IBM

InfoSphere DataStage

Information Server DataStage engine running as a non-root user

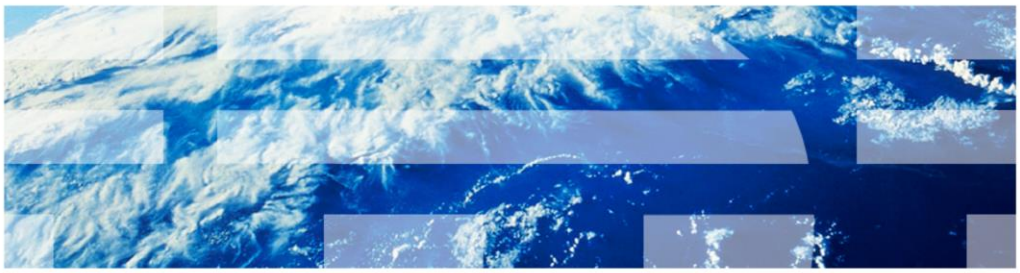

This presentation will discuss how to run the DataStage® engine as a non-root user for all versions of InfoSphere® DataStage.

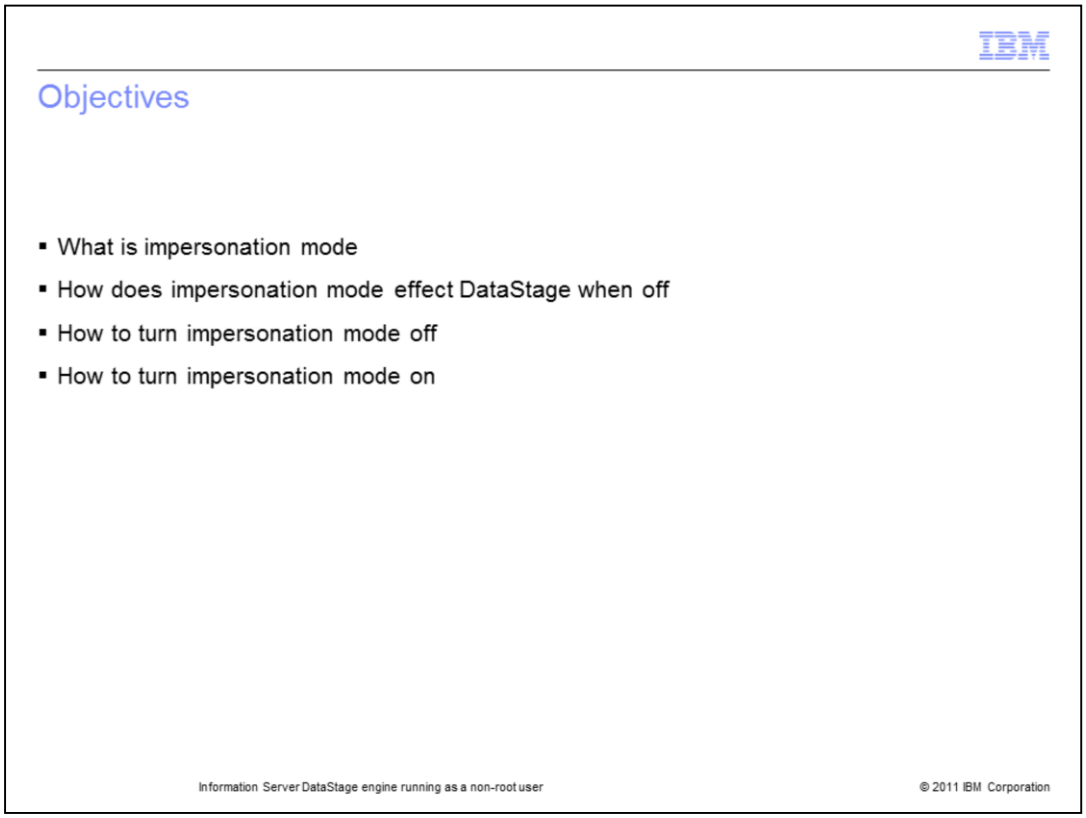

The objectives of this presentation are to discuss what impersonation mode is, and what the differences are in your environment when turning impersonation mode off. This presentation also describes how to turn impersonation mode off and how to turn it back on if needed.

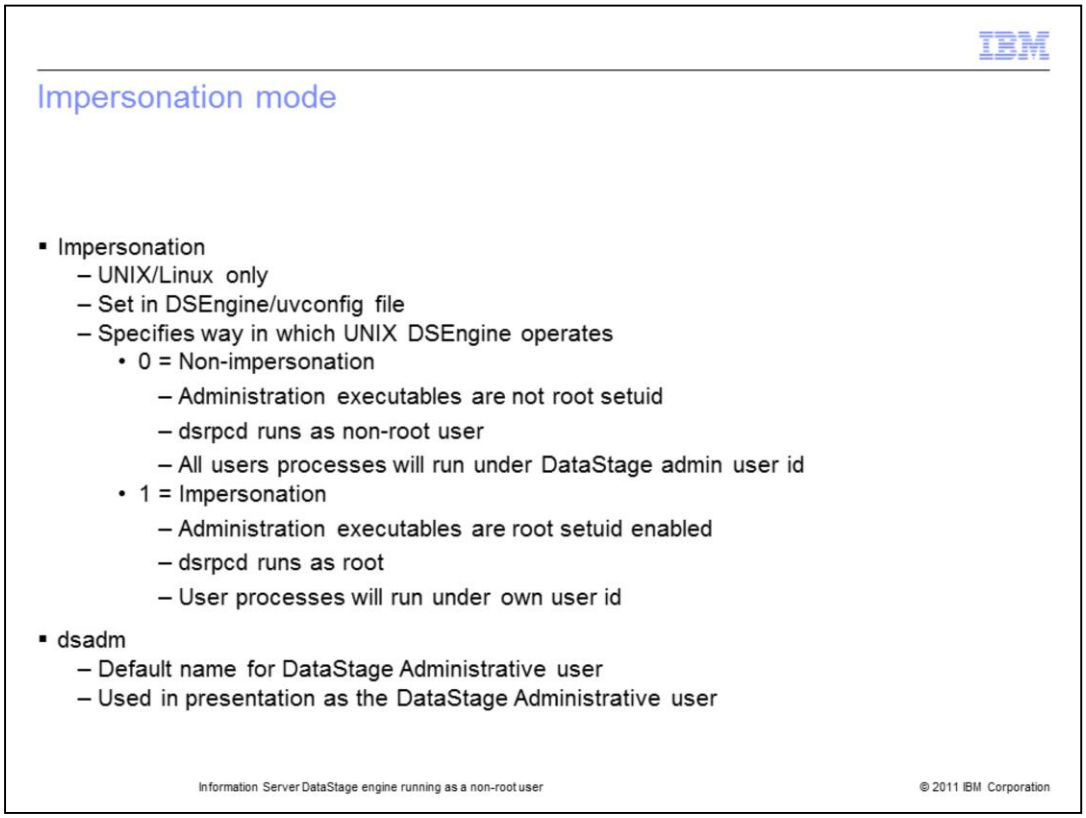

Impersonation mode specifies the way in which the DSEngine operates under UNIX and Linux. Impersonation mode is not applicable to Windows. When this value is set to 0, the engine will run in 'non-impersonation' mode, administration executables are not root setuid, and the dsrpcd will run as a non-root user. All users logging in will execute using the permissions and identification of the owner of the dsrpc connection daemon and will have DataStage administration privileges.

When this value is set to 1, which is the default, the engine will run in 'Impersonation' mode, administration executables are root setuid enabled, and the dsrpcd will run as root. All users logging in will execute using their own permissions and identification. Only the designated administrator and root will have DataStage administration privileges.

The default user ID for the DataStage administrative user is dsadm. For simplicity, the remainder of this presentation will use “dsadm” to represent the DataStage admin user. The name may vary on your server.

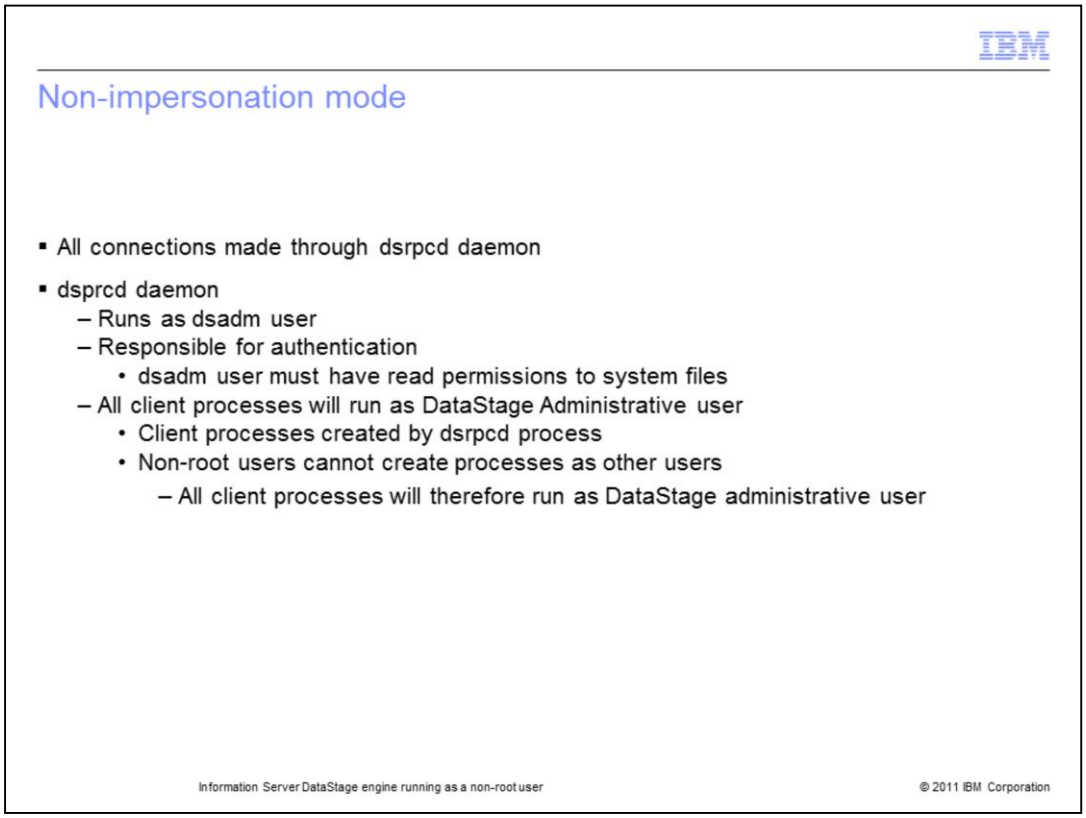

Before turning off impersonation mode, there are a few points that are important to be aware of. When impersonation mode is turned off, the dsrpcd process runs as the dsadm user instead of running as root. The dsrpcd is responsible for authentication of the DataStage user. Since this process is now running as dsadm and not root, permissions on certain system files will need to be changed to allow the dsadm user to complete the authentication. The next slide displays the exact files that need changing for each platform.

Another thing to understand is that when DataStage establishes client connections by way of the gui clients or when a job runs, the client connection is established through the dsrpcd daemon process. The dsrpcd will clone itself to create the client processes. On UNIX and Linux, when a non-root user creates a new process, the process is owned by that user. A non-root user cannot start a process up as a different user, only root is allowed to do this. When the clients connect or a job runs, all of the processes will therefore be owned by dsadm and not by the user that started the gui client or started the job. Since the processes are running as dsadm, the processes will have all the operating system privileges that the dsadm user has.

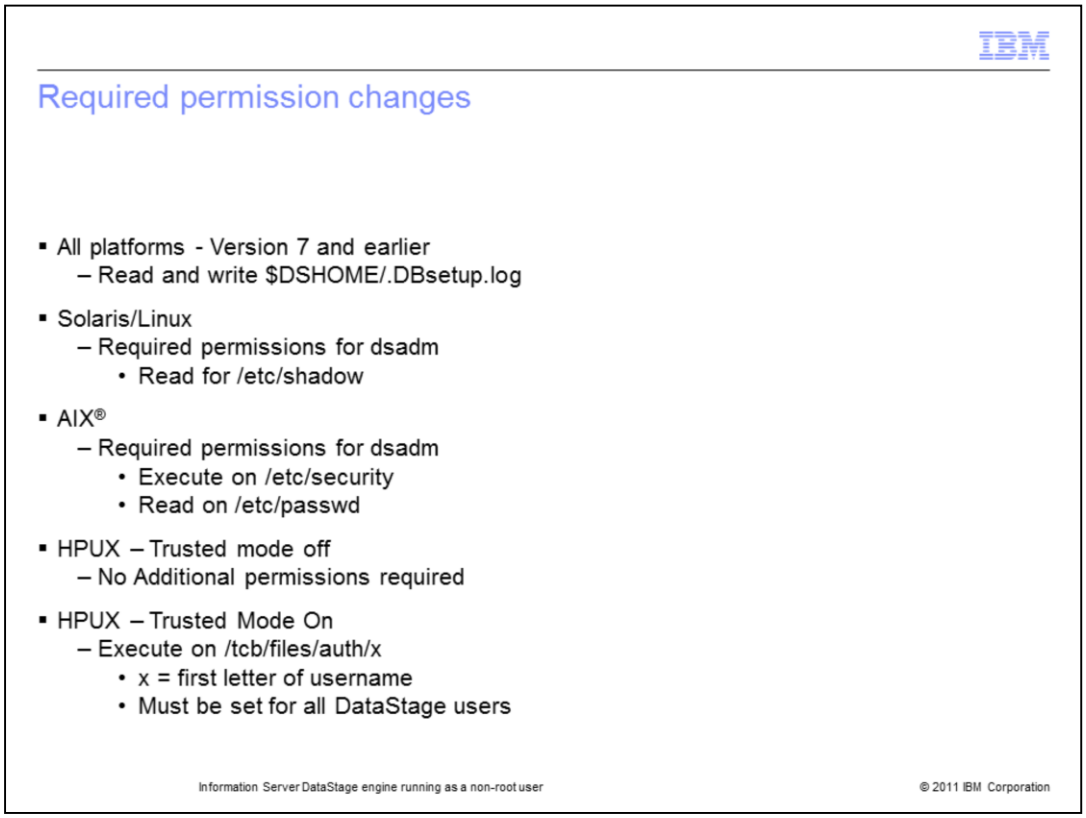

This slide displays the files that will need permission changes so that the dsadm user can authenticate the DataStage users. At version 7 and earlier, the .DBsetup.log file in DSEngine was owned by root with no write permissions for dsadm. Version 8 of DataStage changed the ownership of that file to dsadm so there is no longer a permission issue with that file at version 8 of DataStage.

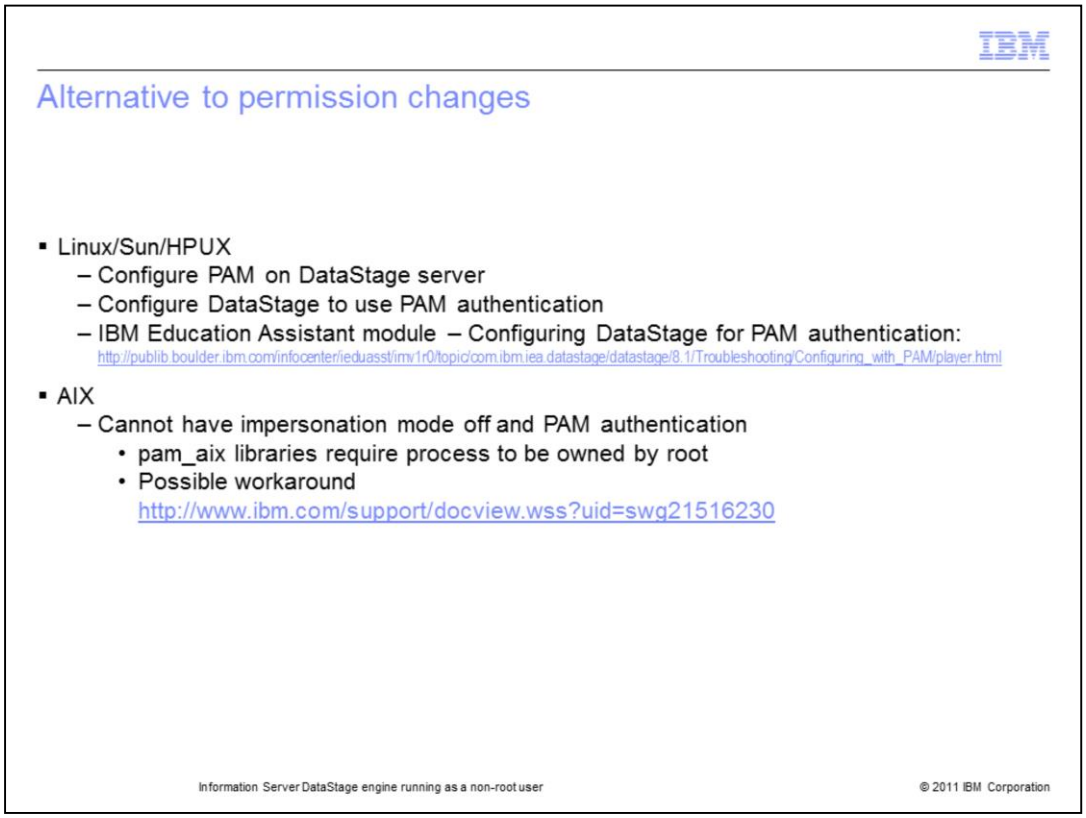

If it is not possible to change permission on the required files due to security policies, there is an alternative solution. You can configure PAM on your DataStage server and then configure DataStage to authenticate using PAM. For more information on how to configure this, see the IBM Education Assistant module on Configuring DataStage for PAM authentication. This module assumes that the DataStage server already has PAM configured properly and will go through the steps necessary to configure DataStage to use PAM authentication.

Changing DataStage to use PAM authentication with impersonation mode off will work for all platforms except AIX. On AIX, the pam_aix library can only be accessed by a process running as root. Since the dsrpcd daemon will run as dsadm with impersonation mode off, the PAM authentication will fail. There is a possible alternative of using a different library that does not have this restriction, such as the Kerberos libraries. See the technote referenced on this slide for more details.

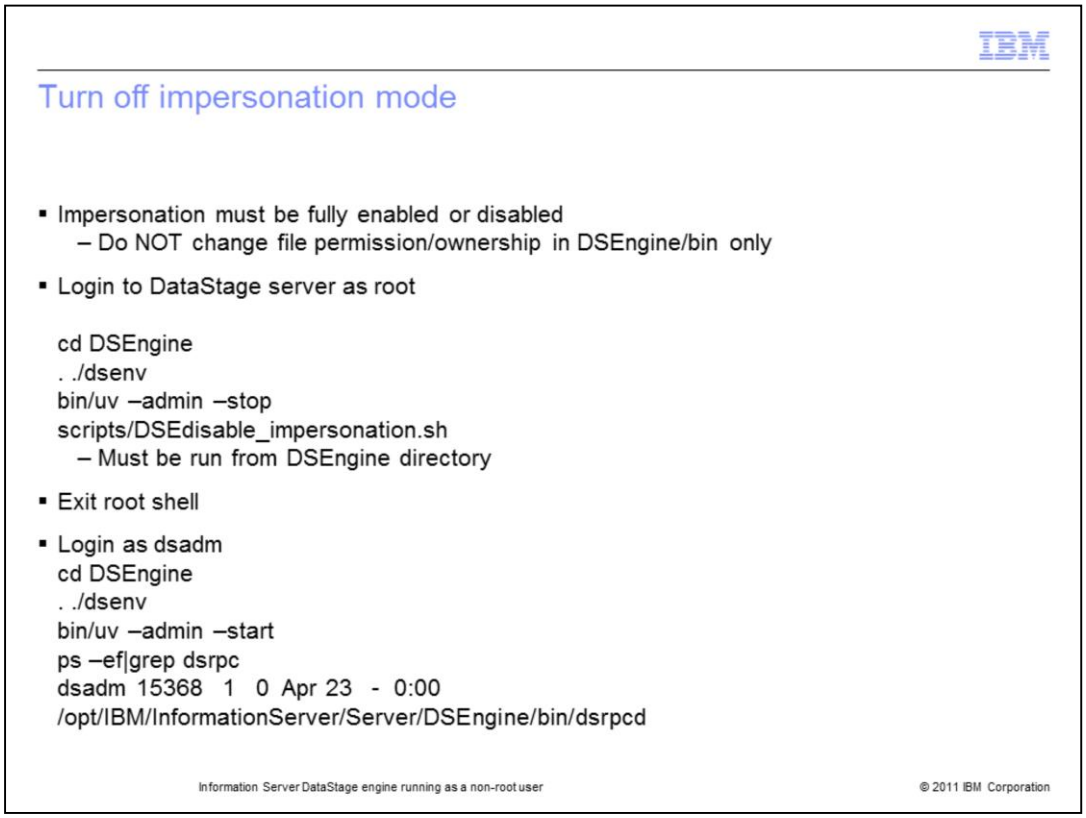

Once the system has been configured as described in the previous slides, the last step is to turn impersonation mode off. It is important that the impersonation be fully enabled or disabled using the supplied scripts. It is not acceptable to manually modify DataStage to be have impersonation mode partially turned off. If this is done by only changing the permissions and ownership of the files in the DSEngine/bin directory for example, the engine will not run properly.

To turn impersonation mode off, login to the DataStage server as the root user. Change directories to the DSEngine directory. Source the dsenv file to set up your environment and then stop the DataStage engine. Next, run:

scripts/DSEdisable_impersonation.sh

This script must be run from within the DSEngine directory and not from the scripts directory. Once the script completes, logout of your root window and log back in as the dsadm user and restart DataStage. Check to be sure that the dsrpcd is now running as the dsadm user.

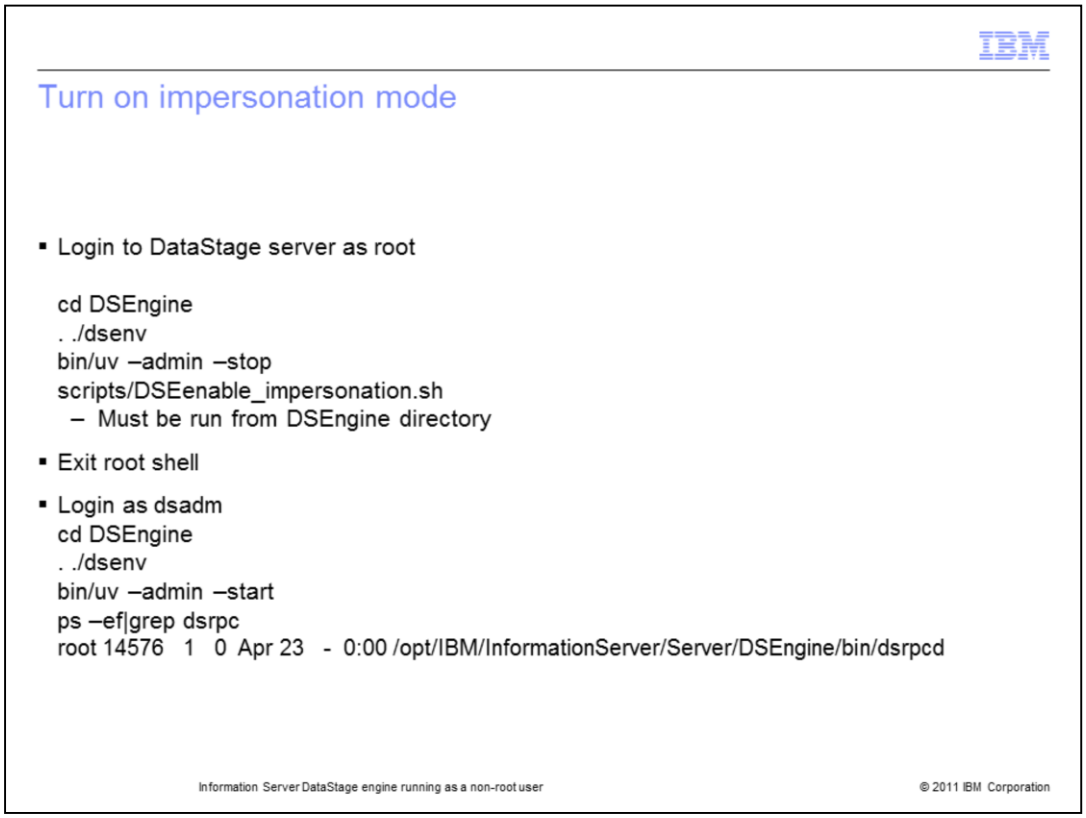

If you want to turn impersonation mode back on, login to the system as the root user. Change directories to the DSEngine directory. Source the dsenv file to set up your environment and then stop the DataStage engine. Next, run:

scripts/DSEenable_impersonation.sh

This script must be run from within the DSEngine directory and not from the scripts directory. Once the script completes, exit the root shell, login as the dsadm user and restart DataStage. Check to be sure that the dsrpcd is now running as the root user.

Trademarks, disclaimer, and copyright information