



---

**Source Code Audit of BIND 9  
for Internet Systems Corporation**  
Final Report and Management Summary

---

2024-02-02

*PUBLIC*

X41 D-SEC GmbH  
Krefelder Str. 123  
D-52070 Aachen  
Amtsgericht Aachen: HRB19989

<https://x41-dsec.de/>  
[info@x41-dsec.de](mailto:info@x41-dsec.de)

<i>Revision</i>	<i>Date</i>	<i>Change</i>	<i>Author(s)</i>
1	2023-06-30	Report First Sprint	E. Sesterhenn, M. Vervier
2	2023-12-05	Final Report and Management Summary	E. Sesterhenn, M. Vervier
3	2024-02-13	Public Release	E. Sesterhenn, M. Vervier



# Contents

<b>1</b>	<b>Executive Summary</b>	<b>4</b>
<b>2</b>	<b>Introduction</b>	<b>6</b>
2.1	Methodology . . . . .	6
2.2	Findings Overview . . . . .	8
2.3	Scope . . . . .	9
2.4	Coverage . . . . .	9
2.5	Recommended Further Tests . . . . .	11
<b>3</b>	<b>Rating Methodology for Security Vulnerabilities</b>	<b>13</b>
<b>4</b>	<b>Results</b>	<b>15</b>
4.1	Findings . . . . .	15
4.2	Informational Notes . . . . .	35
<b>5</b>	<b>About X41 D-Sec GmbH</b>	<b>75</b>
<b>A</b>	<b>Appendix</b>	<b>77</b>
A.1	Fuzz Testing . . . . .	77
A.2	CodeQL Queries . . . . .	92

## Dashboard

### Target

Customer	Internet Systems Corporation
Name	BIND 9
Type	DNS Server
Version	9.19.13 and 9.19.17

### Engagement

Type	Security Source Code Audit
Consultants	2: Eric Sesterhenn and Markus Vervier
Engagement Effort	60 person-days, 2023-05-29 to 2023-11-30

Total issues found 8





Figure 1: Issue Overview (l: Severity, r: CWE Distribution)

# 1 Executive Summary

From June to November 2023, X41 D-Sec GmbH performed a security source code audit against ISC BIND 9 to identify vulnerabilities and weaknesses in the DNS server, tools and libraries.

A total of eight vulnerabilities were discovered during the test by X41. None were rated as critical, none were classified as high severity, two as medium, and six as low. Additionally, 23 issues without a direct security impact were identified.



**Figure 1.1:** Issues and Severity

BIND 9 is an open source DNS server accompanied by libraries and tools that enable users to resolve domain names. It is a critical part of the current Internet infrastructure and therefore subject to scrutiny. Vulnerabilities could have a large impact on the worldwide Internet. Due to this, the codebase was regularly tested before with various fuzzing harnesses and inspected with various static analyzers and analysis tools before this audit and seems to be in good shape.

The test was performed by two experienced security experts between 2023-05-29 and 2023-11-

30. ISC provided the source code and a communication channel with selected developers with whom the communication during the project was flawless.

The most severe issue discovered allows an attacker to crash the named DNS server on some setups by sending maliciously formatted data to the command channel, which will exhaust the available stack memory when parsed. This can be triggered pre-authentication but installations might protect that service using network level access control and is tracked as CVE-2023-3341.

Other security relevant issues are related to the handling of integers such as truncation or overflow of length values. It is recommended to also review the informational issues that are related to hardening or issues that might become security relevant in the future.

The code was audited in a first phase to verify the most obvious attack surface and to validate the absence of commonly present vulnerability classes that are often found in complex C code. In a second phase a deeper investigation occurred for more obscure bugs and logic issues.

In conclusion, the quality of the source code is far above average in comparison with projects of a similar size. While vulnerabilities were found during this audit, the code has been hardened and could serve as an example on how to develop C code with security in mind. X41 attributes the usage of safe patterns for concurrency and safe memory access to be a main contributing factor in elimination or at least reduction of the occurrence of certain bug classes related to unsafe memory operations.

Due to the constantly changing nature of such large code bases and the time-boxed nature of security reviews, X41 recommends to perform independent review on a regular basis. This prevents the manifestation of more complex vulnerabilities and architectural flaws or otherwise unexpected behaviors.

## 2 Introduction

X41 reviewed the source code of the BIND 9 DNS<sup>1</sup> daemon and associated libraries and tools, which perform DNS queries and parse responses to resolve host names.

They are considered sensitive because the DNS is a critical part of the Internet infrastructure and several security mechanisms depend on DNS working correctly. Additionally, the libraries and tools BIND 9 provides are used in a wide range of operating systems and devices.

Attackers could try to abuse implementation-specific issues to gain RCE<sup>2</sup> on the server. Logic bugs might be used for DoS<sup>3</sup> attacks, which attackers could leverage to prevent access to a wide range of systems by disrupting name resolution.

### 2.1 Methodology

The review was based on a source code review against the latest development versions (9.19.13 in the first phase and 9.19.17 in the second phase) of the BIND 9 DNS daemon.

A manual approach for penetration tests and for code reviews is used by X41. This process is supported by tools such as static code analyzers and industry standard web application security tools<sup>4</sup>.

X41 adheres to established standards for source code reviewing and penetration testing. These are in particular the *CERT Secure Coding*<sup>5</sup> standards and the *Study - A Penetration Testing Model*<sup>6</sup> of the German Federal Office for Information Security.

The workflow of source code reviews is shown in figure 2.1. In an initial, informal workshop

---

<sup>1</sup> Domain Name System

<sup>2</sup> Remote Code Execution

<sup>3</sup> Denial of Service

<sup>4</sup> <https://portswigger.net/burp>

<sup>5</sup> <https://wiki.sei.cmu.edu/confluence/display/seccode/SEI+CERT+Coding+Standards>

<sup>6</sup> [https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Studies/Penetration/penetration\\_pdf.pdf?\\_\\_blob=publicationFile&v=1](https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Studies/Penetration/penetration_pdf.pdf?__blob=publicationFile&v=1)

regarding the design and architecture of the application a basic threat model is created. This is used to explore the source code for interesting attack surface and code paths. These are then audited manually and with the help of tools such as static analyzers and fuzzers. The identified issues are documented and can be used in a gap analysis to highlight changes to previous audits.

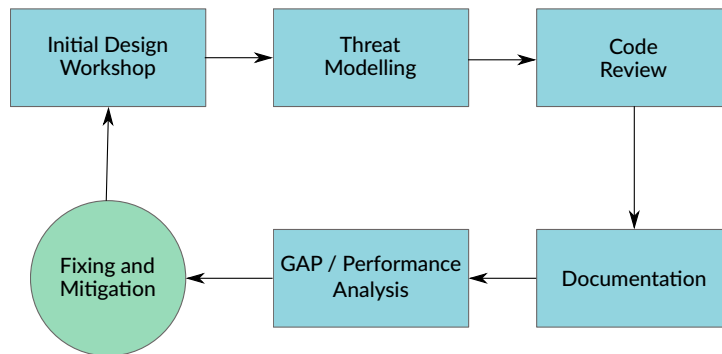


Figure 2.1: Code Review Methodology



## 2.2 Findings Overview

DESCRIPTION	SEVERITY	ID	GL	REF
Integer Overflow in http_alloc()	LOW	BND-CA-23-01	4120	4.1.1
Integer Overflow in malloc()	LOW	BND-CA-23-02	4121	4.1.2
Integer Overflow in resize()	LOW	BND-CA-23-03	4122	4.1.3
Buffer Overflow in process_request()	LOW	BND-CA-23-04	4124	4.1.4
Negative Content-Length Leads to abort()	LOW	BND-CA-23-05	4125	4.1.5
Stack Exhaustion in Command Channel	MEDIUM	BND-CA-23-06	4152	4.1.6
OpenSSL Error Queue Not Emptied	LOW	BND-CA-23-07	4157	4.1.7
Unbounded Token Parsing	MEDIUM	BND-CA-23-08	4339	4.1.8
Memory Leak Due to realloc() Misuse	NONE	BND-CA-23-100	4174	4.2.1
ISC Memory API	NONE	BND-CA-23-101	n/a	4.2.2
Journal File Handling Missing Sanity Checks	NONE	BND-CA-23-102	4175	4.2.3
Unchecked malloc()	NONE	BND-CA-23-103	4077	4.2.4
Stack Exhaustion in Config Parser	NONE	BND-CA-23-104	4176	4.2.5
Connection Flags Mixup	NONE	BND-CA-23-105	4126	4.2.6
Use of Magic Numbers	NONE	BND-CA-23-106	n/a	4.2.7
NULL Pointer Dereference on Wrong API Usage	NONE	BND-CA-23-107	4177	4.2.8
Invalid Free in Low Memory Situation	NONE	BND-CA-23-108	4179	4.2.9
Possible Truncation in dns_keymgr_status()	NONE	BND-CA-23-109	4180	4.2.10
Newline and ANSI Escape Code Injection via CC	NONE	BND-CA-23-110	4181	4.2.11
Name Buffer Truncation	NONE	BND-CA-23-111	4186	4.2.12
Misaligned Structure Causes Exception	NONE	BND-CA-23-112	4187	4.2.13
Race in dns_tsigkey_find()	NONE	BND-CA-23-113	4182	4.2.14
Pointers Dereferenced before Being Checked	NONE	BND-CA-23-114	4432	4.2.15
Files Created with World Read/Write Permissions	NONE	BND-CA-23-115	4443	4.2.16
Locking Inconsistencies in Cache Implementation	NONE	BND-CA-23-116	4340	4.2.17
Endless Loop via GENERATE	NONE	BND-CA-23-117	4353	4.2.18
Endless Loop via INCLUDE	NONE	BND-CA-23-118	4357	4.2.19
Supplied Buffer Too Large	NONE	BND-CA-23-119	4433	4.2.20
Dead Code in DNSTAP Helper	NONE	BND-CA-23-120	4406	4.2.21
Unused AES Functions	NONE	BND-CA-23-121	4421	4.2.22
Usage of isc_safe_memwipe()	NONE	BND-CA-23-122	4435	4.2.23

**Table 2.1:** Security-Relevant Findings, GL lists the ISC GitLab IDs

## 2.3 Scope

The code audit was performed on versions 9.19.13<sup>7</sup> and 9.19.17<sup>8</sup> of the ISC BIND 9 DNS daemon, libraries and tools. The source code consists of around 275.000 lines of C code including tests.

Third party libraries that BIND 9 depends on, such as OpenSSL<sup>9</sup>, LMDB<sup>10</sup>, Libxml<sup>11</sup>, JSON-C<sup>12</sup> or nhttp2<sup>13</sup>, were not covered by this audit. Contributed code in `contrib/` was not inspected in depth.

A Mattermost chat group was created for direct communication between the developers and the testers. Additionally, issues were reported during the test in the ISC GitLab<sup>14</sup> as confidential issues.

## 2.4 Coverage

A security assessment attempts to find the most important or sometimes as many of the existing problems as possible, though it is practically never possible to rule out the possibility of additional weaknesses being found in the future.

The time allocated to X41 for this assessment was sufficient to yield a good coverage of the given scope.

The code was inspected for security issues such as integer overflows, OOB<sup>15</sup> reads and writes, hash table degradation, reference counter overflows, use-after-free and double free bugs. Beyond looking for common C level implementation bugs, the team investigated the codebase for bug patterns that might affect BIND 9 specifically.

The use of `isc_buffer_*`() and `isc_mem_*`() primitives were audited for misuse. Possible log file injections via `isc_log_write()` were checked. The DNS cookie<sup>16</sup> implementation was audited for logic issues. Access control and separation of trust domain issues were investigated. Race conditions and issues leading to unbound recursive calls were investigated. The conditions of the assertion functions such as `ISC_REQUIRE()` were audited for integer overflows and truncations. The handling and deletion of key material was inspected for common errors. Type casts were checked

<sup>7</sup> <https://downloads.isc.org/isc/bind9/9.19.13/>

<sup>8</sup> <https://downloads.isc.org/isc/bind9/9.19.17/>

<sup>9</sup> <https://www.openssl.org/>

<sup>10</sup> <http://www.lmdb.tech/doc/>

<sup>11</sup> <https://gitlab.gnome.org/GNOME/libxml2/-/wikis/home>

<sup>12</sup> <https://github.com/json-c/json-c>

<sup>13</sup> <https://nhttp2.org/>

<sup>14</sup> <https://gitlab.isc.org/isc-projects/bind9>

<sup>15</sup> Out-of-Bounds

<sup>16</sup> <https://www.ietf.org/rfc/rfc7873.txt>

for errors such as truncation or type confusions. EDNS<sup>17</sup> padding as well as IXFR<sup>18</sup> and AXFR<sup>19</sup> handling was inspected for logic issues.

Handling of locks and mutexes was investigated for potential race conditions, especially related to concurrent query processing and other asynchronous operations heavily used in the code. Concurrency and access patterns such as event loops (`libuv`), lockfree access patterns, and userspace RCU<sup>20</sup> (`liburcu`) were audited.

X41 also investigated advanced scenarios and issues such as crashes due to incorrect memory alignment and other more uncommon problems related to compiler optimizations.

Besides manual analysis, dedicated code analysis tools were used during the process that are described in the following subsections.

## 2.4.1 Static Analysis

The source code was analyzed using several static analyzers such as `joern.io`<sup>21</sup>, `semgrep`<sup>22</sup>, `weggli`<sup>23</sup>, `cppcheck`<sup>24</sup>, `gcc 12 analyzer`<sup>25</sup> and the Clang Static Analyzer<sup>26</sup> using the built-in queries and manually created ones. Additionally, the warnings on the Intel oneAPI compilers were analyzed.

To support manual analysis and to make the analysis more efficient and repeatable, the team built CodeQL<sup>27</sup> queries that can find data-flows leading to vulnerabilities. In particular the following were investigated:

- Query to find correct usage of mutexes to verify manual review results (see appendix A.2.1)
- Query to find tainted length values used in `isc_buffer_get*()` function calls (see appendix A.2.2)
- Query to find `printf()` family function calls that use tainted data (see appendix A.2.3) with the intention to identify format string issues
- Query to find tainted calls of ISC memory wrapper functions (see appendix A.2.4)
- Query to find unsafe dereference related to RCU (see appendix A.2.5)
- Query to identify recursive calls of type `A()->B()->A()` (see appendix A.2.6)

---

<sup>17</sup> Extended DNS

<sup>18</sup> Incremental Zone Transfer

<sup>19</sup> Asynchronous Full Transfer Zone

<sup>20</sup> Read, Copy, Update

<sup>21</sup> <https://joern.io/>

<sup>22</sup> <https://semgrep.dev/>

<sup>23</sup> <https://github.com/weggli-rs/weggli>

<sup>24</sup> <https://cppcheck.sourceforge.io/>

<sup>25</sup> <https://gcc.gnu.org/wiki/StaticAnalyzer>

<sup>26</sup> <https://clang-analyzer.llvm.org/>

<sup>27</sup> <https://codeql.github.com>

## 2.4.2 Fuzz Testing

Additionally, fuzz testing using the already existing harnesses shipped with bind<sup>28</sup><sup>29</sup> and honggfuzz<sup>30</sup> was performed and additional fuzz harnesses were created. These could be implemented quickly due to the already working setup and were added whenever a promising function was encountered during the audit.

This led to the creation of the following fuzzing harnesses:

- *isc\_url\_parse()* (see appendix A.1.1)
- *phr\_parse\_request()* (see appendix A.1.2)
- *isc\_regex\_validate()* (see appendix A.1.3)
- *isc\_utf8\_valid()* (see appendix A.1.4)
- *isccc\_cc\_fromwire()* (see appendix A.1.5)
- *irs\_resconf\_load()* (see appendix A.1.6)
- dig/host AFL++ ARGV Harness (see appendix A.1.7)

Furthermore, some speedups were implemented and memory leaks in existing harnesses removed (see appendix A.1.9).

Additionally, as an alternative strategy, template-based fuzzing using Scapy<sup>31</sup> was performed using a straightforward approach (see appendix A.1.8). Some stress-testing was applied to the bad-cache implementation to identify possible errors in concurrent accesses (see appendix A.1.10).

The fuzzing did not trigger security relevant bugs, but caused parsing errors to show up in the logging output. These errors were regarded as being handled gracefully.

Suggestions for next steps in securing this scope can be found in section 2.5.

## 2.5 Recommended Further Tests

After investigating the most exposed and obvious attack surface, X41 recommends to perform a second iteration of the review that concentrates the efforts on more complex attack vectors and a deeper analysis of the code base.

<sup>28</sup> <https://gitlab.isc.org/isc-projects/bind9/-/tree/main/fuzz>

<sup>29</sup> <https://gitlab.isc.org/isc-projects/bind9/-/blob/main/bin/named/fuzz.c>

<sup>30</sup> <https://github.com/google/honggfuzz/tree/master/examples/bind>

<sup>31</sup> <https://scapy.net>

Further steps would be the audit of ACL<sup>32</sup> and limit enforcing when these are configured.

Another avenue for attacks that are hard to identify and mitigate can be compiler or operating system level bugs or unexpected behavior, which can lead to vulnerabilities that affect BIND 9. An example for such bugs was even identified in this initial review and is described in the informational finding 4.2.13.

Code that resides in `contrib/` was not inspected in-depth for this audit. X41 recommends to audit this code as well.

X41 recommends to mitigate the issues described in this report. Afterwards, CVE<sup>33</sup> IDs<sup>34</sup> should be requested and customers should be informed (e.g. via a changelog or a special note for issues with higher severity) to ensure that they can make an informed decision about upgrading or other possible mitigations.

---

<sup>32</sup> Access Control List

<sup>33</sup> Common Vulnerabilities and Exposures

<sup>34</sup> Identifiers

## 3 Rating Methodology for Security Vulnerabilities

Security vulnerabilities are given a purely technical rating by the testers as they are discovered during the test. Business factors and financial risks for Internet Systems Corporation are beyond the scope of a penetration test which focuses entirely on technical factors. Yet technical results from a penetration test may be an integral part of a general risk assessment. A penetration test is based on a limited time frame and only covers vulnerabilities and security issues which have been found in the given time, there is no claim for full coverage.

In total, five different ratings exist, which are as follows:

### Severity Rating

None
Low
Medium
High
Critical

A low rating indicates that the vulnerability is either very hard for an attacker to exploit due to special circumstances, or that the impact of exploitation is limited, whereas findings with a medium rating are more likely to be exploited or have a higher impact. High and critical ratings are assigned when the testers deem the prerequisites realistic or trivial and the impact significant or very significant.

Findings with the rating 'none' are called informational findings and are related to security hardening, affect functionality, or other topics that are not directly related to security. X41 recommends to mitigate these issues as well, because they often become exploitable in the future. Doing so will strengthen the security of the system and is recommended for defense in depth.

## Common Weakness Enumeration

The CWE<sup>1</sup> is a set of software weaknesses that allows the categorization of vulnerabilities and weaknesses in software. If applicable, X41 provides the CWE-ID for each vulnerability that is discovered during a test.

CWE is a very powerful method to categorize a vulnerability and to give general descriptions and solution advice on recurring vulnerability types. CWE is developed by MITRE<sup>2</sup>. More information can be found on the CWE website at <https://cwe.mitre.org/>.

---

<sup>1</sup> Common Weakness Enumeration

<sup>2</sup> <https://www.mitre.org>

## 4 Results

This chapter describes the results of this test. The security-relevant findings are documented in Section 4.1. Additionally, findings without a direct security impact are documented in Section 4.2.

### 4.1 Findings

The following subsections describe findings with a direct security impact that were discovered during the test.

#### 4.1.1 BND-CA-23-01: Integer Overflow in `http_calloc()`

---

<i>Severity:</i>	LOW
<i>CWE:</i>	190 – Integer Overflow or Wraparound
<i>GitLab Issue:</i>	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4120">https://gitlab.isc.org/isc-projects/bind9/-/issues/4120</a>
<i>Affected Component:</i>	lib/isc/netmgr/http.c:http_calloc()

---

##### 4.1.1.1 Description

In `http_calloc()`, two `size_t` variables are multiplied, which might lead to an overflow during the calculation. When that happens, less memory than expected by the caller is allocated, which might lead to a heap-based buffer overflow (see listing 4.1).



---

```
1 static void *
2 http_calloc(size_t n, size_t sz, isc_mem_t *mctx) {
3     const size_t msize = n * sz;
4     void *data = isc_mem_allocate(mctx, msize);
5
6     memset(data, 0, msize);
7     return (data);
8 }
```

---

**Listing 4.1:** Integer Overflow in http\_calloc()

Since this overflow happens in a core function with many possible callers, this is not considered informational.

#### 4.1.1.2 Solution Advice

X41 recommends to add an overflow check for the multiplication and return *NULL* on overflow similar to the check implemented in *isc\_uv\_calloc()*.

## 4.1.2 BND-CA-23-02: Integer Overflow in mallocx()

---

Severity:	LOW
CWE:	190 – Integer Overflow or Wraparound
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4121">https://gitlab.isc.org/isc-projects/bind9/-/issues/4121</a>
Affected Component:	lib/isc/jemalloc_shim.h:mallocx()

---

### 4.1.2.1 Description

In `mallocx()`, an addition is performed to create space to store `size_info` in front of the allocated memory. This addition is not checked and could overflow for large values of `size`. When that happens, less memory than expected by the caller is allocated, which might lead to a heap-based buffer overflow (see listing 4.2).

---

```
1 static inline void *
2 mallocx(size_t size, int flags) {
3     void *ptr = NULL;
4
5     size_info *si = malloc(size + sizeof(*si));
6     INSIST(si != NULL);
7
8     si->size = size;
9     ptr = &si[1];
10
11     if ((flags & MALLOCX_ZERO) != 0) {
12         memset(ptr, 0, size);
13     }
14
15     return (ptr);
16 }
```

---

Listing 4.2: Integer Overflow in mallocx()

Since this overflow happens in a core function, this is not considered informational.

### 4.1.2.2 Solution Advice

X41 recommends to add an overflow check for the addition and return `NULL` on overflow similar to the check implemented in `isc__uv_calloc()`.

### 4.1.3 BND-CA-23-03: Integer Overflow in `resize()`

---

Severity:	LOW
CWE:	190 – Integer Overflow or Wraparound
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4122">https://gitlab.isc.org/isc-projects/bind9/-/issues/4122</a>
Affected Component:	lib/isc/heap.c:resize()

---

#### 4.1.3.1 Description

In `resize()` of the heap implementation, the heap is grown, but no integer overflow checks are performed on the addition for `new_size` or the multiplication for the actual size (see listing 4.3).

---

```
1 static void
2 resize(isc_heap_t *heap) {
3     void **new_array;
4     unsigned int new_size;
5
6     REQUIRE(VALID_HEAP(heap));
7
8     new_size = heap->size + heap->size_increment;
9     new_array = isc_mem_get(heap->mctx, new_size * sizeof(void *));
```

---

Listing 4.3: Integer Overflow in `resize()`

Since the overflow can only occur after previous resizes succeed, which occur in steps of `heap->size_increment`, the likelihood of this being exploitable is low.

#### 4.1.3.2 Solution Advice

X41 recommends to add an overflow check for the addition and multiplication similar to the check implemented in `isc_uv_calloc()`.

## 4.1.4 BND-CA-23-04: Buffer Overflow in process\_request()

---

Severity:	LOW
CWE:	121 – Stack-based Buffer Overflow
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4124">https://gitlab.isc.org/isc-projects/bind9/-/issues/4124</a>
Affected Component:	lib/isc/httpd.c:process_request()

---

### 4.1.4.1 Description

A stack-based buffer overflow exists in `lib/isc/httpd.c`, in function `process_request()`. When the `If-Modified-Since` HTTP<sup>1</sup> header is long enough, the `memmove()` writes out of bounds and corrupts stack data as seen in listing 4.4.

---

```

1 } else if (name_match(header, "If-Modified-Since")) {
2     char timestamp[ISC_FORMATHTTPSTAMP_SIZE + 1];
3     memmove(timestamp, header->value, header->value_len);
4     timestamp[header->value_len] = 0;
5
6     /* Ignore the value if it can't be parsed */
7     (void)isc_time_parsehttpstamp(
8         timestamp, &httpd->if_modified_since);
9 }

```

---

Listing 4.4: Buffer Overflow in process\_request()

With the default compilation settings, the overflow can overwrite the `headers` and `num_headers` values on the stack as well as canaries and return addresses.

The code in listing 4.5 can be used to trigger the overflow:

---

```

1 import socket
2
3 HOST = "127.0.0.1"
4 PORT = 8080
5
6 depth = 100
7 payload = b'A' * depth
8
9 with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as s:
10     data = b'GET / HTTP/1.1\r\nHost: localhost:8080\r\n'

```

---

<sup>1</sup> HyperText Transfer Protocol

```
11     data = b''.join([data, b'If-Modified-Since: ', payload, b'\r\n'])
12     data = b''.join([data, b'\r\n'])
13     s.connect((HOST, PORT))
14     s.sendall(data)
15     s.recv(10)
```

---

#### Listing 4.5: HTTP Overflow PoC

Since the statistics server should only be accessible by trusted parties and requires authentication, this issue is considered having a low impact. This is reflected in the BIND 9 documentation<sup>2</sup> as well:

An issue in the statistics channel would be considered a security issue only if it could be exploited by unprivileged users circumventing the access control list. In other words, any issue in the statistics channel that could be used to access information unavailable otherwise, or to crash named, is not considered a security issue if it can be avoided through the use of a secure configuration.

#### 4.1.4.2 Solution Advice

X41 recommends to add a size check to verify that the source buffer fits into *timestamp*.

---

<sup>2</sup>[https://bind9.readthedocs.io/en/v9\\_18\\_12/reference.html#statistics-channels-block-definition-and-usage](https://bind9.readthedocs.io/en/v9_18_12/reference.html#statistics-channels-block-definition-and-usage)

## 4.1.5 BND-CA-23-05: Negative Content-Length Leads to abort()

---

Severity:	LOW
CWE:	617 – Reachable Assertion
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4125">https://gitlab.isc.org/isc-projects/bind9/-/issues/4125</a>
Affected Component:	lib/isc/httpd.c:process_request()

---

### 4.1.5.1 Description

In `lib/isc/httpd.c`, in `process_request()` a HTTP request is parsed and the HTTP *Content-Length* header evaluated (see listing 4.6). If the integer in that field is negative or large enough, the addition to `httpd->consume` might overflow and trigger an `INSIST(httpd->consume != 0)` assertion in `prepare_response()`. This will cause an `abort()` and stop the process.

---

```

1  ssize_t content_len = 0;
2  bool keep_alive = false;
3
4  isc_time_set(&httpd->if_modified_since, 0, 0);
5
6  for (size_t i = 0; i < num_headers; i++) {
7      struct phr_header *header = &headers[i];
8
9      if (name_match(header, "Content-Length")) {
10         char *endptr;
11         content_len = (size_t) strtoul(header->value, &endptr,
12                                     10);
13         /* Consistency check, if we consumed all numbers */
14         if ((header->value + header->value_len) != endptr) {
15             return (ISC_R_RANGE);
16         }
17     ...
18     /* Consume the request's data, which we do not use. */
19     httpd->consume += content_len;
20     ...

```

---

Listing 4.6: Negative Content-Length Leads to abort()

This should only be possible when the attacker is also able to send enough (`content_len` bytes) data to the service.

Since the statistics server should only be accessible by trusted parties and requires authentication this issue is considered having a low impact.

#### 4.1.5.2 Solution Advice

X41 recommends to add a size sanity check or catch the overflow during the addition.

## 4.1.6 BND-CA-23-06: Stack Exhaustion in Command Channel

---

Severity:	<b>MEDIUM</b>
CWE:	1325 – Improperly Controlled Sequential Memory Allocation
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4152">https://gitlab.isc.org/isc-projects/bind9/-/issues/4152</a>
Affected Component:	lib/isccc:value_fromwire()

---

### 4.1.6.1 Description

The command channel library libisccc parses the full TCP<sup>3</sup> packet before performing the actual authentication. The packet is structured into binary data, tables and lists. By structuring the packet maliciously, repeated recursive calls to *value\_fromwire()* and *table\_fromwire()* can be triggered. These functions are shown in listing 4.7.

---

```

1  static isc_result_t
2  value_fromwire(isc_region_t *source, isccc_sexpr_t **valuep) {
3      unsigned int msgtype;
4      uint32_t len;
5      isccc_sexpr_t *value;
6      isccc_region_t active;
7      isc_result_t result;
8
9      if (REGION_SIZE(*source) < 1 + 4) {
10         return (ISC_R_UNEXPECTEDEND);
11     }
12     GET8(msgtype, source->rstart);
13     GET32(len, source->rstart);
14     if (REGION_SIZE(*source) < len) {
15         return (ISC_R_UNEXPECTEDEND);
16     }
17     active.rstart = source->rstart;
18     active.rend = active.rstart + len;
19     source->rstart = active.rend;
20     if (msgtype == ISCCC_CCMSTYPE_BINARYDATA) {
21         value = isccc_sexpr_frombinary(&active);
22         if (value != NULL) {
23             *valuep = value;
24             result = ISC_R_SUCCESS;
25         } else {
26             result = ISC_R_NOMEMORY;
27         }
28     } else if (msgtype == ISCCC_CCMSTYPE_TABLE) {
29         result = table_fromwire(&active, NULL, 0, valuep);

```

<sup>3</sup> Transmission Control Protocol



```
30     } else if (msgtype == ISCCC_CCMSGTYPE_LIST) {
31         result = list_fromwire(&active, valuep);
32     } else {
33         result = ISCCC_R_SYNTAX;
34     }
35
36     return (result);
37 }
38
39 static isc_result_t
40 table_fromwire(isc_region_t *source, isc_region_t *secret,
41               uint32_t algorithm, iscsexpr_t **alistp) {
42     char key[256];
43     uint32_t len;
44     isc_result_t result;
45     iscsexpr_t *alist, *value;
46     bool first_tag;
47     unsigned char *checksum_rstart;
48
49     REQUIRE(alistp != NULL && *alistp == NULL);
50
51     checksum_rstart = NULL;
52     first_tag = true;
53     alist = isccc_alist_create();
54     if (alist == NULL) {
55         return (ISC_R_NOMEMORY);
56     }
57
58     while (!REGION_EMPTY(*source)) {
59         GET8(len, source->rstart);
60         if (REGION_SIZE(*source) < len) {
61             result = ISC_R_UNEXPECTEDEND;
62             goto bad;
63         }
64         GET_MEM(key, len, source->rstart);
65         key[len] = '\0'; /* Ensure NUL termination. */
66         value = NULL;
67         result = value_fromwire(source, &value);
68         if (result != ISC_R_SUCCESS) {
69             goto bad;
70         }
71     }
...

```

Listing 4.7: Stack Exhaustion in CC

On the test machine, each iteration between two calls to *table\_fromwire()* required 432 bytes of stack memory. If enough calls can be triggered, this might lead to exhaustion of the available stack memory and cause a segmentation fault. The amount of iterative calls is limited by the

parameter to `isccc_ccmsg_setmaxsize()` which is `32768` for BIND named. This is not enough to exhaust the 8MB of process stack usually configured on Linux-based systems. It can be triggered on Microsoft Windows systems where the stack size is 1MB by default. Systems where `ulimit`<sup>4</sup> is used to restrict stack usage are affected as well. BIND 9.18 on FreeBSD is also affected by this.

It was tested with the debug build of the latest<sup>5</sup> BIND 9 version that supports MS Windows, the error is shown in listing 4.8:

---

```
1 Unhandled exception at 0x00007FFCB43676C0 (libisccc.dll) in named.exe: 0xC00000FD: Stack overflow
   ↳ (parameters: 0x0000000000000001, 0x000000F149503F10).
```

---

#### Listing 4.8: Stack Exhaustion on Microsoft Windows

The issue can be triggered by sending the maliciously formatted data to the `rndc` port as shown in listing 4.9:

---

```
1 import socket
2
3 HOST = "127.0.0.1"
4 PORT = 953
5
6 depth = 4500
7 # should not be more than isccc_ccmsg_setmaxsize(&conn->ccmsg, 32768);
8 total_len = 10 + (depth * 7) - 6
9
10 with socket.socket(socket.AF_INET, socket.SOCK_STREAM) as s:
11     data = b''.join([
12         total_len.to_bytes(4, 'big'), # <total length>
13         b'\x00\x00\x00\x01',        # <version>
14         b'\x01\x41',                # <size><name>
15     ])
16
17     for i in range(depth, 0, -1):
18         l = (i - 1) * 7
19         t = b''.join([
20             b'\x02',                # ISCCC_CCMSGTYPE_TABLE
21             l.to_bytes(4, 'big'),    # <size>
22             b'\x01\x41',            # <size><name>
23         ])
24         data = b''.join([data, t])
25
```

---

<sup>4</sup><https://man7.org/linux/man-pages/man3/ulimit.3.html>

<sup>5</sup><https://downloads.isc.org/isc/bind9/cur/9.16/BIND9.16.41.debug.x64.zip>

```
26     s . connect ((HOST, PORT))
27     s . sendall (data)
```

---

#### Listing 4.9: Stack Exhaustion in CC PoC

Usually TCP port 953 should not be reachable by untrusted parties due to firewalling and the BIND 9 configuration, but since authentication is performed as well, this might not always be the case.

#### 4.1.6.2 Solution Advice

X41 recommends to further reduce the size passed to `isccc_ccmsg_setmaxsize()` or limit the number of recursive iterations. Another way to reduce the impact of this would be to dynamically allocate the `key` buffer on the heap instead of the stack.

## 4.1.7 BND-CA-23-07: OpenSSL Error Queue Not Emptied

---

Severity:	LOW
CWE:	401 – Improper Release of Memory Before Removing Last Reference ('Memory Leak')
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4157">https://gitlab.isc.org/isc-projects/bind9/-/issues/4157</a>
Affected Component:	lib/dns/rdata.c:check_private()

---

### 4.1.7.1 Description

OpenSSL 3 uses a queue to track errors<sup>6</sup>. The memory for errors in this queue is allocated via `isc__tls_malloc_ex()` and freed in three cases:

1. When the queue is emptied via `ERR_get_error_all()` and similar<sup>7</sup>, for example via a call to `dst__openssl_toreresult()`
2. Via `OPENSSL_cleanup()` on a clean shutdown during library unload or when called from `isc__tls_shutdown()`
3. `isc__tls_shutdown()` when the memory pool is destroyed via `isc_mem_destroy()`

Some parts of the code clean that queue directly after an error happened (see listing 4.10).

---

```

1 status = EVP_PKEY_fromdata(
2     ctx, retpkey, private ? EVP_PKEY_KEYPAIR : EVP_PKEY_PUBLIC_KEY,
3     params);
4 if (status != 1) {
5     DST_RET(dst__openssl_toreresult2("EVP_PKEY_fromdata",
6                                     DST_R_OPENSSLFAILURE));
7 }
8 ret = ISC_R_SUCCESS;
```

---

**Listing 4.10:** OpenSSL Error Queue Handling

When tracing the binary via `valgrind`<sup>8</sup> these leaks do not show up due to the cleanup on library unload.

<sup>6</sup> <https://subscription.packtpub.com/book/web-development/9781800560345/9/ch091v11sec65/understanding-the-openssl-error-queue>

<sup>7</sup> [https://www.openssl.org/docs/man3.1/man3/ERR\\_get\\_error.html](https://www.openssl.org/docs/man3.1/man3/ERR_get_error.html)

<sup>8</sup> <https://valgrind.org/>

For the following trace (see listing 4.11) no cleanup seems to occur and therefore there should be a memory leak.

---

```

1 ==137103== 16 bytes in 1 blocks are still reachable in loss record 33 of 718
2 ==137103==   at 0x48407B4: malloc (vg_replace_malloc.c:381)
3 ==137103==   by 0x4B4A6A9: mallocx (jemalloc_shim.h:65)
4 ==137103==   by 0x4B4A858: mem_get (mem.c:304)
5 ==137103==   by 0x4B4BF3A: isc_mem_allocate (mem.c:785)
6 ==137103==   by 0x4B5FB11: isc_tls_malloc_ex (tls.c:129)
7 ==137103==   by 0x54F1AF5: CRYPTO_strdup (in /usr/lib/x86_64-linux-gnu/libcrypto.so.3)
8 ==137103==   by 0x54A8167: ERR_set_debug (in /usr/lib/x86_64-linux-gnu/libcrypto.so.3)
9 ==137103==   by 0x53AFACC: ASN1_get_object (in /usr/lib/x86_64-linux-gnu/libcrypto.so.3)
10 ==137103==   by 0x53A9040: d2i_ASN1_OBJECT (in /usr/lib/x86_64-linux-gnu/libcrypto.so.3)
11 ==137103==   by 0x4950D55: check_private (rdata.c:624)
12 ==137103==   by 0x49865A5: fromwire_rrsig (rrsig_46.c:345)
13 ==137103==   by 0x49ADF93: dns_rdata_fromwire (rdata.c:824)

```

---

**Listing 4.11:** Allocation Due to ERR\_set\_debug()

A patch was applied to some fuzzing harnesses to be able to catch these errors during fuzz testing (see listing 4.12).

---

```

1 --- dns_message_parse.c.orig    2023-06-21 06:21:00.172793050 +0200
2 +++ dns_message_parse.c        2023-06-21 20:19:01.586207696 +0200
3 @@ -25,6 +25,10 @@
4
5 #include <dns/message.h>
6
7 +#include <assert.h>
8 +#include <openssl/err.h>
9 +
10 +
11 #include "fuzz.h"
12
13 bool debug = false;
14 @@ -180,6 +184,19 @@ cleanup:
15     if (message != NULL) {
16         dns_message_detach(&message);
17     }
18 +
19 +     if (ERR_peek_error() != 0) {
20 +         const char *file = NULL;
21 +         int line = 0;
22 +         const char *func = NULL;
23 +         const char *data2 = NULL;
24 +         int flags = 0;
25 +         long x = 1;
26 +         while (x != 0) {

```

```
26 +         x= ERR_get_error_all(&file, &line, &func, &data2, &flags);
27 +         printf("%s:%d:%s -> %s\n", file, line, func, data2);
28 +     }
29
30 +     abort();
31 + }
32     return (0);
33 }
```

**Listing 4.12:** Enable Fuzzer to Catch OpenSSL Error Queue Leaks

Another place where similar code could be added to help debugging would be `isc__tls_shutdown()` in `lib/isc/tls.c` before the appropriate cleanup functions are called.

#### 4.1.7.2 Solution Advice

X41 recommends to always empty the OpenSSL error queue after OpenSSL errors occurred by calling `ERR_clear_error()`<sup>9</sup> or similar.

<sup>9</sup>[https://www.openssl.org/docs/manmaster/man3/ERR\\_clear\\_error.html](https://www.openssl.org/docs/manmaster/man3/ERR_clear_error.html)

## 4.1.8 BND-CA-23-08: Unbounded Token Parsing

---

Severity:	<b>MEDIUM</b>
CWE:	789 – Uncontrolled Memory Allocation
GitLab Issue:	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4339">https://gitlab.isc.org/isc-projects/bind9/-/issues/4339</a>
Affected Component:	lib/isc/lex.c:pushandgrow()

---

### 4.1.8.1 Description

During the code review it was found that the lexer code in file `lib/isc/lex.c` is affected by crashes and infinite loops when parsing very long strings.

In particular the function `pushandgrow()`, which is used to enlarge buffers during parsing, can lead to a failed allocation on 32-bit architectures and a subsequent `abort()` call:

---

```
1 static isc_result_t
2 pushandgrow(isc_lex_t *lex, inputsource *source, int c) {
3     if (isc_buffer_availablelength(source->pushback) == 0) {
4         isc_buffer_t *tbuf = NULL;
5         unsigned int oldlen;
6         isc_region_t used;
7         isc_result_t result;
8
9         oldlen = isc_buffer_length(source->pushback);
10        isc_buffer_allocate(lex->mctx, &tbuf, oldlen * 2); // MARK integer overflow
11        isc_buffer_usedregion(source->pushback, &used);
12        result = isc_buffer_copyregion(tbuf, &used);
13        INSIST(result == ISC_R_SUCCESS);
14        tbuf->current = source->pushback->current;
15        isc_buffer_free(&source->pushback);
16        source->pushback = tbuf;
17    }
18    isc_buffer_putuint8(source->pushback, (uint8_t)c);
19    return (ISC_R_SUCCESS);
20 }
```

---

Listing 4.13: Unbounded Buffer Size Increase

When parsing a zone file with a very long token, the lexer will enter the function `pushandgrow()` repeatedly and double the buffer allocation each time the buffer space is exhausted. This quickly leads to a failed allocation on 32-bit architectures, since the maximum single allocation size is at most 4GB due to the limited address space. In practice, it is even much less for most allocators.

An example is shown in the following listing:

```
1  $ ./bin/dnssec/dnssec-signzone-gdb db.long
2  GNU gdb (Debian 10.1-1.7) 10.1.90.20210103-git
3  Copyright (C) 2021 Free Software Foundation, Inc.
4  License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
5  This is free software: you are free to change and redistribute it.
6  There is NO WARRANTY, to the extent permitted by law.
7  Type "show copying" and "show warranty" for details.
8  This GDB was configured as "x86_64-linux-gnu".
9  Type "show configuration" for configuration details.
10 For bug reporting instructions, please see:
11 <https://www.gnu.org/software/gdb/bugs/>.
12 Find the GDB manual and other documentation resources online at:
13     <http://www.gnu.org/software/gdb/documentation/>.
14
15 For help, type "help".
16 Type "apropos word" to search for commands related to "word"...
17 Reading symbols from /home/user/src-new/bind9/bin/dnssec/.libs/dnssec-signzone...
18 gdb-peda$ r
19 Starting program: /home/user/src-new/bind9/bin/dnssec/.libs/dnssec-signzone db.long
20 [Thread debugging using libthread_db enabled]
21 Using host libthread_db library "/lib/x86_64-linux-gnu/libthread_db.so.1".
22 jemalloc_shim.h:68: INSIST(ptr != (void *)0) failed, back trace
23 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(+0x27859) [0xf7f76859]
24 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(isc_assertion_failed+0x26) [0xf7f7678
↵ a]
25 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(+0x4222a) [0xf7f9122a]
26 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(+0x4238b) [0xf7f9138b]
27 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(isc__mem_get+0x58) [0xf7f923d0]
28 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(+0x390f5) [0xf7f880f5]
29 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(+0x39fb8) [0xf7f88fb8]
30 /home/user/src-new/bind9/lib/isc/.libs/libisc-9.19.18-dev.so(isc_lex_gettoken+0x2e0) [0xf7f89334]
31 /home/user/src-new/bind9/lib/dns/.libs/libdns-9.19.18-dev.so(+0x771c9) [0xf7c771c9]
32 /home/user/src-new/bind9/lib/dns/.libs/libdns-9.19.18-dev.so(+0x78fb7) [0xf7c78fb7]
33 /home/user/src-new/bind9/lib/dns/.libs/libdns-9.19.18-dev.so(dns_master_loadfile+0x8b) [0xf7c8066c]
34 /home/user/src-new/bind9/lib/dns/.libs/libdns-9.19.18-dev.so(dns_db_load+0xbd) [0xf7c403e8]
35 /home/user/src-new/bind9/bin/dnssec/.libs/dnssec-signzone(+0xc13c) [0x5656113c]
36 /home/user/src-new/bind9/bin/dnssec/.libs/dnssec-signzone(main+0xfaa) [0x565646a6]
37 /lib/i386-linux-gnu/libc.so.6(__libc_start_main+0x106) [0xf7a32e46]
38 /home/user/src-new/bind9/bin/dnssec/.libs/dnssec-signzone(_start+0x31) [0x56559e51]
39
40 Program received signal SIGABRT, Aborted.
41 [-----registers-----]
42 EAX: 0x0
43 EBX: 0x2
44 ECX: 0xffffbfdc --> 0x0
45 EDX: 0x0
46 ESI: 0x8
47 EDI: 0x0
48 EBP: 0xffffbfdc --> 0x0
```



```

49 ESP: 0xffffbfc0 --> 0xffffbfdc --> 0x0
50 EIP: 0xf7fd0559 (<__kernel_vsyscall+9>: pop    ebp)
51 EFLAGS: 0x200246 (carry PARITY adjust ZERO sign trap INTERRUPT direction overflow)
52 [-----code-----]
53   0xf7fd0553 <__kernel_vsyscall+3>:  mov    ebp,esp
54   0xf7fd0555 <__kernel_vsyscall+5>:  sysenter
55   0xf7fd0557 <__kernel_vsyscall+7>:  int    0x80
56 => 0xf7fd0559 <__kernel_vsyscall+9>:  pop    ebp
57   0xf7fd055a <__kernel_vsyscall+10>: pop    edx
58   0xf7fd055b <__kernel_vsyscall+11>: pop    ecx
59   0xf7fd055c <__kernel_vsyscall+12>: ret
60   0xf7fd055d:  nop
61 [-----stack-----]
62 0000| 0xffffbfc0 --> 0xffffbfdc --> 0x0
63 0004| 0xffffbfc4 --> 0x0
64 0008| 0xffffbfc8 --> 0xffffbfdc --> 0x0
65 0012| 0xffffbfcc --> 0xf7a48e02 (<__GI_raise+194>:  mov    eax,DWORD PTR [esp+0x10c])
66 0016| 0xffffbfd0 --> 0xf7fe26e9 (<_dl_fixup+9>:  add    ebx,0x1a917)
67 0020| 0xffffbfd4 --> 0xf7fc8000 --> 0x78e18
68 0024| 0xffffbfd8 --> 0xf7fad5a8 ("", back trace")
69 0028| 0xffffbfdc --> 0x0
70 [-----]
71 Legend: code, data, rodata, value
72 Stopped reason: SIGABRT
73 0xf7fd0559 in __kernel_vsyscall ()
74 gdb-peda$ bt
75 #0 0xf7fd0559 in __kernel_vsyscall ()
76 #1 0xf7a48e02 in __libc_signal_restore_set (set=0xffffbfdc) at
↳ ../sysdeps/unix/sysv/linux/internal-signals.h:86
77 #2 __GI_raise (sig=0x6) at ../sysdeps/unix/sysv/linux/raise.c:48
78 #3 0xf7a31306 in __GI_abort () at abort.c:79
79 #4 0xf7f76792 in isc_assertion_failed (file=0xf7fb2917 "jemalloc_shim.h", line=0x44,
↳ type=isc_assertiontype_insist,
80   cond=0xf7fb2904 "ptr != ((void *)0)") at assertions.c:49
81 #5 0xf7f9122a in mallocx (size=0x8000002c, flags=0x0) at jemalloc_shim.h:68
82 #6 0xf7f9138b in mem_get (ctx=0x56583b60, size=0x8000002c, flags=0x0) at mem.c:305
83 #7 0xf7f923d0 in isc__mem_get (ctx=0x56583b60, size=0x8000002c, flags=0x0) at mem.c:744
84 #8 0xf7f880f5 in isc_buffer_allocate (mctx=0x56583b60, dbufp=0xffffc334, length=0x80000000) at
↳ ./include/isc/buffer.h:1085
85 #9 0xf7f88fb8 in pushandgrow (lex=0x5658bf30, source=0x565832e0, c=0x41) at lex.c:327
86 #10 0xf7f89334 in isc_lex_gettoken (lex=0x5658bf30, options=0x137, tokenp=0xffffcd44) at lex.c:451
87 #11 0xf7c771c9 in gettoken (lex=0x5658bf30, options=0x137, token=0xffffcd44, eol=0x1,
↳ callbacks=0xffffce84) at master.c:341
88 #12 0xf7c78fb7 in load_text (lctx=0x5658b520) at master.c:1090
89 #13 0xf7c8066c in dns_master_loadfile (master_file=0xffffd549 "db.long", top=0x56582c60,
↳ origin=0x56582c60, zclass=0x1,
90   options=0x0, resign=0x0, callbacks=0xffffce84, include_cb=0x0, include_arg=0x0,
↳ mctx=0x56583b60, format=dns_masterformat_text,
91   maxttl=0x0) at master.c:2637
92 #14 0xf7c403e8 in dns_db_load (db=0x56582c50, filename=0xffffd549 "db.long",
↳ format=dns_masterformat_text, options=0x0) at db.c:316

```

```

93 #15 0x5656113c in loadzone (file=0xffffd549 "db.long", origin=0xffffd549 "db.long", rdclass=0x1,
↳ db=0x5656d5cc <gdb>)
94     at dnssec-signzone.c:2579
95 #16 0x565646a6 in main (argc=0x0, argv=0xffffd38c) at dnssec-signzone.c:3862
96 #17 0xf7a32e46 in __libc_start_main (main=0x565636fc <main>, argc=0x2, argv=0xffffd384,
↳ init=0x565669b0 <__libc_csu_init>,
97     fini=0x56566a10 <__libc_csu_fini>, rtld_fini=0xf7fe3230 <_dl_fini>, stack_end=0xffffd37c) at
↳ ../csu/libc-start.c:308
98 #18 0x56559e51 in _start ()
99 gdb-peda$

```

#### Listing 4.14: Failed Allocation on 32-Bit Architectures

As a proof of concept, a zone file triggering this behavior can be created using the command:

- `perl -e 'print "AA"x0x80000000' > db.long`

The tool `bin/dnssec/ddnssec-signzone` can then be used to trigger the crash. It shares the same code path with `named` and attackers could try to provide malicious zone files to it.

On 64-bit architectures, the parsing code will enter an infinite loop due to unsigned integer truncation.

Arithmetic overflows caused by large amounts of data are notoriously hard to find using general purpose dynamic analysis methods such as fuzzing.

#### 4.1.8.2 Solution Advice

X41 recommends to limit the maximum token size to a sane amount. Furthermore, it is recommended to replace all allocation size calculations in the lexer by using the new `isc_mem_allocate()` function which can detect arithmetic overflows. Besides the examples shown above, there are also other parts of the lexer code that could be prone to similar bugs such as the function shown in listing 4.15.

```

1 static isc_result_t
2 grow_data(isc_lex_t *lex, size_t *remainingp, char **currp, char **prevp) {
3     char *tmp;
4
5     tmp = isc_mem_get(lex->mctx, lex->max_token * 2 + 1);
6     memmove(tmp, lex->data, lex->max_token + 1);
7     *currp = tmp + (*currp - lex->data);
8     if (*prevp != NULL) {
9         *prevp = tmp + (*prevp - lex->data);

```

```
10     }
11     isc_mem_put(lex->mctx, lex->data, lex->max_token + 1);
12     lex->data = tmp;
13     *remainingp += lex->max_token;
14     lex->max_token *= 2;
15     return (ISC_R_SUCCESS);
16 }
```

---

**Listing 4.15:** Allocation Sizes Calculated Using Unsafe Integer Arithmetic

It is recommended to refactor the parsing code and potentially use parser generators to auto generate the parsing code.

## 4.2 Informational Notes

The following observations do not have a direct security impact, but are related to security hardening, affect functionality, or other topics that are not directly related to security. X41 recommends to mitigate these issues as well, because they often become exploitable in the future. Doing so will strengthen the security of the system and is recommended for defense in depth.

### 4.2.1 BND-CA-23-100: Memory Leak Due to `realloc()` Misuse

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4174>  
*Affected Component:* contrib/dlz/modules/bdbhpt/dlz\_bdbhpt\_dynamic.c

---

#### 4.2.1.1 Description

In `dlz_allnodes()` and `dlz_lookup()`, buffers are resized using `realloc()`. When the system is low on memory, `realloc()` will return `NULL` but will not free the original memory. The pointer `tmp` is overwritten by the return value of `realloc()` (as seen in listing 4.16) and no other reference to that buffer exists.

Therefore the memory allocated before is no longer referenced and leaks, which increases the memory pressure on the system even further. The relevant code is shown in the following listing 4.16:

---

```
1 tmp = realloc(tmp, dns_data.size + 1);
2 if (tmp == NULL) {
3     goto allnodes_cleanup;
4 }
```

---

**Listing 4.16:** `realloc()` Misuse

Since this issue is in the contributed and not the main code, it is considered an informational issue. Code in `contrib/` is not actively maintained by Internet Systems Corporation.

#### 4.2.1.2 Solution Advice

X41 recommends to use a temporary pointer to obtain the return value of `realloc()`.

## 4.2.2 BND-CA-23-101: ISC Memory API

---

*GitLab Issue:* n/a  
*Affected Component:* lib/isc/include/isc/mem.h

---

### 4.2.2.1 Description

The ISC memory API<sup>10</sup> functions `isc_mem_get()` and `isc_mem_allocate()` are used to allocate arrays. In several of these cases, the multiplications for these happen without an integer overflow check:

---

```

1 lib/isc/netmgr/netmgr.c:      isc_mem_get(mctx, netmgr->nloops * sizeof(netmgr->workers[0]));
2 lib/isc/lex.c:      tmp = isc_mem_get(lex->mctx, lex->max_token * 2 + 1);
3 lib/isc/symtab.c:      symtab->table = isc_mem_get(mctx, size * sizeof(eltlist_t));
4 lib/isc/symtab.c:      newtable = isc_mem_get(symtab->mctx, newsize * sizeof(eltlist_t));
5 lib/isc/tls.c:      locks = isc_mem_getx(isc__tls_mctx, nlocks * sizeof(locks[0]),
6 lib/isc/commandline.c:      *argvp = isc_mem_get(mctx, n * sizeof(char *));
7 lib/dns/ssu.c:      rule->types = isc_mem_get(mctx, ntypes * sizeof(*rule->types));
8 lib/dns/remote.c:      remote->ok = isc_mem_get(mctx, count * sizeof(bool));
9 lib/dns/rdataset.c:      in = isc_mem_get(cctx->mctx, count * sizeof(*in));
10 lib/dns/rdataset.c:      out = isc_mem_get(cctx->mctx, count * sizeof(*out));
11 lib/dns/diff.c:      v = isc_mem_get(diff->mctx, length * sizeof(dns_difftuple_t *));
12 lib/dns/rdataslab.c:      x = isc_mem_get(mctx, nalloc * sizeof(struct xrdata));
13 lib/dns/rdataslab.c:      offsettable = isc_mem_getx(mctx, nalloc * sizeof(unsigned int),
14 lib/dns/rdataslab.c:      offsettable = isc_mem_getx(mctx, mcount * sizeof(unsigned int),
15 lib/dns/ipkeylist.c:      addrs = isc_mem_get(mctx, n * sizeof(isc_sockaddr_t));
16 lib/dns/ipkeylist.c:      sources = isc_mem_get(mctx, n * sizeof(isc_sockaddr_t));
17 lib/dns/ipkeylist.c:      keys = isc_mem_get(mctx, n * sizeof(dns_name_t *));
18 lib/dns/ipkeylist.c:      tlss = isc_mem_get(mctx, n * sizeof(dns_name_t *));
19 lib/dns/ipkeylist.c:      labels = isc_mem_get(mctx, n * sizeof(dns_name_t *));
20 lib/dns/dispatch.c:      v4ports = isc_mem_get(mgr->mctx, sizeof(in_port_t) * nv4ports);
21 lib/dns/dispatch.c:      v6ports = isc_mem_get(mgr->mctx, sizeof(in_port_t) * nv6ports);
22 lib/dns/dispatch.c:      dset->dispatches = isc_mem_get(mctx, sizeof(dns_dispatch_t *) * n);
23 lib/dns/acl.c:      acl->elements = isc_mem_getx(mctx, n * sizeof(acl->elements[0]),
24 lib/dns/badcache.c:      bc->table = isc_mem_getx(bc->mctx, sizeof(bc->table[0]) * size,
25 lib/dns/badcache.c:      bc->tlocks = isc_mem_getx(bc->mctx, sizeof(bc->tlocks[0]) * size,
26 lib/dns/badcache.c:      newtable = isc_mem_getx(bc->mctx, sizeof(dns_bcentry_t *) * newsize,
27 lib/dns/badcache.c:      newlocks = isc_mem_get(bc->mctx, sizeof(isc_mutex_t) * newsize);
28 lib/dns/master.c:      newlist = isc_mem_get(mctx, new_len * sizeof(*newlist));
29 lib/dns/master.c:      newlist = isc_mem_get(mctx, new_len * sizeof(*newlist));
30 lib/dns/dnssec.c:      data = isc_mem_get(mctx, n * sizeof(dns_rdata_t));
31 lib/dns/zoneverify.c:      dstkeys = isc_mem_get(vctx->mctx, sizeof(*dstkeys) * count);
32 lib/dns/zone.c:      argv = isc_mem_get(zone->mctx, dbargc * sizeof(*argv));
33 lib/ns/update.c:      rules = isc_mem_getx(mctx, sizeof(*rules) * ruleslen,
34 lib/ns/query.c:      aaaaok = isc_mem_get(client->manager->mctx, sizeof(bool) * count);
35 bin/nsupdate/nsupdate.c:      servers = isc_mem_get(gmctx, ns_alloc * sizeof(isc_sockaddr_t));

```

<sup>10</sup> Application Programming Interface

```

36 bin/nsupdate/nsupdate.c:      servers = isc_mem_get(gmctx, ns_alloc * sizeof(isc_sockaddr_t));
37 bin/nsupdate/nsupdate.c:      servers = isc_mem_getx(gmctx, ns_alloc * sizeof(isc_sockaddr_t),
38 bin/dnssec/dnssec-cds.c:      keytable = isc_mem_getx(mctx, sizeof(keytable[0]) * nkey,
   ↪  ISC_MEM_ZERO);
39 bin/dnssec/dnssec-cds.c:      algo = isc_mem_getx(mctx, nkey * sizeof(algo[0]), ISC_MEM_ZERO);
40 bin/dnssec/dnssec-cds.c:      arrdata = isc_mem_get(mctx, n * sizeof(dns_rdata_t));
41 bin/dnssec/dnssec-cds.c:      ds = isc_mem_get(mctx, n * sizeof(dns_rdata_ds_t));
42 bin/dnssec/dnssec-signzone.c:      wassignedby = isc_mem_get(mctx, arraysize * sizeof(bool));
43 bin/dnssec/dnssec-signzone.c:      nowsignedby = isc_mem_get(mctx, arraysize * sizeof(bool));
44 bin/named/zoneconf.c:        types = isc_mem_get(mctx, n * sizeof(*types));

```

Listing 4.17: Example `isc_mem_get()` Calls

A full list (similar to the one shown in listing 4.17) can be found with the `weggli` query shown in listing 4.18:

```

1 weggli --unique -R 'a!=[A-Z_]+$' 'isc_mem_get($a * _);' ~/bind-9.19.13

```

Listing 4.18: `weggli` Query to Identify Multiplication in Allocation

This informational issue was not reported as a separate issue on GitLab but discussed in another issue<sup>11</sup> and addressed in merge request 8007<sup>12</sup>.

#### 4.2.2.2 Solution Advice

X41 recommends to improve the API by adding functions similar to `calloc()` that handle the multiplication and overflow check in a generic way.

<sup>11</sup>[https://gitlab.isc.org/isc-projects/bind9/-/issues/4120#note\\_379884](https://gitlab.isc.org/isc-projects/bind9/-/issues/4120#note_379884)

<sup>12</sup>[https://gitlab.isc.org/isc-projects/bind9/-/merge\\_requests/8007](https://gitlab.isc.org/isc-projects/bind9/-/merge_requests/8007)

## 4.2.3 BND-CA-23-102: Journal File Handling Missing Sanity Checks

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4175>

*Affected Component:* lib/dns/journal.c:journal\_open()

### 4.2.3.1 Description

The journal file handling code is missing sanity checks to prevent OOB reads and integer overflows that might result in OOB writes.

The calculation for *rawbytes* might overflow on 32-bit systems and less memory is allocated for *j->rawindex* and *j->index* than expected as shown in listing 4.19:

---

```

1 rawbytes = j->header.index_size * sizeof(journal_rawpos_t);
2 j->rawindex = isc_mem_get(mctx, rawbytes);
3
4 CHECK(journal_read(j, j->rawindex, rawbytes));
5
6 j->index = isc_mem_get(mctx, j->header.index_size *
7     sizeof(journal_pos_t));
8
9 p = j->rawindex;
10 for (i = 0; i < j->header.index_size; i++) {
11     j->index[i].serial = decode_uint32(p);
12     p += 4;
13     j->index[i].offset = decode_uint32(p);
14     p += 4;
15 }
```

---

Listing 4.19: Journal Handling

Additionally, the calls to *decode\_uint32()* might read more data than available as shown in listing 4.20:

---

```

1 =====
2 ==136920==ERROR: AddressSanitizer: heap-buffer-overflow on address 0x614000000400 at pc
3 ↪ 0x7f66822e43b5 bp 0x7ffe1b4ffa10 sp 0x7ffe1b4ffa08
4 READ of size 1 at 0x614000000400 thread T0
5   #0 0x7f66822e43b4 in decode_uint32 /home/eric/code/bind-9.19.13/lib/dns/journal.c:114:21
6   #1 0x7f66822e43b4 in journal_open /home/eric/code/bind-9.19.13/lib/dns/journal.c:717:25
7   #2 0x7f66822e25a8 in dns_journal_open /home/eric/code/bind-9.19.13/lib/dns/journal.c:775:11
8   #3 0x7f66822ec85f in dns_journal_print /home/eric/code/bind-9.19.13/lib/dns/journal.c:1639:11
```

```

8      #4 0x4ca513 in main /home/eric/code/bind-9.19.13/bin/tools/named-journalprint.c:125:12
9      #5 0x7f66816b2d09 in __libc_start_main csu/../csu/libc-start.c:308:16
10     #6 0x420489 in _start (/usr/bin/named-journalprint+0x420489)
11
12     0x61400000400 is located 0 bytes to the right of 448-byte region [0x61400000240,0x61400000400)
13     allocated by thread T0 here:
14         #0 0x49a15d in malloc (/usr/bin/named-journalprint+0x49a15d)
15         #1 0x7f6681adb290 in mallocx /home/eric/code/bind-9.19.13/lib/isc/./jemalloc_shim.h:65:14
16         #2 0x7f6681adb290 in mem_get /home/eric/code/bind-9.19.13/lib/isc/mem.c:304:8
17         #3 0x7f6681adb290 in isc__mem_get /home/eric/code/bind-9.19.13/lib/isc/mem.c:667:8
18
19     SUMMARY: AddressSanitizer: heap-buffer-overflow
20     ↪ /home/eric/code/bind-9.19.13/lib/dns/journal.c:114:21 in decode_uint32
21     Shadow bytes around the buggy address:
22     0x0c287fff8030: 00 00 00 00 00 00 00 00 00 00 fa fa fa fa fa fa
23     0x0c287fff8040: fa fa fa fa fa fa fa fa 00 00 00 00 00 00 00
24     0x0c287fff8050: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
25     0x0c287fff8060: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
26     =>0x0c287fff8080: [fa]fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
27     0x0c287fff8090: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
28     0x0c287fff80a0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
29     0x0c287fff80b0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
30     0x0c287fff80c0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
31     0x0c287fff80d0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
32     Shadow byte legend (one shadow byte represents 8 application bytes):
33     Addressable:           00
34     Partially addressable: 01 02 03 04 05 06 07
35     Heap left redzone:    fa
36     Freed heap region:    fd
37     Stack left redzone:   f1

```

#### Listing 4.20: ASan Output for Journal OOB Read

Since journal files are considered trusted in general, this is considered an informational note.

#### 4.2.3.2 Solution Advice

X41 recommends to add safeguards and implement fuzz testing for journal file handling.



## 4.2.4 BND-CA-23-103: Unchecked malloc()

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4077>

*Affected Component:* lib/isc/thread.c:thread\_wrap()

---

### 4.2.4.1 Description

In *thread\_wrap()*, the return value of *malloc()* is not checked and can potentially cause a *NULL* pointer dereference in low memory situations as seen in listing 4.21:

---

```
1 static struct thread_wrap *
2 thread_wrap(isc_threadfunc_t func, void *arg) {
3     struct thread_wrap *wrap = malloc(sizeof(*wrap));
4     *wrap = (struct thread_wrap){
5         .func = func,
6         .arg = arg,
7     };
8
9     return (wrap);
10 }
```

---

**Listing 4.21:** Unchecked malloc()

This issue is not likely to be enforceable by an attacker and in case of *malloc()* failing, the server would shutdown anyway. Therefore, this is considered an informational finding.

When reporting this issue it was noticed that this issue was reported in GitLab before as #4077<sup>13</sup>.

### 4.2.4.2 Solution Advice

X41 recommends to handle the low memory situation more gracefully and print the proper error and perform a cleanup before the exit.

---

<sup>13</sup><https://gitlab.isc.org/isc-projects/bind9/-/issues/4077>

## 4.2.5 BND-CA-23-104: Stack Exhaustion in Config Parser

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4176>  
*Affected Component:* isccfg/parser.c:cfg\_parse\_listelt()

---

### 4.2.5.1 Description

When the configuration file contains too many opening brackets, the configuration parser will iterate until the stack is exhausted and the application crashes (see listing 4.22):

---

```

1 AddressSanitizer:DEADLYSIGNAL
2 =====
3 ==881685==ERROR: AddressSanitizer: stack-overflow on address 0x7fff92547ee8 (pc 0x00000049be0e bp
↳ 0x7fff92548730 sp 0x7fff92547ef0 T0)
4   #0 0x49be0e in __asan_memmove (/usr/bin/named-checkconf+0x49be0e)
5   #1 0x7f322de80661 in isc_buffer_compact
↳ /home/eric/code/bind-9.19.13/lib/isc/.include/isc/buffer.h:1076:9
6   #2 0x7f322de80661 in isc_lex_gettoken /home/eric/code/bind-9.19.13/lib/isc/lex.c:399:2
7   #3 0x7f322d406da3 in cfg_gettoken /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:3474:11
8   #4 0x7f322d4256be in cfg_peektoken /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:3537:2
9   #5 0x7f322d4256be in parse_addrmatchelt
↳ /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:1715:2
10  #6 0x7f322d414bf4 in cfg_parse_obj /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:248:11
11  #7 0x7f322d414bf4 in cfg_parse_listelt
↳ /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:1989:11
12  #8 0x7f322d412b58 in parse_list /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:2024:3
13  #9 0x7f322d412b58 in cfg_parse_bracketed_list
↳ /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:2069:2
14  #10 0x7f322d4258a3 in cfg_parse_obj /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:248:11
15  ...
16  #494 0x7f322d412b58 in parse_list /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:2024:3
17  #495 0x7f322d412b58 in cfg_parse_bracketed_list
↳ /home/eric/code/bind-9.19.13/lib/isccfg/parser.c:2069:2
18
19 SUMMARY: AddressSanitizer: stack-overflow (/usr/bin/named-checkconf+0x49be0e) in __asan_memmove
20 ==881685==ABORTING

```

---

**Listing 4.22:** Unbound Recursion in Configuration Parser

Since the configuration is considered trusted, this is considered an informational finding.

#### 4.2.5.2 Solution Advice

X41 recommends to limit the iteration depth to a sane amount.

## 4.2.6 BND-CA-23-105: Connection Flags Mixup

---

**GitLab Issue:** <https://gitlab.isc.org/isc-projects/bind9/-/issues/4126>

**Affected Component:** lib/isc/httpd.c:process\_request()

---

### 4.2.6.1 Description

In lib/isc/httpd.c in function `process_request()`, an HTTP request is parsed and the headers are evaluated.

When the client sets two *Connection* headers, the `HTTPD_CLOSE` and `HTTPD_KEEPAKALIVE` flags might be set on `httpd->flags` as shown in listing 4.23:

---

```
1  #define HTTPD_CLOSE      0x0001 /* Got a Connection: close header */
2  #define HTTPD_FOUNDDHOST 0x0002 /* Got a Host: header */
3  #define HTTPD_KEEPAKALIVE 0x0004 /* Got a Connection: Keep-Alive */
4
5  ...
6
7  } else if (name_match(header, "Connection")) {
8      if (value_match(header, "close")) {
9          httpd->flags |= HTTPD_CLOSE;
10     } else if (value_match(header, "keep-alive")) {
11         keep_alive = true;
12     }
13
14     ...
15
16     switch (httpd->minor_version) {
17     case 0:
18         if (keep_alive == true) {
19             httpd->flags |= HTTPD_KEEPAKALIVE;
20         } else {
21             httpd->flags |= HTTPD_CLOSE;
22         }
23         break;
```

---

**Listing 4.23:** Connection Flags Mixup

This currently has no impact but might lead to logic issues in the future when more features are added to the code.

#### 4.2.6.2 Solution Advice

X41 recommends to clear the opposite flag when a new *Connection* header is read.

## 4.2.7 BND-CA-23-106: Use of Magic Numbers

*GitLab Issue:* n/a

*Affected Component:* lib/dns/catz.c, lib/dns/request.c, lib/dns/db.c

### 4.2.7.1 Description

A few places in the code use magic numbers<sup>14</sup> instead of using `sizeof()` or a named define.

Using defines or `sizeof()` makes the code easier to understand and audit as well as more robust to changes in the future as shown in the following listings:

---

```

1 static isc_result_t
2 catz_process_version(dns_catz_zone_t *catz, dns_rdataset_t *value) {
3 ...
4     char t[16];
5 ...
6     if (rdatastr.length > 15) {
7         result = ISC_R_BADNUMBER;
8         goto cleanup;
9     }

```

---

**Listing 4.24:** Use of Magic Numbers in cat.c

---

```

1 isc_result_t
2 dns_request_creatoraw(dns_requestmgr_t *requestmgr, isc_buffer_t *msgbuf,
3     const isc_sockaddr_t *srcaddr,
4     const isc_sockaddr_t *destaddr,
5     dns_transport_t *transport,
6     isc_tlsctx_cache_t *tlsctx_cache, unsigned int options,
7     unsigned int timeout, unsigned int udptimeout,
8     unsigned int udpretries, isc_loop_t *loop, isc_job_cb cb,
9     void *arg, dns_request_t **requestp) {
10
11 ...
12     isc_buffer_usedregion(msgbuf, &r);
13     if (r.length < DNS_MESSAGE_HEADERLEN || r.length > 65535) {
14         result = DNS_R_FORMERR;
15         goto cleanup;
16     }
17

```

---

<sup>14</sup> [https://en.wikipedia.org/wiki/Magic\\_number\\_\(programming\)](https://en.wikipedia.org/wiki/Magic_number_(programming))

```
18     if ((options & DNS_REQUESTOPT_TCP) != 0 || r.length > 512) {
19         tcp = true;
20         request->timeout = timeout * 1000;
21     } else {
```

---

**Listing 4.25:** Use of Magic Numbers in request.c

```
1  isc_result_t
2  dns_db_getsoaserial(dns_db_t *db, dns_dbversion_t *ver, uint32_t *serialp) {
3  ...
4      INSIST(rdata.length > 20);
5      isc_buffer_init(&buffer, rdata.data, rdata.length);
6      isc_buffer_add(&buffer, rdata.length);
7      isc_buffer_forward(&buffer, rdata.length - 20);
```

---

**Listing 4.26:** Use of Magic Numbers in db.c

```
1  memset(rdatalist->upper, 0xeb, sizeof(rdatalist->upper));
2  ...
3  REQUIRE(rdatalist->upper[0] == 0xea);
```

---

**Listing 4.27:** Use of Magic Numbers in rdatalist.c

#### 4.2.7.2 Solution Advice

X41 recommends to replace magic numbers with defines or calls to *sizeof()*.

## 4.2.8 BND-CA-23-107: NULL Pointer Dereference on Wrong API Usage

---

<i>GitLab Issue:</i>	<a href="https://gitlab.isc.org/isc-projects/bind9/-/issues/4177">https://gitlab.isc.org/isc-projects/bind9/-/issues/4177</a>
<i>Affected Component:</i>	contrib/dlz/modules/sqlite3/dlz_sqlite3_dynamic.c:sqlite3_get_resultset()

---

### 4.2.8.1 Description

In contrib/dlz/modules/sqlite3/dlz\_sqlite3\_dynamic.c in function `sqlite3_get_resultset()`, a wrong use of the API causes a goto `cleanup`. At this point the pointer `dbi` is initialized with `NULL` and dereferenced in the cleanup phase (see listing 4.28).

---

```

1  static isc_result_t
2  sqlite3_get_resultset(const char *zone, const char *record, const char *client,
3                       unsigned int query, void *dbdata, sqlite3_res_t **rsp) {
4
5     isc_result_t result;
6     dbinstance_t *dbi = NULL;
7     sqlite3_instance_t *db = (sqlite3_instance_t *)dbdata;
8     char *querystring = NULL;
9     sqlite3_res_t *rs = NULL;
10    int qres = 0;
11
12    if ((query == COUNTZONE && rsp != NULL) ||
13        (query != COUNTZONE && (rsp == NULL || *rsp != NULL)))
14    {
15        db->log(ISC_LOG_DEBUG(2), "Invalid result set pointer.");
16        result = ISC_R_FAILURE;
17        goto cleanup;
18    }
19    ...
20    cleanup:
21    if (dbi->zone != NULL) {
22        free(dbi->zone);
23        dbi->zone = NULL;
24    }
25    if (dbi->record != NULL) {
26        free(dbi->record);
27        dbi->record = NULL;
28    }
29    if (dbi->client != NULL) {
30        free(dbi->client);
31        dbi->client = NULL;
32    }
33
34    /* release the lock so another thread can use this dbi */
35    (void)dlz_mutex_unlock(&dbi->lock);

```



```
35
36     if (querystring != NULL) {
37         free(querystring);
38     }
39
40     return (result);
41 }
```

---

**Listing 4.28:** NULL Pointer Dereference on Wrong API Usage

Since this issue is in the contributed code and not the main code, it is considered an informational issue.

#### 4.2.8.2 Solution Advice

X41 recommends to add an additional check against *NULL* before dereferencing *dbi*.

## 4.2.9 BND-CA-23-108: Invalid Free in Low Memory Situation

GitLab Issue: <https://gitlab.isc.org/isc-projects/bind9/-/issues/4179>

Affected Component: contrib/dlz/modules/bdbhpt/dlz\_bdbhpt\_dynamic.c:dlz\_allowzonexfr()

### 4.2.9.1 Description

In contrib/dlz/modules/bdbhpt/dlz\_bdbhpt\_dynamic.c in the function `dlz_allowzonexfr()`, the memory of `data` is only initialized after `key` has been set up. When `strdup()` fails, e.g. due to a low memory situation, the value `data.data` might be passed to `free()`, while still being uninitialized and pointing to arbitrary memory on the stack as seen in listing 4.29. This might result in freeing used and allocated memory, in a double free or segfault when non-accessible memory is accessed.

```
1  isc_result_t
2  dlz_allowzonexfr(void *dbdata, const char *name, const char *client) {
3      isc_result_t result;
4      bdbhpt_instance_t *db = (bdbhpt_instance_t *)dbdata;
5      DBT key, data;
6
7      /* check to see if we are authoritative for the zone first. */
8      #if DLZ_DLOPEN_VERSION >= 3
9          result = dlz_findzonedb(dbdata, name, NULL, NULL);
10     #else /* if DLZ_DLOPEN_VERSION >= 3 */
11         result = dlz_findzonedb(dbdata, name);
12     #endif /* if DLZ_DLOPEN_VERSION >= 3 */
13     if (result != ISC_R_SUCCESS) {
14         return (ISC_R_NOTFOUND);
15     }
16
17     memset(&key, 0, sizeof(DBT));
18     key.flags = DB_DBT_MALLOC;
19     key.data = strdup(name);
20     if (key.data == NULL) {
21         result = ISC_R_NOMEMORY;
22         goto xfr_cleanup;
23     }
24     key.size = strlen(key.data);
25
26     memset(&data, 0, sizeof(DBT));
27
28     ...
29
30     xfr_cleanup:
31         /* free any memory duplicate string in the key field */
32         if (key.data != NULL) {
33             free(key.data);
```

```
34     }
35
36     /* free any memory allocated to the data field. */
37     if (data.data != NULL) {
38         free(data.data);
39     }
40
41     return (result);
42 }
```

---

**Listing 4.29:** Invalid Free in Low Memory Situation

Since this issue is in the contributed code and not the main code, it is considered an informational issue.

#### 4.2.9.2 Solution Advice

X41 recommends to initialize *data* earlier in the function.

## 4.2.10 BND-CA-23-109: Possible Truncation in dns\_keymgr\_status()

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4180>

*Affected Component:* lib/dns/keymgr.c: dns\_keymgr\_status()

### 4.2.10.1 Description

The function `isc_buffer_printf()` returns the error value `ISC_R_NOSPACE` when the input buffer is not large enough to contain the result. Callers that do not provide a sufficiently large buffer and do not check the return value might operate on truncated strings without noticing. One such caller is `dns_keymgr_status()` which is called from `named_server_dnssec()` with a 4096 byte buffer as shown in listing 4.30.

```

1  isc_result_t
2  named_server_dnssec(named_server_t *server, isc_lex_t *lex,
3                      isc_buffer_t **text) {
4  ...
5      char output[4096];
6  ...
7      dns_keymgr_status(kasp, &keys, now, &output[0], sizeof(output));
8
9  ...
10
11 void
12 dns_keymgr_status(dns_kasp_t *kasp, dns_dnsseckeylist_t *keyring,
13                  isc_stdtime_t now, char *out, size_t out_len) {
14     isc_buffer_t buf;
15     char timestr[26]; /* Minimal buf as per ctime_r() spec. */
16
17     REQUIRE(DNS_KASP_VALID(kasp));
18     REQUIRE(keyring != NULL);
19     REQUIRE(out != NULL);
20
21     isc_buffer_init(&buf, out, out_len);
22
23     // policy name
24     isc_buffer_printf(&buf, "dnssec-policy: %s\n", dns_kasp_getname(kasp));
25     isc_buffer_printf(&buf, "current time: ");
26     isc_stdtime_tostring(now, timestr, sizeof(timestr));
27     isc_buffer_printf(&buf, "%s\n", timestr);
28
29     for (dns_dnsseckey_t *dkey = ISC_LIST_HEAD(*keyring); dkey != NULL;
30          dkey = ISC_LIST_NEXT(dkey, link))
31     {
32         char algstr[DNS_NAME_FORMATSIZE];
33         bool ksk = false, zsk = false;

```

```
34         isc_result_t ret;
35
36         if (dst_key_is_unused(dkey->key)) {
37             continue;
38         }
39
40         // key data
41         dns_secalg_format((dns_secalg_t)dst_key_alg(dkey->key), algstr,
42                         sizeof(algstr));
43         isc_buffer_printf(&buf, "\nkey: %d (%s), %s\n",
44                         dst_key_id(dkey->key), algstr,
45                         keymgr_keyrole(dkey->key));
46
47         // publish status
48         keytime_status(dkey->key, now, &buf,
49                      " published:      ", DST_KEY_DNSKEY,
50                      DST_TIME_PUBLISH);
51     ...
```

**Listing 4.30:** Possible Truncation in dns\_keymgr\_status()

#### 4.2.10.2 Solution Advice

X41 recommends to either verify the return values or use a dynamic buffer.

## 4.2.11 BND-CA-23-110: Newline and ANSI Escape Code Injection via CC

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4181>  
*Affected Component:* bin/named/control.c:named\_control\_docommand(),  
 bin/delv/delv.c:main()

---

### 4.2.11.1 Description

When invalid commands are received, they are logged into a file, or when `named` is started with the `-g` parameter, to `stdout`. No sanitation is performed before the command data is output into the log as seen in listing 4.31.

---

```

1 } else {
2     isc_log_write(named_g_lctx, NAMED_LOGCATEGORY_GENERAL,
3                   NAMED_LOGMODULE_CONTROL, ISC_LOG_WARNING,
4                   "unknown control channel command '%s'", command);
5     result = DNS_R_UNKNOWNCOMMAND;
6 }

```

---

**Listing 4.31:** Newline and ANSI Escape Character Injection via CC

This can be abused by an authenticated attacker to inject ANSI<sup>15</sup> escape codes<sup>16</sup> into the output by calling `rndc` with a maliciously crafted command such as the one seen in listing 4.32.

---

```

1 /sbin/rndc -s 127.0.0.1 -p 953 -k /etc/bind/rndc.key `echo -e "\x0a\x0d<fakedatehere> Normal
↵ \e[1mBold`

```

---

**Listing 4.32:** Using `rndc` to Inject Newlines and ANSI Escape Characters

This allows attackers to inject newlines for fake log entries or add arbitrary formatting to the data when viewed in a console as shown in figure 4.1.

Since this requires authentication as an already privileged user, this is considered an informational issue.

A similar issue exists in `bin/delv/delv.c` in the function `main()` when tracing is enabled via `+ns` and the `server` value is attacker controlled as shown in listing 4.33.

<sup>15</sup> American National Standards Institute

<sup>16</sup> [https://en.wikipedia.org/wiki/ANSI\\_escape\\_code](https://en.wikipedia.org/wiki/ANSI_escape_code)

```
22-Jun-2023 11:01:46.764 managed-keys-zone: error during managed-keys processing
<fakedatehere> Normal Bold'ceived control channel command ''
22-Jun-2023 11:01:53.423 unknown control channel command ''
□
```

Figure 4.1: Log Output with Injected Bold Code

---

```
1 if (fulltrace && server != NULL) {
2     delv_log(ISC_LOG_WARNING,
3         "WARNING: using internal name server mode: "
4         "'@%s' will be ignored",
5         server);
6 }
```

---

Listing 4.33: Injection in delv.c

Since the *server* parameter should usually not be attacker controlled this is considered informational, but it might be when *delv* is called from a web-interface.

#### 4.2.11.2 Solution Advice

X41 recommends to sanitize the data before sending it to *isc\_log\_write()*, similar to the escaping in *dns\_name\_totext2()*.

## 4.2.12 BND-CA-23-111: Name Buffer Truncation

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4186>

*Affected Component:* lib/isc/loop.c:dns\_zonemgr\_create()

---

### 4.2.12.1 Description

A truncation of the name of memory pools was found which might lead to unintended behavior or incorrect debugging output.

A memory pool structure *isc\_mempool* has a member field *name* with a capacity of 16 bytes as shown in listing 4.34:

---

```
1 struct isc_mempool {
2     /* always unlocked */
3     unsigned int magic;
4     isc_mem_t *mctx;          /*%< our memory context */
5     ISC_LINK(isc_mempool_t) link; /*%< next pool in this mem context */
6     element *items;         /*%< low water item list */
7     size_t size;            /*%< size of each item on this pool */
8     size_t allocated;       /*%< # of items currently given out */
9     size_t freecount;       /*%< # of items on reserved list */
10    size_t freemax;          /*%< # of items allowed on free list */
11    size_t fillcount;        /*%< # of items to fetch on each fill */
12    /*%< Stats only. */
13    size_t gets; /*%< # of requests to this pool */
14    /*%< Debugging only. */
15    char name[16]; /*%< printed name in stats reports */
16 };
```

---

Listing 4.34: Name Structure Member Declaration in File lib/isc/mem.c



In the function `dns_zonemgr_create()`, a string of size 16 without the terminating `NUL` byte is passed on to function `isc_mem_setname()`, leading to silent truncation of the last character in that string as shown in the following listing 4.35:

---

```
1 for (size_t i = 0; i < zmgr->workers; i++) {
2     isc_mem_create(&zmgr->mctxpool[i]);
3     isc_mem_setname(zmgr->mctxpool[i], "zonemgr-mctxpool");
4     ↪ // MARK truncation / off by one (namebuffer is 16 bytes only)
5 }
```

---

**Listing 4.35:** Terminating NUL Byte Truncation

This issue is informational since the truncation has no security implications, but could lead to incorrect assumptions or functionality defects.

#### 4.2.12.2 Solution Advice

X41 recommends to either increase the buffer size or shorten the name value, but to also add an assertion to the `isc_mem_create()` function that ensures the `name` size is larger than zero and less than 16 bytes without the terminating `NUL` byte.

## 4.2.13 BND-CA-23-112: Misaligned Structure Causes Exception

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4187>

*Affected Component:* lib/dns/rbtdb.c:allocate\_version() (and others)

### 4.2.13.1 Description

The `named` process crashes with a segmentation fault when compiled with the optimization flags `-march=native -ftree-slp-vectorize` and `jemalloc` disabled. During tests, X41 found that compilers `clang`<sup>17</sup> version 16 and `GCC`<sup>18</sup> versions 13.2.1 (Fedora 38) and 10.4 (Debian 11) were affected.

On a recent Intel CPU system, a crash can be reproduced reliably when compiling with the following commands shown in listing 4.36:

```
1 export CFLAGS="-O1 -g -march=native -ftree-slp-vectorize"
2 export CXXFLAGS=$CFLAGS
3 make clean
4 ./configure --without-jemalloc && make -j16 2>&1 | tee compile-warnings-log-$(date +%s.%N).txt
```

**Listing 4.36:** Compile with Vector Optimizations Enabled via `-ftree-slp-vectorize`

It was found that for `clang-16`, the issue is present even when compiling with flags `-O1 -march=native` only.

When starting `named` with the example configuration via the command line `./named -f -d10 -M fill`, a crash occurs (see listing 4.37):

```
1 Thread 1 "named" received signal SIGSEGV, Segmentation fault.
2 0x00007ffff7903136 in allocate_version (mctx=mctx@entry=0x5555556da340, serial=serial@entry=1,
  ↳ references=references@entry=1, writer=writer@entry=false) at rbtdb.c:1282
3 1282         ISC_LIST_INIT(version->resigned_list);
4 LEGEND: STACK | HEAP | CODE | DATA | RWX | RODATA
5 -----[ REGISTERS / show-flags off / show-compact-regs off
  ↳ ]-----
6 *RAX 0x5555556dff20 <- 0x0
7 *RBX 0x5555556dfb10 <- 0xbebebebe00000001
8   RCX 0x0
9 *RDX 0x28
10 *RDI 0x5555556dff20 <- 0x0
11   RSI 0x0
```

<sup>17</sup><https://clang.org>

<sup>18</sup><https://gcc.gnu.org>

```

12 *R8  0x5555556dff20 <- 0x0
13 *R9  0x7ffff720cbe0 (main_arena+96) -> 0x5555556ff3f0 <- 0x211b0cca04000100
14 *R10 0x4000000
15   R11 0x0
16 *R12 0x1
17 *R13 0x5555556da340 <- 0x44d656d43
18   R14 0x0
19 *R15 0x1
20 *RBP 0x7fffffff89d0 <- 0x0
21 *RSP 0x7fffffff89a0 -> 0x5555556dad60 -> 0x5555556dfa90 <- 0x5242542b /* '+TBR' */
22 *RIP 0x7ffff7903136 (allocate_version+122) <- vmovdqa ymmword ptr [rbx + 0x20], ymm0
23 -----[ DISASM / x86-64 / set emulate on ]-----
24 > 0x7ffff7903136 <allocate_version+122> vmovdqa ymmword ptr [rbx + 0x20], ymm0
25   0x7ffff790313b <allocate_version+127> mov    qword ptr [rbx + 0x40], -1
26   0x7ffff7903143 <allocate_version+135> mov    qword ptr [rbx + 0x48], -1
27   0x7ffff790314b <allocate_version+143> mov    rax, rbx
28   0x7ffff790314e <allocate_version+146> add    rsp, 8
29   0x7ffff7903152 <allocate_version+150> pop    rbx
30   0x7ffff7903153 <allocate_version+151> pop    r12
31   0x7ffff7903155 <allocate_version+153> pop    r13
32   0x7ffff7903157 <allocate_version+155> pop    r14
33   0x7ffff7903159 <allocate_version+157> pop    r15
34   0x7ffff790315b <allocate_version+159> pop    rbp
35 -----[ SOURCE (CODE) ]-----
36 In file: /home/user/src/bind9-clean/lib/dns/rbtdb.c
37   1277         version->glue_table = isc_mem_getx(mctx, size, ISC_MEM_ZERO);
38   1278
39   1279         version->writer = writer;
40   1280         version->commit_ok = false;
41   1281         ISC_LIST_INIT(version->changed_list);
42   > 1282         ISC_LIST_INIT(version->resigned_list);
43   1283         ISC_LINK_INIT(version, link);
44   1284
45   1285         return (version);
46   1286 }
47   1287

```

Listing 4.37: Segmentation Fault in allocate\_version()

The crash occurs in function `allocate_version()` although the code looks unsuspecting as seen in listing 4.38:

```

1 static rbtdb_version_t *
2 allocate_version(isc_mem_t *mctx, rbtdb_serial_t serial,
3                 unsigned int references, bool writer) {
4     rbtdb_version_t *version;
5     size_t size;
6

```

```
7     version = isc_mem_get(mctx, sizeof(*version));
8     version->serial = serial;
9
10    isc_refcount_init(&version->references, references);
11    isc_rwlock_init(&version->glue_rwlock);
12
13    version->glue_table_bits = ISC_HASH_MIN_BITS;
14    version->glue_table_nodecount = 0U;
15
16    size = ISC_HASHSIZE(version->glue_table_bits) *
17          sizeof(version->glue_table[0]);
18    version->glue_table = isc_mem_getx(mctx, size, ISC_MEM_ZERO);
19
20    version->writer = writer;
21    version->commit_ok = false;
22    ISC_LIST_INIT(version->changed_list);
23    ↔ // MARK incorrect optimization of these consecutive ISC_LIST_INIT statements
24    ISC_LIST_INIT(version->resigned_list);
25    ISC_LINK_INIT(version, link);
26
27    return (version);
28 }
```

**Listing 4.38:** Code Pattern Example in Function `allocate_version()`

After extensive debugging it was found that an unaligned pointer is used in a `x86_64` vector instruction:

```
vmovdqa ymmword ptr [rbx + 0x20], ymm0
```

Further investigation reveals that this seems to be caused by a misaligned address used in instruction `vmovdqa`, which requires a 32-bit or 64-bit alignment depending on the target architecture. The instruction is introduced due to automatic vectorization optimizations caused by the flags `-march=native -ftree-slp-vectorize`, which cause the compiler to use the native instruction set of the detected architecture and to apply auto-vectorization<sup>19</sup> performance optimizations.

<sup>19</sup><https://gcc.gnu.org/projects/tree-ssa/vectorization.html>

Making changes to the order and interleaving of statements in the C code, which should not have any effect on the list operation semantics, makes the crash disappear as shown in listing 4.39:

---

```

1  ISC_LIST_INIT(version->changed_list);
2  volatile int noop = 1; if (noop) { // MARK no effect / NOOP to break the optimization
3      ISC_LIST_INIT(version->resigned_list);
4  }

```

---

**Listing 4.39:** Prevent Optimization

However, additional crashes appear in the code shown in the following listing 4.40:

---

```

1      2586                                     /*
2      2587                                     * We're rolling back this transaction.
3      2588                                     */
4      2589                                     cleanup_list = version->changed_list;
5      2590                                     ISC_LIST_INIT(version->changed_list);
6      2591                                     resigned_list = version->resigned_list;
7  > 2592                                     ISC_LIST_INIT(version->resigned_list);
8      2593                                     rollback = true;
9      2594                                     cleanup_version = version;
10     2595                                     rbtdb->future_version = NULL;
11     2596                                     }
12     2597     } else {

```

---

**Listing 4.40:** Another Crash in File lib/dns/rbtdb.c

When changing the alignments in the `rbtdb_version` struct via compiler attributes, the crash disappears as shown in the following code listing 4.41:

---

```

1  typedef struct rbtdb_version {
2      /* Not locked */
3      rbtdb_serial_t serial;
4      dns_rbtdb_t *rbtdb;
5      /*
6      * Protected in the refcount routines.
7      * XXXJT: should we change the lock policy based on the refcount
8      * performance?
9      */
10     isc_refcount_t references;
11     /* Locked by database lock. */
12     bool writer;
13     bool commit_ok;

```

---

```
14     struct {} __attribute__((aligned (32))); // MARK force alignment
15     rbtdb_changedlist_t changed_list;
16     struct {} __attribute__((aligned (32)));
```

#### Listing 4.41: Patch to Enforce Correct Alignment

X41 considers the root cause to be the byte alignment of size 32, which exceeds the value of `max_align_t`<sup>20</sup>.

Alignment values above `max_align_t` lead to *implementation defined behavior* according to section 6.2.8 of ISO/IEC 9899:201x<sup>21</sup>:

An extended alignment is represented by an alignment greater than `_Alignof(max_align_t)`. It is implementation-defined whether any extended alignments are supported and the contexts in which they are supported. A type having an extended alignment requirement is an over-aligned type.

Memory allocated using `malloc` will be aligned at least as strictly as `max_align_t`, but in this case the alignment required is greater.

Since this issue causes a crash already at startup, it is most likely not security relevant currently. Should it happen that an optimization is done in a part of the code that is not executed at startup, but could be reachable by attacker controller input, this issue could escalate into a denial of service issue.

#### 4.2.13.2 Solution Advice

X41 recommends to either remove the specification of `alignas(ISC_OS_CACHELINE_SIZE)` that leads to extended alignment, or to change the allocation function from `malloc` to `posix_memalign`<sup>22</sup>.

<sup>20</sup> [https://en.cppreference.com/w/c/types/max\\_align\\_t](https://en.cppreference.com/w/c/types/max_align_t)

<sup>21</sup> <https://www.open-std.org/jtc1/sc22/wg14/www/docs/n1570.pdf>

<sup>22</sup> [https://man7.org/linux/man-pages/man3/posix\\_memalign.3.html](https://man7.org/linux/man-pages/man3/posix_memalign.3.html)

## 4.2.14 BND-CA-23-113: Race in dns\_tsigkey\_find()

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4182>

*Affected Component:* lib/dns/tsig.c: dns\_tsigkey\_find()

### 4.2.14.1 Description

There is a race condition which will trigger *ISC\_LINK\_INSIST()* when *dns\_tsigkey\_find()* is called twice for the same generated and expired key.

In the function *dns\_tsigkey\_find()* in lib/dns/tsig.c, a lookup for a key is performed based on the name and algorithm. If the found key is expired, it is removed from the keyring. These operations are protected by the *ring->lock* and depending on the access this lock is held in read or write mode as seen in listing 4.42.

```

1  isc_result_t
2  dns_tsigkey_find(dns_tsigkey_t **tsigkey, const dns_name_t *name,
3                  const dns_name_t *algorithm, dns_tsig_keyring_t *ring) {
4  ...
5      RWLOCK(&ring->lock, isc_rwlocktype_read);
6      key = NULL;
7      result = dns_rbt_findname(ring->keys, name, 0, NULL, (void *)&key);
8  ...
9      if (key->inception != key->expire && isc_serial_lt(key->expire, now)) {
10         /*
11          * The key has expired.
12          */
13         RWUNLOCK(&ring->lock, isc_rwlocktype_read);
14         RWLOCK(&ring->lock, isc_rwlocktype_write);
15         remove_fromring(key);
16         RWUNLOCK(&ring->lock, isc_rwlocktype_write);
17         return (ISC_R_NOTFOUND);
18     }
19 ...
20     RWUNLOCK(&ring->lock, isc_rwlocktype_read);
21     adjust_lru(key);
22     *tsigkey = key;
23     return (ISC_R_SUCCESS);
24 }

```

Listing 4.42: Lock Handling in dns\_tsigkey\_find()

When the process gets interrupted after the call to *RWUNLOCK(&ring->lock, isc\_rwlocktype\_read)*,

and before the call to `RWLOCK(&ring->lock, isc_rwlocktype_write)`, the call to `remove_fromring()` will be performed twice. For generated keys, the function `remove_fromring()` will try to unlink that key for each caller (see listing 4.43), which will trigger `ISC_LINK_INSIST(ISC_LINK_LINKED(elt, link))`.

---

```
1 static void
2 remove_fromring(dns_tsigkey_t *tkey) {
3     if (tkey->generated) {
4         ISC_LIST_UNLINK(tkey->ring->lru, tkey, link);
5         tkey->ring->generated--;
6     }
7     (void)dns_rbt_deletename(tkey->ring->keys, &tkey->name, false);
8 }
```

---

Listing 4.43: `ISC_LIST_UNLINK()` Called in `remove_fromring()`

#### 4.2.14.2 Solution Advice

X41 recommends to ensure that the key still exists in the list in `remove_fromring()` or before calling the function. Additionally, the helper function `UPGRADELOCK()` could be used to clarify the code flow.



## 4.2.15 BND-CA-23-114: Pointers Dereferenced before Being Checked

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4432>

*Affected Component:* lib/isc/netmgr/streamdns.c:isc\_nm\_streamdnsconnect()

### 4.2.15.1 Description

In several places, pointers are dereferenced before being checked against *NULL*. As seen in listing 4.44, the pointer *mgr* is dereferenced to assign *worker* and then checked for validity, which includes a *NULL* pointer check. In case *mgr* is *NULL*, invalid memory is getting read which likely leads to a crash instead of a more controlled abort.

```
1 void
2 isc_nm_streamdnsconnect(isc_nm_t *mgr, isc_sockaddr_t *local,
3     isc_sockaddr_t *peer, isc_nm_cb_t cb, void *cbarg,
4     unsigned int timeout, isc_tlsctx_t *ctx,
5     isc_tlsctx_client_session_cache_t *client_sess_cache) {
6     isc_nmsocket_t *nsock = NULL;
7     isc__networker_t *worker = &mgr->workers[isc_tid()];
8
9     REQUIRE(VALID_NM(mgr));
```

Listing 4.44: mgr Dereferenced Then Validated

Similar code exists in *isc\_nm\_listenstreamdns()*, *isc\_nm\_tcpconnect()*, *isc\_nm\_listentls()*, *isc\_nm\_tlscconnect()*, *isc\_nm\_tcpconnect()* and *isc\_nm\_udpconnect()*.

*sock* is used in a similar pattern in *isc\_\_nm\_udp\_send()*. The *stats* pointer in *dns\_dnssecsignstats\_increment()* and *dns\_dnssecsignstats\_clear()* is accessed in the same way.

### 4.2.15.2 Solution Advice

X41 recommends to change the order of the validation and access of the pointers to ensure correct error messages in these failure cases.

## 4.2.16 BND-CA-23-115: Files Created with World Read/Write Permissions

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4443>

*Affected Component:* lib/isc/file.c:isc\_file\_openunique()

---

### 4.2.16.1 Description

During the code review it was found that the function `isc_file_openunique()` tries to create files with permission mode `0666` as shown in listing 4.45.

---

```
1     isc_result_t
2     isc_file_openunique(char *templet, FILE **fp) {
3         int mode = S_IWUSR | S_IRUSR | S_IRGRP | S_IWGRP | S_IROTH | S_IWOTH;
4         return (isc_file_openunique(mode, templet, mode, fp));
5     }
```

---

**Listing 4.45:** File Created with Mode 0666, World Read/Write

Unless a more restrictive `umask` is set, this results in the created file to be world read- and writable for any user on the system. The function `isc_file_openunique()` is involved in the creation of temporary files, zone files, and configuration files.

On nearly all modern systems, the `umask`<sup>23</sup> will be restrictive, mitigating a security impact because it will turn off corresponding bits requested in the file mode.

### 4.2.16.2 Solution Advice

It is recommended to remove the overly broad file mode permissions. Since on all modern systems a restrictive `umask` is set by default, the code will have not any effect and should be safe to remove.

---

<sup>23</sup><https://man.freebsd.org/cgi/man.cgi?query=umask&sektion=2>

## 4.2.17 BND-CA-23-116: Locking Inconsistencies in Cache Implementation

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4340>

*Affected Component:* lib/dns/cache.c

---

### 4.2.17.1 Description

Several inconsistencies exist in the `lib/dns/cache.c` in regards to locking.

`dns_cache_setcachesize()` locks `cache->lock` to protect `cache->size`. But setting and retrieving this variable is basically a noop, since it is never used for anything meaningful.

`cache->serve_stale_ttl` is protected by `cache->lock` as well in `dns_cache_setservestalettl()` but not in `cache->serve_stale_ttl()`, which can be called via `dns_cache_flush()`.

Furthermore, `dns_cache_getservestalerefresh()` uses `dns_db_getservestalerefresh()` to return the value instead of the locally stored copy in `cache`.

`dns_db_getservestalerefresh()` does not do any locking in the variable it sets and might therefore race against `dns_db_getservestalerefresh()`.

### 4.2.17.2 Solution Advice

X41 recommends to cleanup the locking and check which functions can be removed or need to be reimplemented.

## 4.2.18 BND-CA-23-117: Endless Loop via GENERATE

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4353>

*Affected Component:* lib/dns/master.c:genname()

### 4.2.18.1 Description

In lib/dns/master.c, in function `genname()`, `it` and `delta` are both of type `int`. There is an issue when `delta` is negative (which nothing prevents) when calling `nibbles()`. The implemented check (see listing 4.46) does not create an error, since a negative value is always smaller than `INT_MAX - it` for a positive value of `it`.

```

1  /*
2  * 'it' is >= 0 so we don't need to check for
3  * underflow.
4  */
5  if ((it > 0 && delta > INT_MAX - it)) {
6      return (ISC_R_RANGE);
7  }
8  if (nibblemode) {
9      n = nibbles(numbuf, sizeof(numbuf), width,
10     mode[0], it + delta);
11 } else {
12     n = snprintf(numbuf, sizeof(numbuf), fmt,
13     it + delta);
14 }
```

Listing 4.46: Insufficient Check in genname()

In `nibbles()`, the right shift of `value` will sign extend, and `value` will never reach 0, which causes an infinite loop (see listing 4.47).

```

1  static unsigned int
2  nibbles(char *numbuf, size_t length, unsigned int width, char mode, int value) {
3      unsigned int count = 0;
4
5      /*
6      * This reserve space for the NUL string terminator.
7      */
8      if (length > 0U) {
9          *numbuf = '\0';
10         length--;
```

```
11     }
12     do {
13         char val = hex[(value & 0x0f) + ((mode == 'n') ? 0 : 16)];
14         value >>= 4;
15         if (length > 0U) {
16             *numbuf++ = val;
17             *numbuf = '\0';
18             length--;
19         }
20         if (width > 0) {
21             width--;
22         }
23         count++;
24         /*
25          * If width is non zero then we need to add a label separator.
26          * If value is non zero then we need to add another label and
27          * that requires a label separator.
28          */
29         if (width > 0 || value != 0) {
30             if (length > 0U) {
31                 *numbuf++ = '.';
32                 *numbuf = '\0';
33                 length--;
34             }
35             if (width > 0) {
36                 width--;
37             }
38             count++;
39         }
40     } while (value != 0 || width > 0);
41     return (count);
42 }
```

**Listing 4.47:** Right Shift in nibbles()

Additionally, when *width* is a large number, the processing can take some time as well, but will terminate eventually.

Since GENERATE is only allowed in trusted zone files, it should not be possible for an attacker to trigger this.

#### 4.2.18.2 Solution Advice

X41 recommends to convert *value* to *unsigned int*.

## 4.2.19 BND-CA-23-118: Endless Loop via INCLUDE

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4357>

*Affected Component:* lib/dns/master.c:load\_text()

### 4.2.19.1 Description

When loading a zonefile, e.g. via `./named-compilezone -d -i full -o /dev/null x input` we can enter another endless loop, when the file just contains `$INCLUDE .` statement.

The loop happens via the callchain seen in listing 4.48:

```

1 #0 load_text (lctx=0x7ffff1a54300) at master.c:1083
2 #1 0x00007ffff78042f6 in dns_master_loadfile (master_file=<optimized out>, top=<optimized out>,
↳ origin=<optimized out>, zclass=<optimized out>, options=<optimized out>, resign=<optimized
↳ out>, callbacks=<optimized out>,
3 include_cb=<optimized out>, include_arg=<optimized out>, mctx=<optimized out>,
↳ format=<optimized out>, maxttl=<optimized out>) at master.c:2637
4 #2 0x00007ffff7bbd66b in zone_startload (db=0x7ffff1a18800, zone=<optimized out>,
↳ zone@entry=0x7ffff1a90000, loadtime=...) at zone.c:2641
5 #3 0x00007ffff7b6e362 in zone_load (zone=<optimized out>, flags=<optimized out>,
↳ locked=<optimized out>) at zone.c:2315
6 #4 0x000055555563cd3d in load_zone (mctx=<optimized out>, zonename=<optimized out>,
↳ filename=<optimized out>, fileformat=<optimized out>, classname=<optimized out>,
↳ maxttl=<optimized out>, zonep=<optimized out>)
7 at check-tool.c:639
8 #5 0x000055555563afb0 in main (argc=<optimized out>, argv=<optimized out>) at
↳ named-checkzone.c:546

```

Listing 4.48: Zonefile Loading Callchain

As seen in listing 4.49, the code in `load_text()` contains an infinite loop:

```

1 while (true) {
2 ...
3
4     } else if (strcasecmp(DNS_AS_STR(token), "$INCLUDE") ==
5         0)
6     {
7         COMMITALL;
8     ...
9         GETTOKEN(lctx->lex, ISC_LEXOPT_QSTRING, &token,
10             false);

```

```

11         if (include_file != NULL) {
12             isc_mem_free(mctx, include_file);
13         }
14
15     ...
16 next_line;;
17 }

```

Listing 4.49: load\_text() infinite loop

Since `gettoken()` returns `ISC_R_NOTFILE`, the jump to `next_line` is taken and the parser ends up in an infinite loop as shown in listing 4.50.

```

1  #define GETTOKENERR(lctx, options, token, eol, err) |
2      do { |
3          result = gettoken(lctx, options, token, eol, callbacks); |
4          switch (result) { |
5              case ISC_R_SUCCESS: |
6                  break; |
7              case ISC_R_UNEXPECTED: |
8                  goto insist_and_cleanup; |
9              default: |
10                 if (MANYERRS(lctx, result)) { |
11                     SETRESULT(lctx, result); |
12                     LOGIT(result); |
13                     read_till_eol = true; |
14                     err goto next_line; |
15                 ... |
16             } while (0) |
17 #define GETTOKEN(lctx, options, token, eol) \
18     GETTOKENERR(lctx, options, token, eol, {})

```

Listing 4.50: GETTOKEN Macro

#### 4.2.19.2 Solution Advice

X41 recommends to catch the `ISC_R_NOTFILE` case in the same manner as `ISC_R_UNEXPECTED`.

## 4.2.20 BND-CA-23-119: Supplied Buffer Too Large

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4433>

*Affected Component:* bin/tests/wire\_test.c

---

### 4.2.20.1 Description

The code in `wire_test.c` provides a  $64 * 1024$ -byte buffer to `dns_message_renderbegin()`. This might trigger an error in a corner case, where the code is expecting to receive buffers that are not larger than 65536 bytes as shown in listing 4.51.

---

```
1 if (result != ISC_R_SUCCESS) {
2     INSIST(st.used < 65536);
3     dns_compress_rollback(
4         msg->cctx, (uint16_t)st.used);
5     *(msg->buffer) = st; /* rollback */
6     msg->buffer->length += msg->reserved;
7     msg->counts[sectionid] += total;
8     maybe_clear_ad(msg, sectionid);
9     return (result);
10 }
```

---

Listing 4.51: `realloc()` Misuse

### 4.2.20.2 Solution Advice

X41 recommends to reduce the buffer size by one to prevent the error from happening.



## 4.2.21 BND-CA-23-120: Dead Code in DNSTAP Helper

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4406>

*Affected Component:* bin/tools/dnstap-read.c

### 4.2.21.1 Description

The code in bin/tools/dnstap-read.c allocates memory for *b* in *main()*, never uses it and correctly frees it again.

```
1  int
2  main(int argc, char *argv[]) {
3  ...
4      isc_buffer_t *b = NULL;
5  ...
6      for (;;) {
7  ...
8          if (b != NULL) {
9              isc_buffer_free(&b);
10             }
11             isc_buffer_allocate(mctx, &b, 2048);
12             if (b == NULL) {
13                 fatal("out of memory");
14             }
15         ...
16     cleanup:
17     ...
18         if (b != NULL) {
19             isc_buffer_free(&b);
20         }
21         isc_mem_destroy(&mctx);
22
23         exit(rv);
24     }
```

Listing 4.52: Dead Code in DNSTAP Helper

### 4.2.21.2 Solution Advice

X41 recommends to remove the dead code.

## 4.2.22 BND-CA-23-121: Unused AES Functions

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4421>

*Affected Component:* lib/isc/aes.c

---

### 4.2.22.1 Description

The AES<sup>24</sup> functions *isc\_aes256\_crypt()* and *isc\_aes192\_crypt()* in `lib/isc/aes.c` have no callers besides test code and are therefore dead code.

### 4.2.22.2 Solution Advice

X41 recommends to remove unused functions and dead code.

---

<sup>24</sup> Advanced Encryption Standard

## 4.2.23 BND-CA-23-122: Usage of `isc_safe_memwipe()`

---

*GitLab Issue:* <https://gitlab.isc.org/isc-projects/bind9/-/issues/4435>

*Affected Component:* `lib/isc/safe.c`

---

### 4.2.23.1 Description

The function `isc_safe_memwipe()` is used to wipe memory in a way that is not optimized away by a compiler. An optimizing compiler could remove a call to `memset()` if the memory area is passed to `free()` afterwards. The function seems to be used to remove key material from memory that is passed to BIND 9 via files. It is likely that key material is also left in heap or stack variables of the lexer used to parse the files. This might cause a false sense of security since not all key data is removed from memory. In case the threat model sees this as an actual threat, more memory areas need to be wiped.

Nevertheless, removing key material already lowers the chances of key material leaking.

### 4.2.23.2 Solution Advice

X41 recommends to clarify the threat model and either remove the function `isc_safe_memwipe()` to avoid a false sense of security or enforce proper memory wiping in more places.

## 5 About X41 D-Sec GmbH

X41 D-Sec GmbH is an expert provider for application security and penetration testing services. Having extensive industry experience and expertise in the area of information security, a strong core security team of world-class security experts enables X41 D-Sec GmbH to perform premium security services.

X41 has the following references that show their experience in the field:

- Source code audit of the Git source code version control system<sup>1</sup>
- Review of the Mozilla Firefox updater<sup>2</sup>
- X41 Browser Security White Paper<sup>3</sup>
- Review of Cryptographic Protocols (Wire)<sup>4</sup>
- Identification of flaws in Fax Machines<sup>5,6</sup>
- Smartcard Stack Fuzzing<sup>7</sup>

The testers at X41 have extensive experience with penetration testing and red teaming exercises in complex environments. This includes enterprise environments with thousands of users and vendor infrastructures such as the Mozilla Firefox Updater (Balrog).

Fields of expertise in the area of application security encompass security-centered code reviews, binary reverse-engineering and vulnerability-discovery. Custom research and IT security consulting, as well as support services, are the core competencies of X41. The team has a strong technical background and performs security reviews of complex and high-profile applications such as Google Chrome and Microsoft Edge web browsers.

X41 D-Sec GmbH can be reached via <https://x41-dsec.de> or <mailto:info@x41-dsec.de>.

<sup>1</sup> <https://x41-dsec.de/security/research/news/2023/01/17/git-security-audit-ostif/>

<sup>2</sup> <https://blog.mozilla.org/security/2018/10/09/trusting-the-delivery-of-firefox-updates/>

<sup>3</sup> <https://browser-security.x41-dsec.de/X41-Browser-Security-White-Paper.pdf>

<sup>4</sup> <https://www.x41-dsec.de/reports/Kudelski-X41-Wire-Report-phase1-20170208.pdf>

<sup>5</sup> <https://www.x41-dsec.de/lab/blog/fax/>

<sup>6</sup> <https://2018.zeronights.ru/en/reports/zero-fax-given/>

<sup>7</sup> <https://www.x41-dsec.de/lab/blog/smartcards/>

# Acronyms

<b>ACL</b> Access Control List . . . . .	12
<b>AES</b> Advanced Encryption Standard . . . . .	73
<b>ANSI</b> American National Standards Institute . . . . .	53
<b>API</b> Application Programming Interface . . . . .	36
<b>AXFR</b> Asynchronous Full Transfer Zone . . . . .	10
<b>CVE</b> Common Vulnerabilities and Exposures . . . . .	12
<b>CWE</b> Common Weakness Enumeration . . . . .	14
<b>DNS</b> Domain Name System . . . . .	6
<b>DoS</b> Denial of Service . . . . .	6
<b>EDNS</b> Extended DNS . . . . .	10
<b>HTTP</b> HyperText Transfer Protocol . . . . .	19
<b>ID</b> Identifier . . . . .	12
<b>IXFR</b> Incremental Zone Transfer . . . . .	10
<b>OOB</b> Out-of-Bounds . . . . .	9
<b>RCE</b> Remote Code Execution . . . . .	6
<b>RCU</b> Read, Copy, Update . . . . .	10
<b>TCP</b> Transmission Control Protocol . . . . .	23

# A Appendix

## A.1 Fuzz Testing

Various fuzz harnesses already exist for the BIND 9 nameserver. Among these are those shipped with `bind`<sup>12</sup> and `honggfuzz`<sup>3</sup>. X41 created additional harnesses to either target functions more directly or to test parts that were not covered by the fuzz testing before. Due to the already working fuzzing setup, these could be quickly added whenever a promising target was encountered during the source code audit.

### A.1.1 `isc_url_parse()` Fuzz Harness

The fuzz harness created for `isc_url_parse()` is shown in listing A.1.

---

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14  #include <stddef.h>
15  #include <stdint.h>
16
17  #include <isc/url.h>
18
19  #include "fuzz.h"
20
21  int
22  LLVMFuzzerInitialize(int *argc ISC_ATTR_UNUSED, char ***argv ISC_ATTR_UNUSED) {
```

<sup>1</sup><https://gitlab.isc.org/isc-projects/bind9/-/tree/main/fuzz>

<sup>2</sup><https://gitlab.isc.org/isc-projects/bind9/-/blob/main/bin/named/fuzz.c>

<sup>3</sup><https://github.com/google/honggfuzz/tree/master/examples/bind>

```
23
24     return (0);
25 }
26
27 int
28 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
29
30     isc_url_parser_t p;
31
32     memset(&p, 0, sizeof(p));
33     isc_url_parse(data, size, true, &p);
34     memset(&p, 0, sizeof(p));
35     isc_url_parse(data, size, false, &p);
36
37     return (0);
38 }
```

Listing A.1: isc\_url\_parse() Fuzz Harness

## A.1.2 phr\_parse\_request() Fuzz Harness

The fuzz harness created for `phr_parse_request()` is shown in listing A.2.

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14  #include <stddef.h>
15  #include <stdint.h>
16
17  #include <picohttpparser.h>
18
19  #include "fuzz.h"
20
21  #define HTTP_HEADERS_NUM    100
22
23  int
24  LLVMFuzzerInitialize(int *argc ISC_ATTR_UNUSED, char ***argv ISC_ATTR_UNUSED) {
25
```

```
26     return (0);
27 }
28
29 int
30 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
31     const char *method = NULL;
32     size_t method_len = 0;
33     const char *path;
34     size_t path_len = 0;
35     int minor_version;
36     struct phr_header headers[HTTP_HEADERS_NUM];
37     size_t num_headers = HTTP_HEADERS_NUM;
38
39     phr_parse_request((const char *) data, size, &method,
40                     &method_len, &path, &path_len,
41                     &minor_version, headers,
42                     &num_headers, size);
43
44     return (0);
45 }
```

Listing A.2: phr\_parse\_request() Fuzz Harness

### A.1.3 isc\_regex\_validate() Fuzz Harness

The fuzz harness created for *isc\_regex\_validate()* is shown in listing A.3.

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14  #include <stddef.h>
15  #include <stdint.h>
16
17  #include <isc/regex.h>
18
19  #include "fuzz.h"
20
21  int
```



```
22 LLVMFuzzerInitialize(int *argc ISC_ATTR_UNUSED, char ***argv ISC_ATTR_UNUSED) {
23
24     return (0);
25 }
26
27 int
28 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
29
30
31     char *r = malloc(size + 1);
32     memcpy(r, data, size);
33     r[size] = 0;
34
35     isc_regex_validate(r);
36
37     free(r);
38     return (0);
39 }
```

Listing A.3: isc\_regex\_validate() Fuzz Harness

#### A.1.4 isc\_utf8\_valid() Fuzz Harness

The fuzz harness created for *isc\_utf8\_valid()* is shown in listing A.4.

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14  #include <stddef.h>
15  #include <stdint.h>
16
17  #include <isc/utf8.h>
18
19  #include "fuzz.h"
20
21  int
22  LLVMFuzzerInitialize(int *argc ISC_ATTR_UNUSED, char ***argv ISC_ATTR_UNUSED) {
23
```

```
24     return (0);
25 }
26
27 int
28 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
29
30     isc_utf8_valid(data, size);
31
32     return (0);
33 }
```

---

Listing A.4: isc\_utf8\_valid() Fuzz Harness

### A.1.5 isccc\_cc\_fromwire() Fuzz Harness

The fuzz harness created for *isccc\_cc\_fromwire()* is shown in listing A.5.

---

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14  #include <stdbool.h>
15  #include <stdlib.h>
16
17  #include <isc/buffer.h>
18  #include <isc/mem.h>
19  #include <isc/util.h>
20
21  #include <isccc/alist.h>
22  #include <isccc/cc.h>
23  #include <isccc/ccmsg.h>
24  #include <isccc/seexpr.h>
25  #include <isccc/symtab.h>
26  #include <isccc/util.h>
27
28  #include "fuzz.h"
29
30  bool debug = false;
31
```

```
32 int
33 LLVMFuzzerInitialize(int *argc, char ***argv) {
34     UNUSED(argc);
35     UNUSED(argv);
36     return (0);
37 }
38
39 int
40 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
41
42     isccc_sexpr_t *alistp = NULL;
43     isccc_region_t source;
44
45     source.rstart = (char *) data;
46     source.rend = (char *) data + size;
47
48     // fuzzing table_fromwire() directly would be more effective,
49     // but its currently not exported
50
51     isccc_cc_fromwire(&source, &alistp, 0, NULL);
52
53     if (alistp)
54         isccc_sexpr_free(&alistp);
55     return 0;
56 }
```

Listing A.5: isccc\_cc\_fromwire() Fuzz Harness

## A.1.6 irs\_resconf\_load() AFL Harness

The AFL harness created for *irs\_resconf\_load()* is shown in listing A.6.

```
1  /*
2   * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3   *
4   * SPDX-License-Identifier: MPL-2.0
5   *
6   * This Source Code Form is subject to the terms of the Mozilla Public
7   * License, v. 2.0. If a copy of the MPL was not distributed with this
8   * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9   *
10  * See the COPYRIGHT file distributed with this work for additional
11  * information regarding copyright ownership.
12  */
13
14
15 #include <stdlib.h>
16
```

```
17 #include <isc/mem.h>
18
19 #include <irs/resconf.h>
20
21 /*% Main processing routine for dig */
22 int
23 main(int argc, char **argv) {
24     isc_result_t result;
25     irs_resconf_t *resconf = NULL;
26     static isc_mem_t *mctx = NULL;
27     char *fn;
28
29     if (argc < 2)
30         return -1;
31     fn = argv[1];
32
33     isc_mem_create(&mctx);
34
35     result = irs_resconf_load(mctx, fn, &resconf);
36
37     if (resconf != NULL) {
38         irs_resconf_destroy(&resconf);
39     }
40
41
42     return 0;
43 }
```

Listing A.6: irs\_resconf\_load() AFL Harness

## A.1.7 dig/host AFL++ ARGV Harness

Only very light argument fuzzing was performed using AFL++ on the `dig` and `host` tools. The reasoning for this test was that they might be used on cheap routers with attacker-influenced command line arguments stemming from the web interface. The main parts of that harness can be seen in listing A.7.

```
1 ...
2 #include "/home/eric/tools/AFLplusplus/utils/argv_fuzzing/argv-fuzz-inl.h"
3
4
5 /*% Main processing routine for dig */
6 int
7 main(int argc, char **argv) {
8     isc_result_t result;
9     irs_resconf_t *resconf = NULL;
10    static isc_mem_t *mctx = NULL;
11    char *fn;
```

```

12
13     AFL_INIT_SET0("dig");
14
15     ...

```

Listing A.7: dig/host AFL++ ARGV Harness

## A.1.8 Scapy Template-Based Fuzzing

The following simple template-based fuzzer generates multiple packets to be sent in (nearly) parallel:

```

1  #!/usr/bin/env python3
2
3  from scapy.all import DNS, DNSQR, IPv6, sr1, UDP, fuzz, send
4
5  dns_pkt1 = fuzz(DNS(qd=fuzz(DNSQR(qname="kk.foobar"))))
6  dns_pkt11 = fuzz(DNS(qd=fuzz(DNSQR(qname="kk.foobar\x00"))))
7  dns_pkt12 = fuzz(DNS(qd=fuzz(DNSQR(qname="41.41.41.41"))))
8  dns_pkt13 = fuzz(DNS(qd=fuzz(DNSQR(qname="::ff"))))
9  dns_pkt2 = fuzz(DNS(qd=fuzz(DNSQR())))
10 dns_pkt3 = fuzz(DNS(rd=1, qd=DNSQR()))
11 while 1:
12     send(IPv6(dst='::1')/UDP(sport=53, dport=53)/dns_pkt1, verbose=0)
13     send(IPv6(dst='::1')/UDP(sport=53, dport=53)/dns_pkt11, verbose=0)
14     send(IPv6(dst='::1')/UDP(sport=53, dport=53)/dns_pkt12, verbose=0)
15     send(IPv6(dst='::1')/UDP(sport=53, dport=53)/dns_pkt13, verbose=0)
16     send(IPv6(dst='::1')/UDP(sport=53, dport=53)/dns_pkt2, verbose=0)
17     send(IPv6(dst='::1')/UDP(dport=53)/dns_pkt3, verbose=0)

```

Listing A.8: Scapy Fuzz Script

## A.1.9 Speedup of dns\_rdata\_fromwire\_text.c

The fuzzing harness `fuzz/dns_rdata_fromwire_text.c` can be sped up by a factor of 200 by moving some initialisation routines into `LLVMFuzzerInitialize()` as shown in listing A.9.

```

1  --- bind-9.19.17-orig/fuzz/dns_rdata_fromwire_text.c    2023-09-08 11:09:28.635415105 +0200
2  +++ bind-9.19.17/fuzz/dns_rdata_fromwire_text.c      2023-09-25 08:46:45.028588658 +0200
3  @@ -40,6 +40,8 @@ bool debug = false;
4
5     static isc_mem_t *mctx = NULL;
6     static isc_lex_t *lex = NULL;
7  +unsigned int types = 1;

```

```
8 +dns_rdatatype_t typelist[256] = { 1000 }; /* unknown */
9
10 int
11 LLVMFuzzerInitialize(int *argc ISC_ATTR_UNUSED, char ***argv ISC_ATTR_UNUSED) {
12 @@ -56,6 +58,24 @@ LLVMFuzzerInitialize(int *argc ISC_ATTR_
13     isc_lex_setspecials(lex, specials);
14     isc_lex_setcomments(lex, ISC_LEXCOMMENT_DNSMASTERFILE);
15
16 + unsigned int t;
17 + /*
18 +  * Append known types to list.
19 +  */
20 + for (t = 1; t <= 0x10000; t++) {
21 +     char typebuf[256];
22 +     if (dns_rdatatype_ismeta(t)) {
23 +         continue;
24 +     }
25 +     dns_rdatatype_format(t, typebuf, sizeof(typebuf));
26 +     if (strncmp(typebuf, "TYPE", 4) != 0) {
27 +         /* Assert when we need to grow typelist. */
28 +         assert(types < sizeof(typelist) / sizeof(typelist[0]));
29 +         typelist[types++] = t;
30 +     }
31 + }
32 +
33 +
34     return (0);
35 }
36
37 @@ -73,13 +93,13 @@ nullmsg(dns_rdatacallbacks_t *cb, const
38     }
39 }
40
41 +
42 int
43 LLVMFuzzerTestOneInput(const uint8_t *data, size_t size) {
44     char totext[64 * 1024 * 4];
45     dns_compress_t cctx;
46     dns_rdatatype_t rdtype;
47     dns_rdataclass_t rdclass;
48 - dns_rdatatype_t typelist[256] = { 1000 }; /* unknown */
49     dns_rdataclass_t classlist[] = { dns_rdataclass_in, dns_rdataclass_hs,
50                                     dns_rdataclass_ch, dns_rdataclass_any,
51                                     60 };
52 @@ -92,7 +112,7 @@ LLVMFuzzerTestOneInput(const uint8_t *da
53     unsigned char fromwire[1024];
54     unsigned char towire[1024];
55     unsigned int classes = (sizeof(classlist) / sizeof(classlist[0]));
56 - unsigned int types = 1, flags, t;
57 + unsigned int flags;
58
59     /*
```

```
60     * First 2 bytes are used to select type and class.
61     @@ -102,21 +122,6 @@ LLVMFuzzerTestOneInput(const uint8_t *da
62         return (0);
63     }
64
65     - /*
66     -  * Append known types to list.
67     -  */
68     - for (t = 1; t <= 0x10000; t++) {
69     -     char typebuf[256];
70     -     if (dns_rdatatype_ismeta(t)) {
71     -         continue;
72     -     }
73     -     dns_rdatatype_format(t, typebuf, sizeof(typebuf));
74     -     if (strncmp(typebuf, "TYPE", 4) != 0) {
75     -         /* Assert when we need to grow typelist. */
76     -         assert(types < sizeof(typelist) / sizeof(typelist[0]));
77     -         typelist[types++] = t;
78     -     }
79     - }
80
81     /*
82     * Random type and class from a limited set.
```

---

#### Listing A.9: Speedup of dns\_rdata\_fromwire\_text.c

Additionally, it is recommended to add a call to `isc_lex_close()` to match the calls to `isc_lex_openbuffer()` to avoid memory leaks. The same memory leaks appear in `isc_lex_getmastertoken.c`.

## A.1.10 badcache Stress Testing

Stress testing was performed against the badcache implementation by modifying tests/dns/badcache\_test.c. The modified test makes modifications to a shared cache in parallel.

```
1  /*
2  * Copyright (C) Internet Systems Consortium, Inc. ("ISC")
3  *
4  * SPDX-License-Identifier: MPL-2.0
5  *
6  * This Source Code Form is subject to the terms of the Mozilla Public
7  * License, v. 2.0. If a copy of the MPL was not distributed with this
8  * file, you can obtain one at https://mozilla.org/MPL/2.0/.
9  *
10 * See the COPYRIGHT file distributed with this work for additional
11 * information regarding copyright ownership.
12 */
13
14 #include <inttypes.h>
15 #include <sched.h> /* IWYU pragma: keep */
16 #include <stdarg.h>
17 #include <stdbool.h>
18 #include <stddef.h>
19 #include <stdio.h>
20 #include <stdlib.h>
21 #include <string.h>
22 #include <unistd.h>
23
24 #include <isc/buffer.h>
25 #include <isc/mem.h>
26 #include <isc/os.h>
27 #include <isc/thread.h>
28 #include <isc/urcu.h>
29 #include <isc/util.h>
30 #include <isc/wv.h>
31 #include <isc/loop.h>
32 #include <isc/async.h>
33
34 #include <dns/fixedname.h>
35 #include <dns/badcache.h>
36 #include <dns/name.h>
37 #include <dns/rdatatype.h>
38
39 #define BADCACHE_TEST_FLAG 1 << 3
40
41 static void
42 basic(dns_badcache_t *bc) {
43     dns_fixedname_t fname = { 0 };
44     dns_name_t *name = dns_fixedname_initname(&fname);
45     isc_stdtime_t now = isc_stdtime_now();
46     uint32_t flags = BADCACHE_TEST_FLAG;
```



```
47
48     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
49     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
50
51     flags = 0;
52     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
53     flags = 0;
54     dns_badcache_find(bc, name, dns_rdatatype_a, &flags, now);
55 }
56
57 static void
58 expire(dns_badcache_t *bc) {
59     dns_fixedname_t fname = { 0 };
60     dns_name_t *name = dns_fixedname_initname(&fname);
61     isc_stdtime_t now = isc_stdtime_now();
62     uint32_t flags = BADCACHE_TEST_FLAG;
63
64     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
65     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
66     dns_badcache_add(bc, name, dns_rdatatype_a, false, flags, now + 60);
67     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
68     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
69                       now + 61);
70
71     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
72     dns_badcache_find(bc, name, dns_rdatatype_a, &flags, now);
73     dns_badcache_add(bc, name, dns_rdatatype_a, true, flags, now + 120);
74     dns_badcache_find(bc, name, dns_rdatatype_a, &flags, now + 61);
75 }
76
77 static void
78 print(dns_badcache_t *bc) {
79     dns_fixedname_t fname = { 0 };
80     dns_name_t *name = dns_fixedname_initname(&fname);
81     isc_stdtime_t now = isc_stdtime_now();
82     isc_stdtime_t expire = now + 60;
83     uint32_t flags = BADCACHE_TEST_FLAG;
84     FILE *file = NULL;
85
86     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
87     dns_badcache_add(bc, name, dns_rdatatype_a, false, flags, expire);
88     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, expire);
89
90     file = fopen("/dev/null", "w");
91     dns_badcache_print(bc, "badcache", file);
92     fclose(file);
93 }
94
95 static void
96 flush(dns_badcache_t *bc) {
97     dns_fixedname_t fname = { 0 };
98     dns_name_t *name = dns_fixedname_initname(&fname);
```

```
99     isc_stdtime_t now = isc_stdtime_now();
100     uint32_t flags = BADCACHE_TEST_FLAG;
101
102     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
103     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
104     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
105     dns_badcache_flush(bc);
106     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
107 }
108
109 static void
110 flushname(dns_badcache_t *bc) {
111     dns_fixedname_t fname = { 0 };
112     dns_name_t *name = dns_fixedname_initname(&fname);
113     isc_stdtime_t now = isc_stdtime_now();
114     uint32_t flags = BADCACHE_TEST_FLAG;
115
116     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
117     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
118     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
119     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
120     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
121     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
122     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
123     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
124     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
125     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
126     dns_badcache_flushname(bc, name);
127     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
128     dns_badcache_find(bc, name, dns_rdatatype_a, &flags, now);
129     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
130     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
131     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
132     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
133 }
134
135 static void
136 flushtree(dns_badcache_t *bc) {
137     dns_fixedname_t fname = { 0 };
138     dns_name_t *name = dns_fixedname_initname(&fname);
139     isc_stdtime_t now = isc_stdtime_now();
140     uint32_t flags = BADCACHE_TEST_FLAG;
141
142     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
143     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
144     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
145     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
146     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
147     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
148     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
149     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now + 60);
150     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags, now);
```

```
151     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
152     dns_badcache_flush(b, name);
153     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
154     dns_badcache_find(b, name, dns_rdatatype_aaaa, &flags, now);
155     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
156     dns_badcache_find(b, name, dns_rdatatype_aaaa, &flags, now);
157     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
158     dns_badcache_find(b, name, dns_rdatatype_aaaa, &flags, now);
159
160 }
161
162 static void
163 purge(dns_badcache_t *bc) {
164     dns_fixedname_t fname = { 0 };
165     dns_name_t *name = dns_fixedname_initname(&fname);
166     isc_stdtime_t now = isc_stdtime_now();
167     uint32_t flags = BADCACHE_TEST_FLAG;
168
169     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
170     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now);
171     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now);
172     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
173                       now - 60);
174     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
175     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now);
176     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now);
177     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
178                       now - 60);
179     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
180     dns_badcache_add(bc, name, dns_rdatatype_aaaa, false, flags, now);
181     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
182                       now - 60);
183     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
184                       now + 30);
185     dns_name_fromstring(name, "sub.sub.example.com.", NULL, 0, NULL);
186     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
187                       now + 30);
188     dns_name_fromstring(name, "sub.example.com.", NULL, 0, NULL);
189     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
190                       now + 30);
191     dns_name_fromstring(name, "example.com.", NULL, 0, NULL);
192     dns_badcache_find(bc, name, dns_rdatatype_aaaa, &flags,
193                       now + 30);
194 }
195
196 isc_loopmgr_t *loopmgr = NULL;
197
198 void loop(dns_badcache_t *bc);
199 void loop(dns_badcache_t *bc) {
200     int i = 0;
201     for (i = 0; i < 1000; i++) {
202         basic(bc);
```

```
203     expire(bc);
204     print(bc);
205     flush(bc);
206     flushname(bc);
207     flushtree(bc);
208     purge(bc);
209 }
210 isc_loopmgr_shutdown(loopmgr);
211 }
212 int main(__attribute__((unused)) int argc, __attribute__((unused)) char **argv) {
213     dns_badcache_t *bc = NULL;
214     isc_mem_t *mctx = NULL;
215
216     isc_mem_create(&mctx);
217     isc_loopmgr_create(mctx, 100, &loopmgr);
218
219     bc = dns_badcache_new(mctx);
220
221     isc_loopmgr_setup(loopmgr, (isc_job_cb) loop, bc);
222     isc_loopmgr_run(loopmgr);
223
224
225     dns_badcache_destroy(&bc);
226     isc_loopmgr_shutdown(loopmgr);
227 }
```

---

**Listing A.10:** badcache Stress Testing

## A.2 CodeQL Queries

The following queries were used to inspect the codebase and may be desirable to integrate into automated analysis pipelines:

### A.2.1 Mutex DataFlow Query

```

1  iimport cpp
2  import semmle.code.cpp.dataflow.DataFlow
3
4  class Config extends DataFlow::Configuration {
5    Config() { this = "MutexVarFlow" }
6
7    override predicate isSource(DataFlow::Node source) {
8      exists(FunctionCall alloc, DataFlow::Node v |
9        alloc.getTarget().getName() = "pthread_mutex_init" and
10       v.asExpr() = alloc.getArgument(0).getFullyConverted()
11      )
12    }
13
14    override predicate isSink(DataFlow::Node sink) {
15      exists(FunctionCall dealloc |
16        dealloc.getTarget().getName() = "pthread_mutex_destroy" and
17        sink.asExpr() = dealloc.getArgument(0).getFullyConverted()
18      )
19    }
20  }
21
22  from Config cfg, DataFlow::Node src, DataFlow::Node sink
23  where cfg.hasFlow(src, sink) and src.getLocation() != sink.getLocation()
24  select src, sink

```

Listing A.11: Mutex Dataflow to Find Correct Usage of Mutexes

### A.2.2 Verification of Tainted Length Variable Dataflow

```

1  import cpp
2  import semmle.code.cpp.dataflow.TaintTracking
3
4  class BufferCall extends FunctionCall {
5    BufferCall() {
6      this.getTarget().getName().matches("isc_buffer_get%")
7      or this.getTarget().getName().matches("isc_buffer_peek%")
8      or this.getTarget().getName().matches("%read%")
9      or this.getTarget().getName().matches("%ntohs%")
10     or this.getTarget().getName().matches("%bswap%")

```

```

11     }
12
13     Expr getVal() {
14         //result = this.getArgument(1)
15         result = this.getAnArgument()
16     }
17 }
18
19
20 class MemCall extends FunctionCall {
21     MemCall() {
22         this.getTarget().hasName("memmove")
23         or this.getTarget().hasName("memcpy")
24         or this.getTarget().hasName("strncpy")
25         or this.getTarget().hasName("strcpy")
26     }
27
28     Expr getFileDescriptor() {
29         result = this.getArgument(0)
30     }
31 }
32
33 class Config extends TaintTracking::Configuration {
34     Config() { this = "MemoryTaintFlow" }
35
36     override predicate isSource(DataFlow::Node source) {
37         exists(BufferCall call |
38             source.asExpr() = call.getVal()
39         )
40     }
41
42     override predicate isSink(DataFlow::Node sink) {
43         exists(MemCall call |
44             sink.asExpr() = call.getArgument(2) // or whichever argument you're interested in
45         )
46     }
47 }
48
49 from Config cfg, DataFlow::PathNode source, DataFlow::PathNode sink
50 where cfg.hasFlowPath(source, sink)
51 select sink.getNode(), source, source.getNode().asExpr().getEnclosingElement(), sink,
↪ sink.getNode().asExpr().getActualType(), source.getNode().asExpr().getActualType()

```

Listing A.12: Tainted Operation Length Value Dataflow

## A.2.3 Tainted Printf Function Family Dataflows

```

1  import cpp
2  import semmle.code.cpp.dataflow.TaintTracking
3
4  class BufferCall extends FunctionCall {
5      BufferCall() {
6          this.getTarget().getName().matches("isc_buffer_get%")
7          or    this.getTarget().getName().matches("isc_buffer_peek%")
8          or    this.getTarget().getName().matches("%read%")
9          or    this.getTarget().getName().matches("%ntohs%")
10         or    this.getTarget().getName().matches("%bswap%")
11     }
12
13     Expr getVal() {
14         //result = this.getArgument(1)
15         result = this.getAnArgument()
16     }
17 }
18
19
20 class FormatCall extends FunctionCall {
21     FormatCall() {
22         this.getTarget().getName().matches("%printf%")
23         and not this.getTarget().getName().matches("printf")
24     }
25     Expr getFormatString() {
26         result = this.getArgument(1)
27     }
28 }
29
30 class Config extends TaintTracking::Configuration {
31     Config() { this = "PrintfTaintFlow" }
32
33     override predicate isSource(DataFlow::Node source) {
34         exists(BufferCall call |
35             source.asExpr() = call.getVal()
36         )
37     }
38
39     override predicate isSink(DataFlow::Node sink) {
40         exists(FormatCall call |
41             sink.asExpr() = call.getFormatString() // or whichever argument you're interested in
42         )
43     }
44 }
45
46 from Config cfg, DataFlow::PathNode source, DataFlow::PathNode sink
47 where cfg.hasFlowPath(source, sink)
48 select sink.getNode(), source, source.getNode().asExpr().getEnclosingElement(), sink,
↪ sink.getNode().asExpr().getActualType(), source.getNode().asExpr().getActualType()

```

## Listing A.13: Tainted Printf Function Family Dataflow

## A.2.4 Tainted ISC Memory Wrapper Dataflows

```

1  import cpp
2  import semmle.code.cpp.dataflow.DataFlow
3
4  class Config extends DataFlow::Configuration {
5    Config() { this = "GetAndNotPutFlow" }
6    override predicate isSource(DataFlow::Node source) {
7      exists(FunctionCall alloc, Assignment a, Variable v |
8        (alloc.getTarget().getName() = "isc__mem_get" or alloc.getTarget().getName() =
9         ↪ "isc__mem_getx") and
10       a.getRValue() = alloc and
11       v = a.getLValue().(VariableAccess).getTarget() and
12       source.asExpr() = v.getAnAccess()
13     )
14   }
15
16   override predicate isSink(DataFlow::Node sink) {
17     exists(FunctionCall dealloc |
18       (dealloc.getTarget().getName() = "mem_put" or dealloc.getTarget().getName() = "abort" or
19        ↪ dealloc.getTarget().getName() = "ISC_LIST_APPEND")
20       and
21       sink.asExpr() = dealloc.getArgument(1)
22     )
23   }
24 }
25 from Config cfg, FunctionCall alloc, Assignment a, Variable v
26 where
27   (alloc.getTarget().getName() = "isc__mem_get" or alloc.getTarget().getName() = "isc__mem_getx")
28   ↪ and
29   a.getRValue() = alloc and
30   v = a.getLValue().(VariableAccess).getTarget()
31   and not exists(DataFlow::Node src, DataFlow::Node sink |
32     src.asExpr() = v.getAnAccess() and cfg.hasFlow(src, sink)
33   )
34 select alloc, v, alloc.getLocation(), v.getLocation()

```

## Listing A.14: Tainted ISC Memory Wrapper Dataflow



## A.2.5 Query to Find Unsafe RCU Dereference Dataflows

```

1 import cpp
2 import semmle.code.cpp.dataflow.DataFlow
3
4 class RcuDereferenceFunction extends Function {
5   RcuDereferenceFunction() { this.getName().matches("%r_deref%") }
6   // TODO this is hackish since it overfits, but there
7   // are potentially lots of wrappers for rcu_dereference
8 }
9
10 class Config extends DataFlow::Configuration {
11   Config() { this = "RCUDereferenceFuzzyFlow" }
12
13   override predicate isSource(DataFlow::Node source) {
14     exists(Call call, RcuDereferenceFunction func |
15       call.getTarget() = func and
16       source.asExpr() = call
17     )
18   }
19
20   override predicate isSink(DataFlow::Node sink) {
21     exists(VariableAccess va | va = sink.asExpr())
22   }
23 }
24
25 from Config cfg, DataFlow::Node source, DataFlow::Node sink
26 where cfg.hasFlow(source, sink)
27 select source, sink, "Variable influenced by a call to rcu_dereference() or a similar function"

```

Listing A.15: RCU Dereference Dataflows (Overfitting)

## A.2.6 Query to Find Second Order Recursive Calls

```

1 import cpp
2
3 from Function f1, Function f2, Call call1, Call call2
4 where call1.getEnclosingFunction() = f1 and
5       call1.getTarget() = f2 and
6       call2.getEnclosingFunction() = f2 and
7       call2.getTarget() = f1
8 select call1, call2, call1.getLocation()

```

Listing A.16: Second Order Recursive Calls: A()->B()->A()

## A.2.7 Query (Path-Query) to Find Locks Held during While And For Loops

```
1  /**
2  * @kind path-problem
3  */
4
5  import cpp
6  import semmle.code.cpp.dataflow.TaintTracking
7
8  class LockCall extends FunctionCall {
9      LockCall() {
10         this.getTarget().getName().matches("%\\_lock%")
11     }
12
13     Expr getVal() {
14         //result = this.getArgument(1)
15         result = this.getAnArgument()
16     }
17 }
18
19
20 class UnlockCall extends FunctionCall {
21     UnlockCall() {
22         this.getTarget().getName().matches("%\\_unlock%")
23     }
24
25     Expr getLockArg() {
26         result = this.getArgument(0)
27     }
28 }
29
30 class Config extends TaintTracking::Configuration {
31     Config() { this = "LockTaintFlow" }
32
33     override predicate isSource(DataFlow::Node source) {
34         exists(LockCall call |
35             source.asExpr() = call.getAnArgument().getFullyConverted()
36             or source.asExpr() = call.getAnArgument()
37         )
38     }
39
40     override predicate isSink(DataFlow::Node sink) {
41         exists(UnlockCall call |
42             sink.asExpr() = call.getAnArgument().getFullyConverted()
43             ↪ // or whichever argument you're interested in
44             or sink.asExpr() = call.getAnArgument() // or whichever argument you're interested in
45         )
46     }
47 }
48
49 // UNUSED: this recursive one may allow to filter out paths that
```

```
49     // have an unlock before the while, but it is very slow
50     predicate hasUnlockOnPath(BasicBlock start, BasicBlock end) {
51         // Base case: the start block is a block of a ForStmt
52         start.getANode() instanceof UnlockCall
53         or
54         // Recursive case: there exists a block on the path from
55         // start to end that is a block of a ForStmt
56         exists(BasicBlock mid |
57             start.getASuccessor() = mid and
58             mid != end
59             and hasUnlockOnPath(mid, end)
60         )
61     }
62
63
64
65     query predicate edges(BasicBlock a, BasicBlock b) {
66         b = a.getASuccessor() and
67         not a.getANode() instanceof UnlockCall and
68         not (a.getANode() instanceof WhileStmt or a.getANode() instanceof ForStmt) // and
69     }
70
71     from Config cfg, LockCall start, WhileStmt end, ControlFlowNode entryPoint
72     where
73     entryPoint = start.getEnclosingBlock() and
74     edges+(entryPoint.getBasicBlock(), end.getBasicBlock())
75     select end, entryPoint.getBasicBlock(), end.getBasicBlock(), "Lock taint" +
76     ↪ entryPoint.toString()
```

Listing A.17: Query Finding Paths Between Lock->While, Lock->For, and Lock->Unlock