Project #7: Using `rdist` to Distribute Files

- Distributing Copies of Important Files
- ***`rdist`***
- How ***`rdist`*** Works
- Running ***`rdist`***
- Be Careful When Configuring ***`rdist`***

Distributing Copies of Important Files

- System administrators often have to distribute files to many machines at once
- These files may be text files or scripts or programs
- These files need to be constantly updated
- For example, every server needs an ***`/etc/profile`*** file that contains commands that are run after a user logs in
- Keeping these files up-to-date on a network is a routine task for a system administrator

Distributing Copies of Important Files

- Though ***scp*** can be used to copy individual files, this is not a practical solution for even a moderately sized network
- Fortunately, Linux provides a utility to automate this process

rdist

- ***rdist*** stands for remote distribution
- It is used to distribute copies of files from one machine to many other machines
- ***rdist*** makes identical copies of these files preserving the owner, group, and other permissions and modification times of the files
- ***rdist*** reads the files it is to distribute and the machines to which it needs to copy the files from a configuration text file

rdist

- By default, the name of this text file is ***Distfile***
- By default, ***rdist*** uses ***rsh*** to move files
- ***rsh*** is like ***ssh***, but it is not secure
- We have to tell ***rdist*** to use ***ssh*** instead

How rdist Works

- To use ***rdist***, you set-up one machine as the "master"
- The master contains the current version of all files you need to distribute
- The files can be the contents of a directory and all its sub-directories
- Once the system administrator has made modifications to the master, he or she uses ***rdist*** to distribute these changes to the "slave" machines it is maintaining
- ***rdist*** is a utility program you run from the command line

How rdist Works

- When you run ***rdist***, it uses either ***rsh*** or ***ssh*** to run the ***rdistd*** daemon on the remote machines
- Once the `rdistd` **daemon** is running on the slave machines, the ***rdist*** utility asks it to check the files being synchronized
- Whenever the copy of the file on the slave is different from the copy on the master, the master copy is sent to the slave machine
- The files can differ in size, modification date, or permissions
- This `process` continues until all files are updated, at which point ***rdist*** tells the daemons to quit after which it too quits

Running rdist

- When you run ***rdist*** on the command line, you tell it
  - To run ***ssh*** instead of ***rsh***
  - Where the ***rdistd*** daemon is located on the slave machine
  - What distribution file to use
- The 3 pieces of information are all **absolute pathnames**
- You can get this information using the **which** command
- We tell ***rdist*** the path to ***rdistd*** using the ***-p*** (lower case) option
- We tell ***rdist*** the path to ***ssh*** using the ***-P*** (upper case) option

Running rdist

- We give ***rdist*** the pathname of the distribution file as an argument on the command line
- To run ***rdist*** with ***ssh*** using the distribution file ***~/Distfile*** you run

```
rdist -p $(which rdistd) -P $(which ssh) ~/Distfile
```

- The ***$( )*** construction denotes **command substitution**
- In command substitution, a command is run in a sub-shell and the output of that command replaces what is between the parentheses on the command line
- Writing all this out every time we wanted to run ***rdist*** would be tedious, so we will create a *shell script* to do this for us

Distribution Files

- An entry in a distribution file has the following format

```
LABEL: SOURCE -> TARGET_MACHINE
install TARGET_DIRECTORY ;
```

- LABEL is a convenience and is not strictly necessary
- It will allow you to specify only certain files to be updated on the command line and not all the files being synchronized
- SOURCE an absolute path of a file or directory on the master machine
- TARGET_MACHINE is the hostname of the machine

Distribution Files

- If there is more than one entry in SOURCE, it must be enclosed in parentheses, ( )
- The same is true for TARGET_MACHINE
- To create a distribution file update **/etc/hosts** on **itvm25-1a** and **itvm25-2a**, we would write

```
hosts:   /etc/hosts  -> ( itvm25-5  itvm25-7 )
install  /etc/hosts  ;
```

- Notice the semi-colon ; on the second line of the entry

Be Careful When Configuring rdist

- If you are not careful, using **rdist** can cause major problems
- **rdist** blindly follows the instructions in the distribution file
- If that file tells **rdist** to overwrite some import file, it could make **all** machines on the network inoperable!
- Consider the following distribution file:

```
hosts: /etc/hosts -> ( itvm29-6 )
install /etc ;
```

- This file tells **rdist** to remove the entire **/etc** directory from all slave machines and replace it with the file **hosts**

Be Careful When Configuring rdist

- Since **`/etc`** contains configuration information for many services this change will keep all those services from running properly
- To test your distribution files without actually changing anything use the following option to the "install" instruction

  `install -overify`

- Running **`rdist`** with a distribution file using this option will show you what will happen -- but will not copy any files