SINE NOMINE
ASSOCIATES

Static analysis of OpenAFS code base

Cheyenne Wills
OpenAFS 2019 Workshop

Overview

- What is static analysis
- What are the tools
- Analysis of the OpenAFS code base

What is static analysis

- **Static analysis is performed by analysing the source or object code of a program**
  - **as opposed to dynamic analysis, which is done by analysing a running program**

What are the tools

- **Manual code reviews**
- **Automated tools**
  - **Enhancements to compilers**
  - **Standalone utilities**

Manual Code Reviews

- Traditional method
- Gerrit reviews
  - More than one set of eyes
  - “voting”

Enhancements to compilers

- **GCC 8/9, clang**
  - **Truncation using string functions**
  - **Alignment errors**
  - **Some pointer operations**
  - **Detecting out-of-bounds on arrays**
  - **format overflows and truncations**

Compiler checks

- **--enable-checking on configure**
- **compilers are getting “better” on reporting errors**

gcc compiler warnings

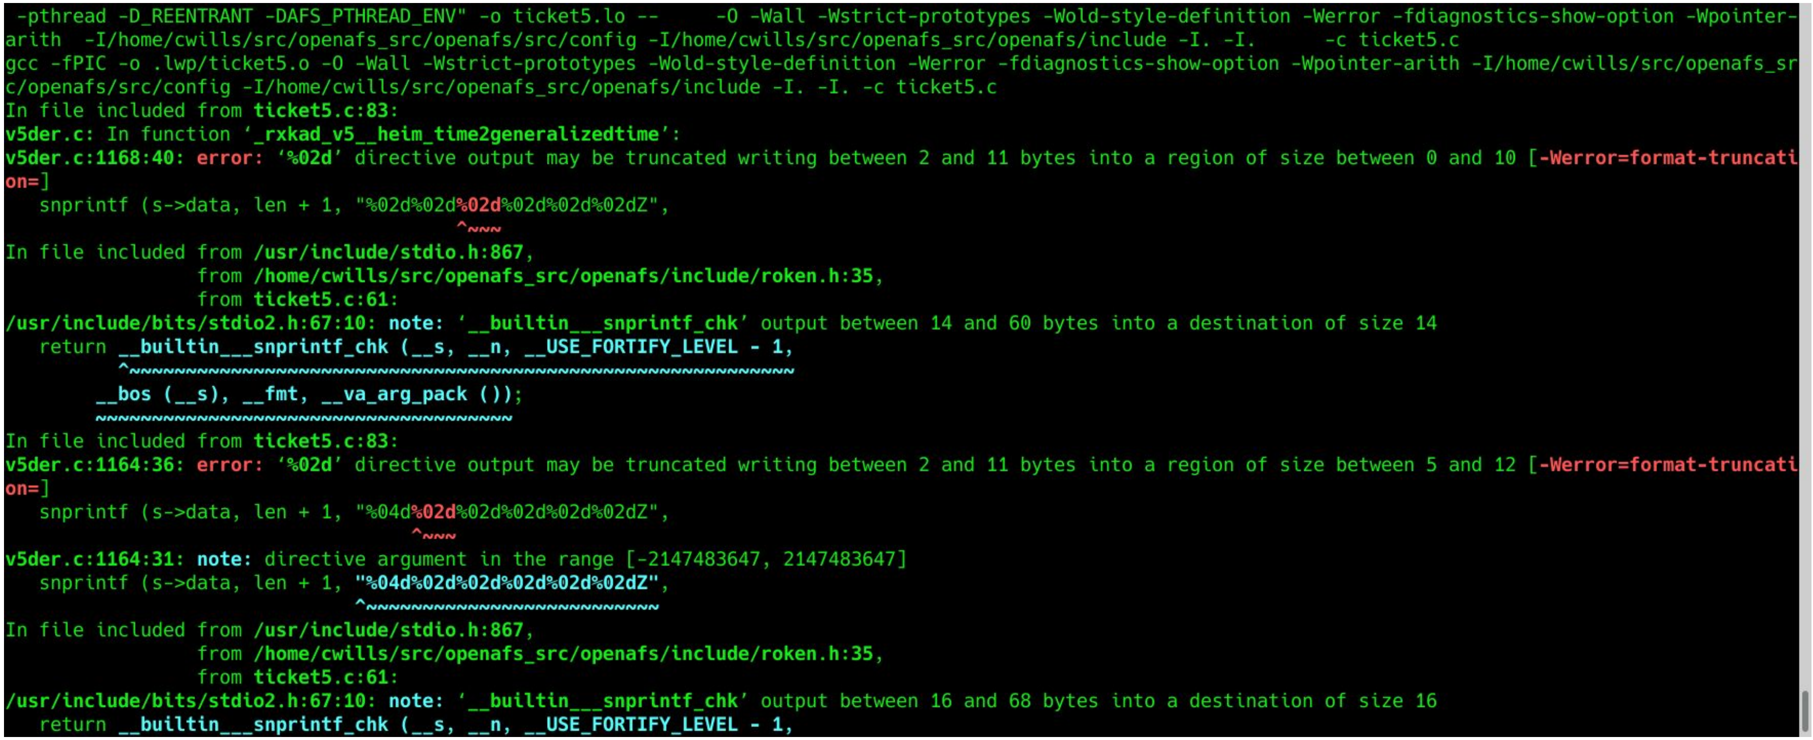

```
 -pthread -D_REENTRANT -DAFS_PTHREAD_ENV" -o ticket5.lo --     -O -Wall -Wstrict-prototypes -Wold-style-definition -Werror -fdiagnostics-show-option -Wpointer-
arith  -I/home/cwills/src/openafs_src/openafs/src/config -I/home/cwills/src/openafs_src/openafs/include -I. -I.      -c ticket5.c
gcc -fPIC -o .lwp/ticket5.o -O -Wall -Wstrict-prototypes -Wold-style-definition -Werror -fdiagnostics-show-option -Wpointer-arith -I/home/cwills/src/openafs_sr
c/openafs/src/config -I/home/cwills/src/openafs_src/openafs/include -I. -I. -c ticket5.c
In file included from ticket5.c:83:
v5der.c: In function ‘_rxkad_v5__heim_time2generalizedtime’:
v5der.c:1168:40: error: ‘%02d’ directive output may be truncated writing between 2 and 11 bytes into a region of size between 0 and 10 [-Werror=format-truncati
on=]
   snprintf (s->data, len + 1, "%02d%02d%02d%02d%02d%02dZ",
                                        ^~~~
In file included from /usr/include/stdio.h:867,
                 from /home/cwills/src/openafs_src/openafs/include/roken.h:35,
                 from ticket5.c:61:
/usr/include/bits/stdio2.h:67:10: note: ‘__builtin___snprintf_chk’ output between 14 and 60 bytes into a destination of size 14
   return __builtin___snprintf_chk (__s, __n, __USE_FORTIFY_LEVEL - 1,
          ^~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
        __bos (__s), __fmt, __va_arg_pack ());
        ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
In file included from ticket5.c:83:
v5der.c:1164:36: error: ‘%02d’ directive output may be truncated writing between 2 and 11 bytes into a region of size between 5 and 12 [-Werror=format-truncati
on=]
   snprintf (s->data, len + 1, "%04d%02d%02d%02d%02d%02dZ",
                                    ^~~~
v5der.c:1164:31: note: directive argument in the range [-2147483647, 2147483647]
   snprintf (s->data, len + 1, "%04d%02d%02d%02d%02d%02dZ",
                               ^~~~~~~~~~~~~~~~~~~~~~~~~~~
In file included from /usr/include/stdio.h:867,
                 from /home/cwills/src/openafs_src/openafs/include/roken.h:35,
                 from ticket5.c:61:
/usr/include/bits/stdio2.h:67:10: note: ‘__builtin___snprintf_chk’ output between 16 and 68 bytes into a destination of size 16
   return __builtin___snprintf_chk (__s, __n, __USE_FORTIFY_LEVEL - 1,
```

Standalone utilities

- **Clang’s static analyzer - scan-build**
  - **part of Clang**
  - **Suite of checkers:**
    - **Core language features and general purpose checks**
    - **Dead Code**
    - **NULL dereferencing**
    - **Security**
    - **Unix/POSIX APIs**

clang scan-build

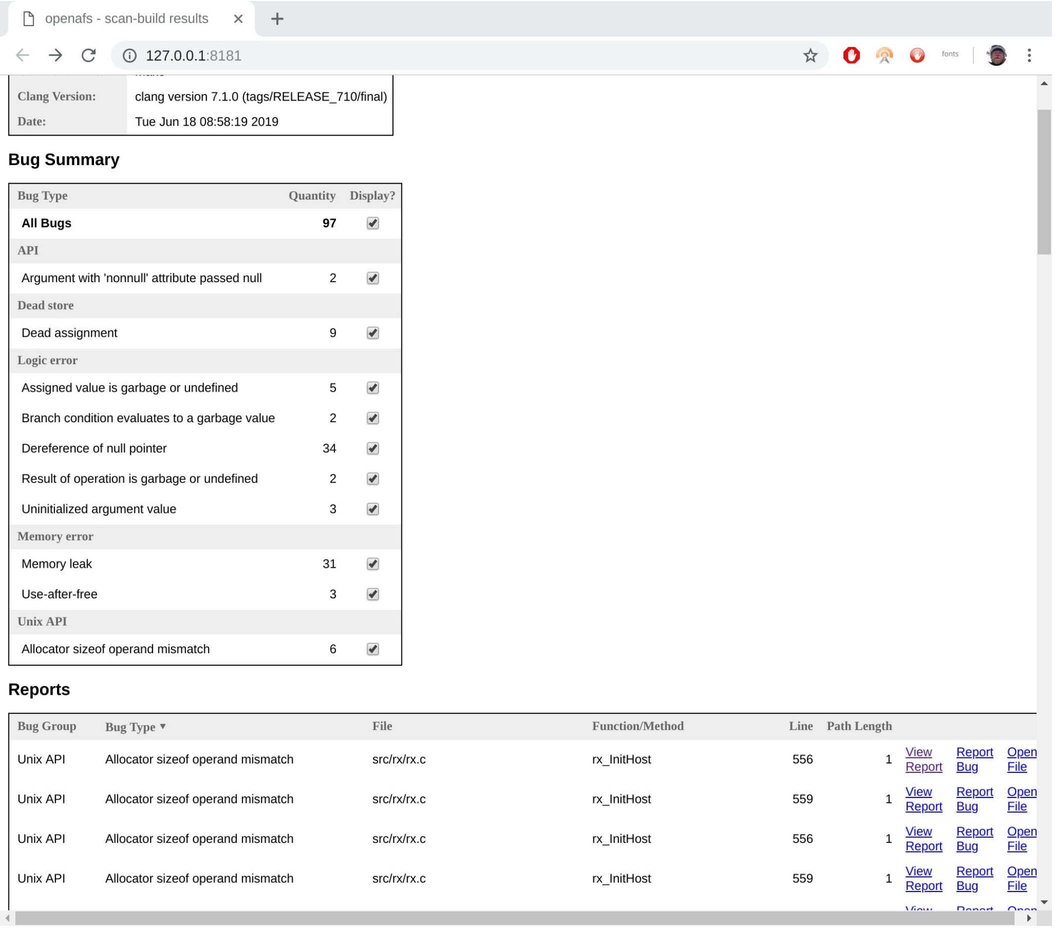

Clang Version:	clang version 7.1.0 (tags/RELEASE_710/final)
Date:	Tue Jun 18 08:58:19 2019

Bug Summary

Bug Type	Quantity	Display?
All Bugs	**97**	☑
API		
Argument with 'nonnull' attribute passed null	2	☑
Dead store		
Dead assignment	9	☑
Logic error		
Assigned value is garbage or undefined	5	☑
Branch condition evaluates to a garbage value	2	☑
Dereference of null pointer	34	☑
Result of operation is garbage or undefined	2	☑
Uninitialized argument value	3	☑
Memory error		
Memory leak	31	☑
Use-after-free	3	☑
Unix API		
Allocator sizeof operand mismatch	6	☑

Reports

Bug Group	Bug Type ▾	File	Function/Method	Line	Path Length			
Unix API	Allocator sizeof operand mismatch	src/rx/rx.c	rx_InitHost	556	1	View Report	Report Bug	Open File
Unix API	Allocator sizeof operand mismatch	src/rx/rx.c	rx_InitHost	559	1	View Report	Report Bug	Open File
Unix API	Allocator sizeof operand mismatch	src/rx/rx.c	rx_InitHost	556	1	View Report	Report Bug	Open File
Unix API	Allocator sizeof operand mismatch	src/rx/rx.c	rx_InitHost	559	1	View Report	Report Bug	Open File

clang scan-build

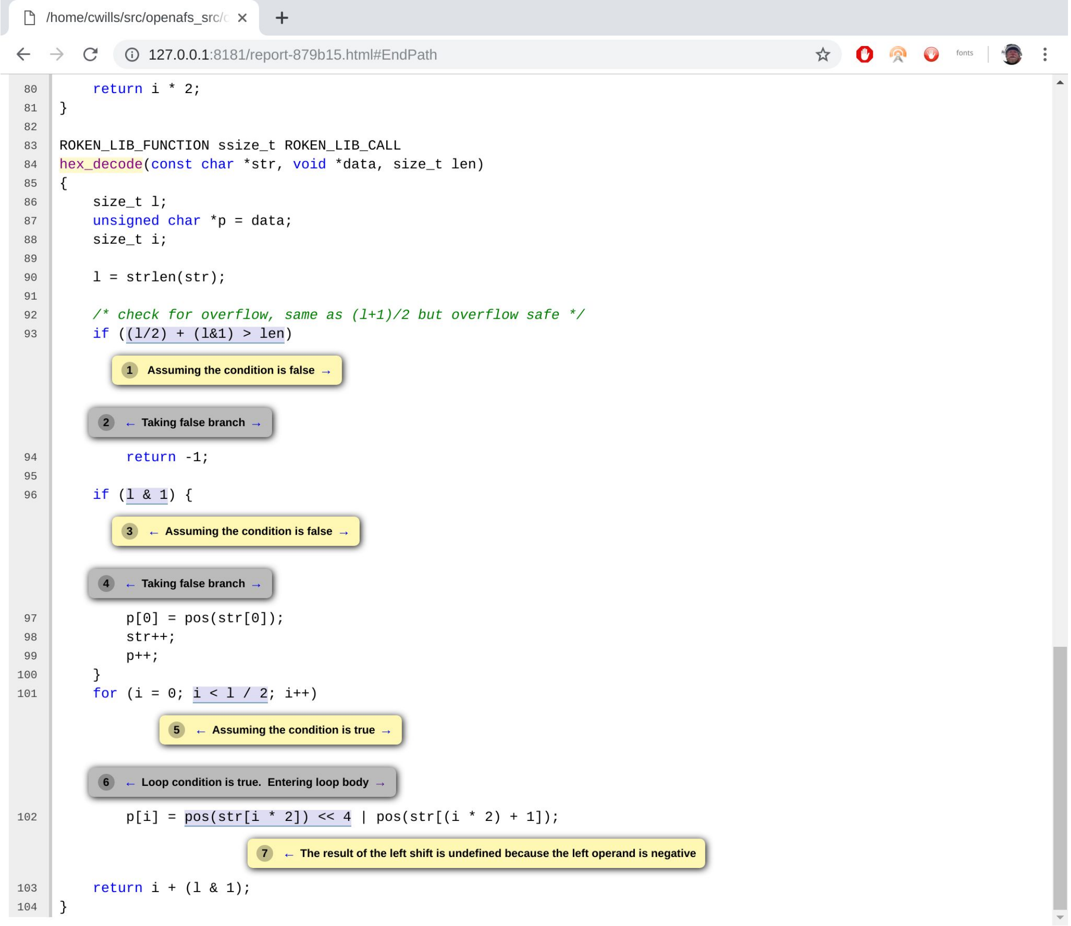
/home/cwills/src/openafs_src/c
127.0.0.1:8181/report-879b15.html#EndPath
fonts

```
    return i * 2;
}

ROKEN_LIB_FUNCTION ssize_t ROKEN_LIB_CALL
hex_decode(const char *str, void *data, size_t len)
{
    size_t l;
    unsigned char *p = data;
    size_t i;

    l = strlen(str);

    /* check for overflow, same as (l+1)/2 but overflow safe */
    if ((l/2) + (l&1) > len)

          1  Assuming the condition is false →

       2  ← Taking false branch →

        return -1;

    if (l & 1) {

          3  ← Assuming the condition is false →

       4  ← Taking false branch →

        p[0] = pos(str[0]);
        str++;
        p++;
    }
    for (i = 0; i < l / 2; i++)

                5  ← Assuming the condition is true →

       6  ← Loop condition is true.  Entering loop body →

        p[i] = pos(str[i * 2]) << 4 | pos(str[(i * 2) + 1]);

                      7  ← The result of the left shift is undefined because the left operand is negative

    return i + (l & 1);
}
```

clang scan-build

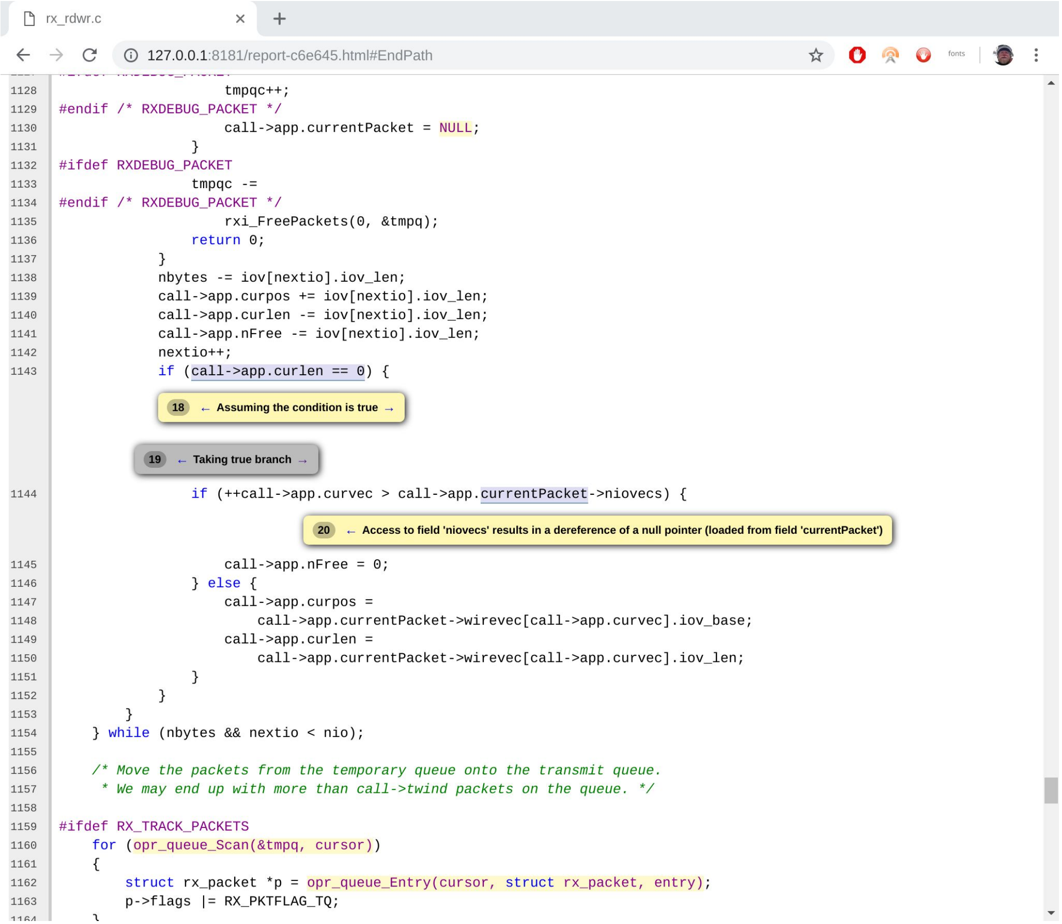

```
                    tmpqc++;
#endif /* RXDEBUG_PACKET */
                    call->app.currentPacket = NULL;
                }
#ifdef RXDEBUG_PACKET
                tmpqc -=
#endif /* RXDEBUG_PACKET */
                    rxi_FreePackets(0, &tmpq);
                return 0;
            }
            nbytes -= iov[nextio].iov_len;
            call->app.curpos += iov[nextio].iov_len;
            call->app.curlen -= iov[nextio].iov_len;
            call->app.nFree -= iov[nextio].iov_len;
            nextio++;
            if (call->app.curlen == 0) {

            18  ← Assuming the condition is true →

         19  ← Taking true branch →

                if (++call->app.curvec > call->app.currentPacket->niovecs) {

                              20  ← Access to field 'niovecs' results in a dereference of a null pointer (loaded from field 'currentPacket')

                    call->app.nFree = 0;
                } else {
                    call->app.curpos =
                        call->app.currentPacket->wirevec[call->app.curvec].iov_base;
                    call->app.curlen =
                        call->app.currentPacket->wirevec[call->app.curvec].iov_len;
                }
            }
        }
    } while (nbytes && nextio < nio);

    /* Move the packets from the temporary queue onto the transmit queue.
     * We may end up with more than call->twind packets on the queue. */

#ifdef RX_TRACK_PACKETS
    for (opr_queue_Scan(&tmpq, cursor))
    {
        struct rx_packet *p = opr_queue_Entry(cursor, struct rx_packet, entry);
        p->flags |= RX_PKTFLAG_TQ;
    }
```

Standalone utilities

- **cppcheck**
  - **variable checking**
  - **out of bounds conditions**
  - **depreciated functions**
  - **memory leaks**
  - **resource leaks**
  - **stylistic and performance errors**

cppcheck

```
...
[src/rxgen/rpc_parse.c:2208]: (warning) %u in format string (no. 3) requires 'unsigned int' but the
argument type is 'signed int'.
[src/rxgen/rpc_parse.c:2208]: (warning) %u in format string (no. 5) requires 'unsigned int' but the
argument type is 'signed int'.
[src/rxgen/rpc_parse.c:690] -> [src/rxgen/rpc_parse.c:696]: (style) Variable 'tailp' is reassigned
a value before the old one has been used.
[src/rxgen/rpc_parse.c:818]: (style) The scope of the variable 'defp' can be reduced.
[src/rxgen/rpc_parse.c:972]: (style) The scope of the variable 'typecontents' can be reduced.
[src/rxgen/rpc_parse.c:997]: (style) The scope of the variable 'typecontents' can be reduced.
[src/rxgen/rpc_parse.c:1175]: (style) The scope of the variable 'noofparams' can be reduced.
[src/rxgen/rpc_parse.c:1175]: (style) The scope of the variable 'i' can be reduced.
[src/rxgen/rpc_parse.c:1245]: (style) The scope of the variable 'i' can be reduced.
[src/rxgen/rpc_parse.c:1347]: (style) The scope of the variable 'defp1' can be reduced.
[src/rxgen/rpc_parse.c:1557]: (style) The scope of the variable 'i' can be reduced.
[src/rxgen/rpc_parse.c:1944]: (style) The scope of the variable 'defp' can be reduced.
[src/rxgen/rpc_util.h:62] -> [src/rxgen/rpc_parse.c:1528]: (style) Local variable zflag shadows
outer variable
...
```

Standalone utilities

- infer
  - null pointer problems
  - memory leaks
  - coding conventions
  - system APIs

infer

```
src/bucoord/config.c:98: error: NULL_DEREFERENCE
  pointer `tentry` last assigned on line 97 could be null and is dereferenced at line 98, column 5.
  96.       /* tlast now points to the next pointer (or head pointer) we should overwrite */
  97.       tentry = calloc(1, sizeof(struct bc_hostEntry));
  98. >     tentry->name = strdup(aname);
  99.       *tlast = tentry;
  100.       tentry->next = (struct bc_hostEntry *)0;

src/afs/afs_cell.c:108: error: UNINITIALIZED_VALUE
  The value read from code was never initialized.
  106.                    timeout);
  107.
  108. >  if (!hostCount || (code && code != EEXIST))
  109.         /* null out the cellname if the lookup failed */
  110.         afsdb_req.cellname = NULL;
```

infer

```
Summary of the reports

  UNINITIALIZED_VALUE: 681
           DEAD_STORE: 325
      NULL_DEREFERENCE: 188
          MEMORY_LEAK: 173
        RESOURCE_LEAK: 20
       USE_AFTER_FREE: 1
```

Analysis of the OpenAFS code base

- "static-analysis" topic in OpenAFS Gerrit
  - Analysis and patches by Pat Riehecky
    - Memory and resource leaks
    - NULL pointer dereferences
    - Problems with "printf" format strings
    - Boundary conditions (arrays and strings)
    - Uninitialized variables
    - Dead code
  - 25 commits, 19 pending merge

Commits pending approvals

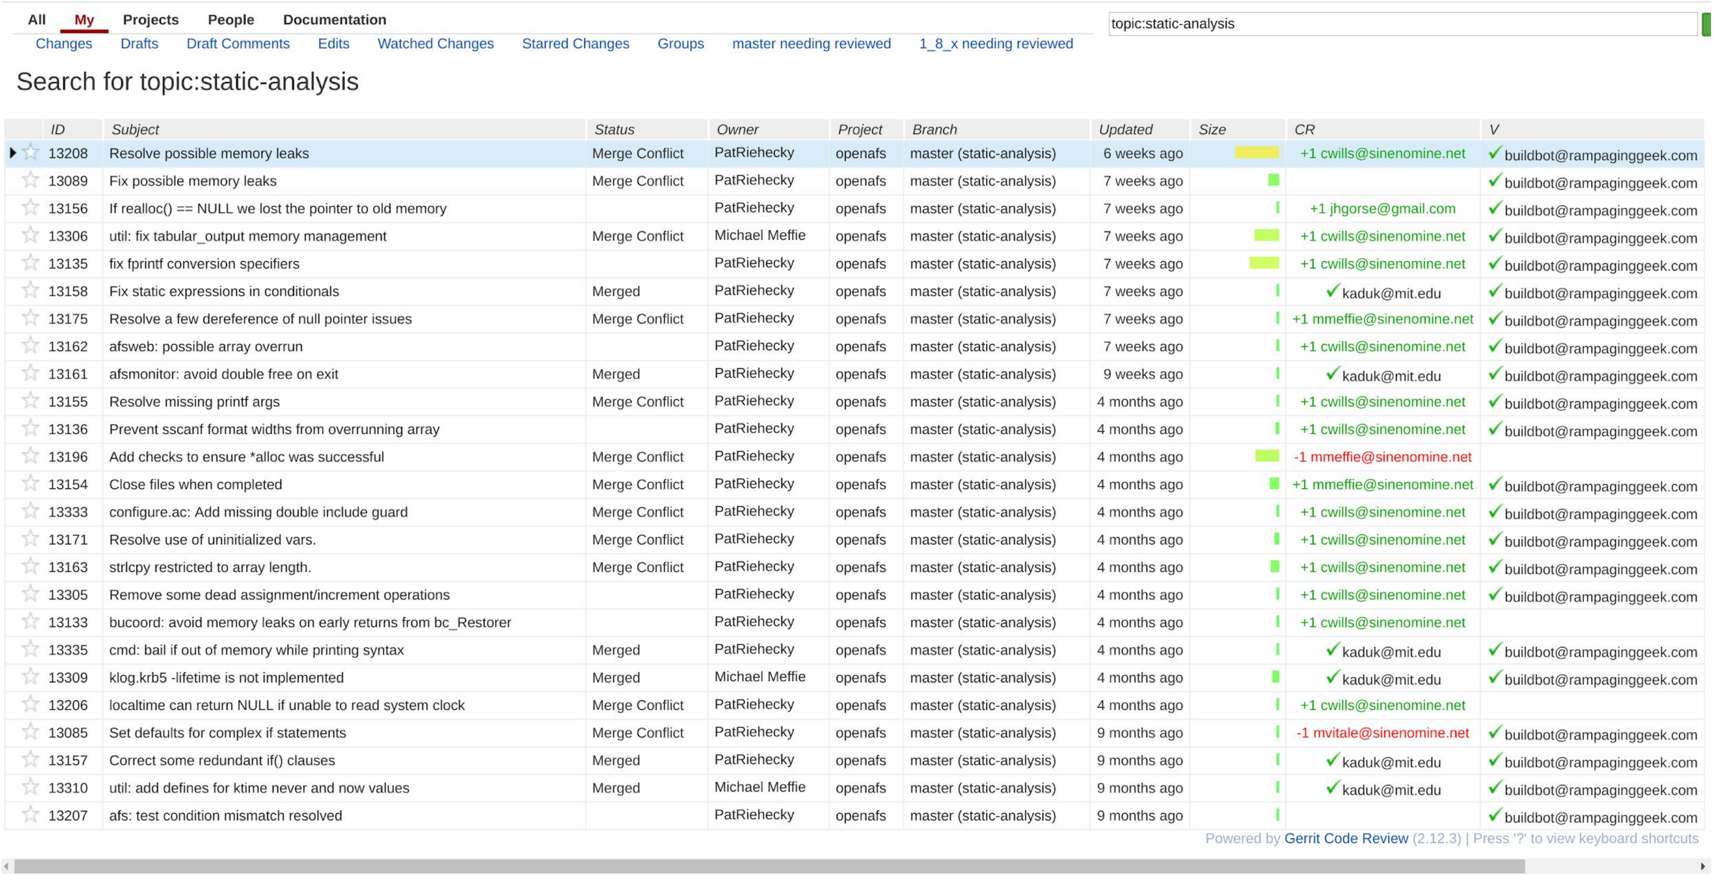

topic:static-analysis

Search for topic:static-analysis

ID	Subject	Status	Owner	Project	Branch	Updated	Size	CR	V
13208	Resolve possible memory leaks	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	6 weeks ago	🟨	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13089	Fix possible memory leaks	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩		✓ buildbot@rampaginggeek.com
13156	If realloc() == NULL we lost the pointer to old memory		PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩	+1 jhgorse@gmail.com	✓ buildbot@rampaginggeek.com
13306	util: fix tabular_output memory management	Merge Conflict	Michael Meffie	openafs	master (static-analysis)	7 weeks ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13135	fix fprintf conversion specifiers		PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13158	Fix static expressions in conditionals	Merged	PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13175	Resolve a few dereference of null pointer issues	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩	+1 mmeffie@sinenomine.net	✓ buildbot@rampaginggeek.com
13162	afsweb: possible array overrun		PatRiehecky	openafs	master (static-analysis)	7 weeks ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13161	afsmonitor: avoid double free on exit	Merged	PatRiehecky	openafs	master (static-analysis)	9 weeks ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13155	Resolve missing printf args	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13136	Prevent sscanf format widths from overrunning array		PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13196	Add checks to ensure *alloc was successful	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	-1 mmeffie@sinenomine.net	
13154	Close files when completed	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 mmeffie@sinenomine.net	✓ buildbot@rampaginggeek.com
13333	configure.ac: Add missing double include guard	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13171	Resolve use of uninitialized vars.	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13163	strlcpy restricted to array length.	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13305	Remove some dead assignment/increment operations		PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	✓ buildbot@rampaginggeek.com
13133	bucoord: avoid memory leaks on early returns from bc_Restorer		PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	
13335	cmd: bail if out of memory while printing syntax	Merged	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13309	klog.krb5 -lifetime is not implemented	Merged	Michael Meffie	openafs	master (static-analysis)	4 months ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13206	localtime can return NULL if unable to read system clock	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	4 months ago	🟩	+1 cwills@sinenomine.net	
13085	Set defaults for complex if statements	Merge Conflict	PatRiehecky	openafs	master (static-analysis)	9 months ago	🟩	-1 mvitale@sinenomine.net	✓ buildbot@rampaginggeek.com
13157	Correct some redundant if() clauses	Merged	PatRiehecky	openafs	master (static-analysis)	9 months ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13310	util: add defines for ktime never and now values	Merged	Michael Meffie	openafs	master (static-analysis)	9 months ago	🟩	✓ kaduk@mit.edu	✓ buildbot@rampaginggeek.com
13207	afs: test condition mismatch resolved		PatRiehecky	openafs	master (static-analysis)	9 months ago	🟩		✓ buildbot@rampaginggeek.com

Powered by Gerrit Code Review (2.12.3) | Press '?' to view keyboard shortcuts

Improving the code base

- **Continued code reviews, adding more eyes.**
  - **Automated tools can’t catch everything**
- **Integrate automated checks into commit and build processes**
  - **adding analysis checks into the buildbot workers**