



# Iceberg

## Automated Verification of DNS Authoritative Engines via Just-in-Time Summarization

Yuxing Xiang, Rilin Huang, Naiqian Zheng, Xin Jin

Peking University

# DNS: critical infrastructure

- Every website visit, email, or API call starts with a DNS lookup.



DNS Authoritative Engines



```
;; QUESTION SECTION:  
;google.com.      IN  A  
  
;; ANSWER SECTION:  
google.com.     15  IN  A 142.250.197.46  
  
;; Query time: 6 msec  
;; SERVER: 127.0.0.1#53(127.0.0.1)  
;; WHEN: Tue Apr 28 17:30:34 CST 2026  
;; MSG SIZE  rcvd: 54
```

# DNS: critical infrastructure

- Every website visit, email, or API call starts with a DNS lookup.
- At the heart of DNS are **authoritative engines** — software that answers queries by looking up records in configured zones.



## DNS Authoritative Engines



**CLOUDFLARE®**



```
;; QUESTION SECTION:  
;google.com.      IN  A  
  
;; ANSWER SECTION:  
google.com.     15  IN  A 142.250.197.46  
  
;; Query time: 6 msec  
;; SERVER: 127.0.0.1#53(127.0.0.1)  
;; WHEN: Tue Apr 28 17:30:34 CST 2026  
;; MSG SIZE  rcvd: 54
```

DNS engine bugs cause real outages

# DNS engine bugs cause real outages

Oct 2025

AWS us-east-1: DNS bug cascades into 15-hour Internet outage

## How Tiny DNS Mistake That Crashed AWS: Inside the Outage That Shook the Internet



Aman Pathak | DevOps | AWS | K8s | Terraform | ML

Follow

5 min read · Oct 21, 2025



---

Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region

# DNS engine bugs cause real outages

Oct 2025

AWS us-east-1: DNS bug cascades into 15-hour Internet outage

## How Tiny DNS Mistake That Crashed AWS: Inside the Outage That Shook the Internet



Aman Pathak | DevOps | AWS | K8s | Terraform | ML

Follow

5 min read · Oct 21, 2025



Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region

Oct 2021

.club / .hsbc SERVFAIL for 3 hours — authoritative engine bug

Oct 2021


Facebook 6-hour outage — DNS unreachable compounded by BGP issue

2018 / 2021

Google, Slack — DNS-related outages




Why is correctness hard?

# Why is correctness hard?



**Complex protocol** — dozens of RFCs; record types, wildcards, redirections, glue all interact

# Why is correctness hard?

 <b>coredns</b> CoreDNS is a D ● Go ★ 14.0k	 <b>hickory-dns</b> A Rust based DNS ● Rust ★ 5.2k	 <b>pdns</b> PowerDNS Authoritative, PowerDNS Recursor, dnsmdist ● C++ ★ 4.4k 🧑 1.0k
--	---	---

Up to 10,000 core LOC



**Complex protocol** — dozens of RFCs; record types, wildcards, redirections, glue all interact

**Large, diverse implementations** — 2K–10K lines of core logic in Go / C / C++ / Rust

# Why is correctness hard?



dns1.example.com

stu.pku.edu.cn

www.not-exist.org

Query

DNS Zone Editor

```
16 dns1 IN A 5.4.3.2
17 dns2 IN A 4.3.2.1
18 server1 IN A 4.4.3.2
19 server2 IN A 5.5.4.3
20 ftp IN A 3.3.2.1
21 | IN A 3.3.3.2
22 mail IN CNAME server1
23 mail2 IN CNAME server2
24 www IN CNAME server1
```

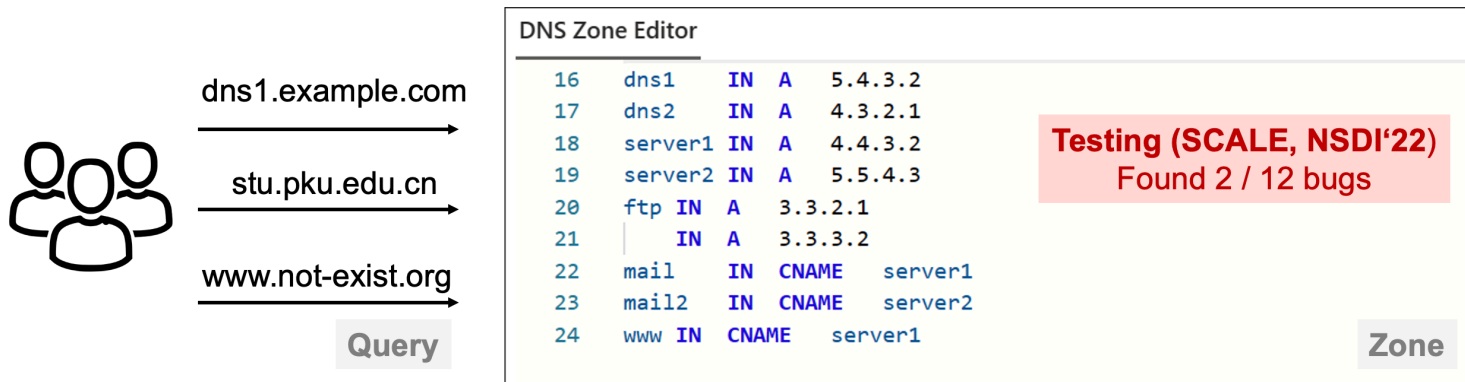
Zone

**Complex protocol** — dozens of RFCs; record types, wildcards, redirections, glue all interact

**Large, diverse implementations** — 2K–10K lines of core logic in Go / C / C++ / Rust

**Testing is limited** — query × zone input space is too large for full coverage

# Why is correctness hard?



**Complex protocol** — dozens of RFCs; record types, wildcards, redirections, glue all interact

**Large, diverse implementations** — 2K–10K lines of core logic in Go / C / C++ / Rust

**Testing is limited** — query × zone input space is too large for full coverage

# Verification can guarantee correctness

[www.cs.columbia.edu/~rgu/assets/media/certikos\\_layer.jpg](http://www.cs.columbia.edu/~rgu/assets/media/certikos_layer.jpg)

- **Formal verification** proves correctness against a **specification**, for **all** inputs

# Verification can guarantee correctness

[www.cs.columbia.edu/~rgu/assets/media/certikos\\_layer.jpg](http://www.cs.columbia.edu/~rgu/assets/media/certikos_layer.jpg)

- **Formal verification** proves correctness against a **specification**, for **all** inputs
- **Refinement proof**: prove by layers, with layer-wise specs

# Verification can guarantee correctness

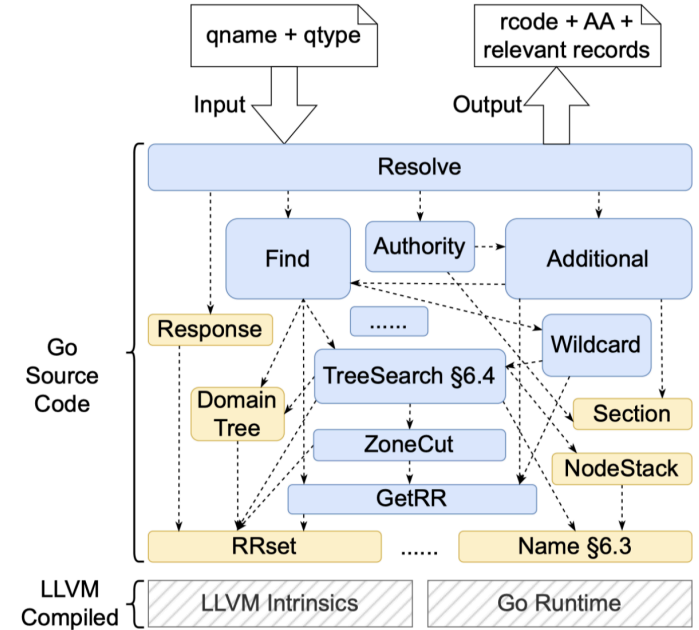
[www.cs.columbia.edu/~rgu/assets/media/certikos\\_layer.jpg](http://www.cs.columbia.edu/~rgu/assets/media/certikos_layer.jpg)

- **Formal verification** proves correctness against a **specification**, for **all** inputs
- **Refinement proof**: prove by layers, with layer-wise specs

**But:** Refinement proof is a lot of manual work!

# Prior work partially automated this

DNS-V (SOSP'23) verified an in-production engine at Alibaba Cloud using **ahead-of-time (AOT) summarization**.



# Prior work partially automated this

DNS-V (SOSP'23) verified an in-production engine at Alibaba Cloud using **ahead-of-time (AOT) summarization**.

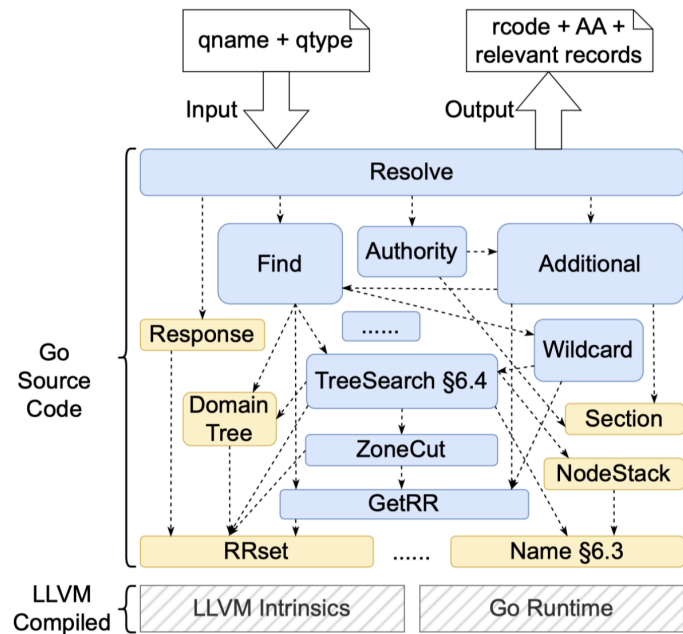
- Higher-level modules → automatically **summarize** via symbolic execution.

```
if (nameLen != 0 && n0 == int("com")){
    match_result := EXACT;
    match_NodePtr := NODE("com.");
}
else if (nameLen != 0 && n0 != int("com")){
    match_result := NOMATCH;
    match_NodePtr := NULL_NODE;
}
else {
    match_result := WILDCARD;
    match_NodePtr := NODE(".");
}
```

Path Condition

Effect

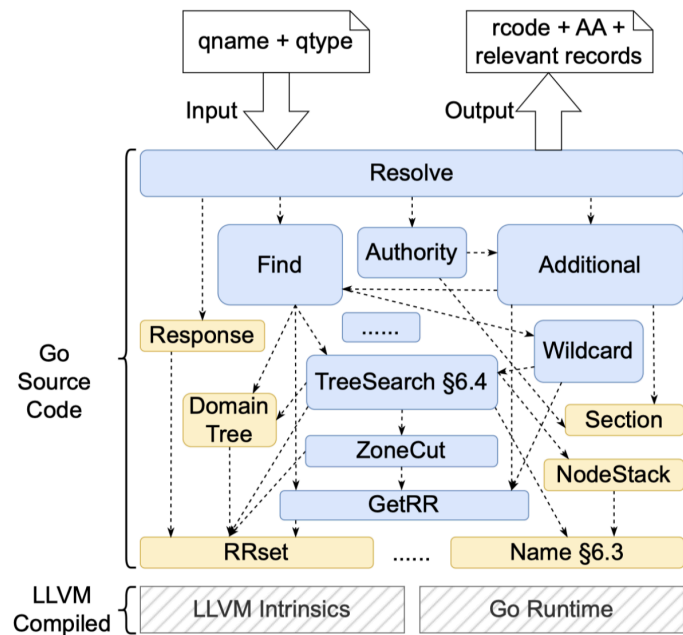
Specification



# Prior work partially automated this

DNS-V (SOSP'23) verified an in-production engine at Alibaba Cloud using **ahead-of-time (AOT) summarization**.

- Higher-level modules → automatically **summarize** via symbolic execution.
- Low-level modules (byte ops, data structures) → **still need manual specifications**.



Manual effort: still too high

# Manual effort: still too high

## Manual Effort Estimation (% of functions)

CoreDNS	Bind 9	PowerDNS
~30 / 79 (38%)	~200 / 314 (64%)	~50 / 114 (44%)

## Bind 9 – entry point

```
5364 /*%
5365  * Starting point for a client query or a chaining query.
5366  *
5367  * Called first by query_setup(), and then again as often as needed to
5368  * follow a CNAME chain. Determines which authoritative database to
5369  * search, then hands off processing to query_lookup().
5370 */
5371 isc_result_t
5372 ns__query_start(query_ctx_t *qctx) {
5373     isc_result_t result = ISC_R_UNSET;
5374     CCTRACE(ISC_LOG_DEBUG(3), "ns__query_start");
5375     ...
5376 }
```

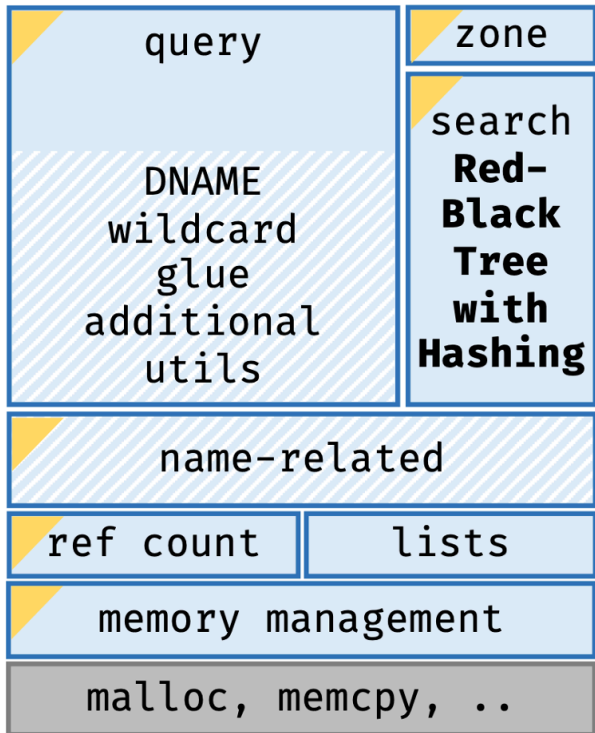
**Core Engine Logic**  
**O(10,000) LOC**

Name
..
include
.gitignore
client.c
hooks.c
interfacemgr.c
listenlist.c
meson.build
notify.c
probes.d
query.c
server.c
stats.c

# Manual effort: still too high

## Many Difficult Specs

$O(10,000)$  LOC



## Changes can break specs

Apr 02, 2026

query.c



Move zone set/get properties to own source file ...

Matthijs Mekking authored 1 month ago



remove deadcode in `query\_addbestns()` ...

Colin Vidal authored 3 weeks ago

Mar 31, 2026



Remove node and db pointer from dns\_rdataset\_t.vec ...

Alessio Podda authored 1 month ago

Mar 30, 2026



test for auth+res server and glues in delegation ...

Colin Vidal authored 1 month ago



remove find\_deepest\_zonecut() from qpcache ...

Evan Hunt authored 1 month ago and Colin Vidal committed 4 weeks ago



Guard against NULL delegset in query\_delegation\_recurse() ...

Ondřej Surý authored 1 month ago and Colin Vidal committed 4 weeks ago



Use delegdb for lookup in query\_delegation\_recurse() ...

Colin Vidal authored 2 months ago

Can we **guarantee** correctness with **low manual effort**?

Can we **guarantee** correctness with **low manual effort**?

Iceberg — automated verification via **just-in-time summarization**

# Insight: invariants from DNS zones

- An engine always operates on **concrete domain names and records**

```
example.com. 3600 IN SOA ns.example.com. a
                2025081301 ; serial
                3600      ; refresh
                900       ; retry
                1209600   ; expire
                60        ; minimum
)

example.com.      3600 IN NS ns.example.

example.com.      3600 IN A   192.0.2.1
example.com.      3600 IN AAAA 2001:db8:10
ns.example.com.   3600 IN A   192.0.2.2
ns.example.com.   3600 IN AAAA 2001:db8:10

www.example.com.  3600 IN CNAME example.com
```

# Insight: invariants from DNS zones

- An engine always operates on **concrete domain names and records**
- Zone data propagates through the code as **invariants**:
  - Fixed function arguments
  - Constant memory contents
  - Bounded data structures

```
example.com. 3600 IN SOA ns.example.com. a
                2025081301 ; serial
                3600      ; refresh
                900       ; retry
                1209600   ; expire
                60        ; minimum
)

example.com.      3600 IN NS ns.example.

example.com.      3600 IN A 192.0.2.1
example.com.      3600 IN AAAA 2001:db8:10
ns.example.com.   3600 IN A 192.0.2.2
ns.example.com.   3600 IN AAAA 2001:db8:10

www.example.com.  3600 IN CNAME example.com
```

# Insight: invariants from DNS zones

- An engine always operates on **concrete domain names and records**
- Zone data propagates through the code as **invariants**:
  - Fixed function arguments
  - Constant memory contents
  - Bounded data structures

These invariants can **dramatically simplify** what we need to summarize.

```
example.com. 3600 IN SOA ns.example.com. a
                2025081301 ; serial
                3600      ; refresh
                900      ; retry
                1209600   ; expire
                60       ; minimum
)

example.com.      3600 IN NS ns.example.

example.com.      3600 IN A   192.0.2.1
example.com.      3600 IN AAAA 2001:db8:10
ns.example.com.   3600 IN A   192.0.2.2
ns.example.com.   3600 IN AAAA 2001:db8:10

www.example.com.  3600 IN CNAME example.com
```

# Example: domain name comparison (Bind 9)

```
1 // RFC-compliant domain name format
2 typedef struct {
3     unsigned char    *ndata;    // binary data
4     unsigned int     labels;    // number of labels
5     unsigned char    *offsets;  // offset of labels
6 } name_t; /* other fields omitted */
7
8 namerefn_t name_fullcomp(name_t *n1, name_t *n2) {
9     /* definitions omitted */
10    nlabels = 0;
11    ofs1    = n1->offsets; // |
12    ofs2    = n2->offsets; // | get offset pointers
13    l       = min(n1->labels, n2->labels);
14    ldiff   = n1->labels - n2->labels;
```

# Example: domain name comparison (Bind 9)

```
1 // RFC-compliant domain name format
2 typedef struct {
3     unsigned char    *ndata;    // binary data
4     unsigned int     labels;    // number of labels
5     unsigned char    *offsets;  // offset of labels
6 } name_t; /* other fields omitted */
7
8 namerefn_t name_fullcomp(name_t *n1, name_t *n2) {
9     /* definitions omitted */
10    nlabels = 0;
11    ofs1    = n1->offsets; //|
12    ofs2    = n2->offsets; //| get offset pointers
13    l       = min(n1->labels, n2->labels);
14    ldiff   = n1->labels - n2->labels;
```

# Example: domain name comparison (Bind 9)

```
23     c2     = *p2++;      //| and move pointer
24     cdiff  = (int)c1 - (int)c2;      //|
25     c1     = cdiff < 0 ? c1: c2;    //|
26     diff   = ascii_lowercmp(p1, p2, c); //| compare
27     if (diff ≠ 0 || cdiff ≠ 0) goto done;
28     nlabels++;
29     }
30     if (ldiff < 0)      return namerefn_contains;
31     else if (ldiff > 0) return namerefn_subdomain;
32     else               return namerefn_equal;
33 done:
34     if (nlabels > 0)   return namerefn_commonancestor;
35     return namerefn_none;
36 }
```


- **Summarizing in general:** loop count symbolic, memory accesses depend on offset array, many interleaving paths.

# Example: domain name comparison (Bind 9)


```
20     p1     = &n1->ndata[*ofs1]; //| go to the
21     p2     = &n2->ndata[*ofs2]; //| next label
22     c1     = *p1++;           //| get label size
23     c2     = *p2++;           //| and move pointer
24     cdiff  = (int)c1 - (int)c2; //|
25     c1     = cdiff < 0 ? c1: c2; //|
26     diff   = ascii_lowercmp(p1, p2, c); //| compare
27     if (diff != 0 || cdiff != 0) goto done;
28     nlabels++;
29 }
30 if (ldiff < 0)     return namerefn_contains;
31 else if (ldiff > 0) return namerefn_subdomain;
32 else              return namerefn_equal;
33 done:
34 if (nlabels > 0)  return namerefn_comparison_error;
```

- **Summarizing in general:** loop count symbolic, memory accesses depend on offset array, many interleaving paths.
- **With zone invariant:** (n1 = "example.com") loop bounded, memory accesses resolve to constants, paths collapse.

# AOT vs. JIT summarization





 Automatically Checked

 Manual Specification

 Buggy Execution

 Correct Execution



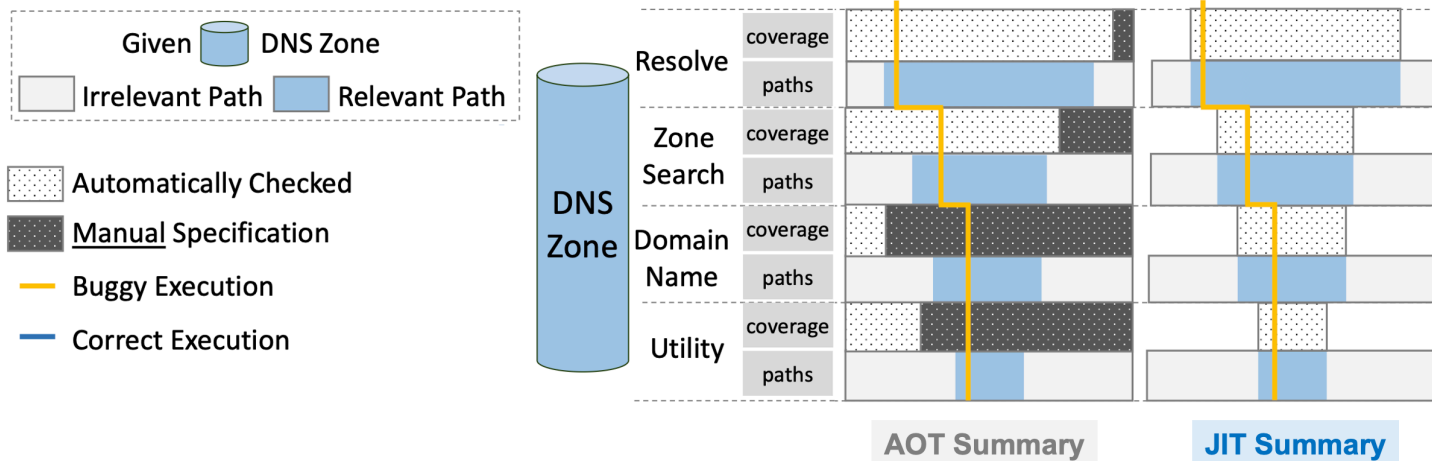
AOT Summary

JIT Summary

## Ahead-of-time (AOT) Summary

- Accounts for **all** inputs
- Many paths unreachable due to invariants
- Summaries unmanageable → manual specs

# AOT vs. JIT summarization



## Ahead-of-time (AOT) Summary

- Accounts for **all** inputs
- Many paths unreachable due to invariants
- Summaries unmanageable → manual specs

## Just-in-time (JIT) Summary

- Summarizes **in context** of actual call
- Zone data prunes impossible paths
- Summaries stay small → **more automated**

# JIT summaries: building top-down

- **Symbolic execution** on LLVM IR, top-down
- At each function call → build summary on the fly with zone invariants

# JIT summaries: building top-down

- Symbolic execution on LLVM IR, top-down
- At each function call → build summary on the fly with zone invariants

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    ↪ name_fullcomp(query, "example.com")
```

```
1  namerefn_t name_fullcomp(name_t *n1, name_t *n2) {  
2  /* definitions omitted */  
3  nlabels = 0;  
4  ofs1    = n1→offsets;  //|  
5  ofs2    = n2→offsets;  //| get offset pointers  
6  l       = min(n1→labels, n2→labels);  
7  ldiff   = n1→labels - n2→labels;  
8  ofs1    += n1→labels;  //|  
9  ofs2    += n2→labels;  //| start from the right  
10 while (l-- > 0) {  
11     ofs1--:                                     //|
```

# JIT summaries: building top-down

- Symbolic execution on LLVM IR, top-down
- At each function call → build summary on the fly with zone invariants

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    ↘ name_fullcomp(query, "example.com")
```

When the code dereferences a pointer into zone data

→ **state splitting**: substitute symbolic (AOT) with concrete value (JIT)

```
5 ofs2 = n2->offsets; //| get offset pointers  
6 l = min(n1->labels, n2->labels);  
7 ldiff = n1->labels - n2->labels;  
8 ofs1 += n1->labels; //|  
9 ofs2 += n2->labels; //| start from the right  
10 while (l-- > 0) {  
11 ofs1--; //|  
12 ofs2--; //|  
13 p1 = &n1->ndata[*ofs1]; //| go to the  
14 p2 = &n2->ndata[*ofs2]; //| next label  
15 c1 = *p1++; //| get label size  
16 c2 = *p2++; //| and move pointer
```

# JIT summaries: building top-down

- Symbolic execution on LLVM IR, top-down
- At each function call → build summary on the fly with zone invariants

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    ↪ name_fullcomp(query, "example.com")
```

When the code dereferences a pointer into zone data

→ **state splitting**: substitute symbolic (AOT) with concrete value (JIT)

```
13     p1 = &n1→ndata[*ofs1]; //| go to the  
14     p2 = &n2→ndata[*ofs2]; //| next label  
15     c1 = *p1++; //| get label size  
16     c2 = *p2++; //| and move pointer  
17     cdiff = (int)c1 - (int)c2; //|  
18     c1 = cdiff < 0 ? c1: c2; //|  
19     diff = ascii_lowercmp(p1, p2, c); //| compare  
20     if (diff ≠ 0 || cdiff ≠ 0) goto done;  
21     nlabels++;  
22 }  
23 if (ldiff < 0) return namerln_contains;  
24 else if (ldiff > 0) return namerln_subdomain;
```

# JIT summaries: building top-down

- Symbolic execution on LLVM IR, top-down
- At each function call → build summary on the fly with zone invariants

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    ↘ name_fullcomp(query, "example.com")
```

When the code dereferences a pointer into zone data

→ **state splitting**: substitute symbolic (AOT) with concrete value (JIT)

```
13     p1 = &n1→ndata[*ofs1]; //| go to the  
14     p2 = &n2→ndata[*ofs2]; //| next label  
15     c1 = *p1++; //| get label size  
16     c2 = *p2++; //| and move pointer  
17     cdiff = (int)c1 - (int)c2; //|  
18     c1 = cdiff < 0 ? c1: c2; //|  
19     diff = ascii_lowercmp(p1, p2, c); //| compare  
20     if (diff ≠ 0 || cdiff ≠ 0) goto done;  
21     nlabels++;  
22 }  
23 if (ldiff < 0) return namereln_contains;  
24 else if (ldiff > 0) return namereln_subdomain;
```

**JIT state splitting**: Replace symbolic values with concrete values, triggered by a set of rules (*see our paper*).

# JIT summaries: applying and re-summation

- Symbolic execution on LLVM IR, top-down

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    → name_fullcomp(query, "example.com")    build  
    → name_fullcomp(query, "example.com")    apply  
    → name_fullcomp(query, "cs.example.com") re-build
```

```
name_compare() # Summary tree (pseudo notation)  
  ↳ [Root]  
    ↳ [Specialize | n1 is "example.com" ]  
      ↳ [Assume | query = "example.com" ]  
        ↳ [Ret | return EQUAL ]  
      ↳ [Assume | query ≠ "example.com" ]  
        ↳ [Ret | return PARTIAL / NONE ]
```

# JIT summaries: applying and re-summarization

- Symbolic execution on LLVM IR, top-down
- First call with "example.com" → build summary

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    → name_fullcomp(query, "example.com")    build  
    → name_fullcomp(query, "example.com")    apply  
    → name_fullcomp(query, "cs.example.com") re-build
```

```
name_compare() # Summary tree (pseudo notation)  
  ↳ [Root]  
    ↳ [Specialize | n1 is "example.com" ]  
      ↳ [Assume | query = "example.com" ]  
        ↳ [Ret | return EQUAL ]  
      ↳ [Assume | query ≠ "example.com" ]  
        ↳ [Ret | return PARTIAL / NONE ]
```

# JIT summaries: applying and re-summarization

- Symbolic execution on LLVM IR, top-down
- First call with "example.com" → build summary
- Second call with "example.com" again → **apply** existing summary
- Second call with "cs.example.com" → **re-build** existing summary

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    → name_fullcomp(query, "example.com")    build  
    → name_fullcomp(query, "example.com")    apply  
    → name_fullcomp(query, "cs.example.com") re-build  
}
```

```
name_compare() # Summary tree (pseudo notation)  
  ↳[Root]  
    ↳[Specialize| n1 is "example.com" ]  
      ↳[Assume| query = "example.com" ]  
        ↳[Ret| return EQUAL ]  
      ↳[Assume| query ≠ "example.com" ]  
        ↳[Ret| return PARTIAL / NONE ]
```

# JIT summaries: applying and re-summarization

- Symbolic execution on LLVM IR, top-down
- First call with "example.com" → build summary
- Second call with "example.com" again → **apply** existing summary
- Next call with "cs.example.com" → not yet covered → **re-summarize and merge**

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    → name_fullcomp(query, "example.com")    build  
    → name_fullcomp(query, "example.com")    apply  
    → name_fullcomp(query, "cs.example.com") re-build  
}
```

```
name_compare() # Summary tree (pseudo notation)  
└─[Root]  
  ├─[Specialize| n1 is "example.com" ]  
  │ └─[Assume| query = "example.com" ]  
  │   └─[Ret| return EQUAL ]  
  │ └─[Assume| query ≠ "example.com" ]  
  │   └─[Ret| return PARTIAL / NONE ]  
  └─[Specialize| n1 is "cs.example.com" ]  
    ├─[Assume| query = "cs.example.com" ]  
    │ └─[Ret| return EQUAL ]  
    └─[Assume| query ≠ "cs.example.com" ]  
      └─[Ret| return PARTIAL / NONE ]
```

# JIT summaries: applying and re-summarization

- Symbolic execution on LLVM IR, top-down
- First call with "example.com" → build summary
- Second call with "example.com" again → **apply** existing summary
- Next call with "cs.example.com" → not yet covered → **re-summarize and merge**

```
5372 ns_query_start(query_ctx_t *qctx) {  
    ...  
    → name_fullcomp(query, "example.com")    build  
    → name_fullcomp(query, "example.com")    apply  
    → name_fullcomp(query, "cs.example.com") re-build
```

- Summary grows just in time to cover necessary behaviors

```
name_compare() # Summary tree (pseudo notation)  
└─[Root]  
  ├─[Specialize| n1 is "example.com" ]  
  │ └─[Assume| query = "example.com" ]  
  │   └─[Ret| return EQUAL ]  
  │ └─[Assume| query ≠ "example.com" ]  
  │   └─[Ret| return PARTIAL / NONE ]  
  └─[Specialize| n1 is "cs.example.com" ]  
    └─[Assume| query = "cs.example.com" ]  
      └─[Ret| return EQUAL ]  
    └─[Assume| query ≠ "cs.example.com" ]  
      └─[Ret| return PARTIAL / NONE ]
```

# Formal model: check out our paper!

JIT summaries are **trees of constraints and effects**,  
managed by two core algorithms:

# Formal model: check out our paper!

JIT summaries are **trees of constraints and effects**, managed by two core algorithms:

- **BuildSummary** — top-down symbolic execution with JIT state splitting
- **ApplySummary** — match existing summary or trigger re-summarization

## Algorithm 2 APPLYSUMMARY

### Algorithm 1 BUILDSUMMARY

**Input:** Function  $f$ , state model  $\mathcal{S}$

**Output:** Summary  $f^{\text{SUM}}$

- 1:  $G \leftarrow$  control flow graph (CFG) of  $f$
- 2:  $C \leftarrow \emptyset$  # Empty constraint set
- 3: Initialize  $f^{\text{SUM}}$  with Root # Empty summary
- 4: Traverse  $G$  in depth-first search (DFS) order

- 8: ▷ **Entering a Specialize node with constraint  $s$ :**
- 9:   Grow Specialize on  $f^{\text{SUM}}$
- 10:   **if**  $s$  is satisfiable given  $C$  **and**  $s$  matches  $\mathcal{S}$  **then**
- 11:      $C \leftarrow C \cup \{s\}$

# Formal model: check out our paper!

JIT summaries are **trees of constraints and effects**, managed by two core algorithms:

- **BuildSummary** — top-down symbolic execution with JIT state splitting
- **ApplySummary** — match existing summary or trigger re-summarization
- Operates on a **unified state/memory model** across all proof layers.

	Effect	Definition	Example
State	Root	Root placeholder	[Root]
	Assume	Path constraint	[Assume  x > 0]
	Specialize	JIT state splitting	[Specialize  x is 1]
	Panic	Runtime error	[Panic  "..."]
ValueState	Return	Binding a return value	[Return  2w64]
	Store	Memory store	null -> Region::0
	Malloc	Heap allocation	[0] := malloc ptr
	Free	Heap deallocation	free [0]
MemoryState	::=	Region :: int	Region with offsets
Region	::=	@func_ident :: %ident	Stack region
		Global :: %ident	Global region
		Const :: %ident	Constant region
		Heap :: %ident	Heap region

Further scale verification

# Further scale verification

```
[Specialize | @name_fullcomp:%n1 is  $\delta$ @f::%name::o ]  
  |-[Assume | @f::%name::o == .. ]--[ .. ]  
  |-[Assume | @f::%name::o != .. ]--[ .. ]
```

Concrete  
Region

Symbolic  
Region

```
[Specialize | @name_fullcomp:%n1 is  $\delta$ [1]::o ]  
  |-[Assume | [1]::o == .. ]--[ .. ]  
  |-[Assume | [1]::o != .. ]--[ .. ]
```

Symbolic regions — reuse summaries across concrete memory addresses.

# Further scale verification

```
[Specialize | @name_fullcomp::%n1 is &[1]::@ ]  
  | [Assume | [1]::@ == .. ]--[ .. ]  
  | [Assume | [1]::@ != .. ]--[ .. ]
```

## Summary Optimization Passes

- MinimizeAssume
- MergeAssume
- ...

**Symbolic regions** — reuse summaries across concrete memory addresses.

**Summary optimization** — compiler-like passes that optimize summaries.

# Further scale verification

```
1 //unique_ptr<DNSPacket> PacketHandler::doQuestion(DNSPacket& p)
2 fn stub(ret: Ptr, h: Ptr, p: Ptr) -> Bool {
3   let qtype = p |> DnsPacket.qtype |> *Int<16>; // read qtype
4   let _ = debug(qtype); // log qtype
5   specialize(h |> PacketHandler.log_dns_details |> *Int<8>)
6   && qtype != 0w16 && qtype < 249w16 // qtype not EMPTY | META
7 }
```

**Stub Function**


(PowerDNS)

**Symbolic regions** — reuse summaries across concrete memory addresses.

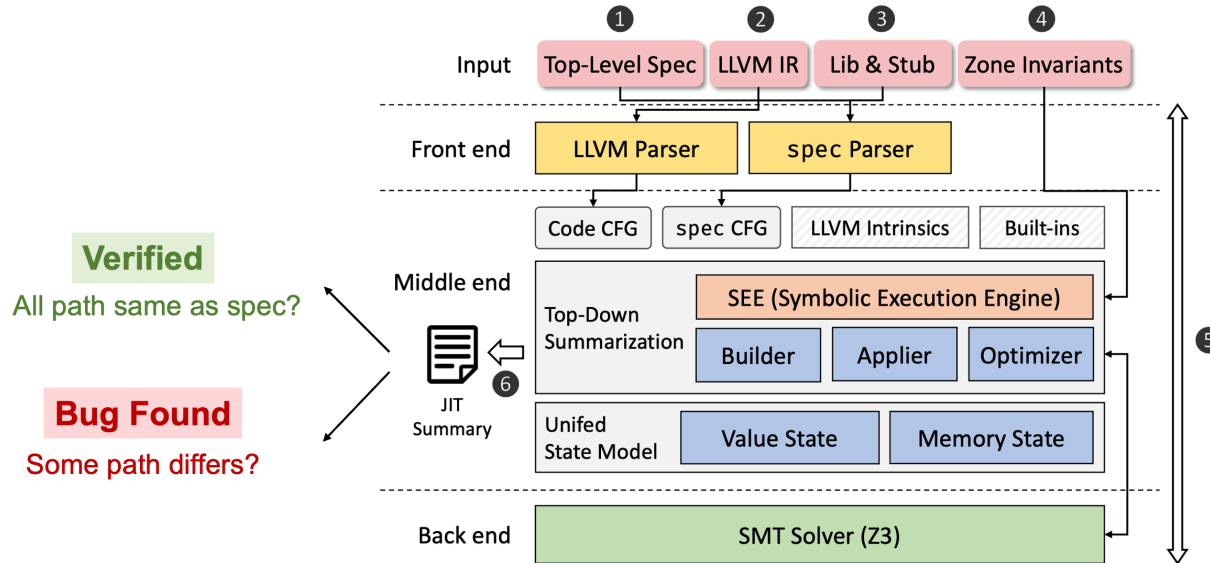
**Summary optimization** — compiler-like passes that optimize summaries.

**Stub functions** — lightweight interposition to customize verification.

# Iceberg overview

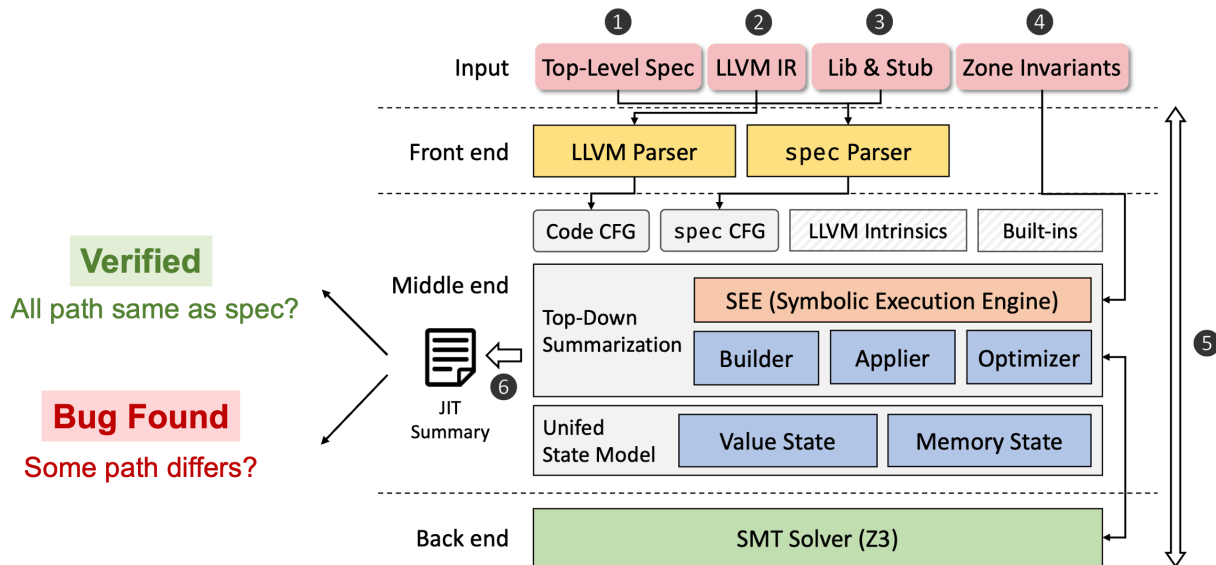


# Iceberg overview



Push-button verification — no per-layer specifications needed.

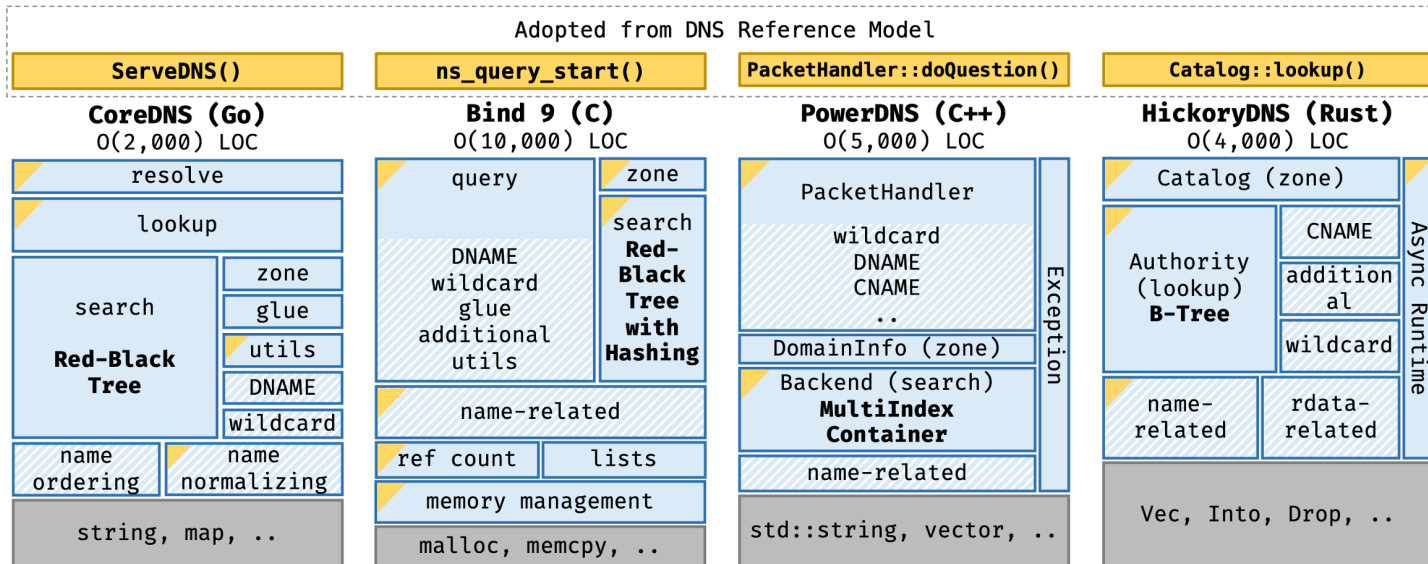
# Iceberg overview



**Push-button verification** — no per-layer specifications needed.

**Implementation:** 16,000 Rust LOC, based on the Z3 SMT solver

# Practical verification



JIT-summarized

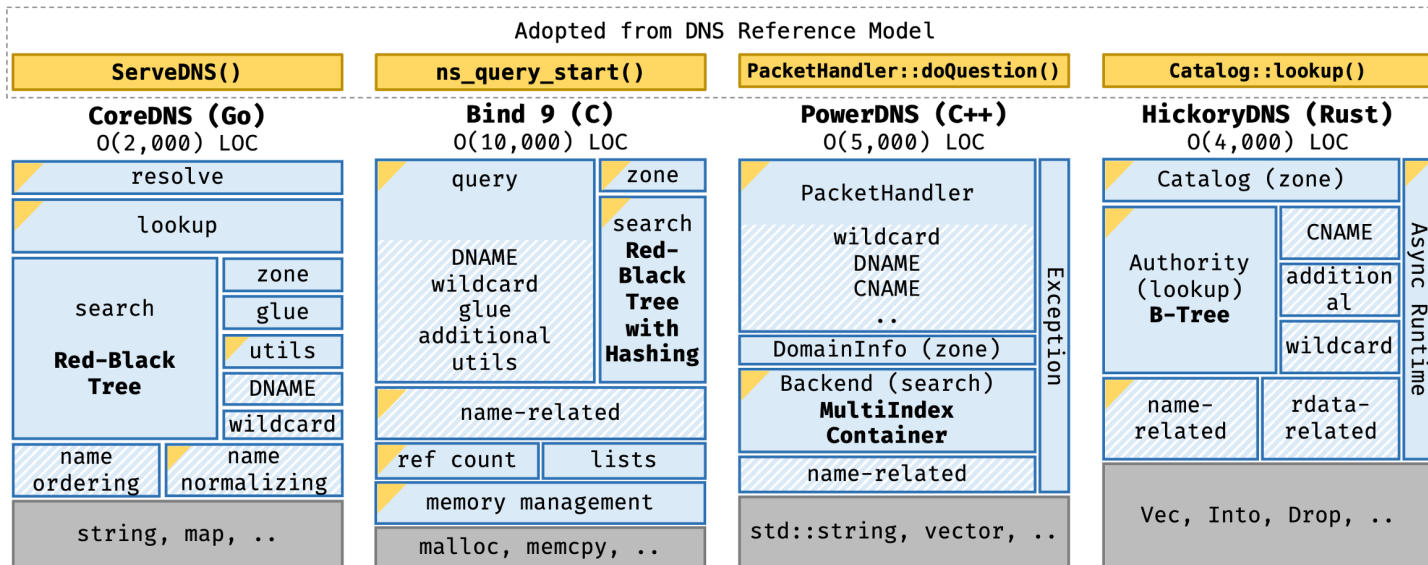


Top-level Spec



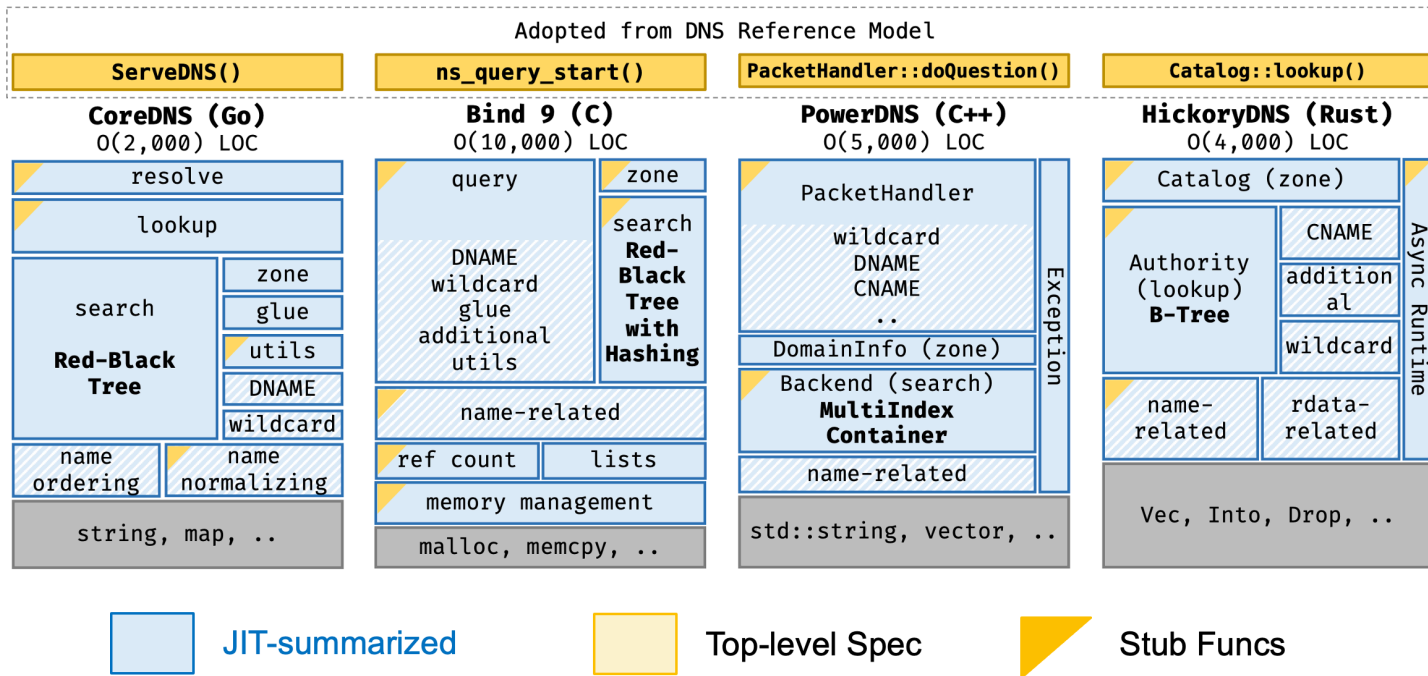
Stub Funcs

# Practical verification



Scope: core lookup logic (query → zone → response). Excludes I/O, networking, concurrency.

# Practical verification

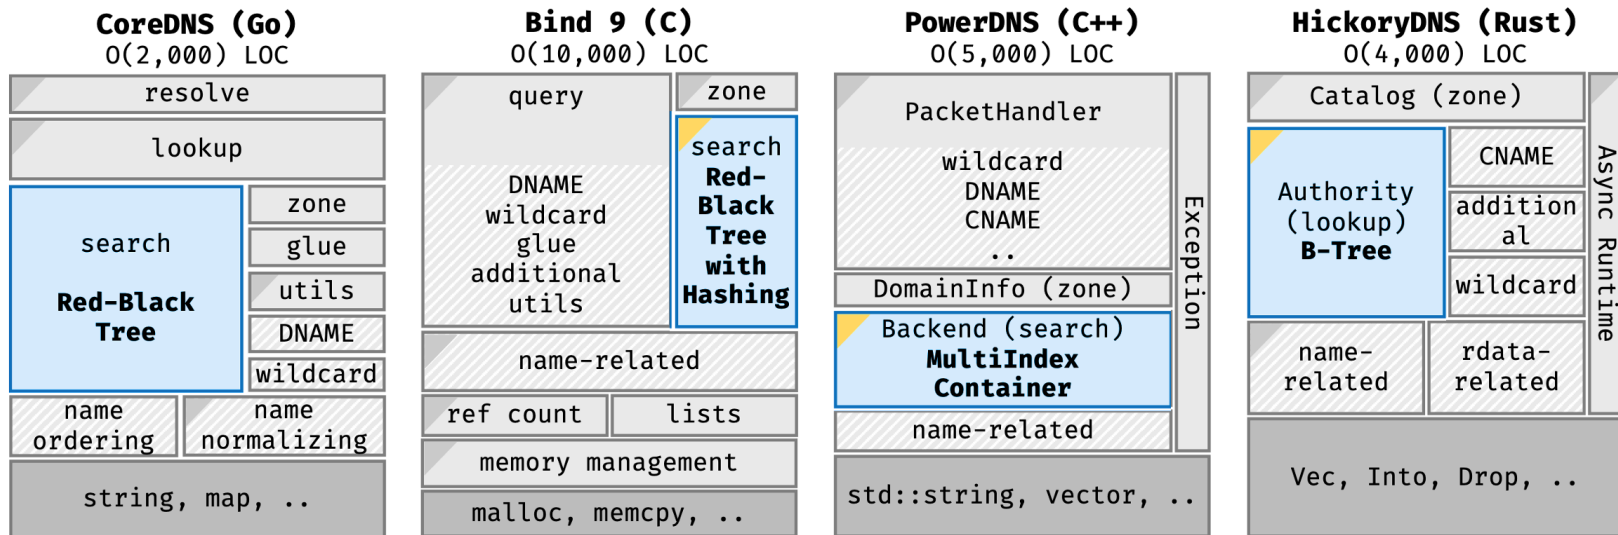


**Scope:** core lookup logic (query → zone → response). Excludes I/O, networking, concurrency.

**Top-level spec:** DNS reference model derived from 7+ RFCs, tailored per engine.

# Verification experience

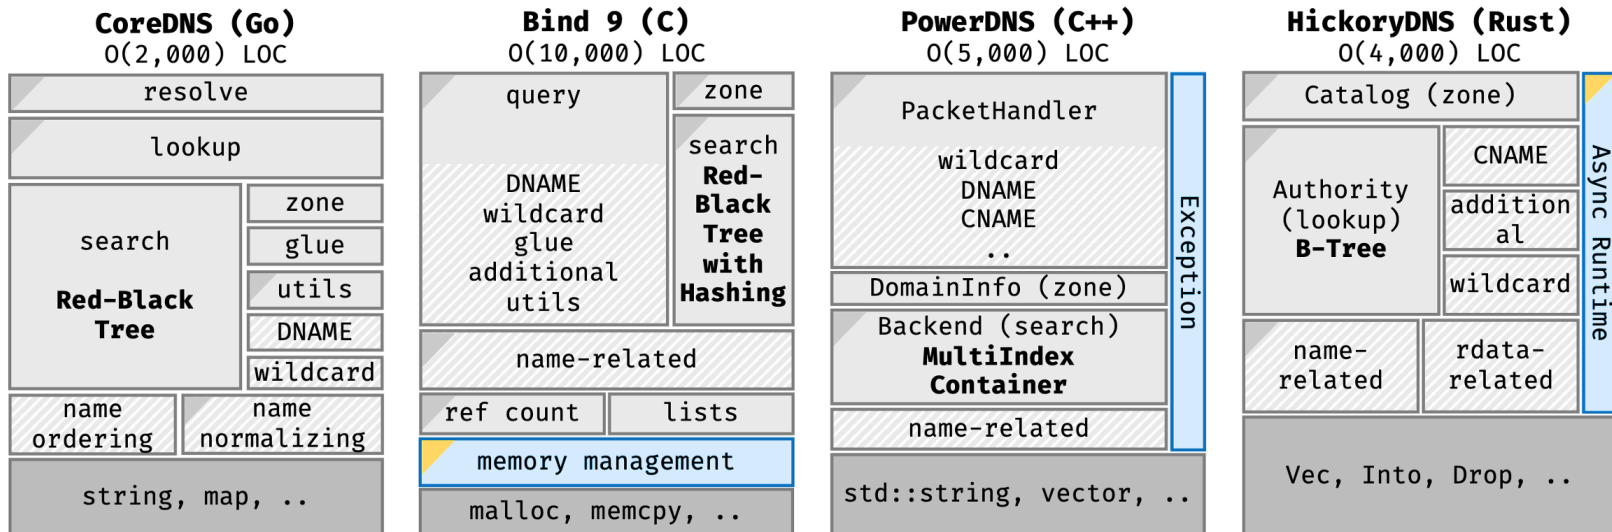
# Verification experience



**Diverse data structures** — red-black trees, B-trees, hash tables, multi-index containers.

JIT doesn't model the data structure — it state-splits on the *concrete zone layout*. The same mechanism handles all four, with no per-data-structure specs.

# Verification experience



Language-specific IR patterns — C++ exceptions, Rust async state machines.

Complex compilation artifacts in LLVM IR. JIT explores only *reachable* states top-down — no manual modeling.

# Evaluation setup

**Zones:** 3 categories for thorough coverage

# Evaluation setup

Zones: 3 categories for thorough coverage

Category	Size	Count	Source
simple	2–4 records	1,000	SCALE
complex	10–12 records	100	randomly generated
real	137 & 393 records	2	public CZDS data

# Evaluation setup

**Zones:** 3 categories for thorough coverage

Category	Size	Count	Source
simple	2–4 records	1,000	SCALE
complex	10–12 records	100	randomly generated
real	137 & 393 records	2	public CZDS data

**Baseline:** SCALE (NSDI'22) — state-of-the-art symbolic test generator for DNS engines

# Evaluation setup

**Zones:** 3 categories for thorough coverage

Category	Size	Count	Source
simple	2–4 records	1,000	SCALE
complex	10–12 records	100	randomly generated
real	137 & 393 records	2	public CZDS data

**Baseline:** SCALE (NSDI'22) — state-of-the-art symbolic test generator for DNS engines

**Metrics:** bugs found / manual effort / verification time

# Bugs Found

Index	Impl	Category	Description	Status	SCALE
1	CoreDNS	Wrong Answer & Flag	Incorrect response when DNAME redirects to a zonecut.	✓	✗
2	CoreDNS	Wrong Answer	Missing SOA/NS records when CNAME redirects to a zone apex. <sup>†</sup>	✓	✗
3	CoreDNS	Wrong Additional	Missing records for wildcard MXs.	✓	✗
4	CoreDNS	Wrong Additional	Extraneous glue records for MX queries.	✓	✗
5	CoreDNS	Wrong Additional	Duplicate records during the additional section handling.	✓	✗
6	CoreDNS	Wrong rcode	Erroneous return code for CNAMEs pointing out of zone. <sup>†</sup>	✓	✗
7	CoreDNS	Wrong Additional	Extraneous glue records for non-referral NS queries.	✓	✓
8	CoreDNS	Wrong Additional	Missing additional records after DNAME substitution.	✓	✗
9	PowerDNS	Wrong Answer	Extraneous records when matching wildcard CNAMEs.	✓	✓
10	PowerDNS	Wrong Authority	Extraneous SOA records after DNAME redirection.	✓	✗
11	HickoryDNS	Wrong Answer & Flag	Incorrect response for a non-existent SOA query.	✓	✗
12	HickoryDNS	Wrong Additional	Wrong record placement for CNAME-related queries.	✓	✗

- SCALE found only 2 of 12 on the same zones.

# Bugs Found

Non-existent CNAME target in the same zone should be returned with NXDOMAIN instead of NOERROR rcode #4288

SCALE  
Bug Report (2020)

Closed #4303

SivaKesava1 opened on Nov 13, 2020

Hi,

This is related to if 'A CNAME B' exists in a zone, the response should be NXDOMAIN. [RFC 6604](#) mentions that when chains are followed the RCODE in the response should be NXDOMAIN.

## [BUG 6] Erroneous return code for CNAMEs pointing out of zone

### What happened:

When the `file` plugin chases a CNAME chain that points out-of-zone (and thus non-existent), the RCODE is `NOERROR`.

### How to reproduce it (as minimally and precisely as possible):

Using the following zone file (`a.txt`):

```
a.                500 IN SOA      b.c.d. e.a.d. 3 604800 86400 2419200 604800
a.                500 IN NS       b.c.d.
b.a.              500 IN CNAME    c.
```

Response for querying `<b.a., A>` (simplified):

```
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 22063
;; ANSWER SECTION:
b.a.                500    IN      CNAME   c.
```

Iceberg Bug Report  
(2024)

Different trigger, same bug  
Incomplete fix!

- Some bugs were in *incomplete fixes* for issues SCALE had previously reported.

# Manual effort & overhead

	top	lib	stub	spec-to-code ratio
CoreDNS	790	233	192	9.6%
Bind 9	860	108	752	7.5%
PowerDNS	752	849	290	5.8%
HickoryDNS	731	249	314	7.8%

## Manual Effort


*spec-to-code ratio measured w/ top specs*

## Manual effort

Spec-to-code ratio: **5.8%–9.6%**

Upfront porting: **~1 person-week**

Version retrofitting: **~1 person-hour**



## Verification Overhead

## Verification time

Simple zones: **minutes**

Complex zones: **~30 min**

Production zones: **few hours**

All verification runs terminate on a standard cloud server (Alibaba Cloud ecs.c9i.4xlarge).

# Beyond DNS

Many systems operate on **pre-configured dataplane:**

# Beyond DNS

Many systems operate on **pre-configured dataplane**:

Network switches → forwarding tables

SDN controllers → flow rules

Database query engines → schemas and indexes

# Beyond DNS

Many systems operate on **pre-configured dataplane**:

Network switches → forwarding tables

SDN controllers → flow rules

Database query engines → schemas and indexes

Whenever behavior is shaped by fixed configuration,  
JIT summarization can exploit those invariants.

# Beyond DNS

Many systems operate on **pre-configured dataplane**:

Network switches → forwarding tables

SDN controllers → flow rules

Database query engines → schemas and indexes

Whenever behavior is shaped by fixed configuration,  
JIT summarization can exploit those invariants.



# Summary

- **Iceberg** uses JIT summarization to automatically verify DNS engines.
- **4 engines** verified, **12 bugs** found, **1 person-week** porting cost.



# Summary

- Iceberg uses JIT summarization to automatically verify DNS engines.
- 4 engines verified, 12 bugs found, 1 person-week porting cost.

Thanks for listening!