

DNS Extensions Working Group  
Internet Draft  
Intended status: Standard Track

Cindy WANG  
Jian JIN  
Feng HAN

Xiaodong LEE

Expires: August 7, 2010  
2010

CNNIC  
February 8,

**A mechanism for synchronization across name servers on zone creation  
draft-wang-dnsextnewzone-notify-00.txt**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 7, 2010.

## Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Abstract

This memo describes the NEWZONE\_NOTIFY opcode for DNS, by which a primary master server advises a set of slave servers that there is a zone has been created and that a query should be initiated to discover the new zone data.

This draft also specifies a mechanism for the slave servers to achieve authenticated synchronization of zone data as well as zone synchronization information with the primary when a zone is created on the primary.

## Table of Contents

<a href="#">1. Introduction.....</a>	<a href="#">4</a>
--------------------------------------	-------------------

<a href="#">2.</a>	Conventions used in this document.....	<a href="#">5</a>
<a href="#">3.</a>	Definitions and Invariants.....	<a href="#">5</a>
<a href="#">4.</a>	NEWZONE_NOTIFY message.....	<a href="#">6</a>
<a href="#">5.</a>	<b>Automating the synchronization on zone creation across multiple</b> name servers.....	<a href="#">7</a>
<a href="#">6.</a>	Security Considerations.....	<a href="#">9</a>
<a href="#">7.</a>	References.....	<a href="#">9</a>

## [1.](#)   **Introduction**

For large and busy domain name registrars, the zone creation operations resulted from frequent domain name registrations are almost daily routines. However, for these operations there are no technical specifications for automatic zone synchronization across multiple name servers. Moreover, the manual operations turn into heavy burden for administrators when there is large number of name servers authoritative for the zones.

The major obstacle to the synchronization in the above situation is that, specified by the original design of DNS, when authority zones are created, they must be declared to have one or more authoritative name servers, usually consisting of one primary name server and several secondary name servers. The configuration of the synchronization relationships among the name servers depends upon out of band information and manual processes, which are normally specified with the zone creation.

This draft specifies a mechanism for the slave servers to achieve authenticated synchronization of both zone data and dependency configuration with the master servers when a zone is created.

## **2. Conventions used in this document**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [1].

## **3. Definitions and Invariants**

The following definitions are used in this document, and intended to be consistent with [[RFC1996](#)] and [[RFC2136](#)]:

- o        Slave: an authoritative server which uses zone transfer to retrieve the zone. All slave servers are named in the NS RRs for the zone.
- o        Master: an authoritative server configured to be the source of zone transfer for one or more slave servers.
- o        Primary Master: the master server at the root of the zone replication dependency graph. The primary master is named in the zone's SOA MNAME field and optionally by an NS RR. There is by definition only one primary master server per zone and the source of zone creation for one or more slave servers.
- o        Notify Set: set of servers to be notified of changes to some new zone. Default is all servers named in the NS RRset of the zone, except for any server also named in the SOA MNAME. Some implementations will permit the name server administrator to override this set or add elements to it (such as, for example, the "also-notify" implemented in BIND [[DNS BIND](#)]).
- o        Trusted-Master Set: set of servers whose notification could be trusted by the slave and the set is specified out-of-bind.
- o        Dependency graph: organization of the zone's name servers, such that there is a primary master, and all other servers must request zone replication either from the primary master or from some slave which is also a master. NO loops are permitted in the dependency graph. The dependency graph is created with the zone by the administrator on the primary master. For example, the set of servers for a zone is {p, s1, s2, s3, s4, s5, s6, s7}, where p is the address of the primary and {s1... s7} addresses of the slaves. The dependency graph could be defined as {{p->s1}, {p->s2}, {p->s3}, {s1->s2}, {s2->s3}, {s2->s4}, {s2->s5}, {s3->s6}, {s3->s7}}, where "->" denotes "is the master of". An example of the "master of" relationship would look like, {1.2.3.4->5.6.7.8}.



#### **4. NEWZONE\_NOTIFY message**

**4.1.** When a primary master has a new zone, the master may send the created zone's name, class, type, and the name of the master from which the slave to request the zone data, to each known slave server using a protocol based on the NEWZONE\_NOTIFY opcode.

**4.2.** NEWZONE\_NOTIFY employs the DNS Message Format [[RFC 1034](#)], although it uses only a subset of the available fields. Fields not otherwise described herein are to be filled with binary zero (0).

**4.3.** NEWZONE\_NOTIFY is similar to QUERY in that it has a request message with the header QR flag "clear" and a response message with QR "set". The response message contains no useful information, but its reception by the master is an indication that the slave has received the NEWZONE\_NOTIFY and that the master can remove the slave from any retry queue for this NEWZONE\_NOTIFY event.

**4.4.** TSIG MUST be enabled between parties exchanging the NEWZONE\_NOTIFY messages.

**4.5.** The NEWZONE\_NOTIFY request has the following characteristics:  
Header:

```
query ID:   (new)
opcode:     NEWZONE_NOTIFY  (5)
resp:       NOERROR
flags:      AA
qdcount:    1
ancount:    1
```

Question Section:

```
qname:      (zone name)
qclass:     (zone class)
qtype:      * (matches all RR types)
```

Answer Section:

Name of one of the masters m which satisfies m->s, where s is the notified slave.

The defined NEWZONE\_NOTIFY event in this situation is that a zone has

been created (QTYPE=\*) on the primary master server.

## **5. Automating the synchronization on zone creation across multiple name servers**

### **5.1. Zone created on the primary master server**

**1.**        The primary master should send a NEWZONE\_NOTIFY request to all the servers in the dependency graph except for itself.

**2.**        The notifying order should be decided by the distances (number of other masters in between) of the slaves to the primary master. A slave CANNOT be notified until at least one of its masters responds back to the primary master with success.

**3.**        For slaves with more than one master, the primary master MUST send multiple notification messages, one for each master.

### **5.2. Slave Receives a NEWZONE\_NOTIFY Request from the Primary Master**

When a slave server receives a NEWZONE\_NOTIFY request enclosing the given QNAME, with QTYPE=\* and QR=0, first of all, it MUST check the authentication of the message by examining,

**1.**        The notifying IP must be present in the slave's 'Trusted-Master Set'; (protecting the slave from being flooded by malicious messages.)

**2.**        By requesting the SOA and NS record sets for the created zone from the notifying master,

**3.**        The notifying master MUST carry a SOA record for the notified zone;

**4.**        The notifying master MUST either appear in the MNAME of the SOA record or goes with one name in the NS record-set for the created zone;

**5.**        The set of NS records for the created zone, as retrieved by the slave from the notifying master, MUST include the name that goes with the IP address of the notified master.

If the notification is justified by all the above conditions, the slave should behave as though the zone given in the QNAME had been created on the primary master. It should respond to the NEWZONE\_NOTIFY message with the following actions,

- 1.** Firstly, it requests the SOA record of the named zone locally to determine whether the zone exists or not.
- 2.** If the zone exists, then the synchronization has been done from other master of the slave; otherwise, it is a zone creation notification and a zone transfer (AXFR) [[AXFR clarify](#)] should be initiated to the master specified in the answer section of the NEWZONE\_NOTIFY message from the primary master.
- 3.** In both cases, the master appearing in the answer section is configured locally to be one of the masters of the slave.
- 4.** Finally, the updates are loaded into memory and the master specified in the answer section of the NEWZONE\_NOTIFY message is added to the master list of the slave.
- 5.** Whether the AXFR succeeds or not, the slave will also send a NEWZONE\_NOTIFY response back to the NEWZONE\_NOTIFY request's source, with the following characteristics:

Header:

query ID: (same)  
opcode: NEWZONE\_NOTIFY (5)  
RCODE: NOERROR (AXFR succeeds) or Server failure (AXFR

fails)

flags: OR AA  
qdcount: 1

Question Section:

qname: (zone name)  
qclass: (zone class)  
qtype: \* (matches all RR types)

Answer Section:

MUST be EMPTY

### **5.3. Primary Master Receives a NEWZONE\_NOTIFY Response from Slave**

- 1.** If a NEWZONE\_NOTIFY response from a slave with RCODE= Server failure arrives, the primary master keeps the NEWZONE\_NOTIFY query in the retry queue.



**2.**        **Otherwise, if the primary master server receives a NEWZONE\_NOTIFY response from a slave with RCODE= NOERROR, it initiates the notification process to all the slaves of that slave. Next it deletes the successful query from the retry queue, thus completing the notification process of the zone creation change to the notifying slave.**

## **6. Security Considerations**

This document is believed to introduce no additional security problems to the current DNS protocol.

## **7. References**

### **7.1. Normative References**

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.

[RFC1035] Mockapetris, P., "Domain Names - Implementation and Specifications", STD 13, [RFC 1035](#), November 1987.

[RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", [RFC 1996](#), August 1996.

[RFC2136] Vixie, P., Ed., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", [RFC 2136](#), April 1997.

[RFC2845] Vixie, P., Gudmundsson, O., Eastlake, D. and Wellington, B., "Secret Key Transaction Authentication for DNS (TSIG)", [RFC 2845](#), May 2000.

### **7.2. Informative References**

[AXFR\_clarify] DNS Zone Transfer Protocol (AXFR). [draft-wang-dnsexst-axfr-clarify-12](#).

[DNS\_BIND] DNS and BIND, 5th Edition.

[PDNS] PowerDNS manual. Chapter 13. Master/Slave operation & replication.

Authors' Addresses

Cindy WANG  
CNNIC  
4, South 4th Street, Zhongguancun  
Beijing 100190  
P.R. China  
Email: wangxin@cnnic.cn

Jian Jin  
CNNIC  
4, South 4th Street, Zhongguancun  
Beijing 100190  
P.R. China  
Email: jinjian@cnnic.cn

Feng Han  
CNNIC  
4, South 4th Street, Zhongguancun  
Beijing 100190  
P.R. China  
Email: hanfeng@cnnic.cn

Xiaodong LEE  
CNNIC  
4, South 4th Street, Zhongguancun  
Beijing 100190  
P.R. China  
Email: lee@cnnic.cn

