

Directory Traversal, File Inclusion, and The Proc File System

Directory traversal and local file inclusion bugs are frequently seen in web applications. Directory traversal is when a server allows an attacker to read a file or directories outside of the normal web server directory. Local file inclusion allows an attacker the ability to include an arbitrary local file (from the web server) in the web server's response. Both of these bugs can be used to read arbitrary files from the server.



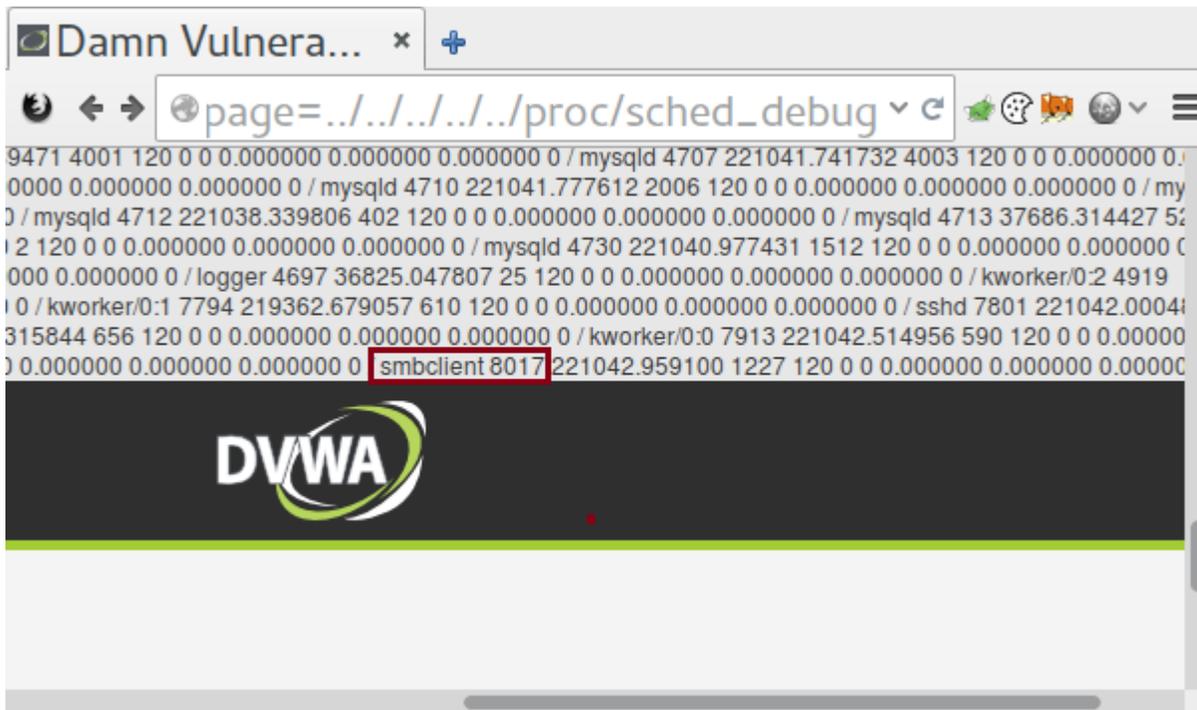
FILE INCLUSION EXAMPLE IN DVWA

In most cases, this means that an attacker can read the `/etc/passwd` file and the shell history files in order to find information leaks. However, an attacker can also use this to read the proc file system. This can provide some interesting insights into what's running on the server.

A few of the more interesting proc entries include:

Directory	Description
<code>/proc/sched_debug</code>	This is usually enabled on newer systems, such as RHEL 6. It provides information as to what process is running on which cpu. This can be handy to get a list of processes and their PID number.
<code>/proc/mounts</code>	Provides a list of mounted file systems. Can be used to determine where other interesting files might be located
<code>/proc/net/arp</code>	Shows the ARP table. This is one way to find out IP addresses for other internal servers.
<code>/proc/net/route</code>	Shows the routing table information.

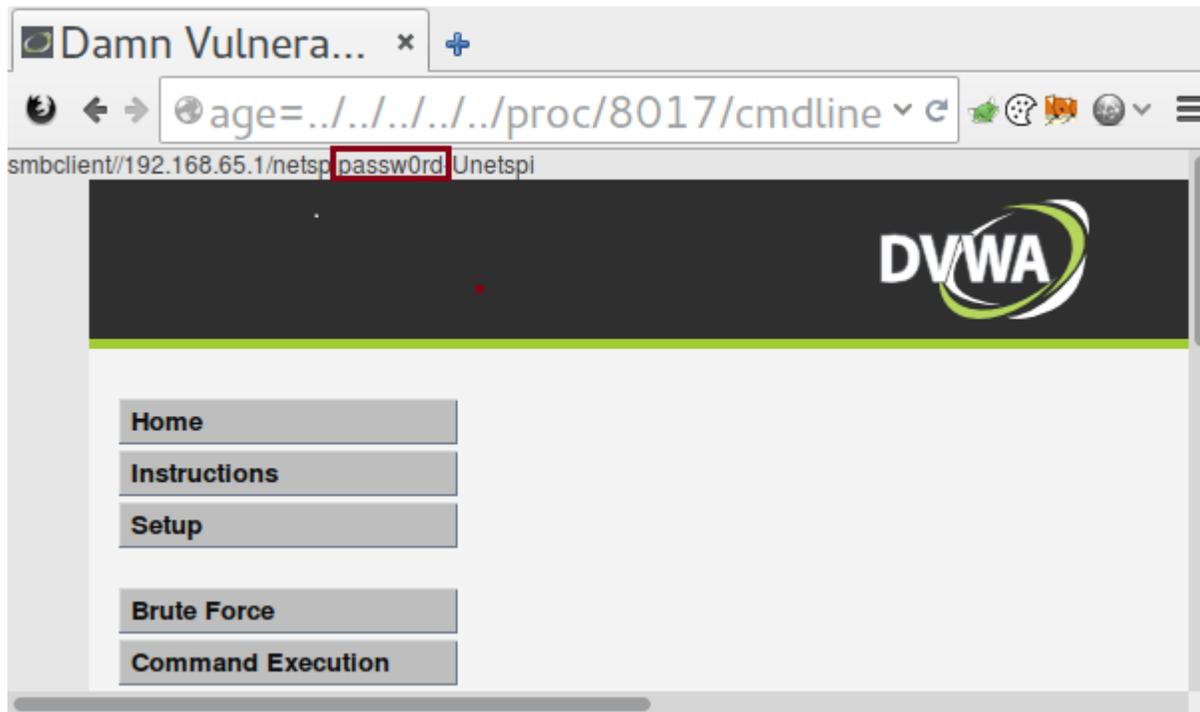
<code>/proc/net/tcp</code> and <code>/proc/net/udp</code>	Provides a list of active connections. Can be used to determine what ports are listening on the server
<code>/proc/net/fib_tribe</code>	This is used for route caching. This can also be used to determine local IPs, as well as gain a better understanding of the target's networking structure
<code>/proc/version</code>	Shows the kernel version. This can be used to help determine the OS running and the last time it's been fully updated.



OUTPUT OF /PROC/SCHED_DEBUG SHOWS SMBCLIENT RUNNING ON PID 8017

Each process also has its own set of attributes. If you have the PID number and access to that process, then you can obtain some useful information about it, such as its environmental variables and any command line options that were run. Sometimes these include passwords. Linux also has a special proc directory called *self* which can be used to query information about the current process without having to know it's PID. In the following examples you can replace [PID] with either self or the PID of the process you wish to examine.

Directory	Description
<code>/proc/[PID]/cmdline</code>	Lists everything that was used to invoke the process. This sometimes contains useful paths to configuration files as well as usernames and passwords.
<code>/proc/[PID]/environ</code>	Lists all the environment variables that were set when the process was invoked. This also sometimes contains useful paths to configuration files as well as usernames and passwords.
<code>/proc/[PID]/cwd</code>	Points to the current working directory of the process. This may be useful if you don't know the absolute path to a configuration file.
<code>/proc/[PID]/fd/[#]</code>	Provides access to the file descriptors being used. In some cases this can be used to read files that are opened by a process.



CMDLINE LEAKS A CIFS SHARE PASSWORD VIA SMBCLIENT

Combining directory traversal and file inclusion vulnerabilities with the proc file system ends up being a great way to gain access to information relating to the running processes on a Linux system. For instance, it's possible to locate database servers by looking at the current connections the machine has. It also makes it easier to fingerprint what software is running on the machines in order to help determine if they are vulnerable to version-specific issues. It's important not to overlook the proc file system when one of these vulnerabilities are found during a penetration test.