

“So, you think your Nameservers are Correct?”

Finding Errors Automatically in Nameserver Implementations

Siva Kesava Reddy Kakarla¹


Ryan Beckett²

Todd Millstein^{1,3}

George Varghese¹

¹  University of California, Los Angeles

²  Microsoft

³  Intentionet

Implementing Nameservers Correctly is Hard!

Implementing Nameservers Correctly is Hard!

Implementing Nameservers Correctly is Hard!

RFCs

1034, 4592, 6672, ...




Implementing Nameservers Correctly is Hard!

RFCs
1034, 4592, 6672, ...



DNS Developers




Implementing Nameservers Correctly is Hard!

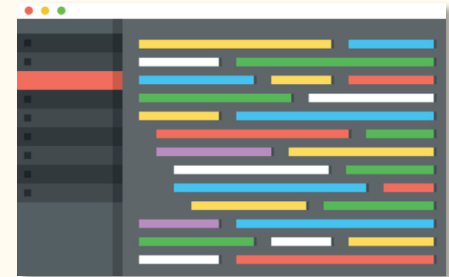
RFCs
1034, 4592, 6672, ...



DNS Developers



Implementations in
C / C++ / Go / ...




Implementing Nameservers Correctly is Hard!

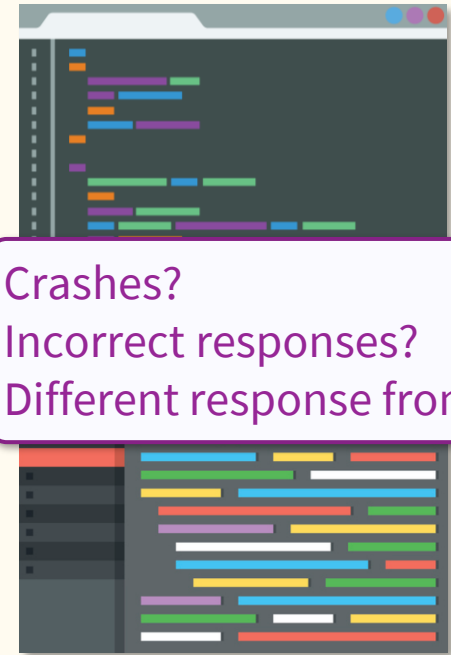
RFCs
1034, 4592, 6672, ...



DNS Developers



Implementations in
C / C++ / Go / ...



Crashes?
Incorrect responses?
Different response from others?


Implementing Nameservers Correctly is Hard!

RFCs
1034, 4592, 6672, ...

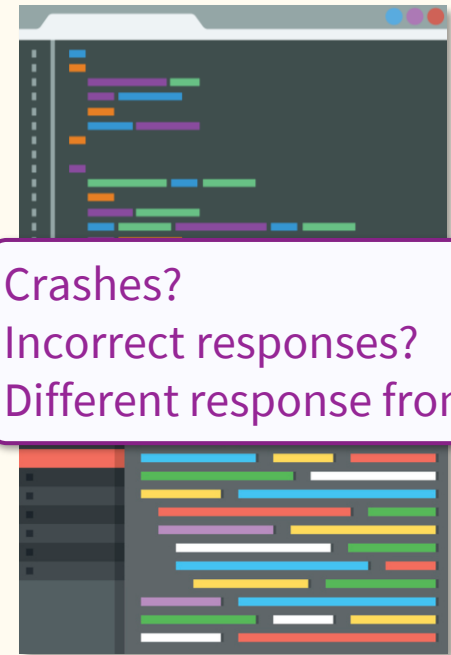


Compliance?

DNS Developers

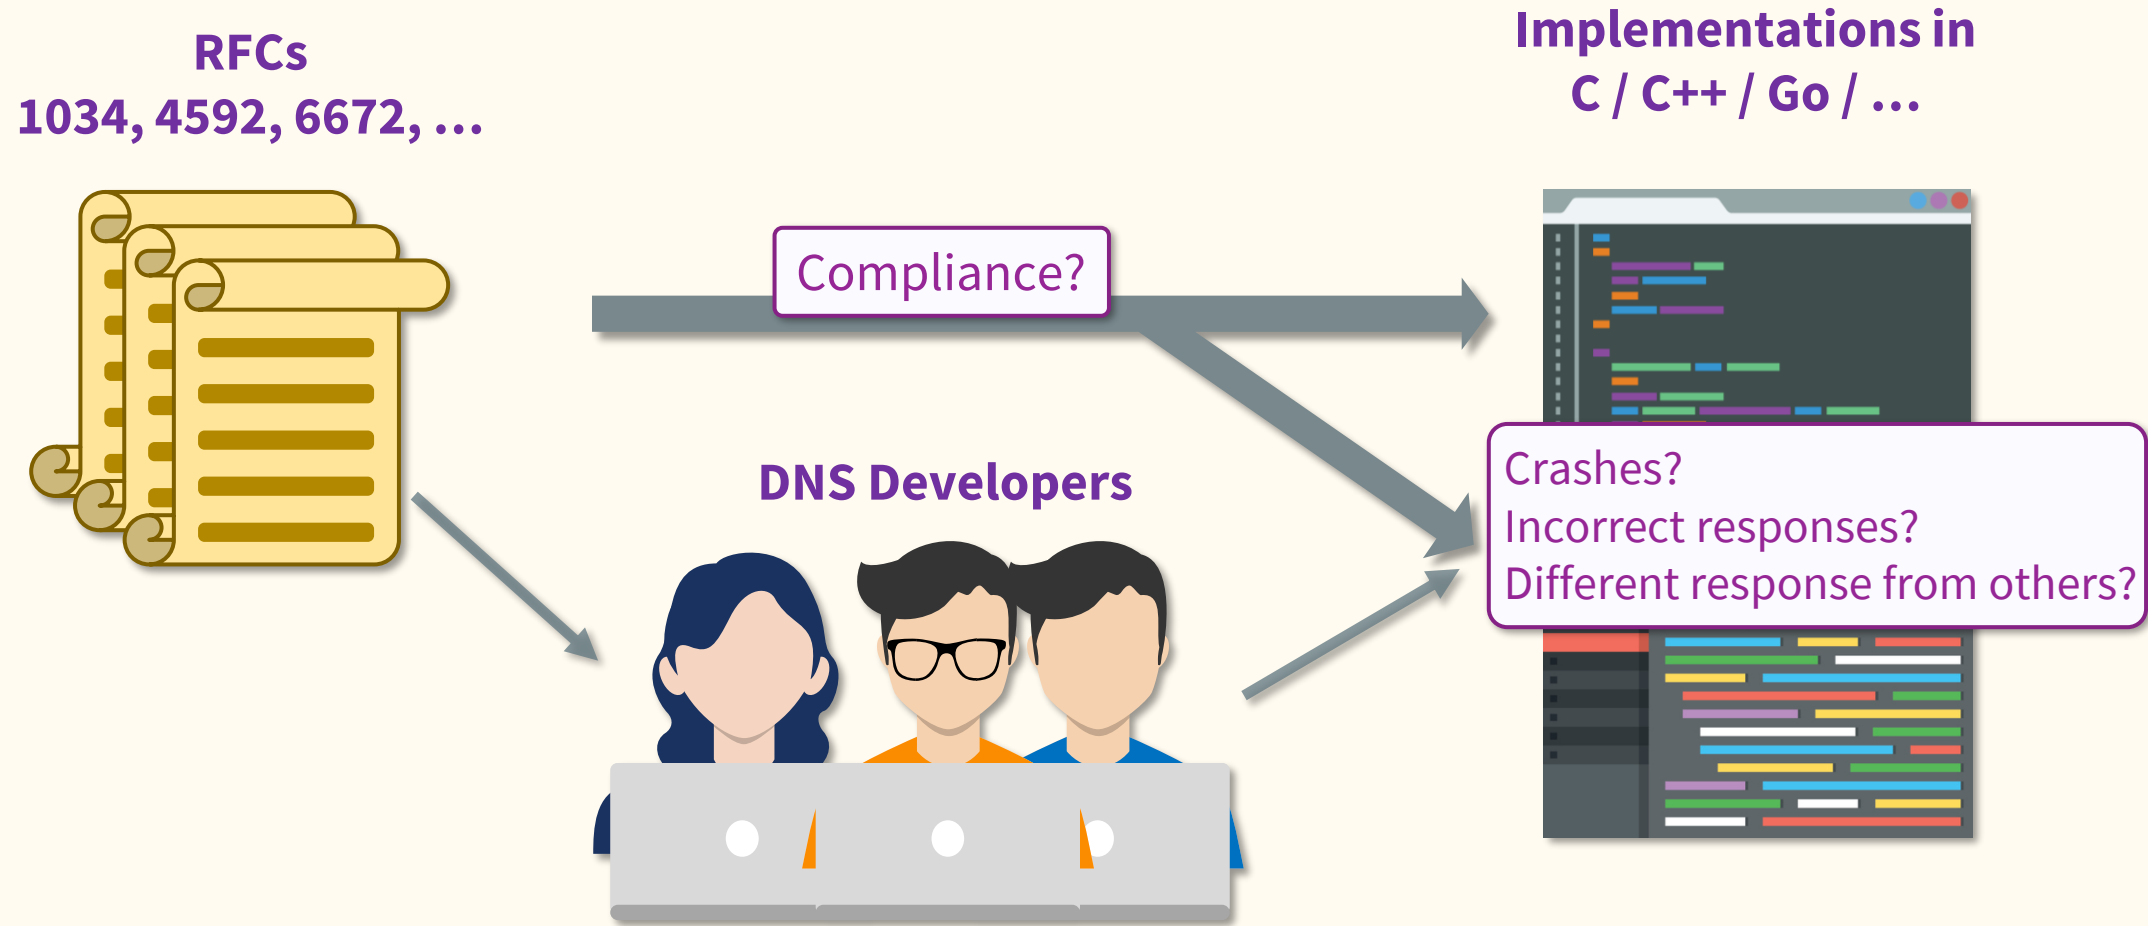


Implementations in
C / C++ / Go / ...

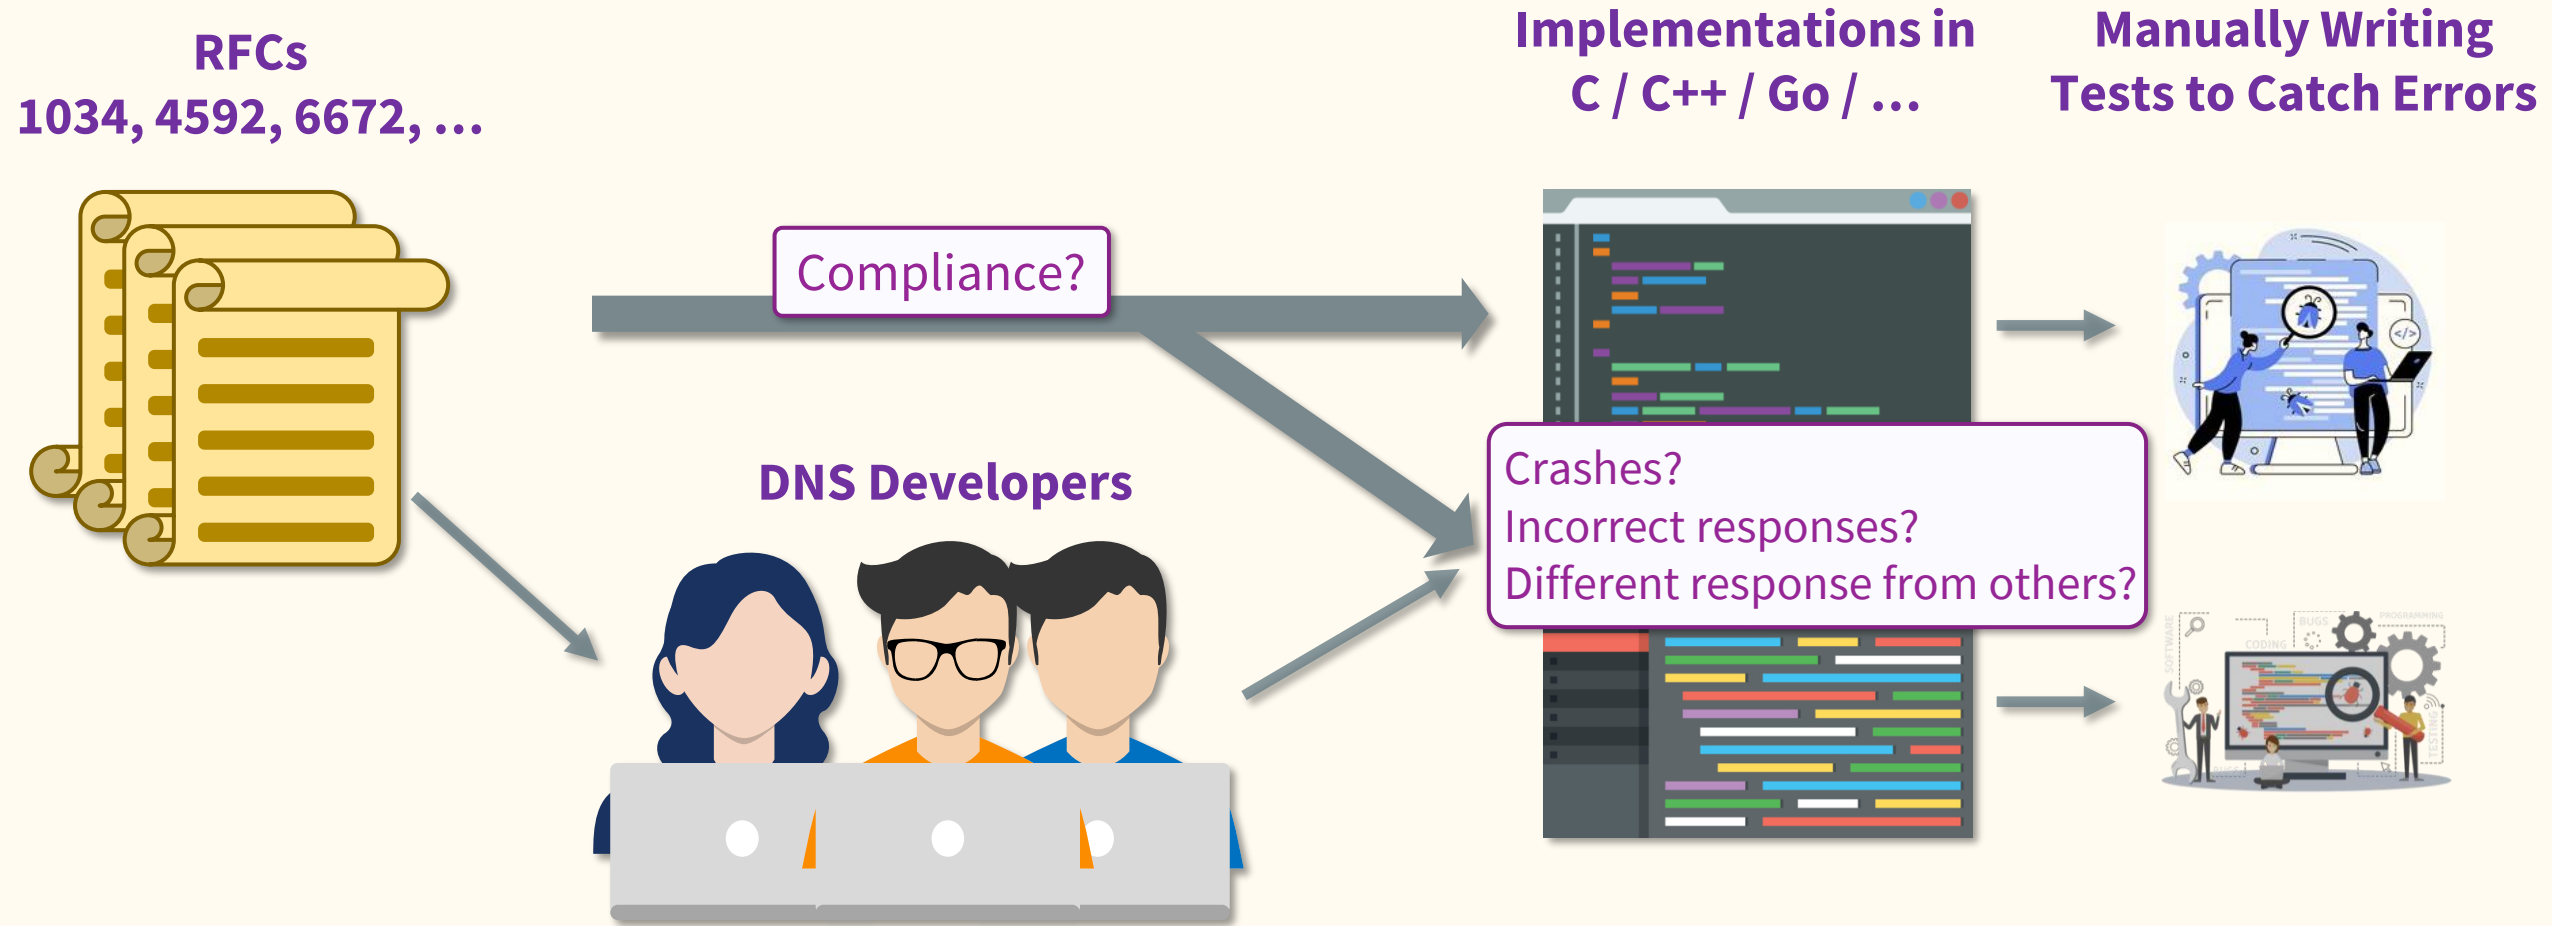


Crashes?
Incorrect responses?
Different response from others?


Current Practice: Ad Hoc Manual Testing



Current Practice: Ad Hoc Manual Testing



Current Practice: Ad Hoc Manual Testing



Current Practice: Ad Hoc Manual Testing



Current Practice: Ad Hoc Manual Testing



Our Idea:

FERRET - Generate tests
automatically and compare
across implementations



Our Idea:

FERRET - Generate tests
automatically and compare
across implementations

How to generate high-
coverage tests that identify
functional correctness errors?



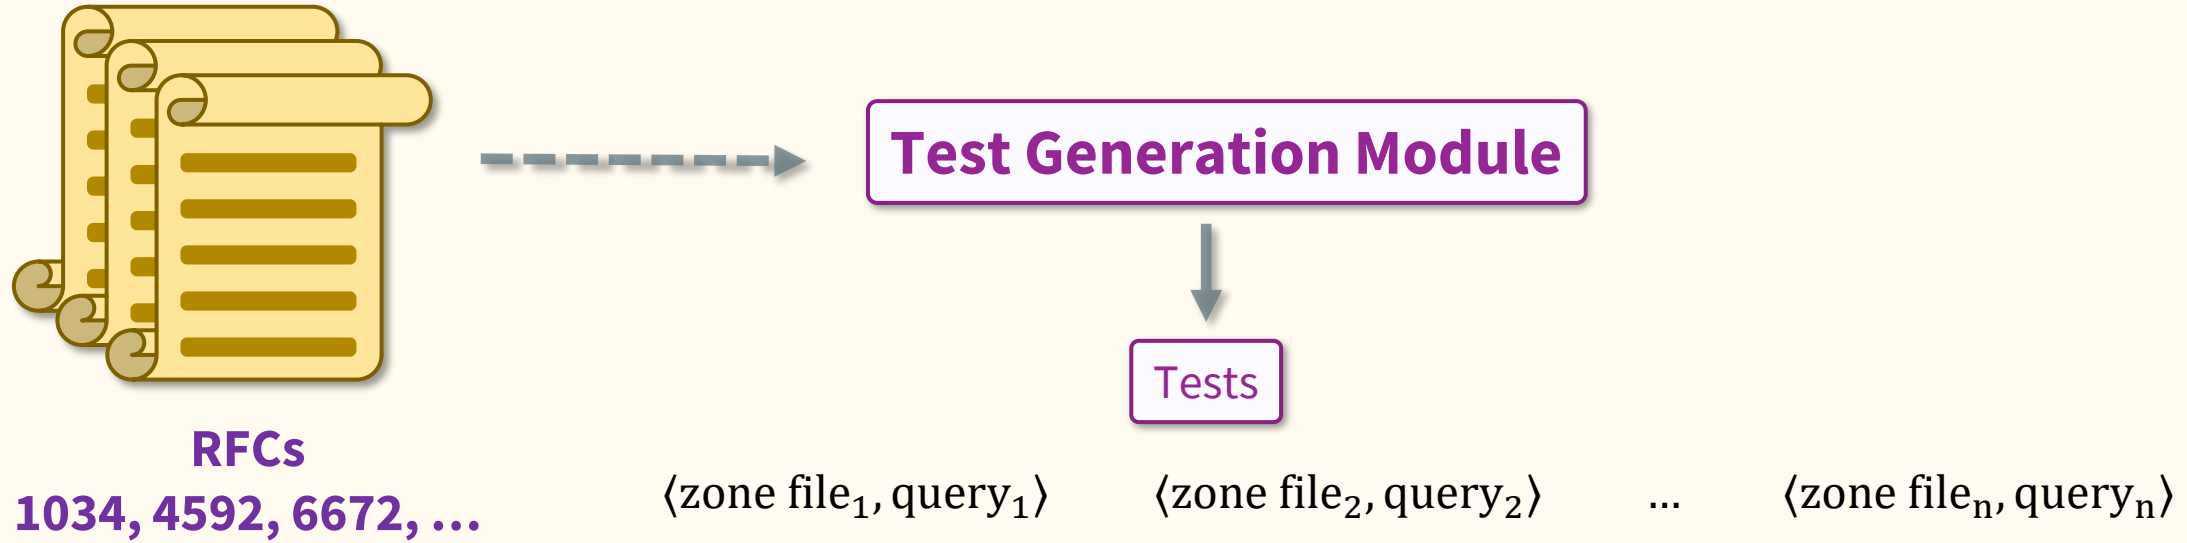
FERRET: End-to-End Design



RFCs

1034, 4592, 6672, ...

FERRET: End-to-End Design



FERRET: End-to-End Design



RFCs

1034, 4592, 6672, ...



Test Generation Module



Tests


$\langle \text{zone file}_1, \text{query}_1 \rangle$

$\langle \text{zone file}_2, \text{query}_2 \rangle$

Domain Name	Type	Data
campus.edu.	SOA	ns1.exm. ...
foo.campus.edu.	NS	ns1.campus.edu
ns1.campus.edu.	A	1.1.1.1

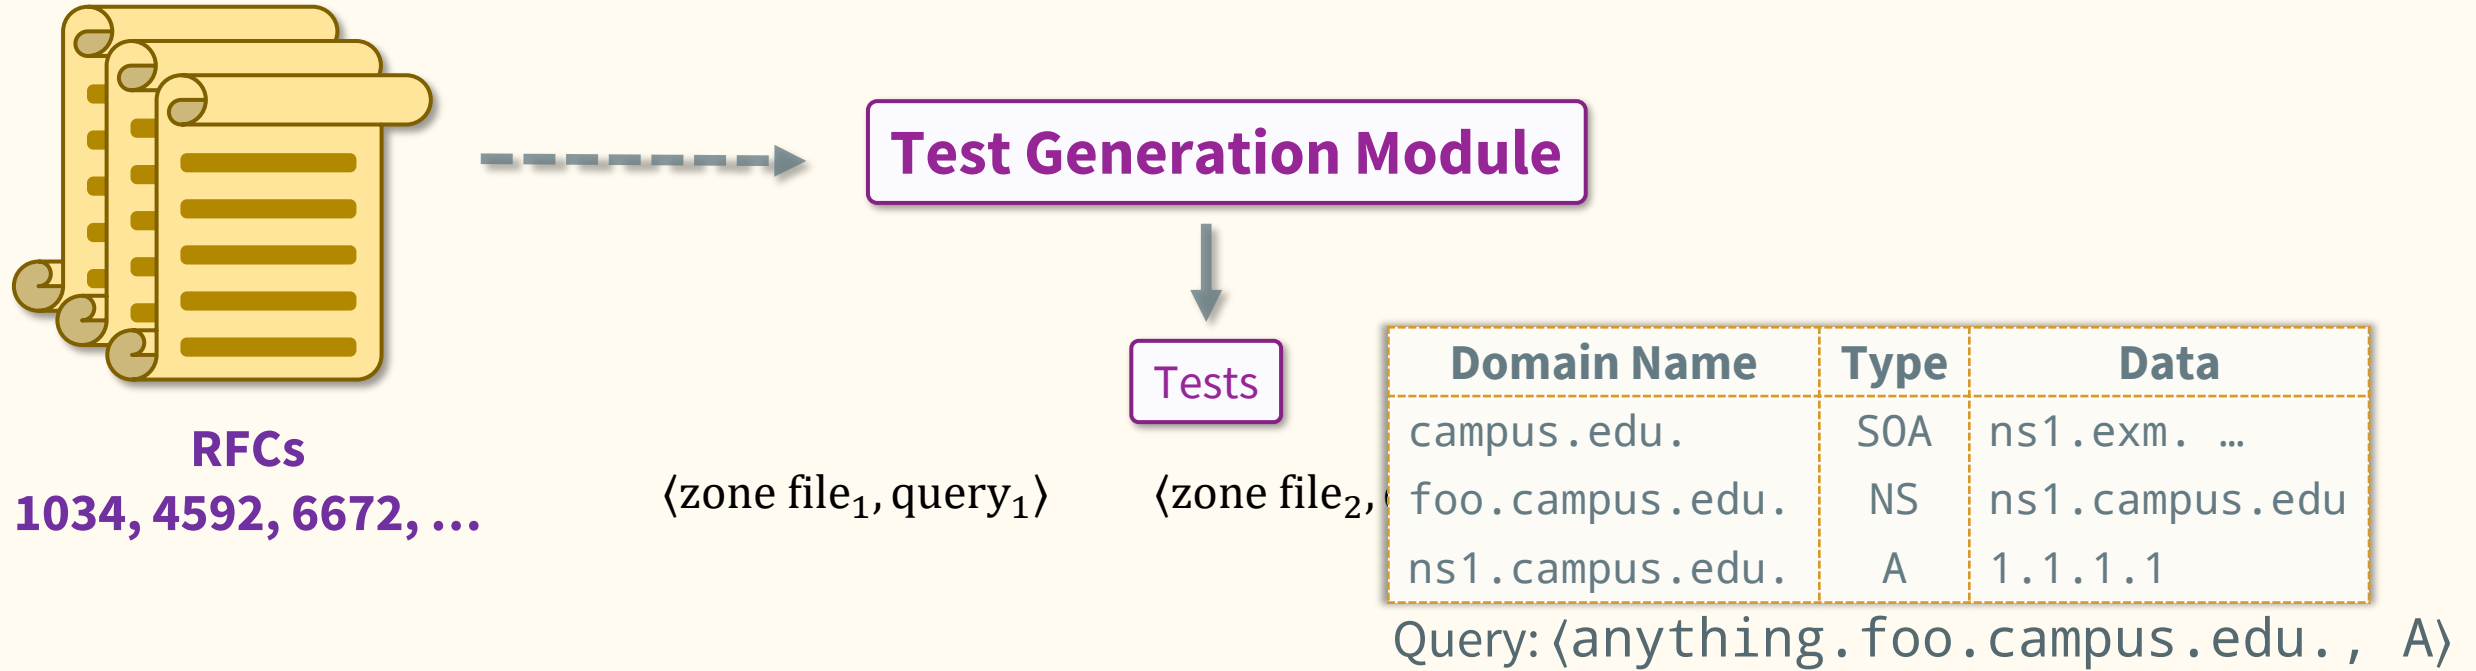
Query: $\langle \text{anything.foo.campus.edu.}, A \rangle$

FERRET: End-to-End Design



Modular Approach: Nameservers keep *no internal state* → A zone file is enough to test the *logic* at an isolated nameserver

FERRET: End-to-End Design



Modular Approach: Nameservers keep *no internal state* → A zone file is enough to test the *logic* at an isolated nameserver

FERRET generates tests that are *independent* of the *source code* →
Can test any nameserver implementation

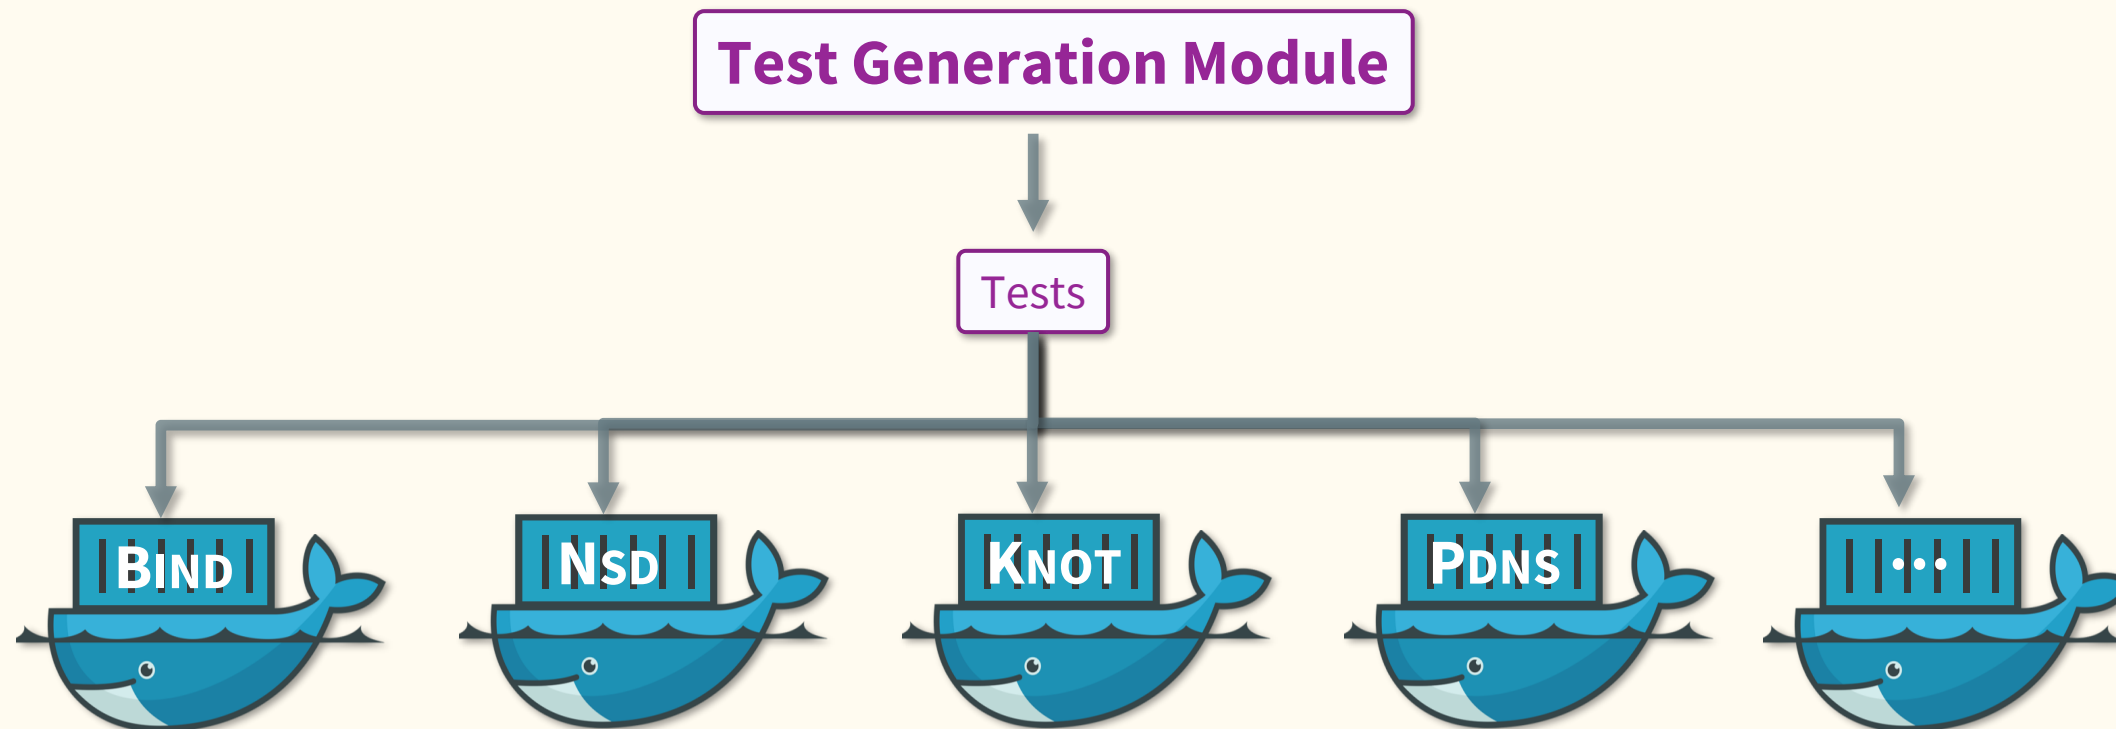
FERRET: End-to-End Design

Test Generation Module

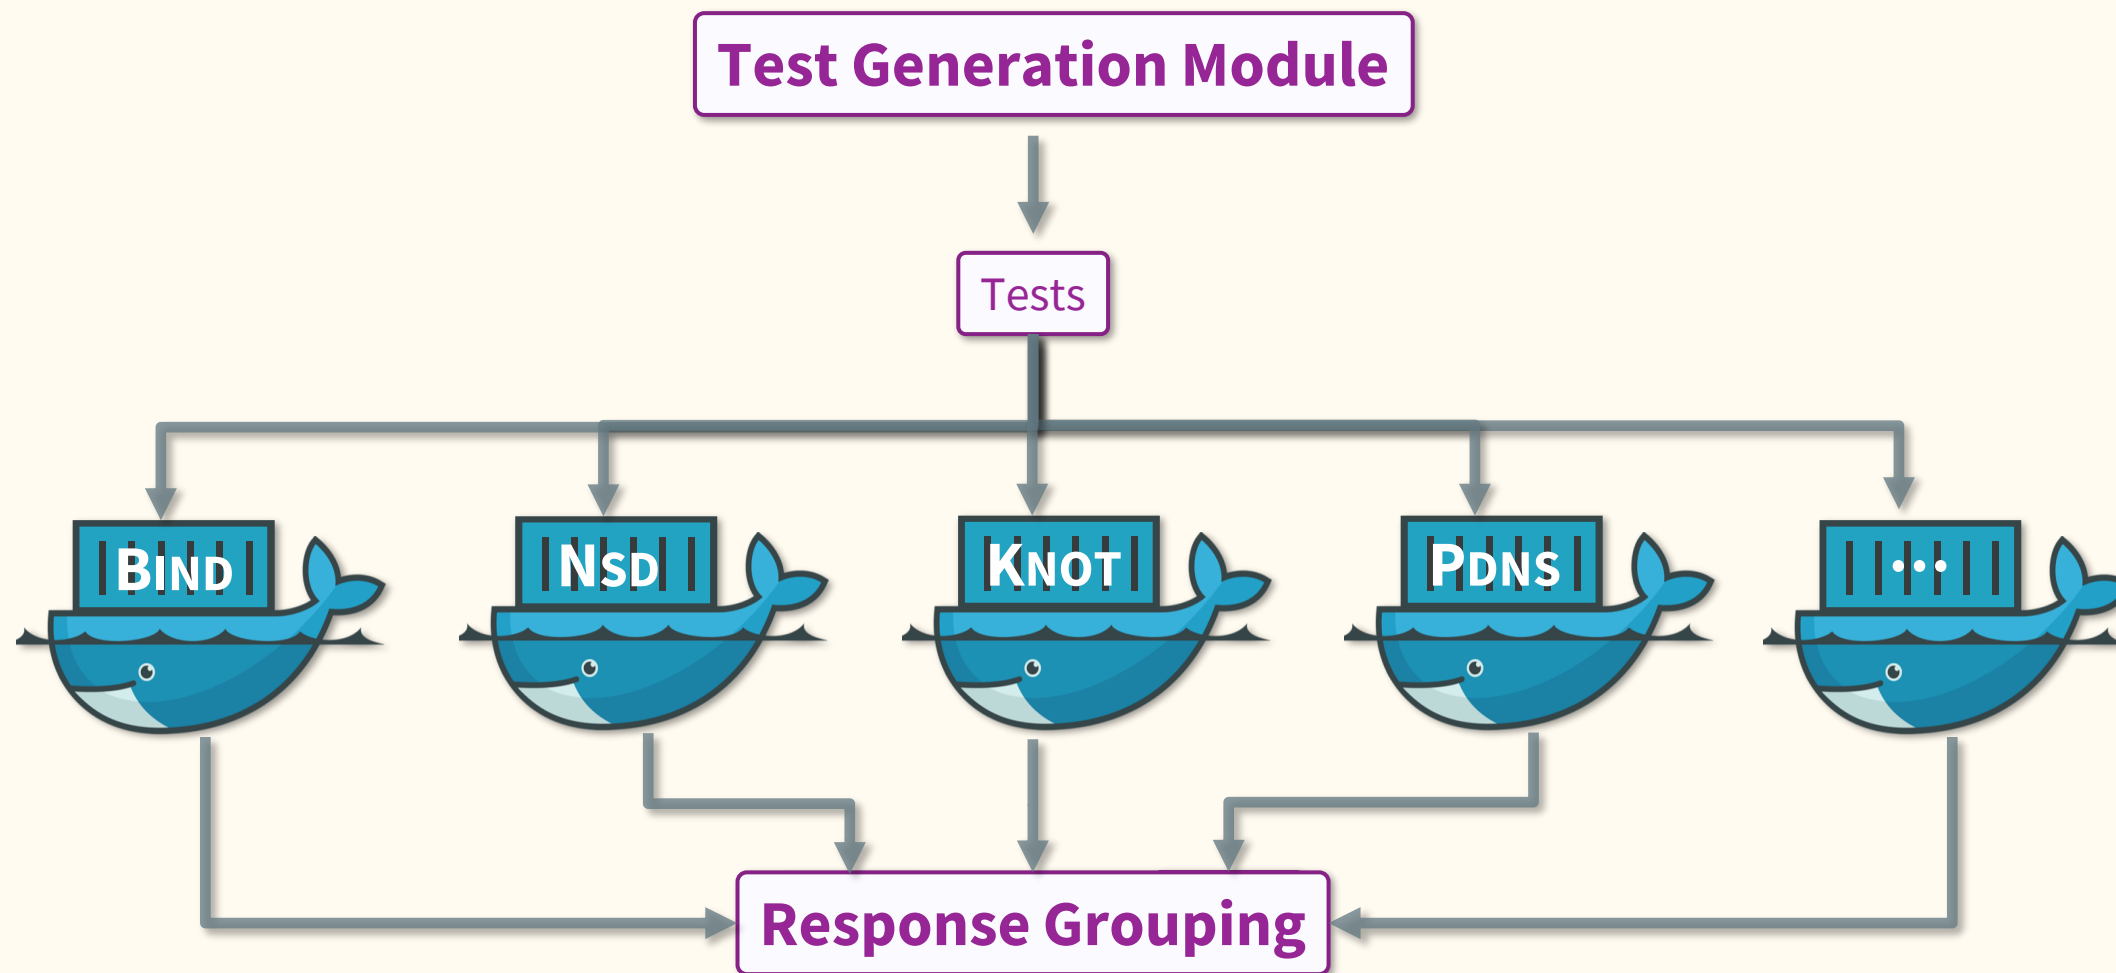


Tests

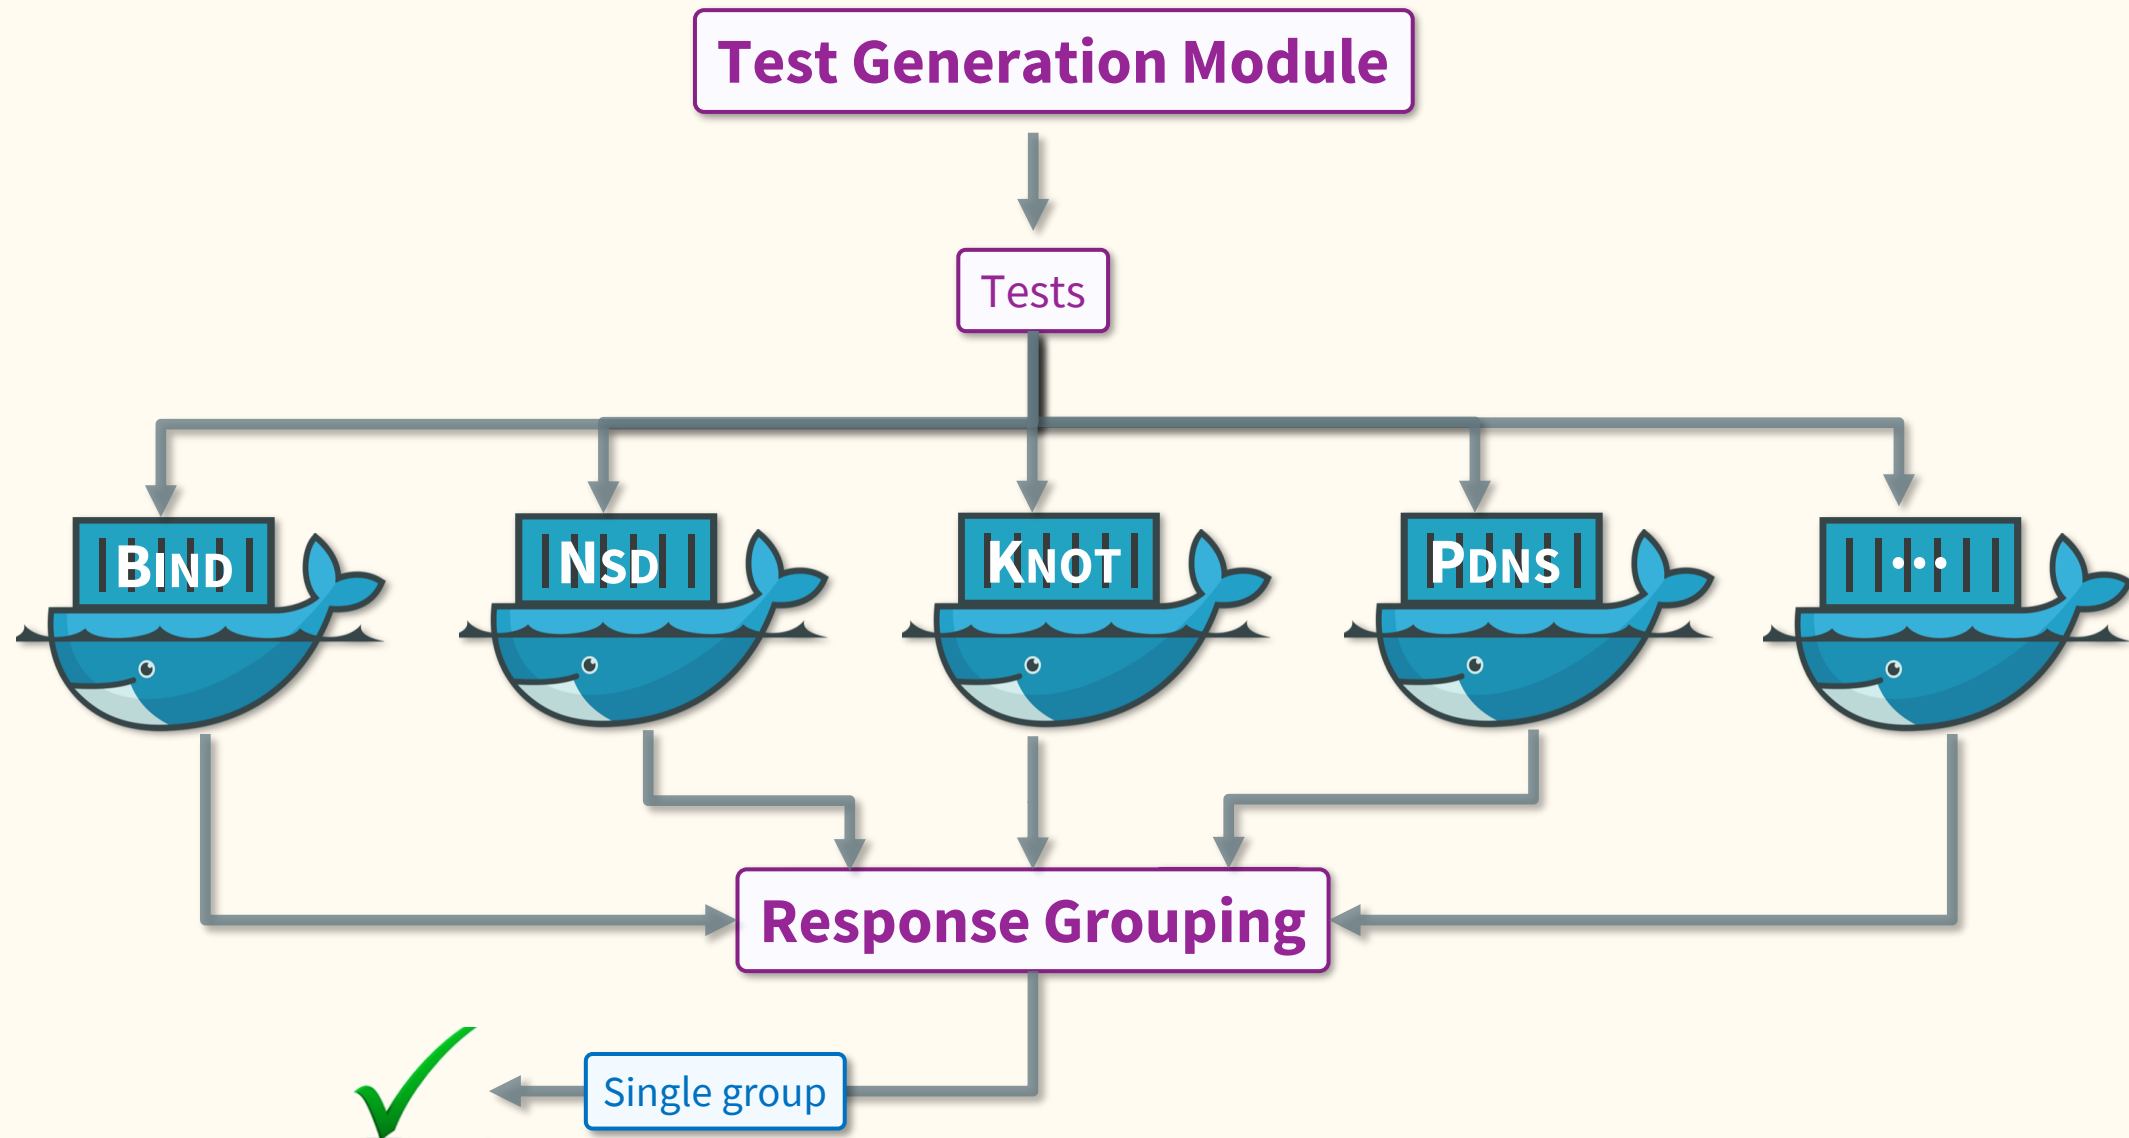
FERRET: End-to-End Design



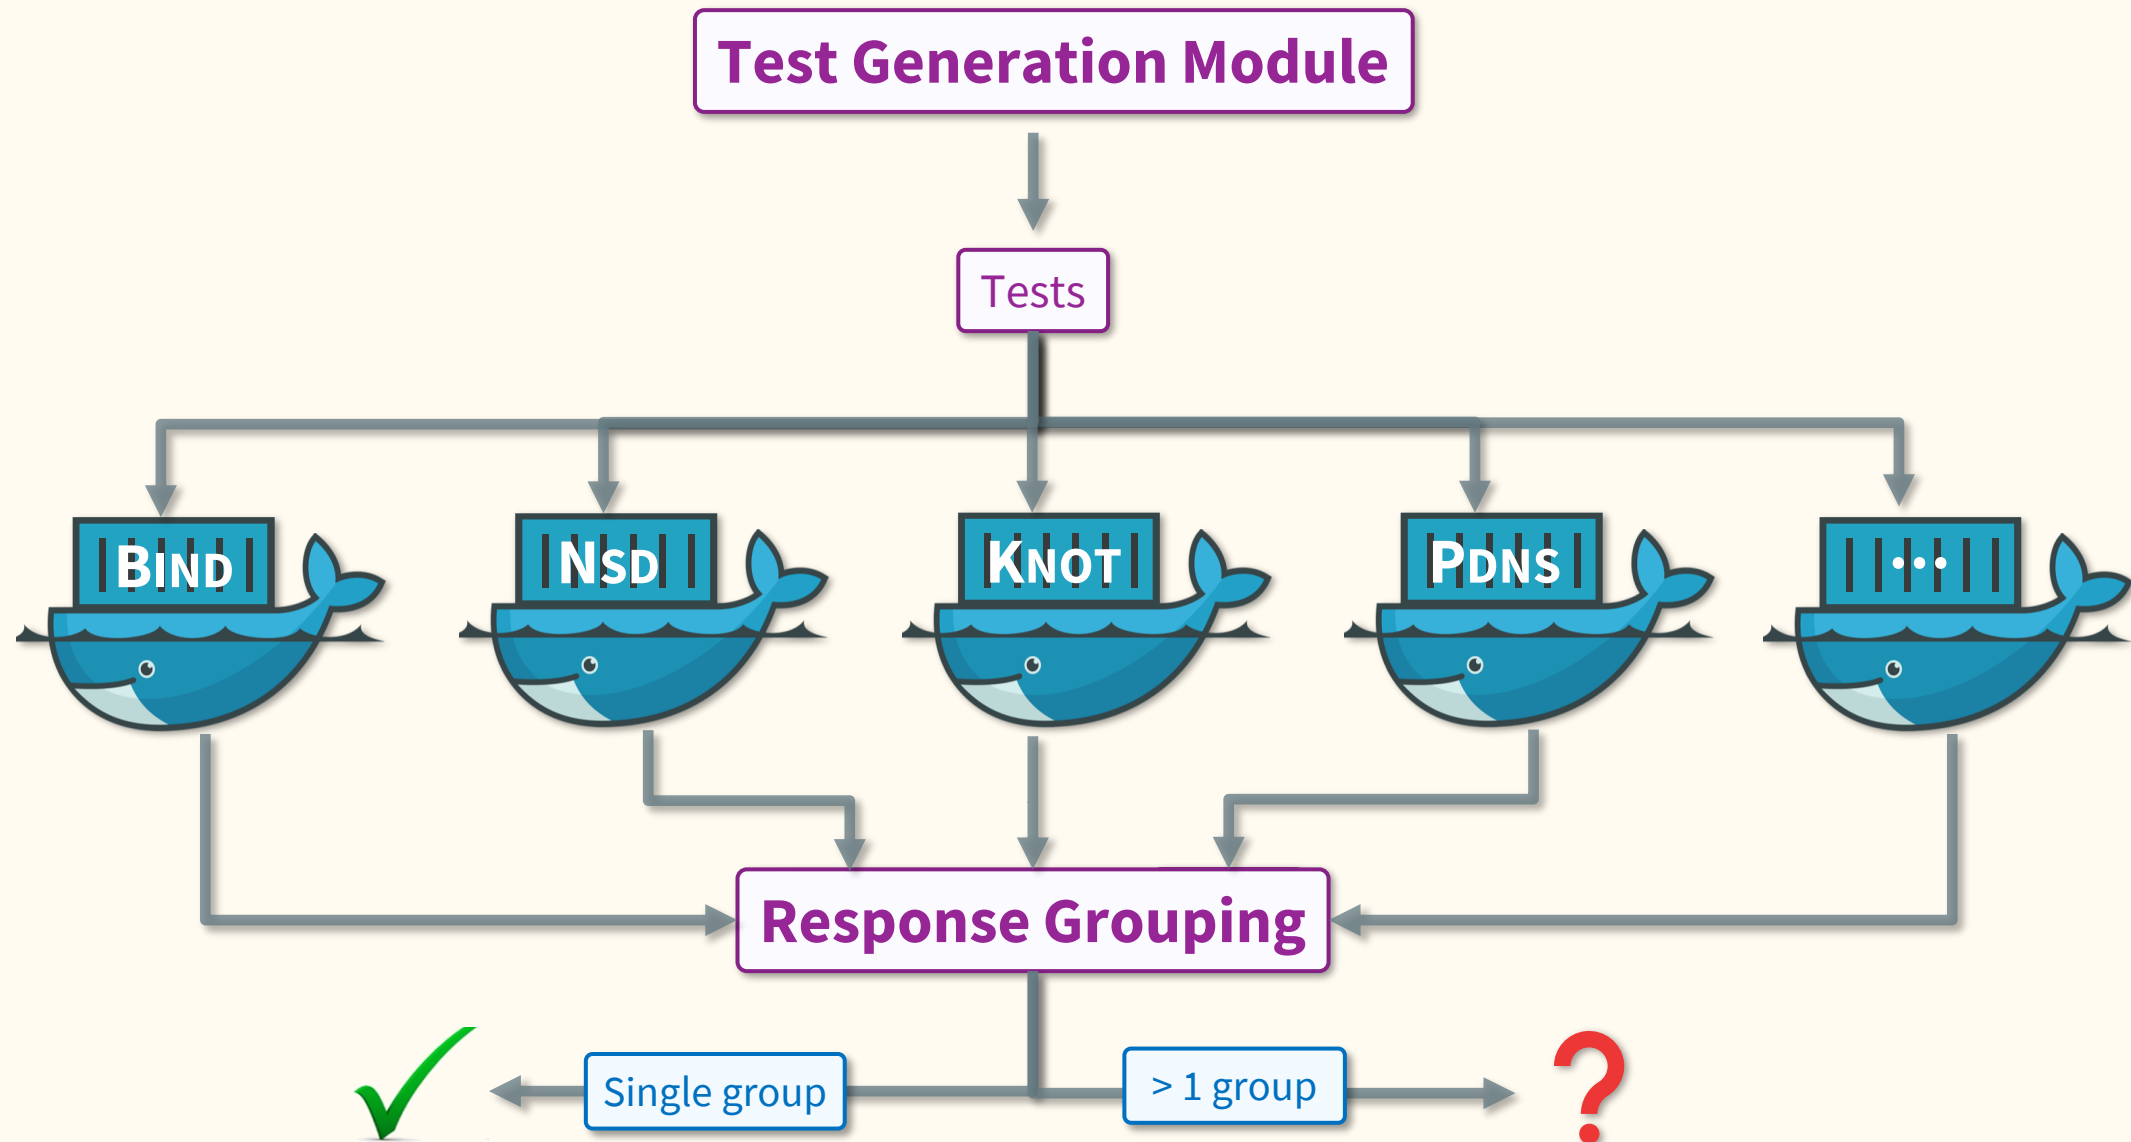
FERRET: End-to-End Design



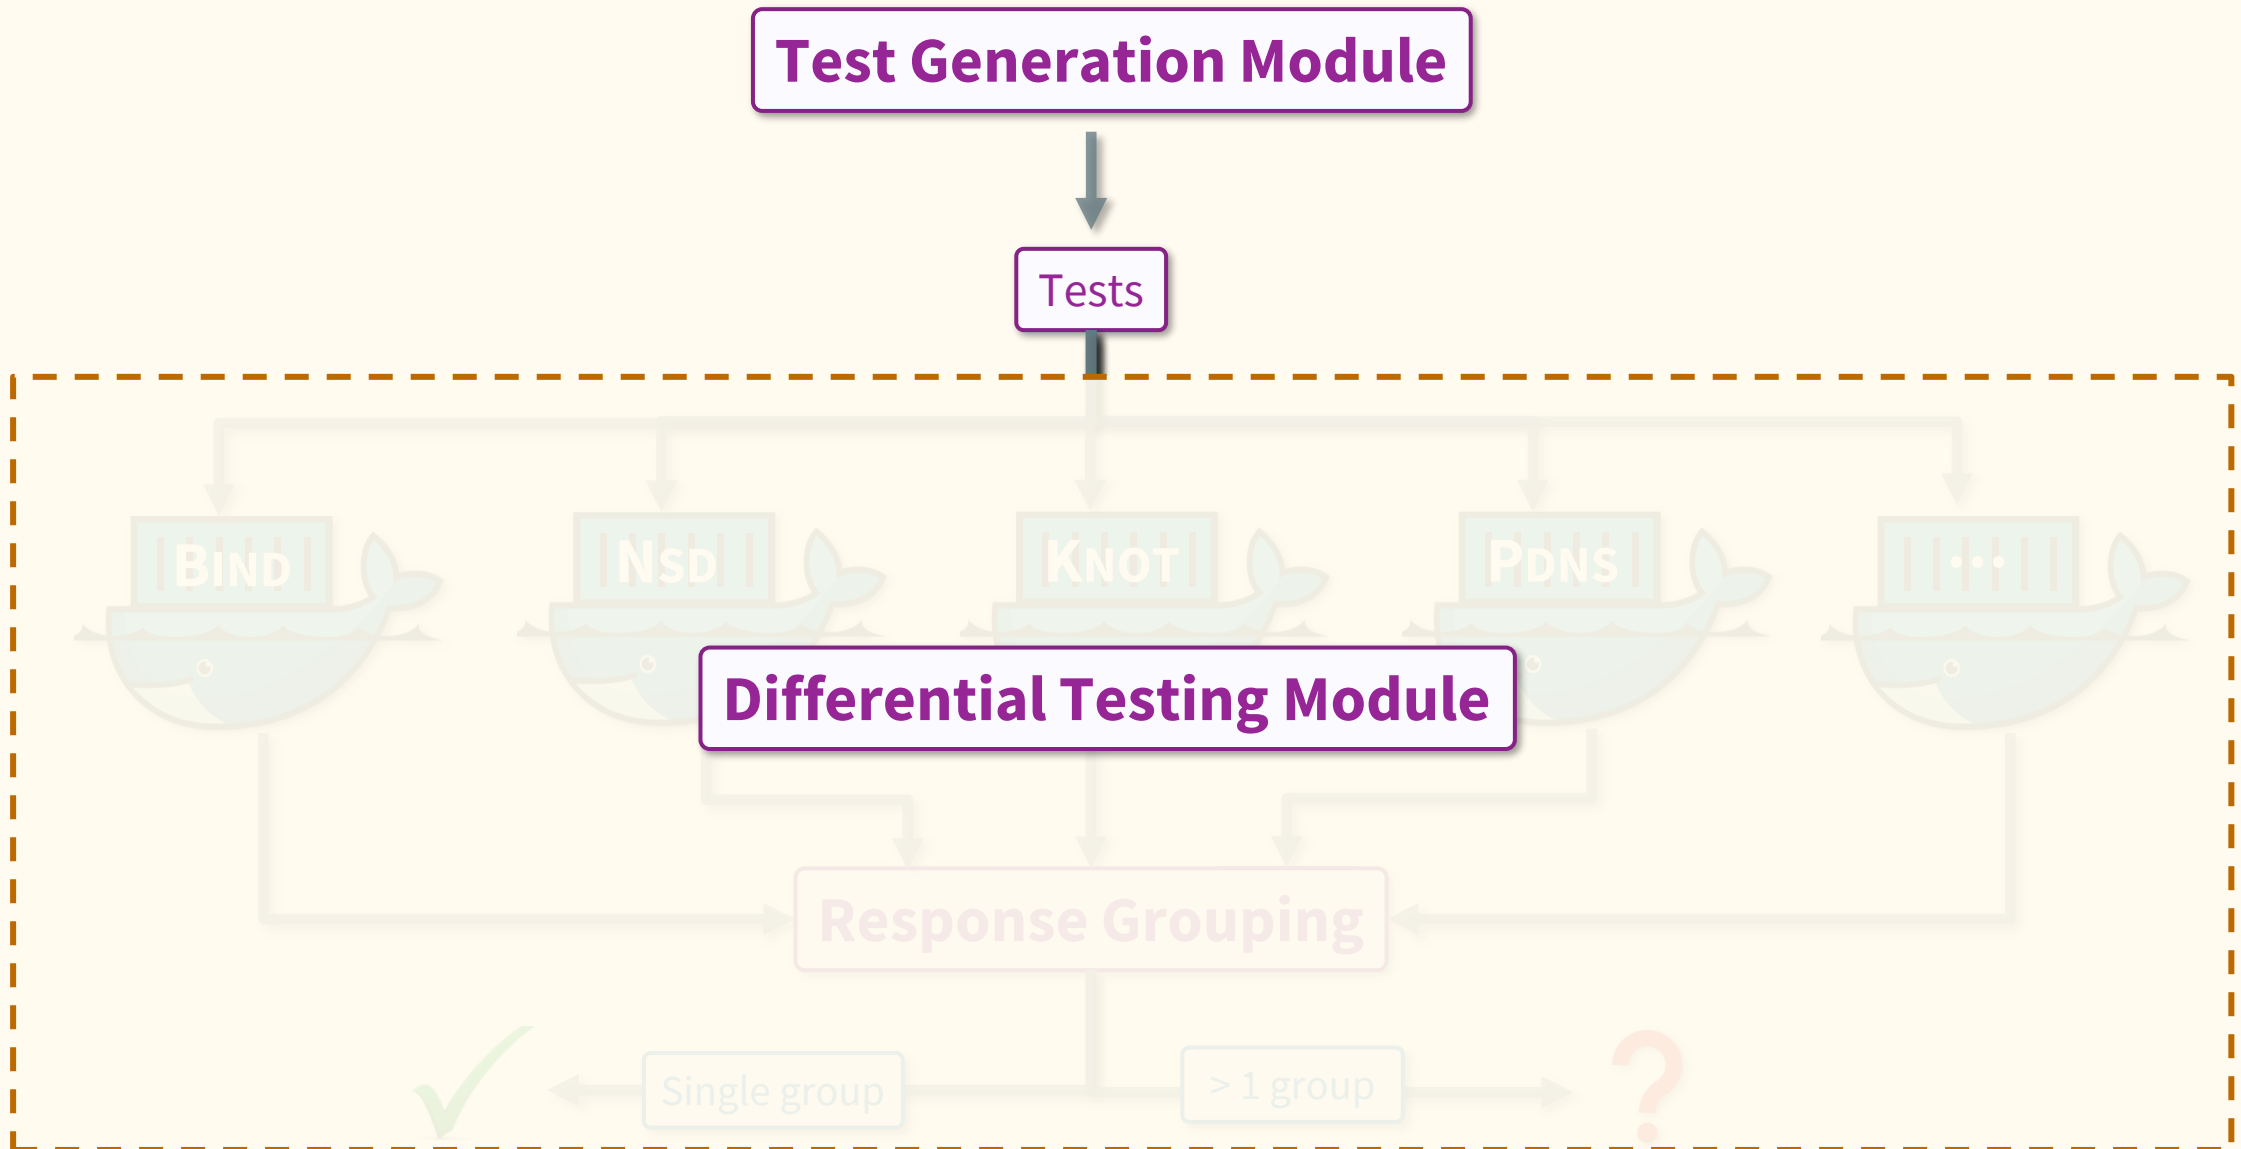
FERRET: End-to-End Design




FERRET: End-to-End Design



FERRET: End-to-End Design




Differential Testing



Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

Differential Testing




Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

© **Docker image** for each implementation

Differential Testing




Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

- ◎ **Docker image** for each implementation
- ◎ FERRET starts a **container** for each image

Differential Testing




Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

- ◎ **Docker image** for each implementation
- ◎ FERRET starts a **container** for each image
- ◎ Unique host port is mapped to **port 53** of the container

Differential Testing




Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

- ◎ **Docker image** for each implementation
- ◎ FERRET starts a **container** for each image
- ◎ Unique host port is mapped to **port 53** of the container
- ◎ Each container servers one zone file at a time as an **authoritative** zone

Differential Testing



Open-source Nameserver Implementations Tested

Implementation	Language	Description
BIND	C	<i>de facto</i> standard
POWERDNS	C++	popular in North Europe
NSD	C	hosts several TLDs
KNOT	C	hosts several TLDs
COREDNS	Go	used in Kubernetes
YADIFA	C	created by EURid (.eu)
TRUSTDNS	Rust	security, safety focused
MARADNS	C	lightweight server

- ◎ **Docker image** for each implementation
- ◎ FERRET starts a **container** for each image
- ◎ Unique host port is mapped to **port 53** of the container
- ◎ Each container servers one zone file at a time as an **authoritative** zone
- ◎ FERRET uses python library **dnspython** to send queries and collect responses

Bugs Found

Implementation	Bugs Found	Bug Type	Confirmed
BIND	Sibling glue records not returned	Wrong Additional	✓
	Zone origin glue records not returned	Wrong Additional	✓
	DNAME recursion denial-of-service	Server Crash	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
NSD	DNAME not applied recursively	Wrong Answer	✓
	Wrong RCODE when * is in Rdata	Wrong RCODE	✓
	Used NS records below delegation	Wrong Answer	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
POWERDNS	CNAME followed when not required	Wrong Answer	✓
	pdnsutil check-zone DNAME-at-apex	Preprocessor Bug	✓
KNOT	incorrect record synthesis	Wrong Answer	✓
	DNAME not applied recursively	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
	Error in DNAME-DNAME loop KNOT test	Faulty KNOT Test	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
COREDNS	NXDOMAIN for existing domain	Wrong RCODE	✓
	Wrong RCODE for CNAME target	Wrong RCODE	✓
	Wildcard CNAME loops & DNAME loops	Server Crash	✓
	Wrong RCODE for synthesized record	Wrong RCODE	?
	CNAME followed when not required	Wrong Answer	?
	Sibling glue records not returned	Wrong Additional	✓
YADIFA	CNAME chains not followed	Wrong Answer	✓
	Wrong RCODE for CNAME target	Wrong RCODE	✓
	Used records below delegation	Wrong Answer	✓
MARADNS [†]	AA flag set for zone cut NS RRs	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
TRUSTDNS [†]	wildcard match only one label	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
	AA flag set for zone cut NS RRs	Wrong Flag	✓
	CNAME loops crash the server	Server Crash	✓

[†]Implementations with unreported issues due to missing or unimplemented features


Example Bugs

FERRET Generated Test Case

Domain Name	Type	Data
uni.edu	SOA	ns1.exm. ...
test.uni.edu.	DNAME	edu.

Query: <test.uni.test.uni.edu., DNAME>

- Query is rewritten using DNAME to:
test.uni.test.uni.edu. CNAME test.uni.edu.
- The rewritten query will match exactly with the DNAME record.



Example Bugs

Crash in BIND



FERRET Generated

```
/lib/x86_64-linux-gnu/libpthread.so.0(+0x76db) [0x7f2094e876db]
/lib/x86_64-linux-gnu/libc.so.6(clone+0x3f) [0x7f209498b71f]
exiting (due to assertion failure)
Aborted (core dumped)
fatal error: stack overflow
```

Domain Name	Type
uni.edu	SOA
test.uni.edu.	DNAME edu.

- Server Crashes !!
- Easily-weaponizable denial-of-service vector
- Remotely Exploitable
- Affected all currently maintained BIND 9 branches

Initiated a responsible disclosure

CVE-2021-25215 (High Severity): An assertion check can fail while answering queries for DNAME records that require the DNAME to be processed to resolve itself

Example Bugs

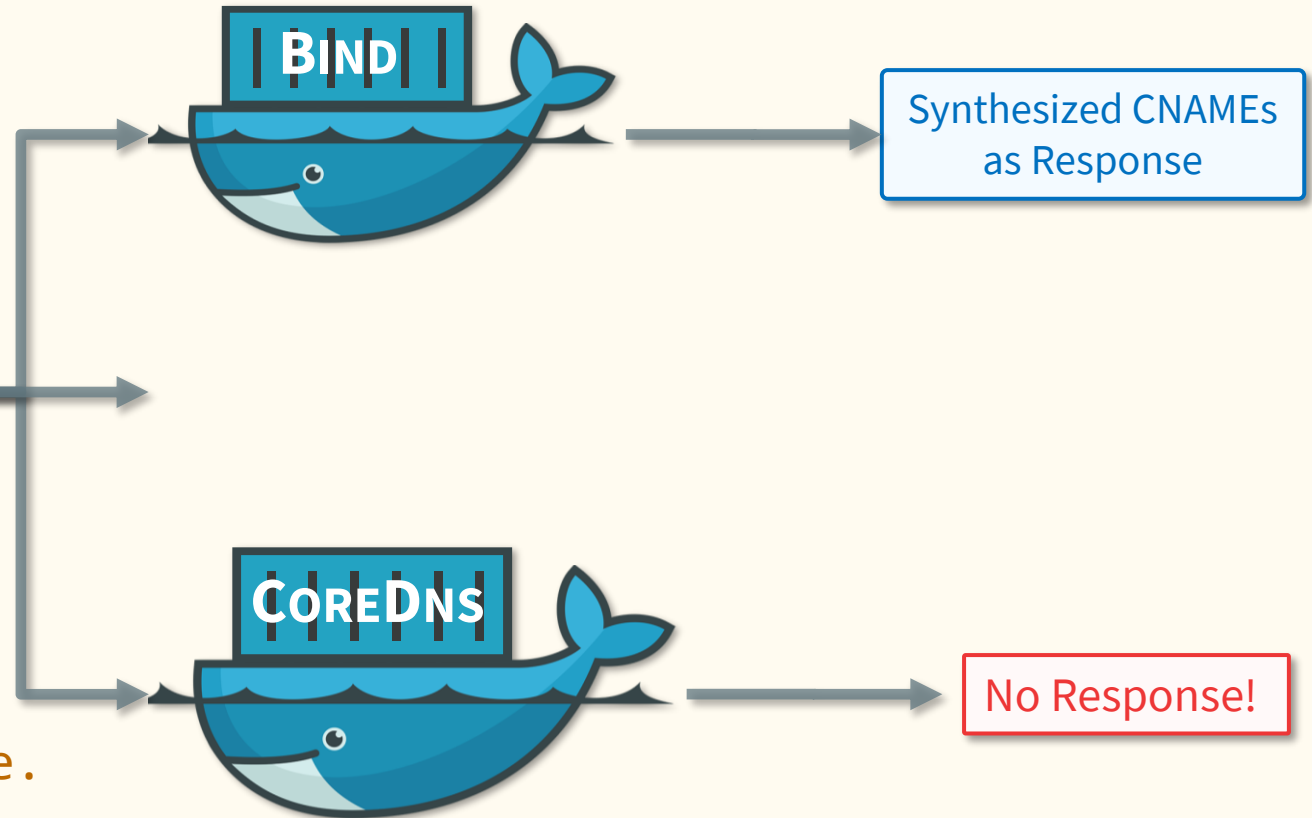
FERRET Generated Test Case

Domain Name	Type	Data
example.	SOA	ns1.exm. ...
*.example.	CNAME	foo.example.

Query: `(baz.bar.example., CNAME)`

🔗 Query is rewritten using CNAME to:
`baz.bar.example. CNAME foo.example.`

🔗 The rewritten query will match the wildcard again !



Popular open-sourced server written in Go
Recommended Server for Kubernetes

Example Bugs

FERRET Generated Test Case

Domain Name	Type	...
example.	SOA	ns1
*.example.	CNAME	foo

```
runtime: goroutine stack exceeds 1000000000-byte limit
runtime: sp=0xc03c6c0378 stack=[0xc03c6c0000, 0xc05c6c0000]
fatal error: stack overflow
```

Crashes !!
Serious Security Vulnerability
(DNS hosting services)

Fixed by adding a loop counter[†] – “For now it’s more important to protect ourselves than to give the client a valid answer”

Popular
Recommended Server for Kubernetes

[†]<https://github.com/coredns/coredns/issues/4378>

Example Bugs

Performance Bug in BIND

Open issue - May 2021 milestone

Domain Name	Type	Data
campus.edu.	SOA	ns1.exm. ...
foo.campus.edu.	NS	ns1.campus.edu
ns1.campus.edu.	A	1.1.1.1

Query: <anything.foo.campus.edu., A>

Response from POWERDNS, KNOT, NSD:

Authority Section:

foo.campus.edu. NS ns1.campus.edu

Additional Section:

ns1.campus.edu. A 1.1.1.1

- ✦ BIND does not return the glue record
- ✦ Response from BIND “This report turns out to be very interesting. Here is what I managed to find out”
- ✦ BIND uses a “glue cache” to speed up the identification of glue records, but it had two unrelated errors.
 - ↳ If the **cache lookup fails**, then glue records are supposed to be searched for in the zone file, but the **latter was never** happening.
 - ↳ glue records for siblings domain nameservers were **accidentally** never searched for at all.

Example Bugs

Data Structure Bug in Nsd

Fixed the issue

Domain Name	Type	Data
booksonline.	SOA	ns1.exm. ...
buy.booksonline.	CNAME	www.*.booksonline.

Query: <buy.booksonline., A>

Response from NSD:

RCODE: **NOERROR**

Answer Section:

buy.booksonline. CNAME www.*.booksonline.

- ✦ BIND, KNOT, POWERDNS return with **NXDOMAIN** as CNAME target does not exist
- ✦ **RCODE is important** as resolvers use it to determine whether domains exist or not
- ✦ NSD responded - “It has to do with the **internal data structure** for storing domains in the memory of NSD, there a domain struct is created for the right hand of the CNAME, and it is set to be non-existing. The **is_existing** was not checked for the wildcard expansion, and this is fixed by the commit. **...Thanks for the report!**”

Example Bugs

CNAME Bug in YADIFA

Fixed the issue

Domain Name	Type	Data
dept.com.	SOA	ns1.exm. ...
www.cs.dept.com.	CNAME	cs.dept.com.
cs.dept.com.	CNAME	dept.com
dept.com.	A	2.2.2.2

Query: <www.cs.dept.com., A>

- ✦ Expected response is to rewrite the query **twice** and return the IP record
- ✦ YADIFA rewrote it only **once** and was **not following** the CNAME chains.
- ✦ CNAME chains are used extensively by **CDNs** so its important to follow
- ✦ YADIFA acknowledged and said – “The **rerun** of the query was **incorrectly disabled**, the issue is fixed and will be updated on github on our next update of the code.”

Example Bugs

DNAME-DNAME Loop Bug in KNOT

Domain Name	Type	Data
corp.	SOA	ns1.exm. ...
corp.	NS	ns1.com.
corp.	DNAME	us.corp.

Query: <www.corp., NS>

↪ Query is rewritten using DNAME to:
www.corp. CNAME www.us.corp.

↪ The rewritten query will again be rewritten using DNAME to:
www.us.corp. CNAME www.us.us.corp.

↪ Leads to an **infinite recursion** !!

- ✦ BIND applies DNAME **multiple times** and stops when limit reaches **17**
- ✦ POWERDNS returns **SERVFAIL**
- ✦ KNOT and Nsd applied DNAME **only once**
 - ↪ Works here but had to be applied multiple times when there is no loop
 - ↪ Both fixed the issue †

† <https://github.com/NLnetLabs/nsd/issues/151>

† <https://gitlab.nic.cz/knot/knot-dns/-/issues/714>

Example Bugs

DNAME-DNAME Loop Bug in KNOT

- ✦ BIND applies DNAME **multiple times** and stops when limit reaches **17**

tests-extra/data/flags.zone

65	65	c.dname-tree	CNAME	dns1
66	66	d.dname-tree	CNAME	cname-wildcard
67	67	e.dname-tree	CNAME	e.dname
68	-	f.dname-tree	DNAME	dname-tree
68	68	+	f.dname-tree	f.f.dname-tree
69	69	www.corp.	CNAME	www.us.corp.

```

200 200 # DNAME-DNAME Loop
201 - resp = knot.dig("f.dname.flags", "A", udp=True)
202 - resp.cmp(bind)
201 + resp = knot.dig("x.f.dname.flags", "A", udp=True)
202 + resp.check(rcode="NOERROR")
203 + resp.check_record(name="dname.flags.", rtype="DNAME", ttl=3600, rdata="dname-tree.flags.")
204 + resp.check_record(name="x.f.dname.flags.", rtype="CNAME", ttl=3600, rdata="x.f.dname-tree.flags.")
205 + resp.check_record(name="f.dname-tree.flags.", rtype="DNAME", ttl=3600, rdata="f.f.dname-tree.flags.")
206 + resp.check_record(name="x.f.dname-tree.flags.", rtype="CNAME", ttl=3600, rdata="x.f.f.dname-tree.flags.")
207 + resp.check_counts(4, 0, 0)
208 + # resp.cmp(bind) BIND responds partially unrolled CNAME Loop

```

✦ The rewritten query will again be rewritten using DNAME to:
www.us.corp. CNAME **www.us.us.corp.**

✦ Leads to an **infinite recursion** !!

- ✦ KNOT had a **test suite** comparing responses with BIND and a test is mentioned as testing the infinite loop as this

- ✧ Test zone file was not **properly constructed**, and that error led to having **no loop**
- ✧ **Fixed it**★ and went with **single response** unlike 17 for a loop

† <https://github.com/NLnetLabs/nsd/issues/151>

† <https://gitlab.nic.cz/knot/knot-dns/-/issues/714>

★ <https://gitlab.nic.cz/knot/knot-dns/-/issues/703>

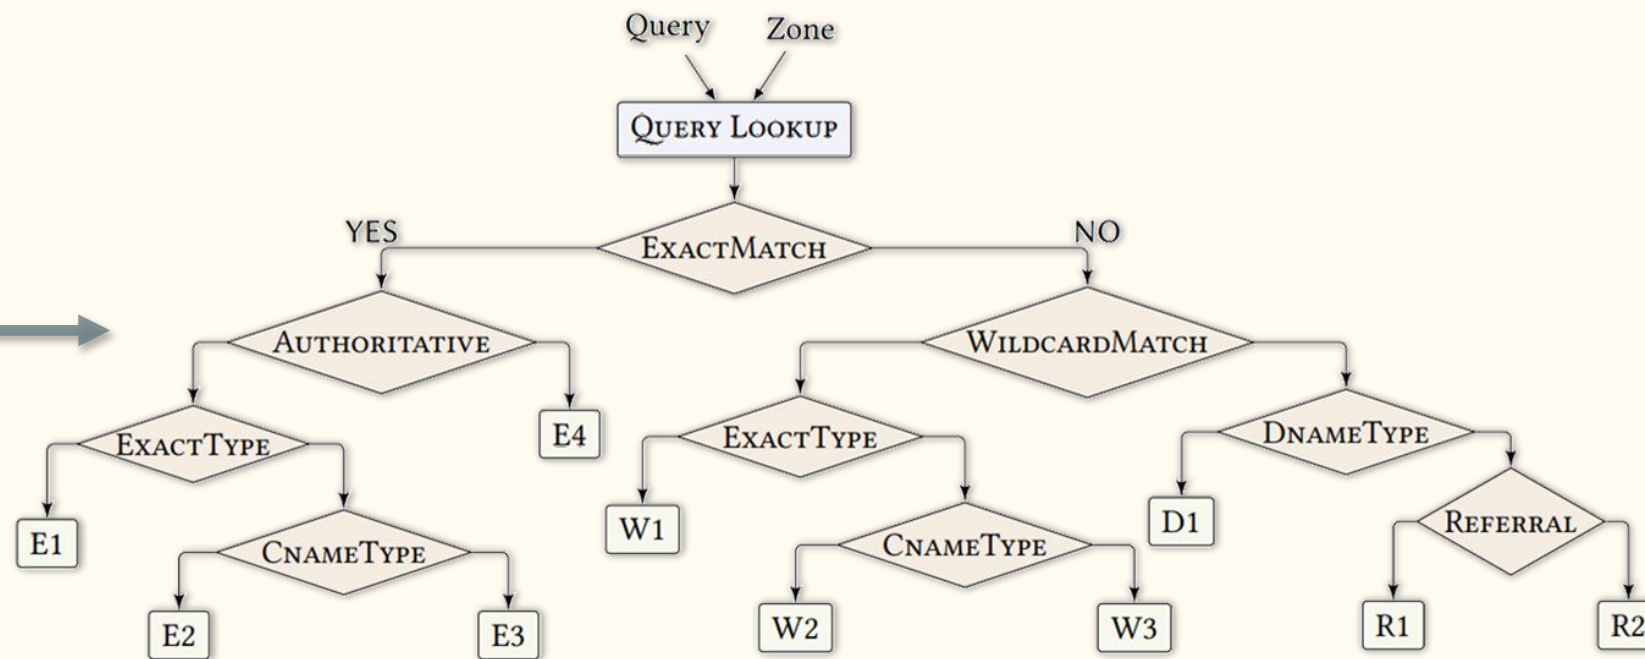
?Test Generation Module?



Test Generation Module

RFCs
1034, 4592, 6672, ...

Formal Model[†]

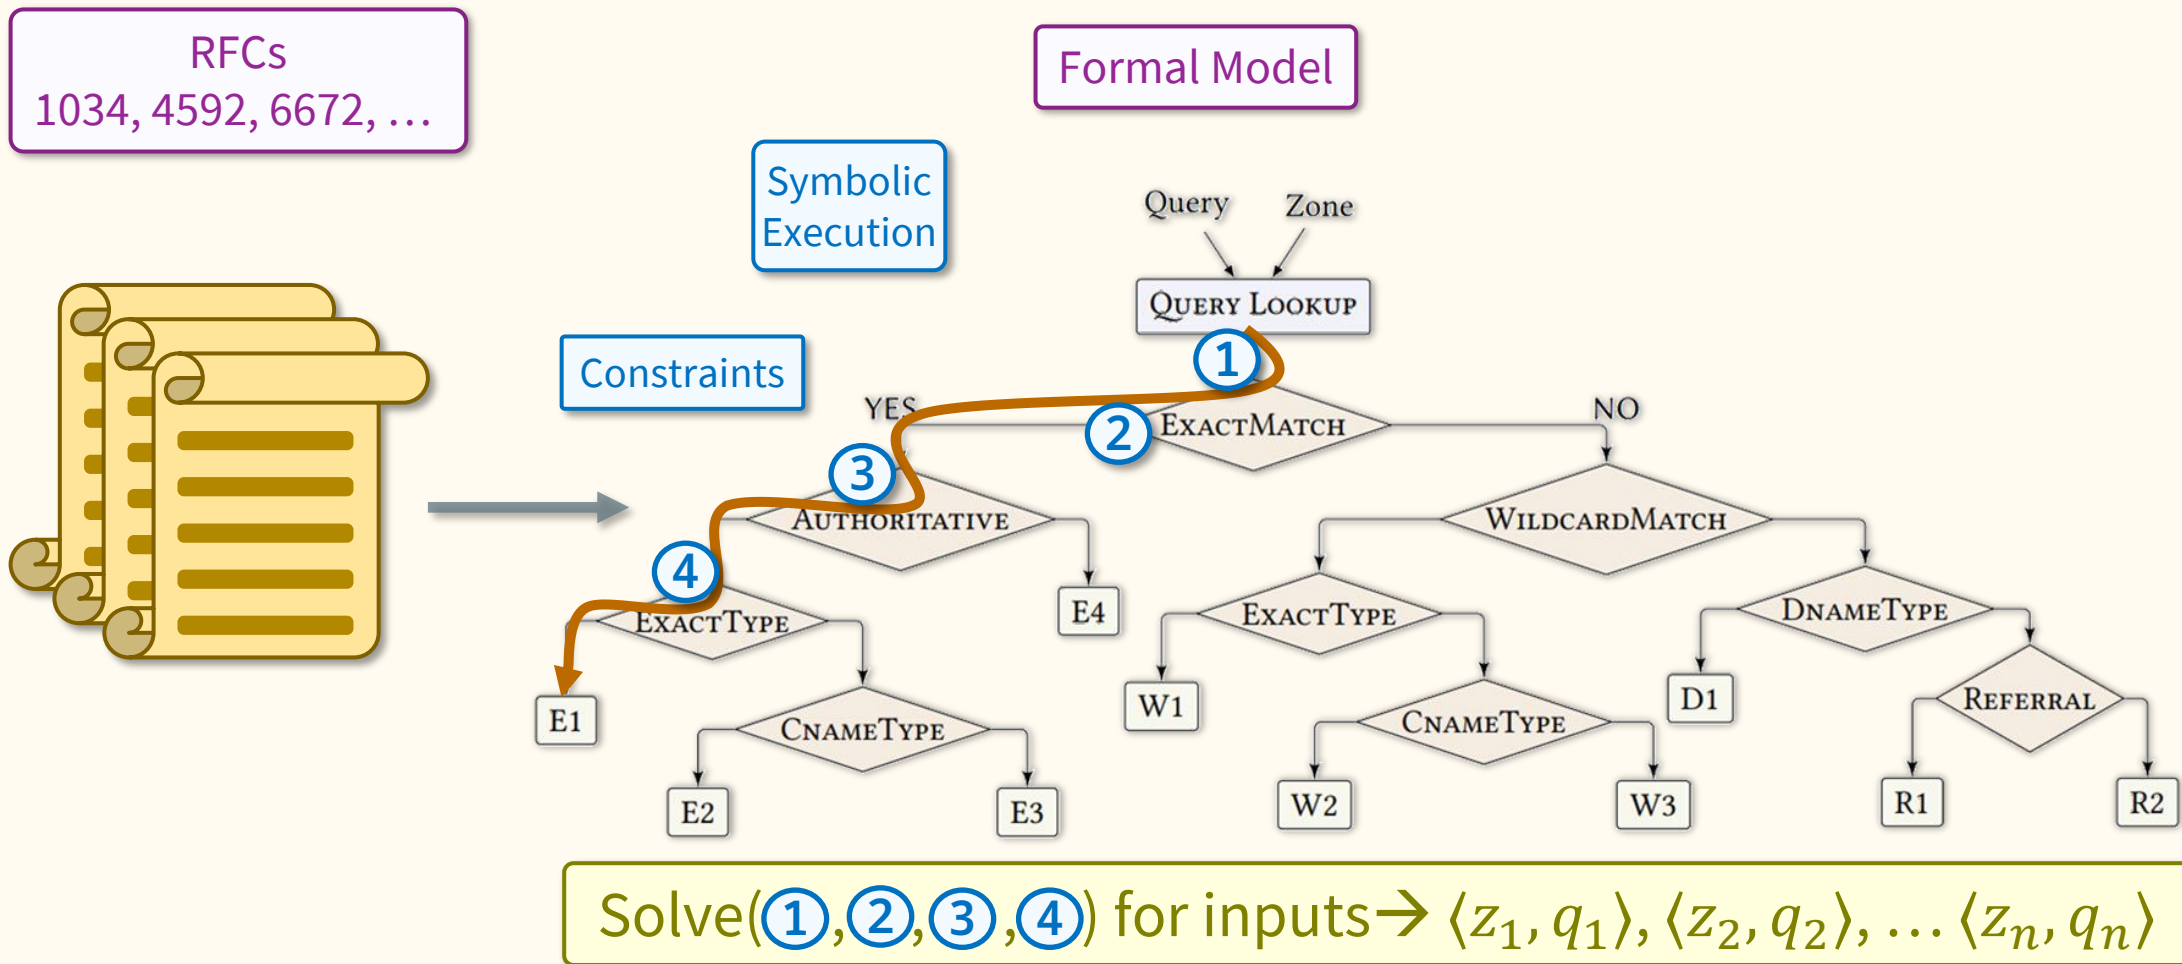


English

Declarative (Mathematical)
specification of the nameserver logic

[†]GRoot: Proactive Verification of DNS Configurations – Siva Kakarla *et al.*, SIGCOMM 2020

Test Generation Module

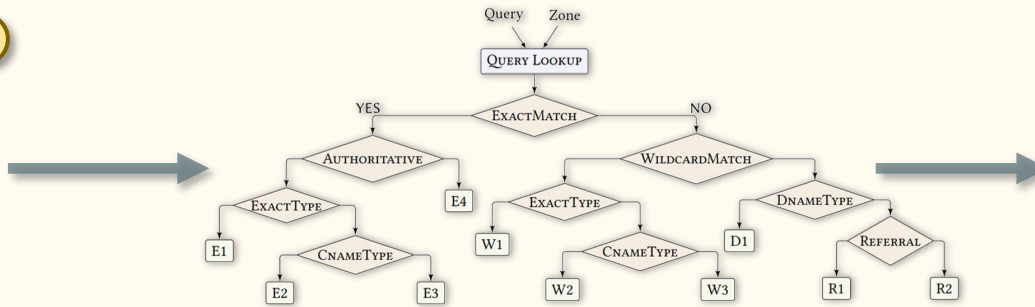


Test Generation Module

RFCs
1034, 4592, 6672, ...

Formal Model

Executable Version in Zen



```

1 Zen<Response> QueryLookup(
2   Zen<Query> q,
3   Zen<Zone> z)
4 {
5   var records = SelectBestRecords(q, z);
6   var rname = records.At(0).Value().Name();
7   var types = records.Select(r => r.Type());
8
9   return If(
10    rname == q.Name(),
11    ExactMatch(records, q, z),
12    If(
13     IsWildcardMatch(q.Name(), rname),
14     WildcardMatch(records, q, z),
15     If(
16      types.Any(t => t == RType.DNAME),
17      Rewrite(records, q),
18      If(
19       And(types.Any(t => t == RType.NS),
20        Not(types.Any(t => t == RType.SOA))),
21      Response(Tag.R1,
22        Delegation(records, z), Null<Query>()),
23      Response(Tag.R2, empty, Null<Query>())
24      ))));
25 }
  
```

An **executable version** of formal model is implemented in **Zen**, a domain-specific modeling language embedded in **C#** with built-in support for **symbolic execution**

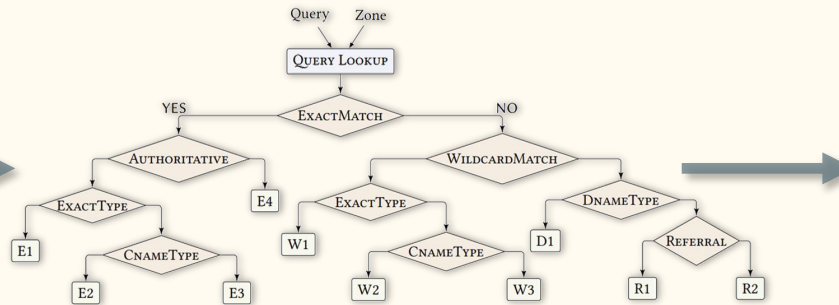

Test Generation Module

RFCs
1034, 4592, 6672, ...

Formal Model

Executable Version in Zen

Tests



```

1 Zen<Response> QueryLookup(
2   Zen<Query> q,
3   Zen<Zone> z)
4 {
5   var records = SelectBestRecords(q, z);
6   var rname = records.At(0).Value().Name();
7   var types = records.Select(r => r.Type());
8
9   return If(
10    rname == q.Name(),
11    ExactMatch(records, q, z),
12    If(
13      IsWildcardMatch(q.Name(), rname),
14      WildcardMatch(records, q, z),
15      If(
16        types.Any(t => t == RType.DNAME),
17        Rewrite(records, q),
18        If(
19          And(types.Any(t => t == RType.NS),
20              Not(types.Any(t => t == RType.SOA))),
21          Response(Tag.R1,
22                  Delegation(records, z), Null<Query>()),
23          Response(Tag.R2, empty, Null<Query>())
24        )));
25 }
  
```

Symbolic
Execution

$\langle \text{zone file}_1, \text{query}_1 \rangle$

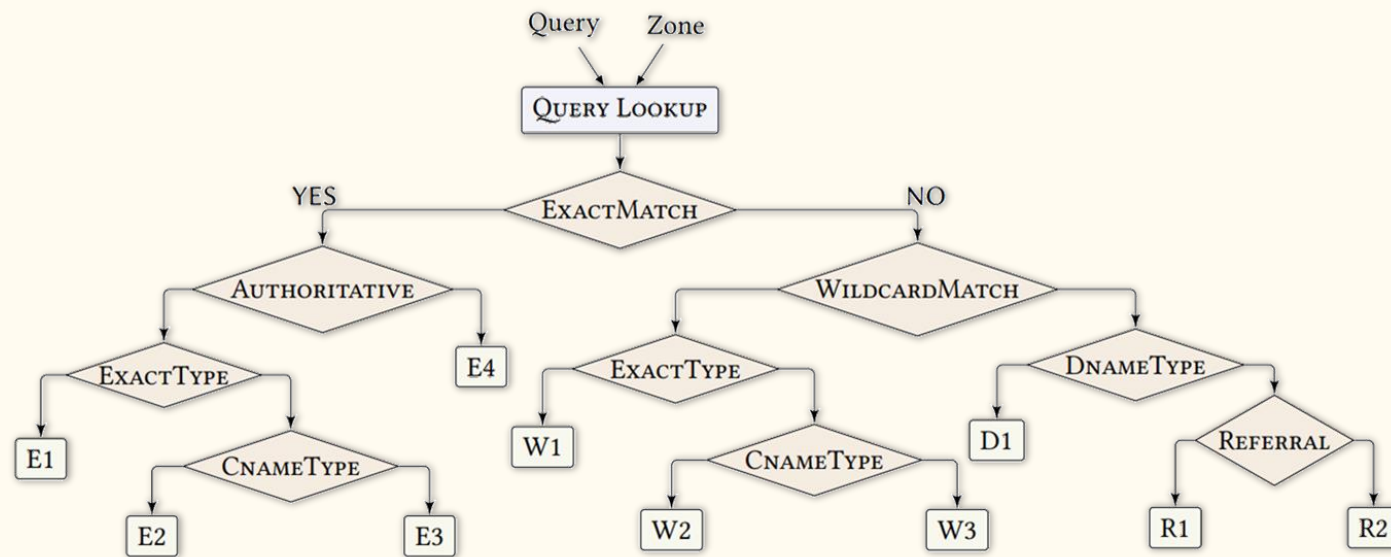
$\langle \text{zone file}_2, \text{query}_2 \rangle$

...

$\langle \text{zone file}_n, \text{query}_n \rangle$

An **executable version** of formal model is implemented in **Zen**, a domain-specific modeling language embedded in **C#** with built-in support for **symbolic execution**

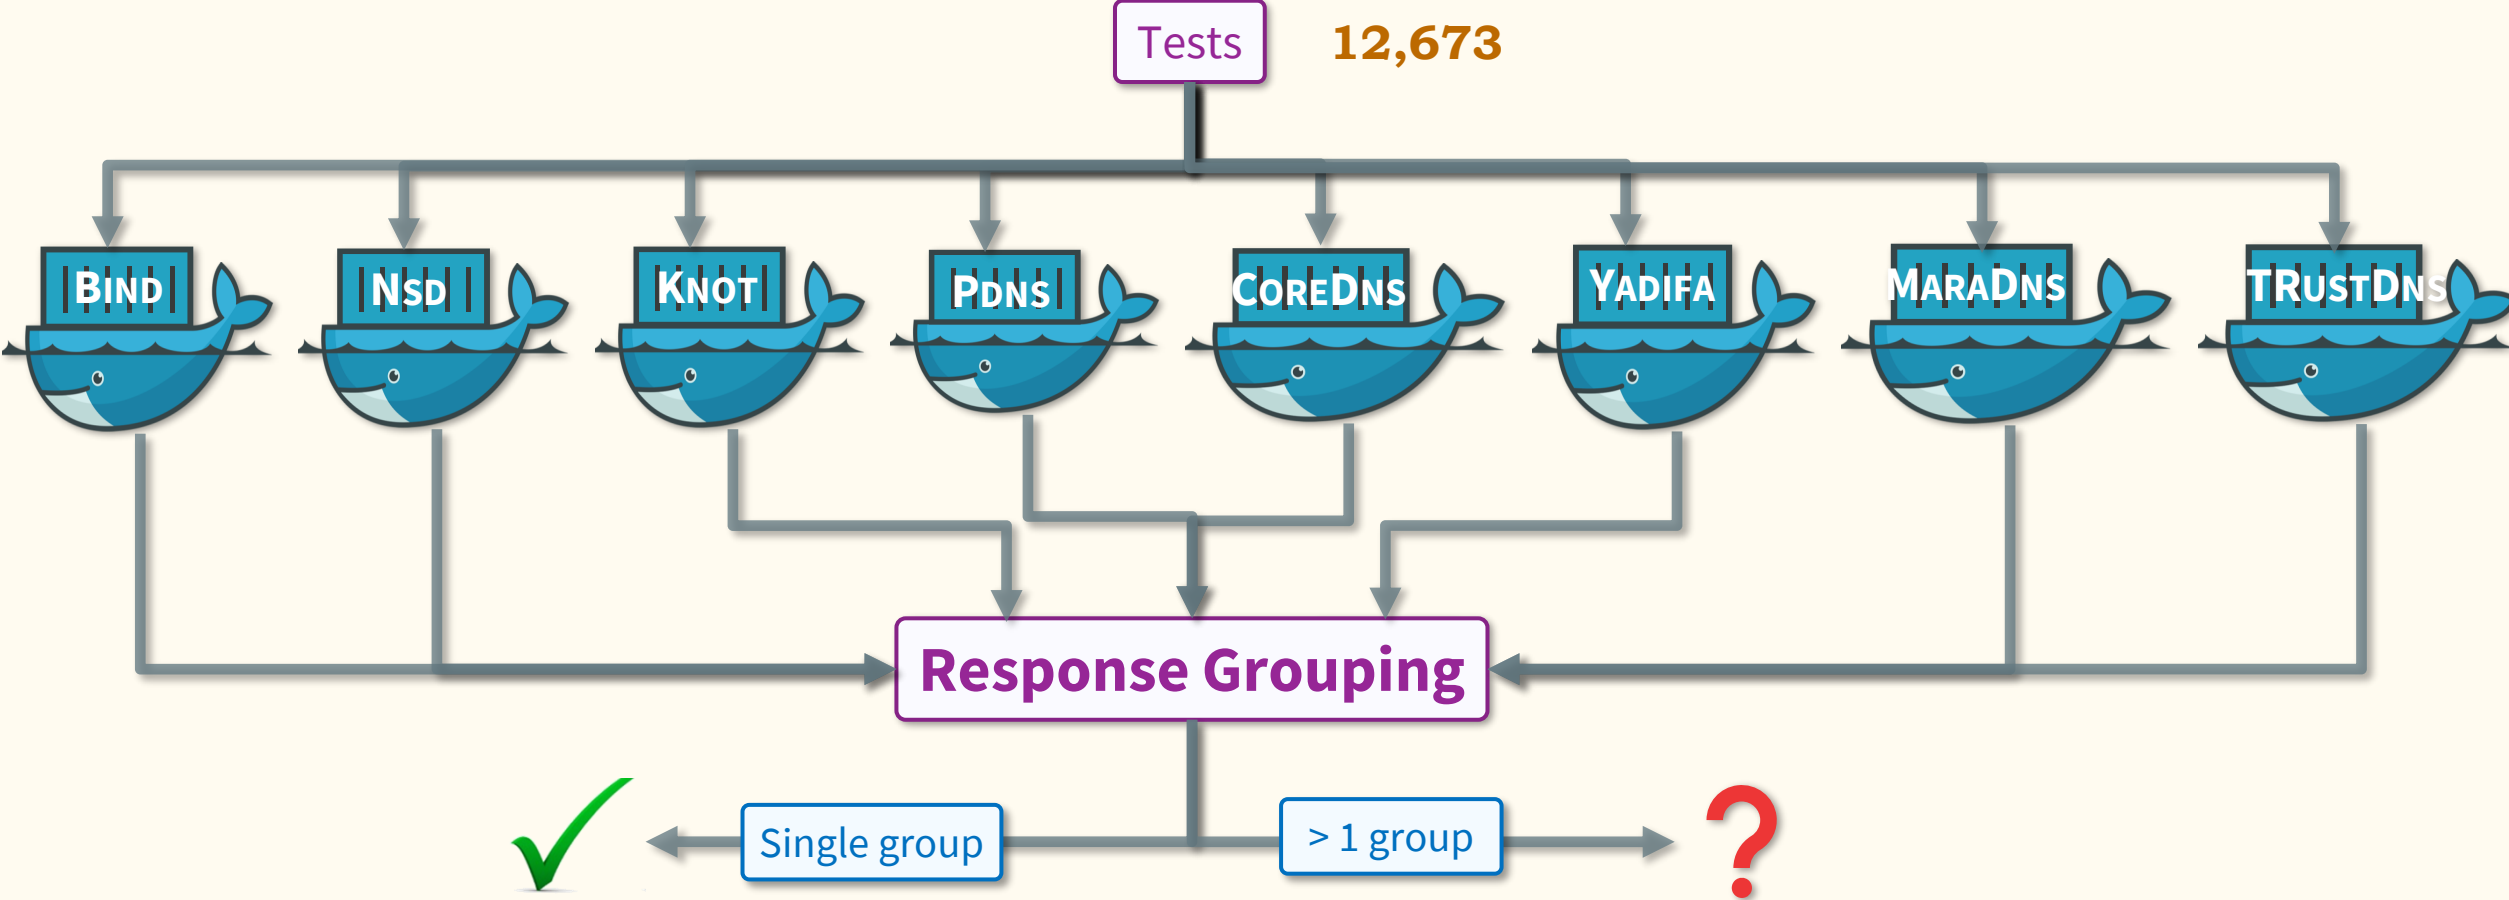
Test Generation Statistics



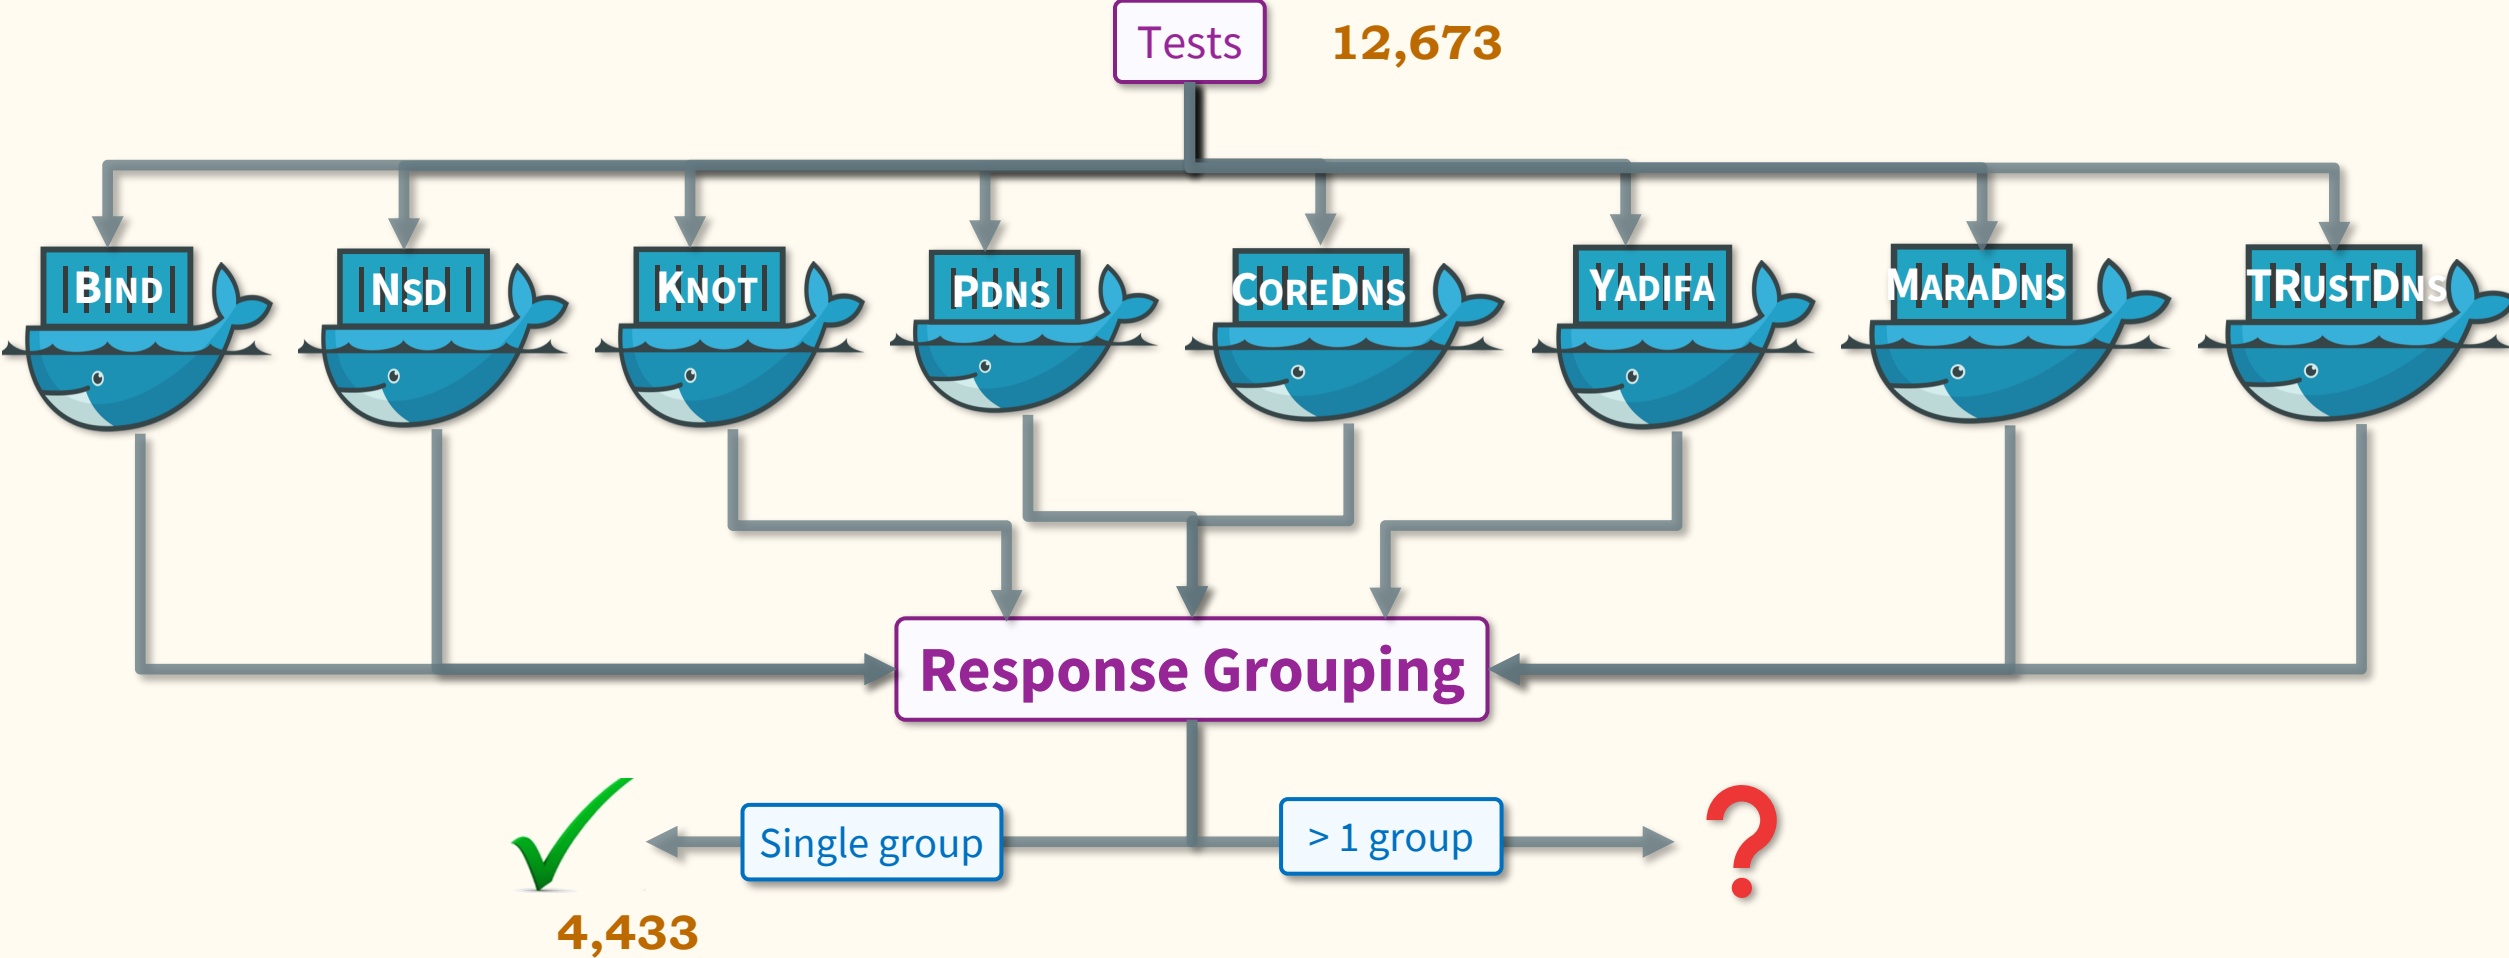
Length of each domain name and the number of records in the zone ≤ 4

Model Case	Number of Tests
E1	3180
E2	12
E3	96
E4	6036
W1	60
W2	24
W3	18
D1	230
R1	2980
R2	37
Total	12,673

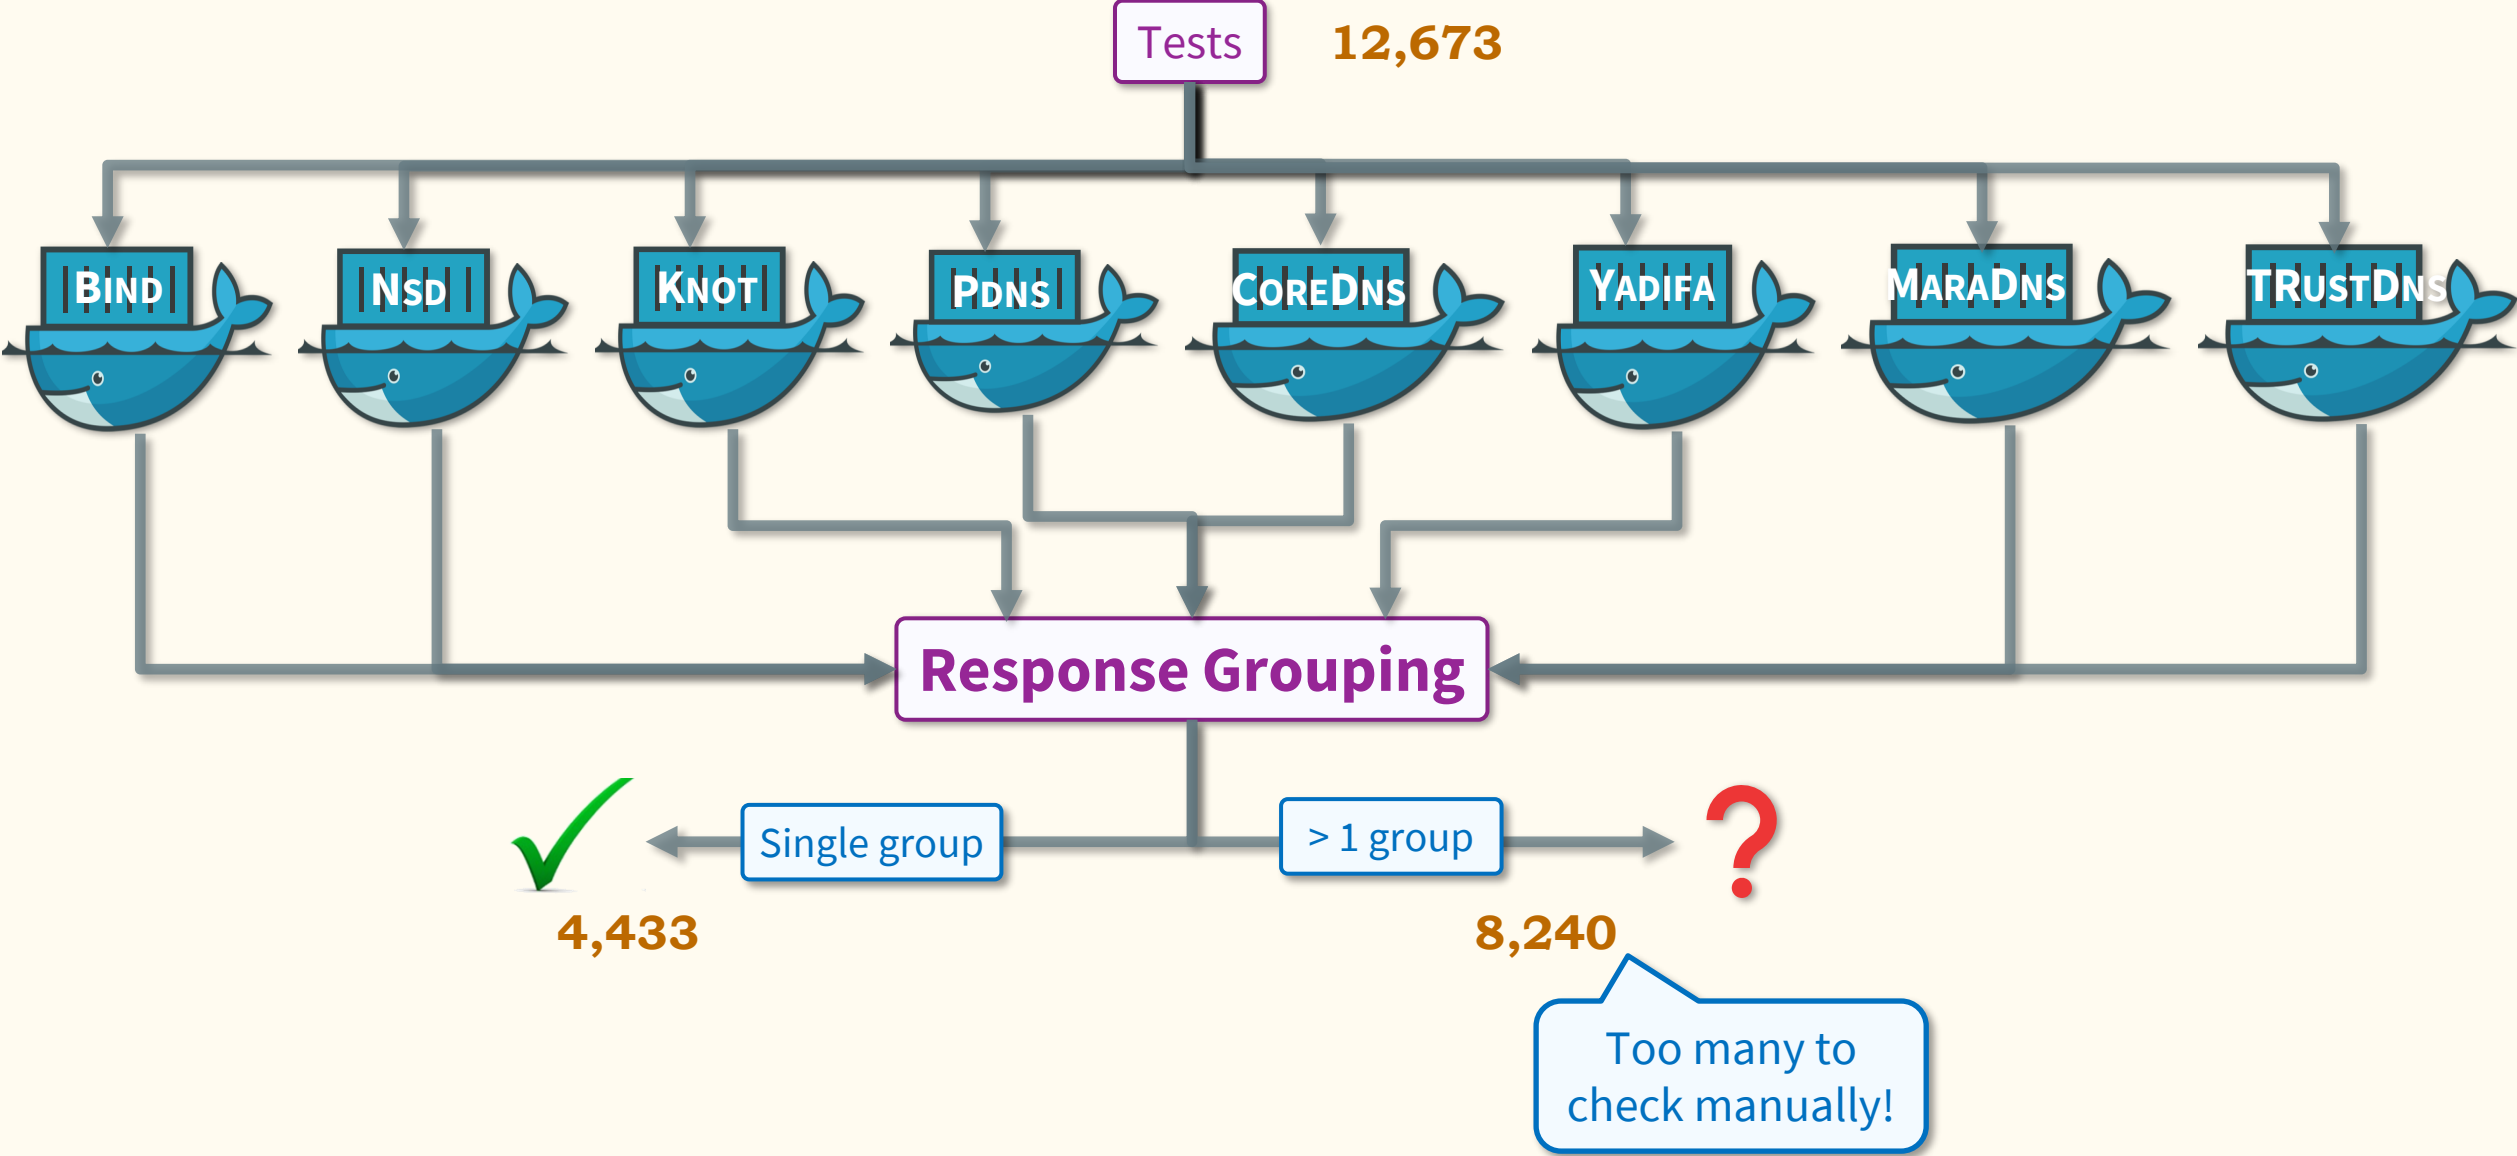
Differential Testing



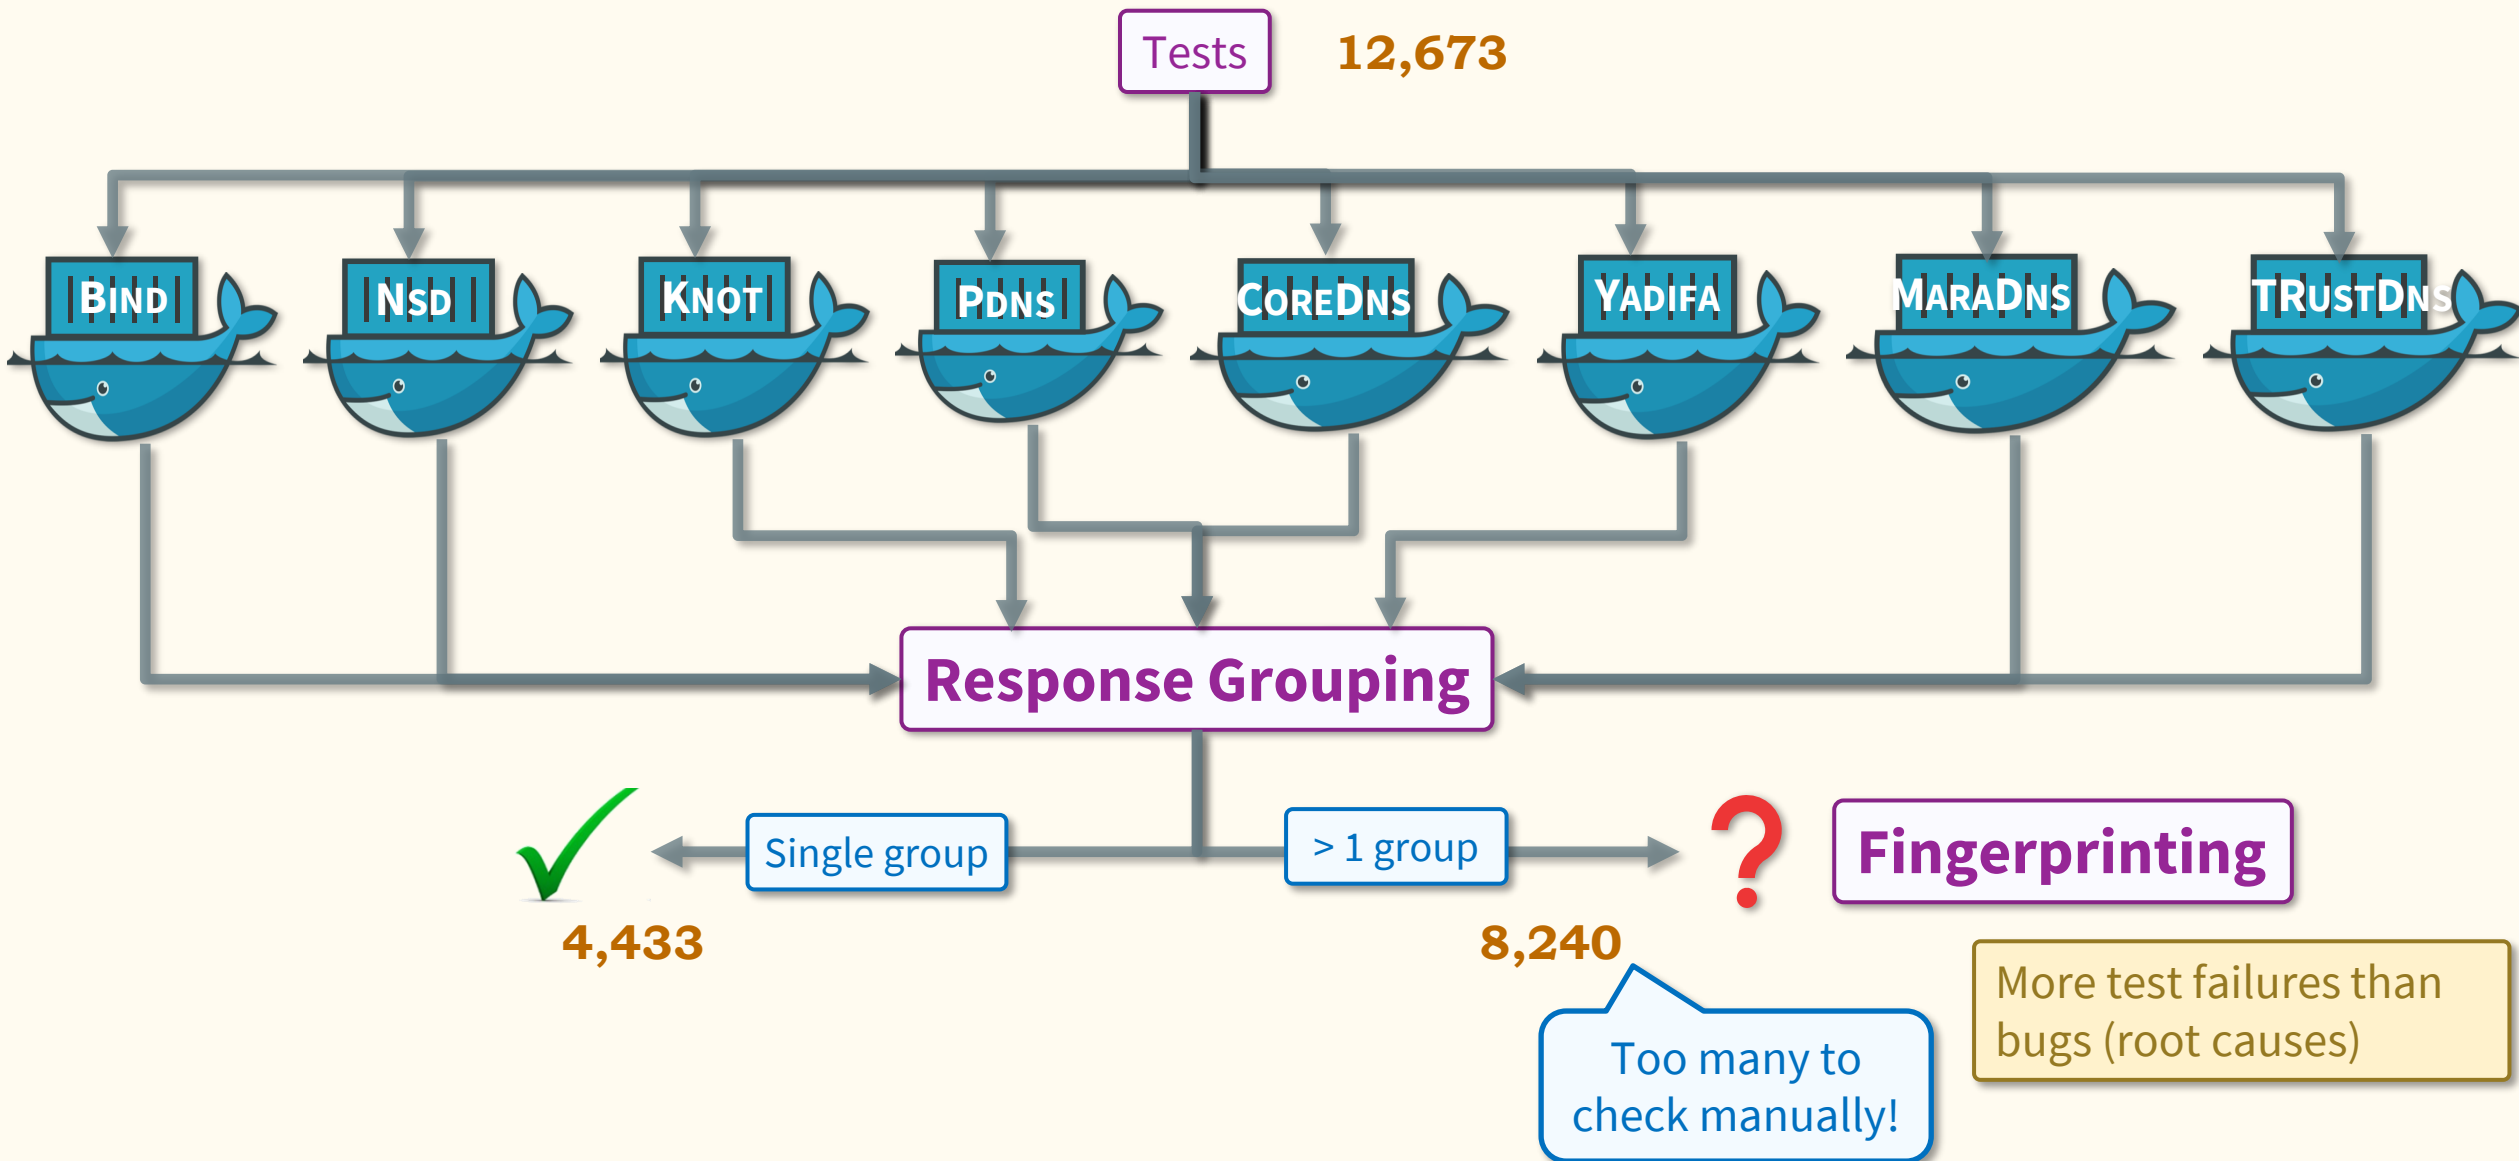
Differential Testing



Differential Testing



Differential Testing



Fingerprinting

Model Case	Number of Tests	Number of Tests Failing
E1	3180	239
E2	12	10
E3	96	12
E4	6036	5312
W1	60	33
W2	24	21
W3	18	16
D1	230	65
R1	2980	2529
R2	37	3

© **Fingerprint** failed tests

Fingerprinting

Model Case	Number of Tests	Number of Tests Failing
E1	3180	239
E2	12	10
E3	96	12
E4	6036	5312
W1	60	33
W2	24	21
W3	18	16
D1	230	65
R1	2980	2529
R2	37	3

- ◎ **Fingerprint** failed tests
- ◎ Based on **model case** and the **unique implementations in each group** from the responses

Fingerprinting

Model Case	Number of Tests	Number of Tests Failing
E1	3180	239
E2	12	10
E3	96	12
E4	6036	5312
W1	60	33
W2	24	21
W3	18	16
D1	230	65
R1	2980	2529
R2	37	3

- ◎ **Fingerprint** failed tests
- ◎ Based on **model case** and the **unique implementations in each group** from the responses
- ◎ Example fingerprint – $\langle \text{R1}, \{\text{NSD, KNOT, POWERDNS, YADIFA}\}, \{\text{BIND, COREDNS}\}, \{\text{TRUSTDNS, MARADNS}\} \rangle$

Fingerprinting

Model Case	Number of Tests	Number of Tests Failing
E1	3180	239
E2	12	10
E3	96	12
E4	6036	5312
W1	60	33
W2	24	21
W3	18	16
D1	230	65
R1	2980	2529
R2	37	3

- ◎ **Fingerprint** failed tests
- ◎ Based on **model case** and the **unique implementations in each group** from the responses
- ◎ Example fingerprint – $\langle \text{R1}, \{\text{NSD}, \text{KNOT}, \text{POWERDNS}, \text{YADIFA}\}, \{\text{BIND}, \text{COREDNS}\}, \{\text{TRUSTDNS}, \text{MARADNS}\} \rangle$
- ◎ Unlikely for different unique bugs to have the same fingerprint

Fingerprinting

Model Case	Number of Tests	Number of Tests Failing	Number of Fingerprints
E1	3180	239	7
E2	12	10	5
E3	96	12	3
E4	6036	5312	11
W1	60	33	8
W2	24	21	9
W3	18	16	1
D1	230	65	4
R1	2980	2529	27
R2	37	3	1

- ◎ **Fingerprint** failed tests
- ◎ Based on **model case** and the **unique implementations in each group** from the responses
- ◎ Example fingerprint – $\langle \mathbf{R1}, \{\mathbf{NSD}, \mathbf{KNOT}, \mathbf{POWERDNS}, \mathbf{YADIFA}\}, \{\mathbf{BIND}, \mathbf{COREDNS}\}, \{\mathbf{TRUSTDNS}, \mathbf{MARADNS}\} \rangle$
- ◎ Unlikely for different unique bugs to have the same fingerprint

Bugs Found

Implementation	Bugs Found	Bug Type	Confirmed
BIND	Sibling glue records not returned	Wrong Additional	✓
	Zone origin glue records not returned	Wrong Additional	✓
	DNAME recursion denial-of-service	Server Crash	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
NSD	DNAME not applied recursively	Wrong Answer	✓
	Wrong RCODE when * is in Rdata	Wrong RCODE	✓
	Used NS records below delegation	Wrong Answer	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
POWERDNS	CNAME followed when not required	Wrong Answer	✓
	pdnsutil check-zone DNAME-at-apex	Preprocessor Bug	✓
KNOT	incorrect record synthesis	Wrong Answer	✓
	DNAME not applied recursively	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
	Error in DNAME-DNAME loop KNOT test	Faulty KNOT Test	✓
	Wrong RCODE for synthesized record	Wrong RCODE	✓
COREDNS	NXDOMAIN for existing domain	Wrong RCODE	✓
	Wrong RCODE for CNAME target	Wrong RCODE	✓
	Wildcard CNAME loops & DNAME loops	Server Crash	✓
	Wrong RCODE for synthesized record	Wrong RCODE	?
	CNAME followed when not required	Wrong Answer	?
	Sibling glue records not returned	Wrong Additional	✓
YADIFA	CNAME chains not followed	Wrong Answer	✓
	Wrong RCODE for CNAME target	Wrong RCODE	✓
	Used records below delegation	Wrong Answer	✓
MARADNS [†]	AA flag set for zone cut NS RRs	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
TRUSTDNS [†]	wildcard match only one label	Wrong Answer	✓
	Used records below delegation	Wrong Answer	✓
	AA flag set for zone cut NS RRs	Wrong Flag	✓
	CNAME loops crash the server	Server Crash	✓


[†]Implementations with unreported issues due to missing or unimplemented features

Testing New Implementations


1. Generate a Docker image
2. Start a container with a host port mapped to port 53 of the container
3. A small Script to:
 - Stop the running server in the container
 - Copy the test zone file
 - Modify the configuration (metadata)
 - Start the server
4. Pick other implementations to compare with

Custom Tests

FERRET




Custom Tests



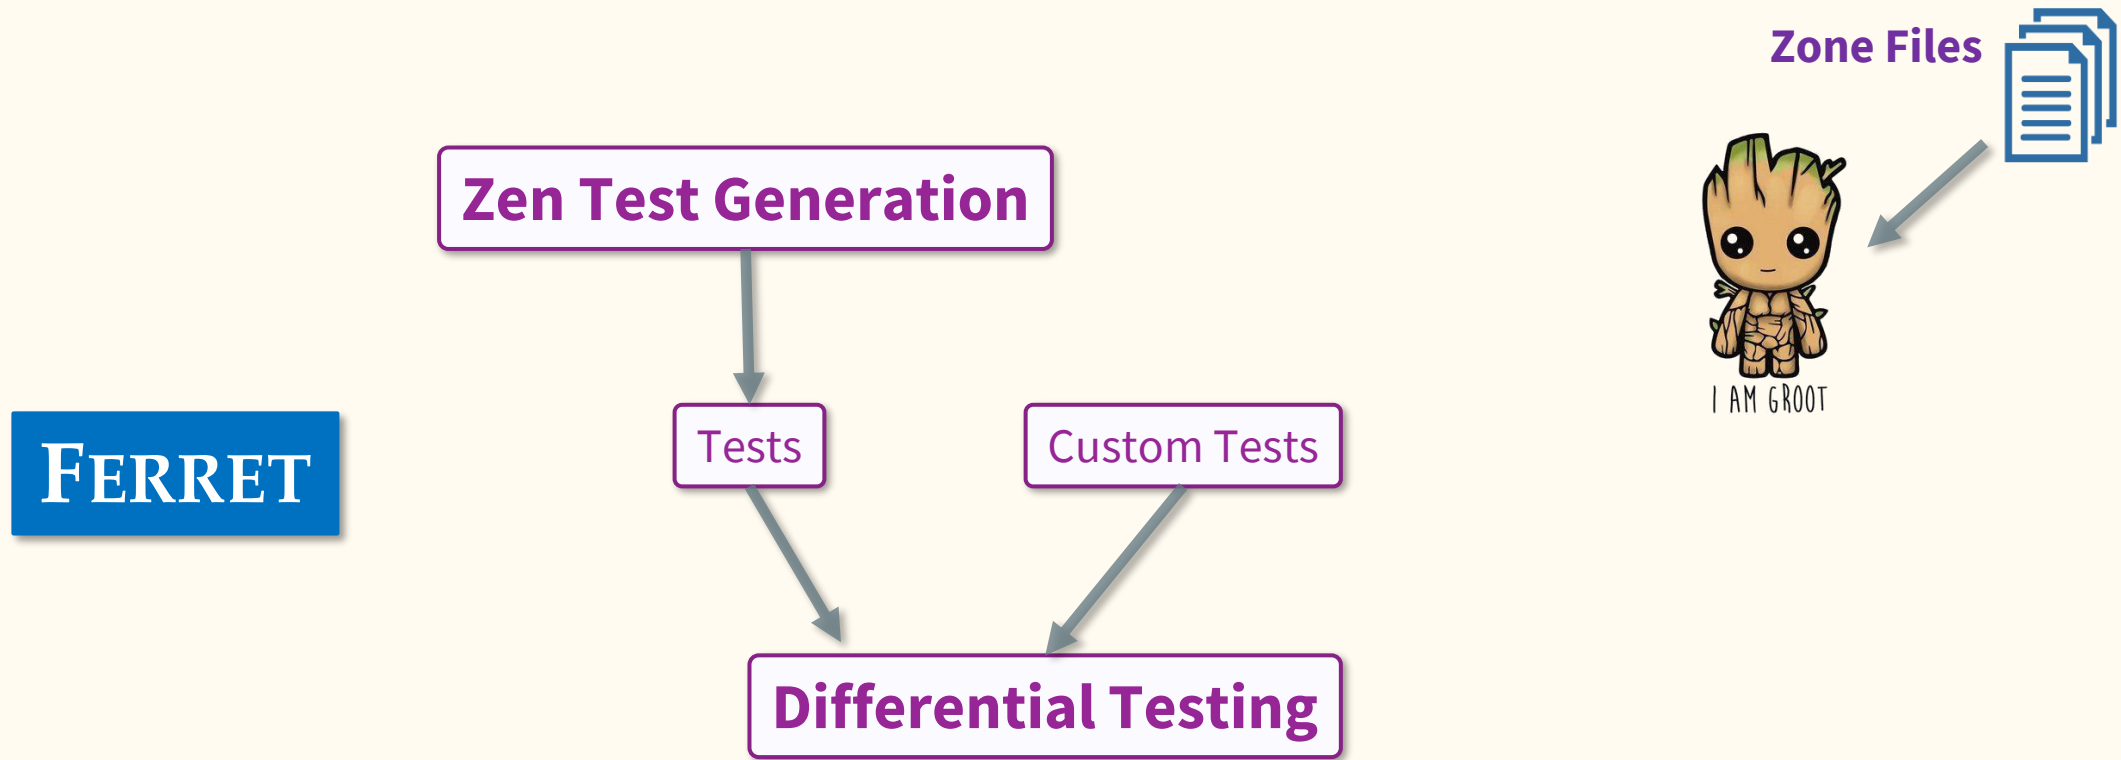
Organization Zone Files

Zone Files



How do we test for any implementation-specific behaviors on *our zone files*?

Organization Zone Files

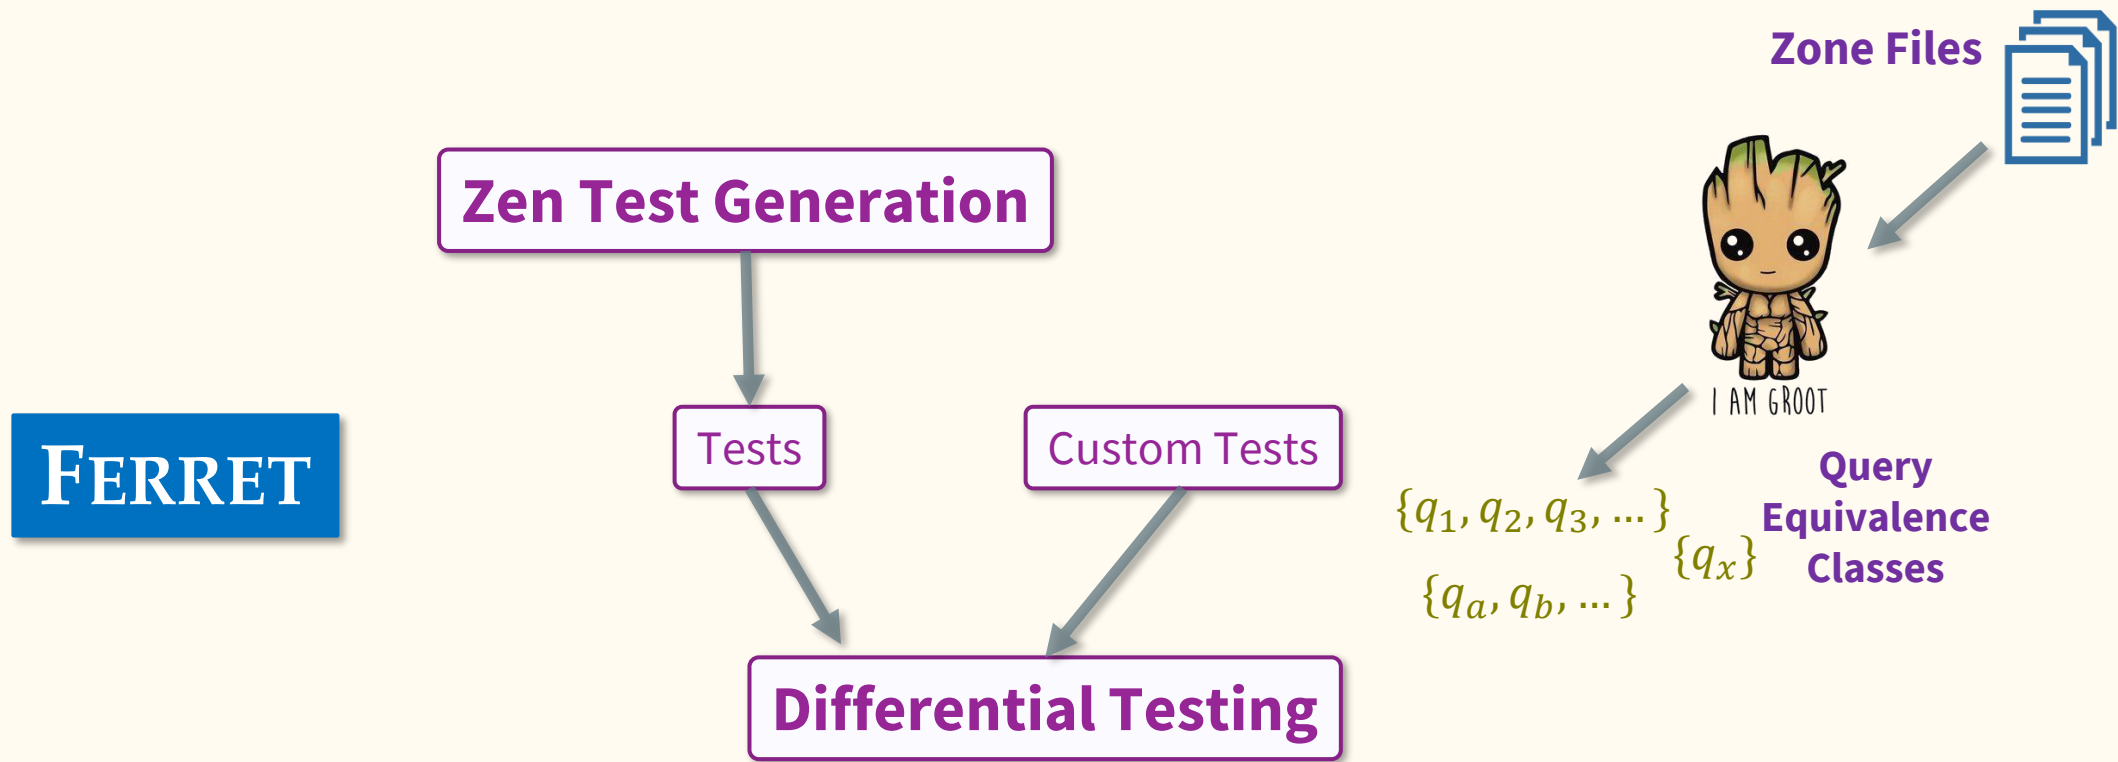


How do we test for any implementation-specific behaviors on *our zone files*?

Use GROOT to generate query *equivalence classes*

(see our paper/tool for details)

Organization Zone Files

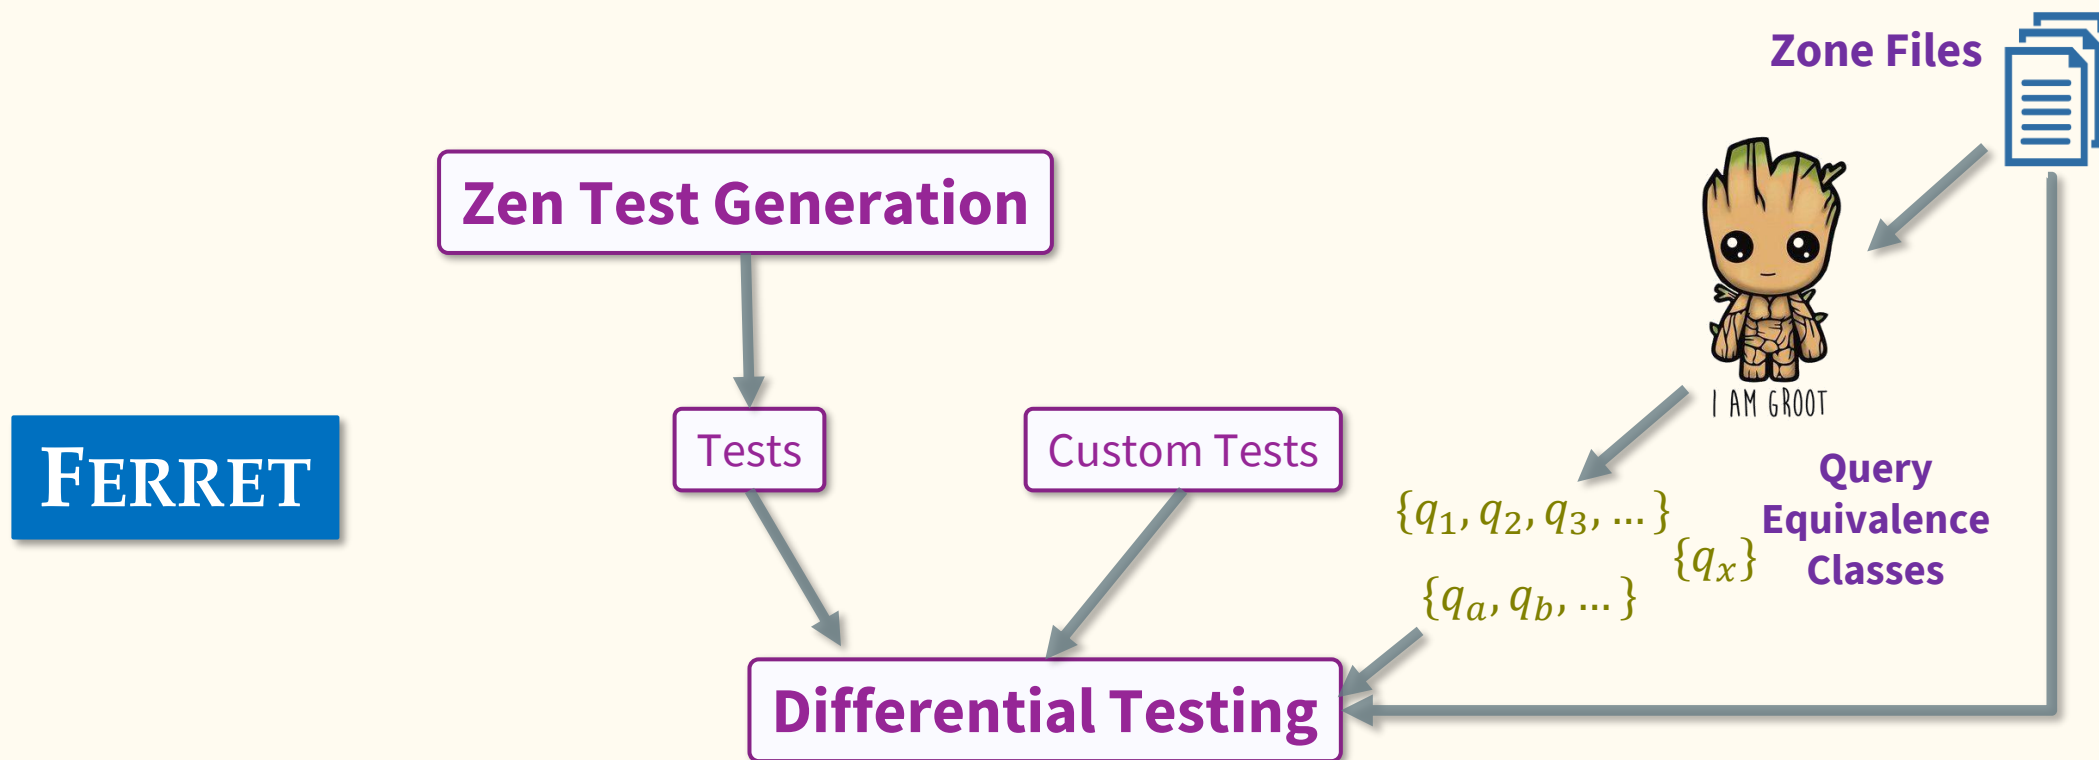


How do we test for any implementation-specific behaviors on *our zone files*?

Use GROOT to generate query *equivalence classes*

(see our paper/tool for details)

Organization Zone Files



How do we test for any implementation-specific behaviors on *our zone files*?

Use GROOT to generate query *equivalence classes*

(see our paper/tool for details)

No Two Nameservers Agree!

- ▣ Nobody agrees with RFCs too!
- ▣ RFCs do the job well but there are gaps and ambiguities
 - ↳ CNAME loops should be signaled as errors (RFC 1034)
 - At what point?
 - Should it be unrolled at all?
 - Should the loop RRs be returned?
 - ↳ Is a synthesized CNAME from DNAME perfect response to a CNAME query?
- ▣ When RFCs are open to interpretation, implementations make choices based on – performance, resource constraints, safety, ...
- ▣ Should resolvers account for different choices? (complex resolvers, interoperability issues)
Or
Should the RFCs be more verbose and stringent?

Conclusion

- ◉ **FERRET** – Our tool for **automatic** test generator for nameserver implementations
- ◉ Generates **high-coverage** test suites stress testing many corner cases of RFCs
- ◉ **Differential testing** to compare multiple implementations
- ◉ Tested **8** implementations
- ◉ Found **30** new bugs
- ◉ <https://github.com/dns-groot/Ferret>,
<https://github.com/dns-groot/groot>
- ◉ Reach me at: sivakesava@cs.ucla.edu

📁 DifferentialTesting	Port fix
📁 TestGenerator	Docker cp and bind commands updated
📄 .gitignore	Files upload
📄 LICENSE	Initial commit
📄 README.md	Codecov badge updated
📄 tool.jpg	Files upload

README.md	
<h2>Ferret</h2>	
License	MIT
codecov	99%