NAME

flawfinder − find potential security flaws ("hits") in source code

SYNOPSIS

flawfinder [**−−help**] [**−−version**] [**−−allowlink**] [**−−inputs**] [**-minlevel=***X*] [**-m** *X*] [**−−neverignore**] [**-n**] [**−−columns**] [**−−context**] [**-c**] [**−−dataonly**] [**−−html**] [**−−immediate**] [**-i**] [**−−singleline]** [**−S**] [**−−omit-time**] [**−−quiet**] [**−−loadhitlist=***F*] [**−−savehitlist=***F*] [**−−diffhitlist=***F*] [**−−**] *[source code file or source root directory]+*

DESCRIPTION

Flawfinder searches through C/C++ source code looking for potential security flaws. To run flawfinder, simply give flawfinder a list of directories or files. For each directory given, all files that have C/C++ filename extensions in that directory (and its subdirectories, recursively) will be examined. Thus, for most projects, simply give flawfinder the name of the source code's topmost directory (use ''.'' for the current directory), and flawfinder will examine all of the project's C/C++ source code.

Flawfinder will produce a list of ''hits'' (potential security flaws), sorted by risk; the riskiest hits are shown first. The risk level is shown inside square brackets and varies from 0, very little risk, to 5, great risk. This risk level depends not only on the function, but on the values of the parameters of the function. For example, constant strings are often less risky than fully variable strings in many contexts, and in those contexts the hit will have a lower risk level. Flawfinder knows about gettext (a common library for internationalized programs) and will treat constant strings passed through gettext as though they were constant strings; this reduces the number of false hits in internationalized programs. Flawfinder correctly ignores most text inside comments and strings. Normally flawfinder shows all hits with a risk level of at least 1, but you can use the −−minlevel option to show only hits with higher risk levels if you wish.

Not every hit is actually a security vulnerability, and not every security vulnerability is necessarily found. Nevertheless, flawfinder can be an aid in finding and removing security vulnerabilities. A common way to use flawfinder is to first apply flawfinder to a set of source code and examine the highest-risk items. Then, use −−inputs to examine the input locations, and check to make sure that only legal and safe input values are accepted from untrusted users.

Once you've audited a program, you can mark source code lines that are actually fine but cause spurious warnings so that flawfinder will stop complaining about them. To mark a line, put a specially-formatted comment either on the same line (after the source code) or all by itself in the previous line. The comment must have one of the two following formats:

- // Flawfinder: ignore
- /* Flawfinder: ignore */

Note that, for compatibility's sake, you can replace "Flawfinder:" with "ITS4:" or "RATS:" in these specially-formatted comments. Since it's possible that such lines are wrong, you can use the ''−−neverignore'' option, which causes flawfinder to never ignore any line no matter what the comments say. Thus, responses that would otherwise be ignored would be included (or, more confusingly, −−neverignore ignores the ignores). This comment syntax is actually a more general syntax for special directives to flawfinder, but currently only ignoring lines is supported.

Flawfinder uses an internal database called the ''ruleset''; the ruleset identifies functions that are common causes of security flaws. The standard ruleset includes a large number of different potential problems, including both general issues that can impact any C/C++ program, as well as a number of specific Unix-like and Windows functions that are especially problematic. As noted above, every potential security flaw found in a given source code file (matching an entry in the ruleset) is called a ''hit,'' and the set of hits found during any particular run of the program is called the ''hitlist.'' Hitlists can be saved (using −−savehitlist), reloaded back for redisplay (using −−loadhitlist), and you can show only the hits that are different from another run (using −−diffhitlist).

Any filename given on the command line will be examined (even if it doesn't have a usual C/C++ filename extension); thus you can force flawfinder to examine any specific files you desire. While searching directories recursively, flawfinder only opens regular files that have C/C++ filename extensions (e..g, .c, .h, .cc,

.cxx, and so on). The filename "−" means the standard input. To prevent security problems, special files (such as device special files and named pipes) are always skipped, and by default symbolic links are skipped,

Flawfinder intentionally works similarly to another program, ITS4, which is not fully open source software (as defined in the Open Source Definition) nor free software (as defined by the Free Software Foundation). The author of Flawfinder has never seen ITS4's source code.

BRIEF TUTORIAL

Here's a brief example of how flawfinder might be used. Imagine that you have the C/C++ source code for some program named xyzzy (which you may or may not have written), and you're searching for security vulnerabilities (so you can fix them before customers encounter the vulnerabilities). For this tutorial, I'll assume that you're using a Unix-like system, such as Linux, OpenBSD, or MacOS X.

If the source code is in a subdirectory named xyzzy, you would probably start by opening a text window and using flawfinder's default settings, to analyze the program and report a prioritized list of potential security vulnerabilities (the "less" just makes sure the results stay on the screen):

```
flawfinder xyzzy | less
```

At this point, you will a large number of entries; each entry begins with a filename, a colon, a line number, a risk level in brackets (where 5 is the most risky), a category, the name of the function, and a description of why flawfinder thinks the line is a vulnerability. If you don't understand the error message, please see documents such as the *Writing Secure Programs for Linux and Unix HOWTO* at http://www.dwheeler.com/secure-programs which provides more information on writing secure programs.

Once you identify the problem and understand it, you can fix it. Occasionally you may want to re-do the analysis, both because the line numbers will change *and* to make sure that the new code doesn't introduce yet a different vulnerability.

If you've determined that some line isn't really a problem, and you're sure of it, you can insert just before or on the offending line a comment like

```
/* Flawfinder: ignore */
```

to keep them from showing up in the output.

Once you've done that, you should go back and search for the program's inputs, to make sure that the program strongly filters any of its untrusted inputs. Flawfinder can identify many program inputs by using the −−inputs option, like this:

```
flawfinder −−inputs xyzzy
```

Flawfinder includes many other options, including ones to create HTML versions of the output (useful for prettier displays). The next section describes those options in more detail.

OPTIONS

Flawfinder has a number of options, which can be grouped into options that control its own documentation, select which hits to display, select the output format, and perform hitlist management.

Documentation

−−help Show usage (help) information.

−−version Shows (just) the version number and exits.

Selecting Hits to Display

−−allowlink Allow the use of symbolic links; normally symbolic links are skipped. Don't use this option if you're analyzing code by others; attackers could do many things to cause problems for an analysis with this option enabled. For example, an attacker could insert symbolic links to files such as /etc/passwd (leaking information about the file) or create a circular loop, which would cause flawfinder to run ''forever''. Another problem with enabling this option is that if the same file is referenced multiple times using symbolic links, it will be analyzed multiple times (and thus reported multiple times). Note that flawfinder already includes some protection against symbolic links to special file types such as device file types (e.g., /dev/zero or C:\mystuff\com1). Note that for flawfinder version 1.01 and before, this was the default.

−−inputs Show only functions that obtain data from outside the program; this also sets minlevel to 0.

−−minlevel=*X*

-m *X* Set minimum risk level to X for inclusion in hitlist. This can be from 0 (''no risk'') to 5 (''maximum risk''); the default is 1.

−−neverignore

-n Never ignore security issues, even if they have an ''ignore'' directive in a comment.

Selecting Output Format

−−columns Show the column number (as well as the file name and line number) of each hit; this is shown after the line number by adding a colon and the column number in the line (the first character in a line is column number 1).

−−context

-c Show context, i.e., the line having the "hit"/potential flaw. By default the line is shown immediately after the warning.

−−dataonly Don't display the header and footer. Use this along with −−quiet to see just the data itself.

−−html Format the output as HTML instead of as simple text.

−−immediate

-i Immediately display hits (don't just wait until the end).

−−singleline

-S Display as single line of text output for each hit. Useful for interacting with compilation tools.

−−omittime Omit timing information. This is useful for regression tests of flawfinder itself, so that the output doesn't vary depending on how long the analysis takes.

−−quiet Don't display status information (i.e., which files are being examined) while the analysis is going on.

Hitlist Management

−−savehitlist=*F*

Save all resulting hits (the "hitlist") to F.

−−loadhitlist=*F*

Load the hitlist from F instead of analyzing source programs.

−−diffhitlist=*F*

Show only hits (loaded or analyzed) not in F. F was presumably created previously using −−savehitlist. If the −−loadhitlist option is not provided, this will show the hits in the analyzed source code files that were not previously stored in F. If used along with −−loadhitlist, this will show the hits in the loaded hitlist not in F. The difference algorithm is conservative; hits are only considered the ‘‘same’’ if they have the same filename, line number, column position, function name, and risk level.

EXAMPLES

flawfinder /usr/src/linux-2.4.12

Examine all the C/C++ files in the directory /usr/src/linux-2.4.12 and all its subdirectories (recursively), reporting on all hits found.

flawfinder −−minlevel=4 .

Examine all the C/C++ files in the current directory and its subdirectories (recursively); only report vulnerabilities level 4 and up (the two highest risk levels).

flawfinder −−inputs mydir

Examine all the C/C++ files in mydir and its subdirectories (recursively), and report functions that take inputs (so that you can ensure that they filter the inputs appropriately).

flawfinder −−neverignore mydir

Examine all the C/C++ files in the directory mydir and its subdirectories, including even the hits marked for ignoring in the code comments.

flawfinder −−quiet −−dataonly mydir

Examine mydir and report only the actual results. This form is useful if the output will be piped into other tools for further analysis.

flawfinder −−quiet −−html −−context mydir > results.html

Examine all the C/C++ files in the directory mydir and its subdirectories, and produce an HTML formatted version of the results. Source code management systems (such as SourceForge and Savannah) might use a command like this.

flawfinder −−quiet −−savehitlist saved.hits *.[ch]

Examine all .c and .h files in the current directory. Don’t report on the status of processing, and save the resulting hitlist (the set of all hits) in the file saved.hits.

flawfinder −−diffhitlist saved.hits *.[ch]

Examine all .c and .h files in the current directory, and show any hits that weren’t already in the file saved.hits. This can be used to show only the ‘‘new’’ vulnerabilities in a modified program, if saved.hits was created from the older version of the program being analyzed.

SECURITY

You should always analyze a *copy* of the source program being analyzed, not a directory that can be modified by a developer while flawfinder is performing the analysis. This is *especially* true if you don't necessily trust a developer of the program being analyzed. If an attacker has control over the files while you're analyzing them, the attacker could move files around or change their contents to prevent the exposure of a security problem (or create the impression of a problem where there is none). If you're worried about malicious programmers you should do this anyway, because after analysis you'll need to verify that the code eventually run is the code you analyzed. Also, do not use the −−allowlink option in such cases; attackers could create malicious symbolic links to files outside of their source code area (such as /etc/passwd).

Source code management systems (like SourceForge and Savannah) definitely fall into this category; if you're maintaining one of those systems, first copy or extract the files into a separate directory (that can't be controlled by attackers) before running flawfinder or any other code analysis tool.

Note that flawfinder only opens regular files, directories, and (if requested) symbolic links; it will never open other kinds of files, even if a symbolic link is made to them. This counters attackers who insert unusual file types into the source code. However, this only works if the filesystem being analyzed can't be modified by an attacker during the analysis, as recommended above. This protection also doesn't work on Cygwin platforms, unfortunately.

Cygwin systems (Unix emulation on top of Windows) have an additional problem if flawfinder is used to analyze programs the analyzer cannot trust due to a design flaw in Windows (that it inherits from MS-DOS). On Windows and MS-DOS, certain filenames (e.g., ''com1'') are automatically treated by the operating system as the names of peripherals, and this is true even when a full pathname is given. Yes, Windows and MS-DOS really are designed this badly. Flawfinder deals with this by checking what a filesystem object is, and then only opening directories and regular files (and symlinks if enabled). Unfortunately, this doesn't work on Cygwin; on at least some versions of Cygwin on some versions of Windows, merely trying to determine if a file is a device type can cause the program to hang. A workaround is to delete or rename any filenames that are interpreted as device names before performing the analysis. These so-called ''reserved names'' are CON, PRN, AUX, CLOCK$, NUL, COM1-COM9, and LPT1-LPT9, optionally followed by an extension (e.g., ''com1.txt''), in any directory, and in any case (Windows is case-insensitive).

BUGS

Flawfinder is currently limited to C/C++. It's designed so that adding support for other languages should be easy.

Flawfinder can be fooled by user-defined functions or method names that happen to be the same as those defined as ''hits'' in its database, and will often trigger on definitions (as well as uses) of functions with the same name. This is because flawfinder is based on text pattern matching, which is part of its fundamental design and not easily changed. This isn't as much of a problem for C code, but it can be more of a problem for some C++ code which heavily uses classes and namespaces. On the positive side, flawfinder doesn't get confused by many complicated preprocessor sequences that other tools sometimes choke on.

Preprocessor commands embedded in the middle of a parameter list of a call can cause problems in parsing, in particular, if a string is opened and then closed multiple times using an #ifdef .. #else construct, flawfinder gets confused. Such constructs are bad style, and will confuse many other tools too. If you must analyze such files, rewrite those lines. Thankfully, these are quite rare.

The routine to detect statically defined character arrays uses simple text matching; some complicated expresions can cause it to trigger or not trigger unexpectedly.

Security vulnerabilities might not be identified as such by flawfinder, and conversely, some hits aren't really security vulnerabilities. This is true for all static security scanners, especially those like flawfinder that use a pattern-based approach to identifying problems. Still, it can serve as a useful aid, and that's the point.

SEE ALSO

See the flawfinder website at http://www.dwheeler.com/flawfinder. You should also see the *Secure Programming for Unix and Linux HOWTO* at http://www.dwheeler.com/secure-programs.

AUTHOR

David A. Wheeler (dwheeler@dwheeler.com).