
FRE Version History

Amy Langenhorst <Amy.Langenhorst@noaa.gov>

This document describes the version history of FRE, the FMS Runtime Environment. For usage information, see <http://www.gfdl.noaa.gov/fms/fre>.

Table of Contents

1. Introduction	2
2. Currently Recommended Versions	2
2.1. FRE Version 2 Recommended Versions	2
2.2. FRE Version 3 Recommended Versions	2
3. fremake Version History	2
3.1. New in fremake023	3
3.2. New in fremake01	4
3.3. Future Work	4
4. frerun Version History	4
4.1. New in frerun024_64_vb	4
4.2. New in frerun024_64	4
4.3. New in frerun024_99	5
4.4. New in frerun024_grid	5
4.5. New in frerun024	5
4.6. New in frerun023_cpio	5
4.7. New in frerun023	6
4.8. New in frerun022	6
4.9. New in frerun021	6
4.10. New in frerun02	6
4.11. New in frerun01	7
4.12. New in rtsrun14	7
4.13. New in rtsrun13	7
4.14. New in rtsrun11, rtsrun12	7
4.15. New in rtsrun10	7
4.16. New in rtsrun9	8
4.17. New in rtsrun8	8
4.18. Future Work	8
5. frepp Version History	9
5.1. New in frepp024_64_vb	9
5.2. New in frepp024_64	9
5.3. New in frepp024_99	9
5.4. New in frepp024_grid	9
5.5. New in frepp024	9
5.6. New in frepp023_ann	10
5.7. New in frepp023_v4	10
5.8. New in frepp023_timavg	10
5.9. New in frepp023	11
5.10. New in frepp022	11
5.11. New in frepp021	12
5.12. New in frepp02	12
5.13. New in frepp01	13
5.14. New in rtspp14	14
5.15. New in rtspp13	14
5.16. New in rtspp12	14

5.17. New in rtsp11	15
5.18. New in rtsp10	15
5.19. New in rtsp9	15
5.20. New in rtsp8	15
5.21. Future Work	16
6. Other Future Work	16

1. Introduction

The following programs all reside in `/home/fms/bin`. You can tell which version is currently default by executing

```
ls -l /home/fms/bin/frerun
```

The command **frerun** is a unix link that changes over time, but the file that the link points to is a frozen version which will not change over time. When a new version of FRE is available for use, an email will be sent to oar.gfdl.forrestal@noaa.gov.

2. Currently Recommended Versions

2.1. FRE Version 2 Recommended Versions

`/home/fms/bin/frerun024_64_vb` is the recommended version, unless that version is incompatible with your models. If so, please let Amy know as soon as possible. The only known problem is with the land model and distributed restarts, where more testing is required.

The current default version is `frerun024_99`, which you will get if you call `/home/fms/bin/frerun`. If you do not need any of the newer features, please use this version.

2.2. FRE Version 3 Recommended Versions

```
/home/fms/bin/fresetup -r fre_20080701 [your chosen FREROOT]
```

This version corresponds to `/home/fms/bin/frerun024_64_vb`, with the following changes:

1. bugfix to making a cpio after combining multi-pe restart files, for regression runs only (no analogous bug exists in version2 of FRE)
2. modified to send shorter email on production crash
3. bugfix on dmget of initCond file

Utilities for FRE version 2 will be released with the FMS Perth release that contain analogous changes.

3. fremake Version History

3.1. New in fremake023

1. Allows mkmf template to be specified in the xml itself.

```
<mkmfTemplate file=""/>
or
<mkmfTemplate>
  FC = f90
  CPPFLAGS = -macro_expand
  etc
</mkmfTemplate>
```

This can go in the <setup> tag, or the setup tag can be overridden if you specify the mkmfTemplate in the <compile> section for your experiment.



Warning

This requires a change in syntax from previous specifications of the mkmf template. One-word mkmf templates will cause fremake023 to print an error message, ie,

```
<mkmfTemplate> /home/fms/bin/mkmf.template.ia64_flt </mkmfTemplate>
```

yields

```
ERROR: mkmfTemplate '/home/fms/bin/mkmf.template.ia64_flt' looks like a file
ERROR: You must use the file attribute: <mkmfTemplate file=''>
```

you would need to change it to

```
<mkmfTemplate file="/home/fms/bin/mkmf.template.ia64_flt"/>
```

Added March 2006:

1. When your machine is upgraded to the new OS, CVS will be updated, and it will no longer generate the path_names file automatically. However, if you use FRE for checking out code from the repository, fremake will generate your path_names file for you automatically, just after it executes your CVS commands. If you use fremake for acquiring and compiling code, the CVS upgrade should be invisible to you.
2. In your xml, you can now use two <mkmfTemplate> tags:

```
<mkmfTemplate file="/home/vb/fms/mkmf.template.ia64_flt"/>
<mkmfTemplate type="debug" file="/home/vb/fms/mkmf.debugtemplate.ia64_flt"/>
```

Then if you use the -t option, fremake will always use the "debug" mkmf template if one is listed in your xml. If one is not listed, it'll use the system default (/home/fms/bin/mkmf.debugtemplate.\$platform). If you don't use -t, it'll use the non-debug template.

3. fremake will put a comment line at the top of any mkmf template it acquires from an external file, ie

```
#from /home/vb/fms/mkmf.template.ia64flt
```

3.2. New in fremake01

1. new version of mkmf allows for calling mkmf from exec dir instead of src dir. fremake uses this to allow you to compile the optimized and debug versions at the same time.
2. new option -p to call list_paths only once, after all the cvs stuff is complete. This is much faster. This works for all experiments except mom4 ones that treat the mom4 paths differently than the shared code paths.

3.3. Future Work

- change behavior of frerun with respect to compiling. Add option (-c?) which always submits a compile script. If the make is successful, the compile shell script should submit the run scripts requested in the frerun command. The 'make' should not be done in the run shell script. (If it is, errors can occur when two runscripts try to 'make' at once.) This would make Balaji happy (he would always use '-c') and satisfy those who never want compiling to be done in the runscript.
- if someone types control-c during cvs checkout, how to know to warn user that checkout didn't get completed?
- currently script decides whether to run cvs commands based on existence of \$expt/src directory. should also test for empty \$expt/src directory.

4. frerun Version History

4.1. New in frerun024_64_vb

1. Supports Bernie Sieber's history data staging. To try it, you must use the following in your xml file:
<postProcess combine="staged">
2. Calls frepp024_64_vb

4.2. New in frerun024_64

1. does not combine restart files; therefore this version supports distributed restart files. NOTE: This may be incompatible with current versions of the land model. More testing is required with the land model.
2. defaults to do mppnccombine on history files in a separate job in between the production run and postprocessing jobs. This is changed from the previous default of combining history files in production jobs. To combine history files in the production jobs in this version, use >postProcess combine="online">
3. creates new 64-bit netcdf format to support large data files
4. uses /home/fms/bin/mppnccombine-2.1.4_ia64 (with -64 option)

5. uses /home/fms/bin/get_fms_data_cxfscp
6. calls frepp024_64

4.3. New in frerun024_99

1. uses "exit 99" for checkpoints to improve checkpointing behavior
2. formats all namelists in scripts with a newline before the end '/'
3. uses qsub soft resources to tell the operators job statistics
4. Calls frepp024_99

4.4. New in frerun024_grid

1. ability to extend a run with -e
2. you can provide a cpio file as a grid spec and frerun knows that it should be extracted rather than just copied to the \$work dir.
3. Calls frepp024_grid

4.5. New in frerun024

1. New option -l: frerun passes the argument on to "qsub -l". Use this to submit to a certain node, or any other qsub -l resource, ie, frerun -l ic1 ...
2. System monitoring commands added: times the mpirun command, calls /home/gcs/bin/cpusetmeminfo before each mpirun command, to print a list of free and used memory on the nodes of the current cpuset.
3. Checkpointing functionality added: If /home/gfdl/flags/fre.checkpoint.* have been created by OPS, the runscript will resubmit itself and exit. The runscript checks for these flags before every mpirun command. Mail will be sent to the user when a job has been checkpointed.
4. Saves a copy of the xml file into each restart cpio file.
5. Uses Hans's newest /home/fms/bin/mppnccombine-2.1.1_ia64
6. Archives a copy of the batch stdout for jobs that failed for the Technical Services group's reference
7. Calls frepp024.

4.6. New in frerun023_cpio

1. If you are running an experiment using the FMS coupler, frerun will look to make sure your fieldTable includes sphum for both land and atmos. If not, it will add them.
2. Adds an extra return character before the '/' ending a namelist. This allows Fortran comments on the last line of a namelist specification in the xml.
3. Cprios data files directly to and from archive with a buffer size of 512K, instead of creating the cpio in \$TMPDIR and using cp/mv to move the file to archive. This should result in a speed improvement, because cp has a small buffer size on the altix platform which slows down file transfers.
4. Calls frepp023_timavg

4.7. New in frerun023

1. calls frepp023

Added March 2006:

1. Now overwrites the fv_core_nml and land_model_nml's layout parameters correctly if you specify the atmos_layout and land_layout attributes in your regression test xml.
2. Fills in the \$baseDate variable correctly if used in namelists
3. It is now possible to create a runscript with neither MPI nor SHMEM.

4.8. New in frerun022

1. creates a runscript which will perform mppnccombine inline in the job script rather than calling frepp to do the combine. This will save one round of queue wait times for the model postprocessing. The 'raw' (uncombined) history files will not be saved to archive.
2. since mppnccombine is online, never call frepp for regression test runs
3. calls frepp022 for postprocessing
4. writes to ~/.rhosts file if necessary for file permissions
5. small improvements to mppnccombine and dmgets in the runscript

4.9. New in frerun021

1. calls frepp021

4.10. New in frerun02

1. new cshell postInit option. There are four categories of csh you can use within the <input> tag:

```
<csh>
#cshell to be run from INPUT/ at the beginning of each job
</csh>
<csh type="always">
#cshell to be run from INPUT/ before each mpirun command
</csh>
<csh type="init">
#cshell to be run from INPUT/ before the very first mpirun command of the
#first job of the experiment only
</csh>
<csh type="postInit">
#cshell to be run from INPUT/ before each mpirun command EXCEPT the very
#first mpirun command of the experiment
</csh>
```

4.11. New in frerun01

1. namelists in runscript will be sorted alphabetically
2. new attribute zetac_layout is processed
3. modified subroutine getexecutable to always check whether an inherited experiment should have its own executable or not
4. bugfix to production runs of less than 1 year; will now set an appropriate combineFreq for these runs
5. new option for reading all the namelists for an experiment and outputting all namelists alphabetically as text in xml tags.
6. bugfix to procedure of finding the debug executable for inherited experiments
7. changed so that project tag gets placed in regression runs as well as production runs

4.12. New in rtsrun14

1. ability to run on altix platform

4.13. New in rtsrun13

1. prints a warning if make_exchange_reproduce is false for regression tests, and prints a warning if make_exchange_reproduce is true for production runs.
2. added code to combine multithreaded write restart files
3. added code to change priority projects to dev projects when submitting a windf job

4.14. New in rtsrun11, rtsrun12

1. just call the corresponding version of the postprocessing program

4.15. New in rtsrun10

1. better wording on namelist warnings
2. supports platform csh for changing version of compiler. To use the new version of the compiler, use xml as follows:

```
<setup>
  <target platform="sgi">
    <csh>
      source /opt/modules/modules/init/csh
      module switch mpt mpt_1900
      module switch mipspro mipspro_741m
    </csh>
  </target>
</setup>
```

and remove the old <targetPlatform>sgi</targetPlatform>.

3. multiple types of input csh supported

```
<input>
  <csh> this is executed in the section which copies input files at the beginning
        of each JOB, same as previously
  </csh>
  <csh type="init"> this is executed in the section which copies input files ONLY ONCE
                    PER EXPERIMENT, only at the initial run time t=0
  </csh>
  <csh type="always"> this is executed in the mpirun loop so that it is executed
                      BEFORE EVERY MPIRUN command
  </csh>
```

4.16. New in rtsrun9

1. developmental changes, calls rtsp9

4.17. New in rtsrun8

1. Doesn't attempt to compile an executable if you specified an existing one in your xml
2. can now set project for jobs with rtspriority

4.18. Future Work

- add totalview support
- frerun should read and edit queue_commands to ensure that jobs do not run out of queue allocations inappropriately, which is a common complaint.
- handle hourly runs (several)
- ability to run no-calendar models (which only use days) in production mode (which only accepts months and years now) (rjw)
- error checking that diagTable corresponds to pp requested
- handle experiments starting with a number (add a letter to the batch job name or something) (several)
- long regression type runs / being able to specify runsPerScript and runsPerExpt (ens)
- for development (LMDT, rsh)
 - if an interactive run, write the ascii output files to a directory in /home/ and do not cpio them.
 - skip or always overwrite the reload_commands file.
 - delete raw cpio
 - wipeftmp at top of runscript

- handle multiple <fmsDataSets> tags (pcm)
- create rhosts file automatically
- add init/postInit options to input ascii tables
- User-level checkpointing. This is a way to get FMS runs to stop gracefully at the end of a "segment" if the system requires.

5. frepp Version History

5.1. New in frepp024_64_vb

1. uses fregrid instead of Mike Herzog's tool for converting cube sphere grids to lat-lon.

5.2. New in frepp024_64

1. supports (and always creates) 64-bit netcdf format files
2. jobs use 4 processors. If jobs are creating a 20-year or longer time series or time average, for any component, they use 8 processors.
3. if the time-averaging program receives an out-of-memory kill error, the command will be retried. If you are still getting a lot of post-processing errors with this version, and you see a lot of results from "grep Killed \$root/scripts/postProcess/stdoutfile", please let Amy know.

5.3. New in frepp024_99

1. updated analysis figure code from Andrew Wittenberg and Fanrong Zeng

5.4. New in frepp024_grid

1. supports cube sphere grid with Michael Herzog's interpolation tool.

5.5. New in frepp024

1. New copy utility, cxfscp, to copy data from archive to TMPDIR for processing. This is about 2x to 3x faster than previous methods of cpio from archive or using "cp", and has resulted in a 25% speedup in most post-processing tests. Performance benefits will vary with system load
2. New external utilities:
 - a. uses v024 version of FRE modules (minor cleanup mods)

- b. uses v8 version of ia64 executables (recompiled under SLES10 with ifort.9.1.040, some evidence on small test cases of reduced variability with these versions.) The timavg and plevel utilities have been statically linked with the newest recommended version of the Intel compiler; therefore, the script does not need to load any modules. It does a module purge, which results in a shorter search path.
 - c. uses Fanrong's newest Analysis.pm, which allows providing script arguments to analysis scripts. Supports Andrew Wittenberg's new analysis scripts.
3. Correct histDir is passed to analysis scripts that need it (bugfix to case where you want to create figures from someone else's postprocessed files and still require the history data.)
 4. Improved method for calculating required vars for converting to pressure levels.
 5. Added checkpointing with saving status to a file, similar to reload_commands for model runscripts. The script writes its status to \$ppRootDir/checkpoint/ as it progresses through a component, and deletes that file when the job is finished. At the beginning of the job script, it checks for the existence of that file and goes to the indicated position in the script.
 6. New option -l: limit static variable processing to the diagnostic source file listed in the postprocessing tags of your xml. Otherwise, frepp will try, for example, to put all atmos-related static variables into post-processing component static files with a "type" string starting with "atmos".
 7. Doesn't attempt to make the "pp/" directory level with frepp -A.
 8. Gives an error message when the user-specified pp directory does not exist when frepp is called with -A (analysis-only.)
 9. Has less verbose stdout messages. To obtain more verbose output, use the "frepp -v" option.

5.6. New in frepp023_ann

1. can generate an annual time series directly from history files (without the intermediate step of monthly data) when the diagnostic file involved has "year" as the third column in the diag table file specification.
3. improved method of determining which variables are required for interpolation to pressure levels
4. increased cpio buffer size on altix
5. the altix platform is the default platform for post-processing. To force post-processing on the SGI Origins with this version of frepp, you need to specify the post-processing platform as "sgi" as follows:

```

<postProcess platform="sgi">
...

```

All versions of frepp023* will be changed to this setting on 10/24/2006.

5.7. New in frepp023_v4

1. Utilities used by FRE have been recompiled for the Altix. They have -fltconsistency and all the rest of the default fms compiler options. This should result in increased performance.

5.8. New in frepp023_timavg

1. Bugfix on labeling the time axis of 100yr timeAverages (or any timeAverage of timeAveraged files, where

the number of files being averaged is odd, as in the 100yr case.)

2. Let jobs covering 1 year of postprocessing have 2 hours of cpu time instead of 1 hour, as hi-res models need this extra time.
3. Additional atmospheric plevel interpolation preset: narcaap
4. Can perform interpolation to pressure levels when temp, sphum, and zsurf are missing (just can't calculate divv, rvort, hght, slp without them.)
5. If this job has no history data files, do not run a blank "dmget \$historyfiles" line.

5.9. New in frepp023

Statistics show that frepp023 is a ~20% reduction in wallclock time on average, but may use a little more cpu time than the default version.

1. new -R option to force regeneration of analysis figures
2. use FFIO layer for all netcdf utilities EXCEPT ncks for best efficiency
3. use Remik's suggestions for NETCDF_FFIOSPEC for best efficiency
4. use Bernie's mvfile utilities for safe and preallocating file moves from \$TMPDIR to /archive
5. use nrcat instead of nccatm for speed
6. submit timeAverages in a separate job from timeSeries if component is ocean and there will be calculates on the 100 year boundary or greater
7. if the component is ocean and there will be calculations of more than 20 years of data, request 30 cpu hours
8. changed order of postprocessing so that the timeAverages are first and the seasonal timeSeries is last

5.10. New in frepp022

1. can submit jobs to both irix and altix. Postprocessing will still be submitted to the AC's by default, as previous versions of frepp have done. If you want your job to submit postprocessing to the IC's instead, you will need to add an XML attribute to your postProcess tag as follows:

```
<postProcess platform="ia64">
...
```

Here are the conditions to consider when choosing a platform for your postprocessing:

- Postprocessed files on the two platforms will not be bitwise identical for any files where averaging or z-level interpolation has been performed. Only switch back and forth between postprocessing platforms if you don't need bitwise reproducibility in your postprocessing files.
 - If the AC queue wait times are small, there is no reason to switch platforms. If the AC queue wait times are long, you may be able to get your postprocessing jobs through the queue faster on the IC's.
2. added -m option from lol to allow creation of timeseries of runs less than one year which may or may not start on Jan. 1.
 3. added maxyrs, a global variable telling the max chunkLength or interval to be calculated in this pp job
 4. uses maxyrs to make the requested cpu time 1 hour if not creating any pp files containing more than one year of data

5. small changes and fixes to Analysis.pm in /home/fms/local/perl/v022/FRE: finds correct history directory when you specified it on the frepp command line, knows how to call rsh-acarch to make directories in / archive from linux, more organized stdout print statements
6. only test for the correct number of time levels in a file if the calendar is noleap
7. use nccatm rather than nrcat for concatenating netcdf files
8. using the options -As (for figures only) will not submit a dummy postprocessing script.

5.11. New in frepp021

1. Added support for segment length of 2, 3, and 4 months (in addition to 1, 6, and 12 which existed already.) frepp still does not perform postprocessing for runs with segments less than one month or greater than one year.
2. Preserves history files in \$TMPDIR for the duration of a pp job
3. Additional dmgets to avoid krecalls
4. Prefer cpios to individual files when assembling timeseries from other timeseries (saves dmgets)
5. In creating a seasonal timeseries by concatenating shorter seasonal timeseries, check that each individual file exists before nrcat, else send error message and skip the variable
6. Instead of doing seasonal ts at same interval as monthly ts and leaving individual seasons in archive until the seasonal chunk length has been reached, wait until least common multiple of monthly and seasonal chunks to do the seasonal ts, and only move chunked files to archive. (writes fewer intermediate files to archive)

5.12. New in frepp02

1. changes to improve the metadata load on the system, buffer size for cpio increased from 64 to 256k, added "dmput -r" throughout
2. supports interpolating to the vertical levels in the atmospheric ERA40 data set. To get them, use

```
<component type="atmos" zInterp="era40" source="atmos_month">
```

3. new zInterp option "zgrid" which calls Bob Hallberg's Resample_on_Z program. Usage is analogous to the atmospheric resampling on pressure levels. Sample component:

```
<component type="ocean_zgrid" source="ocean_month" zInterp="zgrid">
  <timeSeries freq="monthly" chunkLength="5yr">
    <variables> temp salt age u v </variables>
  </timeSeries>
  <timeSeries freq="monthly" chunkLength="20yr"/>
  <timeSeries freq="monthly" chunkLength="100yr"/>
  <timeSeries freq="annual" chunkLength="5yr">
    <variables> temp salt age u v </variables>
  </timeSeries>
  <timeSeries freq="annual" chunkLength="20yr"/>
  <timeSeries freq="annual" chunkLength="100yr"/>
  <timeSeries freq="seasonal" chunkLength="5yr">
    <variables> temp salt age u v </variables>
```

```

</timeSeries>
<timeSeries freq="seasonal" chunkLength="20yr"/>
<timeSeries freq="seasonal" chunkLength="100yr"/>
<timeAverage source="monthly" interval="5yr">
  <variables> temp salt age u v </variables>
</timeAverage>
<timeAverage source="monthly" interval="20yr"/>
<timeAverage source="monthly" interval="100yr"/>
<timeAverage source="annual" interval="1yr">
  <variables> temp salt age u v </variables>
</timeAverage>
<timeAverage source="annual" interval="5yr"/>
<timeAverage source="annual" interval="20yr"/>
<timeAverage source="annual" interval="100yr"/>
</component>

```

If you do not explicitly state the variables as indicated above, then the following variables will be processed by default: temp,salt,age,u,v. Note that usually one doesn't need to specify

```
<timeAverage source="annual" interval="1yr">
```

at all in the xml; it's just done automatically. And it is not required for ocean_zgrid either, however, if you leave it out frepp will give you a message that you didn't specify any variables for annual 1 year averages and it will use the default temp,salt,age,u,v.

5.13. New in frepp01

1. now using Fanrong's Analysis.pm package for figures
2. modified not to cpio files when there is only one file
3. prevent trying to run platform csh intended for the altix on the sgi platform
4. added the ability to average daily data to a monthly timeseries. To take advantage of this capability, you need to add three lines to your xml:

```

<component type="atmos" zInterp="ncep" start="0001" source="atmos_month">
  <timeSeries freq="daily" source="atmos_daily" chunkLength="5yr"/>
  <timeSeries freq="monthly" chunkLength="5yr"/>
  ...
  <timeSeries freq="monthly" averageOf="daily" chunkLength="5yr">
    <variables> t_ref_min t_ref_max </variables>
  </timeSeries>
</component>

```

The new feature is 'timeSeries freq="monthly" averageOf="daily"'. There are several things to note:

- a. You can either calculate this in the same pp job with the 'timeSeries freq="daily"' or do it anytime after the daily timeSeries has been created. (ie, you can run just this piece of postprocessing offline.)
- b. The chunkLength requested for the averaged file must be the same as the chunkLength of the daily timeSeries file, ie, 5yrs. Longer chunks of pp should pick up the new variables and create 20yr timeSeries from the 5yr ones.

- c. This is specifically designed with the noleap calendar in mind and currently does not check the model calendar.
- d. The variables tag is required for this type of timeSeries calculation.
5. gory bugfixy details surrounding making cpio files, calculating one or two variables offline after other variables have been processed, and performing frescrub at any point along the way.
6. the variables tag is now available for all timeSeries options.
7. checks that timeSeries contain the right number of time levels as it calculates them. If the timeSeries file contains the wrong number of levels, you'll get email about it.

5.14. New in rtsp14

1. ability to run on altix platform
2. analysis figure improvements

5.15. New in rtsp13

1. analysis figure bugfix when using <variables> tag inside of <timeSeries>; previously rtsp could mistakenly assume that pp files were missing and not create figures for that data
2. new timeSeries and timeAverage attribute 'from' can override what chunklength or interval pp you want to calculate a timeSeries or timeAverage from, ie

```
<timeSeries freq="monthly" chunkLength="200yr" from="100yr"/>
```

3. annual multi-year averages will only try to calculate the interval they're supposed to calculate. Previously, they looked for the last existing output file and calculated all the intervals after that through the -t year. This change is only useful when you are trying to run offline pp out of order.
4. seasonal averages that need the previous december now know how to get it from the history file as well as from the previously existing monthly timeAverage file. This change is only useful when you are trying to run offline pp out of order.
5. better error message when you try to run seasonal timeSeries without any available monthly timeSeries data
6. -m option: If you give the -m option, you will get an email when your job is complete. Only useful for offline pp. Doesn't send email if option -p is used; this way you only get email when it's done with the stream
7. can put <variables> inside seasonal ts and it will use only those variables. Previously you could only use <variables> in the monthly ts, and those variables would then get used for seasonal and annual ts also. Note that you still need to have the monthly ts for all variables you request a seasonal ts of.
8. minor cleanup: seasonal ts tests whether cpio has already been extracted

5.16. New in rtsp12

1. BUGFIX: timavg.csh was using the wrong weighting interval in some cases
2. truncation of long filenames on timavg'ed files and plevel data avoided

3. bugfix for postprocessing jobs run after an experiment has been frescrubbed
4. dmget improvements to avoid the 'line too long' error

5.17. New in rtsp11

1. doesn't call nc_to_midtime, should improve speed of postprocessing by 10-25%
2. consolidates error notification into fewer actual emails
3. doublechecks that output directories exist so that there are no missing directories
4. additional dmgets implemented
5. static files created with splitncvars, should retain all appropriate attributes
6. time_bounds bugfix to seasonal averages when the first december of the run needs to be recalculated from the end of the first year
7. new analysis figure options:

```

-Y year      = specify a four digit year as your analysis's start year, ex
              This year overrides the startYear specified in the <analysis>
-Z year      = specify a four digit year as your analysis's end year, ex -Z
              This year overrides the endYear specified in the <analysis> t
    
```

5.18. New in rtsp10

1. analysis figures fixes
2. time_bounds fixes
3. bugfixes and improvements to external utilities: list_ncvars, timavg, plevel

5.19. New in rtsp9

1. developmental changes for time_bounds fixes

5.20. New in rtsp8

1. added option to split postprocessing into different jobs: -c component
 - without -c, behavior is same as rtsp6
 - keyword '-c split' means each component will be managed in a separate postprocessing stream; ie, one job will be submitted for each component.
 - using a component name, ie '-c ocean', means that only the postprocessing for that component will be done.
2. nrcat now uses -O option to overwrite files if they already exist. Previously if you were regenerating pp after a system error, you would get emails about nrcat errors. nrcat complained because the output file already existed. This change will prevent those error messages.

3. includes seasonal average functionality.
4. uses ncatted to put the calendar attribute back after plevel.sh called
5. filename attribute corrected in some annual average files
6. writes some version information to ppRootDir/.frepp_history

5.21. Future Work

- perform pp on runs less than 1 year (gamdt)
- perform pp on data not starting jan 1. (gamdt)
- perform pp when segment length (therefore history files) > 1 year (mom, lm3)
- interpolation for ocean data analogous to plevel conversion in atmosphere (bls)
- analysis issues, ie -A incompatible with -p
- do 280 yr timeseries = 100 + 100 + 4x80yrs (cm2)
- option to combine after each segment on production runs (smg)
- automate secondary pp (200yrs plus) (cm2)
- error checking that diagTable corresponds to pp requested
- reduce frepp's metadata load on system
- maintain history file in TMPDIR throughout job so there is no danger of it being dmput when it is still needed.
- add automatic scrub capability
- implement xyInterp: create a subarea with a mask, call a fortran program to do subsampling in xy.
- do not submit pp jobs that won't do anything. check for existence of static file from frepp itself. don't calculate annual averages on an annual basis.
- when calculating seasonal timeSeries from monthly, cpios should be used when individual files don't exist

6. Other Future Work

- alter design enough to allow various pieces to be taken from various places: postprocess someone else's history files, freanalyze two different experiments in two different xml files
- recognize \$root and \$name in namelist file paths
- handle variables in rts.xml
- site specific stuff, porting
- fredb: integrate with model development database
- managing multiple executables in FRE
- frecheck
 - add a timing flag to frecheck to just output the timings of the runs and not do the check
 - enhance performance information printed to frecheck's report

- use chksums of restart files for comparison rather than resdiff
- talk with Koushik about what can be done about incorporating vtune/speedshop
- automate checking of diagnostic history output (currently only restarts are checked)
- freppcheck and frescrub should make use of dmbadf to report files on bad tapes
- update FRE xml schema