

## 20 Using Multiple Machines with WES

These instructions describe a way to run a WES simulation on one “server” machine and connect other machines as clients to the simulation. This can be useful for pairing up forecasters during simulation training. In this configuration, one machine runs the simulation (`start_simulator` and `start_awips`), and the clients just run D2D using `start_awips`. While the following “manual” instructions work, they are intended for users that are relatively comfortable with a basic WES understanding and running simulations. Future versions of WES will have improved plug and play support for running simulations on multiple machines and synchronizing clocks, etc. If you are interested in running WES in a classroom or laboratory environment, contact [wes@infolist.nws.noaa.gov](mailto:wes@infolist.nws.noaa.gov) for more information.

To configure multiple machines to work during a simulation, one machine is a server which runs the simulation and AWIPS decoders. All machines have D2D clients fed from that server. Every machine must have the same WES version installed from the WES install DVD (e.g. WES8.1).

### 20.1 Setup

---

1. Choose a machine to be a server, and install WES8.1 if it isn't already installed.

e.g. “bobcat” will be the server (your server name will likely be different)

e.g. `install-wes81.sh /usr1` on bobcat (if it isn't already installed)

---

**Note:** If you installed WES in the `/usr1` directory, and you do not have `/data` as a directory, then the client data case would be stored in `/usr1/data/awips`.

---

2. Determine the location of the case to mount from the server. One suggestion is to start with the WES test case to make sure everything is set up correctly.

e.g. `cd /data/awips` on bobcat

e.g. `pwd`

---

**Note:** In this example the `pwd` yields `/usr1/data/awips`.

---

3. Choose other machines as client machines, and install WES8.1 on them from the release DVD (do not copy from the server machine). These machines must have different names.

e.g. “wolf” will be a client machine

e.g. `install-wes81.sh /usr1/client` on wolf

---

**Note:** If you installed WES in the `/usr1/client` directory, and you do not have `/data` as a directory, then `/data` is a link that points to `/usr1/client/data`. This means your client data cases would be stored in `/usr1/client/data/awips`.

---

4. After installing WES8.1 on the client machine, change the `FXA_WARNGEN_PRODUCT_ID` variable on the client machine in `/awips/fxa/.environs.<$machinename>` to make it different from the server machine and any of the other client machines.

e.g. on wolf change `${FXA_LOCAL_SITE}WRKW4` to `${FXA_LOCAL_SITE}WRKW5` in `/awips/fxa/.environs.wolf`

---

**Note:** If you do not have a `/awips/fxa/.environs.<$machinename>` file, then you can “`cp /awips/fxa/.environs.localhost /awips/fxa/.environs.<$machinename>`” where `<$machinename>` is the result of “`hostname | cut -d . -f 1`”

---

5. Create the server target directory from step 2 on the client machine in preparation for exporting the case storage directory from the server.

e.g. `mkdir /usr1/data/awips` on wolf (from target in step 2)

---

**Note:** If this directory already exists on the client machine and has contents in it, move any files and directories out of the way (e.g. `mv /usr1/data/awips /usr1/client/cases`) before making the directory. If these contents are WES data cases, these cases can be visible on the client machine by creating links inside `/data/awips` that point to the new case locations (e.g. `ln -s /usr1/client/cases/2002Feb10 /data/awips/2002Feb10`).

---

---

**Note:** If `/data` is a directory (not a link) on the server machine, consider creating a new and different directory in this step (e.g. `mkdir /usr1/servermount`) used to mount the server's data to the client machine in the following steps.

---

6. The system administrator must export the server's directory to all client machines. The clients should auto-mount the exported data directory on system start-up. In the following example, `/usr1/data/awips` on the server, bobcat, is exported to wolf. The mounting instructions below are based on an example from Ken Cook at the ICT WFO

---

**Note: Be careful...**if you haven't done this before, please have your IT do this or you could seriously mess up your machine

---

### On the Server Machine

- From the KDE Desktop Menu choose, **System Settings, Server Settings, Services**. Make sure the **NFS** box is checked.
- As **root**, edit `/etc/exports` and add the following entry:
  - o `your_export_dir ip.address.of.client(rw,no_root_squash,no_all_squash)`
  - o e.g. `/usr1/data/awips 129.15.59.61(rw,no_root_squash,no_all_squash)` where `129.15.59.61` is the client's (e.g. wolf's) ip address.
- As **root** edit `/etc/hosts.allow` and add the following entry:
  - o `ALL: ip.address.of.client`
  - o e.g. `ALL: 129.15.59.61` where `129.15.59.61` is the client's (e.g. wolf's) ip address.
- As **root**, run `exportfs -a` to export the file system
- As **root**, run `exportfs` to check if the file system is listed (i.e. exported correctly)
- As **root**, stop and start the NFS server:
  - o `/etc/rc.d/init.d/nfs stop`
  - o `/etc/rc.d/init.d/nfs start`

## On the Client Machine

- As **root**, edit `/etc/fstab` and add the following entry:
  - o `Server_name:your_export_dir client_dir nfs rw,auto,soft 0 0`
  - o e.g. `bobcat:/usr1/data/awips /usr1/data/awips nfs rw,auto,soft 0 0` where `your_export_dir` is the same exported directory as in the “On the Server Machine” section above, and `client_dir` is the mounted directory on the client machine.

---

**Note:** If `/data` is a directory (not a link) on the server machine, and `/data/awips` is exported to the client machine, then the `your_export_dir` would be `/data/awips` and the `client_dir` would be something like `/usr1/servermount` (see step 5).

---

- As **root**, mount the server: `mount -a`
- 7. Create a symbolic link for the case under `/data/awips` on all machines if it doesn't exist.

e.g. `ln -s /usr1/data/awips/2006Aug24test /data/awips/2006Aug24test` on wolf

---

**Note:** In this example the link already existed on bobcat so nothing was required on bobcat. If the link didn't exist on bobcat, the command is required on bobcat, too.

**Note:** If `/data/awips` was linked to `/usr1/servermount` (Notes in steps 5 and 6), then link `/usr1/servermount/2006Aug24test` to `/data/awips/2006Aug24test` on wolf.

---

- 8. Copy an `/awips/fxa/data/localization/nationalData/ipc.config` file from the server machine into the `localizationDataSets/XXX` directory for the case.

e.g. on bobcat `cp /awips/fxa/data/localization/nationalData/ipc.config /data/awips/2006Aug24test/localizationDataSets/ABR`

---

**Note:** Use the nationalData version of the `ipc.config` file. There are multiple versions of these files in AWIPS, change from build to build. If this file changes in future AWIPS builds, update the `ipc.config` file when creating a new localization.

---

9. Edit the `ipc.config` in the case's `localizationDataSets/XXX` directory, and replace all "localhost" entries with the server name (or ip address). Then save the file.

e.g. replace "localhost" with "bobcat" in  
`/data/awips/2006Aug24test/localizationDataSets/ABR/ipc.config`  
file and save

---

**Note:** In vi this can be done using `:g/localhost/s/localhost/bobcat/` followed by `:wq!`.

**Note:** Once the `ipc.config` file in the case has been hardwired for the server, neither the `start_simulator` nor the `enhanced_case_review` applications will work for this case on the client machine. After deleting this file in the `localizationDataSets/XXX` directory, the client can run `start_simulator` or `enhanced_case_review` because AWIPS defaults to the localhost version in the `awips` directories.

**Note:** Always try to run your simulations on the machine physically containing the data, or the increased disk I/O across machines will significantly slow case preparation down.

---

## 20.2 Testing

---

10. Verify the data for the case is visible from all machines. Don't run `start_simulator` to run a simulation yet.

e.g. `start_awips`, select **2006Aug24test**, and view 0.5 Z/SRM8 radar product

11. Run a simulation on the server machine.

e.g. `start_simulator` on bobcat

12. Verify that the simulation runs correctly on the server machine

e.g. `start_awips` with the Text Workstation Control box checked on bobcat, verify that an all-tilts radar display updates, create a test warning

13. If the server machine works correctly, verify that the all-tilts display updates on the client machine. Do not create a warning on the client machine yet.

e.g. `"start_awips"` on wolf, check the all-tilts display for updates

---

**Note:** At this point, the clock on the client machine isn't set back to the simulation time, so the D2D time is the current date and time, rather than the simulation time. This situation causes WarnGen on the client machine to create "future warnings" situations and prevents the server from correctly generating warnings. The following steps illustrate how to adjust the clock on the client machine.

---

14. Shut down D2D on the client machine.

e.g. `exit D2D` on wolf

15. Determine the current time on the server machine.

e.g. the `date -u` command on bobcat should yield something like Thu May 01 23:46:14 UTC 1997

16. Set the clock back on the client machine to be relatively close to the server (the date format is MMDDHHmmYY).

e.g. `/awips/fxa/DRT/bin/date -u 0501234697`

17. Start D2D on the client machine, and verify data displays update correctly and WarnGen works.

e.g. on wolf, `start_awips` and check the Text Workstation Control box, verify the all-tilts display refreshes and create a test warning in a separate D2D pane.

18. Exit the test simulation and manually set the clock back to the current time on the client machine

e.g. `"date -u"` on bobcat, yields something like Thu Jan 05 15:32:47 UTC 2006

e.g. `"/awips/fxa/DRT/bin/date -u 0105153206"` on wolf

This case is now ready for running a simulation.

## 20.3 Starting a Simulation with D2D Clients on Multiple Machines

Running a simulation with multiple clients requires first completing the “Setup” and “Testing” above. This example shows how to start and stop a simulation after successfully exporting and mounting the case data.

In this example the server (bobcat) holds the case data and runs a simulation using the 2006Aug24test case located in `/usr1/data/awips`. D2D runs on the server machine, “bobcat”. D2D also runs on the client machine, “wolf”.

1. Start the simulation on the server machine.

e.g. `start_simulator` on bobcat

2. After the simulation has been started on the server machine, manually set the clock on the client machine to match the server machine.

e.g. on bobcat `date -u` to find the current simulation time

e.g. on wolf `/awips/fixa/DRT/bin/date -u 0501234697` to synchronize the clock back

3. Start up D2D on the server machine.

e.g. `start_awips` with the Text Workstation Control box checked on bobcat

4. Start D2D on the client machine.

e.g. `start_awips` with the Text Workstation Control box checked on wolf

5. If the simulation is paused and resumed on the server machine, the clock must be manually resynchronized (see step 2) on the client machine

6. When the simulation ends, set the clock back to the current time on the client machine.

e.g. on bobcat `date -u` to find the current time

e.g. on wolf `/awips/fixa/DRT/bin/date -u 0105153206` restores wolf to the current real time