APNIC 36 Conference

DNS Zone Merging

Zheng Wang

wangzheng@conac.cn

China Organizational Name Administration Center

CONAC
政务和公益机构域名注册管理中心

A Corner Case DNS Configuration

May facilitate DNS zone management

Concerns about DNS compatibility and security risks arise

Background and Problem

A set of DNS servers may host many (largely small) zones falling into the same parent zone

- Each zone is individually configured in the conf file
- Complicate conf file management: a single zone update triggers conf file change
- Complicate zone file management: zone content update is located in varied zone file
- Server synchronization needs additional mechanism besides XFR
- Slow startup speed: many complaints about BIND 9 on this until an optimization method is released in July

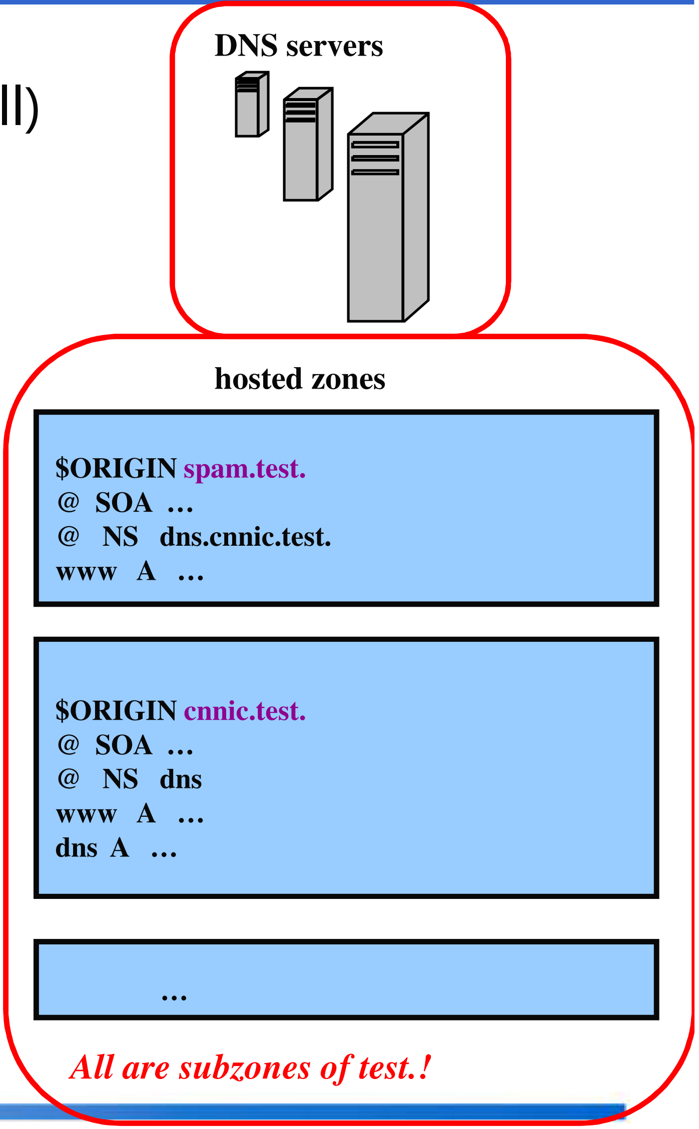

```
$ORIGIN spam.test.
@  SOA  …
@   NS   dns.cnnic.test.
www   A   …
```

```
$ORIGIN cnnic.test.
@  SOA  …
@   NS   dns
www   A   …
dns  A   …
```

```
…
```

All are subzones of test.!

Solution

A simple idea: merge the zones into one?

- The aggregated zone must be the parent zone to embrace all subzones
- Equivalent to rewriting a zone below the delegation of its apex
- The zone contains all subzones’ records along with the SOA and NS records at the apex except for all subzones’ SOA records
- Configure the parent zone in the conf file
- Possible conflicts with DNS specifications?

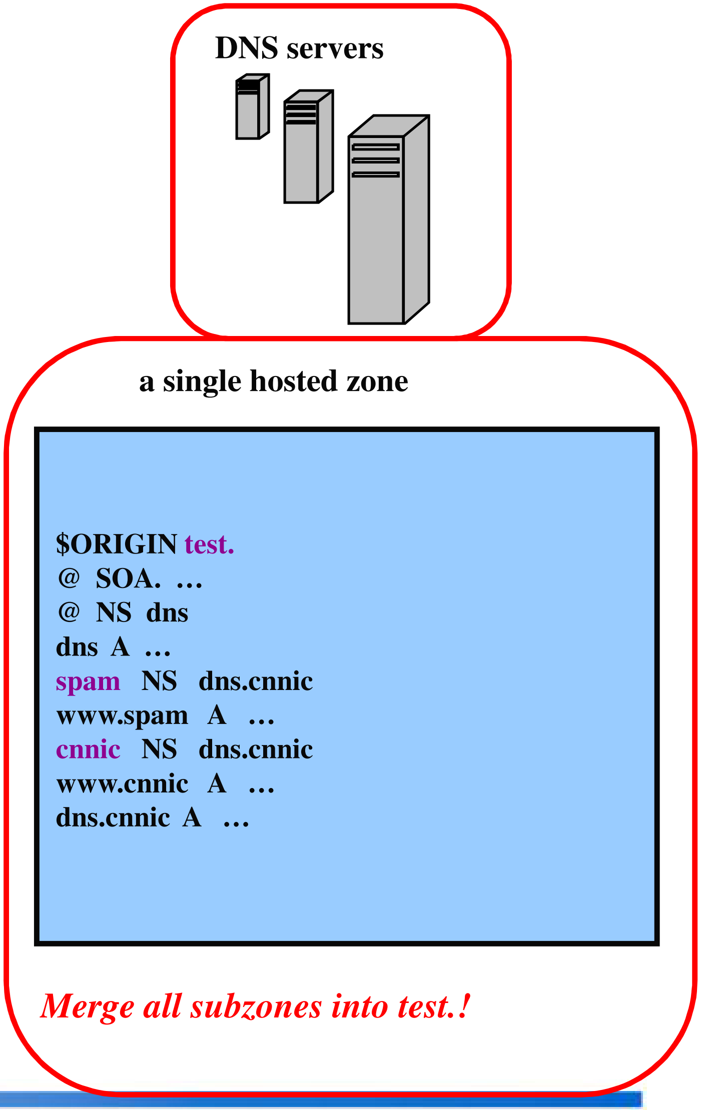

```
$ORIGIN test.
@  SOA.  …
@  NS  dns
dns  A  …
spam   NS   dns.cnnic
www.spam   A   …
cnnic   NS   dns.cnnic
www.cnnic   A   …
dns.cnnic  A   …
```

Merge all subzones into test.!

Implications

From the perspective of the DNS server:

- It believes itself serves the authoritative parent zone
- So when queries from the resolver arrives, it response just as the authoritative server as the parent zone

Only queries for the subzones can be directed to it due to the corresponding delegations in the parent zone

If applicable:

- A single zone is configured in the conf file
- Easy conf file management: conf file remains stable regardless of zone update
- Easy zone file management: all zone content updates are located in one zone file
- Easy server synchronization: XFR is enough
- Fast startup: minimized zone file amounts

Implications (Cont’)

Possible conflicts with DNS specifications:

- The subzone’ s SOA and NS records are missing from the authority section of response which may be not expected by the resolver
- How does the resolver explain it? Or can the resolver accept it?

Tests on DNS Implementations

As viewed from an individual authoritative server, zone configurations and zone content are compliant to DNS specifications, the test is only necessary for the resolver implementation.

Authoritative zone file configuration (BIND 9.6.1cn2-P1)

parent zone

```
$ORIGIN test.

@          IN    SOA          …

@          IN    NS      dns

dns        IN    A       218.241.108.65

spam       IN    NS      dns.cnnic

cnnic      IN    NS      dns.cnnic

dns.cnnic      IN    A      218.241.108.66

cnnic2         IN    A      218.241.108.66
```

merged child zones

```
$ORIGIN test.

@          IN    SOA          …

@          IN    NS      dns

dns        IN    A       218.241.108.66

spam       IN    NS      dns.cnnic
cnnic      IN    NS      dns.cnnic

dns.cnnic      IN    A     218.241.108.66
www.spam       IN    A     218.241.108.66
www.cnnic      IN    A     218.241.108.66
```

Tests on DNS Implementations (Cont’)

dig +trace results show the DNS resolution path

BIND resolver (BIND 9.6.1cn2-P1)

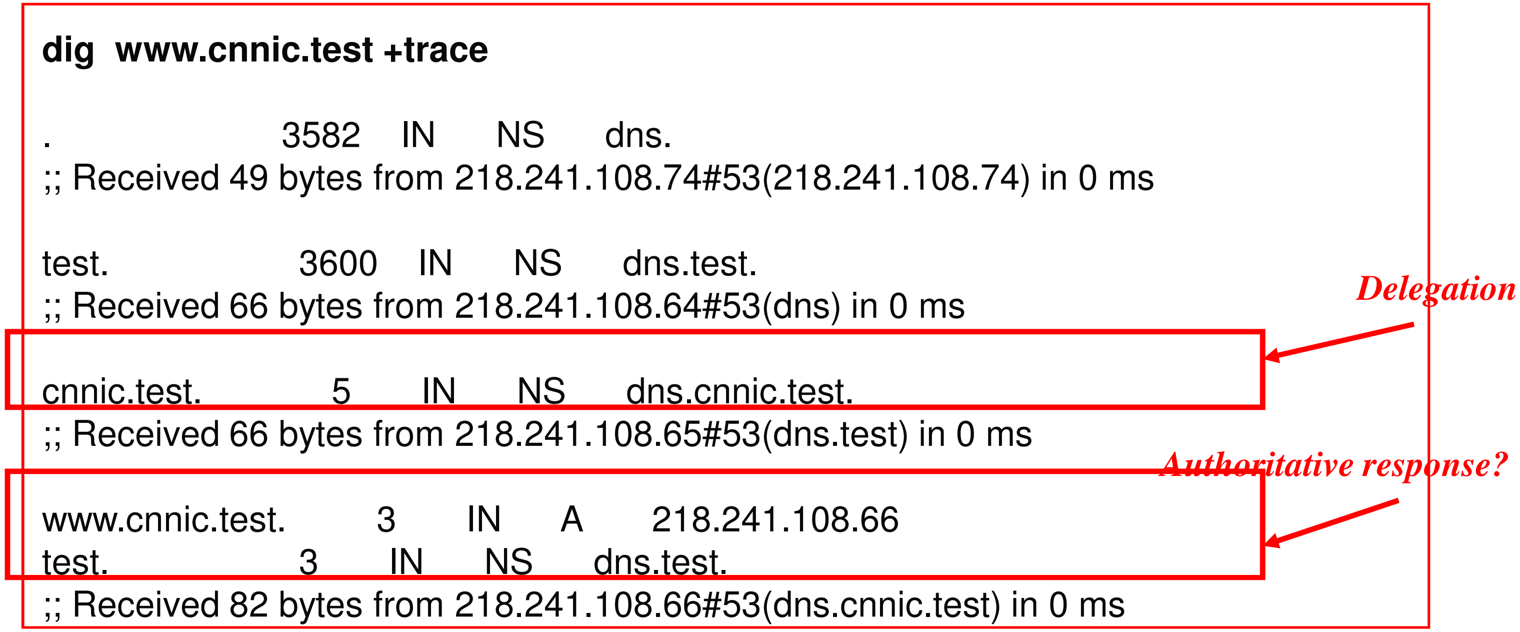

```
dig  www.cnnic.test +trace

.                    3582    IN      NS      dns.
;; Received 49 bytes from 218.241.108.74#53(218.241.108.74) in 0 ms

test.                3600    IN      NS      dns.test.
;; Received 66 bytes from 218.241.108.64#53(dns) in 0 ms

cnnic.test.          5       IN      NS      dns.cnnic.test.
;; Received 66 bytes from 218.241.108.65#53(dns.test) in 0 ms

www.cnnic.test.      3       IN      A       218.241.108.66
test.                3       IN      NS      dns.test.
;; Received 82 bytes from 218.241.108.66#53(dns.cnnic.test) in 0 ms
```

Delegation

Authoritative response?

Tests on DNS Implementations (Cont’)

BIND resolver can successfully return all pertinent records

BIND resolver (BIND 9.6.1cn2-P1)

```
dig www.cnnic.test

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 6181
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;www.cnnic.test.                IN    A

;; ANSWER SECTION:
www.cnnic.test.     3     IN    A     218.241.108.66

;; AUTHORITY SECTION:
cnnic.test.         5     IN    NS    dns.cnnic.test.

;; Query time: 1 msec
;; SERVER: 218.241.108.74#53(218.241.108.74)
```

Synthetized from the delegation

Tests on DNS Implementations (Cont’)

dig +trace results show the DNS resolution path of a negative answer

BIND resolver (BIND 9.6.1cn2-P1)

```
dig ww1.cnnic.test +trace

.                       3576    IN      NS      dns.
;; Received 49 bytes from 218.241.108.74#53(218.241.108.74) in 0 ms

test.                   3600    IN      NS      dns.test.
;; Received 66 bytes from 218.241.108.64#53(dns) in 0 ms

cnnic.test.             5       IN      NS      dns.cnnic.test.
;; Received 66 bytes from 218.241.108.65#53(dns.test) in 0 ms

test.                   3       IN      SOA     dns.test. cert.cnnic.test. 2 20 20 604800 3600
;; Received 76 bytes from 218.241.108.66#53(dns.cnnic.test) in 0 ms
```

Tests on DNS Implementations (Cont’)

Negative answer test: Nothing unusual except for the missing SOA record

BIND resolver (BIND 9.6.1cn2-P1)

```
dig ww1.cnnic.test

;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 50985
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;ww1.cnnic.test.                 IN      A

;; Query time: 1 msec
;; SERVER: 218.241.108.74#53(218.241.108.74)
;; WHEN: Tue Oct 18 09:55:56 2011
;; MSG SIZE  rcvd: 32
```

Tests on DNS Implementations (Cont’)

UNBOUND also supports, but without synthetized authority section

UNBOUND 1.2.0

```
dig www.cnnic.test

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 52082
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.cnnic.test.                    IN      A

;; ANSWER SECTION:
www.cnnic.test.        3       IN      A       218.241.108.66

;; Query time: 4 msec
;; SERVER: 218.241.108.74#53(218.241.108.74)
```

Tests on DNS Implementations (Cont’)

UNBOUND supports negative response, but also without SOA record

UNBOUND 1.2.0

```
dig ww1.cnnic.test

;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 13508
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;ww1.cnnic.test.                    IN      A

;; Query time: 1 msec
;; SERVER: 218.241.108.74#53(218.241.108.74)
```

Tests on DNS Implementations (Cont’)

Does this configuration make it possible for the server administrator to compromise its parent zone?

```
$ORIGIN test.
@          IN    SOA         …
@          IN    NS      dns
dns        IN     A     218.241.108.66
spam       IN     NS    dns.cnnic
cnnic      IN     NS    dns.cnnic
dns.cnnic     IN    A     218.241.108.66
www.spam   IN    A     218.241.108.66
www.cnnic  IN    A     218.241.108.66

ww1.spam   IN   CNAME cnnic2

cnnic2.       IN   A  218.241.108.65
```

Zone infringement!

Seemingly viable through parent zone rewriting, but how to link the subzone records in service to the residual space of the parent zone?

CNAME chain may do this!

- Configure a CNAME record to point to any record in the zone interested
- The response is sure to include the in-zone CNAME chain
- The only problem is whether the resolver would accept the CNAME chain

Tests on DNS Implementations (Cont’)

Authoritative response of the merged zone

```
dig @218.241.108.66 ww1.spam.test

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 13161
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADDITIONAL: 1

;; QUESTION SECTION:
;ww1.spam.test.                 IN      A

;; ANSWER SECTION:
ww1.spam.test.          3       IN      CNAME   cnnic2.test.
cnnic2.test.            3       IN      A       218.241.108.65

;; AUTHORITY SECTION:
test.                   3       IN      NS      dns.test.

;; ADDITIONAL SECTION:
dns.test.               3       IN      A       218.241.108.66

;; Query time: 0 msec
;; SERVER: 218.241.108.66#53(218.241.108.66)
```

Different from the parent zone!

Tests on DNS Implementations (Cont’)

BIND resolver (BIND 9.6.1cn2-P1)

```
dig ww1.spam.test

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 38370
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;ww1.spam.test.            IN     A

;; ANSWER SECTION:
ww1.spam.test.       3      IN      CNAME   cnnic2.test.
cnnic2.test.        5      IN      A       218.241.108.66

;; AUTHORITY SECTION:
test.              3600    IN      NS      dns.test.
```

From the parent zone!

Tests on DNS Implementations (Cont’)

UNBOUND 1.2.0

```
dig ww1.spam.test

;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 38296
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADDITIONAL: 1

;; QUESTION SECTION:
;ww1.spam.test.              IN      A

;; ANSWER SECTION:
ww1.spam.test.          3       IN      CNAME   cnnic2.test.
cnnic2.test.        5       IN      A       218.241.108.66

;; AUTHORITY SECTION:
test.               3600    IN      NS      dns.test.

;; ADDITIONAL SECTION:
dns.test.           3600    IN      A       218.241.108.65

;; Query time: 1 msec
;; SERVER: 218.241.108.74#53(218.241.108.74)
```

From the parent zone!

Tests on DNS Implementations (Cont’)

BIND and UNBOUND are not caught in the trap!

They do not accept the chained results as the final authoritative answer

The canonical name are queried in a dedicated separate request whose response is handle by the parent zone

Zone Loading Tests

Startup performance has been significantly improved!

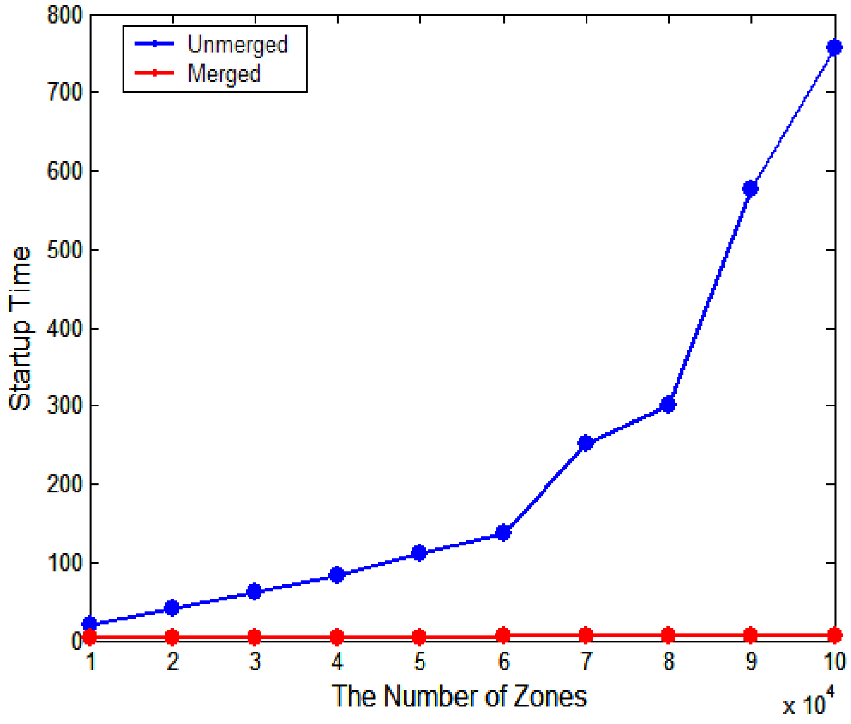

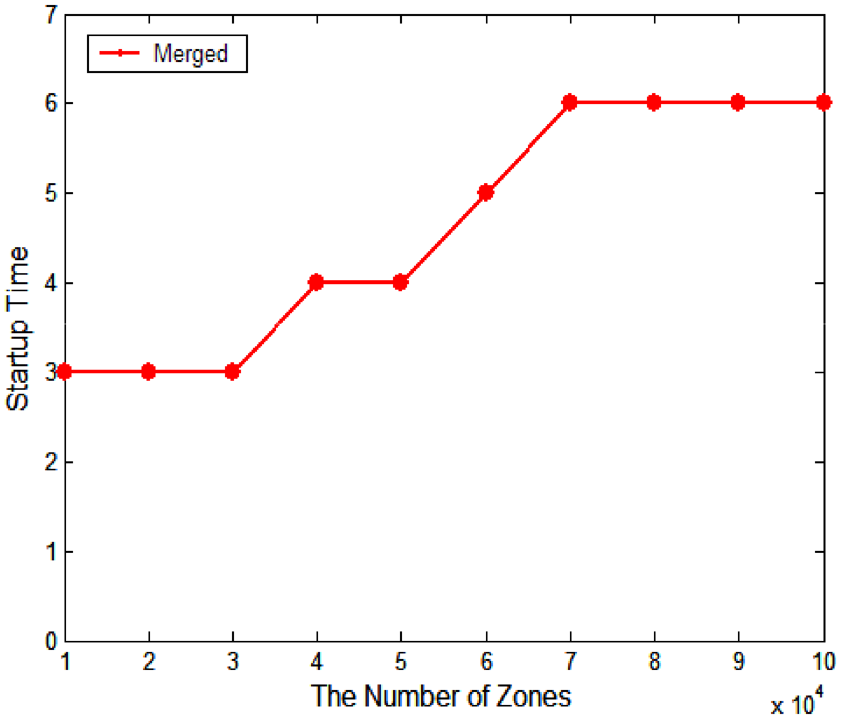

1 GB of RAM, one quad-core processors running at 3.2 GHz, and standard SATA drives configured without any raid or mirroring.

Each test zone was loaded from one of many different physical files. Each file was identical in its content, which included one SOA record, two NS records, and two A records.

Summary

A DNS zone merging method is proposed

Though possibly problematic in DNS compatibilty, it does work in at least BIND and Unbound implementations

Parent zone compromise risks exists but are avoided by BIND and Unbound implementations

CONAC
政务和公益机构域名注册管理中心

Thank You !