .logout
-rw-r--r--          1 root        other        232  Jan  6  1998  .mailrc
drwx------          2 root        other        512  Jan  6  1998  Mail/
drwxr-xr-x          2 root        other        512 Jan 6 1998 bin/
```

In this example, the first block of each row (-rw-r--r--) shows the type and permissions for an individual file within your current directory. The first column tells you the type of file: '-' means it is a normal file, 'd' means it is a directory, and an 'l' means it is a link to another file (much like a shortcut in Windows). The next nine positions show the file's Read, Write, and Execute permissions in blocks of three. The first three characters (often 'rw-') are the permissions for the owner of the file, the second three are group permissions, and the final three are permissions for anyone else besides the owner and anyone in the group.

The next column is a number representing how many physical links to the directory or file exist on the file system. Files generally have one link to them since directories cannot link to files specifically. Notice that the directory up one level, '..', contains a large number of links (102) due to the number of home-directories that reference it as '..'.

The next four columns contain the name of the owner of the file (equivalent to the first three 'rwx' permissions), the group owner (second three 'rwx'), the file size, file date, and file name. Note that by using the '-F' flag (must be a capital F), directories are marked with the / symbol, and the / is not actually part of the directory name. Also, files marked as executable are shown with the '*' symbol following the name (only when the -F flag is set).

cd

The "cd" command is used to change directories. The "cd /" will take you to the root directory of the directory tree; "cd /usr" will take you to the directory named 'usr' that is located off of the root directory; "cd mail" will take you into a directory called 'mail' off from your current directory (note the lack of '/' before the directory name); and "cd .." will take you to the parent directory in the directory tree. Some command-line shells, such as "bash" or "tcsh" expand the ~ symbol to mean a person's home directory. You could then either type "cd ~" to change back to your home directory, or "cd ~another user" to go into another user's home directory.

pwd

The "pwd" command prints/outputs your present working directory to standard output.

mkdir, rmdir

These commands are used to create and remove directories. The "mkdir class" will create a directory called class in your current directory. Whereas "rmdir class" will delete the directory named class. Note that "rmdir" only works on directories that do not contain any files.

cp

To copy a file, use the "cp" command. The syntax is "cp *source destination*", where *source* is a source file and *destination* is either a file name or directory in which to place the file. For example, "cp /usr/local/seminar.c ." will copy a file located in '/usr/local' named 'seminar.c' to '.' which is the symbol that stands for your current directory. "cp /usr/local/seminar.c my.c" will copy the 'seminar.c' file to your current directory under the name 'my.c'.

mv

The "mv" command is used to move *and/or* rename a file. "mv my.c first.c" will rename a file called 'my.c' to 'first.c'. Assuming that 'class' is a directory, "mv my.c class" would move the file 'my.c' into the 'class' directory. If 'class' is a file, this operation will instead delete 'class' and then rename 'my.c' as 'class', which is probably something you didn't want. You can also rename and move at the same time. For example, "mv my.c class/first.c" would move the file 'my.c' into the 'class' directory and rename it to 'first.c' all in one step.

rm

The "rm" command deletes, or removes, files from the files ystem. Typing "rm my.c" would delete the file 'my.c'. You can also use wildcards. For example, "rm my*" will delete every file whose name started with the word 'my', including 'my.c'. Unlike Windows, these wildcards can appear anywhere in the file name, and do not need to be confined to the standard 8.3 notation used by MS-DOS. Directories and their contents can be deleted by using the '-r' flag

which stands for recursiveness. For example, "rm -r class" would delete ALL files and directories under 'class' and
'class' itself regardless if the directory is empty or
full.

cat, less, more, head, tail

These commands are used to view the contents of files. "cat" displays the entire file on the screen, regardless if your terminal can view everything at once (the top portion scrolls off, but can be retrieved by scrolling back in a window or sometimes by holding down the 'Shift' key and pressing 'Page Up'). "less" and "more" both display the contents of a file one full screen at a time. "head" is used to display the first few lines of a file, and "tail" is used to display the last few lines of a file. For example, "tail -25 my.c" displays the last 25 lines of the file 'my.c'. "cat first.c" would display all of 'first.c', and "more first.c" will display the file 'first.c' one screen at a time.

grep

Search a file for a particular pattern by using the “grep” command. "grep" stands for Get Regular Expression Pattern. This command searches a file for all lines matching a particular pattern and prints those matching lines on standard output. For example, "grep fun my.c" prints every line of the file ‘my.c’ that contains the word 'fun'.

man

The “man” command brings up the UNIX Manual Pages for a particular command. For example, to learn about the rest of the options for the "ls" command, type: "man ls"

w, who, whoami, id

These commands are used to find out the list of users currently logged into the system and what they are doing. "w" provides slightly different output than "who". "whoami" returns the username of the person currently logged in. "id" prints your User ID, and Group ID, and a list of all the groups you have access to.

chmod

The “chmod” command is used to change the permission bits, called mode, on a file or directory. Only the owner of a file or directory and the superuser can change the permission bits/mode. See the “ls” command for information on viewing a file or directory mode.

Every file and directory is assigned nine permission bits that comprise the mode. The mode is used to determine what operations may be performed on a file and by whom. The mode consists of Read, Write, and Execute bits in blocks of three (Execute is also called “search” when referring to a directory): one block for the owner of the file or directory, one block for the group that the file or directory belongs to, and one block for others. Each user will fit into only one of the three blocks/categories.

For files, the read bit allows the file to be opened and read. The write bit allows the contents of the file to be modified or truncated. The execute bit allows an executable file to be executed.

For directories, the combination of read and execute (aka “search”) bits allows the contents of the directory to be listed. The combination of write and execute bits allows files within that

directory to be created, deleted, and renamed. Having only the execute bit will allow the directory to be entered but not have any of its contents listed.

File or directory modes may be changed using octal notation or symbolic format. The octal notation format is described below. The basic syntax for the "chmod" command is:

chmod <octal or symbolic notation of read, write, execute permissions> <file or directory name>

The octal notation format sets file or directory permissions based on three octal numbers. The first number defines access for the user, the second number for the group, and the third number defines access for others. Each of the three numbers is a combination of permission bits. The octal notation is between 0 and 7.

Use the UNIX command CHMOD to change access permissions by summarizing the numerical values for Owner, Group and Other.

Digit	Permissions
0	None
1	Execute only
2	Write only
3	Write and execute
4	Read only
5	Read and execute
6	Read and write
7	Read, write, and execute

Examples:
#1) chmod 770 tedn.sas - allows owner and group to have read/write/execute privileges, other no privileges
#2) chmod 764 tedn.sas - allows owner read/write/execute, allows group read/write only, allows other read only

Using the "ls –l" command, example #1 and #2 above would appear as such:
#1) rwxrwx--- example1 groupname 2567 Sep 10 12:15 tedn.sas
#2) rwxrw-r-- example1 groupname 2567 Sep 10 12:15 tedn.sas

ps

The "ps" command is used to find the Process Status of all processes currently running on the machine. Typing "ps" without any options lists all programs currently running in your process group. "ps ux" or "ps -ef" lists all of the processes on the machine (Some UNIX versions prefer

one method over the other. If one way does not work, try the other). Combining the "ps" and "grep" commands through a pipe can give you a list of all the processes you have running, and can be done like this:

ps aux | grep <yourusername>

The '|' symbol, called a pipe, takes the results of the first command (in this case, "ps -aux") and uses it as input to the second command ("grep yourusername"). Let's take a look at an example:

```
USER  PID %CPU %MEM VSZ RSS TTY STAT TIME COMMAND
root   1 0.0 0.1 1048 112 ?  S 0:03 init [3]
gandalf  521   0    0    1652  0 tty2 SW   0:00                    [bash]
gandalf  652   0    0.7  1668  672 tty4 S  0:00                    -bash
gandalf  672   0    0    1516  0 tty4 SW   0:00                    [startx]
gandalf  677   0    0    2224  0 tty4 SW   0:00                    [xinit]
gandalf  4481  0.2  1.6  4104  1560 pts/0 S       0:00       vi
commands.html
gandalf  4452  0    1    1652  1012 pts/3  S 0:00 -bash
gandalf  4545  0    0.9  2496  936  pts/3
R                                  0:00 ps ux
```

The first column contains the user ID of the owner of that process. The second column contains the PID, or Process Identification number. The next two columns show the percentage of CPU Time and Memory the process has consumed during its entire running life. The next two columns show the Virtual Size of the program (total size of all memory in use) and how much of that memory is currently resident in the computer at the moment. Thus, the RSS number is usually much less than VSZ, because data are either shared, or is marked unused and are swapped to disk to free up room for other processes to use resident memory.

TTY shows which terminal the process is currently running on. Usually, tty1 through tty6 are processes controlled by a user at the main console (ALT-F1 through F6 to switch), and pts/# are processes from either inside X-Windows or through telnet in the network. STAT corresponds to the type of process status running at the time: 'R' means the process is currently active (as is the case for our “ps” command gathering the process data), 'S' means the process is sleeping, 'W' means the process is waiting for a signal to trigger swapping from disk and back into memory (note the 0 Resident Set Size), and 'Z' stands for a zombie process (see below). TIME shows the total amount of time the process has taken up CPU. In this case, each process except 'init' has taken up trace amounts of CPU, while 'init' has taken 3 seconds since the computer was started.

The last column of the "ps" command output contains the name of the process. The name can be represented in several different ways depending on the state of the process. If the name begins with a '-' symbol, such as '-bash', then it is recognized as a login shell. If the name is surrounded by [ ]'s, then the process is currently swapped out of memory to disk (with an RSS of zero).

Zombie Processes. Processes marked 'Z' occur during a coding error on behalf of the programmer to not properly *wait* for one of its children to exit after a *fork* call in the C language ("man" pages are available for both of these functions). Such processes are marked <defunct> and will only terminate when their parent process (PPID) exits.

kill, killall

Force a process to exit or terminate by using the “kill” command. This command is used to send signals to processes, most commonly the signal to kill a process (hence its name). Looking back at the "ps" example, to kill the "vi" process, you can type "kill 4481". Only programs that do not block the Terminate signal can be killed in this fashion. You instead must specify the “kill”
command using a '-9' as the first argument to the prompt. For example, to kill the unused,

swapped out 'bash' process,
type: kill -9 521

The "killall" command behaves in the same way as "kill", only each process matching the name you give it will be sent the signal. To kill all of your "bash" programs (and consequently kick yourself back to the login prompt), type:

killall -9 bash

Note: If you own a process that is using close to 99% of the CPU, it is generally a good idea to "kill -9" it, as it most likely won't respond to normal signal handling. The system PID 1, "init", can never be killed.

Note: Do not use kill to kill a SAS job, use the jobQueue –stop or jobQueue –remove commands.

As of July 2026, compress/uncompress functions are no longer available on the MDR/SCE. Use gzip/gunzip instead.

gzip

The “gzip” command compresses data using the DEFLATE algorithm to reduce the size of files. When possible, a compressed file replaces the original file, specified by the file parameter, with a .gz appended to its name. The compressed file retains the same ownership, modes, and modification time of the original file. If no files are specified, the standard input is compressed to the standard output. Files being compressed must have correct permissions to be replaced. For example, to gzip the ‘foo.txt’ file, type:

gzip foo.txt

The foo.txt file is compressed and renamed foo.txt.gz.

For more information about different options and techniques on the formation of the ‘gzip’ command, reference the man pages.

gunzip

The “gunzip” command (or gzip -d) restores original files that were compressed by the “gzip” command. Each compressed file specified by the file parameter is removed and replaced by an uncompressed copy. The uncompressed file has the same name as the compressed version but without the .gz extension. If the user has root authority, the expanded file retains the same owner, group, modes, and modification time as the original file. If the user does not have root authority, the file retains the same modes and modification time, but acquires a new owner and group. If no files are specified, standard input is expanded to standard output. For example, to gunzip the foo.txt.gz file, enter:

gunzip foo.txt.gz

The foo.txt.gz file is uncompressed and renamed foo.txt.

Note that gunzip is also compatible with files that were compressed with the legacy “compress” function (those that have a .Z extension). Many older MDR files were compressed using the “compress” function, e.g. /mdr/pub/deers/lvm/lvm6/fy24.txt.Z, but they can still be accessed using gunzip.

2.20 Linux File Naming Conventions

Absolute Naming

Within the UNIX directory structure there are two ways to name any file: relative naming and absolute naming. An absolute name, or absolute path as it is often called, specifies exactly where in the directory the particular file is. It tells the whole name of the file, starting at the root of the directory. An absolute name starts with /, the root, and names each directory along the path to the file, separating the directories with /. This is very similar to DOS, except for the direction of the slash and the fact that there is no disk drive designation. As an example, the absolute name for your mailbox might be /home/iris2/class9/mbox. The 'pwd' command always reports an absolute pathname.

Relative Naming

The second way to name a file in UNIX is with a relative name. Whereas an absolute name specifies exactly where in the directory a particular file exists, a relative name specifies how to get to it from your current directory. The look of a relative name may vary a lot. Depending on your starting directory there are a number of paths to a particular file.

In the simplest case, just naming a file in your current directory is a relative name. You are specifying how to get to this file from your current directory, and the path is to just open the file in the current directory.

When using relative paths, the special directories '.' and '..' that are contained in every directory are used frequently. Recall that '.' specifies the directory itself, and '..' specifies its parent. So, if the file mbox is contained in your current directory, naming the file with ./mbox and mbox are equivalent. The special directory '..' is used to name a directory at the same level in the tree as your current directory, that is, a sibling of your current directory. The following example illustrates using '..' to look at a sibling directory.

```
% pwd
/home/iris2/class9
% ls
NewBibligraphy.refer        bibs                mbox
ShortFile                   bin                 src
baby.1                      ig.discography      unix.refer
% cd bin
% ls
pwgen
% ls ../src
helloworld.c        pwgen.c
```

Short-cuts for File Naming

Since certain file naming patterns are used over and over, UNIX provides some short-cuts for file naming. Actually it is the shell that provides these, but the distinction is not critical at this point. In particular, many file accesses are either to your own home directory or to the home directory of another user. To make it easier to point to these places, the ~ character is used. Alone, ~ refers to

your home directory. So, the file ~/mbox refers to the file mbox in your home directory. Likewise, ~ username refers to the home directory of that particular user. So,
~dickson/mbox refers to the file mbox in the user dickson's home directory.

File Naming Limitations

Linux allows you great flexibility in naming your files. Many systems allow the name of individual files to be up to 256 characters, and the maximum absolute pathname to be about 1023 characters. This means that files can be given meaningful names quite easily. Also, since Linux is sensitive to the case of the file names, you can use mixed-case names to add clarity. While long, meaningful names can be quite nice, it is also possible to go overboard with file naming. So, try to give your files meaningful names, but try not to overburden yourself or others that have to access the files with overly long names. Since this environment works with AWS S3, file/directory naming should exclude special characters including spaces. Linux may be case sensitive, but to avoid file naming contention with AWS S3, it is recommended not to rely on case sensitivity to distinguish between two different files.

File Name Extensions

While Linux does not actually have the same notion of a file extension as is found in some other systems, notably DOS, many user and applications programs behave as it did. Linux does not consider the '.' character any differently than any other character in a file name. However, applications programs often use it to add an extension to a file name that specifies something about the contents of the file. These extensions sometimes tell what programs were used to create a file. If more than one was used, and is needed to decode the file, multiple extensions are used. Typically, the extension will tell what language the data in the file is in. A good example of this is the C compiler. It assumes that its input files will have an extension of '.c'. Its executable output generally has no extension, unlike DOS which uses an .EXE extension for executables. If we have a program called hello.c, its executable version is likely to be called just hello. For another example, a submitted SAS program (program.sas) will usually create a log and list: program.log and program.lst.

2.21 The Linux Directory Tree

Files in Linux are stored in directories, which are organized into a tree type configuration. At the top of the tree is the root, called /.

Locations in the directory tree are called paths and are specified as a list of directories separated by a '/'. Barney's home directory, for instance, is in the home directory, which is in the root directory, and therefore the path is "/home/med.ds.osd.mil/barney".

If you leave off the leading "/", the path is assumed to start in the directory you are currently in, for instance if you are in "home", "barney/songs" would be the directory songs contained in Barney's home directory. You can also use "..", which specifies the parent directory of the one you are in "../etc" from home would be the "/etc" directory.

There are a few standard directories you will find in pretty much all UNIX systems.

/bin This contains the basic system commands.

/etc	This contains system configuration files and programs used for administrating the system.
/lib	This contains the system libraries.
/tmp	This is used to store temporary files.
/usr/bin	This contains all the commands not in /bin.
/usr/man	This contains manual pages for programs.
/usr/local	This contains local programs, that were installed by your sysadmin and not included with the original system. In particular, /usr/local/bin is local binaries and /usr/local/man is local man pages.
/home	The location of this varies greatly from system to system, but somewhere on the system will be a place where all the home directories are, including yours.

2.22 Getting Around Directories

You have been given a directory in which you can store your personal files, called your *home directory*. When you log in to the system, this is where you start off. You probably don't have any directories to start off with. If you want to create a directory underneath your home directory, use the *mkdir* command. For instance, to create a directory to hold your programs, type '*mkdir programs*'. Please remember that '~/', your home directory, is not meant to store data, any files within your home directory may be removed without notice. This does not affect linked directories like '~/MIPFILE' or '~/mdrfile'.

Now you have a directory you can get into with the *cd* command. For instance, to get into your *programs* directory, type '*cd programs*'. To get back to your home directory from anywhere, simply type '*cd*'. To go up a directory, type '*cd ..*'.

If you want to copy or move a file into a directory, use the *cp* or *mv* command as normal, but instead of specifying a destination file, specify a directory. For instance, if you type '*cp helloworld.c programs*' from your home directory, it will copy the file *helloworld.c* into your *programs* directory.

To remove a directory, use the *rmdir* command. Note that the directory must be fully empty to do this. To remove your *programs* directory, change to that directory using the '*cd*' command and type '*rm **'. This deletes all files in that directory so make sure you are in the right place! You can find out where you are by typing '*pwd*'. Now that all files have been deleted, you can delete the directory. Go back to your home directory and type '*rmdir programs*'.

Appendix A: The NANO Editor (For editing in Linux, a clone of PICO) and VI editor

NANO Editor

(https://www.nano-editor.org/dist/v2.2/nano.html)

NANO is a line-mode text editor used to create files, reports and letters. You will need a line- mode editor if you use your account on a non-X Windows computer.

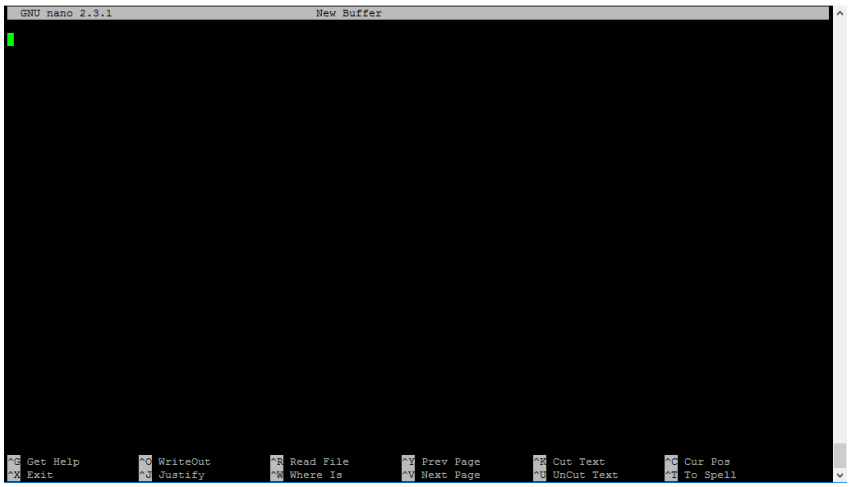
To start using NANO, enter 'nano filename' at the system prompt where filename is either the name of the file you want to edit (one that already exists) or the name of a new file you want to create. If you do not provide a file name, NANO will open a new file, but it will be unnamed until you save it.

If you prefer to use the "pico" command, you can set up an alias where "pico" refers to "nano":

```
alias pico='nano'
```

You can put this alias in the .bash_profile that lives in your home directory if you want it to be active every time you login.

The screen for a new file will look something like this:



In the list of commands at the bottom of the NANO screen, caret (^) stands for the control (ctrl) key. '^G' means to hold down the control key and press the letter 'G'. NANO commands are always a combination of the control key and another key. After you enter a NANO command you do not need to press the return key.

Quitting NANO

Press '^x' to quit or exit the NANO editor. NANO will prompt you with a message asking if you want to save your changes. 'Y' saves the changes and 'N' quits NANO without saving what you've done since the last time you saved.

Entering Text

At the system prompt, type 'nano' to create a new file. You should see NANO in your window. The top line of your screen is called the status line. It shows what version of NANO you are using, the file name (it will say New Buffer if it hasn't been named yet), and whether or not you have made changes (Modified will be in the upper right corner of the NANO screen if something has been entered but not saved). At the bottom of the screen are two lines showing the commands available for your use.

Now that you have an empty file, you can start typing. NANO automatically starts a new line when you run out of room (called "wrapping text," "text wrap," or "autowrap"), so you do not need to watch the screen while you type. Press [return] when you want to start a new line. If you want a blank line between lines of text or data, press [return] twice.

Saving the file

Now that you have a file with something in it, this is a good time to save your work. You should save your files often, if for no other reason than strange things can happen and sometimes do. The commonsense rule is to save often enough so that if you lose your work, you will not have much to re-do. Some people recommend saving every 15 minutes.

To save a file, use '^o', which stands for "write out" because you are writing what you have done to a file as output. When you press '^o', NANO will prompt you for the file name. If the file does not have a name yet, you will need to type one in. If the file does have a name, it will follow the file name prompt, and all you need to do is press [return].

You may also save the file during the process of quitting NANO. '^x' quits NANO after prompting you to save the file. Press 'Y' to save the file. Pressing 'N' will return the file to the state it was in the last time you saved it (if you had not saved the file during this editing session, the file will return to the state it was in when you opened it). If you alter a file and want to save the original version as well as the modified version, enter 'Y' to save the file. When you are prompted for a file name enter a new file name.

Editing in NANO

You can easily edit short files using only the following keys:

- [ret]: Inserts a new line in the file.
- cursor (arrow) keys: The four arrow keys move the cursor up, down, left, and right.
- backspace/delete key: Deletes the character preceding the cursor.

Cutting text

Cutting or deleting text is a major component of any kind of editing. NANO gives you two ways to cut text: (1) cut a character at a time, or (2) cut a line at a time (unfortunately, you cannot cut a word at a time).

To delete a character, place the cursor on the character you want to delete and press ^d. Pressing your keyboard's delete key (some keyboards label it backspace) will delete the character preceding the cursor.

To delete an entire line, place the cursor anywhere within the line and press '^k' (for kill or cut text). The entire line will be cut. If you change your mind about deleting the line before you delete another

one, you may press '^u' to bring back the line. The restored line will be placed wherever the cursor is.

Pasting text

Since '^u' lets you restore the cut line wherever the cursor is and not just to wherever it came from, you may use '^k' and '^u' to "cut and paste" text. In NANO, you move blocks of text similar to if you were using a word processing package. You first need to mark (highlight) the text you want to move. Cut the text, move the cursor to where you want the text inserted, then paste it. Position your cursor at the beginning of the first line of the text block you want to move. Press '^^' (ctrl-shift-6 key). Move the cursor to the end of the block you want to cut. The text should then be highlighted. Once the block is highlighted, press '^k' to cut the text. Place the cursor to where you want the text placed, and press '^u' to paste. The text should then appear.

Reformatting text

To justify or reformat a paragraph (fill in gaps left by starting a new line) press '^j'. NANO defines paragraphs as text surrounded by blank lines or indentation. However, if you had a blank line between each line of text pressing '^j' would not affect the text.

If you do not like the result of the justification, press '^u' to "unjustify" the paragraph before you move the cursor outside the paragraph. Once you have justified text, the '^u' command function changes from undelete to unjustify. However, when you move the cursor from the justified paragraph, the '^u' command function returns to undelete. Therefore, if you need to unjustify the paragraph, you must do so before you move the cursor outside the paragraph.

To join lines that are separated by blank lines use '^k' to delete the blank line and '^j' to fill the paragraph.

Searching a NANO file

If you are proofreading a file by looking at a printout, it may be tricky to find the specific line you want to revise. An easy way to get to the right point in your file is to use NANO's search, or 'where is' command, which is '^w'. This command lets you enter a character(s) or word(s) to search for, and then moves the cursor to the first occurrence. NANO begins the search from the point in the file where the cursor is located. So, if the cursor is at the beginning of the second paragraph, the search begins at the beginning of the second paragraph.

To find a specific character or word, press '^w' (for where is). A black strip with the word 'search:' will replace the two lines of commands at the bottom of the screen. You should then type the character(s) or word(s) you want to find. This search command is not case sensitive. This search command is also not limited to finding only whole words. For example, if you were to tell NANO to find "th", it would find "that", "the", "they", and "thick", etc.

Using the spellchecker

NANO has a spellchecker that will check spelling in your text files.

To start the spell checker, press '^t'. The NANO spellchecker compares the words in your file to words that are in its dictionary. When it finds a word that is not in its dictionary, it sends the prompt 'Edit a

replacement:' followed by the misspelled word. If the word is correct, just press return. If the word is misspelled, you may change it by backspacing over it (to delete it) and then typing the correct word. You may also use the cursor and delete keys to edit parts of the word.

Inserting a file into NANO

To insert an existing file into the NANO file you are currently editing:

- Locate the cursor where you want to place the inserted file
- Press '^r'. After you press '^r', you will receive the prompt 'Insert file:'
- Type in the name of the file you want to insert and press return. The inserted file's contents will be inserted at the cursor.

This command does not delete or in any way alter the file that is to be inserted (a copy of the file is inserted).

If you change your mind about inserting the file after you have pressed '^r', when prompted for the file name enter any character string that does not correspond with a file. The system will attempt to process the command, but when it does not find a corresponding file it will cancel the command and your text will not be altered. You will, however, receive the message 'No such file:' followed by the character string you entered.

Control keys for moving around in a file

If for some reason you cannot use cursor (arrow) keys, you may use the following commands to move around in your file.

- '^p' moves to the Previous line.
- '^n' moves to the Next line.
- '^f' moves Forward a character.
- '^b' moves Backward a character.

Clearing the display

If your display becomes cluttered (most likely because of interactive messages), you may clear or refresh it by entering '^l'. The messages are erased from your screen and the file you are working with remains.

Command list for NANO

The commands you can use in NANO are listed across the bottom of the screen. Remember, '^' symbolizes the control key. '^G' means to hold down the control key and press G.

^G	Help. More information on the commands.
^X	Exit NANO
^O	Save a NANO file
^J	Justify text within a paragraph. This command will fill in paragraphs that are missing, etc.
^R	Insert an existing file into NANO.
^W	Find a word or character string in the NANO file.
^Y	Go back to the previous page.
^V	Go to the next page.

^K Cuts the line of text where the cursor is located.
^U Uncuts the last text cut.
^C Gives you the exact location of your cursor within the file.
^T Run the spell checker.

When you have selected ^O to save a file or ^R to insert a file, you will be prompted for a file name. You will also be given the following new options:

^C Cancels the command.
^T Gives you a list of files in your home directory. If you select a file that already exists, NANO will ask you if you want to overwrite the file.

Basic VI Editor

Beginning Your Editing Session

To edit a file	vi [filename]
To recover an editing session	vi -r [filename]

Notes on vi commands and modal editing

All vi commands are entered in command mode. To enter command mode, press the ESC key. Some vi commands cause vi to enter another mode. For example, the i (insert command) causes vi to enter insert mode after which all keystrokes are inserted as text. To return to command mode from insert mode, press the ESC key. The :set showmode command will cause vi to display the current editing mode in the lower right corner of the editing screen.

Controlling the Screen Display of Your Session

Repaint the current screen	{^l}
Display line #, # of lines	{^}

Moving the Cursor

Beginning of current line	0 or ^
Beginning of first screen line	H
Beginning of last screen line	L
Beginning of middle screen line	M
Down one line	j, {return}, +
End of current line	$
Left one character	h, {ctrl-h}
Left to beginning of word	b, B
Right one character	l, {space}
Right to end of word	e, E
Right to beginning of word	w, W
Up one line	k, -
Beginning of next sentence	)
Beginning of previous sentence	(

Paging Through Text

Back one screen	{^}
Down half a screen	{^} \
Down one screen	{^f}
Forward to end of file	G

Move cursor to specified line	line no.G
Up half a screen	{^j}

Special Pattern Characters

Beginning of line	^
End of line	$
Any character except newline	.
Any number of the preceding character	*
Any set of characters (except newline)	*

Creating Text

Append text after cursor	a
Append text after end of line	A
Insert text before cursor	i
Insert text at beginning of line	I
Open new line after current line	o
Open new line before current line	O
Take next character literally (i.e. control characters...) and display it	{^v}

Ending Your Editing Sessions

Quit (no changes made)	:q
Quit and save changes	ZZ, :wq
Quit and discard changes	:q!

Using ex Commands From Within vi

Copy specified lines	:co, t
Display line numbers	:set nu
Disable display of line numbers	:set
nonu Move lines after specified line	:m
Read file in after specified line	:r
filename Review current editor options	:set
Review editor options	:set all
Set new editor option	:set option
Write changes to original file	:w
Write to specified file	:w filename
Force write to a file	:w! filename

Some Useful ex commands for use in vi

Some useful set options for your ~/.exrc file:

:set all	Display all Set options
:set autoindent	Auto magically indent following lines to the indentation of previous line.
:set ignorecase	Ignore case during pattern matching.

:set list	Show special characters in the file.
:set number	Display line numbers.
:set shiftwidth=n	Width for shifting operators << and >>
:set showmode	Display mode when in Insert, Append, or Replace mode.
:set wrapmargin=n	Set right margin 80-n for autowrapping lines (inserting newlines). 0 turns it off.

Searching Through Text

Backward for pattern	?pattern
Forward for pattern	/pattern
Repeat previous search	n
Reverse direction of previous search	N
Show *all* lines containing pattern	beg,endg/pattern/p

Example
:1,$g/compiler/p Will print all lines with the pattern compiler.

Substitute patt2 for all patt1 found.	:beg,ends/patt1/patt2/g

Example
:%s/notfound/found/g Will change all occurrences of ‘notfound’ to ‘found’.

Appendix B: MIP SCE Quick Reference Guide

***Note: There is no password requirement for the MIP SCE, the new hostname is MIPSAPSASP02

***Reminder: Do not store programs, scripts, data, formats, etc. under your user home directory /home/med.ds.osd.mil/$USER/. If this space becomes full, all users will be locked out of the SCE.
Users can set up their profile in MIP SCE? (e.g. For aliases and environmental variables)
Use NANO or VI to edit: '~/.bash_profile'

To submit a SAS job:
Log on to the server using a SSH client and change to the directory where the SAS program is located in your directory space.

Example:
[mipsapsasp02}:/home/med.ds.osd.mil/$USER > cd /hpae2/

Note: Do not store code or data in your home directory, specifically /home/med.ds.osd.mil/$USER. Files in this area may be deleted without warning.

Passing a SAS argument sas <program_name> -<variable_name> <value>

Commands for submitting SAS jobs

To start a job running in the default workspace: **sas *jobname.sas***

Current status of running jobs: **jobQueue –list**

Job Queue Help: **jobQueue –help**

Current status of running jobs only for the user: **jobQueue –list –user $USER**

Cancel a submitted job, before job starts running: **jobQueue –remove <jobPID>**

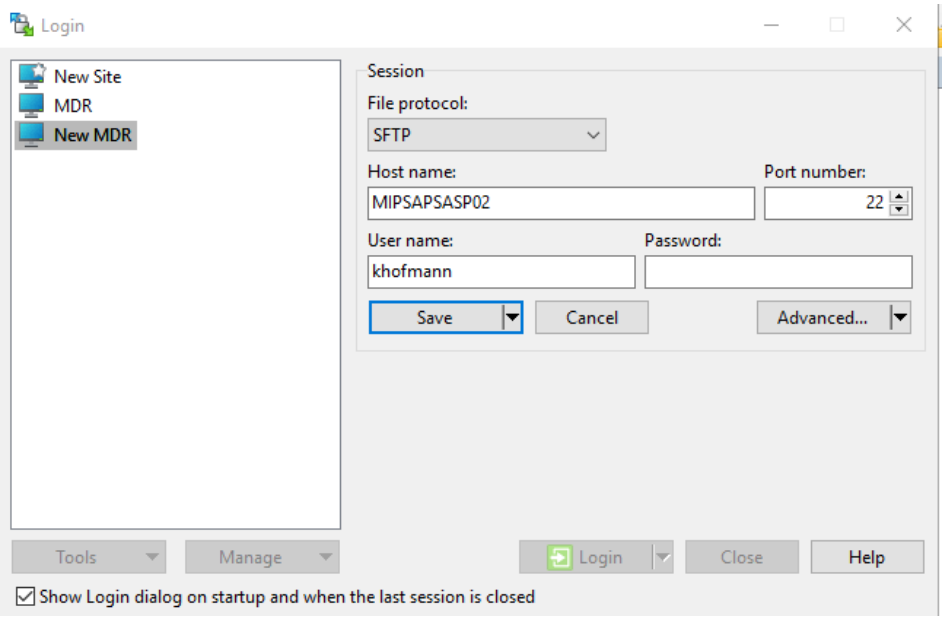
Cancel a submitted job, when job is running: **jobQueue –stop <progPID>**

Monitoring job queue alternative: **ps -ef | grep -i sasexe | grep -i *username* | grep -v grep**

Scheduling SAS Jobs
One time execution: sas /dir/file.sas --datetimeExec '2023-01-01 23:59:59'
Recurring execution: Edit cron: **crontab -e**
Write the entry to crontab **:w**
Quit crontab **:q**
Crontab Man: https://linux.die.net/man/5/crontab

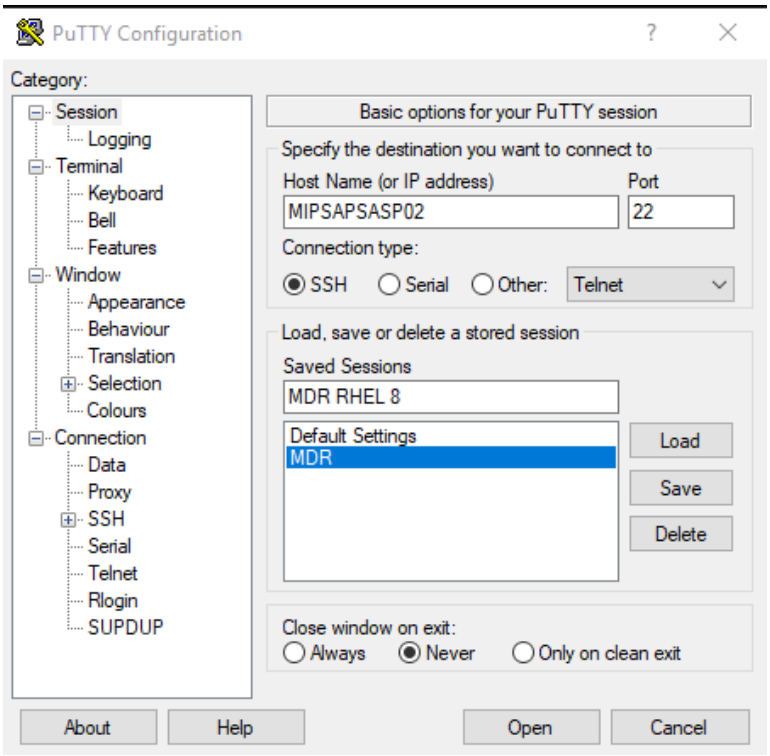
Appendix C: AVHE MIP AWS – WinSCP Login to MIPSAPSASP02

WinSCP - Setup Connection to MIPS



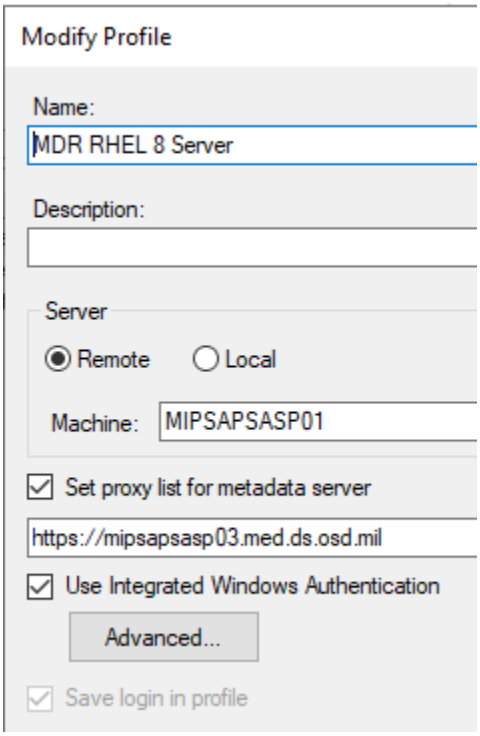
Appendix D: Launch PuTTY SSH Session to MIPSAPSASP02

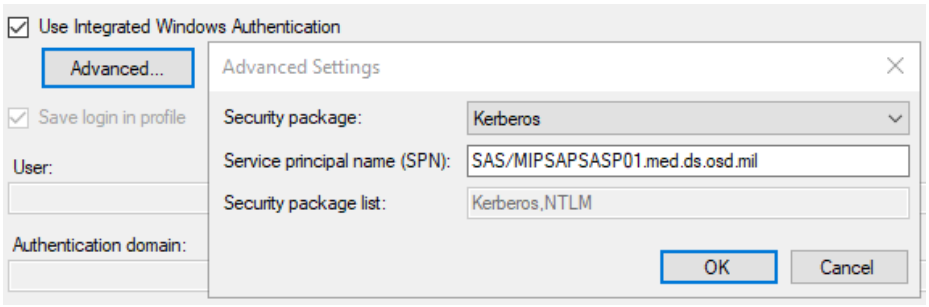
SAS Launch PuTTY SSH Session to MIPS



Appendix E: SAS Create Enterprise Guide Profile to MIPSAPSASP01

SAS Create Enterprise Guide Pro


Modify Profile
Name:
MDR RHEL 8 Server
Description:
Server
Remote
Local
Machine:
MIPSAPSASP01
Set proxy list for metadata server
https://mipsapsasp03.med.ds.osd.mil
Use Integrated Windows Authentication
Advanced...
Save login in profile


Use Integrated Windows Authentication
Advanced...
Save login in profile
User:
Authentication domain:
Advanced Settings
Security package:
Kerberos
Service principal name (SPN):
SAS/MIPSAPSASP01.med.ds.osd.mil
Security package list:
Kerberos,NTLM
OK
Cancel

Appendix F: Connecting to SQL Databases

The MDR/SCE can connect to many databases in the MIP, including SQL-based environments like Redshift and SQL Farm. This greatly enhances the power of working in the MDR/MIP by enabling more data connections and being able to leverage powerful tools available throughout the MIP.

Generally, when these connections are made, SAS passes code to the Database Management System (DBMS) and data is returned to SAS. You must have permissions/access to the DBMS you are trying to access (i.e. for SAS to hit Redshift, you need to have access to Redshift). Please follow these tips to ensure the most efficient performance of queries across systems and to ensure stability of the MDR/SCE platform:

1) When in Rome: As much as possible, try to use SAS within the SAS environment and SQL within the SQL environment.
2) Minimize data traveling through the pipe: Filter the data in the SQL database as much as possible before returning it to SAS. Don't use "select *". Limit to only the variables/columns that you need. Use where statements to limit the number of rows you need.
3) Avoid heterogenous joins: **Never** merge/join a SAS table to a SQL table directly. This will require copying the *entire* SQL table into the SAS environment to perform the merge. Imagine you want to enhance data from the MDR GENESIS Encounter file (SAS data set) with some data from the raw BDE 3.0 Clinical Event table (SQL table in genesis_vw in Redshift). If you try to join them directly, since they live in different environments, the *entire* Clinical Event table (**tens of billions of rows**) will be copied into the SAS environment unnecessarily. Instead, pass the minimum amount of data between the two systems and perform the join within one system or the other. Here are two potential ways to do this:
    a. Filter the SQL table as much as possible before bringing it back to the SAS environment. Do a SAS merge between the original SAS data set and the results from this SAS/SQL step. For the example above, perhaps you are only interested in a certain lab test within a certain time frame. You can limit the Clinical Events table by event_cd and event_end_dt_tm before bringing it into SAS and potentially reduce the number of rows returned through the pipe by 90% or more.
    b. Pass a finder file from SAS into the SQL environment. Use an inner join in SQL to filter the SQL table to that finder file and return the results to the SAS environment. For the example above, perhaps you only need clinical events associated with a certain type of encounter, like same day surgeries performed overseas. By filtering Clinical Events to those particular encounters (by encntr_id), you can greatly reduce the amount of data moving through the pipe.

For more on connection options from SAS, go to:
https://carepoint.health.mil/sites/MIP/mugteam/SitePages/SAS%20EG%20Training.aspx

2.1 F.1: Pulling Data from Redshift with Explicit Pass-Through SQL

```
PROC SQL ;
  CONNECT TO redshift (DATABASE='prodautocreateddb'
  SERVER='mipsdbrdsp01.med.ds.osd.mil'
  AUTHDOMAIN=Redshift_HSDW_AES);
  CREATE TABLE WORK.adm_type AS
  SELECT * FROM CONNECTION TO redshift
    (
      /*Redshift SQL query here*/
select code_set, code_value, display
from genesis_vw.code_value
where code_set = 3
    ) ;
QUIT;
```

Notes:

- Connection details go within the proc sql statement.
- The AUTHDOMAIN used depends on your particular account (your Role Based Access Controls aka RBACs) and on the schema/table you are trying to hit. See here for some potential authdomain options: Location of authdomain options
- The results of the SQL query will be fed into the SAS data set you specify in the create table statement.
- Tip: Debug your Redshift SQL using a SQL tool (like DBeaver) before trying to run it from SAS.
- To use explicit pass-through, you must know the SQL language of the DBMS you are trying to hit (Redshift SQL in this case)
- Using explicit pass-through gives you full control over how the SQL code is generated in the DBMS.

2.2 F.2: Pulling Data from Redshift with Implicit Pass-Through SQL

```
libname bde redshift
      server='mipsdbrdsp01.med.ds.osd.mil'
      db='prodautocreateddb'
      authdomain=Redshift_HSDW_AES
      schema=genesis_vw;
*Option 1: Pull data from the "bde" library using proc sql;
proc sql;
      create table admit_type as
      select code_set, code_value, display
      from bde.code_value
      where code_set = 3;
quit;
*Option 2: Pull data from the "bde" library using a SAS data step;
data admit_type;
      set bde.code_value(keep = code_set code_value display);
      where code_set = 3;
run;
```

Notes:

- Connection details go within the libname statement.
- The AUTHDOMAIN used depends on your particular account (your Role Based Access Controls aka RBACs) and on the schema/table you are trying to hit. See here for some potential authdomain options: Location of authdoman options.
- The results of the SQL query will be fed into the SAS data set you specify in the create table statement in proc sql or in the data statement in a data step.
- SAS/ACCESS and the DBMS will convert your code into the appropriate SQL language to run in the DBMS.
- Using implicit pass-through means you don't need to know the particular SQL language of the DBMS and may require less coding.

2.3 F.3: Pushing a Temp Table from SAS to Redshift Using Bulkload

```
options nocenter ls=132 ps=32000 obs=max validmemname=extend;

*Build a list of people from a SAS data set;
proc sort data = out.deers(keep = edipn) out = people nodupkey;
      by edipn;
      where edipn ne '';
run;

*Push to Redshift as a temp table;
libname rstemp redshift server='mipsdbrdsp01.med.ds.osd.mil' database='prodautocreateddb'
      authdomain=redshift_hsdw_aes insertbuff=32000 dbmstemp=yes connection=global;

*Use SAS data step to push data into Redshift using the bulkload option;
data rstemp.'#cohort'n(bulkload=yes bl_use_manifest=yes
          bl_bucket='awprmips-dl-operations/sas-redshift' bl_region='us-gov-west-1'
          bl_aws_config_file='/scefile/system/redshift/redshiftConfig'
          bl_aws_credentials_file='/scefile/system/redshift redshiftCred');
      set people;
run;

*Pass-through SQL to pull DaVINCI OMOP Person IDs in Redshift;
proc sql;
      connect to redshift (database='prodautocreateddb' readbuff=32000 connection=global
            server='mipsdbrdsp01.med.ds.osd.mil' authdomain=redshift_hsdw_aes);
      create table work.omop as
      select * from connection to redshift
      (
            select p.person_id, p.person_source_value
                  from davinci_omop_v5.person p
                  inner join #cohort col on p.person_source_value = col.edipn
        );
quit;
```

Notes:

- Need validmemname = extend to use the pound sign in a SAS data set name (#cohort), which flags it as a temp table in Redshift
- The **connection details** between the *bulkload* push from SAS to Redshift must match the **connection details** on the pull from Redshift EXACTLY.
- Using the *bulkload* option results in much faster processing for pushing data from SAS to Redshift.
- The inner join to the temp table (#cohort) serves as a filter on the SQL query so that only rows that match the SAS data set are returned through the pipe.

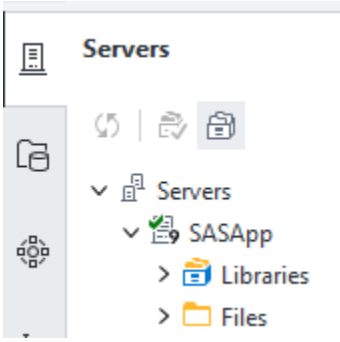
- Notice how this process avoids a heterogenous join. Rather than merging the SAS data set directly to the SQL table, the SAS table is used as a finder file, pushed into Redshift as a temp table, where the join is performed within SQL between SQL tables. The filtered set of results is then returned to SAS.

2.4 F.4: Establishing a Kerberos Ticket to Enable the Connection from SAS to SQL Farm

You must have an active Kerberos ticket for the connection between SAS and SQL Farm to function. To ensure you have an active Kerberos ticket, login to SAS EG and make sure the MDR is the active profile (look in the top right corner):



Click the arrow next to SASApp to establish the connection (left side under servers, notice the green check mark):



Once the Kerberos ticket is established, you can run your SAS code from either SAS EG or PuTTY.

2.5 F.5: Pulling Data from SQL Farm with Explicit Pass-Through SQL

```
PROC SQL ;
  CONNECT TO sqlsvr (dsn="SFM-PRDDWSP02-M2_NETWORK");
  CREATE TABLE WORK.tedi AS
  SELECT * FROM CONNECTION TO sqlsvr
    (
      /*T-SQL query here*/
      select fy, sum(admission_count) as adm
      from M2_NETWORK.dbo.Network_Inpatient
      group by fy
      order by fy
    ) ;
QUIT;
```

Notes:

- Connection details go within the proc sql statement.
- Only select data connections (dsn) have been established. Note that the dsn needs the server name as well as the database name. See here for some potential dsn options: https://carepoint.health.mil/sites/MIP/Documents/SAS%20SQLFARM%20Libnames.txt. If you need a connection from SAS to SQL Farm that does not yet exist, file a GSC ticket (see the Getting Help section).
- The results of the SQL query will be fed into the SAS data set you specify in the create table statement.
- Tip: Debug your T-SQL using a SQL tool (like DBeaver or SQL Management Studio) before trying to run it from SAS.
- To use explicit pass-through, you must know the SQL language of the DBMS you are trying to hit (T-SQL in this case)
- Using explicit pass-through gives you full control over how the SQL code is generated in the DBMS.

2.6 F.6: Pulling Data from SQL Farm with Implicit Pass-Through SQL

```
libname m2sql sqlsvr
      dsn="SFM-PRDDWSP02-M2_NETWORK"
      schema=dbo;
*Option 1: Pull data from the "m2sql" library using proc sql;
proc sql;
      create table tedi as
      select fy, sum(admission_count) as freq
      from m2sql.Network_Inpatient
      group by fy
      order by fy;
quit;
*Option 2: Pull data from the "m2sql" library using a SAS proc summary;
proc summary data = m2sql.Network_Inpatient nway;
      class fy;
      var admission_count;
      output out = tedi sum=;
run;
```

Notes:

- Connection details go within the libname statement.
- Only select data connections (dsn) have been established. Note that the dsn needs the server name as well as the database name. See here for some potential dsn options: Location of some potential dsn options.. If you need a connection from SAS to SQL Farm that does not yet exist, file a GSC ticket (see the Getting Help section).
- The results of the SQL query will be fed into the SAS data set you specify in the create table statement in proc sql or in the output statement in the proc summary (in this example).
- SAS/ACCESS and the DBMS will convert your code into the appropriate SQL language to run in the DBMS.
- Using implicit pass-through means you don't need to know the particular SQL language of the DBMS and may require less coding.

2.7 F.7: Pushing Data from SAS to SQL Farm

```
LIBNAME PROPIOD SQLSVR  READ_ISOLATION_LEVEL=RU  INSERTBUFF=10000 DBCOMMIT=10000
      DATASRC="SFM-PRDDWSP02-PRJ_OPIOID"
      SCHEMA=dbo ;

data propiod.naloxone_qpp_registry;
      set opioids;
run;
```

Notes:

- Connection details go within the libname statement.
- Only select data connections (dsn) have been established. Note that the dsn needs the server name as well as the database name. See here for some potential dsn options: Note that the dsn needs the server name as well as the database name. Review the [MDR Database DSN Configuration Reference Guide] for some potential options.. If you need a connection from SAS to SQL Farm that does not yet exist, file a GSC ticket (see the Getting Help section).
- SAS/ACCESS and the DBMS will convert your code into the appropriate SQL language to run in the DBMS. In this example, the SAS data step will create a table in the SQL Farm