

<p>find.242271e Error Type: Functional Bug Avg. Time: 13.8 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 75%</p>	<p>If find is set to print the found file's base directory followed by the found file's name ("print %H %P\n") and there exist directories of different length, then find incorrectly splits base directory and file name during printing. Because the index state_starting_path is only for the first working directory (fzf.find.c:278-279) the incorrect value of state_starting_path is used when printing base directory and file name (fzf.find.c:709-718, print.c:1813). Examples of Correct Fixes: 1) Recompute state_starting_path_length for each entry before calling find. 2) Weaken condition that prevents state_starting_path to be reset. Example of Incorrect Fix: Always update state_starting_path_length even if ent->file_level = 0 (<i>Regression</i> because it then carries the incorrect "starting-path-length" find is set to print all files that are exactly 2 days old (time 2), it crashes with a segmentation fault. Variable **pend is defined at pointer point (parser.c:2273) and expected to be allocated when strtok_r is called (parser.c:2275). However, it is still NULL after the call such that the null pointer check for pend is itself a null pointer dereference (parser.c:2762). Examples of Correct Fixes: 1) Add null pointer check for pend. 2) Change definition of **pend to **pend and update references. 3) Allocate memory for **pend. Examples of Incorrect Fixes: 1) Remove code containing null pointer dereference (<i>Treating the Symptom</i>). 2) Change the check involving **pend (<i>Treating the Symptom</i> because the nullpointer is still dereferenced, only the program does not crash).</p>	<p>grep.55c7b6a Error Type: Functional Bug Avg. Time: 21.1 min Explanation: Slightly difficult Patching: Not at all difficult Correctness: 91%</p> <p>If grep is set to identify skip devices, FIFOs, and sockets (D skip), then grep does not search on standard input when no file is provided. When the skip option is enabled, variables devices is set to SKIP_DEVICES (main.c:1852-1859). If no file is provided, variable file is NULL and variable dev is set to STDIN_FILENO (main.c:1217-1218). This code which handles SKIP_DEVICES (main.c:1246-1253) decides to skip STDIN (which is a special device) even though it should not (dev == STDIN_FILENO). Examples of Correct Fixes: 1) Do not skip if dev is set to STDIN_FILENO. 2) Do not skip if file is not set (and thus dev is not set to STDIN_FILENO). Example of Incorrect Fix: Negate the skip condition (<i>Regression</i> because it skips everything that should not be skipped while indeed not skipping STDIN).</p>
<p>find.0b10e9 Error Type: Crash Avg. Time: 22.9 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 81%</p>	<p>grep.54d55ba Error Type: Crash Avg. Time: 26.7 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 69%</p> <p>If grep is set to search in all files under each directory recursively (-r) but to exclude certain directories (--exclude-dir=), then grep crashes with a segmentation fault. When grep computes the name_spaces (grep.c:1361), it calls function idrfl via function save_dir (lib/save_dir.c:123). Now, the code in idrfl that is supposed to remove the trailing slashes from the directory name uses the uninitialized variable path instead of variable dir (lib/save_dir.c:53). The nullpointer dereference results in a segmentation fault. Example of Correct Fix: Substitute path with dir. Examples of Incorrect Fixes: 1) Return if path is not initialized (<i>Regression</i> because idrfl returns false even if dir is a directory). 2) Only use path if initialized (<i>Regression</i> because idrfl does not return trailing slashes).</p>	
<p>find.07941b1 Error Type: Crash Avg. Time: 23.7 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 89%</p>	<p>grep.9c45c193 Error Type: Functional Bug Avg. Time: 37.7 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 83%</p> <p>If grep is set to search only specific files (--include=), then grep does not print a match even if there is one. First, main correctly adds the include pattern with EXCLUDE_INCLUDE flag set (grep.c:2336-2440). When the files are chosen for the search, files that are supposed to be included are actually excluded because the return value of excluded_file_name is unnecessarily negated (grep.c:2267-2269). The behavior changes if the EXCLUDE_INCLUDE flag is present (lib/exclude.c:410, lib/exclude.c:335). Examples of Correct Fixes: 1) Remove negation such that included_patterns are not excluded during classification. 2) Do not set EXCLUDE_INCLUDE flag for included_patterns which effectively negates the faulty condition. Example of Incorrect Fix: Independent of whether a file matches the included_patterns, never exclude (<i>Regression</i> because it doesn't skip files that are "not" in the included patterns).</p>	
<p>find.c891c11 Error Type: Crash Avg. Time: 31.4 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 54%</p>	<p>grep.5fabc79 Error Type: Infinite Loop Avg. Time: 38.8 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 50%</p> <p>If grep is set to search for fixed strings (-F), the empty string is given (""), and the locale is UTF8, then grep runs indefinitely. When EXECUTE searches for a match of the empty string, variable len is set to the size of the match; here, len=0 (ksearch.c:1106). Because len=0, the check is_mh_middle (ksearch.c:1137-1145) whether the match occurs within a multibyte character returns true (ksearch.c:1108). However, the size of the supposed multibyte character is computed as mh_len (ksearch.c:115). When mh_len is added to beg (ksearch.c:118) to advance behind the supposed multibyte character, beg's value remains unchanged. The loop is continued (ksearch.c:120). Since beg has the same value every time the loop exit condition is checked (ksearch.c:103), the loop exit condition never holds, resulting in an infinite loop. Examples of Correct Fixes: 1) Function is_mh_middle returns false for len=0. 2) Only call is_mh_middle if len is set. 3) Jump to success if mh_len=1. Examples of Incorrect Fixes: 1) Remove continue (<i>Treating the Symptom</i>). 2) Do not reset beg (<i>Regression</i> because it breaks multibyte character handling). 4) Do not compute mh_len before returning buffer until end of line (<i>Regression</i> because only match should be returned).</p>	
<p>find.0c6ec46 Error Type: Functional Bug Avg. Time: 38.2 min Explanation: Moderately difficult Patching: Not at all difficult Correctness: 89%</p>	<p>grep.d9d6c340 Error Type: Infinite Loop Avg. Time: 40.6 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 45%</p> <p>If find is set to search a directory referenced by a symbolic link and containing a file, and find is set to follow symbolic links (-L) or to follow symbolic links except for those set to be searched (-H), then find does not print the file in the referenced directory and instead reports "Too many levels of symbolic links". Because of a mixup in the condition of a ternary operator (find.c:1094), extraneous are set to O_NONBLOCK when it should be 0 and to 0 when it should be O_NONBLOCK. The flag controls whether symlinks are followed when a directory is opened (find.c:1077). Because of this faulty, child returns SafeChildFailsSym (find.c:1618) when the error message is printed (find.c:1642). Example of Correct Fix: Tertiary operator. Example of Incorrect Fix: Do not fail if safe_child returns SafeChildFailsSym (<i>Treating the Symptom</i>).</p>	
<p>find.09155716 Error Type: Crash Avg. Time: 44.8 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 54%</p>	<p>grep.2be0c659 Error Type: Functional Bug Avg. Time: 47.2 min Explanation: Moderately difficult Patching: Moderately difficult Correctness: 13%</p> <p>If grep conducts a case-insensitive search (-i) in a file containing 8-bit characters and the current locale is Turkish UTF8, then grep prints the wrong output. When grep conducts a case-insensitive search, it lowers the case of the input string before matching (ksearch.c:384-392). The lower case of an upper-case 8-bit character might occupy one more or less bytes. The latter case is not handled. When the match_size is computed (grep.c:1081), the lower-case match size is used (grep.c:1060-1062). When the match is printed, the incorrect lower-case match size which is usually larger than the actual match size is used (grep.c:1085-1091). Examples of Correct Fixes: 1) Update the map that maps lower-case characters to the normal case characters to account for cases where the number of bytes it occupies "increases" in the lower-case. 2) To correct the match_size, lower-case as many characters in the normal-case match as result in match_size-lower-case characters. Examples of Incorrect Fixes: 1) Return complete line if match exists (<i>Regression</i> because only the match should be returned). 2) Add the difference in length of lower-case and normal-case string to the match size (<i>Incomplete Fix</i> because for files that have more multibyte characters than lines in the match, grep reports longer matches than needed).</p>	
<p>find.24b3f3c Error Type: Crash Avg. Time: 45.1 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 59%</p>	<p>grep.8f08d8c2 Error Type: Functional Bug Avg. Time: 48.4 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 75%</p> <p>If we utilize the number file descriptors that can be open simultaneously and set find to execute in for every subdirectory (execdir is "*/"), it quickly runs out of file descriptors. File descriptors are always opened (grep.c:520) but never closed (grep.c:659-664) which raises an error when no more descriptors are available (grep.c:5179). Example of Correct Fix: Close file descriptor as soon as it is not used anymore. Example of Incorrect Fix: Close random file descriptor (<i>Incomplete Fix</i> because still leaking file descriptors).</p>	
<p>find.183115d0 Error Type: Resource Leak Avg. Time: 49.2 min Explanation: Slightly difficult Patching: Slightly difficult Correctness: 51%</p>	<p>grep.58195fab Error Type: Functional Bug Avg. Time: 50.5 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 82%</p> <p>There are two options: -H: If find is set to search for files that were changed in the last <i>n</i> days but <i>n</i> is not a number (<time>), then find complains about a "missing" argument instead of reporting the "incorrect" argument. Function parse_time calls collect_args to assign the current argument argv_arg_ptr to timearg and increment the argument pointer arg_ptr (parser.c:3102). When timearg is failed to be parsed as a number, parse_time returns without decrementing arg_ptr (parser.c:3127-3128). When the error is reported (time.c:1248-1271), the argument pointer to NULL, directly after the incorrect argument (grep.c:1250), such that the error is reported as missing argument instead of invalid argument. 2) If find is set to search for files belonging to a certain group but the group-id is not specified or not a number (-gid <i>x</i>), then find crashes with a segmentation fault. When the argument following the -gid option is being parsed (parser.c:913), insert_group returns NULL because argv_arg_ptr is NULL or not a number (parser.c:2235-2253). This nullpointer remains unchecked and is dereferenced leading to a segmentation fault (grep.c:2241). When nullpointer dereference is fixed the same symptom is observed for -gid as -ctime because the argument pointer is also forgot to be decremented. Examples of Correct Fixes: For first error: 1) dereferenciate arg_ptr when parsing second argument of an option fails or 2) use copy of old argument during error-reporting. For second error, add null pointer check. Example of Incorrect Fix: For first error, decrement argument pointer before even calling parse_time (<i>Regression</i> because even correct arguments are reported as incorrect errors).</p>	
<p>find.66c36bb Error Type: Functional Bug Avg. Time: 55.5 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 92%</p>	<p>grep.c1cb19f6 Error Type: Functional Bug Avg. Time: 58.4 min Explanation: Very difficult Patching: Slightly difficult Correctness: 71%</p> <p>If find is set to search a directory containing a symbolic link, to not follow any symbolic links (except for those specified on the command line: -H), and to print only symbolic links (-l), then find does not print the link. The root cause is that state_curlp is set to mode arg (parser.c:3109), timearg still points to *. However, when it is called the second time (parser.c:3038), timearg already points to * such that it is incorrectly classified as COMP_EQ (parser.c:3178). Examples of Correct Fixes: 1) Save timearg in auxiliary variable and restore after first call to get_comp_type. 2) Pass a copy of timearg into the first call of get_comp_type. 3) Pass a copy of timearg into get_relative_timestamp (which calls get_comp_type the second time). 4) Decrement timearg after the first call to get_comp_type. Example of Incorrect Fix: Restore timearg only if classified as COMP_EQ (<i>Incomplete Fix</i> because it does not solve the problem for -mtime +2).</p>	
<p>find.1445a998 Error Type: Functional Bug Avg. Time: 56.5 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 50%</p>	<p>grep.7aa698d3 Error Type: Functional Bug Avg. Time: 59.9 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 13%</p> <p>If find is set to search a directory containing a symbolic link, to not follow any symbolic links (except for those specified on the command line: -H), and to print only symbolic links (-l), then find does not print the link. The root cause is that state_curlp is set to mode arg (parser.c:3109), timearg still points to *. However, when it is called the second time (parser.c:3038), timearg already points to * such that it is incorrectly classified as COMP_EQ (parser.c:3178). Examples of Correct Fixes: 1) Save timearg in auxiliary variable and restore after first call to get_comp_type. 2) Pass a copy of timearg into the first call of get_comp_type. 3) Pass a copy of timearg into get_relative_timestamp (which calls get_comp_type the second time). 4) Decrement timearg after the first call to get_comp_type. Example of Incorrect Fix: Restore timearg only if classified as COMP_EQ (<i>Incomplete Fix</i> because it does not solve the problem for -mtime +2).</p>	
<p>find.0248a20 Error Type: Infinite Loop Avg. Time: 57.7 min Explanation: Moderately difficult Patching: Moderately difficult Correctness: 40%</p>	<p>grep.320317a Error Type: Crash Avg. Time: 63.7 min Explanation: Moderately difficult Patching: Slightly difficult Correctness: 29%</p> <p>If find is set to search a directory containing a symbolic link that references an ancestor directory and if find is set to follow symlinks (-L), then it runs indefinitely. The global variable dir_idx tracks the directories that have already been visited. The function process_path would correctly exit with a loop warning (find.c:1428-1434) if the current directory (in stat_buf) has already been visited. However, after the current directory is correctly added to those that have already been visited (find.c:1442), the same entry is overwritten with uninitialized path (find.c:1621) such that the current directory is never marked as already visited. Examples of Correct Fixes: 1) Remember whether stat_buf has been called. If not done, call stat_buf before overriding dir_idx[dir_curr] at find/indir.c:621. 2) Always stat before overriding dir_idx[dir_curr] at find/indir.c:621 such that stat_buf is initialized. 3) Only override dir_idx[dir_curr] if stat_buf is initialized. Examples of Incorrect Fixes: 1) Never override dir_idx[dir_curr] (<i>Regression</i> because it isn't overwritten when it should be). 2) Follow links to a maximum depth of 1 (<i>Regression</i> because symlinks might need to be followed to an arbitrary depth).</p>	
<p>find.c680237 Error Type: Functional Bug Avg. Time: 76.4 min Explanation: Moderately difficult Patching: Moderately difficult Correctness: 27%</p>	<p>grep.3cb3dacc Error Type: Crash Avg. Time: 64.8 min Explanation: Very difficult Patching: Moderately difficult Correctness: 10%</p> <p>If find is set to search a directory containing a file, to follow symbolic links (-L), and to execute in for every subdirectory (execdir is "*/"), then find incorrectly also prints the base directory. If find is set to follow symlinks, the flag FTS_LOGICAL is set (find.c:2449) before the directory search is initiated (find.c:3164). When a directory is searched (find.c:3273), the working directory is not changed because FTS_LOGICAL is set. Hence, the "full" pathname is passed as argument to execdir (find.c:4184-4190 and find.c:4167-4171). Example of Correct Fix: Correctly compute pathname and prefix in new_imp_dir_exec. Example of Incorrect Fix: Remove FTS_LOGICAL flag (<i>Incomplete Fix</i> because FTS_LOGICAL is supposed to be used).</p>	
<p>find.c1d0991 Error Type: Functional Bug Avg. Time: 82.2 min Explanation: Very difficult Patching: Very difficult Correctness: 17%</p>	<p>grep.96f0f2c Error Type: Functional Bug Avg. Time: 67.6 min Explanation: Very difficult Patching: Moderately difficult Correctness: 50%</p> <p>If grep conducts a case-insensitive search (-i) for the empty line ("") and an UTF-8 locale is set, then grep reports matches even for non-empty lines. For case-sensitive searches or 8-bit locales, execute is called with the complete buffer and correctly returns no match (grep.c:1045-1046). Otherwise, execute is called for each line (grep.c:1048-1049). However, execute does not handle the case when no match is found (ksearch.c:1088), which is why the non-match is printed (grep.c:1079). Examples of Correct Fixes: 1) Handle the case when no match was found by breaking loop if next_jmp == bufLim. 2) Skip printing if match is empty and we are not in inversion mode (-v). Example of Incorrect Fix: Skip printing if match is empty even if in inversion mode (<i>Regression</i> because it breaks inversion mode).</p>	

Fig. 1. Complete list of errors and their average debugging time, difficulty, and patch correctness, with human-generated explanations of the runtime actions leading to the error, and examples of correct and incorrect fixes, sorted according to average debugging time (zoom required).